The Mathematics of Logic

Richard Kaye

A guide to completeness theorems and their applications

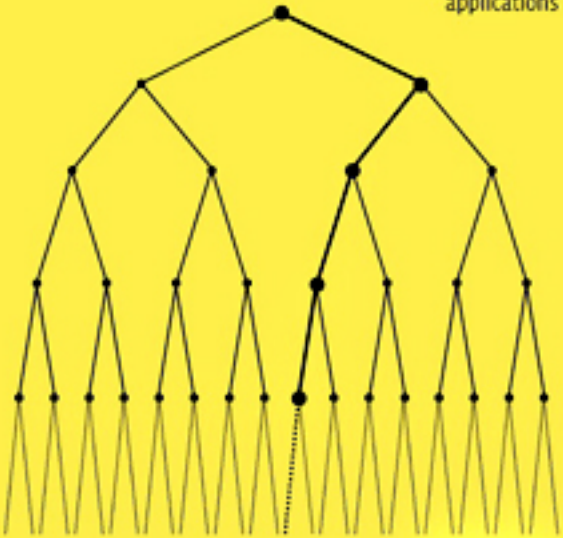

CAMBRIDGE

CAMBRIDGE www.cambridge.org/9780521708777

This page intentionally left blank

The Mathematics of Logic

A guide to completeness theorems and their applications

This textbook covers the key material for a typical first course in logic for undergraduates or first year graduate students, in particular, presenting a full mathematical account of the most important result in logic: the Completeness Theorem for first-order logic.

Looking at a series of interesting systems increasing in complexity, then proving and discussing the Completeness Theorem for each, the author ensures that the number of new concepts to be absorbed at each stage is manageable, whilst providing lively mathematical applications throughout. Unfamiliar terminology is kept to a minimum; no background in formal set-theory is required; and the book contains proofs of all the required set theoretical results.

The reader is taken on a journey starting with König's Lemma, and progressing via order relations, Zorn's Lemma, Boolean algebras, and propositional logic, to Completeness and Compactness of first-order logic. As applications of the work on first-order logic, two final chapters provide introductions to model theory and non-standard analysis.

DR RICHARD KAYE is Senior Lecturer in Pure Mathematics at the University of Birmingham.

The Mathematics of Logic

A guide to completeness theorems and their applications

Richard Kaye

School of Mathematics, University of Birmingham

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK
Published in the United States of America by Cambridge University Press, New York
www.cambridge.org
Information on this title: www.cambridge.org/9780521708777

First published in print format 2007

ISBN-13 978-0-511-33998-1 eBook (Adobe Reader)
ISBN-10 0-511-33998-4 eBook (Adobe Reader)

ISBN-13 978-0-521-70877-7 paperback
ISBN-10 0-521-70877-X paperback

Contents

Preface

Mathematical logic has been in existence as a recognised branch of mathematics for over a hundred years. Its methods and theorems have shown their applicability not just to philosophical studies in the foundations of mathematics (perhaps their original *raison d'être*) but also to 'mainstream mathematics' itself, such as the infinitesimal analysis of Abraham Robinson, or the more recent applications of model theory to algebra and algebraic geometry.

Nevertheless, these logical techniques are still regarded as somewhat 'difficult' to teach, and possibly rather unrewarding to the serious mathematician. In part, this is because of the notation and terminology that still survives as a relic of the original reason for the subject, and also because of the off-putting and didactically unnecessary logical precision insisted on by some of the authors of the standard undergraduate textbooks. This is coupled by the professional mathematician's very reasonable distrust of so much emphasis on 'inessential' non-mathematical details when he or she only requires an insight into the mathematics behind it and straightforward statements of the main mathematical results.

This book presents the material usually treated in a first course in logic, but in a way that should appeal to a suspicious mathematician wanting to see some genuine mathematical applications. It is written at a level suitable for an undergraduate, but with additional optional sections at the end of each chapter that contain further material for more advanced or adventurous readers. The core material in this book assumes as prerequisites only: basic knowledge of pure mathematics such as undergraduate algebra and real analysis; an interest in mathematics; and a willingness to discover and learn new mathematical material. The main goal is an understanding of the mathematical content of the Completeness Theorem for first-order logic, including some of its mathematically more interesting applications. The optional sections often require additional background material and more 'mathematical maturity' and go beyond a

typical first undergraduate course. They may be of interest to beginning postgraduates and others.

The intended readership of this book is mathematicians of all ages and persuasions, starting at third year undergraduate level. Indeed, the 'unstarred' sections of this book form the basis of a course I have given at Birmingham University for third and fourth year students. Such a reader will want a good grounding in the subject, and a good idea of its scope and applications, but in general does not require a highly detailed and technical treatment.

On the other hand, for a full mathematical appreciation of what the Completeness Theorem has to offer, a detailed discussion of some naive set theory, especially Zorn's Lemma and cardinal arithmetic, is essential, and I make no apology for including these in some detail in this book.

This book is unusual, however, since I do not present the main concepts and goals of first-order logic straight away. Instead, I start by showing what the main mathematical idea of 'a completeness theorem' is, with some illustrations that have real mathematical content. The emphasis is on the content and possible applications of such completeness theorems, and tries to draw on the reader's mathematical knowledge and experience rather than any conception (or misconception) of what 'logic' is.

It seems that 'logic' means many things to different people, from puzzles that can be bought at a newsagent's shop, to syllogisms, arguments using Venn diagrams, all the way to quite sophisticated set theory. To prepare the reader and summarise the idea of a completeness theorem here, I should say a little about how I regard 'logic'.

The principal feature of logic is that it should be about reasoning or deduction, and should attempt to provide rules for valid inferences. If these rules are sufficiently precisely defined (and they should be), they become rules for manipulating strings of symbols on a page. The next stage is to attach meaning to these strings of symbols and try to present mathematical justification for the inference rules. Typically, two separate theorems are presented: the first is a 'Soundness Theorem' that says that *no incorrect deductions* can be made from the inference rules (where 'correct' means in terms of the meanings we are considering); the second is a 'Completeness Theorem' which says that *all correct deductions* that can be expressed in the system can actually be made using a combination of the inference rules provided. Both of these are precise mathematical theorems. Soundness is typically the more straightforward of the two to prove; the proof of completeness is usually much more sophisticated. Typically, it requires mathematical techniques that enable one to create a new mathematical 'structure' which shows that a particular supposed deduction that is not derivable in the system is not in fact correct.

Thus logic is not only about such connectives as 'and' and 'or', though the main systems, including propositional and first-order logic, do have symbols for these connectives. The power of the logical technique for the mathematician arises from the way the formal system of deduction can help organise a complex set of conditions that might be required in a mathematical construction or proof. The Completeness Theorem becomes a very general and powerful way of building interesting mathematical structures. A typical example is the application of first-order logic to construct number systems with infinitesimals that can used rigorously to present real calculus. This is the so-called nonstandard analysis of Abraham Robinson, and is presented in the last chapter of this book.

The mathematical content of completeness and soundness is well illustrated by König's Lemma on infinite finitely branching trees, and in the first chapter I discuss this. This is intended as a warm-up for the more difficult mathematics to come, and is a key example that I refer back to throughout the book.

Zorn's Lemma is essential for all the work in this book. I believe that by final year level, students should be starting to master straightforward applications of Zorn's Lemma. This is the main topic in Chapter 2. I do not shy away from the details, in particular giving a careful proof of Zorn's Lemma for countable posets, though the details of how Zorn's Lemma turns out to be equivalent to the Axiom of Choice is left for an optional section.

The idea of a formal system and derivations is introduced in Chapter 3, with a system based on strings of 0s and 1s that turns out to be closely related to König's Lemma. In the lecture theatre or classroom, I find this chapter to be particularly important and useful material, as it provides essential motivation for the Soundness Theorem. Given a comparatively simple system such as this, and asked whether a particular string σ can be derived from a set of assumptions Σ, students are all too ready to answer 'no' without justification. Where justification is offered, it is often of the kind, 'I tried to find a formal proof and this was my attempt, but it does not work.' So the idea of a careful proof by induction on the length of a formal derivation (and a carefully selected induction hypothesis) can be introduced and discussed without the additional complication of a long list of deduction rules to consider. The idea of semantics, and the Soundness and Completeness Theorems, arises from an investigation of general methods to show that certain derivations are not possible, and, to illustrate their power, König's Lemma is re-derived from the Soundness and Completeness Theorems for this system.

The reader will find systems with mathematically familiar derivations for the first time in Chapter 4. Building on previous material on posets, I develop a system for derivations about a poset, including rules such as 'if $a < b$ and

$b < c$ then $a < c$'. The system also has a way of expressing statements of the form 'a is not less than b', and this is handled using a Reductio Ad Absurdum Rule, a rule that is used throughout the rest of the book. By this stage, it should be clear what the right questions to ask about the system are, and the mathematical significance of the Completeness Theorem (the construction of a suitable partial order on a set) is clear. As a bonus, two pretty applications are provided: that any partial order can be 'linearised'; and that from a set of 'positive' assumptions a 'negative' conclusion can always be strengthened to a 'positive' one.

The material normally found in a more traditional course on mathematical logic starts with Chapter 5. Chapters 5 to 8 discuss boolean algebras and propositional logic. My proof system for propositional logic is chosen to be a form of natural deduction, but set out in a linear form on the page with clearly delineated 'subproofs' rather than a tree structure. This seems to be closest to a student's conception of a proof, and also provides clear proof strategies so that exercises in writing proofs can be given in a helpful and relatively painless way. (I emphasise the word 'relatively'. For most students, this aspect of logic is never painless, but at least the system clearly relates to informal proofs they might have written in other areas of mathematics.) I do not avoid explaining the precise connections between propositional logic and boolean algebra; these are important and elegant ideas, and are accessible to undergraduates who should be able to appreciate the analogies with algebra, especially rings and fields. More advanced students will also appreciate the fact that deep results such as Tychonov's Theorem and Stone Duality are only a few pages extra in an optional section. However, if time is short, the chapter on filters and ideals can be omitted entirely.

Chapters 9 and 10 are the central ones that cover first-order logic and the main Completeness Theorem. Apart from the choice of formal system (a development of the natural deduction system already used for propositional logic) they follow the usual pattern. These chapters are the longest in the book and will be found to be the most challenging so I have deliberately avoided many of the technically tricky issues such as: unique readability; the formal definition of the scope of a quantifier; or when a variable may be substituted by a term. An intelligent reader at this level using his or her basic mathematical training and intuition and following the examples is sure to do the 'right thing' and does not want to be bogged down in formal syntactic details. These technical details are of course important later on if one becomes involved in formalising logic in a first-order system such as set theory or arithmetic. But the place for that sort of work is certainly not a *first* course in logic. For those readers that need it, further details are available on the companion web-pages.

The method of proof of the Completeness Theorem is by 'Henkinising' the language and then using Zorn's Lemma to find a maximal consistent set of sentences. This is easier to describe to first-timers than tree-constructions of sets of consistent sentences with their required inductive properties, but is just as general and applicable. Two bonus optional sections for adventurous students with background in point-set topology include a topological view of the Compactness Theorem, and a proof of the full statement of the Omitting Types Theorem via Baire's Theorem, which is proved where needed.

Chapters 11 and 12 (which are independent of each other) provide applications of first-order logic. Chapter 11 presents an introduction to model theory, including the Löwenheim–Skolem Theorems, and (to put these in context) a short survey of categoricity, including a description of Morley's Theorem. This chapter is where infinite cardinals and cardinal arithmetic are used for the first time, and I carefully state all the required ideas and results before using them. Full proofs of these results are given in an optional section, using Zorn's Lemma only. The traditional options of using ordinals or the well-ordering principle are avoided as being likely to beg more questions than they answer to students without any prior knowledge in formal set theory. Chapter 12 presents an introduction to nonstandard analysis, including a proof of the Peano Existence Theorem on first-order differential equations. My presentation of nonstandard analysis is chosen to illustrate the main results of first-order logic and the interplay between the standard and nonstandard worlds, rather than to be optimal for fast proofs of classical results by nonstandard methods.

I have enjoyed writing this book and teaching from it. The material here is, to my mind, much more exciting and varied than the standard texts I learnt from as an undergraduate, and responses from the students who were given preliminary versions of these notes were good too. I can only hope that you, the reader, will derive a similar amount of pleasure from this book.

How to read this book

Chapters are deliberately kept as short as possible and discuss a single mathematical point. The chapters are divided into sections. The first section of each chapter is essential reading for all. The second section generally contains further applications, examples and exercises to test and expand on material presented in the previous section, and is again essential to read and explore. One or more extra 'starred' sections are then added to provide further commentary on the key material of the chapter and develop the material. These other sections are not essential reading and are intended for more inquisitive, ambitious or advanced readers with the background knowledge required. Chapter 8 may be omitted if time is short, and Chapters 11 and 12 are independent of each other.

Mathematical terminology is standard or explained in the text. Bold face entries in the index refer to definitions in the text; other entries provide further information on the term in question.

Additional material, including some technical definitions that I have chosen to omit in the printed text for the sake of clarity, further exercises, discussion, and some hints or answers to the exercises here, will be found on the companion web-site at http://web.mat.bham.ac.uk/R.W.Kaye/logic.

1

König's Lemma

1.1 Two ways of looking at mathematics

It seems that in mathematics there are sometimes two or more ways of proving the same result. This is often mysterious, and seems to go against the grain, for we often have a deep-down feeling that if we choose the 'right' ideas or definitions, there must be only one 'correct' proof. This feeling that there should be just one way of looking at something is rather similar to Paul Erdős's idea of 'The Book' [1], a vast tome held by God, the SF, in which all the best, most revealing and perfect proofs are written.

Sometimes this mystery can be resolved by analysing the apparently different proofs into their fundamental ideas. It often turns out that, 'underneath the bonnet', there is actually just one key mathematical concept, and two seemingly different arguments are in some sense 'the same'. But sometimes there really are two different approaches to a problem. This should not be disturbing, but should instead be seen as a great opportunity. After all, two approaches to the same idea indicates that there are some new mathematics to be investigated and some new connections to be found and exploited, which hopefully will uncover a wealth of new results.

I shall give a rather simple example of just the sort of situation I have in mind that will be familiar to many readers – one which will be typical of the kind of theorem we will be considering throughout this book.

Consider a binary *tree*. A tree is a diagram (often called a *graph*) with a special *point* or *node* called the *root*, and *lines* or *edges* leaving this node downwards to other nodes. These again may have edges leading to further nodes. The thing that makes this a tree (rather than a more general kind of graph) is that the edges all go downwards from the root, and that means the tree cannot have any *loops* or *cycles*. The tree is a *binary* tree if every node is connected to *at most* two lower nodes. If every node is connected to *exactly*

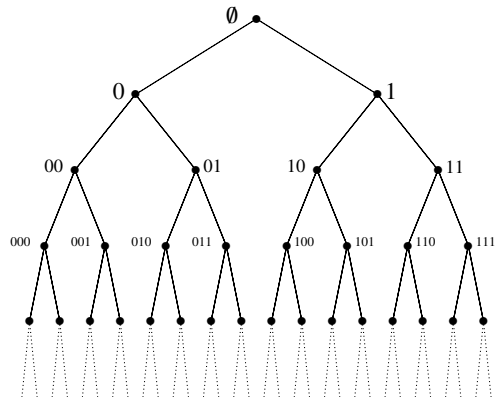

Figure 1.1 The full binary tree.

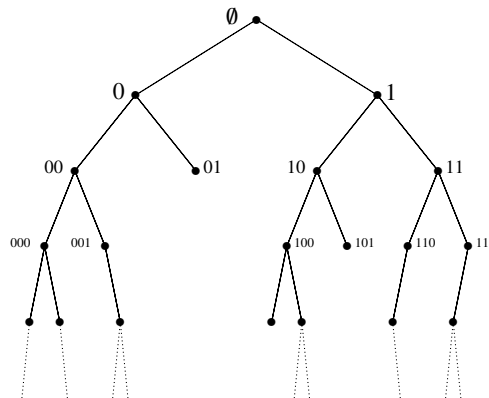

Figure 1.2 A binary tree.

two lower nodes, the tree is called the *full binary tree*. Note that in general, a node in a binary tree may be connected to 0, 1 or 2 lower nodes. We will label the nodes in our trees with sequences of integers. It is convenient to make labels for the nodes below the node that has label x by adding either the digit 0 or 1 to the end of x, giving $x0$ and $x1$. Figure 1.1 illustrates the full binary tree, whereas Figure 1.2 gives a typical (non-full) binary tree.

Trees are very important in mathematics, because many constructions follow trees in some way or other. Binary trees are especially interesting since a *walk* along a tree, following a path that starts at the root, has at most two choices of direction at every node. Binary trees arise quite naturally in many mathematical ideas and proofs and general theorems about them can be quite powerful and useful. One of the better known and more useful of these results is called König's Lemma.

To explain König's Lemma, consider what it means for a tree T to be *infinite*. There are two viewpoints, and two possible definitions.

Firstly, suppose you have somehow drawn the whole of the tree T on paper or on the blackboard and are inspecting it. You are in a fortunate position to be able to take in every one of its features, and to examine every one of its nodes and edges. You will quite naturally say that the tree is infinite if it has infinitely many nodes, or – amounting to the same thing – infinitely many edges. This is a sort of 'definition from perfect information' and is similar to what logicians call semantics, though we will not see the connection with semantics and the theory of 'meaning' for a while.

Now consider you are an ant walking on the binary tree T, which is again drawn in its entirety on paper. You start at the root node, and you follow the edges, like ant tracks, which you hope will take you to something interesting. Unlike the mathematician viewing the tree in its entirety, you can only see the node you are at and the edges leaving it. If you take a walk down the tree, you may have choices of turning left or right at any given node and continuing your path. But it is possible that you have no choice at all, because either there is only one edge out of the node other than the one you entered it by, or possibly there is no such edge at all, in which case your walk has come to an end. To the ant, which cannot perceive the whole of the tree, but just follows paths, there is a quite different idea of what it means for the tree to be infinite: the ant would say that T is infinite if it can find somehow (by guessing the right combination of 'left' and 'right' choices) an infinite path through the tree. The ant's definition of 'infinite' might be thought of as a 'definition from imperfect information' and is similar to the logician's idea of *proof*. If you like, you can think of an infinite path chosen by the ant as a *proof* that the tree is infinite. Like all proofs, it supports the claim made, without giving much extra information – such as what the tree looks like off this path.

König's Lemma is the statement that, for binary trees, these two ideas of a tree being infinite are the same. It is in fact a rather useful statement with many interesting applications. The key feature of this statement is that it relates two definitions, one mathematical definition working from perfect or total

information, and one working from the point of view of much more limited information, and shows that they actually say the same thing.

As with all 'if and only if' theorems, there are two directions that must be proved. The first, that if there is an infinite path through the tree then the tree is infinite, is immediate. This easier direction is called a *Soundness Theorem* since it says the ant's perception based on partial information is *sound*, or in other words will not result in erroneous conclusions. The other direction is the non-trivial one, and its mathematical strength lies in the way it states that a rather general mathematical situation (that the tree is infinite) can always be detected in a special way from partial information. The reason why it is called *Completeness* will be discussed later in relation to some other examples.

This has been a long preliminary discussion, but I hope it has proved illuminating. We shall now turn to the more formal mathematical details and define *tree*, *path*, etc., and then state and prove König's Lemma properly.

Definition 1.1 The set of *natural numbers*, $\mathbb{N}$, will be taken in this book to be $\{0, 1, 2, \ldots\}$.

For those readers who expect the natural numbers to start with 1, I can only say that I appreciate that there are occasions when it is convenient to forget about zero, but for me zero is very natural, probably the most logically natural number of all, so is included here in the set of natural numbers.

Definition 1.2 A *sequence* is a function s whose domain is either the set $\mathbb{N}$ of all natural numbers or a subset of it of the form $\{x \in \mathbb{N} : x < n\}$ for some $n \in \mathbb{N}$. Normally the values of the sequence will be numbers, 0 or 1 say, but the definition above (with $n = 0$) allows the empty sequence with no values at all. We write a sequence by listing its values in order, for example as 00110101001 or 0101010101. The *length* of a sequence is the number of elements in the domain of the function. This will always be a natural number or infinity.

Definition 1.3 If s is a sequence of length l and $n \in \mathbb{N}$ is at most l, then $s \restriction n$ denotes the initial part of s of length n.

For example, if $s = 00100011$ then $s \restriction 4 = 0010$.

Definition 1.4 If s is a sequence of length l and x is 0 or 1 then sx is the sequence of length $l + 1$ whose last element is x and all other elements agree with those of s.

Our definition of a tree is of a set of sequences that is closed under the restriction operation $\restriction$.

Definition 1.5 A *tree* is a set of sequences T such that for any $s \in T$ of length n and for any $l < n$ then $s \restriction l \in T$.

Think of a sequence $s \in T$ as a finite path starting from the root and arriving at some node. The individual digits in the sequence determine which choice of edge is made at each node. The set of nodes of the whole tree is then the set of sequences in the set T and two nodes $s, t \in T$ are connected by a single edge if one can be got from the other by adding a single number to the sequence. In other words, s and t are *connected* if $s \restriction (n-1) = t$ when s is the longer of the two and has length n, or the other way round if t is longer. Then the condition in the definition says, not unreasonably, that each node that this path passes through must also be in the tree. The root of the tree is the empty sequence of length 0.

Definition 1.6 A *subtree* of a tree T is a subset S of T that is a tree in its own right.

A subtree of a tree T might contain fewer nodes, and therefore fewer choices at certain nodes.

Definition 1.7 A *binary tree* is a tree T where all the sequences in it are functions from some $\{n \in \mathbb{N} : n < k\}$ to $\{0, 1\}$.

In other words, at any node, a path from the root of a binary tree has at most two options: to go left (0) or right (1). However, it may turn out that only one, or possibly neither, of these options is available at a particular node.

Definition 1.8 A tree T is *infinite* if it contains infinitely many sequences, or (equivalently) has infinitely many nodes.

A path is a subtree with no branching allowed. That means for any two nodes in the tree, one is a 'predecessor' of the other. More formally, we have the following definition.

Definition 1.9 A *path*, p, in a tree T is a subtree of T such that for any $s, t \in p$ with lengths n, k respectively and $n \leqslant k$, we have $s = t \restriction n$.

A tree T containing an infinite path p is obviously infinite. König's Lemma states that the converse is also true for binary trees.

Theorem 1.10 (König's Lemma) *Let T be an infinite binary tree. Then T contains an infinite path p.*

Proof Suppose T is an infinite binary tree. For a sequence s of length n let T_s be $\{r \in T : r \restriction n = s\} \cup \{s \restriction k : k < n\}$, which we will call the *subtree of T below s*. You will be able to check easily that T_s is a tree. In general it may or may not be infinite.

We are going to find a sequence $s(n)$ of elements of T such that

- $s(n)$ has length n,
- $s(n) = s(n+1) \restriction n$,
- the tree $T_{s(n)}$ below $s(n)$ is infinite.

This construction is by induction, using the third property above as our induction hypothesis. When we have completed the proof the set $\{s(n) : n \in \mathbb{N}\}$ will be our infinite path p in T.

So suppose inductively that we have chosen $s = s(n)$ of length n and T_s is infinite. Then since the tree is binary, made from sequences of 0s and 1s, we have

$$T_s = \{r \in T : r \restriction (n+1) = s0\} \cup \{r \in T : r \restriction (n+1) = s1\} \cup \{s \restriction k : k \leqslant n\}.$$

This is, by the induction hypothesis, infinite. Hence (as the third of these three sets is obviously finite) at least one of the first two sets, corresponding to '0' or '1' respectively, is infinite. If the first of these is infinite we set $s(n+1) = s0$ and in this case we have

$$T_{s(n+1)} = \{r \in T : r \restriction (n+1) = s0\} \cup \{s0\} \cup \{s \restriction k : k \leqslant n\}$$

which is infinite. If not we set $s(n+1) = s1$ which would then be infinite as before. Either way we have defined $s(n+1)$ and proved the induction hypothesis for $n+1$. □

1.2 Examples and exercises

The central theorem of this book, the Completeness Theorem for first-order logic, is not only of the same flavour as König's Lemma, but is in fact a powerful generalisation of it. To give you an idea of the power that this sort of theorem has, it is useful to see a selection of applications of König's Lemma here.

We start by exploring the limits of König's Lemma a little: it turns out that the important thing is not that there are at most two choices at each node but that the number of ways in which the branches divide is always finite.

Definition 1.11 If T is a tree and $s \in T$ is a node of T then the *valency* or *degree* of s is the number of nodes of T connected to s. Thus this is the number

of x such that $sx \in T$ plus one (to cater for the edge back towards the root), or just the number of such x if s is the root node.

Exercise 1.12 Prove the following generalisation of König's Lemma: an infinite tree in which every vertex has finite valency has an infinite path. Assume that the tree has vertices or nodes which are sequences of *natural numbers* of finite length and that for each $s \in T$ there is a bound $B_s \in \mathbb{N}$ on the possible values x such that $sx \in T$.

There are two ways that you might have done the last exercise. You might have modified the proof given above, or you may have tried to reduce the case of arbitrary finite valency trees to the case of binary trees by somehow 'encoding' arbitrary finite branching by a series of binary branches.

Exercise 1.13 Whichever method you used, have a go at proving the extension of König's Lemma by the other method.

Exercise 1.14 By giving an appropriate example of an infinite tree, show that König's Lemma is false for graphs with vertices of infinite valency.

König's Lemma is an elegant but nevertheless not very surprising or difficult result to see. Its truth, it seems, is reasonably clear, though a completely rigorous proof takes a moment or two to come up with. It is all the more surprising, therefore that there should be non-trivial applications. We will look at a few of these now, though nothing later in this book will depend on them.

Example 1.15 The set $X = [0, 1]$ has the property (called *sequential compactness*, equivalent to compactness for metric spaces) that every sequence (a_n) of elements of X has a subsequence converging to some element in X.

Proof Starting with $[0, 1]$ we continually divide intervals into equal halves, but at stage k of the construction we throw away any such interval that contains none of the a_n with $n > k$. More precisely, the nodes of the tree at depth k are identified with intervals $I = [(r-1)2^{-k}, r2^{-k}]$ for which $r \in \mathbb{N}$ and $\{a_n : n > k \text{ and } a_n \in I\}$ is non-empty, and two nodes are connected if one is a subset of the other.

This defines a binary tree. It is infinite because there are infinitely many a_n and each lies in an interval. By König's Lemma there is an infinite path through this tree, and by the construction of the tree we may take an infinite subsequence of a_n in this path, one at each level of the tree. This is the required convergent subsequence. □

Now consider infinite sequences $u_0u_1u_2\ldots$ of the digits $0, 1, 2, \ldots, k-1$. We will call such sequences *k-sequences*. Say a k-sequence s is *x^n-free* if there is no finite sequence, x, of digits $0, 1, 2, \ldots, k-1$, such that the finite sequence x^n (defined to be the result of repeating and concatenating x as $xxxx\ldots x$, where there are n copies of the string x) does not appear as a contiguous block of the sequence s.

Exercise 1.16 (a) Show that there is no x^2-free 2-sequence.

(b) Use König's Lemma to show that there is an x^3-free 2-sequence if and only if there are arbitrarily long finite x^3-free 2-sequences. State and prove a similar result for x^2-free 3-sequences.

(c) Define an operation on finite 2-sequences σ such that $\sigma(0) = 01$, $\sigma(1) = 10$, and $\sigma(u_0u_1\ldots u_m) = \sigma(u_0)\sigma(u_1)\ldots\sigma(u_m)$, where this is concatenation of sequences. Let $\sigma^n(s) = \sigma(\sigma(\ldots(\sigma(s))\ldots))$, i.e. σ iterated n times. Show that each of the finite sequences $\sigma^n(0)$ is x^3-free, and hence there is an infinite x^3-free 2-sequence.

(d) Show there is an x^2-free 3-sequence.

Another example of the use of König's Lemma is for graphs in the plane. A *graph* is a set V of vertices and a set E of edges, which are unordered subsets of V with exactly two vertices in each edge. In a *planar graph* the set of vertices V is a set of points of $\mathbb{R}^2$, and the edges joining vertices are lines which are 'smooth' (formed from finitely many straight-line segments) and may not cross except at a vertex.

A graph with set of vertices V can be *k-coloured* if there is a map $f\colon V \to \{0, 1, \ldots, k-1\}$ such that $f(u) \neq f(v)$ for all vertices u, v that are joined by an edge. You should think of the values $0, 1, \ldots, k-1$ as 'colours' of the vertices; the condition says two adjacent vertices must be coloured differently. Graph colourings, and especially colourings of planar graphs, are particularly interesting and have a long history [12]. A deep and difficult result by Appel and Haken shows that every finite planar graph is 4-colourable [10].

Exercise 1.17 Use König's Lemma to show that an infinite graph can be k-coloured if and only if every finite subgraph of it can be so coloured. (Make the simplification that the vertices of our infinite graph can be ordered as $v_0, v_1, \ldots$ with indices from $\mathbb{N}$. Construct a tree where the nodes at level n are all k-colourings of the subgraph with vertices $v_0, v_1, \ldots, v_{n-1}$, and edges join nodes if one colouring extends another.) Deduce from Appel and Haken's result that every infinite planar graph can be 4-coloured.

Tiling problems provide another nice application of König's Lemma. Con-