

A comparison of two algorithms for the simulation of bending-active structures

Pierre Cuvilliers¹, Justina R Yang¹, Lancelot Coar²
and Caitlin Mueller¹

International Journal of Space Structures
1–13

© The Author(s) 2018

Reprints and permissions:

sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/0266351118779979

journals.sagepub.com/home/sps



Abstract

Bending-active structures, made from elastically bent materials such as fiberglass rods, offer exciting opportunities in architecture because of their broad formal palette and ease of construction. While they have been relevant since Frei Otto's Mannheim Multihalle (1974), recent computational developments that help simulate active-bending processes have renewed interest in them. Such tools are important because they can replace time-consuming and imprecise physical modeling processes. However, physically meaningful simulations, using real materials and full scale, are difficult to create, and there are no good mechanisms to reveal when a simulation is inaccurate. This article offers a conceptual and numerical study of two popular contemporary algorithms for simulations of bending-active structures, mainly through a comparison of their results on the planar elastica. We then offer guidelines on best practice modeling settings and demonstrate possibilities and pitfalls through an architectural-scale case study.

Keywords

bending-active structures, digital design methods, dynamic relaxation, explicit dynamics, numerical experiments, position-based methods

Introduction

According to shell expert Chris Williams,¹ form-active structures are a broad category of structures that react to external loads and constraints with large deformations, in contrast to more common form-passive structures such as frames or rigid shells. These structures are lightweight, since they do not resist forces by direct material rigidity but by geometric deformations, and resilient because they can flex instead of breaking.

One subset of interest is the bending-active structures, where structural elements are linear rod elements or elongated plates that deform in bending, often combined with fabric membranes to create lightweight structures. The construction process of these structures is fast and accommodating to large tolerances, making them a great solution to temporary covering problems, and to problems where adaptivity and reconfigurability are required.² Some examples are shown in Figure 1, like non-regular rod assemblies,³ regular grids of rods,⁴ and hybrid structures.⁵

The computational design process of a bending-active structure always involves finding the equilibrium shape of

the structure. It is crucial that this final shape be predicted accurately, which remains a challenge with available software tools. For example, the final shape could be the support of another element so that errors in the shape would result in incompatibilities with the secondary elements.⁶ In addition, small errors in shape can lead to larger errors in the internal stresses predicted for the structure⁷ and even larger errors in the prediction of non-linear processes such as buckling.⁸

In the simulation process, the topology and connection types of the structure are first defined, along with the length of the bars and the boundary conditions such as anchors in the ground. An iterative algorithm (the specific

¹Massachusetts Institute of Technology, USA

²University of Manitoba, Canada

Corresponding author:

Pierre Cuvilliers, Building Technology Program, Massachusetts Institute of Technology, Room 5-414, Cambridge, MA 02139, USA.

Email: pcuvil@mit.edu



Figure 1. Three examples of bending-active structures. From left to right: bending-active frame, bending-active gridshell with flexible membrane, and hybrid bending-active gridshell integrating a rigid shell. Central picture © Hubert Berberich (CC-BY 3.0).

types available are described in section “Literature review: precedents in simulation of bending-active structures”) then relaxes the rods’ positions until they are at equilibrium.

However, this iterative algorithm is at risk of being too slow for interactive explorative design, unreliable (giving unexpected results), and inaccurate (giving out-of-equilibrium results). It has no guarantee of convergence and can often produce a false sense of accuracy and definiteness. An example of a typical failed design process is represented in Figure 2. An attempt is made to reproduce the shape of a complex bending-active structure made of a two-layer grid. When a solver is run on this problem, it initially seems like it finds an equilibrium. However, running the solver for more iterations—and a longer time—shows that this is not the case. It could be that the solver parameters were not appropriate or that the simulation did not accurately represent the construction process. In any case, the first shapes obtained after smaller runtimes are not to be trusted, as the nodal positions and forces they give are more representative of the initial configuration of the simulation than reality. This raises awareness to the fact that convergence settings need to be carefully selected to produce reliable results. This problem is detailed in section “Larger example: Cocoon.”

Several algorithms and frameworks have tackled this problem, discussed in section “Literature review: precedents in simulation of bending-active structures.” However, these algorithms have not yet been comparatively evaluated specifically for the modeling of bending-active structures. As a result, building a design tool for one structure requires a lot of trial and error across the different possibilities. In addition, the stopping criterion—determining for how long the solver will try to improve its solution—is critical to the quality of the results, yet its setting is often overlooked. In this article, we compare two different algorithms (Dynamic Relaxation¹⁰ in Kangaroo 1,¹¹ and a variant of the projective constraints method¹² in Kangaroo 2¹¹) and their stopping criteria on three benchmark metrics: speed, reliability, and accuracy, and provide guidelines for the parameters to use.

The two software tools selected are widely used in the architectural design community and implement methods that have become standards for form-finding simulations

(section “Literature review: precedents in simulation of bending-active structures”). By comparing their predictions against the analytical result for the planar elastica (section “Methodology: single elastica comparison”), we find guidelines for the simulation setup that produces accurate results without compromising too much on execution time (sections “Kangaroo 1” and “Kangaroo 2”). Finally, we use these guidelines for the simulation of a larger structure and show that they lead to reliable results (section “Larger example: Cocoon”).

Literature review: precedents in simulation of bending-active structures

Simulation of elasticas

At the heart of every bending-active structure is the elastica, the mechanical problem describing the shape of one elastic rod spanning between two supports.¹³ Simulating one elastica is a complex problem, for three reasons:

- It is a highly geometrically non-linear problem; small changes in the boundary conditions leading to large changes in the shape and different stability regimes can coexist;¹⁴
- The interactions between bending and torsion are complex, and simulating them involves keeping track of more than 3 degrees of freedom at each discretization point;¹⁵
- Bending and axial compression in the rod operate at very different stiffnesses, leading to ill-conditioned numerical problems when they are modeled simultaneously. Considering non-extensible rods requires more advanced constrained optimization algorithms.¹⁶

Simulation of complex structures

For active-bending structures, elasticas and other elements are then assembled to form a complete structure. The loads are usually light live loads of wind and impact, given the often-temporary nature of the structures. This results in

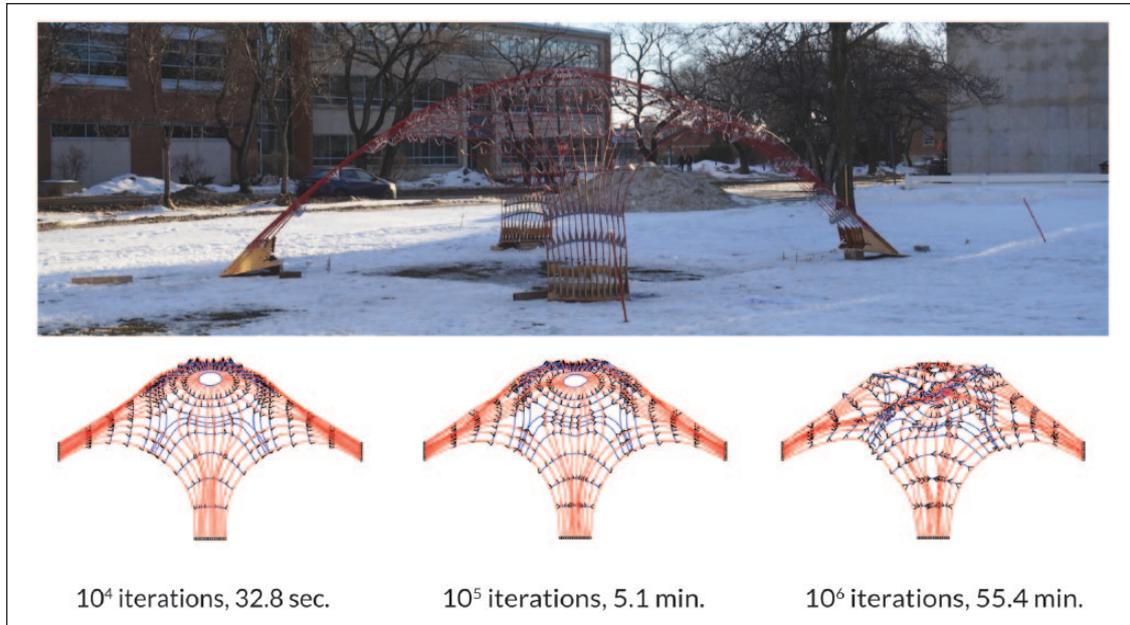


Figure 2. Example of a computational reproduction of the construction process for a bending-active structure.⁹ The runtimes are extracted from one run of the Kangaroo 2 solver, a commonly used architectural design tool, on the structure.

very flexible structures. The assembly and erection processes, in particular, always incorporate large displacements that conventional finite element method (FEM) packages for structural engineering such as Robot or RISA-3D cannot easily represent (although more powerful generalist packages can be used, as described below). As such, these new structural typologies and systems require new approaches for engineering.

The approach we consider here has the formulation of a form-finding problem: given a definition of a bending-active structure, with its rods, connections, boundary conditions, and loads, what is its equilibrium shape? Two solving methods are typically used in this approach:

- Generalist FEM packages such as ABAQUS are available for bending-active structures.¹⁷ The resulting models are often accurate and easily calibrated with physical quantities; however, they tend to be slow by default and do not reliably find the main equilibrium. They are also poorly integrated with the usual architectural design tools.
- Discretized elements in bending have recently attracted considerable attention from the computer graphics community, mainly for the simulation of dynamic systems such as hair.¹⁸ These methods tend to be faster and give predictable results for dynamic problems but are sometimes not accurate enough for the simulation of a static problem and require great care in tuning.¹⁹ Derivatives of these tools have been integrated in architectural design tools, most notably the Shape-Up library.¹² It is not always

clear how each algorithm can be calibrated to give results in meaningful physical units.²⁰

Computational design of bending-active structures

Examples abound for bending-active structures where a digital prototyping tool was critical to the design process. The Centre for Information Technology and Architecture (CITA) at the Danish Royal Academy of Fine Arts and the Complex Modeling project produced several towers made from fiberglass rods in a water-drop shape, stacked and tensed by a tailored-designed membrane,²¹ and elastic gridshells.²² They used specialist tools built on top of Kangaroo (described below) for the design. The Institute of Building Structures and Structural Design (ITKE) at the University of Stuttgart created several examples such as Flectofin[®], a large-size flapping mechanism,²³ and umbrella-shaped bending-active structures,²⁴ using custom-made non-linear FEM procedures. Several recent elastic gridshells also provide interesting examples using design tools based on the Dynamic Relaxation method (described below).^{7,25} Each time, the authors show how only a specialist use of form-finding tools made a more comprehensive design process possible. In addition, all of the examples cited in this section had to incorporate tolerance-correction systems to overcome the accuracy shortcomings of the software. This shows a strong need by designers for accessible tools that can simulate bending-active structures.

Design tool reviews

Reviews of custom-made tools and frameworks for bending-active structures exist. For structures made of membranes and elasticas, Van Mele et al.²⁶ and Ahlquist and Menges²⁷ look at the influence of the quality of the simulation tool on the design process and improve on its speed or ease of use. However, they do not closely consider reliability or accuracy. More comprehensive reviews of modeling and design techniques are also available,^{28,29} detailing how design methods stemming from different generations of design tools created new categories of bending-active structures. These present very detailed analysis of construction systems but do not focus on verifying the accuracy of simulation tools for a range of conditions. A comparison of simulation results to analytical results, however using exclusively a custom-made dynamic relaxation method, is found in the work by Adriaenssens and Barnes.³⁰ Although these studies manage to simulate simple bending-active behavior, they do not address a comparison between simulated and built models, making the accurate fabrication of physical elements from simulated forms an unreliable proposition, showing a need for adequately verified and accessible software.

Bending-active algorithms

There are two algorithms for simulating bending-active structures considered in this article: dynamic relaxation and projective constraint-based solving, respectively, implemented in the Kangaroo 1 and Kangaroo 2 software packages (for Rhinoceros 3D/Grasshopper). These tools are tested in section “Methodology: single elastica comparison.” They are widely used in the architectural design community, with a combined number of downloads of close to 250,000 at present,¹¹ and free. In addition, they both share the same author-developer and programming language (C#), allowing for a comparison of the algorithms, not only their software implementation.

Kangaroo 1¹¹ (this paper uses version 0.099) implements a Dynamic Relaxation solver (DR). DR is a time discretization of the dynamical behavior of physical systems,¹⁰ introduced in the 1960s. Specifically, Kangaroo 1 implements a time integration of Newton’s second law using a semi-implicit Euler method.³¹ At every iteration with time step Δt after time t , the solver computes the nodal velocities $v_i^{t+\Delta t}$ using nodal forces F_i^t from the previous update and masses M_i and then gets positions $x_i^{t+\Delta t}$ using the new velocities

$$\begin{cases} v_i^{t+\Delta t} = v_i^t + \Delta t \frac{F_i^t}{M_i} \\ x_i^{t+\Delta t} = x_i^t + \Delta t v_i^{t+\Delta t} \end{cases}$$

This is a conditionally stable integration scheme with a wide stability region and is symplectic, meaning that it has

a very good energy conservation.³² DR has to add virtual unit masses on the points that do not have one defined so that their dynamic behavior is defined. The time step is chosen to fall in the stability region of the method, given the nodal masses.³¹ The damping of the system is done here with the kinetic damping scheme, resetting the system’s velocities every time the kinetic energy reaches a peak.³³ DR has been used for a wide range of form-active structures,^{34,35} including active-bending structures.^{36,37} These papers also use kinetic damping and use a central-difference time integration scheme that leads to the same update—and properties—as the semi-implicit Euler method albeit considering speeds at a different time

$$\begin{cases} v_i^{t+\Delta t/2} = v_i^{t-\Delta t/2} + \Delta t \frac{F_i^t}{M_i} \\ x_i^{t+\Delta t} = x_i^t + \Delta t v_i^{t+\Delta t/2} \end{cases}$$

This makes of Kangaroo 1 a representative tool for contemporary implementations of the Dynamic Relaxation method in bending-active simulations.

Kangaroo 2¹¹ (this paper uses version 2.1.2) uses projective constraint-based solving, a method developed in the computer graphics community and made available to wider audiences through the general purpose simulation tool Shape-Up.¹² At each iteration, the solver moves closer to the equilibrium position by projecting the positions on the sets of constraints representing the relationships between them. The solution is a physical equilibrium with correct derived forces, if the constraints are physically accurate. The solving process used by Kangaroo 2 is a specialized version of the Shape-Up algorithm: it uses the same projective-based constraints and accelerates movements using a virtual velocity attached to each vertices. This code is not entirely available to the public, slightly obscuring this process. However, simple experiments—observing the movement of a single particle in a singular force field—and explanations by the developer³⁸ give some insight. They show that the algorithm resembles the kinetic damping of dynamic relaxation,³³ resetting the vertices’ velocities when the forces change direction, but using forces derived from the projective constraints method.³⁹

Methodology: single elastica comparison

Before studying the simulation of complex bending-active structures containing many members, we first focus on simulations of simple, single-curve elastica problems, as they represent the elementary problem of all bending-active structures and can be solved exactly for forces and geometry using analytical equations. This allows for a direct comparison between the algorithms’ and analytical results. All algorithms for active bending can represent the

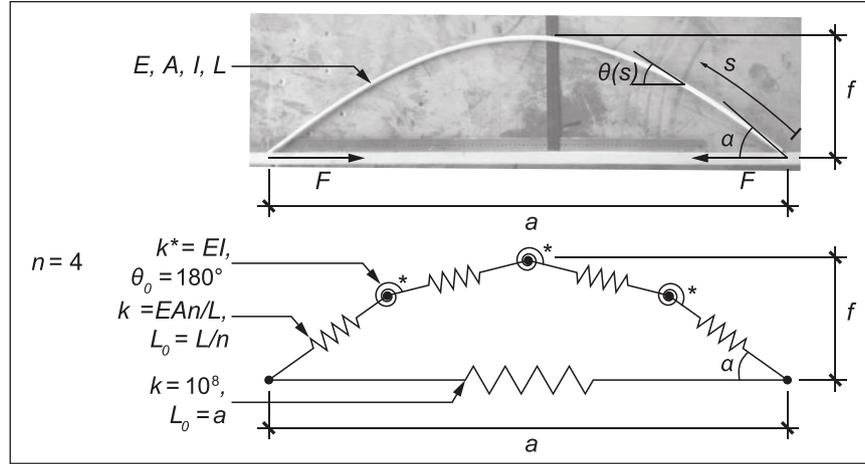


Figure 3. Definition of the planar elastica problem, continuous (top) and after discretization (bottom). The beam is pinned at both ends. The solution is a relationship between a , the distance between the supports; F , the reaction force at the supports; f , the maximum height of the beam over the support line; and α , the angle of the beam at the supports. See text for a description of the pseudo-rotational spring and definition of k^* .

elastica, and the quality of their results on multiple-rod structures depends directly on their results for one elastica.

For each solver (Kangaroo 1 and 2), we look for parameter settings that will give a predictable accuracy in the shortest possible time. For this, we run the solver with varying parameters on a range of boundary conditions for the elastica. For each set of parameters, the goal is to have consistent accuracy and speed across the range of boundary conditions: this represents the different shapes that the designer will encounter when modeling a set of elasticas.

Elastica problem

We consider a beam of length L , with Young modulus E , section A , and moment of inertia I , as described in Figure 3, top. Analytically, the elastica is the solution of the moment equilibrium in the beam⁴⁰

$$EI \frac{d^2\theta}{ds^2} = -F \sin \theta$$

where s is the curvilinear position along the beam and θ is the angle that the beam makes with the horizontal at that point. Integrating this equation gives a , f , and F in terms of α . Solutions usually focus on the non-dimensional parameters a/L , f/L , and F/F_c , where $F_c = \pi^2 EI / L^2$ is the Euler buckling force, to remove scaling and units issues. For example, for f/L ⁴¹

$$\frac{f}{L} = \frac{\sin \frac{\alpha}{2}}{K\left(\frac{\alpha}{2}\right)}, \text{ where } K(x) = \int_0^{\frac{\pi}{2}} \frac{1}{\sqrt{1 - \sin^2 x \sin^2 \varphi}} d\varphi$$

Extracting a/L from physical observations or numerical simulations, and inverting the relationship on a/L , it is then possible to get the angle α corresponding to that beam and the rest of the parameters. This means that for one simulation, there are three possible comparisons to the analytical values (a/L , f/L , and F/F_c) from one measurement (on α).

Numerical simulations

The simulation model adopted for the beam is the same in each of the two solvers. The connectivity model stems from two straight lines of length $L/2$ forming an isosceles triangle with the horizontal axis as the base of the triangle. The lines are connected at the apex of the triangle and form an angle $\alpha = 1^\circ$ with the horizontal. The lines are discretized in $n/2$ segments each, such that we always have an even total number of segments and a vertex at the middle of the discretized beam.

Internal forces are represented by elastic forces, as available in Kangaroo 1 and 2 (see Figure 3, bottom). Position constraints are defined by an elastic linear spring of rest length 0 and given stiffness, attached on one side to a virtual fixed point and on the other to a defined vertex of the model. Distance constraints are represented by a linear spring of given rest length and stiffness, attached to two given vertices of the model. Angle constraints, representing bending forces, are represented by the discrete three points model from the work by Adriaenssens and Barnes,³⁰ describing the shear force acting on the vertices when the interpolating arc going through them is bent.

In Figure 3, this action of the bending moment is represented by a pseudo-rotational spring that will be integrated by the model into a relationship between bending angle and shear forces applied on the three nodes around it. k^* is

Table 1. Parameters used in the studies and presented in section “Results of numerical experiments.”

Solver	Discretization	Compression ratio	Stopping criterion	Time step	Damping parameter
Kangaroo 1	$n \in [2 : 2 : 36]$	$a/L \in [0 : 0.02 : 1]$	$\log_{10} T \in [-12 : -10]$	0.05	10
Kangaroo 2	$n \in [2 : 2 : 36]$	$a/L \in [0 : 0.02 : 1]$	$\log_{10} T \in [-14 : -11]$	N/A	N/A
Constants	$E = 10 \text{ GPa}, R = 5 \text{ cm}, L = 20 \text{ m}, A = \pi R^2, I = \pi R^4 / 4, k_{\text{Tie}} = 10^8, k_{\text{Anchor}} = 10^{14}$				

[start: step: end] is used to represent the set of numbers from start (inclusive) to end (exclusive), stepping by increments of step. [start: end] = [start: 1: end].

the strength of this relationship as defined in Kangaroo; it has units of rotational stiffness times length (force*length²). This value has the advantage of being independent of the discretization length. For angles close to 180°, this is equivalent to a rotational spring of stiffness k^*n/L . Although this model does not directly track torsional effects, results are still correct for initially straight and untwisted uniform isotropic sections where torsion occurs as a result out of plane loads, as shown in the work by Adriaenssens and Barnes.³⁰

The left endpoint is connected to an elastic anchor of stiffness 10^{14} N/m . A distance constraint is added between the endpoints, with stiffness 10^8 N/m and rest length corresponding to the target a/L . Each segment is characterized by a distance constraint of stiffness (called “strength” in Kangaroo) EAn/L and rest length L/n . Each angle between two consecutive segments is characterized by an angle constraint with rest angle 0 and stiffness (or strength in Kangaroo) EI . The last two properties represent a physically correct discretized beam with the given E , A , and I properties.

We used fixed values for L , E , A , and I . We chose the tie stiffness so that it would be around one order of magnitude higher than EA . This way, the target length of the tie would be closely matched without introducing unnecessarily disparate stiffnesses in the model, which tend to make it less likely to converge. The stiffness of the anchor is high but is only used as a safeguard against rigid-body movement of the model.

Variables and observations definitions

From the nodal positions at the end of the simulation, we extract several observations:

- α , the angle of the first segment with the horizontal—this slightly underestimates the real initial angle of the beam interpolating through the vertices but is coherent with the discretized beam model;
- f , the distance from the middle vertex to the horizontal tie;
- a , the distance between the two endpoints;
- F , the support reaction, is computed by multiplying the tie stiffness by the difference between a and the tie’s rest length.

These are the observations we compare against analytical results. Note that because the tie has a finite stiffness, its final length a is not exactly its rest length. Then, we must make our analytical predictions based on the observed a , not on the tie’s rest length.

In the analysis, we always present comparison of observed simulations versus analytical results as an error on non-dimensional parameters. For example, for an observation f_{obs} associated with an analytical result f_{ana} , the “error on f/L ” is as follows

$$\text{error} \left(\frac{f}{L} \right) = \frac{|f_{\text{obs}} - f_{\text{ana}}|}{|f_{\text{ana}}|}$$

This helps in comparing errors across different boundary conditions, solvers, and observation types.

The stopping criterion terminates the simulation when the particles’ total kinetic energy T falls below a fixed threshold, and the final position is the equilibrium configuration. We checked that this happened before the solver reached its maximum number of iterations, in our experiments, on the elastica. Kangaroo 2 does not use an explicit time step as Kangaroo 1 does, this changes how velocities are computed so the kinetic energy cannot be compared between the two. The number of iterations corresponds to the number of times the points were moved on the process of finding the equilibrium. The damping parameter used by Kangaroo 1’s length constraints is chosen to be as close as possible to critical damping such that all results converge on a test case with $a/L = 0.74$. The time step used in Kangaroo 1 is the length of the time discretization interval or time represented by an iteration.

The wall time is the time that the elastica problem took to run, as reported by the program launching the simulations. While wall time can be influenced by other routines running on the computer, efforts were made to minimize these effects during the simulations so that the results can be reasonably compared.

Finally, the combination of parameters that were tested is represented in Table 1, totaling 6300 experiments. Scripts were created to automatically run the simulation in each case, time it, and collect the results.

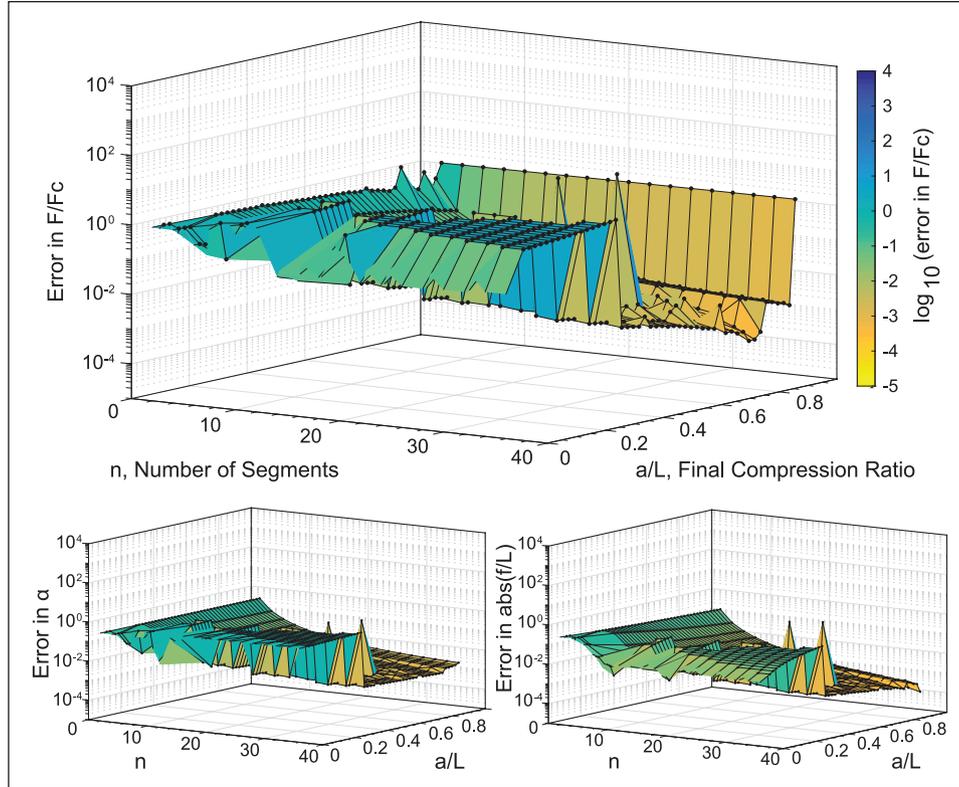


Figure 4. Surface plots of errors in Kangaroo 1 simulations of elastica for different numbers of discretized segments and compression ratio. Kinetic energy threshold used: 10^{-10} .

Speed, accuracy, and reliability

Throughout this article, we use three concepts to evaluate the quality of the tools: speed, accuracy, and reliability. Speed is simply evaluated by the wall time (described above) taken by the algorithm. Speed is nothing if the result is not correct; thus, we also look at accuracy, evaluated with the error measure defined in the previous section. This gives an upper bound achievable by the tool. A reasonable goal that is close to typical construction tolerances could be set at 1% error. Finally, reliability is evaluated by the likeliness of the algorithm to output an incorrect solution. All the elastica results presented converged to the correct solution, but, for example, the introductory example of section “Introduction” shows that this is not always the case for more complicated structures.

Results of numerical experiments

In this section, we present the results of the two software tools, Kangaroo 1 and 2, on the set of numerical experiments described above, focusing on speed and accuracy. Then, we consider a larger example solved using Kangaroo 2 only. There is no analytical solution in that case to

evaluate accuracy of the solvers, but it displays reliability behavior that is of interest in selecting the stopping criterion.

Kangaroo 1

The first set of results we present in Figure 4 are from Kangaroo 1. They show the error on three different measures, in log scale, as we vary the number of segments n and the compression ratio a/L . The errors represented are on α , F/F_c , and f/L . The stopping criterion used was 10^{-10} ; this value displays similar behavior to smaller thresholds as shown in Figure 5.

The three graphs show similar patterns: for small n , the error is close to 100%, meaning that the model failed to predict reality; then, the error decreases asymptotically toward 0.1% as n is increased. This is only true in a limited range of a/L : for a/L smaller than 0.2, the model almost never represents reality, even as n reaches 36. This shows that Kangaroo 1 can only represent limited amounts of bending (when a/L is high, the beam is close to a flat line) and should only be used when the elasticas are compressed by less than 50%.

In the error on F/F_c , the points that have $a/L = 1$ have a high error. This is because at these points, the

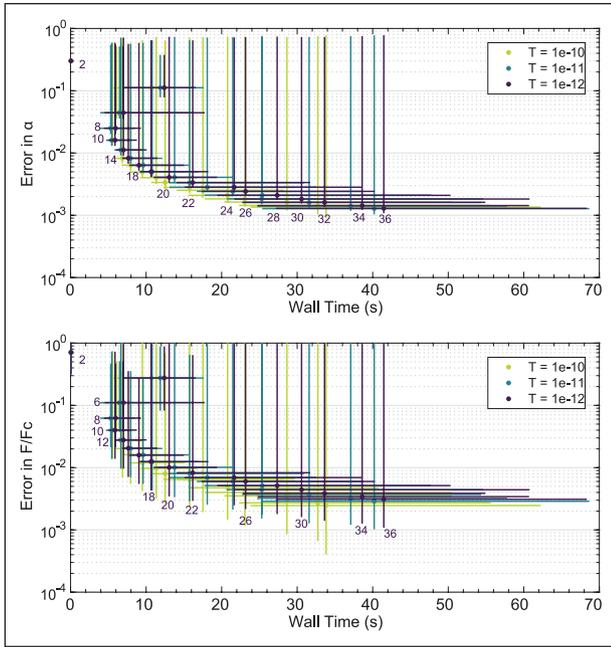


Figure 5. Time versus error in Kangaroo 1 for elastica simulations. Each color represents one threshold for the stopping criterion, each point represents one number of discretized segments (labeled). The point is at the median time and median error for all simulations that have the same number of segments and the same threshold. The extent of the bars represents the spread from first to ninth decile in time and error for these same simulations.

solution beam is nearly flat yet buckled beam, while in reality, it should be exactly flat at the onset of buckling. Then, the force derived from Kangaroo 1 is largely underestimated.

In general, Kangaroo 1 is complex to run reliably, often reaching divergent conditions. In fact, Figure 5 shows that the only way to achieve a precision of 1% or better in F/F_c is with $n \approx 20$, which takes 10–20s to run for one single elastica. Even with small time steps, the simulation is not reliable as a large proportion of the simulation has errors close to 1. This is true for several values of the stopping criterion, and it even seems that increasing it shows in general a degradation in accuracy.

Kangaroo 2

Next, we present similar studies for Kangaroo 2. The stopping criterion used was 10^{-12} (Figure 6). Figure 7 shows how it compares to other possible choices. The plots show the same global trend on n , with the error going from 1% to 0.1% on average when n is increased, in the errors on α and f/L . The error on F/F_c initially starts at much higher levels but then quickly returns to usual when $n > 6$. Here, we find no evidence of

particularly unstable regions in a/L , except on F/F_c when $a/L = 1$, as seen before. This shows that the solver is more reliable, as the error is almost constant across a wide range of boundary conditions. The error tends to be less constant when $n > 24$, especially in α and f/L , but in general remains bounded under the general trend going toward 0.1% error.

In general, it seems that around 20 elements are needed to reliably get an error below 1% in F/F_c . In the other two errors this is achieved even for $n = 10$. Figure 7 confirms this behavior, showing that it is possible to achieve a 1% accuracy in 0.5s. This is using 14–18 elements per elastica, with a threshold of 10^{-12} to 10^{-14} . However, it seems that for thresholds of 10^{-12} or higher, increasing the number of elements risks reducing the accuracy. In addition, accuracies better than 0.1% are almost impossible to obtain reliably for a range of boundary conditions.

In general, these studies show that Kangaroo 1 is not appropriate as a reliable design tool for bending-active structures, as no combination of parameters allows it to reach a predictable level of accuracy or speed. Kangaroo 2, on the contrary, is a good candidate for rapid design iterations in bending-active structures. For a typical rod, it converges in less than a second to accuracies of 1% or better in position and forces. We recommend using a threshold of 10^{-12} or smaller and 15–20 segments per discretized beam. Table 2 summarizes these results.

Larger example: Cocoon

We applied our recommendations to the a posteriori simulation of a bending-active structure. The Cocoon project, built in Winnipeg, Manitoba, is a 12-m-long pavilion made for the “Warming Huts” competition of 2012.³ It is made of fiberglass rebars tied into “double-A-frame” modules and anchored into a frozen river. Dimensions and material properties are shown in Figure 8.

A comparison of the numerical model to actual photos in Figure 9 shows good agreement in shape. We obtained this result by discretizing the beams in 15 elements each, using ties of stiffness 10^7 N/m and rest length 0 m, and anchoring with stiffness 10^9 N/m. Figure 10 (left) shows this assembly for one module of the structure. We ran the simulation in Kangaroo 2 for 10^4 , 10^5 , and 10^6 iterations, with 10^5 iterations the closest to our recommended threshold of 10^{-12} , as demonstrated in Table 3. We found that 10^5 gave the best ratio of accuracy to time of computation. Comparing nodal positions to the reference run of 10^6 iterations, 10^5 iterations is a significant improvement over a threshold of 10^{-11}

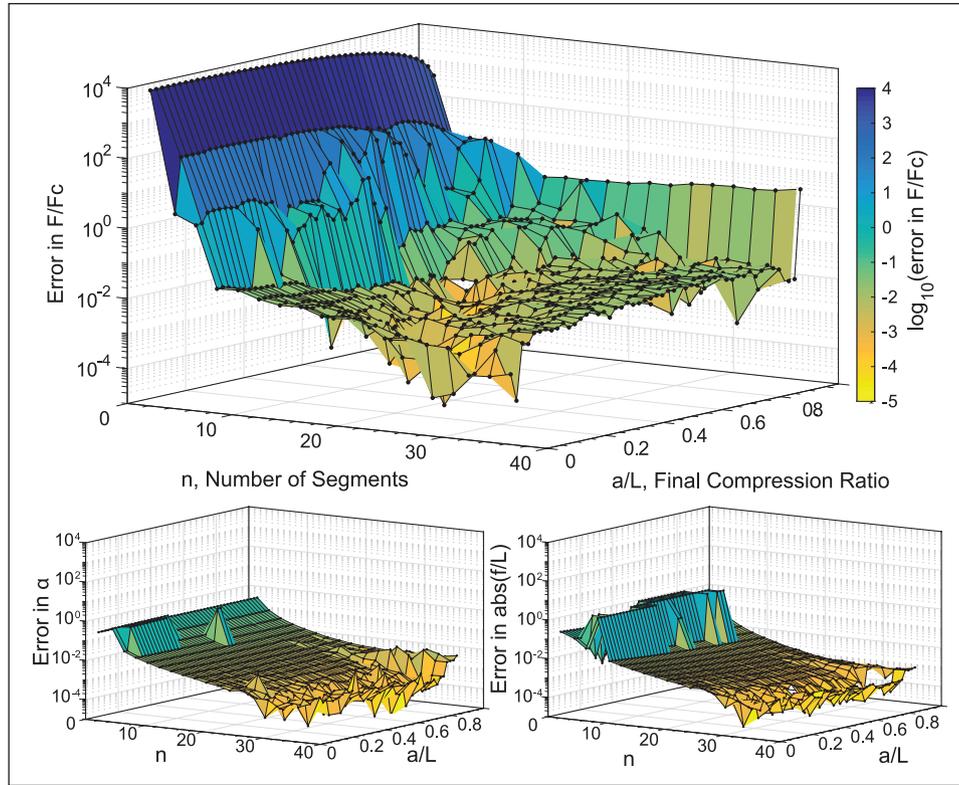


Figure 6. Surface plots of errors in Kangaroo 2 simulations of elastica for different numbers of discretized segments and compression ratio. Kinetic energy threshold used: 10^{-12} .

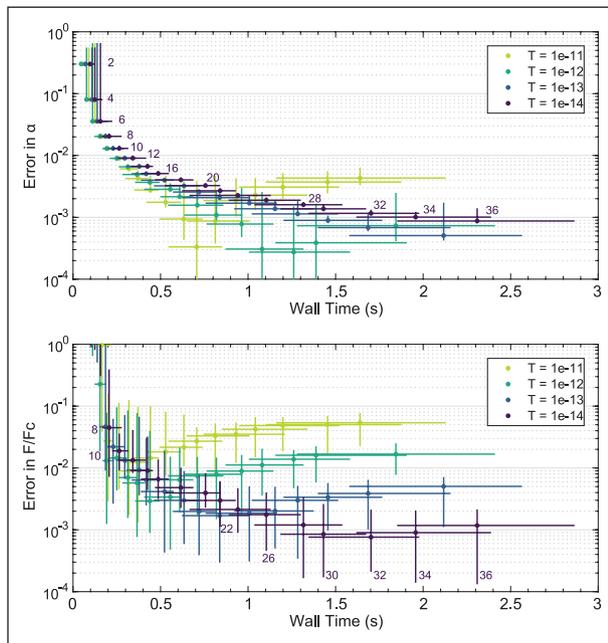


Figure 7. Time versus error in Kangaroo 2 for elastica simulations. See Figure 5 for labels.

(10^4 iterations), without being far from the result found with 10^6 iterations, as shown in Figure 10. We chose to

use the number of iterations as a stopping criterion instead because sometimes a threshold is never reached (see Table 3).

To check the quality of our model, we compared the simulated distance from the top points of the modules to their four anchors, to the same data extracted from photographs with different viewpoints. The results are presented in Figure 11, grouped by anchor position in the module and number of iterations used. Across all measurements but one, the absolute value of the relative error is below 2.5%. There is a clear improvement in errors from 10^4 iterations to 10^5 iterations, with the error getting below 1% on almost all measurements, then very little change as the number of iterations is changed to 10^6 .

This is an encouraging result showing that our recommendations from the previous section work for larger scales of projects. We also find that this is true for reaction forces' prediction (see Figure 12). However, Figure 13 shows that caution is still needed when external forces are applied. In this case, we applied a uniform moderate wind pressure of 0.2 kN/m^2 on the structure, in the pushing direction on the longer side, and pulling direction on the shorter side. As the simulation with 10^6 iterations shows, the structure is failing by buckling on the three external modules on each side (the inner

modules are stiffer because they are connected to more neighbors). This does not happen in the shorter simulation runs, with the 10^5 iterations run only hinting at the phenomenon. This buckling behavior was also observed in other experiments by the authors on prototypes of the structure.

Conclusion

This article has presented new results and guidance for modeling bending-active structures, focusing on key performance metrics: accuracy and speed. Our results show that designers cannot reliably use Kangaroo 1 for simulating such structures, but that Kangaroo 2 can produce good results when tuned properly. In general, there is a trade-off between accuracy and speed; on a single elastica with Kangaroo 2, the error is at best 0.1% for a 1-s runtime. Some combinations of simulation parameters lead to long simulation times and high errors in geometry and forces, especially when thresholds are set too low. This can impede the creative design process both by interrupting a designer's flow and by misrepresenting physical reality. This article recommends using in Kangaroo 2 a threshold of 10^{-12} or smaller, and 15–20 nodes per element in bending, to achieve a spatial

accuracy of 1% or better. This is when using the physical values EA for the axial stiffness and EI for the bending stiffness, which ensures that positions and forces extracted from the model can be directly linked to the physical structure.

Using the guidelines proposed in the article, the Cocoon case study illustrates that high-quality simulation of bending-active structures is possible, but not guaranteed with contemporary, widely available tools. Because the simulation is so sensitive to modeling parameters, critical behavior such as buckling under wind loading can be missed. Furthermore, the physical structure will be impacted by construction tolerances and site variables. As a result, the usefulness of a precise digital model will depend greatly on an equally precise and constrained construction method that can reflect these accuracies.

Future research is needed to understand different bending-active typologies, such as those that use hollow tubes instead of rods as elasticas. In this case, different ratios of bending to stretching stiffness are present, so this article's results may not be directly applicable. In general, simulating tube-based bending-active structures should be faster because of a smaller range of stiffnesses, but not necessarily more accurate.

In conclusion, the exciting possibilities of bending-active structures are expanded by the proliferation of designer-accessible simulation tools such as Kangaroo 1 and 2. However, as shown in this article, a better understanding of how these tools work and perform is needed for physically realistic modeling, and trust in seemingly precise digital results can be easily misplaced. The results presented in this article contribute clarity in this direction

Table 2. Summary of results for the elastica experiments.

Software	Threshold for a 1% accuracy	Typical runtime for a 16 nodes problem
Kangaroo 1	10^{-10}	10s
Kangaroo 2	10^{-12}	0.5s



Figure 8. Cocoon project: dimensions and material properties. Labeled length dimensions are in mm. Photo © Matthieu Léger.



Figure 9. Comparison of actual footage of the construction to numerical simulations. Photos in top row © Matthieu Léger.

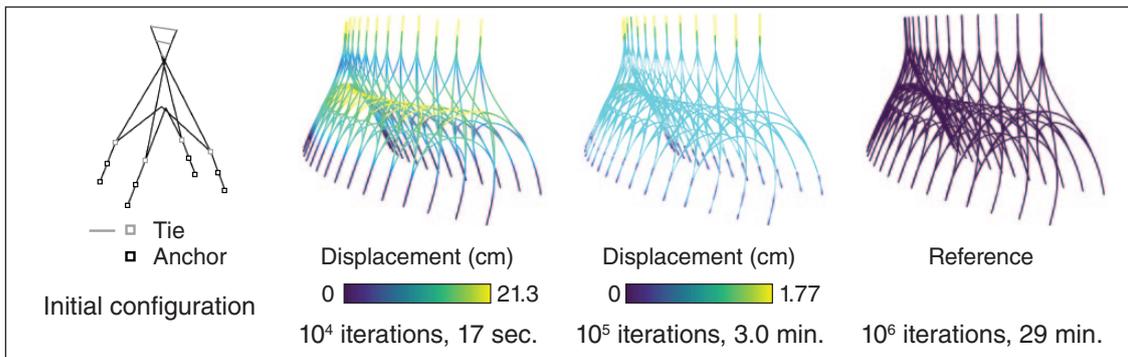


Figure 10. Comparison of results in positions for different numbers of iterations. Displacement shows the positional difference with the reference obtained after 10^6 iterations. In the initial configuration, point-like ties are connected to neighboring modules.

Table 3. Iterations needed for different energy thresholds in the Cocoon model.

Threshold	10^{-10}	10^{-11}	10^{-12}	10^{-13}
Number of iterations needed	6580	10,320	25,810	$> 30 \times 10^4$

The 10^{-13} simulation did not converge after 12-h runtime.

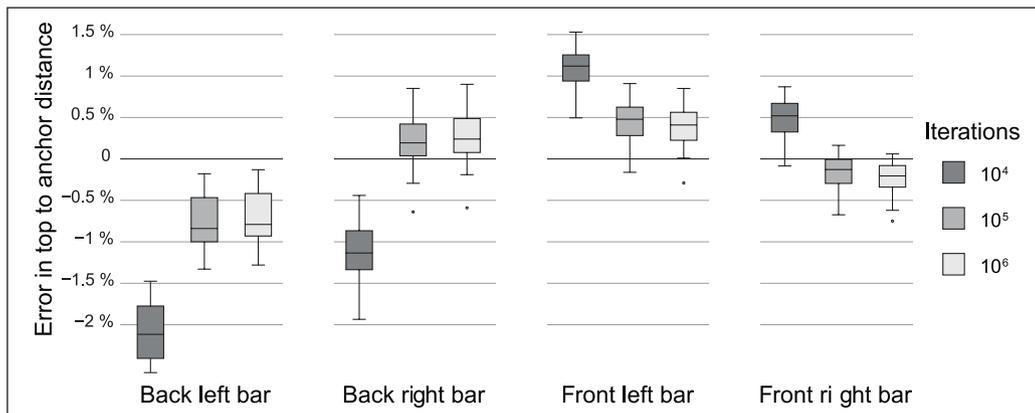


Figure 11. Box plot of the relative errors in the simulated distance from the top points of vertical bars to their anchors, compared to the physical structure, grouped by position of the anchor and number of iterations. Anchors' positions are the four corners of each module, as can be seen in Figure 8. The central bar in the box shows the median of the group of measures, the extent of the box shows the first and third quartile, the whiskers show the minimum and maximum data values, and the points are outliers.

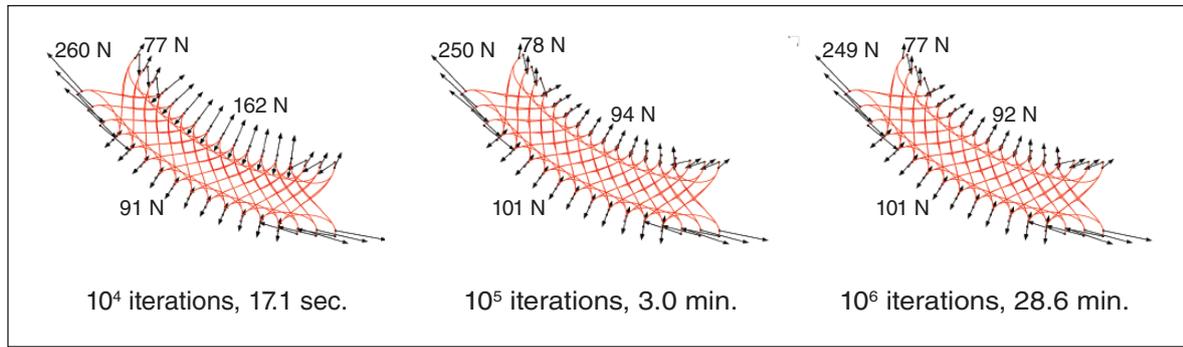


Figure 12. Comparison of results in anchor forces for different numbers of iterations (top view). Maximum force vector error from 10⁴ to 10⁶: 76%, from 10⁵ to 10⁶: 1.5%.

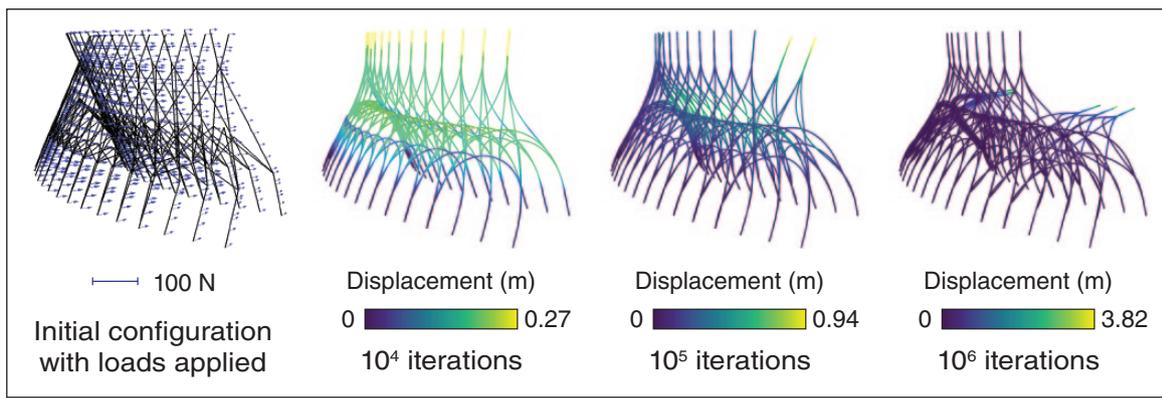


Figure 13. Comparison of results in positions for different numbers of iterations, when a uniform wind pressure is applied on the structure. The displacements are shown from the unloaded configuration with 10⁶ iterations.

and offer steps for improved workflows for designing bending-active structures.

Acknowledgements

Michael Cox, Sigrid Adriaenssens, Lars de Laet, and Mark West are collaborators on the Ice Bloom project shown in Figure 2.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: J.Y. was funded by the MIT Undergraduate Research Opportunities Program (UROP) office's DeFlorez Endowment Fund. Lancelot Coar and the Cocoon project were funded by the Office of the Dean—Faculty of Architecture at the University of Manitoba, the Winnipeg Arts Council, and The Forks Development Corporation.

References

- Williams C. What is a shell? In: Adriaenssens S, Block P, Veenendaal D, et al. (eds) *Shell structures for architecture: form finding and optimization*. New York: Routledge, pp. 21–31, 2014.
- Coar L. On the road/en route, <http://ontheroadenroute.blogspot.ca> (2010, accessed 8 March 2017).
- Coar L. Wobbly structures: exploring the potentials of flexible frames & fabric formed ice structures. In: *Proceedings of the international conference on flexible formwork*, Bath, 1 September 2012.
- McQuaid M, Frei O, Luc E. The Japanese pavilion. In: Shigeru Ban. Paris: Phaidon, pp. 8–11, 2004.
- Cuvilliers P, Douthe C, du Peloux L, et al. Hybrid structural skin: prototype of a GFRP elastic gridshell braced by a fiber-reinforced concrete envelope. *J IASS* 2017; 58: 65–78.
- Olcayto R. Solutions: timber structures–grid shell glazes over the past. *Building Design* 2007; 1776: 14–17.
- Douthe C, Caron J-F and Baverel O. Gridshell structures in glass fibre reinforced polymers. *Constr Build Mater* 2010; 24: 1580–1589.
- Mesnil R, Ochsendorf J and Douthe C. Stability of pseudo-funicular elastic grid shells. *Int J Space Struct* 2015; 30: 27–36.
- Coar L, Cox M, Adriaenssens S, et al. The design and construction of bending active framed fabric formed ice shells utilizing principle stress patterns. In: *Proceedings of the IASS symposium*, Hamburg, 25–27 September 2017.

10. Brew JS and Brotton DM. Non-linear structural analysis by dynamic relaxation. *Int J Numer Method Eng* 1971; 3: 463–483.
11. Piker D. Kangaroo physics, <http://www.food4rhino.com/app/kangaroo-physics> (2016, accessed 3 August 2017).
12. Bouaziz S, Deuss M, Schwartzburg Y, et al. Shape-Up: shaping discrete geometry with projections. *Comput Graph Forum* 2012; 31: 1657–1667.
13. Levien R. The elastica: a mathematical history, 2008, http://www.levien.com/phd/elastica_hist.pdf
14. Goyal S, Perkins NC and Lee CL. Non-linear dynamic intertwining of rods with self-contact. *Int J Nonlinear Mech* 2008; 43: 65–73.
15. du Peloux L, Tayeb F and Lefevre B. Formulation of a 4-DoF torsion/bending element for the formfinding of elastic gridshells. In: *Proceedings of the IASS symposium*, Amsterdam, 17–20 August 2015.
16. Bergou M, Wardetzky M, Robinson S, et al. Discrete elastic rods. *ACM T Graph* 2008; 27: 63.
17. Nabaei S, Baverel O and Weinand Y. Mechanical form-finding of the timber fabric structures with dynamic relaxation method. *Int J Space Struct* 2013; 28: 197–214.
18. Nealen A, Müller M, Keiser R, et al. Physically based deformable models in computer graphics. *Comput Graph Forum* 2006; 25: 809–836.
19. Bergou M, Audoly B, Vouga E, et al. Discrete viscous threads. *ACM T Graph* 2010; 29: 116.
20. Anders C, Deleuran H, Quinn G, et al. Calibrated modelling of form-active hybrid structures, 2016, http://andersholden-deleuran.com/sg2016_HybridStructures.pdf
21. Tamke M, Baranovskaya Y, Deleuran AH, et al. Bespoke materials for bespoke textile architecture. In: *Proceedings of the IASS symposium*, Tokyo, 26–30 September 2016.
22. Nicholas P. Graded territories. In: Nicholas P (ed.) *Designing material materialising design*. Toronto, ON, Canada: Riverside Architectural Press, pp. 49–68, 2013.
23. Lienhard J, Schleicher S, Poppinga S, et al. Flectofin: a hingeless flapping mechanism inspired by nature. *Bioinspiration Biomimetics* 2011; 6: 45001.
24. Lienhard J and Knippers J. Bending-active textile hybrids. *J IASS* 2015; 56: 37–48.
25. Mork HJ, Dyvik HS, Manum B, et al. Introducing the segment lath—a simplified modular timber gridshell built in Trondheim Norway. In: *Proceedings of the world conference on timber engineering*, Vienna, Austria, 21–24 August 2016, pp. 1–8. Trondheim: WCTE.
26. Van Mele T, De Laet L, Veenendaal D, et al. Shaping tension structures with actively bent linear elements. *Int J Space Struct* 2013; 28: 127–135.
27. Ahlquist S and Menges A. Frameworks for computational design of textile micro-architectures and material behavior in forming complex force-active structures. In: Beesley P, Khan O and Stacey M (eds) *Adaptive architecture: proceedings of the 33rd annual conference of the association for computer aided design in architecture*. Cambridge: ACADIA, pp. 281–292, 2013.
28. Lienhard J, Alpermann H, Gengnagel C, et al. Active bending: a review on structures where bending is used as a self-formation process. *Int J Space Struct* 2013; 28: 187–196.
29. Lienhard J. *Bending-active Structures : form-finding strategies using elastic deformation in static and kinetic systems and the structural potentials therein*. Dr.-Ing. Thesis, University of Stuttgart, Stuttgart, 2014.
30. Adriaenssens S and Barnes MR. Tensegrity spline beam and grid shell structures. *Eng Struct* 2001; 23: 29–36.
31. Senatore G and Piker D. Interactive real-time physics. *Comput Des* 2014; 61: 32–41.
32. Hairer E, Lubich C and Wanner G. *Geometric numerical integration: structure-preserving algorithms for ordinary differential equations*. 2nd ed. New York; Berlin; Heidelberg: Springer.
33. Barnes MR. Form-finding and analysis of prestressed nets and membranes. *Comput Struct* 1988; 30: 685–695.
34. Bagrianski S and Halpern AB. Form-finding of compressive structures using prescriptive dynamic relaxation. *Comput Struct* 2014; 132: 65–74.
35. Barnes MR. Form finding and analysis of tension structures by dynamic relaxation. *Int J Space Struct* 1999; 14: 89–104.
36. Barnes MR, Adriaenssens S and Krupka M. A novel torsion/bending element for dynamic relaxation modeling. *Comput Struct* 2013; 119: 60–67.
37. Liew A, Van Mele T and Block P. Vectorised graphics processing unit accelerated dynamic relaxation for bar and beam elements. *Structures* 2016; 8: 111–120.
38. Piker D. Kangaroo solver. Kangaroo group on Grasshopper forums, <http://www.grasshopper3d.com/xn/detail/2985220:Comment:1718553> (2017, accessed 10 May 2017).
39. Piker D. Kangaroo 2 Goals, <https://github.com/Dan-Piker/K2Goals/tree/2a2dacbf63206755ad964c0b0c1a23ee18b0aa4> (2016, accessed 31 May 2016).
40. Audoly B and Pomeau Y. *Elasticity and geometry: from hair curls to the nonlinear response of shells*. New York: Oxford University Press, 2010.
41. Douthe C. *Étude de structures élançées précontraintes en matériaux composites: application à la conception des gridshells*, Thèse de Doctorat, École Nationale des Ponts et Chaussées, Marne-la-Vallée, 2007.