

Optimization of Turbomachinery Airfoils with a Genetic/Sequential-Quadratic-Programming Algorithm

Brian H. Dennis*

The Pennsylvania State University, University Park, Pennsylvania 16802
and

George S. Dulikravich[†] and Zhen-Xue Han[‡]

The University of Texas at Arlington, Arlington, Texas 76019

The objective of this aerodynamic shape design effort is to minimize total pressure loss across the two-dimensional linear-airfoil cascade row while satisfying a number of constraints. They included fixed axial chord, total torque, inlet and exit flow angles, and blade cross-section area, while maintaining thickness distribution greater than a minimum specified value. The aerodynamic shape optimization can be performed by using any available flowfield analysis code. For the analysis of the performance of intermediate cascade shapes, we used an unstructured-grid-based compressible Navier–Stokes flowfield analysis code with a $k-\epsilon$ turbulence model. A robust genetic optimization algorithm was used for optimization, and a constrained sequential quadratic programming was used for enforcement of certain constraints. The airfoil geometry was parameterized using conic section parameters and B-splines, thus keeping the number of geometric design variables to a minimum while achieving a high degree of geometric flexibility and robustness. Significant reductions of the total pressure loss were achieved using this constrained method for a supersonic exit flow axial turbine cascade.

Nomenclature

A	=	cross-sectional area of the airfoil, m^2
c	=	penalty constant
L	=	lift, N
\dot{m}	=	mass flow rate, $kg\ s^{-1}$
P_0	=	total pressure, Pa
θ	=	average exit flow angle, deg

I. Introduction

AERODYNAMIC shape design has a global objective of increasing the aerodynamic efficiency of the flowfield created by the designed object. This means that the design objective should be minimization of entropy generation in the flowfield. In general, entropy generation is caused by viscous dissipation, heat transfer, internal heat sources, chemical reactions, and electromagnetogasdynamics effects. In the case of a turbomachinery aerodynamics, sources of entropy production other than viscous effects and heat transfer could be neglected. This means that the entropy is generated by viscosity and heat transfer. The most prominent regions of entropy production are viscous boundary layer and shock waves.¹ Then, any reliable Navier–Stokes flowfield analysis code could be used to compute local values of total pressure or entropy. For a given set of inlet and exit flow boundary conditions, the amount of entropy generated in the two-dimensional flowfield is determined by the shape of the turbomachinery airfoil cascade.

Minimization of the entropy generation (flow losses) can therefore be achieved by the proper reshaping of the airfoil cascade. This process of achieving a global design goal is called shape optimization. It is different and much more powerful than the shape inverse design^{2–4} where a desired pressure distribution is enforced on the

blade surface. The surface pressure is often enforced without any a priori evaluation criterion⁵ to judge the effects of this pressure distribution on the corresponding global design objectives. There have been numerous inverse shape-design attempts reported where the desired surface-pressure distribution was enforced via an optimization algorithm. However, there are considerably less expensive and more accurate methods for performing the aerodynamic shape inverse design.⁶ In other words, the optimization algorithms should be used for actual optimization of the global aerodynamic and geometric parameters, not for achieving inverse shape design by enforcing the specified surface pressure distributions.^{3,6}

Shape optimization, especially when subjected to numerous user-specified constraints, is considerably more time consuming than the shape inverse design because it requires a large number of calls to the aerodynamic flowfield analysis code. This was one of the main reasons why the aerodynamic shape optimization has not been used until very recently. With the rapidly increasing use of distributed parallel computing, the adequate computing hardware is becoming affordable to perform aerodynamic shape optimization. This trend is also supported by the fact that for optimization purposes the flowfield analysis code does not need to be changed, whereas for inverse shape design it is often necessary to perform extensive modifications to an existing flowfield analysis code.⁴

II. Problem Statement and Constraints

There are basically two types of aerodynamic shape-design problems. One problem is to design an entirely new turbomachine in which case only loose geometric constraints are needed. A considerably more difficult problem is the aerodynamic shape design for retrofitting an existing turbomachine with a more efficient rotor or a stator. This is because the retrofitted design has a number of stringent constraints. In this work the following constraints were enforced.

The axial chord of the new row of blades must be the same or slightly smaller than the axial chord length of the original blade row. Otherwise, the new blade row will not be able to fit in the existing turbomachine. Inlet and exit flow angles must be the same in the redesigned blade row as in the original blade row, or the velocity triangles will not match with the neighboring blade rows. Mass flow rate through the new blade row must be the same as through the original blade row or the entire machine will perform at an

Received 24 July 1999; revision received 11 April 2001; accepted for publication 12 April 2001. Copyright © 2001 by the authors. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission.

*Graduate Research Assistant, Department of Aerospace Engineering; currently Postdoctoral Fellow, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan. Student Member AIAA.

[†]Professor and Director of MAIDO Laboratory, Department of Mechanical and Aerospace Engineering, UTA Box 19018. Fellow AIAA.

[‡]Research Fellow, Department of Mechanical and Aerospace Engineering, P.O. Box 19018.




Fig. 1 Basic airfoil parameterization.

off-design mass flow rate. This can lead to serious drop in overall efficiency and create unsteady flow problems. Torque created on the new rotor blade row must be the same as on the old rotor blade row, or the new rotor will rotate at the wrong angular speed. The cross-section area of the new blade should be the same as the cross-section area of the original rotor blade. This is to ensure that the new blade will be able to sustain the expected loads without performing a detailed elasticity analysis of each candidate blade geometry. Minimum airfoil thickness distribution has to be specified, or the new aerodynamically optimized cascade of airfoils could have (an example will be shown in this paper) an unacceptably thin shape. In the case of an internally cooled turbine blade, such a thin airfoil will not be able to accommodate the internal coolant flow passages. In addition, airfoil trailing-edge thickness should not be smaller than a user-specified value, or it will overheat and burn in the case of a turbine airfoil. This condition can be easily incorporated in the constraint on the specified minimum airfoil thickness distribution. Gap-to-axial chord ratio will be kept fixed in this work because the total number of blades will be treated as fixed. Otherwise, if the gap between the neighboring airfoils is treated as an additional design variable the optimization will involve both shape optimization and the minimization of the number of blades in the rotor or a stator.⁷

In summary, this example of the aerodynamic shape optimization for a retrofit application includes five equality constraints and one inequality constraint.

III. Geometric Design Variables

The shape optimization of a two-dimensional cascade of airfoils will be performed by iteratively determining an optimum airfoil shape that also satisfies the constraints just listed. This means that a large number of candidate airfoil shapes will be created and examined using an appropriate constrained optimization algorithm. To minimize the computing effort, the candidate airfoil shapes will be parameterized with a relatively small number of parameters that will then serve as design variables. In this paper two approaches were used.

In the first approach of using a minimum possible number of the design variables, the shape of an airfoil in a turbomachinery cascade

was chosen to require only nine parameters (Fig. 1). These variables include the tangential and axial chord, the inlet and exit half-wedge angle, the inlet and outlet airfoil angle, the throat, unguided turning angle, and the leading- and trailing-edge radii.⁸⁻¹⁰ One of these parameters (axial chord) will be kept fixed in this exercise.

In the second approach the airfoil shape was allowed significant additional flexibility by adding a continuous arbitrary perturbation in addition to the original nine parameters. This shape perturbation was modeled with a B-spline that had eight control vertices, thus resulting in a total of $9 + 8 = 17$ design variables. The ranges of design variables were set so that the optimizer would have a wide variety of very different airfoil shapes so as to test its robustness. Users with experience in turbomachinery airfoil design could easily narrow these ranges and target a specific area where they think the best design will be found. For example, in the work of Trigg et al.¹¹ the airfoil geometry was parameterized using 17 geometric parameters with very narrow ranges, thus circumventing a need for explicit constraints. Their design application was not a retrofit, but a design of a new cascade.

IV. Flowfield Analysis

With the genetic algorithm variations in geometry from iteration to iteration might not be small especially if a large design space is to be explored. Thus, it is of utmost importance that a robust flowfield analysis code and a robust grid-generation code are used. For this reason we used a two-dimensional compressible turbulent flow Navier-Stokes analysis code that utilizes Runge-Kutta time stepping. The flowfield analysis code uses unstructured triangular grids¹² generated using a Delaunay tessellation with iterative point insertion.¹³ All boundary conditions in the flowfield analysis code were imposed using first-order accuracy. The inlet and outlet boundary conditions were treated with locally one-dimensional nonreflection of waves passing out of the computational domain. Implementation of two-dimensional nonreflection boundary conditions at inlet and exit boundaries would require considerably more computing time. The flowfield periodic conditions were imposed with phantom cells' values equal to the corresponding periodic real cells' values. The flow variables at the cells' vertices were calculated




Fig. 2 Surface isentropic Mach-number distribution on the original VKI turbine airfoil cascade with sonic exit flow condition: comparison of experimental data and numerical results.

from surrounding cells' centers' values with distance weighted averaging method.¹⁴ The $k-\epsilon$ limiter and wall function formulation were used in turbulence modeling.

Accuracy of the flowfield analysis code can be evaluated by comparing computed and experimentally measured surface isentropic Mach-number distribution (Fig. 2) on an axial turbine cascade tested by the VKI (von Karman Institute of Fluid Mechanics).¹⁵

In all test cases described in this paper, the inlet total pressure, and inlet flow angle were set to 440,000 Pa, 1600.0 K, and 30.0 deg, respectively. The exit static pressure was specified as 101,330 Pa, thus making this a supersonic exit cascade. Adiabatic wall conditions were enforced along the airfoil surface. The unstructured grid flowfield analysis code was used by the optimizer to determine the total pressure loss, lift, and exit flow angle for a given airfoil design. The unstructured grids generated around typical airfoil cascade designs contained around 3000 nodes and 6000 triangles. Reduction of the average residual by three and a half orders of magnitude takes the flow solver around 1200 time steps when starting from an initially uniform flowfield and a Courant-Friedrichs-Levy number of 1.8. A single flowfield analysis for a typical airfoil cascade takes around eight minutes on a 175-MHz R10000 processor.

V. Optimization Algorithm

A genetic algorithm (GA)-based optimizer¹⁶⁻¹⁸ was used to design a turbine airfoil cascade shape by varying the geometric parameterization parameters. The objective of the optimization was to determine the airfoil shape that gives the minimum total pressure loss while conforming to the specified constraints of producing 22,000.0 N lift force, an average exit flow angle of -63.0 degrees, a mass flow rate of 15.5 kg s^{-1} , an airfoil cross-sectional area of 0.00223 m^2 , and an airfoil axial chord of 0.1 m. One possibility of enforcing these constraints is to use a penalty function formulation by combining them with the true objective function (the difference between the total pressure computed at the exit and the total pressure imposed at the inlet). The mathematical form of such a combined objective function F could be expressed as

$$F = p_0^{\text{outlet}} - p_0^{\text{inlet}} + c_0(22,000.0 - L)^2 + c_1(-63.0 - \theta')^2 + c_2(15.5 - \dot{m})^2 + c_3(0.00223 - A)^2 + c_4(tde) \quad (1)$$

The variable tde is the largest relative error in the airfoil thickness distribution compared to a specified thickness distribution that is considered to be minimum allowable. This geometric constraint prevents airfoil from becoming too thin so that it would not become mechanically or thermally infeasible. The constants c_i are user-specified penalty terms.

The GA-based optimizer used the micro-GA technique¹⁹ with no mutation. The objective was to perform the entire constrained optimization with as few calls to the flowfield analysis code as possible. Trigg et al.¹¹ did not use the micro-GA approach. Instead, they used a very large randomly generated initial population of designs, which was then progressively reduced during each consecutive optimization cycle. As a result, they typically required 1000 flowfield analysis runs to optimize an airfoil cascade even without enforcing stringent constraints.

A binary string that used eight bits for each design variable represented each design in the population. A tournament selection was used to determine the mating pairs.¹⁶ Each pair produced two children designs that then replaced the parents in the population. Uniform crossover with a 50% probability of crossover was used to produce the children. Elitism was also implemented in the optimizer; the best individual found from the previous generations was placed in the current generation.

A. Penalty Method for Constraints

This example optimization case used penalty terms alone to enforce the constraints. A constant penalty term $c_i = 4.0E6$ was applied to each normalized constraint. For this application the micro-GA approach was found to be rather insensitive to the exact value of the penalty terms. It was found that the micro-GA algorithm performed better at satisfying the constraints when the penalty constants were larger than $1.0E5$.

The first example optimization run was performed without enforcing a minimum thickness distribution constraint. Instead, a minimum allowable trailing-edge radius was used as an inequality constraint. Range of the design variables is given in Table 1 indicating that B-spline shape perturbations were not used in this example where the number of design variables was nine.

The genetic optimizer was run for 45 generations with a population of 15 amounting to $45 \times 15 = 675$ flowfield analysis calculations. It used eight bit strings for each design variables. This implies very high accuracy in the representation of the design variables. In this study we considered this to be necessary in order to eliminate any possibility for premature convergence that might be caused by the insufficient accuracy of the representation of the geometry. The best design in the randomly generated initial generation of airfoil shapes had a total pressure loss of 11,231.2 Pa. The best design of generation 45 had a total pressure loss of 5965.6 Pa. This represents 46.9% reduction in the total pressure loss as compared to the best member of the initial population. Figure 3 shows the difference in airfoil shapes for the best designs from generations 1, 25, and 45. By generation 25 the airfoil had largely converged to the final shape shown for generation 45. Figure 4 shows the percent error in the constraints for the best design of each generation. It is obvious that the optimizer has difficulty satisfying all constraints simultaneously and tends to violate one constraint to satisfy another. This was particularly true for the lift and mass flow rate. By generation 45 all constraints were less than 4% violated. The most intriguing aspect of the optimized geometry is the unacceptably thin final airfoil shape, which cannot be used because it cannot accommodate internal cooling passages. Thus, this example should serve as a demonstration of an unacceptable geometry that can result if no constraint is enforced on the minimum allowable thickness distribution along the airfoil.

Table 1 User-specified range of design variables

Design variable	Minimum	Maximum
Tangential chord, m	0.05	0.16
Unguided turning, deg	-15.0	15.0
Inlet blade angle, deg	-25.0	45.0
Outlet half-wedge angle, deg	-10.0	10.0
Inlet half-wedge angle, deg	1.0	40.0
Leading-edge radius, m	0.001	0.03
Outlet blade angle, deg	-75.0	-57.0
Throat, m	0.05	0.10
Trailing-edge radius, m	0.0015	0.007




Fig. 3 Best designs for generations 1, 25, and 30 when using penalty function only and without minimum thickness distribution constraint.




Fig. 5 Best designs for generations 1 and 30 when using penalty function and thickness constraint.

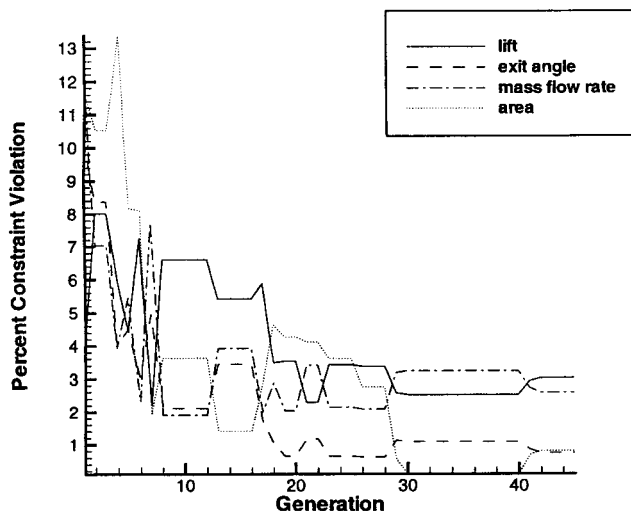


Fig. 4 Violation of constraints when using penalty function only and without minimum thickness distribution constraint.

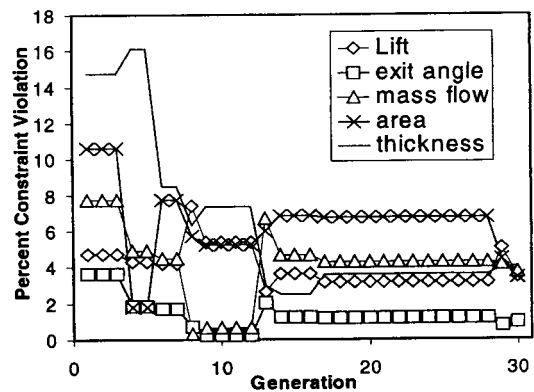


Fig. 6 Violation of constraints when using penalty function and minimum thickness distribution constraint.

The second example optimization run had the objective to demonstrate the difference that enforcement of the minimum allowable thickness distribution can make. The genetic optimizer was run for 30 generations with a population of 20 (thus $30 \times 20 = 600$ flowfield analysis). The calculation consumed 70 hours of CPU time on a 350 MHz Pentium II based personal computer. The best airfoil designs from generation 1 and generation 30 are shown in Fig. 5. The final airfoil shape in this example is significantly different from the final airfoil shape in the preceding example (Fig. 2). The best design from the GA generation 1 had a total pressure loss of 8200.0 Pa, which was after 30 generations reduced to 6850 Pa representing a reduction of 16.4%. The smaller reduction of the total pressure loss in this example is caused by the enforcement of the minimum thickness distribution constraint. Also, after 30 generations a 3.5% violation of the lift constraint, a 1% violation in the exit flow angle constraint, a 3.3% violation of the mass flow rate constraint, a 3.1% violation of the airfoil cross-section area constraint, and a 4% violation in the thickness distribution constraint were achieved (Fig. 6).

B. Gene Correction Sequential Quadratic Programming Method for Constraints

In this example optimization run a gene correction method based on sequential quadratic programming (SQP)²⁰ was used to enforce the cross-sectional area and thickness distribution constraints while penalty terms were used to enforce the lift, mass flow rate, and

average exit angle constraints with a penalty constant of $4.0E6$. The SQP algorithm was used to minimize the thickness distribution error with the specified cross-sectional area of the airfoil as an equality constraint. This was done to every individual in the population for every generation before the flowfield analysis was performed. This procedure was implemented in the GA by calling the SQP algorithm after each individual was created immediately after the crossover operation was completed.

With this treatment of the equality constraints, essentially all of the designs became geometrically feasible before the expensive flowfield analysis was performed. This allows the GA to focus on satisfying the lift, mass flow rate, and exit angle constraints only. The combined genetic and SQP optimizer was then run for 30 generations with a population of 15 (thus $30 \times 15 = 300$ flowfield analysis) and consumed 50 hours of computing time on a 550-MHz AlphaPC workstation. The best airfoil designs from the GA generation 1 and generation 30 are shown in Fig. 7. The best design from generation 1 had a total pressure loss of 7800.0 Pa, which was reduced after 30 generations to 6546.0 Pa representing a reduction of 16.1%. Also, after 30 generations the design had a 5.8% violation of the lift constraint, a 1% violation in the exit flow angle constraint, a 1.9% violation of the mass flow rate constraint, a 0% violation of the area constraint, and a 0% violation in the thickness distribution constraint. Figure 8 shows the violation of the constraints for the best design of each generation. This method consistently produced higher fitness designs over the earlier method that enforced all constraints via a composite penalty function.




Fig. 7 Best designs for generations 1 and 30 when using penalty function with SQP and minimum thickness distribution constraint.




Fig. 9 Best designs for generations 1 and 11 when using penalty function with SQP, minimum thickness distribution constraint, and B-spline geometry perturbation.

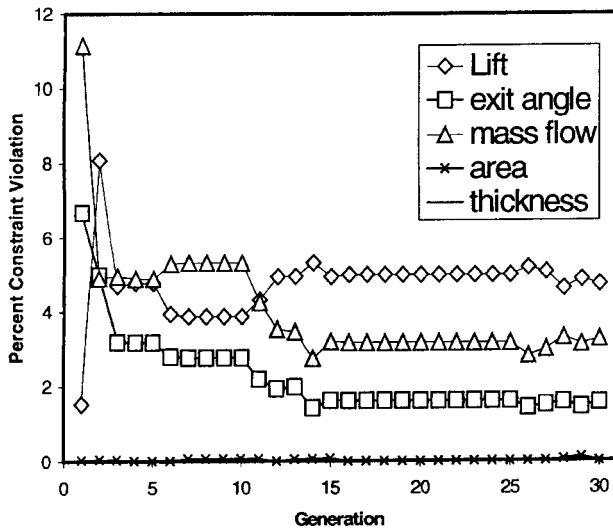


Fig. 8 Violation of constraints when using penalty function with SQP and minimum thickness distribution constraint.

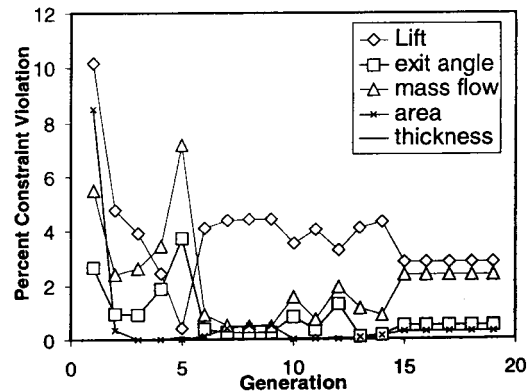


Fig. 10 Violation of constraints when using penalty function with SQP, minimum thickness distribution constraint, and B-spline geometry perturbation.

The computations required to find airfoil cross-section area and thickness distribution are very inexpensive. This proved to be more economical than trying to use SQP to satisfy all of the constraints simultaneously. Besides, for every new design generated by the GA the SQP code converged to a local minimum before finding a feasible solution. When applying it only to area and thickness distribution constraints, the SQP was able to find a design that meets those constraints most of the time. It was also more effective than treating all the constraints with just a penalty function.

The final example optimization run had the objective of demonstrating what benefits can be obtained from utilizing more flexible geometry parameterization. In this example run the eight B-spline control points (four on the pressure surface and four on the suction surface) were used to allow a limited range of perturbations of the airfoil shape in addition to those created by the basic nine parameters. This means that there were 17 design variables in this test case where each variable was represented by an eight-bit string. Like in the preceding example run, penalty formulation was used to enforce mass flow rate, lift, and exit flow angle. The SQP algorithm was used to enforce airfoil cross-sectional area and minimum allowable thickness distribution. The combined micro-GA/SQP optimizer was then run for 11 generations with a population of 20 (thus $11 \times 20 = 220$

flowfield analysis) and consumed 33 hours of computing time on a 550-MHz AlphaPC workstation. The best airfoil designs from the GA generation 1 and generation 11 are shown in Fig. 9. The best design from generation 1 had a total pressure loss of 7800.0 Pa, which was reduced after 11 generations to 6546.0 Pa representing a reduction of 16.1%. After 11 generations the design had a 4.0% violation of the lift constraint, a 0.3% violation in the exit flow angle constraint, a 0.7% violation of the mass flow rate constraint, a 0% violation of the area constraint, and a 0% violation in the thickness distribution constraint (Fig. 10). This was a clear improvement over the case when airfoil geometry was parameterized without the addition of B-spline perturbations.

Because the limits on the B-spline perturbations are user-specified, the optimization in this example run was continued beyond the generation 11 by doubling the user-specified limit on each of the eight B-spline control point motions in the direction outward of the airfoil. This additional geometry flexibility resulted in an even better average satisfaction of all constraints (Fig. 10) after the generation 11. The total pressure loss, which was reducing nonmonotonically because of the conflicting constraints, in this optimization case was also smaller than after generation 11 (Fig. 11). The computed surface-pressure distributions for the best airfoils after generation 1 and generation 11 show that the optimized airfoil is thicker and more loaded in the leading-edge region (Fig. 12) although its curvature in this region is decreased.




Fig. 11 Convergence history of total pressure loss for best designs when using penalty function with SQP, minimum thickness distribution constraint, and B-spline geometry perturbation.

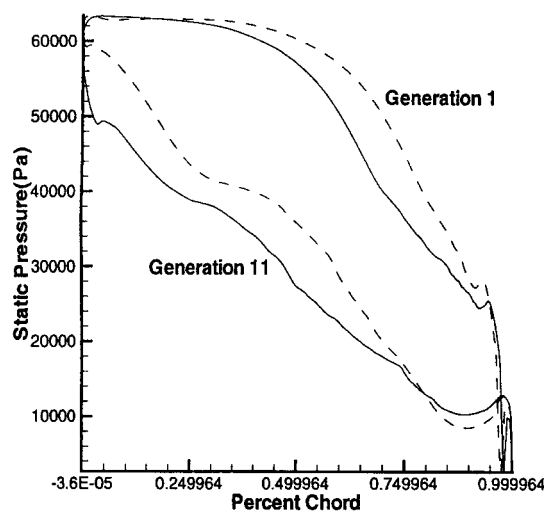


Fig. 12 Surface static-pressure distribution for best design of generation 1 and generation 11 when using penalty function with SQP, minimum thickness distribution constraint, and B-spline geometry perturbation.

VI. Conclusions

We have developed and demonstrated an aerodynamic shape-design optimization algorithm that utilizes a Navier–Stokes flow-field analysis code and a constrained microgenetic optimizer and a sequential-quadratic-programming optimizer for enforcement of certain constraints. The design optimization algorithm was applied in a redesign of an existing two-dimensional cascade of supersonic exit turbine airfoils with the objective to minimize the total pressure loss across the cascade. The airfoil geometry was discretized with either nine or 17 parameters, and a clustered unstructured grid was created for every airfoil cascade configuration generated during the optimization process. The following equality constraints were iteratively enforced: specified aerodynamic lift force, mass flow rate, exit flow angle, and airfoil cross-section area. In addition, axial chord and the gap-to-axial chord were kept fixed, while enforcing an inequality constraint that the airfoil thickness should always be greater or equal than the specified minimum allowable thickness distribution. The optimization code proved to be very robust because it found the narrow feasible domain and converged to a minimum that satisfied all of the constraints within the tolerances specified. The design optimization process results in a significantly more efficient aerodynamic shape of the turbomachinery cascade. The optimization is feasible on a single-processor workstation, it requires no changes to the existing flowfield analysis code and can be operated even by

a semiskilled designer. Improving the results further may require larger population sizes, more flexibility in geometric shape parameterization, and longer bit strings for the design variable encoding. Use of variable penalty terms that vary with generation can also prove to be beneficial. The entire design optimization methodology is extendable to three-dimensional configurations.

Acknowledgments

The authors would like to express their gratitude for the National Science Foundation Grants DMI-9522854 and DMI-9700040 monitored by George A. Hazelrigg, the NASA John H. Glenn Research Center Grant NAG3-1995 facilitated by John K. Lytle and Herbert J. Gladden and monitored by Kestutis Civinskas, and for the Lockheed Martin Skunk Works Grant monitored by Thomas Oatway.

References

- Emanuel, G., *Analytical Fluid Dynamics*, 2nd ed., CRC Press, Boca Raton, FL, 1999, pp. 45–47.
- Dulikravich, G. S., "Aerodynamic Shape Design and Optimization: Status and Trends," *Journal of Aircraft*, Vol. 29, No. 6, 1992, pp. 1020–1026.
- Dulikravich, G. S., "Design and Optimization Tools Development," *New Design Concepts for High Speed Air Transport*, edited by H. Sobieczky, Springer-Verlag, New York, 1997, pp. 159–236.
- Fujii, K., and Dulikravich, G. S. (ed.), *Recent Development of Aerodynamic Design Methodologies—Inverse Design and Optimization*, Vieweg Series on Notes on Numerical Fluid Mechanics, Vol. 68, Vieweg and Sohn, Braunschweig, Germany, 1999.
- Dulikravich, G. S., "A Criteria for Surface Pressure Specification in Aerodynamic Shape Design," AIAA Paper 90-0124, Jan. 1990.
- Dulikravich, G. S., "Shape Inverse Design and Optimization for Three-Dimensional Aerodynamics," AIAA Paper 95-0695, Jan. 1995.
- Dennis, B. H., Egorov, I. N., Han, Z.-X., Dulikravich, G. S., and Poloni, C., "Multi-Objective Optimization of Turbomachinery Cascades for Minimum Loss, Maximum Loading, and Maximum Gap-to-Chord Ratio," American Society of Mechanical Engineers, Paper 2000-GT-0253, June 2000; also *International Journal of Turbo & Jet-Engines* (to be published).
- Pritchard, L. J., "An Eleven Parameter Axial Turbine Aerofoil Geometry Model," American Society of Mechanical Engineers, Paper 85-GT-219, June 1985.
- Dulikravich, G. S., and Martin, T. J., "Aero-Thermal Analysis and Optimization of Internally Cooled Turbine Airfoils," *XIII International Symposium on Airbreathing Engines, ISABE*, Chattanooga, TN, Sept. 7–12, 1997, edited by F. S. Billig, Vol. 2, 1997, pp. 1232–1250.
- Dulikravich, G. S., Martin, T. J., Dennis, B. H., Lee, E.-S., and Han, Z.-X., "Aero-Thermo-Structural Design Optimization of Cooled Turbine Blades," AGARD—AVT Propulsion and Power Systems Symposium on Design Principles and Methods for Aircraft Gas Turbine Engines, NATO-RTO-MP-8 AC/323(AVT)TP/9, Ch. 35, NATO-AGARD, May 1998.
- Trigg, M. A., Tubby, G. R., and Sheard, A. G., "Automatic Genetic Optimization Approach to 2D Blade Profile Design for Steam Turbines," American Society of Mechanical Engineers, Paper 97-GT-392, June 1997.
- Han, Z.-X., and Liu, Z.-J., "Numerical Calculation of Flow Fields with Multi-Unstructured Grids," *Journal of Power*, Vol. 13, No. 2, 1998, pp. 149–152 (in Chinese).
- Shewchuk, R. J., "Triangle: Engineering a 2D Quality Mesh Generator and Delaunay Triangulator," *First Workshop on Applied Computational Geometry*, ACM, Philadelphia, PA, 1996, pp. 124–133.
- Frink, N. T., "Recent Progress Toward a Three-Dimensional Unstructured Navier–Stokes Flow Solver," AIAA Paper 94-0061, Jan. 1994.
- Sieverding, C. H., "Test Case E/CA-8 Transonic Turbine Cascade," AGARD AR 275, *Test Cases for Computation of Internal Flows in Aero Engine Components*, edited by L. Fotner, NATO-AGARD, Brussels, 1990.
- Goldberg, D. E., *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison Wesley Longman, Reading, MA, 1989.
- Carroll, D. L., "Chemical Laser Modeling with Genetic Algorithms," *AIAA Journal*, Vol. 34, No. 2, 1996, pp. 338–346.
- Foster, N. F., and Dulikravich, G. S., "Three-Dimensional Aerodynamic Shape Optimization Using Genetic and Gradient Search Algorithms," *Journal of Spacecraft and Rockets*, Vol. 34, No. 1, 1997, pp. 36–42.
- Krishnakumar, K., "Micro-Genetic Algorithms for Stationary and Non-Stationary Function Optimization," *SPIE: Intelligent Control and Adaptive Systems*, Vol. 1196, Philadelphia, PA, 1989.
- Lawrence, C. T., Zhou, J. L., and Tits, A., "User's Guide for CFSQP Version 2.5: A C Code for Solving (Large Scale) Constrained Nonlinear (Min-max) Optimization Problems, Generating Iterates Satisfying All Inequality Constraints," Inst. for Systems Research, Univ. of Maryland, Technical Rept. TR-94-16r1, College Park, MD, 1997.