1
Introduction

We see structures all around us, bridges, buildings, furniture, sandcastles, glass, bone, cloth, string, tents, rock piles, for example. Some hold up, some don't. Why? Part of the answer has to do with the materials that make the structure, and part of the answer has to do with the geometry of the structure. A door hinge allows a door to open when pushed, while a cardboard box keeps its shape, unless the cardboard buckles or tears.

We will explain a particular model we call a pin-jointed framework. Joints are regarded as points in space or the plane, and they are connected by inextendible, incompressible bars. The bars are allowed to rotate freely about the joints, and we are seldom concerned whether the bars intersect or not. But we are concerned with rigidity and the behaviour of the framework under loads. There is quite a bit that can be said, not only in calculating forces and displacements, but also in whether the framework is rigid or not.

Another special feature is that after explaining some of the basics in Part I, in Part II we describe the effects of symmetry on the rigidity of the framework. Symmetry can be used to simplify the rigidity calculations described in Part I, but it can also create frameworks that do change their shape.

We have also taken a dual point of view, both from a mathematical/geometric perspective and from an engineering/physical perspective. The geometric problem is: Given the nodes (i.e. points) and distance constraints (i.e. the bars), do these constraints allow the configuration to change its shape. This is a very concrete problem and there are a number of tools that can be used to make this decision. But there are certain demands of rigor and clarity that manifest themselves in the style of exposition. Concepts are defined (usually in italics) clearly and unambiguously, and important statements are put in the form of a "Theorem," helper facts are labelled "Lemmas," and fairly direct consequences of some theorems are called "Corollaries." Justification for these statements are the proof, which is usually set aside and has a little square at the end to let the reader know that the proof has finished, like this. □

When this is done well, it helps the reader see the thread of the argument and see what is important. When it is done poorly it can be tedious, distracting, and confusing. Indeed Gordon (1978) declares

"What we find difficult about mathematics is the formal, symbolic presentation of the subject by pedagogues with a taste for dogma, sadism and incomprehensible squiggles".

We aspire to minimize such tastes, while maintaining our self respect.

The engineering problem is: Given a structure, what is its behaviour under loads? How does the deformation determine the strain (i.e. change in length) of each member? What are the stresses (i.e. forces in the bars) involved? The answer in turn depends on the physical characteristics of the bars, such as Young's modulus that determines their stress–strain ratio. Instead of hard distance constraints, softer more physically realistic elastic behaviour is assumed. Mathematical precision is not necessarily such a concern, but the ultimate justification is the results of an experiment. Indeed, even an experiment in a lab may not be sufficient. Witness the numerous examples of apparently carefully planned structures that have "failed," often with loss of life as described in Levy and Salvadori (2002) and Petroski (2008). Nevertheless, our simple discrete model of a bar and joint framework has been shown experimentally to work for many examples of structures with long relatively thin beams with smaller welded joints. Our purpose here is not to provide complete detail or even provide a totally accurate model of some structure, but instead describe how the analysis of frameworks proceeds and apply this knowledge to a very wide range of circumstances, such as cabled structures (called tensegrities), where some bars are replaced by cables that can only carry tension and not compression, or to packings of spherical balls modelling granular materials, or to large models of glass material, for example. These two approaches are not really separate. The mathematical model effectively assumes some elasticity in the bars in order to prove the framework's rigidity.

But there is more to understand than simply the rigidity or stability of a framework. Another basic question to understand is frameworks that do change their shape even given the hard bar constraints. We call these frameworks *flexible*, but this does not mean that the materials one might use to build them are pliable or soft. It means that there is an exact continuous family of solutions to the equations given by the bars, other than rigid motions of the whole framework. These are often called finite mechanisms in the engineering world.

Chapter 2: Frameworks and rigidity. We introduce the basic framework model, and describe the various different ways in which we can describe the rigidity of the framework, anticipating results in future chapters.

Chapter 3: First-order structures. A natural approach to rigidity is to replace the equations that describe the bar constraints with an appropriate system of linear equations, where the coordinates of the nodes are the variables. There are a wide range of computational and conceptual tools that can be applied. When these techniques apply and show the rigidity of the the framework, it is called infinitesimally rigid or equivalently statically rigid, or kinematically rigid. This is the first requirement for the rigidity or stability in structural engineering. On the hand many structures that are built are not infinitesimally rigid, but nevertheless are still rigid, but not bridges and buildings. They don't fall down, but they tend to be a bit "shaky" in that one can usually feel the play in the framework.

Chapter 4: Tensegrities. In addition to the bar-and-joint model, it is useful to put one-sided constraints between some pairs of joints. If the constraint prevents the end joints from getting further apart, it is called a cable. One can think of this as a string or chain (or cable) connected between its joints. It offers no resistance to the joints

https://doi.org/10.1017/9780511843297.002 Published online by Cambridge University Press

coming together, but constrains them from moving apart. It is also useful to have constraints between some joints where the ends are allowed to move apart, but not come together. At first sight this may seem to be strange, but if one imagines hard spherical billiard balls, the centres of any two touching balls form a natural strut. Indeed, this model is natural for some kinds of granular materials and sphere packing problems. See Donev et al. (2004) for how this can be applied. These structures with cables and struts are called tensegrities. Inspired by their "tensional integrity," the name was coined by R. Buckminster Fuller, who saw some of the sculptures/structures that the artist Kenneth Snelson showed him in the 1940s. Snelson (see Snelson, 1948) went on to create many such large tensegrity sculptures all over the world, made of disjoint bars (or struts) suspended with cables in tension – an example is shown in Figure 1.1. The book Edmondson (1987) is an interesting attempt to give a coherent account of how Buckminster Fuller thought about tensegrities. For books with engineering or practical information about tenesgrities there are Pugh (1976); Skelton and de Oliveira (2010); Motro (2003); Gomez-Jauregui (2010); and Juan and Tur (2008).

(a) (b)

Figure 1.1 Two views of the Snelson sculpture "Needle Tower II," built in 1969 at the Kröller-Müller Museum in the Netherlands. Photographs from: (a) Henk Monster https://commons.wikimedia.org/wiki/File:Hoge_Veluwe_Kroller_Muller_sculpture_garden,_the_1971_18_m_tall_Tensegrity_structure_from_Kenneth_Snelson_-_panoramio.eps; (b) Onderwijsgek https://commons.wikimedia.org/wiki/File:Kenneth_Snelson_Needle_Tower.JPG.

https://doi.org/10.1017/9780511843297.002 Published online by Cambridge University Press

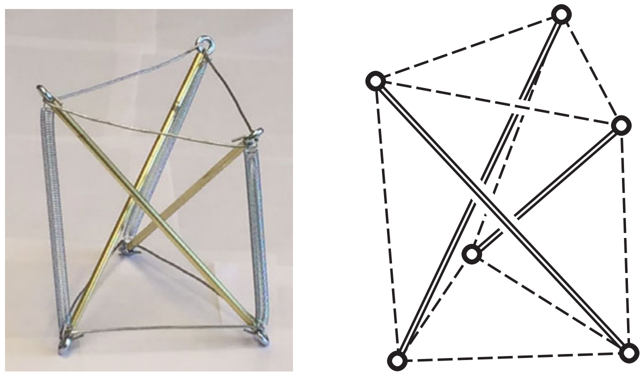

(a) (b)

Figure 1.2 Perhaps the simplest three-dimensional tensegrity. (a) is a model made by one of the authors from brass tubes, springs and wires. (b) is a simplified representation where members in tension are shown by dashed lines, and members in compression are shown by double lines – this anticipates the convention described in Chapter 4.

Chapter 5: The stress matrix. Many of the tensegrity structures created by Snelson and others are not infinitesimally rigid. They are under-braced in that they have too few cables and struts needed for infinitesimal rigidity. A natural question is: Why don't they fall down? One way to answer this question is to use a "higher-order" model that goes beyond the linear approximations in Chapter 3. As a first step a particular quadratic energy function associated with a stress in the tensegrity can be used to solve the geometric question of whether it is rigid. These are defined by a matrix we call a stress matrix. The energy function is quite unrealistic as a physical energy function, considering Hooke's law for springs that is the usual model for most structures. But interestingly this stress–energy function can be used to show global properties with regard to the one-sided tensegrity constraints. For example, the simplified version of the "Needle" tensegrity shown in Figure 1.2 is not only rigid in space, but it is globally rigid in the sense that any other configuration with the cables not longer and the struts not shorter can be translated and rotated to any other or can be reflected and then translated and rotated to the other. Cables are indicated with dashed line segments and bars (which could alternatively be struts) with doubled solid line segments.

Chapter 6: Second-order mechanism and prestress stability. But what is a more realistic physical analysis of an under-braced structure. How does the tension induce the tensional integrity of tensegrities? One way to understand this is to bring in more realistic energy functions than described in Chapter 5. When this is done, interestingly, the stress–energy given by the stress matrix comes in as one component. Because of a change in coordinate systems used in standard engineering analysis, this decomposition is not usually seen there, but the analysis is finally equivalent. This analysis can be used to show the rigidity of many structures, where it is not clear even from the first-order analysis, whether the tensegrity is rigid or not.

https://doi.org/10.1017/9780511843297.002 Published online by Cambridge University Press

Chapter 7: Generic rigidity. Suppose that one has a framework with a large number of nodes, and one wants to know whether it is rigid or not. Stated this way, the question is clearly too hard. Special positions, where the framework is not infinitesimally rigid, but are rigid, nevertheless, can be quite complicated. One can solve the linear equations in principle, but even that can be too time consuming. One way out is simply to assume that the configuration of nodes is not in any sort of special position. The special positions for infinitesimal rigidity are given by polynomial equations with node coordinates as variables and integer coordinates, so why not just assume that these coordinates do not satisfy any such non-zero polynomial equation. We call those node configurations generic. This can be thought of as randomly choosing the coordinates with a natural continuous distribution. Then the chance of hitting a non-generic configuration is zero. If a generic configuration is chosen, in many cases it is possible to determine infinitesimal rigidity for bar frameworks. A very important case is for generic bar frameworks in the plane. It turns out there is a very efficient purely combinatorial algorithm to decide the generic rigidity of bar frameworks, even for finding the rigid parts, the flexible parts, the over-braced parts, and so on. A popular version of this algorithm is called the pebble game (see Jacobs and Hendrickson, 1997) and this has been used to understand a percolation problem inspired by glass networks, where a large over-braced framework gradually has its bars cut until large parts become flexible as in Jacobs and Thorpe (1996). It is also interesting that the generic hypothesis can be used to determine generic global rigidity Connelly (2005); Gortler et al. (2010); and Jackson and Jordán (2005).

Chapter 8: Finite mechanisms. The other side of structural mechanics is (finite) mechanisms. There is a large variety with a wide range of purposes, applications, and unexpected properties. They can be used to draw curves, open robot arms, and draw straight lines, and there are examples of closed polyhedral surfaces that actually flex, but only with their volume constant. These are exact finite mechanisms, that strictly satisfy the distance constraints. Many of the tools that are used to show frameworks are rigid can be turned around to show that they are not rigid.

Part II is concerned with frameworks that are symmetric. In other words there are rigid motions, such as rotations or reflections that take the nodes and bars of a framework to itself.

Chapter 9: Groups and representation theory. We start with a brief introduction to the theory of groups – which is what we need to deal with symmetric tensegrities, and with representation theory, which explains how groups operate as a set of transformations or matrices. Representation theory turns out to be exactly what we need in the remaining two chapters.

Chapter 10: First-order symmetry analysis. Here we look at how representation theory can be applied to the material from Chapter 3, and how this can be used to predict how to make symmetric structures that are rigid, and how some of the properties will be special.

Chapter 11: Generating stable symmetric tensegrities. Finally we show how representation theory allows us to generate whole families of rigid symmetric tensegrities.

1.1 Prerequisites

We assume that the reader is familiar with a first course in linear algebra including vectors, matrices, linear transformations, etc. such as is in Strang (2009) and is not put off by a formula or two. Of standard books that introduce the reader to the analysis of structures from an engineering perspective, we suggest Parkes (1965) or Livesley (1964). We also have included a few exercises at the end of the chapters for a little work-out for the reader.

For the second part of the book, we have included a brief introduction to group theory and some of the relevant parts of representation theory. The books by Bishop (1973) and James and Liebeck (2001) are good sources for that theory.

1.2 Notation

We denote vectors and matrices in bold. Vectors are thought of as columns, and $()^{\mathrm{T}}$ denotes the transpose operation for both vectors and matrices.

We often have occasion to denote a vector of vectors, which is a string of vectors on top of each other, which we call a configuration. So we use the notation

$$[\mathbf{p}_1; \mathbf{p}_2; \ldots; \mathbf{p}_n] = \begin{bmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \\ \vdots \\ \mathbf{p}_n \end{bmatrix},$$

which is notationally more convenient than writing columns everywhere.

Tensegrities, discussed in Chapter 4, are composed of cables, struts, and bars, where non-extendable cables are denoted by dashed lines, non-compressible struts by double lines, and length invariant bars by solid single lines, as shown in Figure 4.1.

The graph associated with a framework or tensegrity is usually denoted by G, and a framework or tensegrity, which consists of a graph G, as well as a configuration of points $\mathbf{p} = [\mathbf{p}_1; \ldots; \mathbf{p}_n]$, is denoted by $(G, \mathbf{p})$.

In Part II a group is usually denoted by $\mathcal{G}$, and a representation of a group as ρ.

https://doi.org/10.1017/9780511843297.002 Published online by Cambridge University Press