

ERDC MP-21-4

Engineering Research
and Development Center



**US Army Corps
of Engineers®**
Engineer Research and
Development Center



Helicopter Rotor Blade Multiple-Section Optimization with Performance Considerations

Luke D. Allen, Joon W. Lim, Robert B. Haehnel,
and Ian D. Dettwiller

June 2021

The U.S. Army Engineer Research and Development Center (ERDC) solves the nation's toughest engineering and environmental challenges. ERDC develops innovative solutions in civil and military engineering, geospatial sciences, water resources, and environmental sciences for the Army, the Department of Defense, civilian agencies, and our nation's public good. Find out more at www.erdclibrary.on.worldcat.org/discovery.

To search for other technical reports published by ERDC, visit the ERDC online library at <https://erdclibrary.on.worldcat.org/discovery>.

Helicopter Rotor Blade Multiple-Section Optimization with Performance Considerations

Luke D. Allen and Robert B. Haehnel

*Cold Regions Research and Engineering Laboratory
U.S. Army Engineer Research and Development Center
72 Lyme Road
Hanover, NH 03775*

Joon W. Lim

*U.S. Army Aviation and Missile Research
Development and Engineering Center
Moffett Field, CA 94035*

Ian D. Detwiller

*Information Technology Laboratory
U.S. Army Engineer Research and Development Center
3090 Halls Ferry Road
Vicksburg, MS 39180*

Final report

Approved for public release; distribution is unlimited.

Prepared for U.S. Army Corps of Engineers
Washington, DC 20134

Under Program Element Number 0603465, Project AL3

Preface

This study was conducted for the Deputy Assistant Secretary of the Army for Research and Technology and supported by the Future Vertical Lift Cross-Functional Team in collaboration with the U.S. Army Combat Capabilities Development Command (CCDC) Aviation and Missile Center (AvMC), under the High-Performance Computing for Rotorcraft Applications Program, Air Systems Rotorcraft Work Unit, Program Element 0603465, Project AL3. The technical monitor was Dr. Robert M. Wallace.

The work was performed by the Engineering Resources Branch (Caitlin Callaghan, Branch Chief) and the Terrestrial and Cryospheric Science Branch (John Weatherly, Branch Chief) of the Research and Engineering Division, (George Calfas, Division Chief), U.S. Army Engineer Research and Development Center, Cold Regions Research and Engineering Laboratory (ERDC-CRREL) and the Computational Analysis Branch (Jeffery Hensley, Branch Chief) of the Computational Science and Engineering Division (Devin Sham, Acting Division Chief), U.S. Army Engineer Research and Development Center, Information Technology Laboratory (ERDC-ITL). At the time of publication, the Deputy Director of ERDC-CRREL was Mr. David Ringelberg and the Director was Dr. Joseph Corriveau, the ITL Deputy Director was Ms. Patti S. Duett and the Director was Dr. David A. Horner.

This paper was originally presented at the *Vertical Flight Society's 77th Annual Forum & Technology Display* (Virtual), 10–14 May 2021.

The Commander of ERDC was COL Teresa A. Schlosser and the Director was Dr. David W. Pittman.

DISCLAIMER: The contents of this report are not to be used for advertising, publication, or promotional purposes. Citation of trade names does not constitute an official endorsement or approval of the use of such commercial products. All product names and trademarks cited are the property of their respective owners. The findings of this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

DESTROY THIS REPORT WHEN NO LONGER NEEDED. DO NOT RETURN IT TO THE ORIGINATOR.

Helicopter Rotor Blade Multiple-Section Optimization with Performance Considerations

ABSTRACT

This paper presents advancements in a surrogate-based, rotor blade design optimization framework for improved helicopter performance. The framework builds on previous successes by allowing multiple airfoil sections to be designed simultaneously to minimize required rotor power in multiple flight conditions. Rotor power in hover and forward flight, at advance ratio $\mu = 0.3$, are used as objective functions in a multi-objective genetic algorithm. The framework is constructed using Galaxy Simulation Builder with optimization provided through integration with Dakota. Three independent airfoil sections are morphed using ParFoil and aerodynamic coefficients for the updated airfoil shapes (i.e., lift, drag, moment) are calculated using linear interpolation from a database generated using C81Gen/ARC2D. Final rotor performance is then calculated using RCAS. Several demonstrative optimization case studies were conducted using the UH-60A main rotor. The degrees of freedom for this case are limited to the airfoil camber, camber crest position, thickness, and thickness crest position for each of the sections. The results of the three-segment case study show improvements in rotor power of 4.3% and 0.8% in forward flight and hover, respectively. This configuration also yields greater reductions in rotor power for high advance ratios, e.g., 6.0% reduction at $\mu = 0.35$, and 8.8% reduction at $\mu = 0.4$.

NOTATION

C_p	Rotor power coefficient
c	Airfoil chord length
D	Feasibility metric (L ₂ norm)
f_m	Camber scaling factor
f_t	Thickness scaling factor
m	Camber in chords
N_{cfd}	Number of CFD evaluations
N_{opt}	Number of optimizer evaluation
PLL	Pitch link load, lb
PLL_{hptp}	Half-peak-to-peak pitch link load, lb
p	Camber crest position in chords
R	Rotor radius, ft
t	Airfoil thickness in chords
S	Estimated speedup factor
t_{cfd}	CFD evaluation time
t_{opt}	Optimizer iteration evaluation time
X	Feasible design space, $X \subset \mathcal{R}$
x	Airfoil thickness crest position in chords

ϑ	Design parameter set/vector
μ	Advance ratio
σ	Rotor solidity

INTRODUCTION

Throughout their history, rotorcraft have demonstrated incredible versatility for civilian and military applications. Their ability to take off and land practically anywhere provides unique advantages over other forms of aircraft. These advantages, however, carry with them many complexities that make designing components with optimal characteristics very difficult. Many of the difficulties can be attributed to the wide variety of conditions through which the rotors are operated throughout the flight envelope. Recently, many researchers have attempted to improve rotorcraft performance through optimization of the main rotor.

Imiela and Wilke (Ref. 1) analyzed the effects of blade twist and anhedral on hover and forward flight performance for the

EC1 and EC2 rotors. Their experiments were a series of single objective optimizations using the Dakota optimization library (Ref. 2) with METAR/Euler flow solvers (Ref. 3) to either maximize the Figure of Merit (FM) or minimize the power requirements of the rotor. They noted that the competing objectives in a combined optimization made improvement in both conditions difficult, with the design space offering little room for increased performance, e.g., 2.2% power reduction at 240 km/h ($\mu = 0.334$), but with a degradation in FM of 0.5% at $C_T/\sigma = 0.093$ (hover was not considered in objective).

Leon et al. (Ref. 4) used a CONMIN/elsA/HOST computational workflow (Refs. 5–7) to carry out a multiobjective optimization of the ERATO rotor (Ref. 8). They coupled the Dakota optimizer with a Nash Game algorithm (Ref. 9) which divided the design space between two non-cooperative “players,” where each attempted to improve its objective. They used blade twist, sweep, and chord as design parameters with models constructed using Bezier curves and cubic splines. With a constraint on the maximum pitch link load imposed, the forward flight rotor power was reduced by as much as 1% with respect to the baseline, while FM was also improved by 2 counts (about 3%) at $C_T/\sigma = 0.0625$.

Wang et al. (Ref. 10) used the sequential quadratic-programming optimizer SNOPT (Ref. 11) to perform constrained multiobjective optimization of the UH-60A main rotor. Deformations of blade twist, airfoil thickness and camber were done using MASSOUD (Ref. 12) and 81 design variables located at intervals along the blade span. Their technique used a linear combination of FM and rotor power for a forward flight condition ($\mu = 0.368$) as the objective function, calculated using FUN3D and DYMORE. With this strategy, they were able to show a reduction in rotor power of 3.91% at $\mu = 0.368$, while improving FM by 1.03%.

Lim (Ref. 13) adopted an ad-hoc approach to improving the UH-60A rotor power. CAMRAD II was used to compare the baseline rotor power using the existing airfoils (SC1094R8 inboard and SC1095 outboard) with advanced airfoils (SC2110 inboard and SSC-A09 outboard). While formal optimization was not used in this study, he found that the ad-hoc approach reduced the rotor power by 7.3% at $\mu = 0.3$ and 12.1% at $\mu = 0.4$. These results suggest that greater improvements in rotor power ought to be possible by employing rigorous optimization methodologies.

Allen et al. (Ref. 14) introduced an automated framework for multiobjective blade optimization that combined the airfoil parameterization code ParFoil (Ref. 13) with the Dakota numerical optimization library. The framework was based on a surrogate modeling approach to generate airfoil performance tables (C81 tables) that substantially reduced the required computational resources. This work separately employed both the CONMIN algorithm and a multiobjective genetic algorithm (Ref. 15) to minimize the power

requirements during hover and forward flight ($\mu = 0.3$). The selected parameters were airfoil thickness and thickness crest position at the blade tip (84–100% of the rotor radius, R). Multiobjective CONMIN was performed using the weighted sum of the individual responses and investigating variations in function weights. Conversely, the multiobjective genetic algorithm (MOGA) is used to populate the Pareto front and then the best designs are selected by the user upon completion. Results from the gradient-based and genetic algorithms were compared, showing that comparable reductions in rotor power are feasible for either method, provided a suitable starting condition for CONMIN. This multiobjective framework reduced the required power by 3.5% in forward flight and 0.4% in hover.

Multi-objective optimization maintains an ever-increasing role in the design of rotorcraft and their components (Refs. 16–18). Furthermore, multi-objective optimization itself is a rapidly advancing field with many established algorithms as well as an ever-expanding list of newly proposed alternatives (Ref. 19). For this work, MOGA was selected over NSGA-II (Ref. 20) and other established approaches (Refs. 21–24) for its robustness and convenient implementation in Dakota.

The presented work builds on the previous results in Ref. 14 by extending the framework to allow for multiple blade sections to morph independently. In addition to airfoil thickness and crest position, the current work also considers airfoil camber and camber crest position. The objectives of this work are 1) to demonstrate the feasibility of the MOGA approach in the Galaxy environment for the application to rotor performance design, 2) to improve upon previous results as measured by further reduction of rotor power, and 3) demonstrate the abilities of the rotor blade optimization framework with multiple airfoil sections.

METHODOLOGY

Galaxy Simulation Builder

Galaxy Simulation Builder (GSB) is a general-purpose tool for controlling the environment and execution of programs on a target host system. It has native integration with Dakota and provides tools for executing programs remotely on high-performance computing (HPC) systems. GSB differentiates itself by its robust integration with the Department of Defense Supercomputing Resource Center (DSRC) network of supercomputers. GSB streamlines the connection and data transfer process to these machines and manages environment setup, job submission, monitoring, and report generation.

Analysis capabilities in GSB are largely provided through its integration with Dakota. GSB provides customized GUI support for setting up Dakota simulations as part of the integration between the two tools. The combination of Galaxy’s parallel job monitoring, management, and execution with Dakota’s extensive analysis capabilities provides a powerful framework for multi-fidelity and multidisciplinary analysis.

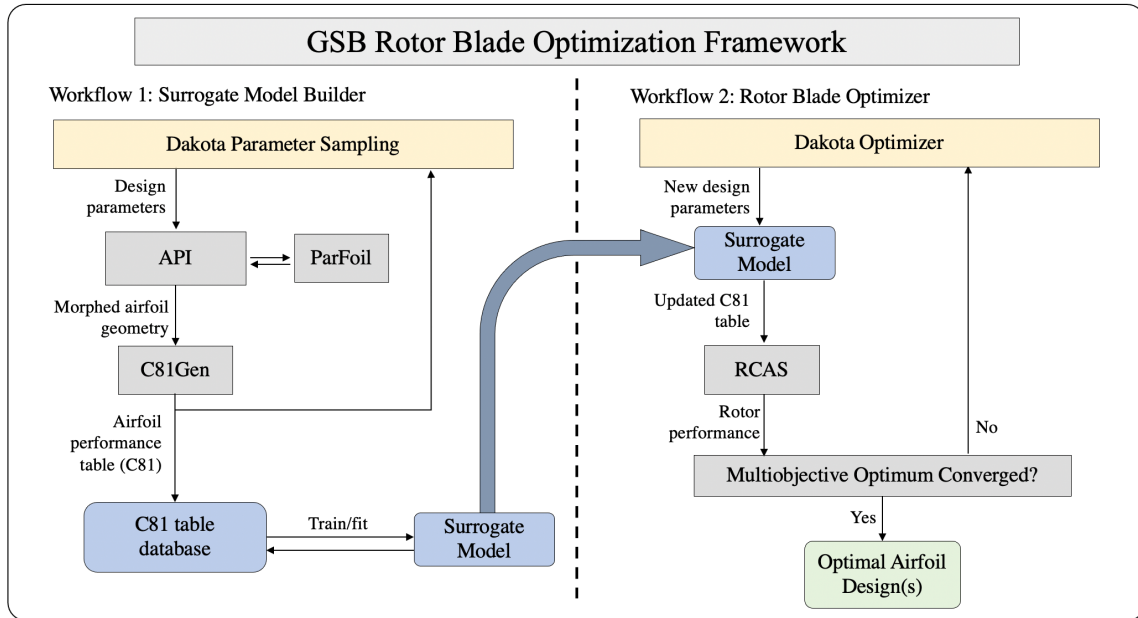


Figure 1. Flowchart for rotor blade optimization implemented in Galaxy Simulation Builder.

This work uses GSB to construct an automated framework for rotor blade optimization. The framework itself is relatively unchanged from its initial implementation in Ref. 14, though much of the underlying scripting has been updated to accommodate additional capabilities. The framework consists of two interdependent workflows: the C81 model builder workflow, and the multiobjective rotor blade optimization workflow. A simplified flowchart illustrating the framework is shown in Figure 1.

C81 Model Builder Workflow

The C81 model builder workflow is built around C81Gen (Refs. 25, 26), a wrapper for 2D airfoil mesh generation and NASA’s ARC2D Navier-Stokes CFD solver (Ref. 27). A user-defined sampling method (e.g., regular, random, Latin hypercube) is used to populate the design space, and then airfoil performance tables such as those used by rotorcraft comprehensive analysis codes (e.g., C81 tables) are generated for each design.

The choice of shape parameterization methodology is one of the most crucial components of optimization. A good parameterization accesses the entirety of the feasible design space using the smallest number of design parameters possible. This work uses the airfoil parameterization code ParFoil (Ref. 13) to accomplish this task. ParFoil starts with a baseline airfoil geometry and morphs it into a new design by augmenting the design parameters. The design parameters are leading-edge radius, leading-edge droop, camber (m), camber crest position (p), thickness (t), thickness crest position (x), trailing-edge camber, trailing-edge camber crest position, and boat-tail angle. Before parameterization, the baseline airfoil coordinates are redistributed using a nonuniform radial basis spline (NURBS). The baseline values for the design parameters are then extracted from the spline curve.

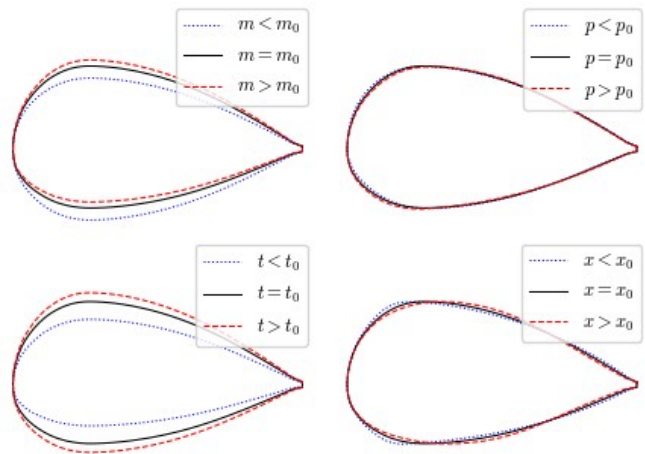


Figure 2. Effects of parameter changes in ParFoil on SC1095 airfoil geometry. Vertical scales are exaggerated.

Parameter augmentation is performed by applying either a scaling factor or delta value to the baseline and then updating the airfoil coordinates. The individual effects of varying the m , p , t , and x parameters (used in this study) are shown in Figure 2. Baseline values are denoted with zero subscripts. A summary of the parameter formulations and their feasible bounds for the SC1095 airfoil is given in Table 1.

Once all of the tables are complete, they are stored in a database. The database is then used to fit a model using interpolation or some form of regression (e.g., a neural network). The surrogate model is saved and can be downloaded for future use. The model fitting and storage tasks are handled by a custom Python API written for this work.

Table 1. ParFoil parameters and augmentation strategies. All values, except the scaling factors f_m and f_t , are given in fractions of the airfoil chord.

Symbol	Baseline ^a	Strategy	Feasible Bounds
m	0.0081	$m = m_0 f_m$	$0.0 \leq f_m \leq 1.5$
p	0.270	$p = p_0 + \Delta p$	$-0.0965 \leq \Delta p \leq 0.1035$
t	0.095	$t = t_0 f_t$	$0.5 \leq f_t \leq 1.3$
x	0.270	$x = x_0 + \Delta x$	$-0.1 \leq \Delta x \leq 0.25$

^aSC1095

Rotor Blade Optimization Workflow

The rotor blade optimization workflow begins according to the user-defined optimization algorithm selected through Dakota. Following parameter substitution, it then leverages the surrogate model described above to replace expensive CFD calculations within the optimization loop. The surrogate model is used to generate new airfoil tables which are inserted into the appropriate input files before running the RCAS rotorcraft comprehensive code (Refs. 28, 29). Details of the rotor model and flight conditions specific to this work are given in subsequent sections. After the RCAS simulations are complete, the desired performance metrics are extracted and returned to the Dakota optimizer. The simulation proceeds according to the chosen optimization method.

An advantage to using the Galaxy framework is the built-in tools for job monitoring. Real-time results are collected by Galaxy and made available via a convenient web interface. This can save the user substantial amounts of time versus manually navigating to the on-disk locations of the output files. In addition to the tabular view of input and output data, Galaxy also provides some basic analysis tools (e.g., statistical summaries, parallel line plots) that can be used without requiring any data transfer. Users can also use the web interface to locate specific results and output files for further investigation.

Rotor Model

The starting point for this work is the standard UH-60A main rotor. The rotor radius, R , is 26.833 ft, and the rotational speed is 27.03 rad/s. The standard rotor is made up of two airfoils: the SC1095 extending from the blade root to 0.48R, the SC1094R8 from 0.48R to 0.84R, and then returning to the SC1095 from 0.84R to the blade tip.

This work performs shape deformation on three distinct sections of the blade. The segments are: the blade tip segment (Tip: 0.92857R to 1.0R), the outboard-1 segment (OB1: 0.85404R to 0.91R) and the outboard-2 segment (OB2: 0.73575R to 0.83851R). Figure 3 shows sketches of the baseline and optimized blade configurations with airfoil distributions and segment labels.

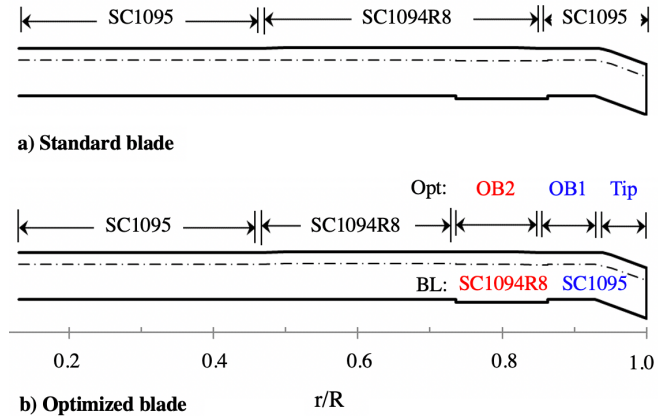


Figure 3. Sketches of a) standard UH-60A rotor blade, and b) the optimized rotor blade with airfoil distributions.

The UH-60A blade model is constructed using the rotorcraft comprehensive code RCAS. Lookup tables are used to define airfoil performance coefficients for each airfoil. For strictly unmodified blade segments (e.g., the initial span of SC1095 stretching from the root to 0.48R), experimentally obtained values are used for the tables. For all other sections, the baseline experimental tables are replaced with tables produced by C81Gen. The differences between experimentally-derived tables and those produced by CFD are different enough that failing to make this substitution results in artificially high reductions in rotor power. Amending the baseline airfoil tables ensures a fair comparison between them and the modified airfoil results.

Note that the segments are morphed independently. To more deeply understand the behavior of the optimizer and the underlying physics, this paper presents the results of three separate optimization scenarios. The first considers only the Tip segment, the second considers the Tip and OB1 segments, and the last considers modification of the Tip, OB1, and OB2 segments.

Flight Conditions

The flight conditions used for the UH-60A in this work are hover and steady forward flight. The forward flight condition corresponds to an advance ratio of $\mu = 0.3$. Trim targets for this and the hover case are based on test results from the USAF National Full-scale Aerodynamics Complex (NFAC) 40 ft by 80 ft wind tunnel at NASA Ames Research Center (Ref. 30). The rotor model is based on the experimental setup for the same test. The model uses an isolated rotor with a 4-DOF propulsive trim. The trim targets are thrust, propulsive force, rotor roll and pitch moments, and the trim variables are the blade collective, lateral and longitudinal cyclic controls, and the rotor shaft angle (i.e., nose-down pitching). The trim variables are given initial values and then updated until equilibrium is reached. The hover condition uses the same model with updated initial conditions suitable for a zero free stream velocity.




Figure 4. Plots of the PLL oscillation (upper) and PLL_{hptp} (lower) histories versus rotor revolution.

A third flight condition used is the Utility Tactical Transport Aerial System (UTTAS) pull-up maneuver. This condition is used to investigate the maximum design load on the pitch link. Flight conditions are based on test C11029 of the UH-60A Airloads Flight Test Program (Ref. 31). The rotorcraft enters the maneuver near its maximum forward flight speed and quickly pulls up, achieving load factors that greatly exceed the steady-state lift limit of the rotor. In the severe case of the C11029 test condition, the vehicle experiences a normal load factor of 2.1g. The maneuver lasts for approximately 40 rotor revolutions before returning to level flight. During the maneuver, the time history of the pitch link load (PLL) is recorded as shown in the upper plot of Figure 4. Additionally, the half-peak-to-peak pitch link load (PLL_{hptp}) is calculated, as shown in the lower plot in Figure 4.

The pitch link load is intended to constrain the optimization such that it may not exceed that of the baseline configuration. While implementation of the maneuvering flight condition into the GSB workflow is complete, validation is ongoing. Therefore, most of the results in this paper are given for unconstrained optimization. Further details are provided in the Constrained Optimization section of the results.

Airfoil Surrogate Model

The development of the surrogate model used for airfoil performance predictions is described in Ref. 14. At this time, the model is reduced to a linear interpolation scheme. Interpolation is easily performed on a regular grid or irregular grid using Delaunay triangulation. Implementation of more sophisticated models (e.g., neural networks) has been limited due to inaccuracies in the drag coefficient, c_d , near the zero angle of attack. Improvement of the modeling technique is an area of active development within the group.

Regardless of the modeling technique used, it must be built upon existing data. Initially, many airfoil tables must be generated using CFD simulations (via C81Gen). Each of these

tables takes an average of approximately 150 CPU hours to complete. By comparison, the surrogate model using linear interpolation on a regular grid can generate a new table in 0.5–0.8 seconds. Let the speedup, S , be defined as the ratio of CPU times for optimization performed with the surrogate model (including generation of the database) to that of optimization performed using only CFD. The ratio can be represented as

$$S = \frac{N_{opt}(t_{cfd} + t_{opt})}{N_{cfd}t_{cfd} + N_{opt}t_{opt}}, \quad (1)$$

where N_{opt} is the number of evaluations required by the optimizer, N_{cfd} is the number of tables generated via CFD for the surrogate model database, t_{opt} is the time for a single optimizer iteration, and t_{cfd} is the time to generate a table using CFD. The time for the surrogate model to generate a table is considered negligible.

A database was generated for the airfoil surrogate model with four design parameters: m , p , t , and x . More accurately, the input parameters for the database are the augmentation values: f_m , Δp , f_t , and Δx . Points were sampled on a regular grid with four partitions (five levels) per parameter. The resulting database contained 625 tables (i.e., $N_{cfd} = 5^4 = 625$). Note that the level spacing within each parameter was not uniform. Instead, the spacing was skewed such that 1) the baseline values were captured for each parameter, 2) the levels covered the full feasible design space, and 3) the desired number of levels is maintained. The grid levels for each of the active parameters are shown here:

$$\begin{aligned} f_m &= \{0.0, 0.5, 1.0, 1.25, 1.5\} \\ \Delta p &= \{-0.0965, -0.045, 0.0, 0.5, 0.1035\} \\ f_t &= \{0.5, 0.75, 1.0, 1.15, 1.3\} \\ \Delta x &= \{-0.1, 0.0, 0.08, 0.16, 0.25\}. \end{aligned}$$

The augmentation values above yield the absolute parameter ranges here (fractions of airfoil chord):

$$\begin{aligned} m &\in [0.0, 0.00121] \\ p &\in [0.1701, 0.3735] \\ t &\in [0.0475, 0.1235] \\ x &\in [0.167, 0.517]. \end{aligned}$$

The average CFD evaluation time (t_{cfd}) during the database generation was 2.6 hours. The total time needed to generate this database was approximately 7.15×10^4 CPU hours. Figure 5 shows the relationship between the two theoretical times as functions of N_{opt} on a log scale. The average optimization function evaluation time (t_{opt}) used was 0.167 hours. The figure shows that generating a database and interpolating, or fitting a model to it, is favorable whenever the anticipated N_{opt} exceeds the projected size of the database, N_{cfd} .

All of the discussion in this section assumes that the model database is populated specifically for a single optimization case. In reality, once a database is created, it can continue to be used indefinitely as long as the conditions under which it was made are still valid (e.g., same active design parameters,

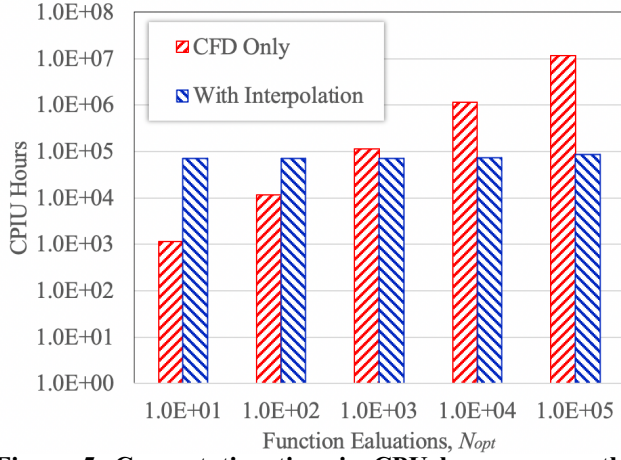


Figure 5. Computation time in CPU hours versus the number of evaluations for optimization with CFD only, and with the surrogate model included.

same baseline airfoil). In this scenario, N_{cfd} is equal to zero, and Eq. (1) can be simplified to show the effective speedup factor, S_{eff} , such that

$$S_{eff} = \frac{N_{opt}(t_{cfd} + t_{opt})}{N_{opt}t_{opt}}. \quad (2)$$

Optimization Problem Statement

A new airfoil design is obtained via variable substitution by Dakota. Let ϑ_i be the set of input parameters such that $\vartheta_i = \{f_m, \Delta p, f_t, \Delta x\}_i$ is the set of parameter augmentation values for the i^{th} airfoil section. For this work, the remaining parameters are fixed at their baseline values. A surrogate model is then used to generate an updated C81 table for the new airfoil. The table is provided to RCAS, and the airfoil performance is evaluated based on the rotorcraft power coefficient at a given flight condition. This study considers two cases: hover and forward flight with advance ratio $\mu = 0.3$. The power coefficients in these two cases, $C_{p,h}$ and $C_{p,f}$ respectively, are therefore functions of ϑ_i . Thus, the optimization formulation is

$$\begin{aligned} &\text{minimize} && f_1 = C_{p,f}(\vartheta_1, \dots, \vartheta_n) \\ & && f_2 = C_{p,h}(\vartheta_1, \dots, \vartheta_n) \\ &\text{subject to} && \vartheta_i \in X, \quad i = 1, \dots, n \end{aligned}$$

where n is the number of airfoil sections being optimized and X is the feasible design space (defined in Table 1 for the SC1095 airfoil).

When the maneuvering case is used, an inequality constraint function, g_1 , is added to the optimization formulation. g_1 is defined as the ratio of the maximum PLL_{hptp} for the modified blade to that of the baseline blade, and is represented as

$$g_1(\vartheta_1, \dots, \vartheta_n) = \frac{\max(PLL_{hptp}(\vartheta_1, \dots, \vartheta_n))}{\max(PLL_{hptp}(\vartheta_0))}. \quad (3)$$

A design violates the constraint if the value of g_1 exceeds that of the baseline blade configuration, i.e., $g_1(\vartheta_1, \dots, \vartheta_n) \leq 1$.

Multiobjective Genetic Algorithm

MOGA is the primary optimization technique use in this work and is based on the work of Eddy and Lewis (Ref. 15). The algorithm uses distinct-point and clustering metrics to generate a distribution of points in the Pareto front that is as close to uniform as possible. Some of the key aspects of the algorithm are summarized below, along with some of the specific settings used in this work.

Optimization begins with an initial population, pop_0 , of points distributed throughout the design space, typically a simple random distribution. The initial population is evaluated and then assessed according to a fitness function. In this case, the fitness function is the “domination count” of each point. By definition, point B has a domination count, n , if there exist n points, A_1, \dots, A_n , that improve the outcome of one or more objectives of B without degradation of another. Put simply, points A_1, \dots, A_n are more optimal than point B. Figure 6 illustrates how domination count is determined for a simple case in which the two objectives are to be minimized. Any points that fall in the lower-left quadrant relative to a point B are said to dominate B, as shown by the red, dashed lines.

A constraint penalty function is also used to measure the feasibility of each design. The constraint metric, D , is the L_2 norm of the constraint functions and is computed using

$$D = \sqrt{\sum_{j=1}^{N_j} (\max(0, g_j(\vartheta)))^2 + \sum_{k=1}^{N_k} |h_k(\vartheta)|^2}, \quad (4)$$

where ϑ is the design parameter vector, g_j for $j = 1, \dots, N_j$ are the inequality constraints, and h_k for $k = 1, \dots, N_k$ are the equality constraints.




Figure 6. Simple multiobjective response plot showing how point B is dominated by A_1 , A_2 , and A_3 . The red, dashed lines indicate the boundary between dominating and non-dominating points.

The genetic algorithm then loops through the following steps until convergence or stopping criteria is reached:

1. **Generate “mating pool”** – The mating pool is generated based on the feasibility metric, D , determined from Eq. (4). All points in the current population, pop_i , are represented at least once in the pool, regardless of feasibility. Points with lower values of D are represented more frequently, increasing the likelihood that they will reproduce. All feasible points are given an equal chance of reproducing that is higher than for infeasible points.
2. **Explore the design space** – New design points are generated via two sequential processes: crossover and mutation.
 - a. **Crossover** – Two individuals in the mating pool are chosen at random to form a “mating pair.” The two points are crossed over using bit switching in the binary encoded genome of two designs, a method known as single point uniform parameterized crossover (Ref. 32). The resulting “children,” $\text{pop}_{i,c}$, are then added to pop_i , yielding the interim population: $\text{pop}_{interim} = \text{pop}_i + \text{pop}_{i,c}$. The number of children produced from each pair may vary and is set to three for this work. The chance of a crossover event occurring for any given point is 80%.
 - b. **Mutation** – Introduce random variation by choosing a random variable from a random point and converting it to a binary string. It then flips a randomly chosen bit in the string from 0 to 1, or vice versa. Thus, the interim population becomes $\text{pop}_{interim} = M(\text{pop}_i + \text{pop}_{i,c})$, where M is the bit mutation operator. The chance of a mutation event occurring for any given point is 8%.
3. **Evaluate the fitness function** – Evaluate objective and constraint functions for the interim population. Note that duplicate points in any generation are not reevaluated. The interim population is then ranked according to domination count.
4. **Check stopping criteria** – Determine whether stopping criteria are satisfied (e.g., maximum function evaluations) and proceed accordingly.
 - a. **Not satisfied** – Down-select the interim population to points with a domination count of six or fewer. Apply additional “niche pressure” by enforcing a minimum Euclidean distance between designs. This controls clustering of very similar points, thus limiting the tendency for the population size to grow very large and encouraging differentiation along the Pareto front. For this work, the minimum distance is 1×10^{-6} units. The units correspond to the units of each objective (non-dimensional in this case). The new population, pop_{i+1} , is returned to step 1.
 - b. **Satisfied** – End optimization at the current generation. The Pareto front is the non-dominated subset of the final population (i.e., domination count is zero).




Figure 7. Evolution of $C_{p,h}$ vs. $C_{p,f}$ for select populations in CS2.

Figure 7 illustrates the evolution of $C_{p,h}$ and $C_{p,f}$ responses for select population numbers in an example case (Case Study 2 in the following section). The plot shows the progression from the relatively wide range of responses in pop_0 toward the final set of non-dominated points making up the Pareto front. Note that the number of function evaluations is much greater than the sum of the points in each population. This is due to the inflated size of each interim population that must be evaluated to ensure coverage of the design space.

RESULTS

Case Study 1: Single Airfoil Section

For the first case study (CS1), the optimization framework described above was used to morph the Tip segment of the UH-60A main rotor to minimize $C_{p,h}$ and $C_{p,f}$. The design parameters were airfoil camber, camber crest position, thickness, and thickness crest position. This simulation resembles the ones reported in Ref. 14, however it differs in that it considers a smaller portion of the blade (0.92R to 1.0R instead of 0.84R to 1.0R), and uses more design parameters.

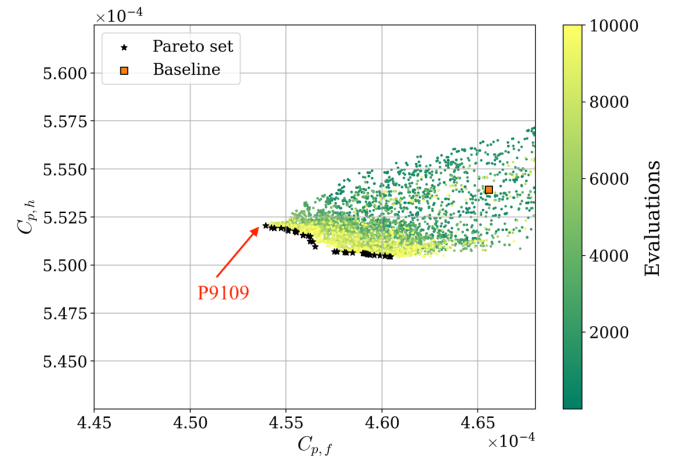


Figure 8. CS1 multiobjective response. Plot shows $C_{p,h}$ vs. $C_{p,f}$ and is colored by the evaluation number.




Figure 9. CS1 optimized airfoil geometry compared with the SC1095 (baseline).

The run was initialized with a population of 440 randomly distributed points. Optimization ended after reaching the maximum number of function evaluations; 10,000 in this case. The result was 9,974 unique designs evaluated over 10 optimizer generations. The Pareto front contained 64 points. The wall-clock time, including queue, for CS1 was 33.7 hours. The estimated theoretical speedup factor yielded by Eq. (1) is $S = 15.9$, indicating that the proposed approach is 15.9 times faster than using CFD alone, including the time needed for database generation.

The optimized rotor performance gains are, expectedly, modest given the limited scope of blade modifications. Figure 8 shows the responses of $C_{p,h}$ and $C_{p,f}$ for this case. The points are colored by the evaluation number. Note that many function evaluations take place within each optimizer generation. A single point was selected from the Pareto front for further analysis and is highlighted in the figure. The point, P9109, was selected since it provided the largest improvement in $C_{p,f}$ of any design in the experiment. This selection criterion is used due to the relative sensitivity of $C_{p,f}$ being higher than $C_{p,h}$. The observed reduction in $C_{p,f}$ for the selected point was 2.38% while it also reduced $C_{p,h}$ by 0.35%. Whereas the highest performing points when favoring $C_{p,h}$ only improve the objective to 0.65%. Further evidence supporting the choice in this P9109 is given later in this section.

The modified airfoil cross-section, denoted CS1-P9109, is shown in Figure 9. Compared to the baseline (SC1095) the optimal design has reduced camber (m_{Tip}) and the maximum camber crest position (p_{Tip}) is moved toward the leading edge. Additionally, the thickness (t_{Tip}) has been reduced and the thickness crest position (x_{Tip}) moved toward the trailing edge.

Case Study 2: Two Airfoil Sections

The next case study (CS2) considers morphing of the Tip and OB1 blade segments. The optimizer was, again, initialized with 440 randomly distributed points. To account for the increase in design variables, the maximum function evaluations were increased to 50,000. This setup resulted in 49,857 unique designs evaluated over 17 MOGA generations. The final Pareto front contained 218 points. Wall-clock time,




Figure 10. CS2 multiobjective response. Plot shows $C_{p,h}$ vs. $C_{p,f}$ and is colored by the evaluation number.




Figure 11. CS2 optimized airfoil geometries compared with the SC1095 (baseline).

including queue, for CS2 was 142.0 hours and Eq. (1) yields the estimated CPU time speedup factor, $S = 77.5$. The increase in wall-clock time compared to CS1 is due to the increased number of evaluations. Conversely, it is the larger number of evaluations that provides the substantial increase in speedup factor. Additionally, CS2 leverages the database that was generated previously. Therefore, Eq. (2) is used to calculate the effective speedup factor, yielding $S_{eff} = 2747$.

The objective responses for all points in CS2 are plotted in Figure 10, with the Pareto front shown as black stars. The point P47522 is selected from the Pareto front using the same criteria as for CS1. This design point reduces $C_{p,f}$ and $C_{p,h}$ by 3.74% and 0.80%, respectively. The optimized airfoil geometries for CS2 are shown in Figure 11. The tip airfoil, denoted CS2-P47522_Tip, has deformations similar to that in CS1. The maximum camber and thickness are both decreased, while the crest positions move toward the leading edge and trailing edge, respectively. The thickness profile (t_{OB1} and x_{OB1}) for the outboard airfoil, CS2-P47522_OB1, is very similar to that of the tip airfoil. However, CS2-P47522_OB1 differs from both the baseline and CS2-P47522_Tip in its camber profile. The maximum camber (m_{OB1}) increased by approximately 27% from the baseline, while the camber crest position (p_{OB1}) moved slightly toward the trailing edge compared to that of the baseline.

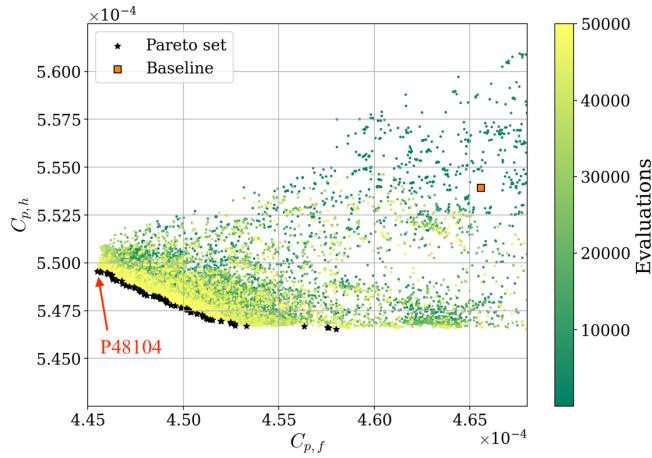


Figure 12. CS3 multiobjective response. Plot shows $C_{p,h}$ vs. $C_{p,f}$ and is colored by the evaluation number.

Case Study 3: Three Airfoil Sections

The final case study (CS3) considers the morphing of all three airfoil segments: Tip, OB1, and OB2. The total number of design variables for CS3 is 12. This case used an initial population of 440 points and limited the maximum evaluations to 50,000. The result was 49,923 unique evaluations over 10 optimizer generations. The final population size was 93 points. The wall-clock time, including queue, for this simulation was 84.8 hours with up to 440 concurrent evaluations. The decrease in wall-clock time compared to CS2 indicates that, on average, much less time was spent waiting in the queue. Equations (1) and (2) yield speedup factors $S = 77.5$ and $S_{eff} = 2747$, respectively. These values are identical to those of CS2 since they use the same database and have nearly the same number of unique evaluations.

The response plot, as well as the selected point for this case (P98104), are shown in Figure 12. P98104 improves $C_{p,f}$ and $C_{p,h}$ by 4.63% and 0.59%, respectively. Optimal airfoil geometries for CS3 are shown in Figure 13. The tip and OB1 airfoils, CS3-P48104_{Tip} and CS3-P48104_{OB1}, respectively, are drawn in Figure 13a with the SC1095. Whereas, the OB2 airfoil, CS3-P48104_{OB2}, is drawn in Figure 13b with the SC1094R8. This division accounts for the differing baseline airfoils at each segment.

Note that while CS3-P48104_{OB2} replaces the SC1094R8 for the indicated section, it is derived parametrically from the SC1095. This distinction makes no difference if all nine design parameters are included. However, since the degrees of freedom are limited to m , p , t , and x in this analysis, there are locations in the design space that cannot be accessed (i.e., it cannot recreate the SC1094R8). This limitation is one possible explanation for the slight degradation of $C_{p,h}$ from CS2 to CS3. If the full design space were available, the optimization framework is expected to match, or improve upon, all objectives from the previous case. Another possible reason for the degradation is that the maximum number of




Figure 13. CS3 optimized airfoil geometries. Plot a) shows the OB1 and Tip segments compared with the SC1095 airfoil, whereas plot b) shows the OB2 segment compared with the SC1094R8 airfoil.

Table 2. Summary of blade parameters and responses for CS1, CS2, and CS3 (fraction of airfoil chord).

	CS1 P9109	CS2 P47522	CS3 P48104	Baseline ^a UH-60A
m_{Tip}	0.0045	0.0040	0.0045	0.0081
p_{Tip}	0.1770	0.1800	0.1840	0.2700
t_{Tip}	0.0712	0.0715	0.0706	0.0950
x_{Tip}	0.3500	0.3610	0.3490	0.2700
m_{OB1}	—	0.0102	0.0108	0.0081
p_{OB1}	—	0.2260	0.2230	0.2700
t_{OB1}	—	0.0712	0.0714	0.0950
x_{OB1}	—	0.3540	0.3460	0.2700
m_{OB2}	—	—	0.0117	0.0210
p_{OB2}	—	—	0.2380	0.2170
t_{OB2}	—	—	0.0725	0.0940
x_{OB2}	—	—	0.3150	0.2780

^a SC1095 for the Tip and OB1 segments, SC1094R8 for the OB2 section.

evaluations were unchanged from CS2 to CS3. More evaluations/generations would likely drive the Pareto front to a point where $C_{p,h}$ is improved further by CS3. An improvement of this nature is expected to be limited.

A summary of the optimized geometries for all three case studies is shown in Table 2. The values for m , p , t , and x for each deformed section, as well as the baseline airfoils, are represented in fractions of airfoil chord length. The crest positions are measured from the leading edge. The CS3

solutions for the Tip and OB1 blade segments are similar to their counterparts in CS1 and CS2. The OB2 maximum camber (m_{OB2}) is the largest of any optimized designs, as is the thickness (t_{OB2}), though each is substantially less than the baseline value. The crest positions exhibit the same trends as noted previously; with the maximum camber (p_{OB2}) moving toward the leading edge and the maximum thickness (x_{OB2}) moving toward the trailing edge.

Speed Sweep Analysis

A speed sweep was conducted for the selected points from CS1, CS2, and CS3. RCAS was used to calculate C_p for several forward flight speeds using each optimized configuration. Trim targets for this analysis are based on the measured values from the NFAC flight test (Ref. 30). The results of this sweep are shown in Figure 14. Despite only two of the conditions being considered during optimization, the rotor performance is improved at all flight speeds, especially at high advance ratios ($\mu > 0.3$).

A second set of the optimal designs favoring hover performance is chosen from the Pareto front for the three different segment designs, denoted as CS1-P9735, CS2-P46472, and CS3-45148. The associated results are shown in Figure 15. This figure supports the decision to analyze the designs optimized favoring $C_{p,f}$. The increase in hover performance gained from favoring $C_{p,h}$ in the optimization is not significant, whereas there is a substantial performance loss at high advance ratios. Using the baseline hover performance as the upper bound to an inequality constraint may also be an attractive solution.

An in-depth examination of the aerodynamic and structural loads for the CS2 and CS3 optimized designs is given in a companion paper (Ref. 33).

Constrained Optimization

Validation of the maneuvering flight condition used for evaluating the pitch link load (PLL) constraint is ongoing. That said, preliminary results show that consideration of PLL is critical to the design optimization process. The 3-section case study, CS3, was repeated using the constrained optimization problem described previously. The histogram in Figure 16 shows that the PLL_{hptp} exceeded that of the baseline for nearly all of the designs. This is the case for all designs in the Pareto set as well, which violated the PLL constraint by 14%-26%.

More work needs to be done to understand this unexpected result. One possibility is that the treatment of constraints by the optimizer (e.g., using Eq. (4) to penalize violators instead of discarding them) may be ill-suited to this particular problem. Additionally, the RCAS model used for the UTTAS maneuver needs to be closely examined and validated. A complete investigation into this is planned for imminent future work.

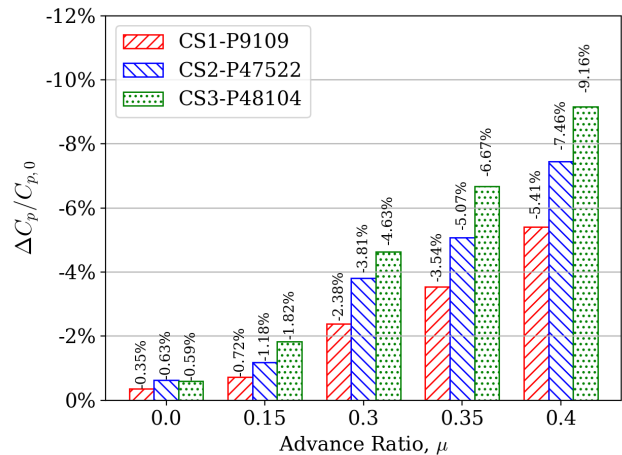


Figure 14. Rotor power percent change vs. advance ratio for CS1, CS2, and CS3 optimized favoring $C_{p,f}$.

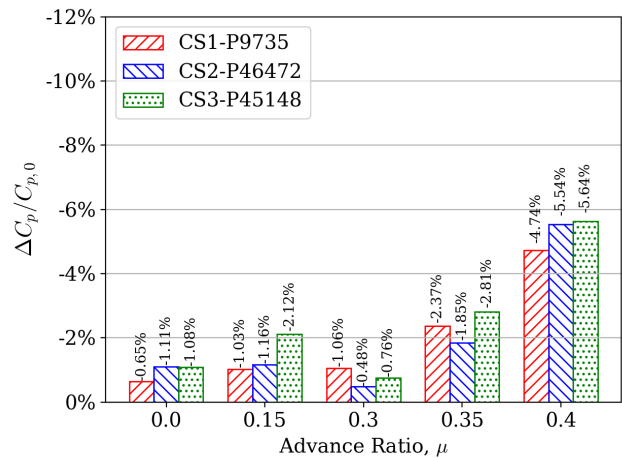


Figure 15. Rotor power percent change vs. advance ratio for CS1, CS2, and CS3 optimized favoring $C_{p,h}$.




Figure 16. Probability density histogram for the maximum PLL_{hptp} .

CONCLUSIONS

This paper demonstrated advancements to the multiobjective rotor blade optimization framework first introduced in Ref. 14. The major take-aways from this work are:

1. An automated framework for multiobjective airfoil shape optimization of multiple airfoil sections has been developed
2. The framework was implemented for unconstrained optimization with multiple airfoil sections and reduced the rotor power by 0.75% in hover and 8.78% for forward flight ($\mu = 0.4$)
3. Large reductions in rotor power are consistently obtained over a wide speed range
4. The automated optimization framework successfully handled several problems of different sizes, though expertise is still needed by the user to determine an appropriate optimizer and settings

ACKNOWLEDGMENTS

Funding for this program was provided by the Deputy Assistant Secretary of the Army for Research and Technology and supported by the Future Vertical Lift Cross-Functional Team. This work was performed by the US Army Engineer Research and Development Center (ERDC) under the Engineered Resilient Systems program, in collaboration with the US Army U.S. Army Combat Capabilities Development Command (CCDC) Aviation and Missile Center (AvMC). This work was made possible by the HPC Modernization Program's allocation of compute time on the ERDC DSRC's Cray XC40, Onyx.

REFERENCES

1. Imiela, M. and Wilke, G., "Passive Blade Optimization and Evaluation in Off-Design Conditions," presented at the 39th European Rotorcraft Forum, Moscow, Russia, Sep. 2013.
2. Adams, B. M., Bohnhoff, W. J., Dalbey, K. R., Ebeida, M. S., Eddy, J. P., Eldred, M. S., Gararci, G., Hooper, R. W., Hough, P. D., Hu, K. T., Jakeman, J. D., Khalil, M., Maupin, K. A., Monschke, J. A., Ridgway, E. M., Rushdi, A. A., Stephens, J. A., Swiler, L. P., Vigil, D. M., Wildey, T. M., and Winokur, J. G., "Dakota, a Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.11 Reference Manual," Sandia National Laboratories SAND2014-5015, 2018. [Online]. Available: <http://dakota.sandia.gov/documentation.html>.
3. Arnaud, G. and Beaumier, P., "Validation of R85/METAR on the Puma Rae Flight Tests," presented at the 18th European Rotorcraft Forum, 1992.
4. León, E. R., Le Pape, A., Désidéri, J.-A., Alfano, D., and Costes, M., "Concurrent Aerodynamic Optimization of Rotor Blades Using a Nash Game Method," presented at the American Helicopter Society 69th Annual Forum, Phoenix, AZ, May 2013.
5. Vanderplaats, G. N., "CONMIN – a FORTRAN Program for Constrained Function Minimization," NASA TM X-62282, 1973.
6. Cambier, L. and Veillot, J., "Status of the ElsA CFD Software for Flow Simulation and Multidisciplinary Applications," presented at the 46th AIAA Aerospace Science Meeting and Exhibit, 2008.
7. Benoit, B., Dequin, A., Kampa, K., Von Grünhagen, W., Basset, P., and Gimonet, B., "Host, a General Helicopter Simulation Tool for Germany and France," presented at the American Helicopter Society 56th Annual Forum, Virginia Beach, Virginia, May 2000.
8. Prieur, J. and Splettstoesser, W. R., "ERATO – An ONERA-DLR Cooperative Programme on Aeroacoustic Rotor Optimization," presented at the 25th European Rotorcraft Forum, Rome, Italy, Sep. 1999.
9. Nash, J., "Non-Cooperative Games," *Annals of Mathematics*, Vol. 54, (2), pp. 286–295, 1951.
10. Wang, L., Diskin, B., Lopes, L. V., Nielson, E. J., Lee-Rausch, E., and Biedron, R. T., "High-Fidelity Multidisciplinary Design Optimization of Low-Noise Rotorcraft," presented at the Vertical Flight Society 75th Annual Forum & Technology Display, Philadelphia, PA, May 2019.
11. Gill, P. E., Murray, W., and Saunders, M. A., "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, (1), pp. 99–131, 2005.
12. Samareh, J. A., "Novel Multidisciplinary Shape Parameterization Approach," *Journal of Aircraft*, Vol. 38, (6), pp. 1015–1024, 2001.
13. Lim, J. W., "Application of Parametric Airfoil Design for Rotor Performance Improvement," presented at the 44th European Rotorcraft Forum, Delft, Netherlands, Sep. 2018.
14. Allen, L. D., Lim, J. W., Haehnel, R. B., and Dettwiler, I. D., "Rotor Blade Design Framework for Airfoil Shape Optimization with Performance Considerations," presented at the AIAA SciTech 2021 Forum, Virtual Event, Jan. 2021, DOI:10.2514/6.2021-0068.
15. Eddy, J. and Lewis, K., "Effective Generation of Pareto Sets Using Genetic Programming," presented at the ASME International Design Technical Conferences, Design Automation Conference, Pittsburgh, PA, 2001.

16. Pözlbauer, P., Desvigne, D., and Breitsamter, C., "Aerodynamic Design Optimization of a Helicopter Rotor Blade-Sleeve Fairing," *CEAS Aeronautical Journal*, Vol. 10, (3), pp. 665–685, 2019, DOI:10.1007/s13272-018-0341-0.
17. Hersey, S., Sridharan, A., and Celi, R., "Multiobjective Performance Optimization of a Coaxial Compound Rotorcraft Configuration," *Journal of Aircraft*, Vol. 54, (4), pp. 1498–1507, 2017, DOI:10.2514/1.C033999.
18. Lim, J., Shin, S., and Kee, Y., "Optimization of Rotor Structural Design in Compound Rotorcraft with Lift Offset," *Journal of the American Helicopter Society*, Vol. 61, (1), pp. 1–14, 2016, DOI:10.4050/JAHS.61.012005.
19. Chugh, T., Sindhya, K., Hakanen, J., and Miettinen, K., "A Survey on Handling Computationally Expensive Multiobjective Optimization Problems with Evolutionary Algorithms," *Soft Computing*, Vol. 23, (9), pp. 3137–3166, 2019, DOI:10.1007/s00500-017-2965-0.
20. Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, Vol. 6, (2), pp. 182–197, 2002, DOI:10.1109/4235.996017.
21. Beume, N., Naujoks, B., and Emmerich, M., "SMS-EMOA: Multiobjective Selection Based on Dominated Hypervolume," *European Journal of Operational Research*, Vol. 181, (3), pp. 1653–1669, 2007, DOI:10.1016/j.ejor.2006.08.008.
22. Zhang, Q. and Li, H., "MOEA/D: A Multiobjective Evolutionary Algorithm Based on Decomposition," *IEEE Transactions on Evolutionary Computation*, Vol. 11, (6), pp. 712–731, 2008, DOI:10.1109/TEVC.2007.892759.
23. Zitzler, E., Laumanns, M., and Thiele, L., "SPEA2: Improving the Strength Pareto Evolutionary Algorithm," TIK Report 103, 2001.
24. Coello, C. A. C. and Lechuga, M. S., "MOPSO: A Proposal for Multiple Objective Particle Swarm Optimization," *Proceedings of the 2002 Congress on Evolutionary Computation*, Vol. 2, Feb. 2002, pp. 1051–1056, DOI:10.1109/CEC.2002.1004388.
25. Rajagopalan, R. G., Baskaran, V., Hollingsworth, A., Lestari, A., Garrick, D., Solis, E., and Hagerty, B., "RotCFD - A Tool For Aerodynamic Interference of Rotors: Validation and Capabilities," presented at the American Helicopter Society Future Vertical Lift Aircraft Design Conference, San Francisco, California, Jan. 2012.
26. Koning, W. J. F., Johnson, W., and Allan, B. G., "Generation of Mars Helicopter Rotor Model for Comprehensive Analyses," presented at the AHS International Technical Meeting on Aeromechanics Design for Transformative Vertical Flight, San Francisco, California, Jan. 2018.
27. Strawn, R., Mayda, E. A., and van Dam, C. P., "Automated CFD for Generation of Airfoil Performance Tables," NASA ARC-15649-1, 2009.
28. Saberi, H., Hasbun, M., Hong, J., Yeo, H., and Ormiston, R. A., "Overview of RCAS Capabilities, Validations, and Rotorcraft Applications," presented at the American Helicopter Society 71st Annual Forum, Virginia Beach, Virginia, May 2015.
29. Jain, R., Yeo, H., Ho, J., and Bhagwat, M., "An Assessment of RCAS Performance Prediction for Conventional and Advanced Rotor Configurations," *Journal of the American Helicopter Society*, Vol. 61, (042005), 2016, DOI:10.4050/JAHS.61.042005.
30. Norman, T. R., Shinoda, P., Peterson, R. L., and Datta, A., "Full-Scale Wind Tunnel Test of the UH-60A Airloads Rotor," presented at the American Helicopter Society 67th Annual Forum, Virginia Beach, Virginia, May 2011.
31. Bousman, W. and Kufeld, R., "UH-60A Airloads Catalog," NASA TM-2005-212827, Aug. 2005.
32. Anderson, K. S. and Hsu, Y., "Crossover Strategy For Improved Solution Space Exploration With Genetic Algorithms," presented at the ASME 1998 Design Engineering Technical Conferences, Atlanta, Georgia, Sep. 1998.
33. Lim, J. W., Allen, L. D., Haehnel, R. B., and Dettwiler, I. D., "An Examination of Aerodynamic and Structural Loads for a Rotor Blade Optimized with Multi-Objective Genetic Algorithm," to be presented at the 77th VFS Annual Forum, Virtual, May 2021.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) June 2021			2. REPORT TYPE Final			3. DATES COVERED (From - To)			
4. TITLE AND SUBTITLE Helicopter Rotor Blade Multiple-Section Optimization with Performance Considerations						5a. CONTRACT NUMBER			
						5b. GRANT NUMBER			
						5c. PROGRAM ELEMENT NUMBER 0603456			
6. AUTHOR(S) Luke D. Allen, Joon W. Lim, Robert B. Haehnel, and Ian D. Dettwiller						5d. PROJECT NUMBER AL3			
						5e. TASK NUMBER			
						5f. WORK UNIT NUMBER			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) See next page.						8. PERFORMING ORGANIZATION REPORT NUMBER ERDC MP-21-4			
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Headquarters, U.S. Army Corps of Engineers 441 G Street NW Washington, DC 20314-1000						10. SPONSOR/MONITOR'S ACRONYM(S) USACE			
						11. SPONSOR/MONITOR'S REPORT NUMBER(S)			
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.									
13. SUPPLEMENTARY NOTES This paper was originally presented at the <i>Vertical Flight Society's 77th Annual Forum & Technology Display</i> , Virtual, May 10-14, 2021.									
14. ABSTRACT This paper presents advancements in a surrogate-based, rotor blade design optimization framework for improved helicopter performance. The framework builds on previous successes by allowing multiple airfoil sections to designed simultaneously to minimize required rotor power in multiple flight conditions. Rotor power in hover and forward flight, at advance ratio $\mu = 0.3$, are used as objective functions in a multi-objective genetic algorithm. The framework is constructed using Galaxy Simulation Builder with optimization provided through integration with Dakota. Three independent airfoil sections are morphed using ParFoil and aerodynamic coefficients for the updated airfoil shapes (i.e., lift, drag, moment) are calculated using linear interpolation from a database generated using C81Gen/ARC2D. Final rotor performance is then calculated using RCAS. Several demonstrative optimization case studies were conducted using the UH-60A main rotor. The degrees of freedom for this case are limited to the airfoil camber, camber crest position, thickness, and thickness crest position for each of the sections. The results of the three-segment case study show improvements in rotor power of 4.3% and 0.8% in forward flight and hover, respectively. This configuration also yields greater reductions in rotor power for high advance ratios, e.g., 6.0% reduction at $\mu = 0.35$, and 8.8% reduction at $\mu = 0.4$.									
15. SUBJECT TERMS Helicopters—Design; Rotors (Helicopters)—Design; Helicopters--Computer simulation; Rotors (Helicopters)--Computer simulation; Genetic algorithms									
16. SECURITY CLASSIFICATION OF:						17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT Unclassified		b. ABSTRACT Unclassified		c. THIS PAGE Unclassified		SAR	18	19b. TELEPHONE NUMBER (include area code)	

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Cold Regions Research and Engineering Laboratory
U.S. Army Engineer Research and Development Center
72 Lyme Road
Hanover, NH 03755

Development and Engineering Center
U.S. Army Aviation and Missile Research
75 B. S. Hood Road
Moffett Field, CA 94035

Information Technology Laboratory
U.S. Army Engineer Research and Development Center
3090 Halls Ferry Road
Vicksburg, MS 39180