

Evolutionary Form-Finding of Tensegrity Structures

Chandana Paul, Hod Lipson, Francisco J. Valero Cuevas
Mechanical and Aerospace Engineering
Cornell University
Ithaca, NY 14853, USA
cp228@cornell.edu

ABSTRACT

Tensegrity structures are stable 3-dimensional mechanical structures which maintain their form due to an intricate balance of forces between disjoint rigid elements and continuous tensile elements. Tensegrity structures can give rise to lightweight structures with high strength-to-weight ratios and their utility has been appreciated in architecture, engineering and recently robotics. However, the determination of connectivity patterns of the rigid and tensile elements which lead to stable tensegrity is challenging. Available methods are limited to the use of heuristic guidelines, hierarchical design based on known components, or mathematical methods which can explore only a subset of the space. This paper investigates the use of evolutionary algorithms in the form-finding of tensegrity structures. It is shown that an evolutionary algorithm can be used to explore the space of arbitrary tensegrity structures which are difficult to design using other methods, and determine new, non-regular forms. It suggests that evolutionary algorithms can be used as the basis for a general design methodology for tensegrity structures.

Categories and Subject Descriptors

Computer Applications, Computer Aided Engineering

General Terms

Algorithms

Keywords

Evolutionary Algorithms, Tensegrity Structures, Evolutionary Robotics.

1. INTRODUCTION

Tensegrity structures were first discovered by Snelson in 1948 and formally patented by Buckminster Fuller in 1962

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '05, June 25–29, 2005, Washington, DC, USA.
Copyright 2005 ACM 1-59593-010-8/05/0006 ...\$5.00.

[7], who coined the word *tensegrity* as an abbreviation of *tensile integrity*. The general definition of a tensegrity structure is a structure that maintains a stable volume in space through the use of discontinuous compressive elements (struts) connected to a continuous network of tensile elements (cables) [19]. Figure 1 shows the schematic and simulation of a simple tensegrity structure. Due to their design, tensegrity structures were discovered to have the ability to form the basis of lightweight and strong mechanical structures using little material. This gained them widespread popularity in architectural design for structures such as bridges and geodesic domes [10], in engineering for lightweight space structures such as deployable masts [8] and reflector antennas [24][12] and in robotics for the design of robot manipulators [22] and recently legged robots [18].

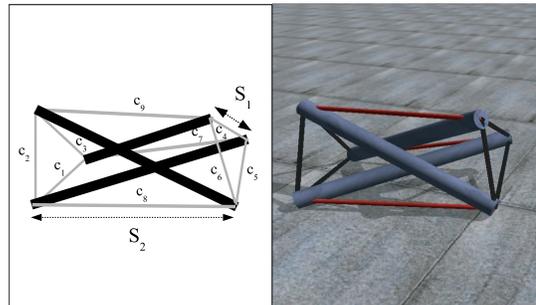


Figure 1: a) Schematic of a simple tensegrity structure. The thick black lines indicate the struts, the grey lines indicate the cables. b) Simulation of the three strut tensegrity prism.

There are various types of tensegrity structures. Connelly and Black [4] describe tensegrity structures in which vertices connected by a cable may be arbitrarily close, but they may not be further than the length of the cable joining them. Similarly two vertices joined by a strut cannot be closer than the length of the strut. This definition works well when the cables are made of inelastic material. However, when elastic cables are used, then the vertices may be further apart than the rest length of the cable, if appropriate force is applied. There are also variations in the level of spatial proximity allowed between the struts. In the definition of a *class I* tensegrity structure [1], the struts cannot share common vertices, that is they must be physically separated in space. However, the definition of a *class II* tensegrity structure allows more than one strut to originate from a vertex. Finally, although the canonical form of a tensegrity

rity structure is composed of rigid struts and tensile cables, tensegrity structures also exist in which the struts are elastic and connected by cables or sheets of material¹. In this work, we are primarily concerned with the form-finding of class I tensegrity structures.

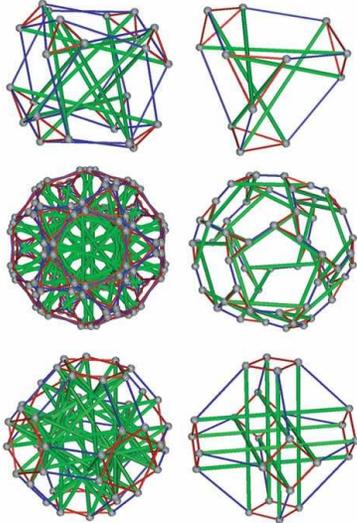


Figure 2: A few examples of regular tensegrities developed based on group theory by Connelly and Black [4]. Reprinted with permission from *American Scientist*.

The problem of *form finding* of a tensegrity structure, or determining the geometrical configuration of a stable 3D tensegrity structure, has received widespread attention. There are two components to the problem of form-finding: Determining a connectivity pattern which will enable a stable form to exist, and determining the lengths of the rigid and tensile elements for a given connectivity pattern which result in a stable form. There have been numerous approaches to form-finding which address one or both of these issues. Early structures developed by Fuller and Snelson, used convex polyhedra as the basis for form-finding. These methods were focused on specifying connectivity patterns which could enable tensegrity. This approach resulted in various configurations which were summarized by Pugh [19]. However, these were all based on regular geometries. More recently other methods were developed which primarily focused on determining the length parameters of the rigid and tensile elements in the stable configuration. These methods involved the use of non-linear programming [20], dynamic relaxation [15] and calculation of force density [21][13][26]. These methods did not provide guidelines for finding connectivity patterns. Finally, Connelly and Black [3][4] developed a new method for determining both connectivity patterns and parameters for structures based on group symmetries. Once again, these resulted in only regular structures (Fig. 2). A good review of these and other form finding methods can be found in [25].

The general problem of determining a connectivity pattern and set of parameter which lead to a stable tensegrity,

¹The pop-up tent is a good example of such a tensegrity structure.

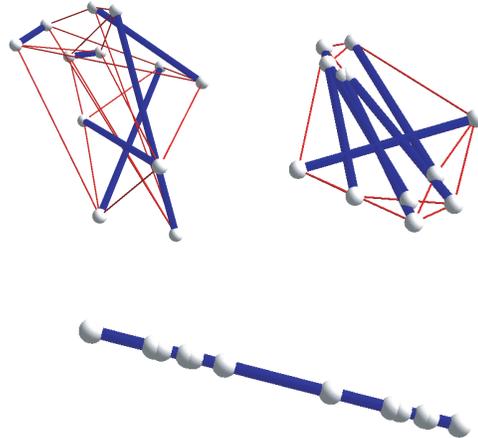


Figure 3: (Top Left) a. A sample initial configuration, before relaxation. (Top Right) b. Relaxation into a 2D tensegrity structure. (Bottom) c. Relaxation into 1D.

for the entire space of structures, is an unsolved problem. There is currently no direct mathematical way to derive the form of an arbitrary tensegrity structure.

2. EVOLUTIONARY ALGORITHM

In this work, a genetic algorithm was used to evolve the connectivity pattern and parameter values of a tensegrity structure. The genotype utilized a direct encoding to specify the initial position of the vertices and the connectivity pattern of the struts and cables (Fig. 3(a)). Following, this a relaxation algorithm was applied such that all the cables simultaneously contracted in length, by setting their rest length to zero, and all struts were set to length 1. In many cases, the structure would contract into a 2-D tensegrity structure (Fig. 3(b)) on applying relaxation, and in others it would collapse into a 1-D bundle of struts (Fig. 3(c)). Only in some cases did the structure contract into a 3-D tensegrity structure. The fitness of the structure was determined by fitting a tight rectangular bounding box around the final configuration after relaxation and computing its volume. The highest volume individuals were selected, and mutation and crossover were applied. Iteration of this process led to stable 3-D tensegrity structures with 6, 7, 8, 10 and 12 struts. The following sections will describe the algorithm and its results.

2.1 Representation

In this study only class I tensegrity structures were of interest. Furthermore, only tensegrity structures with three cables originating from each vertex were used. Three is the minimum number of cables required at each vertex to form a tensegrity. If there are less than three cables, then the vertex is not fully constrained in three dimensions and therefore does not maintain a stable position with respect to the rest of the structure, violating the definition of a tensegrity structure. Here, it was of interest to study whether the evolutionary algorithm could design tensegrities using the minimum number of cables required.

A direct encoding was used to represent the the genotype to phenotype mapping for the tensegrity structure. The genome string was divided into three regions. The first region was used to specify the initial location of the vertices. The number of vertices in a class I tensegrity structure is determined by the number of struts. If there are n struts, then by definition there are $2n$ vertices, as none of the struts can share common vertices. Thus, the first part of the genome string consisted of triples with floating point values between 0.0 and 1.0 corresponding to the initial $[x,y,z]$ positions in space for each of these $2n$ vertices. To translate the genotype to the actual location of the vertices in space, each value was multiplied by a constant k ($=40$).

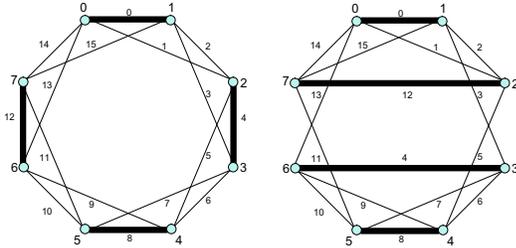


Figure 4: (a) An example of a seed structure with 4 struts. The thick lines indicate struts and the thin lines indicate cables. (b) A shuffle action in which the endpoints of two struts are switched.

The initial location of the vertices was used to define a *seed* structure (Fig. 4(a)). The connectivity pattern of the seed structure was always the same. The vertices were labeled from 1 to n . Each vertex was connected to the two following vertices by links. Thus, for example vertex 1 was connected by links to vertices 2 and 3, and vertex 7 was linked to 0 and 1. The links were labeled in the order they were created from 1 to m , where m was the number of links ($m = 4n$). Following the creation of the links, every link labeled as a multiple of 4 was converted to a strut, and all the others were converted to cables. This ensured the topological validity of the structure: That no two struts shared a common vertex and every vertex had exactly three cables originating from it.

The second region of the genome string encoded the final connectivity of the struts through a sequence of pairwise *shuffle* actions to be performed on struts of the seed structure. In a shuffle action, two struts S_1 and S_2 were selected. If strut S_1 was located between two vertices a_1 and b_1 and strut S_2 was located between two vertices a_2 and b_2 , one of these vertices was selected on each strut as the target of the shuffle action and the locations of these two vertices were switched. Thus for example, as shown in Fig. 4(b), the struts labeled as 4 and 12 have been selected as targets of the shuffle action. The endpoint of strut 4, originally located at vertex 2, is moved to vertex 6. Correspondingly, the endpoint of strut 12, originally located at vertex 6, is moved to vertex 2. As can be seen, such an action gives rise to a new connectivity pattern, which can in turn specify a new tensegrity structure. Thus, the second region of the genome string consisted of 4-tuples representing pairs of struts s_1 and s_2 and vertices p_1 and p_2 which were to be shuffled. The number of shuffle actions represented was equal to two times the number of struts.

The third region of the genome string encoded the final location of the cables through a sequence of similar shuffle actions. Thus, it consisted of 4-tuples of floating point values representing pairs of cables c_1 and c_2 and their vertices p_1 and p_2 that were to be switched. The number of shuffle actions performed on the cables was equal to two times the number of cables. The total length of the genome string was therefore $3n + 8s + 8c$, where $n = 2s$ was the number of vertices, s the number of struts, and c the number of cables.

2.2 Relaxation Algorithm

In order to find the equilibrium state of the tensegrity structure a kinematic relaxation algorithm was applied [14]. It was comprised of two steps, applied iteratively. The first step computed the total force on each joint; this was the residual force, since these forces would be zero if the structure was in perfect equilibrium. The second step adjusted the position of each vertex so as to reduce the residual force acting upon it. These two steps were repeated until the maximum residual force went below a desired threshold.

As each vertex was connected to four links corresponding to three cables and one strut, each of these links applied a force on the vertex at each time step. Each cable was set to contract to zero length, and each strut to a length of 1. Thus, at each time step each cable applied a force proportional to its length, and each strut applied a force proportional to the difference of its length from 1.

The effective stiffness of each joint was estimated by summing all the stiffnesses of all links connected to it. Given the residual force acting on each joint and its effective stiffness, the displacement of each joint was computed as the component-wise ratio between residual force and stiffness. The vertex positions were then adjusted by a fraction of this displacement. For greater computational efficiency, many of the above steps can be calculated concurrently. A detailed description and sample code for the kinematic relaxation solver is provided in [14].

2.3 Fitness Criteria

2.3.1 Bounding Box

The first fitness criterion selected was the volume of the bounding box encompassing the structure. To determine a tight bounding box, the volume spanned by a candidate tensegrity structure is usually computed using the convex hull spanned by its vertices. The convex hull can be found using a number of algorithms, such as the Incremental, Gift Wrap, Divide and Conquer, and QuickHull methods that run in $O(n \log n)$ time [5]. We used a tight bounding-box approach that provides an upper bound on the convex-hull volume, but amplifies any gradient towards generating a volumetric tensegrity because any deviation from a plane is multiplied by the plane's area. The current fastest algorithm for calculating a 3D minimal volume bounding box runs in $O(n^3)$ time [17], although faster approximate methods exist [2]. Here we used an implementation that is based on the property of a bounding box that at least one face of the box contains a face of the convex hull, and at least one edge of the bounding box contains an edge of the convex hull.

We calculated the volume of a bounding box leaning against each triplet of points of the convex hull structure in $O(n^4)$

time. Each triplet of vertices of the structure was considered in turn, without repetitions or permutations in $O(n^3)$ time. Each triplet defines a plane, and each pair defines an orientation of a candidate bounding-box coordinate system. A min-max bounding box was computed for each candidate orientation in $O(n)$ time, and the minimal one selected as the overall tightest bounding box.

2.3.2 Ordered Tetrahedra

Although the bounding box criterion was elegant in its simplicity and ability to demonstrate the power of evolution for the design of tensegrity structures, it was not the best criteria. As the bounding box was only determined by the extremal points of the structure, the actual volume of the structure itself could vary greatly within the space of the bounding box. Thus, structures with significantly different volumes could receive the same fitness value, confounding the gradient necessary for evolutionary search.

To alleviate this problem a second fitness criterion was designed, which involved considering the volumes of tetrahedra formed by points of the structure. The vertices of the structure, which were the endpoints of the struts, were ordered from left to right along the x-axis. Then going through these points in order, tetrahedra were formed by connecting the point to its three following points, and the volumes of these tetrahedra were calculated and summed. This ensured, for example, that if the tensegrity structure was composed of two 2D tensegrity structures attached at an angle, then the fitness would still be low, as considering the points in sets of four would reveal that various parts of the tensegrity were flat. This criterion was used to compare the performance of the algorithm against random search and gradient ascent.

2.4 Evolution

A fixed length genetic algorithm was used to evolve the controllers. Each run of the genetic algorithm was conducted for 20 generations, using a population size of 20. At the end of each generation, the 10 most fit genomes were preserved; the others were deleted. Tournament selection with a tournament size of three, was employed to probabilistically select genotypes from among those remaining for mutation and crossover. 2 pairwise one-point crossings produced 4 new genotypes: the remaining 6 new genotypes were mutated copies of genotypes from the previous generation. The mutation rate was set to generate an average of n mutations for each new genome created, where n was defined as a function of the genome length g_l , as $mn=g_l/100$. Mutation involved the replacement of a single value with a new random value. The floating-point values were rounded to two decimal places and thus ranged between 0.00 and 1.00.

3. RESULTS

Two evolutionary trials were initially performed with 7 struts. The history of the best fitness and the average fitness in each generation is shown in Figure 5. The best fitness at the end of 20 generations on the first trial was a volume of 45876.7 and on the second trial it was 45602.5 indicating that the two trials had found very similar tensegrity structures. The final form of the tensegrity structures found by the algorithm is shown in Figures 6.

Following these trials which indicated that evolution could be used to discover new tensegrity structures, a more systematic evaluation of the performance of the algorithm with

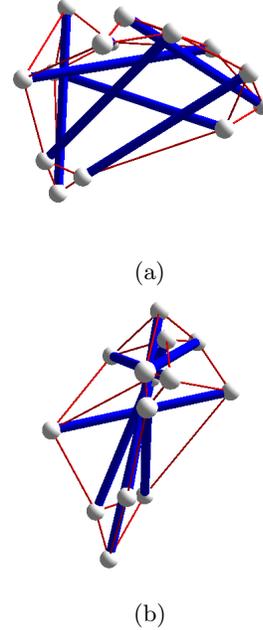


Figure 5: (a) Best 3D tensegrity structure evolved in evolutionary trial 1. Volume = 45876 (b) Best 3D tensegrity structure evolved in evolutionary trial 2. Volume = 45603

respect to the number of struts was carried out. In this set of experiments, the algorithm was run for 6, 8, 10 and 12 struts respectively for 100 generations each. The fitness history that resulted in each of these experiments is presented in Fig. 8. The best fitness achieved at the end of 100 generations was 32040 with 6 struts, 58534 with 8 struts, 71734 with 10 struts, and 95296 with 12 struts. These structures are shown in Fig. 6(a)-(d). The fact that the volume was observed to increase with the number of struts was not surprising as a larger space can usually be spanned with a larger number of elements, if the elements are of the same average size, as is the case in the experiments.

A more interesting characteristic of these graphs is the shape of the curve of best fitness. In the experiment with 6 struts, a large increase in fitness occurs relatively early (approximately at generation 18), followed only by insignificant changes. In the experiment with 8 struts, a large increase in fitness occurs around generation 29, after which no further changes occur. However, the size of this change is smaller than the increase at generation 18 with six struts. In the experiment with 10 struts, the size of the increments are smaller than the large jumps seen in the previous graphs, and continue until generation 50. Finally, in the experiment with 12 struts, the changes are smaller on average and continue until generation 70. Thus, the trend seems to be that at lower dimensions larger but more infrequent changes occur early in the evolutionary timeline, while at higher dimensions the changes are smaller, but more frequent and continue for a longer duration of the evolutionary timeline. The latter is a characteristic pattern of successful evolutionary optimization, while the former is characteristic of performance hindered by isolated relative maxima.

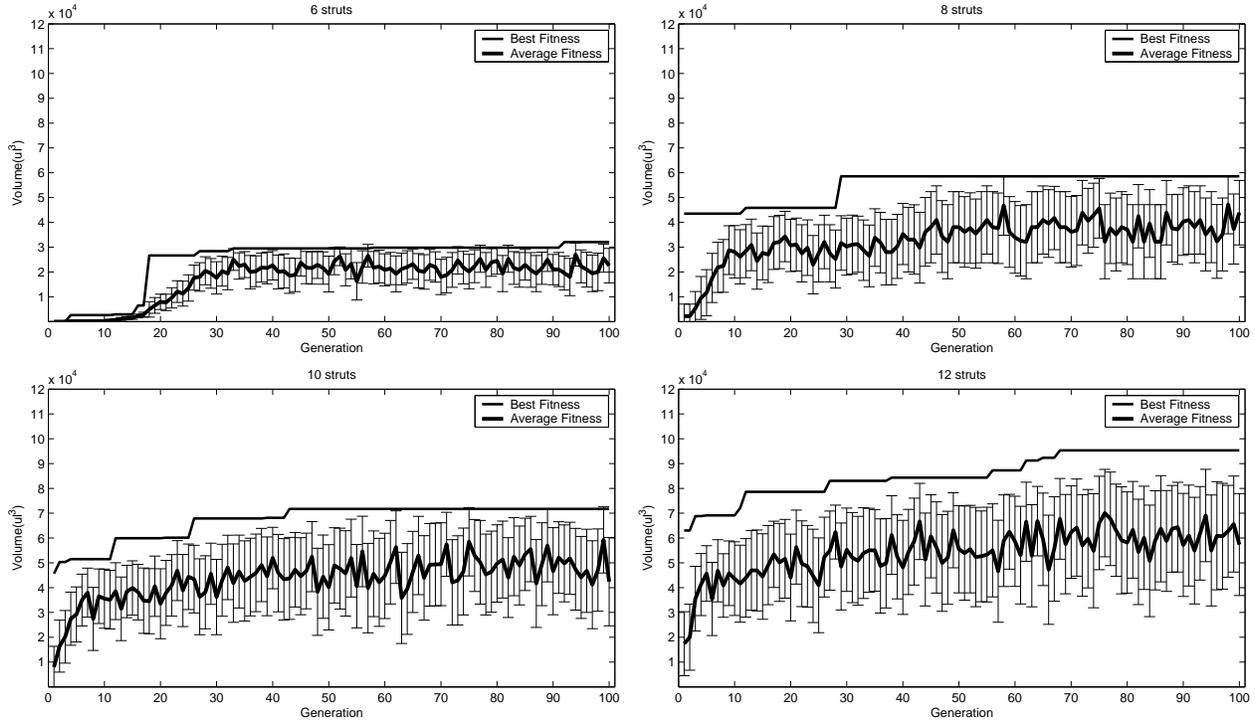


Figure 7: Best and average fitnesses plotted for each generation for evolutionary trials with 6, 8, 10 and 12 struts. The x-axis shows the generation number. The y-axis shows the fitness, which is the volume of a minimum bounding box that can be fitted around the structure, given in terms of unit length (ul). The length of the error bars along the plots of average fitness represent the standard deviation.

Fig. 8 presents bar graphs of the number of 3D, 2D and 1D structures in the population at every generation, during the first 20 generations of the four experiments. The graphs are presented in order to compare the ease with which 3D structures can be generated in the various trials. With 6 struts, it is observed that for the first 15 generations, a large number of 2D tensegrity structures are produced, and the remainder are 1D structures. Only in generation 16 is the first 3D tensegrity structure found. Compared to this, in the graph of the experiment with 8 struts, a 3D tensegrity structure is already found in generation 1, purely as result of random search. Subsequently, in the graph of the experiment with 10 struts, 4 are found in generation 1, and with 12 struts, 7 are found. Thus, the number of 3D tensegrities found purely as a result of random search seems to increase with the number of struts.

The reverse trend is observed for 1D structures. In the graph of the experiment with 6 struts, 14 of the 20 generations have a non-zero value for the number of 1D structures, with the highest value being 5. In the experiment with 8 struts, 11 have a non-zero value, with the highest value being 3. With 10 struts, 11 have a non-zero value, with the highest value equal to 2. Finally, with 12 struts only 1 of the generations produces a 1D structure. This suggests that fewer 1D solutions exist at higher dimensions.

As Fig. 8 showed the total number of tensegrity structures of a given dimensionality it was possible that in each generation some of the individuals shared the same genotype, thus confounding the estimate of the relative frequency of the various dimensional structures in the solution space. Thus,

the number of unique 3D structures created in every generation was extracted for the four experiments, and plotted in Fig. 9. Although, for this set of experiments, the evolutionary trial with 6 struts produced more unique structures than 8, the number of unique structures produced with 10 struts was almost double, and with 12 struts almost triple that of the one with 6 struts, suggesting once again that there may be some positive correlation between the number of struts and the number of unique 3D tensegrity structures that can be found.

In order to determine whether evolution could perform better than random search and gradient ascent, the ordered tetrahedron fitness criterion was used to compare the performance of both algorithms on the design of 12 strut tensegrity structures. Three experiments were run for both conditions. The results are shown in Fig. 10. In these graphs it is clear that evolutionary optimization performs better than both random search and gradient ascent.

Finally, as a proof of concept that the structures designed in simulation translate to the real world, a physical tensegrity structure was implemented based on the connectivity pattern specified by a model obtained from simulation. The resulting tensegrity structure is shown in Fig. 11. The geometry of the structure did not resemble the simulation exactly as the length and pre-stress parameters of the cables were not replicated exactly. However, it showed that the connectivity patterns which gave rise to successful tensegrity structures in simulation, could translate to equivalent structures in the real world.

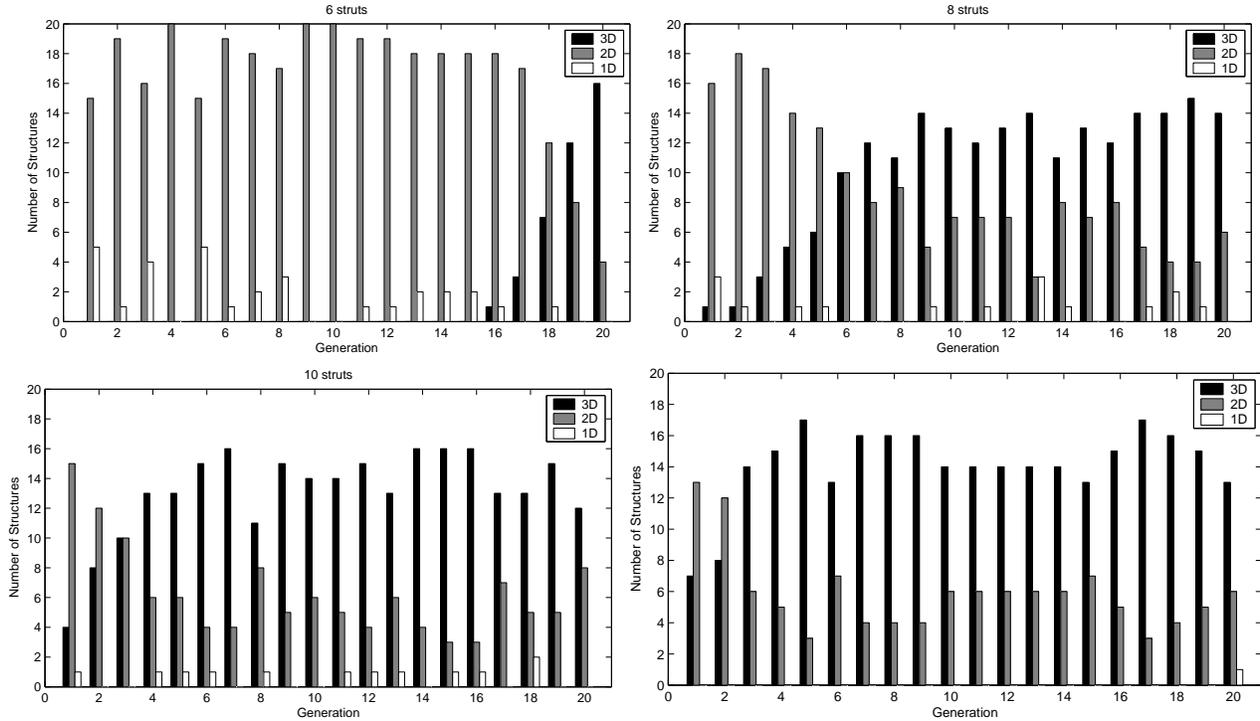


Figure 8: Number of 3D, 2D and 1D structures in the population during the first 20 generations of the evolutionary experiments with 6, 8, 10 and 12 struts.

4. DISCUSSION

The preliminary results are promising as they show that evolutionary search can be useful in the design of new tensegrity structures. In the experiments, the search algorithm was able to find stable 3D tensegrity structures for all the dimensions that were tested, including for 6, 7, 8, 10 and 12 struts. For 6 struts, 81 unique tensegrity structures were found, for 8, 52 were found, for 10, 143 were found and for 12, 240 structures were found as a result of a single evolutionary search. It is likely that running a large number of evolutionary experiments could be used to efficiently search a large region of the space of solutions, potentially much larger than what was possible using previous methods.

However, there are a few areas in which evolutionary search may have difficulties in the form-finding of tensegrity structures. The results suggest that at lower dimensions the fitness landscape becomes more rugged, with isolated relative maxima surrounded by deep troughs. This is also seen by directly performing mutations on a successful genome. For example, taking the genome string which leads to the best fitness of 48576.7 for 7 struts, and inducing a single mutation in a parameter which encodes a cable in a pairwise shuffle, leads to fitnesses of 0.0, 9.541, 42696.9 and 1.435. Three out of four of these represent a drastic decline in performance indicating that the fitness landscape has very abrupt changes. Such landscapes are typically challenging for evolutionary and other gradient based search methods.

The landscape is also determined to some extent by the encoding scheme. As in the current encoding scheme the connectivity is specified by a direct encoding, a mutation induces a change in the position of a vertex or the location of two struts or two cables. While, the position of a single vertex does not significantly affect the final configuration

of the tensegrity structure, a change in the location of two struts or two cables represents a drastic change in the configuration, and can be expected to lead to drastic consequences in terms of fitness. Thus, for lower dimensionality alternate encoding schemes may also prove effective.

However, the results suggest that for higher dimensionality, evolution using a direct encoding scheme is suitable for generating a multitude of unique solutions. This is ideal for the purposes of design, as it is only at higher dimensions that the other design techniques fail in the free-form generation of structures. Evolution can be used to overcome this limitation, and provide an effective tool for design.

The trend observed in the relationship between scale and evolvability is interesting as it is the opposite of what is usually found with a direct encoding. Usually, a direct encoding performs well at lower dimensions and its performance drops off as the dimensionality of the problem increases. In other words, it does not scale well with size. For higher dimensional problems it has often been found that modular and hierarchical encoding schemes are more effective. However, in this particular problem it seems that the direct encoding performs better as the dimensionality of the problem increases. The unexpected result that it is able to find a larger number of 3D structures sooner on the evolutionary timeline suggests that perhaps in this particular problem space, more solutions exist at higher dimensions than at lower dimensions.

5. CONCLUSION

The form-finding of tensegrity structures, particularly of an irregular nature, has been a difficult problem, which has not easily yield itself to direct mathematical analysis. Thus, so far the existing methods for the free-form design of tensegrity

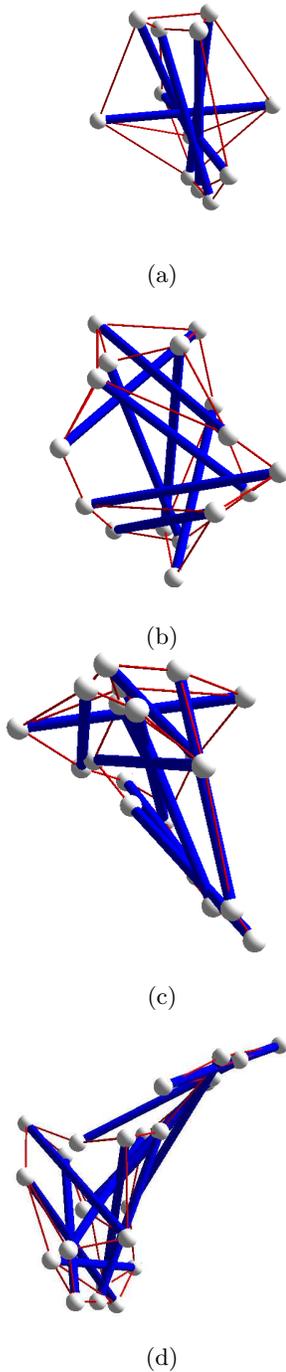


Figure 6: Best 3D tensegrity structures evolved with (a) 6 struts (volume = 32040) (b) 8 struts (volume = 58534) (c) 10 struts (volume = 71734) and (d) 12 (volume = 95296) struts.

rity structures have been limited. The goal of this paper was to attempt to surpass this limitation by applying evolutionary computation to the problem of form-finding. Using a direct encoding to represent the connectivity pattern of

a tensegrity, it was shown that an evolutionary algorithm could successfully explore the space of irregular structures and give rise to new unknown configurations. These preliminary results suggest that evolutionary search may serve as a powerful tool in the free-form design of tensegrity structures.

6. ACKNOWLEDGEMENTS

This research has been supported by the DCI Postdoctoral Research Fellowship Program, award number NMA501-03-1-2013.

7. REFERENCES

- [1] R. Adhikari, R. E. Skelton, and W. J. Helton, Mechanics of Tensegrity Beams, UCSD Structural Systems and Control Lab, Report No. 1998-1, 1998
- [2] Gill Barequet, and Sariel Har-Peled (1999) Efficiently Approximating the Minimum-Volume Bounding Box of a Point Set in Three Dimensions, Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'99)
- [3] Connelly, R. and Terrell M. Globally rigid symmetric tensegrities. *Structural Topology* 21 (1995), pp 59-78
- [4] Connelly, R. and Black, A. "Mathematics and Tensegrity," *American Scientist*, Vol. 86, 1998
- [5] Mark De Berg, Marc Van Kreveld, Mark Overmars, Otfried Schwarzkopf (1997) *Computational Geometry : Algorithms and Applications*, Springer-Verlag
- [6] Fuller, R. (1961). Tensegrity. *Portfolio and Artnews Annual*, 4, 112-127.
- [7] Fuller, R. B., Tensile-integrity structures, United States Patent 3,063,521, November 13, 1962.
- [8] H. Furuya, "Concept of Deployable Tensegrity Structures in Space Applications," *Int. J. Space Structures*, 7:2, pp 143-151, 1992
- [9] Hanaor, A. "Tensegrity: Theory and Application" *Beyond the Cube: The Architecture of Space Frames and Polyhedra*, edited by J. Francois Gabriel, John Wiley and Sons, Ic, 1997 pp 385-408.
- [10] Hanaor, A. "Aspects of Design of Double Layer Tensegrity Domes" *International Journal of Space Structure*, Vol. 7, No. 2, 1992 pp 101-103
- [11] Kenner, H. *Geodesic math and how to use it*. Berkeley, University of California Press, 1976
- [12] Knight, B.F., *Deployable Antenna Kinematics using Tensegrity Structure Design*, Doctor of Philosophy Dissertation, University of Florida, Gainesville, 2000.
- [13] Linkwitz, K., Formfinding by the "direct approach" and pertinent strategies for the conceptual design of pre-stressed and hanging structures, *International Journal of Space Structures*, 14(2), 1999, 73-87
- [14] Lipson H. (2006) "A relaxation method for simulating the kinematics of compound nonlinear mechanisms", *ASME Journal of Mechanical Design*, *in press*.
- [15] Motro, R., *Forms and Forces in Tensegrity Systems*, in H. Nooshin, ed., *Proceedings of the Third International Conference on Space Structures*, Elsevier, Amsterdam, 1984, 180-185
- [16] Motro R., "Tensegrity Systems: The State of the Art", *International Journal of Space Structures*, Vol. 7, No. 2, 1992 pp 75-84

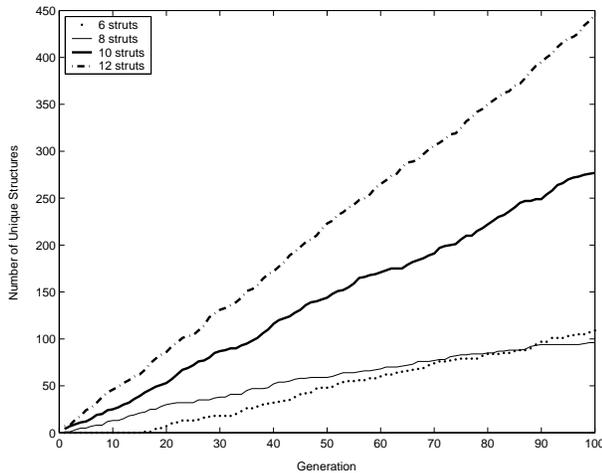


Figure 9: Number of unique tensegrity structures found by Generation n , for the experiments with 6, 8, 10 and 12 struts.

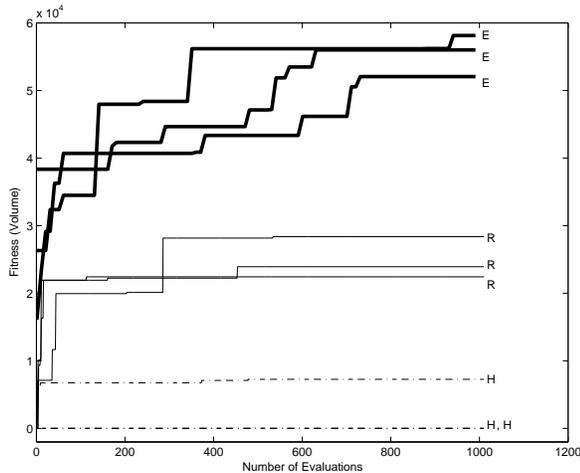


Figure 10: Comparison of performance of evolution vs. random search and gradient ascent. The results of the evolutionary algorithm are indicated by thick lines and labelled as (E), those of random search are indicated by thin lines labelled as (R), and those of gradient ascent (hill climber) are indicated by dashed lines labelled as (H).

[17] J. O'Rourke. Finding minimal enclosing boxes. *International Journal of Computer and Information Sciences*, 14(3):183-199, June 1985

[18] Paul, C., Lipson, H., and Valero-Cuevas, F.J. Gait Production in a Tensegrity Based Robot. In *Proc. Int. Conf. on Advanced Robotics*, to be held in Seattle, WA, USA, July 2005.

[19] Pugh, A., *An Introduction to Tensegrity*, University of California Press, 1976.

[20] Pellegrino, S. *Mechanics of kinematically indeterminate structures*, PhD. dissertation, University of Cambridge, U.K., 1986.

[21] Schek, H. J., *The force density method for form*

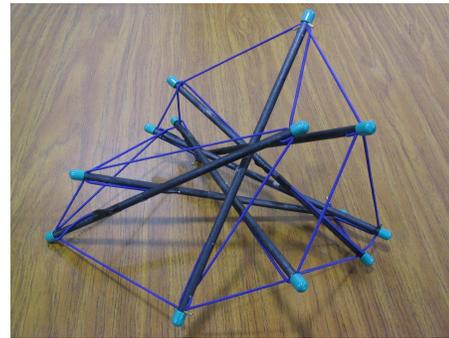
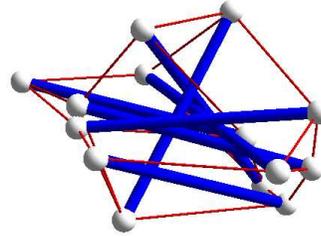


Figure 11: Real world implementation of tensegrity structure based on connectivity pattern specified by simulation. (The geometry of the structure does not match simulation exactly as cable lengths and pre-stress parameters were roughly approximated.)

finding and computation of general networks, *Computer Methods in Applied Mechanics and Engineering*, 3, 1974, 115-134

[22] Aldrich, J. B., Skelton, R. E. and Kreuz-Delgado, K. Control Synthesis for a Class of Light and Agile Robotic Tensegrity Structures, *Proceedings of the IEEE American Control Conference*, Denver, Colorado, USA, June 4-6, 2003

[23] Snelson, K. D., Continuous tension, discontinuous compression structures, United States Patent 3,169,611, February 16, 1965.

[24] Tibert, G. *Deployable Tensegrity Structures for Space Applications*, PhD Thesis, Department of Mechanics, Royal Institute of Technology (KTH), Sweden, 2002

[25] Tibert A. G. and Pellegrino, S. "Review of Form-Finding Methods for Tensegrity Structures", *International Journal of Space Structures*, 18:4(2003), pp 209-223

[26] Vassart, N., and Motro, R., "Multiparametered form-finding method: Application to tensegrity systems. *International Journal of Space Structure*, 14:2(1999), pp 147-154