Modularity and Flexibility at the Small Scale: evolving continuous material variation with stereolithography

Sean Hanna Bartlett School of Architecture, University College London
Siavash Haroun Mahdavi Faculty of Computer Science, University College London

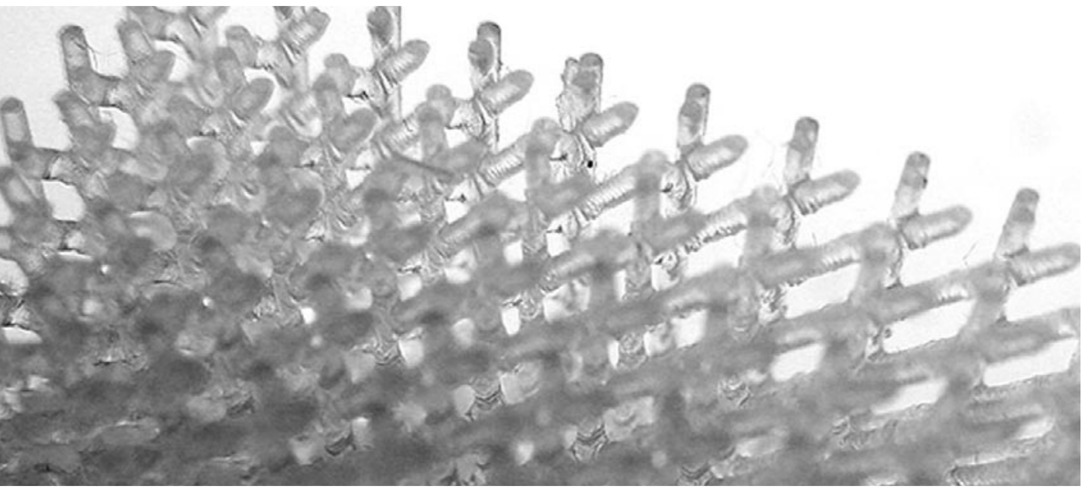

Abstract

In this paper, we introduce a technique by which the internal material properties of an object can be optimised at a microstructural level (5x10-5m) to counteract the forces that are applied to it. These can then be fabricated using the rapid prototyping method of stereolithography. The proposed technique is analogous to principles of mass customization and takes advantage of a flexible module to create complex structures in a manner that is computationally efficient and effective. The process is two-staged, in which a genetic algorithm evolves the topology of the microstructure and a second algorithm incorporating Finite Element Analysis then optimises the geometry. The examples shown are designed specifically for the fabrication technique, but the method and general principles are applicable to structural problems at any scale.

Introduction

Rapid prototyping methods are among the manufacturing techniques with which it is possible to perform mass customisation, fabrication methods unique in that every object may be different, the cost of complexity or change depending only on computation rather than the physical process. This introduces into the modularity inherent in most twentieth century production a flexibility previously found primarily in nature. This paper extends this effect to the smallest possible scale to engineer the material properties of an object. A method of design and fabrication is described that takes advantage of the variability inherent in rapid prototyping technologies to evolve a microstructure for complex objects that varies its properties at each point. This two-stage process incorporates a genetic algorithm and a deterministic search, and differs from many other optimisation methods that simulate the behaviour of the entire structure simultaneously and are therefore restricted to small or simple objects. By creating a single structural unit, a module that can be arrayed indefinitely throughout the fabricated item, the process is computationally tractable for objects of ever increasing size, but by making this module flexible in terms of its geometry it can be applied in a different form at every point. One module thus fills many roles.

All materials, treated as continuous, have complex internal structures that determine their properties: at the cellular level these give wood its strength, at the molecular level they differentiate diamond from graphite. The efficiency with which nature describes the necessary structural properties of a tree's cell in its chromosomes can serve here as a useful analogy, in that this cell is also a repeated module that changes its form at each iteration. All cells have an identical structure, but, depending on the conditions in which they grow, they change in size and shape, yielding familiar alternating rings with seasonal climate changes, the grain of wood, and the structural properties that have made it a superior building material throughout history.

This technique is based on the use of a flexible module of microstructure that is repeated many times over the volume of the object, taking advantage of the computer automated fabrication process to change shape at each iteration. The use of a repeated module provides benefits in terms of cost, taking advantage of the efficiency of mass production, and allows for theoretically limitless extrusion of the volume in any direction. It also facilitates the consideration of a complex design by focusing on the attributes of a single part. The benefits of making some properties of this module changeable include the ability to respond to local conditions, and greatly increased complexity in the kind of overall structure that it is possible to design and fabricate.

The implications on construction are twofold: at the small scale of material properties and at the larger scale of traditional structure. Altering the microstructure over the volume of a structural member would be similar to altering the cellular structure of wood to distribute material where it is needed, as though the tree itself were grown specifically for that structural task. Fabrication techniques such as fused deposition of metal, currently under development, suggest the possibility for this more organic approach to other materials. The continuous variation of properties would allow abrupt discontinuities at joints to be eliminated and a more seamless, efficient structure to result. In addition to material properties, these principles can be applied at the structural scale: although the method presented here is proposed for a specific, small-scale fabrication method (stereolithography), the same engineering principles hold at any size. As such, any automated manufacturing technique can be used, and the types of structures shown here can be scaled up to the size of a roof (see sec. 2.1) or entire building made of more traditional materials like welded steel or timber struts.

1. The microstructure

A series of linear structural members acting in either tension or compression traverse the volume of the object to be made and meet at node points, much like a 3D space frame. Rather than consisting of identical members at fixed angles from one another, their position and orientation vary continuously (Fig. 1). A genetic algorithm is used to determine a module's topology as defined by structural members and their connections to one another, and a deterministic process finds the ultimate shape. This underlying module is evolved to suit the material and machine: its topology is evolved in response to the specific fabrication process

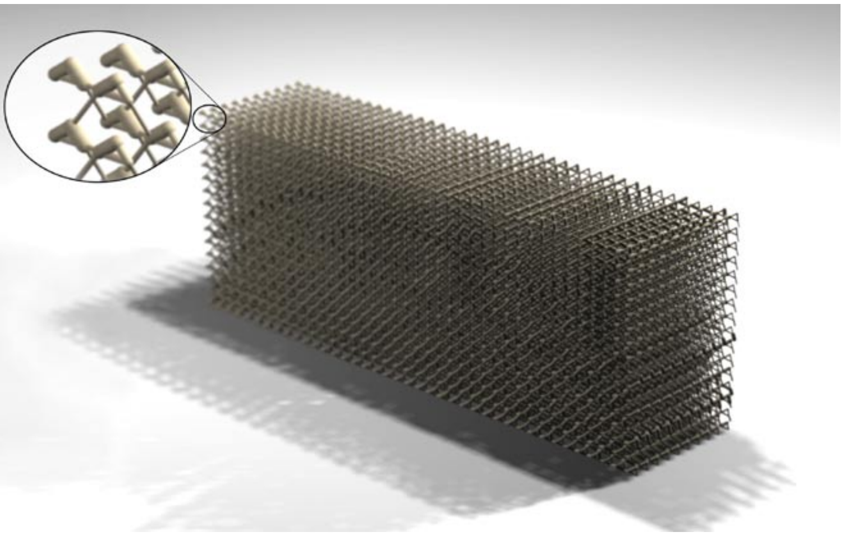

Figure 1. A microstructure space frame

of stereolithography and the range of forces, and its shape is modified to suit the object and specific stress, flexibility, or other required properties as they change from point to point. The module maintains a constant topology and as such connects seamlessly to its neighbours while changing shape. An item that would have been composed of many jointed parts by any other process may be produced as a seamless whole.

2. FEM Analysis

The design method is based on the underlying principles of the Finite Element Method of analysis (FEM), namely, that a continuous and complex surface or volume may be analysed by approximating it with a set of discrete elements arranged in a similar form. All of these elements may be regular beams, plates, or solids that behave according to simple equations, and as such their solution as individual, finite elements is straightforward.

The approximation of continuously variable material properties with a discrete module is to fabrication what the FEM is to analysis, in which a continuous object is represented by a mesh of separate structural elements whose behaviour can be predicted. This similarity is exploited in the technique. Our method uses the FEM as an analytical technique to model the input of a given structure and loading condition, but the same principles also guide the manufactured output. Just as the complex internal structure of any natural material can be assumed to be continuous and homogenous at the macroscopic scale, when specified by coefficients such as modulus of elasticity or Poisson's ratio (the amount of expansion in one axis under compression in another), we fabricate a complex microstructure equivalent to the desired continuous properties of a material.

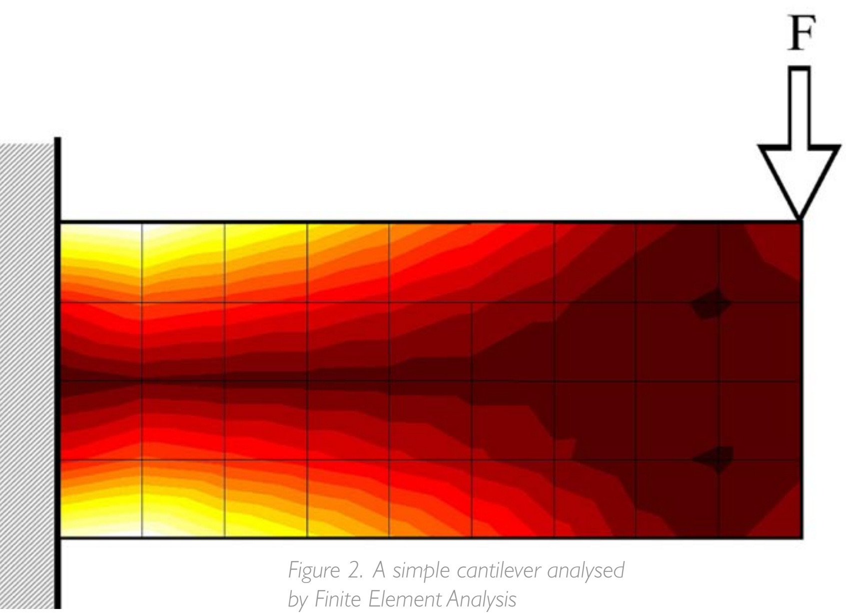

Figure 2. A simple cantilever analysed by Finite Element Analysis

In the design of an object of optimal strength and lightness, for example, one begins with the overall desired shape and a set of loading conditions. A mesh is generated from the CAD model and the resulting stresses at each point in the object are calculated using the Finite Element Method. These are then used to set the varying shape of each unit to accommodate the forces inherent at that particular point. More material is deposited where necessary, and less where it is not, resulting in a more efficient overall structure.

3. A sample structural problem

Our proposed method generates a microstructure for an object of any overall form that significantly reduces the deflection due to overall loads on the object. To demonstrate this, a microstructure will be evolved and shown to achieve the above objective for a simple example problem.

The figure below shows a simple cantilever beam fixed at one end, with a force (F) acting directly downward at the other end (Fig. 2). The stresses and strains within the object can be deduced by using Finite Element Analysis. As can be observed, the forces vary greatly throughout the structure. There is greater stress on the fixed side than there is on the side where the force is applied, mainly concentrated in horizontal tension along the top edge and horizontal compression along the bottom. At the small scale, using a simple homogeneous material, it would be necessary to accommodate the maximum stress everywhere by considering the forces exerted on the structure, and finding a suitable material to counteract these. This material would behave in the same way along the length of the object, even though the forces are clearly different. Scaling the cantilever to a size where a space frame might be used, convention would suggest the use of a space frame and topology that were the same throughout the length of the structure, and designed to withstand the maximum stress condition.

The method in this paper, by contrast, creates a microstructure in which every point on the structure is optimised to the forces acting on that point. The following section gives a brief review of the technologies used, and section 3 will take the reader step by step through the process of analysis, evolution, form generation, and fabrication of our proposed method. A completed structure will then be analysed structurally, and compared to a standard, uniform structure to demonstrate the effectiveness of the method.

Background

1. Mass customisation examples

The design and fabrication methods discussed are part of a wider trend, exemplified by the process of mass customisation in industries traditionally geared toward mass production, and by parametric approaches in architecture and design. Mass customisation is the mass production of products via a method that allows each to be individually customised (Pine 1993). It has been investigated and implemented in the apparel industry, for example, notably with Levi's brief but popular Personal Pair program that provided custom fitted jeans based on measurements taken in selected shops (Lee and Chen 1999). This took advantage of the vast production of the company to maintain a low cost, while providing a customised product for a minimal additional cost.

Architecture differs in that usually only one product, a building, is produced, although it is comprised of many individual parts, as is the case with the microstructures discussed in this paper. In architectural design, parametric approaches tend toward two sometimes overlapping goals. Either a single design is constructed hierarchically with some parameters left open to be changed at will in the future, or components of a complex design are designed to be repeated many times with possible variation. This paper focuses on the second. It is becoming ever more feasible to integrate the same design principle into the fabrication process at a large scale, to accurately manufacture elements with any desired variation.

The British Museum Great Court roof by Foster and Partners serves as one example, particularly by noting that with many other shell structures based on a regular radial or toroidal geometry the process of fabrication is facilitated by the repetition of many identical panels on the surface. In contrast, this triangulated steel grid shell that covers the courtyard of the existing building reconciles the contrasting rectangular and cylindrical geometry of the main building and reading room, and as such, the shape and division of its surface are also complex. The 4,878 steel members and 3,312 glass panels that comprise it are each different in dimension and shape (Sischka et al. 2004). To manufacture these unique elements, manufacturing contractor Waagner Biro used a modified automobile industry robot to cut the steel members to shape directly from the digital model. The elements are different in geometry, but all alike in detail, and the consideration of drainage, connections and material tolerances are handled similarly at every point in the roof, allowing these issues to be considered in detail by the designer once, and the solution replicated. There

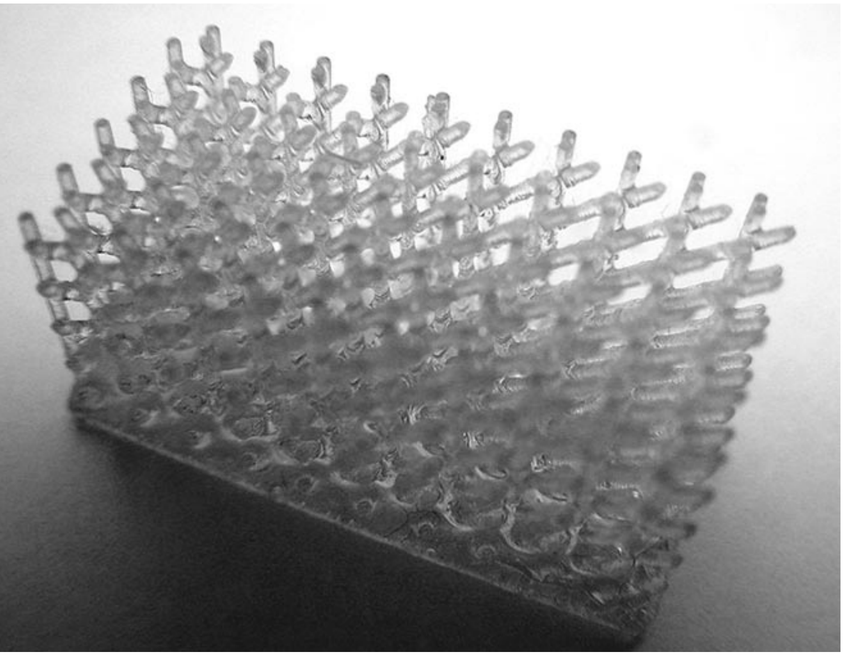

Figure 3. A sample microstructure fabricated by stereolithography

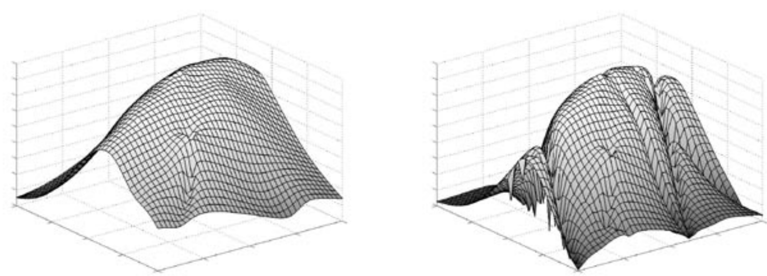

Figure 4. Structural strengths with and without material considerations

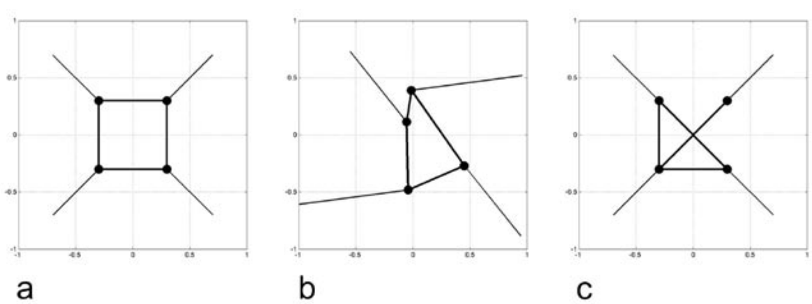

Figure 5. An illustration of a change in the geometry and topology of a structure

is one module in the roof, but its form is different in every instance. With such a manufacturing system in place, an increase in complexity of the design in terms of the variety of individual members does not incur an increased cost, and the manufacture of highly complex geometries is possible.

2. Rapid Prototyping

Stereolithography is a method of creating solid 3D models of CAD drawings (see Brain 2002 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 in the rapid building of prototypes for testing by engineers and designers. However, as the technology has improved dramatically in recent years, it has become evident that this process can be used for more than building prototypes and can itself be a method for constructing parts.

The stereolithography machine consists of a tank filled with liquid photopolymer that 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 (Fig. 3). Once completed, the object is rinsed with a solvent and then baked in an ultraviolet oven that thoroughly cures the plastic.

In construction, the structure is formed in a liquid stereolithography resin as a series of horizontal layers. This results in an inherent horizontal *'grain'* in every part of the model and a limitation on structural possibilities in that horizontal struts cannot be fabricated. Members constructed at differing angles to this *'grain'* therefore have differing strengths and their calculated deflections were modified accordingly in assessing the fitness of the solution. This material consideration significantly affects the complexity of the behaviour of a structure (Fig. 4). These figures are two-dimensional slices through a 3n dimensional space (where n is the number of node points in the structure), showing the relative strengths of structures before and after this material property is considered. The *'valleys'* seen in the figure at right indicate node positions for which the corresponding struts approach horizontal, and thereby weaken the structure considerably. When this property of the material is considered it is clear that these geometries should be avoided.

3. Manufacturing with stereolithography

Stereolithography and other rapid prototyping techniques are now beginning to be investigated as an alternative method of construction for objects of high complexity, particularly with intricate internal structures. This has not yet become commercially viable for mass production, but several researchers are preparing for the increasing accuracy and decreasing cost of the technology in the future. Molecular Geodesisics, Inc., for example, is investigating structures based on a regular tensegrity space frame which would, at a microscopic size, be useful as biological or industrial filters (Molecular Geodesics 1999).

4. Geometry and topology of space frames

The types of structures proposed are known as space frames: a set of linear members 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.

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: Figures 5 a & b show 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 the members differ. 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, as in Figures 5 a & c. The structural modules proposed in this paper share a common topology throughout the object to be fabricated, but change in geometry in each instance.

5. The Genetic Algorithm

The topology of a microstructure is evolved using a genetic algorithm (GA), while the geometry of the microstructures are calculated by using Finite Element Analysis. A GA functions by encoding a set of design parameters into a genetic code, by which various solutions can be defined, and evaluating those solutions by a fitness function that indicates their quality. Populations are created at random, and as in natural evolution the fittest members survive to contribute genetic material, which is recombined to form subsequent generations. Such algorithms are very good at finding solutions to problems that have bumpy or discontinuous search spaces, and so are ideal for searching through a population of topologies, which are defined discretely. The details of the GA ensure that the solutions evolved are as unrestricted as possible while also allowing a simplification of structure (please see Haroun Mahdavi and Hanna 2003 for further details).

Method

The proposed method can be split up into various steps, summarized as follows:

1. The problem must first be analysed. The object is described within a simulated physical environment and has forces applied to it, simulating the real world problem. The problem in this paper is that of a rectangular cantilever, loaded vertically at its extremity.

↓

2. The object is then split up into unit-cubes of a suitable size representing the locations of the structural modules. This allows the method to be used for an object of any complexity and size. Each of these units is analysed under the overall loading condition to determine its local stress.

↓

3. A topology is evolved via a genetic algorithm to accommodate the range of forces present in the problem.

↓

4. For each unit-cube in the volume of the object, the topology is then repeated with its geometry modified based on the stress at each point.

↓

5. The final structure is converted into an SLC file and sent to the stereolithography machine for manufacture.

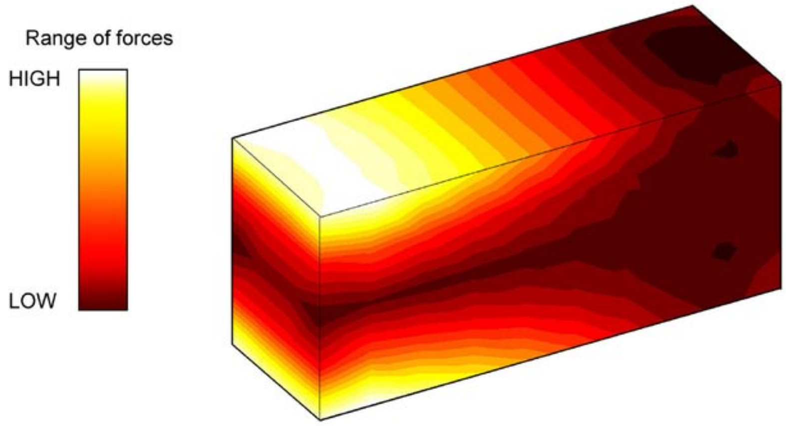

Figure 6. The cantilever as analysed by Finite Element Analysis

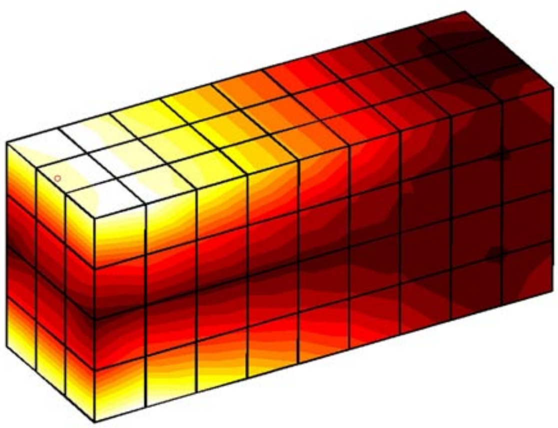

Figure 7. The cantilever beam is split up into unit-cubes

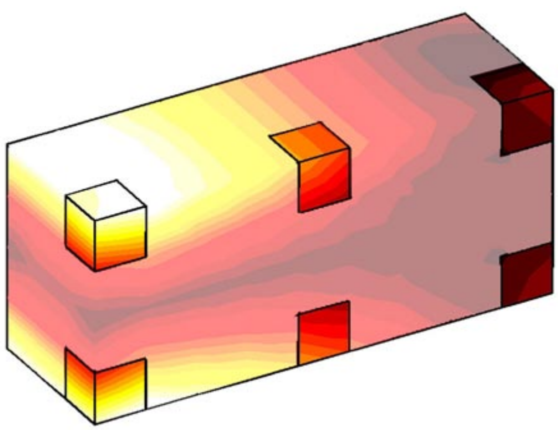

Figure 8. The stress condition within four unit-cubes is extracted

Computation time is roughly proportional to the size of the object measured in units: step 3 is an iterative process that finds better topologies the longer it is run, and step 4 is repeated for each unit-cube in the object. Each step will now be discussed in detail.

1. Analyse the problem

The method proposed in this paper will use the simple cantilever example in the introduction as a proof of concept. An object is thus described within a physical simulator with one end fixed and the other end deflected by a force. The physical simulator can then analyse the resulting forces within the object by using Finite Element Analysis, the result of which is shown in the following figure (Fig. 6). The range in colours from dark to white represent the range of stresses present within the object.

2. Divide the object into unit-cubes

The volume of the cantilever is divided into a three-dimensional grid of cubes, each forming a modular unit that will be referred to as *unit-cubes*. As the work presented here is concerned with relative performance rather than specific real world units of force and size, no real world units are used throughout. This unit-cube defines the basic unit of measurement used in the algorithm and has a volume of one. The actual size of these can vary, but is in the range of one to several millimetres.

In the loaded object the stresses at each of these unit-cubes are calculated using Finite Element Analysis for each of the x, y, and z-axes. These can be seen to vary gradually from unit to adjacent unit. A topology will be evolved that can optimise its geometry for the stresses incurred at each particular point, and its adjacent units will contain a very slightly different structure for a slightly different stress condition. In the object as a whole the arrangement of each cube varies considerably, but does so gradually and continuously over its volume in response to the external forces (see Fig. 7).

The placement of the node points in 3D space varies from point to point, but their connections to one another do not. This accomplishes two things that we find necessary for scalability. First, it avoids breaks in the continuity of the structure by allowing a difference between one unit-cube of structure and the next, which can be scaled to arbitrary units of precision. These break points would be literal break points if the topologies were to change, as adjacent units of incompatible structures would be zones of weakness in the overall object. Second, it allows one single graph of connections to be evolved which can be applicable to all points. This allows an object of any size and any number of units to be evolved once, rather than running the GA many times to generate structures for large objects or complex load conditions.

3. Evolve a suitable topology for a representative sample of stresses

Though the example used in this paper is relatively simple, the same method holds for a structure of any complexity and size. Therefore, after the cantilever beam is analysed using Finite Element Analysis, and split into unit-cubes, the force conditions for the n unit-cubes with the greatest variance are used as force conditions in the fitness function of the GA. The value of n (six in our example; see Fig. 8) should adequately represent the range of forces present within the whole structure.

The GA randomly generates a population of topologies, and exposes each to the n force conditions that were extracted from the cantilever beam. The sum of the average displacements of the nodes of each topology for all n force conditions then forms the fitness of the topology, the objective being to find structures that deform less under the set load. By repeating this process for each topology, the whole population can be evaluated. After many generations, the GA evolves a topology that can best counteract the range of force conditions.

Once a topology has been evolved, it is then repeated throughout the cantilever beam, with variation due to local conditions. For comparison, the examples in the next section will examine both an evolved topology and a standard engineered solution.

4. Repeat the module with variation

The Finite Element Analysis of the structure yielded a set of strains in the x, y, and z-axes for each unit-cube of the simulated continuous volume. These are now used to derive the set of forces that would cause the same local strains when applied to that cube in isolation. The process of finding a geometry for each of the unit-cubes is a deterministic, iterative algorithm that finds an optimal set of node positions and strut diameters to withstand the applied forces for that cube. One of several alternatives may be used, including gradient descent or the iterated movement of nodes to positions satisfying equations of equilibrium (please see Haroun Mahdavi and Hanna 2004 for further details). Because the changes in strain are gradual across the continuous simulation this geometry changes gradually with small changes of the applied forces from unit-cube to adjacent unit-cube.

5. Send to the stereolithography machine

While models constructed by a stereolithography machine can be complex on a macroscopic scale, they are generally solid internally, and so result in relatively small file sizes. The structures evolved by this method are generally very complicated and can result in very large file size if converted via .STL format to .SLC format in the usual way. The software usually used to send an initial CAD model to the stereolithography machine typically involved a computationally intense process of generating support structures beneath a model, and then calculating intersections of the solid geometry with a steadily rising horizontal plane to generate progressive ‘slices’ for fabrication. These processes are unnecessary and costly in the case of these structures, so the algorithm converts the structure directly into .SLC file format for fabrication by the machine.

This direct translation takes advantage of the regularity of the unit-cubes to send data to the machine efficiently, as the units themselves are aligned to the horizontal slice plane. This method could also be used to stream data to the machine, which would prove very useful when working with very complicated or large objects by generating the internal structure online during fabrication. Figure 9 shows how one unit-cube would be sliced up for SLC formatting. The machine instruction consists of a series of 2-dimensional outlines of the geometry that determine the edges of the slice to be filled by solid material. These are easily calculated by centering a standard polygon on the intersection of each member with the slice plane, scaling it uniformly by the diameter of the member, and axially by the tangent of the angle.

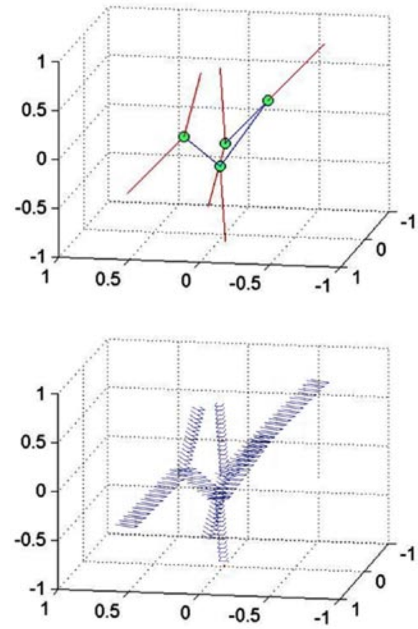

Figure 9. The structure of each unit-cube is sliced to .slc format.

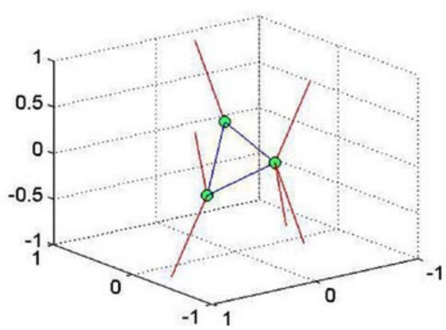

Figure 10. The evolved topology

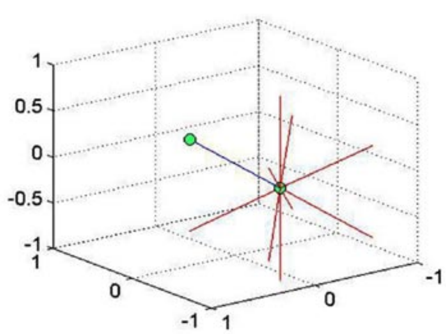

Figure 11. A standard topology

Results and analysis

The performance of the structures formed by the microstructural units is evaluated in this section. For sake of comparison, the overall dimensions of the structure, the total mass of material used, and loading conditions are maintained in each example. The units of measurement for deflection of the structures are given in unit-cubes.

The GA described in section 3.3 has been used to evolve a topology for the cantilever problem example. While this topology can be applied to situations other than the cantilever example, there might be some loads for which the resulting strut angles would be less efficient at transmitting the load. Therefore, as the results will confirm, the best approach is to use a topology suited to the problem. The following diagram shows the topology of this structure best able to accommodate the range of forces present within the cantilever beam, with a geometry suggested by equilibrium loading conditions (Fig. 10). This evolved topology will be compared to a more regular topology designed manually to be efficient, light, and symmetrical (Fig. 11). Both topologies easily satisfy the restriction of the stereolithography machine in that there are no horizontal struts.

Both topologies were replicated along the length of the cantilever beam and the composite structure analysed by Finite Element Analysis to evaluate performance under the loading conditions. First the topologies were replicated without any variation in the geometry of the units or diameter of struts. As can be observed in the following diagrams (Fig. 12 & 13), the uniform repetition of both the engineered and the evolved topology perform relatively poorly under the cantilever loading conditions.

The maximum deflections of the structure for the engineered and evolved topologies are 2.052 and 3.623 units respectively. The evolved topology was designed to be replicated throughout the volume of the beam with its geometry varied and struts thickened in accordance with the force conditions present in each unit-cube. Without this flexibility there was 76.5% more deflection at the maximum point, and the asymmetry of the unit resulted in a twisting motion throughout the structure.

Next, the geometries of both the engineered and evolved topologies were varied to counteract the forces present within each individual unit-cube. The following diagrams (Fig. 14 & 15) show this.

The total deflections for the engineered and evolved topologies were 0.303 and 0.053 respectively, indicating a significant improvement in strength of the beam due to local optimization of each module in both cases. Furthermore, the evolved topology that was optimised specifically for the force conditions now deforms only 17% as much as the regular structure.

A closer look indicates that the geometry of the more regular engineered module only changes in the thickness of the struts, whereas the evolved topology accommodates the variation in stresses also by changing the length and angle of struts throughout the structure. This greater flexibility of the small-scale module increases the overall performance of the entire object.

Conclusions

This paper presented a method that treated a complex microstructure as though it were a continuous material with properties that could be changed to benefit the design. We have demonstrated that materials can be fabricated, using a technique such as stereolithography, which are optimised at every point to the forces present there. This method of fabricating complex structures through the use of a flexible module has several benefits in terms of efficiency that are based on the saving in computation time and inherent to the production process itself. These are:

- The use of a module allows repetition indefinitely, limited only by the physical limitations of fabrication.
- The flexibility of the module allows a high degree of complexity of structure.
- Scalability to larger and more complex structures is easily accomplished. The speed of computation of the .slc files increases at the same rate as the scale of the object to be built and the fabrication time. Evolving a topology with the GA is performed only once at the outset, so the overall computational efficiency actually increases with larger scale structures.

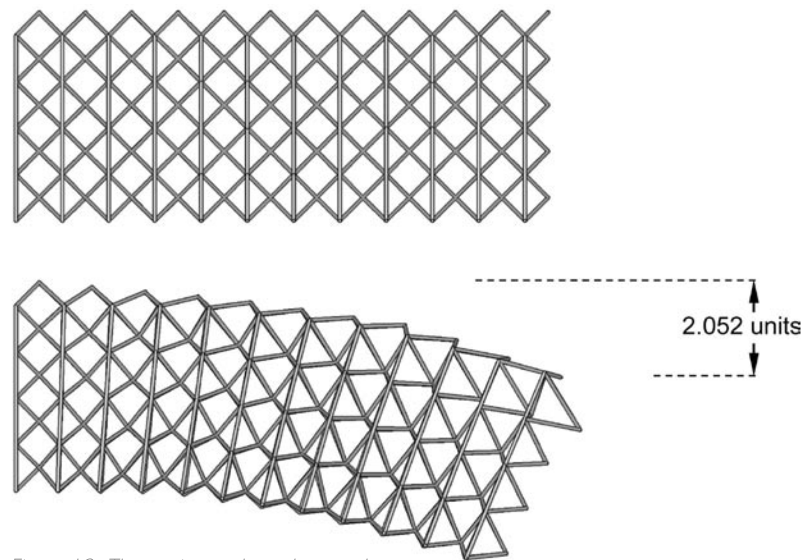

Figure 12. The engineered topology, and the resulting deformation under loading

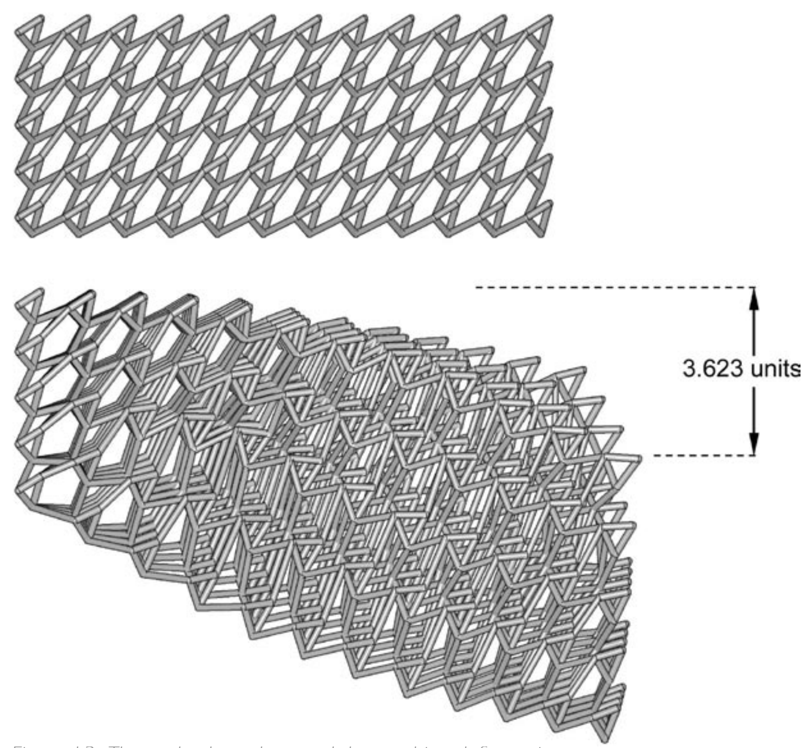

Figure 13. The evolved topology, and the resulting deformation under loading

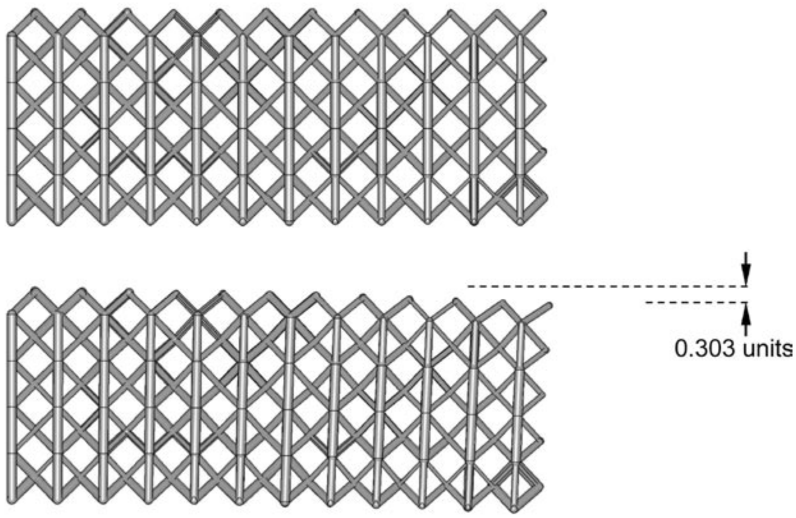

Figure 14. The engineered topology, and the resulting deformation.

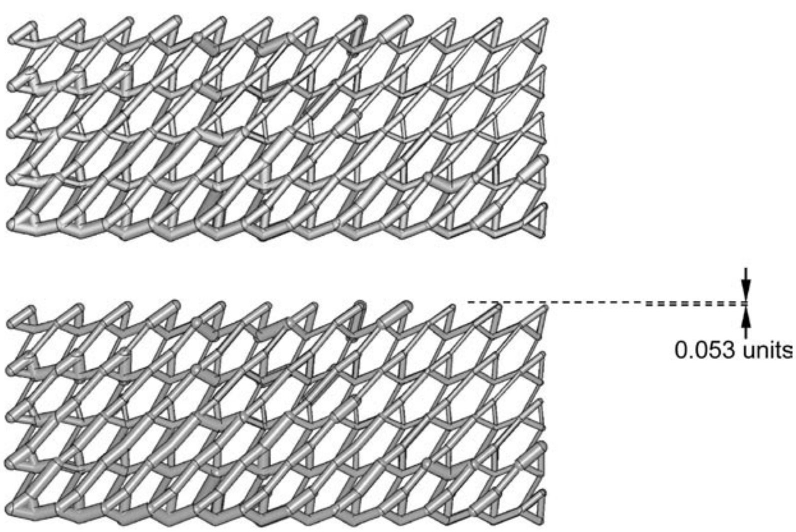

Figure 15. The evolved topology, and the resulting deformation.

The analysis of the generated designs under the specified loading conditions yielded favourable results. Several conclusions can be drawn from the comparative analysis of the performance of optimised and unoptimised structures under identical loading conditions:

- Structures in which the unit module is optimised to accommodate local force conditions outperform those that are uniform throughout. This is the case whether the topology is a generic grid or evolved for the specific stresses of the problem.
- During optimisation it is not necessary to simulate the entire structure in all its complexity, as an abstraction by Finite Element Analysis allows the local stresses to be approximated at any point. Local sampling is an approximation that applies to the structure as a whole, so optimising each unit-cube individually for these local stresses is a far more efficient method, shown by our results to be effective.
- Evolved topologies can enhance the optimisation by being more flexible. Particularly as applied forces become more complex there is a greater need for a non-standard topology. A greater ability to change geometry in terms of node positions and strut diameters allows a module to better accommodate local forces, and the GA can find topologies with this property.

These benefits would similarly apply to techniques at a far larger scale than stereolithography, and the principle of building complexity through the use of a flexible, changeable module can be applied elsewhere. Manufacturing methods for architectural components and large manufacturing industries are already in place for example, and in architectural and engineering applications the treatment and analysis of structure and loading is identical. Scalability functions in much the same way, as initial evolution of a topology is augmented by design time and detailing, but this need only be performed once.

Future work

In addition to structural efficiency the same process can be applied to other properties. A material's dynamic response may be designed for variation at different points: its movement under load in various orientations, flexibility, or Poisson's ratio (the amount of expansion in one axis under compression in another). Applications currently under research by the authors thus include the static, such as architecture, or the dynamic, such as robotics.

References

Brain, Marshall. (2002). *How Stereolithography Works*. www.howstuffworks.com/stereolith.htm.

Haroun Mahdavi, S., and Hanna, S. (2003). An Evolutionary approach to microstructure optimisation of stereolithographic models. *Proceedings of CEC2003, The Congress on Evolutionary Computation.*

Haroun Mahdavi, S., and S. Hanna. (2004). Optimising Continuous Microstructures: A Comparison of Gradient-Based and Stochastic Methods. In *Proceedings of SCIS & ISIS 2004, The Joint 2nd International Conference on Soft Computing and Intelligent Systems and 5th International Symposium on Advanced Intelligent Systems* (Forthcoming).

Lee, S-E., and Chen, J. C. (1999). Mass-customization methodology for an apparel industry with a future. In *Journal of industrial technology* 16(1). The National Association of Industrial Technology.

Molecular Geodesics, Inc. (1999). *Rapid prototyping helps duplicate the structure of life*. April 99 Rapid Prototyping Report.

Pine, B. J. II. (1993). *Mass customization.* Boston: Harvard Business School Press.

Sischka, J., Hensel, M., Menges, A., and Weinstock, M. (2004). Manufacturing complexity. In *Architectural design*, 74(3). London: Wiley-Academy.

Sean Hanna is currently researching optimisation and machine-learning methods as applied to design in an architectural context, with a focus on structural and spatial problems. Siavash Haroun Mahdavi's research involves the creation of self-adapting robots that change their behaviour and shape in response to their environment. Both share an active interest in rapid prototyping technologies. The authors have collaborated on materials and structural research for the past year, developing methods for structural optimisation and design algorithms.