



Strategies for Solving High-Fidelity Aerodynamic Shape Optimization Problems

Zhoujie Lyu* and Joaquim R. R. A. Martins†

Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI

Aerodynamic shape optimization based on high-fidelity models is a computational intensive endeavor. The majority of the computational time is spent in the flow solver, and in the gradient calculation. In this paper, we present two approaches for reducing the overall computational cost of the optimization. The techniques are tested using the Common Research Model wing benchmark defined by the Aerodynamic Design Optimization Discussion Group (ADODG). The aerodynamic model solves the Reynolds-averaged Navier–Stokes equations with a Spalart–Allmaras turbulence model. A gradient-based optimization algorithm is used in conjunction with an adjoint method that computes the required derivatives. The drag coefficient is minimized subject to lift, pitching moment, and geometric constraints. The first approach uses Richardson’s extrapolation to approximate the flow solutions and gradients. The second approach performs multilevel optimization with a series of grid sizes. We found the multilevel approach to be the most effective, achieving a 84.5% reduction in computational cost.

I. Introduction

Advances in high-performance computing hardware and algorithms have enabled the ever-increasing fidelity of the computational fluid dynamics (CFD) models used for evaluating aircraft performance. As the computational time for a given CFD model reduces below a certain level, it becomes feasible to use it together with numerical optimization to perform aircraft design optimization. While there are various possible optimization techniques, the use of gradient-based algorithms together with an adjoint method that computes the required gradients efficiently has proven to be particularly effective. This enables aircraft designers to shorten design cycle times and thus explore the design space more effectively. They can also obtain detailed designs earlier in the design process, allowing them to better understand the design trade-offs and to make more informed design decisions.

Aerodynamic shape optimization based on high-fidelity models is a computational intensive endeavor. The majority of the computational time is spent in the flow solver and in the gradient calculation. There are several possible ways to reduce the overall optimization time. One way is to reduce the flow solution time. This has been extensively researched by the CFD community. Commonly used methods, such as multigrid [1], pre-conditioning [2], and variations on Newton-type methods [3, 4], can improve the convergence of the solver, thus reducing the overall optimization time.

The second way to reduce the overall computational cost is to reduce the gradient computational time, which was pioneered by Jameson [5] through the development of adjoint method, which efficiently compute gradients with respect to large numbers of shape design variables. With an efficient adjoint implementation, the cost of computing the gradient of a single function of interest with respect to hundreds or thousands of shape design variables is roughly of the same order of the cost of one flow solution [6]. Those methods have been successfully applied in recent aerodynamic shape optimizations [7, 8, 9, 10, 11, 12].

As these numerical techniques become mature, there is a need to research new ways to reduce the optimization cost. In this paper, we explore two approaches to further improve the efficiency of aerodynamic shape optimization. The first idea is to use Richardson’s extrapolation to approximate the flow solution and gradients of a fine grid using the results of coarse grids. The second approach is to perform grid sequencing at the optimization level. This is inspired by the multigrid method in CFD. We use smaller grids to accelerate the convergence of a large grid. We perform the

*PhD Candidate, Department of Aerospace Engineering, University of Michigan, AIAA Student Member

†Associate Professor, Department of Aerospace Engineering, University of Michigan, AIAA Associate Fellow

optimization first on a smaller first until a certain level of optimality is achieved. Then, we move on to the next grid level and use the design variables from the previous grid level as the initial design variables. This process is repeated until the last grid level has converged. We compare both approaches with a direct optimization with only the fine grid. The comparison of the optimized designs, convergence and computational cost is presented in the paper.

This study is performed using the benchmark case for aerodynamic design optimization developed by the Aerodynamic Design Optimization Discussion Group (ADODG)^a: the lift-constrained drag minimization of the NASA Common Research Model (CRM) wing [13, 14, 15] with a RANS model that were presented at the 2014 AIAA Science and Technology Forum and Exposition in a special session organized by the ADODG [16, 11, 17, 18, 19].

The paper is organized as follows. The numerical tools used in this work are described in Section II. The problem formulation, the mesh, and the baseline geometry are described in Section III. The aerodynamic shape optimization of the CRM wing with only the fine grid is presented in Section A. The optimization with Richardson's extrapolation is discussed in Section B, and the multilevel optimization acceleration technique is discussed in Section C, followed by the conclusions V.

II. Methodology

This section describes the numerical tools and methods that we used for the shape optimization studies. These tools are components of the framework for multidisciplinary design optimization (MDO) of aircraft configurations with high fidelity (MACH) [20]. MACH can perform simultaneous optimization of aerodynamic shape and structural sizing variables considering aeroelastic deflections [21]. However, in this paper we use only the components of MACH that are relevant for aerodynamic shape optimization: the geometric parametrization, mesh perturbation, CFD solver, and optimization algorithm.

A. Geometric Parametrization

We use a free-form deformation (FFD) volume approach to parametrize the wing geometry [22]. The FFD volume parametrizes the geometry changes rather than the geometry itself, resulting in a more efficient and compact set of geometry design variables, thus making it easier to manipulate complex geometries. Any geometry may be embedded inside the volume by performing a Newton search to map the parameter space to the physical space. All the geometric changes are performed on the outer boundary of the FFD volume. Any modification of this outer boundary indirectly modifies the embedded objects. The key assumption of the FFD approach is that the geometry has constant topology throughout the optimization process, which is usually the case in wing design. In addition, since FFD volumes are trivariate B-spline volumes, the derivatives of any point inside the volume can be easily computed. Figure 1 shows the FFD volume and the geometric control points used in the aerodynamic shape optimization. The shape design variables are the displacement of all FFD control points in the vertical (z) direction.




Figure 1. The shape design variables are the z -displacements of 720 FFD control points (red spheres).

^a<https://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG/default.aspx>, last accessed May 10, 2014

B. Mesh Perturbation

Since FFD volumes modify the geometry during the optimization, we must perturb the mesh for the CFD to solve for the modified geometry. The mesh perturbation scheme used in this work is a hybridization of algebraic and linear elasticity methods, developed by Kenway *et al.* [22]. The idea behind the hybrid scheme is to apply a linear-elasticity-based perturbation scheme to a coarse approximation of the mesh to account for large, low-frequency perturbations, and to use the algebraic warping approach to attenuate small, high-frequency perturbations. For the results in this paper, the additional robustness of the hybrid scheme is not required, so we use only the algebraic scheme.

C. CFD Solver

We use SUmb [23] as the CFD solver; it is a finite-volume, cell-centered multiblock solver for the compressible Euler, laminar Navier–Stokes, and RANS equations (steady, unsteady, and time-periodic). SUmb provides options for a variety of turbulence models with one, two, or four equations and options for adaptive wall functions. The Jameson–Schmidt–Turkel (JST) scheme [24] augmented with artificial dissipation is used for the spatial discretization. The main flow is solved using an explicit multi-stage Runge–Kutta method, along with geometric multi-grid. A segregated Spalart–Allmaras turbulence equation is iterated with the diagonally dominant alternating direction implicit (DDADI) method.

To efficiently compute the gradients required for the optimization, we have developed and implemented a discrete adjoint method for the Euler and RANS equations within SUmb [6, 25]. The adjoint implementation supports both the full-turbulence and frozen-turbulence modes, but in the present work we use the full-turbulence adjoint exclusively. We solve the adjoint equations with preconditioned GMRES [26] using PETSc [27, 28, 29]. We have previously performed extensive Euler-based aerodynamic shape optimization [30, 31] and aerostructural optimization [21, 32]. However, we have observed serious issues with the resulting optimal Euler-based designs due to the missing viscous effects. While Euler-based optimization can provide design insights, we found that the resulting optimal Euler shapes are significantly different from those obtained with RANS [6]. Euler-optimized shapes tend to exhibit a sharp pressure recovery near the trailing edge, which is nonphysical because such flow near the trailing edge would separate. Thus, RANS-based shape optimization is necessary to achieve realistic designs.

D. Optimization Algorithm

Because of the high computational cost of CFD solutions, we must choose an optimization algorithm that requires a reasonably low number of function evaluations. Gradient-free methods, such as genetic algorithms, have a higher probability of getting close to the global minimum for multimodal functions. However, slow convergence and the large number of function evaluations make gradient-free aerodynamic shape optimization infeasible with the current computational resources, especially for large numbers of design variables. Since we require hundreds of design variables, we use a gradient-based optimizer combined with adjoint gradient evaluations to solve the problem efficiently.

The optimization algorithm we use for all the results presented herein is SNOPT (sparse nonlinear optimizer) [33] through the Python interface pyOpt [34]. SNOPT is a gradient-based optimizer that implements a sequential quadratic programming method; it is capable of solving large-scale nonlinear optimization problems with thousands of constraints and design variables. SNOPT uses a smooth augmented Lagrangian merit function, and the Hessian of the Lagrangian is approximated using a limited-memory quasi-Newton method.

III. Problem Formulation

The optimization case we used in this study is to perform lift-constrained drag minimization of the NASA CRM wing using the RANS equations. In this section, we provide a complete description of the problem. The problem formulation has been previously presented in [11].

A. Baseline Geometry

The baseline geometry is a wing with a blunt trailing edge extracted from the CRM wing-body geometry [14, 15]. The NASA CRM geometry was developed for applied CFD validation studies. The CRM is representative of a contemporary transonic commercial transport, with a size similar to that of a Boeing 777. The CRM geometry has been




Figure 2. Baseline CRM wing geometry scaled by its mean aerodynamic chord. [11]

optimized, but several design features, such as an aggressive pressure recovery in the outboard wing, were introduced into the design to make it more interesting for research purposes and to protect intellectual property. This baseline geometry provides a reasonable starting point for the optimization, while leaving room for further performance improvements. In addition, the CRM was designed together with the fuselage of the full CRM configuration, so its performance is degraded when only the wing is considered.

The geometry and specifications are given by the ADODG^b. The fuselage and tail are removed from the original CRM, and the root of the remaining wing is moved to the symmetry plane. This baseline geometry is shown in Fig. 2. All coordinates are scaled by the mean aerodynamic chord (275.8 in). The resulting reference chord is 1.0, and the half span is 3.758151. The moment reference point is at $(x, y, z) = (1.2077, 0.0, 0.007669)$, while the reference area is 3.407014.

B. Mesh Convergence Study

We generate the mesh for the CRM wing using an in-house hyperbolic mesh generator. The mesh is marched out from the surface mesh using an O-grid topology to a farfield located at a distance of 25 times the span (about 185 mean chords). The nominal cruise flow condition is Mach 0.85 with a Reynolds number of 5 million based on the mean aerodynamic chord. The mesh we generated for the test case optimization contains 3.6 million cells. The mesh size and aerodynamic coefficients under the nominal operating condition are listed in Table 1.

Mesh level	Mesh size	C_D	C_L	C_M	α
$h = 0$	∞	0.01990			
L00	230,686,720	0.01992	0.5000	-0.1776	2.2199°
L0	28,835,840	0.01997	0.5000	-0.1790	2.2100°
L1	3,604,480	0.02017	0.5000	-0.1810	2.1837°
L2	450,560	0.02111	0.5000	-0.1822	2.1944°


Table 1. Mesh convergence study for the baseline CRM wing. [11]

We perform a mesh convergence study to determine the resolution accuracy of this mesh. Table 1 lists the drag and moment coefficients for the baseline meshes. We also compute the zero-grid spacing drag using Richardson's extrapolation, which estimates the drag value as the grid spacing approaches zero [35]. The zero-grid spacing drag coefficient is 199.0 counts for the baseline CRM wing. We choose L1 for this study due to the high computational cost of large grid size with direct optimization. The surface and symmetry plane meshes for the L0, L1, and L2 grid levels are shown in Fig. 3. These meshes are publicly available in the CFD general notation system (CGNS) format as supplemental materials in Lyu *et al.* [11].


C. Optimization Problem Formulation

The aerodynamic shape optimization seeks to minimize the drag coefficient by varying the shape design variables subject to a lift constraint ($C_L = 0.5$), and a pitching moment constraint ($C_{My} \geq -0.17$). The shape design variables are the z -coordinate movements of 720 control points on the FFD volume (shown in Fig. 1) and the angle-of-attack. The control points at the trailing edge are constrained to avoid any movement of the trailing edge. Therefore, the twist about the trailing edge can be implicitly altered by the optimizer using the remaining degrees of freedom. The leading-edge control points at the wing root are also constrained to maintain a constant incidence for the root section. There are 750 thickness constraints imposed in a 25 chordwise and 30 spanwise grid covering the full span and from 1% to 99% local chord. The thickness is set to be greater than 25% of the baseline thickness at each location. Finally, the internal volume is constrained to be greater than or equal to the baseline volume. The complete optimization problem is described in Table 2.


^b<https://info.aiaa.org/tac/ASG/APATC/AeroDesignOpt-DG/default.aspx>, last accessed May 10, 2014



a) L0 mesh: 28.8M cells, 199.7 drag counts.



b) L1 mesh: 3.6M cells, 201.7 drag counts.



c) L2 mesh: 450 k cells, 211.1 drag counts.

Figure 3. O-grids of varying sizes were generated using a hyperbolic mesh generator. [11]

	Function/variable	Description	Quantity
minimize	C_D	Drag coefficient	
with respect to	α	Angle of attack	1
	z	FFD control point z -coordinates	720
		Total design variables	721
subject to	$C_L = 0.5$	Lift coefficient constraint	1
	$C_{M_y} \geq -0.17$	Moment coefficient constraint	1
	$t \geq 0.25t_{\text{base}}$	Minimum thickness constraints	750
	$V \geq V_{\text{base}}$	Minimum volume constraint	1
	$\Delta z_{\text{TE,upper}} = -\Delta z_{\text{TE,lower}}$	Fixed trailing edge constraints	15
	$\Delta z_{\text{LE,upper,root}} = -\Delta z_{\text{LE,lower,root}}$	Fixed wing root incidence constraint	1
	Total constraints	769	

Table 2. Aerodynamic shape optimization problem




Figure 4. Sensitivity study of the baseline wing shows which shape changes yield the largest improvements. [11]

D. Surface Sensitivity on the Baseline Geometry

To examine the potential improvements of the baseline geometry, we performed a sensitivity analysis in [11]. The sensitivity of the drag and pitching moment with respect to the airfoil shape is shown in Fig. 4 as a contour plot of the derivatives of C_D and C_{M_y} with respect to shape variations in the z direction. The regions with the highest gradient of C_D are near the shock on the upper surface and near the wing leading edge. This indicates that leading-edge shaping and shock reduction through local shape changes should be the major drivers in C_D reduction at the beginning of the optimization. As for C_{M_y} , the shape changes near the root and tip of the wing are the most effective in adjusting the pitching moment. Since these sensitivity plots are a linearization about the current design point, they provide no information about the constraints. Nonetheless, these sensitivity plots indicate what drives the design at this design point.

IV. Aerodynamic Shape Optimization Results

Three different aerodynamic shape optimization approaches are presented: direct optimization, optimization with Richardson's extrapolation, and optimization with a multilevel approach. The optimized designs and the efficiency are compared in this section.

A. Direct Aerodynamic Shape Optimization

We present the direct aerodynamic design optimization for the CRM wing benchmark problem (described in Table 2) under the nominal flight condition (Mach 0.85, $Re = 5 \times 10^6$). We use the L1 grid (3.6 M cells) directly for the optimization. This is currently the most common way to perform aerodynamic shape optimization due to its simplicity. The optimization is computed with 64 processors. Figure 5 shows a detailed comparison of the baseline wing and the optimized wing using direct optimization.

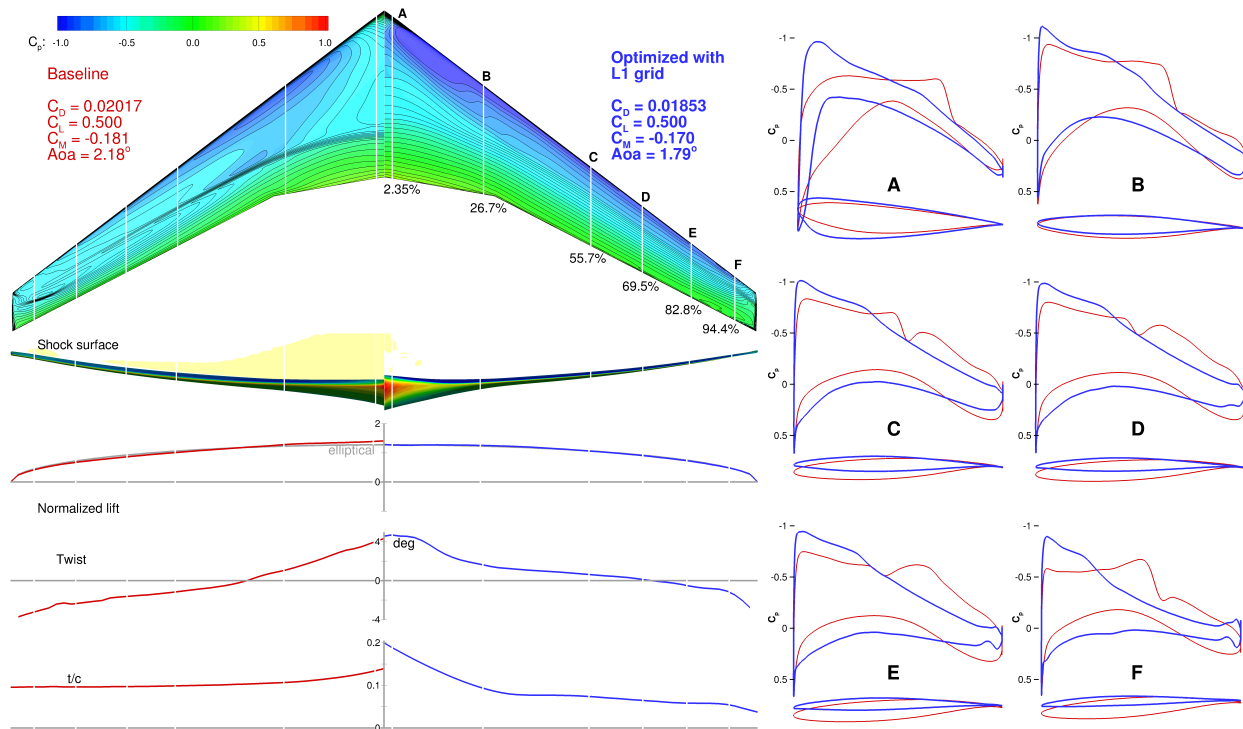


Figure 5. The optimized wing is shock-free and has 8.1% lower drag.

In this figure, the baseline wing results are shown in red and the optimized wing results are shown in blue. At the optimum, the lift coefficient target is met, and the pitching moment is reduced to the lowest allowed value. The lift distribution of the optimized wing is much closer to the elliptical distribution than that of the baseline, indicating induced drag is close to the theoretical minimum. This is achieved by fine-tuning the twist distribution and airfoil shapes. The baseline wing has a near-linear twist distribution. The optimized design has more twist at the root and tip, and less twist near mid-wing. The optimized thickness distribution is significantly different from that of the baseline, since the thicknesses are allowed to decrease to 25% of the original thickness, and there is a strong incentive to reduce the airfoil thicknesses in order to reduce wave drag. The volume is constrained to be greater than or equal to the baseline volume, so the optimizer drastically decreases the thickness of the airfoils on the outboard of the wing to the lower bounds, where there is less volume to be gained, while increasing the thickness near the root (up to 20%), where the chords are larger and the volume-drag trade-off is more favorable. The low outboard thickness would in practice incur a large structural weight penalty, and to trade off the reduction in drag and increase in weight would require aerostructural optimization [21].

The baseline wing exhibits a front of closely spaced pressure contour lines spanning a significant portion of the wing, indicating a shock. The optimized wing shows parallel pressure contour lines with uniform spacing, indicating a shock-free solution under the nominal flight condition. This is confirmed by the shock surface plots: we can see that the baseline wing has a shock on the upper surface, while the optimized wing does not show shocks under the design condition. The shock elimination can also be seen on the airfoil C_p distributions. The sharp increase in local pressure due to the shock becomes a gradual change from the leading edge to the trailing edge.

This optimization uses a relatively large grid size (3.6M). It took a significant amount of computational cost: it converges after 914 optimization iterations, but it takes 616.5 hours (26 days) on 64 processors. Therefore, the computational cost is prohibitive for large grid size. The following two sections presented two new approaches in an effort to reduce the computational cost.

B. Aerodynamic Shape Optimization Using the Richardson's Extrapolation

Richardson's extrapolation is commonly used to determine the errors in the spatial discretization in CFD [35]. The idea is to use Richardson's extrapolation to obtain the solution values at zero-grid spacing from a series of grids with coarser discretization. We first determine the order of accuracy using three solutions with a constant grid refinement ratio r , as follows:

$$p = \ln\left(\frac{f_3 - f_2}{f_2 - f_1}\right) / \ln(r) \quad (1)$$

For a second-order CFD solver, p should be close to 2 when the grids are well resolved. We can then compute the zero-grid spacing values using,

$$f_{h=0} = f_1 + \frac{f_1 - f_2}{r^p - 1} \quad (2)$$

We attempted to extend this idea to aerodynamic shape optimization. Both flow solutions and the adjoint sensitivities are evaluated on the two coarser grids. We then use Richardson's extrapolation to compute the zero-grid spacing objective function, constraints, and gradient values. Since the computational cost of the two coarser grids combined is still much lower than that of the fine grid, the total computational cost should be reduced. However, when we implemented this approach with the CRM wing optimization problem, the optimization was terminated due to numerical difficulties with only 5 iterations. The drag coefficient only had negligible changes. We found that both objective and constraints from the flow solution follows the Richardson's extrapolation's assumption: the solution is globally second-order, in addition to being locally second-order, and the solution functionals were computed using consistent second-order methods [35]. However, the gradients do not have a consistent order.

To further examine the gradients on different grid levels, we plot the shape gradients for each grid level, as shown in Figure 6.

The gradient can be affected by local phenomena such as shocks and separation, which causes the inconsistency between the grid levels. Thus, the zero-grid spacing gradients cannot be computed with Richardson's extrapolation, and the optimization cannot improve the design with incorrect gradients. We notice that the gradients between each grid are close.




Figure 6. The grid convergence for the shape gradients at different locations on the wing.

As an alternative, we perform the same direct optimization as Section A but with gradients computed on coarser L2 mesh in an effect to reduce the computational cost of the gradients. The optimization is able to converge further with the coarse gradients. It terminated with 54 iterations due to numerical difficulties. Once it terminated, we continue the optimization with the fine L1 grid gradients. Figure 7 show the optimized results.

We observed a similar optimal shape as that in Section A. The optimized drag is higher by 0.5 counts. Both lift and moment constraints are met at the optimum. The shock is also eliminated, as shown on the airfoil C_p distributions. The computational cost is reduced when compared to the direct L1 optimization. The total time is 441.8 hours (18.4 days), resulting in a 28% reduction.

C. Aerodynamic Shape Optimization Using Multilevel Technique

In this section, we present an acceleration technique inspired by the multigrid startup procedure in CFD that reduced the overall computational cost of the optimization effectively. Since we have improved the efficiency of our flow and adjoint solvers significantly over the last few years [25, 6, 20], we seek new methods to further reduce the computational cost of the aerodynamic shape optimization.

Since it is less costly to compute both the flow solution and the gradient on a coarser grid, we perform the optimization first on the coarsest grid until a certain level of optimality is achieved. Then, we move to the next grid level and start with the optimal design variables from the coarser grid. Since the drag and lift coefficients are generally different for each grid level, the approximate Hessian (used by the gradient-based optimizer) must be restarted. We

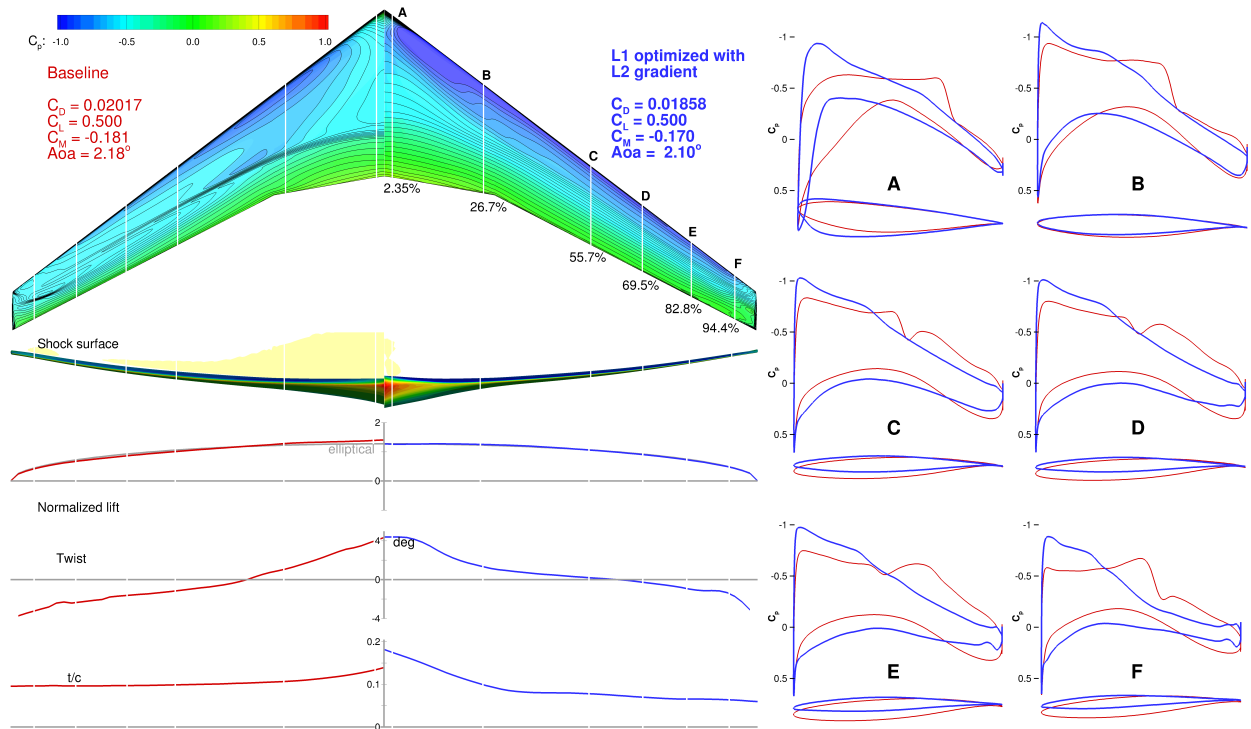


Figure 7. The optimization starts with the coarser L2 gradients.

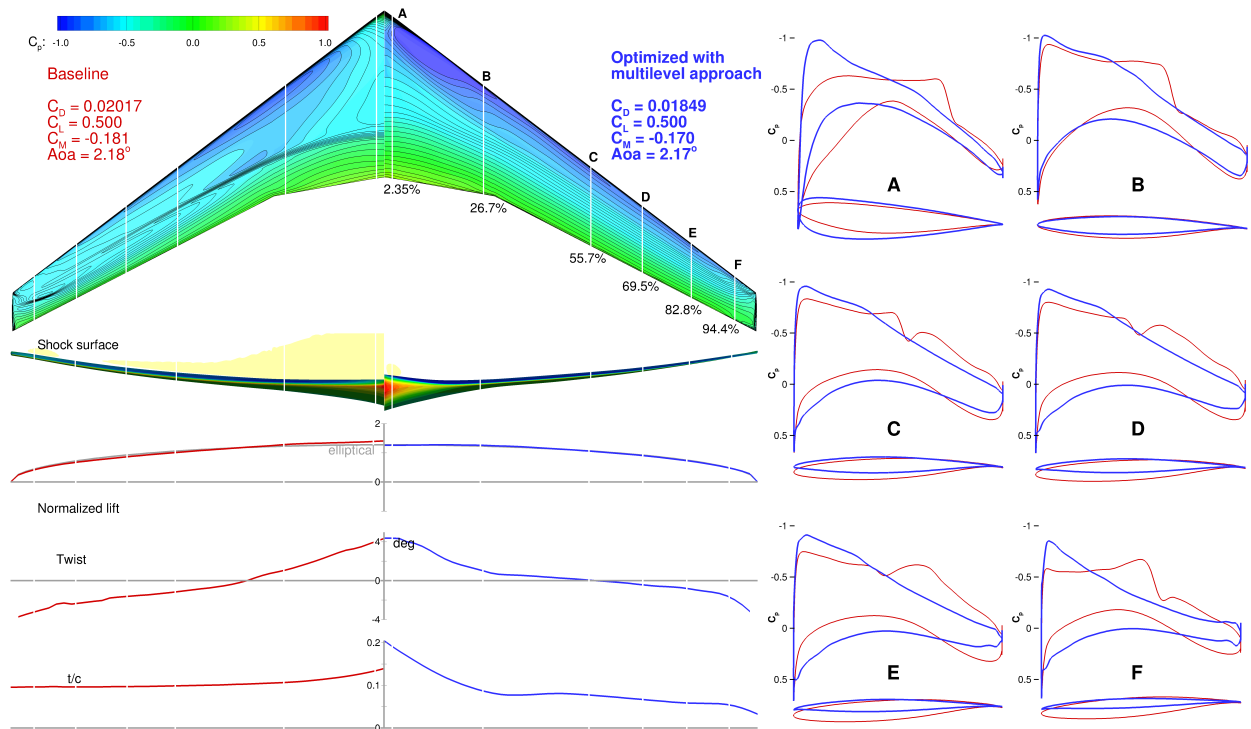


Figure 8. Multilevel optimization with L2 and L1 grids significant reduced computational cost.

repeat this process until the optimization on the finest grid has converged. Note that this procedure is different from traditional multigrid methods, where the coarse levels are revisited via multigrid cycles.

We performed the same CRM wing optimization problem using the approach described above. Figure 8 shows the comparison of the baseline and the optimized results. The optimized design is very similar to that in the previous two sections. The difference between the drag coefficients is within one count. There are visible differences in the airfoils C_p distributions, as shown in Figure 9. This might be caused by local minima that are close to each other, as previously observed by Lyu *et al.* [16, 11]. The multilevel approach has also been successfully applied in [11].

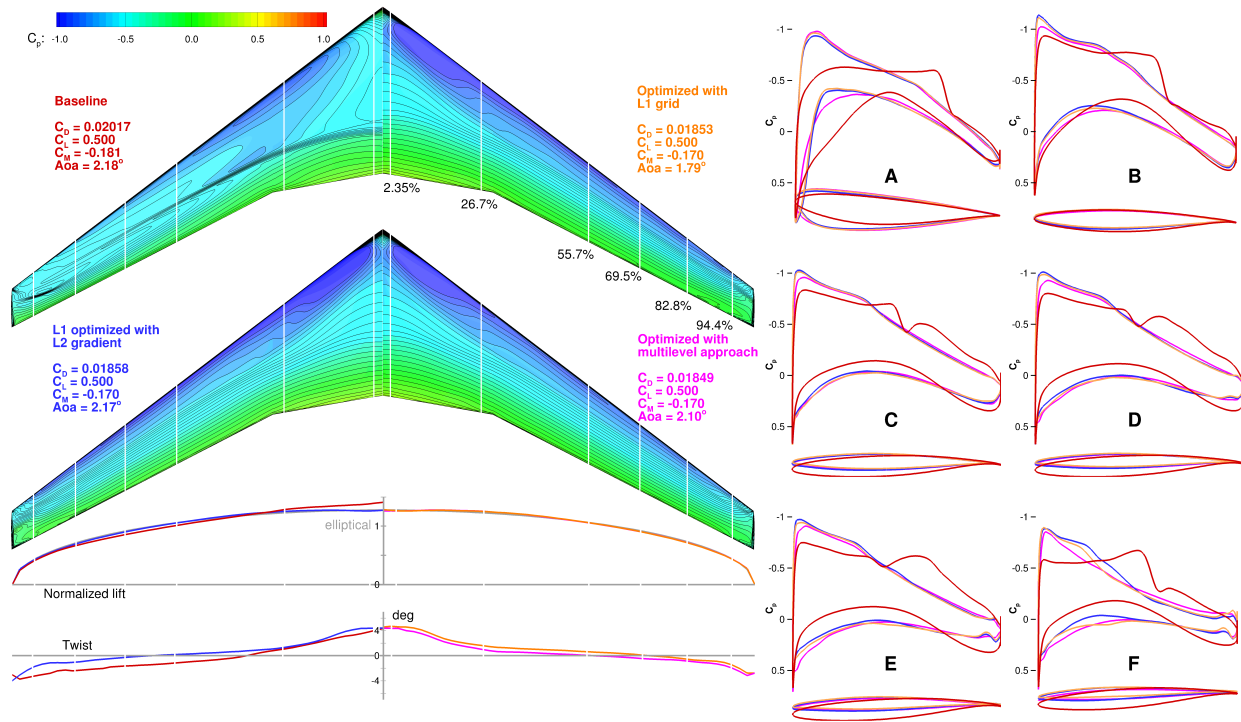


Figure 9. The difference in drag coefficients is within 1 count.

The multilevel approach uses significantly less computational resources. We use two grid levels: L2 (451 k cells) and L1 (3.6M cells). We can see that most of the optimization iterations are performed on the coarse grid, and as a result, the number of the function and gradient evaluations on the successively finer grids is greatly reduced. The optimization has 638 iterations on the L2 grid. Thanks to the optimization with the coarser grid, only 89 iterations are needed on the L1 grid to converge the optimization. The optimization with multilevel approach converges in 95.4 hours (4 days) on 64 processors. Table 3 shows the comparison of the optimization time. The merit function, optimality, and feasibility histories are plotted in Figure 10; detailed definitions of these values can be found in the SNOPT manual [36].

	Iterations	Time (hr)	Reduction
L1 direct	914	616.5	
L1 with L2 gradients	645	441.8	-28.3%
L1 multilevel	727	95.4	-84.5%

Table 3. The optimization with the multilevel approach reduces computational time by 84.5%.

We can see that the multilevel approach achieves the lowest computational cost by a large margin. With the increase

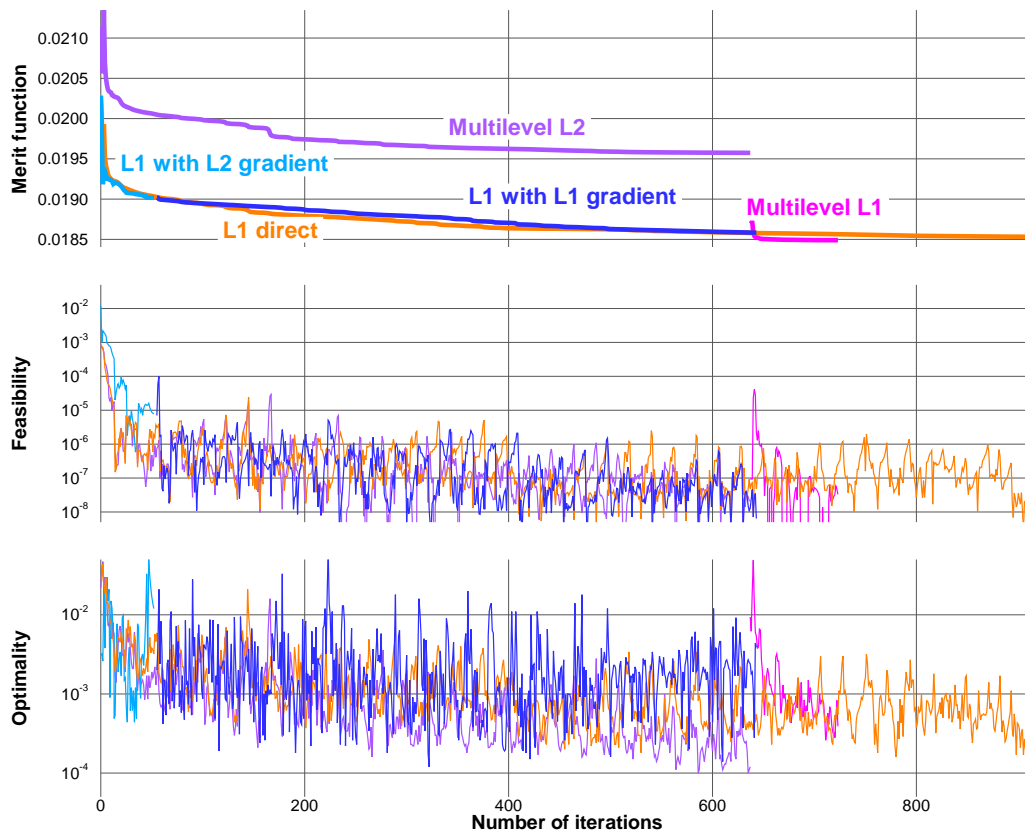


Figure 10. The optimization history of the three approaches.

of the grid size, this benefit becomes even larger. In this approach, we only performed one-way grid sequencing. Additional benefits may be achieved with a true multigrid V- or W-cycle at the optimization level.

V. Conclusions

In this paper, we have presented a study of the CRM wing shape optimization problem using three different optimization approach. The optimization problem is defined by the Aerodynamic Design Optimization Discussion Group (ADODG). The drag coefficient is minimized for one flight condition with respect to 720 shape design variables, subject to lift, pitching moment, and geometric constraints. We compared direct optimization with optimization with Richardson's extrapolation, and with optimization using multilevel approach. We found that the multilevel approach achieved the lowest computational cost. The total computation time was reduced from 616.5 hours to 95.4 hours using two grid levels.

The strategies presented in this paper open a new door to aerodynamic shape optimization. Further development of the techniques at the optimization level, in conjunction with MDO architectures have the potential to make future large-scale optimization more efficient and effective.

VI. Acknowledgments

This research is partially funded by FlexSys, Inc. The computations were performed on the Flux HPC cluster at the University of Michigan Center of Advanced Computing, and on the Gordon cluster of the Extreme Science

and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575. The authors would like to thank Yin Yu for his assistance in generating figures, as well as our colleagues at MDOLab for numerous ideas.

References

- [1] Hackbusch, W., *Multi-Grid Methods and Applications*, Vol. 4, Springer-Verlag Berlin, 1985.
- [2] Turkel, E., “Review of Preconditioning Methods for Fluid Dynamics,” *Applied Numerical Mathematics*, Vol. 12, No. 1–3, 1993, pp. 257 – 284. doi:[10.1016/0168-9274\(93\)90122-8](https://doi.org/10.1016/0168-9274(93)90122-8).
- [3] Gropp, W., Keyes, D., McInnes, L. C., and Tidriri, M., “Globalized Newton-Krylov-Schwarz Algorithms and Software for Parallel Implicit CFD,” *International Journal of High Performance Computing Applications*, Vol. 14, No. 2, 2000, pp. 102–136. doi:[10.1177/109434200001400202](https://doi.org/10.1177/109434200001400202).
- [4] Knoll, D. and Keyes, D., “Jacobian-free Newton–Krylov Methods: a Survey of Approaches and Applications,” *Journal of Computational Physics*, Vol. 193, No. 2, 2004, pp. 357 – 397. doi:[0.1016/j.jcp.2003.08.010](https://doi.org/0.1016/j.jcp.2003.08.010).
- [5] Jameson, A., “Aerodynamic Design via Control Theory,” Vol. 3, No. 3, 1988, pp. 233–260. doi:[10.1007/BF01061285](https://doi.org/10.1007/BF01061285).
- [6] Lyu, Z., Kenway, G., Paige, C., and Martins, J. R. R. A., “Automatic Differentiation Adjoint of the Reynolds-Averaged Navier-Stokes Equations with a Turbulence Model,” *43rd AIAA Fluid Dynamics Conference and Exhibit*, June 2013.
- [7] Anderson, W. K. and Bonhaus, D. L., “Airfoil Design on Unstructured Grids for Turbulent Flows,” *AIAA journal*, Vol. 37, No. 2, 1999, pp. 185–191. doi:[10.2514/2.712](https://doi.org/10.2514/2.712).
- [8] Dwight, R. P. and Brezillon, J., “Efficient and Robust Algorithms for Solution of the Adjoint Compressible Navier–Stokes Equations with Applications,” *International Journal for Numerical Methods in Fluids*, Vol. 60, 2009, pp. 365–389. doi:[10.1002/flid.1894](https://doi.org/10.1002/flid.1894).
- [9] Brezillon, J. and Dwight, R. P., “Applications of a Discrete Viscous Adjoint Method for Aerodynamic Shape Optimisation of 3D Configurations,” *CEAS Aeronautical Journal*, Vol. 3, No. 1, 2012, pp. 25–34. doi:[10.1007/s13272-011-0038-0](https://doi.org/10.1007/s13272-011-0038-0).
- [10] Lyu, Z. and Martins, J. R. R. A., “Aerodynamic Design Optimization Studies of a Blended-Wing-Body Aircraft,” *Journal of Aircraft*, 2014. doi:[10.2514/1.C032491](https://doi.org/10.2514/1.C032491).
- [11] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., “Aerodynamic Shape Optimization Investigations of the Common Research Model Wing Benchmark,” *AIAA Journal*, 2014, (Submitted).
- [12] Lyu, Z. and Martins, J. R. R. A., “Aerodynamic Shape Optimization of an Adaptive Morphing Trailing Edge Wing,” *AIAA Aviation and Aeronautics Forum and Exposition (AVIATION 2014)*, Atlanta, GA, June 2014.
- [13] Vassberg, J., “Introduction: Drag Prediction Workshop,” *Journal of Aircraft*, Vol. 45, No. 3, Jun 2008, pp. 737–737. doi:[10.2514/1.37761](https://doi.org/10.2514/1.37761).
- [14] Vassberg, J., Dehaan, M., Rivers, M., and Wahls, R., “Development of a Common Research Model for Applied CFD Validation Studies,” *26th AIAA Applied Aerodynamics Conference*, American Institute of Aeronautics and Astronautics, August 2008. doi:[10.2514/6.2008-6919](https://doi.org/10.2514/6.2008-6919).
- [15] Vassberg, J., “A Unified Baseline Grid about the Common Research Model Wing/Body for the Fifth AIAA CFD Drag Prediction Workshop (Invited),” *29th AIAA Applied Aerodynamics Conference*, Jul 2011. doi:[10.2514/6.2011-3508](https://doi.org/10.2514/6.2011-3508).
- [16] Lyu, Z., Kenway, G. K. W., and Martins, J. R. R. A., “RANS-based Aerodynamic Shape Optimization Investigations of the Common Research Model Wing,” *AIAA Science and Technology Forum and Exposition (SciTech)*, National Harbor, MD, January 2014.
- [17] Vassberg, J. and Jameson, A., “Influence of Shape Parameterization on Aerodynamic Shape Optimization,” Tech. rep., Von Karman Institute, Brussels, Belgium, April 2014.
- [18] Telidetzki, K., Osusky, L., and Zingg, D. W., “Application of Jetstream to a Suite of Aerodynamic Shape Optimization Problems,” *52nd Aerospace Sciences Meeting*, Feb 2014. doi:[10.2514/6.2014-0571](https://doi.org/10.2514/6.2014-0571).
- [19] Carrier, G., Destarac, D., Dumont, A., Meheut, M., Din, I. S. E., Peter, J., Khelil, S. B., Brezillon, J., and Pestana, M., “Gradient-Based Aerodynamic Optimization with the elsA Software,” *52nd Aerospace Sciences Meeting*, Feb 2014. doi:[10.2514/6.2014-0568](https://doi.org/10.2514/6.2014-0568).
- [20] Kenway, G. K. W., Kennedy, G. J., and Martins, J. R. R. A., “Scalable Parallel Approach for High-Fidelity Steady-State Aeroelastic Analysis and Adjoint Derivative Computations,” *AIAA Journal*, Vol. 52, No. 5, 2014, pp. 935–951. doi:[10.2514/1.J052255](https://doi.org/10.2514/1.J052255).

- [21] Kenway, G. K. W. and Martins, J. R. R. A., "Multipoint High-fidelity Aerostructural Optimization of a Transport Aircraft Configuration," *Journal of Aircraft*, Vol. 51, No. 1, 2014, pp. 144–160. doi:10.2514/1.C032150.
- [22] Kenway, G. K., Kennedy, G. J., and Martins, J. R. R. A., "A CAD-free Approach to High-Fidelity Aerostructural Optimization," *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference, Fort Worth, TX*, 2010. doi:10.2514/6.2010-9231.
- [23] van der Weide, E., Kalitzin, G., Schluter, J., and Alonso, J., "Unsteady Turbomachinery Computations Using Massively Parallel Platforms," *44th AIAA Aerospace Sciences Meeting and Exhibit*, 2006. doi:10.2514/6.2006-421.
- [24] Jameson, A., Schmidt, W., and Turkel, E., "Numerical Solution of the Euler equations by Finite Volume Methods Using Runge Kutta Time Stepping Schemes," *14th AIAA, Fluid and Plasma Dynamics Conference*, 1981.
- [25] Mader, C. A., Martins, J. R. R. A., Alonso, J. J., and van der Weide, E., "ADjoint: An Approach for the Rapid Development of Discrete Adjoint Solvers," *AIAA Journal*, Vol. 46, No. 4, April 2008, pp. 863–873. doi:10.2514/1.29123.
- [26] Saad, Y. and Schultz, M. H., "GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems," *SIAM Journal on Scientific and Statistical Computing*, Vol. 7, No. 3, 1986, pp. 856–869. doi:10.1137/0907058.
- [27] Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F., "Efficient Management of Parallelism in Object Oriented Numerical Software Libraries," *Modern Software Tools in Scientific Computing*, edited by E. Arge, A. M. Bruaset, and H. P. Langtangen, Birkhäuser Press, 1997, pp. 163–202. doi:10.1007/978-1-4612-1986-6_8.
- [28] Balay, S., Brown, J., Buschelman, K., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H., "PETSc Users Manual," Tech. Rep. ANL-95/11 - Revision 3.4, Argonne National Laboratory, 2013.
- [29] Balay, S., Brown, J., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H., "PETSc Web page," 2013, <http://www.mcs.anl.gov/petsc>.
- [30] Mader, C. A. and Martins, J. R. R. A., "Stability-Constrained Aerodynamic Shape Optimization of Flying Wings," *Journal of Aircraft*, Vol. 50, No. 5, September 2013, pp. 1431–1449. doi:10.2514/1.C031956.
- [31] Lyu, Z. and Martins, J. R. R. A., "Aerodynamic Shape Optimization of a Blended-Wing-Body Aircraft," *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2013. doi:10.2514/6.2013-283.
- [32] Liem, R. P., Kenway, G. K. W., and Martins, J. R. R. A., "Multi-mission Aircraft Fuel Burn Minimization via Multi-point Aerostructural Optimization," *AIAA Journal*, 2014, (Submitted).
- [33] Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM journal on optimization*, Vol. 12, No. 4, 2002, pp. 979–1006. doi:10.1137/S1052623499350013.
- [34] Perez, R. E., Jansen, P. W., and Martins, J. R. R. A., "pyOpt: A Python-Based Object-Oriented Framework for Non-linear Constrained Optimization," *Structures and Multidisciplinary Optimization*, Vol. 45, No. 1, 2012, pp. 101–118. doi:10.1007/s00158-011-0666-3.
- [35] Roache, P. J., "Verification of Codes and Calculations," *AIAA Journal*, Vol. 36, No. 5, 1998, pp. 696–702. doi:10.2514/2.457.
- [36] Gill, P. E., Murray, W., and Saunders, M. A., *User's Guide for SNOPT Version 7: Software for Large-Scale Nonlinear Programming*, 2006.