

Multi-objective Optimisation of Turbomachinery Blades Using Tabu Search

Timoleon Kipouros¹, Daniel Jaeggi¹, Bill Dawes², Geoff Parks¹, and Mark Savill³

¹ Engineering Design Centre, Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom

² Computational Fluid Dynamics Laboratory, Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, United Kingdom

³ Computational Aerodynamic Design Group, Department of Aerospace Sciences, Cranfield University, Cranfield MK43 0AL, United Kingdom

Abstract. This paper describes the application of a new multi-objective integrated turbomachinery blade design optimisation system. The system combines an existing geometry parameterisation scheme, a well-established CFD package and a novel multi-objective variant of the Tabu Search optimisation algorithm. Two case studies, in which the flow characteristics most important to the overall performance of turbomachinery blades are optimised, are investigated. Results are presented and compared with a previous (single-objective) investigation of the problem.

1 Introduction

The optimisation of airfoil designs is a challenging, computationally expensive, highly constrained, non-linear problem. As with most real-world problems, there are multiple (usually conflicting) performance metrics that an engineer might seek to improve in optimising, for example, the design of turbomachinery blades, wings or other aerodynamic surfaces. This suggests a multi-objective approach, a notion that is reinforced by the recognition that any consideration of robustness – the retention of performance over a range of operating conditions, in the face of geometry changes (e.g. through creep) etc. – must also inevitably entail multiple objectives.

Despite this obvious motivation, multi-objective aerodynamic optimisation seems to have been somewhat overlooked. However, two recent studies in particular have embraced multi-objective optimisation and show the possible benefits compared to single-objective optimisation with a composite objective function.

Gaiddon *et al.* [9] perform multi-objective optimisation on a supersonic missile inlet. They compare a number of optimisation algorithms using both composite and multiple objective functions, and conclude that “performing real multi-objective optimization and finding a Pareto front is the only effective way to find a set of designs satisfying several performance criteria in an industrial context”.

Nemec *et al.* [17] perform multi-objective optimisation on both a single and a multi-element 2-D aerofoil. Their integrated approach combines a Newton-Krylov adjoint CFD code, a b-splines-based parameterisation scheme and both

a gradient-based optimiser and a Genetic Algorithm (GA). They obtain good results on some simple test problems.

The multi-objective integrated design system used in the present work has been developed and described by Kipouros *et al.* [16] building on the single-objective integrated design optimisation system (BOS3D) developed by Harvey [11] and described by Dawes *et al.* [6]. The system combines an existing, efficient and flexible geometry parameterisation scheme, a well-established CFD package and a novel multi-objective variant of the Tabu Search (TS) optimisation algorithm for continuous problems [13]. The system can readily be run on parallel computers, which can substantially reduce wall-clock run times – a significant benefit when tackling computationally demanding design problems.

In previous work [16] the performance of this system has been investigated considering a compressor blade design test case. The effectiveness of the multi-objective optimisation procedure was verified and the expected trade-offs between the chosen objectives confirmed. In the work presented in this paper we use our system to tackle more realistic turbomachinery design test cases, taking advantage of the greater computational power offered by exploiting its parallel processing capabilities.

2 Description of the Integrated System

Fig. 1 presents a flow diagram showing the stages of the process executed by our integrated multi-objective turbomachinery blade design optimisation system. The first stage is the parameterisation of the initial blade design, input through an initial CAD geometry together with boundary conditions for the flow solution. The geometry is parameterised using a Partial Differential Equation approach [3], giving a compact but flexible representation of the design, in a design vector comprising 26 variables. This design vector is the input to the main loop of the design system, which consists of the flow simulation and optimisation processes. On receipt of a new design vector, a computational mesh is automatically generated from the geometry specification, and then a detailed CFD analysis (blade to blade) is performed. The mesh is a 3D structured grid consisting of $21 \times 87 \times 23$ nodes in each direction. The flow simulation is performed by a CFD code solving the 3D Navier-Stokes equations, and this routine returns all the necessary metrics that describe the flow around the blade [5]. Based on this evaluation, the optimisation routine generates a new design vector that is meshed and evaluated, and this process continues until a stopping criterion is met.

At the end of the optimisation process, the best design vectors identified and their associated flow solutions are converted into a single file, in the final stage of representation. This stage is accomplished by using Non-Uniform Rational B-Splines (NURBS) [18]. The optimal geometries can then be examined in detail through, for instance, contour plots.

At the heart of our system is a multi-objective variant [13] of the well-established Tabu Search optimisation algorithm [10]. There has been substantial

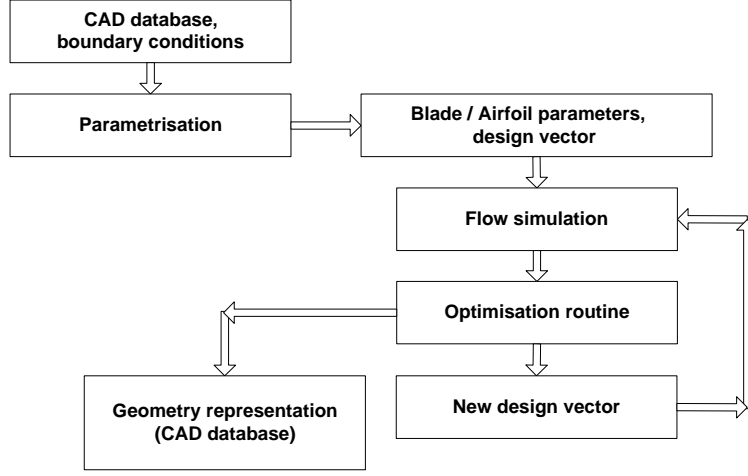


Fig. 1. The structure of the integrated design optimisation system

recent interest in developing multi-objective optimisation algorithms – the vast majority of which concerns multi-objective GAs [7]. GAs have been applied to aerodynamic design problems, but, although able to locate optimal designs, the method proved to be sensitive to constraint handling schemes and required significantly more computational time than a gradient-based method [1]. This is not surprising. The highly constrained nature of most aerodynamic design optimisation problems suggests that algorithms, like GAs, that routinely make large changes to solutions may experience difficulties in trying to negotiate feasible space on such problems, and that optimisation algorithms which progress by making small changes are likely to be more effective. Harvey attributes the local search paradigm at the heart of TS as being one of the reasons for its effectiveness in his work [11]. The execution of a local search at each iteration of the algorithm also offers the potential benefit of being able to estimate the robustness of the current design to variations in geometry without the need for additional flow solutions [2]. For these reasons we have opted to use a multi-objective TS variant in our work.

Our multi-objective TS variant takes as its starting point the single-objective TS implementation of Connor and Tilley [4]. This uses a Hooke and Jeeves local search algorithm (designed for continuous optimisation problems) [12] coupled with short, medium and long term memories to implement search diversification and intensification, as prescribed by Glover and Laguna [10]. This algorithm and analysis of its performance on benchmark constrained optimisation problems from the literature are described in detail in a companion paper [14].

3 Case Studies

The most important flow parameters that affect the performance of a turbomachinery blade are:

- flow separation (blockage),
- losses (any flow feature that reduces the efficiency of the turbomachine), and
- deviation in flow turning.

Here we seek to find, starting from a gas turbine compressor guide vane specification, a blade geometry that efficiently gives a good pressure rise at a particular flow coefficient. Thus, a global performance measure of a given blade geometry is needed for the optimisation process. Efficiency is only one of a number of possible design objectives when undertaking detailed aerodynamic shape optimisation. A good design must also respect mechanical and manufacturing constraints while achieving the required aerodynamic performance (with respect to flow turning, separation and good off-design performance etc.).

The design optimisation of compressor blade geometries has previously been studied by Harvey [11] from a single-objective perspective. In our multi-objective test cases we retain Harvey’s objective function (equation 1) as an essential 1D (throughflow) measure of blade performance. This is a normalised function including penalty function terms for specific flow characteristic and geometry constraints. This objective function considers the span-averaged blockage for a given mass flow rate:

$$f_1 = \frac{B}{B_0} + 250\left(1 - \frac{\dot{m}}{\dot{m}_0}\right)^2 + 0.4max^2\left(0, 1 - \frac{R_{LE}}{R_{LE,0}}\right) + 500max^2\left(0, 1 - \frac{\Delta\theta}{\Delta\theta_0}\right) + 0.5max^2\left(0, 1 - \frac{C}{0.015}\right) \quad (1)$$

In equation 1, B represents the blockage, the extent to which viscous forces restrict the effective flow area in a blade passage, which is probably the most critical quantity in high-speed compressor design. Then, \dot{m} is the mass flow rate; R_{LE} is the minimum radius of the leading edge of the blade; $\Delta\theta$ is the mass-averaged flow turning; and C measures the tip clearance of the blade. The zero subscripts identify the equivalent quantities for the datum blade geometry, the initial design in the optimisation – a real compressor blade design shown in Fig. 2. Harvey established suitable values for the weightings for each of the penalty terms through extensive testing [11].

In highly loaded compressors, the flow tends to separate from the blade under conditions of low mass flow. Flow separation acts as a blockage in the flow path, which limits pressure recovery.

The mass flow associated with the design should be equality constrained for two reasons. First, if it was not, the inlet dynamic head from the rotor would vary, which is not modelled by the boundary conditions. Second, if the axial




Fig. 2. The datum blade geometry showing its axial velocity distribution

velocity drops, then the inlet static pressure must be higher (since inlet pressure and flow angles are prescribed), so that the static pressure rise across the stator will be lower (outlet pressure is fixed), and the blade row will not be an effective diffuser. Equally, it is important that, if the mass flow is fixed, the flow turning in the stage should not be reduced during optimisation, otherwise the static pressure recovery will not be sufficient. Therefore, control of the flow turning is achieved by treating it as a penalty term.

In addition, there are two terms in equation 1 describing the geometrical constraints on the blade. The first limits the sharpness of the blade's leading edge, while the second allows a weighted penalty factor to trade off aerodynamic performance against mechanical proximity. The objective function value is penalised when the blade design has less than 1.5 cm clearance. Both these penalty terms reflect a concern for robust aerodynamic performance from the design, since these geometric characteristics are closely related to the off-design performance of the blade:

1. A sharp leading edge produces a high velocity profile at the front of the blade and, in addition, such a design may well have a flat section at its front. Such a geometry results in a velocity peak at off-design conditions, which means that the velocity distribution produced has a high probability of early transition, laminar or turbulent separation, and consequently poor off-design performance.
2. The tip clearance is responsible for the deviation created on the end walls. In particular, the secondary flows are known to create over-turning very close to the end walls and a region of under-turning some distance away from them. The effects of this deviation on the exit velocity profile and on the inlet incidence angles to the next blade row can be very large and can substantially increase the incidence losses on the next blade row.

Harvey [11] found that it was necessary to use a penalty function approach with these constraints in order to successfully navigate the highly constrained,

nonlinear search space characteristic of aerodynamic design optimisation problems. Other constraints, such as those on the geometric feasibility of blade designs and on their operational feasibility (a design which produces unsteady flow patterns is not acceptable), are handled as hard constraints – designs violating them are not accepted.

3.1 First Multi-objective Test Case

In our first test case, we introduce a second objective *Brms* (the RMS variation in blockage), which provides an additional representation of the spatial variation, and hence homogeneity, of the blockage – effectively a 2D measure of blade performance (equation 2), subject to the same set of penalty function constraints:

$$f_2 = \frac{Brms}{Brms_0} + 250 \left(1 - \frac{\dot{m}}{\dot{m}_0}\right)^2 + 0.4max^2 \left(0, 1 - \frac{R_{LE}}{R_{LE,0}}\right) + 500max^2 \left(0, 1 - \frac{\Delta\theta}{\Delta\theta_0}\right) + 0.5max^2 \left(0, 1 - \frac{C}{0.015}\right) \quad (2)$$

The same penalty terms are included in the formulation of each objective in order to ensure that the final Pareto front identified contains only designs that are (near-)feasible.

3.2 Second Multi-objective Test Case

In our second test case, we introduce an objective function which quantifies losses associated with the design. This is defined in terms of the rate of entropy generation [8], subject again to the same set of penalty function constraints:

$$f_3 = \frac{EntropyRate}{EntropyRate_0} + 250 \left(1 - \frac{\dot{m}}{\dot{m}_0}\right)^2 + 0.4max^2 \left(0, 1 - \frac{R_{LE}}{R_{LE,0}}\right) + 500max^2 \left(0, 1 - \frac{\Delta\theta}{\Delta\theta_0}\right) + 0.5max^2 \left(0, 1 - \frac{C}{0.015}\right) \quad (3)$$

In a turbomachine the isentropic efficiency is defined as the ratio of the actual work to the isentropic work for a work-producing device (such as a turbine) and the ratio of the isentropic work to the actual work for a work-absorbing device (such as a compressor). The only factors that change this efficiency are departures from isentropic flow. These may arise due either to heat transfer or to thermodynamic irreversibility. For most turbomachines the flow is close to adiabatic (no heat transfer) and so only entropy creation by irreversibilities contributes significantly to the loss of efficiency, which means that entropy generation rate is the only rational measure of overall loss in an adiabatic machine.

Any irreversible flow process creates entropy and thus reduces the isentropic efficiency. The sources of loss in a turbomachine can be categorised as profile loss, secondary (or end wall) loss, and tip leakage loss. In many machines the three are comparable in magnitude, each accounting for about 1/3 of the total loss.

4 Results

4.1 First Test Case

The optimisation was initiated from the datum geometry shown in Fig. 2. Fig. 3 shows the progress in objective-space of the search performed by our multi-objective TS algorithm over the 1350 iterations of the run, using the control parameters specified in Table 1. See [14] for a detailed explanation of these parameters. The values of these control parameters were chosen based on experience, but it should be noted that studies in [14] show that the algorithm’s performance is relatively insensitive to the control parameter settings.

Table 1. Tabu Search Parameter Settings

Parameter	Value	Description
<i>intensify</i>	25	Intensify search after <i>intensify</i> iterations without adding to the Medium Term Memory (MTM)
<i>diversify</i>	75	Diversify search after <i>diversify</i> iterations without adding to the MTM
<i>reduce</i>	95	Reduce step size and restart after <i>reduce</i> iterations without adding to the MTM
<i>n_stm</i>	15	Short Term Memory size – the last <i>n_stm</i> visited points are tabu
<i>n_regions</i>	4	In the Long Term Memory each variable is divided into <i>n_regions</i> to determine which regions of the search space have been under-explored
<i>SS</i>	1%	Initial step size as percentage of variable range
<i>SSRF</i>	0.5	Factor by which step sizes are reduced on restarting

The optimisation was run on an 8-node parallel PC cluster of 2.8GHz Pentium 4 machines in order to reduce wall-clock run times by exploiting our system’s parallel capabilities. The CFD flow solution required for evaluation of a single candidate design takes 3 minutes on a single node of the cluster, and up to 52 CFD evaluations (on average 32 evaluations) are required at each iteration of the optimisation algorithm. Thus, on our cluster a 1350 iteration run takes 270 hours (just over 11 days).

Fig. 4 shows the set of Pareto-optimal solutions found during the search. The geometry shown in Fig. 5 represents a compromise design from the middle of the Pareto front, which is clearly quite different to the datum design. The large changes made to the blade geometry during the search demonstrate the




Fig. 3. The optimisation search pattern for the first blading test case

flexibility of the geometry management system used. A high twist along the span is the main characteristic of this blade. Reassuringly, this optimised geometry has a similar leading edge (LE) profile to the single-objective optimised design found by Harvey, shown in Fig. 6 [11].

It can be seen that the performances of the optimised geometries lying at the low blockage end of the Pareto front (Fig. 4) have matched or slightly exceeded that of the optimised geometry found by Harvey (Fig. 6). This has been achieved in an optimisation run that is actually shorter in terms of the number of CFD flow solutions required than that reported by Harvey [11] (using the same initial design), even though we are tackling a multi-objective optimisation problem. In effect the additional information provided to the designer by revealing the trade-off between the main objective (blockage) and the secondary objective (the RMS variation in blockage) costs nothing because the quality of the best designs found with respect to the primary objective have not been compromised at all by the switch to multi-objective optimisation.

4.2 Second Test Case

Optimised geometries found for our second test case are presented in Figs. 7, 8, 9 and 10, which show respectively the lowest blockage design, lowest entropy generation design and two compromise geometries from the trade-off surface (Fig. 11). These blades are again quite different to the datum geometry used to




Fig. 4. The Pareto front found for the first test case

initiate the optimisation, and in addition there are significant differences between them. The LE of these geometries is similar, but differs from the LE of the single-objective optimised blade (Fig. 6). Furthermore, there are considerable differences in the trailing edge (TE) between the multi-objective and single-objective designs.

It is worth remarking that compromise geometry A (Fig. 9) displays geometrical characteristics from both the lowest blockage (Fig. 7) and the lowest entropy generation (Fig. 8) designs. However, the big difference between these blades is in the tip profiles. For the lowest blockage design (Fig. 7) and compromise design B (Fig. 10), for which the blockage is almost as low, there is a rapid change in the tip camber, which results in a thick profile, whereas for the geometries in Figs. 8 and 9 there is a smooth change in the camber of the tip profile. This geometrical characteristic is shared with the blade in Fig. 6.

As regards aerodynamic performance, all the blades have good, i.e. smoothly varying, axial velocity distributions. However, there is noticeable flow separation along the TE of the blade optimised for lowest entropy generation.

The Pareto front shows that significant performance improvements are achievable. For instance, compromise design B (Fig. 10) reduces the blockage significantly compared to the initial datum design (Fig. 2) with very little increase in the rate of entropy generation. The lowest blockage design (Fig. 7) performs




Fig. 5. Test case 1: The optimised geometry for a compromise point on the Pareto front (Fig. 4)




Fig. 6. The optimised geometry for the single-objective test case [11]

comparably to Harvey’s blockage-optimised blade (Fig. 6) but has a significantly lower entropy generation rate.

Interestingly the Pareto front also exhibits a sharp “elbow” around compromise design B. Relative to this design, it is possible to reduce the blockage objective but only at the cost of quite large increases in entropy generation. It is also possible to reduce the rate of entropy generation but only at the cost of significantly increased blockage. In this case the designer might find it quite straightforward to select a good compromise design.

5 Conclusions and Future Work

The foregoing test cases demonstrate that our multi-objective integrated turbo-machinery design optimisation system can successfully tackle realistic real-world




Fig. 7. Test case 2: The optimised geometry for lowest blockage (Fig. 11)

problems, negotiating the highly constrained, nonlinear search space, and presenting the designer with a range of designs showing the trade-offs between the objectives under consideration.

In both test cases the performance of the designs found match or exceed the performance of the optimised blade identified in an earlier single-objective study. The computational effort required to solve these multi-objective problems was no more than that required to solve the single-objective problem and, in addition to equally good or better designs, the designer is also presented with helpful information about the trade-offs between the objectives of interest. This demonstrates very clearly the value of multi-objective optimisation.

The factors influencing efficiency of turbomachinery blades and the trade-offs between them are extremely complex and therefore this area needs further investigation. The next steps will be to define additional loss objectives for the design problem examined above. Thus, new objective functions will evaluate individually the profile losses and the secondary losses in order to improve understanding of the trade-offs between them in design. These investigations will require the tackling of three- and four-objective problems and will therefore also allow us to test the effectiveness of our multi-objective TS variant on these higher dimension problems.

To improve further the effectiveness of our multi-objective TS variant a more sophisticated system for selecting the design variables to be modified in search, based on the work of Kellar [15], will be developed. This will seek to identify the variables that have the greatest impact on the performance of the current design and prioritise them at each local search iteration. It is hoped that this will improve the wall-clock performance of the system substantially.




Fig. 8. Test case 2: The optimised geometry for lowest entropy generation (Fig. 11)

Acknowledgements

The first author gratefully acknowledges the support of the Embiricos Foundation and the Cambridge European Trust. The second author acknowledges the support of the UK Engineering and Physical Sciences Research Council (EPSRC) under grant number GR/R64100/01.

References

1. Aly, S., Ogot, M., Pelz, R.: Stochastic Approach to Optimal Aerodynamic Shape Design. *Journal of Aircraft*. **33** (1996) 956–961
2. Asselin-Miller, C. S.: Robust 2D-Aerofoil Design: Proof of Concept. 4th Year Project Report. University of Cambridge, Department of Engineering (2003)
3. Bloor, M. I. G., Wilson, M. J.: Efficient Parameterisation of Generic Aircraft Geometry. *Journal of Aircraft*. **32** (1995) 1269–1275
4. Connor, A. M., Tilley, D. G.: A Tabu Search Method for the Optimisation of Fluid Power Circuits. *IMEchE Journal of Systems and Control*. **212** (1998) 373–381
5. Dawes, W. N.: Development of a 3D Navier-Stokes solver for application to all types of turbomachinery. ASME Conference Paper 88-GT-70. ASME Gas Turbine Conference, Amsterdam. (1988)
6. Dawes, W. N., Kellar, W. P., Harvey, S. A., Dhanasekaran, P. C., Savill, A. M., Cant, R. S.: Managing the Geometry is Limiting the Ability of CFD to Manage the Flow. 33rd AIAA Fluid Dynamics Conference, Orlando, Florida. AIAA-2003-3732. (2003)
7. Deb K., Multi-Objective Optimization using Evolutionary Algorithms. John Wiley & Sons Ltd., Chichester UK (2001)
8. Denton, J. D.: Loss Mechanisms in Turbomachines. *Journal of Turbomachinery*. **115** (1993) 621–656
9. Gaiddon, A., Knight, D. D., Poloni, C.: Multicriteria Design Optimisation of a Supersonic Inlet Based upon Global Missile Performance. *Journal of Propulsion and Power*. **20** (2004) 542–558




Fig. 9. Test case 2: The optimised geometry for compromise point A on the Pareto front (Fig. 11)

10. Glover, F., Laguna, M.: Tabu Search. Kluwer Academic Publishers, Boston MA (1997)
11. Harvey, S. A.: The Design Optimisation of Turbomachinery Blade Rows. Ph.D. Dissertation. University of Cambridge (2002)
12. Hooke, R., Jeeves, T.: Direct Search Solution of Numerical and Statistical Problems. *Journal of the ACM*. **8** (1961) 212–229
13. Jaeggi, D. M., Asselin-Miller, C. S., Parks, G. T., Kipouros, T., Bell, T., Clarkson, P. J.: Multi-objective Parallel Tabu Search. In: Yao, X., Burke, E., Lozano, J.-A., Smith, J., Merelo-Guervos, J., Bullinaria, J., Rowe, J., Tino, P., Kaban, A., Schwefel, H.-P. (eds.): *Parallel Problem Solving from Nature – PPSN VIII*. Lecture Notes in Computer Science, Vol. 3242. Springer-Verlag, Berlin (2004) 732–741
14. Jaeggi, D. M., Parks, G. T., Kipouros, T., Clarkson, P. J.: A Multi-objective Tabu Search Algorithm for Constrained Optimisation. In: *3rd Int. Conf. Evolutionary Multi-Criterion Optimization*. Lecture Notes in Computer Science, Springer-Verlag, Berlin (2005)
15. Kellar, W. N.: Geometry Modeling in Computational Fluid Dynamics and Design Optimisation. Ph.D. Dissertation. University of Cambridge (2002)
16. Kipouros, T., Parks, G. T., Savill, A. M., Jaeggi, D. M.: Multi-objective Aerodynamic Design Optimisation. In: Giannakoglou, K. C., Haase, W. (eds.): *ERCOFTAC Design Optimization: Methods and Applications Conference Proceedings*. On CD Rom (2004) Paper ERCODO2004_239
17. Nemec, M., Zingg, D. W., Pulliam, T. H.: Multipoint and Multi-objective Aerodynamic Shape Optimization. *AIAA Journal*. **42** (2004) 1057–1065
18. Rogers, D. F.: *An Introduction to NURBS With Historical Perspective*. Morgan Kaufmann Publishers, San Francisco CA (2000)




Fig. 10. Test case 2: The optimised geometry for compromise point B on the Pareto front (Fig. 11)




Fig. 11. The optimisation search pattern and the Pareto front for the second test case