The Future of Gas Turbine Technology
3^{rd} International Conference
11-12 October 2006, Brussels, Belgium

Paper ID Number S1 T1 / 2 Ghisu

MULTI-OBJECTIVE OPTIMISATION OF AERO-ENGINE COMPRESSORS

T. Ghisu, G. T. Parks, J. P. Jarrett and P. J. Clarkson
Engineering Design Centre, Department of Engineering, University of Cambridge
Trumpington Street, Cambridge CB2 1PZ, UK
Phone: +44 1223 760569, Fax: +44 1223 332662
Email: tg269@cam.ac.uk

April 26, 2007

ABSTRACT

The design of a new aero-engine compressor is a complex task: design objectives are almost always conflicting, the design space is large, nonlinear and highly constrained, and the effects of some geometrical changes can be difficult to predict.

Computational fluid dynamics (CFD) is now widely used in real-world applications and especially in the design of turbomachinery. However, the large design space and the time required for the numerical simulation of the whole turbomachine make the use of CFD in the early phases of the design process infeasible: preliminary design relies on a number of physical and empirical relations, still quite similar to those used in the early history of turbomachinery design.

In this study, 87 independent parameters were used to define the geometry of a 7-stage compressor, the performance of which was evaluated using proprietary design codes for mean-line, multi-stage analysis. The effects on efficiency and surge margin of changing 44 design variables were analysed and their optimal values found by means of deterministic (gradient-based) and meta-heuristic (Tabu Search [TS]) optimisation methods.

The results show clearly how the use of meta-heuristic optimisation tools can improve the preliminary design of turbomachinery, allowing a more thorough but still rapid exploration of the design space to identify the most promising regions that will then be verified and further analysed with higher fidelity tools.

The results also reveal the impact of introducing various constraints into the design process, highlighting the effects of design decomposition.

INTRODUCTION

Since the first uses of gas turbine engines for aircraft applications back in the 1930s, the design of turbomachinery has evolved considerably. As is common with innovative technologies, the first engineers had little knowledge of the complex flow structures inside a compressor. The first attempts to gain a wider understanding of these phenomena involved extensive use of wind tunnel experiments on compressor cascades, leading to the development of a number of empirical correlations to be used for 1-dimensional (mean-line) and 2-dimensional (throughflow) design and analysis that remained the basis of turbomachinery design for over 30 years.

The development of powerful computers and CFD introduced a new era of aerodynamics and compressor design. Initially, CFD techniques were used to simulate the 2-dimensional flow behaviour around a profile (like a virtual cascade experiment), then ever increasing computer power made 3-dimensional viscous simulations of the flow around complex geometries a reality, and multi-stage simulations have now become routinely used in industry ([7]).

While initially CFD was mostly a verification tool to identify undesired flow features in an already established design, in recent years there has been much interest in the possibility of including computational analyses into the main design process to take advantage of the higher level of fidelity of CFD simulations. Unfortunately, this is not straightforward in compressor design (and even less so in gas turbine design) where the large number of variables and the still considerable computational time required for the simulation of the whole component make the use of CFD impracticable in the early phases of the process.

The large number of variables involved make the design process highly sequential and fragmented: the preliminary design uses low cost, low fidelity analysis (basically the same empirical rules developed in the early history of turbomachinery design) to perform a rapid exploration of the design space, with the aim of establishing the values for the most global variables to reduce the size of the design space. It is only at this point – when the general compressor outline is decided – that the use of higher fidelity and high cost tools (CFD) becomes feasible. Including CFD in the design process makes possible the optimisation of the more local design variables (such as blade thickness, lean and sweep, tip clearance, etc.) ([18]). While this offers the possibility of improving the performance of the design, the more global variables (which have the biggest influence on compressor performance) have already been decided, relying on empirical relations and simple physical rules. The high fragmentation of the compressor design process implies that each phase has to deal with a different subset of variables and this makes the early choices fundamental to the entire process. No amount of CFD can correct suboptimal preliminary decisions [13].

In addition, even more complication is added by the multi-objective and multi-disciplinary nature of the design: aerothermal analyses tend to be the driving discipline in the design of a turbomachine, but structural and economic considerations also play fundamental roles. Furthermore, even considering only aerodynamics, there are several design objectives which are almost always conflicting (maximum efficiency and surge margin, minimum size and weight). It is almost impossible, even for an experienced designer, to obtain an optimal design by simple trial and error: more sophisticated tools for exploring the design space are needed to obtain a complete overview.

DESIGN OPTIMISATION

All the features of the turbomachinery design process described above, together with the ever present commercial pressure for higher performance in shorter design times, make the use of automated design systems very attractive. Figure 1 shows the three principal components of a typical automated optimisation loop:

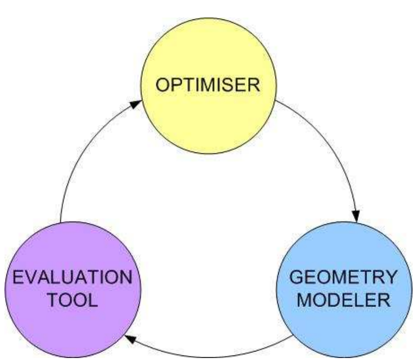

Figure 1: Optimisation loop

- *Optimiser*. The main novelty introduced by the automation of the design process is the use of an optimiser for driving the design towards the best possible compromise. Optimisation can be defined as the process that aims to determine the value of the vector $\vec{x}$ (*control variables*) that minimises (or maximises) an *objective function* $f(\vec{x})$ while satisfying a number of *constraints*. An optimiser essentially takes the place of the designer. In the traditional approach the designer starts from an initial design and, after having assessed its performance, changes it accordingly based on experience and knowledge until satisfying improvements have been achieved. The optimiser does exactly the same, using the information about the design vector and the objective function to produce a new design vector. As stated earlier, the design of turbomachinery is a multi-variate, multi-objective problem with a large number of local optima: the choice of the right optimisation algorithm is essential for the performance of the whole loop.

- *Geometry modeler*. A parametrisation scheme expresses the whole geometry through a number of parameters (a subset of these will form the design vector). These parameters have to define the geometry completely (for the purposes of the evaluation tool) and uniquely (they have to be mutually independent). It is important to ensure that the parametrisation scheme is able to cover all the feasible geometries, in order not to lose potentially good designs, but also that it uses the minimum possible number of parameters, since this affects the convergence time of the optimiser.

- *Evaluation tool*. This calculates the figure of merit for each candidate geometry. In the case of compressor design, it can range from a tool that evaluates expected performance based on empirical correlations to a fully viscous 3D CFD solver.

Aerodynamic design optimisation has received significant attention during the last 15 years. The first attempts to apply optimisation techniques to aerodynamic design were made by means of deterministic optimisers: [19] use Newton's method with approximate second-order derivatives in the multi-disciplinary (aerodynamic and structural) optimisation of a civil aircraft wing. [4] apply a gradient-based method to the minimisation of the drag of a body of revolution in supersonic flow. [15] link together a 16-variable parametrisation, a gradient-based optimiser and a 2D Euler solver for the two-point design of a transonic wing, concluding that numerical optimisation can improve the aerodynamic design while reducing the required design time, but also that the use of deterministic techniques can be impractical in the presence of multiple local optima.

One of the first applications of stochastic techniques to aerodynamic design is due to [11] who apply Genetic Algorithms (GAs) to the minimisation of the drag generated by a two-dimensional wing. They conclude that stochastic methods increase significantly the computational time required for design optimisation, but also that they achieve much better results than deterministic techniques in multimodal and highly constrained optimisation problems. [1] conduct the aerodynamic design of an axisymmetric forebody with the aid of a modified Simulated Annealing optimiser for minimising its drag, obtaining better results than with deterministic optimisers in comparable times. [25] apply GAs to the multidisciplinary design of a helicopter rotor blade, making use of decomposition methods to subdivide this multidisciplinary system into simpler subsystems, and implementing the algorithm on a parallel computer network to reduce calculation times.

An interesting study by [24] uses a combined random search code together with a gradient method for the optimisation of compressor blades for heavy-duty gas turbines. They combine a number of objectives into a single objective function, obtaining a new family of airfoils with improved performance relative to previously used controlled diffusion blades. They reach an important conclusion: the use of automated design procedures, not restricting the design space to the "state of the art" experience of the designer, has the potential to break existing design rules and move the search to unexplored promising regions.

[9] (2001b) compare the results of the application of a deterministic and a stochastic approach (GAs) to the optimisation of a multi-element airfoil for maximum lift and of a hull for minimum drag. Their conclusion is similar to that of [11]: GAs are more efficient in multi-criteria optimisations with many local minima, but they require higher computational times. They suggest the integration of stochastic and deterministic tools: the former are able to break the existing design rules to find new interesting unexpected shapes, while the latter are more efficient for local optimisation.

[3] use Evolutionary Algorithms (EAs) in the aerodynamic and structural optimisation of rotor blades for subsonic applications. Aggregating a number of objectives into a single objective function, they develop a new family of blades with design performance comparable to the original design but improved operating range and structural integrity. A similar study was performed by [28].

Two related studies were conducted in the Cambridge University Engineering Department by [16] and [22]. Harvey develops a system for automated optimisation of turbomachinery blading, with particular care to the choice of the parametrisation and optimisation schemes. His system utilises multiple accuracy levels in CFD calculations to minimise computational cost. Kellar uses a more general parametrisation scheme (free-form deformation) in the optimisation of a Formula 1 rear wing. This scheme has the big advantage of being able to handle arbitrary shapes, but it introduces a large number of variables. To overcome this problem, Kellar introduces a technique for automatically and adaptively selecting the most influential variables at each point of the optimisation.

Multi-objective optimisation

Aerodynamic design, in general, and compressor design, in particular, is typically a multi-objective problem. Deterministic methods are not very efficient in optimising multiple functions simultaneously if the trade-off surface is sought rather than a single optimum, while stochastic methods, apart often from being the best option in the presence of multiple local optima, have widely been applied with success in the solution of true multi-objective cases ([5]).

Despite these benefits, the large computational requirements of both multi-objective stochastic methods and CFD has limited the convergence of these techniques into an integrated design system, and it is only in recent years that, thanks to the development of ever more powerful computers, a serious exploration of this area has started.

One of the first applications of multi-objective stochastic optimisation to aerodynamic design is due to [8], who use GAs to analyse a multi-objective version of the same problems already solved in [9] and [10]. Using 2D CFD to evaluate the objective functions, they perform the optimisation of a three-element airfoil's flap and note how the time required for stochastic optimisation increases significantly. They suggest the use of approximated cost functions to overcome this limitation. A similar study is performed by [2], reaching basically the same conclusions.

[26] develop a multi-objective optimisation tool for multi-stage compressor design, using EAs as the optimiser and an axially symmetric throughflow code to evaluate the cost functions. They use 10 variables for each blade row (2 variables at 5 different radial locations) to express the geometry of a 4-stage axial compressor with inlet guide vanes, and, trying to maximise its pressure ratio and efficiency, obtain a Pareto front of improved designs.

[12] perform the multi-objective optimisation of a supersonic missile inlet, demonstrating the capacity of design optimisation in terms of design time reduction (in three months they obtain an inlet design with performance comparable to those resulting from 2 years standard design) and envisaging a more intensive use of optimisation techniques for multi-objective and multi-disciplinary problems. The study compares different algorithms using both single and multi-objective approaches. A similar study is conducted by [27] for the design of an industrial diffuser.

Finally, related studies were conducted in the Cambridge University Engineering Department by [23] and [20], extending the work of [16]. Kipouros develops a multi-objective automatic optimisation system for 3D blades based on an extension of the TS algorithm previously used by Harvey for single-objective optimisation. He applies it to the same test-case (slightly modified with the introduction of a second objective function) and shows the capability of the system to improve blade performance. Jaeggi improves the TS algorithm and develops a parallelised version of the code, showing performance to be comparable to the leading multi-objective GAs.

OPTIMISATION OF A 7-STAGE COMPRESSOR PRELIMINARY DESIGN

As discussed in the Introduction, the use of CFD in the very early phases of turbomachinery design is still impractical. The first part of the design of a multi-stage compressor has to be addressed with the help of a number of simpler physical laws and empirical relations. The aim of this preliminary design is to facilitate a more rapid exploration of the design space, allowing a reduction of its size by establishing the values for some global variables (annulus shape, number of blades, etc.), leaving CFD to produce further refinements by optimising the values for more local variables (blade shape, maximum thickness, etc.).

In this work, the optimisation of a 7-stage aeroengine compressor preliminary design has been accomplished. The components of the automatic optimisation loop are described in detail below:

Optimiser

Two different optimisation techniques were applied:

- a modified steepest descent algorithm: the gradient of the objective function suggests the direction of the move;
- the TS algorithm developed by [20].

The single-objective TS implementation of [6] is used as a starting point for this multi-objective variant. This uses a Hooke and Jeeves (H&J) local search algorithm (designed for continuous optimisation problems) ([17]) coupled with short, medium and long term memories to implement search intensification and diversification, as prescribed by [14].

TS operates in a sequential, iterative manner: the search starts at a given point and the algorithm selects a new point in the search space to be the next current point. The basic search pattern is a modified version of H&J search. Allowing for points that violate constraints or are tabu (see below), the H&J local search strategy requires approximately $2n\_var$ solution evaluations, where n_var is the number of design variables. A real-world problem may contain a large number of design variables and this strategy can become prohibitively expensive. To address this an element of random sampling is incorporated in the H&J step. The $2n\_var$ potential new points are generated, those that are tabu or infeasible are removed, and $n\_sample \leq 2n\_var$ points from those that remain are evaluated, selecting randomly to avoid introducing any directional bias. If one of these points dominates the current point, it is automatically accepted as the next point. If more than one point dominates the current point, a non-dominated point from these is randomly selected. If no points dominate the current point, a further n_sample points are sampled and the comparison is repeated. If all the feasible, non-tabu points have been sampled without finding a point that dominates the current solution, the standard selection procedure is employed.

Recently visited points are stored in the *Short Term Memory* (STM) on a first-in-first-out basis and are *tabu* – the search is not allowed to revisit these points.

The *Medium Term Memory* (MTM) maintains a record of the Pareto-optimal points found thus far during search, and its contents are the primary output at the end of the optimisation.

An *Intensification Memory* (IM) stores locally Pareto-optimal points that have not been selected as part of the H&J search pattern; the IM is used to select points for *search intensification*, focusing the search on areas of the search space with known good objective function values.

The *Long Term Memory* (LTM) records the regions of the search space which have been explored, and is used on *diversification*, directing the search to regions which are under-explored. This is achieved by dividing the allowed range for each control variable into a certain number of regions and counting the number of solutions evaluated in those regions.

A local iteration counter i_local is used, and reset upon a successful addition to the MTM. When i_local reaches user-specified values, the algorithm will diversify or intensify the search, or reduce the search step size and restart the search from a randomly selected point from the MTM. Thus, TS combines a systematic local search with a stochastic element and an intelligent coverage of the entire search space.

Geometry modeler

The geometry modeler translates the compressor geometry into the design vector that is managed by the optimiser during the optimisation loop. The *parametrisation* is one of the most important and complex phases of the whole process: the entire feasible geometry needs to be described with a minimum number of variables to reduce computational times.

In this study, 87 variables were used to describe the compressor geometry: the annulus shape was parametrised through 8 variables, 5 for the mean line (total length plus radius at 4 control points) and 3 for the area distribution (inlet area plus area ratio at the end of the compressor and in the middle). Rotational speed, number of stages and inlet flow speed, pressure ratio distribution, axial chord, spacing, number of blades, exit flow angle for each stage, blade thickness on chord and clearance on chord were used to describe the remaining features.

Evaluation tool

The evaluation of the compressor performance was done with a proprietary code for compressor mean-line evaluation. Given the compressor geometry, the code evaluates its performance in term of efficiency (total pressure losses for each stage are calculated as the sum of profile losses, secondary losses, tip clearance losses and shock losses) and operating margin (the load on the blade is calculated based on a number of parameters such as Diffusion Factor (DF), Diffusion Ratio, stage loading parameter $\frac{\Delta H}{U^2}$, De Haller number $\frac{V_2}{V_1}$, static pressure rise coefficient, Koch factor).

A prediction of the surge margin is also made on the basis of the DF and its increase for lower mass flows.

Formulation of the optimisation problem

Starting from a datum geometry from a generic 7-stage aeroengine core compressor, the optimisation had the goal of looking for possible improvements in performance and operating margin avoiding large changes in the geometry of the other engine components.

For these reasons, two objective functions were considered: isentropic efficiency (η) and surge margin (SM). The following 44 parameters were chosen as design variables: pressure ratio for each stage (keeping the total pressure ratio constant) (6 variables); annulus shape (compressor length, inlet and outlet mean radius and inlet area fixed) (4); axial chord for the blades (14); number of blades for each blade row (14); flow angle at the exit of each stator (apart from the last one where the exit flow angle was imposed) (6). Blade thickness/chord and clearance/chord were kept fixed as they involve structural, materials and manufacturing considerations that go beyond the present aims of this work.

A lower and an upper bound for each of the 44 design variables was set, mainly to avoid the search of clearly infeasible regions. Some infeasible areas may still exist in the design space (e.g. a DF higher than 0.6 would mean that the compressor is already stalled at the design point): in this case, a negative value is assigned to the two objective functions to move the optimiser towards different regions.

Many loading parameters are commonly used to measure the load carried by a blade row and thus the feasibility of a particular design: DF, Koch factor, De Haller number and pressure rise coefficient are probably the most common. Some of these (e.g. DF and Koch parameter, ones of the two most widely used ones) derive from very different approaches, and thus it is possible for a design to be deemed feasible by one of them and infeasible by the other. To obtain a "safer" design (since the DF is already taken into account in the correlations for calculating the surge margin) the optimisation has been constrained by specifying a maximum allowed value for Koch factor, De Haller number and static pressure rise coefficient.

RESULTS

Gradient-based optimisation

Since a deterministic approach is only applicable to the optimisation of a single objective function, the problem had to be reformulated with the use of a weighting factor. The objective function thus becomes:

$$f(\vec{x}) = a\eta(\vec{x}) + (1-a)SM(\vec{x}) \quad (1)$$

By changing the value of the weighting factor a it is possible to explore different efficiency-surge margin compromises. Figure 2 shows the results of optimising $f(\vec{x})$ for values of a ranging from 0.8 to 1[1]. The "improved efficiency redesign" ([21]) is the result of a previous optimisation performed by means of a combination of human-driven design and gradient-based optimisation. All the results have been expressed relative to datum. The presence of dominated solutions is a clear demonstration of the multi-modal nature of the objective function and thus of the inefficiency of a deterministic method in locating optimal designs for a similar problem.

Figure 2 also compares the results of the unconstrained optimisation with those of an optimisation in which a constraint was imposed to limit the maximum Koch factor to be no higher than that in the datum design. As expected, these differ significantly only in the lower region of the Pareto front, because the disagreement in DF and Koch factor becomes an issue only for highly loaded compressors.

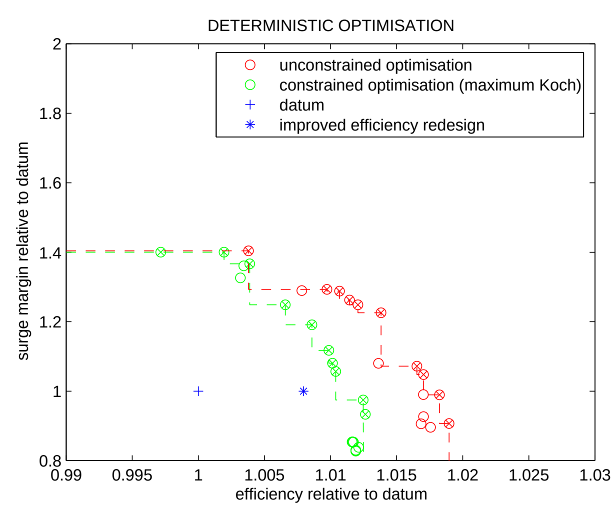

Figure 2: Unconstrained and constrained deterministic optimisation of the 7-stage compressor

[1]Lower values of a return designs that can be considered unrealistic due to the too low value of efficiency

Tabu Search

As demonstrated by the results of the deterministic optimisation, the multi-modality of the objective function, the large number of design variables and the presence of multiple objectives make the use of a stochastic approach attractive.

Figure 3 compares the results of deterministic and stochastic constrained optimisation. The gain in performance is evident.

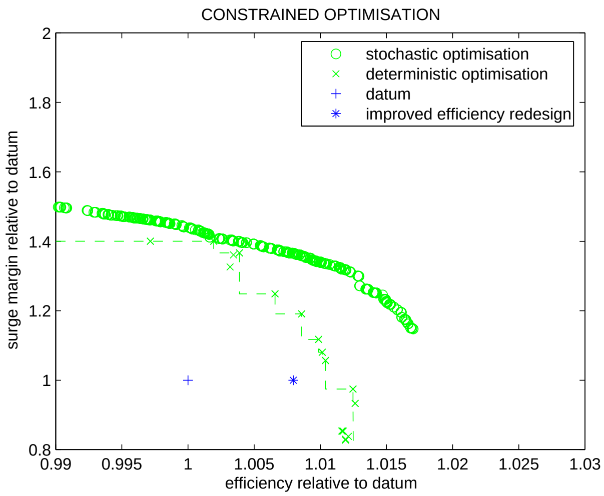

Figure 3: Constrained TS optimisation of the 7-stage compressor

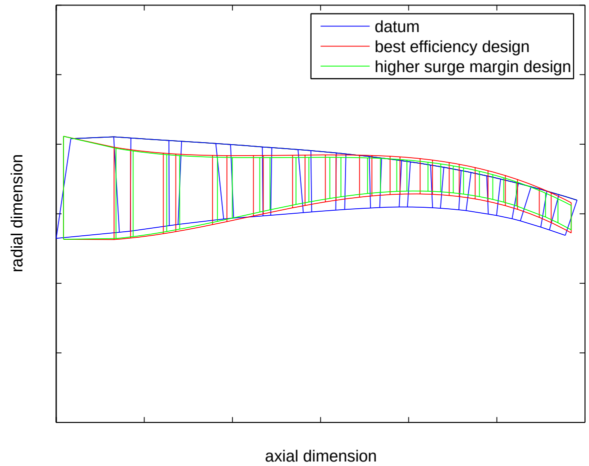

Figure 4: Initial and improved designs

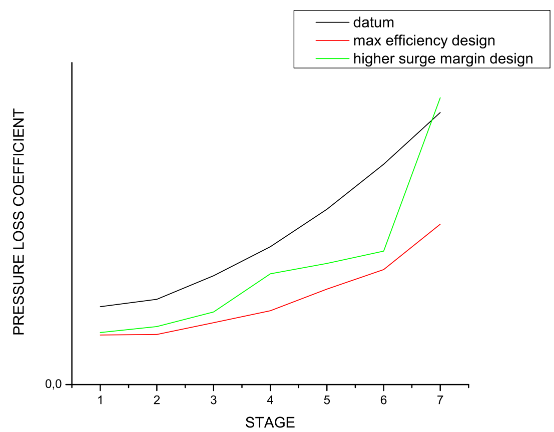

Figure 5: Rotor total pressure losses

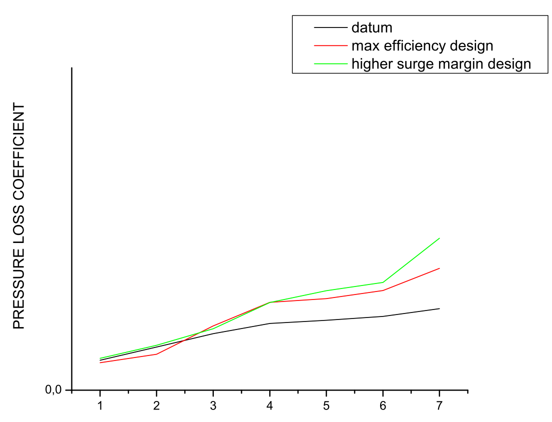

Figure 6: Stator total pressure losses

GEOMETRY COMPARISONS

Since a complete analysis of the whole Pareto front is clearly impractical, two designs will be analysed to understand the main characteristics of the design process: the maximum efficiency design (its surge margin is also much higher than that of the initial design) and a design with higher surge margin and an efficiency similar to that of the datum design.

For reasons of space, in the figures that follow information is provided about both these optimum designs and the datum design, while the characteristics of the optimum designs are discussed in turn. Only qualitative results are shown.

Best efficiency design

The best efficiency compressor (BEC) shows an improvement of 1.5% in efficiency and of 3.3% in surge margin. The main features of these designs are examined to understand where the gain in efficiency was achieved.

Figures 5 and 6 compare total pressure losses in each rotor and stator. For the BEC the losses reduce noticeably in the rotors, while there is a small increase in the stators. The net effect is a higher isentropic efficiency.

Since the total pressure losses are calculated as the sum of different types of losses, an analysis has been conducted to determine where the improvement in performance originates. As expected (both because secondary losses, tip clearance losses and shock losses depend on geometrical parameters that remained basically unchanged during the optimisation and because profile losses are usually about 2/3 of the total losses), the total pressure losses trend originates from a similar trend in profile losses (figures 7 and 8).

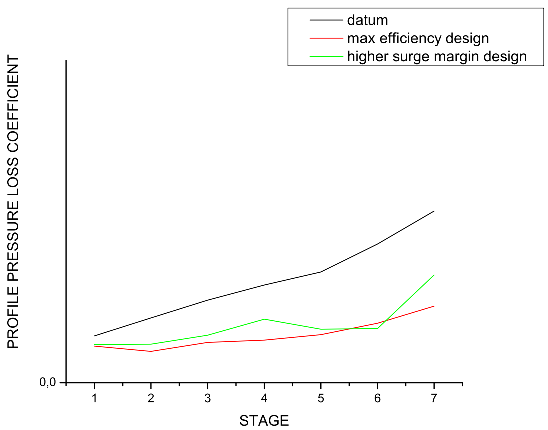

Figure 7: Rotor profile pressure losses

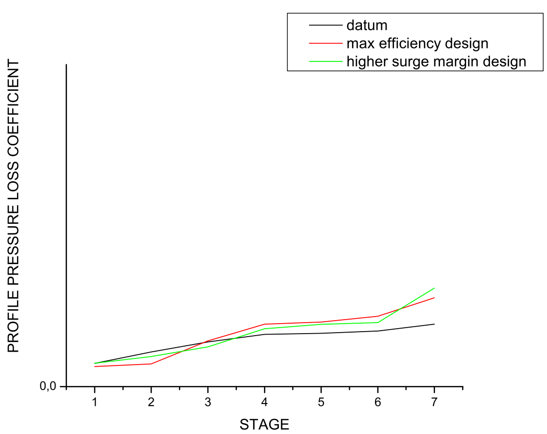

Figure 8: Stator profile pressure losses

Profile losses are the sum of losses due to friction and losses due to flow separation. While the former are proportional to the velocity squared, the latter depend on the development of the boundary layer on the suction surface of the blade due to the adverse pressure gradient. The more a blade is cambered the larger the adverse pressure gradient, the greater the deviation of the flow from the blade exit angle and thus the losses are. Clearly there is an optimal trade-off for minimising the total profile losses.

Despite the large numbers of variables used during the optimisation and of parameters influencing compressor performance, it is still possible to outline some important changes between the geometries shown in figure 4 that led to the improved performance.

In the **rotors**, the decreased annulus area yields a higher flow velocity that reduces the deflection needed to obtain the required pressure ratio (or ΔV_w) and thus the load on the blade row. This leads to lower flow separation but also increased frictional losses that can be controlled by reducing the number of blades (with a positive effect also on compressor weight). Fewer blades mean lower frictional losses but also higher separation losses and loads (reduced stall margin). Losses and loads are also influenced by the chord of the blades. Stage pressure ratio and blade tangential speed (which depends only on the radius in this case since the rotational speed of the spool is fixed in the optimisation) also influence the required ΔV_w or flow deflection. Summarising, flow velocity, number of blades, blade chords and stage pressure ratios influence losses and operating margins. Consideration must be given to the fact that some of these variables influence other blade rows' performance.

The higher flow velocity, together with the different distribution of pressure ratio, number of blades and blade chords, allow the elimination of the strange peak in flow deviation in the 6th stage of the datum compressor (figure 9). There is a small increase in required flow deflection in the first stages due to the reduced mean radius, while the reduction in the last stages depends on the above mentioned changes. The reduction in flow deflection is particularly important also in consideration of the fact that the last stages are the ones that suffer the most in reduced mass flow operation.

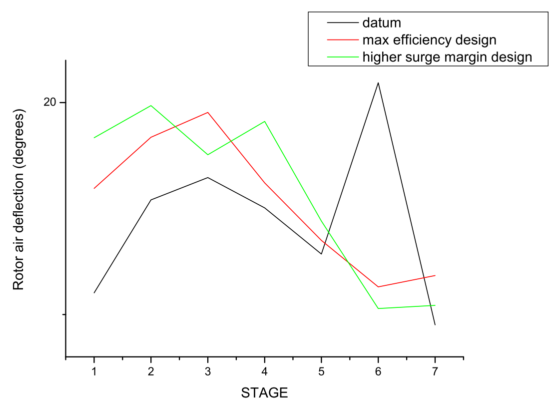

Figure 9: Rotor flow deviation distributions

All these considerations are supported by the loadings on the different blade rows, analysed with different parameters. The only significant change is in the De Haller number (figure 10) (which is not very important because a simple ratio of exit to inlet flow velocity without geometrical considerations and useful only in the very early stages of the design), while DF (figure 12) and pressure rise coefficient (figure 11) (ratio of static pressure rise and stalling pressure rise based on geometrical considerations) show values very similar to those of the initial design. The reduction in load achieved through the increased flow velocity is balanced by the reduced number of blades, with a net gain in efficiency. The increased surge margin is largely due to the reduced flow deflection in the 6th stage.

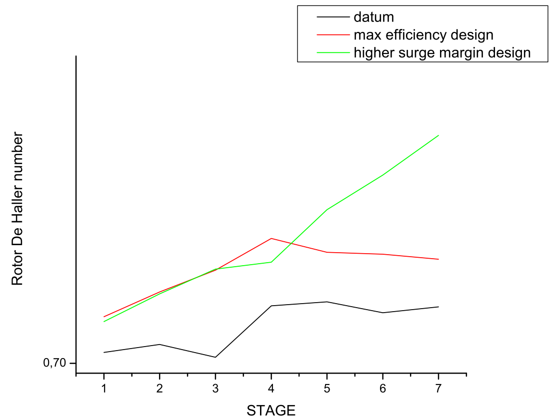

Figure 10: Rotor De Haller number distributions

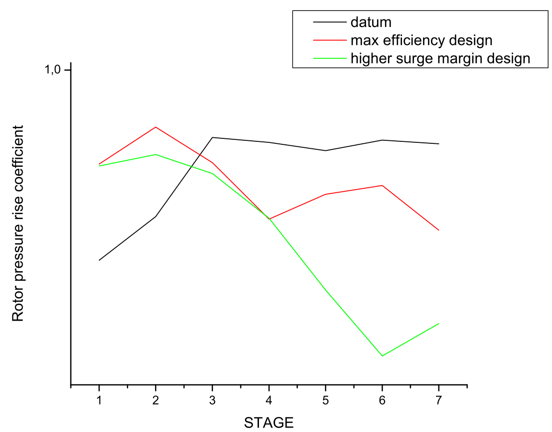

Figure 11: Rotor pressure rise coefficient distributions

The situation is similar in the **stators**, where the increased flow velocity reduces the inlet flow angle and thus the required deflection (figure 13), with a consequent reduction in deviation losses. The profile losses can again be controlled by reducing the number of blades, now possible for the increased flow velocity. Blade chords can control the load too and exit flow angles can be reduced to decrease the load on the following rotor or increased to reduce the load on the stator itself. All these considerations are supported by the loading parameters: once again De Haller number increases (figure 14) because it does not take into account blade geometry, while pressure rise coefficient (figure 15) and DF (figure 16) are basically the same (lower diffusion but also fewer blades).

All the previous considerations are corroborated by the Koch factor distribution along the stages (figure 17). The general level of loading is very similar to the original design,

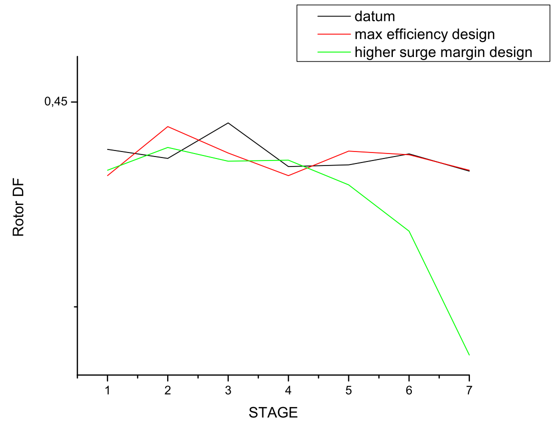

Figure 12: Rotor Diffusion Factor distributions

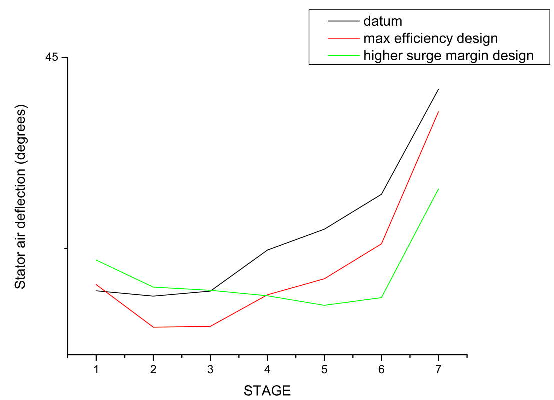

Figure 13: Stator flow deviation distributions

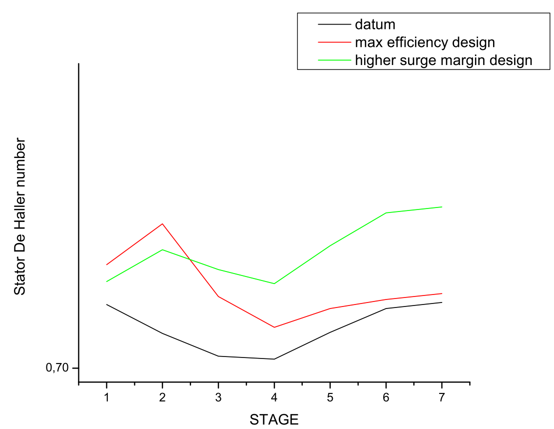

Figure 14: Stator De Haller number distributions

but the load is more evenly distributed among the stages, with the first and last stages a bit underloaded (a positive effect at off-design conditions).

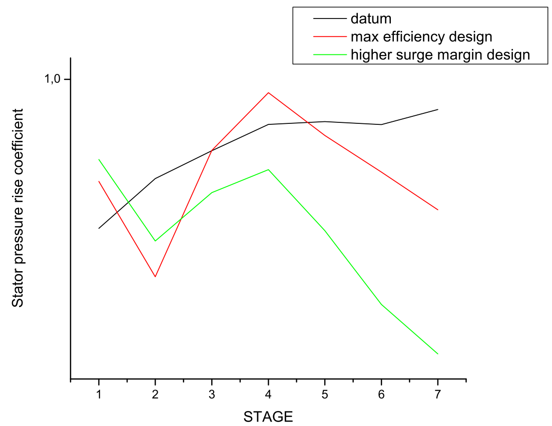

Figure 15: Stator pressure rise coefficient distributions

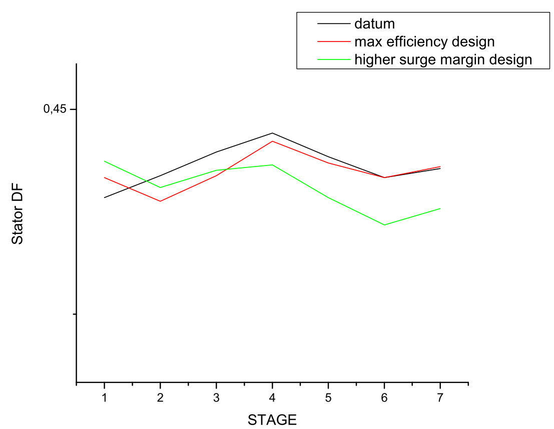

Figure 16: Stator Diffusion Factor distributions

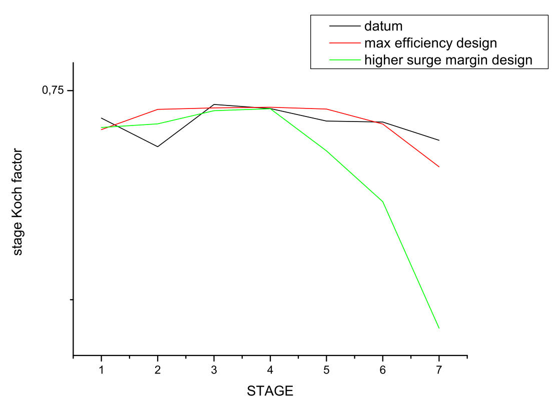

Figure 17: Stage Koch factor distributions

Higher surge margin design

To understand the reasons behind the high spread of points in the computed Pareto front (figure 3), a higher surge margin design (HSMD) has been analysed. The point chosen has an efficiency very similar to the datum design and a surge margin larger by 9.4%.

Figure 4 compares the HSMD with the datum geometry and with the BEC. The annulus is very similar to the latter, but there is a further reduction in area leading to a further increment in flow velocity.

As previously explained, the effect of the increased flow velocity is again a lower blade camber for the last stages (figures 9 and 13). The difference in the first stages is due to a different pressure ratio distribution.

The lower required flow deflection leads to a reduction in blade loading, evident from the lower Diffusion (figures 12 and 16) and Koch factors (figure 17), and hence the increased surge margin.

At the same time, the increased flow velocity causes a non-optimal (from an efficiency point of view) balance between deviation and frictional losses and thus the lower isentropic efficiency. Figures 5 and 6 compare the losses for rotors and stators for the three designs.

The high surge margin and the low loadings on the last stages suggest that the higher flow velocity could allow a reduction in the number of stages (with a consequent reduction in compressor size, weight and probably manufacturing costs).

ADDITIONAL CONSTRAINT: EXIT MACH NUMBER

As shown, quite a significant gain in performance can be achieved through a reduction in annulus area. However, the increased flow velocity at the end of the compressor may have a big impact on the performance of the downstream component.

The effect of constraining the exit Mach number was analysed both to verify its implications on compressor design and to test the feasibility of this approach in a even more constrained design space.

The new Pareto front (figure 18) clearly lies below the one obtained for the free exit Mach number optimisation. It is still possible to design a compressor with an isentropic efficiency similar to the one obtained in the previous optimisation, but at the cost of a reduced operating margin.

CONCLUSIONS

Despite the increasing use of CFD in recent years, the large number of variables involved in compressor design, the nature of the design space and the coexistence of multiple often conflicting objectives do not allow a direct application of CFD throughout the design process. In the early phases, a number of lower fidelity tools need to be used to allow a more rapid exploration of the design space, aiming to identify its most promising regions that will then be analysed in more detail. The large number of variables together with the topology of the design space make it impossible to produce an optimal design by simple trial and error: an automated optimisation system used in conjunction with

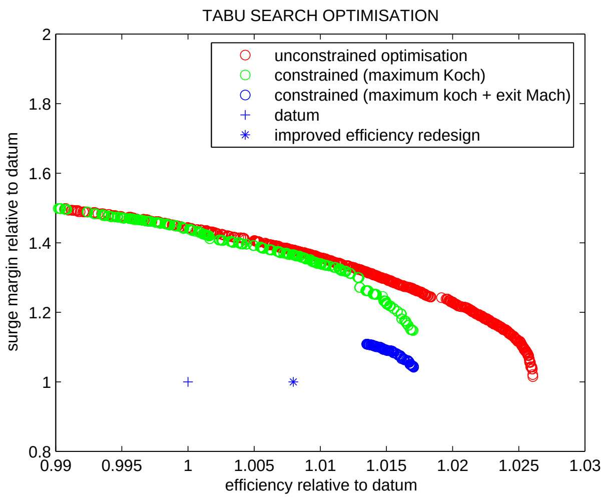

Figure 18: Mach constrained optimisation

preliminary design evaluation tools can allow a rapid and thorough exploration of the design space.

This paper reports the results of the design optimisation of a 7-stage generic compressor. The geometry was parametrised in terms of 87 variables; 44 of these were used as design variables in the optimisation process, while the rest remained fixed at their initial values.

A two-objective optimisation (efficiency and surge margin) was conducted both with a gradient-based method and a metaheuristic one (Tabu Search). To increase the confidence in the predicted operating limits, further measures of the blade loading were added through the introduction of design constraints that involve the value of De Haller number, static pressure rise coefficient and Koch factor. The results show the advantages of using a stochastic approach in the preliminary design of a multi-stage compressor, where the large design space, the multi-objective nature of the design and the presence of multiple local optima make the use of deterministic optimisers inconvenient.

In recognition of the fact that the design of compressor systems is often limited by other engine components' requirements, a further constraint was added on the exit Mach number in order to not overload the following duct. The results show a much smaller Pareto front, but still the possibility of large performance improvements relative to the initial design: the use of TS together with a mean-line performance analysis evaluation tool is a fast approach to the preliminary design of multi-stage compressors.

These results also highlight the effects on component performance of the constraints imposed to facilitate design decomposition – to allow components to be designed in parallel. In the next stage of this research we will seek to optimise a compressor and a downstream duct in combination to establish whether there is scope to further improve performance by relaxing the constraints on the interface between these components.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the support of the Regione Autonoma Sardegna (RAS), the Cambridge European Trust and the UK Engineering and Physical Sciences Research Council (EPSRC) under grant GR/R64100/01.

REFERENCES

[1] S. Aly, M. Ogot, and R. Pelz. Stochastic approach to optimal aerodynamic shape design. *Journal of Aircraft*, 33(5):956–961, 1996.

[2] E. Benini and A. Toffolo. Development of high-performance airfoils for axial flow compressors using evolutionary computation. *Journal of Propulsion and Power*, 18(3):544–554, 2002.

[3] D. Buche, G. Guidati, and P. Stoll. Automated design optimisation of compressor blades for stationary, large scale turbomachinery. *Proceedings of the IGTI03 ASME Turbo Expo 2003, Atlanta, Georgia, USA*, 2003.

[4] S. Cheung, P. Aaronson, and T. Edward. Optimization of a theoretical minimum drag body. *Journal of Aircraft*, 32(1):193–198, 1995.

[5] C. A. Coello Coello, D. A. Von Veldhuizen, and G. A. Lamont. *Evolutionary Algorithms for Solving Multi-objective Problems*. Kluwer Academic Publisher, USA, 2002.

[6] A. M. Connor and D. G. Tilley. A tabu search method for the optimization of fluid power circuits. *IMechE Journal of Systems and Control*, 212(5):373–381, 1998.

[7] J. D. Denton and W. N. Dawes. Computational fluid dynamics for turbomachinery design. *IMechE*, 1999.

[8] R. Duvigneau and M. Visonneau. Single- and multi-objective optimisation for high fidelity CFD using genetic algorithms. *Evolutionary Methods for Design, Optimization and Control, 19-21 September, Athens, Greece*, 2001.

[9] R. Duvigneau and M. Visonneau. Towards a practical design optimization tool for incompressible and turbulent flows. *4th Numerical Towing Tank Symposium, 23-25 September 2001, Hamburg, Germany*, 2001.

[10] R. Duvigneau and M. Visonneau. Shape optimization strategies for complex applications in computational fluid dynamics. *2nd International Conference on Computer Applications and Information Technology in the Maritime Industries, 14-17 May 2003, Hamburg, Germany*, 2003.

[11] P. Gage and I. Kroo. A role for genetic algorithm in preliminary design environment. *AIAA paper 93-3933*, 1993.

[12] A. Gaiddon and D. D. Knight. Multicriteria design optimization of integrated three-dimensional supersonic inlets. *Journal of Propulsion and Power*, 19(3):456–463, 2003.

[13] S. Gallimore. Axial flow compressor design. *IMechE*, 1999.

[14] F. Glover and M. Laguna. *Tabu Search*. Kluwer Academic Publisher, USA, 1997.

[15] J. Hager, S. Eyi, and K. Lee. Two-point transonic airfoil design using optimization for improved off-design performance. *Journal of Aircraft*, 31(1):1143–1147, 1994.

[16] S. A. Harvey. *The design optimisation of turbomachinery blade rows*. PhD thesis, University of Cambridge, 2002.

[17] R. Hooke and T. Jeeves. "Direct search" solution of numerical and statistical problems. *Journal of the Association of Computing Machinery*, 8(2):212–229, 1961.

[18] J. H. Horlock and J. D. Denton. A review of some early design practice using computational fluid dynamics and a current perspective. *ASME Journal of Turbomachinery*, 127, 2005.

[19] M. G. Hutchinson, E. R. Unger, W. H. Mason, B. Grossman, and R. T. Haftka. Variable-complexity aerodynamic optimization of high-speed civil transport wing. *Journal of Aircraft*, 31(5):110–116, 1994.

[20] D. M. Jaeggi, T. Kipouros, G. T. Parks, and P. J. Clarkson. A multi-objective tabu search algorithm for constrained optimisation problems. *The 3rd International Conference of Evolutionary Multi-Criterion Optimisation, Guanajuato, Mexico*, 2005.

[21] J. P. Jarrett. Towards orthogonal turbomachinery design. *ASME International Gas Turbine Institute TURBO EXPO 2006, Barcelona, Spain*, 2006.

[22] W. P. Kellar. *Geometry modeling in computational fluid dynamic and design optimisation*. PhD thesis, University of Cambridge, 2002.

[23] T. Kipouros, D. M. Jaeggi, W. N. Dawes, G. T. Parks, and A. M. Savill. Multi-criteria optimisation of turbomachinery blades: investigating the trade-off surface. *41st AIAA/AME/SAE/ASEE Joint Propulsion Conference and Exhibit, Tucson, Arizona*, 2005.

[24] U. Koller, R. Monig, B. Kusters, and H. A. Shreiber. Development of advanced compressor airfoils for heavy-duty gas turbines - Part I: design and optimisation. *Trans. of the ASME Journal of Turbomachinery*, 122:397–405, 2000.

[25] J. Lee and P. Ayela. Parallel genetic algorithm implementation in multidisciplinary rotor blade design. *Journal of Aircraft*, 33(5):962–969, 1996.

[26] A. Oyama and M. S. Liou. Multiobjective optimization of a multi-stage compressor using evolutionary algorithm. *AIAA 2002-3535*, 2002.

[27] S. Shahpar. SOPHY: an integrated CFD based automatic design optimisation system. *ISABE-2005-1086*, 2005.

[28] F. Sieverding, B Ribi, M. Casey, and M. Mayer. Design of industrial axial compressor blade sections for optimal range and performance. *Trans. of ASME Journal of Turbomachinery*, 126:323–331, 2004.