Robust Aerodynamic Design Optimisation of Gas Turbine Compression Systems

by

Tiziano Ghisu, Geoffrey T. Parks
Jerome P. Jarrett and P. John Clarkson
Engineering Design Centre
Department of Engineering
Univeristy of Cambridge
Cambridge, CB2 1PZ, England
e-mail: tg269@cam.ac.uk
web page: http://www-edc.eng.cam.ac.uk/

A common simplification in the design of any engineering system is to assume nominal values for a number of design parameters (ranging from operating conditions to manufacturing tolerances, damage processes and so on). While these assumptions are often essential to make the design process tractable, they may lead to products that in real-world environments exhibit significant deviations from their nominal performance. This problem is often exacerbated in heavily optimised products that tend to lie in extreme regions of the design space.

For gas turbine engines (and compression systems more specifically) the nominal design point is usually chosen to be the point with the severest conditions or the strongest influence on the figures of merit: this is generally a high speed condition with high non-dimensional mass-flow and rotational speed (often maximum climb). However, as the engine, during its normal operation, is expected to work over a significant range of power settings (such as engine starting, idling, reduced power, maximum power, acceleration, deceleration), satisfactory performance and safe operations must be guaranteed for all the possible working conditions. Off-design performance is usually assessed after the preliminary design (the more the project is advanced, the more complex it is to modify any fundamental design choices).

This approach, while guaranteeing that some minimum off-design performance requirements are met, introduces a further loop in the design process, hence adding a further level of hierarchicality and fragmentation. Moreover, there might be a substantial difference between satisfactory and optimal performance and a heavy optimisation of some design-point figures of merit might actually be detrimental to the overall performance: a tighter integration of some robustness considerations into the design process can reduce the level of decomposition typical of gas turbine design and, potentially, produce a design solution with improved performance over a range of operating conditions rather than at a particular (even if the most important) operating point.

In this work, two methods for dealing efficiently with off-design performance within the optimisation process are suggested: interval analysis, which is the only sensible approach when limited information about engine (and compressor) usage is available, and Polynomial Chaos, which offers a convenient means for producing high order information – mean, variance, skew and successive moments – at a sensibly reduced cost compared to Monte Carlo simulations, which are the only approach ensuring the convergence of the calculated moments to their effective values, but suffer from a very slow convergence rate. A method for dealing with generic PDFs more efficiently is also suggested: orthogonal polynomials are calculated based on the given input PDF, with considerable computational savings relative to the usual approach that makes use of a standard basis of orthogonal polynomials (Hermite, Legendre, Jacobi, etc.).

The results demonstrate how consistent improvements in overall performance can be achieved with both methods: with interval analysis, this requires the introduction of a new figure of merit (unless some approach for condensing the two objectives is used); Polynomial Chaos, while being computationally more expensive than interval analysis, allows the mean performance to be set as the design objective, leading to an optimisation problem equivalent to the non-robust one. A level of mean performance improvement of the same order as the design-point performance improvement obtained with the standard approach can be achieved through a robust optimisation, with the substantial differences that the off-design performance analysis is already embedded inside the optimisation loop and that more robust designs are sought rather than just found as a result of the non-robust design – the difference in mean performance between the two approaches is, for this particular problem, of the order of 0.5% in isentropic efficiency. The clear downside is the substantial increase in computational costs, dictated, more than by the larger number of evaluations needed, by the significantly larger time required by the off-design evaluation tool compared to the design-point one. It must however be remembered that, with off-design performance embedded into the optimisation, the risk of the results being unsatisfactory is substantially reduced, if not eliminated.