

# Comparison of Several Optimization Strategies for Robust Turbine Blade Design

Andy J. Keane\*

*University of Southampton, Southampton, England SO17 1BJ, United Kingdom*

DOI: 10.2514/1.38673

**This paper addresses the problem of turbine blade shape optimization in the presence of geometric uncertainties. Several strategies are tested and compared on a two-dimensional compressor blade optimization process for which performance is assessed using a commercial Reynolds-averaged Navier–Stokes computational fluid dynamics code. In each case, a range of shape errors are considered that attempt to simulate foreign object damage, erosion damage, and manufacturing errors. These lead to stochastic performance measures that, in turn, are considered in a multi-objective optimization framework. Because of the long run times associated with Reynolds-averaged Navier–Stokes codes, use is also made of surrogate or response-surface-based optimization methods to speed up the search processes. The paper shows that a range of techniques can be used to tackle this problem, but that no one method is clearly best overall. The practitioner is therefore cautioned against favoring a single approach for such design problems. Further research may help clarify these issues.**

## I. Introduction

THE use of optimization methods to design aerodynamic sections is now completely routine in many companies and research institutes. Typically a computer aided design (or other geometry design) system, meshing tool, computational fluid dynamics (CFD) solver, and postprocessor are used, coupled to a range of different optimization tools to give a design search and optimization capability. More recently, attention has turned to the issue of design robustness, because designs optimized purely for nominal performance may suffer from significantly degraded performance if they are subject to slight errors in manufacture, wear or damage in operation, or operation in conditions different to those at which they were optimized. This leads naturally to a multi-objective problem for which designers seek to improve the mean or nominal performance of a design while guaranteeing that falloff away from nominal conditions is strictly controlled. At the same time, the significant run times that can be encountered in using high-fidelity Reynolds-averaged Navier–Stokes (RANS) CFD codes for design optimization have been addressed recently by the use of surrogate or metamodel schemes to construct so called “response surfaces” that can help speed up search. In this paper we focus on both sets of approaches using industrial CFD tools and a large parallel computing cluster running the 64-bit Windows Server 2008 compute cluster edition operating system. The focus of this paper is on the two-dimensional optimization of a gas-turbine compressor blade section subject to damage in service and uncertainty in manufacture.

It is now commonplace to use design optimization methods to change the shape of aerodynamic sections when designing new gas turbines; see Keane and Nair [1] for an introduction to this approach. More recently, attention has focused on both robust multi-objective design (see for, example, Welch et al. [2] and Kumar et al. [3]) and surrogate-assisted design optimization (see, for example, Jones et al. [4] and Keane [5]). Approaches combining both these methods have also appeared (Keane [6] and Emmerich et al. [7]). Here we focus on the use of a series of shape-manipulation tools to explore the effects of foreign object damage, erosion, and uncertainty in manufacture

while optimizing mean design performance and variance in a multi-objective framework. We compare and contrast four approaches. First, we simply optimize a baseline geometry for nominal performance using direct [8] and surrogate-based search [4]. Then, we employ a noisy phenotype [9] approach to allow for possible sensitivity in the mean performance but do not explicitly seek a tradeoff between mean performance and robustness. Third, we use a multi-objective evolutionary search [10] to deliberately construct a trade surface (Pareto front) adopting both direct and surrogate-assisted approaches. Finally, we use a surrogate-based multi-objective expected improvement formulation [6] to construct the tradeoff using single-objective search methods. Together these schemes represent the current state of the art of industrial practice in this area; by applying them with a set of powerful industrial analysis codes, it is possible to gain some insight into which approaches show the most promise for use in practical design work. It is of course true that more sophisticated direct stochastic solver methods are being researched in universities but these methods are, as yet, not widely adopted.

## II. Geometry Modification

To carry out any form of automated aerodynamic section optimization, some form of geometry modification is required. Here we use a modified form of the parametric design and rapid meshing (PADRAM) code (Shahpar and Lapworth [11]) that has been described elsewhere (Kumar et al. [3]). In this process, a baseline airfoil section shape is altered by the addition of a series of Hicks–Henne functions in various ways (Hicks and Henne [12]). In the present approach, this scheme is used to make four sets of changes to the baseline airfoil that allow for overall design improvement and the modelling of manufacturing uncertainty, foreign object damage, and flank erosion. Figures 1a–1f illustrate the baseline airfoil section and each of these modification processes in turn. Note that all the changes are controlled such that the trailing-edge form and angle are kept fixed as would be needed to ensure unchanged inflow angles for the next stage in the compressor. Here, 10 Hicks–Henne functions are used for the design shape changes, distributed around the section, and this permits very significant alterations in geometry. In all the tables that follow, the normalized amplitude changes of these design variables are provided. These are defined in such a way that a value of 0.5 generates the base geometry, with values less than this removing material from the section and those greater adding to it; full details may be found in Kumar [13]. The other three sets of changes work in the same way but make much smaller shifts:

1) The manufacturing error changes affect the entire section (again using 10 Hicks–Henne functions distributed around the section, but

Received 2 June 2008; revision received 8 July 2009; accepted for publication 8 July 2009. Copyright © 2009 by Dr. Andy J. Keane. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0748-4658/09 and \$10.00 in correspondence with the CCC.

\*Professor of Computational Engineering, School of Engineering Sciences, Highfield; ajk@soton.ac.uk.