Optimising Continuous Microstructures: A Comparison of Gradient-Based and Stochastic Methods

Siavash Haroun Mahdavi
Department of Computer Science
University College London
Gower St. London WC1E 6BT, UK
mahdavi@cs.ucl.ac.uk

Sean Hanna
Bartlett School of Architecture
University College London
Gower St. London WC1E 6BT, UK
s.hanna@cs.ucl.ac.uk

***Abstract* – This work compares the use of a deterministic gradient based search with a stochastic genetic algorithm to optimise the geometry of a space frame structure. The goal is not necessarily to find a global optimum, but instead to derive a confident approximation of fitness to be used in a second optimisation of topology. The results show that although the genetic algorithm searches the space more broadly, and this space has several global optima, gradient descent achieves similar fitnesses with equal confidence. The gradient descent algorithm is advantageous however, as it is deterministic and results in a lower computational cost.**

I INTRODUCTION

Optimisation techniques are used by engineers to design structures to satisfy many criteria, such as high strength or low weight. Recent advances in computer controlled manufacturing technology have also allowed the construction of such structures to be automated, so that the machine plays a significant role in both design and building processes. The work in this paper investigates optimisation of a microstructure suited to a rapid prototyping technology known as stereolithography that is capable of construction at a high resolution, currently around 0.05mm. Our technique is based on the seamless repetition of a tiny structural module over a large volume such that the overall object behaves as a continuous material. It is, in effect, operating at a scale between traditional large-scale manufacturing and nanotechnology. The optimisation method is analysed with the specific requirements of this technique in mind, but involves generic structural principles that are shared with many other optimisation problems. As such the methods investigated in this paper can be applicable to other types of structures.

The types of structures investigated are known as space frames. Space frames are made of linear members that can be oriented in any direction in 3-dimensional space, and connected at node points either by rigid or flexible connections. To define a particular space frame one must specify both the members themselves, and the locations and orientations of the nodes in 3-dimensional space. We refer to these as the topology and geometry of the structure respectively, and it is these two properties that are considered by a structural optimisation algorithm.

A. The geometry

The distinction between geometry and topology can be described by an example 2-dimensional illustration. Geometry refers specifically to the positions in space of the node points joining the structural members. The following diagrams are of two structures with the same topology but different geometries. As can be observed, the connections and number of members are the same, but the coordinates and orientations of these members differ, see Fig. 1.

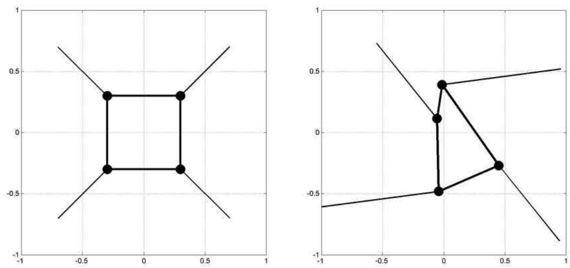

Fig 1. An illustration of a change in the geometry of a structure.

B. The topology

Topology refers to the structural connections between the node points. A change in the topology of a structure is a change in the number, or way in which the members are connected. This is illustrated in the two figures below, in which the two structures could not be made equivalent simply by moving the positions of the nodes, see Fig. 2.

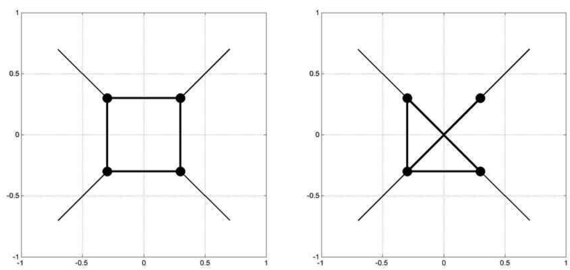

Fig. 2. An illustration of a change in the topology of a structure.

Geometry and topology are very different characteristics to be determined by an optimisation algorithm. Continuous, real-valued coordinates can specify geometry, whereas topology is determined by discrete connections between points. Optimisation procedures, including genetic algorithms (GAs) and also gradient-based methods such as gradient descent and simulated annealing, have been used for problems involving real valued parameters, but for discrete problems gradient-based methods are inappropriate. As in a previous paper we will assume a GA is used to evolve a discrete topology for a structure, and compare the effectiveness of gradient descent (continuous) and a GA (discrete) to optimise the geometry.

Gradient-based methods are guaranteed to be effective in simple problems with only one local optimum; however there may be several such optima in the search for a suitable geometry. For this reason GAs are often used for such problems. The overall optimisation of the structures proposed requires a dual process: to evaluate the fitness of a given *topology*, it is necessary to first find its optimal *geometry* for that topology. It is this second process that is the focus of this paper. This search procedure must be performed for every member of a population and runs as a second, inner loop within the larger GA, and therefore its efficiency is paramount. This double loop can be computationally expensive, and so if found to work well, the speed of a non-stochastic method such as gradient descent would be of great benefit. The results in this paper indicate that even though gradient descent can fall into local optima, the value of fitness is near enough to the global that the topology can be judged by the fitness of a geometry that might not be the global optimum.

II BACKGROUND

A. Genetic Algorithms in structural optimisation

The initial and simplest application of GAs in structural optimisation determined member widths in a set structure. Adeli and Cheng used a GA to optimise the weight of space trusses by determining the width of each member in a given structure. The shape and load points were fixed in advanced, and the cross sectional areas of groups of members were encoded in the genome, then selected to minimize the total weight [1].

Both geometry and topology have also been addressed in optimisation research. Yu-Ming Chen used a non-random iterative process of shifting node points in the FEM representation toward high stress zones, thus optimising the geometry [4]. Yang Jia Ping has developed a GA that determines both shape and topology, which begins with an acceptable unoptimised solution and refines the topology by removing connections [7]. Most recently, and more closely related to the methods in this paper, Peter von Buelow used a two stage algorithm nesting one GA within another. An outer GA evolved a topology for the structure expressed as a matrix representing the structural connections. Another GA found the geometry for each member of the population, expressed as real valued node positions [2].

B. Stereolithography: The Process

Stereolithography is a method of creating physical 3D realisations of CAD models; see [3] for a fuller explanation. It is one of the many types of machines collectively called ‘rapid prototyping machines’. As the name suggests, their primary usage is with the rapid building of prototypes for testing by engineers and designers. However as the technology has been dramatically improving over the past several years, it has become evident that this process can be used for more than building prototypes and can be itself a method for constructing parts.

The stereolithography machine consists of a tank filled with liquid photopolymer which is sensitive to ultraviolet light. An ultraviolet laser ‘paints’ one of the layers, exposing the liquid in the tank and hardening it, a platform then drops down into the tank a fraction of a millimetre and the laser paints the next layer. This process repeats until the model is complete.

Once completed, the object is rinsed with a solvent and then baked in an ultraviolet oven that thoroughly cures the plastic.

C. Stereolithography and Structural Optimisation

Previous work by the authors involved the utilization of an evolutionary algorithm to evolve the microstructure of an object created by a stereolithography machine [5].

This structure was optimised to withstand loads applied to it while at the same time minimizing overall weight. A two part algorithm was proposed that evolved the topology of the structure with a genetic algorithm, while calculating the details of the shape with a separate, deterministic, iterative process derived from standard principles of structural engineering. The division of the method into two separate processes allowed both flexibility to changed design parameters without the need for re-evolution, and scalability of the microstructure to manufacture objects of increasing size.

Ten thousand generations of the GA under equilibrium loading conditions resulted in the structure shown in Fig. 3. Results showed that a structure was evolved that was both light and stable. The overall shape of the evolved lattice resembled a honeycomb structure that also satisfied the restrictions imposed by the stereolithography machine [5].

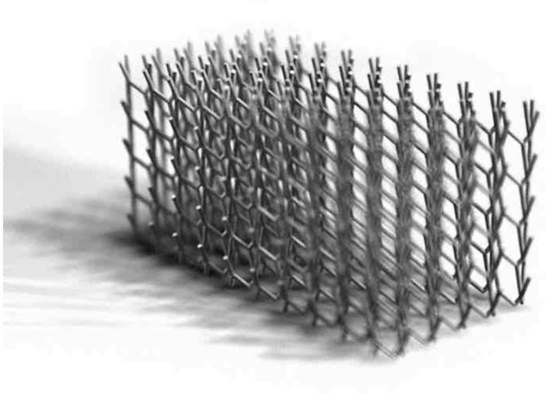

Fig. 3. The stereolithographic model created by the genetic algorithm.

A continuous structure was efficiently designed by arraying a series of smaller modules, referred to as 'unit cubes'. The structures examined in this paper therefore will be formed of repeated units of identical topology in which the structural members of each are smoothly connected to its neighbours. (The size of this unit cube is used as the unit of measurement of 3-dimensional space.) A space frame structure itself is defined by both the node positions (the geometry) and the connections between those nodes that form the structural members (the topology). In the case of the unit cube array, the connecting members may join two nodes within the same unit cube, or may join a node to another in any of the 26 neighbouring units. This topology was represented by a GA as described in [5].

D. Effect of member angles on strength

The stereolithography process causes the strength of material to be affected by the angle at which it is deposited, a fact that contributes to a more complex fitness landscape than would be the case with other manufacturing techniques. Because the model is built up of horizontal layers of resin, a linear member oriented perpendicular to these layers (i.e. a vertical member) has a greater strength than one at a shallower angle, and this strength decreases continuously down to approximately 30º from horizontal, below which the machine is unable to deposit material. The effect of these changes in strength can be seen in the two representations of fitness below, see [5] for more details.

The solution space to be searched consists of the node positions in 3-dimensional space for n nodes, or a 3n-dimensional space. In the samples formed by 6 node points we are therefore searching in 18 dimensions. For visualisation purposes, these plots show the fitness of a 2-dimensional slice through this 3n-dimensional space determined by two random orthogonal 3n-dimensional vectors and centred on an optimal solution found by gradient descent. While not an exhaustive mapping of the entire fitness landscape, they do indicate the added complexity introduced by adjusting the material strength due to angle. The first plot indicates the fitness of solutions in which all member angles are treated equally, (see Fig. 4). Adjusting strength due to angle however produces the characteristic valleys cutting through the fitness landscape of the second plot, corresponding to node point positions that produce weaker, more horizontal members, (see Fig. 5). Several local optima are clearly seen, and can be found in each of the 3n dimensions. The added complexities of this solution space indicate that the search for a global optimum would be more difficult when angle strengths are considered, particularly for gradient-based methods.

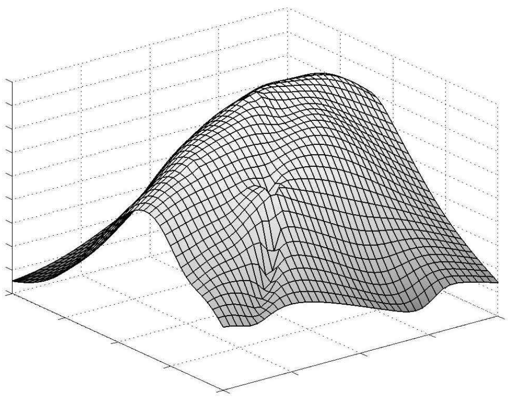

Fig. 4. A 2-dimensional slice through 3n-dimensional solution space.

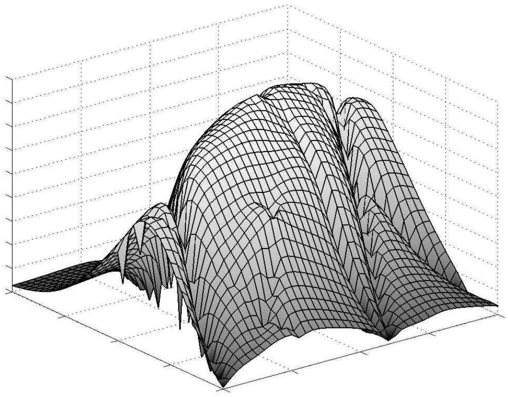

Fig. 5. A 2-dimensional slice through 3n-dimensional solution space. Member strengths are affected by angle.

III Method

A. Search procedures: fitness

The goal of optimisation in this case is to find the node positions for a given topology, which minimise the total stress in the structural under a specified load condition. The structure is defined as a set of linear members that can therefore be simulated by the finite element method [6]. The members are rigid jointed members with identical cross sections and a single unit of the structure is loaded under compression in the z-axis at its boundary nodes to calculate the stresses of each member. These stresses are then used as the fitness of the solution for each of the two search techniques.

B. The internal GA

A genetic algorithm can be used to determine the node point positions by first creating an initial population of random positions within the unit cube for the same number of points (n = 6 in our case). Crossover between two random parents then operates by comparing the corresponding node positions in each. A new point is drawn randomly from an elliptical normal distribution around the nodes of both parents such that the variance along the primary axis is the distance between the two points and the orthogonal variance is fixed at 0.01. Mutation is similarly accomplished by selecting a random point from a spherical normal distribution with a variance of 0.01 around the node to be mutated. These real valued operations are similar to those used by [2] for similar problems and as such serve as a useful comparison to gradient descent.

C. Gradient descent

The shape of a solution to be found is determined by the positions of six nodes along the three (x, y and z) axes, and therefore amounts to a search in 18-dimensional space. Gradient descent is used to find an optimal solution by sampling the fitness in each dimension at every iteration using the finite element method mentioned above. This provides a non-stochastic alternative search procedure to the GA.

D. Experiment

A randomly generated population (75) of topologies was created. These topologies were then optimised individually. For each topology gradient descent was performed five times from five random starting coordinates. Each run of the gradient descent was allowed to continue for 50 iterations as each topology was seen to reach a plateau (i.e. finding its local optimum) within this time.

For comparison, a genetic algorithm was also used to evolve the geometry given the same 75 topologies. It was allowed five runs for 300 generations. For our specific application the same experiment was performed again with the strengths of members due to their angles taken into consideration. (As in sec. II.D)

IV Results and Analysis

The GA was allowed to search for 300 generations with a population of 10, and gradient descent was run for 50 iterations. It was found that gradient descent reached an optimal solution (i.e. showed no further improvement) after an average of 718 fitness calculations (39.9×18 dimensions) when strength due to member angles is not considered, and 703 fitness calculations (39.1×18 dimensions) when it is. The GA found similar solutions after approximately 3000 fitness calculations (300×10 members in the population). The GA was therefore found to require an average of 4.2 times the number of fitness calculations compared to gradient descent for the examples in our set of topologies.

The two graphs below show examples taken from the same randomly generated topology. They represent a typical solution in that the optimal solution was found by the GA after 3000 fitness calculations, whereas gradient descent arrived at a similar solution with its optimum in approximately 630 fitness calculations. See Fig. 6 and 7.

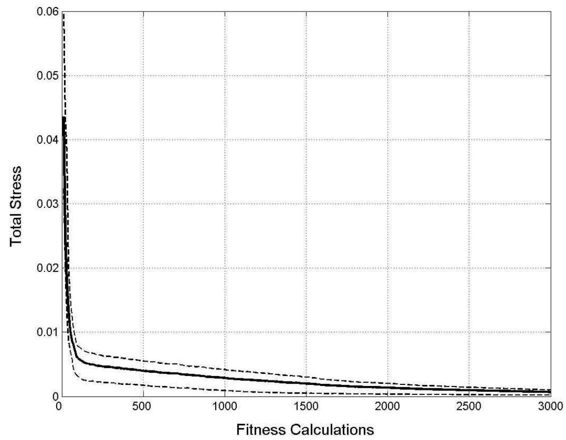

Fig. 6. A graph of the average and variance of the genetic algorithm for 300 generations.

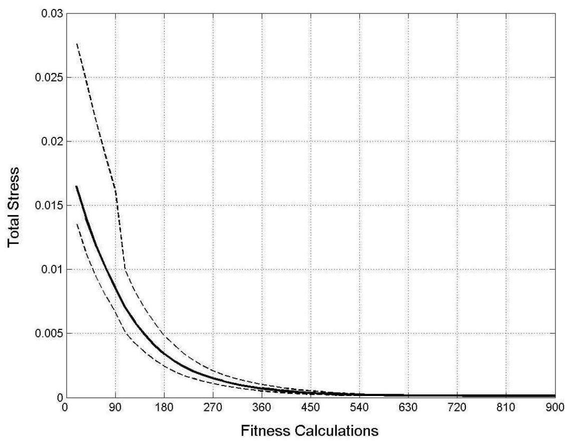

Fig. 7. A graph of the average and variance of the gradient descent for 50 iterations.

The variance in fitness over several runs of the algorithm can be taken as a measure of confidence in the accuracy of the result. A small variance in the final fitnesses would indicate a high level of confidence that these approximate the global optimum. These graphs show the mean and variance of total stresses over time for five runs of the GA and gradient descent respectively. On average it was seen that the variance in the solution decreases from 0.0351 in generation 1 to 0.0020 in generation 300 for the GA, and from 0.0288 in iteration 1 to 0.0015 in iteration 50 for gradient descent. This decrease in variance would indicate that both the GA and the gradient decent algorithm are not sensitive to initial starting conditions, but find similarly fit solutions each time.

A clearer comparison of the GA and gradient descent can be achieved by plotting the average of 5 runs for each method of each topology against each other. It can be seen by the graph below that the fitness of solutions found by either method are nearly equivalent for all topologies tested. If the final total stresses found by the genetic algorithm and gradient descent were equal, they would lie on the x = y plane. As can be observed in Fig 8 below, the 75 points lie very close to the x = y line. (Note that the log of both axes was taken to better illustrate the distribution). Another thing to note is that gradient descent actually slightly out performs the GA even though the GA was run for 4.2 times the number of fitness calculations.

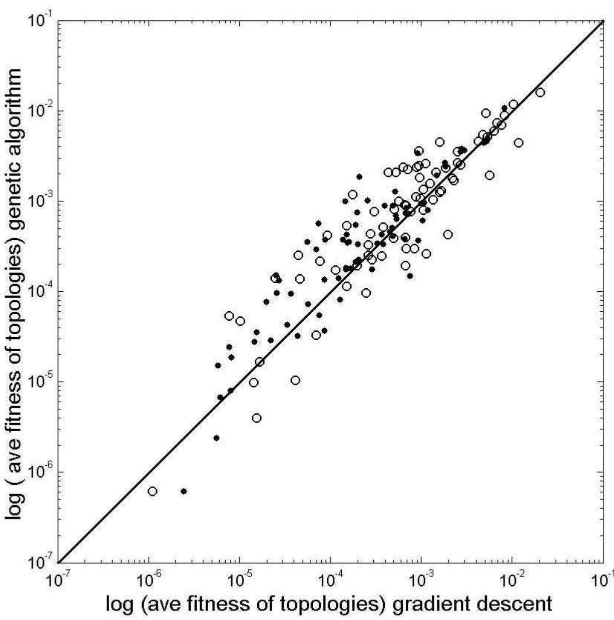

Fig. 8. A graph of the average fitness of the gradient descent vs. the average fitness of the GA for each topology. Open circles represent solutions with strengths adjusted due to angle.

Although the solution fitness was seen to be approximately equal for both the GA and gradient descent, it was found that the two methods behave differently in searching the space. The images below show the mean and variance of node point positions for five runs of the GA and gradient descent on the same topology. See Fig. 9 & 10.

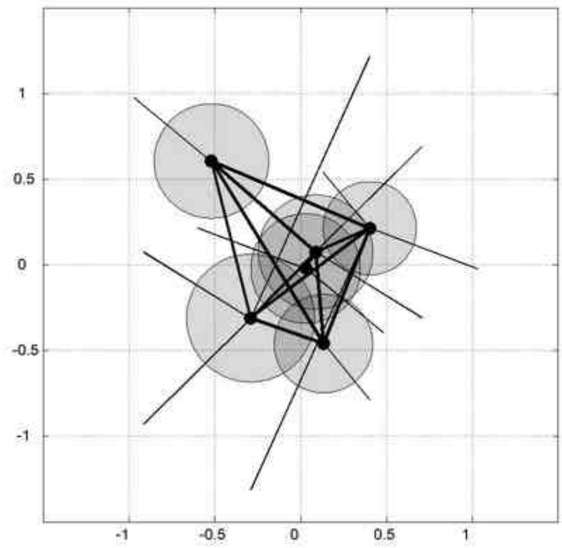

Fig. 9. Mean node points found by the GA. Variance is indicated by the grey circles.

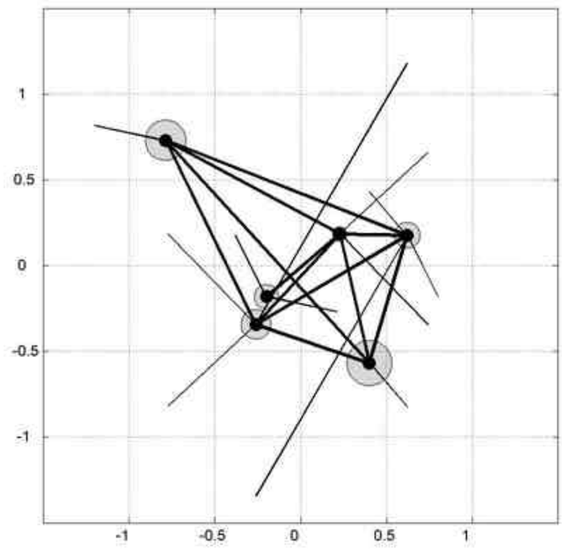

Fig. 10. Mean node points found by gradient descent. Variance is indicated by the grey circles.

For the entire sample set, the average variance for node points in solutions evolved by GA was 0.55 units, compared to an average of 0.37 units for gradient descent. As can be expected, the optimal solutions evolved by the genetic algorithm vary more than those found by gradient descent, indicating that it searched the space more broadly.

If the variance in final fitness were directly correlated with the finding of greatly diverging geometries, this would indicate that only structures with one single, clear optimal would be amenable to this evaluation. The fact that this is not the case can be seen in the following graph, a plot of the variance in point position vs. the variance in final fitness, see Fig. 11.

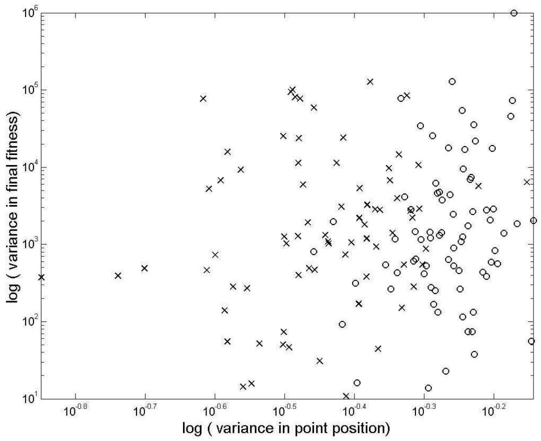

Fig 11. A graph of the variance in point position vs. variance in final fitness.

It can be observed that there is no correlation between variance in point position and variance in the final fitness. Notice also that the variance in point position generated by the genetic algorithm (depicted by o) is generally greater than that produced by the gradient descent (depicted by ´), confirming what was observed in figures 9 and 10. Little difference can be observed in the variance in final fitness.

V Conclusions

The aim of this work was to compare a deterministic gradient-based method with a stochastic genetic algorithm in optimising the geometry of a given topology. Each method yields a fitness in terms of the total member stresses under load, and the variance in these fitnesses over a number of runs has been taken as a measure of confidence in that fitness. The results have shown that a topology can be judged by the fitness of a geometry that might not be the global optimum.

The following conclusions are drawn from the observations:

1. The GA and gradient descent were able to find similar solutions with similar fitnesses.
   To achieve this however, the computational cost of the GA was 4.2 times that of gradient descent.

2. There was little observable difference between the two methods in the variance in the final fitness.
   This indicates an equal confidence in the results of both methods.

3. The geometries found by the GA revealed more variation in shape, indicating that it was performing a broader search of the solution space.
   Although this may be desirable for certain problems, it did not result in a noticeable difference in fitnesses and is therefore not of significant benefit in judging the topology.

4. There is no correlation between variance in fitness and variance in point position.
   This indicates that all topologies could be amenable to evaluation by these methods.

Considering the above conclusions, the choice of one method over another can be made on the basis of a 4.2 fold increase in speed, without any loss in accuracy or reliability. Therefore we conclude that for this problem domain, gradient descent is a better method with which to implement a geometry search.

VI Future Work

Both algorithms were stopped at the point at which they reached a plateau, however it was found that the initial iterations of gradient descent outperformed the GA even more rapidly than the 4.2 increase mentioned above. It is therefore hypothesised that by prematurely stopping the gradient descent at a point before it plateaus may dramatically decrease the cost of computation without affecting the reliability of the results.

Acknowledgements

We would like to thank Dr. David Fenner for his advice and expertise in the analysis of structure by the finite element method.

References

[1] Adeli, H and Cheng, N. (1993) “Integrated Genetic Algorithm for Optimisation of Space Structures,” Journal of Aerospace Engineering, Vol. 6, No. 4.

[2] Von Buelow, P. (2002) “Using Evolutionary Algorithms To Aid Designers of Archictural Structures” in Creative Evolutionary Systems, Bentley, P. J and Corne D. W. (Eds.) Morgan Kaufmann Pub.

[3] Brain, M. (2002) “How Stereolithography Works,” www.howstuffworks.com/stereolith.htm.

[4] Chen, Y. (2002) “Nodal Based Evolutionary Structural Optimisation Methods,” PhD Thesis, University of Southhampton.

[5] Haroun Mahdavi, S and Hanna, S. (2003) “An Evolutionary approach to microstructure optimisation of stereolithographic models.” Proceedings of CEC203. The Congress on Evolutionary Computation.

[6] McGuire, W., Gallagher, R. H. and Ziemian, R. D. (2000) “Matrix Structural Analysis,” 2nd edition.

[7] Ping, Y. (1996) “Development of Genetic Algorithm Based Approach for Structural Optimisation,” PhD. Thesis, Nanyang Technological University.