Finite Element Optimisation For Problems In Acoustics

Trevor Esward (*) & Louise Wright (+)

Centre for Mechanical and Acoustical Metrology (*) and Centre for Mathematics and Scientific Computing (+), National Physical Laboratory

1. Abstract

NPL uses finite element (F.E.) techniques to investigate the performance of transducers for applications in acoustics, and to assist designers and manufacturers of medical ultrasound devices. In both cases there is frequently the need to optimise the performance of the device against an objective function such as a required acoustic field characteristic or electrical response. Frequently these problems involve resonant systems that make measurement and modelling problematic due to the large responses near to resonant frequencies.

This paper describes an optimisation program developed at NPL, which optimises transducer design within given constraints. The program varies the material properties of components in the design, uses commercial finite element vibro-acoustics software to generate results, and minimises either the difference between results and desired behaviour of an ideal device, or the difference between results and measured response data of an existing device as an objective function. The objective function includes a weighting function to help avoid problems with resonance.

F.E. software is effectively a "black box" where the relationship between the input data and the output results is complicated and hidden. This relationship cannot generally be written as an analytical formula and derivative information is not available, and this restricts the range of optimisation techniques that can be used. The paper explains the implementation used at NPL, and presents results which demonstrate the success of the method which was adopted. Similar techniques could be applied to many other optimisation problems involving "black box" type software.

2. Optimisation techniques

Optimisation of a design is a common problem in metrology. Frequently it is necessary to optimise the behaviour of a device against a required output characteristic so as to give as close to ideal performance as possible. Sometimes it is useful to "reverse engineer" a design, if it is of interest which characteristics of a device gave rise to a measured response. These optimisation problems are often a question of choosing the right design geometry or material properties, and can be analysed using F.E. methods.

Consider the problem of designing an acoustic device that produces a particular pressure distribution in some area of space. If $\{\mathbf{x}_i: i = 1, 2, \ldots n\}$ are a series of points in the area we are concerned with, then define the desired pressure at the i^{th} point to be $P(\mathbf{x}_i)$. Suppose there is a model of the device that provides predictions of the pressure at the points. The predictions will depend on a set of model parameters, for example device geometry and material properties. Define $\mathbf{y}=\{y_j: j = 1, 2, \ldots m\}$ to be this set of model parameters and $P_i(\mathbf{y})$ to be the resulting model prediction at the point $\mathbf{x}_i$. Then the aim of the optimisation is to alter the model parameters and hence the predictions so that the desired pressure and the predicted pressure are the same, so that $P(\mathbf{x}_i) - P_i(\mathbf{y}) = 0$ for i = 1, 2, … n. In addition, it is possible that some points are more important than others (for instance, a peak pressure value may be more important than the behaviour away from the peak), so it may be useful to assign an importance weighting w_i to each point. This problem can be rewritten as: Minimise $\sum_{i=1}^{n} w_i \left(P(\mathbf{x}_i) - P_i(\mathbf{y})\right)^2$ with respect to $\{y_j: j = 1, 2, \ldots m\}$. The sum that is being minimised is called the objective function. There may also be constraints on the model parameters, for instance the device may have a maximum size, or the material properties may have to lie within a realistic range of values.

Traditional optimisation methods rely on two pieces of information about the objective function. For a given set of parameters, they need the value of the function, and the values of the derivatives of the function with respect to each of the parameters. This is because the derivative gives information about how altering the parameter value affects the function value, and so knowing the derivative helps to choose a new parameter value that will improve the function value. In the example above, the value of each of the P_i varies when the parameters y_j vary, so the required derivatives would be of each of the P_i with respect to each of the y_j.

For finite element models there is not a clear expression linking the parameters to the results because the calculation of results is done by a “black box” piece of software where the calculation process is not visible to the user. This means that the derivative information is not immediately available. This generally holds for any model involving “black box” software where the solution is produced from a numerical approximation to a partial differential equation, such as CFD or boundary element packages. It is possible to estimate the derivatives by perturbing each of the model parameters slightly and using the perturbed results in an approximation. This requires extra runs of the model, and for a complex finite element model the model run time can be too long to make this a practical solution. An ideal optimisation technique uses the minimum number of model evaluations to find its optimum.

There are several optimisation methods that do not require derivative information, including simulated annealing and evolutionary techniques such as genetic algorithms. These techniques are often at least as efficient as the method chosen, but they can sometimes be complicated to set up, particularly turning the model problem into a form the algorithm can handle. As this software was initially intended to be a proof of concept, a simpler method was chosen.

3. The software

The optimisation method chosen was a modified version of the Nelder-Mead algorithm [1], called COBYLA [2] (from Constrained Optimisation BY Linear Approximation). This was chosen because it is a simple method to understand and adapt to our needs, and because it is designed to optimise constrained problems using only function values. This is particularly appropriate for our application as many problems can be reduced in size by restricting the possible values a material property can take.

The Nelder-Mead algorithm uses a group of m+1 sets of parameter values $\mathbf{y}$ to generate a new set $\mathbf{y}$* that is an improvement on the worst of the group. The COBYLA method uses the idea of generating an improved set of variables from the existing ones, but it also generates some sets that ensure that the search does not prematurely confine itself to one area, which is likely to happen with constrained problems. It uses linear approximations to generate the new set, and includes a check for suitability of the group to ensure the problem does not become degenerate and another check to ensure that the best set obeys the constraints.

The software developed uses the PAFEC [3] finite element software to generate model results. The parameters that can be varied to produce an optimal solution are mostly material model parameters, and include piezoelectric, fluid, orthotropic, and isotropic material properties. As well as being varied independently, these properties can be related so, for example, if a material is known to have identical properties in two directions but different ones in the third this can be enforced. A wide range of model results can be used in the function, and different weightings can be given to different results. The results can come from different excitation frequencies and can be any combination of allowable types. In addition to optimisation problems, the software can be used to carry out automated sensitivity analyses. This can help to identify those parameters that are worthwhile varying in an optimisation and those that do not affect the results so much, thus reducing run time.

The software was initially tested by taking a standard problem with a known solution, altering the material properties, and requiring the software to get back to the original properties. This was done successfully, so further more demanding problems were tried.

4. Results and future developments

The first real problem on which the software was tested was a piezoelectric disc. Three of the physical properties of the disc were varied, two of which were forced to be the same so as to produce an axisymmetric material. The disc was driven at various frequencies and its electrical impedance was calculated at each frequency. The model results were compared with experimental data for the objective function. This problem has difficulties with the experimental measurement and with the modelling, as the disc is a resonant system so data and model predictions will have large peaks at certain frequencies. This means that the driving frequencies have to be chosen carefully to capture the resonant behaviour without allowing it to dominate.

Before the optimisation run was done, a sensitivity analysis of the problem was carried out. The results of this are shown in figure 1. The vertical scale is truncated to emphasise the area of interest. From these results it was clear that both properties strongly affected the function value. Figure 1 shows that there was a line of local

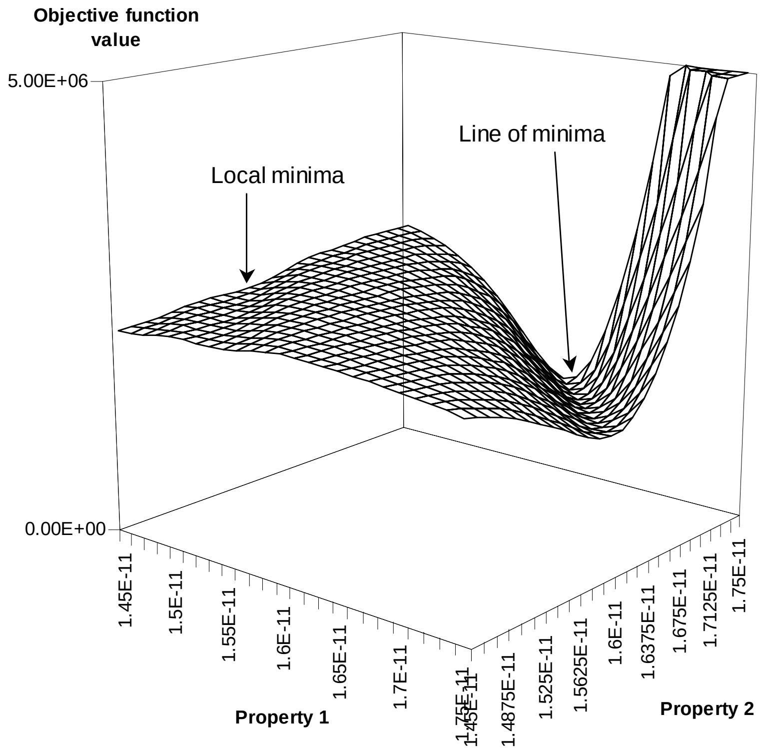

Figure 1 Sensitivity analysis results. Vertical scale is truncated to emphasise area of interest.

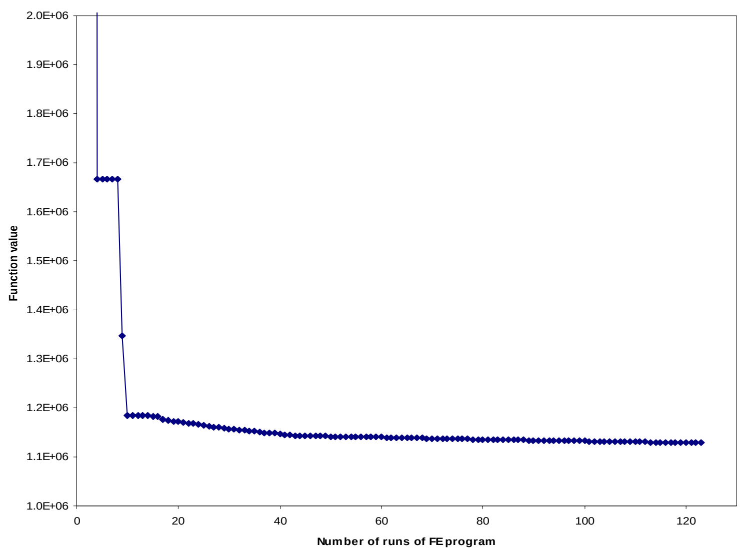

Figure 2 Optimal function values versus number of F.E. runs.

minima, which could potentially lead the software to a false result if the algorithm restricted its search to that region too rapidly, and a deeper line of minima that contains the global minimum.

The optimisation was run using the same limits on material properties as the sensitivity analysis. Figure 2 shows the evolution of the objective function values, again with a truncated vertical scale for clarity. The horizontal scale is the number of calls of the finite element software package, which is analogous to time as the F.E. job run is the slowest step in the program. The job was fully converged after 123 evaluations, but had reached the optimal parameter values after 115 and was within 1% of the optimal function value and parameter values after 55 evaluations. The reason for the slow convergence was that the program was working its way along the line of minima to the global minimum, and the function values change very little along the line so progress was slow. This is illustrated by figure 3, which shows a plot of the optimal pairs of values found during the optimisation superimposed on a two-dimensional plot of the surface shown in figure 1. This shows that the software initially finds the line of minima then gradually works its way along it to the global minimum.

The global minimum found was consistent with the values calculated by the sensitivity analysis. Figure 4 shows the target values of impedance compared with the final values obtained. Although the matching is not exact, the overall shapes of both real and imaginary parts are in agreement. It is possible that better agreement could be achieved outside of the range of values specified. This data also shows the resonant behaviour referred to previously: the data and results both show a truncated peak at about 1.25 MHz where driving frequencies either side of the resonant frequency were chosen. The matching of data and prediction away from this peak are good. An appropriate choice of weighting function would make it possible to either concentrate on or ignore the resonant region.

Several options are being considered as a next development. At the moment, as discussed in section 3, the objective function is a simple weighted sum of squares. There are some problems that may be better tackled with a more complex function, for instance, if the overall shape of a pressure distribution is important but the actual values are less so, a sum of squares is unlikely to be the most suitable option. Similarly the weighting

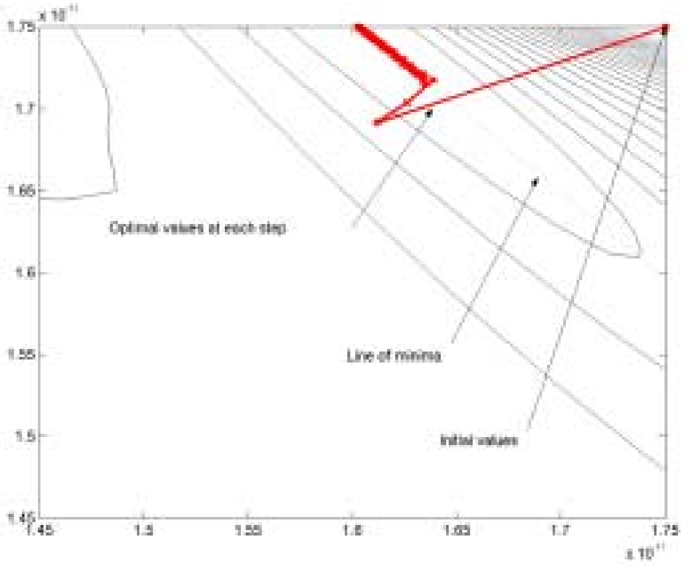

Figure 3 Contour plot of surface as shown in figure 1 and line of optimal parameters at each step, showing how the software reaches the minimum.

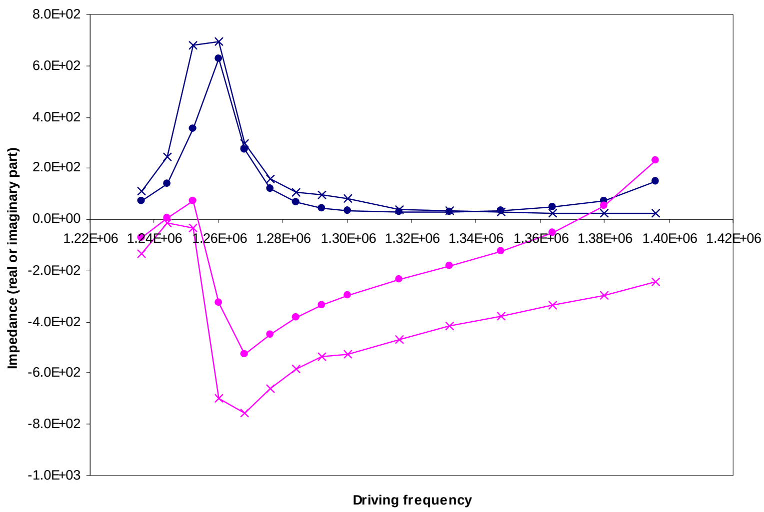

Figure 4 Comparison of objective function target values (crosses) and optimal parameter model results (circles). Note the truncated peak at about 1.25MHz indicating resonant behaviour.

function needs investigating as it can help to prioritise important areas of data. Another potentially useful option would be to alter certain geometric parameters for standard problems. For instance, for a device made up of several layers whose thicknesses can be altered to change the performance, it would be useful to be able to optimise these thicknesses to produce a given performance.

The main barrier to the new developments is not the optimisation algorithm, it is the problem of automating the writing of the F.E. input file and the reading of the output file. For any such new development, this is generally the most difficult stage as “black box” software generally needs a lot of user interaction and cannot easily be automated.

5. Conclusions

Optimisation software has been developed to find material parameters that give a user-specified response when used in a finite element model of an acoustics problem. The software can alter a number of different properties and can link properties together. The user-specified response can be any output of the F.E. model and can be weighted if one part of the response is more important than another. The software can also run sensitivity analyses to identify which parameters affect the relevant model results most strongly.

There is nothing that is specific to acoustics in this work: the techniques could be applied to any “black box” software that can be run and have its results processed automatically. Many fields of measurement involve optimisation problems that could benefit from this type of software.

6. References

1. Nelder J. A. and Mead R. Computer Journal, v. 7, issue 4, p. 308-313 (1965).

2. Powell M.J.D., University of Cambridge Numerical Analysis Report DAMTP 1992/NA5, April 1992, 18p.

3. PAFEC-FE, SER Solutions International, 39 Nottingham Road, Stapleford, Nottingham NG9 8AD.