Multi-disciplinary System Design Optimisation of a 2-Stage Axial Compressor for Aerodynamic Performance, Production Cost and Dynamic Stability

Josep M. Dorca-Luque & Vincent Perrot

Abstract

A simultaneous optimisation of a 2-stage compressor for aerostructural performance, cost and stability margin is carried out by suitable variation of the blade counts, dimensions and metal angles. The design improvement is done through sequential steps. A single-objective optimisation by gradient-based and heuristic techniques in combination serves to establish the upper performance boundary for the machine, in terms of pressure ratio. Subsequently, a Pareto bi-objective optimisation is performed to unveil the relation between surge margin and cost for a given lower bound on aero-performance, as set by the initial study. The cost and surge margin are found to be highly opposing, since they require a minimisation and maximisation of the bladerow solidities, respectively. Three discrete families of designs are defined by the Pareto front, and the one closest to the standard 20% surge margin requirement is chosen to proceed to the final multi-objective optimisation.

The optimised design has a pressure ratio that is 4% superior to the initial design, with the same level of efficiency. The surge margin is increased by 22% to the required level of 20%, whilst the cost is reduced by 24.4%.

As well as outlining a successful methodology for integrated optimisation, the key inter-relation between stability and cost for a given performance is clearly elucidated. This establishes the need for more accurate methods of compressor dynamic stability prediction, which supersede the over-conservative surge-margin. Using these new metrics, operation at much lower surge margins can be safely justified and significant reductions in production costs are to be expected.

1 - Introduction

The incorporation of dynamic stability in axial flow compressor design has traditionally been carried out by enforcing an industry-standard surge margin on the aerodynamic performance characteristic. Current research in this field is focused on the understanding of the phenomena leading to compression system instability. It is envisaged that totally new metrics will be developed in the near future that will rely on dynamic models capable of determining the exact degree of dynamic safety at each operating point.

This report presents a multi-disciplinary optimisation (MDO) study to determine the blade geometry and configuration in a 2-stage axial compressor that is capable of achieving preestablished levels of aerostructural performance and stability within tight bounds in production cost. The purpose of this work is threefold:

- To outline a systematic methodology for the multi-disciplinary optimisation of such systems.
- To demonstrate a practical inverse design with embedded stability considerations, contrary to current practice, which is mainly oriented to the dynamic analysis of existing compressors with given geometries.
- To assess the potential performance and cost benefits of incorporating stability requirements throughout the design process, thus giving the motivation for more advanced dynamic compression system models.

A simulation code has been developed in order to determine the aerodynamic, structural, dynamic and cost characteristics of different machines. The physical basis for these models is discussed at the onset.

An overview of the optimisation strategy is presented subsequently, emphasising the progressive nature with which such a multi-disciplinary task must be carried out; one cannot expect to perform a single, successful multi-objective optimisation run, but rather must gradually approach the target by different methods. The sequence of optimisation phases is detailed, together with the expected progression of the design through each one.

Having set the scene, the analytical procedures of each optimisation step and the results obtained are outlined.

Finally, a detailed discussion of the physical evolution of the compressor throughout the optimisation is given. The improvements in the various areas of performance are calculated, with concentrated attention on the role of the surge margin, as a stability metric, in driving the optimised compressor towards a particular region of the design space.

2 – Nomenclature

C	Cost
c	Blade chord
c_R	Rotor chord, *m*
c_S	Stator chord, *m*
D	Diffusion factor
i	Incidence angle into a blade/vane, °
M	Mach number
m	Machining cost per unit surface area, *$/m²*
$\dot{m}$	Mass flow, *(9 kg/s)*
N_R	Number of rotor blades
N_S	Number of stator vanes
N_{stages}	Number of stages, *(2)*
p	Material cost per unit mass, *$/kg*
PR	Pressure ratio
P_{To}	Inlet total pressure, *(101300 Pa)*
$R_{tip}(z)$	Tip radius distribution, *m*
$R_{hub}(z)$	Hub radius distribution, *m*
$\overline{R}$	Euler radius at a given axial location
SM	Surge margin
s	Blade spacing
T_{To}	Inlet total temperature, *(288.15 K)*
z	Longitudinal coordinate along the compressor, *m*
β_{in}	Inlet swirl angle post-IGV, *(0 °)*
$\beta_{mi,R}$	Rotor inlet metal angle, °
$\beta_{mo,R}$	Rotor exit metal angle, °
$\beta_{mi,S}$	Stator inlet metal angle, °
$\beta_{mo,S}$	Stator exit metal angle, °
ϕ	Blade turning, °
η	Adiabatic efficiency
ρ_{blade}	Blade material density, *kg/m³*
σ	Bladerow solidity
σ_{blades}	Blades maximum bending stress, *Pa*
$\sigma_{u,blade}$	Blade ultimate tensile stress, *Pa*
ω	Nondimensional pressure loss coefficient $\omega = \frac{\Delta P_o}{P_{o,1} - P_1} = \frac{P_{o,1} - P_{o,2}}{P_{o,1} - P_1}$ 1 – Blade inlet 2 – Blade exit
Ω	Rotational speed, (1000 *rad/s)*

3 – Simulation Code and Optimisation Problem

3.1 Optimisation problem statement

The blade count and geometry are to be optimised for one particular operating point. For this particular condition, it is attempted to maximise the aerodynamic performance (*PR* & η) while at the same time minimising the cost *C*. One must ensure that a surge margin *SM* of at least 20% is maintained at all times and that structural integrity is not compromised.*($\sigma_{blades} < \sigma_{u,blades}$)*.

Due to limitations in the number of inputs to the optimisation software and in order to simplify the analysis, it is assumed that both stages share the exact same blade geometry ($Rot_1 = Rot_2$ & $Sta_1 = Sta_2$).

The design variables are therefore:

- Rotor inlet metal angle, $\beta_{mi,R1} = \beta_{mi,R2}$
- Rotor outlet metal angle, $\beta_{mo,R1} = \beta_{mo,R2}$
- Rotor mean chord, $c_{R1} = c_{R2}$
- Number of rotor blades, $N_{R1} = N_{R2}$
- Stator inlet metal angle, $\beta_{mi,S1} = \beta_{mi,S2}$
- Stator outlet metal angle, $\beta_{mo,S1} = \beta_{mo,S2}$
- Stator mean chord, $c_{S1} = c_{S2}$
- Number of stator blades, $N_{S1} = N_{S2}$

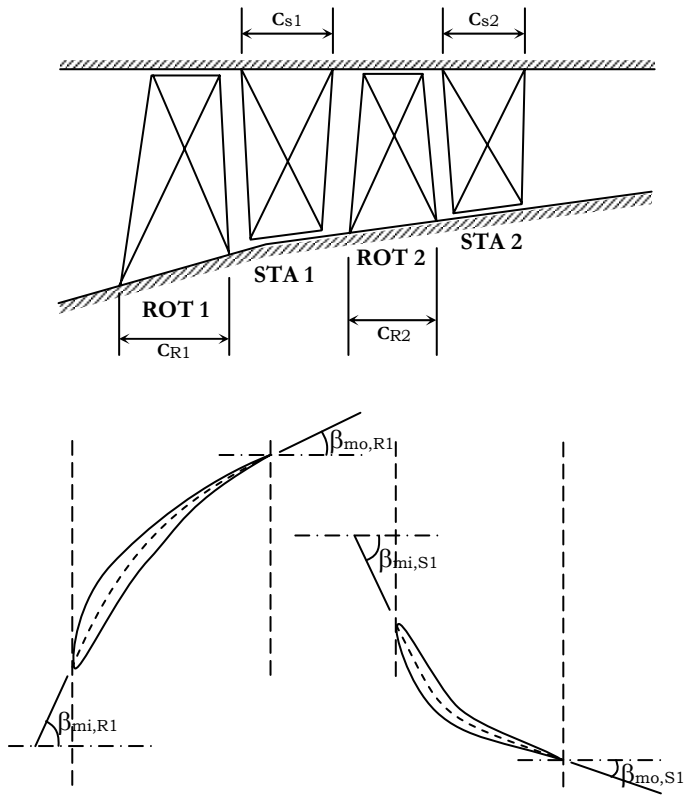

Figure 1

All other thermodynamic and operational parameters are fixed to the values shown in the nomenclature section. We are dealing with an optimisation at design point only.

The formal statement of the problem is given on the following page.

The bounds on blade counts and metal angles are based on experience and restrict the design space somewhat so that only realistic blade configurations are produced.

On the other hand, the limits on chord dimensions are set in order to prevent the development of blades of excessive aspect ratio in the design, which would be unfeasible for aeroelastic reasons.

$$\text{Max } \vec{J}(\vec{x},\vec{p}) = \begin{bmatrix} J_1(\vec{x},\vec{p}) = PR(\vec{x},\vec{p}) \\ J_2(\vec{x},\vec{p}) = \eta(\vec{x},\vec{p}) \\ J_3(\vec{x},\vec{p}) = 1/C(\vec{x},\vec{p}) \end{bmatrix}$$

Subject to:

$$\vec{g}(\vec{x},\vec{p}) = \begin{bmatrix} g_1(\vec{x},\vec{p}) = \sigma_{blades}(\vec{x},\vec{p}) - \sigma_{u,blades} \\ g_3(\vec{x},\vec{p}) = SM(\vec{x},\vec{p}) - 20\% \end{bmatrix} \le 0$$

$$65 < N_R < 69$$

$$63 < N_S < 69$$

$$0.06 > c_R > 0.03$$

$$0.05 > c_S > 0.025$$

$$57 < \beta_{in,R\&S} < 70$$

$$27 < \beta_{out,R} < 35$$

$$0 < \beta_{out,S} < 15$$

$$\vec{x} = \left[\beta_{mi,R} \quad \beta_{mo,R} \quad c_R \quad N_R \quad \beta_{mi,S} \quad \beta_{mo,S} \quad c_S \quad N_S\right]$$

$$\vec{p} = [\dot{m}, R_{tip}(z), R_{tip}(z), R_{hub}(z), N_{stages}, \omega, P_{To}, T_{To}, \beta_{IGV}, \rho_{blade}, p, \sigma_{u,blades}, m(\phi), \omega(M,i), \omega(D, geometry)]$$

3.2 Simulation Model

The simulation model used for the optimisation consists of four distinct parts: a compressible mean flow solver for aero-performance, coupled to a separate loss-prediction tool; a dynamic stability code to determine the surge margin; a structural analysis solver; and a cost model. The figure below shows the architecture, data flow and optimisation interfaces.

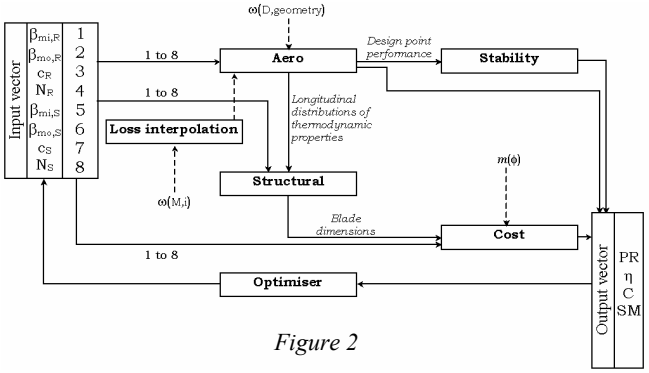

Figure 2

In order to minimise the number of iterations performed in each *simcode* run, an N2 analysis of the various modules is carried out. After rearrangement, the optimum sequence is shown in Figure 3.

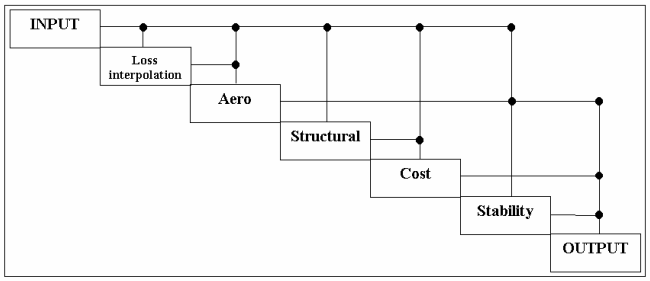

Figure 3

While a complete description of the physical principles behind the fluid and solid mechanics on which the model is based is clearly beyond the scope of this article, a brief description of each module is provided.

3.2.1 Aerodynamics

This module is a 1-D compressible mean line solver for the compressor. The routine is actually also employed in the stability analysis, where the full characteristic at design speed has to be calculated. A brief schematic of the program is given below:

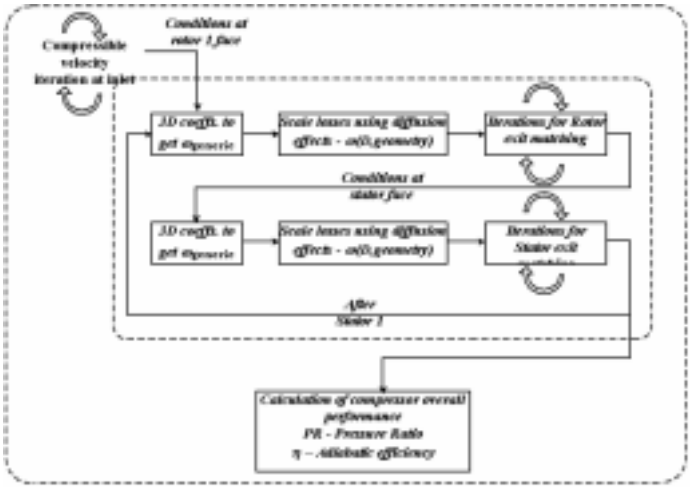

Figure 4

3.2.2 Loss interpolation:

The 3D coefficients referred to in the aero code are obtained from a separate profile loss model. Generic 3D surfaces are employed that relate loss levels to incidence and Mach number for a given technology level. It is assumed that this applies regardless of the blade configuration.

The surface is usually defined as a collection of points in 3D space from empirical results or CFD analyses. The interpolation routine finds the coefficients of a polynomial that represents it in a least-squares sense.

This polynomial is then used in run-time to determine the point in the 3D surface that corresponds to the flow conditions calculated prior to a blade row. This point represents the reference pressure loss for our particular incidence and Mach number, i.e. the loss that the profile would have if it had a diffusion factor *D=0.45*.

It is straightforward to calculate the real D-factor of the blade from:

$$D = 1 - \frac{V_2}{V_1} + \frac{|v_2 - v_1|}{2\sigma V_2} \text{ , where } \sigma = \frac{c}{s} = \frac{cN}{2\pi\bar{R}}$$

Another trend is available that relates the D-factor to a function *f(ω)*, which is used to scale the reference losses.

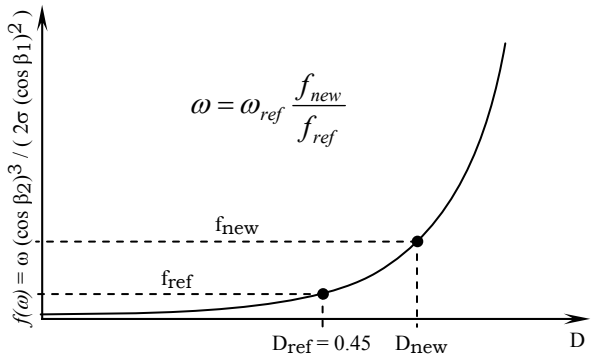

Figure 5

3.2.3 Structural

Once the flowfield has been solved, the loads on each of the four blades are calculated by standard elasticity theory. The stress suffered by each bladerow is never actually calculated. Instead, the code solves for the minimum thickness required to develop a stress level that is below the material limit by a given factor of safety. Different thicknesses are calculated for each of the bladerows. Thus, the rotors will have the same chord, but not necessarily the same thickness. The same applies to the stators.

3.2.4 Stability

The aerodynamics code is run in a sequential fashion. The mass flow is progressively reduced from the design point by arbitrarily small intervals *(0.05kg/s)*. According to the Moore-Greitzer[1] formulation for dynamic stability of lumped, actuator-disk compressor models, neutral stability is reached when:

$$\frac{\partial \psi_{TS}}{\partial \phi} = 0$$

Where ψ_{TS} and ϕ are the nondimensional total-static pressure rise and flow coefficient, respectively. These are calculated automatically at each mass flow step.

The process is repeated until the extremum is reached. That point is considered to be the inception of instability. The surge margin is calculated based on the neutral stability performance (*NS*) and the one obtained at the design point (*DP*), from the expression:

$$SM = \frac{PR_{NS} - PR_{DP}}{PR_{NS}}$$

3.2.5 Cost

The total cost is split between material and manufacturing cost.

- Since the blade dimensions have been established in the structural code, the volume and surface area are easily calculated. This enables a rapid determination of the metal cost.

- The geometric turning of each blade is known. A value of machining cost per unit surface area is read off a correlation for $m(\phi)$ and the machining cost derived accordingly.

The actual numerical values for the manufacturing and material costs have been chosen such that their relative sizes are similar to what would be expected in a real process. The absolute value of the cost is, however, not realistic. The constants have been set so as to obtain a cost of *O(100)*. This makes it possible to illustrate scaling issues later on and does not in any way modify the outcome of the optimisation process.

4 – Optimisation strategy

In this section,an outline is given of optimization strategy that is adopted to go from a well-defined problem, to a well-suited multiobjective optimisation of it.

Instead of beginning in a multi-objective fashion, a sequence of single-objective steps taking the *PR* as the objective function are carried out. These push the design towards a region of high performance, which we want to attain anyway. The different incremental steps in this first section are:

- DOE: In order to establish a starting point for the deterministic optimisation approaches.
- Sequential Quadratic Programming (SQP): Gradient-based optimisation to get the first degree of improvement. At this early stage, a sensitivity analysis is performed to understand the role of each variable.
- Analysis of the Hessian matrix by finite difference towards the application of appropriate scaling factors to the variables.
- Scaled SQP optimisation to try to further improve the design with the new variable definitions.

A heuristic optimisation by means of a Genetic Algorithm (GA) is also carried out to establish a basis for comparison between the techniques.

Once the performance level that can be expected is known, a separate bi-objective Pareto optimisation is carried out on *SM* and *C*, as a clear physical interaction exists between the two, which is discussed later on.

The most suitable design from the Pareto front is the point taken to the final multi-objective optimisation, which constitutes the end to our task.

5 – Design Of Experiments (DOE)

This technique was employed to identify an initial starting point for the numerical optimisation methods.

For this purpose, an orthogonal array of experiments was initially envisaged as a DOE technique to use. However, using 8 variables with 3 levels each (chosen arbitrarily based on our experience), it soon became apparent that the number of experiments to carry out in order to maintain orthogonality throughout would be excessive. A *quasi-orthogonal* array is proposed instead, providing some degree of broad coverage. Half the factors are orthogonal, whereas the other half are taken in sequenced combination so as to minimise repetition.

The results of the DOE experiments are presented in figure 6.

	Orthogonal Region				Sequenced region			
	A	B	C	D	E	F	G	H
Expt #	$\beta_{mi,R}$	$\beta_{mo,R}$	C_r	N_r	$\beta_{mi,S}$	$\beta_{mo,S}$	C_s	N_s
1	A1	B1	C1	D1	E1	F1	G3	H2
2	A1	B2	C2	D2	E1	F2	G3	H2
3	A1	B3	C3	D3	E2	F3	G1	H3
4	A2	B1	C2	D3	E2	F1	G2	H3
5	A2	B2	C3	D1	E3	F1	G3	H1
6	A2	B3	C1	D2	E3	F2	G1	H2
7	A3	B1	C3	D2	E1	F2	G1	H3
8	A3	B2	C1	D3	E2	F3	G2	H1
9	A3	B3	C2	D1	E3	F3	G2	H1

DEFINITION OF FACTOR LEVELS

A - $\beta_{mi,R}$ (deg)		C - C_r (m)		E - $\beta_{mi,S}$ (deg)		G - C_s (m)	
1	60	1	0.06	1	56	1	0.06
2	63	2	0.08	2	59	2	0.08
3	66	3	0.1	3	62	3	0.1
B - $\beta_{mo,R}$ (deg)		**D - N_r**		**F - $\beta_{mo,S}$ (deg)**		**H - N_s**	
1	33	1	67	1	8.66	1	65
2	35	2	72	2	10.66	2	70
3	37	3	77	3	12.66	3	75

Expt #	PR	η	Cost	SM
1	2.281	0.959	255.879	0.110
2	2.221	0.966	305.054	0.219
3	2.152	0.970	320.646	0.318
4	2.320	0.977	288.355	0.228
5	2.270	0.979	302.869	0.283
6	2.191	0.979	220.188	0.274
7	2.306	0.979	295.715	0.304
8	2.236	0.982	239.556	0.299
9	2.181	0.984	254.300	0.360
Overall means	**2.240**	**0.975**	**275.840**	**0.286**

Table 1

The next step consists of the calculation of the individual absolute effects of each of the factors.

To this end, the *percentage* effects are calculated first, simply by taking the overall means as reference values. As the percentage effect on cost was, in magnitude, much larger than the effects on efficiency and pressure ratio, it would have been naïve to add up the percentage effects without some sort of weighting, since clearly a 1% *delta* in efficiency is far more important than a 1% *delta* in cost, the same being true for the efficiency.

Therefore, an appropriate weighting distribution was adopted whereby the resulting effects all had the same order of magnitude, and therefore could be added up.

The effects of all the design variables are presented in the following table.

A - $\beta_{mi,R}$ (deg)		Effect	C - C_r (m)		Effect
1	60	-1.33	1	0.06	0.85
2	63	0.71	2	0.08	-0.10
3	66	0.62	3	0.1	-0.58
B - $\beta_{mo,R}$ (deg)			**D - N_r**		
1	33	1.06	1	67	0.15
2	35	-0.06	2	72	0.02
3	37	-1.00	3	77	-0.17
E - $\beta_{mi,S}$ (deg)		**Effect**	**G - C_s (m)**		**Effect**
1	56	0.07	1	0.06	-0.51
2	59	-0.17	2	0.08	0.74
3	62	0.10	3	0.1	-0.23
F - $\beta_{mo,S}$ (deg)			**H - N_s**		
1	8.66	0.76	1	65	0.32
2	10.66	0.02	2	70	-0.19
3	12.66	-0.78	3	75	-0.13

Table 2

At this point, we have a cumulative effect for each variable setting and the level producing the most favourable impact is chosen as the initial point for the numerical optimisation. It is interesting to notice that the rotor exit metal angle has the largest effect on the overall design performance.

The recommended, initial design to be fed into the full optimisation, therefore, is summarised below:

Symbol	Units	Initial value for numerical optimisaion
$\beta_{mi,R}$	[°]	63
$\beta_{mo,R}$	[°]	33
c_R	[m]	0.06
N_R	[-]	67
$\beta_{mi,S}$	[°]	62
$\beta_{mo,S}$	[°]	8.66
c_S	[m]	0.08
N_S	[-]	65

Table 3

This initially recommended set of design variables is run in the simulation model yielding the following results:

Pressure ratio: 2.316	**Efficiency: 0.975**
Total Cost: 219.251	**Surge Margin: 0.164**

This is a reassuring result, because even though this particular combination had not been run before, it produces better pressure ratio, efficiency and cost than the average of all experiments.

6 – Single-objective gradient-based optimisation and initial sensitivity analysis

6.1 SQP optimisation

SQP is chosen among various gradient-search techniques for this first stage of single-variable optimisation. The following points all favour the use of SQP over other Newton methods.

- Although the number of blades has a discrete magnitude, it is fixed to the value obtained from DOE for now. Hence, all variables are real and continuous.
- The design space is highly nonlinear, which is obvious from the various theoretical models used in the code and also the empirical correlations used in the aerodynamic loss prediction and cost modelling, none of which is by far linear.
- The problem is constrained.

The pressure ratio is selected as the first objective function to optimise, since:

- The pressure ratio is the most significant magnitude related to performance.
- Efficiency is seen to stay within very acceptable ranges in the DOE experiments, so the push for pressure ratio appears more attractive.
- The cost and surge margin seem natural objectives to constrain by experience and legislation, respectively.
- On the other hand, one does not have a priori knowledge of the aerodynamic potential of the machine and a single optimisation of the pressure ratio is the adequate way to explore that.

The remaining objectives are constrained to pre-determined values or bounds as detailed below:

- Efficiency: Higher than 0.96
- Cost: Lower than 200 units
- Surge margin: Higher than 0.2

(20% being around the widely used value in industry)

Having set up this optimisation within the *iSight* commercial package, the SQP method is found to converge to the solution shown in Table 4.

The result is markedly better than the initial value in all respects. However, it cannot be ascertained that this is the global optimum (in fact, a trial run with a Genetic Algorithm was carried out from the same initial data and a *PR* of 2.470 was achieved, which clearly shows that we are still far away from the ideal target).

It is interesting to note that the stator inlet metal angle was not modified by the optimiser in order to achieve this superior performance.

The chords were not altered either, since they only impact the cost and their inital values made it stay within bounds for all the combinations tried.

	Initial	SQP optimised
$\beta_{mi,R}$	63	70.002
$\beta_{mo,R}$	33	29.469
c_R	0.06	0.06
N_R	67	67
$\beta_{mi,S}$	62	62.001
$\beta_{mo,S}$	8.66	8.7075
c_S	0.04	0.04
N_S	65	65
PR	2.317	2.412
η	0.976	0.978
Cost	191.095	180.38
SM	0.164	0.230

Table 4

6.2 Sensitivity analysis

Having established the optimum *PR*, each of the eight variables is increased by 5% of its nominal value and the output from the simulation code is recorded in order to calculate the sensitivities.

The absolute variation in pressure ratio is computed. This allows for a numerical estimation of the derivatives $\partial PR/\partial\beta_{in}$, $\partial PR/\partial\beta_{out}$, etc. These are normalised with respect to the optimum design vector and pressure ratio and the following sensitivity graph is obtained:

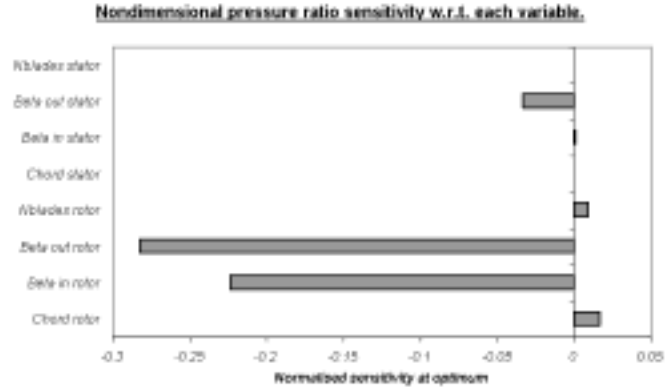

Figure 6

It is reassuring to see that the aerodynamically influential parameters are the most significant ones in this sensitivity analysis, given that we were only trying to optimise pressure ratio. The rotor metal angles were the ones with highest impact, as physically expected.

7 – Scaling

The next step of our optimisation strategy is to compute the Hessian matrix so as to identify the necessary scaling to apply to our design variables. The pressure ratio is considered, again for now, as the only objective function.

To compute the diagonal members of the H-matrix, standard finite difference formulations are employed.

$$\frac{\partial^2 PR}{\partial x_i^{\,2}} = \frac{PR(x+\Delta x) - 2PR(x) + PR(x-\Delta x)}{\Delta x^2}$$

After several runs that enable us to compute the finite differences, the diagonal elements of the Hessian are found to be, at our current optimal solution:

$$H = \begin{bmatrix} -0.0019591 & & & & & & & \\ & -0.0018424 & & & & & & \\ & & 0 & & \ddots & & & \\ & & & 0 & & & & \\ & & & & 0.0204015 & & & \\ & & \ddots & & & 0 & & \\ & & & & & & 0 & \\ & & & & & & & 0 \end{bmatrix}$$

Scaling is then carried out in the following way: if any of the entries computed in the Hessian diagonal is greater than 10^2 or less than 10^{-2}, the corresponding design variable needs to be scaled.

If $H_{ii} \sim 1/x_i^{\,2}$, that is to say if $H_{ii} \sim 10^{-4}$, then the appropriate scaling to be applied is $X_i \sim 10^{-2} x_i$. Using these rules, the scaling required for each design variable is computed to make $H_{ii} \sim O(1)$. Only scalings of the form 10^{-2}, 10^{-1}, 10^{1}, 10^{2},… are considered (i.e. we only worry about the order of magnitude).

From this analysis, it turns out that the rotor inlet and exit metal angles, together with the stator inlet metal angle, all need to be scaled by a factor of 10^{-2}.

It is also seen that all the other variables return second derivatives of 0. This is expected for the blade chord and count, which do not directly affect the pressure ratio in our model, but do impact the other magnitudes of interest – cost, efficiency and surge margin.

Having modified the code to allow for the scaling factors of the aforementioned three angles, the SQP optimisation is re-run in *iSight* starting from the previous optimal solution. A much better optimal solution is attained (see Table 5).

Design variable	Initial values	SQP before scaling	SQP after scaling
$\beta_{mi,R}$	63	70.002	0.6993 (x100)
$\beta_{mo,R}$	33	29.469	0.27066 (x100)
c_R	0.06	0.06	0.06
N_R	67	67	67
$\beta_{mi,S}$	62	62.001	0.7 (x100)
$\beta_{mo,S}$	8.66	8.7075	8.7075
c_S	0.04	0.04	0.04
N_S	65	65	65
PR	2.317	2.412	2.463
η	0.976	0.978	0.975
Cost	191.095	180.38	167.737
SM	0.164	0.230	0.201

Table 5

The performance gain achieved simply by scaling our design variables is very significant (2.11% in *PR*). The efficiency, cost and surge margin all lie within the bounds specified at the onset. The cost is actually much better than the previous best result and the surge margin has now become an active constraint, which was not the case in the unscaled SQP optimisation.

Appropriate scaling, therefore, proves extremely beneficial to the gradient-based optimisation efforts. This scaling is to be used throughout all the remaining steps of our optimisation strategy.

8 – Final Single-Objective Heuristic Optimisation

8.1 Genetic Algorithm Approach

A genetic algorithm is now used to perform a more detailed heuristic optimisation and to determine how much improvement is left on our previous optimum.

A parametric study is performed by means of several GA runs with different mutation rates, population sizes and numbers of generations. The best one is shown below:

Rotor chord	0.05999
Rotor Beta in (x100)	0.69937
Rotor Beta out (x100)	0.28565
Nblades rotor	69
Stator Chord	0.0496
Stator Beta in (x100)	0.6801
Stator Beta out	4.49175
Nblades stator	69
Pressure Ratio	2.475
Efficiency	0.977
Cost	182.749
Surge Margin	0.2
Population size	40
Generations	40
Crossover rate	1
Mutation rate	0.04

Pressure ratio: 2.475

Efficiency: 0.977

Total Cost: 182.749

Surge Margin: 0.200

Table 6

8.2 Some caveats on single-objective optimisation

The genetic algorithm returns a much better answer than the best scaled SQP. Almost a 3% improvement is realised. However, the case for multiobjective optimisation is clearly made – while the pressure ratio is maximised and the surge margin has now become a hard constraint, the total cost of the compressor is significantly higher than the previous one.

A simultaneous *PR* maximisation and cost minimisation is called for. This will be done eventually. At this point, it would be arguable, in an industrial context, whether one would want to stick to the GA design or sacrifice the 3% increase in performance whilst remaining at a cost of 167.737 as opposed to 182.749.

9 – Bi-Objective Optimisation and Pareto Front Generation

Now that the target pressure ratio level is known to be around 2.45, we cease to treat *PR* as an objective and make it a constraint. From here onwards, all feasible solutions must produce a *PR* in excess of 2.4.

Two clearly conflicting objectives are the surge margin and the cost. In order to increase the solidity one has to either use more blades or make their chord larger (or both), which invariably leads to higher expenses.

Before proceeding to the final multi-objective analysis, a bi-objective optimisation with *SM* and *C* is set up to assess how much dynamic stability one has to give up to achieve a certain cost target, or viceversa.

This procedure is carried out by a first-order weighted sum method. The scaling factors used are $C_o=1.5x10^2$ and $SM_o=2x10^{-1}$, which were roughly the baseline values from the best GA solution.

The objective is therefore formulated as:

$$J = \lambda\left(\frac{SM}{SM_o}\right) + (1-\lambda)\left(\frac{C}{C_o}\right)$$

Various values of λ are swept from 0 to 1 and the same SQP algorithm is employed to optimise each weighted function, yielding the Pareto front.

It is important to notice that the points in the Pareto front appear to be very localised in three discrete regions. This probably points to a non-convex topology.

In order to ascertain this, very small λ intervals are tried between the values at which the design *jumped* from one discrete region to the other. It is found that this change happens across a very small interval in λ - one is tempted to say that there is probably *one* limiting λ point between each region of the objective domain, but this cannot be proved mathematically here.

The results of all the runs with different values of λ are summarized in the table below and the Pareto front itself follows it.

λ SM	0.900	0.750	0.600	0.500	0.450
λ Cost	0.100	0.250	0.400	0.500	0.550
SM	0.251	0.251	0.250	0.250	0.248
Cost	169.844	169.608	169.252	168.895	168.356
PR	2.405	2.405	2.405	2.405	2.405

λ SM	0.400	0.350	0.335	0.325
λ Cost	0.600	0.650	0.665	0.675
SM	0.248	0.248	0.247	0.203
Cost	168.356	167.969	167.830	151.129
PR	2.405	2.405	2.405	2.405

λ SM	0.310	0.300	0.200	0.100
λ Cost	0.690	0.700	0.800	0.900
SM	0.202	0.199	0.172	0.171
Cost	149.794	149.715	147.874	148.000
PR	2.400	2.404	2.409	2.440

Table 7

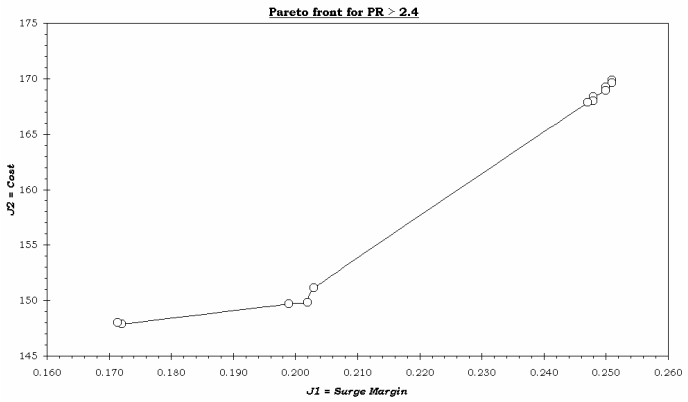

Figure 7

The Pareto front itself shows the opposing character of the two objective functions – not numerically, but in terms of design "goodness". The utopia point is located at the bottom right hand corner of the plot.

Out of the three designs, one must decide which one to take forward to the multi-objective arena. A surge margin of around 17% is unacceptable by existing CAA standards and, while a value of 25% would be very attractive, the cost incurred would be clearly excessive.

Therefore, the design point corresponding to $\lambda_{SM} = 0.3$ and $\lambda_C = 0.7$ is selected as the starting point for the following optimisation.

10 – Total Multi-Objective Optimisation

The summit of this sequence of steps is the multi-objective optimisation of performance, cost and stability at the same time. All the previous steps serve to progressively approach the region in the design space where this final effort is to be tackled.

For this final step, a GA is used, with a view to maximise *PR* and η, minimise *C* and keep *SM* above 20%. The same GA tuning parameters arising from the parametric study from Section 8 are used in this final computation.

A summary of all the results presented thus far is given in the table below.

	After DOE	Single-objective		
	Initial values	**SQP unscaled**	**SQP scaled**	**GA scaled**
$\beta_{mi,R}$	63	70.002	0.6993 (x100)	0.69937 (x100)
$\beta_{mo,R}$	33	29.469	0.27066 (x100)	0.28565 (x100)
c_R	0.06	0.06	0.06	0.05999
N_R	67	67	67	69
$\beta_{mi,S}$	62	62.001	0.7 (x100)	0.6801 (x100)
$\beta_{mo,S}$	8.66	8.7075	8.7075	4.49175
c_S	0.04	0.04	0.04	0.0496
N_S	65	65	65	69
PR	**2.317**	**2.412**	**2.463**	**2.475**
η	**0.976**	**0.978**	**0.975**	**0.977**
C	**191.095**	**180.38**	**167.737**	**183.571**
SM	**0.164**	**0.230**	**0.201**	**0.200**

	Bi-objective	MOO
	Pareto optimal	**GA MOO**
$\beta_{mi,R}$	0.7 (x100)	0.69956 (x100)
$\beta_{mo,R}$	0.31031 (x100)	0.280582
c_R	0.04823	0.05037
N_R	69	68
$\beta_{mi,S}$	0.7 (x100)	0.69999 (x100)
$\beta_{mo,S}$	4.5055	11.247509
c_S	0.025	0.0250167
N_S	69	63
PR	**2.404**	**2.406**
η	**0.972**	**0.972**
C	**149.715**	**144.515**
SM	**0.199**	**0.199**

Table 8

11 – Summary & conclusions

11.1 Optimisation results and discussion

A systematically progressive approach to the optimisation problem has proved successful in arriving at a design that is far superior to the initial machine, while at the same time providing physical insight on the compressor features that most influence each objective function.

After the first run of calculations in single-objective mode, the results are already very encouraging, in a real engineering sense, since an overall 7% improvement is achieved from the baseline design. It seems clear that the relative improvement over the different methods becomes progressively smaller – 4.1% from initial to unscaled SQP result, 2.2% from unscaled to scaled SQP and 0.5% from scaled SQP to GA.

Although the best SQP solution is not the global optimum, the solution seems to asymptote towards the 2.48 – 2.5 limit. Similarly, the blade chord dimensions all converge to the upper bounds, as do the blade counts. These all contribute to forming a high solidity compressor, which is expected if cost considerations are not taken into account. Hence, this is taken as the performance utopia for this compressor.

It should be said that, in reality, one does not expect a 2.5 *PR* on two stages. This machine is very highly loaded and the loss buckets for the profile losses should actually have a much larger bias, which was not input into our model. Apart from this scaling difference, all the arguments used to make decisions throughout the optimisation are equally valid.

Aside from the mere performance boost in single-objective mode, there is an improvement in cost, especially for the scaled-SQP result, which has a cost that is 12% lower than the post-DOE design. However, the GA returns an even higher *PR* and η, but a higher cost, which relates to a larger solidity blading. This illustrates the beauty of the heuristic methods.

The SQP search must have stopped at an extremum, which is obviously not the global optimum. By increasing the diversity in the population, the GA is capable of reaching another peak. Even though the performance is marginally better (0.5% better in PR and 0.2% better in η), the blade counts and chords have been pushed to their upper bounds, achieving the maximum solidity. This enables a high pressure ratio with a smaller blade turning on the stator. Because the diffusion pressure losses are reduced with the higher solidity, the loading on the stators can be relaxed for the same pressure ratio.

The trade-off here is evident. While the GA run gives mathematically the best SOO solution, no real-world engineer would choose a design that is superior to the scaled SQP by 0.5% in performance but incurs 9.5% higher costs. One has no way of knowing the potential reduction in costs that could be achieved while staying within acceptable performance bounds. The case for multi-objective optimisation is made transparent.

The opposing nature of *SM* and η, established in section 9, was basically explained by the fact that a higher solidity increases stability and at the same time requires a higher number of larger blades. It seems natural therefore to

initiate the multi-objective analyses using these two to generate a meaningful Pareto front. One could argue that efficiency seems to be left out of the picture until the very final multi-objective run. However, it is seen that within the design region under consideration, the maximum variation in efficiency is around 0.6%, even between the cases with large difference in *PR*. Therefore, since the single-objective optimisation led us to zero-in on a narrower region of the design space, the efficiency is expected to remain fairly constant, which justifies the choice of *SM* and *C* as the two opposing functions for the Pareto optimisation.

The physical reason behind these small variations in efficiency is the same that explains the unusually high levels of *PR* – the low value of the profile losses. The appropriate bias for the highly nonlinear profile loss effects at high blade turning and loading have not been modelled in detail within the simulation code. As said before, this leads to numerically counter-intuitive answers, but since this is a conceptual optimisation, it does not in any way deter the argument.

The highly discrete nature of the Pareto front, being divided into three distinct families of designs, shows that only a few configurations are capable of achieving the high performance targets set from the single-objective analyses. The design in the middle region, around the 20% *SM* point, has been chosen to comply with currently established stability safety margins. However, if the stability constraint could be relaxed somewhat, a further significant reduction in cost could be attained. Similarly, if the legislation were to impose a slightly higher *SM* requirement, the costs could escalate rapidly.

Once a selection is made from the Pareto front, a last run with all four objectives being optimised simultaneously is performed as a final refinement of the solution. As expected, not much variation is recorded, which is reassuring. The pressure ratio, efficiency and surge margin all stay virtually at the same level but an additional cost reduction of 3.5% is achieved.

Design change visualisation

A graphical representation of the geometric changes performed on the blades is useful to illustrate the profound effect of the optimisation process on the design in a single plot. This is given in figure 8.

Figure 9 shows the aerodynamic characteristics of the final design. The non-dimensional $\psi_{TS} - \phi$ characteristic depicts the process followed by the simulation code, in which the mass flow is reduced until the maximum in ψ_{TS} is reached, at which point the compressor is considered to be at neutral stability. The equivalent point and the *SM* are plotted in the dimensional graph.

It should be pointed out that, even though the inception of instability coincides with the extremum of the nondimensional characteristic, this is not true for the dimensional *PR* – *m* curve and stable operation is actually possible to the left of the *PR* peak.

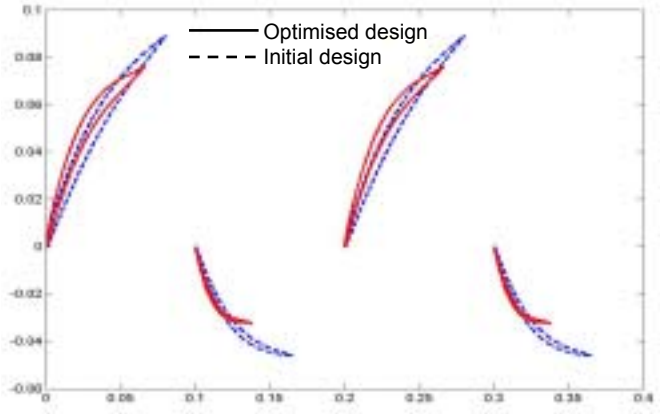

Figure 8

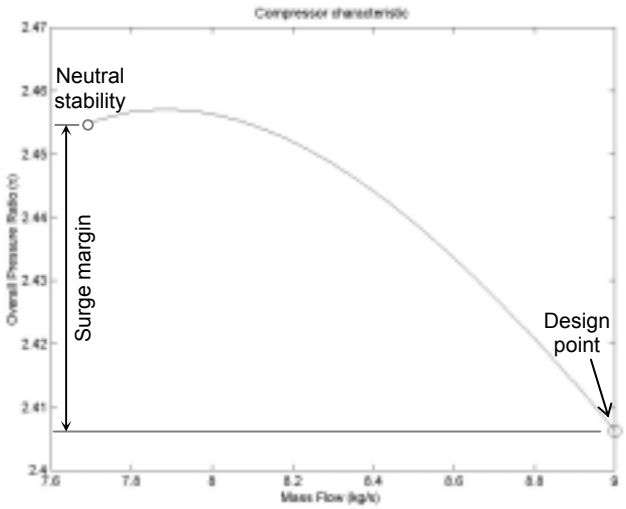

Figure 9

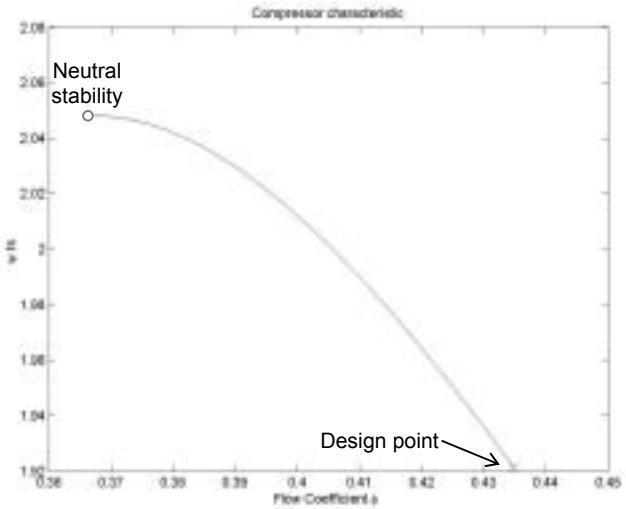

Figure 10

11.2 Motivation for the development of more accurate stability metrics and coupling with optimisation techniques

Having illustrated the deep impact of the *SM* constraint on the cost of optimal designs, the need for a good understanding of the mechanisms leading to unstable compressor operation is easily shown. Surge margin is an almost archaic metric for compressor stability based on keeping a *safe operating distance from the extremum of the characteristic*.

However, state-of-the-art dynamic models for axial flow compressors are capable of predicting the exact degree of dynamic stability in terms of more analytic eigenvalue-type formulations. Development of disturbance-energy methods, based on the accurate pressure and velocity disturbance modeshapes of the 2D compressor flowfield are even capable of providing accurate disturbance-energy distributions along the machine, thus pointing out to designers the areas of the compressor that most contribute to the overall disturbance-energy balance, and therefore need redesign.

The whole point of this ongoing research is to be able to reliably assess the dynamic characteristics of axial-flow compressors, in order to abolish surge-margin as a stability metric, which is much too conservative.

It has been shown already that compressors having the same dynamic stability characteristics (critical mode damping, etc.) as calculated by accurate models may have very different surge margins. The figure below illustrates this.

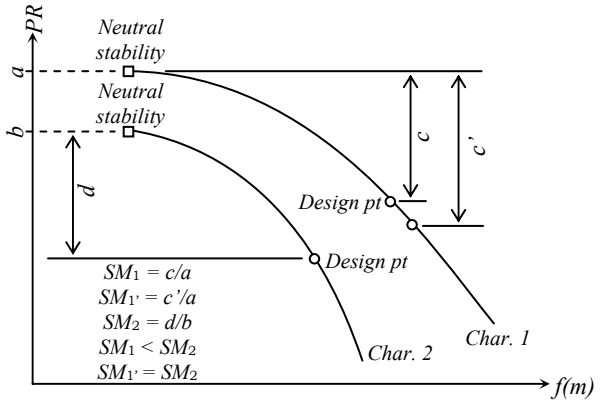

Figure 11

Consider the case where the surge margin of Characteristic 1 is smaller than that of Characteristic 2. By careful dynamic modelling, the two design points are found to have the same critical mode damping, so they are equally stable, despite the fact that they have different surge margins.

If one was to make design 1 comply with the same surge margin level as design 2, the operating point should be driven to c', which has a much lower performance and is overly stable.

As well as permitting more aggressive designs, dynamics-based stability metrics effectively mean that the artificial surge margin concept can be abandoned, and compressors can be pushed to smaller surge margins without compromising operational safety. As seen in the Pareto front generated for this compressor, this immediately translates into reduced costs.

The development and implementation of these new dynamic models is therefore very appealing. New design methodologies incorporating a systematic use of optimisation algorithms and accurate stability predictions will allow the systematic development of inherently stable machines with optimum aerodynamic/structural performance and minimum production costs. Once legislation changes are in effect that permit the utilisation of the more precise stability metrics, such tools will revolutionise compressor design.

12 – References

[1] Greitzer, E.M., 1976a, "Surge and Rotating Stall in Axial Flow Compressors, Part I: Theoretical Compression System Model", ASME *Journal of Engineering for Power*, Vol. 98, pp. 190-198.