

1

Introduction

When on sunny days in the early years of the eighteenth century people hastened through the narrow streets of Königsberg, now Kaliningrad in Russia, it was not for catching a breeze of fresh air, but for being the lucky first to unlock the mystery of its now famous bridges. Not a valuable golden treasure was to be unearthed, but the solution to the simple problem of finding a path over the seven bridges that connect both banks of the river Pregel with the two islands around which that beautiful Prussian town was built. The goal, however, was not to find just any path, but one along which each bridge was crossed only once. A futile endeavour indeed, and it is unknown how many of Königsberg's inhabitants and intrigued travellers from afar succumbed to this formidable challenge before in 1736 the brilliant mind of a young genius named Leonard Euler provided the definite answer. *Such a path does not exist!* [42].

What was Euler's deep insight which allowed him to solve the problem of Königsberg's bridges without setting foot on any of them, the insight which had eluded all mathematical treasure hunters, professional and recreational alike, before him? In short, seven bridges connect the four main geographical areas of Königsberg, but only one or two of these areas can serve as start and finish of the sought-after path through the city. This means that at least two areas must be traversed and, in order to do so, must be connected by an even number of bridges across the river Pregel. However, as a quick look at the map of Königsberg reveals (Fig. 1.1a), each of its four areas was accessible only through an odd number of bridges, thus rendering the existence of a valid path an impossibility.

Although it still took many decades after Euler's simple yet brilliant resolution of the now famous *Königsberg bridge problem* before a coherent and rigorous mathematical framework was established, it undoubtedly marks both historically and conceptually the beginning of what is now known as network or *graph theory*. Indeed, his contribution in this regard cannot be overstated.

2

Introduction




Figure 1.1 The bridges of Königsberg and their graph-theoretical representation. **a:** Seven bridges (*a* through *f*) connect both sides of the river Pregel (*B* and *C*) with the large islands Kneiphof and Lomse (*A* and *D*). Is it possible to reach all four parts of the city by crossing each bridge once and only once? **b:** Euler's ingenious solution rests on the abstraction of the problem, replacing *A*, *B*, *C* and *D* by nodes, and the bridges by links, or edges, interconnecting these nodes. As all nodes are connected by an odd number of edges, the problem's solution is in the negative [42]. **c:** Visualisation of the weighted adjacency relations defining the network at the heart of the Königsberg bridge problem in the form of a matrix. Rows and columns represent the departure and arrival areas, respectively, the weights indicate the number of bridges connecting each of these areas with one another. Panels **a** and **b** are modified from Kraitchik [68, Chapter 8.4].

By stripping colourful yet less relevant details off the given original problem, Euler single-handedly introduced an abstract notion which not just allowed him to uncover its defining essence, but in fact touches upon the very fabric of physical reality. *Each real-world system is finite and discrete*, comprised of discernible parts which share links to, interact with or assert an action on other parts of the system. Be it the plethora of elementary particles interacting with one another through the exchange of virtual bosons, the complex dendritic and axonal trees which carry electrical signals through a network of synaptically interlinked neurons in our brains, or the arrangement of galaxies into large filaments through the action of the gravitational force, all natural phenomena we discovered so far can be abstractly described by and modelled as networks of interconnected objects. These objects are called *nodes* or *vertices*, their connections *edges*, and graph theory, in its most general sense, is the mathematical framework for the study of networks formed by nodes and edges.

Truth be told, it would be an almost futile endeavour to present just a glimpse of the large body of graph-theoretical literature which emerged in the past three centuries since Euler's time. For that reason, we must, and only humbly can, refer the interested reader to the historically first textbook written by König [65], the still-unmatched definite textbook on the subject by Harary [59] and perhaps the most comprehensive presentation of modern applications of graph

theory by Newman [86], and by doing so challenge, somewhat self-servingly, each reader to construct from these initial literature nodes his or her very own ‘network of exploration’ of the formidable world of graph-theoretical literature. From this it will become immediately clear that not even a genius such as Euler could have envisioned the undoubtedly overwhelming success of graph theory, its underpinning conceptual beauty and far-reaching applications in the study, modelling and characterisation of real-world systems. Unfortunately, however, even the greatest success story hides a tiny stain, some form of caveat hinting at a possible limitation or complication, and it is such a stain on the Michelangelonian masterpiece of graph theory which this book intends to identify and eventually target.

As we already asserted above, due to the very makeup of physical reality, most, if not all, real-world phenomena can be described in terms of networks of interconnected objects. Graph theory, as the mathematical framework formalising such a description, has over the past century not just provided us with a vast yet ever-growing number of concrete examples of networks – ranging in size from a few nodes, such as in the case of the Königsberg bridges, to hundreds of million of nodes in the case of the World Wide Web – but also amassed a plethora of tools and methods for analysing and understanding these networks. Most of these tools and applied methods, however, are of a quantitative nature and, thus, inherently rely on numerical evaluations by hand or, more contemporarily, the use of computers. Although big advances in computational hardware over the past few decades have made it possible to characterise and delineate finer structural details of increasingly larger networks, the question of principal limitations of computational approaches looms like the sword of destiny over the head of every researcher who dares to venture into the seemingly endless realms of graph theory. How far can we go, or – more importantly – how far will we be ever able to go?

To illustrate the more fundamental nature of this question, let us return to the seven bridges of Königsberg and ignore for a moment Euler’s definite, and simple, solution of the associated problem in the negative. How many paths through the city are possible? More precisely, how many *Eulerian walks* – that is, paths which traverse each bridge at most once, do exist – and how many of these walks define *Eulerian paths* by crossing each bridge exactly once? It will not take long to draw a tree of all possibilities, a small fraction of which is shown in Fig. 1.2, and count a total of 90 Eulerian walks starting from either bank of the river Pregel (node *B* or *C*, Fig. 1.1), 6 of which leave three bridges uncrossed, 12 of which leave two bridges uncrossed, and 72 of which end with one bridge remaining untraversed. None of these 90 walks – nor, for that matter, any of the possible walks emanating from the other parts of the city – is,




Figure 1.2 Solving the problem of Königsberg's seven bridges through an exhaustive search. Each branch in this tree constructs a possible Eulerian walk, with bridges being crossed marked in grey and bridges already traversed in black. Shown are only Eulerian walks starting from one bank of the river Pregel, omitting all walks which begin by crossing the second bridge to the south. Each branch of this tree of possibilities leaves at least one bridge uncrossed; thus, none of the constructed walks constitutes a sought-after Eulerian path.

thus, an Eulerian path. Without having ever set foot in Königsberg, we just solved the problem of its bridges by brute force on a piece of paper. But today, one can find some 100 bridges and overpasses in Kaliningrad. How long would it take now to draw a tree of all possible walks through the city?

In short, about 10^{150} times as long, give or take some orders of magnitude! A formidable exercise which undoubtedly lies beyond the capacity of even the most modern computers, let alone a wholeheartedly motivated researcher with paper and pencil. The longer and more nuanced answer to this question necessitates the introduction of the notion of *computational complexity*, a measure

which, in the most general sense, quantifies the scaling behaviour of computational algorithms with respect to the magnitude of their input variables, in our case the number of bridges and overpasses. To distinguish further the time it takes to run a given algorithm from the storage space it requires, both defining aspects quantifying the efficiency of a computational algorithm in relation to the stringent constraints imposed by the utilised computational hardware, one typically differentiates *time complexity* and *space complexity*, respectively.

Let us first take a look at the time complexity of our little exercise. Drawing a single walk through the city which traverses each bridge – that is, each edge in the associated graph – at most once certainly scales linearly with the total number of bridges E . In mathematical terms, one expresses such a scaling behaviour with the notion $\mathcal{O}(E)$. However, our exercise requires us to draw all Eulerian walks and identify potential Eulerian paths, a rather brute-force approach which is called *exhaustive search*. In order to do so, we must start at every possible point and traverse all the possible paths from that point onwards (Fig. 1.2), a task with a staggering time complexity of not less than $\mathcal{O}(E!)$. Luckily for us, the original problem absolves us from having to cope with such an unmanageable complexity, as it requires us only to decide whether an Eulerian path exists or not. Indeed, following Euler's inspiring solution, we only need to calculate the number of edges connected to each node and test whether any of these numbers is even, a task which can certainly be performed in $\mathcal{O}(E)$ and leaves us off the hook, this time. Unfortunately, however, many of the quantitative graph-theoretical methods and measures do not allow for such an elegant approach, and they typically scale polynomially with the size of the graph, that is, with a time complexity of $\mathcal{O}(E^k)$ for some $k \gg 1$. Undeniably, such a scaling is certainly better than the factorial scaling we encountered above in a brute-force construction. But even with the development of computational hardware in the foreseeable future or the conception of novel, highly efficient computational algorithms, such a polynomial time complexity is still hard to cope with and likely renders the exact quantification of larger networks difficult or even impossible for many decades to come.

The situation regarding space complexity does not look much less gloomy when dealing with graph-theoretical formulations of real-world phenomena. Although efficient representations for a small subset of special graphs exist, generally one is left with a matrix representation of a graph's adjacency relations (see Fig. 1.1c), an approach with a space complexity of $\mathcal{O}(N^2)$ in the number of nodes N . Certainly, neither the seven bridges of Königsberg during Euler's time nor the more numerous bridges in current-day Kaliningrad pose a representational problem here. But what about the graph describing the World Wide Web with its hundreds of millions of nodes [1, 34]? Even if we assign

only a Boolean number in the form of a computational bit to indicate the presence, or absence, of an edge between any given pair of nodes, it would require a staggering 1000 terabytes to store the complete adjacency matrix of a graph describing only 100 million linked websites. Luckily for us, in this specific case, a linear list of source and target nodes would suffice, as each node of this graph is, at average, connected to only seven or eight other nodes (e.g., see [26, 54] but also [77]). But even such a sparse representation would still require about one gigabyte of memory alone for storing the graph's defining structure, without any computational analysis having yet been performed on it. When considering still larger or less sparsely connected graphs, such as the network of neurons in the human brain with thousands of connections embedding each of its 100 billion nodes [11, 23], we quickly run out of options as the memory requirements for their representation alone will breach the storage capacity made available in modern computers, perhaps even the computational devices of the next few generations.

The question, then, is, where does this leave us? Of course, a computational analysis is just one of several available options when facing a graph-theoretical formulation of a natural phenomenon. Indeed, the very limitation of a purely numerical approach has led to the development of a whole host of new and exciting methods and continues to stimulate novel ideas and techniques beyond the mere optimisation of computational algorithms. Arguably, one of the conceptually most important and fruitful of these ideas is the distinction of real-world networks into classes of simplified, heavily idealised graph models. Examples of such idealised models include perhaps the most-widely known Erdős–Rényi model [22, 38] which generally serves as the prototype when referring to random graphs, the small-world graph of Watts and Strogatz [109] inspired by Stanley Milgram's infamous 'six degrees of separation' [78] and the scale-free graph of Barabási and Albert [12] which captures the property of self-similarity found in complex systems, to mention but a few. Most of these models are conceived by abstracting from many of the finer structural peculiarities we observe in real-world networks, details which are, as we argued above, increasingly harder to delineate through computational approaches in cases where larger or densely connected networks are concerned. With these graph models typically requiring less than a handful of parameters for their construction and characterisation, the simplified nature of the models then not only allows for the application of numerical optimisation techniques which significantly reduce their computational complexity, but eventually opens the door for a more rigorous mathematical exploration by allowing for an algebraic formalisation of their algorithmic construction.

In this book, we will propose and motivate exactly such a rigorous mathematical approach. By embarking on this adventurous journey, however, we will seek to avoid at least one of the major caveats accompanying many of the idealised graph-theoretical models proposed as descriptions of real-world phenomena since the time of Euler. To briefly illustrate this caveat, let us consider perhaps the simplest of these abstract models, that of the Erdős–Rényi random graph. For each point in its two-dimensional parameter space, which is spanned by the number of nodes and the number of edges, or, alternatively but with subtle conceptual differences, the connection probability between pairs of nodes, a whole ensemble of possible realisations exists, each of which describes a distinct connectivity pattern between all possible pairs of nodes. Although the algebraic formalisation of this model allows, at least in principle, access to this ensemble, most of the defining characteristics of the Erdős–Rényi model can still only be assessed numerically, ideally so over a large number of realisations for each parameter set in order to ensure statistical significance of the obtained results.

Naturally, just as in the case of graph-theoretical descriptions of real-world phenomena, such a computational approach is subject to the stringent limitations mentioned earlier, especially when models of larger graphs are considered. Due to its mathematically well-rooted conception, however, the Erdős–Rényi random graph model also invites a more rigorous analytical examination. Some of its characteristics, such as the number and distribution of edges connected to each node, are indeed governed by exact formulae which, due to the random makeup of the Erdős–Rényi model, are probabilistic in nature and deliver the expectation value for those characteristics across the entire parameter space and, thus, the whole ensemble of possible realisations. With other properties – such as the shortest distance between any pair of nodes, measured in number of edges which must at least be traversed to reach a target node, or the clustering, quantified by the abundance of triangular cycles – we are unfortunately not so lucky. Indeed, many if not most of the properties allow for an analytical take only in some parts of the model’s parameter space, that is, under certain and often stringent conditions imposed on the graph’s makeup.

It was the meticulous work of Paul Erdős and Alfréd Rényi in the early second half of the last century which eventually established what is now called *random graph theory*, the mathematical study of properties of the probability space of random graphs, a field which stretches today with applications far beyond the random graph model named after its original architects (for a comprehensive introduction, see [22]). The conceptual core of this theory defines a set of conditions which allow for a mathematically rigorous quantification of a

random graph's properties. But – and herein lies the problem or caveat hinted at above – these very conditions are rooted in asymptotic considerations, specifically and most importantly the heavy-weighting assumption which demands that the network size should become infinitely large. Although for the math-savvy reader such a demand might certainly appear as nothing unusual, as something being widely employed in mathematics and its many applications – a trick if you will, to establish some form of analytical framework for describing a physical phenomenon or, as in the case of random graphs, for formalising a theoretical model – we must stop here and evaluate its potentially pernicious repercussions and pitfalls.

The important question to ask before issuing any demands on our approach is, do we find in nature any system which we know for certain is infinite in one or more of its properties or, alternatively, is genuinely continuous in an analytical or, if you like, Cantorian sense? Or is it that our models, describing these systems, assume such an idealising makeup for reasons of simplicity, or the sheer lack of a more suitable approach in tune with our study of nature and, thus, epistemology? We could go as far as posing the heretical question of whether the fabric of physical reality itself is algebraic, that is finite and discrete, or analytic, that is infinite and continuous, in nature? Although on first sight such questions might appear far removed from an graph-theoretical inquiry of our world – indeed, as being of little or no relevance today as our analytical formulation of real-world phenomena has undeniably demonstrated tremendous success – their gravity continues to penetrate philosophical discussions and, arguably, lies at the very bedrock of modern science and its mathematical foundation (a comprehensive review of arguments from both sides can be found in Hagar [56]).

To illustrate this point, arguably naively so, in the context of this book, let us take a look at a physical phenomenon which many of us encounter on a daily basis, namely preparing, and eventually enjoying, a nice cup of coffee. Hot water is poured on finely grained coffee and finds its way – ‘percolates’ – through the porous brown powder to finally yield the tasty, longed-for beverage. This mundane process is just one example of a whole class of phenomena in condensed matter physics which finds a natural representation in terms of random graph theory. However, due to its analytical foundation, the obtained results are, at best, valid only in the case that the number of nodes in the associated graph-theoretical models approaches infinity – that is, we deal with an impossible-to-imagine amount of coffee in front of us – or if the connection probability between nodes is small enough that the connectivity pattern of the model graph resembles that of a tree, a structure that is devoid of cycles or loops and resembles, to stick with our illustrative example, anything but the

coffee we are all familiar with. Indeed, the reality of making coffee certainly looks much different from the model we use to describe it, as we do not have at our disposal an (almost) infinite quantity of coffee powder for quenching our desire, nor do we find a resemblance between a sparsely connected tree and the gaps separating the coffee grains through which the water finds its way into our cups!

Despite the fact that theoretical models, by their very conceptual definition, can provide only a simplified, more abstract view of natural phenomena, the above highlighted differences are certainly, or at least arguably, not even close to being a viable reflection of the real-world system in question. Thus, reiterating the above questions, we must ask whether random graph theory, indeed graph theory in general, with its analytical core is the right mathematical approach to describe the discrete and finite phenomena, such as percolation, which we find all around us. After all, every system we can discern in nature, in lack of a proof to the opposite, is both finite and discrete. An answer to this question becomes even more pressing when considering systems which, by their very definition, are bound in size, such as the game of chess, or display boundaries which crucially shape the system's defining behavioural characteristics, such as the aromatic powder in a coffee machine. In such cases we have no choice but to strictly adhere to a finite description, as each simplifying asymptotic assumption will necessarily push us outside the defining premise of the studied phenomenon and eventually deliver a model which no longer observes the phenomenon being considered.

Is it possible to conceive of a mathematical framework which is capable of more precisely capturing finite and discrete graph-theoretical models, thus allowing for a more viable and, arguably, more accurate and thus valid description of real-world phenomena? In this book, we will answer this question affirmatively and motivate a rigorous algebraic approach for the construction, analysis and characterisation of finite graphs. This approach, which we will term *operator graph theory*, comes at the price of abandoning the original definition of a network, or graph, which resides at the very core of classical graph theory in favour of a more dynamic, operational viewpoint. A graph will no longer be a mere static collection of nodes and edges, but instead a construct which can evolve and change due to the actions of operators. We will argue that by studying the algebraic properties typically associated with such operators, this approach opens up a new dimension of qualitative and quantitative insights into the very properties which define a given finite graph, insights which find a viable mathematical representation without resorting to the dangerous tool bag of asymptotic approximations and throttling limitations. In order to further stress, on epistemological grounds and in full awareness of its heretical nature,

this very rejection of a continuous, thus analytic, description of real-world phenomena in favour of a finite, thus algebraic, approach, we will restrict, unless otherwise motivated, to the ring of rational numbers \mathbb{Q} instead of the ring of real numbers \mathbb{R} throughout the presentation.

However, before we embark on this adventurous journey, it is of upmost importance to note that the approach presented in this book must and can only be viewed as a coarse introduction, primarily focusing on the mere motivation of a new angle from which to view graph theory and its utilisation as a descriptive tool in the understanding of real-world phenomena. For that reason, the book is divided into two main parts. In the first part, we will lay the necessary theoretical foundation for operator graph theory. To that end, Chapter 2 will provide a necessarily incomplete introduction into classical graph theory by focusing on the presentation of concepts, terminology and mathematical notions needed for the comprehension of the material in the remainder of this book. A similar approach holds for Chapter 3, in which we will take a brief look at the essential fundamentals of the vast field of discrete operator calculus. By fusing the notion of a classical graph with concepts of operator theory introduced in these two chapters, we will then be ready to define the central contribution that this book is aiming for, namely that of an operator graph. In Chapter 4, concluding the first part, not only will this fusion be motivated, but the core conceptual notions of operator graph theory, along with definitions and a presentation of its mathematical framework, will be justified and, hopefully, accessibly illustrated.

The second part of this book focuses on exemplifying the proposed operator graph-theoretical framework by presenting its applications to a few relatable systems and well-known conceptual models of real-world phenomena. Specifically, we will formalise the rigorous algebraic generation of the most-widely used finite random graph models in Chapter 5, and then use this formalisation in Chapter 6 to obtain exact algebraic expressions for a variety of properties characterising these graphs, without resorting to asymptotic approximations. Closing with Chapter 7, we will finish our journey with a brief stroll through the playful realms of game theory by demonstrating the potential usefulness of our operator graph-theoretical framework in the construction and analysis of the game of chess.

We ardently hope that with the chosen examples and applications, the interested reader will be inspired to partake in the furthering of the yet-to-be-fleshed-out theoretical foundations of the operator graph-theoretical framework and be thoroughly motivated not just to consider the latter as a potentially powerful ally in the graph-theoretical description and analysis of real-world phenomena, but to actively utilise this exciting novel approach in the conquest of understanding nature.