ISABE 99-7176

A99-34177

EXPERIENCES WITH OPTIMIZERS IN DESIGN

A.J. Keane and G.M. Robinson
Dept. of Mechanical Engineering,
University of Southampton,
Highfield, Southampton, SO17 1BJ, U.K.
E-mail: andy.keane@soton.ac.uk

ABSTRACT

This paper outlines the authors' experiences on applying optimization search methods to a number of engineering design tasks over the last 10-15 years. These projects have covered a range of domains and have been multi-discipline and multi-level in nature. The search methods deployed have ranged over the full spectrum of methods available in the literature including gradient decent, genetic, stochastic and heuristic approaches. Problems have been tackled in drug design, ship design, civil transport wing design, sonar data processing, communications network operation and satellite structure design. Currently, as part of an initiative with British Aerospace and Rolls-Royce, a new research centre is being created to focus attention on current problems of interest in aerospace and gas-turbine design. The paper closes by listing some of the key topics which must be considered when setting up design search and optimization aids.

INTRODUCTION

The natural desire of engineers to improve or 'optimize' their designs has been the subject of research activity for many decades[1]. Perhaps the earliest systematic work in this field was carried by the operational research community during and after the second world war. More recently the focus of research has moved on from the desire to find optimal designs, since it is rarely the case that optimality can be achieved in any mathematical sense. Instead, researchers have studied ways in which computer tools may be used to aid the engineer in the search for improved designs. This paper outlines some of the authors' experiences in this field over the last 10-15 years.

The earliest work described here concerns naval ship concept design. This is an area where optimization techniques have been applied for many decades since it is not possible to prototype such vessels and also because reasonable empirical modelling data has been available since the 1940's. This work is multi-disciplinary in nature but given the simplicity of the models used, searches addressing powering, weight prediction, structural strength and space demands can be relatively easily set out. The major breakthrough in the usefulness of this work came with adoption of simplified hull form modelling systems that allowed full hull descriptions to be automatically generated for subsequent analysis codes to use. Following this change many of the naval architect's more advanced analysis techniques could then be applied. Even so, almost all work in this area is still performed using classical gradient decent methods. Current research is focussed more on dealing with varied vessel types than better search methods.

More advanced search methods come into their own when dealing with deceptive search spaces and so the next area of work described here concerns vibration control in structural design. This work has ranged from studies of highly idealised classical structures of little direct relevance right through to satellite booms where experimental test have been conducted to verify some of the improvements predicted. This domain is characterised by being difficult and time consuming to analyse and by having very many local optima in the search domain. It therefore is highly suitable for searching with modern stochastic and genetic optimizers and some spectacular design improvements have been achieved in the field of passive vibration control. Future studies are concentrating on mixed active / passive designs.

The final project described deals with recent work on the design of aircraft wings. The wings on a modern civil transport aircraft are already extremely highly optimized. However, as they operate in transonic flows that are extremely difficult to model computationally and also subject to many varied and subtle effects, further improvements in design are continually being sought and steady improvements are being achieved. Given the extremely computationally expensive codes used to model aircraft wings, current research is focused, among other things, on how more approximate or even empirical codes can be used to aid full three dimensional Euler and Navier-Stokes codes. Developments in this area are likely to be based on parallel operation of generational based search engines which selectively call different codes as the search progresses, interleaving full analyses with faster approaches. Although still at an early stage, improvements in wing design have already been achieved following these ideas and further gains can be expected.

To summarize, this paper briefly reviews the history of search and optimization in engineering design and discusses a number of areas where progress will be needed to make such codes more useful as everyday tools to practising engineers.

SHIP HULL-FORM DESIGN

The hulls of ships, particularly high-speed warships, have very complex shapes. They are a typical engineering compromise between the needs for low drag, high strength, good stability and simple manufacture. Moreover, it is very rare to prototype ships before building the final design: in essence a ship must be designed right first time. This has made the concept design of ships a topic of interest to the search and optimization community for very many years[2]. Fortunately, many of the relationships between the parameters a designer can alter and the consequent performance metrics do not vary in ways that are overly complex[3,4]. Consequently, gradient descent optimization methods may be used on problems in this field of design to good effect.

This is not to imply that this problem is trivial or one that can be tackled without some care. In the early days of such research it was impracticable to use very complex methods to analyse drag, stability, strength, etc., and recourse was usually made to empirical relationships derived from the analysis of past designs. This fact tended to limit true optimization studies to those groups with extensive historical data to call upon; often the world's naval design divisions. With increasing computing power the number of viable analysis codes available increased and so eventually a point was reached where it was desirable to combine search methods with full analysis codes based on the underlying physics.

This leads to a recurring problem in design search and optimization (DSO). The designer wishes to work with high level parameters such as length, beam, draught, form coefficients, etc., while the physics based analysis codes need full geometric data to be input. It is usually extremely tedious to generate the full data from the overall parameters (even today with a modern hull fairing system it can take an experienced user several days to generate a full description). In consequence, the ability to apply optimization methods to this problem (and many others in design) is constrained by the need to automatically generate full geometry data from a few key design parameters.

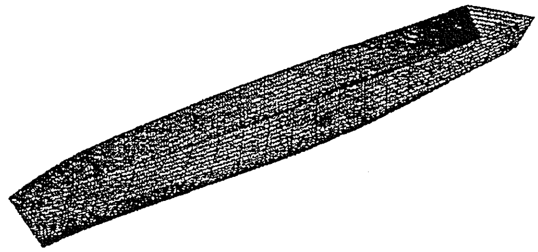

Figure 1 – frigate hull-form.

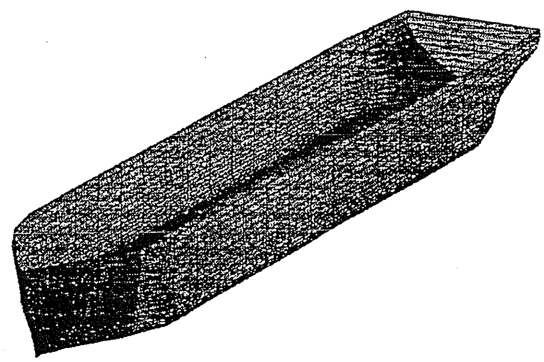

Figure 2 – tanker hull-form.

In ship hull-form design this requirement was tackled in a number of ways including the use of specialised conformal transformation methods[5]. These allowed simplified but realistic hull-forms to be created automatically over wide ranges of parameters, see Figures 1 and 2, which show views of a frigate hull and a tanker created using the same code. One important consequence of such automated geometry generation is the de facto adoption of a 'house style', i.e., all hulls generated by a particular code have a family resemblance. This has both advantages and disadvantages: being similar, such hulls readily allow trade-off studies to be made with confidence; conversely, they do not admit of specialised local modifications to suit the particular problem (such as the inclusion of bulbous bows, etc.), also problems arise when the final concept hull is transferred to the next stage of design where a detailed hull fairing package is to be used. Even so, it is apparent from the figures that very wide ranges of shape can be created with a single code using, in this case, just 19 control parameters (of which a subset of perhaps five or six would be varied during a concept design search). Such searches allow the concept design team to make key parameter choices with increased confidence.

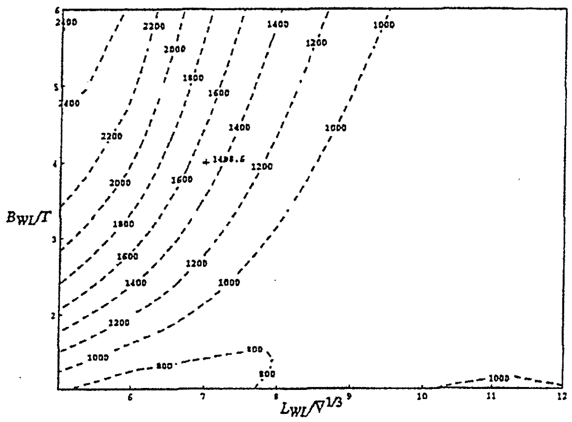

Figure 3 – contour map of hull resistance versus B_{WL}/T and $L_{WL}/\nabla^{1/3}$.

To illustrate this process take the frigate hull of Figure 1 and consider attempts to reduce the drag of this design at a speed of 30 knots by varying the beam to draft ratio (B_{WL}/T) and length to volume ratio ($L_{WL}/\nabla^{1/3}$) while maintaining adequate structural strength and stability[6]. Figure 3 shows a contour map of the hull-form resistance for this case while Figure 4 shows the effects of the constraints. It is clear from the figures that the objective function here varies quite smoothly and, as is often the case, the best design is defined by the constraint boundaries. Application of any reasonably efficient gradient descent code to this problem rapidly results in an optimal design.

Notice that to carry out this process in a realistic fashion requires that the hull structure be considered alongside the hull drag. Not evident in these figures is the fact that the system used here redesigns the hull structure as the hull shape changes so that the longer and thinner designs favoured by resistance calculations are correctly penalised by their rather worse structural efficiency. Thus, although the optimal hull has lower resistance, this is despite the fact that it has a higher structural weight. Moreover, the design case outlined here also accounts for the changes in fuel weight needed to accomplish a given mission profile (since less fuel is needed by the more efficient hull less fuel must be carried and this, to some extent, offsets the added weight of structure in the design). Clearly, such analyses must take into account a great many aspects of the design if they are not to give misleading results.

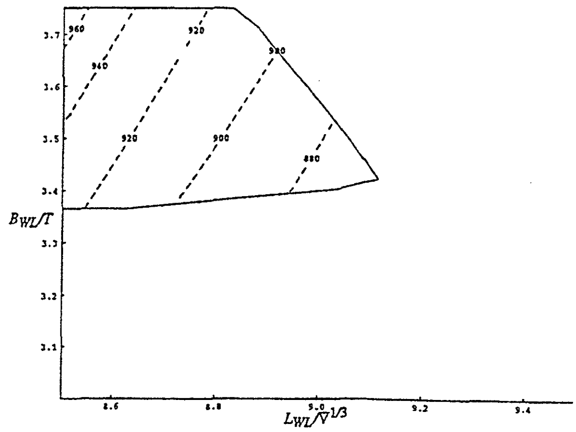

Figure 4 – detail contour map of hull resistance versus B_{WL}/T and $L_{WL}/\nabla^{1/3}$ truncated by constraints.

It is also extremely rare to carry out a single optimization run with such tools. Rather a series of linked studies based on differing assumptions and variations in the severity of constraints would normally be studied. These would then be reviewed by experienced designers who would bring to bear considerations that might not be embedded within the available analysis tools (an example in ship design is manoeuvring, where calculation methods are still far from reliable and a degree of judgement based on experience must be used).

Current research in this area is focussed on the selection of vessel type as well as leading dimensions. For example there is currently much interest in catamaran designs for naval frigates. Fast ferries are another area where various overall hull configurations are routinely considered. Such choices are often subjective in nature and fuzzy logic methods are now being used to aid decision making in this area alongside direct optimization searches.

Systems similar to the one outlined here are in now in routine use in a number of ship design offices.

SATELLITE STRUCTURAL DESIGN

The next problem considered here concerns the problem of structural vibration control in satellite booms. Such booms are common elements in a number of satellite designs, particularly so where various items of equipment must be spatially separated. The design of interferometric instruments represents one such class of design and one where extremely severe alignment requirements must be met: it is not uncommon to need the various elements of a interferometer to be held in alignment to accuracies of nano-metres with separations between the elements of tens of metres. The most demanding of these are the proposed planet finder missions scheduled for launch by NASA and ESA in around 20 years time where satellites hundreds of metres in span are being considered and where stringent vibration targets must be met over wide frequency ranges.

The structural dynamics of such systems have also been studied for many years and sophisticated numerical methods are now commonly used to predict their performance. The most common of these being the finite element analysis (FEA). Now although FEA can be highly accurate in predicting the performance of nominal structural designs, such methods suffer from being computationally expensive. They are also rather strongly affected by detail errors in manufacture such that the nominal and as-built designs may vary enough to significantly change the measured performance from that predicted by FEA.

The most common treatment for these vibration problems is to coat the structural elements with heavy viscoelastic damping materials with consequent weight and cost penalties. Moreover, the effectiveness of such treatments diminishes with the vibration levels which makes continuously improving noise and vibration targets difficult to meet. Clearly, if the vibrational energy could be contained near to the points of excitation there would be a reduced need for damping treatments and, additionally, they could be concentrated in regions where they were most effective. This can be achieved by active methods or by geometric re-design, which is the approach described here.

Figure 5 indicates the structure to be considered. This consists of a simplified, two-dimensional satellite boom made up of forty Euler-Bernoulli beams all having the same properties per unit length. Here, EA is taken as 12.86 kN, EI as 69.8MNm^2 and mass per unit length as 2.74 kg/m. The beams are all either 1 m or 1.414 m long and the joints at (0,0) and (0,1) are taken to be pinned to ground; all other joints are free to move. The structure has been chosen to be two-dimensional for simplicity – real booms are three-dimensional[7]. It is excited by a point transverse force half way between (0,0) and (1,0) and, for this study, the aim is set as the minimization of the vibrational energy level in the right-hand end vertical beam (which in practice might carry an instrumentation package or similar). The damping of the structure is fixed so that the normal modes of the uncoupled beam elements all have a constant bandwidth of 20 s^{-1}.

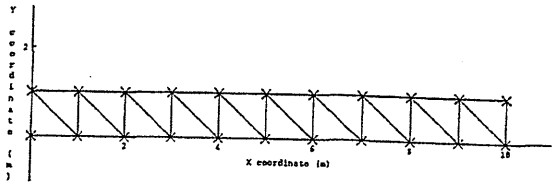

Figure 5 - a simplified, two-dimensional satellite boom.

Figure 6 shows how the vibrational energy level of the end beam varies with frequency for this base design. As can be seen from the figure, the energy level exhibits the many peaks and troughs characteristic of a complex, lightly damped structure. Moreover, there is a large response in the region 100-200 Hz. The calculation of a single frequency point in this plot requires the solution of some 260 complex simultaneous equations and this takes around 20 s on a Sun Sparc 10 workstation; i.e., the calculation of a frequency band average takes around seven minutes using the 21 point rule adopted here. Even so, this is considerably faster than can be achieved using a finite element approach with the same level of accuracy. Nonetheless, it remains a very hard problem for optimizers to deal with in a realistic time span.

Optimization Problem

The aim of this second study is set as the reduction of the frequency averaged response of the end beam in the range 150-250 Hz by the use of optimization, i.e., to trim down the upper half of the major response peak. In practice, such a requirement would reflect known sensitivities of the payload or known excitation frequencies. To meet other design requirements the optimization is constrained to keep the end beam unchanged in length and position with respect to the fixed points. Further, to meet structural requirements, all of the joints within the structure must be kept within fixed distances from their original positions. This ensures that no beam is too long or too short and also restricts the overall envelope of the structure. The free variables in the problem are thus set as the x and y co-ordinates of the 18 mid-span joints, i.e., 36 variables in all. In this case a GA† is used to tackle the optimization

† The GA works by maintaining a pool or population of competing designs which are combined to find improved solutions. In their basic form, each

problem as gradient descent methods are unable to cope with highly multi-modal nature of the problem.

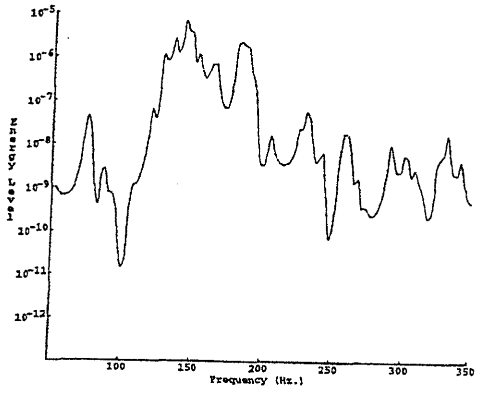

Figure 6 - initial frequency response of the end beam of the boom.

Results

To begin the optimization process relatively tight limits of ±5% are placed on the joint positions and 1,000 trials used, spread over 5 generations (giving 200 members per generation so as to ensure reasonable coverage over the rather large domain being investigated here, at the expense of final convergence). This run gives rise to the configuration shown in Figure 7. With this configuration, the total energy flow in the 100Hz band investigated has been reduced from 0.33×10^{-6} to 0.21×10^{-7} (for unit forcing), i.e., a reduction of 23dB, indicating how purely periodic structures can give rise to very significant noise transmission problems. Next, the limits on the joint positions are relaxed to ±25% but only six of the midspan joints are allowed to move, again using 1,000 trials and 5 generations. This leads to the configuration of Figure 8. It gives a slightly lesser level of isolation with the total energy flow now being 0.31×10^{-7}, i.e., a reduction of 21dB over the base design. As can be seen from Figure 8, a kind of notch has been driven into the structure and this clearly causes reflections of energy carrying waves in the desired frequency ranges

member of the population is represented by a binary string that encodes the variables characterizing the design. The search progresses by manipulating the strings in the pool to provide new generations of designs, hopefully with better properties on average than their predecessors. The processes that are used to seek these improved designs are set-up to mimic those of natural selection: hence the method's name.

Having gained some initial impressions, the GA was next applied with a ±25% constraint on the joint positions applied to the whole structure, but now using around 4,500 evaluations with 15 generations of 300 members resulting in the structure of Figure 9. Here, the energy flow has been reduced to 0.13×10^{-9}, i.e., a reduction of 68dB. The final structure is particularly interesting, showing a strong overall pattern where, the depth of the composite cantilever is increased at both ends and reduced mid-span. Such a design is clearly workable and demonstrates the power of the technique being used. Figure 10 further illustrates this point, showing the full frequency response of this design compared to the base structure, and illustrating that gains over the 150-250 Hz band have not been achieved at the expense of significantly worsened behaviour at other frequencies.

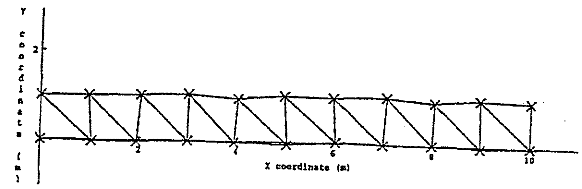

Figure 7 - optimized design with limits of ±5% on all 18 joints, 1,000 evaluations over five generations.

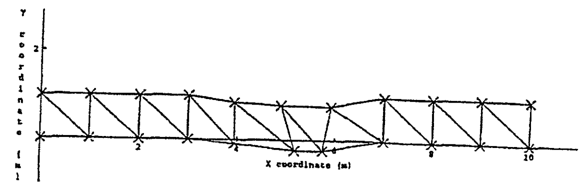

Figure 8 - optimized design with limits of ±25% on six mid-span joints, 1,000 evaluations over five generations.

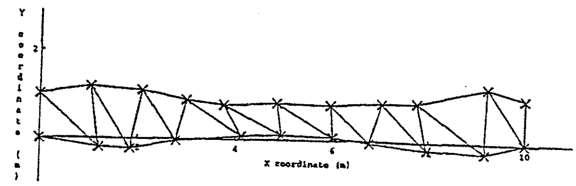

Figure 9 - optimized design with limits of ±25% on all joints, 4,500 evaluations over 15 generations.

This overview has shown that significant noise isolation characteristics can be introduced into simplified satellite structures by modifying them in a controlled way. Moreover, the structures produced are not unworkable and do not suffer from degraded performance at frequencies away from those that have been the subject of optimization. It should be noted that the improved vibration isolation characteristics gained arise from the constructive reflections of travelling waves caused by the discontinuities introduced. Part of the improvement is seen to be the result of moving away from uniform designs. However, even more significant gains arise from the unusual geometries adopted when significant changes are made. Therefore, a simple small-scale randomization, although often beneficial when dealing with a structure that originally exhibits precise geometric periodicity, does not realize all the gains that can be made. It is, however, much harder to select a particular set of large-scale modifications that control the noise performance in a particular way, while still maintaining a reasonable overall shape. Moreover, it should be noted that, even using a four processor work-station, some of the runs discussed here take more than five days of computer time to carry out. Clearly, larger structures would be even more expensive to deal with. However, when considering the cost of satellite or aircraft design and the penalties associated with redundant weight, such calculations should be well worth while. The results obtained in this theoretical study have been validated experimentally where the overall gains predicted where achieved in practice[8].

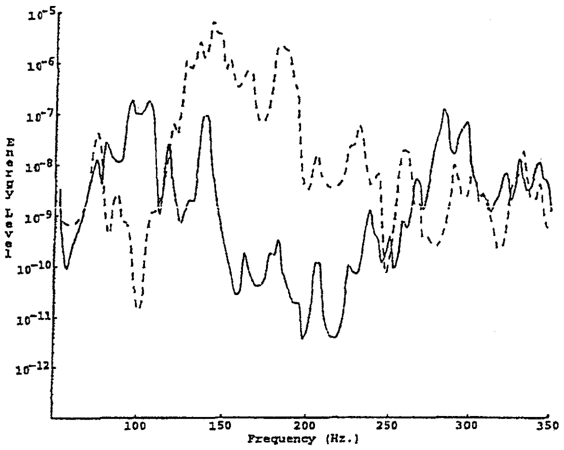

Figure 10 - frequency response of optimized design with limits of ±25% on all joints, 4,500 evaluations over 15 generations with response for initial design shown dotted.

Other studies have also shown that the Genetic Algorithm (GA) performs very well on such tasks when compared to a range of other methods, especially if large numbers of variables are used to describe the designs[9,10]. This arises because the flow of vibrational energy around a complex structure is dominated by the many resonances exhibited by such structures and also the large number of physical parameters needed to specify typical designs (in the cases studied here up to forty, many hundreds in a full ship or aircraft structural model); both features which the GA is seen to deal with well. Nonetheless, caution must still be exercised when using methods like the GA. Although beguilingly simple to describe they are, in fact, quite difficult to operate to their best effect. They are not unlike many other new and advanced techniques in this respect.

TRANSONIC WING DESIGN

Despite their long history, optimisation methods have failed to become established as an essential tool for most aeronautical designers in the way analysis methods such as FEA and CFD have, most designs being the product of engineers using analysis tools directly rather than in conjunction with optimisers.

One of the factors behind this lack of penetration by optimisation methods in this field is their tendency to neglect the finite nature of computational resources. Although the relentless improvement in computer performance allows both designer and optimiser use of increasingly sophisticated models, the available computational effort is always finite, limiting the sophistication of the models used. Even where optimisation is routinely used, there is a lag between the analysis methods adopted within optimisers and those used by human designers. For example, in conceptual wing design, optimisation studies are generally carried out using empirical models with solution times measured in tens of milliseconds. However, human designers supplement such empirical models with the sparing use of Euler and Navier-Stokes CFD codes requiring hours per solution. The take up of optimisation will remain limited while there is such a large disparity in the sophistication of the models used. To become more generally applicable, improved optimisation algorithms are therefore required which, like human designers, make use of all the available analysis methods while paying regard to the overall computational cost.

Multi-level design techniques

The key to designers' effective use of sophisticated and costly analysis methods is their co-ordinated use with cheaper, less sophisticated models. In a typical design process, preliminary design exploration is carried out using cheap analyses, perhaps as simple as 'back of envelope' calculations. As the design progresses, occasional use is made of more sophisticated analyses to guide the exploration and reduce the number of candidate designs. Finally, designs are refined using the most costly tools. Moreover, practical design is not a one way progression in the direction of increasingly sophisticated models. Designers also use simpler models to make quick checks on the results of more sophisticated analyses. The differences in analysis sophistication may result from switches in analysis method, for example between full potential, Euler and Navier-Stokes solutions in CFD. Alternately, a single analysis method may be used with different domain discretisations, for example by varying cell size in finite volume codes. Trade off between accuracy and cost may also take place within a single analysis method by switching between problem representations, for example between beam and shell, and full solid models in finite element analysis. In carrying out typical designs, engineers may use a range of analyses whose computational cost vary by several orders of magnitude.

Designers therefore switch naturally between analyses at different levels of accuracy and cost. In so doing, they balance the computational cost of each method with the confidence they place in the results. Current research effort is therefore focussed on adopting a similar 'multi-level' approach for use with optimizers. The ultimate aim being the development of algorithms which achieve the performance of the most sophisticated models at computational costs approaching those of the simplest models.

Issues in multi-level optimisation

Multi-level optimisation algorithms are still in their infancy. It is therefore difficult to make other than speculative statements as to the form of successful strategies. However, it is possible to highlight the major issues that must be addressed.

Uncertainty and evaluation cost The fundamental difference between multi-level and single level optimisation is the requirement for explicit consideration of computational cost and uncertainty of the objective function. The cost of an evaluation can be predicted with reasonable accuracy and is known retrospectively. Uncertainty is more problematic. The relative accuracy of different methods is usually not known explicitly and may vary over the design space. Much of a designer's expertise lies in the judgement of the level of confidence that can be placed in analysis results. Some of this judgement is based on checks for consistency between analyses at different levels, although most is based on experience on similar problems. Successful multi-level optimisers are also likely to consider consistency across analyses. However, to be truly effective they will be required to maintain an uncertainty model based on experience from previous optimisations and input from experienced designers.

Problem dimensionality The analysis process is one of the expansion of a few design variables to many numbers characterising behaviour within the domain which are then collapsed to the few results of interest to the designer. Choice of design variables has a profound effect on the feasibility of optimisation. In the case of multi-level algorithms, it will be necessary to enforce a consistent representation across analysis levels. This is not simply a book-keeping problem as different analyses may use radically different problem parameterisations. For example, empirical wing drag models may represent the wing by a planform and spanwise thickness distribution. On the other hand, CFD models require a complete surface representation. While it is possible to arrange for many design variables such as aspect ratio and wing area to be common to all models, the CFD models require additional aerofoil section design variables absent from empirical models. In effect the dimensionality of the design space differs according to the analysis method in use.

Multi-level algorithms must therefore allow for each analysis model utilising a different sub-set of design variables. One possible method of allowing for the changing dimensionality is combining global optimisation of those variables common to all methods with sub-optimisations of the remaining variables. For example in the case of wing design, a multilevel optimisation of planform and thickness distribution might use empirical models directly but call CFD models via a sub-process which optimises the aerofoil sections for the given planform and thickness distribution. However, it is likely that algorithms which cope seamlessly with changes in dimensionality will be more powerful.

Distortion and deception It is the fundamental premise of multi-level algorithms that there will be correlation between the topography of the objective function as viewed via different analysis methods, each of which distorts the 'true' landscape. In cases of mild distortion, locating extrema using one method will reduce the effort required to locate corresponding extrema using other methods. However, in practice there are likely to be regions of design space where some methods are subject to gross distortion such that 'true' extrema may be absent or 'false' extrema may be present. Multi-level algorithms will be required to tolerate such deception without disrupting the search.

Allocation of resources between levels A multi-level algorithm must also allocate the limited computational effort between analysis levels efficiently. The most efficient allocation for two hypothetical extreme cases is straightforward. If no distortion exists between levels, the most efficient allocation is exclusive use of the cheapest level. At the opposite extreme of cheaper analysis levels being totally deceptive, the computational effort would be devoted entirely to the most expensive level. In realistic cases between these extremes, it is unlikely that it will be possible to ascertain the most efficient allocation *a priori*. Rather it is to be expected that efficient algorithms will dynamically allocate effort during the optimisation, the allocation between levels being likely to differ between different regions of the design space.

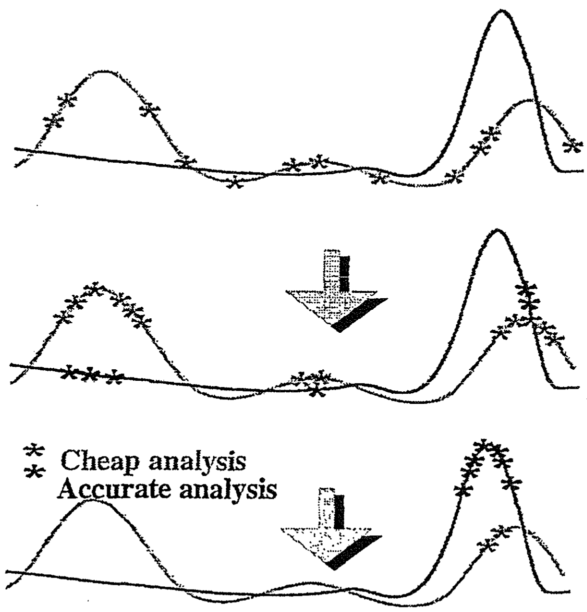

Figure 11 - Idealised progression of a two-level optimization.

It would perhaps be expected that a multi-level optimisation algorithm would initially use mainly cheap evaluations with occasional calls to more expensive methods. As the optimisation progresses we would expect interleaving of an increasing proportion of more expensive solutions until, at convergence, almost all evaluations would use the sophisticated model, with perhaps occasional use of simple models as a reality check (Figure 11 depicts such a process with two levels).

Possible multi-level strategies

Comparatively little work has addressed these issues and it is therefore too early to draw conclusions. Nevertheless, the authors believe that the flexibility of population based algorithms are likely to render them more suitable candidates as a basis for multi-level strategies than gradient-based algorithms. The few investigations so far reported have adopted such an approach with population members farmed out between levels.

Elby *et al.*[11] reported the use of a multi-level injection island genetic algorithm in design of a minimum weight flywheel, with each island population evaluated using finite element models of different element size. They observed a significant advantage over single level algorithms. However, the degree of deception in the problem was probably limited and the algorithm's performance in the presence of significant deception has yet to be reported.

El-Beltagy and Keane[12] have investigated application of three basic multi-level strategies to a variety of population based optimisation algorithms. The aim was maximisation of a test function in which significant distortions had been applied to notionally cheaper levels. The tests showed little gain for the multi-level strategies over exclusive use of the expensive undistorted level, with only genetic algorithms showing any consistent improvement and then only of order 5-20%. Their work indicates that overcoming significant distortion requires more advanced algorithms.

Work on more sophisticated algorithms is continuing at Southampton with investigations of the use of self-organising maps for controlling determination of evaluation level and with the application of ecological models to genetic algorithms[13]. It is too early to report on the effectiveness of these and other strategies.

A Multi-level Wing Design Environment

In order to investigate the issues involved in multi-level optimisation, Southampton University has developed a system to act as a test bed for multi-level algorithms on a realistic design task. This is the preliminary conceptual design of civil transport aircraft wings. One of the principle tools used by British Aerospace in this task is the TADPOLE program[14] which is based on empirical models. The computational cost of FEA and CFD analyses has so far excluded their routine use from this stage of the design process, when large design spaces must be searched. It is hoped that effective multi-level optimisation algorithms would allow the empirical models to be supplemented by more expensive numerical methods.

The simplified design objective is drag minimisation subject to constraints on wing weight, fuel tank volume, pitch-up margin, and root triangle volume, all of which are currently estimated empirically. It is intended that the empirical wing weight model will be supplemented by finite element models, allowing multi-level multidisciplinary structural and drag optimisation.

Drag estimation Three levels of analysis are available (Figure 12). The cheapest is a cut down version of the British Aerospace TADPOLE empirical code, which estimates a drag breakdown at a cost of approximately 50 milliseconds of compute time. The drag values returned are for a 'well designed' wing, the aerofoil sections of which are not specified explicitly.

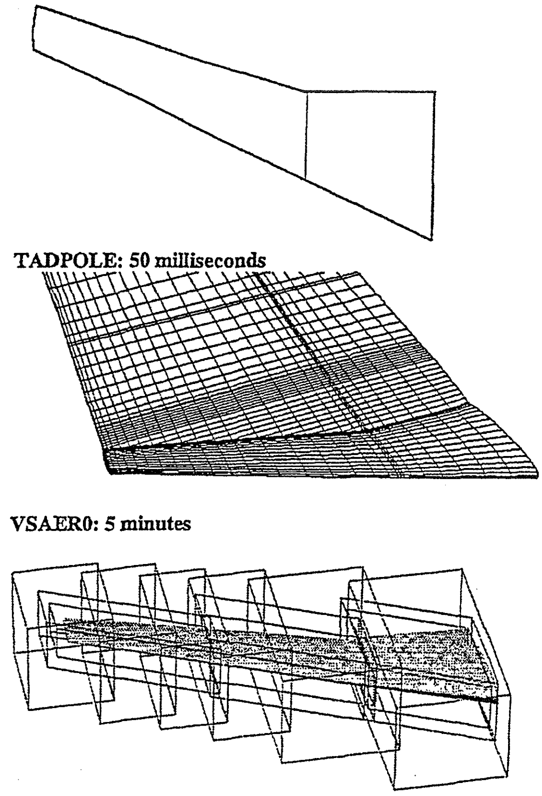

Figure 12 - The three drag analysis levels.

The most expensive method uses the viscous coupled multi-grid Cartesian Euler code MGAERO[5]. Drag estimation using this code costs approximately two hours of compute time. The induced drag is obtained from integration of vorticity immediately downstream of the trailing edge, the wave drag using Lock's second order method[16] and the profile drag using Cooke's implementation of Squire and Young's approximation[17].

In terms of computational cost, the third method lies between these extremes. It uses the viscous coupled boundary element code, VSAERO[18], requiring approximately five minutes of compute time per evaluation. The induced drag is calculated from Trefftz plane integration, and the profile drag again using Cooke's implementation. The linearised potential equations preclude wave drag. However, if the optimisation algorithm requires this absent drag component, a routine applying simple sweep theory and an aerofoil drag model provides a wave drag estimate from the calculated spanwise loading.

Parameterisation The parameterisation of wing planform and thickness is presented in table 1, the parameters being common to all three drag estimation levels. In addition the CFD codes require further parameters specifying aerofoil section. The aerofoil parameterisation utilises orthogonal basis functions derived from an existing aerofoil family[19], the first three of which are illustrated in Figure 13.

Symbol	Definition
S	Area
AR	Aspect ratio
Λ	Leading edge sweep
λ_1	Inner panel taper ratio
λ_2	Outer panel taper ratio
η_k	Trailing edge kink position
t/c_r	Root thickness/chord
t/c_k	Kink thickness/chord
t/c_t	Tip thickness/chord

Table 1: Wing design variables

Although these functions allow less design variation than B-spline or Bézier representations, they are ideal for conceptual design where the aim is selection of sections appropriate for local flow conditions. The spanwise variation of the orthogonal function weights is represented by between four and ten parameters. For each drag code, routines expand from the parameter values to generate appropriate input files. In the case of the boundary element code these specify surface and wake panel geometries. Similarly, for the Euler solver, the files specify the surface discretisation and the Cartesian grids.

Test bed characteristics The combination of the three drag codes provides the multi-level algorithm test bed with some interesting characteristics. The codes span a 150,000 fold range of computational cost, the two CFD codes having a 25 fold cost difference.

The relative accuracy of the three codes is less easily quantifiable. Drag values from MGAERO are subject to the usual inaccuracies associated with drag recovery from Euler codes[20]. However, for the purposes of evaluating multi-level algorithms the predictions of the Euler code are assumed to be accurate. The objective of the optimisations is therefore identification of designs with minimum drag according to the Euler code. TADPOLE is calibrated for conventional wing designs. For unconventional designs the validity of assumptions implicit in the empirical models may be limited, resulting in significant differences from the Euler predictions. Although VSAERO is capable of reasonable drag predictions at low Mach numbers, the use of an essentially subsonic boundary element code for drag prediction at transonic Mach numbers would not normally be contemplated. However, from the point of view of testing of multi-level algorithms this provides an interesting distortion with Mach number.

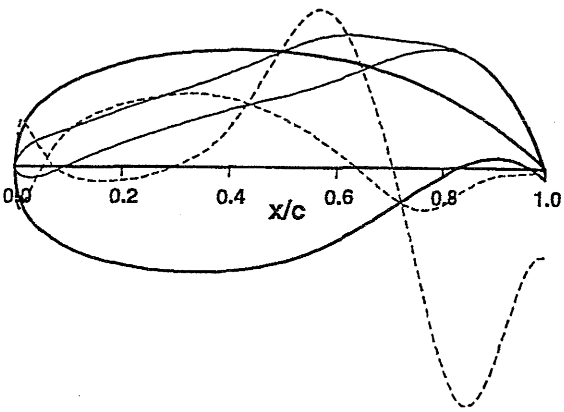

Figure 13 - First three orthogonal basis functions.

Preliminary results The system is still in an early stage of development and therefore only limited results have so far been obtained. However, results from an elementary optimisation serve to illustrate the potential of multi-level optimization. A simple test problem was constructed with the aim of optimising a civil transport aircraft wing for operation at Mach 0.785 and a Reynolds number of 7.3 million. The objective was minimisation of wing D/q as calculated by VSAERO with target lift, wing weight, volume, pitch-up margin and root triangle area chosen to be representative of a 220 seat wide body airliner. Three optimisation runs were carried out, each limited to 500 VSAERO evaluations. This would normally be considered rather a small number of evaluations for meaningful optimization of such a problem, but would be a typical limitation on the use of the most expensive code.

Two single level optimizations were carried out. The first of these ran a genetic algorithm, population size 100, for five generations. The second used the planform and thickness distribution of the DLR F4 wing[11] as a starting point for 500 evaluations of a gradient descent algorithm[12].

The third optimization used a sequential two level strategy in which the result from an explorative TADPOLE optimisation was used to seed a VSAERO optimisation. A preliminary genetic algorithm optimisation was run for 40 generations using the TADPOLE code, the cost of the 2000 evaluations being less than one VSAERO calculation. The optimum point from the TADPOLE search was then used as the starting point for a 500 evaluation gradient descent search. The results from the three searches are summarised in table 2.

	Search and analysis method	VSAERO D/q
1	Gradient search, 500 VSAERO evaluations	3.06 m^2
2	Genetic Algorithm search, 500 VSAERO evaluations	2.88 m^2
3a	Genetic Algorithm search, 2000 TADPOLE evaluations ↓	2.78 m^2
3b	Gradient search, 500 VSAERO evaluations	2.75 m^2

Table 2 - Results of three wing optimizations

It is clear that, on this problem, the sequential two level strategy utilises the limited VSAERO evaluations more effectively than either of the single level strategies. However, the constraints and design objective of the problem presented here were much simplified and it is unlikely that such a simple strategy would be as effective on more realistic problems. Future work will address the requirement for more sophisticated multi-level strategies on these more demanding problems.

CONCLUDING REMARKS

This brief review of design search and optimization problems has spanned three domains of engineering and over twenty years of effort. It demonstrates that computer aided search can be an important tool in the design engineer's armoury. However, it is also clear that such approaches have to be quite sophisticated to deal with the full range of problems encountered in the design of complex systems. This tends to mean that such approaches cannot be inserted into design organizations without considerable thought and planning. A few of the key topics that must be considered when attempting to get the best out of such techniques may be mentioned, however.

1. Search and optimization tools must be seen as *aids* to the designer rather than automated ways of replacing him or her.
2. Such tools must be able to interface to a wide variety of analysis codes, at differing levels of complexity without major effort by the designer.
3. A range of appropriate search engines must be available and assistance provided on how to use them and in what order.
4. Appropriate representation, which is the way a design is viewed by a search engine is of vital importance in producing workable tools.
5. A clear view of the available computational resources must always be kept in mind.

Work continues in this field and there seems little doubt of the increased utility of such methods.

REFERENCES

1. J. N. Siddall, Optimal Engineering Design: Principles and Applications, Marcel Dekker, Inc., New York (1982).
2. P. Mandel and R. Leopold, ``Optimization Methods Applied to Ship Design," Trans. S.N.A.M.E., pp. 477-521 (1966).
3. G. M. Watson and A. W. Gilfillan, ``Some Ship Design Methods," Trans. R.I.N.A. 118 pp. 279-324 (1976).
4. J. Holtrop and G. G. J. Mannen, ``An Approximate Power Prediction Method," International Shipbuilding Progress 29(335) pp. 166-170 (1982).
5. A.J. Keane, ``A Computer Based Method for Hull Form Concept Design: Applications to Stability Analyses," Trans. R.I.N.A. 130(A) pp. 61-75 (1988).
6. A.J. Keane, W. G. Price and R. D. Schachter, 'Optimization Techniques in Ship Concept Design', Trans. R.I.N.A. 132(A) 1990.
7. A.J. Keane, ``Passive vibration control via unusual geometries: the application of genetic algorithm optimization to structural design," J. Sound Vib. 185(3) pp. 441-453 (1995).
8. A. J. Keane and A. P. Bright, ``Passive vibration control via unusual geometries: experiments on model aerospace structures," J. Sound Vib. 190(4) pp. 713-719 (1996).
9. A. J. Keane, ``Experiences with optimizers in structural design," pp. 14-27 in Proceedings of the Conference on Adaptive Computing in Engineering Design and Control 94, ed. I. C. Parmee,P.E.D.C., Plymouth (1994).
10. A. J. Keane and S. M. Brown, ``The design of a satellite boom with enhanced vibration performance using genetic algorithm techniques," pp. 107-113 in Proceedings of the Conference on Adaptive Computing in Engineering Design and Control 96, ed. I. C. Parmee,P.E.D.C., ISBN 0 905227 61 1, Plymouth (1996).
11. Elby, D., Averill, R. C., Punch, W. F., and Goodman, W. F., Evaluation of injection island GA performance on flywheel design optimization, *Adaptive Computing in Design and Manufacture*, Ed. I. Parmee, Springer-Verlag, 1998.
12. El-Beltagy, M. A. and Keane, A. J., A comparison of genetic algorithms with various optimization methods for multilevel problems, *Engineering Applications of Artificial Intelligence*, to appear.
13. Robinson, G. M., Elbeltagy, M. A., and Keane, A. J., Optimization in mechanical design, in *Evolutionary Design by Computers*, Ed. P. Bentley, Academic Press, to be published 1999.
14. Cousin, J., and Metcalfe, M., The BAe Ltd transport aircraft synthesis and optimisation program, AIAA 90-3295, September 1990.
15. Epstein, B., Luntz, A., and Nachson, A., Multigrid Euler solver about aircraft configurations, with cartesian grids and local refinement, AIAA paper 89-1960, 1989.
16. Lock, R. C., Prediction of the drag of aerofoils and wings at high subsonic speeds, *Aeronautical Journal*, June/July 1986.
17. Cooke, T. A., Measurements of the boundary layer and wake of two aerofoil sections at high Reynolds numbers and subsonic Mach numbers, *ARC R&M 3722*, 1973.
18. Maskew, B., Prediction of subsonic aerodynamic characteristics: A case for low-order panel methods, *J. Aircraft*, **19**, No. 2, Feb. 1982.
19. Robinson, G. M. and Keane, A. J. Orthogonal representation of a family of aerofoils, *In preparation*.
20. van Dam, C. P., Nikfetrat, K., Accurate prediction of drag using Euler methods, *J. Aircraft*, **29**, No. 3, March, 1992.
21. Redeker, G., DLR-F4 wing body combination, *A selection of experimental test cases for the validation of CFD codes*, AGARD AR-303, August, 1994.
22. Rosenbrock, H. H., An automatic method for finding the greatest or least value of a function, *Computer Journal*, **3**, No. 3, 1960.