Passive vibration control of a satellite boom structure by geometric optimization using genetic algorithm

M. Moshrefi-Torbati*, A.J. Keane, S.J. Elliott, M.J. Brennan, E. Rogers

School of Engineering Sciences (SES), University of Southampton, Southampton SO17 1BJ, UK
Institute of Sound and Vibration Research (ISVR), UK
Electronics and Computer Science Department, UK

Abstract

In this paper, the superior mid-frequency vibration isolation of a geometrically optimized lightweight structure is demonstrated. The initial structure under test here was a 4.5 m long satellite boom consisting of 10 identical bays with equilateral triangular cross-sections. An unusual geometric variant of this, with inherent isolation characteristics, has been designed by the use of genetic algorithm (GA) methods. In order to obtain the best design, the joints in the boom were allowed to move around by 20% of the length of each bay (i.e., ± 9 cm in all three translational directions). This work is based on results from a Fortran code (which was derived from receptance analysis) that are fully validated against detailed finite element (FE) models of the structure. The experimental forced response of the regular boom structure has been measured and compared with predicted curves. Finally, having obtained the geometrically optimized boom structure, its experimental response is compared with the theoretical results predicted by the receptance method. It is shown that the average of 30 dB isolation in the vibration energy transfer between the ends of the network of beams, over a 100 Hz bandwidth predicted in the design process, is achieved experimentally in an essentially undamped structure.

1. Introduction

A great number of modern engineering structures are fabricated from beam elements, often arranged in some kind of regular lattice type structure. When these structures are subjected to vibrations, it is useful to be able to prevent vibration energy from travelling along them. Perhaps

*Corresponding author. Tel.: + 44-23-8059-2350; fax: + 44-23-8059-3230.

E-mail addresses: m.m.torbati@soton.ac.uk (M. Moshrefi-Torbati), ajk@soton.ac.uk (A.J. Keane), sje@isvr.soton.ac.uk (S.J. Elliott), mjb@isvr.soton.ac.uk (M.J. Brennan), etar@soton.ac.uk (E. Rogers).

the most challenging vibration control issues arise in space missions that involve satellites with highly sensitive instrumentation packages. To function correctly, these packages must be supported on structures where the vibration levels have been reduced to extremely low levels (micro-vibrations). This need becomes more severe when the instrument concerned forms one of the individual sensors of a multi-sensor interferometric telescope or synthetic aperture radar [1]. In such cases, there is a need to support instruments spaced tens of meters apart using structural booms, with the relative motions between their ends being restricted to microns over wide ranges of excitation frequency [2]. A number of design approaches have been proposed to try to meet these demanding requirements but it is still not clear how best to proceed in this field [3]. The most common treatment for such problems is to use anti-vibration mountings or 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 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 is precisely the aim of the vibration isolators used between most pieces of equipment and their supporting structure. However, such isolators cannot deliver the desired behavior in all situations, particularly for sensitive equipment. The upshot of this problem is the need for some kind of widely applicable, generic structural filter design capability that can be used to build desirable characteristics into a structure, retaining its ability to carry static loads while blocking higher frequency motions. This paper reflects part of a research program currently being undertaken into the development of a medium frequency vibration control approach for such systems.

A great deal of work has been reported in the literature on controlling the low-frequency vibrational response of space truss structures. However, as such problems are brought under control other performance issues become of relevance. The program that this research forms part of is examining the higher frequency behaviour of space structures because the designers and manufacturers collaborating in the work consider these to be issues that will need tackling next. Although the structures considered here are geometrically unusual, they are not considered by the industrial manufacturing team to be overly difficult or expensive for use in space applications where the potential benefits outweigh any increased construction costs. This is particularly true given modern five-axis numerically controlled machining facilities where it is straightforward to manufacture the range and variety of structural elements involved.

The passive optimization that is employed here is based on energy flow analysis models combined with genetic algorithm (GA) optimization methods [4–6]. This approach relies on the wavelengths of the vibrations being controlled being of a similar order to the changes introduced in the geometry of the structure and hence is not most effective at very low frequencies. In this paper, the focus is on optimizing a full three-dimensional truss structure by using a parallel, multi-workstation GA. The structure consists of 93 individual beams and allows for bending, axial and torsional vibrations. An initial base-line structure of 4.5 m length with 10 identical bays (representative of a satellite boom) has been constructed. The response of this structure, when suspended in the laboratory, has been measured and the experimental results are presented. Subsequently, a GA designed variant of this was also built and tested to validate the approach proposed. This modified design produced significant vibration isolation. Although the investigation presented in this paper is based on

satellite structures, the overall results may be of interest in all fields where reduced vibration transmission is important, such as aircraft and car design.

2. The structure

Fig. 1 shows the geometry of the boom structure before optimization. All the 93 Euler–Bernoulli beams in the structure have the same basic properties: all the rods have a flexural rigidity EI of 2.184 MN, a bending rigidity of 5.503 N m^2 and mass per unit length of 0.085 kg/m. Based on initial trial experiments, the value of the structural viscous damping coefficient has been taken to be 0.005. The beams were all initially either 0.450 or 0.636 m long and they were joined together by 33 aluminium spheres of 25 mm diameter. The beams are produced in practice using a numerically controlled lathe and the spheres using a five-axis numerically controlled machining centre. The design for connecting the elements of the boom structure is based on the fact that at any stage of the assembly, there is only ever one element that may require exertion of bending force to fit in place. More specifically, each sphere contains six holes that accommodate the six rods coming from various angles. In constructing each bay, the three non-diagonal elements will easily slide into the corresponding holes and another two elements belonging to the following adjacent bay will also fit with ease. In order to place the sixth (i.e., a diagonal element), a hole is drilled all the way through the sphere. This design proved to be very efficient both in terms of accuracy and practicality. Moreover, it is capable of dealing with arbitrary geometries without any additional machining costs. The model was excited by a point transverse force on the fourth joint (0.45, 0.0, 0.0) m and during optimization, the aim was to minimize the vibrational energy level in the right-hand three joints (which in practice might be attached to an instrumentation package) between 150 and 250 Hz.
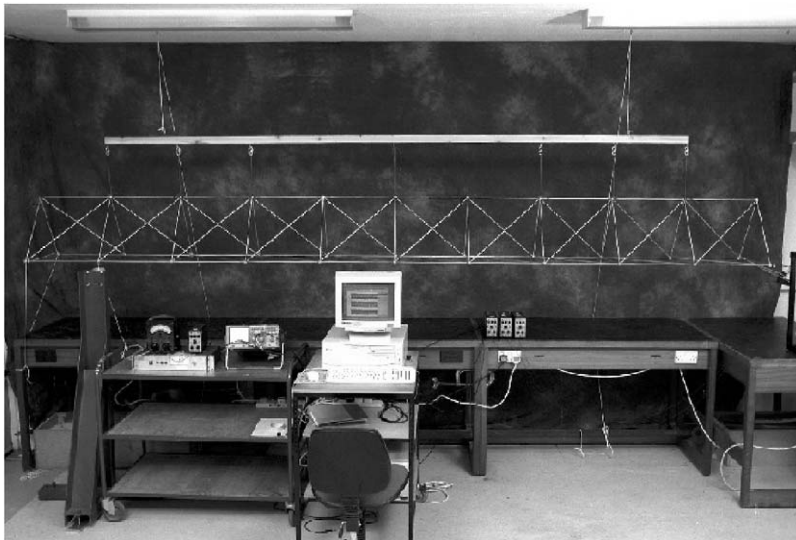

Fig. 1. Experimental set-up for the forced response of the regular boom structure.

3. Theory

The forced response of the above structure may be readily computed by using finite element methods but this is rather cumbersome when the force is broadband and especially when many different configurations need to be considered during optimization. Instead, the well-known characteristics of the Euler–Bernoulli beams together with receptance methods are used. Such methods have been studied for many years and may be set up for solving directly for the energetic quantities of interest [4,7]. Using this approach leads to the problem being defined in terms of the unknown displacements and forces at the joints. The displacements at the ends of each member meeting at a joint are equated and the joint coupling forces and moments summed to zero. The effects of the external driving and the joint forces on any beam are specified in terms of the Green functions of the elements in isolation. This approach needs some care to ensure convergence of the Green functions but allows the selection of various models for the elements to be made with great ease. The references show that energy flows calculated in this way agree very well with those based on finite elements. Furthermore, the calculation of a frequency band average is faster than can be achieved by using a finite element approach with the same level of accuracy.

4. The regular boom experiment

The experimental set-up for the boom structure is shown in Fig. 1 and schematically in Fig. 2. The structure is connected to a wooden panel via several elastic strings and the wooden panel is subsequently suspended from the ceiling by two equal length strings. Free–free boundary conditions were chosen here as they are simple to achieve in practice. A force transducer (B&K 8200) was used in conjunction with a shaker (LDS V201) that was screwed to the structure. The joint with co-ordinates (0.45, 0.0, 0.0) m was chosen as the input force location. The response was measured at the location of one of the end joints with co-ordinates (4.5, 0.3897, −0.225) m, as shown in Fig. 2. A random input force from 1 to 500 Hz, generated by the analyzer was supplied to the shaker via an amplifier (type TPA100-D). An ammeter was then used in series with the amplifier to avoid burning the shaker's coil. The force transducer's output was fed back to the analyzer's first acquisition channel via a charge amplifier (B&K 2635) whose gain was recorded for later data calibrations. The output of the B&K 4374 miniature accelerometer (attached to a sphere by a small dab of wax) was also fed back to the analyzer's second acquisition channel via a charge amplifier. In order to monitor the quality of input and output signals, they were both displayed on an oscilloscope before reaching the analyzer.

5. Results

Fig. 3 shows the experimental forced response of the structure at one of the joints (joint 32 and in the y direction) for the frequency band of 1–500 Hz compared with the same result obtained using the receptance theory [7]. To demonstrate the quality of the acquired experimental data, the recorded coherence function between the input and output data is shown in Fig. 4. From the figures, it may be seen that there is a poor coherence at anti-resonances due to very low output at
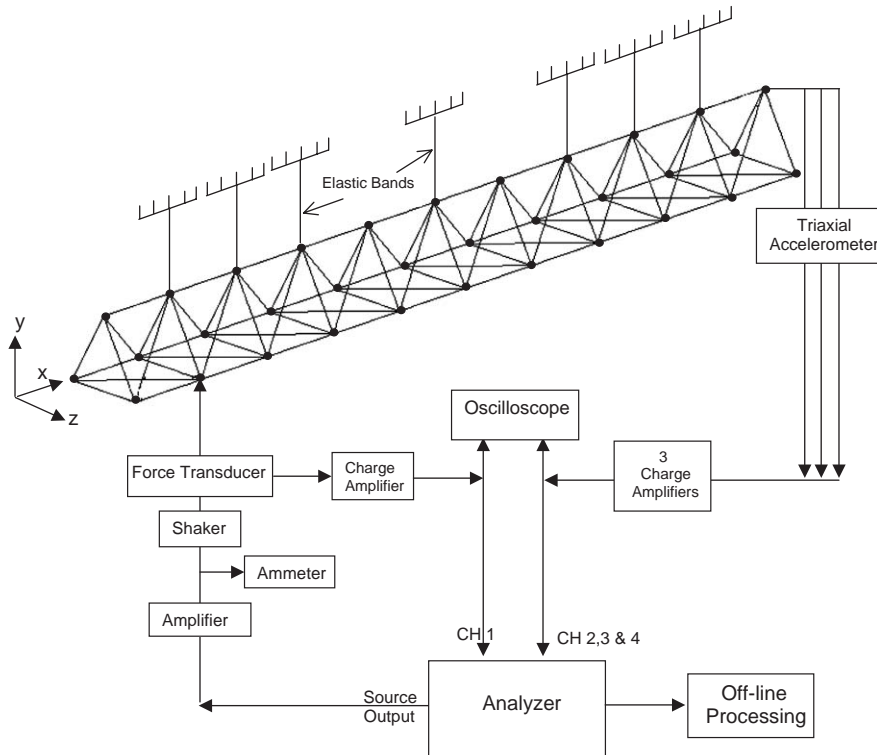
Fig. 2. Experimental set-up for the regular boom structure.

these frequencies. The big discrepancy at very low frequencies is mainly due to the fact that the miniature accelerometer does not start picking up the structural vibration accurately before around 30 Hz. This problem could be overcome by using a bigger accelerometer to measure very low frequencies. In order to maximize the quality of the data around the peaks and troughs within each frequency band, the number of frequency lines and the input gain were increased, respectively. As can clearly be seen from the results shown in Fig. 3 there is a convincing agreement between the experimental results and the theoretical curve. This agreement implies that the theoretical model is capable of predicting the vibrational behaviour of the structure. This gives the confidence to go ahead and search for a geometrically optimized structure with reduced vibration transmission levels.

6. Passive optimization

Having demonstrated the reliability of the receptance method, the code was then interfaced with an optimization package that has been developed by one of the authors [8]. The aim here was to produce a new boom geometry with an improved frequency response curve. More specifically, the aim was to reduce the frequency-averaged response of the boom in the range 150–250 Hz by
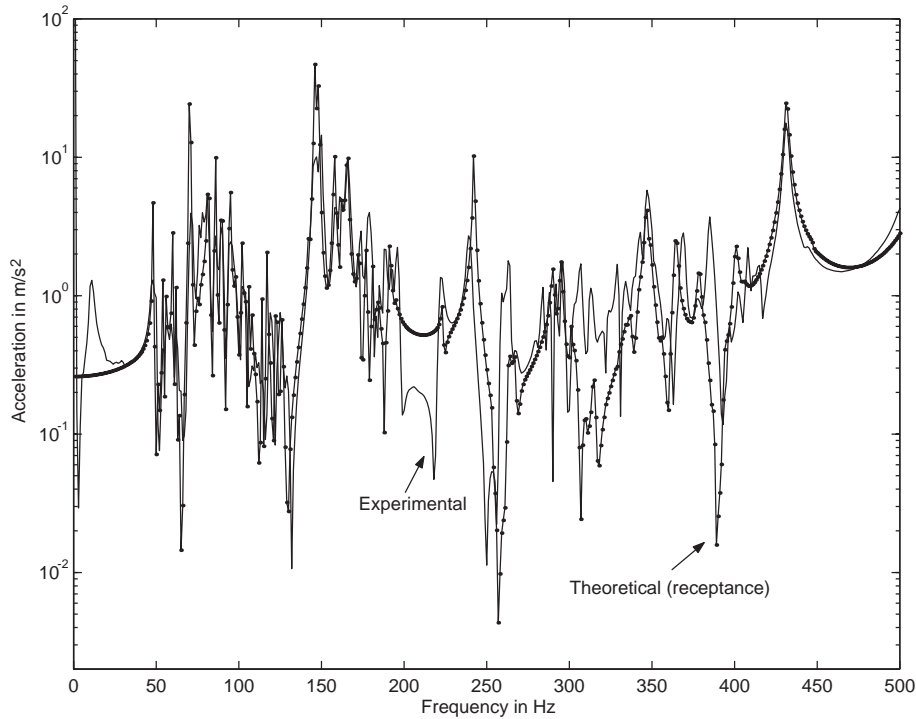
Fig. 3. Comparison of the experimental and theoretical responses of the boom structure at joint 32 and in the y direction.

the use of optimization, i.e.,

$$\text{Objective function } (J) = \int_{150\,\text{Hz}}^{250\,\text{Hz}} \sum_{j=31}^{33} \sqrt{v_{xj}^2 + v_{yj}^2 + v_{zj}^2}, \quad j = \text{joint}.$$

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 three end beams unchanged in length and position with respect to the fixed points. Further, to meet structural requirements, all of the joints within the structure have been 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, y and z co-ordinates of the 27 mid-span joints, i.e., 81 variables in all. To achieve the above objectives, a GA optimizer was selected from the optimization package [8] to produce the improved design. For this particular run of the GA, the number of generations was set as 30 and the population size as 300. This means that during the first generation, the optimizer generates 300 random new geometries. In the creation of a new geometry, the co-ordinates of all joints in the structure were varied within $\pm 20\%$ maximum deviation from their original positions. After analyzing the frequency response of each of these new boom designs, the GA outputs the design
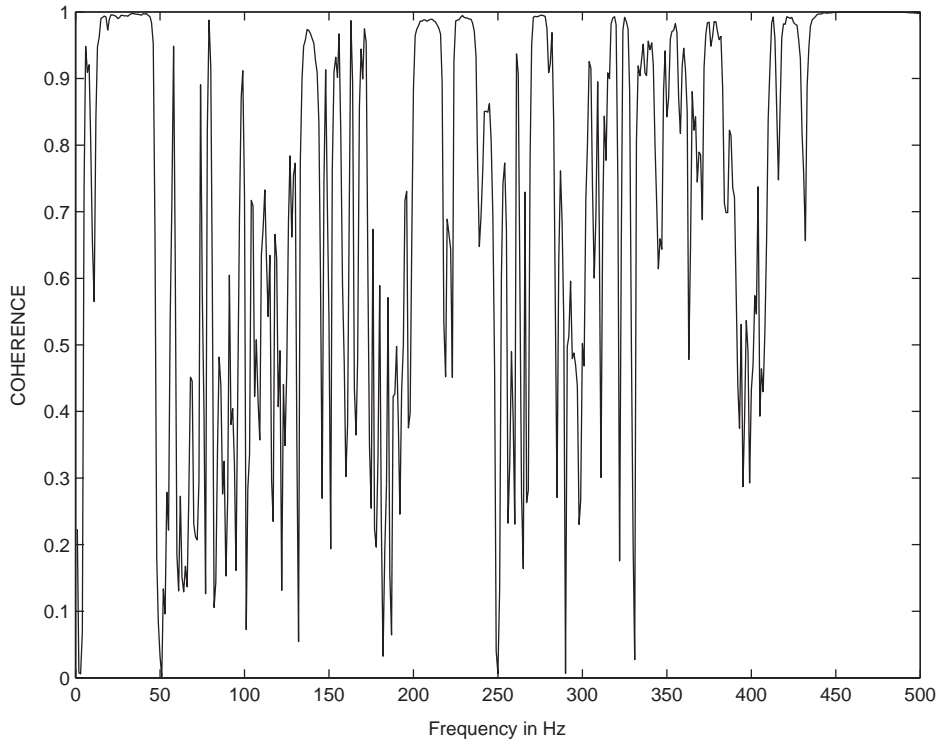
Fig. 4. Coherence function for the experimental response of the regular boom structure to joint 32 and in the y direction.

with the best performance and then continues on to the second generation. The designs carried over to the second generation and the creation of subsequent new designs are governed by the specific *natural selection* algorithm employed by the GA. The geometry producing the best frequency response curve from the 300 random geometries created after the 30th generation is illustrated in Fig. 5a.

As can be seen, the locations of all the 'interior' joints have been significantly changed (note that for clarity of the picture, all the diagonal elements are omitted). In order to confirm that the optimizer has indeed produced a geometry with an improved frequency response curve, the response of the new structure may be compared with the original regular boom structure. This is shown in Fig. 5b, where the frequency response of the optimized structure is plotted along with that produced using the regular geometry. From this figure, it is evident that the new design has greatly improved frequency response within the frequency range of 150–250 Hz. To present a value for the reduction in the vibration transmission, the frequency-averaged performance (the attenuation of the mean of the squares of the velocities) in decibels may be defined by α,

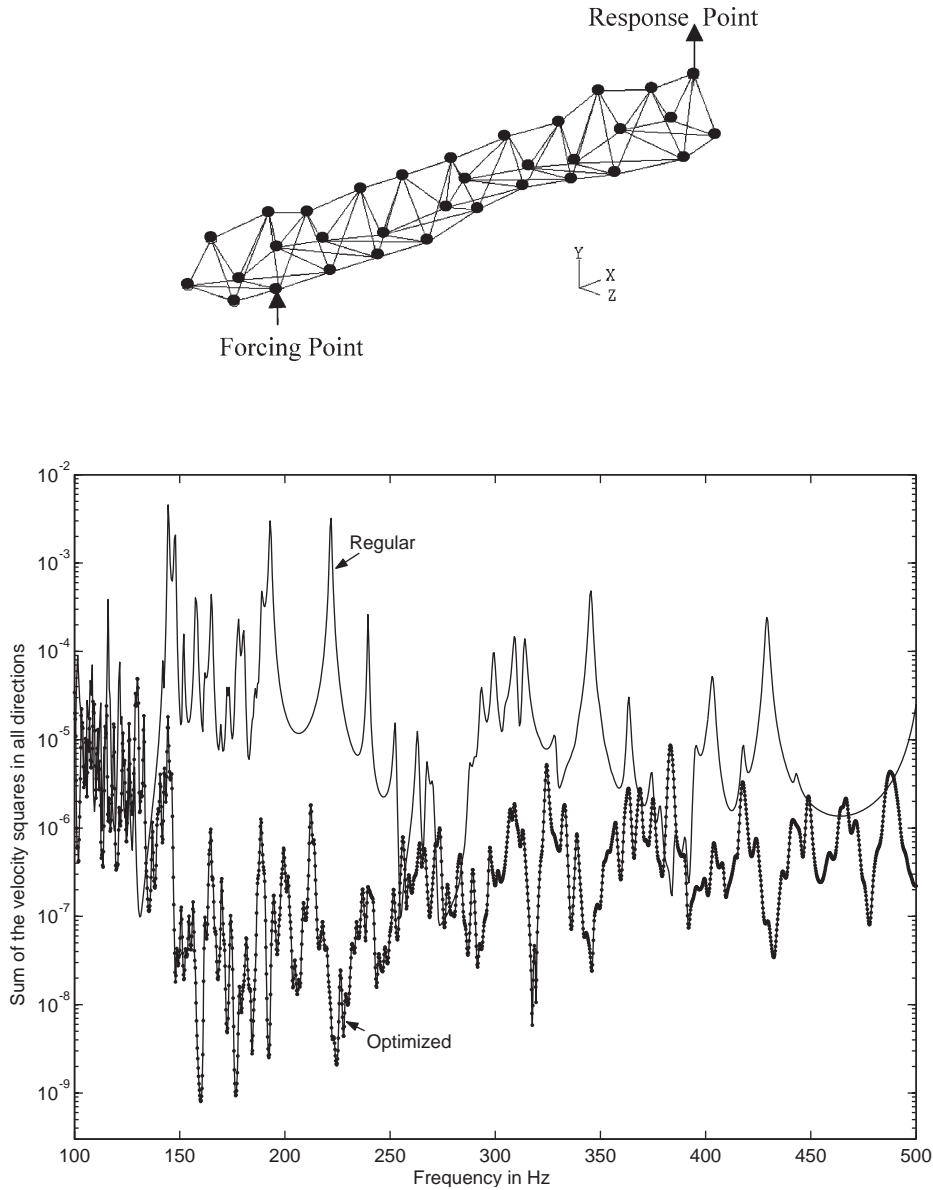$$\alpha = 10 \log_{10} \frac{\overline{J_{nom}}}{\overline{J_{opt}}},$$

Fig. 5. (a) Geometry of the best design after 30 generations. (b) Comparison of the theoretical frequency response of the regular boom structure (solid line) with that of the optimized structure (dotted line).

where the numerator and the denominator refer to the objective functions of the regular and optimized structures, respectively. The theoretical attenuation for the above-optimized structure within the frequency range of 150–250 Hz was found to be 30.95 dB. To verify these results, it was decided to manufacture the optimized design and repeat the forced response tests on it. The results are presented in the following section.

It should be noted that although an optimizer is used here to produce the improved designs, it is not suggested that these are optimal in any sense. In fact, without carrying out an exhaustive search it would not be possible to guarantee such an outcome. Given the number of real valued variables involved and the cost of analysis, this would not be feasible within realistic time scales. Such a situation is quite common in practical engineering design problems and the use of optimization in these circumstances is best termed "design search": the goal is merely to improve the design as much as possible within a given computing budget. A key part of the study presented here is to demonstrate the magnitudes of the improvements that might be expected from this essentially practical standpoint. It is of little practical relevance to know that a further 10 dB of isolation should be achievable according to theoretical analysis if only such a design could be found if the computing budget needed to find such a design is wildly unreasonable. Design teams need to know if design modifications will produce worthwhile benefits at affordable cost; happily this turns out to be the case for the designs considered here.

7. The optimizer

Design search optimization of the type posed here is characterized by having many variables with highly non-linear relationships between them and the objective function, which in turn has many peaks and troughs. Since, any one configuration is time consuming to evaluate, they are very difficult to deal with. The search for methods that can cope with such problems has led to the subject of evolutionary computation. Techniques in this area are characterized by a stochastic approach to the search for improved solutions, guided by some kind of evolutionary control strategy. There are three main methods that are currently in use: (1) simulated annealing [9], where the control strategy is based on an understanding of the kinetics of solidifying crystals; (2) genetic algorithm [10], where the methods of Darwinian evolution are applied to the selection of "fitter" design; (3) evolutionary programming [11], which is a more heuristic approach to the problem but which has an increasing number of adherents.

One of the authors [4] has applied all of these methods to structural problems and found that, for the current case, the GA work best. The GA used here is fairly typical of those discussed in the well-known book by Goldberg [10] but encompasses a number of new ideas that are particularly suited to engineering design problems [12,13]. Such methods work by maintaining a pool or population of competing designs which are combined to find *improved* solution. In their basic form, each number of the population is presented 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 on average better properties 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.

The most commonly used operations in GAs are currently: (1) selection according to fitness, i.e., the most promising designs are given a bigger share of the next generation; (2) crossover, where portions of two good designs, chosen at random, are used to form a new design, i.e., two parents "breed" an "offspring"; (3) inversion, whereby the genetic encoding of a design is modified so that subsequent crossover operations affect different aspects of the design; (4) mutation, where small but random changes are arbitrarily introduced into a design. In addition, the number of generations and

their size must be chosen, as must a method for dealing with constraints (usually by application of a penalty function). The algorithm used here works with 12-bit binary encoding. It uses an elitist survival strategy which ensures that the best of each generation always enters the next generation and has optional niche forming to prevent a few moderately successful designs dominating and so preventing wide ranging searches. Two penalty functions are available to deal with constraints. The main parameters used to control the method may be summarized as follows: N_{gen}, the number of generations allowed (default 10); N_{pop}, the population size or number of trials used per generation which is therefore inversely related to the number of generations given a fixed number of trials in total (default 250); P[best], the proportion of the population that survive to the next generation (default 0.8); P[cross], the proportion of the surviving population that are allowed to breed (default 0.8); P[invert], the proportion of this population that have their genetic material re-ordered (default 0.5); P[mutation], the proportion of the new generation's genetic material that is randomly changed (default 0.01); a proportionality flag, which selects whether the new generation is biased in favour of the most successful members of the previous generation or alternatively if all P[best] survivors are propagated equally (default TRUE); the penalty function choice.

When using the GA to explore large design spaces with many variables, it has also been found that the method must be prevented from being dominated by a few moderately good designs which prevent further innovation. A number of methods have been proposed to deal with this problem and the one used here is based on MacQueen's Adaptive KMEAN algorithm [14] which has recently been applied with some success to multi-peak problems [15]. This algorithm subdivides the population into clusters that have similar properties. The members of each cluster are then penalized according to how many members the cluster has and how far it lies from the cluster centre. It also, optionally, restricts the crossover process that forms the heart of the GA, so that large successful clusters mix solely with themselves. This aids convergence of the method, since radical new ideas are prevented from contaminating such sub-pools.

The version of the clustering algorithm used here is controlled by the following; D_{min}, minimum non-dimensional Euclidean distance between cluster centres, with clusters closer than this being collapsed (default 0.1); D_{max}, maximum non-dimensional Euclidean radius of a cluster, beyond which clusters subdivide (default 0.2); N_{clust}, the initial number of clusters before exclusive inbreeding within the cluster takes place (default 25); N_{breed}, the minimum number of members in a cluster before exclusive in breeding within the cluster takes place (default 5) and α is the penalizing index for cluster members which determines how severely members sharing an over crowded niche will suffer, with small numbers giving greater penalty (default 0.5), i.e., the objective functions of members of a cluster of m solutions are scaled by: $m^{min(\alpha,1)}[1 - (E/D_{max})^{\alpha}] + (E/D_{max})^{\alpha}$, where E is the Euclidean distance of the member from its cluster centre (which is always less than D_{max}; moreover, when $E = D_{max}$ no penalty is applied). In addition, the implementation of the GA used here allows the solution of individual members of the population to be run in parallel if a multiple processor computer is available.

8. The optimized boom experiment

Fig. 6 shows the manufactured passively optimized boom structure. The experimental set-up for this experiment was exactly as before, when the regular structure was tested. The structure was again
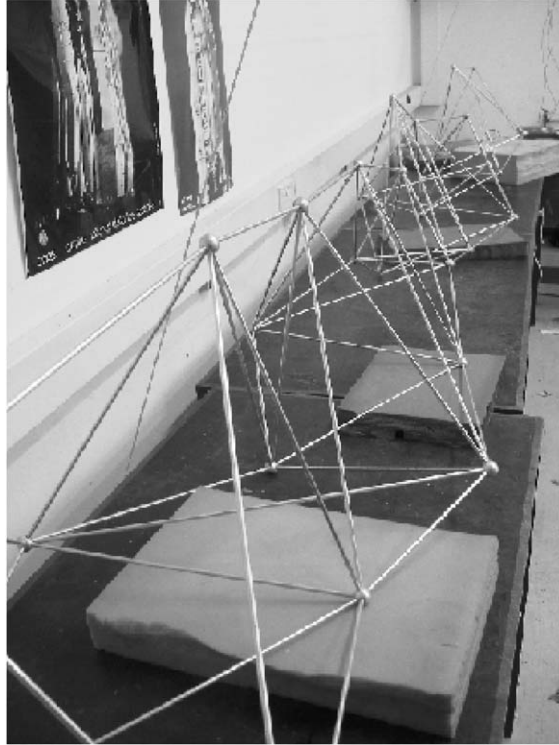
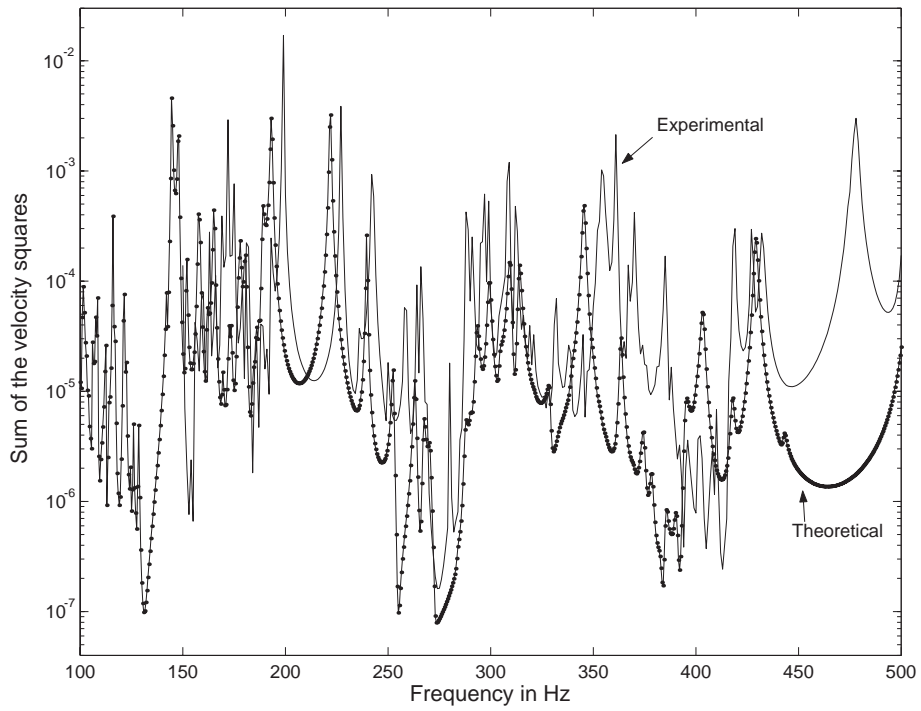Fig. 6. The optimized boom structure.

Fig. 7. Comparison of the experimental and theoretical objective function curves for the regular boom structures.
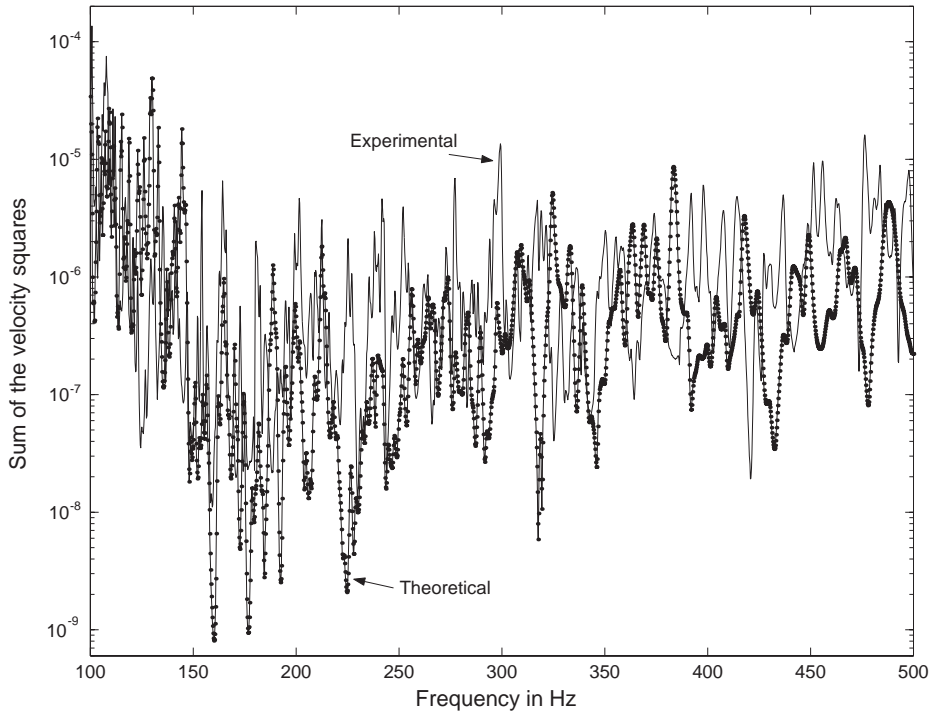
Fig. 8. Comparison of the experimental and theoretical objective function curves for the optimized boom structures.

excited by a unit random input force from 1 to 500 Hz, generated by the analyzer and supplied to the shaker via the amplifier. For the purpose of comparing the objective functions for both boom structures as the criteria for design improvement, the velocities at the three end joints with co-ordinates (4.5, 0.0, 0.0), (4.5, 0.3897, −0.225) and (4.5, 0.0, −0.45) were measured. This experiment was also repeated for the regular boom structure. For this test, the translational response of the structure in the x, y and z directions at these joints were measured with a Bruel & Kjaer triaxial accelerometer (type 4326 A). As this accelerometer weighs around 13 g, at any joint measurement, the other two end joints had masses of equal weights attached to them. The squares of all these velocities were then summed and the results compared with their equivalent theoretical curves (Figs. 7 and 8). Despite the fact that all the quantities are squared and added for nine different measurements, the differences between the curves are acceptable, particularly for the frequency range of 150–250 Hz that is of interest here. A point worth noting is that after the final optimized design was submitted to the workshop, it was realized that two of the rods on the 7th bay intersected each other and therefore the structure could not be assembled. To overcome this problem, the optimization run was repeated but this time starting from the optimized design and with the positions of all the joints kept fixed except the single joint that affected the position of both the problematic rods (i.e., only three optimization variables). After each generation, the position of the rods were monitored until there was sufficient gap between them and the boom in the above figure is the resulting design. This process had only minor impact on the predicted noise isolation.

Note also that the unusual irregular geometry of the optimized structure implies that at some joints, where the ends of six rods meet, the rods may partially block each other's path to the joint. To overcome
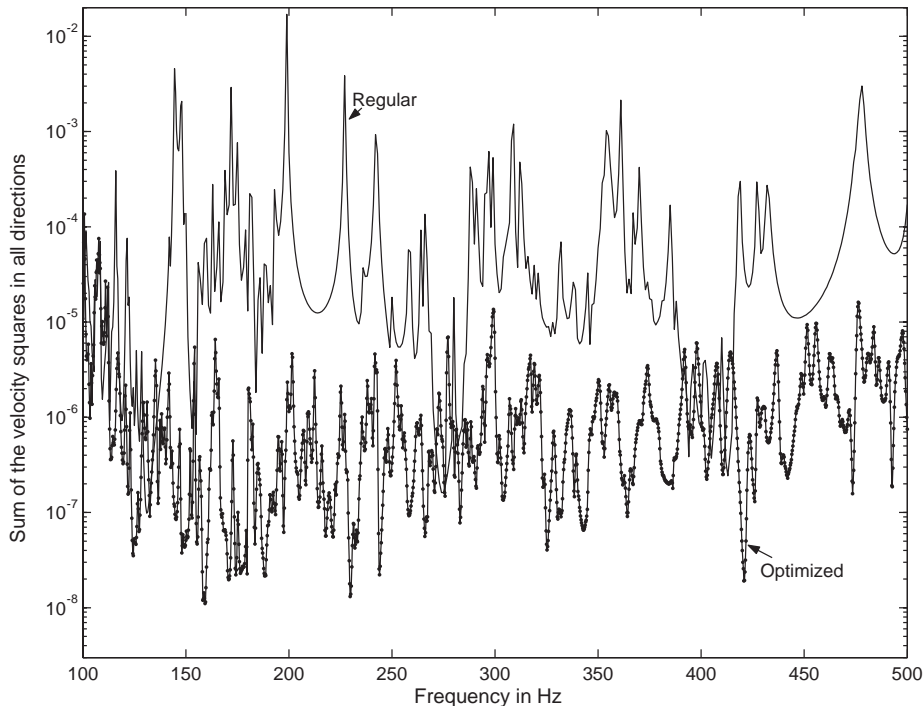
Fig. 9. Comparison of the experimental objective function curves for the regular and optimized boom structures.

this problem, it was decided to shave off the overlapping parts of the rods involved. This will have slightly changed the material properties of the elements involved and these are not accounted for in the theoretical model. The match between the results in Fig. 8 is very satisfactory, despite the inevitable errors arising from manufacturing and modelling inaccuracies. To demonstrate the improvement achieved in practice by the design, the experimental objective functions for the regular and optimized structures are shown in Fig. 9. It can clearly be seen that the optimized design displays much improved vibration characteristics at most frequencies and in particular within the band of interest.

9. Conclusions

This paper illustrates that significant improvements in noise isolation performance can be achieved in practice by modifying a regular structure in a controlled manner. The results of the experiments on a 4.5 m boom structure show that it is possible to build lightweight, lightly damped aluminium truss structures that have up to three decades, or 30 dB of energy transmission isolation between their ends without using additional damping materials or active control methods. Moreover, the structure produced is not unworkable and does not suffer from degraded performance at frequencies away from the band that has 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 geometrical discontinuities introduced. It

should also be noted that, even using 10 dual-processor workstations, the above optimization run took more than 10 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 worthwhile. Finally, it is worth mentioning that the optimized boom structure possesses less resistance to torsion and thus twists (and buckles) slightly more readily than its regular version.

Acknowledgements

References

[1] J.W. Melody, H.C. Briggs, Analysis of structural and optical interactions of the precision optical interferometer, Proceedings of the SPIE 1947 (1993) 44–57.

[2] S.W. Sirlin, R.A. Laskin, Sizing of active piezoelectric struts for vibration suppression on a space-based interferometer, Proceedings of the 1st US/Japan Conference on Adaptive Structures, 1990, pp. 47–63.

[3] J.W. Melody et al., Integrated modeling methodology validation using the micro-precision interferometer testbed, Proceedings of the IEEE CDC Conference, Kobe, 1996.

[4] A.J. Keane, Passive vibration control via unusual geometries: the application of genetic algorithm optimisation to structural design, Journal of Sound and Vibration 185 (3) (1995) 441–453.

[5] A.J. Keane, A.P. Bright, Passive vibration control via unusual geometries: experiments on model aerospace structures, Journal of Sound and Vibration 190 (4) (1996) 713–719.

[6] A.J. Keane, S.M. Brown, The design of a satellite boom with enhanced vibration performance using genetic algorithm techniques, in: I.C. Parmee (Ed.), Proceedings of the Conference on Adaptive Computing in Engineering Design and Control 96, PEDC, Plymouth, 1996, pp. 107–113, ISBN 0 905227 61 1.

[7] K. Shankar, A.J. Keane, Energy flow predictions in a structure of rigidly joined beams using receptance theory, Journal of Sound and Vibration 185 (1995) 867–890.

[8] A.J. Keane, 1994, The OPTIONS Design Exploration System User Guide and Reference Manual, http:// www.eng.ox.ac.uk/people/Andy.Keane/options.ps.

[9] S. Kipkpatrick, C.D. Gelatt, M.P. Vecchi, Optimization by simulated annealing, Science 220 (4598) (1983) 671–680.

[10] D.E. Goldberg, Genetic Algorithm in Search Optimization and Machine Learning, Addison-Wesley, Cambridge, MA, 1989.

[11] D.B. Fogel, Applying evolutionary programming to selected travelling salesman problems, Cybernetics and systems 24 (1) (1993) 27–36.

[12] A.J. Keane, Structural design for enhanced noise performance using Genetic Algorithm and other optimisation techniques, in: R.F. Albrecht, C.R. Reeves, N.C. Steele (Eds.), Proceedings of the International Conference on Artificial Neural Nets and Genetic Algorithms, Springer, Innsbruck, 1993, pp. 536–543.

[13] A.J. Keane, Genetic algorithm optimisation of multi-peak problems: studies in convergence and robustness, Artificial Intelligence in Engineering 9 (2) (1995) 75–83.

[14] M.R. Anderberg, Cluser Analysis for Applications, Academic Press, New York, 1975.

[15] X. Yin, N. Germay, A fast genetic algorithm with sharing scheme using cluster methods in multimodal function optimisation, in: R.F. Albrecht, C.R. Reeves, N.C. Steele (Eds.), Proceedings of the International Conference on Artificial Neural Nets and Genetic Algorithms, Springer, Innsbruck, 1993, pp. 450–457.