Harmonic form-finding for the design of doubly-curved shells

Cecilie Brandt-Olsen, Paul Shepherd*, Paul Richens

Dept. of Architecture & Civil Engineering, University of Bath, Bath BA2 7AY, UK

Abstract

Shell structures are a highly efficient and elegant way of covering large uninterrupted spaces, but their complex geometry is notoriously difficult to model and analyse. This paper describes a novel free-form shell modelling technique based on structural harmonics. The method builds on work using weighted Eigenmodes for three-dimensional mesh modelling in a computer graphics setting, and extends it by specifically adapting the technique to an architectural design context. This not only enables the sculpting of free-form architectural surfaces using only a few control parameters, but also takes advantage of the synergies between Eigenmodes and structural buckling modes, to provide an efficient means of stiffening a shell against failure by buckling. This paper includes a full case study of the iconic British Museum Great Court Roof to demonstrate the applicability of the developed framework to real-world problems, and the software developed to implement the method is available as an open-source download.

Keywords: Harmonics, Fourier analysis, Free-form surfaces, Curvature, Buckling

1. Introduction

The last decade of technological advancements, both in design and construction, has set architects free to explore a wide variety of building shapes, which in effect has given rise to a trend of complex free-form buildings in contemporary architecture. Since the requirements for low-energy buildings are ever more stringent, and designs focus on material economy and efficiency, it is no longer sufficient to design such buildings from a purely aesthetic point of view. As a response, form-finding techniques have been developed to generate efficient shapes, where form follows force. However, these methods are associated with a very restricted design space, expressing a theoretical optimum with little or no appreciation of other design constraints, whereas the best designs in practice

*Corresponding author
Email address: p.shepherd@bath.ac.uk (Paul Shepherd)

http://dx.doi.org/10.1108/EC-12-2018-0557

strike a balance between form and force as part of an integrated design process. Such a process is generally supported by two components: a flexible modelling strategy using a limited number of variables to control the geometry; and a live evaluation of some property that is desirable for design. NURBS, and to a lesser extent Subdivision Surfaces, are typically used for the first part, whilst stresses or deflections are the most common structural quantities to measure success. Surprisingly little research has been carried in the use of buckling as a criteria at the design stage, even though this is often later identified as a critical failure mode for shells.

Observations of historic buildings, such as those of Candela and Dieste for example, combined with a newly conducted study by Malek (2012), point towards the influence of curvature as an essential design parameter, to increase the stiffness of a shell and hence its buckling capacity. Interestingly, a novel free-form modelling approach proposed by Michalatos and Kaijima (2014) generates shapes by a summation of waves, which are inherently rich in curvature. However, the approach has several unexplored properties in relation to Fourier analysis, is not easily understood by the uninitiated, and needs further adaptation to an architectural context.

The research presented here further develops the described free-form modelling approach, and takes advantage of the resulting double curvature to stiffen the shell against buckling failure. It is implemented as an interactive plug-in to the Grasshopper modelling environment (McNeel, 2018), a visual-programming style parametric-modelling software application commonly used in the architecture, engineering and construction (AEC) industry. This ensures the implementation is compatible with existing architectural design work flows seen in practice, and is released under an open-source license (Brandt-Olsen, 2018a) to promote industry uptake. The plug-in aims to assist designers in optimising their response to conflicting design constraints and deliver more economic solutions without restricting artistic creativity.

2. Harmonic modelling

Harmonic modelling was originally developed by Taubin (1995) in the context of surface fairing, but has recently found several other applications, including mesh quadrangulation, mesh segmentation and geometry compression (Zhang et al., 2007). The work of Hildebrandt et al. (2011) is particularly elegant, since it uses the approach both as a means of parametric dimension reduction and to facilitate shape modification in a computer graphics context. In this current paper, the authors develop these ideas specifically for real-world structural design applications, combining the low dimension parametrisation with an understanding of how the structure will resist buckling loads. This provides design flexibility, control of surface smoothness and allows designers to take into account structural engineering considerations during the early stages of design. The variables associated with this approach are numerical scalar values that control the weighted summation of mode shapes with different frequencies. This concept is more abstract than the usual representation of a NURBS-like

control polygon, but has other advantages. For example, it is known from Fourier analysis, that representing a function in the frequency domain provides information that it is impossible to extract from the spatial domain. Separating a shell-surface into a linear combination of its component harmonics can inform an understanding of how it will behave as a structure, as well as allowing surface smoothing by removing undesirable high-frequency components that contribute insignificantly to the final shape, a consideration that is important in architectural design.

2.1. Framework

The translation of the theory behind Fourier analysis to arbitrary discrete meshes would not be directly possible without the missing link observed by Taubin: "The classical Fourier transform of a signal can be seen as the decomposition of the signal into a linear combination of the eigenvectors of the Laplacian operator" (Taubin, 1995). Mathematically, this can be expressed as:

$$L \cdot \overrightarrow{v} = \lambda \cdot \overrightarrow{v} \tag{1}$$

Here L is the Laplacian matrix, λ the eigenvalues and $\overrightarrow{v}$ the eigenvectors. The Laplacian is a second order differential operator in n-dimensional space, which defines the divergence of the gradient. For a mesh, this is equivalent to the difference between a function value at a specific vertex and the (weighted) average of the function at the 1-ring connected vertices. Even though the Laplacian only incorporates local information, it is still capable of acting globally and reveals properties that are unique to the given mesh (Zhang et al., 2007). In matrix form, it is defined as an n *by* n matrix, where n is the number of vertices in the mesh and its elements are given by:

$$L_{ij} = \begin{cases} -w_{ij} & \text{, if } i \neq j \text{ and } v_i \text{ adjacent to } v_j. \\ \sum_{j \in N(i)} w_{ij} & \text{, if } i = j. \\ 0 & \text{, otherwise.} \end{cases} \tag{2}$$

Here $w_{ij} = w_{ji}$ is an edge weight and $N(i)$ is the set of vertices included in the 1-ring neighbourhood of vertex v_i. Masses associated with vertices v_i and v_j are typically included in the edge weight, which generally result in a non-symmetric matrix. However, symmetry is prioritised in this context to guarantee real eigenvectors that create an orthogonal basis equivalent to the complex exponential function used for the Fourier Transform (Zhang et al., 2007).

Since a discrete Laplacian that meets all the properties of the continuous operator cannot exist on general meshes (Vallet and Lévy, 2008), several different weighting schemes exist. The simplest discretisation of the operator is the graph Laplacian, which is defined by an edge weight equal to 1 and with a uniform mass distribution. It means that the matrix has -1 in off-diagonal cells where two vertices are connected by an edge, and the sum of the edge weights (thus the valence) on the diagonals. This definition of the graph Laplacian is

very simple and easy to compute and is applicable to arbitrary mesh topologies, but the drawback is that it solely depends on topology and not geometry. As a result, Pinkall and Polthier (1993) derived the widely used geometric mesh Laplacian with cotangent weights:

$$w_{ij} = cot(\alpha_{ij}) + cot(\beta_{ij}) \tag{3}$$

Where α_{ij} and β_{ij} are the angles at the two vertices opposite to the edge connecting vertices v_i and v_j as illustrated in Figure 1. In order to compute an opposite angle, the mesh must be triangulated. For an edge located at a boundary, there is no opposite angle, and so α_{ij} or β_{ij} is set to zero, corresponding to Neumann boundary conditions (Vallet and Lévy, 2008).

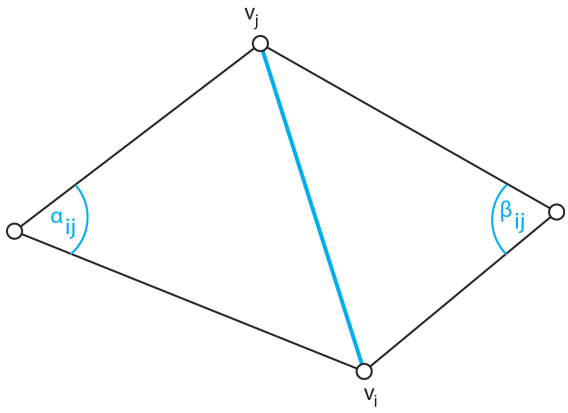

Figure 1: Angles associated with the cotangent weights

This definition lacks a proper mass weighting, which means that the weights are dependent on the mesh density. To account for this, Vallet and Lévy (2008) propose the following definition, which respects the symmetry criteria:

$$w_{ij} = \frac{cot(\alpha_{ij}) + cot(\beta_{ij})}{\sqrt{A_i \cdot A_j}} \tag{4}$$

Here A_i is the area associated with vertex v_i and A_j the area associated with vertex v_j. These areas can be calculated either using the barycell or voronoi cell. The latter ensures that the area is independent of the mesh topology and can be easily calculated from the hybrid approach of Meyer et al. (2003), so it is the method adopted by the authors.

The eigenvalue problem in Equation 1 has many similarities with a finite element modal analysis, which performs an eigen-decomposition of the stiffness matrix of a structure in order to describe the vibrational modes (Cook, 1995). However, more information, such as choice of material and degrees of freedom (DOFs) is necessary to construct this stiffness matrix, which is unavailable in the

context of free-form modelling at the conceptual design stage. The advantage of the Laplacian operator is that it is solely related to geometry, but nevertheless it can be thought of as a simplified stiffness matrix. This coherence gives a physical interpretation of the eigenvectors as the vibrational modes of the mesh, and the eigenvalues as the squared frequencies of vibration (Dong et al., 2006). This physical analogy is important, because it enables the eigenvectors to be sorted according to the magnitude of their eigenvalues. This means that the mode which requires the least energy for the original shape to deform, geometrically speaking the smoothest shape, can be ranked first. This is exactly what makes this particular basis interesting compared to any other orthogonal basis.

In the context of free-form modelling, a single degree of freedom (DOF) design approach has been chosen for simplicity and to avoid in-plane vibrational modes. This means that the eigenvectors represent displacements along the vertex normals. The first 8 eigenvectors resulting from an eigen-decomposition of the graph Laplacian for a flat square mesh are shown in Figure 2 as modes of vibration. Note that there exist as many modes as there are vertices in the mesh.

The key concept behind modelling with harmonics is that any n-dimensional vector (representing the normal displacements) can be constructed as a linear combination of the computed eigenvector basis. The coefficient (referred to as the *weight*) related to each eigenvector determines the proportion of that harmonic in the final shape. A low parametrisation of the mesh is obtained when a number of the higher order frequencies are excluded, in order to generate smoother shapes. Therefore, the number of design variables is vastly reduced from one value (displacement) per node, to one value (weight) per included mode.

2.2. Boundaries

It is essential that control over the boundary of the mesh is available, in order for this method to be useful in an architectural engineering context. Due to the single degree of freedom design approach it is only possible to control whether a vertex is fixed or not, hence no sliding boundary supports can be modelled. To fix vertices, the Laplacian matrix must be manipulated in such a way that this control is acquired. By setting the diagonal matrix elements related to the fixed vertices, to a relatively large number, the connectivity information is maintained (therefore preserving smoothness). The large ratio between the artificial vertex stiffness and the other values in the matrix ensures that the modes associated with a movement of those vertices appear last in the sorted eigenvector list, since they would require most energy to deform. Studies by the authors have shown that setting the diagonal matrix entry of fixed nodes to a value around 100,000 gives a good compromise between maintaining smoothness and suppressing movement for typical architectural problems. This means that these modes still exist, but are excluded when only the first k modes are used, where k is defined as:

$$k \leq n - c \tag{5}$$

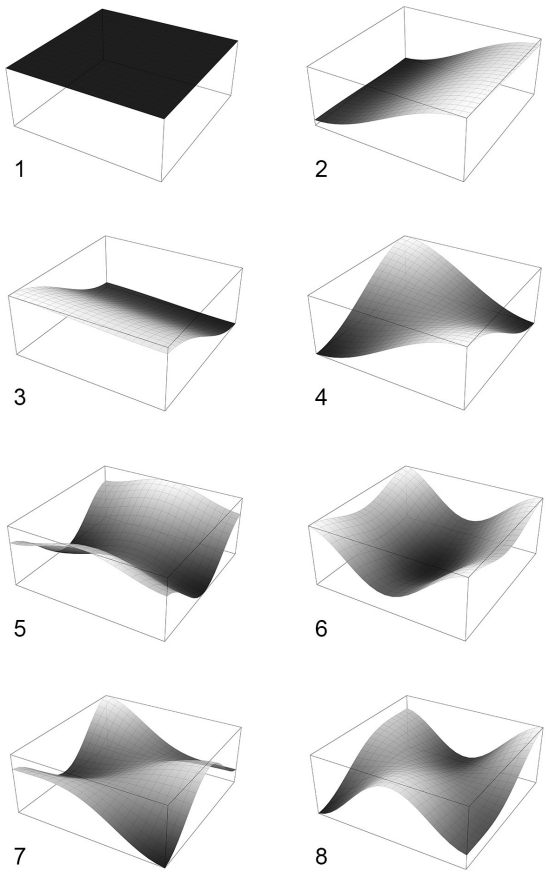

Figure 2: The first 8 mode shapes of a flat square mesh

Here n is the number of vertices and c is the number of imposed constraints. The first 8 modes of the same flat square mesh as Figure 2, but with a fixed perimeter, are shown in Figure 3, demonstrating the desired level of control over surface boundaries.

2.3. Target approximation

It is generally difficult to manually identify the right modes, and assign each the correct weight, to match a certain spatial intent, since designers are not (yet!) familiar with this interactive mode of design exploration. Thus, given a particular starting mesh, it is important to be able to back-calculate the right modes, and their corresponding weights, in order to best represent a particular target surface.

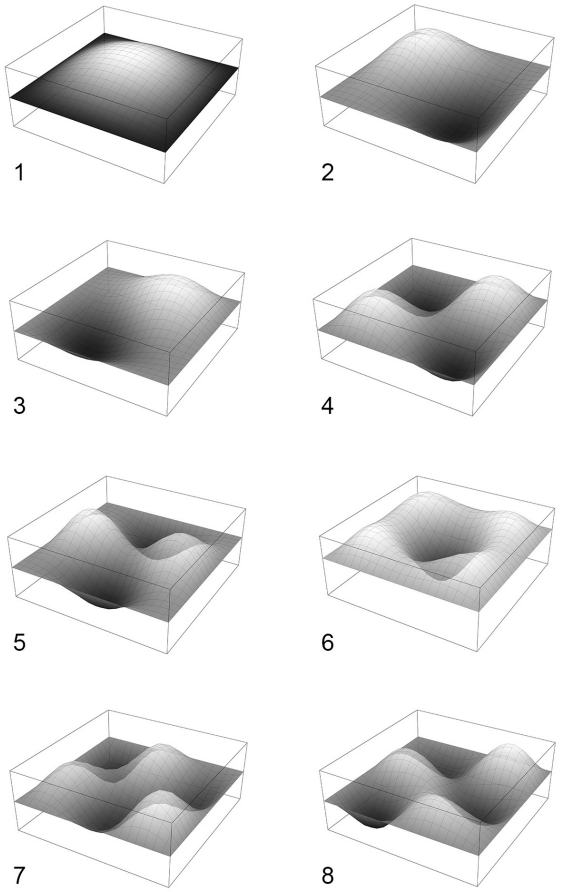

Figure 3: The first 8 mode shapes of a flat square mesh with the perimeter fixed

By expressing the target surface as a vector consisting of n normal-distance values in relation to the n vertices of an initial mesh, and projecting this vector onto each of the eigenvectors using the dot product, it is possible to calculate the weights corresponding to each mode (Rong et al., 2008).

$$\widetilde{X}_i = X \cdot \overrightarrow{e}_i \tag{6}$$

Where $\overrightarrow{e}_i$ is the i^{th} eigenvector, X is the constructed vector of normal distance values and $\widetilde{X}_i$ is the weight corresponding to the i^{th} eigenvector. This weight indicates the proportion of the specific mode that exists in the given target surface.

If all the modes with their corresponding weights are included, a perfect representation is obtained. However, it is desirable to find a suitable cut-off limit (choice of k in Equation (5)) to reduce the number of design variables. By

sorting the modes according to their corresponding weights (absolute values) in descending order, such that the modes with the highest weights are listed first, it is possible to extract useful information about the primary components of the target surface. A pareto-front style plot of how the RMS value of the normal-distances between the generated surface and the target surface reduces as the number of included modes increases, can help to identify this trade-off and select the number of modes to include. In this context, boundary conditions are useful because they help to achieve a better approximation of the target surface with fewer variables.

The method is demonstrated for a simple case where a planar starting geometry (a circle with a radial grid) with a fixed boundary attempts to approximate a cone-shaped target surface, as shown in Figure 4. The corresponding pareto-front style plot (Figure 5) shows how the RMS error rapidly decreases by the time the four most significant modes have been introduced. This suggests a natural cut-off limit to achieve a good approximation with a reduced number of variables. The plot furthermore shows that the best approximation is obtained by including the 10 most significant modes. By this point the RMS error is sufficently small and more modes do not improve the accuracy within 6 significant digits.

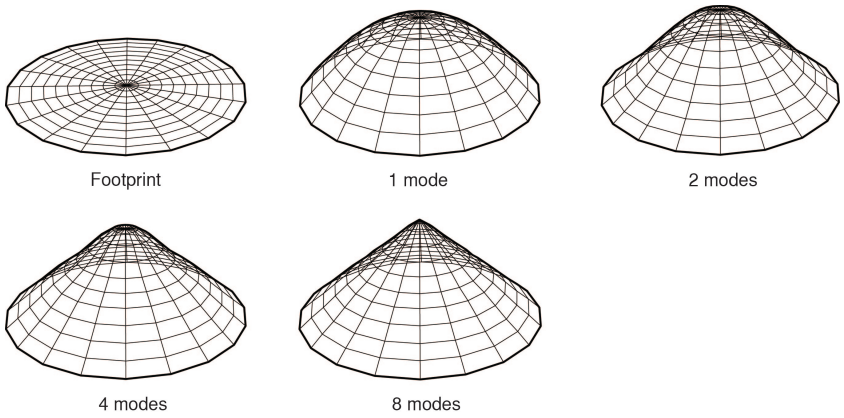

Figure 4: Approximation result for a cone

It may seem counter-productive to go to all this effort to generate an approximation of the target surface, when the target surface itself is already available. However, since the harmonic modelling process outlined here is intended for use at the early conceptual design stage, where experimentation is still in progress and no final shape yet exists, this method can identify new inspirational forms emanating from the initial target, which may otherwise have been overlooked. In the case above for example, a designer might note the surprising aesthetic qualities of the 1, 2 or 4-mode surfaces, and decide to explore them instead of the original cone. Using the harmonic modelling system new shapes can be

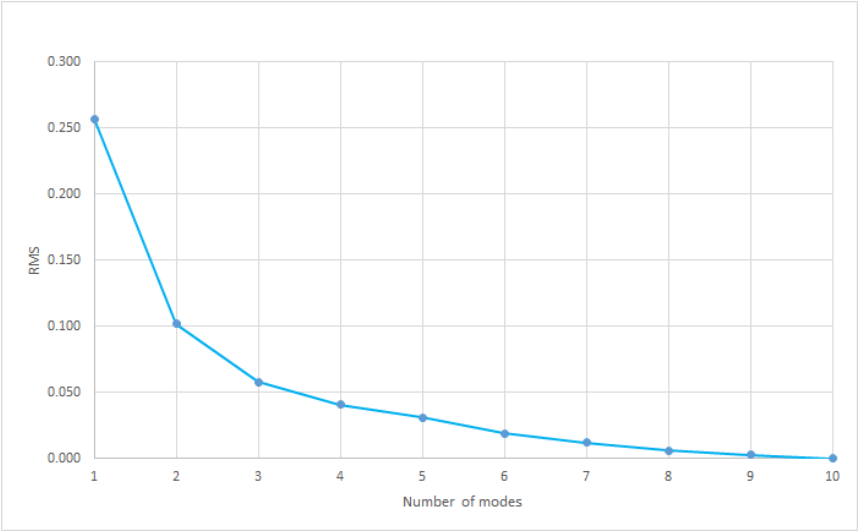

Figure 5: The RMS error value as a function of the number of included modes to approximate a cone

explored quickly within a more restricted design space.

Similarly, it is not immediately obvious why one might want to approximate the target surface with a different starting mesh, instead of just converting the surface into a mesh and using that directly, maintaining its spatial configuration. However, this approximation approach has the advantage of making shape analysis an integral part of the modelling process itself, where high-frequency 'noise' can be removed to create smoother geometries. Decomposing the target into its harmonic components *in relation to a simple known geometry* gives an enriched geometrical understanding of the shape and, perhaps more importantly, could also help simplify the construction process and/or eliminate structural stress concentrations.

2.4. Application in practice

In practice, a mixed approach between arbitrary harmonic modelling and target approximation is most likely to be of interest. This work flow is characterised by the following steps:

1. Model a NURBS surface to represent the spatial intent.
2. Create a simple mesh of desirable density to be used as basis for the surface approximation.
3. Impose boundary constraints (optional), choose the type of the Laplacian matrix, compute the Eigen-decomposition, and back-calculate the, for example 10, most significant modes and their weights to approximate the target surface.

4. Display these modes separately to achieve a better understanding of how the surface is composed. Optionally remove higher frequency (noisy) modes. Interactively adjust the weights to explore alternative design options related to the design space suggested by the original target surface.
5. Look for inspiration among the other less significant modes and extract interesting new modes to introduce.

3. Buckling evaluation

Buckling is an instability problem of a structure in compression, which leads to sudden failure before the ultimate compressive stress of the material is reached. It is characterized by a significant change in displacement resulting from a small load increment, which occurs at the so-called bifurcation point. Since shell structures carry load mainly through compression and have minimal capacity to resist out-of-plane deformation, they are often particularly vulnerable to buckling failure, the sudden nature of which, often without warning, makes it particularly dangerous.

Therefore, an additional benefit of using the design method outlined here is that it is relatively easy to evaluate the buckling capacity of the shapes generated by the harmonic approach. The buckling capacity is usually quantified by means of a buckling load factor (BLF), which specifies the multiple of the applied load that could be supported by the structure before buckling would occur. As it is intended that this method guides shell design during the conceptual stage, it is essential that the evaluation of buckling is an integrated part of the modelling environment (in this case Grasshopper), to enable live feedback and help steer design decisions towards solutions insusceptible to buckling.

There are many methods available to perform buckling analysis for shells, the simplest of which would be to perform a linear analysis using finite element software. This approach is based on the original undeformed configuration, from which the stress state resulting from the applied load is computed, and used to construct a geometric stiffness matrix. The purpose of this matrix is to either increase the conventional stiffness matrix if the structure is in tension, or decrease it if the structure is in compression. The buckling problem is subsequently solved as an Eigenvalue problem. Since the harmonic analysis carried out in the parametrisation stage involves the calculation of Eigenvalues and Eigenvectors, re-using this information to assess buckling would seem particularly efficient. This linear approach would not take the deformed shape into account, which might alter the force distribution in a manner other than purely scaling the initial stress state (Cook, 1995). As a consequence, linear buckling analysis often overestimates the load-carrying capacity of shells, and provides unconservative results. However, this limitation was evaluated by the authors and was found to be acceptable for use in the conceptual design stage, where only qualitative results are needed. Karamba3D (2018) would therefore have been an ideal candidate to calculate the buckling load factor, since it is part of the Grasshopper framework. Unfortunately, at the time this research was

conducted, the Karamba3D consistently crashed when integrated within an automated optimisation work-flow with more than 5 design variables (weights to control the shape). Another candidate for buckling analysis could be isogeometric analysis, for example based on the work of Leonetti et al. (2018). However, at the time this research was conducted, there was no reliable isogeometric analysis engine available within the Grasshopper framework.

In order to demonstrate how buckling considerations can be included as part of the design exploration using this frameowrk, the authors therefore calculated the buckling load factors by performing a non-linear analysis using a physics-based approach. Non-linear analysis in this context is complicated, since the solution has to incorporate information about its actual geometry, which is not fully known until the solution is found. The procedure therefore obtains a solution iteratively, using multiple, linear load-steps, and traces the displacements as a function of load. Various numerical techniques exist to ensure that the displacement stays as close to the correct (in reality unknown) value as possible at each step. The most suitable technique in this case is dynamic relaxation (Cook, 1995) and Kangaroo2 is a physics-engine system for the Grasshopper framework that employs a position-based dynamic relaxation approach (Piker, 2015) with excellent stability, remarkable convergence speed, and includes customisation functionality via scripting. For these reasons it was chosen by the authors to implement a buckling evaluation scheme based on dynamic relaxation, where integration with the developed harmonic modelling framework and computational speed (possibly at the cost of accuracy) were the main priorities.

3.1. Live simulation

Kangaroo2 works by defining a number of goals, which specify the directions and magnitudes of forces acting on the predefined geometry (e.g. a mesh). These forces are summed at each vertex as a weighted average, and the resultant vector dictates the vertex displacement. This is an iterative process, which continues until the residual force at each vertex is zero (within a small tolerance).

The buckling simulation follows the work-flow as outlined in Figure 6. Each edge in the original triangulated mesh is converted into a spring goal by specifying the start and end vertex it spans between, its rest length equal to its current length, and a stiffness value. The latter is divided by the rest length to mimic elastic material behavior, where stress is proportional to strain. Fixed vertices are converted into anchor goals by assigning a very high stiffness value of a zero-length spring to each of them, which connects the vertex with a target particle of infinite mass. This corresponds to a pinned support (all translational degrees of freedom fixed) in a finite element analysis. To simulate bending behaviour, each edge with two adjacent triangular faces is converted into a hinge goal, which applies out-of-plane forces to the four vertices, in an attempt to maintain the angle between the triangular faces at its initial value in the undeformed mesh. The authors have determined that a hinge strength of 1/10 of the average spring strength (stiffness divided by rest length) works well in problems of this sort, to ensure that the structure mainly resists the load by axial force rather than bending, as is usually desirable for material efficiency.

This framework approximates the structural behaviour of a shell with uniform thickness, but it is important to note that since the strengths of the goals are defined with arbitrary values chosen to aid convergence, it is only their relative strength that is of importance. This means that the absolute values of the calculated buckling load factors are not relevant, and cannot be directly compared to the results from a finite element program. However, the change of buckling load factor in relation to the shell geometry, is indeed of relevance.

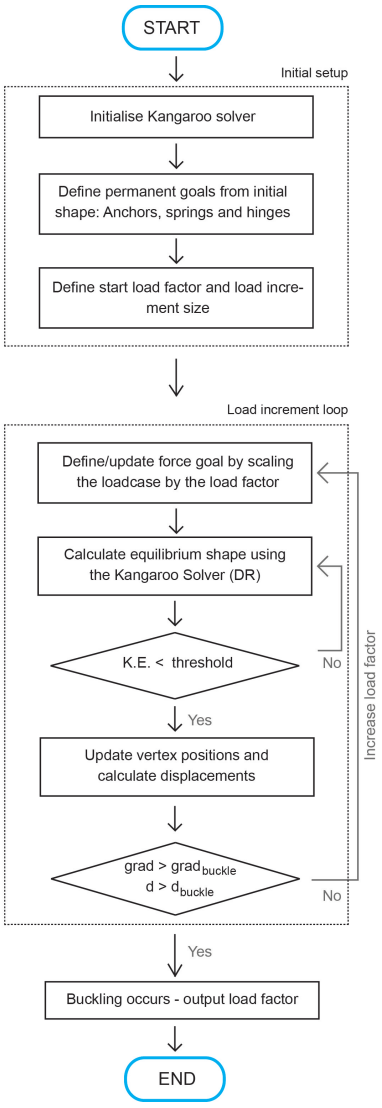

Figure 6: Non-linear buckling flow diagram. K.E. is an abbreviation for kinetic energy.

The load case considered in this study is the self-weight of the structure, calculated as a lumped force on each vertex, represented by a downward vertical vector with a magnitude corresponding to the area of its associated voronoi cell. An iterative loop is then initialised, where the original load case is multiplied by a load factor and converted to force goals. The load factor increases at each iteration, based on a start value and a step size. The structure starts to move when the first load step is applied and converges when the spring and hinge forces equilibrate the load. The displacements are calculated and evaluated against a defined buckling criteria. The loop is repeated until buckling has been identified and the buckling load factor is output.

Automatically identifying when buckling occurs is not straight forward. It requires a formulation of criteria that can detect a sudden change in vertex displacements during the load increments. A number of different criteria have been considered by the authors for this purpose, including individual vertex displacements, an RMS value, or a maximum displacement value. The first measurement enables local buckling modes to be detected, but is computationally slower and requires more data to be stored. The other two measures are more global, and therefore might ignore local effects, but are faster to compute. Since computational speed is a high priority for interactive design exploration, and the main interest with regard to curvature-stiffening of shells is on a global scale, the authors have used the maximum displacement in relation to the undeformed geometry as a proxy for buckling. During the load increments, the load-displacement curve can be traced, and the buckling criteria has been defined as a combination of the gradient of this curve and a limiting displacement value (in case the structure exhibits a very ductile behaviour) as shown in Figure 7. More specifically, buckling is deemed to occur if the gradient corresponding to the current equilibrium configuration is 0.5 times larger than the previous configuration, or the maximum displacement is larger than 1.0m. Note that for snap-through buckling behaviour, it is generally recommended to incrementally prescribe the displacement and trace the corresponding reaction forces, in order to capture the post-buckling behaviour (Cook, 1995). In that case, the buckling criteria would simply be a zero gradient value. However, it is difficult to know which vertex to use to prescribe the displacement when many different shell shapes are investigated at the same time. Additionally, it would require a modification of the dynamic relaxation approach, which calculates unknown displacements from known loads, and not the other way around. For these two reasons, this strategy has not been adopted here.

A buckling analysis of a half-sphere shell with 20m diameter is shown in Figure 8. It shows promising results in terms of computational speed (only 2.1 seconds for this example using a mesh with 353 nodes on a computer with the following specification: Intel(R) Core(TM) i7-455U CPU @ 2.10 GHz, 8.00 GB RAM, 211 GB SSD hard drive) and the shape just before buckling is consistent with structural intuition. Further validation and a comparison with FE results obtained using Autodesk Robot Structural Analysis have also been performed, but such detail is outside the scope of this paper and interested readers are pointed to the work of Brandt-Olsen (2015).

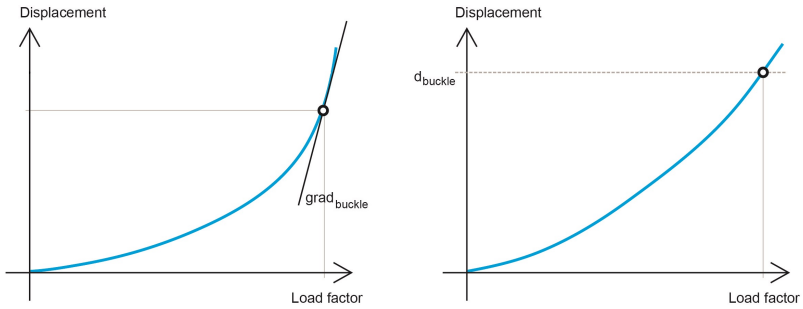

Figure 7: Combined gradient (left) and displacement (right) buckling criteria

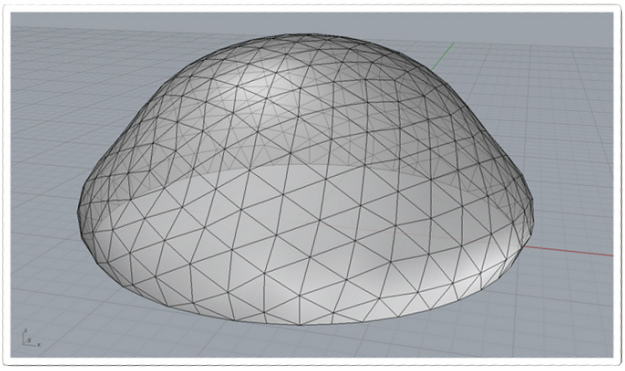

Figure 8: Buckling simulation of a half sphere

4. Case study

The harmonic modelling technique described above, combined with the live buckling evaluation, has been applied to a case study of the British Museum Great Court Roof to demonstrate its potential in an architectural and structural engineering context.

The aim of this case study was threefold: Firstly, to decompose the surface into its harmonic components to gain a better understanding of its underlying shape. Secondly, to investigate the trade-off between the number of included modes and the accuracy of the surface approximation, whilst simultaneously evaluating the effect of each mode with regard to the buckling capacity. Finally, to explore different variations of the shape by modifying the modal weights, based on a combination of considerations involving pure aesthetics and a buckling optimisation.

4.1. Shape analysis

A mesh (1806 vertices and 3372 faces), similar to the one used for the roof of the British Museum Great Court before it was relaxed over the mathematically defined surface (Williams and Shepherd, 2010), was used as the basis for the harmonic modelling approach (see bottom of Figure 9). A NURBS patch, built through the vertices of the British Museum Roof and intentionally extended beyond the perimeter, functioned as the target surface, as shown in the top of Figure 9.

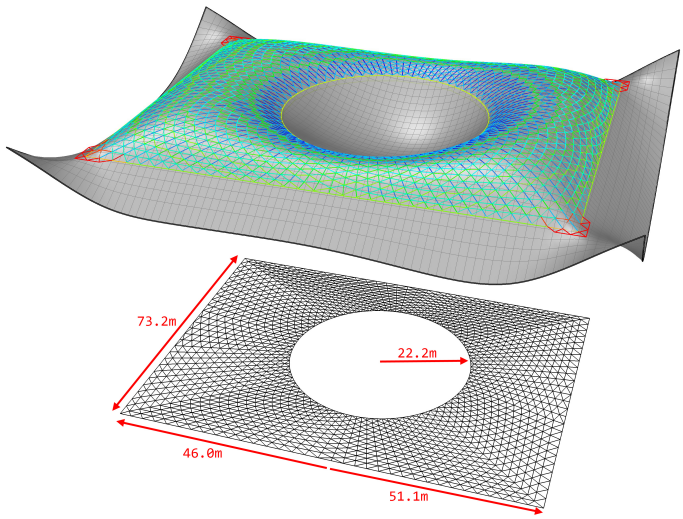

Figure 9: Surface patch approximation and initial mesh

From a construction of the graph Laplacian, followed by an Eigen-decomposition, the ten most significant modes, and their weights to approximate the NURBS surface patch, were back-calculated, as visualised in Figure 10. The approximation result was smooth, had an RMS error value of 0.14m and a maximum deviation from the target surface in the normal direction of 0.72m, which was promising when compared to the size of the structure (72m x 97m) and the fact that only ten modes were used.

4.2. Comparison with subdivision surface modelling

To demonstrate the advantages of this harmonic form-finding approach over more traditional methods of surface modelling, a subdivision surface was used to recreate the British Museum geometry. Using the same plan dimensions (as shown in the bottom of Figure 9), a control polygon was constructed as shown in the left of Figure 11. By fixing all vertices that lie on a boundary, only allowing the other vertices to move vertically, and by taking advantage of the one plane of symmetry, the number of parameters required to define the subdivision

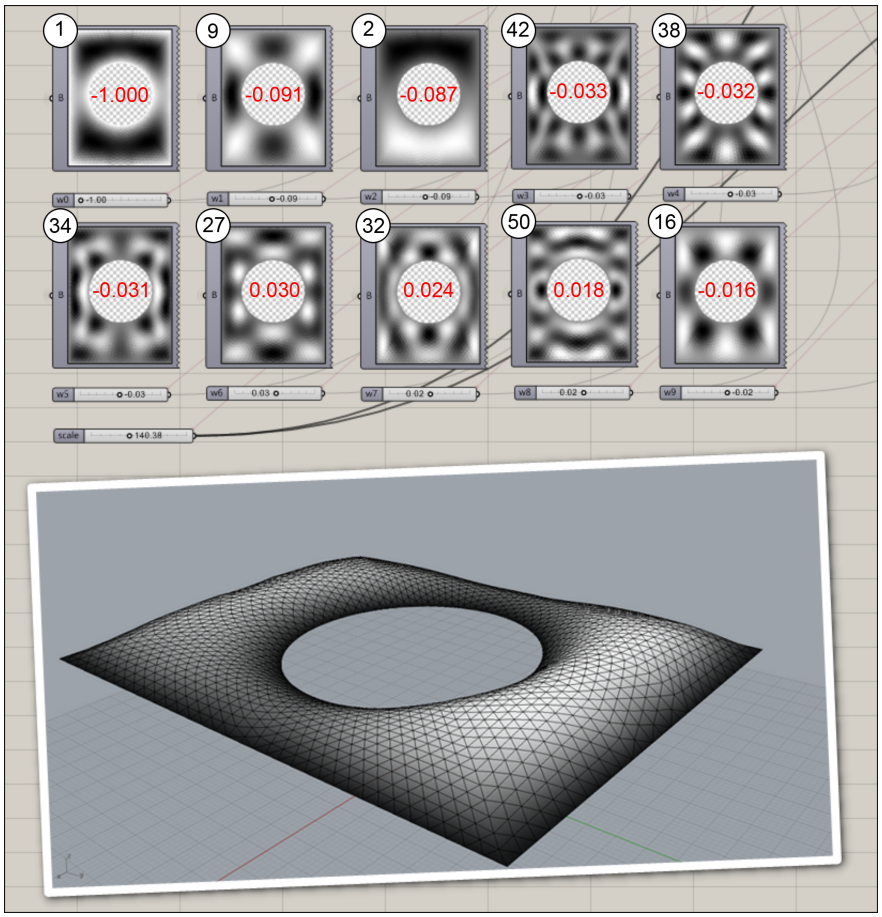

Figure 10: The ten most significant harmonic modes of the British Museum. The circles at the top left of each icon show the mode number (with one being the mode with lowest Eigenvalue). The red numbers in the centre of each icon give the weighting attributed to that mode.

surface was reduced to twelve. This allowed for a fair comparison against the harmonic modelling carried out in the previous section, which required only ten parameters (modal weightings).

The "Galapagos" genetic algorithm component built-in to Grasshopper was used to try to match the subdivision surface to the initial target NURBS patch. The vertical height above the reference plane of each of the twelve control points was encoded as the genome. The objective function was taken as the RMS distance to the target surface of the Level #2 subdivided vertices (shown on the right of Figure 11). The genetic algorithm used a population size of 50 individuals per generation with 5% elitism. After 50 minutes of computation, 55 generations had been evaluated and the solution converged to that shown in

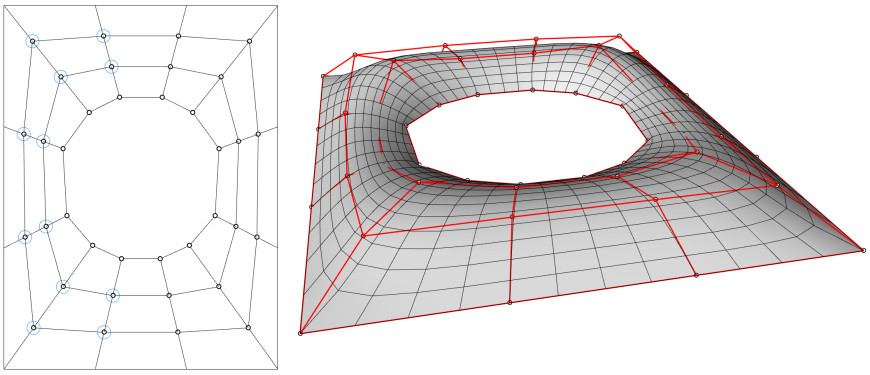

Figure 11: Subdivision surface comparison, showing initial control polygon with blue circles identifying movable vertices (left), and the resulting matched surface (right) with control polygon in red.

Figure 11, which has an RMS error value of 0.27m and a maximum deviation of 1.86m. The near-instantaneous target-fitting using the harmonic form-finding method described in the previous section had comparable RMS and maximum errors of 0.14m and 0.72m respectively, with negligible computational overhead. Thus the harmonic approach was able to approximate the target surface more quickly and more accurately, despite having fewer degrees of freedom.

4.3. Shape refinements

Since the structural performance of a shell can be noticeably affected by only small changes to the geometry, it is possible to refine the shape with regard to its buckling capacity by studying the effect of each mode. Some modes may only be of aesthetic character, whilst others improve the buckling capacity. Thus, a plot of the buckling load factor (BLF) as a function of the number of modes included in the shell definition is shown in Figure 12. Note that pinned supports were used along the boundary for the buckling analysis, which differs from the design constraints of the British Museum itself.

It can be seen that, with the modes sorted in order of decreasing weight, if the first four modes are used to define the shell geometry, the BLF is around 3.75, but this increases to 5.00 when the first eight modes are included. Figure 12 also shows that the addition of the third most significant mode (Mode #2 as identified in the top middle of Figure 10) had a negative effect on the buckling capacity. This observation seems sensible, since the failure mode was a collapse of the dome-like part with the longest span, and Mode #2 decreased the height in this zone. By omitting Mode #2 from the resulting shape (changing its weight from -0.087 to zero), a buckling load factor of 5.85 was obtained, while the RMS error value and the maximum deviation from the target surface were 0.32m and 0.92m respectively. These values have to be compared with a buckling load factor of 5.50, RMS error of 0.14m and maximum deviation of 0.72m when all

ten modes were included (as per Figure 12). So whilst the buckling load factor has been increased, the deviation from the target surface has also increased. This therefore suggests a design compromise, but one which can easily be resolved by visually inspecting the proposed new surface as seen in Figure 13 and comparing the height values against architectural requirements.

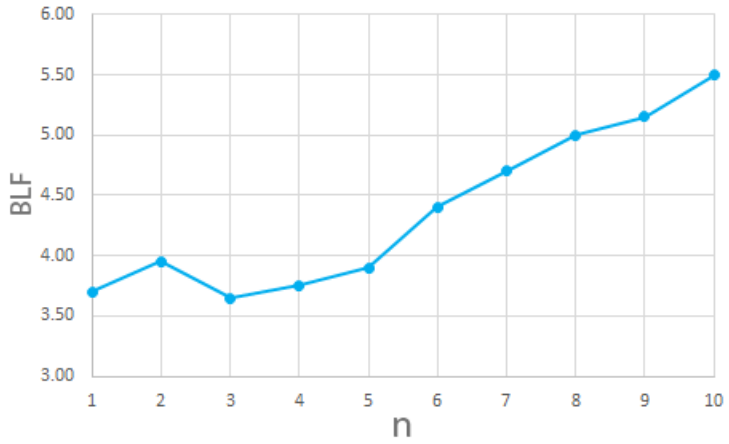

Figure 12: Buckling capacity when 'n' lowest Eigenvalue modes are included

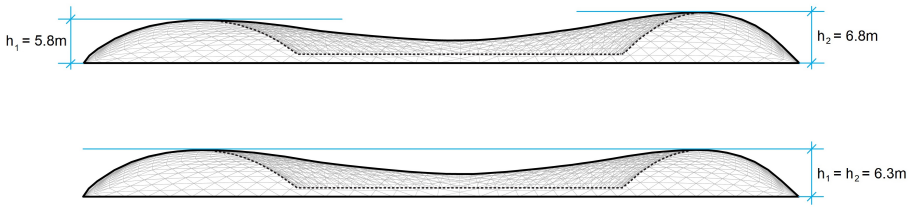

Figure 13: Resulting shape from the first 10 modes with (top) and without (bottom) Mode #3

The decomposition of the target surface into its harmonic components has therefore helped to understand the influence of each mode with regard to structural performance. By suppressing a mode detrimental to the buckling performance, an improved shape with a more robust structural load-carrying capacity has been found with little visible change in shape.

4.4. Shape variations

With the ten short-listed modal components for the British Museum Great Court Roof, it was straight-forward to explore a variety of other design possibilities emanating from the original shape by modifying the weights of these

modes and introducing new interesting modes as well. Figure 14 shows two alternative Great Court Roofs generated using this method, chosen based purely on aesthetics. This demonstrates that this is a useful approach to inspire and drive exploration at the conceptual design stage.

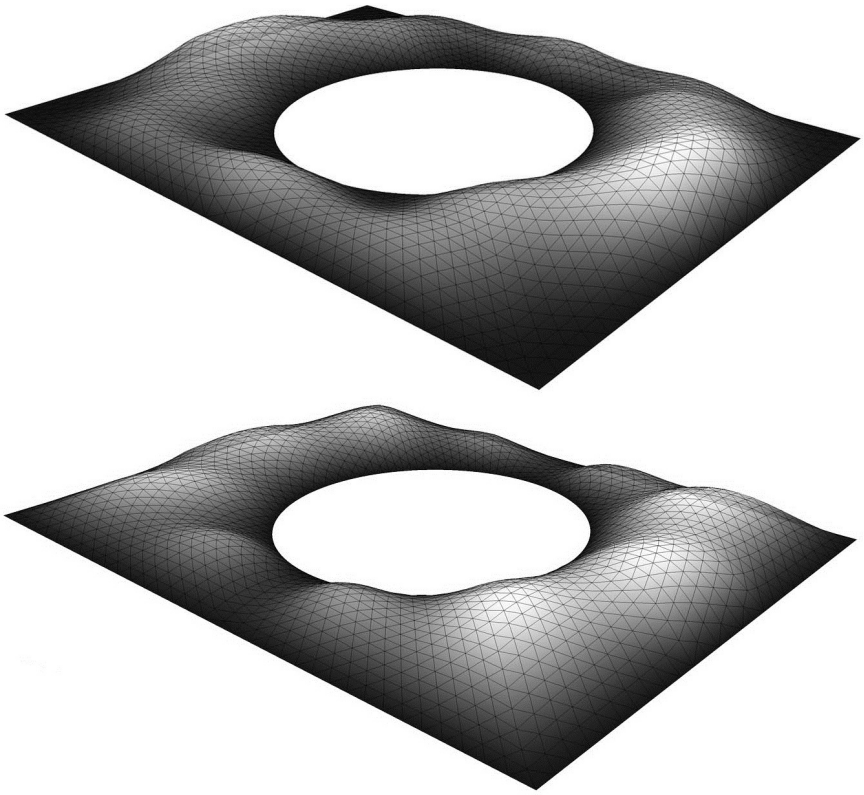

Figure 14: Shape variations based on aesthetics

Automatically adjusting the weights of the mode shapes based on an evaluation of the buckling capacity as part of an optimisation process has also been explored. The efficiency of this work-flow was supported by the low-dimensional parametrisation of the geometry of the mesh and the computational speed of the non-linear Kangaroo2 buckling analysis. Thus, the optimisation variables became the weights of the ten most significant modes, and the objective function was to maximise the buckling load factor. The built-in "Galapagos" component in Grasshopper was selected for this purpose, using a simulated annealing search strategy (Carr, 2018). In order to compare the results objectively, it was essential to ensure that the generated shapes shared the same increase in area from the target surface, such that the best shape reflected geometric stiffness,

rather than stiffness gained by using more material.

Instead of discarding every shell that did not meet the specific area requirement (the majority of the generated shapes and therefore not a very efficient strategy), an algorithm was developed to adjust the weights such that each solution was a potential candidate. The algorithm calculated the area of a newly generated mesh and compared it with the area of a target mesh. Based on the difference, it determined whether the weights had to be decreased or increased, in order to be within a specified percentage of the target area. The weights were subsequently changed in small steps, defined as a factor of the difference between the weights of the target and the weights of the newly generated mesh (using linear interpolation). The area of the adjusted mesh was evaluated after each step, and a bisection strategy was implemented to adjust the size of the last step, which would otherwise cause the area requirement to be exceeded.

The optimal weights are shown in Table 1 and the resulting roof geometry is shown in Figure 15, where it is evident that material has been distributed towards the zones with larger spans, and that the curvature has been amplified. This shell has the same area as the original roof geometry, but with a buckling load factor of 8.55. In comparison, the as-built British Museum Roof has a BLF of 5.50 (referring to Figure 12 with ten modes), which means that the BLF of the optimized geometry is 1.6 times bigger. For comparison, when the subdivision surface model from section 4.2 was re-run to maximise buckling load factor instead of minimising distance to the target surface, a BLF of 5.78 was achieved, better than the as-built surface but much lower than that achieved using the harmonic modelling.

Table 1: Modal Weightings

Mode #	Weight	
	Target Surface	Buckling Optimised
1	-1.000	-0.691
9	-0.091	-0.190
2	-0.087	0.088
42	-0.033	-0.123
38	-0.032	-0.157
27	0.030	0.022
32	0.024	0.081
50	0.018	0.098
16	-0.016	0.070

5. Conclusions

This paper has presented a free-form modelling strategy based on harmonics, with a direct link to Fourier analysis, to achieve a low-dimensional parametrisation of the design space. A number of initiatives, including a single degree of freedom design approach, control over boundary constraints, visualisation aids and guidance towards specific spatial configurations have also been introduced

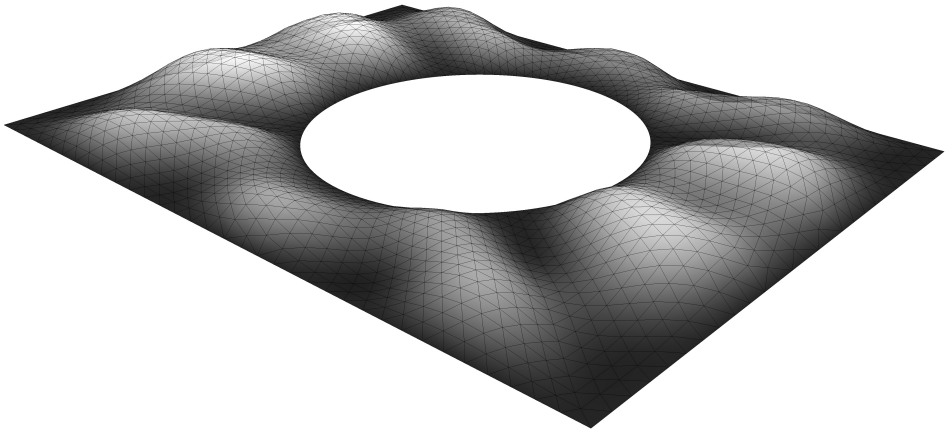

Figure 15: Alternative British Museum Roof, based on an optimisation work flow

to adapt the method to an architectural and civil engineering design context. The approach models real, physical shell structures in a specific parametric design workflow commonly used in building design. In this context, the reduction of the design space to a few parameters is particularly important. WHilst other methods are available to achieve this, using shell harmonics as the basis for dimension reduction brings additional benefits. By removing high frequency "noise", the physical construction of the shell structure can be facilitated though the repetition and rationalisation of building elements. And by assessing the contribution of each mode to the structural performance of the resulting geometry, buckling-inducing low frequency modes can also be removed. This synergy between using Eigenmodes as the basis for design exploration/optimisation, and also for structural performance analysis/optimisation, is particularly elegant. The work reported here is encapsulated within a flexible, free-form modelling tool, the Grasshopper implementation of which has been open-sourced by the authors, and can be downloaded free from Github (Brandt-Olsen, 2018a).

A bespoke script for Kangaroo2 to simulate shell buckling has also been developed as a means of evaluating the buckling response of the generated shapes. The inherent doubly-curved nature of these shapes has been proven to be a useful means of stiffening the shells against buckling failure. A case study of the British Museum Great Court Roof demonstrated this potential, and it was possible to increase the buckling capacity by redistributing material to provide more curvature (for the same surface area). Thus, harmonic form-finding can inform the design of shells at the conceptual stage, and encourages an interaction between the architect and the structural engineer. The authors have open-sourced the non-linear buckling Grasshopper component used for this study, and it can also be downloaded free from Github (Brandt-Olsen, 2018b).

5.1. Discussion

One of the main disadvantages of modelling with harmonics is the lack of tangible spatial control. The numerical parameters are more abstract to work with and it is generally hard to predict the result of adding multiple mode shapes together. However, this is mitigated by the ability to analyse the shape during modelling. This approach is therefore unlikely to replace the two most common free-form modelling strategies (NURBS and Subdivision Surfaces) directly. However, it certainly has great potential to broaden design efficiency when used alongside them in combination.

It is also essential to note that the harmonic modelling approach and the buckling evaluation can be seen as separate parts of a new design work flow. Any other modelling strategy or structural evaluation can replace one of the parts. However, using harmonic shapes and curvature as means of stiffening shells nicely integrates the two approaches.

5.2. Future work

The advantages of the area-weighted cotangent Laplacian was not as significant as expected in the context of free-form modelling, and it sometimes causes undesirable non-smooth transitions near fixed boundaries. The authors aim to investigate this issue further by considering other linear operators (e.g. the Biharmonic operator), to see whether any behavioural differences occur. The Biharmonic operator is especially interesting, since it mimics bending by averaging curvature, and thus takes the 2-ring vertex neighbours into account. This would allow control of tangency at the boundaries, although the single degree of freedom approach would probably need to be sacrificed as a result.

The buckling simulation using Kangaroo2 also needs further calibration against real structural properties in order to determine a more realistic ratio between membrane and bending action. The membrane action would ideally also be modelled using constant-strain triangle elements to more accurately simulate the in-plane stiffness, rather than converting the edges into springs. This calibrated modelling with Kangaroo2 has shown much promise (Brandt-Olsen, 2018b).

References

Brandt-Olsen, C., 2015. Harmonic form-finding for the design of curvature-stiffened shells. Master of philosophy in digital architectonics. University of Bath.

Brandt-Olsen, C., 2018a. Harmonics. URL: `www.github.com/CecilieBrandt/harmonics`. Available online [Accessed 22 November 2018].

Brandt-Olsen, C., 2018b. K2Engineering. URL: `www.github.com/CecilieBrandt/K2Engineering`. Available online [Accessed 22 November 2018].

Carr, R., 2018. Simulated Annealing. URL: www.mathworld.wolfram.com/SimulatedAnnealing.html. Available online [Accessed 22 November 2018].

Cook, R.D., 1995. Finite element modeling for stress analysis. John Wiley & Sons, Inc.

Dong, S., Bremer, P.T., Garland, M., Pascucci, V., Hart, J.C., 2006. Spectral surface quadrangulation. ACM Transactions on Graphics 25, 1057–1066. doi:10.1145/1141911.1141993.

Hildebrandt, K., Schulz, C., Von Tycowicz, C., Polthier, K., 2011. Interactive surface modeling using modal analysis. ACM Transactions on Graphics 30, 1–11. doi:10.1145/2019627.2019638.

Karamba3D, 2018. Karamba. URL: www.karamba3d.com. Available online [Accessed 22 November 2018].

Leonetti, L., Magisano, D., Liguori, F., Garcea, G., 2018. An isogeometric formulation of the Koiter's theory for buckling and initial post-buckling analysis of composite shells. Computer Methods in Applied Mechanics and Engineering , 387–410doi:10.1016/j.cma.2018.03.037.

Malek, S.R., 2012. The Effect of Geometry and Topology on the Mechanics of Grid Shells. Doctor of philosophy in the field of structures and materials. Massachusetts Institute of Technology.

McNeel, 2018. Grasshopper - Algorithmic Modeling for Rhino. URL: www.grasshopper3d.com. Available online [Accessed 22 November 2018].

Meyer, M., Desbrun, M., Schr, P., Barr, A.H., 2003. Discrete Differential Geometry Operators for Triangulated 2 Manifolds. Visualization and Mathematics III , 35–57doi:10.1007/978-3-662-05105-4_2.

Michalatos, P., Kaijima, S., 2014. Shell structures for architecture: Eigenshells. Routledge.

Piker, D., 2015. White paper: Goal-Driven Design. Technical Report. Robert McNeel & Associates.

Pinkall, U., Polthier, K., 1993. Computing Discrete Minimal Surfaces and Their Conjugates. Experimental Mathematics 2, 15–36. doi:10.1080/10586458.1993.10504266.

Rong, G., Cao, Y., Guo, X., 2008. Spectral mesh deformation. The Visual Computer 24, 787–796. doi:10.1007/s00371-008-0260-x.

Taubin, G., 1995. A signal processing approach to fair surface design. Proceedings SIGGRAPH 95 , 351–358doi:10.1145/218380.218473.

Vallet, B., Lévy, B., 2008. Spectral geometry processing with manifold harmonics. Computer Graphics Forum 27, 251–260. doi:10.1111/j.1467-8659.2008.01122.x.

Williams, C.J.K., Shepherd, P., 2010. British museum great court, in: Burry, J., Burry, M. (Eds.), The New Mathematics of Architecture. Thames & Hudson, New York, pp. 122–125.

Zhang, H., Van Kaick, O., Dyer, R., 2007. Spectral Methods for Mesh Processing and Analysis. Proceedings of Eurographics State-of-the-art Report , 1–22.