ICSV14
Cairns • Australia
9-12 July, 2007

14th International Congress on Sound & Vibration

DERIVATIVE FREE OPTIMISATION IN ENGINEERING ACOUSTICS

Rick C. Morgans[1*], Carl Q. Howard[1†], Anthony C. Zander[1], Colin H. Hansen[1]

[1]School of Mechanical Engineering, University of Adelaide, SA, Australia 5005

David J. Murphy[2]

[2]Krix Loudspeakers, Hackham, SA, Australia 5163

*rick.morgans@adelaide.edu.au
†carl.howard@adelaide.edu.au

Abstract

This paper describes recent work on the optimisation of very complex or computationally expensive systems such as those found in many engineering acoustics applications. It uses optimisation techniques that require no knowledge of the derivative of the objective function with respect to the input variables, and hence is suitable for application to problems where the derivative is potentially noisy or expensive to calculate by a finite difference approximation, difficult to calculate analytically, or simply unavailable as is the case in many commercial codes.

The paper begins with a brief review of optimisation methods as applied to acoustic problems, and discusses the limitations of traditional techniques. The theory of two derivative free optimisation methods, a parallel genetic algorithm and a surrogate optimisation technique called Efficient Global Optimisation (EGO) are then described. Two example cases are then discussed: the optimisation of the position and design parameters of vibroacoustic absorbers mounted on the interior of a rocket payload bay to reduce the payload interior pressure fluctuations on launch; and the shape optimisation of an audio loudspeaker to improve sound quality. Finally future directions and challenges in this field are discussed.

1. INTRODUCTION AND LITERATURE REVIEW

The development of modern computational simulation methods, often in conjunction with extensive experimental programmes, has allowed engineers and designers to improve product quality. Simulation methods commonly used in the acoustic domain include volume based discretisation approaches such as Finite Element Analysis (FEA) or a surface based integral equation discretisation such as the Boundary Element Method (BEM). These techniques are often coupled with a structural analysis, usually a shell based Finite Element calculation. The number of elements required by the discretisation to produce an accurate representation of the system can be large, and the response of the system is usually desired over a wide range of frequencies.

Hence the calculations required with simulation based models can often be computationally intensive.

If a suitable measure of product quality can be defined, then mathematical optimisation techniques can potentially automate the design process. A model of the system under consideration is generated (such as an acoustic model of the interior of a vehicle), consisting of inputs (such as parameters that define acoustic lining position and parameters) and outputs (simulation results, such as interior acoustic pressure). The outputs of the model are then reduced to a function that quantifies the objective (for example average sound pressure near the occupants' heads). This objective function is then used by the optimisation routine to systematically change the inputs to the model until the "best" objective has been achieved (reduced sound levels for the occupants).

The use of gradient based optimisation techniques, such as Sequential Quadratic Programming (SQP) [1] is often problematic when optimising objective functions calculated using simulation models. Gradients are often very difficult to calculate analytically, requiring the use of a finite difference approximation and multiple objective function calculations. Other problems include noisy gradients and convergence to local instead of global minima. Global optimisation methods that do not require gradient information from the simulation are thus of great interest.

There are a wide range of potential topics in engineering acoustics, and in this paper we limit ourselves to the interaction of acoustics with mechanical structures (similar to Tinnsten et al. [2]). The review of the literature outlined here can be in no way comprehensive, and for an extensive review of structural acoustic optimisation for passive noise control see Marburg [3].

Genetic Algorithms (GAs) have been used for many vibroacoustic optimisation problems such as finding optimal locations for active control actuators [4–7], the optimisation of acoustic absorber configuration [8, 9] and musical instruments [10]. GAs are used because of their ease of implementation, their ability to take discrete parameter values and deal with multiple minima in the objective function [11].

In the surrogate optimisation technique, the true objective function is replaced with a computationally inexpensive approximation. The current state-of-the-art in acoustic surrogate optimisation is the work of Marsden et al. [12, 13], for the shape optimisation of trailing edge noise, where a surrogate model is placed within a provably convergent pattern search method, Mesh Adaptive Direct Search (MADS). Each computation is a computationally expensive Computational Fluid Dynamics (CFD) calculation which takes approximately $2-3$ weeks of wall clock time. After 1 iteration of the 5 parameter optimisation, an 85% reduction in noise was reported with a total of 8 expensive calculations performed. A total of 21 less expensive steady state CFD calculations were used to filter out infeasible regions of the solution. Interestingly there was also a 65% increase in lift and a 15% decrease in drag.

In the current work, we give two examples of derivative free optimisation methods that have been used to solve acoustic optimisation problems. In the first, a genetic algorithm is used to solve a very large coupled vibro-acoustic problem involving optimising the placement and parameters of up to 500 absorbers on the interior surface of a rocket payload bay. The objective of the optimisation is to reduce payload bay pressure fluctuations on launch. This is a very difficult optimisation problem, with up to 2500 variables to be solved. Analytical gradients are difficult to calculate in this instance, because the position of each absorber is limited to the nodes on the underlying finite element mesh.

The second example is the shape optimisation of an audio product to improve its sound quality. The sound field produced by a horn loaded loudspeaker is calculated using a boundary element like method, with the horn geometry parameterised by splines. Modifications to the shape of the horn influence the sound field projected onto the audience, and a suitable sound quality metric is used as an objective function. The optimisation technique used is a surrogate based, where the true objective function is replaced by a statistical model that interpolates between objective function evaluations, giving both a measure of the value of the objective function between known points and a measure of uncertainty. An auxiliary optimisation problem is then solved using this surrogate, which is computationally inexpensive to evaluate, giving the next most likely place to find an improvement in the true objective function.

The structure of this paper is as follows: first the theory of each optimisation method is described; next details of the application examples are presented; and finally future directions are discussed.

2. THEORY

2.1. Asynchronous parallel genetic algorithm

Genetic Algorithms (GAs) belong to a class of optimisation methods called *search heuristics* which "seek an acceptable improvement rather than a provably optimal solution by methodically searching the feasible region" [14] and they are classified as *incomplete* methods [15]. Despite no provable convergence they are popular optimisation methods possibly because of their ease of implementation, ability to escape local minima, ability to use discrete values of the input parameters and ability to rapidly locate a good solution.

GAs take cues from evolutionary processes. The optimisation takes place using a *population* of *individuals*. Details of each individual that give rise to specific traits (i.e. the parameters to be optimised) are encoded as *chromosomes*. The objective function calculated for each individual provides a *fitness* ranking. The *selection* operator randomly selects a given individual based on its fitness from the population for inclusion in *recombination* (where *parents* reproduce their genetic material in *offspring*) and possibly in *mutation* (random perturbations of the chromosome). The optimisation proceeds iteratively with both the average and maximum fitness increasing progressively as new *generations* are produced.

As the objective function must be calculated for a population of individuals, GAs are well suited to parallel evaluation. One useful paradigm is that of master-worker, where the objective function evaluations are conducted by "worker" computers and a "master" computer is responsible for managing the population by conducting the selection, recombination and mutation processes. An asynchronous GA implemented using the publicly available Condor distributed computing environment [16] is described here and in more detail in Howard et al. [17].

The Genetic Algorithm Toolbox [18] was used as a basis for developing an asynchronous parallel GA. The parameters to be optimised are encoded using either an integer or binary scheme. The asynchronous GA was developed with guidance from Stanley and Mudge [19], and operates as shown in Figure 1.

2.2. Efficient Global Optimisation

Efficient Global Optimisation (EGO) [20, 21] is a gradient free surrogate based global optimisation technique. The EGO technique proceeds as follows and is summarised in Figure 2. A number of different sets of input parameters are randomly generated to give a representative

1. The GA generates a random initial population.
2. The objective functions are evaluated for the initial population using the Condor pool.
3. The GA checks how many jobs are already being calculated by the worker processes and if it has not exceeded a chosen value (it is possible to have more jobs in the Condor queue than the number of computers in the pool), then
   (a) The results from the objective function evaluations are sorted in rank order based on their fitness.
   (b) A new set of individuals is created by performing selection, mutation and recombination operations.
   (c) A file containing the parameters for the chromosomes is generated for submission to the Condor pool. The files are placed in a unique directory, where all the files relevant to that particular objective function evaluation reside. The job is then submitted to the Condor pool.
4. The directory structure is checked for the existence of a file called `success.sub`, which indicates that the objective function evaluation was completed successfully. The jobs that are ready to have their results read back into the GA are formed into a queue, with the oldest jobs at the front of the queue.
5. The results from the objective function evaluation from the oldest job in the queue are retrieved.
6. The new population is inserted into the old population by replacing the chromosomes in the old population that had the worst fitness.
7. The process repeats from step 3 until a predetermined number of iterations is reached.

Figure 1. Asynchronous parallel Genetic Algorithm

sample over the range of potential solutions. Here the random samples are generated by Improved Hypercube Sampling (IHS) [22], which attempts to generate a space filling design, but any suitable design of experiments method could be used.

The objective function is evaluated for each set of input parameters and a surrogate model is fitted to the objective function. This surrogate model describes both the variation of the mean value of the objective function between the sample points and the uncertainty between them, and is much less computationally expensive to evaluate than the original objective function. In this application, a Kriging technique is used [23]. Kriging techniques, developed in the geostatistics and spatial statistics fields, fit a surface to a set of data point values. It models the variation of the unknown function as a constant value plus the variation of a normally distributed stochastic variable. It is essentially a method of interpolation between known points that gives a mean prediction, $\hat{y}(x)$, in addition to a measure of the variability of the prediction, $s(x)$, the estimated standard deviation. Another appropriate optimisation technique such as SQP, simulated annealing [24] or the DIRECT method [25] is then employed to solve an auxiliary problem to find the next best place to sample for a minimum primary objective function. The secondary objective function used to solve the auxiliary problem in this application is the Expected Improvement ($\mathbf{E}[I]$) objective function. The improvement function (I) is defined as the improvement of the current prediction, $\hat{y}(x)$, at point x over the minimum value of the current set of samples, y_{min}, i.e.

$$I = \max\left(y_{min} - \hat{y}(x), 0\right) \tag{1}$$

The expected improvement, defined as the expectation of the improvement, is given by

$$\mathbf{E}[I] = \left(y_{min} - \hat{y}(x)\right) \mathrm{CDF}\left(\frac{y_{min} - \hat{y}(x)}{s(x)}\right) + s(x)\, \mathrm{PDF}\left(\frac{y_{min} - \hat{y}(x)}{s(x)}\right) \tag{2}$$

where CDF is the standard normal cumulative density function, and PDF is the standard normal probability density function. The point at which the value of the expected improvement is maximised gives the best point at which to calculate the true objective function. The Expected Improvement is constructed to search for both local and global minima [20, 21]. The surrogate model is then updated to include the newest sampled point, and the operation repeated until the sampling point does not change and the global minimum of the objective function has been found.

1. An initial set of input parameters is selected using IHS.
2. The true objective function is evaluated for all new members of the set.
3. A Kriging surrogate model is fitted to the values of the objective function.
4. The expected improvement objective function, calculated using values from the computationally inexpensive Kriging model, is minimised using any suitable global optimisation method.
5. The result of the minimisation (the next input parameters most likely to improve the true objective function) is added to the set.
6. The process repeats from step 2 until a predetermined number of iterations is reached.

Figure 2. Efficient Global Optimisation algorithm

One advantage of the EGO method is that it requires a minimal number of true objective function evaluations, and most of the optimisation is done on the computationally inexpensive surrogate. This makes the method very efficient when the objective function is computationally expensive.

3. EXAMPLE APPLICATIONS

3.1. Rocket fairing interior noise

A finite element model of a composite fibre rocket payload fairing was created, as shown in Figure 3a. The fairing has dimensions of 1m diameter and 5m height. One method suggested for reducing the noise levels inside the payload bay during the launch of the rocket is to use passive tuned-mass-dampers and Helmholtz resonators attached to the fairing walls. Optimisations were conducted to determine the parameters for the locations, resonance frequencies and damping of the acoustic and vibration absorbers for 10, 100, and 500 absorbers attached to the fairing. A constant added mass 'budget' of 10% of the fairing mass was divided amongst the absorbers. The optimisations were conducted using a parallel asynchronous genetic algorithm implemented on a distributed computing network, as described in Howard et al. [26]. The results of the optimisation of 500 absorbers after over 200,000 objective function evaluations is shown in Figure 3b. The results indicate that the effect of the absorbers is to create a 'fuzzy' vibro-acoustic response, by 'smearing' the acoustic potential energy across the frequency range, rather than bifurcating the response at resonances, which is typical for passive reactive devices. Figure 3c shows the reduction in the internal acoustic potential energy with increasing number of absorbers. The results show the effect of combined vibration and acoustic absorbers, and the effect of the Helmholtz resonators with a lumped 'blocking' mass, to demonstrate that the sprung mass in the vibration absorbers provides measurable benefit.

These optimisations required significant computational resources for the greater than 200,000 objective function evaluations. If the optimisation had been conducted on a single 3.0GHz Pentium it would have taken 417 days. By using a distributed computing network of about 150 computers of varying processor speeds from 1.8GHz to 2.4GHz, the time taken to conduct the optimisation was less than 14 days.

3.2. Horn shape optimisation

Horn loaded loudspeakers increase the efficiency and control the spatial distribution of the radiated sound. They are often used as components in cinema sound systems where it is desired that the sound be broadcast onto the audience uniformly at all frequencies, improving the listening

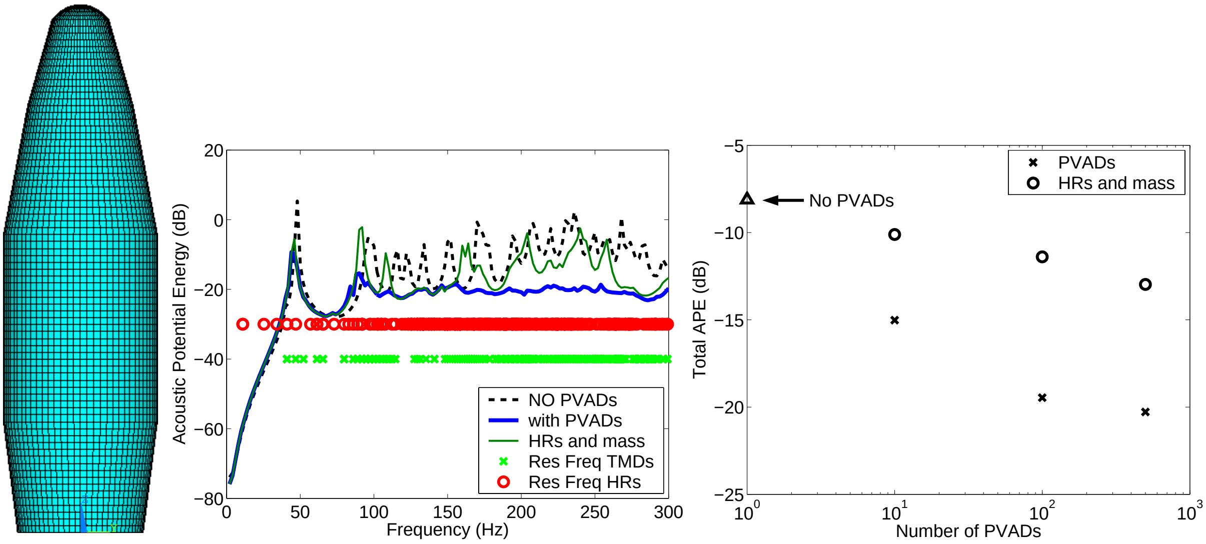

(a) (b) (c)

Figure 3. RSLVF (a) Finite element model of the fairing, (b) Acoustic potential energy versus frequency for 500 PVADs and (c) Acoustic potential energy versus the number of PVADs.

experience. The sound distribution, or beamwidth, is related to the shape of the horn and can be predicted well by the source superposition method [27]. However the cost of evaluating the objective function is high. Gradient calculation is also difficult with finite differences because of the discrete nature of the meshing used such that a small change in horn profile could lead to a jump in the objective function. A spline based parametrisation was used to define the horn geometry, and a two parameter EGO optimisation was performed. Further details on the optimisation including the primary objective function that was used are presented in Morgans [28] and Morgans et al. [29].

The results of the EGO optimisation of the horn appear in Figure 4a. A contour of the Kriging surrogate mean predictions of the objective function is shown, along with the positions of the EGO sample points. The 25 circular markers show the initial sample points, and the 25 square markers show the sample points chosen by the Expected Improvement function, balancing both local and global optimisation. A convergence to the global minimum can be seen with repeated sampling (many square markers) around the global minimum (diamond marker) at $x(1) = 0.49$ and $x(2) = 0.69$. The horn profile corresponding to the global minimum is shown in Figure 4c, and the beamwidth calculated from the optimal horn profile is shown in Figure 4b. Figure 4b reveals that a constant beamwidth as a function of frequency has been achieved above a low frequency limit, thus providing a superior listening experience for the audience.

4. CONCLUSIONS AND FUTURE DIRECTIONS

It has been shown that both derivative free methods examined work well for simulation based optimisation. The GA performs well for problems with a very large number of parameters, enabled by the use of a distributed computing environment and the robustness of the GA technique. The EGO method performs well for smaller dimensional problems where the objective function is very expensive.

Future directions and challenges are many for the field. One obvious improvement not

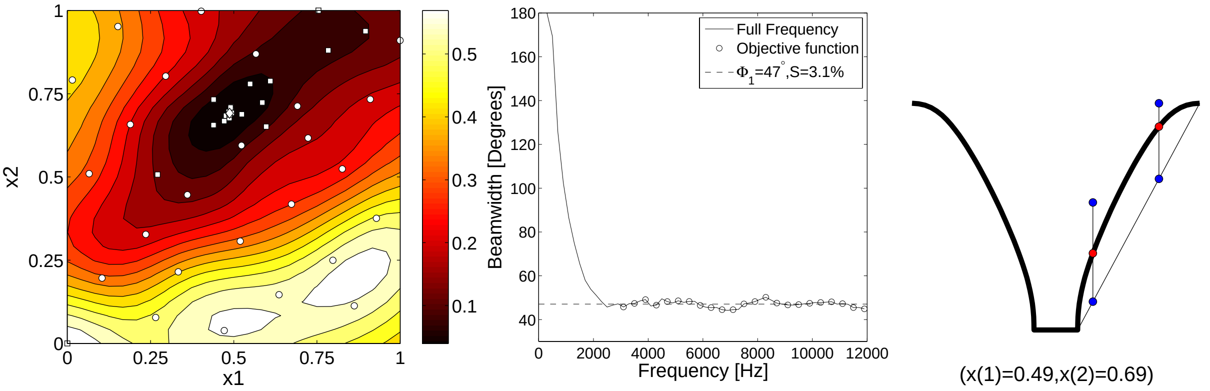

(a) (b) (c)

Figure 4. (a) Optimisation trajectory for the two parameter spline horn geometry, Optimal two parameter spline horn (b) Beamwidth and (c) Geometry profile

directly related to the optimisation method chosen is to improve the computational speed of the simulation, and if analytic gradients can be calculated then their use dramatically improves the speed and robustness of many optimisation algorithms. Both methods described here have only used bound constraints (the parameters are limited between upper and lower values). General constraints can be handled in GAs through various methods including penalty and filter methods. Constrained EGO algorithms have been developed [30], but they have been found by the authors to be fragile. It may be possible to use EGO as a search step inside a provable convergence pattern search algorithm such as MADS and use filter methods for constraints [13]. As the number of optimised parameters increase, the EGO method breaks down because of the sparsity of the data supplied to the Kriging surrogate model. Instead of using a statistical surrogate like Kriging, a simplified physics based surrogate could possibly be used in a similar manner to the Space Mapping technique [31]. Parallel computing has proven to be very successful for the GA calculation, and a similar approach is currently being pursuing for the EGO method. The GA is a heuristic search with no provable convergence and the use of a provable global optimisation method within the parallel function evaluation environment may improve the optimisation result.

REFERENCES

[1] K. Schittkowski. NLPQL: A FORTRAN subroutine solving constrained nonlinear programming problems. *Annals of Operations Research*, 5:485–500, 1985.

[2] M. Tinnsten, B. Esping, and M. Jonsson. Optimization of acoustic response. *Structural Optimization*, 18: 36–47, 1999.

[3] S Marburg. Developments in structural-acoustic optimization for passive noise control. *Archives of the computational methods in engineering*, 9(4):291–370, 2002. Ref: AA094A.

[4] C. H. Hansen, M. T. Simpson, and B. S. Cazzolato. *Active sound and vibration control: theory and applications*, chapter 9: Genetic algorithms for optimising ASVC systems, pages 185–220. Number 62 in IEE control engineering series. Institution of Electrical Engineers, London, UK, 2002.

[5] C. H. Hansen, M. T. Simpson, and B. S. Cazzolato. Genetic algorithms for active sound and vibration control. In *Inter-Active 99*, 1999.

[6] M. T. Simpson and C. H. Hansen. Use of genetic algorithms for optimizing vibration actuator placement for minimizing sound transmission into enclosed spaces. In *Proceedings of SPIE - The International Society for Optical Engineering, Smart Structures and Materials 1996: Smart Structures and Integrated Systems*, 1996.

[7] M. R. Bai and S. Chang. Active noise control of enclosed harmonic fields by using BEM-based optimization techniques. *Applied Acoustics*, 48(1):15–32, 1996.

[8] J. Xu, J. Nannariello, and F. R. Fricke. Optimising flat-walled multi-layered anechoic linings using evolutionary algorithms. *Applied Acoustics*, 65(11):1009–1026, 2004.

[9] O. Tanneau, J. B. Casimir, and P. Lamary. Optimization of multilayered panels with poroelastic components for an acoustical transmission objective. *Journal of the Acoustical Society of America*, 120(3):1227–1238, 2006.

[10] A. Braden. Optimisation techniques for solving design problems in modern trombones. In *Forum Acusticum 2005. Triennial conference of the European Acoustics Association.*, 2005.

[11] P. G. Dylejko, N. J. Kessissoglou, Y. Tso, and C. J. Norwood. Optimisation of a resonance changer to minimise the vibration transmission in marine vessels. *Journal of Sound and Vibration*, 300(1-2):101–116, 2007.

[12] A. L. Marsden, M. Wang, J. E. Dennis, Jr., and P. Moin. Trailing-edge noise reduction using derivative-free optimization and large-eddy simulation. *Journal of Fluid Mechanics*, 572:13–36, 2007.

[13] A. L. Marsden, M. Wang, J. E. Dennis, Jr., and P. Moin. Optimal aeroacoustic shape design using the surrogate management framework. *Optimization and Engineering*, 5(2):235–262, 2004. Special Issue on Surrogate Optimization.

[14] J. E. Dunlap. On the use of surrogate functions for mixed variable optimization of stochastic systems. Master's thesis, Air Force Institute of Technology, 2005.

[15] A. Neumaier. Complete search in continuous global optimization and constraint satisfaction. *Acta Numerica*, pages 271–369, 2004.

[16] Condor. The Condor Project Homepage.
`http://www.cs.wisc.edu/condor/`, 2007.

[17] C. Q. Howard, R. C. Morgans, Hansen C. H., and A. C. Zander. A tool for the optimisation of vibro-acoustic systems using a parallel genetic algorithm and a distributed computing network. *Noise Control Engineering Journal*, 53(6):256–267, 2005.

[18] GA Toolbox. Genetic Algorithm Toolbox for Matlab.
`http://www.shef.ac.uk/acse/research/ecrg/gat.html`, 2007.

[19] T. J. Stanley and T. Mudge. A Parallel Genetic Algorithm for Multiobjective Microprocessor Design. In *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 597–604, 1995.

[20] M. Schonlau. *Computer experiments and global optimization*. PhD thesis, University of Waterloo, 1997.

[21] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optmization of expensive black-box functions. *Journal of Global Optimization*, 13:455–492, 1998.

[22] B. K. Beachkofski and R. V. Grandhi. Improved distributed hypercube sampling. In *43^{rd} AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, Denver*, 2002.

[23] S. N. Lophaven, H. B. Nielsen, and J. Søndergaard. DACE - a Matlab Kriging toolbox, version 2.0. Technical Report IMM-REP-2002-12, Informatics and Mathematical Modelling, Technical University of Denmark, 2002. URL `http://www.imm.dtu.dk/~hbn/publ/TR0212.ps`.

[24] L. Ingber. Adaptive Simulated Annealing (ASA), 1993. URL `http://www.ingber.com/#ASA-CODE`.

[25] D. E. Finkel and C. T. Kelley. Convergence analysis of the DIRECT algorithm. *Submitted to SIAM J. Optimization*, 2004.

[26] C. Q. Howard, C. H. Hansen, and A. C. Zander. Vibro-acoustic noise control treatments for payload bays of launch vehicles: discrete to fuzzy solutions. *Applied Acoustics*, 66:1235–1261, 2005.

[27] G. H. Koopmann and J. B. Fahnline. *Designing quiet structures: a sound power minimization approach*. Academic Press, 1997.

[28] R. C. Morgans. *Optimisation techniques for horn loaded loudspeakers*. PhD thesis, University of Adelaide, 2005.

[29] R. C. Morgans, A. C. Zander, C. H. Hansen, and D.J. Murphy. EGO shape optimization of horn loaded loudspeakers. *submitted to Optimization and Engineering*, 2007.

[30] M. J. Sasena. *Flexibility and Efficiency Enhancements for Constrained Global Design Optimization with Kriging Approximations*. PhD thesis, University of Michigan, 2002.

[31] S. Koziel, J.W. Bandler, and K. Madsen. Space-mapping based interpolation for engineering optimization. *IEEE Trans. Microwave Theory Tech.*, 54(6):2410–2421, 2006.