Multi-variable Geometry Repair and Optimization of Passive Vibration Isolators

Alexander I.J. Forrester* and Andy J. Keane†
University of Southampton, Southampton, Hampshire, SO17 1BJ, UK

A range of techniques are considered for the search of a high dimensional design landscape with extensive, unknown and disjointed regions of infeasibility. We present the use of a hybrid genetic-algorithm / gradient-descent search of the objective function / feasibility problem. The genetic algorithm is used to optimize the vibration isolation of a novel passive structure concept, while the gradient descent method is used to repair infeasible geometries. For a more complicated structure the gradient search fails to find feasible geometries and we resort to a non-dominated sorting multi-objective genetic algorithm which searches the vibration isolation and geometry feasibility simultaneously. Although the complicated geometry is more difficult to optimize, there is potential for significant improvements in vibration isolation.

I. Introduction

Ideally in any optimization procedure the optimizer should visit wide ranges of designs and move between them without interruption. In reality the uniform coverage of a design space by a generic geometry model and simulation procedure is fraught with difficulties. It may be that the geometry model is not generic enough to cope with all possible permutations of inputs or that the simulation process fails, perhaps due to separation in a fluid dynamics problem, for example. Also, some designs could simply be physically unrealisable. In such situations the designer may be forced to reduce the bounds of the problem and risk ignoring potential improvements from radical designs. This negates the possible benefits of global optimization routines, and the designer may have to resort to a local optimization around a known good design.

The problem of unknown regions of infeasibility has been tackled before by using a statistical model to penalise infeasible designs.[1] The same reference also used a modified genetic algorithm (GA) in which the parents survive in cases where offspring are infeasible. Another possible method is to consider the design feasibility as a separate optimization problem. As well as optimizing the designs in terms of the primary objective, we also optimize the feasibility. Sóbester *et al*[2] used a statistical model of the quality of CAD generated designs. This model was then searched to find the smallest possible repair alteration that could be applied to a poor quality design. The example we present here is that of the design of a novel passive vibration isolation structure where, rather than using a statistical model, we are able to search a feasibility metric directly.

We begin in the next section with an outline of our vibration isolator concept and objective function formulation. The following section presents the calculation of a geometry feasibility metric for the isolator. We then go on, in sections IV to VI, to look at search algorithms which optimize the vibration isolation whilst ensuring feasible designs are produced, before drawing conclusions in the final section.

II. Optimization of vibration isolation

Our vibration isolator concept follows on from the the work of Keane *et al*,[3–5] who showed that the geometry of a truss structure could be modified such that it possesses inherent vibration isolation characteristics. An optimized truss is shown in Fig. 1. In the present study a triangular truss is 'folded' so that each section (or 'bay') extends into the previous one. An example 'regular' structure with three bays is shown in

*Member AIAA. Research Fellow, Computational Engineering and Design Group.
†Professor of computational engineering, Computational Engineering and Design Group.

Fig. 2. The geometry of this structure can be optimized in the same way as the truss structures of Keane *et al* to provide vibration isolation between the upper mounting points and the lower base points. The absence of visco-elastic elements in such mountings makes the design lightweight and suitable for applications where accurate instrument alignment and stability are required in adverse environments (e.g. satellite borne equipment).

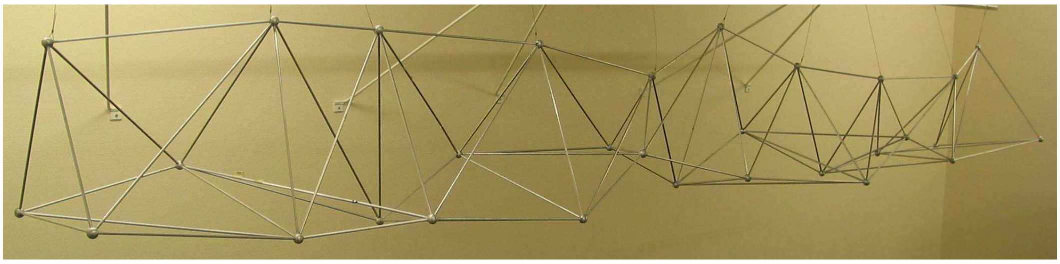

Figure 1. Optimized satellite truss structure.

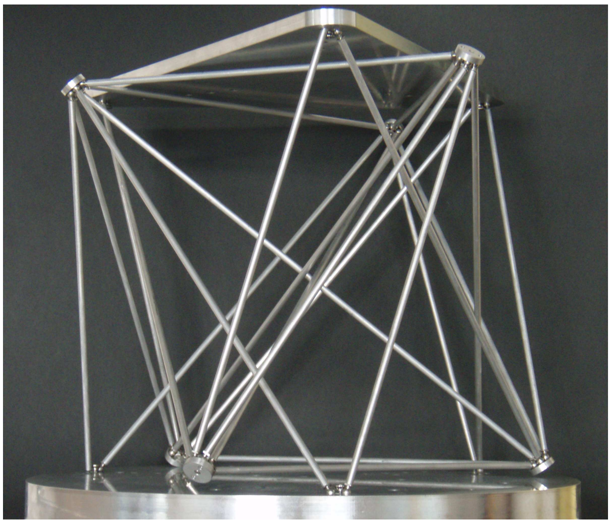

Figure 2. Example three bay structure.

Our optimum structure will be that which minimizes the vibration at the top three mounting points when the structure is subjected to unit force vibration between 150 and 250 Hz at the base three points (a frequency range chosen to meet a class of space flight applications). We calculate the vibration energy in the beams using a receptance theory based code.[6] This provides a significant speed up over finite element analysis when calculating vibrations across a frequency range. With over 200,000 function evaluations required in our optimization, fast computation of the objective is paramount. We calculate the vibration energy in the top three beams of the structure (these three beams are given a very high stiffness and a mass which equals that of the plate seen in Fig. 2) at 21 points across the 150-250 Hz range and take the average. Fewer samples may result in aliasing, while more will, naturally, increase computation time. Our objective function to be minimized is thus given by:

$$\frac{\sum_{i=1}^{21}\sum_{j=n-2}^{n} e_{i,j}}{63}, \tag{1}$$

where $e_{i,j}$ is the energy in the j^{th} rod at frequency i and n is the number of rods in the structure.

Each joint of the structure represents three (x, y, z) optimization variables, with our simplest three bay structure having a total of 18 variables. Given the vibrational behaviour of lightly damped structures and the high dimensionality of the design space, it must be assumed that the objective function landscape will be highly multi-modal. We must therefore use a global search technique such as a genetic algorithm (GA) to get the best quality designs. The Euler-Bernoulli beam analysis does not suffer from failures in the same way as we would expect, for example, in using a complex CAD surface geometry followed by computational fluid dynamics. Here we are concerned purely with physically unrealisable designs due to rod intersections in the geometry.

III. Geometry feasibility

It is apparent from the three bay structure in Fig. 2 that many geometries, particularly those with a higher number of bays, will suffer from intersections of one part of the structure with another. Indeed, from a sample of randomly generated designs, it has been found that only 0.8% of three bay and fewer than 0.02% of five bay geometries are feasible. Clearly we need a systematic way of identifying feasible designs, which are by the nature of the problem disjointed, with pockets of feasibility spread across a wide area of the design space.

Intersections in the geometry are found by first calculating if two rods, defined as lines with ends given by the vectors $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}_3$, $\mathbf{x}_4$, are in the same plane. The lines are in the plane if

$$(\mathbf{x}_1 - \mathbf{x}_3) \cdot [(\mathbf{x}_2 - \mathbf{x}_1) \times (\mathbf{x}_4 - \mathbf{x}_3)] = 0. \tag{2}$$

If so, and the lines are not parallel, we then calculate the point of intersection as:

$$\mathbf{x} = \mathbf{x}_1 + (\mathbf{x}_2 - \mathbf{x}_1) \frac{[(\mathbf{x}_3 - \mathbf{x}_1) \times (\mathbf{x}_4 - \mathbf{x}_3)] \cdot [(\mathbf{x}_2 - \mathbf{x}_1) \times (\mathbf{x}_4 - \mathbf{x}_3)]}{\mid (\mathbf{x}_2 - \mathbf{x}_1) \times (\mathbf{x}_4 - \mathbf{x}_3) \mid^2}. \tag{3}$$

If, as in most cases, the rods are not in a plane, the minimum distance between a pair of skew lines is given by

$$\epsilon = \frac{\mid (\mathbf{x}_3 - \mathbf{x}_1) \cdot [(\mathbf{x}_2 - \mathbf{x}_1) \times (\mathbf{x}_4 - \mathbf{x}_3)] \mid}{\mid (\mathbf{x}_2 - \mathbf{x}_1) \times (\mathbf{x}_4 - \mathbf{x}_3) \mid}. \tag{4}$$

This distance is along the vector

$$\mathbf{v} = \frac{[(\mathbf{x}_2 - \mathbf{x}_1) \times (\mathbf{x}_4 - \mathbf{x}_3)]}{\mid (\mathbf{x}_2 - \mathbf{x}_1) \times (\mathbf{x}_4 - \mathbf{x}_3) \mid}. \tag{5}$$

To find the point of minimum distance, $\mathbf{x}_3$ and $\mathbf{x}_4$ are transformed to $\mathbf{x}'_3$ and $\mathbf{x}'_4$ along vector $\mathbf{v}$ by the distance ϵ. The point of minimum distance is then the point of intersection of the lines $\mathbf{x}_1$, $\mathbf{x}_2$ and $\mathbf{x}'_3$, $\mathbf{x}'_4$.

If the lines are in the plane and intersecting or the distance ϵ is less than the rod diameter (plus a tolerance) *and* the point of minimum distance satisfies

$$\min\{\mathbf{x}_1(1), \mathbf{x}_2(1), \mathbf{x}'_3(1), \mathbf{x}'_4(1)\} < \mathbf{x} < \max\{\mathbf{x}_1(1), \mathbf{x}_2(1), \mathbf{x}'_3(1), \mathbf{x}'_4(1)\} \tag{6}$$

an intersection has occurred. We work though the whole structure finding ϵ for each pair of rods and then calculate our feasibility function as

$$\varepsilon = \sum_j \sum_i (\phi - \epsilon_{i,j}), \tag{7}$$

where $\epsilon_{i,j}$ is the distance between the centrelines of rods i and j (calculated using equation 4), and ϕ is the distance we are aiming for: the diameter plus a tolerance (here we have used a tolerance of 2.5 mm, i.e. we want the rods to be at least 2.5 mm apart). When $\varepsilon = 0$ we have a feasible geometry.

IV. Black box approaches for infeasible designs

While for the structure considered here we have available a direct way of computing the feasibility of a design, often the analysis of a design will fail for a number of different, unpredictable reasons. The geometry may be too complex to build rules describing physical feasibility of the design, where meshing is involved this may fail, or the solution may fail to converge. Such cases where there is no measure of the degree of failure have led to the development of search routines that can learn which areas of the design space are feasible based simply on the success or failure of sampled designs.

Forrester *et al*[1] used a statistical model where regions of feasibility are isolated by penalising failed simulations using Gaussian process based error estimates. We have applied this method to the vibration isolator problem with only two variables – the position of the first joint in the x-plane. Figure 3 shows the feasibility of this two variable design space, calculated from 1681 evaluations of the intersection code, along with the position of the design with optimum vibration isolation (this result is from the search in section V). Figure 4 shows a statistical model built using a sample of 20 designs. Five of these (shown as black dots) are feasible and the 15 infeasible designs (shown in red) have imputed vibration energy values which are penalised with Gaussian process based error estimates. The model correctly predicts the region of the feasible global optimum and an update calculation has been applied here (shown in green).

Although this method will always converge towards the global optimum, it is clear that it will take a considerable time for small isolated regions of feasibility to be sampled. The method has been shown to be effective on problems with up to 80% infeasibility, but beyond this it is likely to be slow, since it relies on an initial sample which contains a number of feasible sample points upon which to build the statistical model. The two variable example here has 79% infeasibility, but varying all 18 variables of the three bay structure leads to 99.2% infeasibility. The time required to find sufficient feasible designs to build an initial model and the amount of training data (this includes the infeasible designs) will make the method prohibitively expensive.

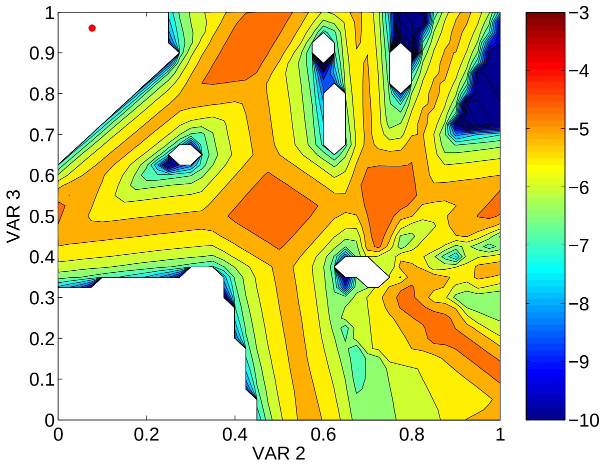

Figure 3. $\log_{10}\varepsilon$ for the three bay structure as joint one is moved in the x-plane. Blank regions indicate $\varepsilon = 0$ (feasible designs).

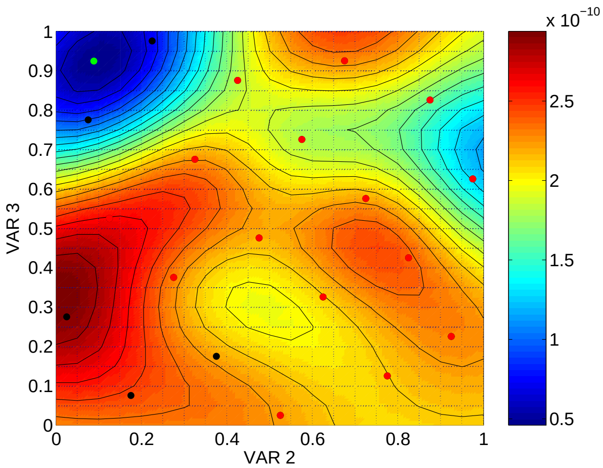

Figure 4. A statistical model based on vibration energies from five feasible designs (black dots) and imputed values for 15 infeasible designs (red dots). An update based on the maximum expected improvement[7] in the vibration energy is shown as a green dot.

Indeed, considering the feasibility as a black box process in problems of very high infeasibility (> 99%) will always lead to problems. As well as using a statistical model based search, we can modify a direct search of the problem to cope with infeasibility. A GA in which the parents survive in cases where offspring are infeasible also performs well for moderate to high regions of feasibility,[1] but it too relies on the selection of a feasible initial population, which must be based on extensive random or space filling sampling. This selection process may be extremely time consuming and, in the same way as the sampling in Fig. 4 omitted areas of feasibility, the initial population of the GA may also do so, leading to a loss of population diversity and an inefficient search. Similar problems will also be encountered if a penalty function is applied to infeasible population members.

Where a measure of the feasibility of designs is available it can be used to significantly enhance the efficiency of a search algorithm. In the following two sections we consider two ways in which this may be achieved.

V. Hybridising the GA with geometry repair

With a large number of variables to be optimized (a minimum of 18 for the simplest three bay structure), the quick run-time of the receptance code, and the highly multi-modal character expected of the design landscape of a lightly damped structure, it is unlikely that the search will benefit from the use of a statistical surrogate of the landscape. We therefore search the problem directly using a GA.

When the GA population contains an infeasible design we wish to find a similar design which is feasible and substitute it into the population. We therefore perform a local search of the feasibility metric, starting from the infeasible member of the population, using a gradient descent method – BFGS.[8] With such a large number of variables, the time taken for finite difference gradient calculation becomes significant. The feasibility function is calculated in Matlab®and the code can be automatically differentiated using MAD[9] to reduce the expense of obtaining gradients.

After the generation of a GA population, each individual is checked for feasibility and repaired if necessary. The new design variables, along with corresponding vibration energy calculations, are then fed back into the GA for selection of the next population. In evolutionary terms this is a hybrid of Darwinian and Lamarkian theory. There is a good deal of literature on such methods, with a survey available in the dissertation of Sóbester.[10] Previous research concentrates on hybrid global/local search of one objective and it is not always apparent how much benefit there is in introducing anti-Darwinian ideas to an evolutionary search. However, here we are considering two objectives, one of which is eminently suited to Darwinian based algorithms (the high dimensional and extremely complex vibration isolation design landscape), the other demanding a gradient search.

A. Results

The results of the method are shown in Fig. 5 and Fig. 6. The first figure shows the progress of the hybrid optimization of a three bay structure over 100 generations with a population of 50. The best design at each generation is displayed as well as the individuals. This shows that there has been no loss of population diversity due to the repair algorithm, and also that there is constant improvement throughout the history of the optimization. The second figure shows the level of geometry infeasibility at each generation. Because we are using a gradient descent method, we cannot be assured of a successful repair – there may be a minima of ε in which the search becomes trapped, but does not represent a feasible geometry (e.g. the basins containing infeasible deigns in Fig. 3). However, as the search progresses the number of infeasible designs diminishes and, since we have included a 2.5 mm tolerance, all designs after the first 18 generations could in fact be built.

Figure 7 shows the optimized structure which has been built for shaker testing to validate the theoretical results. The receptance code analysis predicts an average attenuation in the 150-250 Hz frequency range across the structure of -26 dB. The results of a shaker test of this structure, shown in Fig. 8, indicate that this attenuation is realised over much of the frequency range. However, the presence of a significant resonant peak at 226 Hz reduces the band averaged attenuation to 0.7 dB. The attenuation over the 'regular' structure in Fig. 2 is -2.4 dB. The prototype structure in Fig. 7 has a mean bay length of 0.3 m. The large scale structure facilitates the manufacture and test of many geometries using a generic joint design. It is envisaged that scaled down versions could be used in satellite applications, but would require more advanced manufacturing, such as lost wax casting.

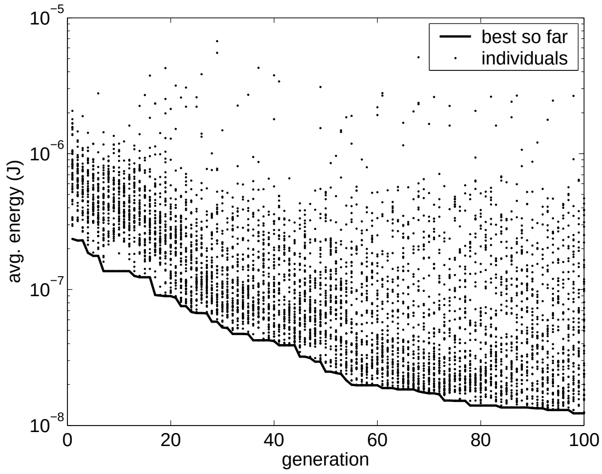

Figure 5. Progress of the hybrid optimization.

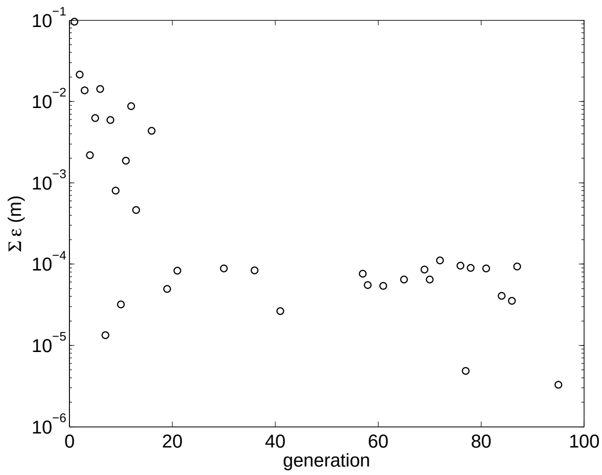

Figure 6. Level of geometry infeasibility, $\sum_{ind=1}^{50} \varepsilon$, throughout the hybrid optimization.

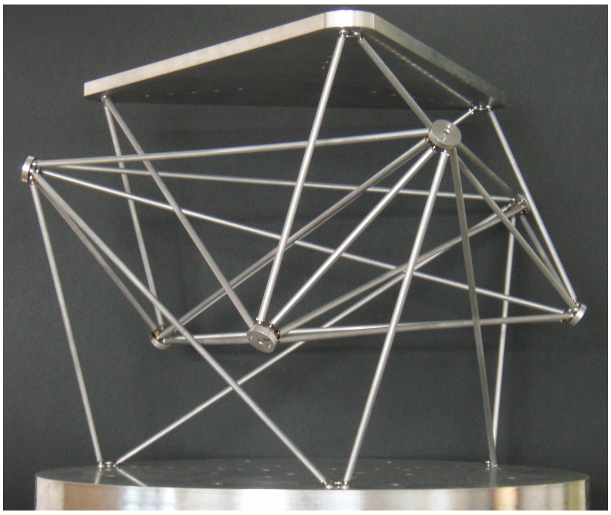

Figure 7. The optimized three bay structure.

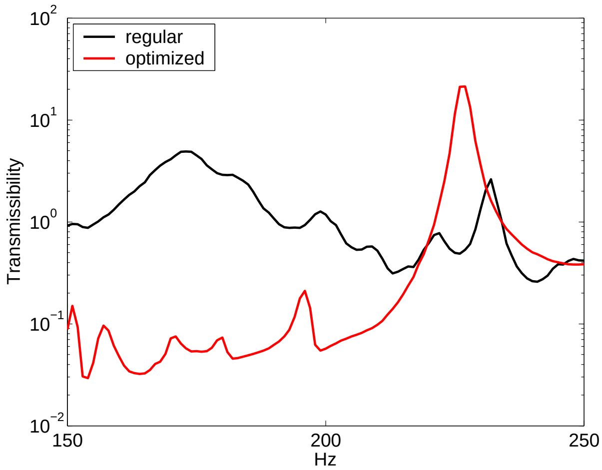

Figure 8. Results of shaker tests (using a random excitation of 0.01 g^2/Hz) of the 'regular' and hybrid optimized three bay structures. Results are averaged readings from three accelerometers placed above the three top plate mounting points.

VI. Multi-objective vibration / feasibility optimization using NSGA-II

The success of the gradient search of feasibility which is hybridised with a GA relies on the basins of attraction of the feasibility landscape containing feasible designs. If there are many local optima which are not in fact feasible the search will be slow. The cross section of the feasibility landscape shown in Fig. 3 reveals that there are seven basins which contain feasible designs and six which do not. Although there are a number of basins without feasible designs, the contours show that a steepest descent search will find a feasible design if started from roughly three quarters of the design space.

While the hybrid method has been seen to perform well on such a landscape (though in 18 dimensions), Fig. 9 shows that the equivalent landscape for a five bay structure is far more complicated. There are 15 local minima which do not contain feasible designs. Although the single region of feasible designs is fairly large, it is clear from the contour plot that many gradient searches would fail to find this area. In practice a 36 variable GA / gradient descent search of the five bay structure fails to find feasible designs. Local feasibility searches are required for all members of the population and the repair process dominates the time taken for the optimization.

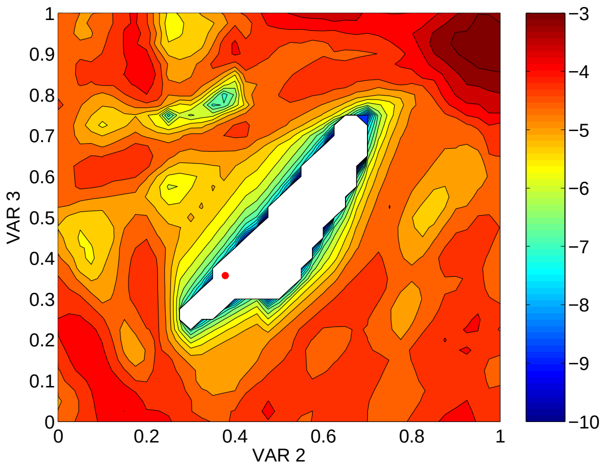

Figure 9. $\log_{10} \varepsilon$ **for the five bay structure as joint one is moved in the** x**-plane.**

With the feasibility landscape now too complicated for local repair, we transform the problem into a dual objective formulation and use the NSGA-II (non-dominated sorting genetic algorithm) method.[11,12] For gradient based geometry repair ε is a suitable continuous metric to search. When using a genetic algorithm it is possible to instead search the number of rod intersections. This is a more useful objective to the designer since, if a significant improvement in vibration could be attained at the expense of a handful of intersections, it may be worthwhile constructing bridges across the intersections.

A. Results

Figure 10 shows the Pareto front of non-dominated solutions of 100 generations of a NSGA-II search with a population size of 100. Although a geometry with up to four intersections offers better isolation than the feasible geometry, considering the most infeasible design of this search had 201 intersections, the two objectives are not in obvious tension with each other, it is purely that so much of the design space is infeasible that we expect that designs with better vibration isolations may contain intersections.

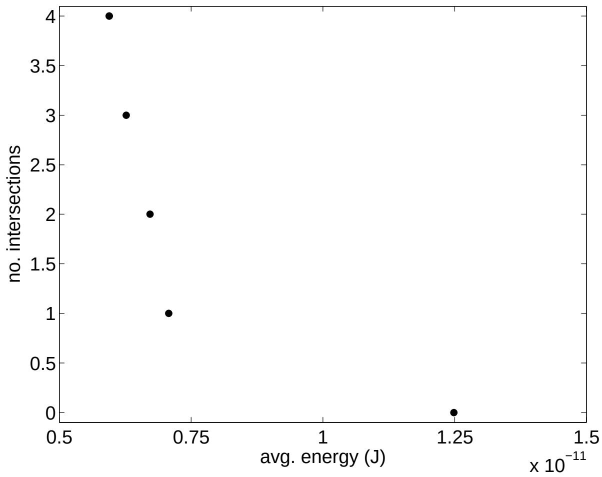

Figure 10. Pareto front showing non-dominated designs from the NSGA-II search.

The designs representing the points at either end of the Pareto front are shown in Fig. 11 (the design with no intersections) and Fig. 12 (the design with the best vibration isolation).

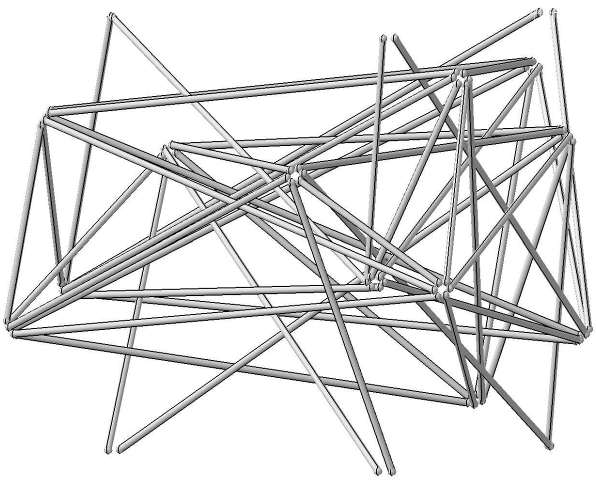

Figure 11. A CAD representation of the best feasible design from the NSGA-II search (joints are omitted for clarity)

Assuming that only designs with no intersections can be made, the design in Fig. 11 still offers a significant improvement in vibration isolation compared to the optimum three bay structure. This follows the findings of Keane *et al*,[4] who saw that each irregular bay of the structure acts as a filter for a certain

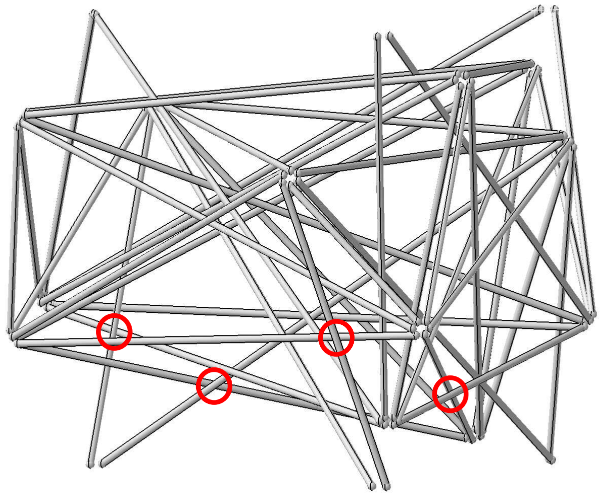

Figure 12. A CAD representation of the design with the minimum vibration energy from the NSGA-II search (joints are omitted for clarity). The four intersections are circled in red.

frequency range, with the whole structure exploiting the cumulative effect of a series of bays. As the number of bays increases beyond five, we expect the vibration isolation of the folded structure to increase likewise. However, there will naturally become a point at which no feasible designs can be found for structures built simply from straight rods.

VII. Conclusions and further work

We have shown the successful application of a hybrid GA/BFGS optimization/repair algorithm and a multi-objective GA to the design of a novel passive vibration isolator. As the complexity of the structure increases geometry repair dominates the optimization problem, but there is potential for greater vibration isolation.

We are continuing this study with the optimization and testing of more complex seven bay structures, putting an even greater emphasis on the geometry repair part of the process.

Acknowledgements

This work has been supported by the Engineering and Physical Sciences Research Council: grant code GR/T19209/01.

References

[1] Forrester, A. I. J., Sóbester, A., and Keane, A. J., "Optimization with missing data," *Proc. R. Soc. A*, Vol. 462, No. 2067, 2006, pp. 935–945.

[2] Sóbester, A. and Keane, A. J., "A Supervised Learning Approach to Parametric CAD Geometry Repair," *AIAA Journal*, Vol. 44, No. 2, 2006, pp. 282–289.

[3] Keane, A. J., "Passive vibration control via unusual geometries: the application of genetic algorithm optimization to structural design," *Journal of Sound and Vibration*, Vol. 185, No. 3, 1995, pp. 441–453.

[4] Keane, A. J. and Bright, A. P., "Passive vibration control via unusual geometries: experiments on model aerospace structures," *Journal of Sound and Vibration*, Vol. 190, No. 4, 1996, pp. 713–719.

[5] Moshrefi-Torbati, M., Keane, A. J., Elliott, S. J., Brennan, M. J., and Rogers, E., "Passive vibration control of a satellite boom structure by geometric optimization using genetic algorithm," *Journal of Sound and Vibration*, Vol. 267, 2003, pp. 879–892.

[6] Shankar, K. and Keane, A. J., "Energy flow predictions in a structure of rigidly joined beams using receptance theory," *Journal of Sound and Vibration*, Vol. 185, 1995, pp. 867–890.

[7] Jones, D. R., Schlonlau, M., and Welch, W. J., "Efficient Global Optimisation of Expensive Black-Box Functions," *Journal of Global Optimization*, Vol. 13, 1998, pp. 455–492.

[8] Kelly, C. T., *Iterative Methods for Optimization*, SIAM, Philadelphia, 1998.

[9]Forth, S. A., "An Efficient Overloaded Implementation of Forward Mode Automatic Differentiation in MATLAB," *ACM Transactions on Mathematical Software (TOMS) – in press*, 2006.

[10]Sóbester, A., *Enhancements to Global Optimisation*, PhD thesis, University of Southampton, Southampton, UK, October 2003.

[11]Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T., "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, Vol. 6, No. 2, April 2002, pp. 182–197.

[12]Voutchkov, I. I., Keane, A. J., and Fox, R., "Robust structural design of a simplified jet engine model, using multiobjective optimization," *11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conf., AIAA 2006-7003, Portsmouth*, 2006.