11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference
6 - 8 September 2006, Portsmouth, Virginia

AIAA 2006-7003

Robust structural design of a simplified jet engine model, using Multiobjective optimization

Ivan Voutchkov[1]*, Andy J. Keane[1], Rob Fox[2]
[1]CEDG, School of Engineering Sciences, University of Southampton, Highfield, Southampton, United Kingdom, SO17 1BJ, *e-mail: iiv@soton.ac.uk
[2]Rolls Royce Plc, Derby DE21 7XX, United Kingdom

Abstract

This study demonstrates advances in multiobjective optimization, supporting a robustness study of a simplified jet engine structural model. The ultimate goal is to find the best structural configuration of shell thicknesses along the engine that will result in reduced reaction force variation under a range of external loads, will be as light as possible and where fuel consumption will be minimal. These are competitive objectives, some of which are of stochastic rather than deterministic in nature. The paper demonstrates that a deep level multiobjective search pays off many times the investment in time and money by providing significant design improvement.

Keywords: multiobjective optimization, robust design, whole engine model, NSGA2

1. Introduction

Until recently it was normal in engineering design to simply use the most important design goal as a single objective during optimization. For instance in aerospace engine design a typical design objective would be to 'make the engine as powerful as possible whilst trying not to exceed the mass limits set by the structural division'. The thrust of the engine would be the primary goal and amongst all design ideas generated, those that exceed the maximum mass would simply be filtered out. After trying to match a few other considerations, such as fuel consumption, overall engine stiffness and internal load distribution, there would be very few designs left to choose from, in some cases none. The design would then be taken back to the drawing board and the whole process started again. Often future improvements were fundamentally based around a design known to work well from previous experience. Many regions of the design space were left unexplored. The generation of design ideas was an *active* process but due to lack of proper tools to direct these ideas they were perhaps somewhat random when looked at over the design space. Such a design process could be classified as *passive*, because the final design was often a product of satisfaction.

With advances of multiobjective algorithms the aim now is to produce a set of designs which in each case are better than the rest for at least one criterion. Such a set is called a Pareto front and is known to be the set of non-dominated solutions in the objective space. Engineers and designers could then sit together and discuss each of these designs, compare different requirements and select the best trade-off. This allows several design goals to be considered simultaneously and actively searched, aiming to obtain as many designs as possible which are evenly distributed and widely spread around the objective space. Multiobjective algorithms can be classified according to these three criteria which define the quality of the Pareto front. As in single objective methods, various algorithms perform differently for various objective functions. There are number of methods published and widely used to do this – MOGA, SPEA, PAES, VEGA, NSGA2, etc. Some are better than others - generally those most popular in the literature are NSGA2 (Deb) and SPEA2 (Zitzler), because they are found to achieve good results for most problems [4-11]. This paper will present the results obtained using NSGA2. It is a well known algorithm, which has been further modified here to obtain improved performance.

2. Surrogates

In real engineering problems the cost of evaluating a design is probably the biggest obstacle that prevents extensive use of optimization procedures. In the multiobjective world the cost is multiplied because there are multiple expensive results to obtain. Evaluating directly a non-linear finite element model can take several days, which makes it impossible to try hundreds or thousands of designs.

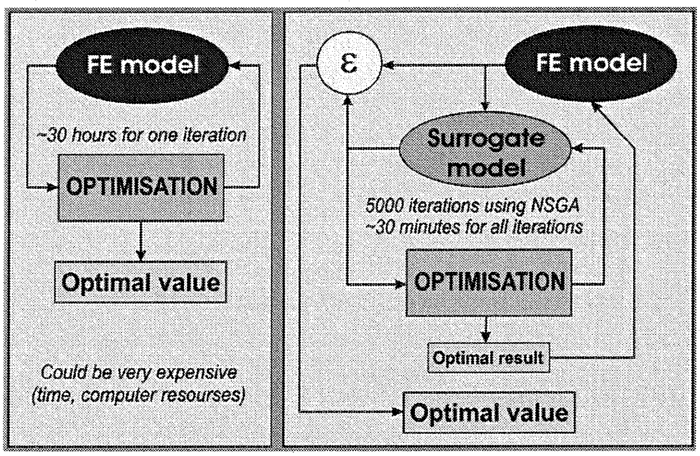

Figure 1 – Surrogate models for optimization

In the single objective world, approaches using surrogate models are fairly well established and have proven to successfully mitigate the problem of computational expense (Figure 1). The key idea is to build a computationally cheaper model, that can be extensively searched instead of the expensive one. When an optimal solution is found, it needs to be verified on the expensive model and then added to the data pool on which the cheaper model is being trained. The procedure iterates until a good agreement between the two models is achieved or an acceptable design is found. More and more companies have adopted surrogate assisted optimization techniques and some are making steps to incorporate this approach in their design cycle as a standard. This makes optimization not only useful, but usable and affordable.

There are a number of approximation and curve fitting methods that can be used as surrogates. Polynomial interpolation, cubic splines, neural networks, fuzzy logic, kriging – all of which offer different level of calculation cost vs. accuracy. Choosing an approximation method can be a multiobjective optimization task in its own right. The ideal choice would be the least expensive method that can best fit the FE output from Figure 1.

Bearing in mind the complexity of the problem considered, kriging is the method of choice for this paper. It is well described elsewhere in the literature [2] and so only limited details are presented here.

3. Robust design

A design is considered robust towards external conditions if it is insensitive to the variation of these conditions. A commonly used robustness measure is the standard deviation of the design criteria. In the past, companies tended to use optimization to obtain the best nominal value of performance metrics. However the resulting designs often turned out not to be the most robust. A slight variation in the design variables or external factors would lead to a significant variation in the performance value. Commercially, a more viable product is one that exhibits robustness towards environment or operating variation. Figure 2 illustrates that the robust design is not necessarily the one having the best nominal value.

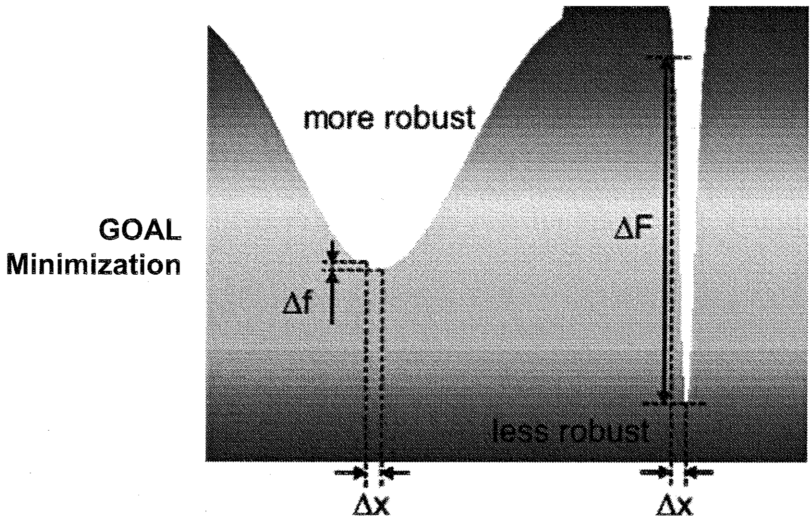

Figure 2 – Robust design

Mathematically robustness is measured in terms of standard deviation:

$$s = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(Obj(x_i) - MeanObj\right)^2} \tag{1}$$

where

$$MeanObj = \frac{1}{n}\sum_{i=1}^{n} Obj(x_i), \tag{2}$$

n is the number of pseudo random variations around a given design point and $Obj(x_i)$ is the value of the objective function for the design x_i. In reality we deal with statistical estimates of standard deviation. The estimate approaches the real value as n approaches infinity. That is why the greater is n the more accurate are the predictions, however the higher the computation costs. Later in this paper the sensitivity of these estimates to the number of variations for the particular physical quantities under consideration will be shown.

In the world of real engineering, it is desirable to search not just for a robust design but for an optimal robust design. However often optimality criteria such as (3) are in competition with robustness criteria such as (4). Therefore an optimal robust design would be a member of a non-dominated solution set, i.e from the Pareto front, for the two objectives.

$$\min\{\text{abs}(Obj(x_0) - TargetObj(x_0))\} \tag{3}$$
$$\min\{s\} \tag{4}$$

Here $Obj(x_0)$ is the nominal value of the objective function. As optimal robust design is based on stochastic process it is widely accepted in practice to use the mean of the objective function with regards to the variable variations, so that criteria (3) is transformed to

$$\min\{\text{abs}(MeanObj - TargetMeanObj)\} \tag{5}$$

4. Simplified shell model of a jet engine and optimization setup

The structural model used in this paper is shown in Figure 3. It is a simplified representation of an entire jet engine. This model was produced for use within the EC funded Vivace project. The engine has been divided into 22 super elements, shown in different colours (shades of gray) and fifteen of these were found to be significant for this study. The structure is modelled in MSC NASTRAN and the casing thicknesses within each super element are defined by a thickness value in the corresponding PSHELL cards. This allows automatic variation of the thicknesses by the optimization algorithm.

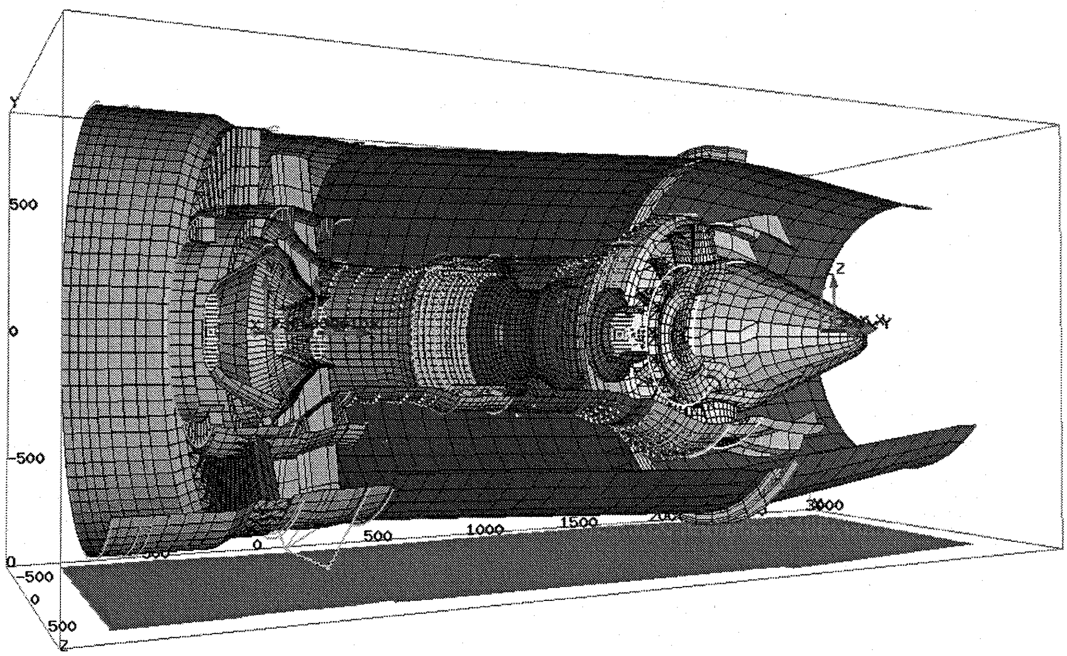

Figure 3 – Simplified Jet Engine model

For each set of 15 thicknesses 200 load variations were performed on the model using multiple right hand side computations in order to minimise the computation time. These 200 runs are a predefined set of loads, selected using a Latin Hypercube algorithm [12], and allow the variance of the reaction forces to be estimated. The following quantities are defined as objective functions:

Objective 1 –	**Minimum** {Standard Deviation of the internal reaction forces over the 200 load variations STDEVFMPX} – minimization of this function would produce the most robust design;
Objective 2 –	**Minimum** {Mean value of the internal reaction forces over 200 load variations, MEANFMPX} – minimization of this function will produce a design with least reaction forces;
Objective 3 –	**Minimum** {Mass} – This will produce the lightest design;
Objective 4 –	**Minimum** {Mean value of the specific fuel consumption, SFC} – this will give us the most economical engine.

Objectives 1 and 2 depend directly on the load variations and indirectly on the thicknesses, Objective 3 depends solely on the thicknesses, while Objective 4 depends directly on the relative movement of the rotors to the skin of the engine and therefore indirectly on load and thickness variations.

An idealized illustration of the effect of thickness on the four objectives is shown in figure 4. It shows that it might be possible to have a robust solution (minimum MEAN and minimum StDev) without compromise. This will occur when using large values for thickness. This is a logical observation, as thicker walls will make the entire structure stiffer – therefore leading to low variations and low reaction forces.

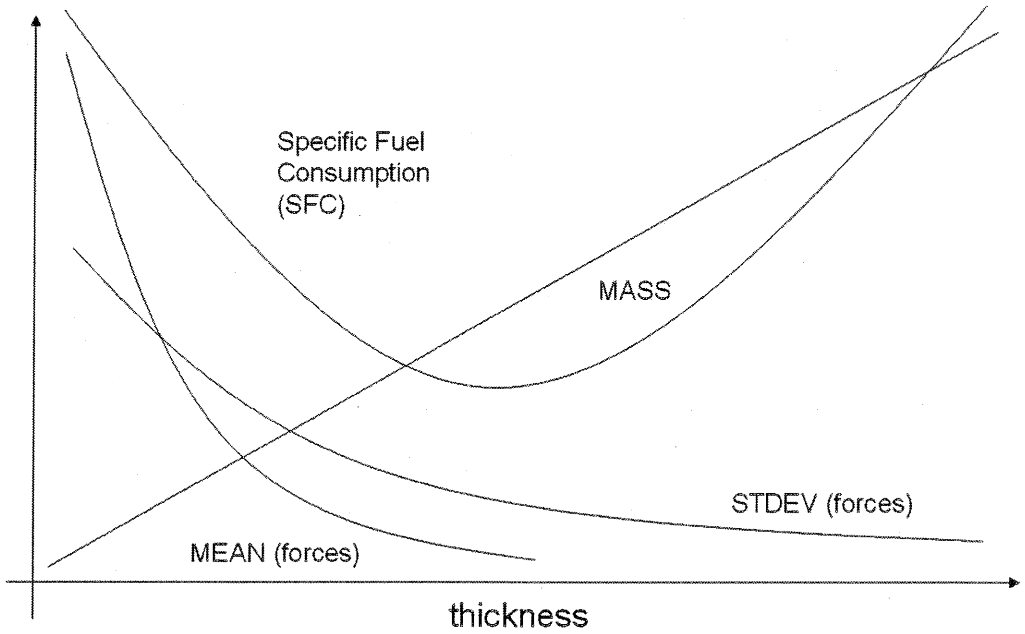

Figure 4 – Principal effect of thickness variation on MASS, SFC, Mean and Stdev, described in this paper

Considering the other two objectives, high values for thickness would increase the mass of the engine whereas fuel consumption will decrease as the engine structure becomes stiffer. This is due to the ability to run with reduced rotor tip clearances. The tip clearances are affected by the external load variations and also affect the fuel consumption. Note also that as the engine becomes heavier with increasing structural thickness, the fuel consumption starts to grow again as the engine has to carry more of its own weight.

It is clear that the task is multiobjective with a relatively high number of dimensions. This is why this particular robustness study was selected. Further degree of complexity could be achieved by introducing uncertainty in the thickness variables themselves, however this will not be a subject of this paper.

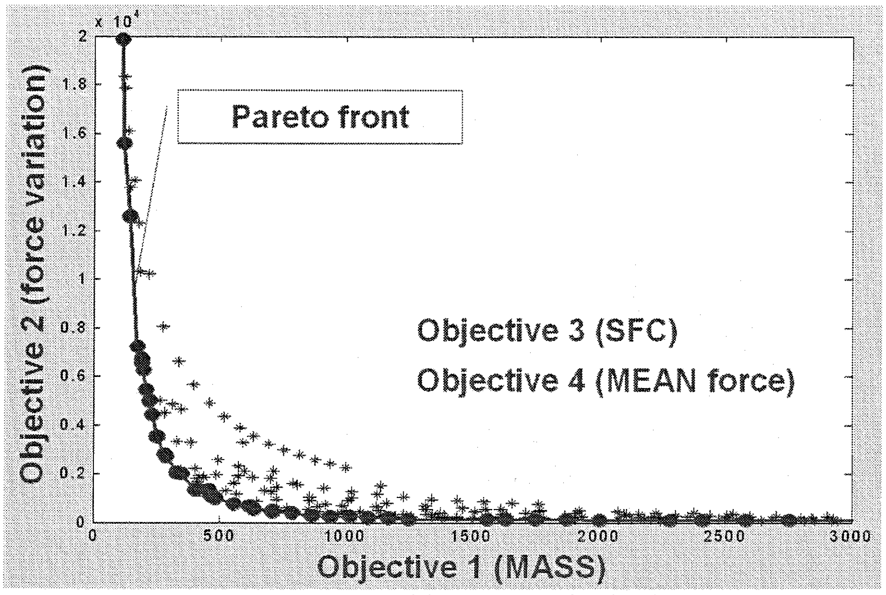

Figure 5 – Competitive nature of objectives

When dealing with multidisciplinary optimization we obtain the solution in the form of a set of optimal designs called a Pareto front, as opposed to the single objective world where the optimal solution is a single design point. The Pareto front is also defined as the set of non-dominated solutions, i.e. there are no better solutions in respect to any of the objectives considered that are better than those on the Pareto front. For example these are represented with circles on Figure 5. All other solutions are at least in respect of one objective worse – shown as star symbols.

5. Estimation of the robust design measure of merit

As mentioned earlier in this paper, mean and standard deviation of the reaction forces subject to external (or internal) variations are used as measures of merit for optimal robustness. It should also be mentioned that the accuracy of these values depends on the number of variation samples. In this section we demonstrate this and justify the selection of 200 variations.

Figure 6 shows the value of the standard deviation and mean over a varying number of points. The data for Figure 6 is associated with a randomly taken design point, after verifying that the general trend of the illustrated curves remains the same over all designs. The estimates are calculated using (1) and (2).

The figure demonstrates graphically that any number of variations less than 200 would yield a rather unsettled value for the estimates. On the other hand the improvement of accuracy due to further increases of stochastic variations would not justify the increased computation costs. It should also be mentioned that as the variation applied in the current task are external to the model variables it is possible to do 200 "right hand side" substitutions instead of 200 complete solves of the FEA, this significantly increases the efficiency of obtaining the stochastic estimates. For a problem with internal variations one would have to produce 200 complete solves, increasing significantly the cost of these computations. Ongoing research is addressing efficient schemes for such problems.

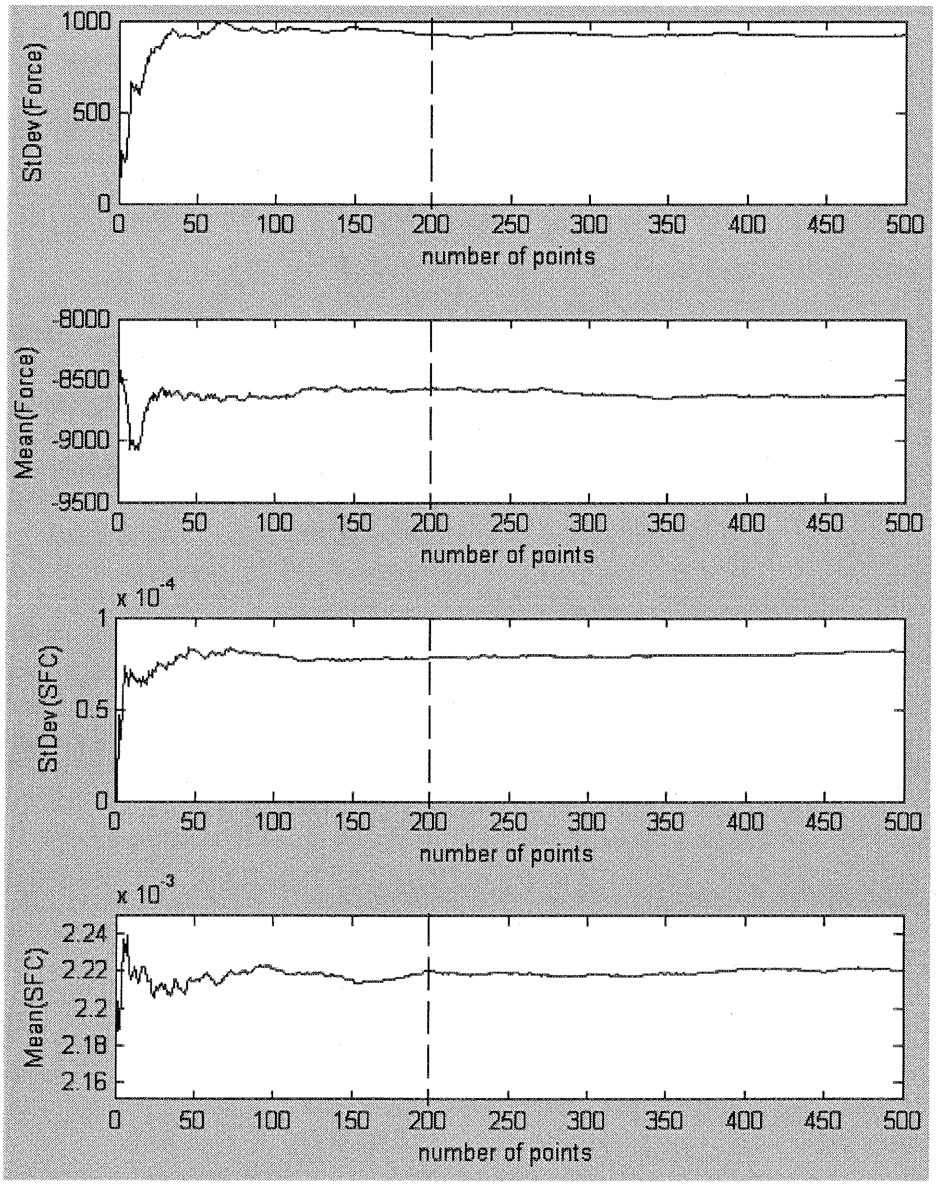

Figure 6 – The effect of number of variations on the mean and standard deviation estimates.

6. Surrogate models in a multiobjective context

As shown earlier in Figure 1, computationally expensive objective functions can be replaced by significantly quicker to compute surrogates models. The price for this is that instead of modelling the entire variable space with all finite element interactions and behaviour, we model only the relationships needed in a much smaller variable space. The aim is to build the model which is accurate around the region of interest. This is usually the area around the optimum. Pursuing this aim, means the surrogate model needs to be updated continuously around the current optimum design. The current optimum is a point suggested after the search of the surrogate model which is less accurate, at least in the beginning. As a result there is a good chance that the current optimum is not near the actual optimum, or it could simply be a local optimum. The conventional update process does not know about this and it will suggest more points at that location. As a result, the surrogate model will become increasingly accurate around the first region identified as optimal and at the same time will be increasingly inadequate elsewhere.

In order to solve this problem methods need to be applied which use samples across the whole solution space, these could be selected randomly or using a specific sampling procedure. These will need to be performed on the full model to retrain the surrogate.

7. Parallel computation facilities v's one year of computations

Those having experience with multiobjective computations may conclude that using 15 design variables and 4 objectives is a relatively complex and tough challenge. Computationally the above task can be very expensive. Each new set of thicknesses takes 20 – 25 minutes to analyze on a 1 GHz CPU. This includes pre and post processing, as well as the finite element solving run. It is known that strategies based on genetic algorithms generally require a large number of function evaluations. For a problem of this dimension it is realistic to say that a minimum of 20000 evaluations would be needed to obtain a good Pareto front. This corresponds to over a year's worth of computations on a single processor. It is for this reason that large multiobjective real life optimization has been avoided until recently.

The same number of function evaluations takes just under two weeks if 30-off 1 GHz CPUs are used in parallel. Instead of receiving results for one set of thicknesses we get results for 30 sets in 25 minutes. Genetic algorithms are very convenient for parallel computations of this sort. The entire population can be run at once, therefore if its size is set as a multiple of the available processor units one can perform the analysis very efficiently.

With the use of surrogate models one does not need to carry out all the 20000 evaluations needed by direct search. Guided by advanced surrogate update strategies [1,2] it is possible to obtain good convergence in just over 1600 evaluations, which take approximately 26 hours on a 30 processor computing cluster.

8. Results

All data obtained by the above setup is stored in a specially designed database. Each database entry contains a full set of input deck files and output files for easier retrieval. Behind each point represented in the figures below there is a full set of such results and geometry files.

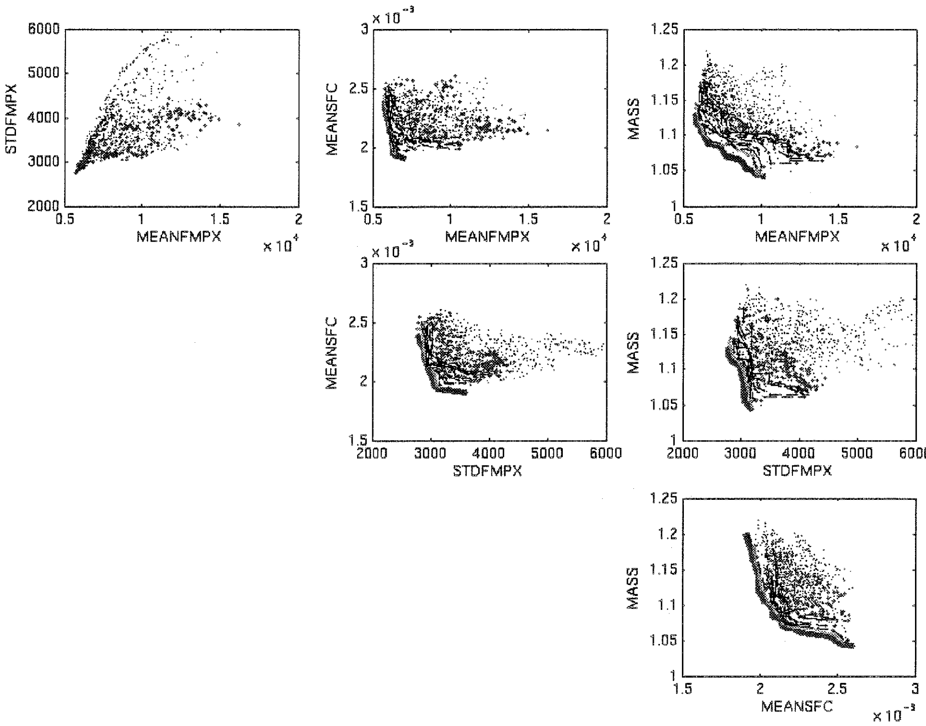

Figure 7 – Two dimensional Pareto front slices

The system used here is built using a tool called OptionsMatlab [17], this acts as an interface between the MATLAB™ environment and the optimization package OPTIONS [3]. Points on the plot are sensitive to mouse clicks to allow fast retrieval of data from the database. For considerations of confidentiality, the axes here are dimensionless and show the general behaviour, instead of actual numbers.

Figure 7 represents all combinations of each of the two objectives. The thick red points are those belonging to the 4 dimensional Pareto front, while the thick red line represents the two dimensional Pareto front for that particular slice. Figure 8 is the robustness set of objective functions.

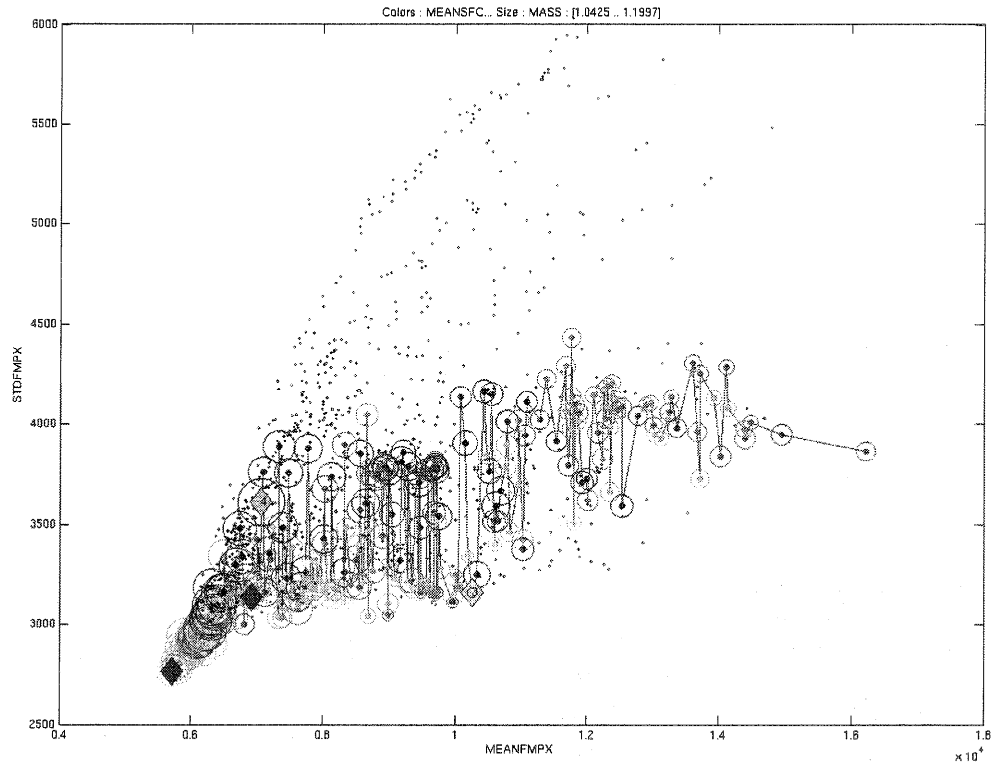

Figure 8 – Robustness measures – Mean(FMPX) vs. St.Dev(FMPX)

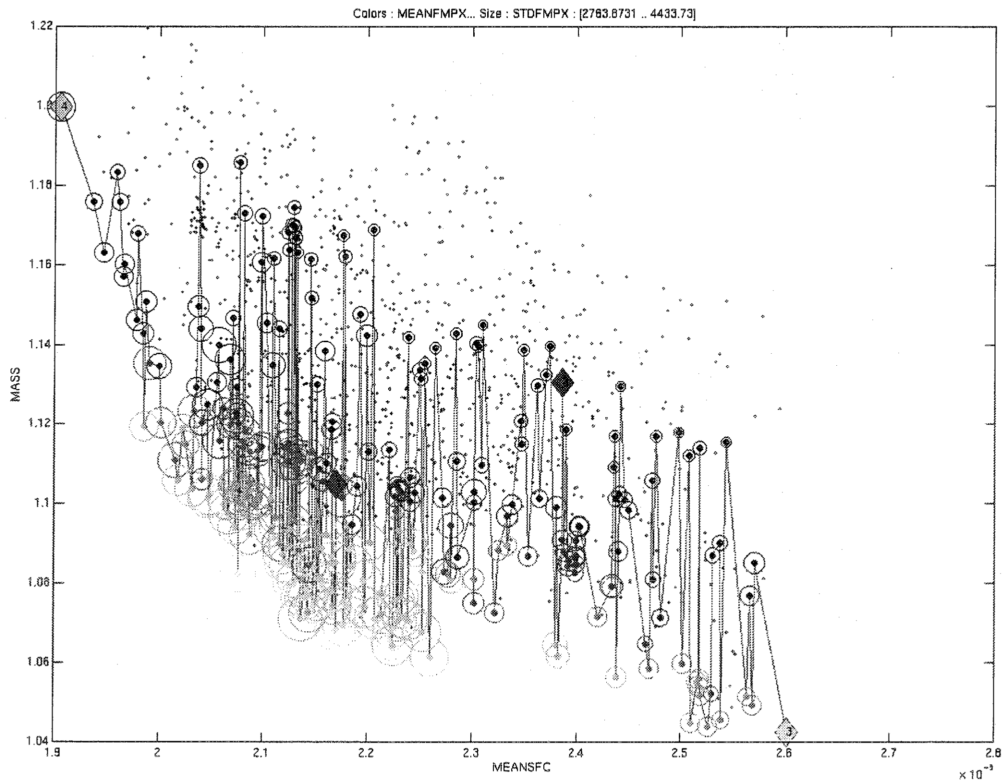

Figure 9 – Performance measures – Mean(SFC) vs MASS

It shows that if no objectives other than mean and standard deviation of the reaction forces are considered – a single optimal solution exists in the bottom left corner of the graph, indicating that the minimum of the two stochastic characteristics can be achieved without the need to compromise. It is designated with a dark blue diamond symbol. The same symbol can be found on Figure 9, where the MASS v's SFC Pareto front is visualized.

It is clear that what was the best in terms of robustness in reaction forces is not good in terms of both MASS or SFC. In other words – one can have a design with low variations and low values of forces, however this would be a very heavy engine which has high fuel consumption. Figure 6 also indicates that there is a significant competition between MASS and SFC. A heavier engine would be stiffer, therefore the tip clearances would be small and unaffected by load conditions, improving the SFC of the engine as reduced tip clearances lead to a more efficient engine.

Figure 8 and 9 are an attempt to show 4 objectives on two dimensional plots. The idea of using glyphs instead of axes is not new [13, 14]. On these graphs one can see that size of the circle and colour indicate the remaining two objectives (good designs are at the blue end of the colour scale and have small circles). These plots may appear chaotic at first sight, however they become a very powerful tool once the designer is trained to interpet them. It is easy to see all 4 objectives simultaneously and one can then choose the best design. Our investigation showed that the point designated with a brown diamond symbol has a good overall performance.

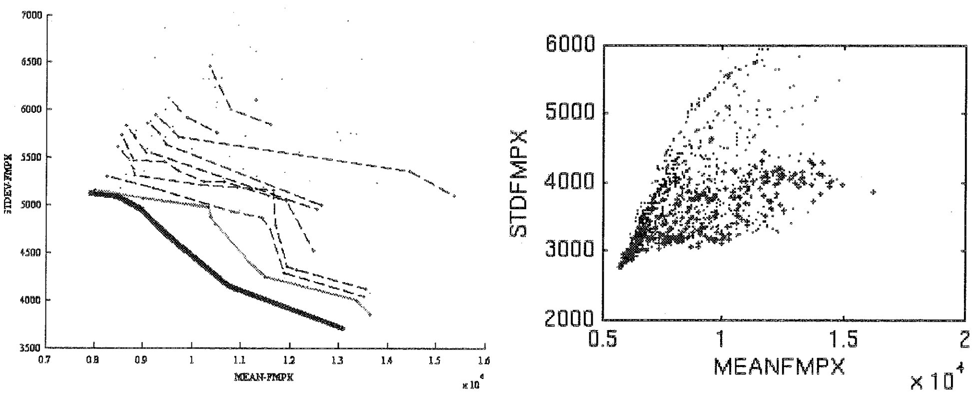

Figure 10 – Pareto front achieved using a low number of function evaluations (200) (left); A simplified version of Figure 8 – 1629 function evaluations (right)

An interesting comparison is shown in Figure 10, where runs of differing lengths are compared: on the left is a short run with just 200 set of solves and on the right a longer one with over 1600. In the right hand plot one can see that the region close to the most robust design is very narrow. The left hand plot shows that the pointed and narrow region was not discovered when fewer evaluations were performed. Both figures represent the same information but lead to completely different conclusions, this highlights the importance of affording sufficient depth to any search.

9. Conclusions

This paper demonstrates that by appropriate combination of the latest computer science advances, efficient surrogate techniques and properly updated models it is feasible to obtain valuable information in a relatively short time. At current prices a 30 CPU computer cluster can be purchased for around $15000, this is a tiny fraction of the benefits that such an investment can provide. Multiobjective optimization is no longer a thing to avoid, it is affordable and gives valuable insights.

It has been shown that is possible to produce an engine that has both low reaction forces and low variation in these forces with regards to external load conditions. Mass and SFC were shown to have a more competitive nature between each other and also in relation to mean reaction force and standard deviation. Using the techniques illustrated within this paper a careful study can locate several designs with relatively low values for all four objectives. These can then be the subject of further more detailed design investigation.

10. Acknowledgements

This work was performed with the support of Rolls-Royce as part of the EC funded Vivace project. For more information on this project please refer to the project website; http://www.vivaceproject.com/

11. References

1. Voutchkov I.I and Keane A.J , "Multiobjective optimization using surrogates", pp. 167-175 in Proc. 7th Int. Conf. Adaptive Computing in Design and Manufacture, (ACDM 2006, ISBN 0-9552885-0-9), Bristol (2006).
2. Keane A. J. (2006), Statistical Improvement Criteria for Use in Multiobjective Design Optimization, *AIAA Journal,* vol, 44, No. 4, pp. 879 – 891, April 2006
3. Keane, A. J, OPTIONS manual, http://www.soton.ac.uk/~ajk/options.ps
4. S. Obayashi, S. Jeong, K. Chiba, "Multi-Objective Design Exploration for Aerodynamic Configurations", AIAA-2005-4666
5. Zitzler et al, Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary computational journal* 8(2), 125-148. 2000
6. Knowles, J. and Corne, D. (1999) The Pareto archived evolution strategy: A new baseline algorithm for multiobjective optimisation. *Proceedings of the 1999 Congress on EvolutionaryComputation,* Piscatway: New Jersey: IEEE Service Center, 98–105.
7. Fonseca, C. M. and Fleming, P. J. (1998) Multiobjective optimization and multiple constraint handling with evolutionary algorithms–Part II: Application example. *IEEE Transactions on Systems, Man, and Cybernetics: Part A: Systems and Humans.* 38–47.
8. Deb, K., *Multi-objective optimization using evolutionary algorithms,* John Wiley & Sons, Ltd., New York, 2003
9. Andy Keane, Prasanth Nair, June 2005, *Computational Approaches for Aerospace Design: The Pursuit of Excellence,* ISBN: 0-470-85540-1
10. S. Leary, A. Bhaskar, and A. J. Keane, "A derivative based surrogate model for approximating and optimizing the output of an expensive computer simulation," *J. Global Optimization* 30 pp. 39-58 (2004)
11. S. Leary, A. Bhaskar, and A. J. Keane, "A Constraint Mapping Approach to the Structural Optimization of an Expensive Model using Surrogates", *Optimization and Engineering* 2 pp. 385-398 (2001)
12. Raymond H. Myers , Douglas C. Montgomery, *Response Surface Methodology: Process and Product in Optimization Using Designed Experiments,* John Wiley & Sons, Inc., New York, NY, 1995
13. Barnett, V. (Ed.) (1981). *Interpreting Multivariate Data.* London: Wiley.

14. Chambers, J. M., Cleveland, W. S., Kleiner, B., & Tukey, P. A. (1983). *Graphical Methods for Data Analysis*. Belmont, CA: Wadsworth
15. M. Sobol' and R. B. Statnikov. The Choice of Optimal Parameters in Multicriteria Problems. Nauka (in Russian), 1981.
16. Roman B. Statnikov and Joseph B. Matusov. Multicriteria Analysis in Engineering. Kluwer Academic Publishers, 2002.
17. http://argos.e-science.soton.ac.uk/blogs/OptionsMatlab/