

John M. Harris
Jeffrey L. Hirst
Michael J. Mossinghoff



UNDERGRADUATE TEXTS IN MATHEMATICS

Combinatorics and Graph Theory

Second Edition

 Springer

Undergraduate Texts in Mathematics

Editors
S. Axler
K.A. Ribet

For other titles published in this series, go to
<http://www.springer.com/series/666>

John M. Harris • Jeffrey L. Hirst • Michael J. Mossinghoff

Combinatorics and Graph Theory

Second Edition

 Springer

John M. Harris
Department of Mathematics
Furman University
Greenville, SC 29613
USA
john.harris@furman.edu

Jeffrey L. Hirst
Mathematical Sciences
Appalachian State University
121 Bodenheimer Dr.
Boone, NC 28608
USA
jlh@math.appstate.edu

Michael J. Mossinghoff
Department of Mathematics
Davidson College
Box 6996
Davidson, NC 28035-6996
USA
mimossinghoff@davidson.edu

Editorial Board

S. Axler
Mathematics Department
San Francisco State University
San Francisco, CA 94132
USA
axler@sfsu.edu

K.A. Ribet
Department of Mathematics
University of California
at Berkeley
Berkeley, CA 94720
USA
ribet@math.berkeley.edu

ISSN: 0172-6056

ISBN: 978-0-387-797710-6

e-ISBN: 978-0-387-79711-3

DOI: 10.1007/978-0-387-79711-3

Library of Congress Control Number: 2008934034

Mathematics Subject Classification (2000): 05-01 03-01

© 2008 Springer Science+Business Media, LLC

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden. The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed on acid-free paper

springer.com

To
Priscilla, Sophie, and Will,
Holly,
Kristine, Amanda, and Alexandra

Preface to the Second Edition

There are certain rules that one must abide by in order to create a successful sequel.

— Randy Meeks, from the trailer to *Scream 2*

While we may not follow the precise rules that Mr. Meeks had in mind for successful sequels, we have made a number of changes to the text in this second edition. In the new edition, we continue to introduce new topics with concrete examples, we provide complete proofs of almost every result, and we preserve the book's friendly style and lively presentation, interspersing the text with occasional jokes and quotations. The first two chapters, on graph theory and combinatorics, remain largely independent, and may be covered in either order. Chapter 3, on infinite combinatorics and graphs, may also be studied independently, although many readers will want to investigate trees, matchings, and Ramsey theory for finite sets before exploring these topics for infinite sets in the third chapter. Like the first edition, this text is aimed at upper-division undergraduate students in mathematics, though others will find much of interest as well. It assumes only familiarity with basic proof techniques, and some experience with matrices and infinite series.

The second edition offers many additional topics for use in the classroom or for independent study. Chapter 1 includes a new section covering distance and related notions in graphs, following an expanded introductory section. This new section also introduces the adjacency matrix of a graph, and describes its connection to important features of the graph. Another new section on trails, circuits, paths, and cycles treats several problems regarding Hamiltonian and Eulerian paths in

graphs, and describes some elementary open problems regarding paths in graphs, and graphs with forbidden subgraphs.

Several topics were added to Chapter 2. The introductory section on basic counting principles has been expanded. Early in the chapter, a new section covers multinomial coefficients and their properties, following the development of the binomial coefficients. Another new section treats the pigeonhole principle, with applications to some problems in number theory. The material on Pólya's theory of counting has now been expanded to cover de Bruijn's more general method of counting arrangements in the presence of one symmetry group acting on the objects, and another acting on the set of allowed colors. A new section has also been added on partitions, and the treatment of Eulerian numbers has been significantly expanded. The topic of stable marriage is developed further as well, with three interesting variations on the basic problem now covered here. Finally, the end of the chapter features a new section on combinatorial geometry. Two principal problems serve to introduce this rich area: a nice problem of Sylvester's regarding lines produced by a set of points in the plane, and the beautiful geometric approach to Ramsey theory pioneered by Erdős and Szekeres in a problem about the existence of convex polygons among finite sets of points in the plane.

In Chapter 3, a new section develops the theory of matchings further by investigating marriage problems on infinite sets, both countable and uncountable. Another new section toward the end of this chapter describes a characterization of certain large infinite cardinals by using linear orderings. Many new exercises have also been added in each chapter, and the list of references has been completely updated.

The second edition grew out of our experiences teaching courses in graph theory, combinatorics, and set theory at Appalachian State University, Davidson College, and Furman University, and we thank these institutions for their support, and our students for their comments. We also thank Mark Spencer at Springer-Verlag. Finally, we thank our families for their patience and constant good humor throughout this process. The first and third authors would also like to add that, since the original publication of this book, their families have both gained their own second additions!

May 2008

John M. Harris
Jeffrey L. Hirst
Michael J. Mossinghoff

Preface to the First Edition

Three things should be considered: problems, theorems, and applications.

— Gottfried Wilhelm Leibniz,
Dissertatio de Arte Combinatoria, 1666

This book grew out of several courses in combinatorics and graph theory given at Appalachian State University and UCLA in recent years. A one-semester course for juniors at Appalachian State University focusing on graph theory covered most of Chapter 1 and the first part of Chapter 2. A one-quarter course at UCLA on combinatorics for undergraduates concentrated on the topics in Chapter 2 and included some parts of Chapter 1. Another semester course at Appalachian State for advanced undergraduates and beginning graduate students covered most of the topics from all three chapters.

There are rather few prerequisites for this text. We assume some familiarity with basic proof techniques, like induction. A few topics in Chapter 1 assume some prior exposure to elementary linear algebra. Chapter 2 assumes some familiarity with sequences and series, especially Maclaurin series, at the level typically covered in a first-year calculus course. The text requires no prior experience with more advanced subjects, such as group theory.

While this book is primarily intended for upper-division undergraduate students, we believe that others will find it useful as well. Lower-division undergraduates with a penchant for proofs, and even talented high school students, will be able to follow much of the material, and graduate students looking for an introduction to topics in graph theory, combinatorics, and set theory may find several topics of interest.

Chapter 1 focuses on the theory of finite graphs. The first section serves as an introduction to basic terminology and concepts. Each of the following sections presents a specific branch of graph theory: trees, planarity, coloring, matchings, and Ramsey theory. These five topics were chosen for two reasons. First, they represent a broad range of the subfields of graph theory, and in turn they provide the reader with a sound introduction to the subject. Second, and just as important, these topics relate particularly well to topics in Chapters 2 and 3.

Chapter 2 develops the central techniques of enumerative combinatorics: the principle of inclusion and exclusion, the theory and application of generating functions, the solution of recurrence relations, Pólya's theory of counting arrangements in the presence of symmetry, and important classes of numbers, including the Fibonacci, Catalan, Stirling, Bell, and Eulerian numbers. The final section in the chapter continues the theme of matchings begun in Chapter 1 with a consideration of the stable marriage problem and the Gale–Shapley algorithm for solving it.

Chapter 3 presents infinite pigeonhole principles, König's Lemma, Ramsey's Theorem, and their connections to set theory. The systems of distinct representatives of Chapter 1 reappear in infinite form, linked to the axiom of choice. Counting is recast as cardinal arithmetic, and a pigeonhole property for cardinals leads to discussions of incompleteness and large cardinals. The last sections connect large cardinals to finite combinatorics and describe supplementary material on computability.

Following Leibniz's advice, we focus on problems, theorems, and applications throughout the text. We supply proofs of almost every theorem presented. We try to introduce each topic with an application or a concrete interpretation, and we often introduce more applications in the exercises at the end of each section. In addition, we believe that mathematics is a fun and lively subject, so we have tried to enliven our presentation with an occasional joke or (we hope) interesting quotation.

We would like to thank the Department of Mathematical Sciences at Appalachian State University and the Department of Mathematics at UCLA. We would especially like to thank our students (in particular, Jae-II Shin at UCLA), whose questions and comments on preliminary versions of this text helped us to improve it. We would also like to thank the three anonymous reviewers, whose suggestions helped to shape this book into its present form. We also thank Sharon McPeake, a student at ASU, for her rendering of the Königsberg bridges.

In addition, the first author would like to thank Ron Gould, his graduate advisor at Emory University, for teaching him the methods and the joys of studying graphs, and for continuing to be his advisor even after graduation. He especially wants to thank his wife, Priscilla, for being his perfect match, and his daughter Sophie for adding color and brightness to each and every day. Their patience and support throughout this process have been immeasurable.

The second author would like to thank Judith Roitman, who introduced him to set theory and Ramsey's Theorem at the University of Kansas, using an early draft

of her fine text. Also, he would like to thank his wife, Holly (the other Professor Hirst), for having the infinite tolerance that sets her apart from the norm.

The third author would like to thank Bob Blakley, from whom he first learned about combinatorics as an undergraduate at Texas A & M University, and Donald Knuth, whose class *Concrete Mathematics* at Stanford University taught him much more about the subject. Most of all, he would like to thank his wife, Kristine, for her constant support and infinite patience throughout the gestation of this project, and for being someone he can always, well, count on.

September 1999

John M. Harris
Jeffrey L. Hirst
Michael J. Mossinghoff

Contents

Preface to the Second Edition	vii
Preface to the First Edition	ix
1 Graph Theory	1
1.1 Introductory Concepts	2
1.1.1 Graphs and Their Relatives	2
1.1.2 The Basics	5
1.1.3 Special Types of Graphs	10
1.2 Distance in Graphs	17
1.2.1 Definitions and a Few Properties	18
1.2.2 Graphs and Matrices	21
1.2.3 Graph Models and Distance	26
1.3 Trees	30
1.3.1 Definitions and Examples	31
1.3.2 Properties of Trees	34
1.3.3 Spanning Trees	38
1.3.4 Counting Trees	43
1.4 Trails, Circuits, Paths, and Cycles	51
1.4.1 The Bridges of Königsberg	52
1.4.2 Eulerian Trails and Circuits	55
1.4.3 Hamiltonian Paths and Cycles	60
1.4.4 Three Open Problems	67
1.5 Planarity	73

1.5.1	Definitions and Examples	74
1.5.2	Euler’s Formula and Beyond	78
1.5.3	Regular Polyhedra	80
1.5.4	Kuratowski’s Theorem	83
1.6	Colorings	85
1.6.1	Definitions	86
1.6.2	Bounds on Chromatic Number	88
1.6.3	The Four Color Problem	93
1.6.4	Chromatic Polynomials	97
1.7	Matchings	101
1.7.1	Definitions	102
1.7.2	Hall’s Theorem and SDRs	104
1.7.3	The König–Egerváry Theorem	109
1.7.4	Perfect Matchings	111
1.8	Ramsey Theory	116
1.8.1	Classical Ramsey Numbers	116
1.8.2	Exact Ramsey Numbers and Bounds	118
1.8.3	Graph Ramsey Theory	124
1.9	References	126
2	Combinatorics	129
2.1	Some Essential Problems	130
2.2	Binomial Coefficients	137
2.3	Multinomial Coefficients	144
2.4	The Pigeonhole Principle	150
2.5	The Principle of Inclusion and Exclusion	156
2.6	Generating Functions	164
2.6.1	Double Decks	166
2.6.2	Counting with Repetition	168
2.6.3	Changing Money	171
2.6.4	Fibonacci Numbers	177
2.6.5	Recurrence Relations	181
2.6.6	Catalan Numbers	185
2.7	Pólya’s Theory of Counting	190
2.7.1	Permutation Groups	191
2.7.2	Burnside’s Lemma	196
2.7.3	The Cycle Index	200
2.7.4	Pólya’s Enumeration Formula	202
2.7.5	de Bruijn’s Generalization	209
2.8	More Numbers	217
2.8.1	Partitions	218
2.8.2	Stirling Cycle Numbers	227
2.8.3	Stirling Set Numbers	231
2.8.4	Bell Numbers	237
2.8.5	Eulerian Numbers	242

2.9	Stable Marriage	248
2.9.1	The Gale–Shapley Algorithm	250
2.9.2	Variations on Stable Marriage	255
2.10	Combinatorial Geometry	264
2.10.1	Sylvester’s Problem	265
2.10.2	Convex Polygons	270
2.11	References	277
3	Infinite Combinatorics and Graphs	281
3.1	Pigeons and Trees	282
3.2	Ramsey Revisited	285
3.3	ZFC	290
3.3.1	Language and Logical Axioms	290
3.3.2	Proper Axioms	292
3.3.3	Axiom of Choice	297
3.4	The Return of der König	301
3.5	Ordinals, Cardinals, and Many Pigeons	304
3.5.1	Cardinality	304
3.5.2	Ordinals and Cardinals	308
3.5.3	Pigeons Finished Off	312
3.6	Incompleteness and Cardinals	318
3.6.1	Gödel’s Theorems for PA and ZFC	318
3.6.2	Inaccessible Cardinals	320
3.6.3	A Small Collage of Large Cardinals	322
3.7	Weakly Compact Cardinals	324
3.8	Infinite Marriage Problems	327
3.8.1	Hall and Hall	328
3.8.2	Countably Many Men	330
3.8.3	Uncountably Many Men	336
3.8.4	Espousable Cardinals	340
3.8.5	Perfect Matchings	343
3.9	Finite Combinatorics with Infinite Consequences	344
3.10	k -critical Linear Orderings	347
3.11	Points of Departure	348
3.12	References	352
	References	355
	Index	369

1

Graph Theory

“Begin at the beginning,” the King said, gravely, “and go on till you come to the end; then stop.”

— Lewis Carroll, *Alice in Wonderland*

The Pregolya River passes through a city once known as Königsberg. In the 1700s seven bridges were situated across this river in a manner similar to what you see in Figure 1.1. The city’s residents enjoyed strolling on these bridges, but, as hard as they tried, no resident of the city was ever able to walk a route that crossed each of these bridges exactly once. The Swiss mathematician Leonhard Euler learned of this frustrating phenomenon, and in 1736 he wrote an article [98] about it. His work on the “Königsberg Bridge Problem” is considered by many to be the beginning of the field of graph theory.




FIGURE 1.1. The bridges in Königsberg.

At first, the usefulness of Euler’s ideas and of “graph theory” itself was found only in solving puzzles and in analyzing games and other recreations. In the mid 1800s, however, people began to realize that graphs could be used to model many things that were of interest in society. For instance, the “Four Color Map Conjecture,” introduced by DeMorgan in 1852, was a famous problem that was seemingly unrelated to graph theory. The conjecture stated that four is the maximum number of colors required to color any map where bordering regions are colored differently. This conjecture can easily be phrased in terms of graph theory, and many researchers used this approach during the dozen decades that the problem remained unsolved.

The field of graph theory began to blossom in the twentieth century as more and more modeling possibilities were recognized — and the growth continues. It is interesting to note that as specific applications have increased in number and in scope, the theory itself has developed beautifully as well.

In Chapter 1 we investigate some of the major concepts and applications of graph theory. Keep your eyes open for the Königsberg Bridge Problem and the Four Color Problem, for we will encounter them along the way.

1.1 Introductory Concepts

A definition is the enclosing a wilderness of idea within a wall of words.

— Samuel Butler, *Higgledy-Piggledy*

1.1.1 Graphs and Their Relatives

A *graph* consists of two finite sets, V and E . Each element of V is called a *vertex* (plural *vertices*). The elements of E , called *edges*, are unordered pairs of vertices. For instance, the set V might be $\{a, b, c, d, e, f, g, h\}$, and E might be $\{\{a, d\}, \{a, e\}, \{b, c\}, \{b, e\}, \{b, g\}, \{c, f\}, \{d, f\}, \{d, g\}, \{g, h\}\}$. Together, V and E are a graph G .

Graphs have natural visual representations. Look at the diagram in Figure 1.2. Notice that each element of V is represented by a small circle and that each element of E is represented by a line drawn between the corresponding two elements of V .




FIGURE 1.2. A visual representation of the graph G .

As a matter of fact, we can just as easily define a graph to be a diagram consisting of small circles, called vertices, and curves, called edges, where each curve connects two of the circles together. When we speak of a graph in this chapter, we will almost always refer to such a diagram.

We can obtain similar structures by altering our definition in various ways. Here are some examples.

1. By replacing our set E with a set of *ordered* pairs of vertices, we obtain a *directed graph*, or *digraph* (Figure 1.3). Each edge of a digraph has a specific orientation.




FIGURE 1.3. A digraph.

2. If we allow repeated elements in our set of edges, technically replacing our set E with a multiset, we obtain a *multigraph* (Figure 1.4).




FIGURE 1.4. A multigraph.

3. By allowing edges to connect a vertex to itself (“loops”), we obtain a *pseudograph* (Figure 1.5).




FIGURE 1.5. A pseudograph.

4. Allowing our edges to be arbitrary subsets of vertices (rather than just pairs) gives us *hypergraphs* (Figure 1.6).




FIGURE 1.6. A hypergraph with 7 vertices and 5 edges.

5. By allowing V or E to be an infinite set, we obtain *infinite graphs*. Infinite graphs are studied in Chapter 3.

In this chapter we will focus on finite, simple graphs: those without loops or multiple edges.

Exercises

- Ten people are seated around a circular table. Each person shakes hands with everyone at the table except the person sitting directly across the table. Draw a graph that models this situation.
- Six fraternity brothers (Adam, Bert, Chuck, Doug, Ernie, and Filthy Frank) need to pair off as roommates for the upcoming school year. Each person has compiled a list of the people with whom he would be willing to share a room.

Adam's list: Doug
 Bert's list: Adam, Ernie
 Chuck's list: Doug, Ernie
 Doug's list: Chuck
 Ernie's list: Ernie
 Frank's list: Adam, Bert

Draw a digraph that models this situation.

- There are twelve women's basketball teams in the Atlantic Coast Conference: Boston College (B), Clemson (C), Duke (D), Florida State (F), Georgia Tech (G), Miami (I), NC State (S), Univ. of Maryland (M), Univ. of North Carolina (N), Univ. of Virginia (V), Virginia Tech (T), and Wake Forest Univ. (W). At a certain point in midseason,

B has played I, T*, W
 C has played D*, G

D has played C*, S, W
 F has played N*, V
 G has played C, M
 I has played B, M, T
 S has played D, V*
 M has played G, I, N
 N has played F*, M, W
 V has played F, S*
 T has played B*, I
 W has played B, D, N

The asterisk(*) indicates that these teams have played each other twice. Draw a multigraph that models this situation.

4. Can you explain why no resident of Königsberg was ever able to walk a route that crossed each bridge exactly once? (We will encounter this question again in Section 1.4.1.)

1.1.2 The Basics

Your first discipline is your vocabulary;

— Robert Frost

In this section we will introduce a number of basic graph theory terms and concepts. Study them carefully and pay special attention to the examples that are provided. Our work together in the sections that follow will be enriched by a solid understanding of these ideas.

The Very Basics

The vertex set of a graph G is denoted by $V(G)$, and the edge set is denoted by $E(G)$. We may refer to these sets simply as V and E if the context makes the particular graph clear. For notational convenience, instead of representing an edge as $\{u, v\}$, we denote this simply by uv . The *order* of a graph G is the cardinality of its vertex set, and the *size* of a graph is the cardinality of its edge set.

Given two vertices u and v , if $uv \in E$, then u and v are said to be *adjacent*. In this case, u and v are said to be the *end vertices* of the edge uv . If $uv \notin E$, then u and v are *nonadjacent*. Furthermore, if an edge e has a vertex v as an end vertex, we say that v is *incident* with e .

The *neighborhood* (or *open neighborhood*) of a vertex v , denoted by $N(v)$, is the set of vertices adjacent to v :

$$N(v) = \{x \in V \mid vx \in E\}.$$

The *closed neighborhood* of a vertex v , denoted by $N[v]$, is simply the set $\{v\} \cup N(v)$. Given a set S of vertices, we define the neighborhood of S , denoted by $N(S)$, to be the union of the neighborhoods of the vertices in S . Similarly, the closed neighborhood of S , denoted $N[S]$, is defined to be $S \cup N(S)$.

The *degree* of v , denoted by $\deg(v)$, is the number of edges incident with v . In simple graphs, this is the same as the cardinality of the (open) neighborhood of v . The *maximum degree* of a graph G , denoted by $\Delta(G)$, is defined to be

$$\Delta(G) = \max\{\deg(v) \mid v \in V(G)\}.$$

Similarly, the *minimum degree* of a graph G , denoted by $\delta(G)$, is defined to be

$$\delta(G) = \min\{\deg(v) \mid v \in V(G)\}.$$

The *degree sequence* of a graph of order n is the n -term sequence (usually written in descending order) of the vertex degrees.

Let's use the graph G in Figure 1.2 to illustrate some of these concepts: G has order 8 and size 9; vertices a and e are adjacent while vertices a and b are nonadjacent; $N(d) = \{a, f, g\}$, $N[d] = \{a, d, f, g\}$; $\Delta(G) = 3$, $\delta(G) = 1$; and the degree sequence is 3, 3, 3, 2, 2, 2, 2, 1.

The following theorem is often referred to as the First Theorem of Graph Theory.

Theorem 1.1. *In a graph G , the sum of the degrees of the vertices is equal to twice the number of edges. Consequently, the number of vertices with odd degree is even.*

Proof. Let $S = \sum_{v \in V} \deg(v)$. Notice that in counting S , we count each edge exactly twice. Thus, $S = 2|E|$ (the sum of the degrees is twice the number of edges). Since S is even, it must be that the number of vertices with odd degree is even. \square

Perambulation and Connectivity

A *walk* in a graph is a sequence of (not necessarily distinct) vertices v_1, v_2, \dots, v_k such that $v_i v_{i+1} \in E$ for $i = 1, 2, \dots, k-1$. Such a walk is sometimes called a v_1 - v_k *walk*, and v_1 and v_k are the *end vertices* of the walk. If the vertices in a walk are distinct, then the walk is called a *path*. If the edges in a walk are distinct, then the walk is called a *trail*. In this way, every path is a trail, but not every trail is a path. Got it?

A *closed path*, or *cycle*, is a path v_1, \dots, v_k (where $k \geq 3$) together with the edge $v_k v_1$. Similarly, a trail that begins and ends at the same vertex is called a *closed trail*, or *circuit*. The *length* of a walk (or path, or trail, or cycle, or circuit) is its number of edges, counting repetitions.

Once again, let's illustrate these definitions with an example. In the graph of Figure 1.7, a, c, f, c, b, d is a walk of length 5. The sequence b, a, c, b, d represents a trail of length 4, and the sequence d, g, b, a, c, f, e represents a path of length 6.




FIGURE 1.7.

Also, g, d, b, c, a, b, g is a circuit, while e, d, b, a, c, f, e is a cycle. In general, it is possible for a walk, trail, or path to have length 0, but the least possible length of a circuit or cycle is 3.

The following theorem is often referred to as the Second Theorem in this book.

Theorem 1.2. *In a graph G with vertices u and v , every u - v walk contains a u - v path.*

Proof. Let W be a u - v walk in G . We prove this theorem by induction on the length of W . If W is of length 1 or 2, then it is easy to see that W must be a path. For the induction hypothesis, suppose the result is true for all walks of length less than k , and suppose W has length k . Say that W is

$$u = w_0, w_1, w_2, \dots, w_{k-1}, w_k = v$$

where the vertices are not necessarily distinct. If the vertices are in fact distinct, then W itself is the desired u - v path. If not, then let j be the smallest integer such that $w_j = w_r$ for some $r > j$. Let W_1 be the walk

$$u = w_0, \dots, w_j, w_{r+1}, \dots, w_k = v.$$

This walk has length strictly less than k , and therefore the induction hypothesis implies that W_1 contains a u - v path. This means that W contains a u - v path, and the proof is complete. \square

We now introduce two different operations on graphs: *vertex deletion* and *edge deletion*. Given a graph G and a vertex $v \in V(G)$, we let $G - v$ denote the graph obtained by removing v and all edges incident with v from G . If S is a set of vertices, we let $G - S$ denote the graph obtained by removing each vertex of S and all associated incident edges. If e is an edge of G , then $G - e$ is the graph obtained by removing only the edge e (its end vertices stay). If T is a set of edges, then $G - T$ is the graph obtained by deleting each edge of T from G . Figure 1.8 gives examples of these operations.

A graph is *connected* if every pair of vertices can be joined by a path. Informally, if one can pick up an entire graph by grabbing just one vertex, then the




FIGURE 1.8. Deletion operations.




FIGURE 1.9. Connected and disconnected graphs.

graph is connected. In Figure 1.9, G_1 is connected, and both G_2 and G_3 are not connected (or *disconnected*). Each maximal connected piece of a graph is called a *connected component*. In Figure 1.9, G_1 has one component, G_2 has three components, and G_3 has two components.

If the deletion of a vertex v from G causes the number of components to increase, then v is called a *cut vertex*. In the graph G of Figure 1.8, vertex d is a cut vertex and vertex c is not. Similarly, an edge e in G is said to be a *bridge* if the graph $G - e$ has more components than G . In Figure 1.8, the edge ab is the only bridge.

A proper subset S of vertices of a graph G is called a *vertex cut set* (or simply, a *cut set*) if the graph $G - S$ is disconnected. A graph is said to be *complete* if every vertex is adjacent to every other vertex. Consequently, if a graph contains at least one nonadjacent pair of vertices, then that graph is not complete. Complete graphs do not have any cut sets, since $G - S$ is connected for all proper subsets S of the vertex set. Every non-complete graph has a cut set, though, and this leads us to another definition. For a graph G which is not complete, the *connectivity* of G , denoted $\kappa(G)$, is the minimum size of a cut set of G . If G is a connected, non-complete graph of order n , then $1 \leq \kappa(G) \leq n - 2$. If G is disconnected, then $\kappa(G) = 0$. If G is complete of order n , then we say that $\kappa(G) = n - 1$.

Further, for a positive integer k , we say that a graph is k -connected if $k \leq \kappa(G)$. You will note here that “1-connected” simply means “connected.”

Here are several facts that follow from these definitions. You will get to prove a couple of them in the exercises.

- i. A graph is connected if and only if $\kappa(G) \geq 1$.
- ii. $\kappa(G) \geq 2$ if and only if G is connected and has no cut vertices.
- iii. Every 2-connected graph contains at least one cycle.
- iv. For every graph G , $\kappa(G) \leq \delta(G)$.

Exercises

1. If G is a graph of order n , what is the maximum number of edges in G ?
2. Prove that for any graph G of order at least 2, the degree sequence has at least one pair of repeated entries.
3. Consider the graph shown in Figure 1.10.




FIGURE 1.10.

- (a) How many different paths have c as an end vertex?
 - (b) How many different paths avoid vertex c altogether?
 - (c) What is the maximum length of a circuit in this graph? Give an example of such a circuit.
 - (d) What is the maximum length of a circuit that does not include vertex c ? Give an example of such a circuit.
4. Is it true that a finite graph having exactly two vertices of odd degree must contain a path from one to the other? Give a proof or a counterexample.
 5. Let G be a graph where $\delta(G) \geq k$.
 - (a) Prove that G has a path of length at least k .
 - (b) If $k \geq 2$, prove that G has a cycle of length at least $k + 1$.

6. Prove that every closed odd walk in a graph contains an odd cycle.
7. Draw a connected graph having at most 10 vertices that has at least one cycle of each length from 5 through 9, but has no cycles of any other length.
8. Let P_1 and P_2 be two paths of maximum length in a connected graph G . Prove that P_1 and P_2 have a common vertex.
9. Let G be a graph of order n that is not connected. What is the maximum size of G ?
10. Let G be a graph of order n and size strictly less than $n - 1$. Prove that G is not connected.
11. Prove that an edge e is a bridge of G if and only if e lies on no cycle of G .
12. Prove or disprove each of the following statements.
 - (a) If G has no bridges, then G has exactly one cycle.
 - (b) If G has no cut vertices, then G has no bridges.
 - (c) If G has no bridges, then G has no cut vertices.
13. Prove or disprove: If every vertex of a connected graph G lies on at least one cycle, then G is 2-connected.
14. Prove that every 2-connected graph contains at least one cycle.
15. Prove that for every graph G ,
 - (a) $\kappa(G) \leq \delta(G)$;
 - (b) if $\delta(G) \geq n - 2$, then $\kappa(G) = \delta(G)$.
16. Let G be a graph of order n .
 - (a) If $\delta(G) \geq \frac{n-1}{2}$, then prove that G is connected.
 - (b) If $\delta(G) \geq \frac{n-2}{2}$, then show that G need not be connected.

1.1.3 Special Types of Graphs

until we meet again ...

— from *An Irish Blessing*

In this section we describe several types of graphs. We will run into many of them later in the chapter.

1. Complete Graphs

We introduced complete graphs in the previous section. A complete graph of order n is denoted by K_n , and there are several examples in Figure 1.11.




FIGURE 1.11. Examples of complete graphs.

2. Empty Graphs

The *empty graph* on n vertices, denoted by E_n , is the graph of order n where E is the empty set (Figure 1.12).




FIGURE 1.12. An empty graph.

3. Complements

Given a graph G , the *complement* of G , denoted by \overline{G} , is the graph whose vertex set is the same as that of G , and whose edge set consists of all the edges that are *not* present in G (Figure 1.13).




FIGURE 1.13. A graph and its complement.

4. Regular Graphs

A graph G is *regular* if every vertex has the same degree. G is said to be *regular of degree r* (or *r -regular*) if $\deg(v) = r$ for all vertices v in G . Complete graphs of order n are regular of degree $n - 1$, and empty graphs are regular of degree 0. Two further examples are shown in Figure 1.14.




FIGURE 1.14. Examples of regular graphs.

5. Cycles

The graph C_n is simply a cycle on n vertices (Figure 1.15).




FIGURE 1.15. The graph C_7 .

6. Paths

The graph P_n is simply a path on n vertices (Figure 1.16).



FIGURE 1.16. The graph P_6 .

7. Subgraphs

A graph H is a *subgraph* of a graph G if $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. In this case we write $H \subseteq G$, and we say that G contains H . In a graph where the vertices and edges are unlabeled, we say that $H \subseteq G$ if the vertices *could* be labeled in such a way that $V(H) \subseteq V(G)$ and $E(H) \subseteq E(G)$. In Figure 1.17, H_1 and H_2 are both subgraphs of G , but H_3 is not.

8. Induced Subgraphs

Given a graph G and a subset S of the vertex set, the *subgraph of G induced by S* , denoted $\langle S \rangle$, is the subgraph with vertex set S and with edge set $\{uv \mid u, v \in S \text{ and } uv \in E(G)\}$. So, $\langle S \rangle$ contains all vertices of S and all edges of G whose end vertices are *both* in S . A graph and two of its induced subgraphs are shown in Figure 1.18.




FIGURE 1.17. H_1 and H_2 are subgraphs of G .




FIGURE 1.18. A graph and two of its induced subgraphs.

9. Bipartite Graphs

A graph G is *bipartite* if its vertex set can be partitioned into two sets X and Y in such a way that every edge of G has one end vertex in X and the other in Y . In this case, X and Y are called the *partite sets*. The first two graphs in Figure 1.19 are bipartite. Since it is not possible to partition the vertices of the third graph into two such sets, the third graph is not bipartite.




FIGURE 1.19. Two bipartite graphs and one non-bipartite graph.

A bipartite graph with partite sets X and Y is called a *complete bipartite graph* if its edge set is of the form $E = \{xy \mid x \in X, y \in Y\}$ (that is, if

every possible connection of a vertex of X with a vertex of Y is present in the graph). Such a graph is denoted by $K_{|X|,|Y|}$. See Figure 1.20.




FIGURE 1.20. A few complete bipartite graphs.

The next theorem gives an interesting characterization of bipartite graphs.

Theorem 1.3. *A graph with at least two vertices is bipartite if and only if it contains no odd cycles.*

Proof. Let G be a bipartite graph with partite sets X and Y . Let C be a cycle of G and say that C is $v_1, v_2, \dots, v_k, v_1$. Assume without loss of generality that $v_1 \in X$. The nature of bipartite graphs implies then that $v_i \in X$ for all odd i , and $v_i \in Y$ for all even i . Since v_k is adjacent to v_1 , it must be that k is even; and hence C is an even cycle.

For the reverse direction of the theorem, let G be a graph of order at least two such that G contains no odd cycles. Without loss of generality, we can assume that G is connected, for if not, we could treat each of its connected components separately. Let v be a vertex of G , and define the set X to be

$$X = \{x \in V(G) \mid \text{the shortest path from } x \text{ to } v \text{ has even length}\},$$

and let $Y = V(G) \setminus X$.

Now let x and x' be vertices of X , and suppose that x and x' are adjacent. If $x = v$, then the shortest path from v to x' has length one. But this implies that $x' \in Y$, a contradiction. So, it must be that $x \neq v$, and by a similar argument, $x' \neq v$. Let P_1 be a path from v to x of shortest length (a shortest v - x path) and let P_2 be a shortest v - x' path. Say that P_1 is $v = v_0, v_1, \dots, v_{2k} = x$ and that P_2 is $v = w_0, w_1, \dots, w_{2t} = x'$. The paths P_1 and P_2 certainly have v in common. Let v' be a vertex on both paths such that the v' - x path, call it P'_1 , and the v' - x' path, call it P'_2 , have only the vertex v' in common. Essentially, v' is the "last" vertex common to P_1 and P_2 . It must be that P'_1 and P'_2 are shortest v' - x and v' - x' paths, respectively, and it must be that $v' = v_i = w_i$ for some i . But since x and x' are adjacent, $v_i, v_{i+1}, \dots, v_{2k}, w_{2t}, w_{2t-1}, \dots, w_i$ is a cycle of length $(2k - i) + (2t - i) + 1$, which is odd, and that is a contradiction.

Thus, no two vertices in X are adjacent to each other, and a similar argument shows that no two vertices in Y are adjacent to each other. Therefore, G is bipartite with partite sets X and Y . \square

We conclude this section with a discussion of what it means for two graphs to be the same. Look closely at the graphs in Figure 1.21 and convince yourself that one could be re-drawn to look just like the other. Even though these graphs




FIGURE 1.21. Are these graphs the same?

have different vertex sets and are drawn differently, it is still quite natural to think of these graphs as being the same. The idea of isomorphism formalizes this phenomenon.

Graphs G and H are said to be *isomorphic* to one another (or simply, isomorphic) if there exists a one-to-one correspondence $f : V(G) \rightarrow V(H)$ such that for each pair x, y of vertices of G , $xy \in E(G)$ if and only if $f(x)f(y) \in E(H)$. In other words, G and H are isomorphic if there exists a mapping from one vertex set to another that preserves adjacencies. The mapping itself is called an *isomorphism*. In our example, such an isomorphism could be described as follows:

$$\{(a, 1), (b, 2), (c, 8), (d, 3), (e, 7), (f, 4), (g, 6), (h, 5)\}.$$

When two graphs G and H are isomorphic, it is not uncommon to simply say that $G = H$ or that “ G is H .” As you will see, we will make use of this convention quite often in the sections that follow.

Several facts about isomorphic graphs are immediate. First, if G and H are isomorphic, then $|V(G)| = |V(H)|$ and $|E(G)| = |E(H)|$. The converse of this statement is not true, though, and you can see that in the graphs of Figure 1.22. The vertex and edge counts are the same, but the two graphs are clearly not iso-




FIGURE 1.22.

morphic.

A second necessary fact is that if G and H are isomorphic then the degree sequences must be identical. Again, the graphs in Figure 1.22 show that the converse of this statement is not true. A third fact, and one that you will prove in Exercise 8, is that if graphs G and H are isomorphic, then their complements \overline{G} and \overline{H} must also be isomorphic.

In general, determining whether two graphs are isomorphic is a difficult problem. While the question is simple for small graphs and for pairs where the vertex counts, edge counts, or degree sequences differ, the general problem is often tricky to solve. A common strategy, and one you might find helpful in Exercises 9 and 10, is to compare subgraphs, complements, or the degrees of adjacent pairs of vertices.

Exercises

1. For $n \geq 1$, prove that K_n has $n(n-1)/2$ edges.
2. If K_{r_1, r_2} is regular, prove that $r_1 = r_2$.
3. Determine whether K_4 is a subgraph of $K_{4,4}$. If yes, then exhibit it. If no, then explain why not.
4. Determine whether P_4 is an induced subgraph of $K_{4,4}$. If yes, then exhibit it. If no, then explain why not.
5. List all of the unlabeled connected subgraphs of C_{34} .
6. The concept of complete bipartite graphs can be generalized to define the *complete multipartite graph* K_{r_1, r_2, \dots, r_k} . This graph consists of k sets of vertices A_1, A_2, \dots, A_k , with $|A_i| = r_i$ for each i , where all possible “inter-set edges” are present and no “intra-set edges” are present. Find expressions for the order and size of K_{r_1, r_2, \dots, r_k} .
7. The *line graph* $L(G)$ of a graph G is defined in the following way: the vertices of $L(G)$ are the edges of G , $V(L(G)) = E(G)$, and two vertices in $L(G)$ are adjacent if and only if the corresponding edges in G share a vertex.
 - (a) Let G be the graph shown in Figure 1.23. Find $L(G)$.




FIGURE 1.23.

- (b) Find the complement of $L(K_5)$.
 - (c) Suppose G has n vertices, labeled v_1, \dots, v_n , and the degree of vertex v_i is r_i . Let m denote the size of G , so $r_1 + r_2 + \dots + r_n = 2m$. Find formulas for the order and size of $L(G)$ in terms of n , m , and the r_i .
8. Prove that if graphs G and H are isomorphic, then their complements \overline{G} and \overline{H} are also isomorphic.
9. Prove that the two graphs in Figure 1.24 are not isomorphic.




FIGURE 1.24.

10. Two of the graphs in Figure 1.25 are isomorphic.




FIGURE 1.25.

- (a) For the pair that is isomorphic, give an appropriate one-to-one correspondence.
- (b) Prove that the remaining graph is not isomorphic to the other two.

1.2 Distance in Graphs

'Tis distance lends enchantment to the view ...

— Thomas Campbell, *The Pleasures of Hope*

How far is it from one vertex to another? In this section we define distance in graphs, and we consider several properties, interpretations, and applications. Distance functions, called metrics, are used in many different areas of mathematics, and they have three defining properties. If M is a metric, then

- i. $M(x, y) \geq 0$ for all x, y , and $M(x, y) = 0$ if and only if $x = y$;
- ii. $M(x, y) = M(y, x)$ for all x, y ;
- iii. $M(x, y) \leq M(x, z) + M(z, y)$ for all x, y, z .

As you encounter the distance concept in the graph sense, verify for yourself that the function is in fact a metric.

1.2.1 Definitions and a Few Properties

I prefer the term ‘eccentric.’

— Brenda Bates, *Urban Legend*

Distance in graphs is defined in a natural way: in a connected graph G , the *distance* from vertex u to vertex v is the length (number of edges) of a shortest u - v path in G . We denote this distance by $d(u, v)$, and in situations where clarity of context is important, we may write $d_G(u, v)$. In Figure 1.26, $d(b, k) = 4$ and $d(c, m) = 6$.




FIGURE 1.26.

For a given vertex v of a connected graph, the *eccentricity* of v , denoted $\text{ecc}(v)$, is defined to be the greatest distance from v to any other vertex. That is,

$$\text{ecc}(v) = \max_{x \in V(G)} \{d(v, x)\}.$$

In Figure 1.26, $\text{ecc}(a) = 5$ since the farthest vertices from a (namely k, m, n) are at a distance of 5 from a .

Of the vertices in this graph, vertices c, k, m and n have the greatest eccentricity (6), and vertices e, f and g have the smallest eccentricity (3). These values and types of vertices are given special names. In a connected graph G , the *radius* of G , denoted $\text{rad}(G)$, is the value of the smallest eccentricity. Similarly, the *diameter* of G , denoted $\text{diam}(G)$, is the value of the greatest eccentricity. The *center* of the graph G is the set of vertices, v , such that $\text{ecc}(v) = \text{rad}(G)$. The *periphery* of G is the set of vertices, u , such that $\text{ecc}(u) = \text{diam}(G)$. In Figure 1.26, the radius is 3, the diameter is 6, and the center and periphery of the graph are, respectively, $\{e, f, g\}$ and $\{c, k, m, n\}$.

Surely these terms sound familiar to you. On a disk, the farthest one can travel from one point to another is the disk's diameter. Points on the rim of a disk are on the periphery. The distance from the center of the disk to any other point on the disk is at most the radius. The terms for graphs have similar meanings.

Do not be misled by this similarity, however. You may have noticed that the diameter of our graph G is twice the radius of G . While this does seem to be a natural relationship, such is not the case for all graphs. Take a quick look at a cycle or a complete graph. For either of these graphs, the radius and diameter are equal!

The following theorem describes the proper relationship between the radii and diameters of graphs. While not as natural, tight, or "circle-like" as you might hope, this relationship does have the advantage of being correct.

Theorem 1.4. *For any connected graph G , $\text{rad}(G) \leq \text{diam}(G) \leq 2 \text{rad}(G)$.*

Proof. By definition, $\text{rad}(G) \leq \text{diam}(G)$, so we just need to prove the second inequality. Let u and v be vertices in G such that $d(u, v) = \text{diam}(G)$. Further, let c be a vertex in the center of G . Then,

$$\text{diam}(G) = d(u, v) \leq d(u, c) + d(c, v) \leq 2 \text{ecc}(c) = 2 \text{rad}(G). \quad \square$$

The definitions in this section can also be extended to graphs that are not connected. In the context of a single connected component of a disconnected graph, these terms have their normal meanings. If two vertices are in different components, however, we say that the distance between them is infinity.

We conclude this section with two interesting results. Choose your favorite graph. It can be large or small, dense with edges or sparse. Choose anything you like, as long as it is your favorite. Now, wouldn't it be neat if there existed a graph in which your favorite graph was the "center" of attention? The next theorem (credited to Hedetniemi in [44]) makes your wish come true.

Theorem 1.5. *Every graph is (isomorphic to) the center of some graph.*

Proof. Let G be a graph (your favorite!). We now construct a new graph H (see Figure 1.27) by adding four vertices (w, x, y, z) to G , along with the following edges:

$$\{wx, yz\} \cup \{xa \mid a \in V(G)\} \cup \{yb \mid b \in V(G)\}.$$

Now, $\text{ecc}(w) = \text{ecc}(z) = 4$, $\text{ecc}(y) = \text{ecc}(x) = 3$, and for any vertex $v \in V(G)$,




FIGURE 1.27. G is the center.

$\text{ecc}(v) = 2$. Therefore, G is the center of H . □

Suppose you don't like being the center of attention. Maybe you would rather your favorite graph avoid the spotlight and stay on the periphery. The next theorem (due to Bielak and Sysło, [25]) tells us when that can happen.

Theorem 1.6. *A graph G is (isomorphic to) the periphery of some graph if and only if either every vertex has eccentricity 1, or no vertex has eccentricity 1.*

Proof. Suppose that every vertex of G has eccentricity 1. Not only does this mean that G is complete, it also means that every vertex of G is in the periphery. G is the periphery of itself!

On the other hand, suppose that no vertex of G has eccentricity 1. This means that for every vertex u of G , there is some vertex v of G such that $uv \notin E(G)$. Now, let H be a new graph, constructed by adding a single vertex, w , to G , together with the edges $\{wx \mid x \in V(G)\}$. In the graph H , the eccentricity of w is 1 (w is adjacent to everything). Further, for any vertex $x \in V(G)$, the eccentricity of x in H is 2 (no vertex of G is adjacent to everything else in G , and everything in G is adjacent to w). Thus, the periphery of H is precisely the vertices of G .

For the reverse direction, let us suppose that G has some vertices of eccentricity 1 and some vertices of eccentricity greater than 1. Suppose also (in anticipation of a contradiction) that G forms the periphery of some graph, say H . Since the eccentricities of the vertices in G are not all the same, it must be that $V(G)$ is a proper subset of $V(H)$. This means that H is not the periphery of itself and that $\text{diam}(H) \geq 2$. Now, let v be a vertex of G whose eccentricity in G is 1 (v is therefore adjacent to all vertices of G). Since $v \in V(G)$ and since G is the periphery of H , there exists a vertex w in H such that $d(v, w) = \text{diam}(H) \geq 2$. The vertex w , then, is also a peripheral vertex (see Exercise 4) and therefore must be in G . This contradicts the fact that v is adjacent to everything in G . \square

Exercises

1. Find the radius, diameter and center of the graph shown in Figure 1.28.




FIGURE 1.28.

2. Find the radius and diameter of each of the following graphs: P_{2k} , P_{2k+1} , C_{2k} , C_{2k+1} , K_n , $K_{m,n}$.
3. For each graph in Exercise 2, find the number of vertices in the center.
4. If x is in the periphery of G and $d(x, y) = \text{ecc}(x)$, then prove that y is in the periphery of G .

5. If u and v are adjacent vertices in a graph, prove that their eccentricities differ by at most one.
6. A graph G is called *self-centered* if $C(G) = V(G)$. Prove that every complete bipartite graph, every cycle, and every complete graph is self-centered.
7. Given a connected graph G and a positive integer k , the k th power of G , denoted G^k , is the graph with $V(G^k) = V(G)$ and where vertices u and v are adjacent in G^k if and only if $d_G(u, v) \leq k$.
 - (a) Draw the 2nd and 3rd powers of P_8 and C_{10} .
 - (b) For a graph G of order n , what is $G^{\text{diam}(G)}$?
8.
 - (a) Find a graph of order 7 that has radius 3 and diameter 6.
 - (b) Find a graph of order 7 that has radius 3 and diameter 5.
 - (c) Find a graph of order 7 that has radius 3 and diameter 4.
 - (d) Suppose r and d are positive integers and $r \leq d \leq 2r$. Describe a graph that has radius r and diameter d .
9. Suppose that u and v are vertices in a graph G , $\text{ecc}(u) = m$, $\text{ecc}(v) = n$, and $m < n$. Prove that $d(u, v) \geq n - m$. Then draw a graph G_1 where $d(u, v) = n - m$, and another graph G_2 where $d(u, v) > n - m$. In each case, label the vertices u and v , and give the values of m and n .
10. Let G be a connected graph with at least one cycle. Prove that G has at least one cycle whose length is less than or equal to $2 \text{diam}(G) + 1$.
11.
 - (a) Prove that if G is connected and $\text{diam}(G) \geq 3$, then \overline{G} is connected.
 - (b) Prove that if $\text{diam}(G) \geq 3$, then $\text{diam}(\overline{G}) \leq 3$.
 - (c) Prove that if G is regular and $\text{diam}(G) = 3$, then $\text{diam}(\overline{G}) = 2$.

1.2.2 Graphs and Matrices

Unfortunately no one can be told what the Matrix is. You have to see it for yourself.

— Morpheus, *The Matrix*

What do matrices have to do with graphs? This is a natural question — nothing we have seen so far has suggested any possible relationship between these two types of mathematical objects. That is about to change!

As we have seen, a graph is a very visual object. To this point, we have determined distances by looking at the diagram, pointing with our fingers, and counting edges. This sort of analysis works fairly well for small graphs, but it quickly breaks down as the graphs of interest get larger. Analysis of large graphs often requires computer assistance.

Computers cannot just look and point at graphs like we can. Instead, they understand graphs via matrix representations. One such representation is an adjacency matrix. Let G be a graph with vertices v_1, v_2, \dots, v_n . The *adjacency matrix* of G is the $n \times n$ matrix A whose (i, j) entry, denoted by $[A]_{i,j}$, is defined by

$$[A]_{i,j} = \begin{cases} 1 & \text{if } v_i \text{ and } v_j \text{ are adjacent,} \\ 0 & \text{otherwise.} \end{cases}$$

The graph in Figure 1.29 has six vertices. Its adjacency matrix, A , is

$$A = \begin{bmatrix} 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 1 & 0 \end{bmatrix}.$$




FIGURE 1.29.

Note that for simple graphs (where there are no loops) adjacency matrices have all zeros on the main diagonal. You can also see from the definition that these matrices are symmetric.¹

A single graph can have multiple adjacency matrices — different orderings of the vertices will produce different matrices. If you think that these matrices ought to be related in some way, then you are correct! In fact, if A and B are two different adjacency matrices of the same graph G , then there must exist a permutation of the vertices such that when the permutation is applied to the corresponding rows and columns of A , you get B .

This fact can be used in reverse to determine if two graphs are isomorphic, and the permutation mentioned here serves as an appropriate bijection: Given two graphs G_1 and G_2 with respective adjacency matrices A_1 and A_2 , if one can apply

¹Can you think of a context in which adjacency matrices might not be symmetric? Direct your attention to Figure 1.3 for a hint.

a permutation to the rows and columns of A_1 and produce A_2 , then G_1 and G_2 are isomorphic.

Let's take a closer look at the previous example. The fact that the $(1, 6)$ entry is 0 indicates that v_1 and v_6 are not adjacent. Consider now the $(1, 6)$ entry of the matrix A^2 . This entry is just the dot product of row one of A with column six of A :

$$\begin{aligned} [A^2]_{1,6} &= (0, 0, 0, 1, 1, 0) \cdot (0, 0, 1, 1, 1, 0) \\ &= (0 \cdot 0) + (0 \cdot 0) + (0 \cdot 1) + (1 \cdot 1) + (1 \cdot 1) + (0 \cdot 0) \\ &= 2. \end{aligned}$$

Think about what makes this dot product nonzero. It is the fact that there was at least one place (and here there were two places) where a 1 in row one corresponded with a 1 in column six. In our case, the 1 in the fourth position of row one (representing the edge v_1v_4) matched up with the 1 in the fourth position of column six (representing the edge v_4v_6). The same thing occurred in the fifth position of the row and column (where the edges represented were v_1v_5 and v_5v_6).

Can you see what is happening here? The entry in position $(1, 6)$ of A^2 is equal to the number of two-edge walks from v_1 to v_6 in G . As the next theorem shows us, this is not a coincidence.

Theorem 1.7. *Let G be a graph with vertices labeled v_1, v_2, \dots, v_n , and let A be its corresponding adjacency matrix. For any positive integer k , the (i, j) entry of A^k is equal to the number of walks from v_i to v_j that use exactly k edges.*

Proof. We prove this by induction on k . For $k = 1$, the result is true since $[A]_{i,j} = 1$ exactly when there is a one-edge walk between v_i and v_j .

Now suppose that for every i and j , the (i, j) entry of A^{k-1} is the number of walks from v_i to v_j that use exactly $k-1$ edges. For each k -edge walk from v_i to v_j , there exists an h such that the walk can be thought of as a $(k-1)$ -edge walk from v_i to v_h , combined with an edge from v_h to v_j . The total number of these k -edge walks, then, is

$$\sum_{v_h \in N(v_j)} (\text{number of } (k-1)\text{-edge walks from } v_i \text{ to } v_h).$$

By the induction hypothesis, we can rewrite this sum as

$$\sum_{v_h \in N(v_j)} [A^{k-1}]_{i,h} = \sum_{h=1}^n [A^{k-1}]_{i,h} [A]_{h,j} = [A^k]_{i,j},$$

and this proves the result. \square

This theorem has a straightforward corollary regarding distance between vertices.

Corollary 1.8. *Let G be a graph with vertices labeled v_1, v_2, \dots, v_n , and let A be its corresponding adjacency matrix. If $d(v_i, v_j) = x$, then $[A^k]_{i,j} = 0$ for $1 \leq k < x$.*

Let's see if we can relate these matrices back to earlier distance concepts. Given a graph G of order n with adjacency matrix A , and given a positive integer k , define the matrix sum S_k to be

$$S_k = I + A + A^2 + \cdots + A^k,$$

where I is the $n \times n$ identity matrix. Since the entries of I and A are ones and zeros, the entries of S_k (for any k) are nonnegative integers. This implies that for every pair i, j , we have $[S_k]_{i,j} \leq [S_{k+1}]_{i,j}$.

Theorem 1.9. *Let G be a connected graph with vertices labeled v_1, v_2, \dots, v_n , and let A be its corresponding adjacency matrix.*

1. *If k is the smallest positive integer such that row j of S_k contains no zeros, then $\text{ecc}(v_j) = k$.*
2. *If r is the smallest positive integer such that all entries of at least one row of S_r are positive, then $\text{rad}(G) = r$.*
3. *If m is the smallest positive integer such that all entries of S_m are positive, then $\text{diam}(G) = m$.*

Proof. We will prove the first part of the theorem. The proofs of the other parts are left for you as exercises.²

Suppose that k is the smallest positive integer such that row j of S_k contains no zeros. The fact that there are no zeros on row j of S_k implies that the distance from v_j to any other vertex is at most k . If $k = 1$, the result follows immediately. For $k > 1$, the fact that there is at least one zero on row j of S_{k-1} indicates that there is at least one vertex whose distance from v_j is greater than $k - 1$. This implies that $\text{ecc}(v_j) = k$. \square

We can use adjacency matrices to create other types of graph-related matrices. The steps given below describe the construction of a new matrix, using the matrix sums S_k defined earlier. Carefully read through the process, and (before you read the paragraph that follows!) see if you can recognize the matrix that is produced.

Creating a New Matrix, M

Given: A connected graph of order n , with adjacency matrix A , and with S_k as defined earlier.

1. For each $i \in \{1, 2, \dots, n\}$, let $[M]_{i,i} = 0$.
2. For each pair i, j where $i \neq j$, let $[M]_{i,j} = k$ where k is the least positive integer such that $[S_k]_{i,j} \neq 0$.

²You're welcome.

Can you see what the entries of M will be? For each pair i, j , the (i, j) entry of M is the distance from v_i to v_j . That is,

$$[M]_{i,j} = d(v_i, v_j).$$

The matrix M is called the *distance matrix* of the graph G .

Exercises

- Give the adjacency matrix for each of the following graphs.
 - P_{2k} and P_{2k+1} , where the vertices are labeled from one end of the path to the other.
 - C_{2k} and C_{2k+1} , where the vertices are labeled consecutively around the cycle.
 - $K_{m,n}$, where the vertices in the first partite set are labeled v_1, \dots, v_m .
 - K_n , where the vertices are labeled any way you please.
- Without computing the matrix directly, find A^3 where A is the adjacency matrix of K_4 .
- If A is the adjacency matrix for the graph G , show that the (j, j) entry of A^2 is the degree of v_j .
- Let A be the adjacency matrix for the graph G .
 - Show that the number of triangles that contain v_j is $\frac{1}{2}[A^3]_{j,j}$.
 - The *trace* of a square matrix M , denoted $\text{Tr}(M)$, is the sum of the entries on the main diagonal. Prove that the number of triangles in G is $\frac{1}{6} \text{Tr}(A^3)$.
- Find the $(1, 5)$ entry of A^{2009} where A is the adjacency matrix of C_{10} and where the vertices of C_{10} are labeled consecutively around the cycle.
- Prove the second statement in Theorem 1.9.
 - Prove the third statement in Theorem 1.9.
- Use Theorem 1.9 to design an algorithm for determining the center of a graph G .
- The graph G has adjacency matrix A and distance matrix D . Prove that if $A = D$, then G is complete.
- Give the distance matrices for the graphs in Exercise 1. You should create these matrices directly — it is not necessary to use the method described in the section.

1.2.3 Graph Models and Distance

Do I know you?

— Kevin Bacon, in *Flatliners*

We have already seen that graphs can serve as models for all sorts of situations. In this section we will discuss several models in which the idea of distance is significant.

The Acquaintance Graph

“Wow, what a small world!” This familiar expression often follows the discovery of a shared acquaintance between two people. Such discoveries are enjoyable, for sure, but perhaps the frequency with which they occur ought to keep us from being as surprised as we typically are when we experience them.

We can get a better feel for this phenomenon by using a graph as a model. Define the *Acquaintance Graph*, A , to be the graph where each vertex represents a person, and an edge connects two vertices if the corresponding people know each other. The context here is flexible — one could create this graph for the people living in a certain neighborhood, or the people working in a certain office building, or the people populating a country or the planet. Since the smaller graphs are all subgraphs of the graphs for larger populations, most people think of A in the largest sense: The vertices represent the Earth’s human population.³

An interesting question is whether or not the graph A , in the large (Earth) sense, is connected. Might there be a person or a group of people with no connection (direct *or* indirect) at all to another group of people?⁴ While there is a possibility of this being the case, it is most certainly true that if A is in fact disconnected, there is one *very* large connected component.

The graph A can be illuminating with regard to the “six degrees of separation” phenomenon. Made popular (at least in part) by a 1967 experiment by social psychologist Stanley Milgram [204] and a 1990 play by John Guare [142], the “six degrees theory” asserts that given any pair of people, there is a chain of no more than six acquaintance connections joining them. Translating into graph theorese, the assertion is that $\text{diam}(A) \leq 6$. It is, of course, difficult (if not impossible) to confirm this. For one, A is enormous, and the kind of computation required for confirmation is nontrivial (to say the least!) for matrices with six billion rows. Further, the matrix A is not static — vertices and edges appear all of the time.⁵ Still, the graph model gives us a good way to visualize this intriguing phenomenon.

Milgram’s experiment [204] was an interesting one. He randomly selected several hundred people from certain communities in the United States and sent a

³The graph could be made even larger by allowing the vertices to represent all people, living *or* dead. We will stick with the living people only — six billion vertices is large enough, don’t you think?

⁴Wouldn’t it be interesting to meet such a person? Wait — it wouldn’t be interesting for long because as soon as you meet him, he is no longer disconnected!

⁵Vertices will disappear if you limit A to living people. Edges disappear when amnesia strikes.

packet to each. Inside each packet was the name and address of a single “target” person. If the recipient knew this target personally, the recipient was to mail the packet directly to him. If the recipient did not know the target personally, the recipient was to send the packet to the person he/she thought had the best chance of knowing the target personally (perhaps someone in the same state as the target, or something like that). The new recipient was to follow the same rules: Either send it directly to the target (if known personally) or send it to someone who has a good chance of knowing the target. Milgram tracked how many steps it took for the packets to reach the target. Of the packets that eventually returned, the median number of steps was 5! Wow, what a small world!

The Hollywood Graph

Is the actor Kevin Bacon the center of Hollywood? This question, first asked by a group of college students in 1993, was the beginning of what was soon to become a national craze: The Kevin Bacon Game. The object of the game is to connect actors to Bacon through appearances in movies. For example, the actress Emma Thompson can be linked to Bacon in two steps: Thompson costarred with Gary Oldman in *Harry Potter and the Prisoner of Azkaban* (among others), and Oldman appeared with Bacon in *JFK*. Since Thompson has not appeared with Bacon in a movie, two steps is the best we can do. We say that Thompson has a *Bacon number* of 2.

Can you sense the underlying graph here?⁶ Let us define the *Hollywood Graph*, H , as follows: The vertices of H represent actors, and an edge exists between two vertices when the corresponding actors have appeared in a movie together. So, in H , Oldman is adjacent to both Bacon and Thompson, but Bacon and Thompson are not adjacent. Thompson has a Bacon number of 2 because the distance from her vertex to Bacon’s is 2. In general, an actor’s Bacon number is defined to be the distance from that actor’s vertex to Bacon’s vertex in H . If an actor cannot be linked to Bacon at all, then that actor’s Bacon number is infinity. As was the case with the Acquaintance Graph, if H is disconnected we can focus our attention on the single connected component that makes up most of H (Bacon’s component).

The ease with which Kevin Bacon can be connected to other actors might lead one to conjecture that Bacon is the unique center of Hollywood. In terms of graph theory, that conjecture would be that the center of H consists only of Bacon’s vertex. Is this true? Is Bacon’s vertex even *in* the center at all? Like the Acquaintance Graph, the nature of H changes frequently, and answers to questions like these are elusive. The best we can do is to look at a snapshot of the graph and answer the questions based on that particular point in time.

Let’s take a look at the graph as it appeared on December 25, 2007. On that day, the Internet Movie Database [165] had records for nearly 1.3 million actors. Patrick Reynolds maintains a website [234] that tracks Bacon numbers, among other things. According to Reynolds, of the 1.3 million actors in the database on

⁶or, “Can you smell the Bacon?”

that day, 917,007 could be linked to Bacon in some way via chains of shared movie appearances. The maximum distance from Bacon to any of the actors in his component was 8 (and so Bacon's eccentricity is 8). What about eccentricities of other actors? Are there any that are less than 8? According to Reynolds, the answer is no — 8 is the smallest eccentricity, and so Kevin Bacon is in the center of H . But it is very crowded there — thousands and thousands of other actors have eccentricity 8 as well.

The Mathematical Collaboration Graph

The Hungarian Paul Erdős (1913–1996) was one of the greatest and most prolific mathematicians of the twentieth century. Erdős authored or coauthored over 1500 mathematical papers covering topics in graph theory, combinatorics, set theory, geometry, number theory, and more. He collaborated with hundreds of other mathematicians, and this collaboration forms the basis of a Bacon-like ranking system. While not as widely popular as Bacon numbers, almost all mathematicians are familiar with the concept of Erdős numbers.

Erdős himself is assigned Erdős number 0. Any mathematician who coauthored a paper with Erdős has Erdős number 1. If a person has coauthored a paper with someone who has an Erdős number of 1 (and if that person himself/herself doesn't have Erdős number 1), then that person has an Erdős number of 2. Higher Erdős numbers are assigned in a similar manner.

The underlying graph here should be clear. Define the *Mathematical Collaboration Graph*, C , to have vertices corresponding to researchers, and let an edge join two researchers if the two have coauthored a paper together. A researcher's Erdős number, then, is the distance from the corresponding vertex to the vertex of Erdős. If a researcher is not in the same connected component of C as Erdős, then that researcher has an infinite Erdős number.

As you might imagine, new vertices and edges are frequently added to C . Jerry Grossman maintains a website [140] that keeps track of Erdős numbers. At one point in 2007, there were over 500 researchers with Erdős number 1 and over 8100 with Erdős number 2. You might surmise that because Erdős died in 1996, the number of people with Erdős number 1 has stopped increasing. While this is surely to be true sometime in the near future, it hasn't happened yet. A number of papers coauthored by Erdős have been published since his death. Erdős has not been communicating with collaborators from the great beyond (at least as far as we know) — it is simply the case that his collaborators continue to publish joint research that began years ago.

Small World Networks

As we saw earlier, the Acquaintance Graph provides a way to model the famous “small world phenomenon” — the sense that humans are connected via numerous recognized and unrecognized connections. The immense size and dynamic nature of that graph make it difficult to analyze carefully and completely, and so smaller models can prove to be more useful. In order for the more manageable graphs to

be helpful, though, it is important that they enjoy some fundamental small world properties.

So what makes a small world small? What properties should a graph have if it is to be a model of a small world? Let's list a few. As you read through the list below, think about your own acquaintance network and see if these properties make sense to you.

1. There should be plenty of mutual acquaintances (shared neighbors). If this were the only property, then complete graphs would surely fit the bill — lots of mutual neighbors there. A complete graph, though, is not a realistic model of acquaintances in the world.
2. The graph should be sparse in edges. In a realistic model, there should be relatively few edges compared to the number of vertices in the graph.
3. Distances between pairs of vertices should be relatively small. The *characteristic path length* of a graph G , denoted L_G , is the average distance between vertices, where the average is taken over all pairs of distinct vertices. In any graph of order n , there are $|E(K_n)|$ distinct pairs of vertices, and in Exercise 1 of Section 1.1.3, you showed that $|E(K_n)| = n(n-1)/2$. So for a graph G of order n ,

$$L_G = \frac{\sum_{u,v \in V(G)} d(u,v)}{|E(K_n)|} = \frac{2}{n(n-1)} \sum_{u,v \in V(G)} d(u,v).$$

One way of obtaining this value for a graph is to find the mean of the non-diagonal entries in the distance matrix of the graph.

4. There should be a reasonable amount of *clustering* in a small world graph. In actual acquaintance networks, there are a number of factors (geography, for instance) that create little clusters of vertices — small groups of vertices among which a larger than typical portion of edges exists. For example, there are likely to be many edges among the vertices that represent the people that live in your neighborhood.

Given a vertex v in a graph of order n , we define its *clustering coefficient*, denoted $cc(v)$, as follows (recall that $\langle N[v] \rangle$ is the subgraph induced by the closed neighborhood of v).

$$cc(v) = \frac{|E(\langle N[v] \rangle)|}{|E(K_{1+\deg(v)})|} = \frac{2|E(\langle N[v] \rangle)|}{(1 + \deg(v)) \deg(v)}.$$

For each vertex v , this is the percentage of edges that exist among the vertices in the closed neighborhood of v . For a graph G of order n , we define the *clustering coefficient of the graph* G , denoted $CC(G)$ to be the average of the clustering coefficients of the vertices of G . That is,

$$CC(G) = \frac{1}{n} \sum_{v \in V(G)} cc(v).$$

Small world networks have the property that characteristic path lengths are low and clustering coefficients are high. Graphs that have these properties can be used as models in the mathematical analyses of the small world phenomenon and its associated concepts. It is interesting to note that other well known networks have exhibited small world traits — the internet, electric power grids, and even neural networks are examples — and this increases even further the applicability of graph models.

Exercises

1. Compute the characteristic path length for each of each of the following graphs: P_{2k} , P_{2k+1} , C_{2k} , C_{2k+1} , K_n , $K_{m,n}$.
2. Compute the clustering coefficient for each of each of the following graphs: P_{2k} , P_{2k+1} , C_{2k} , C_{2k+1} , K_n , $K_{m,n}$.
3. (a) In the Acquaintance Graph, try to find a path from your vertex to the vertex of the President of the United States.
 (b) Your path from the previous question may not be your shortest such path. Prove that your actual distance from the President is at most one away from the shortest such distance to be found among your classmates.

Interesting Note: There are several contexts in which Bacon numbers can be calculated. While Bacon purists only use movie connections, others include shared appearances on television and in documentaries as well. Under these more open guidelines, the mathematician Paul Erdős actually has a Bacon number of 3! Erdős was the focus of the 1993 documentary *N is a Number* [63]. British actor Alec Guinness made a (very) brief appearance near the beginning of that film, and Guinness has a Bacon number of 2 (can you find the connections?). As far as we know, Bacon has not coauthored a research article with anyone who is connected to Erdős, and so while Erdős' Bacon number is 3, Bacon's Erdős number is infinity.

1.3 Trees

“O look at the trees!” they cried, “O look at the trees!”
 — Robert Bridges, *London Snow*

In this section we will look at the trees—but not the ones that sway in the wind or catch the falling snow. We will talk about graph-theoretic trees. Before moving on, glance ahead at Figure 1.30, and try to pick out which graphs are trees.

1.3.1 Definitions and Examples

Example, the surest method of instruction.

— Pliny the Younger

In Figure 1.30 graphs A , B , and E are trees, while graphs C and D are not.




FIGURE 1.30. Which ones are trees?

A *tree* is a connected graph that contains no cycles. Graph-theoretic trees resemble the trees we see outside our windows. For example, graph-theoretic trees do not have cycles, just as the branches of trees in nature do not split and rejoin. The descriptive terminology does not stop here.

Graph D in Figure 1.30 is not a tree; rather, it is a *forest*. A forest is a collection of one or more trees. A vertex of degree 1 in a tree is called a *leaf*.

As in nature, graph-theoretic trees come in many shapes and sizes. They can be thin (P_{10}) or thick ($K_{1,1000}$), tall (P_{1000}) or short (K_1 and K_2). Yes, even the graphs K_1 and K_2 are considered trees (they are certainly connected and acyclic). In the spirit of our arboreal terminology, perhaps we should call K_1 a *stump* and K_2 a *wig*!

While we are on the subject of small trees, we should count a few of them. It is clear that K_1 and K_2 are the only trees of order 1 and 2, respectively. A moment's thought will reveal that P_3 is the only tree of order 3. Figure 1.31 shows the different trees of order 6 or less.

Trees sprout up as effective models in a wide variety of applications. We mention a few brief examples.

Examples

1. Trees are useful for modeling the possible outcomes of an experiment. For example, consider an experiment in which a coin is flipped and a 6-sided die is rolled. The leaves in the tree in Figure 1.32 correspond to the outcomes in the probability space for this experiment.




FIGURE 1.31. Trees of order 6 or less.




FIGURE 1.32. Outcomes of a coin/die experiment.

2. Programmers often use tree structures to facilitate searches and sorts and to model the logic of algorithms. For instance, the logic for a program that finds the maximum of four numbers (w, x, y, z) can be represented by the tree shown in Figure 1.33. This type of tree is a *binary decision tree*.
3. Chemists can use trees to represent, among other things, saturated hydrocarbons—chemical compounds of the form C_nH_{2n+2} (propane, for example). The bonds between the carbon and hydrogen atoms are depicted in the trees of Figure 1.34. The vertices of degree 4 are the carbon atoms, and the leaves represent the hydrogen atoms.
4. College basketball fans will recognize the tree in Figure 1.35. It displays final results for the “Sweet 16” portion of the 2008 NCAA men’s basketball tournament. Each vertex represents a single game.




FIGURE 1.33. Logic of a program.




FIGURE 1.34. A few saturated hydrocarbons.

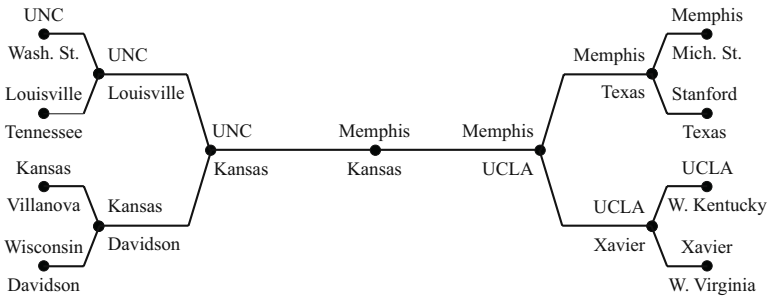


FIGURE 1.35. The 2008 Men's Sweet 16.

Exercises

1. Draw all unlabeled trees of order 7. Hint: There are a prime number of them.
2. Draw all unlabeled forests of order 6.
3. Let T be a tree of order $n \geq 2$. Prove that T is bipartite.

4. Graphs of the form $K_{1,n}$ are called *stars*. Prove that if $K_{r,s}$ is a tree, then it must be a star.
5. Match the graphs in Figure 1.36 with appropriate names: a palm tree, autumn, a path through a forest, tea leaves.




FIGURE 1.36. What would you name these graphs?

1.3.2 Properties of Trees

And the tree was happy.

— Shel Silverstein, *The Giving Tree*

Let us try an experiment. On a piece of scratch paper, draw a tree of order 16. Got one? Now count the number of edges in the tree. We are going to go out on a limb here and predict that there are 15. Since there are nearly 20,000 different trees of order 16, it may seem surprising that our prediction was correct. The next theorem gives away our secret.

Theorem 1.10. *If T is a tree of order n , then T has $n - 1$ edges.*

Proof. We induct on the order of T . For $n = 1$ the only tree is the stump (K_1), and it of course has 0 edges. Assume that the result is true for all trees of order less than k , and let T be a tree of order k .

Choose some edge of T and call it e . Since T is a tree, it must be that $T - e$ is disconnected (see Exercise 7) with two connected components that are trees themselves (see Figure 1.37). Say that these two components of $T - e$ are T_1 and T_2 , with orders k_1 and k_2 , respectively. Thus, k_1 and k_2 are less than n and $k_1 + k_2 = k$.

Since $k_1 < k$, the theorem is true for T_1 . Thus T_1 has $k_1 - 1$ edges. Similarly, T_2 has $k_2 - 1$ edges. Now, since $E(T)$ is the disjoint union of $E(T_1)$, $E(T_2)$, and $\{e\}$, we have $|E(T)| = (k_1 - 1) + (k_2 - 1) + 1 = k_1 + k_2 - 1 = k - 1$. This completes the induction. \square



FIGURE 1.37.

The next theorem extends the preceding result to forests. The proof is similar and appears as Exercise 4.

Theorem 1.11. *If F is a forest of order n containing k connected components, then F contains $n - k$ edges.*

The next two theorems give alternative methods for defining trees. Two other methods are given in Exercises 5 and 6.

Theorem 1.12. *A graph of order n is a tree if and only if it is connected and contains $n - 1$ edges.*

Proof. The forward direction of this theorem is immediate from the definition of trees and Theorem 1.10. For the reverse direction, suppose a graph G of order n is connected and contains $n - 1$ edges. We need to show that G is acyclic. If G did have a cycle, we could remove an edge from the cycle and the resulting graph would still be connected. In fact, we could keep removing edges (one at a time) from existing cycles, each time maintaining connectivity. The resulting graph would be connected and acyclic and would thus be a tree. But this tree would have fewer than $n - 1$ edges, and this is impossible by Theorem 1.10. Therefore, G has no cycles, so G is a tree. \square

Theorem 1.13. *A graph of order n is a tree if and only if it is acyclic and contains $n - 1$ edges.*

Proof. Again the forward direction of this theorem follows from the definition of trees and from Theorem 1.10. So suppose that G is acyclic and has $n - 1$ edges. To show that G is a tree we need to show only that it is connected. Let us say that the connected components of G are G_1, G_2, \dots, G_k . Since G is acyclic, each of these components is a tree, and so G is a forest. Theorem 1.11 tells us that G has $n - k$ edges, implying that $k = 1$. Thus G has only one connected component, implying that G is a tree. \square

It is not uncommon to look out a window and see leafless trees. In graph theory, though, leafless trees are rare indeed. In fact, the stump (K_1) is the only such tree, and every other tree has at least two leaves. Take note of the proof technique of the following theorem. It is a standard graph theory induction argument.

Theorem 1.14. *Let T be the tree of order $n \geq 2$. Then T has at least two leaves.*

Proof. Again we induct on the order. The result is certainly true if $n = 2$, since $T = K_2$ in this case. Suppose the result is true for all orders from 2 to $n - 1$, and consider a tree T of order $n \geq 3$. We know that T has $n - 1$ edges, and since we can assume $n \geq 3$, T has at least 2 edges. If every edge of T is incident with

a leaf, then T has at least two leaves, and the proof is complete. So assume that there is some edge of T that is not incident with a leaf, and let us say that this edge is $e = uv$. The graph $T - e$ is a pair of trees, T_1 and T_2 , each of order less than n . Let us say, without loss of generality, that $u \in V(T_1)$, $v \in V(T_2)$, $|V(T_1)| = n_1$, and $|V(T_2)| = n_2$ (see Figure 1.38). Since e is not incident with any leaves of T ,




FIGURE 1.38.

we know that n_1 and n_2 are both at least 2, so the induction hypothesis applies to each of T_1 and T_2 . Thus, each of T_1 and T_2 has two leaves. This means that each of T_1 and T_2 has at least one leaf that is not incident with the edge e . Thus the graph $(T - e) + e = T$ has at least two leaves. \square

We saw in the previous section that the center of a graph is the set of vertices with minimum eccentricity. The next theorem, due to Jordan [170], shows that for trees, there are only two possibilities for the center.

Theorem 1.15. *In any tree, the center is either a single vertex or a pair of adjacent vertices.*

Proof. Given a tree T , we form a sequence of trees as follows. Let $T_0 = T$. Let T_1 be the graph obtained from T_0 by deleting all of its leaves. Note here that T_1 is also a tree. Let T_2 be the tree obtained from T_1 by deleting all of the leaves of T_1 . In general, for as long as it is possible, let T_j be the tree obtained by deleting all of the leaves of T_{j-1} . Since T is finite, there must be an integer r such that T_r is either K_1 or K_2 .

Consider now a consecutive pair T_i, T_{i+1} of trees from the sequence $T = T_0, T_1, \dots, T_r$. Let v be a non-leaf of T_i . In T_i , the vertices that are at the greatest distance from v are leaves (of T_i). This means that the eccentricity of v in T_{i+1} is one less than the eccentricity of v in T_i . Since this is true for all non-leaves of T_i , it must be that the center of T_{i+1} is exactly the same as the center of T_i .

Therefore, the center of T_r is the center of T_{r-1} , which is the center of T_{r-2}, \dots , which is the center of $T_0 = T$. Since (the center of) T_r is either K_1 or K_2 , the proof is complete. \square

We conclude this section with an interesting result about trees as subgraphs.

Theorem 1.16. *Let T be a tree with k edges. If G is a graph whose minimum degree satisfies $\delta(G) \geq k$, then G contains T as a subgraph. Alternatively, G contains every tree of order at most $\delta(G) + 1$ as a subgraph.*

Proof. We induct on k . If $k = 0$, then $T = K_1$, and it is clear that K_1 is a subgraph of any graph. Further, if $k = 1$, then $T = K_2$, and K_2 is a subgraph of any graph whose minimum degree is 1. Assume that the result is true for all trees with $k - 1$ edges ($k \geq 2$), and consider a tree T with exactly k edges. We know from Theorem 1.14 that T contains at least two leaves. Let v be one of them, and let w be the vertex that is adjacent to v . Consider the graph $T - v$. Since $T - v$





FIGURE 1.39.

has $k - 1$ edges, the induction hypothesis applies, so $T - v$ is a subgraph of G . We can think of $T - v$ as actually sitting inside of G (meaning w is a vertex of G , too). Now, since G contains at least $k + 1$ vertices and $T - v$ contains k vertices, there exist vertices of G that are not a part of the subgraph $T - v$. Further, since the degree in G of w is at least k , there must be a vertex u not in $T - v$ that is adjacent to w (Figure 1.40). The subgraph $T - v$ together with u forms the tree T

FIGURE 1.40. A copy of T inside G .

as a subgraph of G . □

Exercises

1. Draw each of the following, if you can. If you cannot, explain the reason.
 - (a) A 10-vertex forest with exactly 12 edges
 - (b) A 12-vertex forest with exactly 10 edges
 - (c) A 14-vertex forest with exactly 14 edges
 - (d) A 14-vertex forest with exactly 13 edges
 - (e) A 14-vertex forest with exactly 12 edges
2. Suppose a tree T has an even number of edges. Show that at least one vertex must have even degree.
3. Let T be a tree with max degree Δ . Prove that T has at least Δ leaves.

4. Let F be a forest of order n containing k connected components. Prove that F contains $n - k$ edges.
5. Prove that a graph G is a tree if and only if for every pair of vertices u, v , there is exactly one path from u to v .
6. Prove that T is a tree if and only if T contains no cycles, and for any new edge e , the graph $T + e$ has exactly one cycle.
7. Show that every edge in a tree is a bridge.
8. Show that every nonleaf in a tree is a cut vertex.
9. Find a shorter proof to Theorem 1.14. Hint: Start by considering a longest path in T .
10. Let T be a tree of order $n > 1$. Show that the number of leaves is

$$2 + \sum_{\deg(v_i) \geq 3} (\deg(v_i) - 2),$$

where the sum is over all vertices of degree 3 or more.

11. For a graph G , define the average degree of G to be

$$\text{avgdeg}(G) = \frac{\sum_{v \in V(G)} \deg(v)}{|V(G)|}.$$

If T is a tree and $\text{avgdeg}(T) = a$, then find an expression for the number of vertices in T in terms of a .

12. Let T be a tree such that every vertex adjacent to a leaf has degree at least 3. Prove that some pair of leaves in T has a common neighbor.

1.3.3 Spanning Trees

Under the spreading chestnut tree . . .

— Henry W. Longfellow, *The Village Blacksmith*

The North Carolina Department of Transportation (NCDOT) has decided to implement a rapid rail system to serve eight cities in the western part of the state. Some of the cities are currently joined by roads or highways, and the state plans to lay the track right along these roads. Due to the mountainous terrain, some of the roads are steep and curvy; and so laying track along these roads would be difficult and expensive. The NCDOT hired a consultant to study the roads and to assign difficulty ratings to each one. The rating accounted for length, grade, and curviness of the roads; and higher ratings correspond to greater cost. The graph




FIGURE 1.41. The city graph.

in Figure 1.41, call it the “city graph,” shows the result of the consultant’s investigation. The number on each edge represents the difficulty rating assigned to the existing road.

The state wants to be able to make each city accessible (but not necessarily directly accessible) from every other city. One obvious way to do this is to lay track along every one of the existing roads. But the state wants to minimize cost, so this solution is certainly not the best, since it would result in a large amount of unnecessary track. In fact, the best solution will not include a cycle of track anywhere, since a cycle would mean at least one segment of wasted track.

The situation above motivates a definition. Given a graph G and a subgraph T , we say that T is a *spanning tree* of G if T is a tree that contains every vertex of G .

So it looks as though the DOT just needs to find a spanning tree of the city graph, and they would like to find one whose overall rating is as small as possible. Figure 1.42 shows several attempts at a solution.

Of the solutions in the figure, the one in the upper right has the least total weight—but is it the best one overall? Try to find a better one. We will come back to this problem soon.

Given a graph G , a *weight function* is a function W that maps the edges of G to the nonnegative real numbers. The graph G together with a weight function is called a *weighted graph*. The graph in Figure 1.41 is a simple example of a weighted graph. Although one might encounter situations where negative valued weights would be appropriate, we will stick with nonnegative weights in our discussion.

It should be fairly clear that every connected graph has at least one spanning tree. In fact, it is not uncommon for a graph to have many different spanning trees. Figure 1.42 displays three different spanning trees of the city graph.

Given a connected, weighted graph G , a spanning tree T is called a *minimum weight spanning tree* if the sum of the weights of the edges of T is no more than the sum for any other spanning tree of G .




FIGURE 1.42. Several spanning trees.

There are a number of fairly simple algorithms for finding minimum weight spanning trees. Perhaps the best known is Kruskal's algorithm.

Kruskal's Algorithm

Given: A connected, weighted graph G .

- i. Find an edge of minimum weight and mark it.
- ii. Among all of the unmarked edges that do not form a cycle with any of the marked edges, choose an edge of minimum weight and mark it.
- iii. If the set of marked edges forms a spanning tree of G , then stop. If not, repeat step ii.

Figure 1.43 demonstrates Kruskal's algorithm applied to the city graph. The minimum weight is 210.

It is certainly possible for different trees to result from two different applications of Kruskal's algorithm. For instance, in the second step we could have chosen the edge between Marion and Lenoir instead of the one that was chosen. Even so, the total weight of resulting trees is the same, and each such tree is a minimum weight spanning tree.




FIGURE 1.43. The stages of Kruskal's algorithm.

It should be clear from the algorithm itself that the subgraph built is in fact a spanning tree of G . How can we be sure, though, that it has minimum total weight? The following theorem answers our question [183].

Theorem 1.17. *Kruskal's algorithm produces a spanning tree of minimum total weight.*

Proof. Let G be a connected, weighted graph of order n , and let T be a spanning tree obtained by applying Kruskal's algorithm to G . As we have seen, Kruskal's algorithm builds spanning trees by adding one edge at a time until a tree is formed. Let us say that the edges added for T were (in order) e_1, e_2, \dots, e_{n-1} . Suppose T is not a minimum weight spanning tree. Among all minimum weight spanning trees of G , choose T' to be a minimum weight spanning tree that agrees with the construction of T for the longest time (i.e., for the most initial steps). This

means that there exists some k such that T' contains e_1, \dots, e_k , and no minimum weight spanning tree contains all of e_1, \dots, e_k, e_{k+1} (notice that since T is not of minimum weight, $k < n - 1$).

Since T' is a spanning tree, it must be that $T' + e_{k+1}$ contains a cycle C , and since T contains no cycles, C must contain some edge, say e' , that is *not* in T . If we remove the edge e' from $T' + e_{k+1}$, then the cycle C is broken and what remains is a spanning tree of G . Thus, $T' + e_{k+1} - e'$ is a spanning tree of G , and it contains edges e_1, \dots, e_k, e_{k+1} . Furthermore, since the edge e' must have been available to be chosen when e_{k+1} was chosen by the algorithm, it must be that $w(e_{k+1}) \leq w(e')$. This means that $T' + e_{k+1} - e'$ is a spanning tree with weight no more than T' that contains edges e_1, \dots, e_{k+1} , contradicting our assumptions. Therefore, it must be that T is a minimum weight spanning tree. \square

Exercises

1. Prove that every connected graph contains at least one spanning tree.
2. Prove that a graph is a tree if and only if it is connected and has exactly one spanning tree.
3. Let G be a connected graph with n vertices and at least n edges. Let C be a cycle of G . Prove that if T is a spanning tree of G , then \overline{T} , the complement of T , contains at least one edge of C .
4. Let G be connected, and let e be an edge of G . Prove that e is a bridge if and only if it is in every spanning tree of G .
5. Using Kruskal's algorithm, find a minimum weight spanning tree of the graphs in Figure 1.44. In each case, determine (with proof) whether the minimum weight spanning tree is unique.




FIGURE 1.44. Two weighted graphs.

6. Prim's algorithm (from [228]) provides another method for finding minimum weight spanning trees.

Prim's Algorithm

Given: A connected, weighted graph G .

- i. Choose a vertex v , and mark it.
- ii. From among all edges that have one marked end vertex and one unmarked end vertex, choose an edge e of minimum weight. Mark the edge e , and also mark its unmarked end vertex.
- iii. If every vertex of G is marked, then the set of marked edges forms a minimum weight spanning tree. If not, repeat step ii.

Use Prim's algorithm to find minimum weight spanning trees for the graphs in Figure 1.44. As you work, compare the stages to those of Kruskal's algorithm.

7. Give an example of a connected, weighted graph G having (i) a cycle with two identical weights, which is neither the smallest nor the largest weight in the graph, and (ii) a unique minimum weight spanning tree which contains exactly one of these two identical weights.

1.3.4 Counting Trees

As for everything else, so for a mathematical theory: beauty can be perceived but not explained.

— Arthur Cayley [214]

In this section we discuss two beautiful results on counting the number of spanning trees in a graph. The next chapter studies general techniques for counting arrangements of objects, so these results are a sneak preview.

Cayley's Tree Formula

Cayley's Tree Formula gives us a way to count the number of different labeled trees on n vertices. In this problem we think of the vertices as being fixed, and we consider all the ways to draw a tree on those fixed vertices. Figure 1.45 shows three different labeled trees on three vertices, and in fact, these are the only three.

There are 16 different labeled trees on four vertices, and they are shown in Figure 1.46.

As an exercise, the ambitious student should try drawing all of the labeled trees on five vertices. The cautious ambitious student might wish to look ahead at Cayley's formula before embarking on such a task.




FIGURE 1.45. Labeled trees on three vertices.




FIGURE 1.46. Labeled trees on four vertices.

Cayley proved the following theorem in 1889 [50]. The proof technique that we will describe here is due to Prüfer⁷ [229]. Prüfer’s method is almost as noteworthy as the result itself. He counted the labeled trees by placing them in one-to-one correspondence with a set whose size is easy to determine—the set of all sequences of length $n - 2$ whose entries come from the set $\{1, \dots, n\}$. There are n^{n-2} such sequences.

Theorem 1.18 (Cayley’s Tree Formula). *There are n^{n-2} distinct labeled trees of order n .*

The algorithm below gives the steps that Prüfer used to assign a particular sequence to a given tree, T , whose vertices are labeled $1, \dots, n$. Each labeled tree is assigned a unique sequence.

⁷With a name like that he was destined for mathematical greatness!

Prüfer's Method for Assigning a Sequence to a Labeled Tree

Given: A tree T , with vertices labeled $1, \dots, n$.

1. Let $i = 0$, and let $T_0 = T$.
2. Find the leaf on T_i with the smallest label and call it v .
3. Record in the sequence the label of v 's neighbor.
4. Remove v from T_i to create a new tree T_{i+1} .
5. If $T_{i+1} = K_2$, then stop. Otherwise, increment i by 1 and go back to step 2.

Let us run through this algorithm with a particular graph. In Figure 1.47, tree $T = T_0$ has 7 vertices, labeled as shown. The first step is finding the leaf with smallest label: This would be 2. The neighbor of vertex 2 is the vertex labeled 4. Therefore, 4 is the first entry in the sequence. Removing vertex 2 produces tree T_1 . The leaf with smallest label in T_1 is 4, and its neighbor is 3. Therefore, we put 3 in the sequence and delete 4 from T_1 . Vertex 5 is the smallest leaf in tree $T_2 = T_1 - \{4\}$, and its neighbor is 1. So our sequence so far is 4, 3, 1. In $T_3 = T_2 - \{5\}$ the smallest leaf is vertex 6, whose neighbor is 3. In $T_4 = T_3 - \{6\}$, the smallest leaf is vertex 3, whose neighbor is 1. Since $T_5 = K_2$, we stop here. Our resulting sequence is 4, 3, 1, 3, 1.

Notice that in the previous example, none of the leaves of the original tree T appears in the sequence. More generally, each vertex v appears in the sequence exactly $\deg(v) - 1$ times. This is not a coincidence (see Exercise 1). We now present Prüfer's algorithm for assigning trees to sequences. Each sequence gets assigned a unique tree.

Prüfer's Method for Assigning a Labeled Tree to a Sequence

Given: A sequence $\sigma = a_1, a_2, \dots, a_k$ of entries from the set $\{1, \dots, k+2\}$.

1. Draw $k+2$ vertices; label them v_1, v_2, \dots, v_{k+2} . Let $S = \{1, 2, \dots, k+2\}$.
2. Let $i = 0$, let $\sigma_0 = \sigma$, and let $S_0 = S$.
3. Let j be the smallest number in S_i that does not appear in the sequence σ_i .
4. Place an edge between vertex v_j and the vertex whose subscript appears first in the sequence σ_i .
5. Remove the first number in the sequence σ_i to create a new sequence σ_{i+1} . Remove the element j from the set S_i to create a new set S_{i+1} .
6. If the sequence σ_{i+1} is empty, place an edge between the two vertices whose subscripts are in S_{i+1} , and stop. Otherwise, increment i by 1 and return to step 3.




FIGURE 1.47. Creating a Prüfer sequence.

Let us apply this algorithm to a particular example. Let $\sigma = 4, 3, 1, 3, 1$ be our initial sequence to which we wish to assign a particular labeled tree. Since there are five terms in the sequence, our labels will come from the set $S = \{1, 2, 3, 4, 5, 6, 7\}$. After drawing the seven vertices, we look in the set $S = S_0$ to find the smallest subscript that does not appear in the sequence $\sigma = \sigma_0$. Subscript 2 is the one, and so we place an edge between vertices v_2 and v_4 , the first subscript in the sequence. We now remove the first term from the sequence and the label v_2 from the set, forming a new sequence $\sigma_1 = 3, 1, 3, 1$ and a new set $S_1 = \{1, 3, 4, 5, 6, 7\}$. The remaining steps in the process are shown in Figure 1.48.




FIGURE 1.48. Building a labeled tree.

You will notice that the tree that was created from the sequence σ in the second example is the very same tree that created the sequence σ in the first example. Score one for Prüfer!

Matrix Tree Theorem

The second major result that we present in this section is the Matrix Tree Theorem, and like Cayley's Theorem, it provides a way of counting spanning trees of labeled graphs. While Cayley's Theorem in essence gives us a count on the number of spanning trees of complete labeled graphs, the Matrix Tree Theorem applies to labeled graphs in general. The theorem was proved in 1847 by Kirchhoff [175], and it demonstrates a wonderful connection between spanning trees and matrices.

The theorem involves two special matrices. One is the adjacency matrix (defined back in Section 1.2.2), and the other is defined as follows. Let G be a graph with vertices v_1, v_2, \dots, v_n . The *degree matrix* of G is the $n \times n$ matrix D whose (i, j) entry, denoted by $[D]_{i,j}$, is defined by

$$[D]_{i,j} = \begin{cases} \deg(v_i) & \text{if } i = j, \\ 0 & \text{otherwise.} \end{cases}$$

So, the diagonal entries of D are the vertex degrees, and the off-diagonal entries are all zero.

Given an $n \times n$ matrix M , the i, j *cofactor* of M is defined to be

$$(-1)^{i+j} \det(M(i|j)),$$

where $\det(M(i|j))$ represents the determinant of the $(n - 1) \times (n - 1)$ matrix formed by deleting row i and column j from M .

We are now ready to state the Matrix Tree Theorem, due to Kirchhoff. The proof that we give imitates those presented in [148] and [52].

Theorem 1.19 (Matrix Tree Theorem). *If G is a connected labeled graph with adjacency matrix A and degree matrix D , then the number of unique spanning trees of G is equal to the value of any cofactor of the matrix $D - A$.*

Proof. Suppose G has n vertices (v_1, \dots, v_n) and k edges (f_1, \dots, f_k) . Since G is connected, we know that k is at least $n - 1$. Let N be the $n \times k$ matrix whose (i, j) entry is defined by

$$[N]_{i,j} = \begin{cases} 1 & \text{if } v_i \text{ and } f_j \text{ are incident,} \\ 0 & \text{otherwise.} \end{cases}$$

N is called the *incidence matrix* of G . Since every edge of G is incident with exactly two vertices of G , each column of N contains two 1's and $n - 2$ zeros. Let M be the $n \times k$ matrix that results from changing the topmost 1 in each column to -1 . To prove the result, we first need to establish two facts, which we call Claim A and Claim B.

Claim A. $MM^T = D - A$ (where M^T denotes the transpose of M).

First, notice that the (i, j) entry of $D - A$ is

$$[D - A]_{i,j} = \begin{cases} \deg(v_i) & \text{if } i = j, \\ -1 & \text{if } i \neq j \text{ and } v_i v_j \in E(G), \\ 0 & \text{if } i \neq j \text{ and } v_i v_j \notin E(G). \end{cases}$$

Now, what about the (i, j) entry of MM^T ? The rules of matrix multiplication tell us that this entry is the dot product of row i of M and column j of M^T . That is,

$$\begin{aligned} [MM^T]_{i,j} &= ([M]_{i,1}, [M]_{i,2}, \dots, [M]_{i,k}) \cdot ([M^T]_{1,j}, [M^T]_{2,j}, \dots, [M^T]_{k,j}) \\ &= ([M]_{i,1}, [M]_{i,2}, \dots, [M]_{i,k}) \cdot ([M]_{j,1}, [M]_{j,2}, \dots, [M]_{j,k}) \\ &= \sum_{r=1}^k [M]_{i,r} [M]_{j,r}. \end{aligned}$$

If $i = j$, then this sum counts one for every nonzero entry in row i ; that is, it counts the degree of v_i . If $i \neq j$ and $v_i v_j \notin E(G)$, then there is no column of M in which both the row i and row j entries are nonzero. Hence the value of the sum in this case is 0. If $i \neq j$ and $v_i v_j \in E(G)$, then the only column in which both the row i and the row j entries are nonzero is the column that represents the edge $v_i v_j$. Since one of these entries is 1 and the other is -1 , the value of the sum is -1 . We have shown that the (i, j) entry of MM^T is the same as the (i, j) entry of $D - A$, and thus Claim A is proved.

Let H be a subgraph of G with n vertices and $n - 1$ edges. Let p be an arbitrary integer between 1 and n , and let M' be the $(n - 1) \times (n - 1)$ submatrix of M formed by all rows of M except row p and the columns that correspond to the edges in H .

Claim B. If H is a tree, then $|\det(M')| = 1$. Otherwise, $\det(M') = 0$.

First suppose that H is not a tree. Since H has n vertices and $n - 1$ edges, we know from earlier work that H must be disconnected. Let H_1 be a connected component of H that does not contain the vertex v_p . Let M'' be the $|V(H_1)| \times (n - 1)$ submatrix of M' formed by eliminating all rows other than the ones corresponding to vertices of H_1 . Each column of M'' contains exactly two nonzero entries: 1 and -1 . Therefore, the sum of all of the row vectors of M'' is the zero vector, so the rows of M'' are linearly dependent. Since these rows are also rows of M' , we see that $\det(M') = 0$.

Now suppose that H is a tree. Choose some leaf of H that is not v_p (Theorem 1.14 lets us know that we can do this), and call it u_1 . Let us also say that e_1 is the edge of H that is incident with u_1 . In the tree $H - u_1$, choose u_2 to be some leaf other than v_p . Let e_2 be the edge of $H - u_1$ incident with u_2 . Keep removing leaves in this fashion until v_p is the only vertex left. Having established the list of vertices u_1, u_2, \dots, u_{n-1} , we now create a new $(n - 1) \times (n - 1)$ matrix M^* by rearranging the rows of M' in the following way: row i of M^* will be the row of M' that corresponds to the vertex u_i .

An important (i.e., useful!) property of the matrix M^* is that it is lower triangular (we know this because for each i , vertex u_i is *not* incident with any of $e_{i+1}, e_{i+2}, \dots, e_{n-1}$). Thus, the determinant of M^* is equal to the product of the main diagonal entries, which are either 1 or -1 , since every u_i is incident with e_i . Thus, $|\det(M^*)| = 1$, and so $|\det(M')| = 1$. This proves Claim B.

We are now ready to investigate the cofactors of $D - A = MM^T$. It is a fact from matrix theory that if the row sums and column sums of a matrix are all 0, then the cofactors all have the same value. (It would be a nice exercise—and a nice review of matrix skills—for you to try to prove this.) Since the matrix MM^T satisfies this condition, we need to consider only one of its cofactors. We might as well choose i and j such that $i + j$ is even—let us choose $i = 1$ and $j = 1$. So, the $(1, 1)$ cofactor of $D - A$ is

$$\begin{aligned} \det((D - A)(1|1)) &= \det(MM^T(1|1)) \\ &= \det(M_1 M_1^T) \end{aligned}$$

where M_1 is the matrix obtained by deleting the first row of $D - A$.

At this point we make use of the Cauchy–Binet Formula, which says that the determinant above is equal to the sum of the determinants of $(n - 1) \times (n - 1)$ submatrices of M_1 (for a more thorough discussion of the Cauchy–Binet Formula, see [40]). We have already seen (in Claim B) that any $(n - 1) \times (n - 1)$ submatrix that corresponds to a spanning tree of G will contribute 1 to the sum, while all others contribute 0. This tells us that the value of $\det(D - A) = \det(MM^T)$ is precisely the number of spanning trees of G . \square

Figure 1.49 shows a labeled graph G and each of its eight spanning trees.




FIGURE 1.49. A labeled graph and its spanning trees.

The degree matrix D and adjacency matrix A are

$$D = \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 3 & 0 \\ 0 & 0 & 0 & 3 \end{bmatrix}, \quad A = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{bmatrix},$$

and so

$$D - A = \begin{bmatrix} 2 & 0 & -1 & -1 \\ 0 & 2 & -1 & -1 \\ -1 & -1 & 3 & -1 \\ -1 & -1 & -1 & 3 \end{bmatrix}.$$

The $(1, 1)$ cofactor of $D - A$ is

$$\det \begin{bmatrix} 2 & -1 & -1 \\ -1 & 3 & -1 \\ -1 & -1 & 3 \end{bmatrix} = 8.$$

Score one for Kirchhoff!

Exercises

1. Let T be a labeled tree. Prove that the Prüfer sequence of T will not contain any of the leaves' labels. Also prove that each vertex v will appear in the sequence exactly $\deg(v) - 1$ times.
2. Determine the Prüfer sequence for the trees in Figure 1.50.




FIGURE 1.50. Two labeled trees.

3. Draw and label a tree whose Prüfer sequence is

$$5, 4, 3, 5, 4, 3, 5, 4, 3.$$
4. Which trees have constant Prüfer sequences?
5. Which trees have Prüfer sequences with distinct terms?
6. Let e be an edge of K_n . Use Cayley's Theorem to prove that $K_n - e$ has $(n - 2)n^{n-3}$ spanning trees.
7. Use the Matrix Tree Theorem to prove Cayley's Theorem. Hint: Look back at the discussion prior to the statement of the Matrix Tree Theorem.

1.4 Trails, Circuits, Paths, and Cycles

Takes a real salesman, I can tell you that. Anvils have a limited appeal, you know.

— Charlie Cowell, anvil salesman, *The Music Man*

Charlie Cowell was a door to door anvil salesman, and he dragged his heavy wares down every single street in each town he visited. Not surprisingly, Charlie became quite proficient at designing routes that did not repeat many streets. He certainly did not want to drag the anvils any farther than necessary, and he especially liked it when he could cover every street in the town without repeating a single one.

After several years of unsuccessful sales (he saw more closed doors than closed deals), the Acme Anvil Company did the natural thing — they promoted him. Charlie moved from salesman to regional supplier. This meant that Charlie would

be in charge of driving the anvil truck from town to town, delivering each town's supply of anvils. Still efficiency-minded, he wanted to plan driving routes that did not repeat any town along the way. He had been very good at route planning during his door to door days, and avoiding the repetition of towns was basically the same as avoiding the repetition of streets, right?

As you read through this section, see if you can answer that question for yourself.

1.4.1 The Bridges of Königsberg

One should make a serious study of a pastime.

— Alexander the Great

At the very beginning of this chapter, we referred to the legendary Königsberg Bridge Problem. As you will recall, this problem concerned the existence (or non-existence) of a certain type of route across a group of bridges (Figure 1.1). Could one design a route that crossed each bridge exactly once? The residents of seventeenth and eighteenth century Königsberg passed the time making valiant efforts, but no route could be found.

In 1736, the Swiss mathematician Euler addressed the problem ([98], translated in [26]). Near the beginning of his article, Euler described his thoughts as he embarked on the search for a solution.

As far as the problem of the seven bridges of Königsberg is concerned, it can be solved by making an exhaustive list of all possible routes, and then finding whether or not any route satisfies the conditions of the problem. Because of the number of possibilities, this method of solution would be too difficult and laborious, and in other problems with more bridges it would be impossible . . . Hence I rejected it, and looked for another method concerned only with the problem of whether or not the specified route could be found; I considered that such a method would be much simpler. [26]

This passage is enlightening on several levels. For one, it gives us a glimpse of the teacher/expositor side of the master mathematician Euler. It is doubtful that he would have seriously considered listing all possible routes in search of a satisfactory one. His mention of the possibility, though, does help the reader progress along a natural arc of thought regarding the solution. The passage also gives a clue as to what Euler is really after — not just a solution to the problem in Königsberg, but a general solution that could be applied in other land/bridge formations.

Using a figure similar to Figure 1.51, he used sequences of letters to describe routes — routes where no bridges were repeated. For instance, the sequence *ABDACAB* represented a route that started at *A*, crossed a bridge to *B*, crossed a bridge to *D*, crossed a bridge back to *A*, crossed a bridge to *C*, crossed a bridge back to *A*, and then crossed a bridge to *B*. This seven letter sequence includes




FIGURE 1.51. Layout of the bridges in Königsberg.

six of the seven bridges in Königsberg, and these bridges can be identified by the consecutive pairs in the sequence: AB, BD, DA, AC, CA, AB .

For Euler the Königsberg Bridge Problem boiled down to finding a certain sequence of letters. He described it in this way:

The problem is therefore reduced to finding a sequence of eight letters, formed from the four letters A, B, C , and D , in which the various pairs of letters occur the required number of times. Before I turn to the problem of finding such a sequence, it would be useful to find out whether or not it is even possible to arrange the letters in this way, for if it were possible to show that there is no such arrangement, then any work directed towards finding it would be wasted. I have therefore tried to find a rule which will be useful in this case, and in others, for determining whether or not such an arrangement can exist. [26]

Euler argued that since land area D is connected to three bridges, then D must appear in the sequence two times. (If it appeared only once, this would not account for all of the bridges; if it appeared more than twice, this would represent multiple crossing of at least one bridge.) Similarly, Euler argued, B and C must each appear twice in the eight letter sequence. Further, since land area A is connected to five bridges, the letter A must appear three times in the sequence (you will verify this in Exercise 1a). This means that the necessary *eight* letter sequence would have three A s, two B s, two C s and two D s. In Euler's words,

It follows that such a journey cannot be undertaken across the seven bridges of Königsberg. [26]

Once he had settled the problem of the Königsberg bridges, Euler used the same ideas and methods to present a more general result, and we will see that result in the next section. As you see, Euler did not use terms like *graph*, *vertex* or *edge*. Today's graph terminology did not appear until many years later. Still, this article by Euler was the seed from which the field of graph theory grew.

Euler himself recognized that he was working in relatively uncharted territory. We close this section with the passage with which Euler opened his seminal article. See if you can read "graph theory" between the lines.

In addition to that branch of geometry which is concerned with magnitudes, and which has always received the greatest attention, there is another branch, previously almost unknown, which Leibniz first mentioned, calling it the *geometry of position*. This branch is concerned only with the determination of position and its properties; it does not involve measurements, nor calculations made with them. It has not yet been satisfactorily determined what kind of problems are relevant to this geometry of position, or what methods would be used in solving them. Hence, when a problem was recently mentioned [the Königsberg problem!], which seemed geometrical but was so constructed that it did not require the measurement of distances, nor did calculation help at all, I had no doubt that it was concerned with the geometry of position — especially as its solution involved only position, and no calculation was of any use. I have therefore decided to give here the method which I have found for solving this kind of problem, as an example of the geometry of position. [26]

Mathematics is richer because Euler took up the problem of the Königsberg bridges. We are grateful, but perhaps not as grateful as the residents of Königsberg whose feet had grown tired from the search for a route that did not exist.

Exercises

1. In the context of Euler's letter sequences, prove that . . .
 - (a) if a land mass L is connected to 5 bridges, then L will occur 3 times in any representation of a route that crosses all of the bridges once.
 - (b) if a land mass L is connected to n bridges, where n is odd, then L will occur $\frac{n+1}{2}$ times in any representation of a route that crosses all of the bridges once.
2. An eighth bridge was built in Königsberg — an additional bridge joining land masses B and C . Did this addition make the desired route possible? Prove your answer.
3. Euler's 1736 article included a second example of a land/bridge system (see Figure 1.52). Does a route exist that crosses each bridge exactly once? If so, give one. If not, prove it.
4. The streets in River City are shown in Figure 1.53. Is it possible for Charlie the anvil salesman to plan a route that covers every street exactly once? If so, give one. If not, prove it.




FIGURE 1.52.




FIGURE 1.53.

1.4.2 Eulerian Trails and Circuits

If only part of the line-system has been traversed, then every node in the remaining part remains even or odd, just as it was in the original system.

— Carl Hierholzer [159], translated in [26]

Recall from Section 1.1.2 that a trail in a graph is a walk that does not repeat any edges, and that a closed trail (one that begins and ends at the same vertex) is called a circuit.

If a trail in a graph G includes every edge of G , then that trail is said to be an *Eulerian trail*. Similarly, an *Eulerian circuit* in a graph is a circuit that includes every edge of the graph. A graph that contains an Eulerian circuit is said to be an *Eulerian graph*.

What are some examples of Eulerian graphs? The cycles, C_n , have prominent Eulerian circuits. The paths, P_n , have no circuits at all, and so they are certainly not Eulerian. Look at the graphs in Figure 1.54 and try to determine which ones, if any, are Eulerian.

There are two well-known characterizations of Eulerian graphs. One involves vertex degrees, and the other concerns the existence of a special collection of cycles. The following theorem establishes both of these characterizations by asserting the logical equivalence of three statements. The theorem represents work by Euler [98], Hierholzer [159], and Veblen [274].




FIGURE 1.54. Are any of these Eulerian?

Theorem 1.20. *For a connected graph G , the following statements are equivalent.*

1. G is Eulerian.
2. Every vertex of G has even degree.
3. The edges of G can be partitioned into (edge-disjoint) cycles.

Proof. To prove the logical equivalence of these statements, we prove that the first statement implies the second, the second implies the third, and the third implies the first.

For the first implication, suppose that G contains an Eulerian circuit C . Let v be an arbitrary vertex of G . Every time the circuit enters v on an edge, it must leave on a different edge. Since C never repeats an edge, there must be an even number of edges incident with v and hence the degree of v is even.

For the second implication, suppose that every vertex of G has even degree. We use induction on the number of cycles in G . Since G is connected and has no vertices of degree 1, G is not a tree, and therefore G must have at least one cycle. If G has exactly one cycle, the G must be a cycle graph C_n for some n , and so the desired partition contains just the one cycle itself. Suppose now (using strong induction) that the implication is valid for graphs containing at most k cycles, and suppose G has $k + 1$ cycles. Let C be one of the cycles of G , and let G' be the graph obtained from G by deleting the edges of C . With this deletion, each vertex of C loses exactly two edges, and hence the vertices of G' all have even degree. Further, the graph G' (which is possibly disconnected) has connected components that have no more than k cycles each. Each component, then, satisfies the induction hypothesis and has edges that can be partitioned into cycles. These cycles, together with the cycle C , partition the edges of G into cycles. The induction is complete, and the implication is established.

For the third implication, suppose that the edges of G can be partitioned into cycles. Call these cycles S_1, S_2, \dots, S_k . Let C be the largest circuit in G such that the set of edges of C is exactly

$$E(S_{j_1}) \cup E(S_{j_2}) \cup \dots \cup E(S_{j_m})$$

for some collection of the cycles $S_{j_1}, S_{j_2}, \dots, S_{j_m}$. (We note here that this implies that for each cycle S_i ($1 \leq i \leq k$), either all of the edges of S_i are on C or none

of them are.) Now, suppose e is an edge of G that is (a) *not* an edge of C , and (b) incident with a vertex, say v , that *is* on C . Since e is not an edge of C , it must be that e is an edge of cycle S_i , for some i , where no edge of S_i is on C . The vertex v must also be on S_i . Let C' be the circuit in G obtained by patching S_i into C at the vertex v (since no edge of S_i is a member of C , there is no repetition of edges caused by this patching). Since the edges of C' consist of the edges of C together with the edges of S_i , we have contradicted the maximality of C' . This means that no such edge e can exist and therefore that C is an Eulerian circuit of G . The final implication is established. \square

So, Eulerian circuits exist in connected graphs precisely when the degrees of these graphs are all even. What about Eulerian trails? Certainly if an Eulerian circuit exists, then so does an Eulerian trail (the circuit is just a closed trail). But are there graphs which are not Eulerian but which do contain an Eulerian trail? The following corollary gives the complete answer. Its proof is left for you as an exercise.

Corollary 1.21. *The connected graph G contains an Eulerian trail if and only if there are at most two vertices of odd degree.*

Now that we know precisely when Eulerian circuits and trails exist, how easy is it to find them? The algorithm given below, named for nineteenth century mathematician Carl Hierholzer [159], gives a simple way of identifying such routes. While not identical, you may notice a similarity between this algorithm and the method used to prove the third implication in the proof of Theorem 1.20.

Before reading on, take a look back at the quotation given at the beginning of this section. It describes the primary reason for the success of Hierholzer's algorithm.

Hierholzer's Algorithm for Identifying Eulerian Circuits

Given: An Eulerian graph G .

- i. Identify a circuit in G and call it R_1 . Mark the edges of R_1 . Let $i = 1$.
- ii. If R_i contains all edges of G , then stop (since R_i is an Eulerian circuit).
- iii. If R_i does not contain all edges of G , then let v_i be a vertex on R_i that is incident with an unmarked edge, e_i .
- iv. Build a circuit, Q_i , starting at vertex v_i and using edge e_i . Mark the edges of Q_i .
- v. Create a new circuit, R_{i+1} , by patching the circuit Q_i into R_i at v_i .
- vi. Increment i by 1, and go to step ii.

An example of this process is shown in Figure 1.55. You should note that the




FIGURE 1.55. The stages of Hierholzer's algorithm.

process will succeed no matter what the initial circuit, R_1 , is chosen to be. Another algorithm for finding Eulerian circuits is given in Exercise 3.

The even degree characterization of Eulerian graphs is really quite nice. All one needs to do to determine if a graph is Eulerian is simply look at the degrees of the vertices. Once we know a graph is Eulerian, Hierholzer's algorithm will give us an Eulerian circuit. Maybe Charlie Cowell, our anvil salesman, used these ideas to plan his door to door routes!

Exercises

- For each of the following, draw an Eulerian graph that satisfies the conditions, or prove that no such graph exists.

(a) An even number of vertices, an even number of edges.

- (b) An even number of vertices, an odd number of edges.
 (c) An odd number of vertices, an even number of edges.
 (d) An odd number of vertices, an odd number of edges.
2. Use Hierholzer's algorithm to find an Eulerian circuit in the graph of Figure 1.56. Use $R_1 : a, b, c, g, f, j, i, e, a$ as your initial circuit.




FIGURE 1.56.

3. What follows is another algorithm (from [195]) for finding Eulerian circuits. The method used here is to build the circuit, one edge at a time, making sure to make good choices along the way.

Fleury's Algorithm for Identifying Eulerian Circuits

Given: An Eulerian graph G , with all of its edges *unmarked*.

- i. Choose a vertex v , and call it the "lead vertex."
- ii. If all edges of G have been marked, then stop. Otherwise continue to step iii.
- iii. Among all edges incident with the lead vertex, choose, if possible, one that is not a bridge of the subgraph formed by the unmarked edges. If this is not possible, choose any edge incident with the lead. Mark this edge and let its other end vertex be the new lead vertex.
- iv. Go to step ii.

Use Fleury's algorithm to find an Eulerian circuit for the graph in Figure 1.57. Let a be your initial vertex.

4. Prove Corollary 1.21.
5. Prove that if every edge of a graph G lies on an odd number of cycles, then G is Eulerian.
6. Let G be a connected graph which is regular of degree r . Prove that the line graph of G , $L(G)$, is Eulerian.




FIGURE 1.57.

7. Let $G = K_{n_1, n_2}$.
- Find conditions on n_1 and n_2 that characterize when G will have an Eulerian trail.
 - Find conditions that characterize when G will be Eulerian.
8. Let $G = K_{n_1, \dots, n_k}$, where $k \geq 3$.
- Find conditions on n_1, \dots, n_k that characterize when G will have an Eulerian trail.
 - Find conditions that characterize when G will be Eulerian.

1.4.3 Hamiltonian Paths and Cycles

In this new Game (... named Icosian, from a Greek work signifying 'twenty') a player is to place the whole or part of a set of twenty numbered pieces or men upon the points or in the holes of a board ... in such a manner as always to proceed along the lines of the figure ...

— from the instructions which accompanied The Icosian Game [26]

In 1859 the English game company Jaques and Son bought the rights to manufacture and market “The Icosian Game.” The game involved twenty pegs (numbered 1 through 20) and a gameboard with holes (see Figure 1.58). Players were required to place the pegs in the holes in such a way that every pair of consecutive integers, along with the pair (1, 20), were directly connected by one of the lines on the board. As intriguing and action-packed as the game sounds, it did not sell well.⁸

Another version of the game had the board design situated on the exterior of a dodecahedron, a 12-sided solid. The object here was to find a round-trip route around the solid, traveling only on the edges and touching each vertex once. This version was named “A Voyage Round the World” since the vertices represented important cities of the time. Like its twin “The Icosian Game,” this game’s voyage was short lived.

⁸Jaques and Son managed to get over this particular setback. The company, still in business today, had much better success in popularizing Tiddledy Winks (now Tiddly Winks), Snakes and Ladders (now Chutes and Ladders), and Whiff Whaff (now Table Tennis).




FIGURE 1.58. The Icosian Game board.

The inventor who sold the game to Jaques and Son was the prominent mathematician Sir William Rowan Hamilton. Even though his ideas did not take root in a recreational sense, they did become the seed for what would become a major branch of inquiry within the field of graph theory. Let's take a look at some of these ideas.

If a path P spans the vertices of G (that is, if $V(P) = V(G)$), then P is said to be a *Hamiltonian path* of G . Any graph containing a Hamiltonian path is called *traceable*. If a cycle C spans the vertices of a graph G , such a cycle is called a *Hamiltonian cycle*, and any graph containing a Hamiltonian cycle is called, simply, a *Hamiltonian graph*. Hamiltonian graphs are clearly traceable, but the reverse is not always true. Look at the graphs in Figure 1.59 and try to determine which ones are traceable, Hamiltonian, or neither.




FIGURE 1.59. Which ones are Hamiltonian? Which are traceable?

We saw in the previous section that whether or not a connected graph was Eulerian depended completely on degree parity. Unfortunately, this is not the case for Hamiltonicity. Hamiltonian graphs can have all even degrees (C_{10}), all odd degrees (K_{10}), or a mixture (G_1 in Figure 1.59). Similarly, non-Hamiltonian graphs can have varying degree parities: all even (G_2 in Figure 1.59), all odd ($K_{5,7}$), or mixed (P_9).

If degree parity does not have much to do with Hamiltonicity, then what does? Researchers have worked for decades on this question, and their efforts have produced many interesting results. A complete summary of these developments would require many pages,⁹ and we do not attempt to give a thorough treatment here. Rather, we present several of the classic ideas and results.

The first result that we examine is due to Dirac [77]. It does concern degrees in a graph — but their magnitude rather than their parity. Recall that $\delta(G)$ is the minimum degree of G .

Theorem 1.22. *Let G be a graph of order $n \geq 3$. If $\delta(G) \geq n/2$, then G is Hamiltonian.*

Proof. Let G be a graph satisfying the given conditions, and suppose that G is not Hamiltonian. Let P be a path in G with maximum length, and say the vertices of P , in order, are v_1, v_2, \dots, v_p . Because of the maximality of P , we know that all of the neighbors of v_1 and of v_p are on P . And since $\delta(G) \geq n/2$, each of v_1 and v_p has at least $n/2$ neighbors on P .

We now claim that there must exist some j ($1 \leq j \leq p - 1$) such that $v_j \in N(v_p)$ and $v_{j+1} \in N(v_1)$. Suppose for the moment that this was not the case. Then for every neighbor v_i of v_p on P (and there are at least $n/2$ of them), v_{i+1} is not a neighbor of v_1 . This means that

$$\deg(v_1) \leq p - 1 - \frac{n}{2} < n - \frac{n}{2} = \frac{n}{2},$$

contradicting the fact that $\delta(G) \geq n/2$. Therefore, such a j exists (see Fig-




FIGURE 1.60.

ure 1.60).

Let C be the cycle $v_1, v_2, \dots, v_j, v_p, v_{p-1}, \dots, v_{j+1}, v_1$. We have assumed that G is not Hamiltonian, and so there must be at least one vertex of G that is not on P . Further, since $\delta(G) \geq n/2$, we know that G is connected (see Exercise 16a in Section 1.1.2). Therefore there must be a vertex w not on P that is adjacent to a vertex, say v_i , on P . But then the path in G that begins with w , travels to v_i , and then travels around the cycle C is a longer path than our maximal path P — a contradiction. Our initial assumption must have been incorrect. Therefore G is Hamiltonian. \square

⁹In 2003 Ron Gould published an article [129] which summarized related results. His survey was 45 pages long, and even that only covered developments that took place within the previous dozen years!

There are two interesting things to note here. First, the lower bound in this theorem is best possible. To see this, consider the graph $G = K_{r,r+1}$. This graph is not Hamiltonian (you will show this as part of Exercise 12b) and $\delta(G)$ is strictly between $\frac{|V(G)|}{2} - 1$ and $\frac{|V(G)|}{2}$. Second, the theorem does not provide a characterization of Hamiltonian graphs. That is, there are plenty of Hamiltonian graphs that have relatively small minimum (and even maximum) degree. The cycles C_n are obvious examples.

Dirac's theorem is a corollary to the following general result of Ore [217]. The proof of Ore's theorem is similar to that of the above, and it is left for you as an exercise.

Theorem 1.23. *Let G be a graph of order $n \geq 3$. If $\deg(x) + \deg(y) \geq n$ for all pairs of nonadjacent vertices x, y , then G is Hamiltonian.*

A set of vertices in a graph is said to be an *independent set* of vertices if they are pairwise nonadjacent. The *independence number* of a graph G , denoted by $\alpha(G)$, is defined to be the largest size of an independent set of vertices from G . As an example, consider the graphs in Figure 1.61. The only independent set of size 2 in G_1 is $\{c, d\}$, so $\alpha(G_1) = 2$. There are two independent sets of size 3 in G_2 : $\{a, c, e\}$ and $\{b, d, f\}$, and none of size 4, so $\alpha(G_2) = 3$.




FIGURE 1.61.

The next theorem, due to Chvátal and Erdős [54], relates Hamiltonicity to independence number and connectivity. Before stating and proving this result, let us introduce some helpful notation. If x and y are vertices on a path P , let $[x, y]_P$ denote the portion of P that runs from x to y . Further, given a cycle C with its vertices labeled in a specific orientation (say, clockwise), let $[x, y]_{C^+}$ denote the portion of C that runs clockwise from x to y . Similarly, $[x, y]_{C^-}$ would denote the portion of C that runs counter-clockwise from x to y .

Theorem 1.24. *Let G be a connected graph of order $n \geq 3$ with vertex connectivity $\kappa(G)$ and independence number $\alpha(G)$. If $\kappa(G) \geq \alpha(G)$, then G is Hamiltonian.*

Proof. If G is as described, then $\kappa(G) \geq 2$ — for if $\kappa(G) = 1$, then $\alpha(G) = 1$ and thus G is either K_1 or K_2 , contradicting the fact that $n \geq 3$.

Let C be a longest cycle in G . Suppose that C is not a Hamiltonian cycle, and let v be a vertex of G that is not on C . Let H be the connected component of

$G - V(C)$ that contains v . Let c_1, c_2, \dots, c_r be the vertices of C that are adjacent to some vertex of H , and suppose that these vertices are labeled in a clockwise direction around the cycle C . For each i ($1 \leq i \leq r$), let h_i be a vertex of H that is adjacent to c_i , and let d_i be the immediate (clockwise) successor of c_i on C .

We now observe several things. First, it must be that $r \geq \kappa(G)$. If the vertices c_1, c_2, \dots, c_r were removed from G , then H would be disconnected from the rest of the graph. Since $\kappa(G)$ is the size of the smallest cut set, it follows that $r \geq \kappa(G)$. The observation in the first paragraph then implies that $r \geq 2$.

Second, no two of the vertices in the set $\{c_1, c_2, \dots, c_r\}$ are consecutive vertices on C . To see this, suppose that there is some i such that c_i and c_{i+1} are consecutive vertices on C . Let P be a path from h_i to h_{i+1} in H , and consider the cycle formed by replacing the edge $c_i c_{i+1}$ on C with the path $c_i, [h_i, h_{i+1}]_P, c_{i+1}$. This cycle is longer than our maximal cycle C , a contradiction. This observation implies that the sets $\{c_1, c_2, \dots, c_r\}$ and $\{d_1, d_2, \dots, d_r\}$ are disjoint.

Third, for each i ($1 \leq i \leq r$), d_i is not adjacent to v . To see this, suppose $d_i v \in E(G)$ for some i , and let Q be a path from h_i to v in H . In this case, the cycle formed by replacing the edge $c_i d_i$ on C with the path $c_i, [h_i, v]_Q, d_i$ is longer than C , again a contradiction.

Now, let $S = \{v, d_1, d_2, \dots, d_r\}$. The first observation above implies that $|S| \geq \kappa(G) + 1 > \alpha(G)$. This means that some pair of vertices in S must be adjacent. Our third observation implies that d_i must be adjacent to d_j for some $i < j$. If R is a path from h_i to h_j in H , then the cycle $c_i, [h_i, h_j]_R, [c_j, d_i]_{C^-}, [d_j, c_i]_{C^+}$ is a longer cycle than C (see Figure 1.62). Our assumption that C was not a Hamil-




FIGURE 1.62.

tonian cycle has led to a contradiction. Therefore G is Hamiltonian. □

As was the case for Dirac’s theorem, the inequality in this theorem is sharp. That is, graphs G where $\kappa(G) \geq \alpha(G) - 1$ are not necessarily Hamiltonian. The complete bipartite graphs $K_{r,r+1}$ provide proof of this. The Petersen graph, shown in Figure 1.63, is another example.¹⁰

¹⁰The Petersen graph is well known among graph theorists for its surprising connections to many areas of the field, and for its penchant for being a counterexample to many conjectures. You have seen this graph already — it is the graph you should have obtained as the complement of the line graph of K_5 in Exercise 7b of Section 1.1.3.




FIGURE 1.63. The Petersen Graph.

The next theorem belongs to a category of results that relate Hamiltonicity to forbidden subgraphs. Given graphs G and H , if G does not contain a copy of H as an *induced* subgraph, then we say that G is H -free. If S is a collection of graphs, and if G does not contain any of the graphs in S as induced subgraphs, then we say that G is S -free.

In 1974, Goodman and Hedetniemi [127] noticed something regarding two of the graphs shown in Figure 1.64.




FIGURE 1.64. Three special graphs.

Theorem 1.25. *If G is a 2-connected, $\{K_{1,3}, Z_1\}$ -free graph, then G is Hamiltonian.*

Proof. Suppose G is 2-connected and $\{K_{1,3}, Z_1\}$ -free, and let C be a longest cycle in G (we know that G contains at least one cycle — see Exercise 14 in Section 1.1.2). If C is not a Hamiltonian cycle, then there must exist a vertex v , not on C , which is adjacent to a vertex, say w , on C . Let a and b be the immediate predecessor and successor of w on C .

A longer cycle would exist if either a or b were adjacent to v , and so it must be that both a and b are nonadjacent to v . Now, if a is not adjacent to b , then

the subgraph induced by the vertices $\{w, v, a, b\}$ is $K_{1,3}$, and we know that G is $K_{1,3}$ -free. So it must be that $ab \in E(G)$. But if this is the case, then the subgraph induced by $\{w, v, a, b\}$ is Z_1 , a contradiction. Therefore, it must be that C is a Hamiltonian cycle. \square

The fact that this result requires 2-connectivity should not be surprising, for 2-connectivity is required for all Hamiltonian graphs. As you will prove in Exercise 2, graphs G where $\kappa(G) = 1$ cannot have a spanning cycle.

Another classic forbidden subgraph theorem involves $K_{1,3}$ and the third graph shown in Figure 1.64. This is a result of Duffus, Gould, and Jacobson [81].

Theorem 1.26. *Let G be a $\{K_{1,3}, N\}$ -free graph.*

1. *If G is connected, then G is traceable.*
2. *If G is 2-connected, then G is Hamiltonian.*

It is interesting to note that the graph $K_{1,3}$ is involved in both of these theorems. This graph, affectionately referred to as the “claw,” appears in many forbidden subgraph results. Claw-free graphs have received a great deal of attention in recent years, especially within the context of Hamiltonicity problems. The claw will bare itself again in the next section in the context of unsolved problems.

Exercises

1. Give a solution to The Icosian Game.
2. Prove that if G is Hamiltonian, then G is 2-connected.
3. Prove Theorem 1.23.
4. Give the connectivity and independence number of the Petersen graph.
5. Prove or disprove: The independence number of a bipartite graph is equal to the cardinality of one of its partite sets.
6. Prove that if G is of order n and is regular, then $\alpha(G) \leq n/2$.
7. Prove that each of the 18-vertex graphs in Figure 1.65 is 2-connected, claw-free and nontraceable.¹¹
8. For any graph G , prove that the line graph $L(G)$ is claw-free.
9. Let G be a K_3 -free graph. Prove that its complement, \overline{G} , is claw-free.
10. Let G be a graph and let S be a nonempty subset of $V(G)$.

¹¹In [154] the authors show that 2-connected, claw-free graphs of order less than 18 are traceable. They also show that the graphs in Figure 1.65 are the only 2-connected, claw-free, nontraceable graphs with order 18 and size at most 24.




FIGURE 1.65.

- (a) Prove that if G is Hamiltonian, then $G - S$ has at most $|S|$ connected components.
 - (b) Prove that if G is traceable, then $G - S$ has at most $|S| + 1$ connected components.
11. Prove that if G is Eulerian, then $L(G)$ is Hamiltonian.
12. Let $G = K_{n_1, n_2}$.
- (a) Find conditions on n_1 and n_2 that characterize the traceability of G .
 - (b) Find conditions that characterize the Hamiltonicity of G .
13. Let n be a positive integer.
- (a) Prove that $K_{n, 2n, 3n}$ is Hamiltonian.
 - (b) Prove that $K_{n, 2n, 3n+1}$ is not Hamiltonian.

1.4.4 Three Open Problems

Nothing can stop the claw!

— Fletcher Reede, *Liar, Liar*

We close our discussion of paths and cycles with several questions. These problems were posed years ago, and they have received a great deal of attention. While there has been progress on each of them, the original questions remain unanswered.

Intersecting Detour Paths

Given a graph G , the *detour order* of G , denoted $\tau(G)$, is the number of vertices in a longest path in G . If a path P in G has $\tau(G)$ vertices, then we call P a *detour path* in G .

In Exercise 8 of Section 1.1.2, you proved that if G is a connected graph and if P_1 and P_2 are detour paths, then the intersection $V(P_1) \cap V(P_2)$ must be nonempty. As you (hopefully!) saw, this result is not terribly difficult to prove. Consider the following question.

Question A: If G is connected and $P_1, P_2,$ and P_3 are detour paths in G , then must the intersection $V(P_1) \cap V(P_2) \cap V(P_3)$ be nonempty?

As you see, the only difference between this question and the earlier exercise is that this one involves three paths rather than two. But this difference makes all the difference, because Question A remains unsolved!

The origin of this question can be traced to a related question asked by Gallai [118] in 1966: Is it true that in every connected graph there is at least one vertex that lies on every detour path? In 1969, Walther [278] gave a negative answer to the question in the form of the graph in Figure 1.66. The detour order of this graph




FIGURE 1.66. Walther's example.

is 21, and every vertex is missed by at least one of the detour paths. Within the next several years, Walther and Zamfirescu, working independently (see [141], [279], [290]), had produced a smaller example, the graph in Figure 1.67. This




FIGURE 1.67. A smaller example, given by Walther and Zamfirescu.

graph has 12 vertices, has detour order 10, and every vertex is missed by at least one of the detour paths. The graph in Figure 1.67 is the smallest known graph where the detour paths have an empty intersection.

Consider the following more general version of Question A.

Question B: If G is connected and P_1, \dots, P_n are distinct detour paths in G , then must the intersection $\bigcap_{i=1}^n V(P_i)$ be nonempty?

The graph of Figure 1.67 demonstrates that for $n = 12$, the answer to Question B is no.

In 1975, Schmitz [250] presented the graph in Figure 1.68. The detour order




FIGURE 1.68. Example given by Schmitz.

of this graph is 13. There are exactly seven detour paths, and every vertex of the graph is missed by at least one of these paths. This tells us that for $n = 7$, the answer to Question B is no.

We have already mentioned that for $n = 3$, the answer to Question B is unknown. The same is true for $n = 4, 5$ and 6 . When asked, most researchers would probably lean toward believing the result to be true for $n = 3$, although no proof is known as of yet. For now, it is simply a conjecture.

Conjecture 1. *If G is connected, then the intersection of any three distinct detour paths in G is nonempty.*

Matthews and Sumner’s Conjecture

We met the claw, $K_{1,3}$, in the previous section. We saw two results in which the claw, when forbidden with another graph, implied Hamiltonicity in 2-connected graphs. There are other such pairs. In [20], [39], and [130], the respective authors showed that the pairs $\{K_{1,3}, W\}$, $\{K_{1,3}, P_6\}$, and $\{K_{1,3}, Z_2\}$ (see Figure 1.69) all imply Hamiltonicity when forbidden in 2-connected graphs. Do you see a pat-




FIGURE 1.69. Two additional forbidden subgraphs.

tern here? The claw seems to be prominent in results like this. In 1997, Faudree and Gould [102] showed that this was no coincidence. The graph N appears in Figure 1.64.

Theorem 1.27. *If being $\{R, S\}$ -free (where R, S are connected and neither is P_3) implies that 2-connected graphs are Hamiltonian, then one of R, S is the claw, and the other is an induced subgraph of $P_6, Z_2, W,$ or N .*

Is the claw powerful enough to imply Hamiltonicity when forbidden by itself? Well, as you will prove in the exercises, the graph P_3 is the only connected graph that, when forbidden, implies Hamiltonicity in 2-connected graphs. But what if the level of connectivity is increased?

Question C: If G is claw-free and k -connected ($k \geq 3$), must G be Hamiltonian?

The graph in Figure 1.70 is 3-connected, claw-free and non-Hamiltonian,¹² and so the answer to Question C for $k = 3$ is no.




FIGURE 1.70. A lovely example.

In 1984 Matthews and Sumner [199] made the following, still unresolved, conjecture.

Conjecture 2. *If G is 4-connected and claw-free, then G is Hamiltonian.*

There has been some progress with regard to this conjecture. Most notably, in 1997 Ryjáček [245] proved the following theorem.

Theorem 1.28. *If G is 7-connected and claw-free, then G is Hamiltonian.*

At this time, Question C for $k = 4, 5$ and 6 is still unanswered.

¹²This graph, demonstrated by Matthews ([198], see also [199]) in 1982, is the smallest such graph, and it is the line graph of the graph obtained from the Petersen graph (what else?) by replacing each of the five “spoke” edges with a P_3 .

The Path Partition Conjecture

Recall that $\tau(G)$, the detour order of G , is the number of vertices in a longest path of G . Recall also that given a subset S of $V(G)$, the notation $\langle S \rangle$ represents the subgraph of G induced by S .

Given a graph G and positive integers a and b , if the vertices of G can be partitioned into two sets A and B in such a way that $\tau(\langle A \rangle) \leq a$ and $\tau(\langle B \rangle) \leq b$, then we say that G has an (a, b) -partition.

As an example, consider the graph G in Figure 1.71. The partition (A_1, B_1) ,




FIGURE 1.71. Partitioning the vertex set.

where $A_1 = \{a, b, c, d, h, i\}$ and $B_1 = \{e, f, g, j, k, l\}$, is not a valid $(4, 7)$ -partition. This is because every longest path in both $\langle A_1 \rangle$ and $\langle B_1 \rangle$ has 6 vertices. On the other hand, the partition (A_2, B_2) , where

$$A_2 = \{c, d, h, i\}, B_2 = \{a, b, e, f, g, j, k, l\}$$

is a valid $(4, 7)$ -partition, since $\tau(\langle A_2 \rangle) \leq 4$ and $\tau(\langle B_2 \rangle) \leq 7$.

If a graph G has an (a, b) -partition for every pair (a, b) of positive integers such that $a + b = \tau(G)$, then we say that G is τ -partitionable. In order to show that the graph in Figure 1.71 is τ -partitionable (since its detour order is 11), we would also need to show that the graph had $(1, 10)$ -, $(2, 9)$ -, $(3, 8)$ -, and $(5, 6)$ -partitions. Another way to demonstrate this would be to prove the following, still unresolved, conjecture.

The Path Partition Conjecture. *Every graph is τ -partitionable.*

The Path Partition Conjecture was first mentioned by Lovász and Mihók in 1981, and it has received a great deal of attention since then. Much of the progress has been the establishment that certain families of graphs are τ -partitionable. The following theorem gives a small sampling of these results.

Theorem 1.29. *The graph G is τ -partitionable if any one of the following is true:*

1. $\Delta(G) \leq 3$; (follows from [189])
2. $\Delta(G) \geq |V(G)| - 8$; (follows from [38])
3. $\tau(G) \leq 13$; [83]
4. $\tau(G) \geq |V(G)| - 1$. [38]

So, if the max degree of a graph G is either very small or relatively large, or if the longest path in a graph is either rather short or relatively long, then G is τ -partitionable.

In 2007, Dunbar and Frick [84] proved the following interesting result.

Theorem 1.30. *If G is claw-free, then G is τ -partitionable.*

There is the claw again! In the same article, the authors prove that in order to prove that the Path Partition Conjecture is true, it is sufficient to prove that every 2-connected graph is τ -partitionable.

With each new result, researchers add to the arsenal of weapons that can be used to attack the Path Partition Conjecture. So far, though, the conjecture is holding strong.

With all of this era's computing power, how can it be that the Path Partition Conjecture and the other conjectures in this section remain unsolved? Computers are now doing more things than ever, faster than ever, so why can't we just get a computer cranking away at these problems? These are reasonable questions to ask. But the conjectures here are not really questions of computation. These are problems that will require a combination of insight, cleverness and patience.

Exercises

1. If the word "connected" were removed from Conjecture 1, could you settle the resulting conjecture?
2. Show that every vertex of the graph in Figure 1.67 is missed by at least one detour path in the graph.
3. In the graph of Figure 1.68, find the seven distinct detour paths and show that they have an empty intersection.
4. Show that if G is 2-connected and P_3 -free, then G is Hamiltonian.
5. Show that if being H -free implies Hamiltonicity in 2-connected graphs (where H is connected), then H is P_3 .

6. Verify that the graph in Figure 1.70 is actually the line graph of the graph obtained from the Petersen graph by replacing each of the five “spoke” edges with a P_3 .
7. Prove that the graph in Figure 1.67 is τ -partitionable by listing all necessary partitions.
8. Prove that all bipartite graphs are τ -partitionable.
9. Prove that all traceable graphs are τ -partitionable.
10. Prove that every graph G has a $(1, \tau(G) - 1)$ partition.
11. Prove that if a graph is $(1, 1)$ -partitionable, then it is (a, b) -partitionable for all positive integers a and b .
12. Show that all graphs are (a, b) -partitionable when $a \leq 3$.
13. EXTRA CREDIT: Settle any of the conjectures in this section.

1.5 Planarity

*Three civil brawls, bred of an airy word
By thee, old Capulet, and Montague,
Have thrice disturb'd the quiet of our streets . . .*

— William Shakespeare, *Romeo and Juliet*

The feud between the Montagues and the Capulets of Verona has been well documented, discussed, and studied. A fact that is lesser known, though, is that long before Romeo and Juliet’s time, the feud actually involved a third family—the Hatfields.¹³ The families’ houses were fairly close together, and chance meetings on the street were common and quite disruptive.

The townspeople of Verona became very annoyed at the feuding families. They devised a plan to create separate, nonintersecting routes from each of the houses to each of three popular places in town: the square, the tavern, and the amphitheater. They hoped that if each family had its own route to each of these places, then the fighting in the streets might stop.

Figure 1.72 shows the original layout of the routes. Try to rearrange them so that no route crosses another route. We will come back to this shortly.

¹³The Hatfields eventually grew tired of feuding, and they left Verona in search of friendlier territory. They found a nice spot in the mountains of West Virginia, right across the river from a really nice family named McCoy.




FIGURE 1.72. Original routes.

1.5.1 Definitions and Examples

Define, define, well-educated infant.

— William Shakespeare, *Love's Labour's Lost*

A graph G is said to be *planar* if it can be drawn in the plane in such a way that pairs of edges intersect only at vertices, if at all. If G has no such representation, G is called *nonplanar*. A drawing of a planar graph G in the plane in which edges intersect only at vertices is called a *planar representation* (or a *planar embedding*) of G .

Figure 1.73 shows examples of planar graphs. Notice that one of the drawings




FIGURE 1.73. Examples of planar graphs.

is not a planar representation—try to visualize untangling it.

Proving a graph to be planar is in some cases very simple—all that is required is to exhibit a planar representation of the graph. This is certainly quite easy to do with paths, cycles, and trees. What about complete graphs? K_1 , K_2 , and K_3 are clearly planar; Figure 1.74 shows a planar representation of K_4 . We will consider K_5 shortly.




FIGURE 1.74. A planar representation of K_4 .

The Montague/Capulet/Hatfield problem essentially amounts to finding a planar representation of $K_{3,3}$. Unfortunately, the townspeople of Verona just had to learn to deal with the feuding families, for $K_{3,3}$ is nonplanar, and we will see an explanation shortly.

What is involved in showing that a graph G is nonplanar? In theory, one would have to show that every possible drawing of G is not a planar representation. Since considering every individual drawing is out of the question, we need some other tools.

Given a planar representation of a graph G , a *region* is a maximal section of the plane in which any two points can be joined by a curve that does not intersect any part of G .

Informally, if a cookie cutter has the shape of a planar representation of G , then the cookies are the regions (see Figure 1.75). The big region, R_7 , is called the *exterior* (or *outer*) region.




FIGURE 1.75. Six small cookies and one very large cookie.

It is quite natural to think of the regions as being bounded by the edges. A single edge can come into contact with either one or two regions. In Figure 1.76, edge e_1 is only in contact with one region, S_1 , and edges e_2 and e_3 are only in contact with S_2 . Each of the other edges in Figure 1.76 is in contact with two regions. Let us say that an edge e *bounds* a region R if e comes into contact with




FIGURE 1.76. Edges e_1 , e_2 , and e_3 touch one region only.

R and with a region different from R . Define the *bound degree* of R , denoted by $b(R)$, to be the number of edges that bound region R . For example, in Figure 1.75, $b(R_1) = b(R_4) = 4$, $b(R_2) = b(R_3) = b(R_5) = b(R_6) = 3$, and $b(R_7) = 12$. In Figure 1.76, $b(S_1) = b(S_3) = 3$ and $b(S_2) = 6$. Note that in this graph, the edges e_1, e_2 , and e_3 do not contribute to the bound degree of any region.

Figure 1.77 displays six planar graphs along with the numbers of vertices, edges, and regions. Before continuing to the next section, study these numbers and try to find a pattern. You might also notice that two of drawings are actually the same graph. *This brings up an important point: The number of regions in a planar representation of a graph does not depend on the representation itself!*




FIGURE 1.77. Is there a pattern?

Exercises

- Find planar representations for each of the planar graphs in Figure 1.78.
- Give planar representations of the graph in Figure 1.79 such that each of the following is the exterior region.
 - R_1
 - R_2
 - R_3
 - R_4
 - R_5




FIGURE 1.78.




FIGURE 1.79.

3. Explain why embedding a graph in a plane is essentially the same as embedding a graph on a sphere.
4. Write a nice proof of the fact that every tree is planar.
5. Draw a planar graph in which every vertex has degree exactly 5.
6. Suppose that e is a bridge of a planar graph G . Prove that e does not bound a region in any planar representation of G .
7. In [101] and [277], Fáry and Wagner proved independently that every planar graph has a planar representation in which every edge is a straight line segment. Find such a representation for the graph in Figure 1.80.




FIGURE 1.80.

8. If planar graphs G_1 and G_2 each have n vertices, q edges, and r regions, must the graphs be isomorphic? Justify your answer with a proof or a counterexample.

1.5.2 Euler's Formula and Beyond

Now I will have less distraction.

— Leonard Euler, upon losing sight in his right eye [100]

Euler discovered a relationship between the numbers of vertices, edges, and regions of a graph, and his discovery is often called Euler's Formula [99].

Theorem 1.31 (Euler's Formula). *If G is a connected planar graph with n vertices, q edges, and r regions, then*

$$n - q + r = 2.$$

Proof. We induct on q , the number of edges. If $q = 0$, then G must be K_1 , a graph with 1 vertex and 1 region. The result holds in this case. Assume that the result is true for all connected planar graphs with fewer than q edges, and assume that G has q edges.

Case 1. Suppose G is a tree. We know from our work with trees that $q = n - 1$; and of course, $r = 1$, since a planar representation of a tree has only one region. Thus $n - q + r = n - (n - 1) + 1 = 2$, and the result holds.

Case 2. Suppose G is not a tree. Let C be a cycle in G , let e be an edge of C , and consider the graph $G - e$. Compared to G , this graph has the same number of vertices, one edge fewer, and one region fewer, since removing e coalesces two regions in G into one in $G - e$. Thus the induction hypothesis applies, and in $G - e$,

$$n - (q - 1) + (r - 1) = 2,$$

implying that $n - q + r = 2$.

The result holds in both cases, and the induction is complete. \square

Euler's Formula is useful for establishing that a graph is nonplanar.

Theorem 1.32. $K_{3,3}$ is nonplanar.

Proof. Suppose that $K_{3,3}$ were planar and that we had a planar representation. Since $n = 6$ and $q = 9$, Euler's Formula implies that such a planar representation of $K_{3,3}$ would have $r = 5$ regions. Now consider the sum

$$C = \sum_R b(R),$$

where the sum is over all regions R in the representation of the graph. Since every edge of G can be on the boundary of at most two regions, we get $C \leq 2q = 18$. On the other hand, since each region of $K_{3,3}$ has at least four edges on the boundary (there are no triangles in bipartite graphs), we see that $C \geq 4r = 20$. We have reached a contradiction. Therefore, $K_{3,3}$ is nonplanar. \square

Theorem 1.33. *If G is a planar graph with $n \geq 3$ vertices and q edges, then $q \leq 3n - 6$. Furthermore, if equality holds, then every region is bounded by three edges.*

Proof. Again consider the sum

$$C = \sum_R b(R).$$

As previously mentioned, $C \leq 2q$. Further, since each region is bounded by at least 3 edges, we have that $C \geq 3r$. Thus

$$3r \leq 2q \Rightarrow 3(2 + q - n) \leq 2q \Rightarrow q \leq 3n - 6.$$

If equality holds, then $3r = 2q$, and it must be that every region is bounded by three edges. \square

We can use Theorem 1.33 to establish that K_5 is nonplanar.

Theorem 1.34. K_5 is nonplanar.

Proof. K_5 has 5 vertices and 10 edges. Thus $3n - 6 = 9 < 10 = q$, implying that K_5 is nonplanar. \square

Exercise 5 in Section 1.5.1 asked for a planar graph in which every vertex has degree exactly 5. This next result says that such a graph is an extreme example.

Theorem 1.35. If G is a planar graph, then G contains a vertex of degree at most five. That is, $\delta(G) \leq 5$.

Proof. Suppose G has n vertices and q edges. If $n \leq 6$, then the result is immediate, so we will suppose that $n > 6$. If we let D be the sum of the degrees of the vertices of G , then we have

$$D = 2q \leq 2(3n - 6) = 6n - 12.$$

If each vertex had degree 6 or more, then we would have $D \geq 6n$, which is impossible. Thus there must be some vertex with degree less than or equal to 5. \square

Exercises

1. G is a connected planar graph of order 24, and it is regular of degree 3. How many regions are in a planar representation of G ?
2. Let G be a connected planar graph of order less than 12. Prove $\delta(G) \leq 4$.
3. Prove that Euler's formula fails for disconnected graphs.
4. Let G be a connected, planar, K_3 -free graph of order $n \geq 3$. Prove that G has no more than $2n - 4$ edges.
5. Prove that there is no bipartite planar graph with minimum degree at least 4.

6. Let G be a planar graph with k components. Prove that

$$n - q + r = 1 + k.$$

7. Let G be of order $n \geq 11$. Show that at least one of G and \overline{G} is nonplanar.

8. Show that the average degree (see Exercise 11 in Section 1.3.2) of a planar graph is less than six.

9. Prove that the converse of Theorem 1.33 is not true.

10. Find a 4-regular planar graph, and prove that it is unique.

11. A planar graph G is called *maximal planar* if the addition of any edge to G creates a nonplanar graph.

(a) Show that every region of a maximal planar graph is a triangle.

(b) If a maximal planar graph has order n , how many edges and regions does it have?

1.5.3 Regular Polyhedra

We are usually convinced more easily by reasons we have found ourselves than by those which have occurred to others.

— Blaise Pascal, *Pensées*

A polyhedron is a solid that is bounded by flat surfaces. Dice, bricks, pyramids, and the famous dome at Epcot Center in Florida are all examples of polyhedra. Polyhedra can be associated with graphs in a very natural way. Think of the polyhedra as having faces, edges, and corners (or vertices). The vertices and edges of the solid make up its skeleton, and the skeleton can be viewed as a graph. An interesting property of these skeleton graphs is that they are planar. One way to see this is to imagine taking hold of one of the faces and stretching it so that its edges form the boundary of the exterior region of the graph. The regions of these planar representations directly correspond to the faces of the polyhedra. Figure 1.81 shows a brick-shaped polyhedron, its associated graph, and a planar representation of the graph.




FIGURE 1.81. A polyhedron and its graph.

Because of the natural correspondence, we are able to apply some of what we know about planar graphs to polyhedra. The next theorem follows directly from Euler's Formula for planar graphs, Theorem 1.31.

Theorem 1.36. *If a polyhedron has V vertices, E edges, and F faces, then*

$$V - E + F = 2.$$

This next theorem is similar to Theorem 1.35.

Given a polyhedron P , define $\rho(P)$ to be

$$\rho(P) = \min\{b(R) \mid R \text{ is a region of } P\}.$$

Theorem 1.37. *For all polyhedra P , $3 \leq \rho(P) \leq 5$.*

Proof. Since one or two edges can never form a boundary, we know that $\rho(P) \geq 3$ for all polyhedra P . So we need to prove only the upper bound.

Let P be a polyhedron and let G be its associated graph. Suppose P has V vertices, E edges, and F faces. For each k , let V_k be the number of vertices of degree k , and let F_k be the number of faces of P (or regions of G) of bound degree k . From our earlier remarks, if $k < 3$, then $V_k = F_k = 0$. Since every edge of P touches exactly two vertices and exactly two faces, we find that

$$\sum_{k \geq 3} kV_k = 2E = \sum_{k \geq 3} kF_k.$$

If every face of P were bounded by 6 or more edges, then we would have

$$2E = \sum_{k \geq 3} kF_k \geq \sum_{k \geq 6} 6F_k = 6 \sum_{k \geq 6} F_k = 6F,$$

implying that $E \geq 3F$. Furthermore,

$$2E = \sum_{k \geq 3} kV_k \geq 3V,$$

implying that $V \leq \frac{2}{3}E$. Thus

$$E = V + F - 2 \leq \frac{2}{3}E + \frac{1}{3}E - 2 = E - 2,$$

and this, of course, is a contradiction. Therefore, some face of P is bounded by fewer than 6 edges. Hence, $\rho(P) \leq 5$. \square

We now apply this result to derive a geometric fact known to the ancient Greeks.

A regular polygon is one that is equilateral and equiangular. We say a polyhedron is *regular* if its faces are mutually congruent, regular polygons and if the number of faces meeting at a vertex is the same for every vertex. The cube, whose faces are congruent squares, and the tetrahedron, whose faces are congruent equilateral triangles, are regular polyhedra. A fact that has been known for at least 2000 years is that there are only five regular polyhedra: the tetrahedron, the cube, the octahedron, the dodecahedron, and the icosahedron (see Figure 1.82). We can use a graph-theoretic argument to prove this.

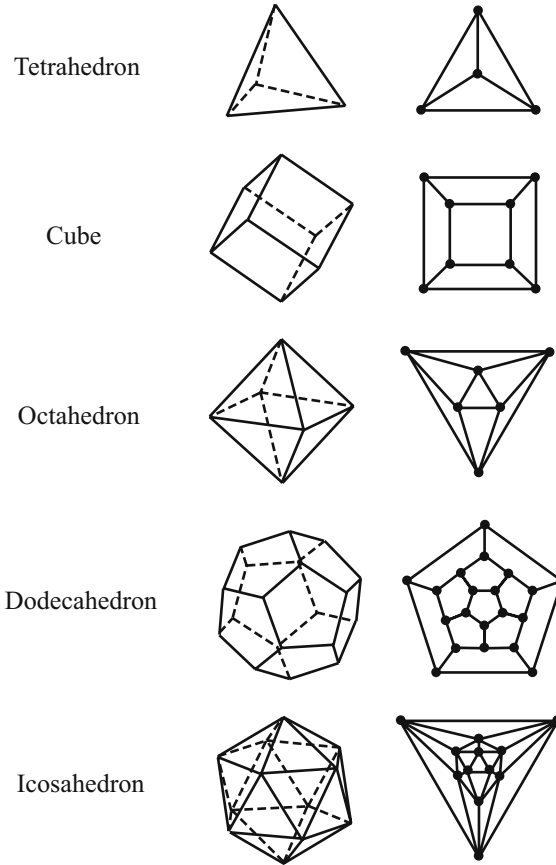


FIGURE 1.82. The five regular polyhedra and their graphical representations.

Theorem 1.38. *There are exactly five regular polyhedra.*

Proof. Let P be a regular polyhedron, and let G be its associated planar graph. Let V , E , and F be the number of vertices, edges, and faces (regions) of P . Since the faces of P are congruent, each is bordered by the same number of edges, say k . Theorem 1.37 tells us that $3 \leq k \leq 5$. Further, since the polyhedron P is regular, it follows that the graph G is also regular. Let us say that G is regular of degree r where $r \geq 3$. From Theorem 1.37, we obtain $rV = 2E = kF$. Now, Theorem 1.36 implies that

$$\begin{aligned}
 8 &= 4V - 4E + 4F \\
 &= 4V - 2E + 4F - 2E \\
 &= 4V - rV + 4F - kF \\
 &= (4 - r)V + (4 - k)F.
 \end{aligned}$$

V and E are of course both positive, and since $3 \leq k \leq 5$ and $r \geq 3$, there are only five possible cases.

Case 1. Suppose $r = 3$ and $k = 3$. In this case, $V = F$ and $8 = V + F$, implying that $V = F = 4$. This is the tetrahedron. (The fact that the tetrahedron is the only regular polygon with $V = F = 4$ is based on a geometrical argument. This applies to the remaining four cases as well.)

Case 2. Suppose $r = 3$ and $k = 4$. Here we have $V = 8$ and $3V = 4F$. Thus $F = 6$, and P is a cube.

Case 3. Suppose $r = 3$ and $k = 5$. In this case we have $8 = V - F$ and $3V = 5F$. Solving this system yields $V = 20$ and $F = 12$. This is a dodecahedron.

Case 4. Suppose $r = 4$ and $k = 3$. Here we have $F = 8$ and $4V = 3F$. Thus $V = 6$ and P is an octahedron.

Case 5. Suppose $r = 5$ and $k = 3$. In this case we have $8 = -V + F$ and $5V = 3F$. Solving this system yields $V = 12$ and $F = 20$. This is an icosahedron. \square

Exercises

- (From [52].) Show that the octahedron is a complete multipartite graph: K_{r_1, \dots, r_n} for some n and for some values r_1, \dots, r_n .
- Find an example of a polyhedron different from the ones discussed in this section. Sketch the polyhedron, and draw the associated graph.
- See if you can find an alternative proof (not necessarily graph-theoretic) of the fact that there are only five regular polyhedra.

1.5.4 Kuratowski's Theorem

... a pair so famous.

— William Shakespeare, *Anthony and Cleopatra*

Our goal in this section is to compile a list of all nonplanar graphs. Since the list will be infinite (and since this book is not), we will make use of a clever characterization due to Kuratowski.

We have already established that both $K_{3,3}$ and K_5 are nonplanar, so we should put them at the top of our list. What other graphs should we include? Suppose G is a graph that contains $K_{3,3}$ as a subgraph. This graph G would have to be nonplanar, since a planar representation of it would have to contain a planar representation of $K_{3,3}$. So we can add to our list of nonplanar graphs all graphs that contain $K_{3,3}$ or K_5 as a subgraph.

The graph in Figure 1.83 shows us that our list of nonplanar graphs is not yet complete. This graph is not planar, but it does not contain K_5 or $K_{3,3}$ as a subgraph. Of course, if we were to replace the two edges labeled a and b with a single edge e , then the graph would contain K_5 as a subgraph. This motivates the following definition.




FIGURE 1.83.

Let G be a graph. A *subdivision* of an edge e in G is a substitution of a path for e . We say that a graph H is a *subdivision* of G if H can be obtained from G by a finite sequence of subdivisions.

For example, the graph in Figure 1.83 contains a subdivision of K_5 , and in Figure 1.84, H is a subdivision of G .




FIGURE 1.84. A graph and a subdivision.

We leave the proof of the following theorem to the exercises (see Exercise 1).

Theorem 1.39. *A graph G is planar if and only if every subdivision of G is planar.*

Our list of nonplanar graphs now includes $K_{3,3}$, K_5 , graphs containing $K_{3,3}$ or K_5 as subgraphs, and all graphs containing a subdivision of $K_{3,3}$ or K_5 . The list so far stems from only two specific graphs: $K_{3,3}$ and K_5 . A well-known theorem by Kuratowski [185] tells us that there are *no other* graphs on the list! The bottom line is that $K_{3,3}$ and K_5 are the only two real enemies of planarity.

Kuratowski proved this beautiful theorem in 1930, closing a long-open problem.¹⁴ In 1954, Dirac and Schuster [78] found a proof that was slightly shorter than the original proof, and theirs is the proof that we will outline here.

Theorem 1.40 (Kuratowski’s Theorem). *A graph G is planar if and only if it contains no subdivision of $K_{3,3}$ or K_5 .*

Sketch of Proof

We have already discussed that if a graph G is planar, it contains no subgraph that is a subdivision of $K_{3,3}$ or K_5 . Thus we need to discuss only the reverse direction of the theorem.

¹⁴We should note here that Frink and Smith also discovered a proof of this fact in 1930, independently of Kuratowski. Since Kuratowski’s result was published first, his name has traditionally been associated with the theorem (and the names Frink and Smith have traditionally been associated with footnotes like this one.)

Suppose G is a graph that contains no subdivision of $K_{3,3}$ or K_5 . Here are the steps that Dirac and Schuster used to prove the result.

1. Prove that G is planar if and only if each block of G is planar. (A *block* of G is a maximal connected subgraph of G that has no cut vertex).
2. Explain why it suffices to show that a block is planar if and only if it contains no subdivision of $K_{3,3}$ or K_5 . Assume that G is a block itself (connected with no cut vertex).
3. Suppose that G is a nonplanar block that contains no subdivision of $K_{3,3}$ or K_5 (and search for a contradiction).
4. Prove that $\delta(G) \geq 3$.
5. Establish the existence of an edge $e = uv$ such that the graph $G - e$ is also a block.
6. Explain why $G - e$ is a planar graph containing a cycle C that includes both u and v , and choose C to have a maximum number of interior regions.
7. Establish several structural facts about the subgraphs inside and outside the cycle C .
8. Use these structural facts to demonstrate the existence of subdivisions of $K_{3,3}$ or K_5 , thus establishing the contradiction.

Exercises

1. Prove that a graph G is planar if and only if every subdivision of G is planar.
2. Use Kuratowski's Theorem to prove that the Petersen graph (Figure 1.63) is nonplanar.
3. Prove the first step of the proof of Kuratowski's Theorem.
4. Determine all complete multipartite graphs (of the form K_{r_1, \dots, r_n}) that are planar.

1.6 Colorings

One fish, two fish, red fish, blue fish.

— Dr. Seuss

The senators in a particular state sit on various senate committees, and the committees need to schedule times for meetings. Since each senator must be present at each of his or her committee meetings, the meeting times need to be scheduled carefully. One could certainly assign a unique meeting time to each of the

committees, but this plan may not be feasible, especially if the number of committees is large. We ask ourselves, given a particular committee structure, what is the fewest number of meeting times that are required? We can answer this question by studying graph coloring.

1.6.1 Definitions

Given a graph G and a positive integer k , a k -coloring is a function $K : V(G) \rightarrow \{1, \dots, k\}$ from the vertex set into the set of positive integers less than or equal to k . If we think of the latter set as a set of k “colors,” then K is an assignment of one color to each vertex.

We say that K is a *proper* k -coloring of G if for every pair u, v of adjacent vertices, $K(u) \neq K(v)$ — that is, if adjacent vertices are colored differently. If such a coloring exists for a graph G , we say that G is k -colorable.

For example, the graph C_5 as shown in Figure 1.85 is 3-colorable: $P(a) = P(c) = 1$, $P(b) = P(d) = 2$, $P(e) = 3$. Since C_5 is 3-colorable, a direct




FIGURE 1.85. C_5 is 3-colorable.

consequence of the definition is that C_5 is k -colorable for all $k \geq 3$. Is C_5 2-colorable?

Another way of viewing a proper k -coloring is as an assignment of vertices to sets, called *color classes*, where each set represents vertices that all receive the same color. For the coloring to be proper, each color class must be an independent set of vertices.

It is natural to wonder how many colors are necessary to color a particular graph G . For instance, we know that three colors are enough for the graph in Figure 1.85, but is this the least required? A quick check of C_5 reveals that coloring with two colors is impossible. So three colors are necessary. This idea motivates a definition.

Given a graph G , the *chromatic number* of G , denoted by $\chi(G)$, is the smallest integer k such that G is k -colorable. In our example, we can say that $\chi(C_5) = 3$. What about odd cycles in general? (Try one!) What about even cycles? (Try one!) Here is a list of chromatic numbers for some common graphs. Verify them!

$$\chi(C_n) = \begin{cases} 2 & \text{if } n \text{ is even,} \\ 3 & \text{if } n \text{ is odd,} \end{cases}$$

$$\chi(P_n) = \begin{cases} 2 & \text{if } n \geq 2, \\ 1 & \text{if } n = 1, \end{cases}$$

$$\chi(K_n) = n,$$

$$\chi(E_n) = 1,$$

$$\chi(K_{m,n}) = 2.$$

Exercises

1. Find the chromatic number of each of the following graphs. Explain your answers completely.
 - (a) Trees.
 - (b) Bipartite graphs.
 - (c) Complete multipartite graphs, K_{r_1, r_2, \dots, r_t} .
 - (d) The Petersen Graph.
 - (e) The graph in Figure 1.86, called the Birkhoff Diamond.
 - (f) The graphs of the regular polyhedra in Figure 1.82.




FIGURE 1.86. The Birkhoff Diamond.

2. Senate committees C_1 through C_7 consist of the members as indicated:
 $C_1 = \{\text{Adams, Bradford, Charles}\}$, $C_2 = \{\text{Charles, Davis, Eggers}\}$,
 $C_3 = \{\text{Davis, Ford}\}$, $C_4 = \{\text{Adams, Gardner}\}$, $C_5 = \{\text{Eggers, Howe}\}$,
 $C_6 = \{\text{Eggers, Bradford, Gardner}\}$, $C_7 = \{\text{Howe, Charles, Ford}\}$.
 Use the ideas of this section to determine the fewest number of meeting times that need to be scheduled for these committees.
3. When issuing seating assignments for his third grade students, the teacher wants to be sure that if two students might interfere with one another, then they are assigned to different areas of the room. There are six main troublemakers in the class: John, Jeff, Mike, Moe, Larry, and Curly. How many different areas are required in the room if John interferes with Moe and

Curly; Jeff interferes with Larry and Curly; Mike interferes with Larry and Curly; Moe interferes with John, Larry, and Curly; Larry interferes with Jeff, Mike, Moe, and Curly; and Curly interferes with everyone?

4. Prove that adding an edge to a graph increases its chromatic number by at most one.
5. Prove that a graph G of order at least two is bipartite if and only if it is 2-colorable.
6. A graph G is called *k-critical* if $\chi(G) = k$ and $\chi(G - v) < k$ for each vertex v of G .
 - (a) Find all 1-critical and 2-critical graphs.
 - (b) Give an example of a 3-critical graph.
 - (c) If G is k -critical, then show that G is connected.
 - (d) If G is k -critical, then show that $\delta(G) \geq k - 1$.
 - (e) Find all of the 3-critical graphs. Hint: Use part (d).

1.6.2 Bounds on Chromatic Number

The point is, ladies and gentlemen, that greed, for lack of a better word, is good. Greed is right. Greed works.

— Gordon Gekko, in *Wall Street*

In general, determining the chromatic number of a graph is hard. While small or well-known graphs (like the ones in the previous exercises) may be fairly easy, the number of possibilities in large graphs makes computing chromatic numbers difficult. We therefore often rely on bounds to give some sort of idea of what the chromatic number of a graph is, and in this section we consider some of these bounds.

If G is a graph on n vertices, then an obvious upper bound on $\chi(G)$ is n , since an n -coloring is always possible on a graph with n vertices. This bound is exact for complete graphs, as it takes as many colors as there are vertices to color a complete graph. In fact, complete graphs are the only graphs for which this bound is sharp (see Exercise 5). We set this aside as Theorem 1.41.

Theorem 1.41. *For any graph G of order n , $\chi(G) \leq n$.*

Let us now discuss a very basic graph coloring algorithm, the *greedy algorithm*. To color a graph having n vertices using this algorithm, first label the vertices in some order—call them v_1, v_2, \dots, v_n . Next, order the available colors in some way. We will denote them by the positive integers $1, 2, \dots, n$. Then start coloring by assigning color 1 to vertex v_1 . Next, if v_1 and v_2 are adjacent, assign color 2 to vertex v_2 ; otherwise, use color 1 again. In general, to color vertex v_i , use

the first available color that has not been used for any of v_i 's previously colored neighbors.

For example, the greedy algorithm produces the coloring on the right from the graph on the left in Figure 1.87. First, v_1 is assigned color 1; then v_2 is assigned color 1, since v_2 is not adjacent to v_1 . Then v_3 is assigned color 1 since it is not adjacent to v_1 or v_2 . Vertex v_4 is assigned color 2, then v_5 is assigned 2, and finally v_6 is assigned 2.




FIGURE 1.87. Applying the greedy algorithm.

It is important to realize that the coloring obtained by the greedy algorithm depends heavily on the initial labeling of the vertices. Different labelings can (and often do) produce different colorings. Figure 1.88 displays the coloring obtained from a different original labeling of the same graph. More colors are used in this




FIGURE 1.88. Applying it again.

second coloring. We see that while “greed works” in that the algorithm always gives a legal coloring, we cannot expect it to give us a coloring that uses the fewest possible colors.

The following bound improves Theorem 1.41.

Theorem 1.42. *For any graph G , $\chi(G) \leq \Delta(G) + 1$, where $\Delta(G)$ is the maximum degree of G .*

Proof. Running the greedy algorithm on G produces a legal coloring that uses at most $\Delta(G) + 1$ colors. This is because every vertex in the graph is adjacent to at most $\Delta(G)$ other vertices, and hence the largest color label used is at most $\Delta(G) + 1$. Thus, $\chi(G) \leq \Delta(G) + 1$. \square

Notice that we obtain equality in this bound for complete graphs and for cycles with an odd number of vertices. As it turns out, these are the only families of graphs for which the equality in Theorem 1.42 holds. This is stated in Brooks’s

Theorem [41]. The proof that we give is a modification of the one given by Lovász [190].

Theorem 1.43 (Brooks’s Theorem). *If G is a connected graph that is neither an odd cycle nor a complete graph, then $\chi(G) \leq \Delta(G)$.*

Proof. Let G be a connected graph of order n that is neither a complete graph nor an odd cycle. Let $k = \Delta(G)$. We know that $k \neq 0$ and $k \neq 1$, since otherwise G is complete. If $k = 2$, then G must be either an even cycle or a path. In either case, $\chi(G) = 2 = \Delta(G)$. So assume that $k = \Delta(G) \geq 3$.

We are now faced with three cases. In each case we will establish a labeling of the vertices of G in the form v_1, v_2, \dots, v_n . We will then use the greedy algorithm to color G with no more than k colors.

Case 1. Suppose that G is not k -regular. Then there exists some vertex with degree less than k . Choose such a vertex and call it v_n . Let $S_0 = \{v_n\}$ and let $S_1 = N(v_n)$, the neighborhood of v_n . Further, let

$$\begin{aligned} S_2 &= N(S_1) - \{v_n\} - S_1, \\ S_3 &= N(S_2) - S_1 - S_2, \\ &\vdots \\ S_i &= N(S_{i-1}) - S_{i-2} - S_{i-1}, \end{aligned}$$

for each i (Figure 1.89). Since G is finite, there is some t such that S_t is not empty,




FIGURE 1.89. The sets S_i .

and S_r is empty for all $r > t$.

Next, label the vertices in S_1 with the labels $v_{n-1}, v_{n-2}, \dots, v_{n-|S_1|}$. Label the vertices in S_2 with the labels $v_{n-|S_1|-1}, \dots, v_{n-|S_1|-|S_2|}$. Continue labeling in this decreasing fashion until all vertices of G have been labeled. The vertex with label v_1 is in the set S_t .

Let u be a vertex in some $S_i, i \geq 1$. Since u has at least one neighbor in S_{i-1} , it has at most $k - 1$ adjacencies with vertices whose label is less than its own. Thus, when the greedy algorithm gets to u , there will be at least one color from $\{1, 2, \dots, k\}$ available. Further, since $\deg(v_n) < k$, there will be a color from $\{1, 2, \dots, k\}$ available when the greedy algorithm reaches v_n . Thus, in this case the greedy algorithm uses at most k colors to properly color G .

Case 2. Suppose that G is k -regular and that G has a cut vertex, say v . The removal of v from G will form at least two connected components. Say the components are G_1, G_2, \dots, G_t . Consider the graph $H_1 = G_1 \cup \{v\}$ (the component G_1 with v added back—see Figure 1.90). H_1 is a connected graph, and the degree




FIGURE 1.90. The graph H_1 .

of v in H_1 is less than k . Using the method in Case 1, we can properly color H_1 with at most k colors. Similarly, we can properly color each $H_i = G_i - \{v\}$ with at most k colors. Without loss of generality, we can assume that v gets the same color in all of these colorings (if not, just permute the colors to make it so). These colorings together create a proper coloring of G that uses at most k colors. Case 2 is complete.

Case 3. Suppose that G is k -regular and that it does not contain a cut vertex. This means that G is 2-connected.

Subcase 3a. Suppose that G is 3-connected. This means that for all v , the graph $G - v$ is 2-connected. Let v be a vertex of G with neighbors v_1 and v_2 such that $v_1v_2 \notin E(G)$ (such vertices exist since G is not complete). By the assumption in this subcase, the graph $G - \{v_1, v_2\}$ is connected.

Subcase 3b. Suppose that G is not 3-connected. This means that there exists a pair of vertices v, w such that the graph $G - \{v, w\}$ is disconnected. Let the components of $G - \{v, w\}$ be G_1, G_2, \dots, G_t . Since $k \geq 3$, it must be that each G_i has at least two vertices. It also must be that v is adjacent to at least one vertex in each G_i , since w is not a cut vertex of G . Let $u \in V(G_1)$ be a neighbor of v . Suppose for the moment that u is a cut vertex of the graph $G - v$. If this is the case, then there must be another vertex y of G_1 such that (i) y is not a cut vertex of the graph $G - v$, and (ii) the only paths from y to w in $G - v$ go through vertex u . Since u is not a cut vertex of G itself, it must be that y is adjacent to v . In either case, it must be that v has a neighbor in G_1 (either u or y) that is not a cut vertex of $G - v$. The vertex v has a similar such neighbor in G_2 . For convenience, let us rename: For $i = 1, 2$, let $v_i \in V(G_i)$ be a neighbor of v that is not a cut vertex of the graph $G - v$. Vertices v_1 and v_2 are nonadjacent, and since they were in different components of $G - \{v, w\}$, it must be that $G - \{v_1, v_2\}$ is connected.

In each subcase, we have identified vertices v, v_1 , and v_2 such that $vv_1, vv_2 \in E(G)$, $v_1v_2 \notin E(G)$, and $G - \{v_1, v_2\}$ is connected. We now proceed to label the vertices of G in preparation for the greedy algorithm.

Let v_1 and v_2 be as labeled. Let v be labeled v_n . Now choose a vertex adjacent to v_n that is not v_1 or v_2 (such a vertex exists, since $\deg(v_n) \geq 3$). Label this vertex v_{n-1} . Next choose a vertex that is adjacent to either v_n or v_{n-1} and is not v_1, v_2, v_n , or v_{n-1} . Call this vertex v_{n-2} . We continue this process. Since $G - \{v_1, v_2\}$ is connected, then for each $i \in \{3, \dots, n-1\}$, there is a vertex $v_i \in V(G) - \{v_1, v_2, v_n, v_{n-1}, \dots, v_{i+1}\}$ that is adjacent to at least one of v_{i+1}, \dots, v_n .

Now that the vertices are labeled, we can apply the greedy algorithm. Since $v_1 v_2 \notin E(G)$, the algorithm will give the color 1 to both v_1 and v_2 . Since each $v_i, 3 \leq i < n$, is adjacent to at most $k - 1$ predecessors, and since v_n is adjacent to v_1 and v_2 , the algorithm never requires more than $k = \Delta(G)$ colors. Case 3 is complete. \square

The next bound involves a new concept.

The *clique number* of a graph, denoted by $\omega(G)$, is defined as the order of the largest complete graph that is a subgraph of G . For example, in Figure 1.91, $\omega(G_1) = 3$ and $\omega(G_2) = 4$.




FIGURE 1.91. Graphs with clique numbers 3 and 4, respectively.

A simple bound that involves clique number follows. We leave it to the reader to provide a (one or two line) explanation.

Theorem 1.44. *For any graph G , $\chi(G) \geq \omega(G)$.*

It is natural to wonder whether we might be able to strengthen this theorem and prove that $\chi(G) = \omega(G)$ for every graph G . Unfortunately, this is false. Consider the graph G shown in Figure 1.92. The clique number of this graph is 5, and the




FIGURE 1.92. Is $\chi(G) = \omega(G)$?

chromatic number is 6 (see Exercise 2).

The upper and lower bounds given in Theorem 1.45 concern $\alpha(G)$, the independence number of G , defined back in Section 1.4.3. The proofs are left as an exercise (see Exercise 6).

Theorem 1.45. *For any graph G of order n ,*

$$\frac{n}{\alpha(G)} \leq \chi(G) \leq n + 1 - \alpha(G).$$

Exercises

1. Recall that $\text{avgdeg}(G)$ denotes the average degree of vertices in G . Prove or give a counterexample to the following statement:

$$\chi(G) \leq 1 + \text{avgdeg}(G).$$

2. If G is the graph in Figure 1.92, prove that $\chi(G) = 6$ and $\omega(G) = 5$.
3. Determine a necessary and sufficient condition for a graph to have a 2-colorable line graph.
4. Recall that $\tau(G)$ denotes the number of vertices in a detour path (a longest path) of G , prove that $\chi(G) \leq \tau(G)$.
5. Prove that the only graph G of order n for which $\chi(G) = n$ is K_n .
6. Prove that for any graph G of order n ,

$$\frac{n}{\alpha(G)} \leq \chi(G) \leq n + 1 - \alpha(G).$$

7. If G is bipartite, prove that $\omega(\overline{G}) = \chi(\overline{G})$.
8. Let G be a graph of order n . Prove that
 - (a) $n \leq \chi(G)\chi(\overline{G})$;
 - (b) $2\sqrt{n} \leq \chi(G) + \chi(\overline{G})$.

1.6.3 The Four Color Problem

That doesn't sound too hard.

— Princess Leia, *Star Wars*

The Four Color Problem. *Is it true that the countries on any given map can be colored with four or fewer colors in such a way that adjacent countries are colored differently?*

The seemingly simple Four Color Problem was introduced in 1852 by Francis Guthrie, a student of Augustus DeMorgan. The first written reference to the problem is a letter from DeMorgan to Sir William Rowan Hamilton. Despite Hamilton's indifference¹⁵, DeMorgan continued to talk about the problem with other mathematicians. In the years that followed, many of the world's top mathematical minds attempted either to prove or disprove the conjecture, and in 1879 Alfred Kempe announced that he had found a proof. In 1890, however, P. J. Heawood discovered an error in Kempe's proof. Kempe's work did have some positive features, though, for Heawood made use of Kempe's ideas to prove that five colors always suffice. In this section, we translate the Four Color Problem into a graph theory problem, and we prove the Five Color Theorem.

Any map can be represented by a planar graph in the following way: Represent each country on the map by a vertex, and connect two vertices with an edge whenever the corresponding countries share a nontrivial border (more than just a point). Some examples are shown in Figure 1.93.




FIGURE 1.93. Graph representations of maps.

The regions on the map correspond to vertices on the graph, so a graph coloring yields a map coloring with no bordering regions colored the same. This natural representation allows us to see that a map is 4-colorable if and only if its associated graph is 4-colorable.

The Four Color Conjecture is equivalent to the following statement. A thorough discussion of this equivalence can be found in [52].

¹⁵Perhaps he was too busy perfecting plans for a cool new game that he would release a few years later. See Section 1.4.3.

Theorem 1.46 (Four Color Theorem). *Every planar graph is 4-colorable.*

When Heawood pointed out the error in Kempe's proof, researchers flocked back to the drawing board. People worked on the Four Color Problem for years and years trying numerous strategies. Finally, in 1976, Kenneth Appel and Wolfgang Haken, with the help of John Koch, announced that they had found a proof [12]. To complete their proof, they verified thousands of cases with computers, using over 1000 hours of computer time. As you might imagine, people were skeptical of this at first. Was this really a proof? How could an argument with so many cases be verified?

While the Appel–Haken proof is accepted as being valid, mathematicians still search for alternative proofs. Robertson, Sanders, Seymour, and Thomas [239] have probably come the closest to finding a short and clever proof, but theirs still requires a number of computer calculations.

In a 1998 article [267], Robin Thomas said the following.

For the purposes of this survey, let me telescope the difficulties with the A&H proof into two points: (1) part of the proof uses a computer and cannot be verified by hand, and (2) even the part that is supposedly hand-checkable has not, as far as I know, been independently verified in its entirety. . . . Neil Robertson, Daniel P. Sanders, Paul Seymour, and I tried to verify the Appel–Haken proof, but soon gave up and decided that it would be more profitable to work out our own proof. . . . We were not able to eliminate reason (1), but we managed to make progress toward (2).

As mentioned earlier, Heawood [156] provided a proof of the Five Color Theorem in the late 1890s, and we present his proof here. Some of the ideas in his proof came from Kempe's attempt [174] to solve the Four Color Problem.

Theorem 1.47 (Five Color Theorem). *Every planar graph is 5-colorable.*

Proof. We induct on the order of G . Let G be a planar graph of order n . If $n \leq 5$, then the result is clear. So suppose that $n \geq 6$ and that the result is true for all planar graphs of order $n - 1$. From Theorem 1.35, we know that G contains a vertex, say v , having $\deg(v) \leq 5$.

Consider the graph G' obtained by removing from G the vertex v and all edges incident with v . Since the order of G' is $n - 1$ (and since G' is of course planar), we can apply the induction hypothesis and conclude that G' is 5-colorable. Now, we can assume that G' has been colored using the five colors, named 1, 2, 3, 4, and 5. Consider now the neighbors of v in G . As noted earlier, v has at most five neighbors in G , and all of these neighbors are vertices in (the already colored) G' .

If in G' fewer than five colors were used to color these neighbors, then we can properly color G by using the coloring for G' on all vertices other than v , and by coloring v with one of the colors that is not used on the neighbors of v . In doing this, we have produced a 5-coloring for G .

So, assume that in G' exactly five of the colors were used to color the neighbors of v . This implies that there are exactly five neighbors, call them w_1, w_2, w_3, w_4, w_5 , and assume without loss of generality that each w_i is colored with color i (see Figure 1.94).




FIGURE 1.94.

We wish to rearrange the colors of G' so that we make a color available for v . Consider all of the vertices of G' that have been colored with color 1 or with color 3.

Case 1. Suppose that in G' there does not exist a path from w_1 to w_3 where all of the colors on the path are 1 or 3. Define a subgraph H of G' to be the union of all paths that start at w_1 and that are colored with either 1 or 3. Note that w_3 is not a vertex of H and that none of the neighbors of w_3 are in H (see Figure 1.95).




FIGURE 1.95.

Now, interchange the colors in H . That is, change all of the 1's into 3's and all of the 3's into 1's. The resulting coloring of the vertices of G' is a proper coloring, because no problems could have possibly arisen in this interchange. We now see that w_1 is colored 3, and thus color 1 is available to use for v . Thus, G is 5-colorable.

Case 2. Suppose that in G' there does exist a path from w_1 to w_3 where all of the colors on the path are 1 or 3. Call this path P . Note now that P along with v forms a cycle that encloses either w_2 or w_4 (Figure 1.96).




FIGURE 1.96. Two possibilities.

So there does not exist a path from w_2 to w_4 where all of the colors on the path are 2 or 4. Thus, the reasoning in Case 1 applies! We conclude that G is 5-colorable. \square

Exercises

1. Determine the chromatic number of the graph of the map of the United States.
2. Determine the chromatic number of the graph of the map of the countries of South America.
3. Determine the chromatic number of the graph of the map of the countries of Africa.
4. Determine the chromatic number of the graph of the map of the countries of Australia. Hint: This graph will be quite small!
5. Where does the proof of the Five Color Theorem go wrong for four colors?

1.6.4 Chromatic Polynomials

Everything should be made as simple as possible, but not simpler.

— Albert Einstein

Chromatic polynomials, developed by Birkhoff in the early 1900s as he studied the Four Color Problem, provide us with a method of counting the number of different colorings of a graph.

Before we introduce the polynomials, we should clarify what we mean by different colorings. Given a graph G , suppose that its vertices are labeled v_1, v_2, \dots, v_n . A coloring of G is an assignment of colors to these vertices, and we call two colorings C_1 and C_2 *different* if at least one v_i receives a different color in C_1 than it does in C_2 . For instance, the two colorings of K_4 shown in Figure 1.97 are considered different, since v_3 and v_4 receive different colorings.




FIGURE 1.97. Two different colorings.

If we restrict ourselves to four colors, how many different colorings are there of K_4 ? Since there are four choices for v_1 , then three for v_2 , etc., we see that there are $4 \cdot 3 \cdot 2 \cdot 1$ different colorings of K_4 using four colors. If six colors were available, there would be $6 \cdot 5 \cdot 4 \cdot 3$ different colorings. If only two were available, there would be no proper colorings of K_4 .

In general, define $c_G(k)$ to be the number of different colorings of a graph G using at most k colors. So we have $c_{K_4}(4) = 24$, $c_{K_4}(6) = 360$, and $c_{K_4}(2) = 0$. In fact, if k and n are positive integers where $k \geq n$, then

$$c_{K_n}(k) = k(k-1)(k-2) \cdots (k-n+1).$$

Further, if $k < n$, then $c_{K_n}(k) = 0$. We also note that $c_{E_n}(k) = k^n$ for all positive integers k and n .

A simple but important property of $c_G(k)$ is that G is k -colorable if and only if $c_G(k) > 0$. Equivalently, $c_G(k) > 0$ if and only if $\chi(G) \leq k$.

Finding values of $c_G(k)$ is relatively easy for some well-known graphs. Computing this function in general, though, can be hard. Birkhoff and Lewis [27] developed a way to reduce this hard problem to an easier one. Before we see their method, we need a definition.

Let G be a graph and let e be an edge of G . Recall that $G - e$ denotes the graph where e is removed from G . Define the graph G/e to be the graph obtained from G by removing e , identifying the end vertices of e , and leaving only one copy of any resulting multiple edges.

As an example, a graph G and the graphs $G - bc$ and G/bc are shown in Figure 1.98.

Theorem 1.48. *Let G be a graph and e be any edge of G . Then*

$$c_G(k) = c_{G-e}(k) - c_{G/e}(k).$$

Proof. Suppose that the end vertices of e are u and v , and consider the graph $G - e$.

How many k -colorings are there of $G - e$ where u and v are assigned the same color? If C is a such a coloring of $G - e$, then C can be thought of as a coloring of G/e , since u and v are colored the same. Similarly, any coloring of G/e can also be thought of as a coloring of $G - e$ where u and v are colored the same. Thus, the answer to this question is $c_{G/e}(k)$.




FIGURE 1.98. Examples of the operations.

Now, how many k -colorings are there of $G - e$ where u and v are assigned different colors? If C is a such a coloring of $G - e$, then C can be considered as a coloring of G , since u and v are colored differently. Similarly, any coloring of G can also be thought of as a coloring of $G - e$ where u and v are colored differently. Thus, the answer to this second question is $c_G(k)$.

Thus, the total number of k -colorings of $G - e$ is

$$c_{G-e}(k) = c_{G/e}(k) + c_G(k),$$

and the result follows. □

For example, suppose we want to find $c_{P_4}(k)$. That is, how many ways are there to color the vertices of P_4 with k colors available? We label the edges of P_4 as shown in Figure 1.99.




FIGURE 1.99. The labeled edges of P_4 .

The theorem implies that

$$c_{P_4}(k) = c_{P_4 - e_1}(k) + c_{P_4/e_1}(k).$$

For convenience, let us denote $P_4 - e_1$ and P_4/e_1 by G_{11} and G_{12} , respectively (see Figure 1.100).



FIGURE 1.100. The first application.

Applying the theorem again, we obtain

$$c_{P_4}(k) = c_{G_{11}-e_2}(k) - c_{G_{11}/e_2}(k) - c_{G_{12}-e_2}(k) + c_{G_{12}/e_2}(k).$$

Denote the graphs $G_{11} - e_2$, G_{11}/e_2 , $G_{12} - e_2$, and G_{12}/e_2 by G_{21} , G_{22} , G_{23} , and G_{24} , respectively (see Figure 1.101).




FIGURE 1.101. The second application.

Applying the theorem once more yields

$$c_{P_4}(k) = c_{G_{21}-e_3}(k) - c_{G_{21}/e_3}(k) - c_{G_{22}-e_3}(k) + c_{G_{22}/e_3}(k) - c_{G_{23}-e_3}(k) + c_{G_{23}/e_3}(k) + c_{G_{24}-e_3}(k) - c_{G_{24}/e_3}(k).$$

That is,

$$c_{P_4}(k) = c_{E_4}(k) - c_{E_3}(k) - c_{E_3}(k) + c_{E_2}(k) - c_{E_3}(k) + c_{E_2}(k) + c_{E_2}(k) - c_{E_1}(k).$$

Thus,

$$\begin{aligned} c_{P_4}(k) &= k^4 - k^3 - k^3 + k^2 - k^3 + k^2 + k^2 - k \\ &= k^4 - 3k^3 + 3k^2 - k. \end{aligned}$$

We should check a couple of examples. How many colorings of P_4 are there with one color?

$$c_{P_4}(1) = 1^4 - 3(1)^3 + 3(1)^2 - 1 = 0.$$

This, of course, makes sense. And how many colorings are there with two colors?

$$c_{P_4}(2) = 2^4 - 3(2)^3 + 3(2)^2 - 2 = 2.$$

Figure 1.102 shows these two colorings. Score one for Birkhoff!



FIGURE 1.102. Two 2-colorings of P_4 .

As you can see, chromatic polynomials provide a way to count colorings, and the Birkhoff–Lewis theorem allows you to reduce a problem to a slightly simpler one. We should note that it is not always necessary to work all the way down to empty graphs, as we did in the previous example. Once a graph G is obtained for which the value of $c_G(k)$ is known, there is no need to reduce that one further.

We now present some properties of $c_G(k)$.

Theorem 1.49. *Let G be a graph of order n . Then*

1. $c_G(k)$ is a polynomial in k of degree n ,
2. the leading coefficient of $c_G(k)$ is 1,
3. the constant term of $c_G(k)$ is 0,
4. the coefficients of $c_G(k)$ alternate in sign, and
5. the absolute value of the coefficient of the k^{n-1} term is the number of edges in G .

We leave the proof of this theorem as an exercise (Exercise 3). One proof strategy is to induct on the number of edges in G and use the Birkhoff–Lewis reduction theorem (Theorem 1.48).

Before leaving this section, we should note that Birkhoff considered chromatic polynomials of planar graphs, and he hoped to find one of them that had 4 as a root. If he had found one, then the corresponding planar graph would not be 4-colorable, and hence would be a counterexample to the Four Color Conjecture. Although he was unsuccessful in proving the Four Color Theorem, he still deserves credit for producing a very nice counting technique.

Exercises

1. Find chromatic polynomials for each of the following graphs. For each one, determine how many 5-colorings exist.
 - (a) $K_{1,3}$
 - (b) $K_{1,5}$
 - (c) C_4
 - (d) C_5
 - (e) $K_4 - e$
 - (f) $K_5 - e$
2. Show that $k^4 - 4k^3 + 3k^2$ is not a chromatic polynomial for any graph.
3. Prove Theorem 1.49.
4. Determine the chromatic polynomial for a tree of order n .

1.7 Matchings

Pardon me, do you have a match?

— James Bond, in *From Russia with Love*

The Senate committees that we discussed earlier need to form an executive council. Each committee needs to designate one of its members as an official representative to sit on the council, and council policy states that no senator can be the official representative for more than one committee. For example, let us suppose there are four committees: Senators $A, B, C,$ and D are on Committee 1; Senators $A, E,$ and F are members of Committee 2; Committee 3 consists of $E, F,$ and D ; and Senator A is the only member of Committee 4. In this example, the executive council could consist of $A, E, F,$ and C —representing Committees 4, 3, 2, and 1, respectively.

As another example, suppose Committee 1 consists of W, X and Y ; Committee 2 of $W, X,$ and Z ; Committee 3 of $W, Y,$ and Z ; Committee 4 of $X, Y,$ and Z ; and Committee 5 of W and Z . It does not take long to see that it is impossible in this case to select official representatives according to the policy.

So a natural question arises: Under what circumstances can the executive council be formed successfully? In the sections that follow, we will see how graphs can be used to help answer this question.

1.7.1 Definitions

And as to the meaning . . .

— C. S. Calverly, *Ballad*

A *matching* in a graph is a set of independent edges. That is, it is a set of edges in which no pair shares a vertex. Given a matching M in a graph G , the vertices belonging to the edges of M are said to be *saturated* by M (or *M -saturated*). The other vertices are *M -unsaturated*.

Consider the graph G shown in Figure 1.103. An example of a matching in G is $M_1 = \{ab, ce, df\}$. $M_2 = \{cd, ab\}$ is also a matching, and so is $M_3 = \{df\}$. We can see that a, b, c, d are M_2 -saturated and $e, f,$ and g are M_2 -unsaturated. The only M_1 -unsaturated vertex is g .




FIGURE 1.103. The matching M_1 .

If a matching M saturates every vertex of G , then M is said to be a *perfect matching*. In Figure 1.104, G_1 has a perfect matching, namely $\{ab, ch, de, fg\}$. None of $G_2, G_3,$ and G_4 has a perfect matching. Why is this? We will talk more about perfect matchings in Section 1.7.4.

A *maximum matching* in a graph is a matching that has the largest possible cardinality. A *maximal matching* is a matching that cannot be enlarged by the




FIGURE 1.104. Only G_1 has a perfect matching.

addition of any edge. In Figure 1.105, $M_1 = \{ae, bf, cd, gh\}$ is a maximum matching (since at most one of $gh, gi,$ and gj can be in any matching). The matching $M_2 = \{dg, af, bc\}$ is maximal, but not maximum.




FIGURE 1.105.

Exercises

1. Determine whether the graph of Figure 1.106 has a perfect matching. If so, then exhibit it. If not, explain why.




FIGURE 1.106. Is there a perfect matching?

2. Find the minimum size of a maximal matching in each of the following graphs.

(a) C_{10}

- (b) C_{11}
 - (c) C_n
3. (From [52].) The *matching graph* $M(G)$ of a graph G has the maximum matchings of G as its vertices, and two vertices M_1 and M_2 of $M(G)$ are adjacent if M_1 and M_2 differ in only one edge. Show that each cycle C_n , $n = 3, 4, 5$, or 6 , is the matching graph of some graph.

1.7.2 Hall’s Theorem and SDRs

I’ll match that!

— Monty Hall, *Let’s Make a Deal*

In this section we consider several classic results concerning matchings. We begin with a few more definitions.

Given a graph G and a matching M , an M -*alternating path* is a path in G where the edges alternate between M -edges and non- M -edges. An M -*augmenting path* is an M -alternating path where both end vertices are M -unsaturated.

As an example, consider the graph G and the matching M indicated in Figure 1.107. An example of an M -alternating path is c, a, d, e, i . An example of an M -augmenting path is j, g, f, a, c, b . The reason for calling such a path “ M -augmenting” will become apparent soon.




FIGURE 1.107. The graph G and matching M .

The following result is due to Berge [23].

Theorem 1.50 (Berge’s Theorem). *Let M be a matching in a graph G . M is maximum if and only if G contains no M -augmenting paths.*

Proof. First, assume that M is a maximum matching, and suppose that $P : v_1, v_2, \dots, v_k$ is an M -augmenting path. Due to the alternating nature of M -augmenting paths, it must be that k is even and that the edges $v_2v_3, v_4v_5, \dots, v_{k-2}v_{k-1}$ are all edges of M . We also see that the edges $v_1v_2, v_3v_4, \dots, v_{k-1}v_k$ are not edges of M (Figure 1.108).

But then if we define the set of edges M_1 to be

$$M_1 = (M \setminus \{v_2v_3, \dots, v_{k-2}v_{k-1}\}) \cup \{v_1v_2, \dots, v_{k-1}v_k\},$$




FIGURE 1.108. An M -augmenting path.

then M_1 is a matching that contains one more edge than M , a matching that we assumed to be maximum. This is a contradiction, and we can conclude that G contains no M -augmenting paths.

For the converse, assume that G has no M -augmenting paths, and suppose that M' is a matching that is larger than M . Define a subgraph H of G as follows: Let $V(H) = V(G)$ and let $E(H)$ be the set of edges of G that appear in exactly one of M and M' . Now consider some properties of this subgraph H . Since each of the vertices of G lies on at most one edge from M and at most one edge from M' , it must be that the degree (in H) of each vertex of H is at most 2. This implies that each connected component of H is either a single vertex, a path, or a cycle. If a component is a cycle, then it must be an even cycle, since the edges alternate between M -edges and M' -edges. So, since $|M'| > |M|$, there must be at least one component of H that is a path that begins and ends with edges from M' . But this path is an M -augmenting path, contradicting our assumption. Therefore, no such matching M' can exist—implying that M is maximum. \square

Before we see Hall’s classic matching theorem, we need to define one more term. If G is a bipartite graph with partite sets X and Y , we say that X can be *matched into* Y if there exists a matching in G that saturates the vertices of X .

Consider the two examples in Figure 1.109. In the bipartite graph on the left,




FIGURE 1.109.

we see that X can be matched into Y . In the graph on the right, though, it is impossible to match X into Y (why is this?). What conditions on a bipartite graph must exist if we want to match one partite set into the other? The answer to this question is found in the following result of Hall [147] (Philip, not Monty).

Recall that the neighborhood of a set of vertices S , denoted by $N(S)$, is the union of the neighborhoods of the vertices of S .

Theorem 1.51 (Hall's Theorem). *Let G be a bipartite graph with partite sets X and Y . X can be matched into Y if and only if $|N(S)| \geq |S|$ for all subsets S of X .*

Proof. First suppose that X can be matched into Y , and let S be some subset of X . Since S itself is also matched into Y , we see immediately that $|S| \leq |N(S)|$ (see Figure 1.110). Now suppose that $|N(S)| \geq |S|$ for all subsets S of X , and




FIGURE 1.110.

let M be a maximum matching. Suppose that $u \in X$ is not saturated by M (see Figure 1.111). Define the set A to be the set of vertices of G that can be joined to u




FIGURE 1.111.

by an M -alternating path. Let $S = A \cap X$, and let $T = A \cap Y$ (see Figure 1.112). Notice now that Berge's Theorem implies that every vertex of T is saturated by




FIGURE 1.112.

M and that u is the only unsaturated vertex of S . That is, every vertex of T is saturated, and every vertex of $S \setminus \{u\}$ is saturated. This implies that $|T| = |S| - 1$. It follows from Berge's Theorem and the definition of T that $N(S) = T$. But then we have that $|N(S)| = |S| - 1 < |S|$, and this is a contradiction. We conclude that such a vertex u cannot exist in X and that M saturates all of X . \square

Given some family of sets X , a *system of distinct representatives*, or SDR, for the sets in X can be thought of as a "representative" collection of distinct elements from the sets of X . For instance, let S_1, S_2, S_3, S_4 , and S_5 be defined as follows:

$$\begin{aligned} S_1 &= \{2, 8\}, \\ S_2 &= \{8\}, \\ S_3 &= \{5, 7\}, \\ S_4 &= \{2, 4, 8\}, \\ S_5 &= \{2, 4\}. \end{aligned}$$

The family $X_1 = \{S_1, S_2, S_3, S_4\}$ does have an SDR, namely $\{2, 8, 7, 4\}$. The family $X_2 = \{S_1, S_2, S_4, S_5\}$ does not have an SDR.

So under what conditions will a finite family of sets have an SDR? We answer this question with the following theorem.

Theorem 1.52. *Let S_1, S_2, \dots, S_k be a collection of finite, nonempty sets. This collection has an SDR if and only if for every $t \in \{1, \dots, k\}$, the union of any t of these sets contains at least t elements.*

Proof. Since each of the sets is finite, then of course $S = S_1 \cup S_2 \cup \dots \cup S_k$ is finite. Let us say that the elements of S are a_1, \dots, a_n .

We now construct a bipartite graph with partite sets $X = \{S_1, \dots, S_k\}$ and $Y = \{a_1, \dots, a_n\}$ (Figure 1.113). We place an edge between S_i and a_j if and only if $a_j \in S_i$.




FIGURE 1.113. Constructing a bipartite graph.

Hall's Theorem now implies that X can be matched into Y if and only if $|A| \leq |N(A)|$ for all subsets A of X . In other words, the collection of sets has an SDR if and only if for every $t \in \{1, \dots, k\}$, the union of any t of these sets contains at least t elements. \square

Hall's Theorem is often referred to as Hall's Marriage Theorem. We will see more about this in Section 2.9.

Exercises

1. (From [56].) For the graphs of Figure 1.114, with matchings M as shaded, find
 - (a) an M -alternating path that is not M -augmenting;
 - (b) an M -augmenting path if one exists; and, if so, use it to obtain a bigger matching.




FIGURE 1.114.

2. For each of the following families of sets, determine whether the condition of Theorem 1.52 is met. If so, then find an SDR. If not, then show how the condition is violated.
 - (a) $\{1, 2, 3\}, \{2, 3, 4\}, \{3, 4, 5\}, \{4, 5\}, \{1, 2, 5\}$
 - (b) $\{1, 2, 4\}, \{2, 4\}, \{2, 3\}, \{1, 2, 3\}$
 - (c) $\{1, 2\}, \{2, 3\}, \{1, 2, 3\}, \{2, 3, 4\}, \{1, 3\}, \{3, 4\}$
 - (d) $\{1, 2, 5\}, \{1, 5\}, \{1, 2\}, \{2, 5\}$
 - (e) $\{1, 2, 3\}, \{1, 2, 4\}, \{1, 3, 4\}, \{1, 2, 3, 4\}, \{2, 3, 4\}$
3. Let G be a bipartite graph. Show that G has a matching of size at least $|E(G)|/\Delta(G)$.
4. Let $\Theta = \{S_1, S_2, \dots, S_r\}$ be a family of distinct nonempty subsets of the set $\{1, 2, \dots, n\}$. If the S_i are all of the same cardinality, then prove that there exists an SDR of Θ .

5. Let M_1 and M_2 be matchings in a bipartite graph G with partite sets X and Y . If $S \subseteq X$ is saturated by M_1 and $T \subseteq Y$ is saturated by M_2 , show that there exists a matching in G that saturates $S \cup T$.
6. (From [139].) Let G be a bipartite graph with partite sets X and Y . Let δ_X denote the minimum degree of the vertices in X , and let Δ_Y denote the maximum degree of the vertices in Y . Prove that if $\delta_X \geq \Delta_Y$, then there exists a matching in G that saturates X .

1.7.3 The König–Egerváry Theorem

What I tell you three times is true.

— Lewis Carroll, *The Hunting of the Snark*

The main theorem that we present in this section is very important, for it is closely related to several results from other areas of graph theory. We will discuss a few of these areas after we have proven the theorem.

A set C of vertices in a graph G is said to *cover* the edges of G if every edge of G is incident with at least one vertex of C . Such a set C is called an *edge cover* of G .

Consider the graphs G_1 and G_2 in Figure 1.115. In G_1 , the set $\{b, d, e, a\}$ is an edge cover, as is the set $\{a, e, f\}$. In fact, you can see by a little examination that there is no edge cover G_1 with fewer than three vertices. So we can say that $\{a, e, f\}$ is a minimum edge cover of G_1 . In G_2 , each of the following sets is an edge cover: $\{v_1, v_2, v_3, v_4, v_5, v_6\}$ (obviously) and $\{u_2, v_6, u_1\}$. What is the size of a minimum edge cover here?




FIGURE 1.115.

We are now ready to prove the following result of König [180] and Egerváry [87].

Theorem 1.53 (König–Egerváry Theorem). *Let G be a bipartite graph. The maximum number of edges in a matching in G equals the minimum number of vertices in an edge cover of G .*

Proof. Let M be a maximum matching of G . Let X and Y be the partite sets of G , and let W be the set of all M -unsaturated vertices of X (see Figure 1.116). Note that $|M| = |X| - |W|$.




FIGURE 1.116.

Now let A be the set of vertices of G that can be reached via an M -alternating path from some vertex of W . Let $S = A \cap X$, and let $T = A \cap Y$. We can note two things now: First, $S \setminus W$ is matched to T (implying that $|W| = |S| - |T|$), and second, $N(S) = T$.

If we let $C = (X \setminus S) \cup T$, we see that C covers the edges of G . So C is an edge cover of G , and $|C| = |X| - |S| + |T| = |X| - |W| = |M|$. Now suppose that C' is any edge cover. Since each vertex of C' can cover at most one edge of M , it must be that $|C'| \geq |M|$. We conclude then that C is a minimum edge cover. \square

The König–Egerváry Theorem is one of several theorems in graph theory that relate the minimum of one thing to the maximum of something else. What follows are some examples of theorems that are very closely related to the König–Egerváry Theorem.

Menger's Theorem

Let G be a connected graph, and let u and v be vertices of G . If S is a subset of vertices that does not include u or v , and if the graph $G - S$ has u and v in different connected components, then we say that S is a u, v -separating set.

The following result is known as Menger's Theorem [202].

Theorem 1.54. *Let G be a graph and let u and v be vertices of G . The maximum number of internally disjoint paths from u to v equals the minimum number of vertices in a u, v -separating set.*

Max Flow Min Cut Theorem

A graph can be thought of as a flow network, where one vertex is specified to be the source of the flow and another is specified to be the receiver of the flow. As an amount of material flows from source to receiver, it passes through other intermediate vertices, each of which has a particular flow capacity. The *total flow* of a network is the amount of material that is able to make it from source to receiver. A *cut* in a network is a set of intermediate vertices whose removal completely cuts

the flow from the source to the receiver. The *capacity of the cut* is simply the sum of the capacities of the vertices in the cut.

Theorem 1.55. *Let N be a flow network. The maximum value of total flow equals the minimum capacity of a cut.*

Independent Zeros

If A is an $m \times n$ matrix with real entries, a set of *independent zeros* in A can be thought of as a set of ordered pairs $\{(i_1, j_1), (i_2, j_2), \dots, (i_t, j_t)\}$ with the following properties:

- a. the (i_k, j_k) entry of A is 0 for $k = 1, 2, \dots, t$;
- b. if $a \neq b$, then $i_a \neq i_b$ and $j_a \neq j_b$.

That is, none of the zeros in the set are in the same row or column.

Now, in this matrix A one can draw lines through each row and column that contains a zero. Such a set of lines is said to cover the zeros of A .

Theorem 1.56. *The maximum number of independent zeros in A is equal to the minimum number of lines through rows or columns that together cover all the zeros of A .*

Exercises

1. Use the König–Egerváry Theorem to prove Hall’s Theorem.
2. Let k be some fixed integer, $1 \leq k \leq n$, and let G be some subgraph of $K_{n,n}$ with more than $(k - 1)n$ edges. Prove that G has a matching of size at least k .
3. Use the original statement of the König–Egerváry Theorem to prove Theorem 1.56.

1.7.4 Perfect Matchings

It’s a perfect ending.

— Sophie, in *Anastasia*

We end this section on matchings by discussing perfect matchings. Recall that a perfect matching is a matching that saturates the entire vertex set of a graph. What kinds of graphs have perfect matchings? One thing that is clear is that a graph must be of even order in order to have a chance at having a perfect matching. But being of even order is certainly not enough to guarantee a perfect matching (look back at Figure 1.105).

We do know that K_{2n} , C_{2n} , and P_{2n} have perfect matchings. The following result regarding perfect matchings in bipartite graphs is a corollary to Hall’s Theorem. The proof is left as an exercise (Exercise 5).

Corollary 1.57. *If G is a bipartite graph that is regular of degree k , then G contains a perfect matching.*

It seems very natural to think that the more edges a graph has, the more likely it is that the graph will have a perfect matching. The following theorem verifies this thought, to a degree.

Theorem 1.58. *If G is a graph of order $2n$ such that $\delta(G) \geq n$, then G has a perfect matching.*

Proof. Let G be a graph of order $2n$ with $\delta(G) \geq n$. Dirac’s theorem (Theorem 1.22) guarantees the existence of a Hamiltonian cycle, C . A perfect matching of G is formed by using alternate edges of C . □

In 1947 Tutte [269] provided perhaps the best known characterization of graphs with perfect matchings. A number of proofs of Tutte’s Theorem have been published since then. The proof that we present is due to Anderson [7].

A definition first: Given a graph G , define $\Omega(G)$ to be the number of connected components of G with odd order. Also, define $\Sigma(G)$ to be the number of connected components of G with even order.

Theorem 1.59 (Tutte’s Theorem). *Let G be a graph of order $n \geq 2$. G has a perfect matching if and only if $\Omega(G - S) \leq |S|$ for all subsets of S of $V(G)$.*

Proof. We begin with the forward direction. Let G be a graph that has a perfect matching. Suppose S is a set of vertices and that O_1, O_2, \dots, O_k are the odd components of $G - S$. For each i , the vertices in O_i can be adjacent only to other vertices in O_i and to vertices in S . Since G has a perfect matching, at least one vertex out of each of the O_i ’s has to be matched with a different vertex in S . If $k > |S|$, then some O_i would be left out (Figure 1.117). Thus, $k \leq |S|$.




FIGURE 1.117.

For the reverse direction of the theorem, suppose that $|S| \geq \Omega(G - S)$ for all S . In particular, if $S = \emptyset$, then $\Omega(G - \emptyset) \leq 0$. This implies that there are no odd components of G —every component of G is even. More generally, we make the following claim.

Claim A. For any proper subset S , $|S|$ and $\Omega(G - S)$ are either both even or both odd.

Let C be some component of G . We know from earlier that C has even order. If an even number of vertices is removed from C , then the number of odd components remaining must also be even. If an odd number of vertices is removed from C , then the number of odd components remaining must be odd. Since this is true for every component of G , it is true for all of G . Hence Claim A is proved.

We now proceed by induction on n , the order of the graph. If $n = 2$, then G is K_2 , which certainly has a perfect matching. Suppose now that the result is true for all graphs of even order up to n , and let G be a graph of even order n . We now have two cases.

Case 1. Suppose that for every proper subset S , $\Omega(G - S) < |S|$. (That is, the strict inequality holds.) Claim A implies that $|S|$ and $\Omega(G - S)$ have the same parity, so we can say in this case that for all subsets S , $\Omega(G - S) \leq |S| - 2$. Let $uv \in E(G)$, and consider the graph $G - u - v$ (a graph with two fewer vertices than G). We would like to apply the induction hypothesis to $G - u - v$, so we need the following claim.

Claim B. For all subsets S' of $V(G - u - v)$, $\Omega(G - u - v - S') \leq |S'|$.

If Claim B were not true, then $\Omega(G - u - v - S_1) > |S_1|$ for some subset S_1 . But since $|S_1| = |S_1 \cup \{u, v\}| - 2$, we get $\Omega(G - u - v - S_1) > |S_1 \cup \{u, v\}|$, and this contradicts the assumption in this case. Claim B is proved.

Since Claim B is true, we can apply the induction hypothesis to $G - u - v$. That is, we can conclude that $G - u - v$ has a perfect matching. This matching, together with the edge uv , forms a perfect matching of G . Case 1 is complete.

Case 2. Suppose there exists a subset S such that $\Omega(G - S) = |S|$. There may be a number of subsets S that satisfy this condition—suppose without loss of generality that S is a largest such set. Let O_1, O_2, \dots, O_k be the components of $G - S$ of odd order.

Claim C. $\Sigma(G - S) = 0$. That is, there are no even-ordered components of $G - S$.

Let E be an even ordered component of $G - S$, and let x be a vertex of E . The graph $G - S - x$ has exactly one more odd component than $G - S$. Thus, $|S \cup \{x\}| = \Omega(G - S - x)$. But this means that $S \cup \{x\}$ is a set larger than S that satisfies the assumption of this case. Since we chose S to be the largest, we have a contradiction. Therefore there are no even-ordered components of $G - S$. Claim C is proved.

Claim D. There exist vertices $s_1, s_2, \dots, s_k \in S$ and vertices v_1, v_2, \dots, v_k , where for each i $v_i \in O_i$, such that $\{v_1 s_1, v_2 s_2, \dots, v_k s_k\}$ is a matching.

For each $i \in \{1, \dots, k\}$, define the set S_i to be the set of vertices in S that are adjacent to some vertex in O_i . Note that if $S_i = \emptyset$ for some i , then O_i is

completely disconnected from anything else in G , implying that G itself has an odd component. Since this contradicts our assumption in this case, we can assume that each S_i is nonempty. Furthermore, our initial assumptions imply that the union of any r of the S_i 's has size at least r . Thus, the condition in Theorem 1.52 is satisfied, implying that there exists a system of distinct representatives for the family of sets S_1, S_2, \dots, S_k . If we let these representatives be s_1, s_2, \dots, s_k , and their adjacencies in the O_i 's be v_1, v_2, \dots, v_k , then Claim D is proved.

The situation in G is depicted in Figure 1.118, where $k = |S|$.




FIGURE 1.118.

At this point, each vertex in S has been matched to a vertex in an O_i . The goal at this point is to show that each $O_i - v_i$ has a perfect matching.

Let W be some subset of vertices of (the even-ordered) $O_i - v_i$.

Claim E. $\Omega(O_i - v_i - W) \leq |W|$.

If $\Omega(O_i - v_i - W) > |W|$, then, by Claim A, $\Omega(O_i - v_i - W) \geq |W| + 2$. But then,

$$\begin{aligned} \Omega(G - S - v_i - W) &\geq |S| - 1 + |W| + 2 \\ &= |S| + |W| + 1 \\ &= |S \cup W \cup \{v_i\}|. \end{aligned}$$

This contradicts our assumption, and thus Claim E is proved.

Since Claim E is true, each $O_i - v_i$ satisfies the induction hypothesis, and thus has a perfect matching. These perfect matchings together with the perfect matching shown in Figure 1.118 form a perfect matching of G , and so Case 2 is complete. □

We conclude this section by considering perfect matchings in regular graphs. If a graph G is 1-regular, then G itself is a perfect matching. If G is 2-regular, then G is a collection of disjoint cycles; as long as each cycle is even, G will have a perfect matching.

What about 3-regular graphs? A graph that is 3-regular must be of even order, so is it possible that every 3-regular graph contains a perfect matching? In a word, no. The graph in Figure 1.119 is a connected 3-regular graph that does not have a perfect matching. Thanks to Petersen [221], though, we do know of a special class of 3-regular graphs that do have perfect matchings. Recall that a bridge in a graph is an edge whose removal would disconnect the graph. The graph in Figure 1.119 has three bridges.




FIGURE 1.119.

Theorem 1.60 (Petersen's Theorem). *Every bridgeless, 3-regular graph contains a perfect matching.*

Proof. Let G be a bridgeless, 3-regular graph, and suppose that it does not contain a perfect matching. By Tutte's Theorem, there must exist a subset S of vertices where the number of odd components of $G - S$ is greater than $|S|$. Denote the odd-ordered components of $G - S$ by O_1, O_2, \dots, O_k .

First, each O_i must have at least one edge into S . Otherwise, there would exist an odd-ordered, 3-regular subgraph of G , and this is not possible, by Theorem 1.1. Second, since G is bridgeless, there must be at least two edges joining each O_i to S . Moreover, if there were only two edges joining some O_i to S , then O_i would contain an odd number of vertices with odd degree, and this cannot happen.

We can therefore conclude that there are at least three edges joining each O_i to S . This implies that there are at least $3k$ edges coming into S from the O_i 's. But since every vertex of S has degree 3, the greatest number of edges incident with vertices in S is $3|S|$, and since $3k > 3|S|$, we have a contradiction. Therefore, G must have a perfect matching. \square

It is probably not surprising that the Petersen of Theorem 1.60 is the same person for whom the Petersen graph (Figure 1.63) is named.

Petersen used this special graph as an example of a 3-regular, bridgeless graph whose edges cannot be partitioned into three separate, disjoint matchings.

Exercises

1. Find a maximum matching of the graph shown in Figure 1.119.
2. Use Tutte's Theorem to prove that the graph in Figure 1.119 does not have a perfect matching.
3. Draw a connected, 3-regular graph that has both a cut vertex and a perfect matching.

4. Determine how many different perfect matchings there are in $K_{n,n}$.
5. Prove Corollary 1.57.
6. Characterize when K_{r_1, r_2, \dots, r_k} has a perfect matching.
7. Prove that every tree has at most one perfect matching.
8. Let G be a subgraph of $K_{20,20}$. If G has a perfect matching, prove that G has at most 190 edges that belong to no perfect matching.
9. Use Tutte's Theorem to prove Hall's Theorem.

1.8 Ramsey Theory

*I have to go and see some friends of mine, some that I don't know,
and some who aren't familiar with my name.*

— John Denver, *Goodbye Again*

We begin this section with a simple question: How many people are required at a gathering so that there must exist either three mutual acquaintances or three mutual strangers? We will answer this question soon.

Ramsey theory is named for Frank Ramsey, a young man who was especially interested in logic and philosophy. Ramsey died at the age of 26 in 1930—the same year that his paper *On a problem of formal logic* was published. His paper catalyzed the development of the mathematical field now known as Ramsey theory. The study of Ramsey theory has burgeoned since that time. While many results in the subject are published each year, there are many questions whose answers remain elusive. As the authors of [136] put it, “the field is alive and exciting.”

1.8.1 Classical Ramsey Numbers


An innocent looking problem often gives no hint as to its true nature.

— Paul Erdős [92]

A *2-coloring* of the *edges* of a graph G is any assignment of one of two colors to each of the edges of G . Figure 1.120 shows a 2-coloring of the edges of K_5 using red (thick) and blue (thin).

Let p and q be positive integers. The (*classical*) *Ramsey number* associated with these integers, denoted by $R(p, q)$, is defined to be the smallest integer n such that every 2-coloring of the edges of K_n either contains a red K_p or a blue K_q as a subgraph.

Read through that definition at least one more time, and then consider this simple example. We would like to find the value of $R(1, 3)$. According to the definition, this is the least value of n such that every 2-coloring of the edges of K_n

FIGURE 1.120. A 2-coloring of K_5 .


either contains as a subgraph a K_1 all of whose edges are red, or a K_3 all of whose edges are blue. How many vertices are required before we know that we will have one of these objects in every coloring of a complete graph? If you have just one vertex, then no matter how you color the edges (ha-ha) of K_1 , you will always end up with a red K_1 . Thus, $R(1, 3) = 1$. We have found our first Ramsey number!

We should note here that the definition given for Ramsey number is in fact a good definition. That is, given positive integers p and q , $R(p, q)$ does in fact exist. Ramsey himself proved this fact, and we will learn more about the proof of “Ramsey’s Theorem” in Chapter 2.

Back to examples. We just showed that $R(1, 3) = 1$. Similar reasoning shows that $R(1, k) = 1$ for all positive integers k (see Exercise 2).

How about $R(2, 4)$? We need to know the smallest integer n such that every 2-coloring of the edges of K_n contains either a red K_2 or a blue K_4 . We claim that $R(2, 4) = 4$. To show this, we must demonstrate two things: first, that there exists a 2-coloring of K_3 that contains neither a red K_2 nor a blue K_4 , and second, that any 2-coloring of the edges of K_4 contains at least one of these as a subgraph.

We demonstrate the first point. Consider the 2-coloring of K_3 given in Figure 1.121 (recall that red is thick and blue is thin—the edges in this coloring are all blue). This coloring of K_3 does not contain a red K_2 , and it certainly does not

FIGURE 1.121. The edges of K_3 colored blue.

contain a blue K_4 . Thus $R(2, 4) > 3$.

For the second point, suppose that the edges of K_4 are 2-colored in some fashion. If any of the edges are red, then we have a red K_2 . If none of the edges are red, then we have a blue K_4 . So, no matter the coloring, we always get one of the two. This proves that $R(2, 4) = 4$.

What do you think is the value of $R(2, 5)$? How about $R(2, 34)$? As you will prove in Exercise 3, $R(2, k) = k$ for all integers $k \geq 2$.

Exercises

1. How many different 2-colorings are there of K_3 ? of K_4 ? of K_5 ? of K_{10} ?
2. Write a nice proof of the fact that $R(1, k) = 1$ for all positive integers k .
3. Write a nice proof of the fact that $R(2, k) = k$ for integers $k \geq 2$.
4. Prove that for positive integers p and q , $R(p, q) = R(q, p)$.
5. If $2 \leq p' \leq p$ and $2 \leq q' \leq q$, then prove that $R(p', q') \leq R(p, q)$. Also, prove that equality holds if and only if $p' = p$ and $q' = q$.

1.8.2 Exact Ramsey Numbers and Bounds

Take me to your leader.

— proverbial alien

How many people are required at a gathering so that there must exist either three mutual acquaintances or three mutual strangers? We can rephrase this question as a problem in Ramsey theory: How many vertices do you need in an (edge) 2-colored complete graph for it to be necessary that there be either a red K_3 (people who know each other) or a blue K_3 (people who do not know each other)? As the next theorem states, the answer is 6.

Theorem 1.61. $R(3, 3) = 6$.

Proof. We begin the proof by exhibiting (in Figure 1.122) a 2-coloring of the edges of K_5 that produces neither a red (thick) K_3 nor a blue (thin) K_3 . This




FIGURE 1.122. A 2-coloring of the edges of K_5 .

2-coloring of K_5 demonstrates that $R(3, 3) > 5$. Now consider K_6 , and suppose that each of its edges has been colored red or blue. Let v be one of the vertices of K_6 . There are five edges incident with v , and they are each colored red or blue, so it must be that v is either incident with at least three red edges or at least three blue edges (think about this; it is called the Pigeonhole Principle—more on this in later chapters). Without loss of generality, let us assume that v is incident with at least three red edges, and let us call them vx , vy , and vz (see Figure 1.123).




FIGURE 1.123.




FIGURE 1.124.

Now, if none of the edges xy, xz, yz is colored red, then we have a blue K_3 (Figure 1.124).

On the other hand, if at least one of xy, xz, yz is colored red, we have a red K_3 (Figure 1.125).




FIGURE 1.125.

Therefore, any 2-coloring of the edges of K_6 produces either a red K_3 or a blue K_3 . □

Let us determine another Ramsey number.

Theorem 1.62. $R(3, 4) = 9$.

Proof. Consider the 2-coloring of the edges of K_8 given in Figure 1.126.

A bit of examination reveals that this coloring produces no red (thick) K_3 and no blue (thin) K_4 . Thus, $R(3, 4) \geq 9$. We now want to prove that $R(3, 4) \leq 9$, and we will use the facts that $R(2, 4) = 4$ and $R(3, 3) = 6$.

Let G be any complete graph of order at least 9, and suppose that the edges of G have been 2-colored arbitrarily. Let v be some vertex of G .

Case 1. Suppose that v is incident with at least four red edges. Call the end vertices of these edges “red neighbors” of v , and let S be the set of red neighbors of v (see Figure 1.127).




FIGURE 1.126. A 2-coloring of the edges in K_8 .




FIGURE 1.127.

Since S contains at least four vertices, and since $R(2, 4) = 4$, the 2-coloring of the edges that are within S must produce either a red K_2 or a blue K_4 within S itself. If the former is the case, then we are guaranteed a red K_3 in G (see Figure 1.128). If the latter is the case, then we are clearly guaranteed a blue K_4 in G .




FIGURE 1.128.

Case 2. Suppose that v is incident with at least six blue edges. Call the other end vertices of these edges “blue neighbors” of v , and let T be the set of blue neighbors of v (see Figure 1.129).

Since T contains at least six vertices, and since $R(3, 3) = 6$, the 2-coloring of the edges that are within T must produce either a red K_3 or a blue K_3 within T itself. If the former is the case, then we are obviously guaranteed a red K_3 in G .




FIGURE 1.129.

If the latter is the case, then we are guaranteed a blue K_4 in G (see Figure 1.130).




FIGURE 1.130.

Case 3. Suppose that v is incident with fewer than four red edges and fewer than six blue edges. In this case there must be at most nine vertices in G altogether, and since we assumed at the beginning that the order of G is at least 9, we can say that G has order exactly 9. Further, we can say that v is incident with exactly three red edges and exactly five blue edges. And since the vertex v was chosen arbitrarily, we can assume that this holds true for every vertex of G .

Now if we consider the underlying “red” subgraph of G , we have a graph with nine vertices, each of which has degree 3. But this cannot be, since the number of vertices in G with odd degree is even (the First Theorem of Graph Theory). Therefore, this case cannot occur.

We have therefore proved that any 2-coloring of the edges of a complete graph on 9 vertices (or more) produces either a red K_3 or a blue K_4 . Hence, $R(3, 4) = 9$. \square

Some known Ramsey numbers are listed below.

$$\begin{aligned} R(1, k) &= 1, \\ R(2, k) &= k, \\ R(3, 3) &= 6, R(3, 4) = 9, R(3, 5) = 14, R(3, 6) = 18, \\ R(3, 7) &= 23, R(3, 8) = 28, R(3, 9) = 36, \\ R(4, 4) &= 18, R(4, 5) = 25. \end{aligned}$$

Bounds on Ramsey Numbers

Determining exact values of Ramsey numbers is extremely difficult in general. *In fact, the list given above is not only a list of some known Ramsey numbers, it is a list of all known Ramsey numbers.* Many people have attempted to determine other values, but to this day no other numbers are known.

However, there has been progress in finding bounds, and we state some important ones here. The proofs of the first two theorems will be discussed in Section 2.10.2 (see Theorem 2.28 and Corollary 2.29). The first bound is due to Erdős and Szekeres [94], two major players in the development of Ramsey theory. Their result involves a quotient of factorials: Here, $n!$ denotes the product $1 \cdot 2 \cdots n$.

Theorem 1.63. *For positive integers p and q ,*

$$R(p, q) \leq \frac{(p+q-2)!}{(p-1)!(q-1)!}.$$

The next theorem gives a bound on $R(p, q)$ based on “previous” Ramsey numbers.

Theorem 1.64. *If $p \geq 2$ and $q \geq 2$, then*

$$R(p, q) \leq R(p-1, q) + R(p, q-1).$$

Furthermore, if both terms on the right of this inequality are even, then the inequality is strict.

The following bound is for the special case $p = 3$.

Theorem 1.65. *For every integer $q \geq 3$,*

$$R(3, q) \leq \frac{q^2 + 3}{2}.$$

The final bound that we present is due to Erdős [90]. It applies to the special case $p = q$. In the theorem, $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x .

Theorem 1.66. *If $p \geq 3$, then*

$$R(p, p) > \lfloor 2^{n/2} \rfloor.$$

A number of other specific bounds are known:

$$\begin{aligned} 35 &\leq R(4, 6) \leq 41, \\ 43 &\leq R(5, 5) \leq 49, \\ 58 &\leq R(5, 6) \leq 87, \\ 102 &\leq R(6, 6) \leq 165. \end{aligned}$$

Even with the sophisticated computing power that is available to us today, we are not able to compute values for more than a handful of Ramsey numbers. Paul Erdős once made the following comment regarding the difficulty in finding exact values of Ramsey numbers [63]:

Suppose an evil alien would tell mankind “Either you tell me [the value of $R(5, 5)$] or I will exterminate the human race.” ... It would be best in this case to try to compute it, both by mathematics and with a computer.

If he would ask [for the value of $R(6, 6)$], the best thing would be to destroy him before he destroys us, because we couldn't [determine $R(6, 6)$].

Exercises

1. Prove that $R(3, 5) \geq 14$. The graph in Figure 1.131 will be very helpful.




FIGURE 1.131. A 2-coloring of K_{13} .

2. Use Theorem 1.64 and the previous exercise to prove that $R(3, 5) = 14$.

3. Construct a graph and a 2-coloring that proves $R(4, 4) \geq 18$.
4. Use Theorem 1.64 and the previous exercise to prove that $R(4, 4) = 18$.
5. Use Theorem 1.64 to prove Theorem 1.65.

1.8.3 Graph Ramsey Theory

All generalizations are dangerous, even this one.

— Alexandre Dumas

Graph Ramsey theory is a generalization of classical Ramsey theory. Its development was due in part to the search for the elusive classical Ramsey numbers, for it was thought that the more general topic might shed some light on the search. The generalization blossomed and became an exciting field in itself. In this section we explain the concept of graph Ramsey theory, and we examine several results. These results, and more like them, can be found in [136].

Given two graphs G and H , define the graph Ramsey number $R(G, H)$ to be the smallest value of n such that any 2-coloring of the edges of K_n contains either a red copy of G or a blue copy of H . The classical Ramsey number $R(p, q)$ would in this context be written as $R(K_p, K_q)$.

The following simple result demonstrates the relationship between graph Ramsey numbers and classical Ramsey numbers.

Theorem 1.67. *If G is a graph of order p and H is a graph of order q , then*

$$R(G, H) \leq R(p, q).$$

Proof. Let $n = R(p, q)$, and consider an arbitrary 2-coloring of K_n . By definition, K_n contains either a red K_p or a blue K_q . Since $G \subseteq K_p$ and $H \subseteq K_q$, there must either be a red G or a blue H in K_n . Hence, $R(G, H) \leq n = R(p, q)$. \square

Here is a result due to Chvátal and Harary [55] that relates $R(G, H)$ to the chromatic number of G , $\chi(G)$, and the order of the largest component of H , denoted by $C(H)$.

Theorem 1.68. $R(G, H) \geq (\chi(G) - 1)(C(H) - 1) + 1$.

Proof. Let $m = \chi(G) - 1$ and let $n = C(H) - 1$. Consider the graph S formed by taking m copies of K_n and adding all of the edges in between each copy (Figure 1.132). Actually, $S = K_{mn}$. Now color all of the edges within each K_n blue, and color all other edges red. From the way we have constructed the coloring, every red subgraph can be vertex colored with m colors. Since $m < \chi(G)$, there can be no red G present. Furthermore, any blue subgraph has at most $n = C(H) - 1$ vertices in its largest component. Hence, there can be no blue H present. We have exhibited a 2-coloring of K_{mn} that contains neither a red G nor a blue H . \square




FIGURE 1.132. The graph S .

The next few theorems give exact graph Ramsey numbers for specific classes of graphs. The first is due to Chvátal [53], and the proof uses a few ideas from previous sections.

Theorem 1.69. *If T_m is a tree with m vertices, then*

$$R(T_m, K_n) = (m - 1)(n - 1) + 1.$$

Proof. If $m = 1$ or $n = 1$, then $R(T_m, K_n) = 1$ and the result holds. Assume then that m and n are both at least 2.

Claim A. $R(T_m, K_n) \geq (m - 1)(n - 1) + 1$.

Consider the graph that consists of $n - 1$ copies of K_{m-1} , with all possible edges between the copies of K_{m-1} . This graph is actually $K_{(m-1)(n-1)}$. Color the edges in each K_{m-1} red, and color all of the other edges blue. Since each of the red subgraphs has order $m - 1$, no red T_m can exist. Also, by this construction, no blue K_n can exist. Since this 2-coloring contains no red T_m and no blue K_n , it must be that $R(T_m, K_n) \geq (m - 1)(n - 1) + 1$.

Claim B. $R(T_m, K_n) \leq (m - 1)(n - 1) + 1$.

Let G be $K_{(m-1)(n-1)+1}$, and suppose that its edges have been arbitrarily 2-colored. Let G_r denote the subgraph of G formed by the red edges, and let G_b denote the subgraph of G formed by the blue edges. If there is no blue K_n , then $\omega(G_b) \leq n - 1$, and if so, then $\alpha(G_r) \leq n - 1$, since G_r is the complement of G_b . Thus by Theorem 1.45, $\chi(G_r) \geq m$. Let H be a subgraph of G_r that is m -critical. By part (d) of Exercise 6 in Section 1.6.1, $\delta(H) \geq m - 1$. By Theorem 1.16, H contains T_m as a subgraph, and therefore G has a red T_m . □

The next theorem is due to Burr [46].

Theorem 1.70. *If T_m is a tree of order m and if $m - 1$ divides $n - 1$, then*

$$R(T_m, K_{1,n}) = m + n - 1.$$

In the following theorem, mK_2 denotes the graph consisting of m copies of K_2 , and nK_2 has a similar meaning.

Theorem 1.71. *If $m \geq n \geq 1$, then*

$$R(mK_2, nK_2) = 2m + n - 1.$$

As we mentioned earlier, these results apply to specific classes of graphs. In general, determining values of $R(G, H)$ is quite difficult. So the generalization that was intended to solve hard classical Ramsey problems has produced hard problems of its own!

Exercises

1. Find $R(P_3, P_3)$.
2. Find $R(P_3, C_4)$.
3. Find $R(C_4, C_4)$.
4. Prove that $R(K_{1,3}, K_{1,3}) = 6$.
5. Prove that $R(2K_3, K_3) = 8$.

1.9 References

Prince John: Are you finished?

Sir Robin of Locksley: I'm only just beginning.

— *Robin Hood*

We have only just begun our walk through the field of graph theory. In this section we will provide references for those who are interested in further study.

The books by Chartrand and Lesniak [52], Buckley and Lewinter [43] and West [281] provide very thorough introductions to a large number of topics in graph theory. The graduate-level texts by Diestel [75] and Bollobás [29], along with offering further study of the concepts covered in this chapter, also cover network flows, extremal graph theory, and random graphs. Gould's book [128] covers a number of graph algorithms, from finding maximum matchings to testing planarity. Many interesting applications of graph theory can be found in texts by Gross and Yellen [139], Foulds [106], and Roberts and Tesman [238]. A good source for connections between graph theory and other mathematical topics is [21], edited by Beineke and Wilson. The text [148] by Harary is a thorough discussion of counting various types of graphs. A wonderful source for the history of graph theory and some of its famous problems is the book [26] by Biggs, Lloyd, and Wilson.

Buckley and Harary [42] have a nice text which focuses on distance in graphs. For more on the development of the Four Color Problem, see the books by Wilson

[285] and Aigner [3]. Much more information regarding Ramsey theory can be found in the book [136] by Graham, Rothschild, and Spencer. Also, Radziszowski [231] maintains a list of current bounds on small Ramsey numbers.

The book by Barabási [17] is a nice general treatment of graphs (networks) and their relationship to all sorts of phenomena. Finally, the books by Hoffman [163] and Schechter [247] tell the story of Paul Erdős, a twentieth century giant in the field of graph theory.

2

Combinatorics

[Combinatorics] has emerged as a new subject standing at the crossroads between pure and applied mathematics, the center of bustling activity, a simmering pot of new problems and exciting speculations.

— Gian-Carlo Rota, [243, p. vii]

The formal study of combinatorics dates at least to Gottfried Wilhelm Leibniz's *Dissertatio de Arte Combinatoria* in the seventeenth century. The last half-century, however, has seen a huge growth in the subject, fueled by problems and applications from many fields of study. Applications of combinatorics arise, for example, in chemistry, in studying arrangements of atoms in molecules and crystals; biology, in questions about the structure of genes and proteins; physics, in problems in statistical mechanics; communications, in the design of codes for encryption, compression, and correction of errors; and especially computer science, for instance in problems of scheduling and allocating resources, and in analyzing the efficiency of algorithms.

Combinatorics is, in essence, the study of arrangements: pairings and groupings, rankings and orderings, selections and allocations. There are three principal branches in the subject. *Enumerative combinatorics* is the science of counting. Problems in this subject deal with determining the number of possible arrangements of a set of objects under some particular constraints. *Existential combinatorics* studies problems concerning the existence of arrangements that possess some specified property. *Constructive combinatorics* is the design and study of algorithms for creating arrangements with special properties.

Combinatorics is closely related to the theory of graphs. Many problems in graph theory concern arrangements of objects and so may be considered as combinatorial problems. For example, the theory of matchings and Ramsey theory, both studied in the previous chapter, have the flavor of existential combinatorics, and we continue their study later in this chapter. Also, combinatorial techniques are often employed to address problems in graph theory. For example, in Section 2.5 we determine another method for finding the chromatic polynomial of a graph.

We focus on topics in enumerative combinatorics through most of this chapter, but turn to some questions in existential combinatorics in Sections 2.4 and 2.10, and to some problems in constructive combinatorics in Sections 2.9 and 2.10. Throughout this chapter we study arrangements of finite sets. Chapter 3 deals with arrangements and combinatorial problems involving infinite sets. Our study in this chapter includes the investigation of the following questions.

- Should a straight beat a flush in the game of poker? What about a full house?
- Suppose a lazy professor collects a quiz from each student in a class, then shuffles the papers and redistributes them randomly to the class for grading. How likely is it that no one receives his or her own quiz to grade?
- How many ways are there to make change for a dollar?
- How many different necklaces with twenty beads can be made using rhodinite, rose quartz, and lapis lazuli beads, if a necklace can be worn in any orientation?
- How many seating arrangements are possible for n guests attending a wedding reception in a banquet room with k round tables?
- Suppose 100 medical students rank 100 positions for residencies at hospitals in order of preference, and the hospitals rank the students in order of preference. Is there a way to assign the students to the hospitals in such a way that no student and hospital prefer each other to their assignment? Is there an efficient algorithm for finding such a matching?
- Is it possible to find a collection of $n \geq 3$ points in the plane, not all on the same line, so that every line that passes through two of the points in fact passes through a third? Or, if we require instead that no three points lie on the same line, can we arrange a large number of points so that no subset of them forms the vertices of a convex octagon?

2.1 Some Essential Problems

The mere formulation of a problem is far more essential than its solution...

— Albert Einstein

We begin our study of combinatorics with two essential observations that underlie many counting strategies and techniques. The first is a simple observation about counting when presented with a number of alternative choices.

Sum Rule. *Suppose S_1, S_2, \dots, S_m are mutually disjoint finite sets, and $|S_i| = n_i$ for $1 \leq i \leq m$. Then the number of ways to select one object from any of the sets S_1, S_2, \dots, S_m is the sum $n_1 + n_2 + \dots + n_m$.*

We often use the sum rule implicitly when solving a combinatorial problem when we break the set of possible outcomes into several disjoint cases, each of which can be analyzed separately. For example, suppose a coy college athlete tells us that his two-digit jersey number is divisible by 3, its first digit is odd, and its second digit is less than its first. How many numbers satisfy these criteria? A natural approach is to break the problem into five cases based on the first digit. Analyzing each of 1, 3, 5, 7, and 9 in turn, we find the possibilities are $\{\}$, $\{30\}$, $\{51, 54\}$, $\{72, 75\}$, or $\{90, 93, 96\}$, so there are eight possible jersey numbers in all.

The second essential observation concerns counting problems where selections are made in sequence.

Product Rule. *Suppose S_1, S_2, \dots, S_m are finite sets, and $|S_i| = n_i$ for $1 \leq i \leq m$. Then the number of ways to select one element from S_1 , followed by one element from S_2 , and so on, ending with one element from S_m , is the product $n_1 n_2 \cdots n_m$, provided that the selections are independent, that is, the elements chosen from S_1, \dots, S_{i-1} have no bearing on the selection from S_i , for each i .*

For example, consider the number of m -letter acronyms that can be formed using the full alphabet. To construct such an acronym, we make m choices in sequence, one for each position, and each choice has no effect on any subsequent selection. Thus, by the product rule, the number of such acronyms is 26^m .

We can apply a similar strategy to count the number of valid phone numbers in the U.S. and Canada. Under the North American Numbering Plan, a phone number has ten digits, consisting of an area code, then an exchange, then a station code. The three-digit area code cannot begin with 0 or 1, and its second digit can be any number except 9. The three-digit exchange cannot begin with 0 or 1, and the station code can be any four-digit number. Using the product rule, we find that the number of valid phone numbers under this plan is $(8 \cdot 9 \cdot 10) \cdot (8 \cdot 10^2) \cdot 10^4 = 5\,760\,000\,000$.

One might object that certain three-digit numbers are service codes reserved for special use in many areas, like 411 for information and 911 for emergencies. Let's compute the number of valid phone numbers for which neither the area code nor the exchange end with the digits 11. The amended number of area codes is then $8(9 \cdot 10 - 1) = 712$, and for exchanges we obtain $8 \cdot 99 = 792$. Thus, the number of valid phone numbers is $712 \cdot 792 \cdot 10^4 = 5\,639\,040\,000$.

We can use the product rule to solve three basic combinatorial problems.

Problem 1. How many ways are there to order a collection of n different objects?

For example, how many ways are there to arrange the cards in a standard deck of 52 playing cards by shuffling? How many different batting orders are possible among the nine players on a baseball team? How many ways are there to arrange ten books on a shelf?

To order a collection of n objects, we need to pick one object to be first, then another one to be second, and another one third, and so on. There are n different choices for the first object, then $n - 1$ remaining choices for the second, and $n - 2$ for the third, and so forth, until just one choice remains for the last object. The total number of ways to order the n objects is therefore the product of the integers between 1 and n . This number, called n factorial, is written $n!$. An ordering, or rearrangement, of n objects is often called a *permutation* of the objects. Thus, the number of permutations of n items is $n!$.

Our second problem generalizes the first one.

Problem 2. How many ways are there to make an ordered list of k objects from a collection of n different objects?

For example, how many ways can a poll rank the top 20 teams in a college sport if there are 100 teams in the division? How many ways can a band arrange a play list of twelve songs if they know only 25 different songs?

Applying the same reasoning used in the first problem, we find that the answer to Problem 2 is the product $n(n - 1)(n - 2) \cdots (n - k + 1)$, or $n!/(n - k)!$. This number is sometimes denoted by $P(n, k)$, but products like this occur frequently in combinatorics, and a more descriptive notation is often used to designate them.

We define the *falling factorial power* $x^{\underline{k}}$ as a product of k terms beginning with x , with each successive term one less than its predecessor:

$$x^{\underline{k}} = x(x - 1)(x - 2) \cdots (x - k + 1) = \prod_{i=0}^{k-1} (x - i). \quad (2.1)$$

The expression $x^{\underline{k}}$ is pronounced “ x to the k falling.” Similarly, we define the *rising factorial power* $x^{\overline{k}}$ (“ x to the k rising”) by

$$x^{\overline{k}} = x(x + 1)(x + 2) \cdots (x + k - 1) = \prod_{i=0}^{k-1} (x + i). \quad (2.2)$$

Thus, we see that $P(n, k) = n^{\underline{k}} = (n - k + 1)^{\overline{k}}$, and $n! = n^{\underline{n}} = 1^{\overline{n}}$. Also, the expressions $n^{\underline{0}}$, $n^{\overline{0}}$, and $0!$ all represent products having no terms at all. Multiplying any expression by such an empty product should not disturb the value of the expression, so the value of each of these degenerate products is taken to be 1.

Our third problem concerns unordered selections.

Problem 3. How many ways are there to select k objects from a collection of n objects, if the order of selection is irrelevant?

For example, how many different hands are possible in the game of poker? A poker hand consists of five cards drawn from a standard deck of 52 different cards. The order of the cards in a hand is unimportant, since players can rearrange their cards freely.

The solution to Problem 3 is usually denoted by $\binom{n}{k}$, or sometimes $C(n, k)$. The expression $\binom{n}{k}$ is pronounced “ n choose k .”

We can find a formula for $\binom{n}{k}$ by using our solutions to Problems 1 and 2. Since there are $k!$ different ways to order a collection of k objects, it follows that the product $\binom{n}{k}k!$ is the number of possible ordered lists of k objects selected from the same collection of n objects. Therefore,

$$\binom{n}{k} = \frac{n^k}{k!} = \frac{n!}{k!(n-k)!}. \quad (2.3)$$

The numbers $\binom{n}{k}$ are called *binomial coefficients*, for reasons discussed in the next section. The binomial coefficients are ubiquitous in combinatorics, and we close this section with a few applications of these numbers.

1. The number of different hands in poker is $\binom{52}{5} = 52^5/5! = 2\,598\,960$. The number of different thirteen-card hands in the game of bridge is $\binom{52}{13} = 635\,013\,559\,600$.
2. To play the Texas lottery game Lotto Texas, a gambler selects six different numbers between 1 and 54. The order of selection is unimportant. The number of possible lottery tickets is therefore $\binom{54}{6} = 25\,827\,165$.
3. Suppose we need to travel m blocks east and n blocks south in a regular grid of city streets. How many paths are there to our destination if we travel only east and south?

We can represent a path to our destination as a sequence b_1, b_2, \dots, b_{n+m} , where b_i represents the direction we are traveling during the i th block of our route. Exactly m of the terms in this sequence must be “east,” and there are precisely $\binom{m+n}{m}$ ways to select m positions in the sequence to have this value. The remaining n positions in the sequence must all be “south,” so the number of possible paths is $\binom{m+n}{m} = \frac{(m+n)!}{m!n!}$.

4. A standard deck of playing cards consists of four suits (spades, hearts, clubs, and diamonds), each with thirteen cards. Each of the cards in a suit has a different face value: a number between 2 and 10, or a jack, queen, king, or ace. How many poker hands have exactly three cards with the same face value?

We can answer this question by considering how to construct such a hand through a sequence of simple steps. First, select one of the thirteen different

face values. Second, choose three of the four cards in the deck having this value. Third, pick two cards from the 48 cards having a different face value. By the product rule, the number of possibilities is

$$\binom{13}{1} \binom{4}{3} \binom{48}{2} = 58\,656.$$

Poker aficionados will recognize that this strategy counts the number of ways to deal either of two different hands in the game: the “three of a kind” and the stronger “full house.” A full house consists of a matched triple together with a matched pair, for example, three jacks and two aces; a three of a kind has only a matched triple. The number of ways to deal a full house is

$$\binom{13}{1} \binom{4}{3} \binom{12}{1} \binom{4}{2} = 3744,$$

since choosing a matched pair involves first selecting one of twelve different remaining face values, then picking two of the four cards having this value. The number of three of a kind hands is therefore $58\,656 - 3744 = 54\,912$.

We can also compute this number directly by modifying our first strategy. To avoid the possibility of selecting a matched pair in the last step, we can replace the term $\binom{48}{2} = 48 \cdot 47/2$ by $48 \cdot 44/2$, since the face value of the last card should not match any other card selected. Indeed, we calculate $13 \cdot 4 \cdot 48 \cdot 44/2 = 54\,912$. Notice that dividing by 2 is required in the last step, since the last two cards may be selected in any order.

Exercises

1. In the C++ programming language, a variable name must start with a letter or the underscore character (`_`), and succeeding characters must be letters, digits, or the underscore character. Uppercase and lowercase letters are considered to be different characters.
 - (a) How many variable names with exactly five characters can be formed in C++?
 - (b) How many are there with at most five characters?
 - (c) How many are there with at most five characters, if they must read exactly the same forwards and backwards? For example, `kayak` and `T55T` are admissible, but `Kayak` is not.
2. Assume that a vowel is one of the five letters A, E, I, O, or U.
 - (a) How many eleven-letter sequences from the alphabet contain exactly three vowels?
 - (b) How many of these have at least one repeated letter?

3. There are 30 teams in the National Basketball Association: 15 in the Western Conference, and 15 in the Eastern Conference.
 - (a) Suppose each of the teams in the league has one pick in the first round of the NBA draft. How many ways are there to arrange the order of the teams selecting in the draft?
 - (b) Suppose that each of the first three positions in the draft must be awarded to one of the fourteen teams that did not advance to the playoffs that year. How many ways are there to assign the first three positions in the draft?
 - (c) How many ways are there for eight teams from each conference to advance to the playoffs, if order is unimportant?
 - (d) Suppose that every team has three centers, four guards, and five forwards. How many ways are there to select an all-star team with the same composition from the Western Conference?
4. According to the *Laws of the Game* of the International Football Association, a full football (soccer) team consists of eleven players, one of whom is the goalkeeper. The other ten players fall into one of three outfield positions: defender, midfielder, and striker. There is no restriction on the number of players at each of these positions, as long as the total number of outfield players is ten.
 - (a) How many different configurations are there for a full football team? For example, one team may field four strikers, three midfielders, and three defenders, in addition to the goalkeeper. Another may play five strikers, no midfielders, and five defenders, plus the goalkeeper.
 - (b) Repeat the previous problem if there must be at least two players at each outfield position.
 - (c) How many ways can a coach assign eleven different players to one of the four positions, if there must be exactly one goalkeeper, but there is no restriction on the number of players at each outfield position?
5. A political science quiz has two parts. In the first, you must present your opinion of the four most influential secretaries-general in the history of the United Nations in a ranked list. In the second, you must name ten members of the United Nations security council in any order, including at least two permanent members of the council. If there have been eight secretaries-general in U.N. history, and there are fifteen members of the U.N. security council, including the five permanent members, how many ways can you answer the quiz, assuming you answer both parts completely?
6. A midterm exam in phenomenology has two parts. The first part consists of ten multiple choice questions. Each question has four choices, labeled (a), (b), (c), and (d), and one may pick any combination of responses on each

of these questions. For example, one could choose just (a) alone on one question, or both (b) and (c), or all four possibilities, or none of them. In the second part, one may choose either to answer eight true/false questions, or to select the proper definition of each of seven terms from a list of ten possible definitions. Every question must be answered on whichever part is chosen, but one is not allowed to complete both portions. How many ways are there to complete the exam?

7. A ballot lists ten candidates for city council, eight candidates for the school board, and five bond issues. The ballot instructs voters to choose up to four people running for city council, rank up to three candidates for the school board, and approve or reject each bond issue. How many different ballots may be cast, if partially completed (or empty) ballots are allowed?
8. Compute the number of ways to deal each of the following five-card hands in poker.
 - (a) Straight: the values of the cards form a sequence of consecutive integers. A jack has value 11, a queen 12, and a king 13. An ace may have a value of 1 or 14, so A 2 3 4 5 and 10 J Q K A are both straights, but K A 2 3 4 is not. Furthermore, the cards in a straight cannot all be of the same suit (a flush).
 - (b) Flush: All five cards have the same suit (but not in addition a straight).
 - (c) Straight flush: both a straight and a flush. Make sure that your counts for straights and flushes do not include the straight flushes.
 - (d) Four of a kind.
 - (e) Two distinct matching pairs (but not a full house).
 - (f) Exactly one matching pair (but no three of a kind).
 - (g) At least one card from each suit.
 - (h) At least one card from each suit, but no two values matching.
 - (i) Three cards of one suit, and the other two of another suit, like three hearts and two spades.
9. In the lottery game Texas Two Step, a player selects four different numbers between 1 and 35 in step 1, then selects an additional “bonus ball” number in the same range in step 2. The latter number is not considered to be part of the set selected in step 1, and in fact it may match one of the numbers selected there.
 - (a) A resident of College Station always selects a bonus ball number that is different from any of the numbers he picks in step 1. How many of the possible Texas Two Step tickets have this property?

- (b) In Rhode Island's lottery game Wild Money, a gambler picks a set of five numbers between 1 and 35. Is the number of possible tickets in this game the same as the number of tickets in Texas Two Step where the bonus ball number is different from the other numbers? Determine the ratio of the number of possible tickets in Wild Money to the number in the restricted Texas Two Step.
10. (a) A superstitious resident of Amarillo always picks three even numbers and three odd numbers when playing Lotto Texas. What fraction of all possible lottery tickets have this property?
- (b) Suppose in a more general lottery game one selects six numbers between 1 and $2n$. What fraction of all lottery tickets have the property that half the numbers are odd and half are even?
- (c) What is the limiting value of this probability as n grows large?
11. Suppose a positive integer N factors as $N = p_1^{n_1} p_2^{n_2} \cdots p_m^{n_m}$, where p_1, p_2, \dots, p_m are distinct prime numbers and n_1, n_2, \dots, n_m are all positive integers. How many different positive integers are divisors of N ?
12. Assume that a positive integer cannot have 0 as its leading digit.
- (a) How many five-digit positive integers have no repeated digits at all?
- (b) How many have no consecutive repeated digits?
- (c) How many have at least one run of consecutive repeated digits (for example, 23324, 45551, or 151155, but not 12121)?
13. How many positive integers are there whose representation in base 8 has exactly eight octal digits, at most one of which is odd? An octal digit is a number between 0 and 7, inclusive. Assume that the octal representation of a positive integer cannot start with a zero.
14. Let Δ be the *difference operator*: $\Delta(f(x)) = f(x+1) - f(x)$. Show that

$$\Delta(x^n) = nx^{n-1},$$

and use this to prove that

$$\sum_{k=0}^{m-1} k^n = \frac{m^{n+1}}{n+1}.$$

2.2 Binomial Coefficients

*About binomial theorem I'm teeming with a lot o' news,
With many cheerful facts about the square of the hypotenuse.*

— Gilbert and Sullivan, *The Pirates of Penzance*

The binomial coefficients possess a number of interesting arithmetic properties. In this section we study some of the most important identities associated with these numbers. Because binomial coefficients occur so frequently in this subject, knowing these essential identities will be helpful in our later studies.

The first identity generalizes our formula (2.3).

Expansion. *If n is a nonnegative integer and k is an integer, then*

$$\binom{n}{k} = \begin{cases} \frac{n!}{k!(n-k)!} & \text{if } 0 \leq k \leq n, \\ 0 & \text{otherwise.} \end{cases} \quad (2.4)$$

Designating the value of $\binom{n}{k}$ to be 0 when $k < 0$ or $k > n$ is sensible, for there are no ways to select fewer than zero or more than n items from a collection of n objects.

Notice that every subset of k objects selected from a set of n objects leaves a complementary collection of $n - k$ objects that are not selected. Counting the number of subsets with k objects is therefore the same as counting the number of subsets with $n - k$ objects. This observation leads us to our second identity, which is easy to verify using the expansion formula.

Symmetry. *If n is a nonnegative integer and k is an integer, then*

$$\binom{n}{k} = \binom{n}{n-k}. \quad (2.5)$$

Before presenting the next identity, let us consider again the problem of counting poker hands. Suppose the ace of spades is the most desirable card in the deck (it certainly is in American Western movies), and we would like to know the number of five-card hands that include this card. The answer is the number of ways to select four cards from the other 51 cards in the deck, namely, $\binom{51}{4}$. We can also count the number of hands that do not include the ace of spades. This is the number of ways to pick five cards from the other 51, that is, $\binom{51}{5}$. But every poker hand either includes the ace of spades or does not, so

$$\binom{52}{5} = \binom{51}{5} + \binom{51}{4}.$$

More generally, suppose we distinguish one particular object in a collection of n objects. The number of unordered collections of k of the objects that include the distinguished object is $\binom{n-1}{k-1}$; the number of collections that do not include this special object is $\binom{n-1}{k}$. We therefore obtain the following identity.

Addition. *If n is a positive integer and k is any integer, then*

$$\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}. \quad (2.6)$$

We can prove this identity more formally using the expansion identity. It is easy to check that the identity holds for $k \leq 0$ or $k \geq n$. If $0 < k < n$, we have

$$\begin{aligned} \binom{n-1}{k} + \binom{n-1}{k-1} &= \frac{(n-1)!}{k!(n-1-k)!} + \frac{(n-1)!}{(k-1)!(n-k)!} \\ &= \frac{((n-k) + k)(n-1)!}{k!(n-k)!} \\ &= \frac{n!}{k!(n-k)!} \\ &= \binom{n}{k}. \end{aligned}$$

We can use this identity to create a table of binomial coefficients. Let $n \geq 0$ index the rows of the table, and let $k \geq 0$ index the columns. Begin by entering 1 in the first position of each row, since $\binom{n}{0} = 1$ for $n \geq 0$; then use (2.6) to compute the entries in successive rows of the table. The resulting pattern of numbers is called *Pascal's triangle*, after Blaise Pascal, who studied many of its properties in his *Traité du Triangle Arithmétique*, written in 1654. (See [85] for more information on its history.) The first few rows of Pascal's triangle are shown in Figure 2.1.

$\binom{n}{k}$	$k = 0$	1	2	3	4	5	6	7	8	9	10	2^n
$n = 0$	1											1
1	1	1										2
2	1	2	1									4
3	1	3	3	1								8
4	1	4	6	4	1							16
5	1	5	10	10	5	1						32
6	1	6	15	20	15	6	1					64
7	1	7	21	35	35	21	7	1				128
8	1	8	28	56	70	56	28	8	1			256
9	1	9	36	84	126	126	84	36	9	1		512
10	1	10	45	120	210	252	210	120	45	10	1	1024

TABLE 2.1. Pascal's triangle for binomial coefficients, $\binom{n}{k}$.

The next identity explains the origin of the name for these numbers: They appear as coefficients when expanding powers of the binomial expression $x + y$.

The Binomial Theorem. *If n is a nonnegative integer, then*

$$(x + y)^n = \sum_k \binom{n}{k} x^k y^{n-k}. \tag{2.7}$$

The notation \sum_k means that the sum extends over all integers k . Thus, the right side of (2.7) is formally an infinite sum, but all terms with $k < 0$ or $k > n$ are zero by the expansion identity, so there are only $n + 1$ nonzero terms in this sum.

Proof. We prove this identity by induction on n . For $n = 0$, both sides evaluate to 1. Suppose then that the identity holds for a fixed nonnegative integer n . We need to verify that it holds for $n + 1$. Using our inductive hypothesis, then distributing the remaining factor of $(x + y)$, we obtain

$$\begin{aligned}(x + y)^{n+1} &= (x + y) \sum_k \binom{n}{k} x^k y^{n-k} \\ &= \sum_k \binom{n}{k} x^{k+1} y^{n-k} + \sum_k \binom{n}{k} x^k y^{n+1-k}.\end{aligned}$$

Now we reindex the first sum, replacing each occurrence of k by $k - 1$. Since the original sum extends over all values of k , the reindexed sum does, too. Thus

$$\begin{aligned}(x + y)^{n+1} &= \sum_k \binom{n}{k-1} x^k y^{n+1-k} + \sum_k \binom{n}{k} x^k y^{n+1-k} \\ &= \sum_k \left(\binom{n}{k-1} + \binom{n}{k} \right) x^k y^{n+1-k} \\ &= \sum_k \binom{n+1}{k} x^k y^{n+1-k},\end{aligned}$$

by the addition identity. This completes the induction, and we conclude that the identity holds for all $n \geq 0$. \square

We note two important consequences of the binomial theorem. First, setting $x = y = 1$ in (2.7), we obtain

$$\sum_k \binom{n}{k} = 2^n. \quad (2.8)$$

Thus, summing across the n th row in Pascal's triangle yields 2^n , and there are therefore exactly 2^n different subsets of a set of n elements. These row sums are included in Table 2.1.

Second, setting $x = -1$ and $y = 1$ in (2.7), we find that the alternating sum across any row of Pascal's triangle is zero, except of course for the top row:

$$\sum_k (-1)^k \binom{n}{k} = \begin{cases} 0 & \text{if } n \geq 1, \\ 1 & \text{if } n = 0. \end{cases} \quad (2.9)$$

This is obvious from the symmetry relation when n is odd, but less clear when n is even.

These two consequences of the binomial theorem concern sums over the lower index of binomial coefficients. The next identity tells us the value of a sum over the upper index.

Summing on the Upper Index. *If m and n are nonnegative integers, then*

$$\sum_{k=0}^n \binom{k}{m} = \binom{n+1}{m+1}. \quad (2.10)$$

Proof. We use induction on n to verify this identity. For $n = 0$, each side equals 1 if $m = 0$, and each side is 0 if $m > 0$. Suppose then that the identity holds for some fixed nonnegative integer n . We must show that it holds for the case $n + 1$. Let m be a nonnegative integer. We obtain

$$\begin{aligned} \sum_{k=0}^{n+1} \binom{k}{m} &= \binom{n+1}{m} + \sum_{k=0}^n \binom{k}{m} \\ &= \binom{n+1}{m} + \binom{n+1}{m+1} \\ &= \binom{n+2}{m+1}. \end{aligned}$$

By induction, the identity holds for all $n \geq 0$. □

To illustrate one last identity, we study the Lotto Texas game in more detail. Recall that a player selects six different numbers between 1 and 54 to enter the lottery. The largest prize is awarded to anyone matching all six numbers picked in a random drawing by lottery officials, but smaller prizes are given to players matching at least three of these numbers. To determine fair amounts for these smaller prizes, the state lottery commission needs to know the number of possible tickets that match exactly k of the winning numbers, for every k .

Clearly, there is just one way to match all six winning numbers. There are $\binom{6}{5} = 6$ ways to pick five of the six winning numbers, and 48 ways to select one losing number, so there are $6 \cdot 48 = 288$ tickets that match five numbers. Selecting four of the winning numbers and two of the losing numbers makes $\binom{6}{4} \binom{48}{2} = 16\,920$ possible tickets, and in general we see that the number of tickets that match exactly k of the winning numbers is $\binom{6}{k} \binom{48}{6-k}$. By summing over k , we count every possible ticket exactly once, so

$$\binom{54}{6} = \sum_k \binom{6}{k} \binom{48}{6-k}.$$

More generally, if a lottery game requires selecting m numbers from a set of $m+n$ numbers, we obtain the identity

$$\binom{m+n}{m} = \sum_k \binom{m}{k} \binom{n}{m-k}.$$

That is, the number of possible tickets equals the sum over k of the number of ways to match exactly k of the m winning numbers and $m - k$ of the n losing numbers. More generally still, suppose a lottery game requires a player to select ℓ numbers on a ticket, and each drawing selects m winning numbers. Using the same reasoning, we find that

$$\binom{m+n}{\ell} = \sum_k \binom{m}{k} \binom{n}{\ell-k}.$$

Now replace ℓ by $\ell + p$ and reindex the sum, replacing k by $k + p$, to obtain the following identity.

Vandermonde's Convolution. *If m and n are nonnegative integers and ℓ and p are integers, then*

$$\binom{m+n}{\ell+p} = \sum_k \binom{m}{p+k} \binom{n}{\ell-k}. \quad (2.11)$$

Notice that the lower indices in the binomial coefficients on the right side sum to a constant.

Exercises

1. Use a combinatorial argument to prove that there are exactly 2^n different subsets of a set of n elements. (Do not use the binomial theorem.)
2. Prove the absorption/extraction identity: If n is a positive integer and k is a nonzero integer, then

$$\binom{n}{k} = \frac{n}{k} \binom{n-1}{k-1}. \quad (2.12)$$

3. Use algebraic methods to prove the cancellation identity: If n and k are nonnegative integers and m is an integer with $m \leq n$, then

$$\binom{n}{k} \binom{k}{m} = \binom{n}{m} \binom{n-m}{k-m}. \quad (2.13)$$

This identity is very useful when the left side appears in a sum over k , since the right side has only a single occurrence of k .

4. Suppose that a museum curator with a collection of n paintings by Jackson Pollack needs to select k of them for display, and needs to pick m of these to put in a particularly prominent part of the display. Show how to count the number of possible combinations in two ways so that the cancellation identity appears.

5. Prove the parallel summation identity: If m and n are nonnegative integers, then

$$\sum_{k=0}^n \binom{m+k}{k} = \binom{m+n+1}{n}. \quad (2.14)$$

6. Prove the hexagon identity: If n is a positive integer and k is an integer, then

$$\binom{n-1}{k-1} \binom{n}{k+1} \binom{n+1}{k} = \binom{n-1}{k} \binom{n}{k-1} \binom{n+1}{k+1}. \quad (2.15)$$

Why is it called the hexagon identity?

7. Compute the value of the following sums. Your answer should be an expression involving one or two binomial coefficients.

(a) $\sum_k \binom{80}{k} \binom{k+1}{31}$.

(b) $\sum_{k \geq 0} \frac{1}{k+1} \binom{99}{k} \binom{200}{120-k}$.

(c) $\sum_{k=100}^{201} \sum_{j=100}^k \binom{201}{k+1} \binom{j}{100}$.

(d) $\sum_k \binom{n}{k}^2$, for a nonnegative integer n .

(e) $\sum_{k \leq m} (-1)^k \binom{n}{k}$, for an integer m and a nonnegative integer n .

8. Prove the binomial theorem for falling factorial powers,

$$(x+y)^{\overline{n}} = \sum_k \binom{n}{k} x^{\overline{k}} y^{\overline{n-k}},$$

and for rising factorial powers,

$$(x+y)^{\overline{n}} = \sum_k \binom{n}{k} x^{\overline{k}} y^{\overline{n-k}}.$$

9. Let n be a nonnegative integer. Suppose $f(x)$ and $g(x)$ are functions defined for all real numbers x , and that both functions are n times differentiable. Let $f^{(k)}(x)$ denote the k th derivative of $f(x)$, so $f^{(0)}(x) = f(x)$, $f^{(1)}(x) = f'(x)$, and $f^{(2)}(x) = f''(x)$. Let $h(x) = f(x)g(x)$. Show that

$$h^{(n)}(x) = \sum_k \binom{n}{k} f^{(k)}(x) g^{(n-k)}(x).$$

10. In the Virginia lottery game Win For Life, an entry consists of a selection of six different numbers between 1 and 42, and each drawing selects seven different numbers in this range. How many different entries can match at least three of the drawn numbers?
11. The state of Florida administers several lottery games. In Florida Lotto, a player picks a set of six numbers between 1 and 53. In Fantasy 5, a gambler chooses a set of five numbers between 1 and 36. In which game is a player more likely to match at least two numbers against the ones drawn?

2.3 Multinomial Coefficients

*Alba, alma, ball, balm, bama, blam, lama, lamb, ma'am, mall,
malm, mama, ...*

— Words constructed from letters in “Alma, Alabama”

Suppose we want to know the number of ways to place n different objects into two boxes, one marked A and the other marked B , in such a way that box A receives a specified number a of the objects, and box B gets the remaining b objects, so $a + b = n$. Assume that the order of placement of the objects in each box is immaterial, and denote the total number of such arrangements by $\binom{n}{a,b}$. We can compute this number easily by using our knowledge of binomial coefficients. Since each valid distribution corresponds to a different subset of a objects for box A , we see that $\binom{n}{a,b}$ is simply the binomial coefficient $\binom{n}{a}$ (or $\binom{n}{b}$). Thus, $\binom{n}{a,b} = \frac{n!}{a!b!}$.

Now imagine we have three boxes, labeled A , B , and C , and suppose we want to know the number of ways to place a prescribed number a of the objects in box A , a given number b in box B , and the remaining $c = n - a - b$ in box C . Again, assume the order of placement of objects in each box is irrelevant, and denote this number by $\binom{n}{a,b,c}$. Since each arrangement can be described by first selecting a elements from the set of n for box A , and then picking b objects from the remaining $n - a$ for box B , we see by the product rule that

$$\begin{aligned} \binom{n}{a,b,c} &= \binom{n}{a} \binom{n-a}{b} \\ &= \frac{n!}{a!(n-a)!} \cdot \frac{(n-a)!}{b!(n-a-b)!} \\ &= \frac{n!}{a!b!c!}. \end{aligned} \tag{2.16}$$

The number $\binom{n}{a,b,c}$ is called a *trinomial coefficient*.

We can generalize this problem for an arbitrary number of boxes. Suppose we have n objects, together with m boxes labeled $1, 2, \dots, m$, and suppose k_1, k_2, \dots, k_m are nonnegative integers satisfying $k_1 + k_2 + \dots + k_m = n$. We define the

multinomial coefficient $\binom{n}{k_1, k_2, \dots, k_m}$ to be the number of ways to place k_1 of the objects in box 1, k_2 in box 2, and so on, without regard to the order of the objects in each box. Then an argument similar to our analysis for trinomial coefficients shows that

$$\begin{aligned} \binom{n}{k_1, \dots, k_m} &= \binom{n}{k_1} \binom{n-k_1}{k_2} \binom{n-k_1-k_2}{k_3} \dots \\ &\quad \cdot \binom{n-k_1-\dots-k_{m-2}}{k_{m-1}} \\ &= \frac{n!}{k_1! k_2! \dots k_m!}. \end{aligned} \quad (2.17)$$

Multinomial coefficients often arise in a natural way in combinatorial problems. While we can always reduce questions about multinomial coefficients to problems about binomial coefficients or factorials by using (2.17), it is often useful to handle them directly. We derive some important formulas for multinomial coefficients in this section. These generalize some of the statements about binomial coefficients from Section 2.2. We begin with a more general formula for expanding multinomial coefficients in terms of factorials.

Expansion. *If n is a nonnegative integer, and k_1, \dots, k_m are integers satisfying $k_1 + \dots + k_m = n$, then*

$$\binom{n}{k_1, \dots, k_m} = \begin{cases} \frac{n!}{k_1! \dots k_m!} & \text{if each } k_i \geq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (2.18)$$

Taking $\binom{n}{k_1, \dots, k_m} = 0$ when at least one of the k_i is negative is certainly sensible, since it is impossible to place a negative number of objects in a box.

Second, it is clear that rearranging the numbers k_1, \dots, k_m does not affect the value of the multinomial coefficient $\binom{n}{k_1, \dots, k_m}$, since this just corresponds to relabeling the boxes. We can state this in the following way.

Symmetry. *Suppose $\pi(1), \dots, \pi(m)$ is a permutation of $\{1, \dots, m\}$. Then*

$$\binom{n}{k_1, \dots, k_m} = \binom{n}{k_{\pi(1)}, \dots, k_{\pi(m)}}. \quad (2.19)$$

Third, we can observe a simple addition law. Let α be one of the objects from the set of n . It must be placed in one of the boxes. If we place α in box 1, then there are $\binom{n-1}{k_1-1, k_2, \dots, k_m}$ ways to arrange the remaining $n-1$ objects to create a valid arrangement. If we set α in box 2, then there are $\binom{n-1}{k_1, k_2-1, k_3, \dots, k_m}$ to complete the assignment of objects to boxes. Continuing in this way, we obtain the following identity.

Addition. If n is a positive integer and $k_1 + \dots + k_m = n$, then

$$\binom{n}{k_1, \dots, k_m} = \binom{n-1}{k_1-1, k_2, \dots, k_m} + \binom{n-1}{k_1, k_2-1, k_3, \dots, k_m} + \dots + \binom{n-1}{k_1, k_2, \dots, k_{m-1}, k_m-1}. \tag{2.20}$$

In the last section, the addition identity for $m = 2$ produced Pascal's triangle for the binomial coefficients. We can use a similar strategy to generate a geometric arrangement of the trinomial coefficients when $m = 3$, which we might call *Pascal's pyramid*. The top level of the pyramid corresponds to $n = 0$, just as in Pascal's triangle, and here we place a single 1, for $\binom{0}{0,0,0}$. The next level holds the numbers for $n = 1$, and we place the three 1s in a triangular formation, just below the $n = 0$ datum at the apex, for the numbers $\binom{1}{1,0,0}$, $\binom{1}{0,1,0}$, and $\binom{1}{0,0,1}$. In general, we use the addition formula (2.20) to compute the numbers in level n from those in level $n - 1$, and we place the value of $\binom{n}{a,b,c}$ in level n just below the triangular arrangement of numbers $\binom{n-1}{a-1,b,c}$, $\binom{n-1}{a,b-1,c}$, and $\binom{n-1}{a,b,c-1}$ in level $n - 1$. Figure 2.1 shows the first few levels of the pyramid of trinomial coefficients. Here, the position of each number in level n is shown relative to the positions of the numbers in level $n - 1$, each of which is marked with a triangle (Δ).




FIGURE 2.1. The first five levels of Pascal's pyramid.

We can use the addition identity to obtain an important generalization of the binomial theorem for multinomial coefficients.

The Multinomial Theorem. *If n is a nonnegative integer, then*

$$(x_1 + \cdots + x_m)^n = \sum_{k_1 + \cdots + k_m = n} \binom{n}{k_1, \dots, k_m} x_1^{k_1} \cdots x_m^{k_m}. \quad (2.21)$$

Here, the notation $\sum_{k_1 + \cdots + k_m = n}$ means that the sum extends over all integer m -tuples (k_1, \dots, k_m) whose sum is n . Of course, there are infinitely many such m -tuples, but only finitely many produce a nonzero term by the Expansion identity, so this is in effect a finite sum. We prove (2.21) for the case $m = 3$; the general case is left as an exercise.

Proof. The formula

$$(x + y + z)^n = \sum_{a+b+c=n} \binom{n}{a, b, c} x^a y^b z^c \quad (2.22)$$

certainly holds for $n = 0$, so suppose that it is valid for n . We compute

$$\begin{aligned} (x+y+z)^{n+1} &= (x+y+z) \sum_{a+b+c=n} \binom{n}{a, b, c} x^a y^b z^c \\ &= \sum_{a+b+c=n} \binom{n}{a, b, c} x^{a+1} y^b z^c + \sum_{a+b+c=n} \binom{n}{a, b, c} x^a y^{b+1} z^c \\ &\quad + \sum_{a+b+c=n} \binom{n}{a, b, c} x^a y^b z^{c+1} \\ &= \sum_{a+b+c=n+1} \left(\binom{n}{a-1, b, c} + \binom{n}{a, b-1, c} + \binom{n}{a, b, c-1} \right) x^a y^b z^c \\ &= \sum_{a+b+c=n+1} \binom{n+1}{a, b, c} x^a y^b z^c, \end{aligned}$$

so (2.22) holds for all $n \geq 0$. \square

Some additional formulas for multinomial coefficients are developed in the exercises. Some of these may be obtained by selecting particular values for x_1, \dots, x_m in the multinomial theorem.

We close this section by describing a common way that multinomial coefficients appear in combinatorial problems. Suppose we need to count the number of ways to order a collection of n objects. If all the objects are different, then the answer is simply $n!$, but what if our collection includes some duplicate objects? Such a collection is called a *multiset*. Certainly we expect fewer different arrangements when there are some duplicate objects. For example, there are just six ways

to line up four poker chips, two of which are red and the other two blue: $rrbb$, $rbrb$, $rbbr$, $brrb$, $brbr$, and $bbrr$.

Suppose we have a multiset of size n that includes exactly k_1 identical copies of one object, k_2 instances of another, and so on, ending with k_m duplicates of the last object, so $k_1 + \cdots + k_m = n$. In any ordering of these n objects, we may rearrange the k_i copies of object i in any way without disturbing the arrangement. Since we can do this for any of the m objects independently, it follows that each distinct ordering of the items occurs $k_1!k_2!\cdots k_m!$ times among the $n!$ ways that one could arrange the objects if they had been distinguishable. Therefore, the number of distinct arrangements of our multiset is

$$\frac{n!}{k_1! \cdots k_m!} = \binom{n}{k_1, \dots, k_m}.$$

We could also obtain this formula by using our first combinatorial model for the multinomial coefficients. Suppose we have n ping-pong balls, numbered 1 through n , and m boxes, each labeled with a different object from our multiset. The number of ways to distribute the balls among the boxes, with k_1 in box 1, k_2 in box 2, and so on, is the multinomial coefficient $\binom{n}{k_1, \dots, k_m}$. But each arrangement corresponds to an ordering of the elements of our multiset: The numbers in box i indicate the positions of object i in the listing.

We have thus answered the analogue of Problem 1 from Section 2.1 for multisets. We can also study a generalization of Problem 2: How many ways are there to make an ordered list of r objects from a multiset of n objects, if the multiset comprises k_i copies of object i for $1 \leq i \leq m$? Our approach to this problem depends on the k_i and r , so we'll study an example. Suppose a contemplative resident of Alma, Alabama, wants to know the number of ways to rearrange the letters of her home town and state, ignoring differences in case. There are eleven letters in all: six As, one B, two Ls, and two Ms, so she computes the total number to be $\binom{11}{6,1,2,2} = \frac{11!}{6!2!2!} = 13\,860$.

Suppose she also wants to know the number of four-letter sequences of letters that can be formed from the same string, ALMAALABAMA, like the ones in the list that open this section, only they do not have to be English words. This is the multiset version of Problem 2 with $n = 11$, $r = 4$, $m = 4$, $k_1 = 6$, $k_2 = 1$, and $k_3 = k_4 = 2$. We can solve this by constructing each sequence in two steps: first, select four elements from the multiset; second, count the number of ways to order that subcollection. We can group the possible sub-multisets according to their pattern of repeated elements. For example, consider the subcollections that have two copies of one object, and two copies of another. Denote this pattern by $wvxx$. There are $\binom{3}{2} = 3$ ways to select values for w and x , since we must pick two of the three letters A, L, and M. Each of these subcollections can be ordered in any of $\binom{4}{2,2} = 6$ ways, so the pattern $wvxx$ produces $3 \cdot 6 = 18$ possible four-letter sequences in all. There are five possible patterns for a four-element multiset, which we can denote $wwww$, $wwvx$, $wvxx$, $wvxy$, and $wxyz$. The analysis of each one is summarized in the following table.

Pattern	Sub-multisets	Orderings per sub-multiset	Total
www	1	1	1
$wwwx$	$\binom{3}{1}$	$\binom{4}{3,1}$	12
wwx	$\binom{3}{2}$	$\binom{4}{2,2}$	18
wxy	$\binom{3}{1}\binom{3}{2}$	$\binom{4}{2,1,1}$	108
$wxyz$	1	4!	24

Summing the values in the rightmost column, we find that there are exactly 163 ways to form a four-letter sequence from the letters in Alma, Alabama.

Exercises

1. Prove the addition identity for multinomial coefficients (2.20) by using the expansion identity (2.18).
2. For nonnegative integers a , b , and c , let $P(a, b, c)$ denote the number of paths in three-dimensional space that begin at the origin, end at (a, b, c) , and consist entirely of steps of unit length, each of which is parallel to a coordinate axis. Prove that $P(a, b, c) = \binom{a+b+c}{a, b, c}$.
3. Prove the multinomial theorem (2.21) for an arbitrary positive integer m .
4. Prove the following identities for sums of multinomial coefficients, if m and n are positive integers.

$$(a) \quad \sum_{k_1 + \dots + k_m} \binom{n}{k_1, \dots, k_m} = m^n.$$

$$(b) \quad \sum_{k_1 + \dots + k_m} \binom{n}{k_1, \dots, k_m} (-1)^{k_2 + k_4 + \dots + k_{2\ell}} = \begin{cases} 0 & \text{if } m = 2\ell, \\ 1 & \text{if } m = 2\ell + 1. \end{cases}$$

5. Prove that if n is a nonnegative integer and k is an integer, then

$$\sum_j \binom{n}{j, k, n - j - k} = 2^{n-k} \binom{n}{k}.$$

6. Prove the multinomial theorem for falling factorial powers,

$$(x_1 + \dots + x_m)^{\underline{n}} = \sum_{k_1 + \dots + k_m = n} \binom{n}{k_1, \dots, k_m} x_1^{\underline{k_1}} \dots x_m^{\underline{k_m}},$$

and for rising factorial powers,

$$(x_1 + \dots + x_m)^{\overline{n}} = \sum_{k_1 + \dots + k_m = n} \binom{n}{k_1, \dots, k_m} x_1^{\overline{k_1}} \dots x_m^{\overline{k_m}}.$$

You may find it helpful to consider the trinomial case first.

7. Use a combinatorial argument to establish the following analogue of Vandermonde's convolution for trinomial coefficients. If m and n are nonnegative integers, and $a + b + c = m + n$, then

$$\sum_{\alpha+\beta+\gamma=m} \binom{m}{\alpha, \beta, \gamma} \binom{n}{a-\alpha, b-\beta, c-\gamma} = \binom{m+n}{a, b, c}.$$

8. State an analogue of Vandermonde's convolution for multinomial coefficients, and use a combinatorial argument to establish it.
9. Compute the number of r -letter sequences that can be formed by using the letters in each location below, for each given value of r . Ignore differences in case.
- Bug Tussle, TX: $r = 3, r = 4, r = 11$.
 - Cooleemee, NC: $r = 4, r = 10, r = 11$.
 - Oconomowoc, WI: $r = 4, r = 11, r = 12$.
 - Unalaska, Alaska: $r = 3, r = 4, r = 14$.
 - Walla Walla, WA: $r = 4, r = 5, r = 12$.
10. Certainly there are more four-letter sequences that can be formed by using the letters in Bobo, Mississippi, than can be formed by using the letters in Soso, Mississippi. Is the difference more or less than the distance between these two cities in miles, which is 267?
11. A band of combinatorial ichthyologists asserts that the number of five-letter sequences that can be formed using the letters of the Hawaiian long-nosed butterfly fish, the lauwiwilinukunuku'oi'oi, is more than twice as large as the number of five-letter sequences that can be created using the name of the state fish of Hawaii, the painted triggerfish humuhumunukunukuapua'a. Prove or disprove their claim by computing the exact number in each case.

2.4 The Pigeonhole Principle

I am just here for anyone that's for the pigeons.

— Mike Tyson, Phoenix City Council meeting, June 1, 2005,
reported in *The Arizona Republic*

We now turn to a simple, but powerful, idea in combinatorial reasoning known as the *pigeonhole principle*. We can state it in the following way.

Theorem 2.1 (Pigeonhole Principle). *Let n be a positive integer. If more than n objects are distributed among n containers, then some container must contain more than one object.*

The proof is simple—if each container held at most one object, then there would be at most n objects in all.

This mathematical idea is also called the *box principle* (especially in number theory texts), which is sensible enough, since we can imagine the containers as boxes. In German, it is the *drawer principle*, logically enough, after Dirichlet's original term, the *Schubfachprinzip*. It may seem odd to think of our containers as pigeon roosts, but the name probably originally referred to the “pigeonholes” one sees in those old desks with lots of square nooks for squirreling away papers. (One imagines however that the origins of the term may be the subject of some, well, squabbling. . . .) So while the traditional name may be somewhat antiquated, at least the avian nomenclature saves us from talking about Dirichlet's drawers.

The pigeonhole principle is very useful in establishing the existence of a particular configuration or combination in many mathematical contexts. We begin with a few simple examples.

1. Suppose 400 freshmen enroll in introductory calculus one term. Then two must have the same birthday. Here, the pigeonholes are calendar days, so $n = 366$.
2. In honor of champion pugilist (and pigeon enthusiast) Mike Tyson, suppose that n boxers schedule a round-robin tournament, so each fighter meets every other in a bout, and afterwards no contestant is undefeated. Then each boxer has between 1 and $n - 1$ wins, so two boxers must have the same record in the tournament.
3. It is estimated that the average full head of hair has 100 000 to 150 000 strands of hair. Let's assume that the most hirsute among us has less than 250 000 strands of hair on their head. The city of Phoenix has over 1.5 million residents, so it follows that there must be at least two residents with exactly the same number of hairs on their head. Moreover, since only a fraction of the population is bald, the statement surely remains true if we exclude those with no hair at all. (Sorry, Iron Mike.)

In this last problem, we can in fact conclude considerably more. The population of Phoenix is more than six times the maximum number of hairs per head, and a moment's thought reveals that there must in fact exist at least *six* people in Phoenix with identical hair counts. We can thus state a more powerful pigeonhole principle.

Theorem 2.2 (Generalized Pigeonhole Principle). *Let m and n be positive integers. If more than mn objects are distributed among n containers, then at least one container must contain at least $m + 1$ objects.*

The proof is again easy—if each container held at most m objects then the total number of objects would be at most mn . An alternative formulation of this statement appears in the exercises. Next, we establish the following arithmetic variation on the pigeonhole principle.

Theorem 2.3. *Suppose a_1, a_2, \dots, a_n is a sequence of real numbers with mean μ , so $\mu = (a_1 + \dots + a_n)/n$. Then there exist integers i and j , with $1 \leq i, j \leq n$, such that $a_i \leq \mu$ and $a_j \geq \mu$.*

The proof is again straightforward—if every element of the sequence were strictly greater than μ , then we would have $a_1 + \dots + a_n > n\mu$, a contradiction. Thus, an integer i must exist with $a_i \leq \mu$. A similar argument establishes the existence of j .

While the pigeonhole principle and the variations we describe here are all quite simple to state and verify, this idea plays a central role in establishing many decidedly nontrivial statements in mathematics. We conclude this section with two examples.

Monotonic Subsequences

We say a sequence a_1, \dots, a_n is *increasing* if $a_1 \leq a_2 \leq \dots \leq a_n$, and *strictly increasing* if $a_1 < a_2 < \dots < a_n$. We define *decreasing* and *strictly decreasing* in the same way. Consider first an arrangement of the integers between 1 and 10, for example,

$$3, 5, 8, 10, 6, 1, 9, 2, 7, 4. \quad (2.23)$$

Scan the list for an increasing subsequence of maximal length. Above, we find $(3, 5, 8, 10)$, $(3, 5, 8, 9)$, $(3, 5, 6, 7)$, and $(3, 5, 6, 9)$ all qualify with length 4. Next, scan the list for a decreasing subsequence of maximal length. Here, the best we can do is length 3, achieved by $(8, 6, 1)$, $(8, 6, 2)$, $(8, 6, 4)$, $(10, 6, 2)$, $(10, 6, 4)$, $(10, 7, 4)$, and $(9, 7, 4)$. Is it possible to find an arrangement of the integers from 1 to 10 that simultaneously avoids both an increasing subsequence of length 4 and a decreasing subsequence of length 4? The following theorem asserts that this is not possible. Its statement dates to an early and influential paper of Erdős and Szekeres [94], the same one cited in Section 1.8 for its contribution to the development of Ramsey theory.

Theorem 2.4. *Suppose m and n are positive integers. A sequence of more than mn real numbers must contain either an increasing subsequence of length at least $m + 1$, or a strictly decreasing subsequence of length at least $n + 1$.*

Proof. Suppose that $r_1, r_2, \dots, r_{mn+1}$ is a sequence of real numbers which contains neither an increasing subsequence of length $m + 1$, nor a strictly decreasing subsequence of length $n + 1$. For each integer i with $1 \leq i \leq mn + 1$, let a_i denote the length of the longest increasing subsequence in this sequence of numbers whose first term is r_i , and let d_i denote the length of the longest strictly decreasing subsequence beginning with this term. For example, for the sequence (2.23) we see that $a_2 = 3$ (for $5, 8, 10$ or $5, 8, 9$), and $d_2 = 2$ (for $5, 1$ or $5, 2$). By our hypothesis, we know that $1 \leq a_i \leq m$ and $1 \leq d_i \leq n$ for each i , and thus there are only mn different possible values for the ordered pair (a_i, d_i) . However, there are $mn + 1$ such ordered pairs, so by the pigeonhole principle there exist two integers j and k with $j < k$ such that $a_j = a_k$ and $d_j = d_k$. Denote this pair

by (α, δ) , so $\alpha = a_j = a_k$ and $\delta = d_j = d_k$. Now let $r_k, r_{i_2}, \dots, r_{i_\alpha}$ denote a maximal increasing subsequence beginning with r_k and let $r_k, r_{i'_2}, \dots, r_{i'_\delta}$ denote a maximal strictly decreasing subsequence beginning with this term. If $r_j \leq r_k$, then $r_j, r_k, r_{i_2}, \dots, r_{i_\alpha}$ is an increasing subsequence of length $\alpha + 1$ beginning with r_j . On the other hand, if $r_j > r_k$, then $r_j, r_k, r_{i'_2}, \dots, r_{i'_\delta}$ is a strictly decreasing subsequence of length $\delta + 1$ beginning with r_j . In either case, we reach a contradiction. \square

Of course, we can replace “increasing” with “strictly increasing” and simultaneously “strictly decreasing” with “decreasing” in this statement.

Approximating Irrationals by Rationals

Let α be an irrational number. Since every real interval $[a, b]$ with $a < b$ contains infinitely many rational numbers, certainly there exist rational numbers arbitrarily close to α . Suppose however we restrict the rationals we may select to the set of fractions with bounded denominator. How closely can we approximate α now? More specifically, given an irrational number α and a positive integer Q , does there exist a rational number p/q with $1 \leq q \leq Q$ and $\left| \alpha - \frac{p}{q} \right|$ especially small? How small can we guarantee?

At first glance, if we select a random denominator q in the range $[1, Q]$, then certainly α lies in some interval $(\frac{k}{q}, \frac{k+1}{q})$, for some integer k , so its distance to the nearest multiple of $1/q$ is at most $1/2q$. We might therefore expect that on average we would observe a distance of about $1/4q$, for randomly selected q . In view of Theorem 2.3, we might then expect that approximations with distance at most $1/4q$ must exist. In fact, however, we can establish a much stronger result by using the pigeonhole principle. The following important theorem is due to Dirichlet and his *Schubfachprinzip*.

We first require some notation. For a real number x , let $\lfloor x \rfloor$ denote the *floor* of x , or *integer part* of x . It is defined to be the largest integer m satisfying $m \leq x$. Similarly, the *ceiling* of x , denoted by $\lceil x \rceil$, is the smallest integer m satisfying $x \leq m$. Last, the *fractional part* of x , denoted by $\{x\}$, is defined by $\{x\} = x - \lfloor x \rfloor$. For example, for $x = \pi$ we have $\lfloor \pi \rfloor = 3$, $\lceil \pi \rceil = 4$, and $\{\pi\} = 0.14159\dots$; for $x = 1$ we obtain $\lfloor 1 \rfloor = \lceil 1 \rceil = 1$ and $\{1\} = 0$.

Theorem 2.5 (Dirichlet’s Approximation Theorem). *Suppose α is an irrational real number, and Q is a positive integer. Then there exists a rational number p/q with $1 \leq q \leq Q$ satisfying*

$$\left| \alpha - \frac{p}{q} \right| < \frac{1}{q(Q+1)}.$$

Proof. Divide the real interval $[0, 1]$ into $Q + 1$ subintervals of equal length:

$$\left[0, \frac{1}{Q+1} \right), \left[\frac{1}{Q+1}, \frac{2}{Q+1} \right), \dots, \left[\frac{Q-1}{Q+1}, \frac{Q}{Q+1} \right), \left[\frac{Q}{Q+1}, 1 \right].$$

Since each of the $Q + 2$ real numbers

$$0, \{\alpha\}, \{2\alpha\}, \dots, \{Q\alpha\}, 1 \quad (2.24)$$

lies in $[0, 1]$, by the pigeonhole principle at least two of them must lie in the same subinterval. Each of the numbers in (2.24) can be written in a unique way as $r\alpha - s$ with r and s integers and $0 \leq r \leq Q$, so it follows that there exist integers r_1, r_2, s_1 , and s_2 , with $0 \leq r_1, r_2 \leq Q$, such that

$$|(r_2\alpha - s_2) - (r_1\alpha - s_1)| < \frac{1}{Q+1}.$$

Since only 0 and 1 in our list have the same r -value, we can assume that $r_1 \neq r_2$, so suppose $r_1 < r_2$. Let $q = r_2 - r_1$, so $1 \leq q \leq Q$, and let $p = s_2 - s_1$. Then p and q satisfy

$$|q\alpha - p| < \frac{1}{Q+1},$$

and the conclusion follows upon dividing through by q . \square

Since $q \leq Q$, we immediately obtain that the rational number p/q guaranteed by the theorem satisfies

$$\left| \alpha - \frac{p}{q} \right| < \frac{1}{q^2 + q}. \quad (2.25)$$

Exercise 11 asks you to show that there exist infinitely many rational numbers p/q that satisfy this inequality for a fixed irrational number α .

Exercises

1. Show that at any party with at least two people, there must exist at least two people in the group who know the same number of other guests at the party. Assume that each pair of people at the party are either mutual friends or mutual strangers.
2. Prove the following version of the pigeonhole principle. Let m and n be positive integers. If m objects are distributed in some way among n containers, then at least one container must hold at least $1 + \lfloor (m-1)/n \rfloor$ objects.
3. Prove the following more general version of the pigeonhole principle. Suppose that m_1, m_2, \dots, m_n are all positive integers, let $M = m_1 + m_2 + \dots + m_n - n + 1$, and suppose each of n containers is labeled with an integer between 1 and n . Prove that if M objects are distributed in some way among the n containers, then there exists an integer i between 1 and n such that the container labeled with i contains at least m_i objects.
4. The top four pitchers on a college baseball team combine for 297 strikeouts over the course of a season. If each pitcher had at least 40 strikeouts over the

course of the season, and the fourth-best pitcher had less than 50 strikeouts, how many strikeouts could the best pitcher have made over the season? Your answer should be a range of possible numbers.

5. Find the smallest value of m so that the following statement is valid: Any collection of m distinct positive integers must contain at least two numbers whose sum or difference is a multiple of 10. Prove that your value is best possible.
6. Suppose $A = (a_1, a_2, \dots, a_n)$ is a sequence of positive real numbers. Let $H(A)$ denote the *harmonic mean* of A , defined by

$$H(A) = n \left(\sum_{i=1}^n \frac{1}{a_i} \right)^{-1}.$$

Show there exist integers i and j , with $1 \leq i, j \leq n$, satisfying

$$a_i \leq H(A) \leq a_j.$$

7. Suppose the integers from 1 to n are arranged in some order around a circle, and let k be an integer with $1 \leq k \leq n$. Show that there must exist a sequence of k adjacent numbers in the arrangement whose sum is at least $\lceil k(n+1)/2 \rceil$.
8. Suppose the integers from 1 to n are arranged in some order around a circle, and let k be an integer with $1 \leq k \leq n$. Show that there must exist a sequence of k adjacent numbers in the arrangement whose product is at least $\lceil (n!)^{k/n} \rceil$.
9. Let n be a positive integer. Exhibit an arrangement of the integers between 1 and n^2 which has no increasing or decreasing subsequence of length $n+1$.
10. Let m and n be positive integers. Exhibit an arrangement of the integers between 1 and mn which has no increasing subsequence of length $m+1$, and no decreasing subsequence of length $n+1$.
11. Let α be an irrational number. Prove that there exist infinitely many rational numbers p/q satisfying (2.25).
12. Let n be a positive integer, and let $b \geq 2$ be an integer.
 - (a) Show that there exists a nonzero multiple N of n whose base- b representation consists entirely of 0s and 1s. (No partial credit will be awarded for the case $b = 2$!) Hint: Consider the sequence of numbers $\sum_{i=0}^k b^i$ for a number of values of k .
 - (b) Show that there exists a multiple N of n whose base- b representation consists entirely of 1s if and only if no prime number p which divides b is a factor of n .

- (c) Suppose that the greatest common divisor of b and n is 1, and let d_1, \dots, d_m be a sequence of integers with $0 \leq d_i < b$ for each i and $d_1 \neq 0$. Show that there exists a multiple N of n whose base- b representation is obtained by juxtaposing some integral number of copies of the base- b digit sequence $d_1 d_2 \cdots d_m$.
13. Let a_1, a_2, \dots, a_n be a sequence of integers. Show that there exist integers j and k with $1 \leq j \leq k \leq n$ such that the sum $\sum_{i=j}^k a_i$ is a multiple of n .
14. (Bloch and Pólya [28].) For a positive integer d , let \mathcal{N}_d denote the set of polynomials with degree at most $d - 1$ whose coefficients are all 0 or 1. For example, $\mathcal{N}_3 = \{0, 1, x, x^2, 1 + x, 1 + x^2, x + x^2, 1 + x + x^2\}$.
- (a) Let $f^{(k)}(x)$ denote the k th derivative of $f(x)$. Show that if $f \in \mathcal{N}_d$ then $f^{(k-1)}(1) \leq d^k/k$.
- (b) Let m be a positive integer. Determine an upper bound on the number of different possible m -tuples $(f(1), f'(1), \dots, f^{(m-1)}(1))$ achieved by polynomials $f(x) \in \mathcal{N}_d$.
- (c) Prove that if $d > 1$ and

$$\frac{d}{\log_2 d} > \binom{m+1}{2},$$

then there exists a polynomial $h(x)$ of degree at most $d - 1$ whose coefficients are all 0, 1, or -1 , and which is divisible by $(x - 1)^m$.

2.5 The Principle of Inclusion and Exclusion

What we here have to do is to conceive, and invent a notation for, all the possible combinations which any number of class terms can yield; and then find some mode of symbolic expression which shall indicate which of these various compartments are empty or occupied...

— John Venn, [275, p. 23]

Suppose there are 50 beads in a drawer: 25 are glass, 30 are red, 20 are spherical, 18 are red glass, 12 are glass spheres, 15 are red spheres, and 8 are red glass spheres. How many beads are neither red, nor glass, nor spheres?

We can answer this question by organizing all of this information using a Venn diagram with three overlapping sets: G for glass beads, R for red beads, and S for spherical beads. See Figure 2.2. We are given that there are eight red glass spheres, so start by labeling the common intersection of the sets G , R , and S in the diagram with 8. Then the region just above this one must have ten elements, since there are 18 red glass beads, and exactly eight of these are spherical. Continuing in this way, we determine the size of each of the sets represented in the diagram, and we

conclude that there are exactly twelve beads in the drawer that are neither red, nor glass, nor spheres.




FIGURE 2.2. A solution using a Venn diagram.

Alternatively, we can answer this question by determining the size of the set $G \cup R \cup S$ (does this make us counting GURUS?). Summing the number of elements in the sets G , R , and S produces a number that is too large, since this sum counts the beads that are in more than one of these sets at least twice. We can try to compensate by subtracting the number of elements in the sets $G \cap R$, $G \cap S$, and $R \cap S$ from the sum. This produces a total that is too small, since the beads that have all three attributes are counted three times in the first step, then subtracted three times in the second step. Thus, we must add the number of elements in $G \cap R \cap S$ to the sum, and we find that

$$|G \cup R \cup S| = |G| + |R| + |S| - |G \cap R| - |G \cap S| - |R \cap S| + |G \cap R \cap S|.$$

Letting N_0 denote the number of beads with none of the three attributes, we then compute

$$\begin{aligned} N_0 &= 50 - |G \cup R \cup S| \\ &= 50 - |G| - |R| - |S| + |G \cap R| + |G \cap S| + |R \cap S| - |G \cap R \cap S| \\ &= 50 - 25 - 30 - 20 + 18 + 12 + 15 - 8 \\ &= 12. \end{aligned}$$

This suggests a general technique for solving some similar combinatorial problems. Suppose we have a collection of N distinct objects, and each object may satisfy one or more properties that we label a_1, a_2, \dots, a_r . Let $N(a_i)$ denote the number of objects having property a_i , let $N(a_i a_j)$ signify the number having

both property a_i and property a_j , and in general let $N(a_{i_1} a_{i_2} \dots a_{i_m})$ represent the number satisfying the m properties a_{i_1}, \dots, a_{i_m} . Let N_0 denote the number of objects having none of the properties. We prove the following theorem.

Theorem 2.6 (Principle of Inclusion and Exclusion). *Using the notation above,*

$$\begin{aligned} N_0 = N &- \sum_i N(a_i) + \sum_{i < j} N(a_i a_j) - \sum_{i < j < k} N(a_i a_j a_k) + \dots \\ &+ (-1)^m \sum_{i_1 < \dots < i_m} N(a_{i_1} \dots a_{i_m}) + \dots + (-1)^r N(a_1 a_2 \dots a_r). \end{aligned} \quad (2.26)$$

Proof. Suppose an object satisfies none of the properties. Then the expression on the right side counts it precisely once, in the N term. On the other hand, suppose an object satisfies precisely m of the properties, with m a positive number. Then it is counted once in the N term, m times in the $\sum N(a_i)$ term, $\binom{m}{2}$ times in the second sum, and in general $\binom{m}{k}$ times in the k th sum. Therefore, the total contribution on the right side from this object is

$$\sum_k (-1)^k \binom{m}{k} = 0$$

by (2.9). This completes the proof. \square

We consider four applications of this counting principle.

The Euler φ Function

Two integers are said to be *relatively prime* if their greatest common divisor is 1. If n is a positive integer, let $\varphi(n)$ be the number of positive integers $m \leq n$ that are relatively prime to n . This function, called the Euler φ function or the Euler totient function, is important in number theory. We can derive a formula for this function by using the principle of inclusion and exclusion.

We must name a set and list a number of properties such that the number of elements in the set satisfying none of the properties is $\varphi(n)$. Suppose n is divisible by precisely r different primes, which we label p_1 through p_r . Select $\{1, 2, \dots, n\}$ as the set, and let a_i be the property “is divisible by p_i .” Then $N_0 = \varphi(n)$, and it is easy to compute the terms on the right side of the equation in Theorem 2.6: $N = n$, $N(a_i) = n/p_i$, $N(a_i a_j) = n/(p_i p_j)$, and so on. Therefore,

$$\begin{aligned} \varphi(n) &= n - \sum_i \frac{n}{p_i} + \sum_{i < j} \frac{n}{p_i p_j} - \sum_{i < j < k} \frac{n}{p_i p_j p_k} + \dots + (-1)^r \frac{n}{p_1 p_2 \dots p_r} \\ &= n \prod_{i=1}^r \left(1 - \frac{1}{p_i}\right). \end{aligned}$$

Exercise 7 asks you to verify the last step.

For example, the primes dividing 24 are 2 and 3, so $\varphi(24) = 24(1 - \frac{1}{2})(1 - \frac{1}{3}) = 8$. The eight numbers between 1 and 24 that are relatively prime to 24 are 1, 5, 7, 11, 13, 17, 19, and 23.

Counting Prime Numbers

Suppose m is a composite positive integer, so m can be written as a product of two integers that are both greater than 1: $m = ab$ with $1 < a \leq b$. Then $a^2 \leq m$, so $a \leq \sqrt{m}$, and so m must be divisible by a prime number p with $p \leq \sqrt{m}$.

We can use this observation, together with Theorem 2.6, to count the prime numbers between 1 and a given positive integer n . We start with the set of integers $\{1, 2, \dots, n\}$, and use the theorem to count the number of elements that remain when multiples of prime numbers $p \leq \sqrt{n}$ are excluded from the set. Since every composite number $m \leq n$ has a prime factor $p \leq \sqrt{m}$, excluding all of these numbers removes all the composite numbers from the set.

For example, for $n = 120$, the largest prime less than or equal to \sqrt{n} is the fourth prime number, 7, so we require just four properties in Theorem 2.6 to exclude all the composite numbers in the set $\{1, 2, \dots, 120\}$. The four properties are $a_1 =$ “is even,” $a_2 =$ “is divisible by 3,” $a_3 =$ “is divisible by 5,” and $a_4 =$ “is divisible by 7.” We compute $N(a_1) = 120/2 = 60$, $N(a_2) = 120/3 = 40$, $N(a_3) = 120/5 = 24$, and $N(a_4) = \lfloor 120/7 \rfloor = 17$. (The quantity $\lfloor x \rfloor$ was defined on page 153.)

Continuing our calculation, we compute $N(a_1 a_2) = \lfloor 120/6 \rfloor = 20$, then $N(a_1 a_3) = \lfloor 120/10 \rfloor = 12$, etc., and find that $N_0 = 120 - (60 + 40 + 24 + 17) + (20 + 12 + 8 + 8 + 5 + 3) - (4 + 2 + 1 + 1) + 0 = 27$. But this is not the number of prime numbers between 1 and 120, for our method excludes the primes 2, 3, 5, and 7, and includes the nonprime 1. Accounting for these exceptions, we find that the number of primes between 1 and 120 is $N_0 + 4 - 1 = 30$.

Chromatic Polynomials

Let G be a graph. Recall that its chromatic polynomial $c_G(x)$ measures the number of ways to color the vertices of G using at most x colors in such a way that no two vertices connected by an edge have the same color. We can use Theorem 2.6 to compute chromatic polynomials.

Suppose G has n vertices, and consider the set of colorings of the vertices of G using at most x colors, so the number of colorings in this set is $N = x^n$. To find $c_G(x)$, we must exclude all of the inadmissible colorings from this set. For each edge e_i in the graph, select property a_i to be “edge e_i connects two vertices that have the same color.” In this way, the colorings in the set that satisfy none of the properties are precisely the admissible colorings, so $N_0 = c_G(x)$.

For example, we compute the chromatic polynomial for the complete graph K_3 using this strategy. This graph has three edges, so we take $r = 3$ in the theorem. We compute $N(a_1) = N(a_2) = N(a_3) = x^2$, since every coloring satisfying one of the properties has two vertices with the same color, and the third vertex may be any color. Also, $N(a_1 a_2) = N(a_2 a_3) = N(a_1 a_3) = N(a_1 a_2 a_3) = x$, as

every coloring satisfying more than one property must have all vertices colored identically. Thus, $c_{K_3}(x) = N_0 = x^3 - 3x^2 + 3x - x = x(x-1)(x-2) = x^{\underline{3}}$.

Derangements

Suppose a lazy professor gives a quiz to a class of n students, then collects the papers, shuffles them, and redistributes them randomly to the class for grading. The professor would prefer that no student receives his or her own paper to grade. What is the probability that this occurs? Is this probability substantially different for different class sizes? What do you think the limiting probability is as $n \rightarrow \infty$? Notice that as n grows larger, there are more ways for at least one person to receive his or her own quiz back, but perhaps this increase is swamped by the growth of the total number of permutations possible.

Suppose we have n objects in an initial configuration. A permutation of these objects in which the position of each object differs from its initial position is called a *derangement* of the objects. Since $n!$ denotes the number of permutations of n objects, following [133] we denote the number of derangements of n objects by n_i (and since $n!$ is often pronounced “ n bang,” perhaps n_i should be pronounced “ n gnab”).

We compute n_i for some small values of n . For $n = 0$, there is just one permutation, and it vacuously satisfies the derangement condition, so $0_i = 1$. There is only one permutation of a single object, and it is not a derangement, so $1_i = 0$. Only one of the two permutations of two objects is a derangement, so $2_i = 1$, and exactly two of the six permutations of three objects satisfies the condition: If our original arrangement is $[1, 2, 3]$, then the derangements are $[2, 3, 1]$ and $[3, 1, 2]$. Thus $3_i = 2$. We find that $4_i = 9$: The derangements of $[1, 2, 3, 4]$ are $[2, 1, 4, 3]$, $[2, 3, 4, 1]$, $[2, 4, 1, 3]$, $[3, 1, 4, 2]$, $[3, 4, 1, 2]$, $[3, 4, 2, 1]$, $[4, 1, 2, 3]$, $[4, 3, 1, 2]$, and $[4, 3, 2, 1]$. Thus, the probability that a random permutation of a fixed number n of objects is a derangement is respectively 1, 0, $\frac{1}{2}$, $\frac{1}{3}$, and $\frac{3}{8}$ for $n = 0$ through 4.

We can use Theorem 2.6 to determine a formula for n_i . We select the original set to be the collection of all permutations of n objects, and for $1 \leq i \leq n$ let a_i denote the property that element i remains in its original position in a permutation. Then N_0 is the number of permutations where no elements remain in their original position, so $N_0 = n_i$.

To compute $N(a_i)$, we see that element i is fixed, but the other $n - 1$ elements may be arranged arbitrarily, so $N(a_i) = (n - 1)!$. Similarly, $N(a_i a_j) = (n - 2)!$ for $i < j$, $N(a_i a_j a_k) = (n - 3)!$ for $i < j < k$, and so on. Therefore,

$$\begin{aligned} n_i &= n! - \sum_i (n - 1)! + \sum_{i < j} (n - 2)! - \cdots \\ &\quad + (-1)^m \sum_{i_1 < \cdots < i_m} (n - m)! + \cdots + (-1)^n. \end{aligned}$$

Since the number of different m -tuples (i_1, i_2, \dots, i_m) with $1 \leq i_1 < i_2 < \dots < i_m \leq n$ is $\binom{n}{m}$, we obtain

$$\begin{aligned} n_i &= n! - \binom{n}{1}(n-1)! + \binom{n}{2}(n-2)! - \dots \\ &\quad + (-1)^m \binom{n}{m}(n-m)! + \dots + (-1)^n \\ &= \sum_m (-1)^m \binom{n}{m} (n-m)! \\ &= n! \sum_{m=0}^n \frac{(-1)^m}{m!}. \end{aligned}$$

Thus, the probability that a permutation of n objects is a derangement is

$$\frac{n_i}{n!} = \sum_{m=0}^n \frac{(-1)^m}{m!},$$

and in the limit,

$$\lim_{n \rightarrow \infty} \frac{n_i}{n!} = \sum_{m \geq 0} \frac{(-1)^m}{m!} = e^{-1}. \quad (2.27)$$

Our lazy professor obtains a desired configuration about 36.8% of the time, for sizable classes.

Exercises

1. A noted vexillologist tells you that 30 of the 50 U.S. state flags have blue as a background color, twelve have stripes, 26 exhibit a plant or animal, nine have both blue in the background and stripes, 23 have both blue in the background and feature a plant or animal, and three have both stripes and a plant or animal. One of the flags in this last category (California) does not have any blue in the background. How many state flags have no blue in the background, no stripes, and no plant or animal featured?
2. Suppose 50 socks lie in a drawer. Each one is either white or black, ankle-high or knee-high, and either has a hole or doesn't. 22 socks are white, four of these have a hole, and one of these four is knee-high. Ten white socks are knee-high, ten black socks are knee-high, and five knee-high socks have a hole. Exactly three ankle-high socks have a hole.
 - (a) Use Theorem 2.6 to determine the number of black, ankle-high socks with no holes.
 - (b) Draw a Venn diagram that shows the number of socks with each combination of characteristics.

3. The buffet line at a local steakhouse has 35 dishes. Sixteen dishes contain meat, fourteen dishes are fried, and of the dishes with meat, eight contain vegetables and seven are fried. Of the fried dishes, five contain a vegetable. Just two dishes are fried and contain both meat and a vegetable, and ten dishes (principally in the dessert section) contain neither meat nor a vegetable and are not fried. Use Theorem 2.6 to determine how many dishes contain vegetables.
4. A sneaky registrar reports the following information about a group of 400 students. There are 180 taking a math class, 200 taking an English class, 160 taking a biology class, and 250 in a foreign language class. 80 are enrolled in both math and English, 90 in math and biology, 120 in math and a foreign language, 70 in English and biology, 140 in English and a foreign language, and 60 in biology and a foreign language. Also, there are 25 in math, English, and a foreign language, 30 in math, English, and biology, 40 in math, biology, and a foreign language, and fifteen in English, biology, and a foreign language. Finally, the sum of the number of students with a course in all four subjects, plus the number of students with a course in none of the four subjects, is 100. Use Theorem 2.6 to determine the number of students that are enrolled in all four subjects simultaneously: math, biology, English, and a foreign language.
5. On a busy evening a number of guests visit a gourmet restaurant, and everyone orders something. 140 guests order a beverage, 190 order an entree, 100 order an appetizer, 90 order a dessert, 65 order a beverage and an appetizer, 125 order a beverage and an entree, 60 order a beverage and a dessert, 85 order an entree and an appetizer, 75 order an entree and a dessert, 60 order an appetizer and a dessert, 40 order a beverage, appetizer, and dessert, 55 order a beverage, entree, and dessert, 45 order an appetizer, entree, and dessert, 35 order a beverage, entree, and appetizer, and ten order all four types of items. Use Theorem 2.6 to determine the number of guests who visited the restaurant that evening.
6. Use Theorem 2.6 to determine the number of five-card hands drawn from a standard deck that contain at least one card from each of the four suits.
7. Let $\alpha_1, \alpha_2, \dots, \alpha_r$ be real numbers. Show that

$$\prod_{i=1}^r (1 - \alpha_i) = 1 - \sum_i \alpha_i + \sum_{i < j} \alpha_i \alpha_j - \sum_{i < j < k} \alpha_i \alpha_j \alpha_k + \cdots + (-1)^r \alpha_1 \alpha_2 \cdots \alpha_r.$$

8. (a) Show that $\varphi(mn) = \varphi(m)\varphi(n)$ if m and n are relatively prime.
 (b) Show that $\varphi(mn) \neq \varphi(m)\varphi(n)$ if m and n are not relatively prime. Is one quantity always larger than the other in this case?

- (c) Determine all integers n satisfying $\varphi(n) = 12, 13, \text{ or } 14$.
9. Use Theorem 2.6 to count the number of prime numbers less than 168.
10. Use Theorem 2.6 to determine the chromatic polynomial for each of the following graphs.
- The yield sign (add a single edge to the bipartite graph $K_{1,3}$).
 - The bipartite graph $K_{2,3}$.
11. Is the probability that a permutation of n objects is a derangement substantially different for $n = 12$ and $n = 120$? Quantify your answer.
12. (Deranged twins.) Suppose $n + 2$ people are seated behind a long table facing an audience to staff a panel discussion. Two of the people are identical twins, wearing identical clothing. At intermission, the panelists decide to rearrange themselves so that it will be apparent to the audience that everyone has moved to a different seat when the panel reconvenes. Each twin can therefore take neither her own former place, nor her twin's. Let T_n denote the number of different ways to derange the panel in this way.
- Compute $T_0, T_1, T_2, \text{ and } T_3$.
 - Compute T_4 .
 - Determine a formula for T_n , and check that your formula produces $T_{10} = 72\,755\,370$.
 - Compute the value of $\lim_{n \rightarrow \infty} \frac{T_n}{(n+2)!}$.
13. Suppose our lazy professor collects a quiz and a homework assignment from a class of n students one day, then distributes both the quizzes and the homework assignments back to the class in a random fashion for grading. Each student receives one quiz and one homework assignment to grade.
- What is the probability that every student receives someone else's quiz to grade, and someone else's homework to grade?
 - What is the probability that no student receives both their own quiz and their own homework assignment to grade? In this case, some students may receive their own quiz, and others may receive their own homework assignment.
 - Compute the limiting probability as $n \rightarrow \infty$ in each case.
14. Let N_m denote the number of objects from a collection of N objects that possess exactly m of the properties a_1, a_2, \dots, a_r . Generalize the principle of inclusion and exclusion by showing that

$$N_m = \sum_{k=m}^r (-1)^{k-m} \binom{k}{m} s_k, \quad (2.28)$$

where

$$s_k = \sum_{i_1 < \dots < i_k} N(a_{i_1} \dots a_{i_k}). \quad (2.29)$$

2.6 Generating Functions

*And own no other function: each your doing,
So singular in each particular,
Crowns what you are doing in the present deed,
That all your acts are queens.*

— William Shakespeare, *The Winter's Tale*, Act IV, Scene IV

Given a sequence $\{a_k\}$ with $k \geq 0$, its *generating function* $G(x)$ is defined by

$$G(x) = \sum_{k \geq 0} a_k x^k. \quad (2.30)$$

Thus, $G(x)$ is a polynomial if $\{a_k\}$ is a finite sequence, and a power series if $\{a_k\}$ is infinite. For example, if $a_k = (-1)^k/k!$, then $G(x)$ is the Maclaurin series for e^{-x} , and if $a_k = \binom{n}{k}$ for a fixed nonnegative integer n , then $G(x) = (1+x)^n$, by the binomial theorem.

To illustrate how generating functions can be used to solve combinatorial problems, let us consider again the problem of determining the number of k -element subsets of an n -element set. Fix n , and let a_k denote this number. Of course, we showed in Section 2.1 that $a_k = n^{\underline{k}}/k!$; here we derive this formula again using generating functions.

Suppose we wish to enumerate all subsets of the n -element set. To construct one subset, we must pick which elements to include in the subset and which to exclude. Let us denote the choice to omit an element by x^0 , and the choice to include it by x^1 . Using “+” to represent “or,” the choice to include or exclude one element then is denoted by $x^0 + x^1$. We must make n such choices to construct a subset, so using multiplication to denote “and,” the expression $(x^0 + x^1)^n$ models the choices required to make a subset.

Since “and” distributes over “or” just as multiplication distributes over addition, we may expand this expression using standard rules of arithmetic to obtain representations for all 2^n subsets. For example, when $n = 3$, we obtain

$$\begin{aligned} (x^0 + x^1)^3 &= x^0 x^0 x^0 + x^0 x^0 x^1 + x^0 x^1 x^0 + x^0 x^1 x^1 \\ &\quad + x^1 x^0 x^0 + x^1 x^0 x^1 + x^1 x^1 x^0 + x^1 x^1 x^1. \end{aligned}$$

The first term represents the empty subset, the second signifies the subset containing just the third item in the original set, etc. Writing 1 for x^0 and x for x^1 and treating the expression as a polynomial, we find that $(1+x)^3 = 1+3x+3x^2+x^3$, and the coefficient of x^k is the number of subsets of a three-element set having exactly k items.

In general, we find that the generating function for the sequence $\{a_k\}$ is $(1+x)^n$, so $a_k = \binom{n}{k} = n^{\underline{k}}/k!$, by the binomial theorem. Since our proof of the binomial theorem relies only on basic facts of arithmetic, this argument provides an independent derivation for the number of k -element subsets of a set with n elements.

This example illustrates the general strategy for using generating functions to solve combinatorial problems. First, express the problem in terms of determining one or more values of an unknown sequence $\{a_k\}$. Second, determine a generating function for this sequence, writing the monomial x^k to represent selecting an object k times, then using addition to represent alternative choices and multiplication to represent sequential choices. Third, use analytic methods to expand the generating function and determine the values of the encoded sequence.

For example, suppose a drawer contains twelve beads: three red, four blue, and five green, and suppose we wish to determine the number of ways to select six beads from a drawer, if beads of the same color are indistinguishable, and the order of selection is irrelevant. Let a_k denote the number of ways to select k beads from the drawer. Then the generating function for this sequence is

$$\begin{aligned} G(x) &= (1+x+x^2+x^3)(1+x+x^2+x^3+x^4) \\ &\quad \cdot (1+x+x^2+x^3+x^4+x^5) \\ &= 1+3x+6x^2+10x^3+14x^4+17x^5+18x^6 \\ &\quad +17x^7+14x^8+10x^9+6x^{10}+3x^{11}+x^{12}. \end{aligned}$$

For example, we see from this that there are exactly $a_6 = 18$ ways to select six beads from the drawer. Indeed, we can check this by constructing all such selections:

$$\begin{aligned} rrrggg, \quad rrrggb, \quad rrrgbb, \quad rrrbbb, \quad rrgggg, \quad rrgggb, \\ rrggbb, \quad rrgbbb, \quad rrbbbb, \quad rggggb, \quad rgggbb, \quad rggbbb, \\ rgbbbb, \quad rbbbb, \quad ggggbb, \quad gggbbb, \quad ggbbbb, \quad gbbbbb. \end{aligned} \quad (2.31)$$

In the following sections, we explore the power of this method by studying several combinatorial problems.

Exercises

- In this problem, we verify that the arithmetic operations performed in generating functions model the logical selections made in combinatorial problems. Write a^k to denote selecting k copies of object a .
 - Clearly, there are exactly four different subsets of the set $\{a, b\}$. We can model the construction of the different possible subsets of this two-element set by considering two choices: Pick a or not, and then pick b or not. Thus, we can denote all the possible choices by writing: $(a^0$ or $a^1)$ and $(b^0$ or $b^1)$. Expand this expression using the logical rule “ $(P$ or $Q)$ and $R \equiv (P$ and $R)$ or $(Q$ and $R)$ ”. Continue expanding

until you obtain an expression of the form “ C_1 or C_2 or C_3 or C_4 ,” where each C_i is a logical expression involving only and’s.

- (b) Rewrite exactly the same logical computation, but now use x^0 in place of a^0 or b^0 , x^1 in place of a^1 or b^1 , $+$ instead of “or”, and $*$ instead of “and”. Then simplify the expression by combining exponents in the usual way.
- (c) Repeat this procedure for the three-element set $\{a, b, c\}$.
- (d) Repeat this procedure for sub-multisets of the three-element multiset $\{a, a, b\}$.
2. Suppose a drawer contains three red beads, four blue beads, and five green beads. Use a generating function to determine the number of ways to select six beads if one must select at least one red bead, an odd number of blue beads, and an even number of green beads. Then check your answer using the combinations shown in (2.31). Assume that beads of the same color are indistinguishable, and that the order of selection is irrelevant.
3. Suppose a drawer contains ten red beads, eight blue beads, and eleven green beads. Determine a generating function that encodes the answer to each of the following problems.
- (a) The number of ways to select k beads from the drawer.
- (b) The number of ways to select k beads if one must obtain an even number of red beads, an odd number of blue beads, and a prime number of green beads.
- (c) The number of ways to select k beads if one must obtain exactly two red beads, at least five blue beads, and at most four green beads.

2.6.1 Double Decks

I don't like the games you play, professor.

— Roger Thornhill, in *North by Northwest*

How many five-card poker hands can be dealt from a double deck? Assume that the two decks are identical. More generally, how many ways are there to select m items from n different items, where each item can be selected at most twice? Let us denote this number by $t_{n,m}$, and let $G_n(x)$ be the generating function for the sequence $\{t_{n,m}\}$ for $m \geq 0$ and n fixed.

We find that $G_n(x) = (1 + x + x^2)^n$, since each object may be selected zero times, one time, or two times. To find $t_{n,m}$, we must determine a formula for the coefficient of x^m in $G_n(x)$. This is simply a matter of applying the binomial

theorem twice:

$$\begin{aligned}
 G_n(x) &= (1 + (x + x^2))^n \\
 &= \sum_k \binom{n}{k} (x + x^2)^k \\
 &= \sum_k \binom{n}{k} x^k \sum_j \binom{k}{j} x^j \\
 &= \sum_k \sum_j \binom{n}{k} \binom{k}{j} x^{j+k} \\
 &= \sum_m \left(\sum_j \binom{n}{m-j} \binom{m-j}{j} \right) x^m,
 \end{aligned}$$

where we obtained the last line by substituting m for $j + k$. Therefore,

$$t_{n,m} = \sum_{j=0}^{\lfloor m/2 \rfloor} \binom{n}{m-j} \binom{m-j}{j}. \quad (2.32)$$

The number of five-card poker hands that can be dealt from a double deck is then

$$\binom{52}{5} \binom{5}{0} + \binom{52}{4} \binom{4}{1} + \binom{52}{3} \binom{3}{2} = 3\,748\,160.$$

There is a simple combinatorial explanation for this expression. A five-card hand dealt from a double deck may have zero, one, or two cards repeated. There are $\binom{52}{5}$ hands with no cards repeated, $\binom{52}{4} \binom{4}{1}$ hands with exactly one card repeated, and $\binom{52}{3} \binom{3}{2}$ hands with exactly two cards repeated. A similar explanation applies for the general formula (2.32).

Exercises

1. Derive (2.32) by using the multinomial theorem to expand $G_n(x)$.
2. Use a combinatorial argument to count the number of different five-card hands that can be dealt from a triple deck, then the number of five-card hands that can be dealt from a quadruple deck.
3. Use a combinatorial argument to count the number of different six-card hands that can be dealt from r combined decks, for each positive integer r .
4. Use a generating function to determine the number of ways to select a hand of m cards from a triple deck, if there are n distinct cards in a single deck. Verify that your expression produces the correct answers when $n = 52$ and $m = 5$ or $m = 6$.

2.6.2 Counting with Repetition

Then, shalt thou count to three, no more, no less. Three shalt be the number thou shalt count, and the number of the counting shall be three. Four shalt thou not count, nor either count thou two, excepting that thou then proceed to three. Five is right out.

— Monty Python and the Holy Grail

Suppose there is an inexhaustible supply of each of n different objects. How many ways are there to select m objects from the n different objects, if you are allowed to select each object as many times as you like?

Let $a_{n,m}$ denote this number. Evidently, for fixed n , the generating function for $\{a_{n,m}\}_{m \geq 0}$ is

$$G_n(x) = (1 + x + x^2 + \cdots)^n = \left(\frac{1}{1-x}\right)^n,$$

since the sum is just a geometric series in x . This raises questions on convergence, for this formula is valid only for $|x| < 1$. We largely ignore these analytic issues, since we treat generating functions as formal series.

Thus, to find a formula for $a_{n,m}$, we must find the coefficient of x^m in $G_n(x)$.

Let us consider a more general problem. Let $f(x) = (1+x)^\alpha$, where α is a real number. Then $f'(0) = \alpha$, $f''(0) = \alpha(\alpha-1)$, and in general, $f^{(k)}(0) = \alpha^{\underline{k}}$. Therefore, the Maclaurin series for $f(x)$ is

$$(1+x)^\alpha = \sum_{k \geq 0} \frac{\alpha^{\underline{k}}}{k!} x^k.$$

Define the *generalized binomial coefficient* by

$$\binom{\alpha}{k} = \begin{cases} \alpha^{\underline{k}}/k! & \text{if } k \geq 0, \\ 0 & \text{if } k < 0. \end{cases} \quad (2.33)$$

Note that $\binom{\alpha}{k}$ equals the ordinary binomial coefficient whenever α is a nonnegative integer. We have the following theorem.

Theorem 2.7 (Generalized Binomial Theorem). *If $|x| < 1$ or α is a nonnegative integer, then*

$$(1+x)^\alpha = \sum_k \binom{\alpha}{k} x^k. \quad (2.34)$$

The proof of convergence may be found in many analysis texts, where it is often proved as a consequence of Bernstein's theorem on convergence of Taylor series (see for instance [11]). We do not supply the proof here.

Before solving our problem concerning selection with unlimited repetition, we note a useful identity for generalized binomial coefficients.

Negating the Upper Index. If α is a real number and k is an integer, then

$$\binom{\alpha}{k} = (-1)^k \binom{k - \alpha - 1}{k}. \quad (2.35)$$

Proof. For $k < 0$, the identity is clear. For $k \geq 0$, we have

$$\binom{\alpha}{k} = \frac{1}{k!} \prod_{i=0}^{k-1} (\alpha - i).$$

Reindex this product, replacing each i by $k - 1 - i$, to obtain

$$\begin{aligned} \binom{\alpha}{k} &= \frac{1}{k!} \prod_{i=0}^{k-1} (\alpha - (k - i - 1)) \\ &= \frac{(-1)^k}{k!} \prod_{i=0}^{k-1} (k - 1 - i - \alpha) \\ &= (-1)^k \binom{k - \alpha - 1}{k}. \quad \square \end{aligned}$$

We may now solve our problem of determining $a_{n,m}$. We compute

$$\begin{aligned} G_n(x) &= (1 - x)^{-n} \\ &= \sum_m \binom{-n}{m} (-x)^m \\ &= \sum_m \binom{n + m - 1}{m} x^m, \end{aligned}$$

and therefore the number of ways to select m objects from a collection of n different objects, with repetition allowed, is

$$a_{n,m} = \binom{n + m - 1}{m}. \quad (2.36)$$

For example, the number of five-card poker hands that can be dealt from a stack of five or more decks is $\binom{56}{5} = 3\,819\,816$.

Finally, suppose we lay all 52 cards of a standard deck face up on a table. How many ways can we place five identical poker chips on the cards if we allow more than one chip to be placed on each card? To solve this, notice that each possible placement of chips corresponds to a hand of five cards, where repeated cards are allowed: If k chips lie on a particular card, place that card into the hand k times. Further, every such five-card hand can be represented by a judicious placement of chips. Therefore, the answer is the same as that of the previous example, $\binom{56}{5}$.

In general, the number of ways to place m identical objects into n distinguishable bins is the same as the number of ways to select m objects from a set of n objects with repetition allowed: The answer to both problems is $\binom{n+m-1}{m}$.

Exercises

1. Prove the addition identity for generalized binomial coefficients: If α is a real number and k is an integer, then

$$\binom{\alpha}{k} = \binom{\alpha-1}{k} + \binom{\alpha-1}{k-1}.$$

2. Prove the absorption/extraction identity for generalized binomial coefficients: If α is a real number and k is a nonzero integer, then

$$\binom{\alpha}{k} = \frac{\alpha}{k} \binom{\alpha-1}{k-1}.$$

3. Prove the cancellation identity for generalized binomial coefficients: If α is a real number and k and m are integers, then

$$\binom{\alpha}{k} \binom{k}{m} = \binom{\alpha}{m} \binom{\alpha-m}{k-m}.$$

4. Prove the parallel summation identity for generalized binomial coefficients: If α is a real number and n is an integer, then

$$\sum_{k=0}^n \binom{\alpha+k}{k} = \binom{\alpha+n+1}{n}.$$

5. Suppose that an unlimited number of jelly beans is available in each of five different colors: red, green, yellow, white, and black.

- How many ways are there to select twenty jelly beans?
 - How many ways are there to select twenty jelly beans if we must select at least two jelly beans of each color?
6. A catering company brings fifty identical hamburgers to a party with twenty guests.
- How many ways can the hamburgers be divided among the guests, if none is left over?
 - How many ways can the hamburgers be divided among the guests, if every guest receives at least one hamburger, and none is left over?
 - Repeat these problems if there may be burgers left over.

7. A zodiac sign is one of twelve constellations that the sun travels through (from the vantage point of the earth) over the course of a year. Each person has a zodiac sign based on the position of sun on their birth date. The *astrological configuration* of a party with n guests is a list of twelve numbers that records the number of guests with each sign, so the first number records the number of people with the sign Capricorn, the second, Aquarius, ..., the last, Sagittarius.

- (a) How many different astrological configurations are possible for $n = 100$?
- (b) How many astrological configurations are possible for $n = 100$, if each component is at least 5?
8. Two lottery systems are proposed for a new state lottery. In the first system, players select six *different* numbers from $\{1, 2, \dots, 50\}$. In the second system, players select six numbers from $\{1, 2, \dots, 45\}$, *and* may select any number as many times as they want. (In the second system, each ball selected in the lottery drawing is replaced before another ball is selected.) Which system has more possible tickets?
9. Suppose 100 identical tickets for rides are distributed among 40 children at a carnival.
- (a) How many ways can the tickets be distributed, if each child receives at least two tickets, and all the tickets are distributed?
- (b) How many ways can the tickets be distributed, if each child receives at least one ticket, and some tickets may be left over?
- (c) Suppose one child has twelve tickets, and each ticket may be used on any of six different rides. How many ways can the child spend her tickets, if she can choose any ride any number of times, and the order of choice is unimportant?

2.6.3 Changing Money

Jesus went into the temple, and began to cast out them that sold and bought in the temple, and overthrew the tables of the moneychangers

...

— *Mark 11:15*

We now turn to a problem popularized by the analyst and combinatorialist George Pólya [225]: How many ways are there to change a dollar? That is, how many combinations of pennies, nickels, dimes, quarters, half-dollars, and dollar coins total \$1? Our discussion of this problem follows the treatment of Graham, Knuth, and Patashnik [133].

Let us define a_k to be the number of ways to make k cents in change, and let $A(x)$ be a generating function for a_k : $A(x) = \sum_k a_k x^k$. Before analyzing this problem, pause a moment and make a guess. Do you think a_{50} is more than 50 or less than 50? Is a_{100} more than 100 or less than 100? How fast do you think a_k grows as a function of k ? Is it a polynomial in k ? Exponential in k ? Perhaps something between these?

To create a pile of change, we must make six choices, selecting a number of pennies, then nickels, then dimes, quarters, half-dollars, and dollars. We can

model our choice of pennies by the sum

$$1 + x + x^2 + x^3 + \cdots = \frac{1}{1-x}.$$

One might be tempted to use the same expression to model the different choices for each of the other coins, since we can pick any number of nickels, and any number of dimes, etc., but this would be incorrect. This would yield a generating function for the number of ways to select k coins from a set of six different coins, not the number of ways to form k cents. Instead, when choosing nickels, we select either zero cents, or five cents, or ten cents, and so on, so this selection is modeled as

$$1 + x^5 + x^{10} + x^{15} + \cdots = \frac{1}{1-x^5}.$$

Therefore, the number of ways to make k cents using either pennies or nickels is given by the generating function

$$\frac{1}{(1-x)(1-x^5)}.$$

Continuing in this way, we find that

$$A(x) = \frac{1}{(1-x)(1-x^5)(1-x^{10})(1-x^{25})(1-x^{50})(1-x^{100})}, \quad (2.37)$$

so we merely need to find the coefficient of a_k in the Maclaurin series for $A(x)$! This sounds rather daunting, so let us determine a few values of a_k by hand first.

Let $\{p_k\}$ denote the number of ways to make k cents using only pennies, so $p_k = 1$ for all k . Let $P(x)$ be the generating function for $\{p_k\}$, so $P(x) = 1/(1-x)$. Let n_k be the number of ways to make k cents using either pennies or nickels, so its generating function is

$$N(x) = \frac{P(x)}{1-x^5}.$$

Thus $N(x) = P(x) + x^5N(x)$, and by equating coefficients we find that

$$n_k = \begin{cases} p_k & \text{if } 0 \leq k \leq 4, \\ p_k + n_{k-5} & \text{if } k \geq 5. \end{cases}$$

In the same way, let d_k denote the number of ways to make k cents using pennies, nickels, or dimes, and let $D(x)$ be its generating function. We then have $D(x) = N(x) + x^{10}D(x)$, and so

$$d_k = \begin{cases} n_k & \text{if } 0 \leq k \leq 9, \\ n_k + d_{k-10} & \text{if } k \geq 10. \end{cases}$$

There is a simple combinatorial interpretation for this equation. If $k < 10$, then we can choose only nickels and pennies to form k cents, so $d_k = n_k$ in this case. If $k \geq 10$, we may form k cents using only nickels and pennies, or we can choose one dime, then form the remaining $k - 10$ cents using dimes, nickels, and pennies. Thus $d_k = n_k + d_{k-10}$ in this case.

Similarly, using q_k for allowing quarters, h_k for half dollars, and finally a_k for dollar coins, we have

$$\begin{aligned}
 q_k &= \begin{cases} d_k & \text{if } 0 \leq k \leq 24, \\ d_k + q_{k-25} & \text{if } k \geq 25; \end{cases} \\
 h_k &= \begin{cases} q_k & \text{if } 0 \leq k \leq 49, \\ q_k + h_{k-50} & \text{if } k \geq 50; \end{cases} \\
 a_k &= \begin{cases} h_k & \text{if } 0 \leq k \leq 99, \\ h_k + a_{k-100} & \text{if } k \geq 100. \end{cases}
 \end{aligned}$$

We may use these formulas to construct Table 2.2 below showing the number of ways to make k cents with the different coin sets.

k	0	5	10	15	20	25	30	35	40	45	50
p_k	1	1	1	1	1	1	1	1	1	1	1
n_k	1	2	3	4	5	6	7	8	9	10	11
d_k	1	2	4	6	9	12	16	20	25	30	36
q_k	1					13					49
h_k	1										50
a_k	1										
k		55	60	65	70	75	80	85	90	95	100
p_k		1	1	1	1	1	1	1	1	1	1
n_k		12	13	14	15	16	17	18	19	20	21
d_k		42	49	56	64	72	81		100		121
q_k						121					242
h_k											292
a_k											293

TABLE 2.2. Computing the number of ways to make k cents in change.

We find that there are precisely 50 ways to make 50 cents in change, and 293 ways to make one dollar in change.

This is a fairly efficient method to determine a_k , since apparently we can calculate this number using at most $5k$ arithmetic operations. But we can do much better! We can compute a_k using at most a constant number of arithmetic operations, regardless of the value of k . To show this, let us first simplify $A(x)$ by

exploiting the fact that all but one of the exponents in (2.37) is a multiple of 5. Let

$$B(x) = \frac{1}{(1-x)^2(1-x^2)(1-x^5)(1-x^{10})(1-x^{20})},$$

so that

$$A(x) = (1+x+x^2+x^3+x^4)B(x^5).$$

Writing b_k for the coefficient of x^k in the Maclaurin series for $B(x)$ and equating coefficients, we find that

$$b_k = a_{5k} = a_{5k+1} = a_{5k+2} = a_{5k+3} = a_{5k+4}.$$

But this makes sense (or perhaps cents?), since the last few cents can be represented only using pennies. Now

$$B(x) = \frac{C(x)}{(1-x^{20})^6},$$

where

$$\begin{aligned} C(x) &= (1+x+\cdots+x^{19})^2(1+x^2+\cdots+x^{18})(1+x^5+x^{10}+x^{15}) \\ &\quad \cdot (1+x^{10}) \\ &= x^{81} + 2x^{80} + 4x^{79} + 6x^{78} + 9x^{77} + 13x^{76} + 18x^{75} + 24x^{74} + 31x^{73} \\ &\quad + 39x^{72} + 50x^{71} + 62x^{70} + 77x^{69} + 93x^{68} + 112x^{67} + 134x^{66} \\ &\quad + 159x^{65} + 187x^{64} + 218x^{63} + 252x^{62} + 287x^{61} + 325x^{60} + 364x^{59} \\ &\quad + 406x^{58} + 449x^{57} + 493x^{56} + 538x^{55} + 584x^{54} + 631x^{53} + 679x^{52} \\ &\quad + 722x^{51} + 766x^{50} + 805x^{49} + 845x^{48} + 880x^{47} + 910x^{46} + 935x^{45} \\ &\quad + 955x^{44} + 970x^{43} + 980x^{42} + 985x^{41} + 985x^{40} + 980x^{39} + 970x^{38} \\ &\quad + 955x^{37} + 935x^{36} + 910x^{35} + 880x^{34} + 845x^{33} + 805x^{32} + 766x^{31} \\ &\quad + 722x^{30} + 679x^{29} + 631x^{28} + 584x^{27} + 538x^{26} + 493x^{25} + 449x^{24} \\ &\quad + 406x^{23} + 364x^{22} + 325x^{21} + 287x^{20} + 252x^{19} + 218x^{18} + 187x^{17} \\ &\quad + 159x^{16} + 134x^{15} + 112x^{14} + 93x^{13} + 77x^{12} + 62x^{11} + 50x^{10} \\ &\quad + 39x^9 + 31x^8 + 24x^7 + 18x^6 + 13x^5 + 9x^4 + 6x^3 + 4x^2 + 2x + 1. \end{aligned}$$

We know from the previous section that

$$\frac{1}{(1-z)^n} = \sum_k \binom{n+k-1}{n-1} z^k,$$

so

$$B(x) = C(x) \sum_k \binom{k+5}{5} x^{20k}. \quad (2.38)$$

Therefore, writing $C(x) = \sum_k c_k x^k$, we have $a_{100} = b_{20} = c_0 \binom{6}{5} + c_{20} \binom{5}{5} = 6 + 287 = 293$, and

$$\begin{aligned} a_{1000} &= b_{200} \\ &= c_0 \binom{15}{5} + c_{20} \binom{14}{5} + c_{40} \binom{13}{5} + c_{60} \binom{12}{5} + c_{80} \binom{11}{5} \\ &= 2\,103\,596. \end{aligned}$$

Our expression for computing a_k is a sum having at most five terms, so this method allows us to compute a_k using only a constant number of operations.

Finally, consider the crazy system of coinage where there is a coin minted worth n cents for every $n \geq 1$. Let p_n denote the number of ways to make n cents in change in this system. For example, $p_4 = 5$, since we can make four cents by using four pennies, or one two-cent piece and two pennies, or one three-cent piece and one penny, or two two-cent pieces, or one four-cent piece. By representing these five possibilities as the sums $1 + 1 + 1 + 1$, $2 + 1 + 1$, $3 + 1$, $2 + 2$, and 4 , we see that p_n is the number of ways to write n as a sum of one or more positive integers, disregarding the order of the summands. Such a representation is called a *partition* of n . Evidently the generating function $P(x)$ for the sequence of partitions is given by the infinite product

$$P(x) = \prod_{k \geq 1} \frac{1}{1 - x^k}. \quad (2.39)$$

We explore this generating function and the sequence $\{p_n\}$ in Section 2.8.1.

Exercises

1. Use (2.38) to compute a_{2009} , the number of ways to make \$20.09 in change.
2. How many ways are there to select 100 coins from an inexhaustible supply of pennies, nickels, dimes, quarters, half-dollars, and dollar coins?
3. Show that the number of ways to make $10m$ cents in change using only pennies, nickels, and dimes is $(m + 1)^2$.
4. Show that a_k can be computed using equation (2.38) using at most 60 arithmetic operations. Optimize your method to show that a_k can be computed using at most 31 arithmetic operations.
5. Prove that a_k grows like k^5 by showing that there exist positive constants c and C such that $ck^5 < a_k < Ck^5$ for sufficiently large k .
6. The following coins were in circulation in the United States in 1875: the Indian-head penny, a bronze two-cent piece (last minted in 1873), a silver three-cent piece (also last minted in 1873), a nickel three-cent piece, the

shield nickel (worth five cents), the seated liberty half-dime, dime, twenty-cent piece (produced for only four years beginning in 1875), quarter, half-dollar, and silver dollar, and the Indian-head gold dollar. (We ignore the trade dollar, minted for circulation between 1873 and 1878, as it was issued for overseas trade. This coin holds the distinction of being the only U.S. coin to be demonetized.)

- (a) How many ways were there to make twenty cents in change in 1875? How about twenty-five cents? Compute these values using the tabular method of this section.
 - (b) Write down a generating function in the form of a rational function for the number of ways to make k cents in change in 1875, then use a computer algebra system to find the number of ways to make one dollar in change in 1875.
7. (Inspired in part by [133, ex. 7.21].) A ransom note demands:
- (i) \$10000 in unmarked fifty- and hundred-dollar bills, and
 - (ii) the number of ways to award the cash.

You realize that both old-fashioned and redesigned anticounterfeit bills are available in both denominations.

- (a) Answer the second demand of the ransom note. For extra credit, answer the first demand ☺.
 - (b) Find a closed form for the number of ways to make $50m$ dollars using the two kinds of fifty- and hundred-dollar bills, for a nonnegative integer m .
8. In 2010, there are six different kinds of nickels in general circulation in the U.S., and six different kinds of pennies. Four of the varieties of nickels were issued in 2004 and 2005 and commemorate the bicentennial of the Lewis and Clark expedition—their respective designs on the reverse show a handshake, a boat, a bison, and an ocean view; the other two show president Jefferson's home, Monticello. Four of the pennies were issued in 2009 to commemorate the bicentennial of Lincoln's birth, with each design evoking a different period of the life of the U.S. president.
- (a) Determine a generating function in closed form for the number of ways a_k to make k cents in change using only pennies and nickels available in 2010, counting each design as a different coin.
 - (b) Determine a finite sequence c_0, c_1, \dots, c_n so that

$$a_k = \sum_{j=\lceil (k-n)/5 \rceil}^{\lfloor k/5 \rfloor} c_{k-5j} \binom{j+11}{11}.$$

- (c) Use the formula to determine a_5 , and verify that your answer is correct.
- (d) Use the formula to determine a_{23} , a_{24} , and a_{25} .
9. In 2010, there are fifty commemorative quarters in general circulation in the U.S., one for each state, and sixteen different presidential dollar coins, showing Washington through Lincoln on the obverse. Prove that the number of ways to make $25k$ cents in change using just these 66 different coins is

$$\sum_{a+b+2c=k} (-1)^{b+c} \binom{65+a}{a} \binom{15+b}{b} \binom{15+c}{c}.$$

Then use this formula to determine the number of ways to change one dollar using just these coins.

10. A hungry math major visits the school's cafeteria and wants to know the number of ways s_k to take k servings of food, including at least one main course, an even number (possibly zero) of side vegetables, an odd number of rolls, and at least two desserts. The cafeteria's food can be distinguished only in the coarsest way: Every dish is either a main course, a side vegetable, a roll, or a dessert. There is an unlimited supply of each kind of dish available.
- (a) Determine a closed form for the generating function $\sum_k s_k x^k$.
- (b) Show that

$$s_k = \binom{\lfloor \frac{k+1}{2} \rfloor}{3} + \binom{\lceil \frac{k+1}{2} \rceil}{3}.$$

The quantities $\lfloor x \rfloor$ and $\lceil x \rceil$ are defined on page 153.

2.6.4 Fibonacci Numbers

Attention! Attention! Ladies and gentlemen, attention! There is a herd of killer rabbits headed this way and we desperately need your help!

— *Night of the Lepus*

Hey, shouldn't that be a *colony* of killer rabbits?

Leonardo of Pisa, better known as Fibonacci, proposed the following hare problem in 1202. Assume that the rabbit population grows according to the following rules.

1. Every pair of adult rabbits produces a pair of baby rabbits, one of each gender, every month.
2. Baby rabbits become adult rabbits at age one month and produce their first offspring at age two months.

3. Rabbits are immortal.

Starting with a single pair of baby rabbits at the start of the first month, how many pairs of rabbits are there after k months?

Let F_k denote this number. In the first month, there is one pair of baby rabbits, so $F_1 = 1$. Likewise, $F_2 = 1$, as there is one pair of adult rabbits in the second month. In the third month, we have one baby pair and one adult pair, so $F_3 = 2$, and in the fourth month, the babies become adults and the adults produce another pair of offspring, so there is one pair of babies and two pairs of adults: $F_4 = 3$. Continuing in this way, we record the population in the following table.

k	Baby pairs	Adult pairs	F_k
0	0	0	0
1	1	0	1
2	0	1	1
3	1	1	2
4	1	2	3
5	2	3	5
6	3	5	8
7	5	8	13

Notice that the number of pairs of adults in month k equals the total number of pairs of rabbits in month $k - 1$. This is F_{k-1} . Also, the number of pairs of baby rabbits in month k equals the number of adult pairs in month $k - 1$, which is the total number of pairs in month $k - 2$. This is F_{k-2} . Therefore,

$$F_k = F_{k-1} + F_{k-2}, \quad k \geq 2. \quad (2.40)$$

This recurrence, together with the initial conditions $F_0 = 0$ and $F_1 = 1$, determines the Fibonacci sequence $\{F_k\} = \{0, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, \dots\}$. This sequence appears frequently in combinatorial problems.

In this section we determine a closed form for F_k by analyzing its generating function. We will adapt this technique to solve other recurrences later in this chapter.

Let $G(x)$ be the generating function for $\{F_k\}$. Then

$$\begin{aligned} G(x) &= \sum_{k \geq 0} F_k x^k \\ &= F_0 + F_1 x + \sum_{k \geq 2} F_k x^k \\ &= x + \sum_{k \geq 2} (F_{k-1} + F_{k-2}) x^k \\ &= x + x \sum_{k \geq 2} F_{k-1} x^{k-1} + x^2 \sum_{k \geq 2} F_{k-2} x^{k-2} \\ &= x + x \sum_{k \geq 1} F_k x^k + x^2 \sum_{k \geq 0} F_k x^k, \end{aligned}$$

and so

$$G(x) = x + xG(x) + x^2G(x).$$

Therefore,

$$G(x) = \frac{-x}{x^2 + x - 1},$$

and thus F_k is the coefficient of x^k in the Maclaurin series for this rational function. How can we determine this series without all the messy differentiation?

The trick is using partial fractions to write $G(x)$ as a sum of simpler rational functions. Write $x^2 + x - 1 = (x + \varphi)(x + \hat{\varphi})$, where φ is the golden ratio, $\varphi = (1 + \sqrt{5})/2$, and $\hat{\varphi} = (1 - \sqrt{5})/2$. Write

$$\frac{-x}{x^2 + x - 1} = \frac{A}{x + \varphi} + \frac{B}{x + \hat{\varphi}}$$

and solve to find that $A = -\varphi/\sqrt{5}$ and $B = \hat{\varphi}/\sqrt{5}$. Thus

$$\begin{aligned} G(x) &= \frac{1}{\sqrt{5}} \left(\frac{\hat{\varphi}}{x + \hat{\varphi}} - \frac{\varphi}{x + \varphi} \right) \\ &= \frac{1}{\sqrt{5}} \left(\frac{1}{1 + x/\hat{\varphi}} - \frac{1}{1 + x/\varphi} \right) \\ &= \frac{1}{\sqrt{5}} \left(\frac{1}{1 - \varphi x} - \frac{1}{1 - \hat{\varphi} x} \right), \end{aligned}$$

since $\varphi\hat{\varphi} = -1$. Now the two terms on the right are closed forms for simple geometric series, so

$$G(x) = \frac{1}{\sqrt{5}} \sum_{k \geq 0} (\varphi^k - \hat{\varphi}^k) x^k,$$

and therefore

$$F_k = \frac{\varphi^k - \hat{\varphi}^k}{\sqrt{5}}. \quad (2.41)$$

Notice that $|\hat{\varphi}| < 1$, so $F_k \sim \varphi^k/\sqrt{5}$: a large number of rabbits indeed.

Exercises

1. In each of the following problems, first compute the value the expression for a few small values of n . Then use your data to conjecture a general formula. Last, prove that your formula is correct.

(a) $\sum_{k=0}^n F_k.$

(b) $\sum_{k=0}^n F_{2k}.$

$$(c) \sum_{k=1}^n F_{2k-1}, \text{ if } n \geq 1.$$

$$(d) F_{n+1}F_{n-1} - F_n^2, \text{ if } n \geq 1.$$

2. Solve the recurrence $a_k = 2a_{k-1} + 3a_{k-2}$, if $a_0 = 0$ and $a_1 = 8$.
3. Suppose $a_0 = 0$, $a_1 = 5$, and $a_k = a_{k-1} + 6a_{k-2}$ for $k \geq 2$. Compute a closed form for the generating function of the sequence $\{a_k\}$. Then use this to determine a formula for a_k .

4. Solve the recurrence $a_k = 2a_{k-1} + 2a_{k-2}$, if $a_0 = 0$ and $a_1 = 1$.

5. Prove the following identities involving Fibonacci numbers.

$$(a) F_{m+n} = F_m F_{n+1} + F_{m-1} F_n, \text{ if } m \geq 1 \text{ and } n \geq 0.$$

$$(b) F_n^2 + F_{n+1}^2 = F_{2n+1}.$$

$$(c) F_{n+1}^2 - F_{n-1}^2 = F_{2n}, \text{ if } n \geq 1.$$

$$(d) \sum_{k=0}^n F_k^2 = F_n F_{n+1}.$$

$$(e) \sum_{k=0}^n (-1)^{n-k} F_k = F_{n-1} - (-1)^n, \text{ if } n \geq 1.$$

$$(f) \sum_{k=0}^n (-1)^{n-k} k F_k = (n+1)F_{n-1} - F_{n-2} - 2(-1)^n, \text{ if } n \geq 2.$$

6. Prove that if m and n are nonnegative integers, then F_m divides F_{mn} .

7. The *Lucas numbers* are defined by $L_0 = 2$, $L_1 = 1$, and $L_k = L_{k-1} + L_{k-2}$ for $k \geq 2$. Find a formula for L_k in terms of φ and $\hat{\varphi}$.

8. Prove the following identities involving Lucas and Fibonacci numbers.

$$(a) L_n = F_{n+1} + F_{n-1}, \text{ if } n \geq 1.$$

$$(b) \sum_{k=0}^n L_k^2 = L_n L_{n+1} + 2.$$

$$(c) \sum_{k=0}^n (-1)^k L_{n-k} = L_{n-1} + 3(-1)^n, \text{ if } n \geq 1.$$

$$(d) F_{2n} = F_n L_n.$$

$$(e) L_{2n} = L_n^2 - 2(-1)^n.$$

9. The *Perrin sequence* is defined by $a_0 = 3$, $a_1 = 0$, $a_2 = 2$, and $a_k = a_{k-2} + a_{k-3}$ for $k \geq 3$. The *Padovan sequence* is defined by $b_0 = 0$, $b_1 = 1$, $b_2 = 1$, and $b_k = b_{k-2} + b_{k-3}$ for $k \geq 3$.

- (a) Find generating functions in the form of rational functions for the Perrin sequence and the Padovan sequence.
- (b) Prove that $a_k = r^k + \alpha^k + \bar{\alpha}^k$, where r , α , and $\bar{\alpha}$ are the three complex roots of $x^3 - x - 1$. Conclude that $a_k \sim r^k$.

The Perrin sequence has an interesting property: If p is a prime number, then p divides the p th term in the Perrin sequence, $p \mid a_p$. This was first noted by Lucas in 1878 [192–194] (perhaps Lucas would have been interested in Exercise 6 of Section 2.6.3). Thus we obtain a test for composite numbers: If n does not divide a_n , then n is not prime. Unfortunately, the converse is false: There are infinitely many composite n with the property that $n \mid a_n$. This was proved by Grantham [137].

10. In the children's game of hopscotch, a player hops across an array of squares drawn on the ground, landing on only one foot whenever there is just one square at a position, and landing on both feet when there are two. If every position has either one or two squares, how many different hopscotch games have exactly n squares? Figure 2.3 shows the five different hopscotch games having four squares.




FIGURE 2.3. Hopscotch games with four squares.

11. Use a combinatorial argument and Exercise 10 to prove that

$$F_n = \sum_k \binom{n-k-1}{k}.$$

2.6.5 Recurrence Relations

O me! O life! . . . of the questions of these recurring:

— Walt Whitman, *Leaves of Grass*

In the “Tower of Hanoi” puzzle, one begins with a pyramid of k disks stacked around a center pole, with the disks arranged from largest diameter on the bottom to smallest diameter on top. There are also two empty poles that can accept disks. The object of the puzzle is to move the entire stack of disks to one of the other poles, subject to three constraints:

1. Only one disk may be moved at a time.
2. Disks can be placed only on one of the three poles.
3. A larger disk cannot be placed on a smaller one.

How many moves are required to move the entire stack of k disks onto another pole? Let a_k denote this number. Clearly, $a_1 = 1$. To move k disks, we must first move the $k - 1$ top disks to one of the other poles, then move the bottom disk to the third pole, then move the stack of $k - 1$ disks to that pole, so $a_k = 2a_{k-1} + 1$ for $k \geq 1$. Thus, $a_2 = 3$, $a_3 = 7$, $a_4 = 15$, and it appears that $a_k = 2^k - 1$.

We can certainly verify this formula by induction, but we wish to show how recurrences of this form can be solved by using generating functions. Consider the more general recurrence

$$a_k = ba_{k-1} + c, \quad k \geq 1,$$

where b and c are constants. This is a linear recurrence relation, since a_k is a linear function of the preceding values of the sequence. (The Fibonacci recurrence is also a linear recurrence relation.) If c is zero, we call the recurrence *homogeneous*; otherwise, it is *inhomogeneous*.

Let $G(x)$ be the generating function for $\{a_k\}$. Then

$$\begin{aligned} G(x) &= \sum_{k \geq 0} a_k x^k \\ &= a_0 + \sum_{k \geq 1} (ba_{k-1}x^k + cx^k) \\ &= a_0 + bx \sum_{k \geq 0} a_k x^k + cx \sum_{k \geq 0} x^k \\ &= a_0 + bxG(x) + \frac{cx}{1-x}, \end{aligned}$$

and so

$$G(x) = \frac{cx}{(1-bx)(1-x)} + \frac{a_0}{1-bx}.$$

Assuming $b \neq 1$, we compute

$$\frac{cx}{(1-bx)(1-x)} = \frac{c}{b-1} \left(\frac{1}{1-bx} - \frac{1}{1-x} \right),$$

so

$$\begin{aligned} G(x) &= \left(a_0 + \frac{c}{b-1} \right) \left(\frac{1}{1-bx} \right) - \frac{c}{b-1} \left(\frac{1}{1-x} \right) \\ &= \left(a_0 + \frac{c}{b-1} \right) \sum_{k \geq 0} b^k x^k - \frac{c}{b-1} \sum_{k \geq 0} x^k, \end{aligned}$$

and therefore

$$a_k = \left(a_0 + \frac{c}{b-1} \right) b^k - \frac{c}{b-1}. \quad (2.42)$$

For example, to find the number of moves needed to solve the Tower of Hanoi puzzle, we set $a_0 = 0$, $b = 2$, and $c = 1$ to obtain $a_k = 2^k - 1$. Also, if we set $b = -\frac{1}{2}$ and $c = 2$, we find that $a_k = (-1)^k(a_0 - \frac{4}{3})/2^k + \frac{4}{3}$, so a_k approaches $\frac{4}{3}$ as k grows large, independent of the initial value a_0 .

We conclude with a short list of useful generating functions. Since

$$\frac{1}{1-x} = \sum_{k \geq 0} x^k, \quad (2.43)$$

we differentiate both sides to find that

$$\frac{1}{(1-x)^2} = \sum_{k \geq 1} kx^{k-1},$$

and so

$$\frac{x}{(1-x)^2} = \sum_{k \geq 0} kx^k. \quad (2.44)$$

Thus we obtain a closed form for the generating function of the identity sequence $\{k\}$. We take up the problem of determining a generating function for $\{k^n\}$, for any fixed positive integer n , in Section 2.8.5.

Finally, we integrate both sides of (2.43) to obtain the generating function for $\{1/k\}$:

$$-\ln(1-x) = \sum_{k \geq 1} \frac{x^k}{k}. \quad (2.45)$$

Exercises

- Find a recurrence relation for the maximal number of regions of the plane separated by k straight lines, then solve it.
- Solve for a_k in terms of a_0 and the other parameters in each of the following recurrence relations.
 - $a_k = a_{k-1} + c$.
 - $a_k = ba_{k-1} + cb^k$.
 - $a_k = ba_{k-1} + cr^k$, assuming $b \neq r$.
 - $a_k = ba_{k-1} + cr^k + d$, assuming $b \notin \{1, r\}$.
 - $a_k = ba_{k-1} + ck$, assuming $b \neq 1$.
 - $a_k = ba_{k-1} + ck + d$, assuming $b \neq 1$.
- Find a closed form for the generating function of the sequence $\{k^2\}_{k \geq 0}$.

4. Let v_n denote the number of ways that $3n$ different people can split up into n three-person teams for a volleyball tournament, and let $v_0 = 1$. Assume that team members are unordered, so the team $\{a, b, c\}$ is the same as the team $\{c, a, b\}$, and assume that the teams are unordered, so putting $\{a, b, c\}$ on the first team and $\{d, e, f\}$ on the second is the same as putting $\{d, e, f\}$ on the first team and $\{a, b, c\}$ on the second. Determine a recurrence relation for v_n , then use it to compute v_4 .
5. Let d_k denote the minimal degree of a polynomial with $\{0, 1\}$ coefficients that is divisible by $(x + 1)^k$. For example, certainly $d_1 = 1$, since $f_1(x) = x + 1$ has the required properties, and $d_2 \leq 4$, since $f_2(x) = (x + 1)(x^3 + 1) = x^4 + x^3 + x + 1$ is permissible (in fact, $d_2 = 4$).
- (a) Determine an upper bound on d_3 by multiplying $f_2(x)$ by a suitable binomial of the form $x^r + 1$, choosing r as small as possible. Then iterate this process to obtain upper bounds for d_4 and d_5 .
- (b) Observe that one can obtain an upper bound on d_k in general by constructing a polynomial of the form

$$f_k(x) = \prod_{i=1}^k (x^{r_i} + 1)$$

for a judiciously selected sequence $\{r_i\}$. Describe how to calculate $\{r_i\}$, and compute the values of this sequence for $i \leq 7$.

- (c) Determine a linear, homogeneous recurrence relation for the sequence $\{r_i\}$.
- (d) Compute a closed formula for r_i .
- (e) Determine an upper bound for d_k .
6. A *binary sequence* is a sequence in which each term is 0 or 1. Determine a recurrence relation for the number of binary sequences of length n that do not contain two adjacent 1s, then find a simple expression for this number.
7. Let t_n denote the number of binary sequences of length n that do not contain three adjacent 1s.

- (a) Determine a recurrence relation for t_n , and enough initial values to generate the sequence.
- (b) Determine a closed form for the generating function

$$T(x) = \sum_{n \geq 0} t_n x^n.$$

- (c) Define t_n^* by $t_0^* = t_1^* = 0$, $t_2^* = 1$, and $t_n^* = t_{n-3}$ for $n \geq 3$. Determine a closed form for $T^*(x) = \sum_{n \geq 0} t_n^* x^n$. The numbers $\{t_n^*\}$ are known as the *tribonacci numbers*.

8. For a fixed positive integer m , let $s_{m,n}$ denote the number of binary sequences of length n that do not contain m adjacent 1s.
- (a) Determine a recurrence relation in n for $s_{m,n}$, and enough initial values to generate the sequence.
- (b) Show that the generating function $S_m(x)$ for $\{s_{n,m}\}_{n \geq 0}$ is

$$S_m(x) = \frac{1 - x^m}{x^{m+1} - 2x + 1}.$$

Hint: First define a sequence $s_{m,n}^*$ from $s_{m,n}$ in the same manner as Exercise 7c. Then find its generating function $S_m^*(x)$, and use this to determine $S(x)$. The numbers $\{s_{m,n}^*\}_{n \geq 0}$ are known as the *generalized Fibonacci numbers of order m* , or the *m -generalized Fibonacci numbers*.

2.6.6 Catalan Numbers

zero, un, dos, tres, quatre, cinc, sis, set, vuit, nou, deu, onze, dotze, tretze, catorze, quinze, setze, disset, divuit, dinou, vint.

How many ways are there to compute a product of $k + 1$ matrices? Matrix multiplication is associative but not commutative, so this is the number of ways to place $k - 1$ pairs of parentheses in the product $x_0 x_1 \dots x_k$ in such a way that the order of multiplications is completely specified. Let C_k denote this number.

Let us first compute a few values of C_k . There is only one way to compute the product of one or two matrices. There are two ways to group a product of three matrices, $(x_0 x_1) x_2$ and $x_0 (x_1 x_2)$, and there are five ways for a product of four matrices: $((x_0 x_1) x_2) x_3$, $(x_0 (x_1 x_2)) x_3$, $(x_0 x_1) (x_2 x_3)$, $x_0 ((x_1 x_2) x_3)$, and $x_0 (x_1 (x_2 x_3))$. A bit more work gives us 14 ways to compute a product of five matrices: There are five ways if one pair of parentheses is $x_0 (x_1 x_2 x_3 x_4)$, another five for $(x_0 x_1 x_2 x_3) x_4$, two for $(x_0 x_1) (x_2 x_3 x_4)$, and two more for $(x_0 x_1 x_2) (x_3 x_4)$. We record these numbers in the following table.

k	C_k
0	1
1	1
2	2
3	5
4	14

Can we determine a recurrence relation for C_k ?

Suppose we group the terms so that the last multiplication occurs between x_i and x_{i+1} :

$$(x_0 x_1 \dots x_i)(x_{i+1} \dots x_k).$$

Then there are C_i ways to group the terms in the first part of the product, and C_{k-1-i} ways for the second part, so there are $C_i C_{k-1-i}$ ways to group the remaining terms in this case. Summing over i , we obtain the following formula for the total number of ways to group the $k + 1$ terms:

$$C_k = \sum_{i=0}^{k-1} C_i C_{k-1-i}, \quad k \geq 1. \quad (2.46)$$

We compute

$$\begin{aligned} C_1 &= C_0 C_0 = 1, \\ C_2 &= C_0 C_1 + C_1 C_0 = 2, \\ C_3 &= C_0 C_2 + C_1 C_1 + C_2 C_0 = 5, \\ C_4 &= C_0 C_3 + C_1 C_2 + C_2 C_1 + C_3 C_0 = 14, \\ C_5 &= C_0 C_4 + C_1 C_3 + C_2 C_2 + C_3 C_1 + C_4 C_0 = 42. \end{aligned}$$

We would like to solve this recurrence to find a formula for C_k , so let us define the generating function for this sequence,

$$G(x) = \sum_{k \geq 0} C_k x^k.$$

Unlike other recurrences we have studied, this one is not linear, and has a variable number of terms. To solve it, we require one fact concerning products of generating functions.

If $A(x) = \sum_{k \geq 0} a_k x^k$ and $B(x) = \sum_{k \geq 0} b_k x^k$, then

$$A(x)B(x) = \sum_{k \geq 0} \left(\sum_{i=0}^k a_i b_{k-i} \right) x^k.$$

Let $c_k = \sum_{i=0}^k a_i b_{k-i}$. The sequence $\{c_k\}$ is called the *convolution* of the sequences $\{a_k\}$ and $\{b_k\}$. Thus, the generating function of the convolution of two sequences is the product of the generating functions of the sequences.

Using this fact, we find that

$$\begin{aligned} G(x) &= \sum_{k \geq 0} C_k x^k \\ &= C_0 + \sum_{k \geq 1} \left(\sum_{i=0}^{k-1} C_i C_{k-1-i} \right) x^k \\ &= 1 + x \sum_{k \geq 0} \left(\sum_{i=0}^k C_i C_{k-i} \right) x^k \\ &= 1 + xG(x)^2, \end{aligned}$$

since $\{\sum_{i=0}^k C_i C_{k-i}\}$ is the convolution of $\{C_k\}$ with itself. Thus,

$$xG(x)^2 - G(x) + 1 = 0,$$

and so

$$G(x) = \frac{1 \pm \sqrt{1 - 4x}}{2x}.$$

Only one of these functions can be the generating function for $\{C_k\}$, and it must satisfy

$$\lim_{x \rightarrow 0} G(x) = C_0 = 1.$$

It is easy to check that the correct function is

$$G(x) = \frac{1 - \sqrt{1 - 4x}}{2x}.$$

We now expand $G(x)$ as a Maclaurin series to find a formula for C_k . Using the generalized binomial theorem and the identity for negating the upper index, we find that

$$\begin{aligned} (1 - 4x)^{1/2} &= \sum_{k \geq 0} \binom{1/2}{k} (-4x)^k \\ &= \sum_{k \geq 0} \binom{k - 3/2}{k} 4^k x^k \\ &= 1 + \sum_{k \geq 1} \binom{k - 3/2}{k} 4^k x^k \\ &= 1 + 4x \sum_{k \geq 0} \binom{k - 1/2}{k + 1} 4^k x^k. \end{aligned}$$

Therefore,

$$G(x) = -2 \sum_{k \geq 0} \binom{k - 1/2}{k + 1} 4^k x^k,$$

and so

$$C_k = -2^{2k+1} \binom{k - 1/2}{k + 1}.$$

We can find a much simpler form for C_k . Expanding the generalized binomial coefficient and multiplying each term in the product by 2, we compute that

$$\begin{aligned} C_k &= -\frac{2^{2k+1}}{(k + 1)!} \prod_{i=0}^k \left(k - \frac{1}{2} - i\right) \\ &= -\frac{2^k}{(k + 1)!} \prod_{i=0}^k (2k - 1 - 2i). \end{aligned}$$

The product consists of all the odd numbers between -1 and $2k - 1$, so

$$\begin{aligned} C_k &= \frac{2^k}{(k+1)!} \prod_{i=1}^k (2i-1) \\ &= \frac{2^k}{(k+1)!} \prod_{i=1}^k \frac{(2i-1)(2i)}{2i} \\ &= \frac{1}{k!(k+1)!} \prod_{i=1}^k (2i-1)(2i). \end{aligned}$$

The remaining product is simply $(2k)!$, so

$$C_k = \frac{(2k)!}{k!(k+1)!} = \frac{1}{k+1} \binom{2k}{k}. \quad (2.47)$$

C_k is called the k th *Catalan number*.

Incidentally, since C_k is an integer, we have shown that $k+1$ always divides the binomial coefficient $\binom{2k}{k}$. Can you find an independent arithmetic proof of this fact?

Sloane and Plouffe [258] remark that the Catalan numbers are perhaps the second most frequently occurring numbers in combinatorics, after the binomial coefficients. Indeed, Stanley [262, ex. 6.19] lists 66 different combinatorial interpretations of these numbers! We close with another problem whose solution involves the Catalan numbers.

A *rooted tree* is a tree with a distinguished vertex called the *root*. The vertices in a rooted tree form a hierarchy, with the root at the highest level, and the level of every other vertex determined by its distance from the root. Some familiar terms are often used to describe relationships between vertices in a rooted tree: If v and w are adjacent vertices and v lies closer to the root than w , then v is the *parent* of w , and w is a *child* of v . Likewise, one may define siblings, grandparents, cousins, and other family relationships in a rooted tree.

We say that a rooted tree is *strictly binary* if every parent vertex has exactly two children. How many strictly binary trees are there with k parent vertices? Do not take symmetry into account: If two trees are mirror images of one another, count both configurations. Figure 2.4 shows that there are five trees with three parent vertices.

It is easy to see that the number of strictly binary trees with k parent vertices is C_k . By Exercise 2, every such tree has $k+1$ leaves. Label these vertices with x_0 through x_k from left to right in the tree. Then the tree determines an order of multiplication for the x_i . For example, the five trees in Figure 2.4 correspond to the multiplications $((x_0x_1)x_2)x_3$, $(x_0(x_1x_2))x_3$, $(x_0x_1)(x_2x_3)$, $x_0((x_1x_2)x_3)$, and $x_0(x_1(x_2x_3))$, respectively. Binary trees like these are often used in computer science to designate the order of evaluation of arithmetic expressions.





FIGURE 2.4. Strictly binary trees with three parent vertices.

Exercises

1. Show that every vertex in a rooted tree has at most one parent.
2. Show that a strictly binary tree having exactly k parent vertices has exactly $k + 1$ leaves.
3. A *diagonal* of a convex polygon is a line segment connecting two non-adjacent vertices of the polygon. Let p_n denote the number of ways to decompose a convex polygon having n vertices into triangles by drawing $n - 3$ diagonals that do not cross inside the polygon. Assume that the vertices of the polygon are labeled, so that triangulations with different orientations are counted separately.
 - (a) Determine p_3 , p_4 , p_5 , and p_6 by showing all the possible triangulations.
 - (b) Let v be a fixed vertex of a polygon with $n = 7$ sides. Count all the triangulations of the heptagon by considering two cases: (i) v is not an endpoint of any of the four diagonals added in a triangulation, and (ii) v is an endpoint of at least one of the diagonals. Use this to determine the value of p_7 without drawing every possible triangulation.
 - (c) Determine a formula for p_n .
4. A *staircase* of size n is a path in the plane from the origin to the point (n, n) consisting of exactly n horizontal and n vertical steps, each of length 1, with the added condition that the path never rises above the line $y = x$. Let s_n denote the number of staircases of size n . For example, $s_1 = 1$ since the only staircase is \square . Also, $s_2 = 2$ since the only possible staircases are



- (a) Determine s_3 and s_4 by drawing all the staircases of these sizes.
 - (b) For $1 \leq k \leq n$, let $s_{n,k}$ denote the number of staircases of size n for which the first intersection of the path with the $y = x$ line (after the origin) is the point (k, k) . Determine a formula for $s_{n,k}$ in terms of the numbers s_m .
 - (c) Determine a recurrence relation for s_n , then a formula.
5. Let r_n denote the number of mountain ridgelines one can draw using n ascending steps and n descending steps. A ridgeline must start and end on the horizon, and may never dip below the horizon. For example, $r_1 = 1$ since the only ridgeline is \wedge , and the following illustrates a valid ridgeline with $n = 11$.



Two ridgelines that are mirror images of one another count as different arrangements.

- (a) Determine $r_2, r_3,$ and r_4 by drawing all of the possible ridgelines.
 - (b) Use a combinatorial argument to determine a recurrence relation for r_n , then find a formula for r_n .
6. Suppose $2k$ people are seated around a table. How many ways are there for the k pairs of people to shake hands simultaneously across the table in such a way that no arms cross?
7. Show that the coefficient of x^k in the Maclaurin series expansion of $(1 - (1 - 3x)^{1/3})/x$ is

$$\frac{1}{(k + 1)!} \prod_{i=1}^k (3i - 1).$$

8. Use an arithmetic argument to show that $(2k)!$ is divisible by $k!(k + 1)!$.
Hint: First compute the number of times a prime number p divides $m!$.

2.7 Pólya’s Theory of Counting

Who are you who are so wise in the ways of science?

— Sir Bedivere, in *Monty Python and the Holy Grail*

How many ways can King Arthur and his knights sit at the round table? How many different necklaces with n beads can be formed using m different kinds of beads?

Both these questions ask for a number of combinations in the presence of symmetry. Since there is no distinguished position at a round table, seating Arthur first, then Gawain, Percival, Bedivere, Tristram, and Galahad clockwise around the table yields the same configuration as seating Tristram first, then Galahad, Arthur, Gawain, Percival, and Bedivere in clockwise order. Similarly, we should consider two necklaces to be identical if we can transform one into the other by rotating the necklace or by turning it over.

Before answering these questions, let us first rephrase them in the language of group theory.

2.7.1 Permutation Groups

I haven't fought just one person in a long time. I've been specializing in groups.

— Fezzik, in *The Princess Bride*

A *group* consists of a set G together with a binary operator \circ defined on this set. The set and the operator must satisfy four properties.

- **Closure.** For every a and b in G , $a \circ b$ is in G .
- **Associativity.** For every a , b , and c in G , $a \circ (b \circ c) = (a \circ b) \circ c$.
- **Identity.** There exists an element e in G that satisfies $e \circ a = a \circ e = a$ for every a in G . The element e is called the *identity* of G .
- **Inverses.** For every element a in G , there exists an element b in G such that $a \circ b = b \circ a = e$. The element b is called the *inverse* of a .

In addition, if $a \circ b = b \circ a$ for every a and b in G , we say that G is an *abelian*, or *commutative*, group.

For example, the set of integers forms a group under addition. The identity element is 0, since $0 + i = i + 0 = i$ for every integer i , and the inverse of the integer i is the integer $-i$. Similarly, the set of nonzero rational numbers forms a group under multiplication (with identity element 1), as does the set of nonzero real numbers.

We can also construct groups of permutations. A permutation of n objects may be described by a function π defined on the set $\{1, 2, \dots, n\}$ by ordering the objects in some fashion, then taking $\pi(i) = j$ if the i th object in the ordering occupies the j th position in the permutation. For example, the permutation $[c, d, a, e, b]$ of the list $[a, b, c, d, e]$ is represented by the function π defined on the set $\{1, 2, 3, 4, 5\}$, with $\pi(1) = 3$, $\pi(2) = 5$, $\pi(3) = 1$, $\pi(4) = 2$, and $\pi(5) = 4$. Notice that a function $\pi : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ arising from a permutation has the property that $\pi(i) \neq \pi(j)$ whenever $i \neq j$. Such a function is called an

injective, or *one-to-one*, function. The map π also has the property that for every m with $1 \leq m \leq n$, there exists a number i such that $\pi(i) = m$. A function like this that maps to every element in its range is called *surjective*, or *onto*, and a function that is both injective and surjective is said to be a *bijection*. Thus, every permutation of n objects corresponds to a bijection π on the set $\{1, 2, \dots, n\}$, and every such bijection corresponds to a permutation.

Let S_n denote the set of all bijections on the set $\{1, 2, \dots, n\}$. Exercise 4 asks you to verify that this set forms a group under the operation of composition of functions. For example, the identity element of the group is the identity map π_0 , defined by $\pi_0(k) = k$ for each k , since $\pi \circ \pi_0 = \pi_0 \circ \pi = \pi$ for every π in S_n . This group is called the *symmetric group* on n elements.

The size of the group S_n is the number of permutations of n objects, so $|S_n| = n!$. Because of our correspondence, we normally refer to an element of S_n as a permutation, rather than a bijection.

To specify a particular permutation π in S_n , we need to name the value of $\pi(k)$ for each k . This is often written in two rows as follows:

$$\begin{pmatrix} 1 & 2 & 3 & \dots & n \\ \pi(1) & \pi(2) & \pi(3) & \dots & \pi(n) \end{pmatrix}.$$

For example,

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 3 & 5 & 1 & 2 & 4 \end{pmatrix}$$

denotes the permutation described earlier.

We can describe the permutations in a more succinct manner by using *cycle notation*. For example, in the permutation above, π sends 1 to 3 and 3 to 1, and sends 2 to 5, 5 to 4, and 4 to 2. So we can think of π as a combination of two *cycles*, $1 \rightarrow 3 \rightarrow 1$ and $2 \rightarrow 5 \rightarrow 4 \rightarrow 2$, and denote it by

$$(13)(254).$$

Of course, we could also denote this same permutation by the cycles $(542)(31)$, so to make our notation unique, we make two demands. First, the cycle containing 1 must appear first, followed by the cycle containing the smallest number not appearing in the first cycle, and so on. Second, we require the first number listed in each cycle to be the smallest number appearing in that cycle. To simplify the notation, cycles of length 1 are usually omitted, so $(1253)(4)$ is written more simply as (1253) . The identity permutation is denoted by (1) .

The composition of two permutations is computed from right to left. For example, suppose $\pi_1 = (13)(254)$ and $\pi_2 = (15423)$. We determine the composition $\pi_1 \circ \pi_2$ by applying π_2 first, then π_1 . Since π_2 sends 1 to 5, and π_1 sends 5 to 4, the composition $\pi_1 \circ \pi_2$ then sends 1 to 4. In the same way, we see that $\pi_1 \circ \pi_2$ sends 4 to 5, 5 to 2, 2 to 1, and 3 to 3. Thus, $\pi_1 \circ \pi_2 = (1452)$. In cycle notation, we denote the composition of two permutations by juxtaposing their cycles, so

$$\pi_1 \circ \pi_2 = (13)(254)(15423) = (1452).$$

Notice that the cycles for π_1 appear first, so products of cycles are always computed from right to left. Also, we calculate that $\pi_2 \circ \pi_1 = (15423)(13)(254) = (2435)$, so in general S_n is not an abelian group.

A subset H of a group G is called a *subgroup* of G if H is itself a group under the same binary operation. The group S_n contains many subgroups; for example, $\{(1), (12)\}$ is a subgroup of S_n for every $n \geq 2$. We investigate three particularly important subgroups of S_n .

The Cyclic Group

If π is a permutation in S_n and m is a nonnegative integer, let π^m denote the permutation obtained by composing π with itself m times, so $\pi^0 = (1)$, and $\pi^3 = \pi \circ \pi \circ \pi$. Let

$$\langle \pi \rangle = \{\pi^m : m \geq 0\}, \quad (2.48)$$

so that $\langle \pi \rangle$ is a subset of S_n . In fact (Exercise 5), $\langle \pi \rangle$ is a subgroup of S_n , and we call this group the cyclic subgroup generated by π in S_n .

The *cyclic group* C_n is the subgroup of the symmetric group S_n generated by the permutation $(123 \cdots n)$, so

$$C_n = \langle (123 \cdots n) \rangle. \quad (2.49)$$

Clearly, C_n contains n elements, since n applications of the generating permutation are required to return to the identity permutation. For example, $(1234)^2 = (13)(24)$, $(1234)^3 = (1432)$, and $(1234)^4 = (1)$, so

$$C_4 = \{(1), (1234), (13)(24), (1432)\}. \quad (2.50)$$

The group C_n may be realized as the group of rotational symmetries of a regular polygon having n sides. For example, each of the permutations of (2.50) corresponds to a permutation of the vertices of Figure 2.5 obtained by rotating the square by 0, 90, 180, or 270 degrees.

The Dihedral Group

The *dihedral group* D_n is the group of symmetries of a regular polygon with n sides, including reflections as well as rotations. Since C_n consists of just the rotational symmetries of such a figure, evidently C_n is a subgroup of D_n .

Referring to Figure 2.5, we see that D_4 consists of the four rotations of C_4 , plus the four reflections $(12)(34)$, $(14)(23)$, (13) , and (24) . The first two permutations represent reflections about the vertical and horizontal axes of symmetry of the square; the last two represent flips about the diagonal axes of symmetry. In general, if n is even, we obtain $n/2$ reflections through axes of symmetry that pass through opposite vertices, and $n/2$ reflections through axes that pass through midpoints of opposite edges. Combining these with the n rotations of C_n , we find that $|D_n| = 2n$ in this case.




FIGURE 2.5. A square with labeled vertices.




FIGURE 2.6. A regular pentagon with labeled vertices.

Using Figure 2.6, we find that D_5 consists of five rotations and five reflections,

$$D_5 = \{(1), (12345), (13524), (14253), (15432), (25)(34), \\ (13)(45), (15)(24), (12)(35), (14)(23)\}.$$

It is easy to see that we always obtain n reflections if n is odd, so $|D_n| = 2n$ for every $n \geq 1$.

The Alternating Group

Every permutation can be expressed as a product of *transpositions*, which are cycles of length 2. For example, the cycle (123) can be written as the product $(12)(23)$, and the permutation $(1234)(567)$ can be expressed as the product of six transpositions: $(12)(23)(34)(56)(67)$. Such a decomposition is not unique; for instance, (123) may also be written as $(23)(13)$, or $(12)(23)(13)(13)$. However, the number of transpositions in any representation of one permutation is either always an even number, or always an odd number. Exercise 6 outlines a proof of this fact. If a permutation π always decomposes into an even number of transpositions, we say that π is an *even* permutation; otherwise, it is an *odd* permutation. Notice that the identity permutation is even, since it is represented by a product of zero transpositions.

The *alternating group* A_n consists of the even permutations of S_n . For example, $A_3 = \{(1), (123), (132)\} = C_3$, and

$$A_4 = \{(1), (123), (132), (124), (142), (134), (143), \\ (234), (243), (12)(34), (13)(24), (14)(23)\}.$$

Exercises 7 and 8 ask you to verify that A_n is a group of size $|A_n| = n!/2$ for $n \geq 2$, and that A_n is not abelian for $n \geq 5$.

Exercises

1. Show that the identity element of a group is unique.
2. (a) Suppose that M is a finite set and $f : M \rightarrow M$ is an injective function. Show that f is a bijection.
 (b) Suppose that M is a finite set and $f : M \rightarrow M$ is a surjective function. Show that f is a bijection.
 (c) Show that neither of these statements is necessarily true if M is an infinite set.
3. In each part, determine all values of n that satisfy the statement.
 - (a) C_n is a subgroup of A_n .
 - (b) D_n is a subgroup of A_n .
 - (c) C_n is a subgroup of D_{n+1} .
 - (d) C_n is a subgroup of S_{n+1} .
4. Verify that S_n forms a group under composition of functions by checking that each of the required properties is satisfied.
 - (a) Closure. If π_1 and π_2 are bijections on $\{1, 2, \dots, n\}$, show that $\pi_1 \circ \pi_2$ is also a bijection on $\{1, 2, \dots, n\}$.
 - (b) Associativity. If π_1, π_2 , and π_3 are in S_n , show that $\pi_1 \circ (\pi_2 \circ \pi_3)$ and $(\pi_1 \circ \pi_2) \circ \pi_3$ represent the same function in S_n .
 - (c) Identity. Check that $\pi_0 \circ \pi = \pi \circ \pi_0 = \pi$, for every π in S_n . Here, π_0 is the identity map on $\{1, 2, \dots, n\}$.
 - (d) Inverses. Given a bijection π in S_n , construct a bijection π^{-1} in S_n satisfying $\pi \circ \pi^{-1} = \pi^{-1} \circ \pi = \pi_0$.
5. Suppose that G is a group and g is an element of G .
 - (a) Show that $\langle g \rangle$ is a subgroup of G .
 - (b) Show that $\langle g \rangle$ is abelian.

6. Let \mathbf{x} denote the vector of n variables (x_1, x_2, \dots, x_n) . Define

$$P(\mathbf{x}) = \prod_{1 \leq i < j \leq n} (x_i - x_j),$$

and if $\pi \in S_n$, let

$$P_\pi(\mathbf{x}) = \prod_{1 \leq i < j \leq n} (x_{\pi(i)} - x_{\pi(j)}).$$

- (a) Show that $P_\pi(\mathbf{x}) = \pm P(\mathbf{x})$.
 - (b) Show that $P_\pi(\mathbf{x}) = -P(\mathbf{x})$ if π is a transposition.
 - (c) Conclude that no permutation π in S_n can be represented both as a product of an even number of transpositions and as a product of an odd number of transpositions.
7. (a) Prove that A_n is a group.
 (b) Show that A_n is not abelian for $n \geq 5$.
8. Let $n \geq 2$, let B_n denote the set of odd permutations in S_n , and let τ be a transposition in S_n .
- (a) Show that the map $T : S_n \rightarrow S_n$ defined by $T(\pi) = \tau \circ \pi$ is a bijection.
 - (b) Show that T maps A_n to B_n , and B_n to A_n .
 - (c) Conclude that $|A_n| = n!/2$.
9. Determine the group of symmetries of each of the following objects.
- (a) The vertices of a regular tetrahedron.
 - (b) The vertices of a cube.
 - (c) The vertices of a regular octahedron.

2.7.2 Burnside's Lemma

Burnside had submitted the scheme to Meade and myself, and we both approved of it, as a means of keeping the men occupied.

— *Personal Memoirs of U. S. Grant*

Armed with our knowledge of permutation groups, we now develop a general method for counting combinations in the presence of symmetry. In general, we are given a set of objects S , a set of colorings of these objects C , and a group of permutations G representing symmetries possessed by configurations of the objects. We consider two colorings in C to be equivalent if one of the permutations in G transforms one coloring to the other, and we would like to determine the number of nonequivalent colorings in C .

For example, suppose $S = \{1, 2, 3, 4\}$ is the set of vertices of the square in Figure 2.5, and C is the set of all possible colorings of these vertices using two colors, red and green. Let $rrgr$ denote the coloring where vertices 1, 2, and 4 are red and vertex 3 is green. Then

$$C = \{gggg, gggr, ggrg, ggrr, grgg, grgr, grrg, grrr, rggg, rggr, rgrg, rgrr, rrrg, rrgr, rrrg, rrrr\}. \quad (2.51)$$

We consider two colorings in C to be equivalent if one can be transformed to the other by a rotation of the square. For example, rotating the coloring $rrgr$ yields the equivalent colorings $rrrg$, $grrr$, and $rgrr$. So we choose G to be the group of rotations, C_4 . A permutation π in C_4 is a function defined on the set $\{1, 2, 3, 4\}$, but π induces a map π^* defined on the set of colorings C in a natural way. For example, if π is the 180-degree rotation $(13)(24)$, then the induced map π^* rotates a coloring by the same amount, so $\pi^*(rrgr) = grrr$, and $\pi^*(grgr) = grgr$.

If c_1 and c_2 are two equivalent colorings in C , so $\pi^*(c_1) = c_2$ for some $\pi \in G$, we write $c_1 \sim c_2$. Using the fact that G is a group, it is easy to verify (Exercise 1) that the relation \sim on the set of colorings is

- reflexive: $c \sim c$ for all colorings c ,
- symmetric: $c_1 \sim c_2$ implies $c_2 \sim c_1$, and
- transitive: $c_1 \sim c_2$ and $c_2 \sim c_3$ implies $c_1 \sim c_3$.

A relation possessing these three properties is called an *equivalence relation*. By grouping together collections of mutually equivalent elements, an equivalence relation on a set partitions the set into a number of disjoint subsets, called *equivalence classes*. Our goal then is to determine the number of equivalence classes of C under the relation \sim .

In our example, the group C_4 partitions our set of colorings (2.51) into six equivalence classes:

$$\begin{aligned} &\{gggg\}, \\ &\{gggr, ggrg, grgg, rggg\}, \\ &\{ggrr, grrg, rggr, rrrg\}, \\ &\{grgr, rgrg\}, \\ &\{grrr, rgrr, rrgr, rrrg\}, \\ &\{rrrr\}. \end{aligned}$$

Therefore, there are just six ways to color the vertices of a square using two colors, after discounting rotational symmetries.

We can now translate the problems from the introduction to this section into this more abstract setting. In the round table problem, S is the set of n places at the table, G is C_n , and C is the collection of the $n!$ seating assignments. In the

necklace problem, S is the set of n bead positions, G is D_n , and C is the collection of the m^n possible arrangements of the m kinds of beads on the necklace.

Before presenting a general method to solve problems like these, we introduce three sets that will be useful in our analysis. Given a permutation π in G , define C_π to be the set of colorings that are invariant under action by the induced map π^* ,

$$C_\pi = \{c \in C : \pi^*(c) = c\}. \tag{2.52}$$

This set is called the *invariant set* of π in C . Similarly, given a coloring c in C , define G_c to be the set of permutations π in G for which c is a fixed coloring,

$$G_c = \{\pi \in G : \pi^*(c) = c\}. \tag{2.53}$$

This set is called the *stabilizer* of c in G . It is always a subgroup of G . Finally, let \bar{c} be the set of colorings in C that are equivalent to c under the action of the group G ,

$$\bar{c} = \{\pi^*(c) : \pi \in G\}. \tag{2.54}$$

The set \bar{c} is thus the equivalence class of c under the relation \sim . It is also called the *orbit* of c under the action of G .

For example, if C is given by (2.51) and G is the dihedral group D_4 , we have

$$\overline{gggr} = \{gggr, ggrg, grgg, rggg\}$$

and

$$G_{gggr} = \{(1), (13)\}.$$

Also,

$$\overline{grgr} = \{grgr, rgrg\}$$

and

$$G_{grgr} = \{(1), (13)(24), (13), (24)\}.$$

Notice that in both cases, the product of the size of the stabilizer of a coloring with the size of the equivalence class of the same coloring equals the number of elements in the group. The following lemma proves that this is always the case.

Lemma 2.8. *Suppose a group G acts on a set of colorings C . For any coloring c in C , we have $|G_c| |\bar{c}| = |G|$.*

Proof. We prove this by showing that every permutation in G may be represented in a unique way as a composition of a permutation in G_c with a permutation in a particular set P , where $|P| = |\bar{c}|$. Suppose there are m colorings in the equivalence class of c , $\bar{c} = \{c_1, c_2, \dots, c_m\}$. For each i between 1 and m , select a permutation $\pi_i \in G$ such that $\pi_i^*(c) = c_i$, and let $P = \{\pi_1, \pi_2, \dots, \pi_m\}$.

Now let π be an arbitrary permutation in G . Then $\pi^*(c) = c_i$ for some i , so $\pi^*(c) = \pi_i^*(c)$. Thus $(\pi_i^{-1} \circ \pi)^*(c) = c$, and so $\pi_i^{-1} \circ \pi \in G_c$. Since $\pi_i \circ (\pi_i^{-1} \circ \pi) = \pi$, we see that π has at least one representation in the desired form. Suppose now that $\pi = \pi_i \circ \sigma = \pi_j \circ \tau$, for some π_i and π_j in P and some σ and τ in G_c . Then $\pi_i(\sigma(c)) = \pi_i(c) = c_i$ and $\pi_j(\tau(c)) = c_j$, so $c_i = c_j$, and hence $i = j$. Therefore, $\sigma = \tau$, so the representation of π is unique. \square

The following formula for the number of equivalence classes of C under the action of a group G is usually named for Burnside (the English mathematician, not the American Civil War general), as it was popularized by his book [45]. This result was first proved by Frobenius [115], however, and Burnside even attributes the formula to Frobenius in the first edition of his textbook [45]. Further details on the history of this result appear in Neumann [213] and Wright [288].

Briefly, Burnside's Lemma states that the number of equivalence classes of colorings is the average size of the invariant sets.

Theorem 2.9 (Burnside's Lemma). *The number of equivalence classes N of the set C in the presence of the group of symmetries G is given by*

$$N = \frac{1}{|G|} \sum_{\pi \in G} |C_\pi|. \tag{2.55}$$

Proof. If P is a logical expression, let $[P]$ be 1 if P is true and 0 if P is false. Then

$$\begin{aligned} \frac{1}{|G|} \sum_{\pi \in G} |C_\pi| &= \frac{1}{|G|} \sum_{\pi \in G} \sum_{c \in C} [\pi^*(c) = c] \\ &= \frac{1}{|G|} \sum_{c \in C} \sum_{\pi \in G} [\pi^*(c) = c] \\ &= \frac{1}{|G|} \sum_{c \in C} |G_c| \\ &= \sum_{c \in C} \frac{1}{|\bar{c}|} \\ &= \sum_{\bar{c}} \sum_{c \in \bar{c}} \frac{1}{|\bar{c}|} \\ &= \sum_{\bar{c}} 1 \\ &= N. \end{aligned}$$

We applied Lemma 2.8 to obtain the fourth line. □

We may apply Burnside's Lemma to solve the problems we described earlier. In the round table problem, $|G| = n$. The invariant set of the identity permutation is the entire set of colorings, $C_{(1)} = C$, and the invariant set of any nontrivial rotation π is empty, $C_\pi = \{ \}$. Therefore, the number of nonequivalent seating arrangements is $|C|/n = (n - 1)!$.

To determine the number of nonequivalent necklaces with four beads using two different kinds of beads, we calculate $|C_{(1)}| = 16$, $|C_{(13)}| = |C_{(24)}| = 8$, $|C_{(12)(34)}| = |C_{(13)(24)}| = |C_{(14)(23)}| = 4$, and $|C_{(1234)}| = |C_{(1432)}| = 2$. Therefore, $N = (16 + 2 \cdot 8 + 3 \cdot 4 + 2 \cdot 2)/8 = 6$. Last, we calculate the number

of nonequivalent three-bead necklaces using three different kinds of beads. Here, $|C_{(1)}| = 27$, $|C_{(12)}| = |C_{(13)}| = |C_{(23)}| = 9$, and $|C_{(123)}| = |C_{(132)}| = 3$, so $N = 60/6 = 10$.

Exercises

1. Show that \sim is an equivalence relation on C .
2. Prove that G_c is a subgroup of G .
3. How many different necklaces having five beads can be formed using three different kinds of beads if we discount:
 - (a) Both flips and rotations?
 - (b) Rotations only?
 - (c) Just one flip?
4. The commander of a space cruiser wishes to post four sentry ships arrayed around the cruiser at the vertices of a tetrahedron for defensive purposes, since an attack can come from any direction.
 - (a) How many ways are there to deploy the ships if there are two different kinds of sentry ships available, and we discount all symmetries of the tetrahedral formation?
 - (b) How many ways are there if there are three different kinds of sentry ships available?
5.
 - (a) How many ways are there to label the faces of a cube with the numbers 1 through 6 if each number may be used more than once?
 - (b) What if each number may only be used once?

2.7.3 The Cycle Index

Lance Armstrong (7), Jacques Anquetil (5), Bernard Hinault (5), Miguel Indurain (5), Eddy Merckx (5), Louison Bobet (3), Greg LeMond (3), Philippe Thys (3).

— Multiple Tour de France winners

To use Burnside's Lemma to count the number of equivalence classes of a set of colorings C , we must compute the size of the invariant set C_π associated with every permutation π in a group of symmetries G . A simple observation allows us to compute the size of this set easily in many situations.

Suppose we wish to determine the number of ways to color n objects using up to m colors, discounting symmetries on the objects described by a group G . If a coloring is invariant under the action of a permutation π in G , then every object permuted by one cycle of π must have the same color. Therefore, if π

has k disjoint cycles, the number of colorings invariant under the action of π is $|C_\pi| = m^k$. For example, if S is the set of vertices of a square and $G = D_4$, then $|C_{(1234)}| = m$, $|C_{(12)(34)}| = m^2$, $|C_{(13)(2)(4)}| = m^3$, and $|C_{(1)(2)(3)(4)}| = m^4$. Notice that it is essential to include the cycles of length 1 in these calculations.

With this in mind, we define the *cycle index* of a group G of permutations on n objects. For a permutation π in G , define a monomial M_π associated with π in the following way. If π is a product of k cycles, and the i th cycle has length ℓ_i , let

$$M_\pi = M_\pi(x_1, x_2, \dots, x_n) = \prod_{i=1}^k x_{\ell_i}. \tag{2.56}$$

Here, x_1, x_2, \dots, x_n are indeterminates. The cycle index of G is defined by

$$P_G(\mathbf{x}) = \frac{1}{|G|} \sum_{\pi \in G} M_\pi(\mathbf{x}), \tag{2.57}$$

where \mathbf{x} denotes the vector (x_1, x_2, \dots, x_n) .

For example, for $G = D_4$, we find that

$$\begin{aligned} M_{(1)(2)(3)(4)} &= x_1^4, \\ M_{(13)(2)(4)} &= M_{(1)(24)(3)} = x_1^2 x_2, \\ M_{(12)(34)} &= M_{(13)(24)} = M_{(14)(23)} = x_2^2, \\ M_{(1234)} &= M_{(1432)} = x_4. \end{aligned}$$

Therefore,

$$P_{D_4}(x_1, x_2, x_3, x_4) = \frac{1}{8} (x_1^4 + 2x_1^2 x_2 + 3x_2^2 + 2x_4), \tag{2.58}$$

and

$$P_{C_4}(x_1, x_2, x_3, x_4) = \frac{1}{4} (x_1^4 + x_2^2 + 2x_4). \tag{2.59}$$

By Burnside's Lemma, the number of ways to color n objects using up to m colors, discounting the symmetries of G , is $P_G(m, m, \dots, m)$. For example, the number of equivalence classes of four-bead necklaces composed using m different kinds of beads is

$$P_{D_4}(m, m, m, m) = \frac{1}{8} (m^4 + 2m^3 + 3m^2 + 2m).$$

Substituting $m = 2$, we find there are six different colorings, as before.

Finally, let us compute the number of twenty-bead necklaces composed of rhodonite, rose quartz, and lapis lazuli beads. We must determine the cycle index for the group D_{20} . We find that eight of the rotations, those by $18k$ degrees with $k = 1, 3, 7, 9, 11, 13, 17$, or 19 , are a single cycle of length 20, yielding the term $8x_{20}$ in the cycle index. Four rotations, $k = 2, 6, 14$, and 18 , make two cycles of length 10, contributing $4x_{10}^2$. Rotations with $k = 4, 8, 12$, or 16 make

four cycles of length 5, adding $4x_5^4$, and $k = 5$ or 15 contributes $2x_4^5$. The rotation with $k = 10$ yields x_2^{10} , and the identity adds x_1^{20} . Ten of the reflections, the ones about axes of symmetry that pass through midpoints of edges, are each represented by ten transpositions, contributing $10x_2^{10}$. The other ten reflections, flipping about opposite vertices, yield $10x_1^2x_2^9$. Therefore,

$$P_{D_{20}}(x_1, \dots, x_{20}) = \frac{1}{40} (x_1^{20} + 10x_1^2x_2^9 + 11x_2^{10} + 2x_4^5 + 4x_5^4 + 4x_{10}^2 + 8x_{20}), \quad (2.60)$$

and the number of different twenty-bead necklaces that can be made using three kinds of beads is $P_{D_{20}}(3, \dots, 3) = 87\,230\,157$.

Exercises

1. Show that the monomial M_π defined in (2.57) has the property that the sum $\sum_{i=1}^k \ell_i = n$.
2. (a) Determine the cycle index for S_4 and for A_4 .
 (b) Show that $P_{S_4}(m, m, m, m)$ may be written as a binomial coefficient.
 (c) Determine the smallest value of m for which $P_{A_4}(m, m, m, m) > P_{S_4}(m, m, m, m)$.
3. Determine the number of different necklaces with 21 beads that can be made using four kinds of beads. Your equivalence classes should account for both rotations and flips.
4. Determine the number of eight-bead necklaces that can be made using red, green, blue, and white beads under each of the following groups of symmetries.
 - (a) D_8 .
 - (b) A subgroup of D_8 having four elements. How does the answer depend on the subgroup you choose?
5. Determine the cycle index for the group of symmetries of the faces of a cube, and use this to determine the number of different six-sided dice that can be manufactured using m different labels for the faces of the dice. Assume that each label may be used any number of times.

2.7.4 Pólya's Enumeration Formula

I have yet to see any problem, however complicated, which, when looked at in the right way, did not become still more complicated.

— Poul Anderson

We can use the cycle index to solve more complicated problems on arrangements in the presence of symmetry. Suppose we need to determine the number of equivalence classes of colorings of n objects using the m colors y_1, y_2, \dots, y_m , where

each color y_i occurs a prescribed number of times. For example, how many different necklaces can be made using exactly two rhodonite, nine rose quartz, and nine lapis lazuli beads?

Let us define the *pattern inventory* of the different ways to color n objects using m colors with respect to a symmetry group G as a generating function in m variables,

$$F_G(y_1, y_2, \dots, y_m) = \sum_{\mathbf{v}} a_{\mathbf{v}} y_1^{n_1} y_2^{n_2} \cdots y_m^{n_m}, \quad (2.61)$$

where the sum runs over all vectors $\mathbf{v} = (n_1, n_2, \dots, n_m)$ of nonnegative integers satisfying $n_1 + n_2 + \cdots + n_m = n$, and $a_{\mathbf{v}}$ represents the number of nonequivalent colorings of the n objects where the color y_i occurs precisely n_i times. For example, if we denote a rhodonite bead by r , a rose quartz bead by q , and a lapis lazuli bead by l , we see that the answer to our question above is the coefficient of $r^2 q^9 l^9$ in the generating function

$$F_{D_{20}}(r, q, l) = \sum_{\substack{i+j+k=20 \\ i,j,k \geq 0}} a_{(i,j,k)} r^i q^j l^k.$$

In his influential paper [224] (translated into English by Read [226]), Pólya found that the cycle index can be used to compute the pattern inventory in a simple way. Recall that each occurrence of x_k in the cycle index arises from a permutation having a cycle of length k , and if a coloring is invariant under this permutation, then these k elements must have the same color. So either each of the k objects permuted by this cycle has color y_1 , or each one has color y_2 , etc. In the spirit of generating functions, this choice can be represented by the formal sum $y_1^k + y_2^k + \cdots + y_m^k$. Pólya found that substituting this expression for x_k for each k in the cycle index yields the pattern inventory for the coloring.

Theorem 2.10 (Pólya's Enumeration Formula). *Suppose S is a set of n objects and G is a subgroup of the symmetric group S_n . Let $P_G(\mathbf{x})$ be the cycle index of G . Then the pattern inventory for the nonequivalent colorings of S under the action of G using colors y_1, y_2, \dots, y_m is*

$$F_G(\mathbf{y}) = P_G \left(\sum_{i=1}^m y_i, \sum_{i=1}^m y_i^2, \dots, \sum_{i=1}^m y_i^n \right). \quad (2.62)$$

The proof we present follows Stanley [262, sec. 7.24].

Proof. Let $\mathbf{v} = (n_1, n_2, \dots, n_m)$ be a vector of nonnegative integers of length m whose components sum to n , and let $C_{\mathbf{v}}$ denote the set of colorings of S where exactly n_i of the objects have the color y_i , for each i . Let $C_{\mathbf{v},\pi}$ denote the invariant set of $C_{\mathbf{v}}$ under the action of a permutation π .

If a permutation π in G does not disturb a particular coloring, then every object permuted by one cycle of π must have the same color. Therefore, $|C_{\mathbf{v},\pi}|$ is the coefficient of $y_1^{n_1} y_2^{n_2} \cdots y_m^{n_m}$ in $M_{\pi}(\sum y_i, \sum y_i^2, \dots, \sum y_i^n)$, where M_{π} is

the monomial defined by (2.56). Let $\mathbf{y}^{\mathbf{v}}$ denote the term $y_1^{n_1} y_2^{n_2} \cdots y_m^{n_m}$. Then, summing over all permissible vectors \mathbf{v} , we obtain

$$\sum_{\mathbf{v}} |C_{\mathbf{v}, \pi}| \mathbf{y}^{\mathbf{v}} = M_{\pi} \left(\sum_{i=1}^m y_i, \sum_{i=1}^m y_i^2, \dots, \sum_{i=1}^m y_i^n \right).$$

Now we sum both expressions over all $\pi \in G$ and divide by $|G|$. On the left side, we have

$$\begin{aligned} \frac{1}{|G|} \sum_{\pi \in G} \sum_{\mathbf{v}} |C_{\mathbf{v}, \pi}| \mathbf{y}^{\mathbf{v}} &= \sum_{\mathbf{v}} \left(\frac{1}{|G|} \sum_{\pi \in G} |C_{\mathbf{v}, \pi}| \right) \mathbf{y}^{\mathbf{v}} \\ &= \sum_{\mathbf{v}} a_{\mathbf{v}} \mathbf{y}^{\mathbf{v}} \end{aligned}$$

by Burnside’s Lemma, and this is the pattern inventory (2.61). On the right side, using (2.57), we obtain (2.62), the cycle index of G evaluated at $x_k = \sum_i y_i^k$:

$$\begin{aligned} F_G(\mathbf{y}) &= \frac{1}{|G|} \sum_{\pi \in G} M_{\pi} \left(\sum_{i=1}^m y_i, \sum_{i=1}^m y_i^2, \dots, \sum_{i=1}^m y_i^n \right) \\ &= P_G \left(\sum_{i=1}^m y_i, \sum_{i=1}^m y_i^2, \dots, \sum_{i=1}^m y_i^n \right). \quad \square \end{aligned}$$

For example, the pattern inventory for nonequivalent four-bead necklaces under D_4 using colors red (r), green (g), and blue (b) is

$$\begin{aligned} F_{D_4}(r, g, b) &= P_{D_4} (r + g + b, r^2 + g^2 + b^2, r^3 + g^3 + b^3, r^4 + g^4 + b^4) \\ &= r^4 + g^4 + b^4 + r^3g + rg^3 + r^3b + rb^3 + g^3b + gb^3 \\ &\quad + 2r^2g^2 + 2r^2b^2 + 2g^2b^2 + 2r^2gb + 2rg^2b + 2rgb^2. \end{aligned}$$

The pattern inventory for nonequivalent four-bead necklaces under C_4 using the same three colors is

$$\begin{aligned} F_{C_4}(r, g, b) &= P_{C_4} (r + g + b, r^2 + g^2 + b^2, r^3 + g^3 + b^3, r^4 + g^4 + b^4) \\ &= r^4 + g^4 + b^4 + r^3g + rg^3 + r^3b + rb^3 + g^3b + gb^3 \\ &\quad + 2r^2g^2 + 2r^2b^2 + 2g^2b^2 + 3r^2gb + 3rg^2b + 3rgb^2. \end{aligned}$$

Notice that there are three nonequivalent necklaces with two red beads, one green bead, and one blue bead under C_4 , but only two under D_4 . Can you explain this?

Using (2.60) and Theorem 2.10, we may compute the pattern inventory for twenty-bead necklaces composed of rhodonite (r), rose quartz (q), and lapis lazuli (l) beads. This pattern inventory is shown in Figure 2.7, where we see that there are exactly 231 260 different necklaces with two rhodonite, nine rose quartz, and nine lapis lazuli beads.

$$\begin{aligned}
 F_{D_{20}}(r, q, l) = & r^{20} + r^{19}q + r^{19}l + 10r^{18}q^2 + 10r^{18}ql + 10r^{18}l^2 + 33r^{17}q^3 + 90r^{17}q^2l \\
 & + 90r^{17}ql^2 + 33r^{17}l^3 + 145r^{16}q^4 + 489r^{16}q^3l + 774r^{16}q^2l^2 + 489r^{16}ql^3 + 145r^{16}l^4 \\
 & + 406r^{15}q^5 + 1956r^{15}q^4l + 3912r^{15}q^3l^2 + 3912r^{15}q^2l^3 + 1956r^{15}ql^4 + 406r^{15}l^5 \\
 & + 1032r^{14}q^6 + 5832r^{14}q^5l + 14724r^{14}q^4l^2 + 19416r^{14}q^3l^3 + 14724r^{14}q^2l^4 \\
 & + 5832r^{14}ql^5 + 1032r^{14}l^6 + 1980r^{13}q^7 + 13608r^{13}q^6l + 40824r^{13}q^5l^2 \\
 & + 67956r^{13}q^4l^3 + 67956r^{13}q^3l^4 + 40824r^{13}q^2l^5 + 13608r^{13}ql^6 + 1980r^{13}l^7 \\
 & + 3260r^{12}q^8 + 25236r^{12}q^7l + 88620r^{12}q^6l^2 + 176484r^{12}q^5l^3 + 221110r^{12}q^4l^4 \\
 & + 176484r^{12}q^3l^5 + 88620r^{12}q^2l^6 + 25236r^{12}ql^7 + 3260r^{12}l^8 + 4262r^{11}q^9 \\
 & + 37854r^{11}q^8l + 151416r^{11}q^7l^2 + 352968r^{11}q^6l^3 + 529452r^{11}q^5l^4 + 529452r^{11}q^4l^5 \\
 & + 352968r^{11}q^3l^6 + 151416r^{11}q^2l^7 + 37854r^{11}ql^8 + 4262r^{11}l^9 + 4752r^{10}q^{10} \\
 & + 46252r^{10}q^9l + 208512r^{10}q^8l^2 + 554520r^{10}q^7l^3 + 971292r^{10}q^6l^4 + 1164342r^{10}q^5l^5 \\
 & + 971292r^{10}q^4l^6 + 554520r^{10}q^3l^7 + 208512r^{10}q^2l^8 + 46252r^{10}ql^9 + 4752r^{10}l^{10} \\
 & + 4262r^9q^{11} + 46252r^9q^{10}l + 231260r^9q^9l^2 + 693150r^9q^8l^3 + 1386300r^9q^7l^4 \\
 & + 1940568r^9q^6l^5 + 1940568r^9q^5l^6 + 1386300r^9q^4l^7 + 693150r^9q^3l^8 + 231260r^9q^2l^9 \\
 & + 46252r^9ql^{10} + 4262r^9l^{11} + 3260r^8q^{12} + 37854r^8q^{11}l + 208512r^8q^{10}l^2 \\
 & + 693150r^8q^9l^3 + 1560534r^8q^8l^4 + 2494836r^8q^7l^5 + 2912112r^8q^6l^6 + 2494836r^8q^5l^7 \\
 & + 1560534r^8q^4l^8 + 693150r^8q^3l^9 + 208512r^8q^2l^{10} + 37854r^8ql^{11} + 3260r^8l^{12} \\
 & + 1980r^7q^{13} + 25236r^7q^{12}l + 151416r^7q^{11}l^2 + 554520r^7q^{10}l^3 + 1386300r^7q^9l^4 \\
 & + 2494836r^7q^8l^5 + 3326448r^7q^7l^6 + 3326448r^7q^6l^7 + 2494836r^7q^5l^8 + 1386300r^7q^4l^9 \\
 & + 554520r^7q^3l^{10} + 151416r^7q^2l^{11} + 25236r^7ql^{12} + 1980r^7l^{13} + 1032r^6q^{14} \\
 & + 13608r^6q^{13}l + 88620r^6q^{12}l^2 + 352968r^6q^{11}l^3 + 971292r^6q^{10}l^4 + 1940568r^6q^9l^5 \\
 & + 2912112r^6q^8l^6 + 3326448r^6q^7l^7 + 2912112r^6q^6l^8 + 1940568r^6q^5l^9 + 971292r^6q^4l^{10} \\
 & + 352968r^6q^3l^{11} + 88620r^6q^2l^{12} + 13608r^6ql^{13} + 1032r^6l^{14} + 406r^5q^{15} + 5832r^5q^{14}l \\
 & + 40824r^5q^{13}l^2 + 176484r^5q^{12}l^3 + 529452r^5q^{11}l^4 + 1164342r^5q^{10}l^5 + 1940568r^5q^9l^6 \\
 & + 2494836r^5q^8l^7 + 2494836r^5q^7l^8 + 1940568r^5q^6l^9 + 1164342r^5q^5l^{10} + 529452r^5q^4l^{11} \\
 & + 176484r^5q^3l^{12} + 40824r^5q^2l^{13} + 5832r^5ql^{14} + 406r^5l^{15} + 145r^4q^{16} + 1956r^4q^{15}l \\
 & + 14724r^4q^{14}l^2 + 67956r^4q^{13}l^3 + 221110r^4q^{12}l^4 + 529452r^4q^{11}l^5 + 971292r^4q^{10}l^6 \\
 & + 1386300r^4q^9l^7 + 1560534r^4q^8l^8 + 1386300r^4q^7l^9 + 971292r^4q^6l^{10} + 529452r^4q^5l^{11} \\
 & + 221110r^4q^4l^{12} + 67956r^4q^3l^{13} + 14724r^4q^2l^{14} + 1956r^4ql^{15} + 145r^4l^{16} + 33r^3q^{17} \\
 & + 489r^3q^{16}l + 3912r^3q^{15}l^2 + 19416r^3q^{14}l^3 + 67956r^3q^{13}l^4 + 176484r^3q^{12}l^5 \\
 & + 352968r^3q^{11}l^6 + 554520r^3q^{10}l^7 + 693150r^3q^9l^8 + 693150r^3q^8l^9 + 554520r^3q^7l^{10} \\
 & + 352968r^3q^6l^{11} + 176484r^3q^5l^{12} + 67956r^3q^4l^{13} + 19416r^3q^3l^{14} + 3912r^3q^2l^{15} \\
 & + 489r^3ql^{16} + 33r^3l^{17} + 10r^2q^{18} + 90r^2q^{17}l + 774r^2q^{16}l^2 + 3912r^2q^{15}l^3 \\
 & + 14724r^2q^{14}l^4 + 40824r^2q^{13}l^5 + 88620r^2q^{12}l^6 + 151416r^2q^{11}l^7 + 208512r^2q^{10}l^8 \\
 & + \mathbf{231260r^2q^9l^9} + 208512r^2q^8l^{10} + 151416r^2q^7l^{11} + 88620r^2q^6l^{12} + 40824r^2q^5l^{13} \\
 & + 14724r^2q^4l^{14} + 3912r^2q^3l^{15} + 774r^2q^2l^{16} + 90r^2ql^{17} + 10r^2l^{18} + rq^{19} + 10rq^{18}l \\
 & + 90rq^{17}l^2 + 489rq^{16}l^3 + 1956rq^{15}l^4 + 5832rq^{14}l^5 + 13608rq^{13}l^6 + 25236rq^{12}l^7 \\
 & + 37854rq^{11}l^8 + 46252rq^{10}l^9 + 46252rq^9l^{10} + 37854rq^8l^{11} + 25236rq^7l^{12} \\
 & + 13608rq^6l^{13} + 5832rq^5l^{14} + 1956rq^4l^{15} + 489rq^3l^{16} + 90rq^2l^{17} + 10rq^{18} + rl^{19} \\
 & + q^{20} + q^{19}l + 10q^{18}l^2 + 33q^{17}l^3 + 145q^{16}l^4 + 406q^{15}l^5 + 1032q^{14}l^6 + 1980q^{13}l^7 \\
 & + 3260q^{12}l^8 + 4262q^{11}l^9 + 4752q^{10}l^{10} + 4262q^9l^{11} + 3260q^8l^{12} + 1980q^7l^{13} \\
 & + 1032q^6l^{14} + 406q^5l^{15} + 145q^4l^{16} + 33q^3l^{17} + 10q^2l^{18} + ql^{19} + l^{20}
 \end{aligned}$$

FIGURE 2.7. Pattern inventory for necklaces with twenty beads formed using three kinds of beads.

Pólya's enumeration formula has many applications in several fields, including chemistry, physics, and computer science. Pólya devotes a large portion of his paper [224] to applications involving enumeration of graphs, trees, and chemical isomers.

Exercises

1. What is the pattern inventory for coloring n objects using the m colors y_1, y_2, \dots, y_m if the group of symmetries is S_n ?
2. Use Pólya's enumeration formula to determine the number of six-sided dice that can be manufactured if each of three different labels must be placed on two of the faces.
3. The hydrocarbon benzene has six carbon atoms arranged at the vertices of a regular hexagon, and six hydrogen atoms, with one bonded to each carbon atom. Two molecules are said to be *isomers* if they are composed of the same number and types of atoms, but have different structure.
 - (a) Show that exactly three isomers (ortho-dichlorobenzene, meta-dichlorobenzene, and para-dichlorobenzene) may be constructed by replacing two of the hydrogen atoms of benzene with chlorine atoms.
 - (b) How many isomers may be obtained by replacing two of the hydrogen atoms with chlorine atoms, and two others with bromine atoms?
4. The hydrocarbon naphthalene has ten carbon atoms arranged in a double hexagon as in Figure 2.8, and eight hydrogen atoms attached at each of the positions labeled 1 through 8.




FIGURE 2.8. Naphthalene.

- (a) Naphthol is obtained by replacing one of the hydrogen atoms of naphthalene with a hydroxyl group (OH). How many isomers of naphthol are there?
- (b) Tetramethylnaphthalene is obtained by replacing four of the hydrogen atoms of naphthalene with methyl groups (CH_3). How many isomers of tetramethylnaphthalene are there?

- (c) How many isomers may be constructed by replacing three of the hydrogen molecules of naphthalene with hydroxyl groups, and another three with methyl groups?
- (d) How many isomers may be constructed by replacing two of the hydrogen molecules of naphthalene with hydroxyl groups, two with methyl groups, and two with carboxyl groups (COOH)?
5. The hydrocarbon anthracene has fourteen carbon atoms arranged in a triple hexagon as in Figure 2.9, with ten hydrogen atoms bonded at the numbered positions.




FIGURE 2.9. Anthracene.

- (a) How many isomers of trimethylantracene can be formed by replacing three hydrogen atoms with methyl groups?
- (b) How many isomers can be formed by replacing four of the hydrogen atoms with chlorine, and two others with hydroxyl groups?
6. The molecule triphenylamine has three rings of six carbon atoms attached to a central nitrogen atom, as in Figure 2.10, and fifteen hydrogen atoms, with one attached to each carbon atom except the three carbons attached to the central nitrogen atom.




FIGURE 2.10. Triphenylamine.

- (a) How many isomers can be formed by replacing six hydrogen atoms with hydroxyl groups?

- (b) How many isomers can be formed by replacing five hydrogen atoms with methyl groups, and five with fluorine atoms?
7. The hydrocarbon tetraphenylmethane consists of four rings of six carbon atoms, each bonded to a central carbon atom, as in Figure 2.11, together with twenty hydrogen atoms, with one hydrogen atom attached to each carbon atom in the rings except for those attached to the carbon at the center.




FIGURE 2.11. Tetraphenylmethane.

- (a) How many isomers can be formed by replacing five hydrogen atoms of tetraphenylmethane with chlorine?
- (b) How many isomers can be formed by replacing five hydrogen atoms with bromine, and six others with hydroxyl groups?
8. Suppose a medical relief agency plans to design a symbol for their organization in the shape of a regular cross, as in Figure 2.12. To symbolize the purpose of the organization and emphasize its international constituency, its board of directors decides that the cross should be white in color, with each of the twelve line segments outlining the cross colored red, green, blue, or yellow, with an equal number of lines of each color. If we discount rotations and flips, how many different ways are there to design the symbol?




FIGURE 2.12. Symbol of a relief agency.

2.7.5 de Bruijn's Generalization

*It doesn't matter what color, well that gets a nope!
Be it pink, purple, or heliotrope!*

— *Boundin'*, Pixar Films

Suppose a jewelry company plans to market a new line of unisex bracelets under the brand name *OPPOSITES ATTRACT*. The bracelets are sold in pairs, for a couple to share. Each bracelet consists of n beads, some gold and some silver, and the two bracelets in a pair are opposites, in the sense that one can be obtained from the other by changing each silver bead to a gold one and each gold to a silver. For example, if one bracelet has two adjacent gold beads and $n - 2$ silver beads, then its mate has two adjacent silver beads and $n - 2$ gold beads. The companion then of the all-gold bracelet is the all-silver one. How many different pairs of n -bead bracelets are possible in the *OPPOSITES ATTRACT* product line?

We have seen that there are exactly six different bracelets for the case $n = 4$, if we discount both rotations and flips. These are represented by the configurations $gggg$, $gggs$, $ggss$, $gsgs$, $gsss$, and $ssss$ of gold and silver beads. This produces just four different (unordered) pairs of bracelets for the product line when $n = 4$:

$$gggg + ssss, \quad gggs + gsss, \quad ggss + ggss, \quad gsgs + gsgs. \quad (2.63)$$

Recall that each of the configurations we listed for $n = 4$ in fact represent an equivalence class of the set of two-colorings of the vertices of a square, where we consider two colorings to be equivalent if one can be obtained from the other by the action of some element of the group of symmetries of the square, D_4 . In the same way, we may consider each of the pairs of bracelets in our product line as representing a single set of two-colorings of the square—the union of the equivalence classes of the two bracelets in the set. For example, the four pairs listed in (2.63) correspond to the following partition of the sixteen ways to color the vertices of a square using at most two colors:

$$\begin{aligned} &\{gggg, ssss\}, \\ &\{gggs, ggsg, gsgg, sggs, sssg, ssgs, sgss, gsss\}, \\ &\{ggss, gssg, ssgg, sggs\}, \\ &\{gsgs, sgsg\}. \end{aligned}$$

This partition is precisely the collection of equivalence classes of two-colorings under a different equivalence relation. Now we consider two colorings to be equivalent if one can be obtained from the other by first performing some geometric transformation corresponding to a symmetry of the bracelet, then possibly inverting all the colors. It is easy to check that this is indeed an equivalence relation.

We can generalize this problem in the following way. Given a set of objects S , a set of colors R , a group G acting on S , and a group H acting on R . Let C denote the set of colorings of S using the colors in R , so this is the set of all functions from S into R . We consider two colorings in C to be equivalent if one can be

obtained from the other by first applying a permutation from G on the objects, then applying a permutation from H on the colors. Exercise 1 asks you to verify that this does in fact form an equivalence relation on C . We would like to know the number of equivalence classes of C with respect to G and H .

In our example with four-bead bracelets, we have that S is the set of vertices of a square, $R = \{g, s\}$ for gold and silver beads, $G = D_4$, and $H = S_2$, since we may either leave the beads unchanged, or swap them. A permutation $\pi \in G$ induces a map π^* on C in the usual way. For instance, if π is the 90-degree rotation (1234) , then $\pi^*(gggs) = sggg$. Similarly, a permutation $\rho \in H$ induces a map ρ^* on C . For example, if $\rho = (12)$ then $\rho^*(gggs) = sssg$.

Of course, if H is the trivial group consisting only of the identity permutation (1) , then we can use the cycle index and the enumeration formula of Pólya to determine the answer. The Dutch mathematician Nicolaas Govert de Bruijn generalized the method of Pólya for arbitrary color groups H , and we describe this theory here. The first step is computing the set of equivalence classes of C with respect to the object group G which are invariant with respect to a given permutation of the colors. Our proof follows de Bruijn's paper [69].

Theorem 2.11. *Suppose S is a set of n objects, $R = \{y_1, \dots, y_m\}$ is a set of m colors, G is a subgroup of the symmetric group S_n , and $\rho \in S_m$. Let $P_G(\mathbf{x})$ denote the cycle index of G . Then the pattern inventory for the colorings of S which are nonequivalent with respect to the action of G on S , but invariant with respect to the action of ρ on R , is*

$$F_{G,\rho}(\mathbf{y}) = P_G(\alpha_1(\rho), \alpha_2(\rho), \dots, \alpha_n(\rho)), \tag{2.64}$$

where

$$\alpha_k(\rho) = \sum_{\rho^k(j)=j} \prod_{i=0}^{k-1} y_{\rho^i(j)}$$

for $1 \leq k \leq n$.

Proof. Let C denote the set of all colorings of S , so C is the set of maps from S into R . For a particular coloring $c \in C$, let \bar{c} denote its orbit with respect to the group G , so $\bar{c} = \{\pi^*(c) : \pi \in G\}$. Also, let $\mathbf{v}(c) = (n_1, n_2, \dots, n_m)$, where for each i the integer n_i records the number of elements of S assigned the color y_i in c , and let $\mathbf{y}^{\mathbf{v}(c)}$ denote the monomial $y_1^{n_1} y_2^{n_2} \dots y_m^{n_m}$. Since $\mathbf{v}(\pi^*(c)) = \mathbf{v}(c)$ for any $\pi \in G$, we may define $\mathbf{y}^{\mathbf{v}(\bar{c})}$ by $\mathbf{y}^{\mathbf{v}(\bar{c})} = \mathbf{y}^{\mathbf{v}(c)}$.

Suppose that c is invariant under the action of ρ , that is $\rho^*(c) \in \bar{c}$. Since we want to find the pattern inventory for the classes of colorings of S that are nonequivalent with respect to G , but invariant with respect to ρ , we need to study the generating function

$$F_{G,\rho}(\mathbf{y}) = \sum_{\rho(\bar{c})=\bar{c}} \mathbf{y}^{\mathbf{v}(\bar{c})}.$$

Since G is a group, it is straightforward to show that the set of all colorings that are invariant under ρ is the union of all the orbits \bar{c} where $\rho(\bar{c}) = \bar{c}$. Thus, using

Lemma 2.8 we find that

$$F_{G,\rho}(\mathbf{y}) = \sum_{\rho^*(c) \in \bar{c}} \frac{\mathbf{y}^{\mathbf{v}(c)}}{|\bar{c}|} = \frac{1}{|G|} \sum_{\rho^*(c) \in \bar{c}} |G_c| \mathbf{y}^{\mathbf{v}(c)},$$

where G_c is the stabilizer of c in G . Now since $\rho^*(c) \in \bar{c}$, there exists a permutation $\pi_c \in G$ such that $\rho^*(c) = \pi_c^*(c)$. Also, the set of permutations $\{\pi_c \circ \pi : \pi \in G_c\}$ is exactly the same as the set $\{\pi \in G : \pi^*(c) = \pi_c^*(c)\}$, so $|G_c|$ equals the number of permutations in G which have the same effect as ρ on c :

$$|G_c| = |\{\pi \in G : \pi^*(c) = \rho^*(c)\}|.$$

Let U_π denote the set of colorings c for which π and ρ have the same effect,

$$U_\pi = \{c \in C : \pi^*(c) = \rho^*(c)\}.$$

Note that if $c \in U_\pi$ then automatically $\rho^*(c) \in \bar{c}$. Thus, we find that

$$F_{G,\rho}(\mathbf{y}) = \frac{1}{|G|} \sum_{\pi \in G} \sum_{c \in U_\pi} \mathbf{y}^{\mathbf{v}(c)}. \tag{2.65}$$

Now suppose $\pi \in G$, and π has λ_i cycles of length i , for each i with $1 \leq i \leq n$. Let ℓ_i denote the length of the i th cycle (when π is written in cycle notation in the canonical way), and let s_i denote the smallest element of the i th cycle. For example, if $n = 7$ and $\pi = (1245)(37)(6)$, then $\lambda_1 = \lambda_2 = \lambda_4 = 1$, $\ell_1 = 4$, $\ell_2 = 2$, $\ell_3 = 1$, $s_1 = 1$, $s_2 = 3$, and $s_3 = 6$. Also, let $M_\pi(x_1, \dots, x_n)$ denote the monomial obtained from π as in (2.56), so in the example we have $M_\pi(x_1, \dots, x_7) = x_1x_2x_4$.

Suppose that $c \in U_\pi$, so that applying π to c has the same effect as applying ρ . If position s_i has color y_j in c , it follows that position $\pi^{-1}(s_i)$ has color $y_{\rho(j)}$, position $\pi^{-2}(s_i)$ has color $y_{\rho^2(j)}$, ..., position $\pi^{-(\ell_i-1)}(s_i)$ has color $y_{\rho^{(\ell_i-1)}(j)}$, and we require that $\rho^{\ell_i}(j) = j$. It therefore follows that

$$\sum_{c \in U_\pi} \mathbf{y}^{\mathbf{v}(c)} = M_\pi(\alpha_1(\rho), \alpha_2(\rho), \dots, \alpha_n(\rho)),$$

and the theorem follows by combining this with (2.57) and (2.65). □

We can apply this theorem to our original example on bracelets, where $n = 4$, $m = 2$, $G = D_4$, and $\rho = (12)$. Write $y_1 = g$ for a gold bead, and $y_2 = s$ for a silver one. Then $\alpha_1(\rho) = 0$, since no color is left unchanged by ρ . Next, $\alpha_2(\rho) = y_1y_2 + y_2y_1 = 2gs$, since $\rho^2(j) = j$ for both $j = 1$ and $j = 2$. We then find that $\alpha_3(\rho) = 0$ since $\rho^3(1) = 2$ and $\rho^3(2) = 1$, and $\alpha_4(\rho) = y_1y_2y_1y_2 + y_2y_1y_2y_1 = 2g^2s^2$. Using (2.58), we obtain then that

$$F_{D_4,(12)}(g, s) = P_{D_4}(0, 2gs, 0, 2g^2s^2) = 2g^2s^2,$$

and we verify that there are indeed just two four-bead bracelets which are invariant under bead swapping, discounting rotations and flips: $gsgs$ and $ggss$.

If we introduce another type of bead in this example, say $y_3 = b$ for bronze, and keep $\rho = (12)$, then we obtain $\alpha_1(\rho) = b$, $\alpha_2(\rho) = 2gs + b^2$, $\alpha_3(\rho) = b^3$, and $\alpha_4(\rho) = 2g^2s^2 + b^4$, and we calculate that

$$F_{D_4,(12)}(g, s, b) = P_{D_4}(b, 2gs + b^2, b^3, 2g^2s^2 + b^4) = 2g^2s^2 + 2gsb^2 + b^4.$$

The five different configurations in this case are represented by $ggss$, $gsgs$, $gbsb$, $gsbb$, and $bbbb$.

We may now use Theorem 2.11 to solve our original problem. We would like to obtain the pattern inventory for a set of colorings C when we account for both a group of symmetries G on the objects, and a group H of symmetries on the colors. We compute this pattern inventory by averaging the patterns $F_{G,\rho}(\mathbf{y})$ over all permutations ρ in H , then combining the terms that correspond to equivalent patterns of colors.

Theorem 2.12 (de Bruijn’s Enumeration Formula). *Suppose S is a set of n objects, $R = \{y_1, \dots, y_m\}$ is a set of m colors, G is a subgroup of the symmetric group S_n , and H is a subgroup of S_m . Then the pattern inventory $\hat{F}_{G,H}(\mathbf{y})$ for the colorings of S which are nonequivalent with respect to both the action of G on S and the action of H on R is obtained by identifying equivalent color patterns in the polynomial*

$$F_{G,H}(\mathbf{y}) = \frac{1}{|H|} \sum_{\rho \in H} F_{G,\rho}(\mathbf{y}), \tag{2.66}$$

where $F_{G,\rho}(\mathbf{y})$ is given by (2.64).

We describe one example before providing the proof. With $n = 4$, $m = 2$, $R = \{g, s\}$, $G = D_4$, and $H = S_2$, we compute

$$\begin{aligned} F_{D_4,S_2}(g, s) &= \frac{1}{2}(P_{D_4}(g + s, g^2 + s^2, g^3 + s^3, g^4 + s^4) \\ &\quad + P_{D_4}(0, 2gs, 0, 2g^2s^2)) \\ &= \frac{1}{2}(g^4 + s^4) + \frac{1}{2}(g^3s + gs^3) + 2g^2s^2. \end{aligned} \tag{2.67}$$

The color patterns g^4 and s^4 are equivalent under the color group $H = S_2$, so we let $[g^4]$ denote either one of these patterns. Likewise, we let $[g^3s]$ denote either of the equivalent patterns g^3s or gs^3 . The last pattern, g^2s^2 , is not equivalent to any of the others, so we let $[g^2s^2]$ designate this single pattern. We obtain the pattern inventory by combining the equivalent terms:

$$\hat{F}_{D_4,S_2}(g, s) = [g^4] + [g^3s] + 2[g^2s^2].$$

Proof of Theorem 2.12. Let C denote the set of all colorings of S using the colors of R , and let \bar{c} denote the orbit of the coloring c under the action of G , so $\bar{c} = \{\pi^*(c) : \pi \in G\}$. The group H acts on the set of equivalence classes $\{\bar{c} : c \in C\}$, and we let $\bar{\bar{c}}$ denote the orbit of \bar{c} under this action, so

$$\bar{\bar{c}} = \{\rho^*(\bar{c}) : \rho \in H\}.$$

In the example above, if $c = gggs$, then

$$\bar{c} = \{gggs, ggs g, g s g g, s g g g\}$$

and

$$\bar{\bar{c}} = \{\{gggs, ggs g, g s g g, s g g g\}, \{s s s g, s s g s, s g s s, g s s s\}\};$$

if $c = gsgs$, then $\bar{c} = \{gsgs, sgs g\}$ and $\bar{\bar{c}} = \{\{gsgs, sgs g\}\}$.

Employing the notation we introduced in the proof of Theorem 2.9, and using Lemma 2.8, we compute

$$\begin{aligned} \frac{1}{|H|} \sum_{\rho \in H} F_{G, \rho}(\mathbf{y}) &= \frac{1}{|H|} \sum_{\rho \in H} \sum_{\bar{c}} [\rho^*(\bar{c}) = \bar{c}] \mathbf{y}^{\mathbf{v}(\bar{c})} \\ &= \frac{1}{|H|} \sum_{\bar{c}} \mathbf{y}^{\mathbf{v}(\bar{c})} \sum_{\rho \in H} [\rho^*(\bar{c}) = \bar{c}] \\ &= \frac{1}{|H|} \sum_{\bar{c}} |H_{\bar{c}}| \mathbf{y}^{\mathbf{v}(\bar{c})} \tag{2.68} \\ &= \sum_{\bar{c}} |\bar{c}|^{-1} \mathbf{y}^{\mathbf{v}(\bar{c})} \\ &= \sum_{\bar{\bar{c}}} |\bar{\bar{c}}|^{-1} \sum_{\bar{c} \in \bar{\bar{c}}} \mathbf{y}^{\mathbf{v}(\bar{c})}. \end{aligned}$$

Since the color patterns in the set $\{\mathbf{y}^{\mathbf{v}(\bar{c})} : \bar{c} \in \bar{\bar{c}}\}$ are equivalent under H , we select one pattern from this set to represent the class $\bar{\bar{c}}$, and denote this equivalence class of patterns by $[\mathbf{y}^{\mathbf{v}(\bar{c})}]$. By replacing each term $\mathbf{y}^{\mathbf{v}(\bar{c})}$ in the last line of (2.68) by its representative class $[\mathbf{y}^{\mathbf{v}(\bar{c})}]$, we obtain the pattern inventory,

$$\hat{F}_{G, H}(\mathbf{y}) = \sum_{\bar{\bar{c}}} [\mathbf{y}^{\mathbf{v}(\bar{c})}]. \quad \square$$

We can use Theorem 2.12 to determine the number different ten-bead pairs of bracelets in the OPPOSITES ATTRACT product line having a given configuration of colors. Since

$$P_{D_{10}}(\mathbf{x}) = \frac{1}{2} (x_1^{10} + x_2^5 + 4x_5^2 + 4x_{10} + 5x_2^5 + 5x_1^2 x_2^4),$$

we compute

$$\begin{aligned} F_{D_{10}, S_2}(g, s) &= \frac{1}{2} (P_{D_{10}}(g + s, g^2 + s^2, \dots, g^{10} + s^{10}) \\ &\quad + P_{D_{10}}(0, 2gs, \dots, 0, 2g^5 s^5)) \\ &= \frac{1}{2} (g^{10} + s^{10}) + \frac{1}{2} (g^9 s + g s^9) + \frac{5}{2} (g^8 s^2 + g^2 s^8) \\ &\quad + 4(g^7 s^3 + g^3 s^7) + 8(g^6 s^4 + g^4 s^6) + 13g^5 s^5, \end{aligned}$$

so the pattern inventory for these pairs of bracelet is

$$\hat{F}_{D_{10}, S_2}(g, s) = [g^{10}] + [g^9 s] + 5[g^8 s^2] + 8[g^7 s^3] + 16[g^6 s^4] + 13[g^5 s^5].$$

The situation is much simpler if we need only compute the total number of distinct colorings with respect to G and H , and we do not need the finer information provided by the pattern inventory. For this case, we need only set each $y_i = 1$ in $F_{G,H}(\mathbf{y})$, and there is no need to compute $\hat{F}_{G,H}(\mathbf{y})$. Since this case is so common, we describe its solution as a corollary to Theorem 2.12. Its proof is left as an exercise.

Corollary 2.13. *Suppose S is a set of n objects, R is a set of m colors, G is a subgroup of the symmetric group S_n , and H is a subgroup of S_m . Then the number of colorings of S using the colors in R which are nonequivalent with respect to both the action of G on S and the action of H on R is*

$$N_{G,H}(n, m) = \frac{1}{|H|} \sum_{\rho \in H} P_G(\beta_1(\rho), \beta_2(\rho), \dots, \beta_n(\rho)), \tag{2.69}$$

where $\beta_k(\rho) = \sum_{j|k} j \lambda_j(\rho)$, with the sum extending over all the positive divisors j of k , and $\lambda_j(\rho)$ is the number of cycles of ρ of length j .

For example, for our ten-bead bracelet problem with $m = 2$ and $H = S_2$, we find that the only nonzero values of the $\lambda_j(\rho)$ are $\lambda_1((1)(2)) = 2$ and $\lambda_2((12)) = 1$. It follows that $\beta_k((1)(2)) = 2$ for $1 \leq k \leq 10$, and

$$\beta_k((12)) = \begin{cases} 2 & \text{if } k \text{ is even,} \\ 0 & \text{if } k \text{ is odd.} \end{cases}$$

Therefore,

$$N_{D_{10}, S_2}(10, 2) = \frac{1}{2}(P_{D_{10}}(2, 2, \dots, 2) + P_{D_{10}}(0, 2, \dots, 0, 2)) = 44.$$

Last, we return to the problem from earlier sections concerning twenty-bead necklaces using rhodonite, rose quartz, and lapis lazuli beads. Using $H = \langle(123)\rangle$, we find that

$$\beta_k((123)) = \beta_k((132)) = \begin{cases} 3 & \text{if } 3 \mid k, \\ 0 & \text{if } 3 \nmid k, \end{cases}$$

and $\beta_k((1)) = 3$ for each k . Thus,

$$\begin{aligned} N_{D_{20}, C_3}(20, 3) &= \frac{1}{3}(P_{D_{20}}(3, \dots, 3) + 2P_{D_{20}}(0, 0, 3, \dots, 0, 0, 3, 0, 0)) \\ &= \frac{1}{3}P_{D_{20}}(3, \dots, 3) = 29\,076\,719, \end{aligned}$$

since none of the variables x_{3k} appears in $P_{D_{20}}(\mathbf{x})$. This then is the number of different 20-bead necklaces if we discount rotations, flips, and the bead substitutions rhodonite \rightarrow rose quartz \rightarrow lapis lazuli \rightarrow rhodonite, or rhodonite \rightarrow lapis lazuli \rightarrow rose quartz \rightarrow rhodonite.

Using $H = S_3$ instead, we obtain

$$\beta_k((12)) = \beta_k((13)) = \beta_k((23)) = \begin{cases} 1 & \text{if } k \text{ is odd,} \\ 3 & \text{if } k \text{ is even,} \end{cases}$$

and so

$$\begin{aligned} N_{D_{20}, S_3}(20, 3) &= \frac{1}{6}(P_{D_{20}}(3, \dots, 3) + 2P_{D_{20}}(0, 0, 3, \dots, 0, 0, 3, 0, 0) \\ &\quad + 3P_{D_{20}}(1, 3, \dots, 1, 3)) \\ &= \frac{1}{6}(87\,230\,157 + 63\,519) = 14\,548\,946. \end{aligned}$$

This is therefore the number of different 20-bead necklaces if we discount rotations, flips, and any permutation of the bead types.

Exercises

1. Suppose that a group G acts on a set S of objects, and a group H acts on a set R of colors. Let C denote the set of functions from S into R , that is, the number of colorings of S using the colors in R . If c_1 and c_2 are two colorings in C , write $c_1 \sim c_2$ if there exists an element $g \in G$ and an element $h \in H$ such that applying g to the underlying objects of c_1 , then h to its colors, produces c_2 . Show that \sim induces an equivalence relation on C .
2. Suppose that G is a group acting on a set of objects S , and that C is the set of colorings of elements of S using the colors in a set R . Let \bar{c} denote the orbit of c in C with respect to the action of G . Let ρ be a permutation acting on R . Prove that $\{c \in C : \rho^*(c) \in \bar{c}\} = \bigcup_{\rho^*(\bar{c})=\bar{c}} \bar{c}$.
3. Compute the number of different pairs of bracelets in the OPPOSITES ATTRACT product line for $n = 6$, $n = 7$, and $n = 8$.
4. Our jewelry company plans to extend their line of bracelets by introducing sets of m bracelets formed using m different colors of beads, so that a set may be shared among a group of m people. If one bracelet in a package has the coloring c , then the others in the package have the coloring $\rho^*(c)$, $(\rho^*)^2(c)$, \dots , $(\rho^*)^{m-1}(c)$, where ρ is the cyclic permutation $(1\,2\,\dots\,m)$. Use D_n for the object group G in each of the following problems.
 - (a) Compute the number of different packages of bracelets for $m = 3$ when $n = 6$, then $n = 7$, then $n = 9$.
 - (b) Compute the number of different packages of bracelets for $m = 4$ when $n = 10$, then when $n = 12$.
 - (c) Determine the pattern inventory $\hat{F}_{D_n, C_m}(\mathbf{y})$ for the case $m = 3$ and $n = 6$, then $n = 9$.
5. Compute the pattern inventory $\hat{F}_{D_n, S_2}(x, y)$ for $n = 6$, $n = 7$, and $n = 8$.
6. Compute the pattern inventory $\hat{F}_{C_n, S_2}(x, y)$ for $n = 6$, $n = 7$, and $n = 8$.

7. Verify that the pattern inventory $\hat{F}_{D_{20}, S_3}(r, q, l)$ for 20-bead necklaces with three kinds of beads, using the full symmetric group S_3 for H , is

$$\begin{aligned} \hat{F}_{D_{20}, S_3}(r, q, l) = & [r^{20}] + [r^{19}q] + 10[r^{18}q^2] + 10[r^{18}ql] + 33[r^{17}q^3] \\ & + 90[r^{17}q^2l] + 145[r^{16}q^4] + 489[r^{16}q^3l] + 430[r^{16}q^2l^2] \\ & + 406[r^{15}q^5] + 1956[r^{15}q^4l] + 3912[r^{15}q^3l^2] + 1032[r^{14}q^6] \\ & + 5832[r^{14}q^5l] + 14724[r^{14}q^4l^2] + 9924[r^{14}q^3l^3] + 1980[r^{13}q^7] \\ & + 13608[r^{13}q^6l] + 40824[r^{13}q^5l^2] + 67956[r^{13}q^4l^3] + 3260[r^{12}q^8] \\ & + 25236[r^{12}q^7l] + 88620[r^{12}q^6l^2] + 176484[r^{12}q^5l^3] \\ & + 111270[r^{12}q^4l^4] + 4262[r^{11}q^9] + 37854[r^{11}q^8l] \\ & + 151416[r^{11}q^7l^2] + 352968[r^{11}q^6l^3] + 529452[r^{11}q^5l^4] \\ & + 2518[r^{10}q^{10}] + 46252[r^{10}q^9l] + 208512[r^{10}q^8l^2] \\ & + 554520[r^{10}q^7l^3] + 971292[r^{10}q^6l^4] + 583784[r^{10}q^5l^5] \\ & + 116398[r^9q^9l^2] + 693150[r^9q^8l^3] + 1386300[r^9q^7l^4] \\ & + 1940568[r^9q^6l^5] + 782141[r^8q^8l^4] + 2494836[r^8q^7l^5] \\ & + 1458578[r^8q^6l^6] + 1665912[r^7q^7l^6]. \end{aligned}$$

8. Prove Corollary 2.13.
9. Consider the symbol of the medical relief agency shown in Figure 2.12. Each of the twelve line segments outlining the cross shape must be colored red, green, blue, or yellow.
- How many ways are there to design the symbol, if we consider two configurations equivalent if one can be obtained from the other by some combination of a rotation, flip, and color reversal? A color reversal exchanges red and green, and exchanges blue and yellow.
 - How many of these configurations have the same number of edges of each color?
 - Repeat the first two problems, but this time consider two colorings to be equivalent if one can be obtained from the other by either exchanging red and green, or exchanging blue and yellow, or both.
 - Repeat the first two problems, but now consider two colorings to be equivalent if one can be obtained from the other by an iterate of the cyclic permutation red \rightarrow green \rightarrow blue \rightarrow yellow \rightarrow red.
 - Suppose now that black is added as a possible color for a segment of the border. How many ways are there to design the symbol, if we consider two configurations equivalent if one can be obtained from the other by some combination of a rotation, flip, and color reversal?

A color reversal exchanges red and green, exchanges blue and yellow, and leaves black fixed.

- (f) Repeat the previous problem, but this time consider two colorings to be equivalent if one can be obtained from the other by either exchanging red and green, or exchanging blue and yellow, or both.
10. Determine the number of ways to color the faces of a cube using the three colors maroon, cardinal, and burnt orange, if two colorings are considered to be equivalent if one can be obtained from the other by rotating the cube in some way in three-dimensional space, and possibly exchanging maroon and burnt orange. Then determine the number of such colorings in which maroon and burnt orange appear the same number of times.
11. Determine the number of ways to color the faces of an octahedron using the four colors heliotrope, lavender, thistle, and wisteria, if two colorings are considered to be equivalent if one can be obtained from the other by rotating the octahedron in some way, and possibly exchanging heliotrope and lavender, or thistle and wisteria, or both. Then determine the number of such colorings in which the number of faces colored heliotrope matches the number colored lavender, and at the same time the number of faces colored thistle matches the number colored wisteria.

2.8 More Numbers

Truly, I thought there had been one number more. . .

— William Shakespeare, *The Merry Wives of Windsor*,
Act IV, Scene I

Many questions in combinatorics can be answered by analyzing the number of ways to arrange a particular collection of objects into a number of bins, without regard to the order of placement. There are four basic kinds of problems of this form: The objects may be identical or distinguishable, and similarly for the bins. Problems of this form in combinatorics are called *occupancy problems*.

We have already studied occupancy problems for the case of distinguishable bins. If the objects are identical, then we saw in Section 2.6.2 that the number of ways to distribute n objects among k bins is the binomial coefficient $\binom{n+k-1}{n}$. This is the same as the number of ways to select n objects from a set of k different objects with repetition allowed, and we described the correspondence between these two problems in the earlier section. On the other hand, if the objects are distinguishable, then the number of ways to distribute n objects among k bins is simply k^n by the product rule, since each object can be placed in any of the bins. For example, consider the problem of determining the number of n -letter words that can be formed using an k -letter alphabet. We can model this as an occupancy problem by taking the integers between 1 and n as our objects, and the k letters of

the alphabet as our bins. Each placement of the objects in the bins corresponds to an n -letter word: The placement of 1 indicates the first letter, 2 the second letter, etc. Furthermore, it is clear that every possible n -letter word can be obtained in this way.

In subsequent sections, we consider some occupancy problems where the bins are indistinguishable. We call the bins *groups* or *piles* in this case, with the understanding that they are always unlabeled. The remaining two basic types of occupancy problems each produce important sequences of numbers in combinatorics. The problem of arranging a number *identical* objects into piles gives rise to *partitions*, which are studied in Section 2.8.1. The case of distributing a collection of *distinguishable* objects into groups produces the Stirling set numbers, discussed in Section 2.8.3, and the Bell numbers of Section 2.8.4. We also study two other important combinatorial sequences here: the Stirling cycle numbers in Section 2.8.2, and the Eulerian numbers in Section 2.8.5. Both of these are connected to the structure of permutations.

We study some important properties of each of these classes numbers, aided by generating functions. We also introduce some different kinds of generating functions to assist with our derivations. Some analysis illuminates for instance some interesting connections between ordinary powers, rising and falling factorial powers, and binomial coefficients.

2.8.1 Partitions

Whew! Don't try to eat these so-called chips!

— Homer Simpson, after choking during a poker game,
The Simpsons, episode 103, *Secrets of a Successful Marriage*

Suppose a winning hand in poker nets you a pot of n identical poker chips, and you want to organize your winnings into a number of neat stacks, in order to intimidate your opponents. Individual stacks are not labeled or distinguishable in any way, except for the number of chips they contain, so an arrangement of chips simply corresponds to a collection of positive numbers that sums to n . How many ways are there to organize your winnings?

An arrangement of n identical objects into a number of (unlabeled) piles is called a *partition* of the objects, so we want to know the number of partitions of the n objects, or, for short, the number of partitions of n . Let p_n denote this number. We might also investigate the number of ways to divide n identical objects into a specific number k of piles. Let $p_{n,k}$ denote this number. Since the piles are unlabeled, we can discount the possibility of an empty pile, so it follows that $p_n = p_{n,1} + p_{n,2} + \cdots + p_{n,n}$ for $n \geq 1$. For example, Figure 2.13 exhibits the fifteen ways to divide $n = 7$ poker chips into stacks. Thus $p_7 = 15$, and we see for instance that $p_{7,3} = 4$ and $p_{7,4} = 3$. Each configuration here is also displayed with a list showing the size of the stacks in descending order. We will always denote partitions in this way. It follows that we can define p_n as the number of ways to write n as a sum of positive integers, with the summands listed in descending




FIGURE 2.13. The fifteen ways to stack seven poker chips.

order. For example, the partitions of $n = 4$ are (4) , $(3, 1)$, $(2, 2)$, $(2, 1, 1)$, and $(1, 1, 1, 1)$.

We first note some particular values for the $p_{n,k}$. As a special case, we set

$$p_{0,k} = \begin{cases} 1 & \text{if } k = 0, \\ 0 & \text{if } k \neq 0, \end{cases} \tag{2.70}$$

so $p_0 = 1$. Also, we set $p_{n,k} = 0$ for all k if $n < 0$, so $p_n = 0$ if n is negative. For $n \geq 1$, certainly there is just one way to write n using a single summand, and just one way using n summands, so $p_{n,1} = p_{n,n} = 1$ for $n \geq 1$. Further, it is impossible to express a positive integer as a sum with zero terms, or more than n terms, or a negative number of terms, so we set

$$p_{n,k} = 0, \quad \text{if } k \leq 0 \text{ or } k > n. \tag{2.71}$$

Thus, for each integer n we have the identity

$$p_n = \sum_k p_{n,k}. \tag{2.72}$$

We can now derive a recurrence relation for $p_{n,k}$. Suppose that (a_1, \dots, a_k) is a partition of n , with the summands in descending order. If $a_k = 1$, then (a_1, \dots, a_{k-1}) is a partition of $n - 1$, and every partition of $n - 1$ into $k - 1$ parts can be obtained in this way. Thus, the number of partitions of n into k parts, where the smallest part is 1, is precisely $p_{n-1,k-1}$. Suppose then that $a_k \geq 2$. In this case, we see that $(a_1 - 1, \dots, a_k - 1)$ is a partition of $n - k$ into exactly k parts, and every partition of $n - k$ can be obtained in this way. It follows that the number of partitions of n into k parts, where the smallest part is at least 2, is $p_{n-k,k}$. Therefore, we find that

$$p_{n,k} = p_{n-1,k-1} + p_{n-k,k} \tag{2.73}$$

for $n \geq 1$. This recurrence relation, together with the initial condition $p_{0,0} = 1$, allows us to compute the value of $p_{n,k}$, for any n and k . A table of these values for $n \leq 10$ appears in Table 2.3.

$p_{n,k}$	$k = 0$	1	2	3	4	5	6	7	8	9	10	p_n
$n = 0$	1											1
1	0	1										1
2	0	1	1									2
3	0	1	1	1								3
4	0	1	2	1	1							5
5	0	1	2	2	1	1						7
6	0	1	3	3	2	1	1					11
7	0	1	3	4	3	2	1	1				15
8	0	1	4	5	5	3	2	1	1			22
9	0	1	4	7	6	5	3	2	1	1		30
10	0	1	5	8	9	7	5	3	2	1	1	42

TABLE 2.3. Number of partitions $p_{n,k}$ of n into k parts, and the number of partitions p_n of n .

We would like to determine a more efficient way of computing p_n , without using (2.73) to determine all of the $p_{n,k}$. In order to do this, we first introduce a useful way to visualize a partition known as a *Young diagram*. The Young diagram of a partition (a_1, \dots, a_k) of n consists of n boxes arranged in k rows, with a_1 boxes in the top row, a_2 boxes in the second row, and so on, and each row is aligned on the left. For example, Figure 2.14(a) illustrates the Young diagram for the partition $(6, 4, 4, 2, 1)$ of $n = 17$. This is then much like our stacks of poker chips of Figure 2.13, only turned sideways.

Many texts use arrays of dots instead of arrays of boxes for illustrating partitions, and in this case the diagrams are known as *Ferrers diagrams*. We find

the Young diagrams more convenient to use. (Young diagrams earned a distinct name due to their use in visualizing more complicated structures known as *Young tableaux*, where the boxes are filled with integers according to particular rules.)




FIGURE 2.14. The Young diagram for a partition λ , and its conjugate λ' .

We now define the *conjugate* λ' of a given partition λ of n as the partition of n obtained by counting the stacks of boxes in the columns of the Young diagram for λ . For example, that the conjugate partition of $\lambda = (6, 4, 4, 2, 1)$ in Figure 2.14(a) is $\lambda' = (5, 4, 3, 3, 1, 1)$. The diagram for λ' is displayed in Figure 2.14(b). Also, the conjugate of the partition of n that consists of all 1s is the trivial partition (n) .

Clearly, different partitions cannot have the same conjugate, and every partition is the conjugate of some partition, so the conjugation mapping is a permutation on the set of partitions of n . This fact is very useful in establishing properties of the numbers $p_{n,k}$ and p_n . For example, it is immediate that the number of partitions of n which have largest summand $a_1 = k$ is simply $p_{n,k}$, since conjugating the partitions with this property yields precisely the set of partitions of n into exactly k parts.

Next, we consider some generating functions. From our work on the money-changing problems of Section 2.6.3, we know that the generating function $P(x)$ for the sequence p_n is given by an infinite product,

$$P(x) = \prod_{k \geq 1} \frac{1}{1 - x^k}. \tag{2.74}$$

Let $\Phi(x) = 1/P(x)$, so

$$\Phi(x) = \prod_{k \geq 1} (1 - x^k). \tag{2.75}$$

Then $\Phi(x)$ is itself the generating function for some sequence $\{c_n\}$. If we imagine expanding enough terms of this product to determine c_n , we see that each partition (a_1, \dots, a_k) of n into *distinct* parts $a_1 > \dots > a_k$ contributes $(-1)^k$ to c_n , and these terms determine c_n . Define $q_e(n)$ to be the number of partitions of n into an even number of distinct parts, and let $q_o(n)$ be the number of partitions of n into an odd number of distinct parts. It follows that $c_n = q_e(n) - q_o(n)$, and so

$$\Phi(x) = \sum_{n \geq 0} (q_e(n) - q_o(n))x^n, \tag{2.76}$$

with the understanding that $q_e(0) = 1$ and $q_o(0) = 0$.

By expanding a number of terms of the product for $\Phi(x)$, we can compute the values of these coefficients up to $n = 100$:


$$\begin{aligned} \Phi(x) = & 1 - x - x^2 + x^5 + x^7 - x^{12} - x^{15} + x^{22} + x^{26} \\ & - x^{35} - x^{40} + x^{51} + x^{57} - x^{70} - x^{77} + x^{92} + x^{100} - \dots \end{aligned}$$

Thus, it appears that $q_e(n)$ and $q_o(n)$ are often equal, and moreover differ by at most 1. Euler first established this fact; the proof we exhibit here employs Young diagrams and is due to Franklin in 1881 [111]. The reason for the curious name of this theorem is explored in Exercise 7.

Theorem 2.14 (Euler’s Pentagonal Number Theorem). *Let n be a nonnegative integer, and let $q_e(n)$ and $q_o(n)$ be defined as above. Then*

$$q_e(n) - q_o(n) = \begin{cases} (-1)^k & \text{if } n = \frac{k(3k \pm 1)}{2}, \\ 0 & \text{otherwise.} \end{cases}$$

Proof. Let λ be a partition of n into distinct parts. Let $s(\lambda)$ denote the smallest part of λ , and let $r(\lambda)$ be the number of consecutive integers in λ , starting with its largest part. In the Young diagram for λ , the number of squares on the bottom row is $s(\lambda)$, and $r(\lambda)$ is the number of boxes in the diagram that lie on a 45° line anchored at the rightmost box. For example, Figure 2.15(a) exhibits a partition of $n = 23$ into five distinct parts. Here $r(\lambda) = 3$ and $s(\lambda) = 2$, and the relevant boxes for these quantities are marked respectively with x’s and +’s.



(a) $\lambda = (7, 6, 5, 3, 2)$.

(b) $\mu = (8, 7, 5, 3)$.

FIGURE 2.15. Constructing μ when $s(\lambda) \leq r(\lambda)$.

We aim to transform λ into another partition μ of n with distinct parts. The number of parts of μ will be either one more or one less than the number of parts of λ , so one of these two partitions will have an even number of parts, and the other will have an odd number. The transformation is described in terms of the Young diagram for λ , and depends on the relative sizes of $r(\lambda)$ and $s(\lambda)$.

If $s(\lambda) \leq r(\lambda)$, then we move the boxes in the bottom row of the Young diagram for λ to the ends of the top $s(\lambda)$ rows of the diagram. Figure 2.15(b) shows the resulting partition μ obtained from the partition λ of Figure 2.15(a). On the other hand, if $s(\lambda) > r(\lambda)$, then we move the rightmost boxes of the top $r(\lambda)$ rows

of the diagram for λ to make a new row at the bottom of the diagram. Figure 2.16 shows this procedure for $\lambda = (9, 7, 5, 2)$, yielding $\mu = (8, 7, 5, 2, 1)$.




FIGURE 2.16. Constructing μ when $s(\lambda) > r(\lambda)$.

The procedure for creating μ from λ fails in some special cases. The first case breaks down precisely when $s(\lambda) = r(\lambda)$ and the corresponding boxes in the Young diagram overlap, as in Figure 2.17(a). In this case, writing k for $r(\lambda)$, we compute that the total number of boxes in the diagram is

$$n = \sum_{j=k}^{2k-1} j = \frac{k(3k-1)}{2}.$$

The second case fails precisely when $s(\lambda) = r(\lambda) + 1$ and the boxes overlap, as in Figure 2.17(b). Again writing k for $r(\lambda)$, we find that

$$n = \sum_{j=k+1}^{2k} j = \frac{k(3k+1)}{2}$$

in this case.




FIGURE 2.17. Exceptional partitions.

Since our mapping on Young diagrams is its own inverse (see Exercise 6), it follows that it defines a bijection between the set of partitions of n into a distinct odd number of parts, and the set of partitions of n into a distinct even number of parts, provided that $n \neq k(3k \pm 1)/2$. When n is one of these exceptional values, there is exactly one additional partition with an even number of parts if k

is even, and exactly one extra partition into an odd number of parts if k is odd. The statement then follows. \square

By combining (2.76) with Theorem 2.14, we see that

$$\Phi(x) = 1 + \sum_{k \geq 1} (-1)^k \left(x^{k(3k-1)/2} + x^{k(3k+1)/2} \right), \tag{2.77}$$

and so

$$\begin{aligned} 1 &= P(x)\Phi(x) \\ &= \left(\sum_{k \geq 0} p_k x^k \right) \left(1 + \sum_{k \geq 1} (-1)^k \left(x^{k(3k-1)/2} + x^{k(3k+1)/2} \right) \right). \end{aligned} \tag{2.78}$$

It follows that the coefficient of x^n on the right side of (2.78) is 0 for $n \geq 1$. We therefore immediately obtain the following result.

Theorem 2.15. *Let n be a positive integer. Then*

$$p_n + \sum_{k \geq 1} (-1)^k \left(p_{n-k(3k-1)/2} + p_{n-k(3k+1)/2} \right) = 0,$$

that is,

$$p_n = p_{n-1} + p_{n-2} - p_{n-5} - p_{n-7} + p_{n-12} + p_{n-15} - \cdots \tag{2.79}$$

We can now use (2.79) to determine values of p_n without using the recurrence (2.73) for the $p_{n,k}$. For example, using the values of p_n computed in Table 2.3, we compute

$$p_{11} = p_{10} + p_9 - p_6 - p_4 = 42 + 30 - 11 - 5 = 56,$$

then

$$p_{12} = p_{11} + p_{10} - p_7 - p_5 + p_0 = 56 + 42 - 15 - 7 + 1 = 77,$$

and so on. Table 2.4 displays the values of p_n computed in this way up to $n = 50$, where

$$p_{50} = p_{49} + p_{48} - p_{45} - p_{43} + p_{38} + p_{35} - p_{28} - p_{24} + p_{15} + p_{10} = 204\,226.$$

We close this section with another interesting fact about the partition sequence. In 1918, Hardy and Ramanujan [152] established a remarkable nonrecursive formula for p_n as the value of a certain convergent series. Their formula was refined by Rademacher in 1937 [230]. We do not reproduce this formula here, but we mention only that it involves the number π , a certain complex root of the polynomial $x^{24} - 1$, and the hyperbolic sine function. From this formula, however, one

n	p_n	n	p_n	n	p_n	n	p_n
11	56	21	792	31	6842	41	44583
12	77	22	1002	32	8349	42	53174
13	101	23	1255	33	10143	43	63261
14	135	24	1575	34	12310	44	75175
15	176	25	1958	35	14883	45	89134
16	231	26	2436	36	17977	46	105558
17	297	27	3010	37	21637	47	124754
18	385	28	3718	38	26015	48	147273
19	490	29	4565	39	31185	49	173525
20	627	30	5604	40	37338	50	204226

TABLE 2.4. The number of partitions of n .

can obtain information on the rate of growth of the sequence p_n . Asymptotically, the number of partitions of n satisfies

$$p_n \sim \frac{e^{\pi\sqrt{2n/3}}}{4n\sqrt{3}}, \quad (2.80)$$

where $a_n \sim b_n$ means that $\lim_{n \rightarrow \infty} a_n/b_n = 1$. See the book by Andrews [9] for the details and a proof, as well as much more information on this rich topic.

Exercises

1. Establish formulas for $p_{n,2}$, $p_{n,n-1}$, and $p_{n,n-2}$.
2. Use (2.73) and Table 2.3 to compute the values of $p_{11,k}$, $p_{12,k}$, and $p_{13,k}$ for each k .
3. Use (2.79) and Table 2.4 to compute the value of p_{51} , then p_{52} .
4. Use Young diagrams to prove that $q_0(n)$ equals the number of partitions λ of n which are invariant under conjugation, that is, for which $\lambda = \lambda'$.
5. Use generating functions to prove that the number of partitions of n into distinct parts equals the number of partitions of n where each part is odd.
6. Suppose that λ is a partition of n , and that λ is not an exceptional partition like those shown in Figure 2.17. Let μ be the partition obtained by applying the procedure described in the proof of Theorem 2.14 on λ . Show that $r(\mu) < s(\mu)$ if and only if $r(\lambda) \geq s(\lambda)$. Then conclude that this map defines a permutation on the set of non-exceptional partitions of n into distinct parts, and that this permutation is its own inverse.
7. (a) Show that (2.77) may be written more simply as

$$\Phi(x) = \sum_k (-1)^k x^{k(3k-1)/2}.$$

- (b) The k th *pentagonal number* α_k is the number of disks in a pentagonal shape formed by stacking a triangular arrangement of $1 + 2 + \cdots + (k - 1)$ disks atop a square arrangement of $k \times k$ disks, as shown in Figure 2.18. Determine a closed formula for the k th pentagonal number. Why is Theorem 2.14 called the Pentagonal Number Theorem?



FIGURE 2.18. Pentagonal numbers: $\alpha_1 = 1$, $\alpha_2 = 5$, and $\alpha_3 = 12$.

8. Let $s_{n,k}$ denote the number of partitions of n whose smallest element is k , so $p_n = s_{n,1} + s_{n,2} + \cdots + s_{n,n}$. Prove that

$$s_{n,k} = \begin{cases} p_{n-1} & \text{if } k = 1, \\ s_{n-1,k-1} - s_{n-k,k-1} & \text{if } k \geq 2. \end{cases}$$

Then use this recurrence, together with the base values $s_{n,k} = 0$ for $k > n$ and $p_{0,0} = 0$, to produce a table of values for the $s_{n,k}$ for $1 \leq n \leq 10$, similar to Table 2.3.

9. Prove that $p_n \leq p_{n-1} + p_{n-2}$ for $n \geq 1$ by considering first the number of partitions of n that have at least two parts equal to 1, then the other partitions. Then use this to establish that $p_n \leq F_{n+1}$ for $n \geq 0$, where F_k denotes the k th Fibonacci number.
10. A *composition* of n is a list of positive integers $\langle a_1, a_2, \dots, a_k \rangle$ whose sum is n , where the order of the integers matters. For example, there are four different compositions of $n = 3$: $\langle 3 \rangle$, $\langle 2, 1 \rangle$, $\langle 1, 2 \rangle$, and $\langle 1, 1, 1 \rangle$. Let c_n denote the number of compositions of n , and let $c_{n,k}$ denote the number of compositions of n into exactly k parts.
- (a) Compute the value of $c_{n,k}$ for each k and n with $1 \leq k \leq n$ and $1 \leq n \leq 5$ by listing all the compositions, and then calculate the value of c_n for $1 \leq n \leq 5$.
- (b) Using these examples, conjecture formulas for $c_{n,k}$ and c_n , for arbitrary positive integers n and k . Then prove that your formulas are correct.

2.8.2 Stirling Cycle Numbers

The Round Table soon heard of the challenge, and of course it was a good deal discussed...

— Mark Twain, *A Connecticut Yankee in King Arthur's Court*

Suppose King Arthur decides to divide his knights into committees in order to better govern Britain. True to his egalitarian nature, he crafts k identical round tables for this purpose. How many ways are there to seat n knights at these tables, if each table can seat any number of knights, and no table can be empty? Here, we count two seating arrangements as different only if some knight has a different neighbor on his left side (or his right) in each one. Since the tables are identical, the particular table occupied by a group of knights is immaterial. Thus, once a group of knights is assigned to a table, we must account for all the possible seating arrangements there. From Section 2.7.2, we know that there are $(m-1)!$ different ways to seat m people at one round table.

Let us represent the n knights by the integers 1 through n , and denote the seating of knights K_1, K_2, \dots, K_m in clockwise order around one table by $(K_1 K_2 \dots K_m)$. Of course, $(K_2 K_3 \dots K_m K_1)$ denotes the same arrangement of knights around the table, so to make our notation unique we demand that the knight represented by the smallest number appear first in the list. An arrangement of knights at the k tables is then uniquely represented by a list of k strings of integers in parentheses, where each integer between 1 and n appears exactly once. For example, with six knights and three tables, we might seat knights 1, 3, and 5 in clockwise order around one table, knights 2 and 6 at another table, and knight 4 alone at the third table. This arrangement is denoted by $(135)(26)(4)$. This is precisely the cycle notation we used to describe a permutation on six objects. We see that each seating arrangement of n knights at k tables corresponds to a unique permutation $\pi \in S_n$ having exactly k cycles, and every such permutation corresponds to a unique seating arrangement.

We define the *Stirling cycle number*, denoted by $\left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right]$, to be the number of ways to seat n knights at k identical tables, or, equivalently, the number of permutations $\pi \in S_n$ having exactly k cycles. These numbers are also known as the *signless Stirling numbers of the first kind*. A signed version of these numbers is also often defined by

$$s(n, k) = (-1)^{n-k} \left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right], \quad (2.81)$$

but we will employ only the signless numbers here.

We derive a few properties of the Stirling cycle numbers. First, it is impossible to seat n knights at zero tables, unless there are no knights, so

$$\left[\begin{smallmatrix} n \\ 0 \end{smallmatrix} \right] = \begin{cases} 1 & \text{if } n = 0, \\ 0 & \text{if } n > 0. \end{cases} \quad (2.82)$$

Second, if there is only one table, then

$$\begin{bmatrix} n \\ 1 \end{bmatrix} = (n-1)!, \quad n \geq 1. \quad (2.83)$$

Next, if there are n tables, then each knight must sit at his own table, and if there are $n-1$ tables, then one pair of knights must sit at one table, and the others must each sit alone. Thus

$$\begin{bmatrix} n \\ n \end{bmatrix} = 1, \quad (2.84)$$

and

$$\begin{bmatrix} n \\ n-1 \end{bmatrix} = \binom{n}{2}. \quad (2.85)$$

There are no arrangements possible if there are more tables than knights, or a negative number of tables, so

$$\begin{bmatrix} n \\ k \end{bmatrix} = 0 \quad \text{if } k < 0 \text{ or } k > n. \quad (2.86)$$

Further, because of the correspondence between seating arrangements and permutations, we have

$$\sum_k \begin{bmatrix} n \\ k \end{bmatrix} = n!. \quad (2.87)$$

Consider now the case $n = 4$ and $k = 2$. Suppose one knight, delayed by an armor adjustment, picks his place after the first three knights are already seated. If the first three knights are seated at one table, then the last knight must sit at the second table by himself. The number of arrangements in this case is the number of ways to seat the first three knights at one table, so $\begin{bmatrix} 3 \\ 1 \end{bmatrix} = 2$. On the other hand, if two of the first three knights occupy one table, and the third sits at the second table, then the last knight may then either join the single knight, or the table with two knights. There are two possibilities in the latter case, since the fourth knight may sit on the left side of either of the knights already at the table. Thus, there are $3\begin{bmatrix} 3 \\ 2 \end{bmatrix} = 9$ possibilities in this case, and we find that $\begin{bmatrix} 4 \\ 2 \end{bmatrix} = 3\begin{bmatrix} 3 \\ 2 \end{bmatrix} + \begin{bmatrix} 3 \\ 1 \end{bmatrix} = 11$. Figure 2.19 shows these eleven arrangements when Tristram joins Bedivere, Lancelot, and Percival at two tables.

This technique generalizes to produce a recurrence relation for these numbers. To seat n knights at k tables, we can first seat $n-1$ knights at $k-1$ tables, then seat the last knight alone at the k th table. Alternatively, we can seat the first $n-1$ knights at k tables, then insert the last knight at one of these tables. This knight must sit on the left side of one of the other $n-1$ knights, so there are $n-1$ different places to seat the last knight. Therefore,

$$\begin{bmatrix} n \\ k \end{bmatrix} = (n-1) \begin{bmatrix} n-1 \\ k \end{bmatrix} + \begin{bmatrix} n-1 \\ k-1 \end{bmatrix}, \quad n \geq 1. \quad (2.88)$$




FIGURE 2.19. Seating Bedivere, Lancelot, Percival, and Tristram at two tables.

We can use this formula to compute a triangle of Stirling cycle numbers, just as we used the addition identity for binomial coefficients to obtain Pascal’s triangle. These computations appear in Table 2.5.

Recall that for fixed n the generating function for the sequence of binomial coefficients has a particularly nice form: $\sum_k \binom{n}{k} x^k = (x + 1)^n$. We can use the identity (2.88) to obtain an analogous representation for the sequence of Stirling cycle numbers. Let $G_n(x) = \sum_k \left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right] x^k$. Clearly, $G_0(x) = 1$, and for $n \geq 1$,

$$\begin{aligned} G_n(x) &= \sum_k \left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right] x^k \\ &= (n - 1) \sum_k \left[\begin{smallmatrix} n - 1 \\ k \end{smallmatrix} \right] x^k + \sum_k \left[\begin{smallmatrix} n - 1 \\ k - 1 \end{smallmatrix} \right] x^k \\ &= (n - 1)G_{n-1}(x) + xG_{n-1}(x), \end{aligned}$$

so $G_n(x) = (x + n - 1)G_{n-1}(x)$. It is easy to verify by induction that this implies that $G_n(x) = x(x + 1)(x + 2) \cdots (x + n - 1) = x^{\overline{n}}$. Thus,

$$x^{\overline{n}} = \sum_k \left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right] x^k. \tag{2.89}$$

$\begin{bmatrix} n \\ k \end{bmatrix}$	$k = 0$	1	2	3	4	5	6	7	8	$n!$
$n = 0$	1									1
1	0	1								1
2	0	1	1							2
3	0	2	3	1						6
4	0	6	11	6	1					24
5	0	24	50	35	10	1				120
6	0	120	274	225	85	15	1			720
7	0	720	1764	1624	735	175	21	1		5040
8	0	5040	13068	13132	6769	1960	322	28	1	40320

TABLE 2.5. Stirling cycle numbers, $\begin{bmatrix} n \\ k \end{bmatrix}$.

for $n \geq 0$. Therefore, the Stirling cycle numbers allow us to express rising factorial powers as linear combinations of ordinary powers. Exercise 7 establishes a similar connection for the falling factorial powers.

Exercises

- Use (2.88) and Table 2.5 to compute the values of $\begin{bmatrix} 9 \\ k \end{bmatrix}$ and $\begin{bmatrix} 10 \\ k \end{bmatrix}$ for each k .
- Prove that

$$\sum_k (-1)^k \begin{bmatrix} n \\ k \end{bmatrix} = \begin{cases} 1 & \text{if } n = 0, \\ -1 & \text{if } n = 1, \\ 0 & \text{if } n \geq 2. \end{cases}$$

- Use a combinatorial argument to show that

$$\begin{bmatrix} n \\ 2 \end{bmatrix} = \frac{n!}{2} \sum_{m=1}^{n-1} \frac{1}{m(n-m)}.$$

- Use a combinatorial argument to determine a simple formula for $\begin{bmatrix} n \\ n-2 \end{bmatrix}$.
- Use a combinatorial argument to show that

$$\begin{bmatrix} n+1 \\ m \end{bmatrix} = \sum_{k=0}^n \begin{bmatrix} n-k \\ m-1 \end{bmatrix} n^k$$

for nonnegative integers n and m .

- Prove that if n and m are nonnegative integers then

$$\begin{bmatrix} n+1 \\ m+1 \end{bmatrix} = \sum_k \begin{bmatrix} n \\ k \end{bmatrix} \binom{k}{m}.$$

7. Prove that if $n \geq 0$ then

$$x^n = \sum_k (-1)^{n-k} \begin{bmatrix} n \\ k \end{bmatrix} x^k = \sum_k s(n, k) x^k. \quad (2.90)$$

8. Use (2.89) to prove that if $n \geq 0$ then

$$\sum_k \begin{bmatrix} n \\ k \end{bmatrix} y^{n-k} = \prod_{k=1}^n (1 + ky).$$

Then use this to prove that $\begin{bmatrix} n \\ k \end{bmatrix}$ equals the sum of all products of $n - k$ distinct integers selected from $\{1, \dots, n - 1\}$. For example, $\begin{bmatrix} 6 \\ 3 \end{bmatrix} = 1 \cdot 2 \cdot 3 + 1 \cdot 2 \cdot 4 + 1 \cdot 2 \cdot 5 + 1 \cdot 3 \cdot 4 + 1 \cdot 3 \cdot 5 + 1 \cdot 4 \cdot 5 + 2 \cdot 3 \cdot 4 + 2 \cdot 3 \cdot 5 + 2 \cdot 4 \cdot 5 + 3 \cdot 4 \cdot 5 = 225$.

9. Let $d(n, k)$ denote the number of ways to place n knights at k identical tables, with at least *two* knights at each table. For example, Figure 2.19 shows that $d(4, 2) = 3$. Set $d(0, 0) = 1$.

(a) Use a combinatorial argument to show that $d(n, k)$ satisfies the recurrence relation

$$d(n, k) = (n - 1)(d(n - 1, k) + d(n - 2, k - 1))$$

for $n \geq 1$.

(b) Compute the table of values of $d(n, k)$ for $0 \leq n \leq 8$, similar to Table 2.5.

(c) Prove that if $n \geq 0$ then

$$\sum_k d(n, k) = n_i,$$

where n_i denotes the number of derangements of n .

2.8.3 Stirling Set Numbers

36 (*Roger Federer, 2006–07*), 35 (*John McEnroe, 1984*),

26 (*Stefan Edberg, 1991–92*), 25 (*Ilie Nastase, 1972–73*).

— Most consecutive sets won in Grand Slam matches in men's tennis

How many ways are there to divide n guests at a party into exactly k groups, if we disregard the arrangement of people within each group? Rephrased, this problem asks for the number of ways to partition a set of n objects into exactly k nonempty subsets, so that each element in the original set appears exactly once among the k subsets. For example, there are three ways to partition the set $\{a, b, c\}$ into two nonempty subsets: $\{a, b\}, \{c\}$; $\{a, c\}, \{b\}$; and $\{b, c\}, \{a\}$. There is just one

way to partition $\{a, b, c\}$ into one subset: $\{a, b, c\}$, and just one way to partition $\{a, b, c\}$ into three subsets: $\{a\}$, $\{b\}$, $\{c\}$.

The number of ways to divide n objects into exactly k groups is denoted by $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$. Thus, $\left\{ \begin{smallmatrix} 3 \\ 2 \end{smallmatrix} \right\} = 3$, and $\left\{ \begin{smallmatrix} 3 \\ 1 \end{smallmatrix} \right\} = \left\{ \begin{smallmatrix} 3 \\ 3 \end{smallmatrix} \right\} = 1$. These numbers are called the *Stirling set numbers*, or the *Stirling numbers of the second kind*. The notation $S(n, k)$ is also often used to denote these numbers.

We begin by listing a few properties of these numbers. First, for $n \geq 1$ we have

$$\left\{ \begin{smallmatrix} n \\ 1 \end{smallmatrix} \right\} = \left\{ \begin{smallmatrix} n \\ n \end{smallmatrix} \right\} = 1, \quad (2.91)$$

since there is only one way to place n people into a single group, and only one way to split them into n groups. Second,

$$\left\{ \begin{smallmatrix} n \\ 0 \end{smallmatrix} \right\} = \begin{cases} 1 & \text{if } n = 0, \\ 0 & \text{if } n > 0, \end{cases} \quad (2.92)$$

since one cannot divide n people into zero groups, unless there are no people. Third, to divide n people into $n - 1$ groups, we must pick two people to be in one group, then place the rest of the people in groups by themselves, so

$$\left\{ \begin{smallmatrix} n \\ n-1 \end{smallmatrix} \right\} = \binom{n}{2}. \quad (2.93)$$

Next, we set

$$\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} = 0, \quad \text{if } k < 0 \text{ or } k > n. \quad (2.94)$$

Also, the Stirling cycle number $\left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right]$ distinguishes among the different ways to arrange n people within k groups, and the Stirling set number $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$ does not, so

$$\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} \leq \left[\begin{smallmatrix} n \\ k \end{smallmatrix} \right] \quad (2.95)$$

for all $n \geq 0$ and all k .

We now derive a recurrence relation for $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$. Suppose we plan to divide $n \geq 1$ people into k groups for a party, and we know that one person will arrive late. We could divide the first $n - 1$ people into $k - 1$ groups, then place the last person in her own group when she arrives, or we can arrange the first $n - 1$ people into k groups, then pick a group for the last person to join. There are $\left\{ \begin{smallmatrix} n-1 \\ k-1 \end{smallmatrix} \right\}$ different ways to arrange the guests in the first case, and $k \left\{ \begin{smallmatrix} n-1 \\ k \end{smallmatrix} \right\}$ different possibilities in the second. Therefore,

$$\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} = k \left\{ \begin{smallmatrix} n-1 \\ k \end{smallmatrix} \right\} + \left\{ \begin{smallmatrix} n-1 \\ k-1 \end{smallmatrix} \right\}, \quad n \geq 1. \quad (2.96)$$

For example, to partition the set $\{a, b, c, d\}$ into two subsets, we can place d in its own set, yielding $\{a, b, c\}$, $\{d\}$, or we can split $\{a, b, c\}$ into two sets, then add d

to one of these sets. The latter possibility yields the six different partitions

$$\begin{aligned} &\{a, b, d\}, \{c\}; \{a, c, d\}, \{b\}; \{b, c, d\}, \{a\}; \\ &\{a, b\}, \{c, d\}; \{a, c\}, \{b, d\}; \{b, c\}, \{a, d\}; \end{aligned} \tag{2.97}$$

and $\left\{ \begin{smallmatrix} 4 \\ 2 \end{smallmatrix} \right\} = 2\left\{ \begin{smallmatrix} 3 \\ 2 \end{smallmatrix} \right\} + \left\{ \begin{smallmatrix} 3 \\ 1 \end{smallmatrix} \right\} = 7$.

Using identity (2.96), we can generate the triangle of Stirling set numbers shown in Table 2.6. The sequence $\{b_n\}$ that appears in this table as the sum across the rows of the triangle is studied in the next section.

$\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$	$k = 0$	1	2	3	4	5	6	7	8	b_n
$n = 0$	1									1
1	0	1								1
2	0	1	1							2
3	0	1	3	1						5
4	0	1	7	6	1					15
5	0	1	15	25	10	1				52
6	0	1	31	90	65	15	1			203
7	0	1	63	301	350	140	21	1		877
8	0	1	127	966	1701	1050	266	28	1	4140

TABLE 2.6. Stirling set numbers, $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$, and Bell numbers, b_n .

Exercise 8 analyzes the generating function for the sequence of Stirling set numbers $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$ with n fixed. We can obtain a more useful relation, however, if we replace the ordinary powers of x in this generating function with falling factorial powers. For fixed n , let

$$F_n(x) = \sum_k \left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} x^{\underline{k}},$$

so $F_0(x) = 1$. If $n \geq 1$, then

$$\begin{aligned} F_n(x) &= \sum_k \left(k \left\{ \begin{smallmatrix} n-1 \\ k \end{smallmatrix} \right\} + \left\{ \begin{smallmatrix} n-1 \\ k-1 \end{smallmatrix} \right\} \right) x^{\underline{k}} \\ &= \sum_k k \left\{ \begin{smallmatrix} n-1 \\ k \end{smallmatrix} \right\} x^{\underline{k}} + \sum_k \left\{ \begin{smallmatrix} n-1 \\ k \end{smallmatrix} \right\} x^{\underline{k+1}} \\ &= \sum_k k \left\{ \begin{smallmatrix} n-1 \\ k \end{smallmatrix} \right\} x^{\underline{k}} + \sum_k (x-k) \left\{ \begin{smallmatrix} n-1 \\ k \end{smallmatrix} \right\} x^{\underline{k}} \\ &= xF_{n-1}(x), \end{aligned}$$

so by induction we obtain

$$x^n = \sum_k \left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} x^{\underline{k}}, \quad n \geq 0. \tag{2.98}$$

Therefore, the Stirling set numbers allow us to express ordinary powers as combinations of falling factorial powers.

We can derive another useful formula by considering the generating function for the numbers $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$ with k fixed. Let

$$H_k(x) = \sum_{n \geq 0} \left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} x^n,$$

so $H_0(x) = 1$. For $k \geq 1$, we obtain

$$\begin{aligned} H_k(x) &= \sum_{n \geq 1} \left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} x^n \\ &= \sum_{n \geq 1} \left(k \left\{ \begin{smallmatrix} n-1 \\ k \end{smallmatrix} \right\} + \left\{ \begin{smallmatrix} n-1 \\ k-1 \end{smallmatrix} \right\} \right) x^n \\ &= kx \sum_{n \geq 0} \left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} x^n + x \sum_{n \geq 0} \left\{ \begin{smallmatrix} n \\ k-1 \end{smallmatrix} \right\} x^n \\ &= kxH_k(x) + xH_{k-1}(x), \end{aligned}$$

so

$$H_k(x) = \frac{x}{1-kx} H_{k-1}(x),$$

and therefore

$$H_k(x) = \frac{x^k}{(1-x)(1-2x) \cdots (1-kx)}. \quad (2.99)$$

Next, we use partial fractions to expand this rational function. Our calculations are somewhat simpler if we multiply by $k!$ first, so we wish to find constants A_1, A_2, \dots, A_k such that

$$\frac{k!x^k}{\prod_{m=1}^k (1-mx)} = \sum_{m=1}^k \frac{A_m}{1-mx}.$$

Clearing denominators, we have

$$k!x^k = \sum_{m=1}^k A_m \prod_{j=1}^{m-1} (1-jx) \prod_{j=m+1}^k (1-jx),$$

and setting $x = 1/m$, we obtain

$$\frac{k!}{m^k} = A_m \prod_{j=1}^{m-1} \left(1 - \frac{j}{m} \right) \prod_{j=m+1}^k \left(1 - \frac{j}{m} \right),$$

so

$$\begin{aligned} k! &= mA_m \prod_{j=1}^{m-1} (m-j) \prod_{j=m+1}^k (m-j) \\ &= mA_m (m-1)! (-1)^{k-m} \prod_{j=m+1}^k (j-m) \\ &= (-1)^{k-m} m! (k-m)! A_m, \end{aligned}$$

and

$$A_m = (-1)^{k-m} \binom{k}{m}.$$

Thus

$$\begin{aligned} H_k(x) &= \frac{1}{k!} \sum_{m=1}^k (-1)^{k-m} \frac{\binom{k}{m}}{1-mx} \\ &= \frac{1}{k!} \sum_{m=1}^k (-1)^{k-m} \binom{k}{m} \sum_{n \geq 0} (mx)^n \\ &= \sum_{n \geq 0} \left(\frac{1}{k!} \sum_{m=1}^k (-1)^{k-m} \binom{k}{m} m^n \right) x^n, \end{aligned}$$

and therefore

$$\left\{ \begin{matrix} n \\ k \end{matrix} \right\} = \frac{1}{k!} \sum_{m=0}^k (-1)^{k-m} \binom{k}{m} m^n, \quad (2.100)$$

for any nonnegative integers n and k . This produces a formula for the Stirling set numbers. For example, we may compute $\left\{ \begin{matrix} 6 \\ 3 \end{matrix} \right\} = \frac{1}{3!} (3 \cdot 1^6 - 3 \cdot 2^6 + 1 \cdot 3^6) = 90$.

Exercises

1. Use (2.96) and Table 2.6 to compute the values of $\left\{ \begin{matrix} 9 \\ k \end{matrix} \right\}$ and $\left\{ \begin{matrix} 10 \\ k \end{matrix} \right\}$ for each k .
2. A hungry fraternity brother stops at the drive-through window of a fast-food restaurant and orders twelve different items. The server plans to convey the items using either three or four identical cardboard trays, and empty trays are never given to a customer. Use (2.96) and your augmented table from Exercise 1 to determine the number of ways that the server can arrange the items on the trays.
3. Use combinatorial arguments to determine simple formulas for $\left\{ \begin{matrix} n \\ 2 \end{matrix} \right\}$ and $\left\{ \begin{matrix} n \\ n-2 \end{matrix} \right\}$.
4. A new casino game takes ten ping-pong balls, each labeled with a different number between 1 and 10, and drops each one at random into one of three identical buckets. A bucket may be empty after the ten balls are dropped.

- (a) Suppose a bet consists of identifying which balls have landed together in each bucket. For example, a bet may state that one bucket is empty, another has just the balls numbered 2, 3, and 7, and the rest are in the other bucket. How many bets are possible?
- (b) Suppose instead that a bet consists of identifying only the number of balls that land in the buckets. For example, a bet might state that one bucket is empty, another has three balls, and the other has seven. The numbers on the balls have no role in the bet. How many bets are possible?
5. How many different fifty-character sequences use every character of the 26-letter alphabet at least once? More generally, how many ways can one place n distinguishable objects into k distinguishable bins, if no bin may be empty?
6. Use (2.99) to prove that $\left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\}$ equals the sum of all products of $n - k$ integers selected from $\{1, \dots, k\}$. For example, $\left\{ \begin{smallmatrix} 6 \\ 3 \end{smallmatrix} \right\} = 1 \cdot 1 \cdot 1 + 1 \cdot 1 \cdot 2 + 1 \cdot 1 \cdot 3 + 1 \cdot 2 \cdot 2 + 1 \cdot 2 \cdot 3 + 1 \cdot 3 \cdot 3 + 2 \cdot 2 \cdot 2 + 2 \cdot 2 \cdot 3 + 2 \cdot 3 \cdot 3 + 3 \cdot 3 \cdot 3 = 90$.
7. Let $r_{n,k}$ denote the number of ways to divide n people into k groups, with at least *two* people in each group. For example, the list (2.97) shows that $r_{4,2} = 3$. Set $r_{0,0} = 1$.

- (a) Use a combinatorial argument to show that $r_{n,k}$ satisfies the recurrence relation

$$r_{n,k} = kr_{n-1,k} + (n-1)r_{n-2,k-1}$$

for $n \geq 1$.

- (b) Define r_n for $n \geq 0$ by $r_n = \sum_k r_{n,k}$. Compute the table of values of $r_{n,k}$ and r_n for $0 \leq n \leq 8$, similar to Table 2.6.
- (c) Determine a formula for $r_{2n,n}$, for a positive integer n .
- (d) A *rhyming scheme* describes the pattern of rhymes in a poem. For example, the rhyming scheme of a limerick is (a, a, b, b, a) , since a limerick has five lines, with the first, second, and last line exhibiting one rhyme, and the third and fourth showing a different rhyme. Also, a sonnet is a poem with fourteen lines. Shakespearean sonnets have the rhyming scheme $(a, b, b, a, b, c, c, d, e, f, e, f, g, g)$; many Petrarchan sonnets exhibit the scheme $(a, b, b, a, a, b, b, a, c, d, e, c, d, e)$. Argue that r_n counts the number of possible rhyming schemes for a poem with n lines, if each line must rhyme with at least one other line.
8. Let $G_n(x) = \sum_k \left\{ \begin{smallmatrix} n \\ k \end{smallmatrix} \right\} x^k$, so $G_0(x) = 1$. Show that $G_n(x) = x(G_{n-1}(x) + G'_{n-1}(x))$ for $n \geq 1$, and use this recurrence to compute $G_4(x)$.

9. Show that

$$x^n = \sum_k \left\{ \begin{matrix} n \\ k \end{matrix} \right\} (-1)^{n-k} x^{\bar{k}}. \quad (2.101)$$

10. Use (2.90) and (2.98), or (2.89) and (2.101), to prove the following identities.

$$\sum_k \left[\begin{matrix} n \\ k \end{matrix} \right] \left\{ \begin{matrix} k \\ m \end{matrix} \right\} (-1)^{(n-k)} = \begin{cases} 1 & \text{if } n = m, \\ 0 & \text{otherwise.} \end{cases} \quad (2.102)$$

$$\sum_k \left\{ \begin{matrix} n \\ k \end{matrix} \right\} \left[\begin{matrix} k \\ m \end{matrix} \right] (-1)^{(n-k)} = \begin{cases} 1 & \text{if } n = m, \\ 0 & \text{otherwise.} \end{cases} \quad (2.103)$$

11. Prove that

$$\sum_{k \geq 0} k^n x^k = \sum_k \left\{ \begin{matrix} n \\ k \end{matrix} \right\} \frac{k! x^k}{(1-x)^{k+1}}$$

for any nonnegative integer n .

12. Suppose $\{r_1, \dots, r_\ell\}$ and $\{s_1, \dots, s_\ell\}$ are two sets of positive integers, $f(x) = \sum_{j=1}^{\ell} (x^{r_j} - x^{s_j})$, and N is a positive integer. Prove that

$$\sum_{j=1}^{\ell} r_j^n = \sum_{j=1}^{\ell} s_j^n$$

for every n with $1 \leq n \leq N$ if and only if $f^{(n)}(1) = 0$ for every n with $1 \leq n \leq N$. Here, $f^{(n)}(x)$ denotes the n th derivative of $f(x)$.

For example, select $\{1, 5, 9, 17, 18\}$ and $\{2, 3, 11, 15, 19\}$ as the two sets, and select $N = 4$. Then $1 + 5 + 9 + 17 + 18 = 2 + 3 + 11 + 15 + 19 = 50$, $1^2 + 5^2 + 9^2 + 17^2 + 18^2 = 2^2 + 3^2 + 11^2 + 15^2 + 19^2 = 720$, $1^3 + 5^3 + 9^3 + 17^3 + 18^3 = 2^3 + 3^3 + 11^3 + 15^3 + 19^3 = 11600$, and $1^4 + 5^4 + 9^4 + 17^4 + 18^4 = 2^4 + 3^4 + 11^4 + 15^4 + 19^4 = 195684$; and $f(x) = x - x^2 + x^5 - x^3 + x^9 - x^{11} + x^{17} - x^{15} + x^{18} - x^{19}$ has $f^{(n)}(1) = 0$ for $1 \leq n \leq 4$.

2.8.4 Bell Numbers

Silence that dreadful bell: it frights the isle...

— William Shakespeare, *Othello*, Act II, Scene III

The *Bell number* b_n is the number of ways to divide n people into any number of groups. It is therefore a sum of Stirling set numbers,

$$b_n = \sum_k \left\{ \begin{matrix} n \\ k \end{matrix} \right\}. \quad (2.104)$$

The first few values of this sequence appear in Table 2.6.

We can derive a recurrence relation for the Bell numbers. To divide n people into groups, consider the different ways to form a group containing one particular person. We must choose some number k of the other $n - 1$ people to join this person in one group, then divide the other $n - 1 - k$ people into groups. It follows that

$$b_n = \sum_k \binom{n-1}{k} b_{n-1-k}.$$

Reindexing the sum by replacing k with $n - 1 - k$, then applying the symmetry identity for binomial coefficients, we find the somewhat simpler relation

$$b_n = \sum_k \binom{n-1}{k} b_k, \quad n \geq 1. \quad (2.105)$$

Rather than analyze the ordinary generating function for the sequence of Bell numbers, we introduce another kind of generating function that is often useful in combinatorial analysis. The *exponential generating function* for the sequence $\{a_n\}$ is defined as the ordinary generating function for the sequence $\{a_n/n!\}$. For example, the exponential generating function for the constant sequence $a_n = c$ is $\sum_{n \geq 0} cx^n/n! = ce^x$, and for the sequence $a_n = (-1)^n n!$, it is $1/(1+x)$. The exponential generating function for the sequence of Bell numbers is therefore

$$E(x) = \sum_{n \geq 0} \frac{b_n}{n!} x^n. \quad (2.106)$$

We can compute a closed form for this series. Differentiating, we find

$$\begin{aligned} E'(x) &= \sum_{n \geq 1} \frac{b_n}{(n-1)!} x^{n-1} \\ &= \sum_{n \geq 1} \frac{1}{(n-1)!} \left(\sum_k \binom{n-1}{k} b_k \right) x^{n-1} \\ &= \sum_{n \geq 1} \sum_{k=0}^{n-1} \frac{b_k}{k!(n-1-k)!} x^{n-1} \\ &= \sum_{k \geq 0} \sum_{n \geq k+1} \frac{b_k}{k!(n-1-k)!} x^{n-1} \\ &= \sum_{k \geq 0} \sum_{n \geq 0} \frac{b_k}{k!n!} x^{n+k} \\ &= \left(\sum_{n \geq 0} \frac{x^n}{n!} \right) \left(\sum_{k \geq 0} \frac{b_k}{k!} x^k \right) \\ &= e^x E(x). \end{aligned}$$

Therefore,

$$(\ln E(x))' = e^x,$$

and so

$$\ln E(x) = e^x + c$$

for some constant c . Since $E(0) = b_0 = 1$, we must have $c = -1$. Thus,

$$E(x) = e^{e^x - 1}. \quad (2.107)$$

We can use this closed form to determine a formula for b_n . Using the Maclaurin series for the exponential function twice, we find that

$$\begin{aligned} E(x) &= \frac{1}{e} e^{e^x} \\ &= \frac{1}{e} \sum_{k \geq 0} \frac{(e^x)^k}{k!} \\ &= \frac{1}{e} \sum_{k \geq 0} \frac{1}{k!} \sum_{n \geq 0} \frac{(kx)^n}{n!} \\ &= \frac{1}{e} \sum_{n \geq 0} \left(\sum_{k \geq 0} \frac{k^n}{k!} \right) \frac{x^n}{n!}. \end{aligned}$$

Therefore,

$$b_n = \frac{1}{e} \sum_{k \geq 0} \frac{k^n}{k!}. \quad (2.108)$$

This formula is sometimes called *Dobiński's formula* [79].

Exercises

- How many ways are there to put ten different dogs into pens, if each pen can hold any number of dogs, and every pen is exactly the same?
- Determine a closed form for the exponential generating function for each of the following sequences.
 - $a_k = c^k$, with c a constant.
 - $a_k = 1$ if k is even and 0 if k is odd.
 - $a_k = k$.
 - $a_k = k^n$, for a fixed nonnegative integer n . The number of terms in the answer may depend on n .
- Verify that equation (2.108) for b_n produces the correct value for b_0 , b_1 , and b_2 .
- Show that the series in equation (2.108) converges for every $n \geq 0$.

5. Use a combinatorial argument to show that

$$\begin{aligned} \left\{ \begin{matrix} n \\ m \end{matrix} \right\} &= \sum_k \binom{n-1}{k} \left\{ \begin{matrix} n-k-1 \\ m-1 \end{matrix} \right\} \\ &= \sum_k \binom{n-1}{k} \left\{ \begin{matrix} k \\ m-1 \end{matrix} \right\}, \end{aligned} \quad (2.109)$$

for $n \geq 1$, and use this to derive the recurrence (2.105) for Bell numbers.

6. Define the *complementary Bell number* \tilde{b}_n for $n \geq 0$ by

$$\tilde{b}_n = \sum_k (-1)^k \left\{ \begin{matrix} n \\ k \end{matrix} \right\}.$$

Wilf asked if $\tilde{b}_n = 0$ for infinitely many n , or if there even exists an integer $n > 2$ where $\tilde{b}_n = 0$. The first few complementary Bell numbers are 1, -1, 0, 1, 1, -2, -9, -9, 50, 267, 413, -2180, -17731, -50533, and 110176.

- Describe a combinatorial interpretation of \tilde{b}_n .
- Use (2.109) to determine a recurrence for the complementary Bell numbers. Then determine a closed form for their exponential generating function, $\tilde{E}(x)$. How is this function related to the function $E(x)$ of this section?
- Use $\tilde{E}(x)$ to determine a formula for \tilde{b}_n , similar to the expression (2.108) for b_n .

It is known that the sequence \tilde{b}_n changes sign infinitely often, and that $\tilde{b}_n \neq 0$ for almost all values of n . See Yang [289] and de Wannemacker, Laffey, and Osburn [71] for more information on this problem.

- Suppose $P(x)$ is the exponential generating function for the sequence $\{p_n\}$, and $Q(x)$ is the exponential generating function for $\{q_n\}$. Prove that the product $P(x)Q(x)$ is the exponential generating function for the sequence $\{\sum_k \binom{n}{k} p_k q_{n-k}\}$.
- Let r_n denote the number of rhyming schemes for a poem with n lines, if each line must rhyme with at least one other line, as in Exercise 7d of Section 2.8.3. Recall that $r_0 = 1$.

(a) Prove that

$$r_n = \sum_{k=0}^{n-2} \binom{n-1}{k} r_k.$$

- Determine a closed form similar to (2.107) for the exponential generating function $R(x)$ for the sequence $\{r_n\}$.

(c) Use this generating function, together with Exercise 7, to show that

$$r_n = \sum_k \binom{n}{k} (-1)^{n-k} b_k.$$

(d) Prove that the number of rhyming schemes for $n + 1$ lines in which each line rhymes with at least one other line equals the number of rhyming schemes for n lines in which at least one line rhymes with no other line. Note that b_n is the total number of rhyming schemes on n lines, including schemes where some lines rhyme with no others.

9. Let $E_k(x)$ denote the exponential generating function for the sequence of Stirling cycle numbers with k fixed,

$$E_k(x) = \sum_{n \geq 0} \left[\begin{matrix} n \\ k \end{matrix} \right] \frac{x^n}{n!}.$$

Prove that

$$E'_k(x) = \frac{E_{k-1}(x)}{1-x},$$

for $k \geq 1$, and use this to derive a closed form for $E_k(x)$,

$$\sum_{n \geq 0} \left[\begin{matrix} n \\ k \end{matrix} \right] \frac{x^n}{n!} = \frac{(-1)^k}{k!} (\ln(1-x))^k. \quad (2.110)$$

Comtet [60] uses this identity, together with (2.100) and (2.113), to derive a complicated formula due to Schlömilch for the Stirling cycle numbers. We include it here without proof:

$$\left[\begin{matrix} n \\ k \end{matrix} \right] = \sum_{m=0}^{n-k} (-1)^{n-k-m} \binom{n-1+m}{k-1} \binom{2n-k}{n-k-m} \left\{ \begin{matrix} n-k+m \\ m \end{matrix} \right\} \quad (2.111)$$

$$= \sum_{m=0}^{n-k} \sum_{j=0}^m (-1)^{n-k-j} \binom{n-1+m}{k-1} \binom{2n-k}{n-k-m} \binom{m}{j} \frac{j^{n-k+m}}{m!}. \quad (2.112)$$

10. Use an argument similar to that of Exercise 9 to prove that

$$\sum_{n \geq 0} \left\{ \begin{matrix} n \\ k \end{matrix} \right\} \frac{x^n}{n!} = \frac{1}{k!} (e^x - 1)^k \quad (2.113)$$

for every $k \geq 0$.

2.8.5 Eulerian Numbers

3 (*Al Hamilton*), 7 (*Paul Coffey*), 11 (*Mark Messier*), 17 (*Jari Kurri*), 31 (*Grant Fuhr*), 99 (*Wayne Gretzky*).

— Retired jersey numbers, Edmonton Oilers

Suppose that a pipe organ having n pipes needs to be installed at a concert hall. Each pipe has a different length, and the pipes must be arranged in a single row. Let us say that two adjacent pipes in an arrangement form an *ascent* if the one on the left is shorter than the one on the right, and a *descent* otherwise. Arranging the pipes from shortest to tallest yields an arrangement with $n - 1$ ascents and no descents; arranging them from tallest to shortest results in no ascents and $n - 1$ descents.

Whether for aesthetic or acoustical reasons, the eccentric director of the concert hall demands that there be exactly k ascents in the arrangement of the n pipes. How many ways are there to install the organ? The answer is the *Eulerian number* $\langle \begin{smallmatrix} n \\ k \end{smallmatrix} \rangle$. Stated in more abstract terms, $\langle \begin{smallmatrix} n \\ k \end{smallmatrix} \rangle$ is the number of permutations π of the integers $\{1, \dots, n\}$ having $\pi(i) < \pi(i + 1)$ for exactly k numbers i between 1 and $n - 1$.

We list a few properties of these numbers. It is easy to see that there is only one arrangement of n pipes with no ascents, and only one with $n - 1$ ascents, so

$$\left\langle \begin{smallmatrix} n \\ 0 \end{smallmatrix} \right\rangle = 1, \quad n \geq 0, \quad (2.114)$$

and

$$\left\langle \begin{smallmatrix} n \\ n - 1 \end{smallmatrix} \right\rangle = 1, \quad n \geq 1. \quad (2.115)$$

The Eulerian numbers have a symmetry property similar to that of the binomial coefficients. An arrangement of n pipes with k ascents has $n - 1 - k$ descents, so reversing this arrangement yields a complementary configuration with $n - 1 - k$ ascents and k descents. Thus,

$$\left\langle \begin{smallmatrix} n \\ k \end{smallmatrix} \right\rangle = \left\langle \begin{smallmatrix} n \\ n - 1 - k \end{smallmatrix} \right\rangle. \quad (2.116)$$

Next, by summing over k we count every possible arrangement of pipes precisely once, so

$$\sum_k \left\langle \begin{smallmatrix} n \\ k \end{smallmatrix} \right\rangle = n!. \quad (2.117)$$

We also note the degenerate cases

$$\left\langle \begin{smallmatrix} n \\ k \end{smallmatrix} \right\rangle = 0, \quad \text{if } n > 0, \text{ and } k < 0 \text{ or } k \geq n, \quad (2.118)$$

and

$$\left\langle \begin{smallmatrix} 0 \\ k \end{smallmatrix} \right\rangle = 0, \quad \text{if } k \neq 0. \quad (2.119)$$

We can derive a recurrence relation for the Eulerian numbers. To arrange n pipes with exactly k ascents, suppose we first place every pipe except the tallest into a configuration with exactly k ascents. Then the tallest pipe can be inserted either in the first position, or between two pipes forming any ascent. Any other position would yield an additional ascent. There are therefore $k + 1$ different places to insert the tallest pipe in this case. Alternatively, we can line up the $n - 1$ shorter pipes so that there are $k - 1$ ascents, then insert the last pipe either at the end of the row, or between two pipes forming any descent. There are $n - 2 - (k - 1) = n - k - 1$ descents, so there are $n - k$ different places to insert the tallest pipe in this case. It is impossible to create a permissible configuration by inserting the tallest pipe into any other arrangement of the $n - 1$ shorter pipes, so

$$\langle n \rangle_k = (k + 1) \langle n - 1 \rangle_k + (n - k) \langle n - 1 \rangle_{k - 1}, \quad n \geq 1. \tag{2.120}$$

For example, $\langle 3 \rangle_1 = 2 \langle 2 \rangle_1 + 2 \langle 2 \rangle_0 = 4$, and $\langle 4 \rangle_2 = 3 \langle 3 \rangle_2 + 2 \langle 3 \rangle_1 = 3 + 8 = 11$. Figure 2.20 shows these eleven arrangements of four pipes with two ascents.

We can use the recurrence (2.120) to compute the triangle of Eulerian numbers, shown in Table 2.7.

$\langle n \rangle_k$	$k = 0$	1	2	3	4	5	6	7	$n!$
$n = 0$	1								1
1	1								1
2	1	1							2
3	1	4	1						6
4	1	11	11	1					24
5	1	26	66	26	1				120
6	1	57	302	302	57	1			720
7	1	120	1191	2416	1191	120	1		5040
8	1	247	4293	15619	15619	4293	247	1	40320

TABLE 2.7. Eulerian numbers, $\langle n \rangle_k$.

Next, we study some generating functions involving the Eulerian numbers. Recall that in Section 2.6.5 we computed the generating function for the sequence $\{0, 1, 2, 3, \dots\}$ by differentiating both sides of the identity $\sum_{k \geq 0} x^k = \frac{1}{1-x}$, then multiplying by x :

$$\sum_{k \geq 0} kx^k = x \cdot \frac{d}{dx} \left(\sum_{k \geq 0} x^k \right) = x \cdot \frac{d}{dx} \left(\frac{1}{1-x} \right) = \frac{x}{(1-x)^2}. \tag{2.121}$$

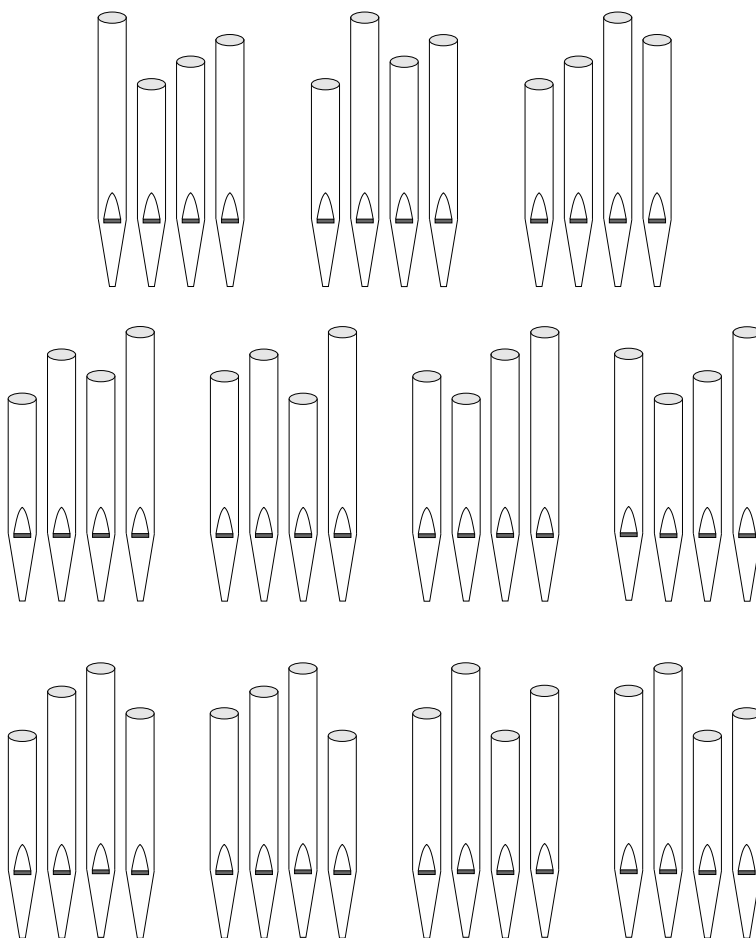


FIGURE 2.20. Four organ pipes with two ascents.

Clearly, we can obtain a generating function for the sequence of squares $\{k^2\}$ by applying the same differentiate-and-multiply operator to (2.121). We find that

$$\begin{aligned}
 \sum_{k \geq 0} k^2 x^k &= x \cdot \frac{d}{dx} \left(\frac{x}{(1-x)^2} \right) \\
 &= x \left(\frac{2x}{(1-x)^3} + \frac{1}{(1-x)^2} \right) \\
 &= \frac{x(1+x)}{(1-x)^3}.
 \end{aligned}
 \tag{2.122}$$

In the same way, we may use this operator to calculate the generating function for the sequence of cubes, then fourth powers and fifth powers. After a bit of

simplifying, we find that

$$\sum_{k \geq 0} k^3 x^k = \frac{x(1 + 4x + x^2)}{(1-x)^4}, \quad (2.123)$$

$$\sum_{k \geq 0} k^4 x^k = \frac{x(1 + 11x + 11x^2 + x^3)}{(1-x)^5}, \quad (2.124)$$

$$\sum_{k \geq 0} k^5 x^k = \frac{x(1 + 26x + 66x^2 + 26x^3 + x^4)}{(1-x)^6}. \quad (2.125)$$

A glance at Table 2.7 shows that the coefficients appearing on the right side of these formulas are all Eulerian numbers, and we would suspect that the numbers $\langle n \rangle_k$ will appear in the generating function for the sequence of n th powers of integers. This is in fact the case.

Theorem 2.16. *If $n \geq 0$ then*

$$\sum_{k \geq 1} k^n x^k = \frac{x}{(1-x)^{n+1}} \sum_k \langle n \rangle_k x^k. \quad (2.126)$$

Proof. We use induction on n . The formula is easy to verify when $n = 0$, so we assume it holds for a nonnegative integer n . We calculate

$$\begin{aligned} \sum_{k \geq 1} k^{n+1} x^k &= x \cdot \frac{d}{dx} \left(\sum_{k \geq 1} k^n x^k \right) \\ &= x \cdot \frac{d}{dx} \left(\frac{x}{(1-x)^{n+1}} \sum_k \langle n \rangle_k x^k \right) \\ &= x \left(\frac{1}{(1-x)^{n+1}} \sum_k \langle n \rangle_k (k+1) x^k + \frac{n+1}{(1-x)^{n+2}} \sum_k \langle n \rangle_k x^{k+1} \right) \\ &= \frac{x}{(1-x)^{n+2}} \left((1-x) \sum_k \langle n \rangle_k (k+1) x^k + (n+1) \sum_k \langle n \rangle_{k-1} x^k \right) \\ &= \frac{x}{(1-x)^{n+2}} \left(\sum_k (k+1) \langle n \rangle_k x^k + \sum_k (n+1-k) \langle n \rangle_{k-1} x^k \right) \\ &= \frac{x}{(1-x)^{n+2}} \sum_k \langle n+1 \rangle_k x^k. \end{aligned}$$

The last step follows from the recurrence relation (2.120). \square

We can use (2.126) to obtain a formula for $\langle n \rangle_k$ in terms of binomial coefficients and powers. We calculate

$$\begin{aligned} \sum_k \langle n \rangle_k x^k &= \frac{(1-x)^{m+1}}{x} \sum_{m \geq 1} m^n x^m \\ &= \sum_{m \geq 0} (m+1)^n x^m \sum_j \binom{n+1}{j} (-1)^j x^j \\ &= \sum_{m \geq 0} \sum_{j \geq 0} (-1)^j \binom{n+1}{j} (m+1)^n x^{j+m} \\ &= \sum_{k \geq 0} \sum_{j=0}^k (-1)^j \binom{n+1}{j} (k+1-j)^n x^k. \end{aligned} \tag{2.127}$$

Now the first and last expressions in (2.127) are power series in x , so we can equate coefficients to obtain a formula for the Eulerian number $\langle n \rangle_k$. We find that

$$\langle n \rangle_k = \sum_{j=0}^k (-1)^j \binom{n+1}{j} (k+1-j)^n. \tag{2.128}$$

Last, we derive one more interesting identity involving Eulerian numbers, binomial coefficients, and ordinary powers. Consider a sort of generating function for the sequence $\{\langle n \rangle_k\}$ with n fixed, where we use the binomial coefficient $\binom{x+k}{n}$ in place of x^k . Let

$$F_n(x) = \sum_k \langle n \rangle_k \binom{x+k}{n},$$

so that $F_0(x) = 1$. For $n \geq 1$, we calculate

$$\begin{aligned} F_n(x) &= \sum_k \left((k+1) \langle n-1 \rangle_k + (n-k) \langle n-1 \rangle_{k-1} \right) \binom{x+k}{n} \\ &= \sum_k (k+1) \langle n-1 \rangle_k \binom{x+k}{n} + \sum_k (n-k) \langle n-1 \rangle_{k-1} \binom{x+k}{n} \\ &= \sum_k (k+1) \langle n-1 \rangle_k \binom{x+k}{n} + \sum_k (n-k-1) \langle n-1 \rangle_k \binom{x+k+1}{n}. \end{aligned}$$

Combining the two sums on the right, and replacing the term $\binom{x+k+1}{n}$ by the sum $\binom{x+k}{n} + \binom{x+k}{n-1}$, we find that

$$\begin{aligned} F_n(x) &= \sum_k \langle n-1 \rangle_k \left(n \binom{x+k}{n} + (n-k-1) \binom{x+k}{n-1} \right) \\ &= \sum_k \langle n-1 \rangle_k \frac{(x+k)^{n-1}}{(n-1)!} ((x+k-n+1) + (n-k-1)) \end{aligned}$$

$$\begin{aligned}
 &= x \sum_k \left\langle \begin{matrix} n-1 \\ k \end{matrix} \right\rangle \binom{x+k}{n-1} \\
 &= xF_{n-1}(x).
 \end{aligned}$$

Therefore, $F_n(x) = x^n$, so we obtain

$$x^n = \sum_k \left\langle \begin{matrix} n \\ k \end{matrix} \right\rangle \binom{x+k}{n}, \quad n \geq 0. \quad (2.129)$$

This is known as *Worpitzky's identity* [287]. Thus, Eulerian numbers allow us to write ordinary powers as linear combinations of certain generalized binomial coefficients. For example, $x^4 = \binom{x}{4} + 11\binom{x+1}{4} + 11\binom{x+2}{4} + \binom{x+3}{4}$.

Exercises

1. Use an ordinary generating function to find a simple formula for $\left\langle \begin{matrix} n \\ 1 \end{matrix} \right\rangle$, and verify your formula using (2.128).
2. Let $E_n(x)$ denote the polynomial

$$E_n(x) = \sum_k \left\langle \begin{matrix} n \\ k \end{matrix} \right\rangle x^k.$$

Use (2.126) to show that the exponential generating function for the sequence of polynomials $\{E_n(x)\}_{n \geq 0}$ is

$$E(x, t) = \frac{1-x}{e^{t(x-1)} - x}.$$

That is, show that

$$E(x, t) = \sum_{n \geq 0} \frac{E_n(x)t^n}{n!}.$$

3. (From [282].) Use (2.126) and Exercise 11 of Section 2.8.3 to prove that

$$\sum_k \left\langle \begin{matrix} n \\ k \end{matrix} \right\rangle 2^k = \sum_k \left\{ \begin{matrix} n \\ k \end{matrix} \right\} k!$$

for any nonnegative integer n .

4. Use (2.128) to establish the following identity for $n \geq 1$:

$$\sum_{j=0}^n (-1)^j \binom{n}{j} (j+1)^{n-1} = 0.$$

5. A neurotic running back for an American football team will run between two offensive linemen only if the jersey number of the player on the left is less than the jersey number of the player on the right. The player will not run outside the last player on either end of the offensive line. The coach wants to be sure that the running back has at least three options on every play. If the coach always puts seven players on the offensive line, and there are fifteen players on the team capable of playing any position on the offensive line, each of whom has a different jersey number, how many formations of linemen are possible?

2.9 Stable Marriage

How do I love thee? Let me count the ways.

— Elizabeth Barrett Browning, Sonnet 43,
Sonnets from the Portuguese

Most of the problems we have considered in this chapter are questions in enumerative combinatorics, concerned with counting arrangements of objects subject to various constraints. In this section we consider a very different kind of combinatorial problem.

Suppose we must arrange n marriages between n men and n women. Each man supplies us with a list of the women ranked according to his preference; each woman does the same for the men. Is there always a way to arrange the marriages so that no unmatched man and woman prefer each other to their assigned spouses? Such a pairing is called a *stable matching*.

Consider a simple example with $n = 2$. Suppose Aaron prefers Yvonne over Zoë, and Björn prefers Zoë over Yvonne. We denote these preferences by

$$\begin{aligned} A : Y > Z, \\ B : Z > Y. \end{aligned}$$

Suppose also that Yvonne and Zoë both prefer Aaron over Björn, so

$$\begin{aligned} Y : A > B, \\ Z : A > B. \end{aligned}$$

Then the matching of Aaron with Zoë and Björn with Yvonne is unstable, since Aaron and Yvonne prefer each other over their partners. The preferences of Björn and Zoë are irrelevant: Indeed, Zoë would prefer to remain with Aaron in this case. On the other hand, the matching of Aaron with Yvonne and Björn with Zoë is stable, for no unmatched pair prefers to be together over their assigned partners.

The stable marriage problem is a question of existential combinatorics, since it asks whether a particular kind of arrangement exists. We might also consider it as a problem in constructive combinatorics, if we ask for an efficient algorithm for

finding a stable matching whenever one does exist. In fact, we develop just such an algorithm in Section 2.9.1.

The stable marriage problem and its variations have many applications in problems involving scheduling and assignments. We mention three examples.

1. Stable Roommates.

Suppose $2n$ students at a university must be paired off and assigned to n dorm rooms. Each student ranks all of the others in order of preference. A pairing is stable if no two unmatched students prefer to room with each other over their assigned partners. Must a stable pairing always exist? This variation of the stable marriage problem, known as *stable roommates*, is considered in Exercise 1.

2. College Admissions.

Suppose a number of students apply for admission to a number of universities. Each student ranks the universities, and each university ranks the students. Is there a way to assign the students to universities in such a way that no student and university prefer each other over their assignment? This problem is similar to the original stable marriage question, since we are matching elements from two sets using information on preferences. However, there are some significant differences—probably not every student applies to every university, and each university needs to admit a number of students. Some variations on the stable marriage problem that cover extensions like these are considered in Section 2.9.2.

3. Hospitals and Residents.

The problem of assigning medical students to hospitals for residencies is similar to the problem of matching students and universities: Each medical student ranks hospital residency programs in order of preference, and each hospital ranks the candidates. In this case, however, a program has been used to make most of the assignments in the U.S. since 1952. The National Resident Matching Program was developed by a group of hospitals to try to ensure a fair method of hiring residents. Since medical students are not obligated to accept the position produced by the matching program, it is important that the algorithm produce a stable matching. (Since the program's inception, a large majority of the medical students have accepted their offer.) We describe this matching algorithm in the next section.

Exercises

1. Suppose that four fraternity brothers, Austin, Bryan, Conroe, and Dallas, need to pair off as roommates. Each of the four brothers ranks the other three brothers in order of preference. Prove that there is a set of rankings for which no stable matching of roommates exists.

2. Suppose M_1 and M_2 are two stable matchings between n men and n women, and we allow each woman to choose between the man she is paired with in M_1 and the partner she receives in M_2 . Each woman always chooses the man she prefers. Show that the result is a stable matching between the men and the women.
3. Suppose that in the previous problem we assign each woman the man she likes less between her partners in the two matchings M_1 and M_2 . Show that the result is again a stable matching.
4. The following preference lists for four men, $\{A, B, C, D\}$, and four women, $\{W, X, Y, Z\}$, admit exactly ten different stable matchings.

	1	2	3	4
A	W	X	Y	Z
B	X	W	Z	Y
C	Y	Z	W	X
D	Z	Y	X	W

	1	2	3	4
W	D	C	B	A
X	C	D	A	B
Y	B	A	D	C
Z	A	B	C	D

- (a) Prove that the matching $\{(A, X), (B, Z), (C, W), (D, Y)\}$ is stable.
- (b) Determine the remaining nine stable matchings.

2.9.1 The Gale–Shapley Algorithm

Matchmaker, matchmaker, make me a match!

— Chava and Hodel, *Fiddler on the Roof*

In 1962, Gale and Shapley [117] proved that a stable matching between n men and n women always exists by describing an algorithm for constructing such a matching. Their algorithm is essentially the same as the one used by the hospitals to select residents, although apparently no one realized this for several years [143, chap. 1].

In the algorithm, we first choose either the men or the women to be the *proposers*. Suppose we select the men; the women will have their chance soon. Then the men take turns proposing to the women, and the women weigh the offers that they receive. More precisely, the Gale–Shapley algorithm has three principal steps.

Algorithm 2.17 (Gale–Shapley). *Construct a stable matching.*

Input. A set of n men, a set of n women, a ranked list of the n women for each man, and a ranked list of the n men for each woman.

Output. A stable matching that pairs the n men and n women.

Description.

Step 1. Label every man and woman as free.

Step 2. While some man m is free, do the following.

Let w be the highest-ranked woman on the preference list of m to whom m has not yet proposed. If w is free, then label m and w as engaged to each other. If w is engaged to m' and w prefers m over m' , then label m' as free and label m and w as engaged to one another. Otherwise, if w prefers m' over m , then w remains engaged to m' and m remains free.

Step 3. Match all of the engaged couples.

For example, consider the problem of arranging marriages between five men, Mack, Mark, Marv, Milt, and Mort, and five women, Walda, Wanda, Wendy, Wilma, and Winny. The men's and women's preferences are listed in Table 2.8.

	1	2	3	4	5
Mack	Winny	Wilma	Wanda	Walda	Wendy
Mark	Wanda	Winny	Wendy	Wilma	Walda
Marv	Winny	Walda	Wanda	Wilma	Wendy
Milt	Winny	Wilma	Wanda	Wendy	Walda
Mort	Wanda	Winny	Walda	Wilma	Wendy
Walda	Milt	Mort	Mack	Mark	Marv
Wanda	Milt	Marv	Mort	Mark	Mack
Wendy	Mort	Mack	Milt	Mark	Marv
Wilma	Mark	Mort	Milt	Mack	Marv
Winny	Marv	Mort	Mark	Milt	Mack

TABLE 2.8. Preferences for five men and women.

First, Mack proposes to Winny, who accepts, and Mark proposes to Wanda, who also accepts. Then Marv proposes to Winny. Winny likes Marv much better than her current fiancé, Mack, so Winny rejects Mack and becomes engaged to Marv. This leaves Mack without a partner, so he proceeds to the second name on his list, Wilma. Wilma currently has no partner, so she accepts. Our engaged couples are now

(Mack, Wilma), (Mark, Wanda), and (Marv, Winny).

Next, Milt proposes to his first choice, Winny. Winny prefers her current partner, Marv, so she rejects Milt. Milt proceeds to his second choice, Wilma. Wilma rejects Mack in favor of Milt, and Mack proposes to his third choice, Wanda. Wanda prefers to remain with Mark, so Mack asks Walda, who accepts. Our engaged couples are now

(Mack, Walda), (Mark, Wanda), (Marv, Winny), and (Milt, Wilma).

Now our last unmatched man, Mort, asks his first choice, Wanda. Wanda accepts Mort over Mark, then Mark asks his second choice, Winny. Winny rejects Mark in favor of her current partner, Marv, so Mark proposes to his third choice, Wendy. Wendy is not engaged, so she accepts. Now all the men and women are engaged, so we have our matching:

(Mack, Walda), (Mark, Wendy), (Marv, Winny),
(Milt, Wilma), and (Mort, Wanda).

We prove that this is in fact a stable matching.

Theorem 2.18. *The Gale–Shapley algorithm produces a stable matching.*

Proof. First, each man proposes at most n times, so the procedure must terminate after at most n^2 proposals. Thus, the procedure is an algorithm. Second, the algorithm always produces a matching. This follows from the observations that a woman, once engaged, is thereafter engaged to exactly one man, and every man ranks every woman, so the last unmatched man must eventually propose to the last unmatched woman. Third, we prove that the matching is stable. Suppose m prefers w to his partner in the matching. Then m proposed to w , and was rejected in favor of another suitor. This suitor is ranked higher than m by w , so w must prefer her partner in the matching to m . Therefore, the matching is stable. \square

We remark that the Gale–Shapley algorithm is quite efficient: A stable matching is always found after at most n^2 proposals. (Exercise 8 establishes a better upper bound.)

Suppose that we choose the women as the proposers. Does the algorithm produce the same stable matching? We test this by using the lists of preferences in Table 2.8. First, Walda proposes to Milt, who accepts. Next, Wanda proposes to Milt, and Milt prefers Wanda over Walda, so he accepts. Walda must ask her second choice, Mort, who accepts. Then Wendy proposes to Mort, who declines, so she asks Mack, and Mack accepts. Last, Wilma asks Mark, and Winny proposes to Marv, and both accept. We therefore obtain a different stable matching:

(Walda, Mort), (Wanda, Milt), (Wendy, Mack),
(Wilma, Mark), and (Winny, Marv).

Only Winny and Marv are paired together in both matchings; everyone else receives a higher-ranked partner precisely when he or she is among the proposers. Table 2.9 illustrates this for the two different matchings. The pairing obtained with the men as proposers is in boldface; the matching resulting from the women as proposers is underlined.

The next theorem shows that this is no accident. The proposers always obtain the best possible stable matching, and those in the other group, which we call the *proposees*, always receive the worst possible stable matching. We define two terms before stating this theorem. We say a stable matching is *optimal* for a person p if p can do no better in any stable matching. Thus, if p is matched with q in an optimal matching for p , and p prefers r over q , then there is no stable matching

	1	2	3	4	5
Mack	Winnie	Wilma	Wanda	Walda	<u>Wendy</u>
Mark	Wanda	Winnie	Wendy	<u>Wilma</u>	Walda
Marv	Winnie	Walda	Wanda	Wilma	Wendy
Milt	Winnie	Wilma	<u>Wanda</u>	Wendy	Walda
Mort	Wanda	Winnie	<u>Walda</u>	Wilma	Wendy
Walda	Milt	<u>Mort</u>	Mack	Mark	Marv
Wanda	<u>Milt</u>	Marv	Mort	Mark	Mack
Wendy	Mort	<u>Mack</u>	Milt	Mark	Marv
Wilma	<u>Mark</u>	Mort	Milt	Mack	Marv
Winnie	Marv	Mort	Mark	Milt	Mack

TABLE 2.9. Two stable matchings.

where p is paired with r . Similarly, a stable matching is *pegsimal* for p if p can do no worse in any stable matching. So if p is matched with q in a pegsimal matching for p , and p prefers q over r , then there is no stable matching where p is paired with r . Finally, a stable matching is optimal for a set of people P if it is optimal for every person p in P , and likewise for a pegsimal matching.

Theorem 2.19. *The stable matching produced by the Gale–Shapley algorithm is independent of the order of proposers, optimal for the proposers, and pegsimal for the proposees.*

Proof. Suppose the men are the proposers. We first prove that the matching produced by the Gale–Shapley algorithm is optimal for the men, regardless of the order of the proposers. Order the men in an arbitrary manner, and suppose that a man m and woman w are matched by the algorithm. Suppose also that m prefers a woman w' over w , denoted by $m : w' > w$, and assume that there exists a stable matching M with m paired with w' . Then m was rejected by w' at some time during the execution of the algorithm. We may assume that this was the first time a potentially stable couple was rejected by the algorithm. Say w' rejected m in favor of another man m' , so $w' : m' > m$. Then m' has no stable partner he prefers over w' , by our assumption. Let w'' be the partner of m' in the matching M . Then $w'' \neq w'$, since m is matched with w' in M , and so $m' : w' > w''$. But then m' and w' prefer each other to their partners in M , and this contradicts the stability of M .

The optimality of the matching for the proposers is independent of the order of the proposers, so the first statement in the theorem follows immediately.

Finally, we show that the algorithm is pegsimal for the proposees. Suppose again that the men are the proposers. Assume that m and w are matched by the algorithm, and that there exists a stable matching M where w is matched with a man m' and $w : m > m'$. Let w' be the partner of m in M . Since the Gale–Shapley algorithm produces a matching that is optimal for the men, we have $m :$

$w > w'$. Therefore, m and w prefer each other over their partners in M , and this contradicts the stability of M . \square

Exercises

1. Our four fraternity brothers, Austin, Bryan, Conroe, and Dallas, plan to ask four women from the neighboring sorority, Willa, Xena, Yvette, and Zelda, to a dance on Friday night. Each person's preferences are listed in the following table.

	1	2	3	4
Austin	Yvette	Xena	Zelda	Willa
Bryan	Willa	Yvette	Xena	Zelda
Conroe	Yvette	Xena	Zelda	Willa
Dallas	Willa	Zelda	Yvette	Xena
Willa	Austin	Dallas	Conroe	Bryan
Xena	Dallas	Bryan	Austin	Conroe
Yvette	Dallas	Bryan	Conroe	Austin
Zelda	Austin	Dallas	Conroe	Bryan

- (a) What couples attend the dance, if each man asks the women in his order of preference, and each woman accepts the best offer she receives?
- (b) Suppose the sorority hosts a "Sadie Hawkins" dance the following weekend, where the women ask the men out. Which couples attend this dance?
2. Determine the total number of stable matchings that pair the four men Axel, Buzz, Clay, and Drew with the four women Willow, Xuxa, Yetty, and Zizi, given the following preference lists.

	1	2	3	4
Axel	Yetty	Willow	Zizi	Xuxa
Buzz	Yetty	Xuxa	Zizi	Willow
Clay	Zizi	Yetty	Xuxa	Willow
Drew	Xuxa	Zizi	Willow	Yetty
Willow	Buzz	Drew	Axel	Clay
Xuxa	Buzz	Axel	Clay	Drew
Yetty	Drew	Clay	Axel	Buzz
Zizi	Axel	Drew	Buzz	Clay

3. Determine a list of preferences for four men and four women where no one obtains his or her first choice, regardless of who proposes.

4. Determine a list of preferences for four men and four women where one proposer receives his or her lowest-ranked choice.
5. Determine a list of preferences for four men and four women where one proposer receives his or her lowest-ranked choice, and the rest of the proposers receive their penultimate choice.
6. Suppose that all the men have identical preference lists in an instance of the stable marriage problem. Show that there exists exactly one stable matching by completing the following argument. Let M be the matching obtained by the Gale-Shapley algorithm using the men as proposers, and suppose another stable matching M' exists. Among all women who change partners between M and M' , let w be the woman who ranks lowest on the men's common preference list. Suppose m and w are matched in M , and m and w' in M' . Determine a contradiction.
7. Suppose that the preference lists of the men m_1, \dots, m_n and the women w_1, \dots, w_n have the property that m_i ranks w_i ahead of each of the women w_{i+1}, \dots, w_n , and w_i ranks m_i ahead of each of the men m_{i+1}, \dots, m_n , for each i .
 - (a) Show that the matching $(m_1, w_1), \dots, (m_n, w_n)$ is stable.
 - (b) (Eeckhout [86].) Show that this is the unique stable matching in this case.
 - (c) Prove that there are $(n!)^{n-1}$ different sets of preference lists for m_1, \dots, m_n that have the property that m_i ranks w_i ahead of each of the women w_{i+1}, \dots, w_n , for each i .
 - (d) Prove that at least $1/n!$ of the possible instances of the stable marriage problem for n couples admits a unique solution.
8. (Knuth [178].) Prove that the Gale-Shapley algorithm terminates after at most $n^2 - n + 1$ proposals by showing that at most one proposer receives his or her lowest-ranked choice.
9. Suppose that more than one woman receives her lowest-ranked choice when the men propose. Prove that there exist at least two stable matchings between the men and the women.

2.9.2 Variations on Stable Marriage

I want what any princess wants—to live happily ever after, with the ogre I married.

— Princess Fiona, *Shrek 2*

The stable marriage problem solves matching problems of a rather special sort. Each member of one set must rank all the members of the other set, and the two

sets must have the same number of elements. In this section, we consider several variations of the stable marriage problem, in order to apply this theory much more broadly. In each case, we study two main questions. First, how does the change affect the existence and structure of the stable pairings? Second, can we amend the Gale-Shapley algorithm to construct a stable matching in the new setting?

Unacceptable Partners

Suppose each of n men and n women ranks only a subset of their potential mates. Potential partners omitted from a person's list are deemed *unacceptable* to that person, and we do not allow any pairing in which either party is unacceptable to the other. Clearly, we cannot in general guarantee even a complete matching, since for instance a confirmed bachelor could mark all women as unacceptable. This suggests a modification of our notion of a stable matching for this problem. We say a matching (or partial matching) M is *unstable* if there exists a man m and woman w who are unmatched in M , each of whom is acceptable to the other, and each is either single in M , or prefers the other to their partner in M . We will show that every such problem admits a matching that is stable in this sense, and further that every stable matching pairs the same subcollection of men and women. We first require a preliminary observation. We say a person p prefers a matching M_1 over a matching M_2 if p strictly prefers his or her partner in M_1 to p 's match in M_2 .

Lemma 2.20. *Suppose M_1 and M_2 are stable matchings of n men and n women, whose preference lists may include unacceptable partners. If m and w are matched in M_1 but not in M_2 , then one of m or w prefers M_1 over M_2 , and the other prefers M_2 over M_1 .*

Proof. Suppose m_0 and w_0 are paired in M_1 but not M_2 . Then m_0 and w_0 cannot both prefer M_1 , since otherwise M_2 would not be stable. Suppose that both prefer M_2 . Then both have partners in M_2 , so suppose (m_0, w_1) and (m_1, w_0) are in M_2 . Both m_0 and w_1 cannot prefer M_2 , since M_1 is stable, so w_1 must prefer M_1 , and likewise m_1 must prefer M_1 . These two cannot be paired in M_1 , so denote their partners in M_1 by m_2 and w_2 . By the same reasoning, both of these people must prefer M_2 , but cannot be matched together in M_2 , so we obtain m_3 and w_3 , who prefer M_1 , but are not paired to each other in M_1 . We can continue this process indefinitely, obtaining a sequence $m_0, w_0, m_2, w_2, m_4, w_4, \dots$ of distinct men and women who prefer M_2 over M_1 , and another sequence $m_1, w_1, m_3, w_3, \dots$ of different people who prefer M_1 over M_2 . This is impossible, since there are only finitely many men and women. \square

We can now establish an important property of stable matchings when some unacceptable partners may be included: For a given set of preferences, every stable matching leaves the same group of men and women single.

Theorem 2.21. *Suppose each of n women ranks a subset of n men as potential partners, with the remaining men deemed unacceptable, and suppose each of the*

men rank the women in the same way. Then there exists a subset X_0 of the women and a subset Y_0 of the men such that every stable matching of the n men and n women leaves precisely the members of X_0 and Y_0 unassigned.

Proof. Suppose M_1 and M_2 are distinct stable matchings, and suppose m_1 is matched in M_1 but not in M_2 . Let w_1 be the partner of m_1 in M_1 . Since m_1 clearly prefers M_1 over M_2 , by Lemma 2.20 w_1 must prefer M_2 over M_1 . Let m_2 be the partner of w_1 in M_2 . Then m_2 prefers M_1 , and so his partner w_2 in M_1 must prefer M_2 over M_1 . Continuing in this way, we obtain an infinite sequence $(m_1, w_1), (m_2, w_2), (m_3, w_3), \dots$ of distinct couples in M_1 (and another sequence $(m_2, w_1), (m_3, w_2), (m_4, w_3), \dots$ in M_2), which is impossible. \square

We still need to show that at least one stable matching exists, and we can do this by altering the Gale-Shapley algorithm for preference lists that may include unacceptable partners. We require just two modifications. First, we terminate the loop either when all proposers are engaged, or when no free proposer has any remaining acceptable partners to ask. Second, proposals from unacceptable partners are always rejected. It is straightforward to show that this amended procedure always produces a stable matching (see Exercise 1). We can illustrate it with an example. Suppose the four men Iago, Julius, Kent, and Laertes each rank a subset of the four women Silvia, Thaisa, Ursula, and Viola, and each of the women ranks a subset of the men, as shown in Figure 2.21. Potential partners omitted from a person’s list are deemed unacceptable to that person, so for example Iago would not consider marrying Thaisa or Ursula.

	1	2	3
I	V	S	
J	S	V	
K	U	T	S
L	S	T	V

	1	2	3	4
S	I	K	L	J
T	J	K		
U	L	I	J	
V	K	J		

FIGURE 2.21. Preferences with unacceptable partners.

Suppose the men propose. Iago first asks Viola, but she rejects him as an unacceptable partner, so he asks Silvia, who happily accepts. Next, Julius asks Silvia, who rejects him in favor of Iago, so he proposes to Viola, who now accepts. Ursula then rejects Kent, then Thaisa accepts his proposal. Finally, Laertes proposes to Silvia, then Thaisa, then Viola, but each rejects him. Our stable matching is then (Iago, Silvia), (Julius, Viola), and (Kent, Thaisa). The set X_0 of unmatched bachelorettes contains only Ursula, and $Y_0 = \{\text{Laertes}\}$.

We have shown how to adapt the Gale-Shapley algorithm to handle incomplete preference lists, but we can also describe a way to alter the data in such a way that we can apply the Gale-Shapley algorithm without any modifications. To do this, we introduce a fictitious man to mark the boundary between the acceptable and unacceptable partners on each woman’s list, and similarly introduce a fictitious

woman for the men’s lists. We’ll call our invented man the *ogre*, and our fictitious woman, the *ogress*. Append the ogre to each woman’s ranked list of acceptable partners, then add her unacceptable partners afterwards in an arbitrary order. Thus, each woman would sooner marry an ogre than one of her unacceptable partners. Do the same for the men with the ogress. The ogre prefers any woman over the ogress, and the ogress prefers any man over the ogre (people are tastier!), but the rankings of the humans on the ogre’s and ogress’ lists are immaterial. For example, we can augment the preference lists of Figure 2.21 to obtain the 5×5 system of Figure 2.22, using **M** to denote the ogre and **W** for the ogress.

	1	2	3	4	5
I	V	S	W	T	U
J	S	V	W	T	U
K	U	T	S	W	V
L	S	T	V	W	U
M	S	T	U	V	W

	1	2	3	4	5
S	I	K	L	J	M
T	J	K	M	I	L
U	L	I	J	M	K
V	K	J	M	I	L
W	I	J	K	L	M

FIGURE 2.22. Augmented preference lists.

We can now characterize when the original configuration has a complete stable matching, that is, a stable pairing where no one is left single.

Theorem 2.22. *Suppose each of n men ranks some subset of n women as acceptable partners, and each of the women does the same for the men. Suppose further that we obtain an instance of the standard stable marriage problem on $n + 1$ men and women by adding an ogre **M** and ogress **W**, and augmenting the preference lists in the manner described above. Then the original system has a complete stable matching if and only if the augmented system has a stable matching where **M** is paired with **W**.*

Proof. Suppose the original system has a complete stable matching. Then each woman prefers her partner in this matching to the ogre under the augmented preferences, and likewise no man would leave his partner for the ogress. Thus, adding (\mathbf{M}, \mathbf{W}) to this pairing produces a stable matching for the augmented system. Next, suppose the augmented system has a stable matching P' that includes (\mathbf{M}, \mathbf{W}) , and let $P = P' \setminus \{(\mathbf{M}, \mathbf{W})\}$. Suppose $(m, w) \in P$. If m is unacceptable to w , then w would prefer the ogre **M** over m , and certainly **M** prefers w over **W**. This contradicts the stability of P' . Similarly, w must be acceptable to m . Thus, P is a complete matching of mutually acceptable partners, and stability follows at once from the stability of P' . □

Exercise 2 asks you to show that **M** and **W** must be paired together in all stable matchings of the augmented system, if they are paired in any particular stable matching. Thus, we can determine if a complete stable matching exists by running the original Gale-Shapley algorithm on the augmented preference lists, choosing either set as the proposers.

While applying the Gale-Shapley algorithm in this way always produces a matching that is stable with respect to the augmented preferences, it is important to note that restricting such a pairing back to the original preferences might not produce a stable matching! For example, when the men propose using the augmented lists of Figure 2.22, we obtain the stable matching

$$\begin{aligned} &(\text{Iago, Silvia}), (\text{Julius, Viola}), (\text{Kent, Ursula}), \\ &(\text{Laertes, Ogress}), (\text{Ogre, Thaisa}). \end{aligned} \tag{2.130}$$

However, Kent is not acceptable to Ursula, so we must disband this pair when we restrict to the original preference lists. The surviving pairs are (Iago, Silvia) and (Julius, Viola), and now Kent and Thaisa are unmatched but mutually acceptable.

Indifference

In the original stable marriage problem, we required that all preferences be *strictly ordered*, since each person needed to assign each potential partner a different rank. However, rankings often contain items that are valued equally. What happens if we allow *weakly ordered* rankings, that is, rankings that may contain some elements of the same rank? Suppose that each of n men supplies a weak ordering of a set of n women, and each of the women does the same for the men. We'll assume for now that all rankings are complete, so there are no unacceptable partners. Must a stable ranking exist? Can we construct one?

We first require a clarification of our notion of stability for this situation. We say a matching M of the men and women is *unstable* if there exists an unmatched couple m and w , each of whom strictly prefers the other to his or her partner in M . For example, if m strictly prefers w to his partner, but w ranks m equal to her partner, then the pair m and w do not violate stability under this definition.

One can certainly study this problem with other notions of stability. For instance, one could demand that no unmatched man and woman weakly prefer each other to their assigned partners. A matching with no such couples is called *super-stable*. Or one could require that no unmatched couple prefer each other, one in a strict sense and the other in a weak manner. Such a matching is said to be *strongly stable*. Since the notion that we employ is the least restrictive, matchings with this property are often called *weakly stable*.

Given a collection of weakly ordered preference lists for n men and n women, we can certainly create a corresponding set of strongly ordered preference lists by breaking each tie in an arbitrary way. We call the strongly ordered preferences a *refinement* of the original weak preferences. A stable matching for the refined lists certainly exists, and it is easy to see that this matching is also a (weakly) stable matching for the original, weakly ordered lists. Furthermore, every stable matching for the original preferences can be obtained in this way. We can summarize these facts in the following theorem.

Theorem 2.23. *Suppose each of n men ranks a collection of n women, with tied rankings allowed, and each woman does the same for the men. Then a stable*

matching for these preferences exists, and further every such stable matching is a stable matching for some refinement of these weakly ordered rankings.

Proof. For the first part, let P' be a refinement of the given list of preferences P , and let M be a stable matching for P' . If m and w are unmatched in M , and according to P strictly prefer each other to their partners in this matching, then they also strictly prefer each other according to P' . This is impossible, since M is stable with respect to P' . Thus, M is stable with respect to P .

For the second part, suppose M is a stable matching with respect to P . We need to construct a refinement P' of P where M is stable. If $(m, w) \in M$, and m ranks w' equal to w in P , then let m rank w ahead of w' in P' . Likewise, if w ranks m' equal to m in P , then w ranks m ahead of m' in P' . Any remaining tied rankings are broken arbitrarily to complete P' . Suppose then that m_0 and w_0 are unmatched in P' , but prefer each other (according to P') to their partners in M . Since M is stable with respect to P , then either m_0 ranks w_0 equal to his partner in M , or w_0 ranks m_0 equal to her partner in M (or both). We obtain a contradiction in either case, by the construction of P' . \square

$$\begin{array}{ll}
 G : D > A = C > F & A : I > G = H = K \\
 H : A = F > C = D & C : H > G = I > K \\
 I : F > C > D > A & D : I > K > H > G \\
 K : D > A = C = F & F : H = I > G = K
 \end{array}$$

FIGURE 2.23. Preference lists with indifference.

The Gale-Shapley algorithm requires no modifications for this variation, once a refinement is selected. Of course, the algorithm may produce different matchings for different refinements, even when the same group proposes. For example, suppose the four men Gatsby, Hawkeye, Ishmael, and Kino, and four women Apolonia, Cora, Daisy, and Fayaway, submit the preference lists shown in Figure 2.23. Using the refinement obtained by replacing each = in these lists with $>$, the Gale-Shapley algorithm produces the following matching when the men propose:

$$\begin{array}{l}
 (\text{Gatsby, Apolonia}), (\text{Hawkeye, Fayaway}), \\
 (\text{Ishmael, Cora}), (\text{Kino, Daisy}).
 \end{array} \tag{2.131}$$

However, if we reverse the order of Apolonia and Cora in the refinement of Gatsby's list, and the order of Apolonia and Fayaway in Hawkeye's, we then obtain a very different stable matching:

$$\begin{array}{l}
 (\text{Gatsby, Cora}), (\text{Hawkeye, Fayaway}), \\
 (\text{Ishmael, Daisy}), (\text{Kino, Apolonia}).
 \end{array} \tag{2.132}$$

Finally, we may also ask about combining this extension of the stable marriage problem with the prior one. Suppose the men and women supply weakly ordered

rankings, and may also declare some potential partners as unacceptable. The stable matching problem becomes much more complicated in this case. Even the size of a stable matching may vary, in contrast to the case of unacceptable partners with strict rankings, where Theorem 2.21 guarantees that all stable matchings have not only the same size, but match exactly the same men and women. For example, consider the following 2×2 system from [196], where A finds Y acceptable but not Z , and Z finds B acceptable but not A .

$$\begin{array}{ll} A : Y & Y : A = B \\ B : Y > Z & Z : B \end{array}$$

These preferences admit exactly two stable matchings, which have different sizes: $\{(A, Y), (B, Z)\}$ and $\{(B, Y)\}$.

We might ask if we could determine a stable matching of maximal size in a problem like this, since this would often be desirable. However, no fast algorithm is known for computing this in the general $n \times n$ case. (Here, a “fast” algorithm would have its running time bounded by a polynomial in n .) In fact, it is known [196] that this problem belongs to a family of difficult questions known as NP-complete problems. The problem remains hard even if ties are allowed in only the men’s or only the women’s preferences, and all ties occur at the end of each list, even if each person is allowed at most one tied ranking.

Sets of Different Sizes

Every stable marriage problem we have considered so far required an equal number of men and women. Suppose now that one group is larger than the other. Of course, we could not possibly match everyone with a partner now, but can we find a stable matching that pairs everyone in the smaller set? Here, we say a matching (or partial matching) M is *unstable* if there exists a man m and woman w , unmatched in M , such that each is either single in M , or prefers the other to his or her partner in M .

We can solve this variation by considering it to be a special case of the problem with unacceptable partners. Suppose we have k men and n women, with $n > k$. Suppose also that each of the men rank each of the women in strict order, and each of the women reciprocate for the men. We introduce $n - k$ ghosts to the set of men. Each ghost finds no woman to be an acceptable partner, and each women would not accept any ghost. Then a stable matching exists by the modified Gale-Shapley algorithm for unacceptable partners, and by Theorem 2.21 there exists a set X_0 of women and Y_0 of ghosts and men such that the members of X_0 and Y_0 are precisely the unassigned parties in any stable matching. Certainly Y_0 includes all the ghosts, since they have no acceptable partners. But no man can be unassigned in a stable matching, since each man is acceptable to all the women. Thus, X_0 is empty and Y_0 is precisely the set of ghosts, and we obtain the following theorem.

Theorem 2.24. *Suppose each of k men ranks each of n women in a strict ordering, and each of the women ranks the men in the same way. Then*

- (i) *a stable matching exists,*

- (ii) every stable matching pairs every member of the smaller set, and
- (iii) there exists a subset X of the larger set such that every stable matching leaves the members of X unassigned, and the others all matched.

An example with groups of different sizes appears in Exercise 6. Some other interesting variations (and combinations of variations) on the stable marriage problem are introduced in the exercises too. We will study marriage problems further in Chapter 3, where in Section 3.8 we investigate matchings for various *infinite* sets.

Exercises

1. Prove that the Gale-Shapley algorithm, amended to handle unacceptable partners, always produces a stable matching.
2. Prove that if the ogre and ogress are paired in some stable matching for an augmented system of preferences as in Theorem 2.22, then they must be paired in every such stable matching.
3.
 - (a) Verify the stable matching (2.130) produced by the Gale-Shapley algorithm when the men propose using the preferences in Figure 2.22.
 - (b) Compute the stable matching obtained when the women propose using these preferences. Does this pairing restrict to a stable matching for Figure 2.21?
 - (c) In the augmentation procedure for the case of unacceptable partners, we can list the unacceptable partners for each person in any order after the ogre or ogress, and we can list the humans in any order in the lists for the ogre and ogress. Show that one can select orderings when augmenting the preferences of Figure 2.21 so that when the men propose in the Gale-Shapley algorithm, one obtains a pairing that restricts to a stable matching of Figure 2.21.
4. The following problems all refer to the weakly ordered preference lists of Figure 2.23.
 - (a) Verify the matching (2.131) obtained from the refinement obtained by replacing each = with $>$, when the men propose in the Gale-Shapley algorithm. Then determine the matching obtained when the women propose.
 - (b) Verify (2.132) using the refinement obtained from the previous one by reversing the order of Apolonia and Cora in Gatsby's list, and Apolonia and Fayaway in Hawkeye's. Then determine the matching obtained when the women propose.
 - (c) Construct another refinement by ranking any tied names in reverse alphabetical order. Compute the stable matchings constructed by the

Gale-Shapley algorithm when the men propose, then when the women propose.

- Construct three refinements of the following preference lists so that the Gale-Shapley algorithm, amended for unacceptable partners, produces a stable matching of a different size in each case.

$$\begin{array}{ll}
 A : W & W : A = B = C = D \\
 B : W > X & X : B = C = D \\
 C : W > X > Y & Y : C = D \\
 D : W > X > Y > Z & Z : D
 \end{array}$$

- Suppose the five men Arceneaux, Boudreaux, Comeaux, Duriaux, and Gautreaux, each rank the three women Marteaux, Robichaux, and Thibodeaux in order of preference, and the women each rank the men, as shown in the following tables.

	1	2	3
A	R	T	M
B	T	R	M
C	M	T	R
D	T	M	R
G	R	T	M

	1	2	3	4	5
M	A	D	B	C	G
R	D	G	A	C	B
T	G	A	D	C	B

Determine the stable matching obtained when the men propose, then the matching found when the women propose. What is the set X of Theorem 2.24 for these preferences?

- Suppose we allow weakly ordered rankings in the hypothesis of Theorem 2.24. Determine which of the conclusions still hold, and which do not necessarily follow. Supply a proof for any parts that do hold, and supply a counterexample for any parts that do not.
- Suppose that each of n students, denoted S_1, S_2, \dots, S_n , ranks each of m universities, U_1, U_2, \dots, U_m , and each university does the same for the students. Suppose also that university U_k has p_k open positions. We say an assignment of students to universities is *unstable* if there exists an unpaired student S_i and university U_j such that S_i is either unassigned, or prefers U_j to his assignment, and U_j either has an unfilled position, or prefers S_i to some student in the new class.
 - Assume that $\sum_{k=1}^m p_k = n$. Explain how to amend the preference lists so that the Gale-Shapley algorithm may be used to compute a stable assignment of students to universities, with no university exceeding its capacity.

- (b) Repeat this problem without assuming that the number of students matches the total number of open positions.
- (c) Suppose each student ranks only a subset of the universities, and each university ranks only a subset of the students who apply to that school. Assume that unranked possibilities are unacceptable choices. Modify the definition of stability for this case, then describe how to use the Gale-Shapley algorithm to determine a stable assignment.
9. Suppose that each of n students, denoted S_1, S_2, \dots, S_n , needs to enroll in a number of courses from among m possible offerings, denoted C_1, C_2, \dots, C_m . Assume that student S_i can register for up to q_i courses, and course C_j can admit up to r_j students. An *enrollment* is a set of pairs (S_i, C_j) where each student S_i appears in at most q_i such pairs, and each course C_j appears in at most r_j pairs. Suppose each student ranks a subset of acceptable courses in order of preference, and the supervising professor of each course ranks a subset of acceptable students. Define a *stable enrollment* in an appropriate way.

2.10 Combinatorial Geometry

We should expose the student to some material that has strong intuitive appeal, is currently of research interest to professional mathematicians, and in which the student himself may discover interesting problems that even the experts are unable to solve.

— Victor Klee, from the translator's preface to
Combinatorial Geometry in the Plane [144]

The subject of *combinatorial geometry* studies combinatorial problems regarding arrangements of points in space, and the geometric figures obtained from them. Such figures include lines and polygons in two dimensions, planes and polyhedra in three, and hyperplanes and polytopes in n -dimensional space. This subject has much in common with the somewhat broader subject of *discrete geometry*, which treats all sorts of geometric problems on discrete sets of points in Euclidean space, especially extremal problems concerning quantities such as distance, direction, area, volume, perimeter, intersection counts, and packing density.

In this section, we provide an introduction to the field of combinatorial geometry by describing two famous problems regarding points in the plane: a question of Sylvester concerning the collection of lines determined by a set of points, and a problem of Erdős, Klein, and Szekeres on the existence of certain polygons that can be formed from large collections of points in the plane. The latter problem leads us again to Ramsey's theorem, and we prove this statement in a more general form than what we described in Section 1.8. (Ramsey theory is developed further in Chapter 3.) In particular, we establish some of the bounds on the Ramsey numbers $R(p, q)$ that were cited in Section 1.8.

2.10.1 Sylvester's Problem

Thufferin' thuccotash!

— Sylvester the cat, *Looney Tunes*

James Joseph Sylvester, a British-born mathematician, spent the latter part of his career at Johns Hopkins University, where he founded the first research school in mathematics in America, and established the first American research journal in the subject, *The American Journal of Mathematics*. Toward the end of his career, Sylvester posed the following problem in 1893, in the “Mathematical Questions” column of the British journal, *Educational Times* [265].

Sylvester's Problem. *Given $n \geq 3$ points in the plane which do not all lie on the same line, must there exist a line that passes through exactly two of them?*

Given a collection of points in the plane, we say a line is *ordinary* if it passes through exactly two of the points. Thus, Sylvester's problem asks if an ordinary line always exists, as long as the points are not all on the same line.

This problem remained unsolved for many years, and seemed to have been largely forgotten until Erdős rediscovered it in 1933. Tibor Gallai, a friend of Erdős' who is also known as T. Grünwald, found the first proof in the same year. Erdős helped to revive the problem by posing it in the “Problems” section of the *American Mathematical Monthly* in 1933 [89], and Gallai's solution was published in the solution the following year [264].

Kelly also produced a clever solution, which was published in a short article by Coxeter in 1948 [62], along with a version of Gallai's argument. Forty years later, the computer scientist Edsger Dijkstra derived a similar proof, but with a more algorithmic viewpoint [76]. The proof we present here is based on Dijkstra's algorithm. Given any collection of three or more points which do not all lie on the same line, it constructs a line with the required property.

In this method, we start with an arbitrary line ℓ_1 connecting at least two points of the set, and some point S_1 from the set that does not lie on ℓ_1 . If ℓ_1 contains just two of the points, we are done, so suppose that at least three of the points lie on ℓ_1 . The main idea of the method is to construct from the current line ℓ_1 and point S_1 another line ℓ_2 and point S_2 , with S_2 not on ℓ_2 . Then we iterate this process, constructing ℓ_3 and S_3 , then ℓ_4 and S_4 , etc., until one is assured of obtaining a line that connects exactly two of the points of the original collection. In order to ensure that the procedure does not cycle endlessly, we introduce a *termination argument*: a strictly monotone function of the state of the algorithm. A natural candidate is the distance d_k from the point S_k to the line ℓ_k , so $d_k = d(S_k, \ell_k)$. We therefore aim to construct ℓ_{k+1} and S_{k+1} from ℓ_k and S_k in such a way that $d_{k+1} < d_k$. Since there are only finitely many points, there are only finitely many possible values for d_k , so if we can achieve this monotonicity, then it would follow that the procedure must terminate.

We derive a procedure that produces a strictly decreasing sequence $\{d_k\}$. Suppose the line ℓ_k contains the points P_k , Q_k , and R_k from our original collection, and S_k is a point from the set that does not lie on ℓ_k . We need to choose ℓ_{k+1} and

S_{k+1} so that $d_{k+1} < d_k$. Suppose we set S_{k+1} to be one of the points that we labeled on ℓ_k , say $S_{k+1} = Q_k$. Certainly Q_k does not lie on either of the lines $\overline{P_k S_k}$ or $\overline{R_k S_k}$, so we might choose one of these two lines for our ℓ_{k+1} . Can we guarantee that one of these choices will produce a good value for d_{k+1} ? To test this, let

$$p_k = d(Q_k, \overline{P_k S_k})$$

and

$$r_k = d(Q_k, \overline{R_k S_k}).$$

We require then that

$$\min(p_k, r_k) < d_k. \tag{2.133}$$




FIGURE 2.24. Similar triangles in the construction.

Using similar triangles in Figure 2.24, we see that the inequality $p_k < d_k$ is equivalent to the statement

$$d(P_k, Q_k) < d(P_k, S_k), \tag{2.134}$$

and likewise $r_k < d_k$ is equivalent to the inequality

$$d(Q_k, R_k) < d(S_k, R_k). \tag{2.135}$$

Now at least one of (2.134) or (2.135) must hold if

$$d(P_k, Q_k) + d(Q_k, R_k) < d(P_k, S_k) + d(S_k, R_k).$$

Further, since S_k does not lie on ℓ_k , by the triangle inequality we know that

$$d(P_k, R_k) < d(P_k, S_k) + d(S_k, R_k).$$

Therefore, inequality (2.133) follows from the statement

$$d(P_k, Q_k) + d(Q_k, R_k) \leq d(P_k, R_k).$$

However, by the triangle inequality, we know that

$$d(P_k, Q_k) + d(Q_k, R_k) \geq d(P_k, R_k).$$

Thus, we require that

$$d(P_k, Q_k) + d(Q_k, R_k) = d(P_k, R_k).$$

Clearly, this latter condition holds if and only if Q_k lies between P_k and R_k on ℓ_k . We therefore obtain the following algorithm for solving Sylvester's problem.

Algorithm 2.25. *Construct an ordinary line.*

Input. A set of $n \geq 3$ points in the plane, not all on the same line.

Output. A line connecting exactly two of the points.

Description.

Step 1. Let ℓ_1 be a line connecting at least two of the points in the given set, and let S_1 be a point from the collection that does not lie on ℓ_1 . Set $k = 1$, then perform Step 2.

Step 2. If ℓ_k contains exactly two points from the original collection, then output ℓ_k and stop. Otherwise, perform Step 3.

Step 3. Let P_k , Q_k , and R_k be three points from the given set that lie on ℓ_k , with Q_k lying between P_k and R_k . Set $S_{k+1} = Q_k$, and set $\ell_{k+1} = \overline{P_k S_k}$ if $d(Q_k, \overline{P_k S_k}) < d(Q_k, \overline{P_k R_k})$; otherwise set $\ell_{k+1} = \overline{R_k S_k}$. Then increment k by 1 and repeat Step 2.

Now Sylvester's problem is readily solved: The monotonicity of the sequence $\{d_k\}$ guarantees that the algorithm must terminate, so it must produce a line connecting just two points of the given set. An ordinary line must therefore always exist.

We can illustrate Dijkstra's algorithm with an example. Figure 2.25 shows a collection of thirteen points that produce just six ordinary lines (shown in bold), along with 21 lines that connect at least three of the points. Figure 2.26 illustrates the action of Algorithm 2.25 on these points, using a particular initial configuration. Each successive diagram shows the line ℓ_k , the point S_k off the line, and the points P_k , Q_k , and R_k on the line.

Much more is now known about Sylvester's problem. For example, Csima and Sawyer [64, 65] proved that every arrangement of $n \geq 3$ points in the plane, not all on the same line, must produce at least $6n/13$ ordinary lines, except for certain arrangements of $n = 7$ points. Figure 2.25 shows that this bound is best possible, and Exercise 2 asks you to determine an exceptional configuration for $n = 7$. Also, it has long been conjectured that there are always at least $\lceil n/2 \rceil$ ordinary lines for a set of n non-colinear points, except for $n = 7$ and $n = 13$, but this remains unresolved. For additional information on Sylvester's problem and several of its generalizations, see the survey article by Borwein and Moser [34], or the book by Brass, Moser, and Pach [37, sec. 7.2].




FIGURE 2.25. A collection of thirteen points with just six ordinary lines.

Exercises

1. Exhibit an arrangement of six points in the plane that produce exactly three ordinary lines.
2. Exhibit an arrangement of seven points in the plane that produce exactly three ordinary lines.
3. Exhibit an arrangement of eight points in the plane that produce exactly four ordinary lines.
4. Exhibit an arrangement of nine points in the plane that produce exactly six ordinary lines.
5. Suppose $n \geq 3$ points in the plane do not all lie on the same line. Show that if one joins each pair of points with a straight line, then one must obtain at least n distinct lines.
6. We say a set of points B is *separated* if there exists a positive number δ such that the distance $d(P, Q) \geq \delta$ for every pair of points P and Q in B . Describe an infinite, separated set of points in the plane, not all on the same line, for which no ordinary line exists. What happens if you apply Dijkstra's algorithm to this set of points?
7. Repeat problem 6, if each of the points (x, y) must in addition satisfy $|y| \leq 1$.

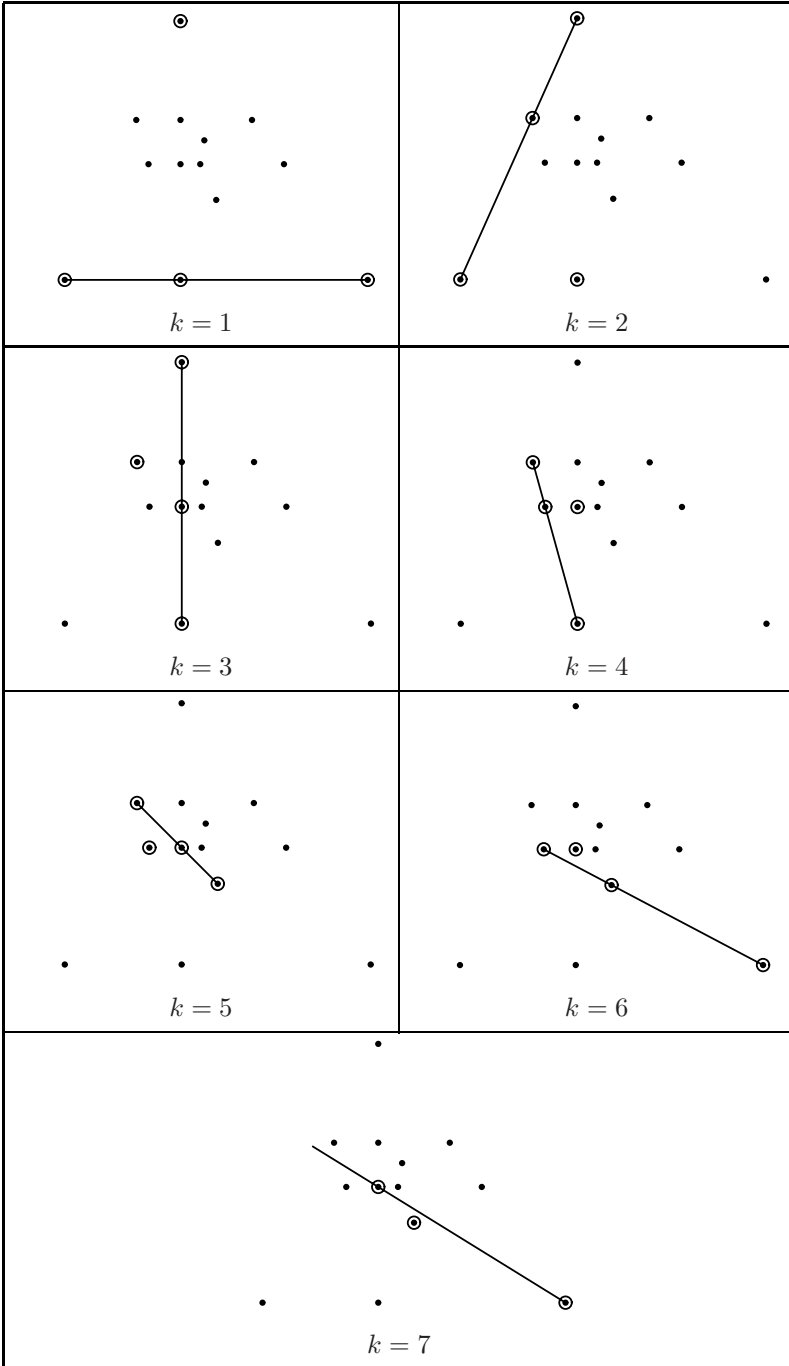


FIGURE 2.26. Dijkstra's algorithm.

8. Let the set S consist of the point $(0, 0)$, together with all the points in the plane of the form $(\frac{1}{3k-1}, \frac{1}{3k-1})$, $(\frac{-1}{3k-1}, \frac{1}{3k-1})$, or $(0, \frac{2}{3k-2})$, where k is an arbitrary integer. Show that every line connecting two points of S must intersect a third point of S .
9. Consider the following collection T of three-element subsets of the seven-element set $S = \{a, b, c, d, e, f, g\}$:
- $$T = \{\{a, b, c\}, \{a, d, e\}, \{a, f, g\}, \{b, d, f\}, \{b, e, g\}, \{c, d, g\}, \{c, e, f\}\}.$$
- (a) Verify that each two-element subset of S is in fact a subset of one of the members of T , and that any two distinct sets in T have at most one element in common.
- (b) Explain how this example is germane to Sylvester's problem. Hint: Try thinking of the elements of S as points, and the elements of T as lines.

2.10.2 Convex Polygons

I would certainly pay \$500 for a proof of Szekeres' conjecture.

— Paul Erdős, [92, p. 66]

A set of points S in the plane is said to be *convex* if for each pair of points a and b in S , the line segment joining a to b lies entirely in S . Loosely, then, a convex set has no "holes" in its interior, and no "dents" in its boundary. Line segments, triangles, rectangles, and ellipses are thus all examples of convex sets.

The *convex hull* of a finite collection of points T in the plane is defined as the intersection of all closed convex sets which contain T . Less formally, if one imagines T represented by a set of pushpins in a bulletin board, then the convex hull of T is the shape enclosed by a rubber band when it is snapped around all the pushpins. The convex hull of a set of three points then is either a triangle or a line segment, and for four points we may obtain one of these shapes, or a convex quadrilateral.

In order to avoid degenerate cases, we will assume in this section that our given collection of points is in *general position*, which means that no three points lie on the same line, or, using the term from the previous section, that each line connecting two of the points is ordinary. Thus, the convex hull of a set of four points in general position forms either a quadrilateral, or a triangle whose interior contains the fourth point of the collection. In the early 1930s, Esther Klein observed that one can always find a convex quadrilateral in a collection of five points in general position.

Theorem 2.26. *Any collection of five points in the plane in general position contains a four-element subset whose convex hull is a quadrilateral.*

Proof. Suppose we are given a collection of five points in the plane, with no three on the same line. If their convex hull is a pentagon or a quadrilateral, then the

statement follows, so suppose that it forms a triangle. Let a and b be the two points of the collection lying inside the triangle, and let ℓ be the line connecting a and b . Since the points are in general position, two of the vertices of the triangle lie on one side of ℓ . Label them c and d . Then the convex hull of $\{a, b, c, d\}$ is a quadrilateral. See Figure 2.27. \square




FIGURE 2.27. A convex quadrilateral may always be found among five points in general position.

Klein then asked about a natural generalization. How many points in the plane (in general position) are required in order to be certain that some subset forms the convex hull of a polygon with n sides? Does such a number exist for each n ? For example, Figure 2.28 illustrates a collection of eight points, no five of which produce a convex pentagon, and a set of sixteen points, no six of which forms a convex hexagon. Thus, at least nine points are needed for $n = 5$, and at least seventeen for $n = 6$.




FIGURE 2.28. Eight points with no convex pentagon, and sixteen points with no convex hexagon.

Erdős and Szekeres studied this problem in their first joint paper, in 1935 [94]. There they independently developed a version of Ramsey's theorem, and the proof we describe in this section is based on their argument. The statement we develop here is much more general than the special case of Ramsey's theorem that we described in Section 1.8, although Ramsey in fact established a still more general

result in his seminal paper of 1930 [232] (see Exercise 7). We will also derive the bounds on the ordinary Ramsey numbers $R(m, n)$ stated in Theorems 1.63 and 1.64 of Section 1.8 as special cases.

Let $ES(n)$ denote the minimal number of points in the plane in general position that are required so that there must exist a subcollection of n points whose convex hull is a polygon with n sides (an n -gon). Thus, we have seen that $ES(3) = 3$, $ES(4) = 5$, and, from Figure 2.28, that $ES(5) \geq 9$ and $ES(6) \geq 17$. We aim to show that $ES(n)$ exists for each n by obtaining an upper bound on its value, in terms of n . As a first step, we show that it is enough to find a collection of n points, each of whose four-element subsets forms a convex quadrilateral.

Theorem 2.27. *Suppose S is a set of n points in the plane in general position with the property that each four-element subset of S is the vertex set of a convex quadrilateral. Then S is the set of vertices of a convex n -gon.*

Proof. Let H denote the convex hull of S , and suppose $a \in S$ lies in the interior of H . Let $b \in S$ with $a \neq b$. Divide H into triangles by joining b to each vertex of H . Then a lies in the interior of one of these triangles, and we label its vertices b , c , and d . But then $\{a, b, c, d\}$ is a four-element subset of S whose convex hull is a triangle, contradicting our assumption. \square

Next, we develop the more general version of Ramsey's theorem. Recall that in Section 1.8 we defined $R(m, n)$ to be the smallest positive integer N such that any 2-coloring of the edges of the complete graph K_N (using the colors red and blue) must produce either a red K_m or a blue K_n as a subgraph. Coloring each edge of K_N is certainly equivalent to assigning a color to each of the $\binom{N}{2}$ subsets of size 2 of the set $\{1, 2, \dots, N\}$, and so we might consider what happens more generally when we assign a color to each of the $\binom{N}{k}$ subsets of size k , for a fixed positive integer k . We call such a subset a k -subset of the original set. Ramsey's theorem extends in a natural way to this setting. For convenience, we let $[N]$ denote the set $\{1, 2, \dots, N\}$, and we define the generalized Ramsey numbers in the following way.

Definition. For positive integers k , m , and n , with $m \geq k$ and $n \geq k$, the *Ramsey number* $R_k(m, n)$ is defined as the smallest positive integer N such that in any 2-coloring of the k -subsets of $[N]$ (using the colors red and blue) there must exist either a subset of m elements, each of whose k -subsets is red, or a subset of n elements, each of whose k -subsets is blue.

Thus, the Ramsey numbers $R(m, n)$ of Section 1.8 are denoted by $R_2(m, n)$ here. Also, just as the ordinary Ramsey numbers can be described in terms of coloring edges of complete graphs, so too can we describe $R_k(m, n)$ in terms of coloring edges of certain hypergraphs (see Exercise 1).

The next theorem establishes that the Ramsey numbers $R_k(m, n)$ always exist, and provides an upper bound on their values.

Theorem 2.28 (Ramsey’s Theorem). *Let k , m , and n be positive integers, with $\min\{m, n\} \geq k$. Then the Ramsey number $R_k(m, n)$ exists. Furthermore, for each such k , m , and n , we have*

$$R_1(m, n) = m + n - 1, \tag{2.136}$$

$$R_k(k, n) = n, \tag{2.137}$$

$$R_k(m, k) = m, \tag{2.138}$$

and, if $\min\{m, n\} > k \geq 2$, then

$$R_k(m, n) \leq R_{k-1}(R_k(m - 1, n) + R_k(m, n - 1)) + 1. \tag{2.139}$$

Proof. First, consider the case $k = 1$. If the elements of $[N]$ are each colored red or blue, and there are fewer than m red elements and fewer than n blue elements, then certainly $N \leq m + n - 2$, and (2.136) follows.

Second, suppose $k = m$, and suppose that each k -subset of $[N]$ is colored red or blue. If any is red then we have a qualifying m -subset, so suppose all are blue. Then we have a qualifying n -subset precisely when $N \geq n$. Thus, the formula (2.137) follows, and by symmetry so does (2.138).

To establish (2.139), suppose $\min\{m, n\} > k \geq 2$. Using induction on k , we may assume that $R_{k-1}(a, b)$ exists for all integers a and b with $\min\{a, b\} \geq k - 1$, and further by induction on $m + n$ we may assume that $R_k(m - 1, n)$ and $R_k(m, n - 1)$ both exist. Let $m' = R_k(m - 1, n)$, $n' = R_k(m, n - 1)$, and $N = R_{k-1}(m', n') + 1$, and consider an arbitrary 2-coloring C of the k -subsets of $[N]$ using the colors red and blue. Create a coloring C' of the $(k - 1)$ -subsets of $[N - 1]$ by assigning a subset X of size $k - 1$ the color of the set $X \cup \{N\}$ in C . Since $N - 1 = R_{k-1}(m', n')$, the coloring C' must produce either a subset of $[N - 1]$ of cardinality m' , each of whose $(k - 1)$ -subsets is red, or a subset of $[N - 1]$ of cardinality n' , each of whose $(k - 1)$ -subsets is blue. Suppose the first possibility occurs (the argument for the second case is symmetric), and let S be a qualifying subset of $[N - 1]$. Since S has $m' = R_k(m - 1, n)$ elements, there must exist either a subset of size $m - 1$ of S , each of whose k -subsets is red in the original coloring C , or a subset of size n of S , each of whose k -subsets is blue in C . In the latter case, we are done, so suppose the former case occurs, and let T be such a subset of $[N - 1]$. Let $T' = T \cup \{N\}$, and suppose X is a k -subset of T' . If $N \notin X$, then $X \subseteq S$, so X is red in C . If $N \in X$, then $X \setminus \{N\}$ is a $(k - 1)$ -subset of S and so is red in C' , and thus X is red in C . \square

Using this result, we can now establish the upper bound for the original Ramsey numbers $R_2(m, n)$ that was cited in Section 1.8.

Corollary 2.29. *Suppose m and n are integers with $\min\{m, n\} \geq 2$. Then*

$$R_2(m, n) \leq R_2(m - 1, n) + R_2(m, n - 1) \tag{2.140}$$

and

$$R_2(m, n) \leq \binom{m + n - 2}{m - 1}. \tag{2.141}$$

Proof. The inequality (2.140) follows at once from (2.136) and (2.139). The formulas (2.137) and (2.138) produce equality in (2.141) for the cases $m = 2$ and $n = 2$ respectively, and the general inequality follows by induction on $m + n$ (see Exercise 3). \square

Armed with Ramsey's theorem, we may now prove that a sufficiently large collection of points in the plane in general position must contain a subset that forms the vertices of a convex n -gon, for any positive integer n .

Theorem 2.30. *If $n \geq 3$ is an integer, then $\text{ES}(n) \leq R_4(5, n)$.*

Proof. Let S be a collection of $N = R_4(5, n)$ points in the plane in general position. For each four-element subset T of S , assign T the color red if its convex hull is a triangle, and assign it the color blue if it is a quadrilateral. By Ramsey's Theorem, there must exist either a five-element subset of S whose 4-subsets are all red, or an n -element subset of S whose 4-subsets are all blue. The former case is impossible by Theorem 2.26, so the latter case must occur, and this implies that the n points form the vertex set of a convex n -gon by Theorem 2.27. \square

Much more is known about the quantity $\text{ES}(n)$. In the same article [94], Erdős and Szekeres employ a separate geometric argument to show that in fact

$$\text{ES}(n) \leq \binom{2n-4}{n-2} + 1.$$

Since then, this bound has been improved several times. For example, in 2005 Tóth and Valtr [268] proved that

$$\text{ES}(n) \leq \binom{2n-5}{n-2} + 1$$

for $n \geq 5$.

Few exact values of $\text{ES}(n)$ have been determined. In [94], Erdős and Szekeres noted that Makai first proved that $\text{ES}(5) = 9$, so Figure 2.28 exhibits an extremal configuration. Proofs of this statement were published later in [171] and [30]. In 2006, Szekeres and Peters [266] employed a computational strategy to establish that $\text{ES}(6) = 17$. Thus, again Figure 2.28 illustrates an optimal arrangement. Erdős and Szekeres conjectured that in fact $\text{ES}(n) = 2^{n-2} + 1$ for all $n \geq 3$, and this problem remains open. This is the \$500 conjecture that Erdős was referring to in the quote that opens this section.

It is known that $\text{ES}(n)$ cannot be any smaller than the conjectured value. In 1961, Erdős and Szekeres [95] described a method for placing 2^{n-2} points in the plane in general position so that no convex n -gon appears. Their construction was later corrected by Kalbfleisch and Stanton [172]. Thus, certainly

$$\text{ES}(n) \geq 2^{n-2} + 1$$

for $n \geq 7$. For additional information on this problem and many of its generalizations, see for instance the books by Brass, Moser, and Pach [37, sec. 8.2] and

Matoušek [200, chap. 3], the survey article by Morris and Soltan [208], or the note by Dumitrescu [82].

Exercises

1. State Ramsey's theorem in terms of coloring edges of certain hypergraphs.
2. Exhibit a collection of eight points in general position in the plane whose convex hull is a triangle, so that no subset of four points forms the vertex set of a convex quadrilateral.
3. Complete the proof of Corollary 2.29.
4. (Johnson [169].) If S is a finite set of points in the plane in general position, and T is a subset of S of size 3, let $\psi_S(T)$ denote the number of points of S that lie in the interior of the triangle determined by T . Complete the following argument to establish a different upper bound on $\text{ES}(n)$.

- (a) Let $n \geq 3$ be an integer. Prove that if S is sufficiently large, then there exists a subset U of S of size n such that either every 3-subset T of U has $\psi_S(T)$ even, or every such subset has $\psi_S(T)$ odd.
- (b) If U does not form the vertex set of a convex n -gon, then by Theorem 2.27 there exist four points a, b, c , and d of U , with d lying inside the triangle determined by a, b , and c . Show that

$$\psi_S(\{a, b, c\}) = \psi_S(\{a, b, d\}) + \psi_S(\{b, c, d\}) + \psi_S(\{a, c, d\}) + 1.$$

- (c) Establish a contradiction and conclude that $\text{ES}(n) \leq R_3(n, n)$.
5. (Tarsy [188].) If a, b , and c form the vertices of a triangle in the plane, let $\theta(a, b, c) = 1$ if the path $a \rightarrow b \rightarrow c \rightarrow a$ induces a clockwise orientation of the boundary, and let $\theta(a, b, c) = -1$ if it is counterclockwise. Thus, for example, $\theta(a, b, c) = -\theta(a, c, b)$. Complete the following argument to establish an upper bound on $\text{ES}(n)$.

- (a) Let $n \geq 3$ be an integer, and let $S = \{v_1, v_2, \dots, v_N\}$ be a set of labeled points in the plane in general position. Prove that if N is sufficiently large, then there exists a subset U of S of size n such that either every 3-subset $\{v_i, v_j, v_k\}$ of U with $i < j < k$ has $\theta(v_i, v_j, v_k) = 1$, or every such subset has $\theta(v_i, v_j, v_k) = -1$.
- (b) Prove that if S contains a 4-subset whose convex hull is a triangle, then this subset must contain triangles of both orientations with respect to the ordering of the vertices.
- (c) Conclude that $\text{ES}(n) \leq R_3(n, n)$.

6. Complete the proof of Theorem 1.64 by proving that if m and n are positive integers with $\min\{m, n\} \geq 2$, and $R_2(m-1, n)$ and $R_2(m, n-1)$ are both even, then

$$R_2(m, n) \leq R_2(m-1, n) + R_2(m, n-1) - 1.$$

Use the following strategy. Let $r_1 = R_2(m-1, n)$, $r_2 = R_2(m, n-1)$, and $N = r_1 + r_2 - 1$. Suppose that the edges of K_N are 2-colored, using the colors red and blue, in such a way that no red K_m nor blue K_n appears.

- (a) Show that the red degree of any vertex in the graph must be less than r_1 .
 - (b) Show that the red degree of any vertex in the graph must equal $r_1 - 1$.
 - (c) Compute the number of red edges in the graph, and establish a contradiction.
7. Prove the following more general version of Ramsey's theorem. Let k, n_1, n_2, \dots, n_r be positive integers, with $\min\{n_1, \dots, n_r\} \geq k$, and let c_1, c_2, \dots, c_r denote r different colors. Then there exists a positive integer $R_k(n_1, \dots, n_r)$ such that in any r -coloring of the k -subsets of a set with $N \geq R_k(n_1, \dots, n_r)$ elements, there must exist a subset of n_i elements, each of whose k -subsets has color c_i , for some i with $1 \leq i \leq r$.
8. (Schur [251].) If C is an r -coloring of the elements of $[N]$, then let C' be the r -coloring of 2-subsets of $[N] \cup \{0\}$ obtained by assigning the pair $\{a, b\}$ the color of $|b - a|$ in C .
- (a) Use the generalized Ramsey's Theorem of Exercise 7 to assert that if N is sufficiently large then in $[N] \cup \{0\}$ there must exist a set of three nonnegative integers, each of whose 2-subsets has the same color in C' .
 - (b) Conclude that if N is sufficiently large then there exist integers a and b in $[N]$, with $a + b \leq N$, such that a, b , and $a + b$ all have the same color in C .
9. Let S be a finite set of points in the plane, and let P be a convex polygon whose vertices are all selected from S . We say P is *empty* (with respect to S) if its interior contains no points of S . Erdős asked if for each integer $n \geq 3$ there exists a positive integer $ES_0(n)$ such that any set of at least $ES_0(n)$ points in general position in the plane must contain an empty n -gon, but this need not be the case for sets with fewer than $ES_0(n)$ points.
- (a) Compute $ES_0(3)$ and $ES_0(4)$.
 - (b) (Ehrenfeucht [91].) Prove that $ES_0(5)$ exists by completing the following argument. Let S be a set of $ES(6)$ points in general position in the plane, and let P be a convex hexagon whose vertices lie in S , selected so that its interior contains a minimal number of points of S . Denote this number by m .

- i. Complete the proof if $m = 0$ or $m = 1$.
- ii. If $m \geq 2$, let H be the convex hull of the points of S lying inside P , and let ℓ be a line determined by two points on the boundary of H . Finish the proof for this case.

The argument above establishes that $ES_0(5) \leq 17$; in 1978 Harborth [151] showed that in fact $ES_0(5) = 10$. Horton [164] in 1983 proved the surprising result that $ES_0(n)$ does not exist for $n \geq 7$. More recently, Gerken [121] and Nicolás [215] solved the problem for $n = 6$: A sufficiently large set of points in the plane in general position must contain an empty convex hexagon. The precise value of $ES_0(6)$ remains unknown, though it must satisfy $30 \leq ES_0(6) \leq ES(9) \leq 1717$. (An example by Overmars [219] establishes the lower bound; additional information on the upper bound can be found in [182, 271].)

2.11 References

You may talk too much on the best of subjects.

— Benjamin Franklin, *Poor Richard's Almanack*

We list several additional references for the reader who wishes to embark on further study.

General References

The text by van Lint and Wilson [273] is a broad and thorough introduction to the field of combinatorics, covering many additional topics. Classical introductions to combinatorial analysis include Riordan [235] and Ryser [246], and many topics in discrete mathematics and enumerative combinatorics are developed extensively in Graham, Knuth, and Patashnik [133]. The text by Pólya, Tarjan, and Woods [227] is a set of notes from a course in enumerative and constructive combinatorics. A problems-oriented introduction to many topics in combinatorics and graph theory can be found in Lovász [191]. The book by Nijenhuis and Wilf [216] describes efficient algorithms for solving a number of problems in combinatorics and graph theory, and a constructive view of the subject is developed in Stanton and White [263]. Texts by Aigner [4, 5], Berge [24], Comtet [60], Hall [146], and Stanley [261, 262] present more advanced treatments of many aspects of combinatorics.

Combinatorial Identities

The history of binomial coefficients and Pascal's triangle is studied in Edwards [85], and some interesting patterns in the rows of Pascal's triangle are observed by Granville [138]. Combinatorial identities are studied in Riordan [236], and automated techniques for deriving and proving identities involving binomial coefficients and other quantities are developed in Petkovšek, Wilf, and Zeilberger

[222]. Combinatorial proofs for many identities are also developed in the book by Benjamin and Quinn [22].

Pigeonhole Principle

More nice applications of the pigeonhole principle, together with many other succinct proofs in combinatorics and other subjects, are described in Aigner and Ziegler [6]. An interesting card trick based in part on a special case of Theorem 2.4 is described by Mulcahy [210]. Polynomials with $\{-1, 0, 1\}$ coefficients and a root of prescribed order m at $x = 1$, as in Exercise 14 of Section 2.4, are studied by Borwein and Mossinghoff [35].

Generating Functions

More details on generating functions and their applications can be found for instance in the texts by Wilf [284] and Graham, Knuth, and Patashnik [133], and in the survey article by Stanley [260]. The problem of determining the minimal degree d_k of a polynomial with $\{0, 1\}$ coefficients that is divisible by $(x + 1)^k$, as in Exercise 5 of Section 2.6.5, is studied by Borwein and Mossinghoff [36]. Some properties of the generalized Fibonacci numbers (Exercise 8b of Section 2.6.5) are investigated by Miles [203].

Pólya's Theory of Counting

Pólya's seminal paper on enumeration in the presence of symmetry is translated into English by Read in [226]. Redfield [233] independently devised the notion of a cycle index for a group, which he termed the *group reduction formula*, ten years before Pólya's paper. As a result, many texts call this topic *Pólya-Redfield theory*. This theory, along with the generalization incorporating a color group, is also described in the expository article by de Bruijn [68], and his research article [69]. Further generalizations of this theory are explored by de Bruijn in [70], culminating in a "monster theorem." Another view of de Bruijn's theorem is developed by Harary and Palmer in [149; 150, chap. 6].

Applications of this theory in chemistry are described in the text by Fujita [116], and additional references for enumeration problems in this field are collected in the survey article [13]. Some applications of Pólya's and de Bruijn's theorems in computer graphics appear for example in articles by Banks, Linton, and Stockmeyer [15, 16].

More Numbers

The book [10] by Andrews and Eriksson is an introduction to the theory of partitions of integers, directed toward undergraduates. A more advanced treatment is developed by Andrews [9]. Euler's original proof of the pentagonal number theorem, along with some of its additional ramifications, is described by Andrews in [8].

The history of Stirling numbers, the notations developed for them, and many interesting identities they satisfy are discussed by Knuth in [177]. Rhyming schemes, as in Exercise 7d of Section 2.8.3 and Exercise 8 of Section 2.8.4, are analyzed by Riordan [237]. Stirling set numbers arise in a natural way in an interesting problem on juggling in an article by Warrington [280]. Some identities involving the complementary Bell numbers (Exercise 6 of Section 2.8.4) are established in the article by Uppuluri and Carpenter [270].

Eulerian numbers appear in the computation of the volume of certain slabs of n -dimensional cubes in articles by Chakerian and Logothetti [51] and Marichal and Mossinghoff [197], and in the solution to a problem concerning a novel graduation ceremony in an article by Gessel [122].

The reference book by Sloane and Plouffe [258] and website by Sloane [257] catalog thousands of integer sequences, many of which arise in combinatorics and graph theory, and list references to the literature for almost all of these sequences. The book by Conway and Guy [61] is an informal discussion of several kinds of numbers, including many common combinatorial sequences.

Stable Marriage

The important results of Gale and Shapley appeared in [117]. A fast algorithm that solves the “stable roommates” problem whenever a solution exists was first described by Irving in [166]. Stable matching problems are studied in Knuth [178] as motivation for the mathematical analysis of algorithms, and the structure of stable matchings in marriage and roommate problems is described in detail by Gusfield and Irving [143], along with algorithms for their computation. A matching algorithm for the “many-to-many” variation of the stable marriage problem, as in Exercise 9 of Section 2.9.2, is developed by Baïou and Balinski [14]. The monograph by Feder [103] studies extensions of the stable matching problem to more general settings.

Combinatorial Geometry

A survey on Sylvester’s problem regarding ordinary lines for collections of points, as well as related problems, appears in Borwein and Moser [34]. A variation of Sylvester’s theorem for an infinite sequence of points lying within a bounded region in the plane is investigated by Borwein [33]. The influential paper of Erdős and Szekeres on convex polygons, first published in [94], also appears in the collection by Gessel and Rota [123]. The survey by Morris and Soltan [208] summarizes work on this problem and several of its variations. Dozens of problems in combinatorial geometry, both solved and unsolved, are described in the books by Brass, Moser, and Pach [37], Hadwiger, Debrunner, and Klee [144], Herman, Kučera, and Šimša [158], and Matoušek [200], as well as the survey article by Erdős and Purdy [93].

Collected Papers

The collection [123] by Gessel and Rota contains many influential papers in combinatorics and graph theory, including the important articles by Erdős and Szekeres [94], Pólya [225], and Ramsey [232]. The two-volume set edited by Graham and Nešetřil [134, 135] is a collection of articles on the mathematics of Paul Erdős, including many contributions regarding his work in combinatorics and graph theory. The *Handbook of Combinatorics* [131, 132] provides an overview of dozens of different areas of combinatorics and graph theory for mathematicians and computer scientists.

3

Infinite Combinatorics and Graphs

... the definitive clarification of the nature of the infinite has become necessary...

— David Hilbert [160]

Infinite sets are very peculiar, and remarkably different from finite sets. This can be illustrated with a combinatorial example.

Suppose we have four pigeons and two pigeonholes. If we place the pigeons in the pigeonholes, one of the pigeonholes must contain at least two pigeons. This crowding will always occur, regardless of the arrangement we choose for the pigeons. Furthermore, the crowding will occur whenever there are more pigeons than holes. In general, if P (pigeons) is a finite set, and H (pigeonholes) is a proper subset of P , then there is no matching between the elements of P and H .

Now suppose that we have a pigeon for each real number in the closed interval $P = [0, 2]$. Put a leg tag on each pigeon with its real number. Also suppose that we have a pigeonhole for each real number in the interval $H = [0, 1]$. Put an address plate on each pigeonhole with its real number. Note that $H \subsetneq P$, so the set of address plate numbers is a proper subset of the set of leg tag numbers. For each $x \in [0, 2]$, place the pigeon tagged x in the pigeonhole with address $x/2$. Using this arrangement, no two pigeons will be assigned to the same pigeonhole. Thus, if P is infinite and H is a proper subset of P , there may be a matching between the elements of P and those of H .

Infinite sets behave differently from finite sets, and we have used ideas from graph theory and combinatorics to illustrate this difference. One of the justifications for studying infinite versions of combinatorial and graph-theoretic theorems is to gain more insight into the behavior of infinite sets and, by contrast, more




FIGURE 3.1. A matching.

insight into finite sets. Sections 3.1 and 3.2 follow this agenda, culminating in a proof of a finite combinatorial statement using infinite tools. We can also use combinatorial properties to distinguish between different sizes of infinite sets, as is done in Section 3.7. This requires the deeper understanding of the axioms for manipulating infinite sets provided by Sections 3.3 and 3.4, and a precise notion of size that appears in Section 3.5. Combinatorial and graph-theoretic properties can also illuminate the limitations of our axiom systems, as shown in Sections 3.6 and 3.9. The chapter concludes with a hint at the wealth of related topics and a list of references.

3.1 Pigeons and Trees

I wonder about the trees.

— Robert Frost, *The Sound of Trees*

The chapter introduction shows one way to use pigeons to distinguish between some finite and infinite sets. We could use this as a basis for defining finite sets, but this approach has some drawbacks that we will see in Section 3.4. It is more straightforward to say that a set is *infinite* if its not finite, and that a set is *finite* if its elements can be matched with a bounded initial segment of \mathbb{N} . For example, the set $\{A, B, C, D\}$ is finite, because the matching in Figure 3.1 exists. Note that the least integer not used in this matching is 4, which is also the size of the set $\{A, B, C, D\}$. This nifty trick, the result of using 0 in our matchings, reappears in Section 3.5.

Using the preceding notion of infinite and finite sets, we can propose another pigeonhole principle. Suppose we have an infinite number of pigeons that we stuff into a finite number of pigeonholes. Momentarily disregarding physical considerations, we must have at least one pigeonhole that contains an infinite number of pigeons. Letting P be the set of pigeons, H the set of holes, and f the stuffing function, we obtain the following theorem.

Theorem 3.1 (Infinite Pigeonhole Principle). *Suppose P is infinite, H is finite, and $f : P \rightarrow H$. Then there is an element $h \in H$ such that the pre-image set $\{p \in P \mid f(p) = h\}$ is infinite.*

Proof. Let P , H , and f be as in the hypothesis of the theorem. In particular, let $H = \{h_0, h_1, \dots, h_n\}$. Suppose, by way of contradiction, that for each $h_i \in H$,




FIGURE 3.2. A tree with labels.

the set $P_i = \{p \in P \mid f(p) = h_i\}$ has s_i elements. Because P can be written as $P = P_0 \cup P_1 \cup \dots \cup P_n$, we see that $\sum_{i \leq n} s_i$ is the size of P . Thus P is finite, providing the desired contradiction. \square

A physicist might suggest that the density of matter resulting from cramming an unbounded number of pigeons into a bounded pigeonhole would yield a fusion explosion, obliterating any evidence that could be used by litigious animal rights advocates. Home experiments with actual live pigeons are strongly discouraged. Despite the physical impracticality of our theorem, it is handy for proving a very nice theorem about trees.

As stated in Chapter 1, a tree is a connected acyclic graph. For big trees, it is handy to designate a root and label the vertices. Figure 3.2 is an example. As a convenient convention, we always think of the root r as the bottom of the tree and vertices farther from r as being higher in the tree. A *path* through a tree is a path leading up and away from the root. For example, r, a_1, b_2 and r, a_1, b_1, c_0 are paths in the tree above. The sequence r, a_1, b_0 is not a path, because $a_1 b_0$ is not an edge in the graph. If we add the edge $a_1 b_0$, the resulting graph is not a tree. (Find the cycle!)

A *level* in a tree is the collection of all vertices at a fixed distance from the root. The levels in our sample tree are $\{r\}$, $\{a_0, a_1\}$, $\{b_0, b_1, b_2\}$ and $\{c_0, c_1, c_2\}$. If v is a vertex and w is a neighboring vertex in the next higher level, then we call w an immediate successor of v . In the sample, b_1 is an immediate successor of a_1 , and b_0 is not. We can even say that c_1 is a successor of a_1 , but not an immediate successor.

The vertex labels in the sample tree are arbitrary; if we want more than 26 levels, we could use a different labeling scheme. It is even possible to reuse labels in some circumstances, as shown in Exercise 2.

Now we are ready to state König's Lemma. The result concerns infinite trees, that is, trees with an infinite number of vertices. Essentially, König's Lemma says that big skinny trees are tall.

Theorem 3.2 (König's Lemma). *If T is an infinite tree and each level of T is finite, then T contains an infinite path.*

Proof. Let T be an infinite tree in which every level is finite. Let $L_0 = \{r\}$, L_1 , L_2, \dots be the levels of T . We will construct a path as follows. Let r be the first element of the path. There are infinitely many vertices in T above r . Each of these vertices is either in L_1 or above a unique vertex in L_1 . Map each of the vertices above r to the vertex of L_1 that it is equal to or above. We have mapped infinitely many vertices (pigeons) to the finitely many vertices of L_1 (pigeonholes). By Theorem 3.1, there is at least one vertex of L_1 that is above r and has infinitely many vertices above it; pick one and call it v_1 . The path so far is r, v_1 . Since there are infinitely many vertices above v_1 , we can replace r by v_1 in the preceding argument and select v_2 . Similarly, for each $n \in \mathbb{N}$, when we have found v_n we can find v_{n+1} . Thus T contains an infinite path. \square

König's Lemma appears in the 1927 paper of Dénes König [179]. Some authors (e.g [218]) refer to the lemma as König's Infinity Theorem. The name König's Theorem is usually reserved for an unrelated result on cardinal numbers proved by Julius König, another (earlier) famous Hungarian mathematician.

Exercises

1. Suppose we arrange finitely many pigeons in infinitely many pigeon holes. Use the Infinite Pigeonhole Principle to prove that there are infinitely many pigeonholes that contain no pigeons.
2. Reusing labels in trees.
Figure 3.3 shows an example of a tree where labels are reused.




FIGURE 3.3. A tree with reused labels.

Note that in this tree, each vertex can be reached by a path corresponding to a unique sequence of labels. For example, there is exactly one vertex corresponding to $r, 0, 1$.

- (a) Give an example of a tree with badly assigned labels, resulting in two vertices that have the same sequence of labels.
- (b) Prove that if the immediate successors of each vertex in a tree have distinct labels, then no two vertices can have matching sequences of labels.
- (c) Prove the converse of part 2b.
3. 2-coloring an infinite graph.
Suppose G is a graph with vertices $V = \{v_i \mid i \in \mathbb{N}\}$ and every finite subgraph of G can be 2-colored. Use König's Lemma to prove that G is 2-colorable. (Hint: Build a tree of partial colorings. Put the vertex *root*, *red*, *blue*, *blue* in the tree if and only if assigning red to v_0 , blue to v_1 , and blue to v_2 yields a 2-coloring of the subgraph with vertices $\{v_0, v_1, v_2\}$. An infinite path through such a tree will be a coloring of G . You must prove that the tree is infinite and that each level is finite.)
4. Construct an infinite graph where each finite subgraph can be colored using a finite number of colors, but where infinitely many colors are needed to color the entire graph. (Hint: Use lots of edges.)
5. Heine–Borel Theorem on compactness of the real interval $[0, 1]$.
Use König's Lemma to prove that if $(a_0, b_0), (a_1, b_1), \dots$ are open intervals in \mathbb{R} and $[0, 1] \subset (a_0, b_0) \cup (a_1, b_1) \cup \dots$, then for some finite value n , $[0, 1] \subset (a_0, b_0) \cup (a_1, b_1) \cup \dots \cup (a_n, b_n)$. (Hint: Build a tree where the labels in the i th level are the closed intervals obtained by removing $(a_0, b_0) \cup (a_1, b_1) \cup \dots \cup (a_i, b_i)$ from $[0, 1]$ and the successors of a vertex v are labeled with subintervals of the interval for v . Use the fact that the intersection of any sequence of nested closed intervals is nonempty to show that the tree contains no infinite paths. Apply the contrapositive of König's Lemma.)

3.2 Ramsey Revisited

Ah! the singing, fatal arrow,

Like a wasp it buzzed and stung him!

— H. W. Longfellow, *The Song of Hiawatha*

Suppose that we 2-color the edges of K_6 , the complete graph with six vertices, using the colors red and blue. As we proved in Chapter 1, the colored graph must contain a red K_3 or a blue K_3 . Since we can 2-color K_5 in a way that prevents monochromatic triangles, K_6 is the smallest graph that must contain a monochromatic triangle. Thus, the Ramsey number $R(3, 3)$ is 6, as noted in Theorem 1.61. If we want to guarantee a monochromatic K_4 subgraph then we must 2-color K_{18} , because $R(4, 4) = 18$. Exact values for $R(p, p)$ when $p \geq 5$ are not known, but by

the Erdős–Szekeres bound (Theorem 1.63) we know that these Ramsey numbers exist.

Suppose that G is the complete graph with vertices $V = \{v_i \mid i \in \mathbb{N}\}$. If we 2-color the edges of G , what can we say about monochromatic complete subgraphs? Since G contains K_6 , it must contain a monochromatic K_3 . Similarly, since G contains K_{18} , it must contain a monochromatic K_4 . For $p \geq 5$, we know that $R(p, p)$ is finite and that G contains $K_{R(p,p)}$ as a subgraph, so G must contain a monochromatic K_p . So far we know that G must contain arbitrarily large finite monochromatic complete subgraphs. As a matter of fact, G contains an infinite complete monochromatic subgraph, though this requires some proof.

Theorem 3.3. *Let G be a complete infinite graph with vertices $V = \{v_i \mid i \in \mathbb{N}\}$. Given any 2-coloring of the edges, G will contain an infinite complete monochromatic subgraph.*

Proof. Suppose the edges of G are colored using red and blue. We will build an infinite subsequence $\langle w_i \mid i \in \mathbb{N} \rangle$ of V by repeatedly applying the pigeonhole principle (Theorem 3.1). Let $w_0 = v_0$. For each $i > 0$, the edge $v_0 v_i$ is either red or blue. Since this assigns v_i to one of two colors for each $i > 0$, there is an infinite set of vertices V_0 such that all the edges $\{v_0 v \mid v \in V_0\}$ are the same color. Suppose we have selected w_n and V_n . Let w_{n+1} be the lowest-numbered vertex in V_n , and let V_{n+1} be an infinite subset of V_n such that the edges in the set $\{w_{n+1} v \mid v \in V_{n+1}\}$ are the same color. This completes the construction of the sequence.

This sequence $\langle w_i \mid i \in \mathbb{N} \rangle$ has a very interesting property. If $i < j < k$, then w_j and w_k are both in V_i , and consequently $w_i w_j$ and $w_i w_k$ are the same color! We will say that a vertex w_i is blue-based if $j > i$ implies $w_i w_j$ is blue, and red-based if $j > i$ implies $w_i w_j$ is red. Each vertex in the infinite sequence $\langle w_i \mid i \in \mathbb{N} \rangle$ is blue-based or red-based, so by the pigeonhole principle there must be an infinite subsequence $\langle w_{i_0}, w_{i_1}, \dots \rangle$ where each element has the same color base. As a sample case, suppose the vertices in the subsequence are all blue-based. Then for each $j < k$, since w_{i_j} is blue-based, the edge $w_{i_j} w_{i_k}$ is blue. Thus all the edges of the complete subgraph with vertices $\{w_{i_0}, w_{i_1}, \dots\}$ are blue. If the subsequence vertices are red-based, then the edges of the associated infinite complete subgraph are red. □

Using the preceding theorem, we can prove that the finite Ramsey numbers exist without relying on the Erdős–Szekeres bound.

Theorem 3.4. *For each $n \in \mathbb{N}$ there is an $m \in \mathbb{N}$ such that $R(n, n) = m$.*

Proof. By way of contradiction, suppose that there is an n such that for every m there is a 2-coloring of the edges of K_m that contains no monochromatic K_n subgraph. Let G be the complete graph with vertices $V = \{v_i \mid i \in \mathbb{N}\}$. Suppose $E = \{e_i \mid i \in \mathbb{N}\}$ is an enumeration of the edges of G . Construct a tree T of partial edge colorings of G as follows. Include the sequence $root, c_0, c_1, c_2, \dots, c_k$ in T if and only if whenever edge e_i is colored color c_i for all $i \leq k$, the subgraph of

G containing e_0, e_1, \dots, e_k contains no monochromatic K_n . The k th level of T contains at most 2^k vertices, so each level is finite. Since we have assumed that there is a way of coloring any K_m so that no monochromatic K_n appears, T is infinite. By König's Lemma (Theorem 3.2), T has an infinite path. This infinite path provides a 2-coloring of G that contains no monochromatic K_n . Thus for this coloring, G has no infinite complete monochromatic subgraph, contradicting the preceding theorem. Our initial supposition must be false, and so for each n , there is an m such that $R(n, n) = m$. \square

We just used the infinite pigeonhole principle, infinite trees, and colorings of infinite graphs to prove a result about finite graphs! (In doing so, we are imitating Ramsey [232].) Besides being inherently fascinating, infinite constructions are very handy. Furthermore, the arguments are easily generalized. In order to take full advantage of our work, we need some new notation.

Here come the arrows! The notation $\kappa \rightarrow (\lambda)_c^2$ means that every c -colored complete graph on κ vertices contains a monochromatic complete subgraph with λ vertices. Most people pronounce $\kappa \rightarrow (\lambda)_c^2$ as "kappa arrows lambda 2 c." The statement that $R(3, 3) = 6$ combines the facts that $6 \rightarrow (3)_2^2$ (K_6 is big enough) and $5 \not\rightarrow (3)_2^2$ (K_5 is not big enough). If we imitate set theorists and write ω for the size of the set $V = \{v_i \mid i \in \mathbb{N}\}$, we can rewrite Theorem 3.3 as $\omega \rightarrow \omega_2^2$. Abbreviating "for all n " by $\forall n$ and "there exists an m " by $\exists m$, Theorem 3.4 becomes $\forall n \exists m m \rightarrow (n)_2^2$.

Arrow notation is particularly useful if we want to use lots of colors. It is easy to check that if every use of two colors is replaced by some finite value c in the proof of Theorem 3.3, the result still holds. The same can be said for Theorem 3.4. Consequently, for any $c \in \mathbb{N}$ we have

$$\omega \rightarrow (\omega)_c^2 \quad \text{and} \quad \forall n \exists m m \rightarrow (n)_c^2.$$

Note that when c is largish, the arrow notation is particularly convenient. For example, the statement " m is the least number such that $m \rightarrow (3)_9^2$ " translates into Ramsey number notation as the unwieldy formula $R(3, 3, 3, 3, 3, 3, 3, 3, 3) = m$. Nobody would want to translate $m \rightarrow (3)_{1000}^2$. On the other hand, $R(3, 4) = 9$ does not translate into our arrow notation.

The 2 in $\kappa \rightarrow (\lambda)_c^2$ indicates that we are coloring *unordered pairs* of elements taken from a set of size κ . When we edge color a graph, we are indeed assigning colors to the pairs of vertices corresponding to the edges. However, we can extend Ramsey's theorem by coloring larger subsets. The resulting statements are still very combinatorial in flavor, though they no longer refer to edge colorings. For example, the notation $\kappa \rightarrow (\lambda)_c^n$ means that for any assignment of c colors to the unordered n -tuples of κ , there is a particular color (say lime) and a subset $X \subset \kappa$ of size λ such that no matter how we select n elements from X , the corresponding n -tuple will be lime colored. The proofs of Theorems 3.3 and 3.4 can be modified to prove the following theorems.

Theorem 3.5 (Infinite Ramsey's Theorem). *For all $n \in \mathbb{N}$ and $c \in \mathbb{N}$, $\omega \rightarrow (\omega)_c^n$.*

Proof. By induction on n . Exercise 2 gives hints. \square

Theorem 3.6 (Finite Ramsey's Theorem). *For all $k, n, c \in \mathbb{N}$, there is an $m \in \mathbb{N}$ such that $m \rightarrow (k)_c^n$.*

Proof. Follows from Theorem 3.5. Exercise 3 gives hints. \square

Throughout this section we have been very picky about our infinite sets. For example, $V = \{v_i \mid i \in \mathbb{N}\}$ has a built-in matching with \mathbb{N} . What happens if we look at graphs with a vertex for each real number? In Section 3.7 we will learn that the analog of Theorem 3.3 fails for an infinite graph of this sort. For what sorts of infinite graphs does Theorem 3.3 hold? To answer this question, we need a deeper understanding of the infinite.

Exercises

1. Let $X = \{x_i \mid i \in \mathbb{N}\}$ be a set. Suppose that the relation \leq is a partial ordering on X . That is, for any $a, b, c \in X$, suppose that

- $a \leq a$,
- if $a \leq b$ and $b \leq a$, then $a = b$, and
- if $a \leq b \leq c$, then $a \leq c$.

Use Theorem 3.3 to prove that there is an infinite subset $Y \subset X$ such that either

- for every $a, b \in Y$, either $a \leq b$ or $b \leq a$, or
- for every $a, b \in Y$, both $a \not\leq b$ and $b \not\leq a$.

A subset of the first type is called a *chain*, and a subset of the second type is called an *antichain*.

2. Prove Theorem 3.5. Begin by proving Theorem 3.5 for 2 colors. Proceed by induction on n . For $n = 1$, use the pigeonhole principle as a base case. For the induction step, assume that $\omega \rightarrow (\omega)_2^n$, and prove that $\omega \rightarrow (\omega)_2^{n+1}$ by imitating the proof of Theorem 3.3, substituting applications of $\omega \rightarrow (\omega)_2^n$ for the use of the pigeonhole principle.

Given the theorem for 2 colors, there are many ways to prove it for other finite numbers of colors. You could replace 2 by c everywhere in the proof you just did, or you could try proving the theorem for c colors and n -tuples by using the theorem for 2 colors and $2n$ -tuples.

3. Prove Theorem 3.6. Imitate the proof of Theorem 3.4, using Theorem 3.5 in place of Theorem 3.3.




FIGURE 3.4. A tripartite graph representing a 2-coloring.

4. One way to visualize coloring triples.

We can represent a coloring of triples by attaching a claw to a triple that points in a particular direction. For example, the tripartite graph in Figure 3.4 represents coloring $\{0, 1, 2\}$ red and $\{1, 3, 4\}$ blue.

Figure 3.5 represents a 2-coloring of the ten triples that can be formed from the set $V = \{0, 1, 2, 3, 4\}$. You can check that every four-element subset of V contains a triple with a claw on the blue side and a triple with a claw on the red side. Thus, Figure 3.5 illustrates that $5 \not\rightarrow (4)_2^3$.




FIGURE 3.5. A 2-coloring of the triples from $\{0, 1, 2, 3, 4\}$.

- (a) Find a different coloring that shows that $5 \not\rightarrow (4)_2^3$ and represent it as a tripartite graph. (How do you know that your coloring is significantly different?)
- (b) Find a tripartite graph that shows that $5 \not\rightarrow (3)_2^2$.
- (c) Devise a way to draw a similar graph that shows that $6 \not\rightarrow (3)_3^2$.
- (d) Find a tripartite graph that shows that $6 \not\rightarrow (4)_2^3$. Since every triple gets a claw, make your life easier by drawing only the red claws.

3.3 ZFC

No one shall be able to drive us from the paradise that Cantor created for us.

— David Hilbert [160]

Paraphrasing Hilbert, in Cantor's paradise mathematicians can joyfully prove new and rich results by employing infinite sets. Since we have been living reasonably comfortably in this paradise since the beginning of the chapter, Hilbert's anxiety about eviction may seem misplaced. However, Russell and other mathematicians discovered some set-theoretic paradoxes that made the naïve use of infinite sets very questionable. Hilbert responded by calling for a careful investigation with the goal of completely clarifying the nature of the infinite.

One could argue that Hilbert's call (made in 1925) had already been answered by Zermelo in 1908. In the introduction to [293], Zermelo claimed to have reduced the entire theory created by Cantor and Dedekind to seven axioms and a few definitions. Although we now use formulations of the axioms of separation and replacement that more closely resemble those of Fraenkel and Skolem, the most commonly used axiomatization of set theory, ZFC, consists primarily of axioms proposed by Zermelo. The letters ZFC stand for **Z**ermelo, **F**raenkel, and **A**xiom of **C**hoice. Although Skolem does not get a letter, it would be hard to overestimate his influence in recasting ZFC as a first order theory.

ZFC succinctly axiomatizes what has become the de facto foundation for standard mathematical practice. With sufficient diligence, it would be possible to formalize every theorem appearing so far in this book and prove each of them from the axioms of ZFC. Since these proofs can be carried out in a less formal setting, foundational concerns are insufficient motivation for adopting an axiomatic approach. However, many of the results in Sections 3.4 through 3.10 cannot even be stated without referring to ZFC. We will use ZFC as a base theory to explore the relative strength of some very interesting statements about sets. In particular, ZFC will be central to our discussion of large cardinals and infinite combinatorics.

3.3.1 Language and Logical Axioms

The comfort of the typesetter is certainly not the summum bonum.

— Gottlob Frege [112]

Before we discuss the axioms of ZFC, we need to list the symbols we will use. Although some of these symbols may be unfamiliar, they can be used as a very convenient shorthand.

Variables can be uppercase or lowercase letters with subscripts tacked on if we please. Good examples of variables include A , B , x , and y_3 . The symbol \emptyset denotes the empty set, and \mathcal{P} and \cup are function symbols for the power set and union. The exact meaning of \emptyset , $\mathcal{P}(x)$, and $\cup x$ are determined by the axioms in the next section. ($\cup x$ is *not* a typographical error; a discussion appears later.) A

Formula	Translation
$\neg\theta$	not θ
$\theta \wedge \psi$	θ and ψ
$\theta \vee \psi$	θ or ψ
$\theta \rightarrow \psi$	if θ then ψ
$\theta \leftrightarrow \psi$	θ if and only if ψ
$\forall x\theta$	for all sets x , θ holds
$\exists x\theta$	there is a set x such that θ holds

TABLE 3.1. Translations of connectives and quantifiers.

term is a variable, the symbol \emptyset , or the result of applying a function to a term. In ZFC, terms always denote sets. Consequently, all the objects discussed in ZFC are sets. Some early formalizations of set theory include distinct objects with no elements. These objects are usually called atoms or urelements. They do not show up in ZFC.

The atomic formulas of ZFC are $x \in y$ and $x = y$, where x and y could be any terms. As one would expect, the formula $x \in y$ means x is an element of y . The connection between \in and $=$ is partly determined by the axiom of extensionality (in the next section) and partly determined by the fact that $=$ really does denote the familiar equality relation.

All other formulas of ZFC are built up by repeatedly applying logical connectives and quantifiers to the atomic formulas. Table 3.1 lists typical formulas and their translations. The letters θ and ψ denote formulas of ZFC.

Specifying that ZFC is a first order theory implicitly appends the axioms for predicate calculus with equality to the axioms for ZFC. In a nutshell, these logical axioms tell us that the connectives and quantifiers have the meanings shown in Table 3.1, and that $=$ is well behaved. In particular, we can substitute equal terms. Thus, if $x = y$ and $\theta(x)$ both hold, then $\theta(y)$ holds, too. As a consequence, we can prove the following theorem.

Theorem 3.7. *Equal sets have the same elements. Formally,*

$$x = y \rightarrow \forall t(t \in x \leftrightarrow t \in y).$$

Proof. Suppose $x = y$. Fix t . If $t \in x$, then by substitution, $t \in y$. Similarly, if $t \in y$, then $t \in x$. Our choice of t was arbitrary, so $\forall t(t \in x \leftrightarrow t \in y)$. \square

We could completely formalize the preceding argument as a symbolic logic proof in any axiom system for predicate calculus with equality. Some good formal axiom systems can be found in Mendelson [201] or Kleene [176] by readers with a frighteningly technical bent.

It is very convenient to write $x \subset y$ for the formula $\forall t(t \in x \rightarrow t \in y)$. Using this abbreviation and only the axioms of predicate calculus, we could prove that $\forall x(x \subset x)$, showing that every set is a subset of itself. We could also prove that

containment is a transitive relation, which can be formalized as

$$\forall x \forall y \forall z ((x \subset y \wedge y \subset z) \rightarrow x \subset z).$$

The preceding results (which appear in the exercises) rely on logical axioms rather than on the actual nature of sets. To prove meaty theorems, we need more axioms.

3.3.2 Proper Axioms

... I tasted the pleasures of Paradise, which produced these Hell torments...

— Pangloss, in *Candide*

The axiom system ZFC consists of nine basic axioms plus the axiom of choice. Typically, the nine basic axioms are referred to as ZF. In this section, we will examine the axioms of ZF, including their formalizations, some immediate applications, and a few random historical comments. This should be less painful than the affliction of Pangloss.

1. Axiom of extensionality: If a and b have the same elements, then $a = b$. Formally,

$$\forall x (x \in a \leftrightarrow x \in b) \rightarrow a = b.$$

This axiom is the converse of Theorem 3.7, so ZF can prove that $a = b$ if and only if a and b have exactly the same elements. Using this, we can prove the following familiar theorem about the connection between subsets and equality.

Theorem 3.8. *For all sets a and b , $a = b$ if and only if $a \subset b$ and $b \subset a$. Formally,*

$$\forall a \forall b (a = b \leftrightarrow (a \subset b \wedge b \subset a)).$$

Proof. First suppose that $a = b$. Since $a \subset a$ (see Exercise 1), by substitution we have $a \subset b$ and $b \subset a$. Thus, $a = b \rightarrow (a \subset b \wedge b \subset a)$.

To prove the converse, suppose $a \subset b$ and $b \subset a$. Since $a \subset b$, for every x we have that if $x \in a$ then $x \in b$. Similarly, since $b \subset a$, $x \in b$ implies $x \in a$. Summarizing, for all x , $x \in a \leftrightarrow x \in b$. By the axiom of extensionality, $a = b$. \square

The axiom of extensionality and the preceding theorem give us strategies for proving that sets are equal. Most proofs of set equality apply one of these two approaches.

2. Empty set axiom: \emptyset has no elements. Formally, $\forall x (x \notin \emptyset)$.

The empty set has some unusual containment properties. For example, it is a subset of every set.

Theorem 3.9. *\emptyset is a subset of every set. Formally, $\forall t (\emptyset \subset t)$.*

Proof. The proof relies on the mathematical meaning of implication. Suppose t is a set. Pick any set x . By the empty set axiom, $x \notin \emptyset$, so $x \in \emptyset$ implies $x \in t$. (When the hypothesis is false, the implication is automatically true. If I am the king of the world, then you will send me all your money. The statement is true, but no checks have arrived.) Formally, $\forall x(x \in \emptyset \rightarrow x \in t)$, so $\emptyset \subset t$. \square

The preceding proof also implies that $\emptyset \subset \emptyset$, although Exercise 1 provides a more direct proof.

3. Pairing axiom: For every x and y , the pair set $\{x, y\}$ exists. Formally,

$$\forall x \forall y \exists z \forall t (t \in z \leftrightarrow (t = x \vee t = y)).$$

In the formal version of the axiom, the set z has x and y as its only elements. Thus, z is $\{x, y\}$. The pair sets provided by the pairing axiom are unordered, so $\{x, y\} = \{y, x\}$. The pairing axiom can be used to prove the existence of single-element sets, which are often called singletons.

Theorem 3.10. *For every x , the set $\{x\}$ exists. That is, $\forall x \exists z \forall t (t \in z \leftrightarrow t = x)$.*

Proof. Fix x . Substituting x for y in the pairing axiom yields a set z such that $\forall t (t \in z \leftrightarrow (t = x \vee t = x))$. By the axiom of extensionality, $z = \{x\}$. \square

The empty set axiom, the pairing axiom, and Theorem 3.10 on the existence of singletons are all combined in Zermelo's original axiom of elementary sets [293]. As an immediate consequence he solves Exercise 4, showing that singleton sets have no proper subsets.

The statement of the next axiom uses the union symbol in an unusual way. In particular, we will write $\cup\{x, y\}$ to denote the familiar $x \cup y$. This prefix notation is very convenient for writing unions of infinite collections of sets. For example, if $X = \{x_i \mid i \in \mathbb{N}\}$, then the infinite union $x_0 \cup x_1 \cup x_2 \cup \dots$ can be written as $\cup X$, eliminating the use of pesky dots. The union axiom says that $\cup X$ contains the appropriate elements.

4. Union axiom: The elements of $\cup X$ are precisely those sets that are elements of the elements of X . Formally,

$$\forall t (t \in \cup X \leftrightarrow \exists y (t \in y \wedge y \in X)).$$

Exercise 5 is a verification that $\cup\{x, y\}$ is exactly the familiar set $x \cup y$. The notion of union extends naturally to collections of fewer than two sets, also. By the union axiom, $t \in \cup\{x\}$ if and only if there is a $y \in \{x\}$ such that $t \in y$, that is, if and only if $t \in x$. Thus, $\cup\{x\} = x$. For an exercise in wildly vacuous reasoning, try out Exercise 6, showing that $\cup \emptyset = \emptyset$.

Like the union axiom, the power set axiom defines one of our built-in functions.

5. Power set axiom: The elements of $\mathcal{P}(X)$ are precisely the subsets of X . Formally,

$$\forall t(t \in \mathcal{P}(X) \leftrightarrow t \subset X).$$

This is the same power set operator that appears in the first chapter of dozens of mathematics texts. For example,

$$\mathcal{P}(\{a, b\}) = \{\emptyset, \{a\}, \{b\}, \{a, b\}\}.$$

If X is a finite set of size n , then $\mathcal{P}(X)$ has 2^n elements. Thus for finite sets, the size of $\mathcal{P}(X)$ is always larger than the size of X . In Section 3.5 we will prove that this relation continues to hold when X is infinite.

It may seem odd that we do not have other built-in functions, like intersection, set-theoretic difference, or Cartesian products. However, all these operations can be defined using the next axiom and are omitted in order to reduce redundancy in the axioms. Our version of the separation axiom is the invention of Skolem [255]. Both Skolem and Fraenkel [109] proposed emendations to Zermelo's version of the separation axiom.

6. Separation axiom: If $\psi(x)$ is a formula and X is a set, then the set denoted by $\{x \in X \mid \psi(x)\}$ exists. More formally, given any set X and any formula $\psi(x)$ in the language of ZFC, if $\psi(x)$ does not contain the variable S , then

$$\exists S \forall x(x \in S \leftrightarrow (x \in X \wedge \psi(x))).$$

Note that $\psi(x)$ may contain unquantified variables that can be viewed as parameters. Thus S can be defined in terms of X and other given sets.

We can use the separation axiom to prove that intersections exist. It is nice to use intersection notation that is parallel to our union notation, so we write $\cap\{a, b\}$ for $a \cap b$. In general, an element should be in $\cap X$ precisely when it is in every element of X .

Theorem 3.11. *For any nonempty set X , $\cap X$ exists. That is, for any set X there is a set Y such that*

$$\forall x(x \in Y \leftrightarrow \forall t(t \in X \rightarrow x \in t)).$$

Proof. Fix X . Let $Y = \{x \in \cup X \mid \forall t(t \in X \rightarrow x \in t)\}$. By the separation axiom, Y exists. We still need to show that Y is the desired set. By the definition of Y , if $x \in Y$, then $\forall t(t \in X \rightarrow x \in t)$. Conversely, if $\forall t(t \in X \rightarrow x \in t)$, then since X is nonempty, $\exists t(t \in X \wedge x \in t)$. Thus $x \in \cup X$. Because $x \in \cup X$ and $\forall t(t \in X \rightarrow x \in t)$, we have $x \in Y$. Summarizing, $x \in Y$ if and only if $\forall t(t \in X \rightarrow x \in t)$. \square

It is also possible to show that $\cap X$ is unique. (See Exercise 8.) Since we can show that for all X the set $\cap X$ exists and is unique, we can add the function

symbol \cap to the language of ZFC. Of course, the symbol itself could be subject to misinterpretation, so we need to add a defining axiom. The formula

$$\forall x(x \in \cap X \leftrightarrow \forall t(t \in X \rightarrow x \in t))$$

will do nicely. The resulting extended theory is more convenient to use, but proves exactly the same theorems, except for theorems actually containing the symbol \cap . Mathematical logicians would say that ZFC with \cap is a *conservative extension* of ZFC.

Using the same process, we can introduce other set-theoretic functions. For example, we can specify a set that represents the ordered pair (x, y) , and define the Cartesian product $X \times Y = \{(x, y) \mid x \in X \wedge y \in Y\}$. Ordered n -tuples can be defined in a number of reasonable ways from ordered pairs. We could define the relative complement of Y in X by $X - Y = \{x \in X \mid x \notin Y\}$. See Exercises 9, 10, and 11 for more discussion of these operations.

There are some significant restrictions in the sorts of functions that could be conservatively added to ZFC. For example, as above it is acceptable to introduce the relative complement, but not a full-blown general complement. (Books usually use \overline{X} or X^c to denote a general complement.) Given a general complement, we could prove that $X \cup X^c$ existed. This would give us a set of all sets, but that is prohibited by the separation axiom.

Theorem 3.12. *There is no universal set. That is, there is no set U such that $\forall x(x \in U)$.*

Proof. Suppose by way of contradiction that $\forall x(x \in U)$. Applying the separation axiom, there is a set X such that $X = \{z \in U \mid z \notin z\}$. Note that for all z , $z \in X$ if and only if $z \in U$ and $z \notin z$. Furthermore, $z \in U$ and $z \notin z$ if and only if $z \notin z$. Thus, $z \in X$ if and only if $z \notin z$ for any z we care to choose. In particular, substituting X for z gives us $X \in X$ if and only if $X \notin X$, yielding the desired contradiction. \square

The preceding proof contains the gist of Russell's paradox. Briefly, Russell's paradox says that the existence of $\{z \mid z \notin z\}$ leads inexorably to contradictions. Note that the existence of $\{z \mid z \notin z\}$ is **not** proved by the separation axiom, because the specified set is not bounded. For any bound X , we can prove that $\{z \in X \mid z \notin z\}$ exists; it is just a harmless subset of X . By requiring bounds on definable sets, we cleverly sidestep paradoxes that ensnare the users of naïve set theory. For another experiment with Russell's style of argument, try Exercise 12.

Part of Hilbert's motivation for the rigorous study of set theory was to gain a deeper understanding of infinite sets. So far, our axioms do not guarantee the existence of a single infinite set. (Readers who love technical details may want to construct a model of axioms 1 through 6 in which every set is finite. The universe for this model will be infinite, but each element in the universe will be finite.)

One way to construct an infinite set is to start with \emptyset and successively apply Theorem 3.10. If we let $x_0 = \emptyset$ and $x_{n+1} = \{x_n\}$ for each n , this yields a set for

each natural number. In particular, $x_0 = \emptyset$, $x_1 = \{\emptyset\}$, $x_2 = \{\{\emptyset\}\}$, and so on. The next axiom affirms the existence of a set containing all these sets as elements.

7. Infinity axiom: There is a set Z such that (i) $\emptyset \in Z$ and (ii) if $x \in Z$, then $\{x\} \in Z$. Formally,

$$\exists Z(\emptyset \in Z \wedge \forall x(x \in Z \rightarrow \exists y(y \in Z \wedge \forall t(t \in y \leftrightarrow t = x))).$$

The axiom of infinity guarantees the existence of *some* set satisfying properties (i) and (ii). By applying the power set axiom, the separation axiom, and taking an intersection, we can find the smallest set with this property. For details, see Exercise 13.

Zermelo's axiomatization of set theory consists of axioms 1 through 7 plus the axiom of choice. We will discuss the axiom of choice shortly. In the meantime, there are two more axioms that have been appended to ZF that should be mentioned. The first of these is the axiom of replacement, proposed in various versions by Fraenkel ([107], [108], and [110]), Skolem [255], and Lennes [186].

8. Replacement axiom: Ranges of functions restricted to sets exist. That is, if $f(x)$ is a function and D is a set, then the set $R = \{f(x) \mid x \in D\}$ exists. More formally, if $\psi(x, y)$ is a formula of set theory such that

$$\forall x \forall y \forall z ((\psi(x, y) \wedge \psi(x, z)) \rightarrow y = z),$$

then for every set D there is a set R such that

$$\forall y(y \in R \leftrightarrow \exists x(x \in D \wedge \psi(x, y))).$$

Note that the formula $\psi(x, y)$ in the formal statement of the axiom can be viewed as defining the relation $f(x) = y$. The replacement axiom is useful for proving the existence of large sets. In particular, if we assume that ZFC and the continuum hypothesis are consistent, in the absence of the replacement axiom it is impossible to prove that any sets of size greater than or equal to \aleph_ω exist. (To find out what \aleph_ω is, you have to stick around until Section 3.5.)

The final axiom of ZF is the regularity axiom. In a nutshell, it outlaws some rather bizarre behavior, for example having $x \in y \in x$. Attempts to avoid these strange constructs can be found in the work of Mirimanoff [207], but Skolem [255] and von Neumann [276] are usually given credit for proposing the actual axiom.

9. Regularity axiom: Every nonempty set x contains an element y such that $x \cap y = \emptyset$. Formally,

$$\forall x(x \neq \emptyset \rightarrow \exists y(y \in x \wedge x \cap y = \emptyset)).$$

The idea here is that \in can be viewed as a partial ordering on any set by letting $x < y$ mean $x \in y$. The regularity axiom says that every set has a minimal element in this ordering. This rules out loops (like $x \in y \in x$) and infinite descending chains (like $\cdots \in x_3 \in x_2 \in x_1 \in x_0$). The following theorem shows that tight loops are outlawed.

Theorem 3.13. *For all x , $x \notin x$.*

Proof. By way of contradiction, suppose $x \in x$. By Theorem 3.10, we know the set $X = \{x\}$ exists. The set X is nonempty, so by the regularity axiom, there is an element $y \in X$ such that $X \cap y = \emptyset$. The only element of X is x , so $y = x$ and $X \cap x = \emptyset$. However, $x \in X$ and $x \in x$, so $x \in X \cap x = \emptyset$, a contradiction. \square

Summarizing this section, the proper axioms of ZF are:

1. Axiom of extensionality,
2. Empty set axiom,
3. Pairing axiom,
4. Union axiom,
5. Power set axiom,
6. Separation axiom,
7. Infinity axiom,
8. Replacement axiom, and
9. Regularity axiom.

We are still missing one axiom from Zermelo's list, the axiom of choice.

3.3.3 Axiom of Choice

Vizzini: *... so I can clearly not choose the wine in front of me.*

Man in black: *You've made your decision then?*

Vizzini: [happily] *Not remotely!*

— *The Princess Bride*

Suppose that we, like Vizzini, are faced with the task of selecting one glass from a set of two glasses. Since the set of glasses is nonempty, we can select one element and get on with our lives, which hopefully will be much longer than Vizzini's. To be very technical, the justification for our selection is the logical principle of existential instantiation. Similarly, using only axioms of ZF, we can always select one element from any nonempty set, without regard for the size of the set. Furthermore, we could repeat this process any finite number of times, so we can choose one element from each set in any finite list of nonempty sets.

By contrast, making an infinite number of choices simultaneously can often be problematic, depending on the circumstances. Suppose that we have an infinite collection of pairs of boots. We can pick one boot from each pair by specifying that we will select the left boot from each pair. Because each nonempty set (pair of boots) has a designated element (left boot), ZF suffices to prove the existence of the set of selected boots. Working in ZF, we cannot carry out the same process with an infinite collection of pairs of socks, because socks are not footed. We need a new axiom. In [244], Russell discusses this boot problem, though rather than selecting socks, he considers the case where “the left and right boots in each pair are indistinguishable.” Cruel shoes indeed!

The axiom of choice guarantees the existence of a set of selected socks. The following version of the axiom is very close to that of Zermelo [293].

10. Axiom of choice (AC): If T is a set whose elements are all sets that are nonempty and mutually disjoint, then $\cup T$ contains at least one subset with exactly one element in common with each element of T .

Most recent works use a formulation of the axiom of choice that asserts the existence of choice functions. In terms of socks, when a choice function is applied to a pair of socks, it outputs a designated sock. In the following statement, if T is a set of pairs of socks, t would be a pair of socks, and $f(t)$ would be a sock.

- 10'. Axiom of choice (AC2): If T is a set of nonempty sets, then there is a function f such that for every $t \in T$, $f(t) \in t$.

We use ZFC to denote ZF plus either version of AC. This is not imprecise, since we can prove that the two versions of the axiom of choice are interchangeable.

Theorem 3.14. *ZF proves that AC holds if and only if AC2 holds.*

Proof. First assume all the axioms of ZF plus AC. Let T be a set of nonempty sets. Define the function g with domain T by setting

$$g(t) = \{(t, y) \mid y \in t\}$$

for each $t \in T$. Essentially, $g(t)$ looks like the set t with a flag saying “I’m in t ” attached to each element. By the replacement axiom, the set $Y = \{g(t) \mid t \in T\}$ exists. The elements of Y are nonempty and disjoint, so by AC there is a set S that contains exactly one element from each element of Y . Thus S is a set of ordered pairs of the form (t, y) , where exactly one pair is included for each $t \in T$. Let $f(t)$ be the unique y such that $(t, y) \in S$. Then f is the desired choice function.

To prove the converse, assume ZF plus AC2. Let T be a set whose elements are nonempty and disjoint. By AC2, there is a function f such that for each $t \in T$, $f(t) \in t$. By the replacement axiom, $S = \{f(t) \mid t \in T\}$ exists. S is the desired subset of $\cup T$. \square

Zermelo ([291], [292]) used AC to prove that every set can be well-ordered. Hartogs [155] extended Zermelo's result by proving that AC is equivalent to this well-ordering principle. Hartogs' result is identical in format to the equivalence result that we just proved. What really makes Hartogs' result and our equivalence theorem interesting is the fact that AC can neither be proved nor disproved in ZF. (Technically, we just implicitly assumed that ZF is consistent. I assure you that many people make much more bizarre assumptions in their daily lives.) Gödel [125] proved that ZF cannot disprove AC, and Cohen ([58], [59]) showed that ZF cannot prove AC. Thus our equivalence theorem and the theorem of Hartogs list statements that we can add interchangeably to *strengthen* ZF. In later sections we will see more examples of equivalence theorems and more examples of statements that strengthen ZF and ZFC.

Exercises

1. Prove that set containment is reflexive. That is, prove $\forall x(x \subset x)$. (This requires only logical properties.)
2. Prove that set containment is transitive. That is, prove

$$\forall x \forall y \forall z ((x \subset y \wedge y \subset z) \rightarrow x \subset z).$$

(This requires only logical properties.)

3. Prove that the empty set is unique. That is, if $\forall x(x \notin y)$, then $y = \emptyset$.
4. Prove that if $y \subset \{x\}$, then $y = \emptyset$ or $y = \{x\}$.
5. Prove that $\cup\{x, y\}$ is exactly the familiar set $x \cup y$. That is, prove that $t \in \cup\{x, y\}$ if and only if $t \in x$ or $t \in y$.
6. Prove that $\cup\emptyset = \emptyset$.
7. Find $\mathcal{P}(\emptyset)$, $\mathcal{P}(\mathcal{P}(\emptyset))$, and $\mathcal{P}(\mathcal{P}(\mathcal{P}(\emptyset)))$. (To make your answers look really bizarre and drive your instructor nuts, write $\{ \}$ for \emptyset .)
8. Prove that $\cap X$ is unique. That is, show that if Y is a set that satisfies the formula $\forall x(x \in Y \leftrightarrow \forall t(t \in X \rightarrow x \in t))$ and Z is a set that satisfies the formula $\forall x(x \in Z \leftrightarrow \forall t(t \in X \rightarrow x \in t))$, then $Y = Z$. (Proving the existence of Y and Z requires the separation axiom (see Theorem 3.11), but this problem uses the axiom of extensionality.)
9. Let $X - Y$ denote the set $\{x \in X \mid x \notin Y\}$.
 - (a) Prove that for every X and Y , $X - Y$ exists.
 - (b) Prove that for every X and Y , $X - Y$ is unique.
 - (c) Under what circumstances does $X - Y = Y - X$?

10. Representations of ordered pairs.

- (a) Kuratowski [184] suggested that the ordered pair (a, b) can be represented by the set $\{\{a, b\}, a\}$. (This encoding is still in use.) Using this definition, prove that $(a, b) = (c, d)$ if and only if $a = c$ and $b = d$.
- (b) Using Kuratowski's encoding, show that if X and Y are sets, then the set $X \times Y$ defined by $X \times Y = \{(x, y) \mid x \in X \wedge y \in Y\}$ exists and is uniquely determined by X and Y .
- (c) Wiener [283] suggested that the ordered pair (x, y) can be represented by the set $\{\{\{x\}, \emptyset\}, \{\{y\}\}\}$. If you dare, repeat parts 10a and 10b using this encoding.
- (d) Show that encoding (a, b) by $\{a, \{b\}\}$ leads to difficulties. (Find two distinct ordered pairs that have the same representation in this encoding.)

11. Representations of n -tuples.

- (a) Usually, set theorists represent (a, b, c) by $((a, b), c)$, where pairs are represented using the Kuratowski encoding from Exercise 10. Using this representation prove the following:
 - (i) $(a, b, c) = (d, e, f) \leftrightarrow (a = d \wedge b = e \wedge c = f)$,
 - (ii) $X \times Y \times Z$ exists, and
 - (iii) $X \times Y \times Z$ is unique.
- (b) To address type-theoretic concerns, Skolem [256] suggested representing (a, b, c) by $((a, c), (b, c))$. Repeat part 11a with this encoding.
- (c) Using parts 11a and 11b as the base cases in an induction argument, extend the statements in part 11a to n -tuples for each natural number n . (If you do this, you clearly have a great love for long technical arguments. You might as well repeat the whole mess with the Wiener encoding.)
- (d) Show that encoding (a, b, c) by $\{\{a, b, c\}, \{a, b\}, \{a\}\}$ leads to difficulties. (You can find distinct triples with the same representation, or you can find an ordered pair that has the same representation as an ordered triple.)

12. Prove that for all X , $\mathcal{P}(X) \not\subset X$. (Hint: Suppose that for some X , $\mathcal{P}(X) \subset X$. Define $Y = \{t \in X \mid t \notin t\}$. Show that $Y \in X$ and shop for a contradiction.)

13. Let Z be the set provided by the infinity axiom. Let T be the set of subsets of Z that satisfy properties (i) and (ii) of the infinity axiom. Let $Z_0 = \cap T$.

- (a) Prove that T exists. (Hint: $T \subset \mathcal{P}(Z)$.)

- (b) Prove that Z_0 exists.
- (c) Prove that if X satisfies properties (i) and (ii) of the infinity axiom, then $Z_0 \subset X$.
14. Use the regularity axiom to prove that for all x and y either $x \notin y$ or $y \notin x$.

3.4 The Return of der König

And Aragorn planted the new tree in the court by the fountain.

— J. R. R. Tolkien, *The Return of the King*

It may seem that the discussion of the last section strayed from our original topics in graph theory and combinatorics. However, AC is actually a statement about infinite systems of distinct representatives (SDR). As defined in Section 1.7.2, an SDR for a family of sets T is a set that contains a distinct element from each set in T . For disjoint families, we have the following theorem.

Theorem 3.15. *ZF proves that the following are equivalent:*

1. AC.
2. *If T is a family of disjoint nonempty sets, then there is a set Y that is an SDR for T .*

Proof. First, assume ZF and AC and suppose T is a family of disjoint nonempty sets. By AC, there is a set $Y \subset \cup T$ that has exactly one element in common with each element of T . Since the elements of T are disjoint, Y is an SDR for T .

To prove the converse, suppose T is a family of disjoint nonempty sets. Let Y be an SDR for T . Then $Y \subset \cup T$, and Y has exactly one element in common with each element of T , as required by AC. \square

What if T is not disjoint? For finite families of sets, it is sufficient to know that every union of k sets has at least k elements. This is still necessary for infinite families, but no longer sufficient. Consider the family of sets $T = \{X_0, X_1, X_2, \dots\}$ defined by $X_0 = \{1, 2, 3, \dots\}$, $X_1 = \{1\}$, $X_2 = \{2\}$, and so on. The union of any k sets from T has at least k elements. As a matter of fact, any collection of k sets from T has an SDR. However, the whole of T has no SDR. To build an SDR for T , we must pick some n as a representative for X_0 . This immediately leaves us with no element to represent X_n . We are out of luck. Note that if we chuck X_0 , we can find an SDR for the remaining sets. (There are not many options for the representatives; it is hard to go wrong.) The infinite set X_0 is the source of all our problems. If we allow only finite sets in the family, then we get a nice SDR existence theorem originally proved by Marshall Hall [145].

Theorem 3.16. *Suppose $T = \{X_0, X_1, X_2, \dots\}$ is a family of finite sets. T has an SDR if and only if for every $k \in \mathbb{N}$ and every collection of k sets from T , the union of these sets contains at least k elements.*

Proof. Let $T = \{X_0, X_1, X_2, \dots\}$ and suppose that each X_i is finite. If T has an SDR, then for any collection of k sets, their representatives form a k element subset of their union.

To prove the converse, assume that for every $k \in \mathbb{N}$, the union of any k elements of T contains at least k elements. By Theorem 1.52, for each k the subfamily $\{X_0, X_1, \dots, X_k\}$ has an SDR. Let Y be the tree whose paths are of the form r, x_0, x_1, \dots, x_k , where $x_i \in X_i$ for $i \leq k$ and $\{x_0, x_1, \dots, x_k\}$ is an SDR for $\{X_0, X_1, \dots, X_k\}$. Since arbitrarily large finite subfamilies of T have SDRs, the tree Y is infinite. Furthermore, the size of the k th level of the tree Y is at most $|X_0| \cdot |X_1| \cdot \dots \cdot |X_k|$, where $|X_i|$ denotes the size of X_i . Since these sets are all finite, each level is finite. By König's Lemma, Y has an infinite path, and that path is an SDR for T . \square

In the preceding proof we made no immediately *obvious* use of AC. Here is a question: Have we actually avoided the use of AC, or did we merely disguise it? The answer is that we have used some of the strength of AC in a disguised form.

There are two very natural ways to restrict AC. Recall that AC considers a family of sets. We can either restrict the size of the sets or restrict the size of the family. If we require that each set is finite, we get the following statement.

Axiom of choice for finite sets (ACF): If T is a family of finite, nonempty, mutually disjoint sets, then $\cup T$ contains at least one subset having exactly one element in common with each element of T .

If we specify that the family can be enumerated, we get the following statement. (We say that an infinite set is countable if it can be written in the form $\{x_0, x_1, x_2, \dots\}$.)

Countable axiom of choice (CAC): If $T = \{X_0, X_1, X_2, \dots\}$ is a family of nonempty, mutually disjoint sets, then $\cup T$ contains at least one subset having exactly one element in common with each element of T .

Combining both restrictions gives us CACF, the countable axiom of choice for finite sets. The statement of CACF looks like CAC with the added hypothesis that each X_i is finite. This weak version of AC is exactly what we used in proving Theorem 3.16.

Theorem 3.17. *ZF proves that the following are equivalent:*

1. König's Lemma.
2. Theorem 3.16.
3. CACF.

Proof. The proof of Theorem 3.16 shows 1 implies 2. The proofs that 2 implies 3 and 3 implies 1 are Exercises 1 and 2. \square

The relationships between our various versions of choice are very interesting. It is easy to see that ZF proves $AC \rightarrow CAC$, $AC \rightarrow ACF$, $CAC \rightarrow CACF$, and $ACF \rightarrow CACF$. It is not at all obvious, but can be shown, that not a single one of the converses of these implications is provable in ZF, and also that CACF is not a theorem of ZF. To prove that ZF cannot prove these statements we would assume that ZF is consistent and build special models where each particular statement fails. The models are obtained by lifting results from permutation models or by forcing. Jech's book *The Axiom of Choice* [167] is an excellent reference.

Since ZF proves that König's Lemma is equivalent to CACF and CACF is not a theorem of ZF, we know that König's Lemma is not a theorem of ZF. Of course, ZFC can prove König's Lemma, so it is perfectly reasonable to think of it as a theorem of mathematics. Also, ZF can prove some restrictions of König's Lemma, for example if all the labels in the tree are natural numbers. Many applications of König's Lemma can be carried out with a restricted version.

Our proof closely ties König's Lemma to countable families of sets. As we will see in the next section, not all families are countable. We will see bigger sets where the lemma fails, and still bigger sets where it holds again. This is not the last return of König. (König means "king" in German.)

In the introduction to this chapter we noted that if P is finite, then whenever H is a proper subset of P there is no matching between P and H . A set X is called *Dedekind finite* if no proper subset of X can be matched with X . Thus, the introduction shows that if X is finite, then X is Dedekind finite. Exercise 4 shows that CAC implies the converse. Thus, in ZFC, the finite sets are exactly the Dedekind finite sets. This characterization of the finite sets requires use of a statement that is weaker than CAC, but not provable in ZF [167].

Exercises

1. Prove in ZF that Theorem 3.16 implies CACF. (Hint: Use disjointness to show that the union of any k sets contains at least k elements.)
2. Challenging exercise. Prove König's Lemma using ZF and CACF. To do this, let S be the set of nodes in the tree that have infinitely many successors. Find an enumeration for S . (It is easy to slip up and use full AC when finding the enumeration.) For each $s \in S$, let X_s be the set of immediate successors of s that have infinitely many successors. Apply CACF to the family $\{X_s \mid s \in S\}$. Use the selected vertices to construct a path through the tree.
3. Disjointification trick. Suppose that $\{X_n \mid n \in \mathbb{N}\}$ is a family of sets. For each $n \in \mathbb{N}$ let $\bar{X}_n = \{(n, x) \mid x \in X_n\}$. Show that $\{\bar{X}_n \mid n \in \mathbb{N}\}$ exists and is a disjoint family of sets.
4. Use CAC to prove that every infinite set has a countable subset. (Hint: Suppose that W is infinite. For each $k \in \mathbb{N}$ let W_k be the set of all subsets of W of size k . Apply CAC to a disjointified version of the family $\{W_k \mid k \in \mathbb{N}\}$. Show that the union of the selected elements is a countable subset of W .)

5. Assume that every infinite set has a countable subset. Prove that if X cannot be matched with any proper subset of itself, then X is finite. (Hint: Suppose X is infinite and use a countable subset of X to find a matching between X and a proper subset of X .)

3.5 Ordinals, Cardinals, and Many Pigeons

Whenever Gutei Oshō was asked about Zen, he simply raised his finger. Once a visitor asked Gutei's boy attendant, "What does your master teach?" The boy too raised his finger. Hearing of this, Gutei cut off the boy's finger with a knife. The boy, screaming with pain, began to run away. Gutei called to him, and when he turned around, Gutei raised his finger. The boy suddenly became enlightened.

— Mumon Ekai, *The Gateless Gate*

The previous section contains some references to infinite sets of different sizes. To make sense of this we need to know what it means for sets to be the same size. We can illustrate two approaches by considering some familiar sets. Thanks to the gentleness of my religious training, I have the same number of fingers on my left and right hands. This can be verified in two ways. I can count the fingers on my left hand, count the fingers on my right hand, and check that the results match. Note that in the process of counting, I am matching fingers with elements of a canonical ordered set, probably $\{1, 2, 3, 4, 5\}$. By emphasizing the matching process, I can verify the equinumerousness of my fingers without using any canonical set middleman. To do this, I match left thumb with right thumb, left forefinger with right forefinger, and so on. When my pinkies match, I know that I have the same number of fingers on my left and right hands. One advantage of this technique is that it works without modification for people with six or more fingers on each hand.

For infinite sets, either method works well. We will start by comparing sets directly, then study some canonical ordered sets, and finish the section off with some applications to pigeons and trees.

3.5.1 Cardinality

The big one!

— Connie Conehead

Suppose we write $X \lesssim Y$ if there is a one-to-one function from X into Y , and $X \sim Y$ if there is a one-to-one function from X onto Y . Thus $X \sim Y$ means that there is a matching between X and Y . If $X \lesssim Y$ and $X \approx Y$, so X can be embedded into but not onto Y , we will write $X \prec Y$. With this notation, we can describe relative sizes of some infinite sets.

First consider the sets $\mathbb{N} = \{0, 1, 2, \dots\}$ and $\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$. Define the function $f : \mathbb{N} \rightarrow \mathbb{Z}$ by

$$f(n) = (-1)^{n+1} \left(\frac{2n+1}{4} \right) + \frac{1}{4}.$$

It is not hard to verify that if $f(j) = f(k)$, then $j = k$, proving that f is a one-to-one function. Additionally, if $m > 0$, then $f(2m-1) = m$, and if $t \leq 0$, then $f(-2t) = t$, so f maps the odd natural numbers onto the positive integers and the even natural numbers onto the negative integers and 0. Thus, f witnesses that $\mathbb{N} \sim \mathbb{Z}$, and we now know that \mathbb{N} and \mathbb{Z} are the same size.

If X is a set satisfying $\mathbb{N} \sim X$, then we say that X is *countable* (or *countably infinite* if we are being very precise.) We just prove that \mathbb{Z} is countable. Not every infinite set is countable, as shown by the following theorem of Cantor.

Theorem 3.18 (Cantor's Theorem). *For any set X , $X \prec \mathcal{P}(X)$. In particular, $\mathbb{N} \prec \mathcal{P}(\mathbb{N})$.*

Proof. Define $f : X \rightarrow \mathcal{P}(X)$ by setting $f(t) = \{t\}$ for each $t \in X$. Since f is one-to-one, it witnesses that $X \lesssim \mathcal{P}(X)$. It remains to show that $X \approx \mathcal{P}(X)$. Suppose $g : X \rightarrow \mathcal{P}(X)$ is any one-to-one function. We will show that g is not onto. Let $y = \{t \in X \mid t \notin g(t)\}$. Suppose by way of contradiction that for some $x \in X$, $g(x) = y$. Because $g(x) = y$, $x \in g(x)$ if and only if $x \in y$, and by the definition of y , $x \in y$ if and only if $x \notin g(x)$. Concatenating, we get $x \in g(x)$ if and only if $x \notin g(x)$, a clear contradiction. Thus y is not in the range of g , completing the proof. \square

One consequence of Cantor's Theorem is that any function from $\mathcal{P}(\mathbb{N})$ into \mathbb{N} must not be one-to-one. More combinatorially stated, if we try to ram a pigeon for each element of $\mathcal{P}(\mathbb{N})$ into pigeonholes corresponding to the elements of \mathbb{N} , some pigeonhole must contain at least two pigeons. Another consequence is that by sequentially applying Cantor's Theorem to an infinite set, we get lots of infinite sets, including some very big ones.

Corollary 3.19. *There are infinitely many infinite sets of different sizes.*

Proof. \mathbb{N} is infinite, and $\mathbb{N} \prec \mathcal{P}(\mathbb{N}) \prec \mathcal{P}(\mathcal{P}(\mathbb{N})) \cdots$ by Cantor's Theorem. \square

Using only the definition of \sim and chasing some functions, we can prove that \sim is an equivalence relation. (See Exercise 2.) In particular, for any sets A , B , and C , we have

- $A \sim A$,
- if $A \sim B$ then $B \sim A$, and
- if $A \sim B$ and $B \sim C$, then $A \sim C$.

These are handy shortcuts, and it would be nice to have analogous statements for the \lesssim relation. We can easily show that $A \lesssim A$ and that if $A \lesssim B$ and $B \lesssim C$, then $A \lesssim C$. (See Exercise 3.) Symmetry does not hold for the \lesssim relation, but we can prove that if $A \lesssim B$ and $B \lesssim A$ then $A \sim B$. This last statement is the Cantor–Bernstein Theorem, an incredibly handy shortcut for showing that sets are the same size. After discussing the proof and history of the theorem, we will look at some nice applications.

Theorem 3.20 (Cantor–Bernstein Theorem). *If both $X \lesssim Y$ and $Y \lesssim X$, then $X \sim Y$.*

Proof. Suppose $f : A \rightarrow B$ and $g : B \rightarrow A$ are one-to-one functions. We will sketch the construction of a function $h : A \rightarrow B$ that is one-to-one and onto. Define a set of subsets of A as follows. Let $A_0 = A$, $A_1 = g(B)$, and $A_n = g(f(A_{n-2}))$ for $n \geq 2$. In particular, writing $g \circ f(A)$ for $g(f(A))$, we have

$$\begin{aligned} A_0 &= A, \\ A_1 &= g(B), \\ A_2 &= g \circ f(A), \\ A_3 &= g \circ f \circ g(B), \\ A_4 &= g \circ f \circ g \circ f(A), \text{ and} \\ A_5 &= g \circ f \circ g \circ f \circ g(B). \end{aligned}$$

Note that A_n is defined with n function applications. It goes “back and forth” n times. Using induction as described in Exercise 4a, it is fairly easy to prove the following claim.

Claim 1: For all n , $A_n \supset A_{n+1}$.

Given Claim 1, define the sets $A'_n = A_n - A_{n+1}$ for each n . Also define the set $A'_\omega = \bigcap_{n \in \mathbb{N}} A_n$. These sets form a partition of A into disjoint pieces, as claimed below. Hints for the proof of this claim appear in Exercise 4b.

Claim 2: For every $x \in A$ there is a unique $n \in \{\omega, 0, 1, 2, \dots\}$ such that $x \in A'_n$.

Define the function $h : A \rightarrow B$ by the following formula:

$$h(x) = \begin{cases} f(x) & \text{if } x \in \cup\{A'_\omega, A'_0, A'_2, \dots\}, \\ g^{-1}(x) & \text{if } x \in \cup\{A'_1, A'_3, A'_5, \dots\}. \end{cases}$$

By Claim 2, $h(x)$ is well-defined and has all of A as its domain. It remains to show that $h(x)$ is one-to-one and onto. This can be accomplished by defining $B_0 = B$, and $B_{n+1} = f(A_n)$ for each $n \geq 0$. Imitating our work with the A_n , in Exercise 4c we prove the following.

Claim 3: For all n , $B_n \supseteq B_{n+1}$. Furthermore, if we define the prime sets $B'_n = B_n - B_{n+1}$ and $B'_\omega = \bigcap_{n \in \mathbb{N}} B_n$, then for every $y \in B$, there is a unique $n \in \{\omega, 0, 1, 2, \dots\}$ such that $y \in B'_n$.

The partitions of A and B are closely related. In particular, Exercise 4d gives hints for proving the following claim.

Claim 4: For each $n \in \mathbb{N}$, $h(A'_{2n}) = B'_{2n+1}$ and $h(A'_{2n+1}) = B'_{2n}$. Also, $h(A'_\omega) = B'_\omega$.

Since h matches the A' pieces with the B' pieces and is one-to-one and onto on these pieces, h is the desired one-to-one and onto function. \square

One indication of the importance of the Cantor–Bernstein Theorem is the number of mathematicians who have produced proofs of it. The following is a partial listing. According to Levy [187], Dedekind proved the theorem in 1887. Writing in 1895, Cantor [47] described the theorem as an easy consequence of a version of the axiom of choice. In the endnotes of [49], Jourdain refers to an 1896 proof by Schröder. Some texts, [209] for example, refer to the theorem as the Schröder–Bernstein Theorem. Bernstein proved the theorem without using the axiom of choice in 1898; this proof appears in Borel’s book [32]. Additional later proofs were published by Peano [220], J. König [181], and Zermelo [292]. It is good to remember that the axioms for set theory were in flux during this period. These mathematicians were making sure that this very applicable theorem was supported by the axioms *du jour*.

Now we will examine a pair of applications of the Cantor–Bernstein Theorem. Note that we are freed of the tedium of constructing onto maps.

Corollary 3.21. $\mathbb{N} \sim \mathbb{Z}$.

Proof. Define $f : \mathbb{N} \rightarrow \mathbb{Z}$ by $f(n) = n$. Note that f is one-to-one, so $\mathbb{N} \lesssim \mathbb{Z}$. Define $g : \mathbb{Z} \rightarrow \mathbb{N}$ by

$$g(z) = \begin{cases} 2z & \text{if } z \geq 0, \\ 2|z| + 1 & \text{if } z < 0. \end{cases}$$

Note that g is one-to-one, so $\mathbb{Z} \lesssim \mathbb{N}$. Applying the Cantor–Bernstein Theorem, $\mathbb{N} \sim \mathbb{Z}$. \square

Corollary 3.22. $\mathbb{N} \sim \mathbb{N} \times \mathbb{N}$.

Proof. The function $f : \mathbb{N} \rightarrow \mathbb{N} \times \mathbb{N}$ defined by $f(n) = (0, n)$ is one-to-one, so $\mathbb{N} \lesssim \mathbb{N} \times \mathbb{N}$. The function $g : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ defined by $g(m, n) = 2^{m+1} \cdot 3^{n+1}$ is also one-to-one, so $\mathbb{N} \times \mathbb{N} \lesssim \mathbb{N}$. By the Cantor–Bernstein Theorem, $\mathbb{N} \sim \mathbb{N} \times \mathbb{N}$. \square

Corollary 3.23. $\mathcal{P}(\mathbb{N}) \sim \mathbb{R}$. Consequently, \mathbb{R} is uncountable.

Proof. First we will prove that $\mathcal{P}(\mathbb{N}) \lesssim \mathbb{R}$. Define $f : \mathcal{P}(\mathbb{N}) \rightarrow \mathbb{R}$ by setting $f(X) = \sum_{n \in X} 10^{-n}$ for each $X \in \mathcal{P}(\mathbb{N})$. As an example of how this map works, if X consists of the odd natural numbers, then $f(X) = 0.10\overline{10}$. If X and Y are distinct subsets of \mathbb{N} , then they differ at some least natural number n , and $f(X)$ and $f(Y)$ will differ in the n th decimal place. Thus, f is one-to-one, and so $\mathcal{P}(\mathbb{N}) \lesssim \mathbb{R}$.

Now we must construct a one-to-one function $g : \mathbb{R} \rightarrow \mathcal{P}(\mathbb{N})$. Let r be a real number. If we avoid decimal expansions that terminate in an infinite sequence of 9's, we can assume that r has a unique decimal expansion of the form

$$(-1)^\epsilon \left(\sum_{i \in X_1} k_i 10^i + \sum_{j \in X_2} d_j 10^{-j} \right),$$

where $\epsilon \in \{0, 1\}$, each k_i and d_j is between 1 and 9, X_1 is a set of natural numbers, and X_2 is a set of nonzero natural numbers. In this representation, $(-1)^\epsilon$ is the sign of r , $\sum_{i \in X_1} k_i 10^i$ is the integer portion of r , and $\sum_{j \in X_2} d_j 10^{-j}$ is the fractional portion of r . Define the function g by setting

$$g(r) = \{\epsilon\} \cup \{10^{2i+1} + k_i \mid i \in X_1\} \cup \{10^{2j} + d_j \mid j \in X_2\}$$

for each $r \in \mathbb{R}$. As a concrete example of the behavior of this map, consider $g(-12.305) = \{1, 1001, 12, 103, 1000005\}$. Since different reals differ in some decimal place, g is one-to-one. By the Cantor–Bernstein Theorem, $\mathcal{P}(\mathbb{N}) \sim \mathbb{R}$.

By Theorem 3.18, $\mathbb{N} \prec \mathcal{P}(\mathbb{N})$. Together with $\mathcal{P}(\mathbb{N}) \sim \mathbb{R}$, this implies that $\mathbb{N} \prec \mathbb{R}$, so \mathbb{R} is uncountable. □

Note that we did not construct a one-to-one function from $\mathcal{P}(\mathbb{N})$ onto \mathbb{R} . The Cantor–Bernstein Theorem tells us that such a function must exist, so we are not obligated to construct it. (If you are not already convinced that existence theorems are tremendously convenient, try doing a direct construction for the preceding corollary. This is intentionally *not* listed in the exercises.)

3.5.2 Ordinals and Cardinals

The aleph was heavy, like trying to carry a small engine block.
— William Gibson, *Mona Lisa Overdrive*

For Gibson, an aleph is a huge biochip of virtually infinite storage capacity. For a linguist, aleph is א, the first letter of the Hebrew alphabet. For a set theorist, an aleph is a cardinal number. Saying that there are א₀ natural numbers is like saying that there are five fingers on my right hand. Alephs are special sorts of ordinals, and ordinals are special sorts of well-ordered sets.

Suppose that X is set and \leq is an ordering relation on X . We write $x < y$ when $x \leq y$ and $x \neq y$. The relation \leq is a linear ordering on X if the following properties hold for all x, y , and z in X .

Antisymmetry: $(x \leq y \wedge y \leq x) \rightarrow x = y$.

Transitivity: $x \leq y \rightarrow (y \leq z \rightarrow x \leq z)$.

Trichotomy: $x < y \vee x = y \vee y < x$.

Familiar examples of linear orderings include \mathbb{N} , \mathbb{Z} , \mathbb{Q} , and \mathbb{R} with the typical orderings. We say that a linear ordering is a *well-ordering* if every nonempty subset has a least element. Since every subset of \mathbb{N} has a least element, \mathbb{N} is a well-ordering (using the usual ordering). Since the open interval $(0, 1)$ has no least element, the usual ordering does not well-order \mathbb{R} . An analyst would say that 0 is the greatest lower bound of $(0, 1)$, but 0 is not the least element of $(0, 1)$ because $0 \notin (0, 1)$. The following theorem gives a handy characterization of well-ordered sets.

Theorem 3.24. *Suppose X with \leq is a linearly ordered set. X is well-ordered if and only if X contains no infinite descending sequences.*

Proof. We will prove the contrapositive version, that is, X is not well-ordered if and only if X contains an infinite descending sequence.

First suppose X is not well-ordered. Then X has a nonempty subset Y with no least element. Pick an element x_0 in Y . Since x_0 is not the least element of Y , there is an element x_1 in Y such that $x_0 > x_1$. Continuing in this fashion, we obtain $x_0 > x_1 > x_2 > \cdots$, an infinite descending sequence.

To prove the converse, suppose X contains $x_0 > x_1 > x_2 > \cdots$, an infinite descending sequence. Then the set $Y = \{x_i \mid i \in \mathbb{N}\}$ is a nonempty subset of X with no least element. Thus, X is not well-ordered. \square

For any set X , the \in relation defines an ordering on X . To see this, for each $x, y \in X$, let $x \leq_{\in} y$ if $x \in y$ or $x = y$. In general, this is not a particularly pretty ordering. For example, if $a \neq b$ then the set $X = \{a, b, \{a, b\}\}$ is not linearly ordered by the \leq_{\in} relation. On the other hand, $Y = \{a, \{a, b\}, \{a, \{a, b\}\}\}$ is well-ordered by the \leq_{\in} relation. In a moment, we will use this property as part of the definition of an ordinal number.

A set X is *transitive* if for all $y \in X$, if $x \in y$ then $x \in X$. A transitive set that is well-ordered by \leq_{\in} is called an *ordinal*. The ordinals have some interesting properties.

Theorem 3.25. *Suppose X is a set of ordinals and α and β are ordinals. Then the following hold:*

1. $\cup X$ is an ordinal.
2. $\alpha \cup \{\alpha\}$ is an ordinal.
3. $\alpha \in \beta$ or $\alpha = \beta$ or $\beta \in \alpha$.

Proof. See Exercises 10, 13, 14, 15, and 16. \square

The first two properties in the preceding theorem give good ways to build new ordinals from old ones. For example, a little vacuous reasoning shows that \emptyset is an ordinal. By the theorem, the sets $\emptyset \cup \{\emptyset\} = \{\emptyset\}$, $\{\emptyset\} \cup \{\{\emptyset\}\} = \{\emptyset, \{\emptyset\}\}$, and $\{\emptyset, \{\emptyset\}\} \cup \{\{\emptyset, \{\emptyset\}\}\} = \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\}$ are all ordinals. Set theorists have special names for these finite ordinals. They write

$$\begin{aligned} \emptyset &= 0, \\ \{\emptyset\} &= \{0\} = 1, \\ \{\emptyset, \{\emptyset\}\} &= \{0, 1\} = 2, \\ \{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\} &= \{0, 1, 2\} = 3, \end{aligned}$$

and so on for all $k \in \mathbb{N}$. Since each k is a set, we can define ω by $\omega = \cup_{k \in \mathbb{N}} k$, and use the first property in the theorem to see that ω is an ordinal. We do not have to stop here. Since ω is an ordinal, so is $\omega \cup \{\omega\} = \{\omega, 0, 1, 2, \dots\}$, and we start all over. Sometimes, texts write $\alpha + 1$ for the ordinal $\alpha \cup \{\alpha\}$ and call ordinals like 1, 2, 3, and $\omega + 1$ successor ordinals. Ordinals that are not successors are called *limit ordinals*. The set ω is a good example of a limit ordinal.

Traditionally, greek letters are used to denote ordinals. Also, we usually write $\alpha \leq \beta$ rather than $\alpha \leq_{\in} \beta$. Consequently, for ordinals the formula $\alpha < \beta$ means the same thing as $\alpha \in \beta$. Because ordinals are transitive, $\alpha \in \beta$ implies $\alpha \subset \beta$, although the converse is not always true.

There are three ways to think about ordinals and well-orderings. First, every ordinal is a well-ordered set under the \leq relation. Second, the class of all ordinals is well-ordered by the \leq relation. Third, every well-ordered set looks just like an ordinal. The next theorem is a precise expression of the way that ordinals act as canonical well-orderings.

Theorem 3.26. *Every nonempty well-ordered set is order isomorphic to an ordinal. That is, if X is well-ordered by \leq , then there is an ordinal α and a function $h : X \rightarrow \alpha$ such that h is one-to-one and onto, and for all x and y in X , $x \leq y$ implies $h(x) \leq h(y)$.*

Proof. Let X be a well-ordered set. For each $x \in X$, define the initial segment for x by $I_x = \{t \in X \mid t \leq x\}$. Let W be the subset of X consisting of all elements x such that I_x is order isomorphic to an ordinal. Note that for each $x \in W$, I_x is order isomorphic to a unique ordinal. By the replacement axiom, we can construct a set A of all the ordinals isomorphic to initial segments of X . Let $\alpha = \cup A$; by Theorem 3.25, α is an ordinal. If $x, y \in W$, $x \leq y$, γ and δ are ordinals, and h_x and h_y are order isomorphisms such that $h_x : I_x \rightarrow \gamma$ and $h_y : I_y \rightarrow \delta$, then for all $t < x$, $h_y(t) = h_x(t)$. Using the replacement axiom, we can form the set of all the order isomorphisms corresponding to the initial segments, and concatenate them to build a new function h . This new function is an order isomorphism from W onto α . To complete the proof, we claim that $W = X$. Suppose not; since X is well-ordered, we can find a least t in X such that $t \notin W$. If we extend h by setting $h(t) = \alpha$, then h witnesses that I_t is order isomorphic to $\alpha + 1$. Thus $t \in W$, yielding a contradiction and completing the proof. \square

The next theorem shows that using AC we can well-order any set, widening the applicability of the preceding theorem. Our proof of the “well-ordering principle” uses ideas from Zermelo’s [291] original proof. The proof can also be viewed as a special case of Zorn’s Lemma. See Exercise 18 for more about Zorn’s Lemma.

Theorem 3.27. *Every set can be well-ordered.*

Proof. Let X be a set. We will construct a one-to-one map $h : \alpha \rightarrow X$ from an ordinal α onto X . This suffices to prove the theorem, since the elements of α are well-ordered and h matches elements of α with elements of X .

By AC we can pick $x \in X - t$ for each nonempty $t \subset X$. There are two things to note here. First, x_t is never an element of t . This is important later in the proof. Second, this is the only use of AC in this entire proof. This is handy for the exercises.

Suppose that $f : \alpha \rightarrow Y$ is a one-to-one map of an ordinal α onto a set $Y \subset X$. For each $\beta < \alpha$, let $f[\beta]$ denote $\{f(\delta) \mid \delta \in \beta\}$. (Remember, since β and α are ordinals, $\beta < \alpha$ is the same thing as $\beta \in \alpha$.) We say that f is a γ -function if $f(\beta) = x_{f[\beta]}$ for every $\beta \in \alpha$. Let Γ be the set of all γ -functions. The γ -functions cohere nicely; if f and g are γ -functions and β is in both of their domains, then $f(\beta) = g(\beta)$. (See Exercise 17.) If we view the functions in Γ as sets of ordered pairs, $\cup\Gamma$ is also a set of ordered pairs. Since the functions cohere and are one-to-one, $\cup\Gamma$ is actually a one-to-one function; call it h . By Theorem 3.25, the union of the ordinals that are domains of the functions in Γ is also an ordinal, so for some ordinal α , $h : \alpha \rightarrow X$. Furthermore, h is a γ -function.

It gets better. Suppose that h does not map α onto X , so $h[\alpha] \subsetneq X$. Then we can define an extension h' by setting $h'(\beta) = h(\beta)$ for $\beta < \alpha$ and $h'(\alpha) = x_{h[\alpha]}$. This extension h' is also a γ -function, so $h' \in \Gamma$. Applying the definition of h , we find that $h'(\alpha)$ is in the range of h . But $h'(\alpha) = x_{h[\alpha]}$ and the range of h is $h[\alpha]$, so we have $x_{h[\alpha]} \in h[\alpha]$, contradicting the statement two paragraphs back. Summarizing, h is a one-to-one map of α onto X , so X is well-ordered. \square

Combining the last two theorems yields the following corollary.

Corollary 3.28. *For every X there is a unique least ordinal α such that $X \sim \alpha$.*

Proof. Fix X . By Theorem 3.27, X can be well-ordered. By Theorem 3.26, $X \sim \beta$ for some ordinal β . Let $A = \{\gamma \leq \beta \mid \gamma \sim X\}$ be the set of ordinals less than or equal to β that are equinumerous with X . Since A is a nonempty set of ordinals well-ordered by \leq , A contains a least element, α . Let δ be any ordinal such that $X \sim \delta$. By Theorem 3.25, $\delta < \alpha$ or $\alpha \leq \delta$. If $\delta < \alpha$, then $\delta \leq \beta$ and we have $\delta \in A$, contradicting the minimality of α . Thus $\alpha \leq \delta$, and α is unique. \square

Since every set has a unique least equinumerous ordinal, we can define $|X|$ as the least ordinal α such that $X \sim \alpha$. We say that an ordinal κ is a *cardinal number* if $|\kappa| = \kappa$. In slogan form, a cardinal number is the least ordinal of its cardinality. The finite pigeonhole principle asserts that every finite ordinal is a cardinal. Thus, 0, 1, and 17324 are all cardinals. The infinite pigeonhole principle shows that ω

cannot be mapped one-to-one into any finite cardinal, so ω is a cardinal number; indeed, it is the least infinite cardinal. On the other hand, $\omega+1 \sim \omega$ and $\omega+2 \sim \omega$, so $\omega+1$ and $\omega+2$ are not cardinals. The elements of the next larger cardinal cannot be matched with the elements of ω , so the next larger cardinal is uncountable.

Even though every cardinal number is an ordinal, we have special notation to distinguish the cardinals. When we are thinking of ω as a cardinal, we denote it with an aleph, so $\omega = \aleph_0$. The next larger (and consequently uncountable) cardinal is \aleph_1 . Proceeding in this way, and using unions at limit ordinals, we can define \aleph_α for every ordinal number α . For example, the least cardinal bigger than $\aleph_0, \aleph_1, \aleph_2, \dots$ is \aleph_ω . Assuming AC, for every infinite set X , there is an ordinal α such that $|X| = \aleph_\alpha$.

The ordinals are like a long string of beads. The infinite cardinals, which are the alephs, appear like infrequent pearls along the string. The ordinals are good for counting steps in order (like rosary beads), and the cardinals are ideal for summing up sizes (like abacus beads). For finite sets, cardinals and ordinals are identical. Thus $|\{A, B, C, D\}| = 4 = \{0, 1, 2, 3\}$ and $\{A, B, C, D\} \sim \{0, 1, 2, 3\}$. In general, the matching approach to measuring the sizes of sets agrees with the cardinality approach. This is formalized in the following theorem.

Theorem 3.29. *For all sets X and Y , $|X| = |Y|$ if and only if $X \sim Y$.*

Proof. Suppose $|X| = |Y| = \kappa$. Then $X \sim \kappa$ and $Y \sim \kappa$, so $X \sim Y$. Conversely, suppose $X \sim Y$, and let $\kappa_1 = |X|$ and $\kappa_2 = |Y|$. Since $\kappa_1 \sim X \sim Y \sim \kappa_2$, we have $\kappa_1 \sim \kappa_2$. Since κ_1 and κ_2 are cardinals, $\kappa_1 \sim \kappa_2$ implies $\kappa_1 = \kappa_2$. Thus $|X| = |Y|$. □

3.5.3 Pigeons Finished Off

*Every Sunday you'll see
My sweetheart and me,
As we poison the pigeons in the park.*

— Tom Lehrer

At this point, we know quite a bit about stuffing pigeons into pigeonholes. For example, if p and h are both finite cardinal numbers, we have the following finite pigeonhole principle.

- If we put p pigeons into h pigeonholes and $h < p$, then some pigeonhole contains at least two pigeons.

The idea here is that any function from the set of larger cardinality into the set of smaller cardinality must fail to be one-to-one. By the Cantor–Bernstein Theorem, this holds for infinite cardinals as well. Thus for any cardinals κ and λ , we get the following analogue of the finite pigeonhole principle.

- If we put κ pigeons into λ pigeonholes and $\lambda < \kappa$, then some pigeonhole contains at least two pigeons.

The preceding infinite analogue of the finite pigeonhole principle is not the same as the infinite pigeonhole principle of Theorem 3.1. Here is a restatement of Theorem 3.1 using our notation for cardinals.

- If we put \aleph_0 pigeons into h pigeonholes and $h < \aleph_0$, then some pigeonhole contains \aleph_0 pigeons.

The infinite pigeonhole principle says that some pigeonhole is infinitely crowded. This does not transfer directly to higher cardinalities. For example, we can put \aleph_ω pigeons into \aleph_0 pigeonholes in such a way that every pigeonhole has fewer than \aleph_ω pigeons in it. To do this, put \aleph_0 pigeons in the 0th hole, \aleph_1 pigeons in the 1st hole, \aleph_2 pigeons in the 2nd hole, and so on. The total number of pigeons is $\aleph_\omega = \bigcup_{n \in \omega} \aleph_n$, but each hole contains \aleph_n pigeons for some $n < \omega$. This peculiar behavior stems from the singular nature of \aleph_ω .

A cardinal κ is called *singular* if there is a cardinal $\lambda < \kappa$ and a function $f : \lambda \rightarrow \kappa$ such that $\bigcup_{\alpha < \lambda} f(\alpha) = \kappa$. (Remember, κ is transitive, so if $f(\alpha) \in \kappa$, then $f(\alpha) \subset \kappa$.) As an example, if we define $f : \aleph_0 \rightarrow \aleph_\omega$ by $f(n) = \aleph_n$, then $\bigcup_{\alpha < \aleph_0} f(\alpha) = \aleph_\omega$, showing that \aleph_ω is singular. Any infinite cardinal number that is not singular is called *regular*. One good example of a regular cardinal is \aleph_0 ; it is not equal to any finite union of finite cardinals. We can generalize the infinite pigeonhole principle for regular cardinals, but to prove the new result, we will need the following theorem.

Theorem 3.30. *For every infinite cardinal κ , $|\kappa \times \kappa| = \kappa$.*

We will postpone the proof of Theorem 3.30 for a while and jump right to the avian corollary. Since this is the last pigeonhole principle in this book, we will call it ultimate. Of course our list of pigeonhole principles is not all inclusive. For example, more set theoretic pigeonhole principles are given in [72].

Corollary 3.31 (Ultimate Pigeonhole Principle). *The following are equivalent:*

1. κ is a regular cardinal.
2. If we put κ pigeons into $\lambda < \kappa$ pigeonholes, then some pigeonhole must contain κ pigeons.

Proof. First suppose that $\lambda < \kappa$ and κ is regular. Suppose that $g : \kappa \rightarrow \lambda$ is an assignment of κ pigeons to λ pigeonholes. Define $f(\alpha) = |\{x \in \kappa \mid g(x) = \alpha\}|$ for each $\alpha < \lambda$, so $f(\alpha)$ is the population of the α th pigeonhole. Suppose, by way of contradiction, that $f(\alpha) < \kappa$ for each α . Because each $f(\alpha)$ is a cardinal, $\mu = \bigcup_{\alpha < \lambda} f(\alpha)$ is a cardinal. Furthermore, $\mu < \kappa$ because κ is regular. For each α , $f(\alpha) \leq \mu$, so the population of the α th pigeonhole can be matched with a subset of μ . Since there are λ pigeonholes, the entire pigeon population can be matched with a subset of $\mu \times \lambda$, so $\kappa \leq |\mu \times \lambda|$. Let $\nu = \max\{\mu, \lambda\}$. Since $\mu \leq \nu$ and $\lambda \leq \nu$, $|\mu \times \lambda| \leq |\nu \times \nu|$. By Theorem 3.30, $|\nu \times \nu| = \nu$. Since $\mu < \kappa$ and $\lambda < \kappa$, we have $\nu < \kappa$, and concatenating inequalities yields

$$\kappa \leq |\mu \times \lambda| \leq |\nu \times \nu| = \nu < \kappa,$$

a contradiction. Thus, for some α , $f(\alpha) = \kappa$ and the α th pigeonhole contains κ birds.

To prove the converse, suppose that κ is a singular cardinal. Then there is a cardinal $\lambda < \kappa$ and a function $f : \lambda \rightarrow \kappa$ such that $\cup_{\alpha < \lambda} f(\alpha) = \kappa$. Define $g : \kappa \rightarrow \lambda$ by letting $g(\beta)$ be the least α such that $\beta \in f(\alpha)$. Since $\cup_{\alpha < \lambda} f(\alpha) = \kappa$ and λ is well-ordered, g is well-defined and maps each element of κ to an element of λ . Furthermore, for each $\alpha < \lambda$,

$$|\{\beta \in \kappa \mid g(\beta) = \alpha\}| \leq |f(\alpha)| < \kappa.$$

Thus g can be viewed as an assignment of κ pigeons to λ pigeonholes so that the population of each pigeonhole is less than κ . □

The first part of the preceding proof can be adapted to prove that lots of cardinals are regular. This is a nice fact, since it means that we can apply the pigeonhole principle in lots of situations.

Corollary 3.32. *For each ordinal α , the cardinal $\aleph_{\alpha+1}$ is regular.*

Proof. We will sketch the argument. Suppose $f : \lambda \rightarrow \aleph_{\alpha+1}$, where λ is any cardinal such that $\lambda < \aleph_{\alpha+1}$. Then $\lambda \leq \aleph_\alpha$, and for each $\beta < \lambda$, $|f(\beta)| \leq \aleph_\alpha$. Applying Theorem 3.30 yields $|\cup_{\alpha < \lambda} f(\alpha)| \leq |\aleph_\alpha \times \aleph_\alpha| = \aleph_\alpha < \aleph_{\alpha+1}$. □

We should list the regular cardinals we have found. \aleph_0 is regular, and by the preceding corollary so are $\aleph_1 (= \aleph_{0+1})$, $\aleph_2 (= \aleph_{1+1})$, \aleph_3 , \aleph_4 , and so on. We have seen that the limit cardinal \aleph_ω is singular; the subscript cannot be written as $\alpha + 1$, so this does not contradict Corollary 3.32. However, $\aleph_{\omega+1}$ is regular, as are $\aleph_{\omega+2}$, $\aleph_{\omega+3}$, and so on. Our only good example of a regular limit cardinal is \aleph_0 . We do not have an example of an uncountable regular limit cardinal. The reason for this is explained in Section 3.6.

It seems that Theorem 3.30 is a handy way to bound the sizes of unions. Here is a nice way to capsulize that.

Corollary 3.33. *If κ is an infinite cardinal and $|X_\alpha| \leq \kappa$ for each $\alpha < \kappa$, then $|\cup_{\alpha < \kappa} X_\alpha| \leq \kappa$. In particular, a countable or finite union of at most countable sets is at most countable.*

Proof. Suppose $|X_\alpha| \leq \kappa$ for each $\alpha < \kappa$. For each α , let $g_\alpha : X_\alpha \rightarrow \kappa$ be a one-to-one map. Define $f : \cup_{\alpha < \kappa} X_\alpha \rightarrow \kappa \times \kappa$ by $f(x) = (\alpha, g_\alpha(x))$, where α is the least ordinal such that $x \in X_\alpha$. The function f is one-to-one, so $\cup_{\alpha < \kappa} X_\alpha \lesssim \kappa \times \kappa$. Thus by Theorem 3.30, $|\cup_{\alpha < \kappa} X_\alpha| \leq |\kappa \times \kappa| = \kappa$. To prove the particular case, let $\kappa = \aleph_0$. □

We have used Theorem 3.30 repeatedly, but still have not proved it. It is time to pay the piper.

Proof of Theorem 3.30. We will use induction to prove that $|\kappa \times \kappa| = \kappa$ for every infinite cardinal κ . For the base case, apply Corollary 3.22 to get $|\aleph_0 \times \aleph_0| = \aleph_0$.

As the induction hypothesis, assume $|\lambda \times \lambda| = \lambda$ for every infinite cardinal $\lambda < \kappa$. Since $\kappa \lesssim \kappa \times \kappa$, by the Cantor–Bernstein Theorem, it suffices to show that $\kappa \times \kappa \lesssim \kappa$.

Define the ordering $<$ on $\kappa \times \kappa$ as follows. Let $(\alpha, \beta), (\alpha', \beta') \in \kappa \times \kappa$ and let $\mu = \max\{\alpha, \beta\}$ and $\mu' = \max\{\alpha', \beta'\}$. We say that $(\alpha, \beta) < (\alpha', \beta')$ if and only if

- $\mu < \mu'$, or
- $\mu = \mu'$ and $\alpha < \alpha'$, or
- $\mu = \mu'$ and $\alpha = \alpha'$ and $\beta < \beta'$.

Informally, this relation sorts $\kappa \times \kappa$ by looking at maxima, then first elements, and then second elements. A routine but technical argument shows that $<$ is a well-ordering of $\kappa \times \kappa$. By Theorem 3.26, $\kappa \times \kappa$ under $<$ is order isomorphic to some ordinal. Let δ denote that ordinal, and let $f : \kappa \times \kappa \rightarrow \delta$ be the order isomorphism. If $\delta \leq \kappa$, then $\kappa \times \kappa \lesssim \delta \lesssim \kappa$ and the proof is complete.

Suppose by way of contradiction that $\kappa < \delta$. Since δ is an ordinal, we know that $\kappa \in \delta$, so there is an element $(\sigma, \tau) \in \kappa \times \kappa$ such that $f(\sigma, \tau) = \kappa$. Let μ denote $\max\{\sigma, \tau\}$ and note that $\sigma < \kappa, \tau < \kappa$, and consequently $\mu < \kappa$. Furthermore, by definition of the well-ordering on $\kappa \times \kappa$,

$$\{(\alpha, \beta) \in \kappa \times \kappa \mid f(\alpha, \beta) < \kappa\} \subset \mu \times \mu,$$

so $\kappa \lesssim \mu \times \mu$ and $\kappa \leq |\mu \times \mu|$. Let $\lambda = |\mu|$. Since $\lambda \sim \mu$, we have $|\mu \times \mu| = |\lambda \times \lambda|$. Since $\mu < \kappa$, λ is a cardinal less than κ , so by the induction hypothesis, $|\lambda \times \lambda| = \lambda < \kappa$. Concatenating inequalities yields $\kappa \leq |\mu \times \mu| = |\lambda \times \lambda| < \kappa$, a contradiction that completes the proof. □

Exercises

1. Define $f : \mathbb{N} \rightarrow \mathbb{Z}$ by $f(n) = (-1)^{n+1}(\frac{1}{4})(2n + 1) + (\frac{1}{4})$.
 - (a) Show that f is one-to-one. (Assume $f(j) = f(k)$ and prove $j = k$.)
 - (b) Show that f is onto. (Show that if $m > 0$, then $f(2m - 1) = m$ and if $t \leq 0$ then $f(-2t) = t$.)
2. Prove that \sim is an equivalence relation. That is, show that for all sets A, B and C , the following hold:
 - (a) $A \sim A$,
 - (b) $A \sim B \rightarrow B \sim A$, and
 - (c) $(A \sim B \wedge B \sim C) \rightarrow A \sim C$.
3. Show that for all sets A, B , and C , the following hold:
 - (a) $A \lesssim A$ (so \lesssim is reflexive), and

(b) $(A \lesssim B \wedge B \lesssim C) \rightarrow A \lesssim C$ (so \lesssim is transitive).

4. Details of the Cantor–Bernstein proof. This problem uses notation from the proof of Theorem 3.20.

(a) Use induction to prove Claim 1. As a base case, show $A_0 \supset A_1 \supset A_2$. For the induction step, assume $A_n \supset A_{n+1} \supset A_{n+2}$ and show that $A_{n+2} \supset A_{n+3}$, using the fact that $A_n \supset A_{n+1}$ implies $g \circ f(A_n) \supset g \circ f(A_{n+1})$.

(b) Prove Claim 2. Use the fact that either $x \in A'_\omega$ or there is a least j such that $x \notin A_j$ to show that each x is in some A'_n . (Here n is a natural number or ω .) To prove that x is in a unique A'_n , suppose that x is in two such sets, and seek a contradiction.

(c) Prove Claim 3. Use Claim 1 to get a short proof that $B_n \supset B_{n+1}$. The remainder of the argument parallels the proof of Claim 2.

(d) Prove Claim 4. To show that $h(A'_{2n}) = B'_{2n+1}$, note that because f is one-to-one we have $f(A_{2n}) - f(A_{2n-1}) = f(A_{2n} - A_{2n-1})$, and so

$$B'_{2n+1} = B_{2n+1} - B_{2n} = f(A_{2n}) - f(A_{2n-1}) = f(A'_{2n}) = h(A'_{2n}).$$

The proof that $B'_{2n} = h(A'_{2n+1})$ is similar. For the limit, the proof of $h(A'_\omega) = B'_\omega$ relies on the fact that since f is one-to-one, we must have that $f(\cap_{n \in \omega} A_n) = \cap_{n \in \omega} f(A_n)$.

5. Let \mathbb{Q} denote the set of rationals. Prove that $\mathbb{Q} \sim \mathbb{N}$.

6. Let Seq denote the set of all finite sequences of natural numbers. Prove that $Seq \sim \mathbb{N}$.

7. Without using Theorem 3.24, prove that \mathbb{Z} with the usual ordering is not well-ordered.

8. Using Theorem 3.24, prove that \mathbb{Z} with the usual ordering is not well-ordered.

9. Repeat Exercises 7 and 8 for the set $\mathbb{Q}^+ = \mathbb{Q} \cap [0, \infty)$ with the usual ordering.

10. Suppose that X is a transitive set.

(a) Prove that $X' = X \cup \{X\}$ is transitive.

(b) Prove that $\cup X$ is transitive.

11. Give an example of a nontransitive set where \leq_\in is a transitive relation. (Hint: The relation \leq_\in is vacuously transitive on every two-element set.)

12. Give an example of a transitive set where \leq_\in is not a transitive relation. (Hint: There is an example with three elements.)

13. Prove that if α is an ordinal, then so is $\alpha' = \alpha \cup \{\alpha\}$. (Hint: Exercise 10a shows that α' is transitive. Use the fact that α is well-ordered by \leq_\in to show that α' is too.)
14. Prove that if α and β are ordinals, then $\alpha \cap \beta = \alpha$ or $\alpha \cap \beta = \beta$. (Hint: Let $C = \alpha \cap \beta$ and suppose that $C \neq \alpha$ and $C \neq \beta$. Let γ be the least element of α such that $\gamma \notin C$. Show that $\gamma = C$, so $C \in \alpha$. Similarly, $C \in \beta$, so $C \in \alpha \cap \beta = C$, contradicting Theorem 3.13.)
15. Prove that if X is a set of ordinals, then $\cup X$ is an ordinal. (Hint: To show that $\cup X$ is transitive, note that if $x \in y \in \cup X$, then for some ordinal $z \in X$ we have $y \in z$. The fact that \leq_\in well-orders each element of X helps in proving antisymmetry and transitivity. Exercise 14 is useful in showing that trichotomy holds. Use the axiom of regularity to show that $\cup X$ has no infinite descending sequences.)
16. Prove that if α and β are ordinals, then $\alpha \in \beta$, $\alpha = \beta$, or $\beta \in \alpha$.
17. Details from the proof of Theorem 3.27.

Let $f : \beta_1 \rightarrow X$ and $g : \beta_2 \rightarrow X$ be γ -functions as defined in the proof of Theorem 3.27. Prove that if $\beta \in \beta_1 \cap \beta_2$, then $f(\beta) = g(\beta)$. (If f and g disagree, then there is a least β such that $f(\beta) \neq g(\beta)$. For this β , $f[\beta] = g[\beta]$. Apply the definition of a γ -function.)

18. Zorn's Lemma and AC.

Prove in ZF that the following are equivalent:

1. AC.
 2. Zorn's Lemma: Let P be a partial ordering (transitive and antisymmetric) such that every chain (linearly ordered subset) has an upper bound in P . Then P contains a maximal element (an element with no elements above it.)
 3. Every set can be well-ordered.
- (a) Prove that 1 implies 2. (Hint: Emulate the proof of Theorem 3.27. Suppose P has no maximal elements. For each chain C , let x_C be an upper bound of C that is not an element of C . Call C a γ -chain if for every $p \in C$, $x_{\{y \in C \mid y < p\}} = p$. Use the union of the γ -chains to derive a contradiction.)
- (b) Prove that 2 implies 3. (Hint: Fix a set X to well-order. Let P be the set of all one-to-one maps from ordinals to subsets of X . P is partially ordered by function extension. Show that every chain has an upper bound. Show that a maximal element maps an ordinal one-to-one onto X .)
- (c) Prove that 3 implies 1.

19. Describe a way to place \aleph_0 pigeons in \aleph_0 pigeonholes so that each pigeonhole contains at most one pigeon and \aleph_0 of the pigeonholes are empty. (Hint: $|\mathbb{Z}| = |\mathbb{N}| = \aleph_0$.)
20. Describe a way to place \aleph_0 pigeons in \aleph_0 pigeonholes so that each pigeonhole contains \aleph_0 pigeons. (Hint: $|\aleph_0 \times \aleph_0| = \aleph_0$.)
21. Describe a way to place \aleph_0 pigeons in \aleph_0 pigeonholes so that for each $\kappa \leq \aleph_0$ exactly one pigeonhole contains exactly κ pigeons.
22. Describe a way to place \aleph_0 pigeons in \aleph_0 pigeonholes so that for each $\kappa \leq \aleph_0$ exactly \aleph_0 pigeonholes contain exactly κ pigeons.
23. Show that if we put κ pigeons into λ pigeonholes and $\kappa < \lambda$, then λ pigeonholes will remain empty. (Hint: The set of all pigeonholes is the union of the empty holes and the occupied holes.)

3.6 Incompleteness and Cardinals

... we never assumed that (ZFC) included all the “true” facts.

— Levy [187]

In the last section we noted the absence of an example of an uncountable regular limit cardinal. We do not have an example because in ZFC we cannot prove that uncountable regular limit cardinals exist, assuming that ZFC is consistent. We have to *assume* that ZFC is consistent, because ZFC cannot prove that either. The last sentence is essentially Gödel’s Second Incompleteness Theorem and the starting point for our exploration of large cardinals.

3.6.1 Gödel’s Theorems for PA and ZFC

Gödel’s 1931 (paper) was undoubtedly the most exciting and the most cited article in mathematical logic and foundations to appear in the first eighty years of the (twentieth) century.

— Kleene [126]

Gödel’s First Incompleteness Theorem [124] is a statement about the provability of a formula in formal Peano Arithmetic (PA). In a nutshell, Gödel’s first theorem says that if PA is ω -consistent, then there is a formula G such that PA does not prove G and PA does not prove $\neg G$. In order to make this clear, we should discuss PA, ω -consistency, and the formula G .

The axioms of PA are an attempt to describe the important properties of the natural numbers under the operations of successor (adding one), addition, and multiplication. PA includes predicate calculus and axioms that say that

- 0 is not the successor of any element,

- $x + 0 = x$ and $x + (n + 1) = (x + n) + 1$,
- $x \cdot 0 = 0$ and $x \cdot (n + 1) = x \cdot n + x$, and
- the distributive laws hold.

PA also includes an induction scheme that can be used to prove a wealth of facts about the natural numbers. At one point, it was thought that PA might be able to prove every true statement about \mathbb{N} , but then Gödel's work ruled that possibility out.

We say that a theory is *consistent* if there is no formula A such that the theory proves both A and $\neg A$. A theory that is inconsistent and includes predicate calculus can prove every formula. Thus a theory is consistent if and only if there is some formula the theory cannot prove. When we say PA is ω -consistent, we mean PA cannot prove both $\exists x A(x)$ and every formula in the list $\neg A(0)$, $\neg A(1)$, $\neg A(2)$, and so on. Assuming ω -consistency is very reasonable, but a little stronger than assuming regular old consistency. Rosser [242] devised a way to prove Gödel's first theorem assuming only the consistency of PA. His replacement for Gödel's sentence G is a slightly more complicated formula.

Informally, Gödel's formula G says "there is no proof in PA of the formula G ." This is encoded in the language of arithmetic. Given our daily exposure to word processors and automated spelling and grammar checkers, we are used to the idea that formulas and lists of formulas (like proofs) can be represented as strings of zeros and ones, and that such strings can be viewed as integers and described by arithmetical formulas. It is very remarkable that Gödel devised and utilized an encoding scheme in 1931, long before the advent of electronic computers.

The method for making G refer to G is very entertaining. Let $G_0(x)$ be the formula that says "there is no number that encodes a proof in PA of the formula obtained by substituting the number x for the free variable in the formula encoded by the number x ." Suppose that n is the number that encodes $G_0(x)$. Note that the formula obtained by substituting n for the free variable in the formula encoded by n is exactly $G_0(n)$. Informally, $G_0(n)$ says "there is no number that encodes a proof in PA of $G_0(n)$." Thus, $G_0(n)$ is the desired formula G .

Once we have the encoding procedures in hand and have created the formula G , the remainder of the proof of Gödel's First Incompleteness Theorem is straightforward. Suppose that PA is ω -consistent. First, suppose that PA proves G . Then this proof is encoded by a number n and PA proves that "there is a number that encodes a proof in PA of G ." Thus, PA proves $\neg G$, contradicting the consistency of PA. Now suppose that PA proves $\neg G$. Then PA proves that "there is a number that encodes a proof of G ." By the ω -consistency of PA, we can find some number that actually does encode such a proof, and so PA proves G . Again, we have contradicted the consistency of PA.

Gödel's Second Incompleteness Theorem [124] says that if PA is consistent, then there is no proof in PA that PA is consistent. Much of the machinery used for the first theorem applies here also. The formula that asserts that PA is consistent, Con_{PA} , is an encoding of the sentence "there are no numbers x and y such that x

encodes a proof in PA of a formula and y encodes a proof in PA of the negation of that formula.” It is possible to prove in PA that $Con_{PA} \rightarrow G$. Thus, if PA proved Con_{PA} , then PA would prove G , and that contradicts Gödel’s First Incompleteness Theorem.

Perhaps the most remarkable quality of the incompleteness theorems is the ubiquity of their applicability. The theorems utilize only a few important features of PA and therefore apply to a wide variety of formal theories. In particular, the proofs of the theorems rely heavily on the ability to carry out a modest amount of arithmetic and the ability to check proofs in a mechanical fashion. Consequently, if a theory has enough axioms (to prove facts about encoding) but not too many axioms (so proof checking is not incomprehensibly complicated), then both incompleteness theorems apply. For example, both incompleteness theorems hold for ZFC. Thus, assuming that ZFC is consistent, there is a formula Z such that ZFC proves neither Z nor $\neg Z$, and ZFC does not prove Con_{ZFC} . The incompleteness theorems also hold for ZF and for any extensions of ZF by a finite number of axioms.

3.6.2 Inaccessible Cardinals

Better to reign in L, than serve in Heav’n.

— Milton, *Paradise Lost* (slightly misquoted)

If κ is an uncountable regular limit cardinal, then we say κ is *weakly inaccessible*. Our goal is to prove that the existence of weakly inaccessible cardinals is not provable in ZFC. At the end of this section we will link this back to our study of pigeonhole principles. The plan for achieving the goal is straightforward. The first step is to prove in ZFC that if there is a weakly inaccessible cardinal, then ZFC is consistent. Then we apply Gödel’s Second Incompleteness Theorem and get the desired result. The first step requires a journey to L , the constructible universe.

We will build the constructible universe in stages. Let $L_0 = \emptyset$. If L_α is defined, let $L_{\alpha+1}$ be the set of all subsets of L_α that are definable by restricted formulas with parameters from L_α . To be precise, a set X will be placed in $L_{\alpha+1}$ if all of the following conditions hold.

- $X \subset L_\alpha$,
- $u_1, u_2, \dots, u_n \in L_\alpha$ is a finite list of parameters,
- ψ is a formula in the language of set theory,
- ψ does not contain the power set or union symbols,
- each quantifier in ψ is of the form $\exists x \in L_\alpha$ or $\forall x \in L_\alpha$, and
- $X = \{y \in L_\alpha \mid \psi(y, u_1, u_2, \dots, u_n)\}$.

If α is a limit ordinal, let $L_\alpha = \cup_{\beta < \alpha} L_\beta$. This suffices to define L_α for each ordinal number α . Note that for each α , L_α is a set. The constructible universe, L , is the class defined by $L = \cup L_\alpha$, where the union ranges over all ordinal numbers. Neither $\{L_\alpha \mid \alpha \text{ is an ordinal number}\}$ nor L itself are sets, but it is convenient to refer to them using set-theoretic notation.

The finite levels of the constructible universe are simple in structure. For each $k < \omega$, each L_k is finite and $L_{k+1} = \mathcal{P}(L_k)$. By definition $L_0 = \emptyset$, and so $L_1 = \mathcal{P}(\emptyset) = \{\emptyset\}$ and $L_2 = \mathcal{P}(\{\emptyset\}) = \{\emptyset, \{\emptyset\}\}$. Rewriting with ordinal notation, $L_0 = 0$, $L_1 = 1$, and $L_2 = 2$. At the next level, L_3 breaks this pattern, since $L_3 = \mathcal{P}(L_2) = \{0, 1, \{1\}, 2\}$, which is not an ordinal. L_ω is defined by a union, so every element of L_ω is an element of L_k for some $k < \omega$. Beyond L_ω , the sets become vastly more complicated very rapidly. If κ is weakly inaccessible, the L_κ is an abridged version of the entire universe. This is stated more formally in the following theorem.

Theorem 3.34. *If κ is weakly inaccessible, then L_κ is a model of ZFC. That is, if we restrict the quantifiers in the axioms to L_κ , then the axioms of ZFC all hold.*

Comments on the proof. For a detailed treatment, see [80] or [73]. We provide only a hint of some of the main issues in the proof. The basic idea is to verify that the axioms of ZFC hold in V_κ in much the same way that one would show that the axioms defining vector spaces hold in \mathbb{R}^3 . Some axioms are very easy to manage. For example, $\emptyset \in L_1$, so the empty set axiom holds. If a and b are in L_α , then $\{a, b\}$ is in $L_{\alpha+1}$, so some instances of the pairing axiom are easy to verify. Verification of the infinity axiom relies on κ being larger than \aleph_0 , since every set in L_{\aleph_0} is finite. The verification of the power set axiom is particularly tricky and relies on the assumption that κ is a regular limit cardinal. □

Now we can finish our proof of the unprovability of the existence of a weakly inaccessible cardinal. The proof relies on the fact that any set of axioms with a model must be consistent. If M is any model for a set of axioms T , then every theorem that can be proved from T is true in M . (This is actually what makes proofs useful. If you prove a theorem from the axioms for vector spaces, then it has to be true in every vector space.) If T is inconsistent, then it proves a contradiction that would have to be true in M . But models are concrete (think of \mathbb{R}^3 as a model for the vector space axioms), so no contradiction can be true in a model. For a more technical discussion of models, truth, and provability, see [201].

Theorem 3.35. *If ZFC is consistent, then ZFC does not prove the existence of a weakly inaccessible cardinal.*

Proof. Assume that ZFC is consistent. Suppose, by way of contradiction, that ZFC proves that there is a cardinal κ that is weakly inaccessible. Then ZFC proves that the set L_κ exists. By Theorem 3.34 we know that L_κ is a model of ZFC, so ZFC is consistent. Thus ZFC proves the consistency of ZFC, contradicting Gödel’s Second Incompleteness Theorem and completing the proof. □

The unprovability of the existence of weakly inaccessible cardinals is not like the unprovability of AC in ZF. If we write I for the statement “there is a weakly inaccessible cardinal” and write $ZFC \vdash I$ for “ZFC proves I ,” then the preceding theorem says that $ZFC \not\vdash I$, provided that ZFC is consistent. To show that I is *independent* of ZFC (like AC is for ZF) we would also need to prove $ZFC \not\vdash \neg I$ assuming that ZFC is consistent. Thanks to Gödel, we know that this is an unattainable goal. To see this, suppose (for an eventual contradiction) that from Con_{ZFC} we can prove in ZFC that $ZFC \not\vdash \neg I$. If $ZFC \not\vdash \neg I$, then Con_{ZFC+I} . Thus, our hypothesis boils down to $ZFC \vdash Con_{ZFC} \rightarrow Con_{ZFC+I}$. Since $ZFC+I$ is an extension of ZFC, we have $ZFC+I \vdash Con_{ZFC} \rightarrow Con_{ZFC+I}$. Theorem 3.34 shows that $ZFC+I \vdash Con_{ZFC}$. Concatenating the last two lines, we obtain $ZFC+I \vdash Con_{ZFC+I}$, contradicting Gödel’s Second Incompleteness Theorem for $ZFC+I$.

Finally, we should summarize the combinatorial implications of this section. We know that if ZFC is consistent, then it cannot prove the existence of a weakly inaccessible cardinal. Also, κ is weakly inaccessible if and only if it is an uncountable regular limit cardinal. By the Ultimate Pigeonhole Principle, κ is regular if and only if whenever κ pigeons are placed in fewer than κ pigeonholes, then some hole contains κ pigeons. So a weakly inaccessible cardinal is an uncountable limit cardinal with this pigeonhole property. In Section 3.5 we proved that \aleph_0 and $\aleph_{\alpha+1}$ for each ordinal α have this pigeonhole property. In this section we have proved that we cannot prove the existence of any more cardinals with this property.

3.6.3 A Small Collage of Large Cardinals

All for one . . . and more for me.

— Cardinal Richelieu in *The Three Musketeers*

A cardinal number is said to be *large* if there is no proof in ZFC of its existence. We just met our first large cardinal, the weakly inaccessible cardinal. There are many other large cardinals related to combinatorial principles. This section lists the ones we need for the next two sections.

Recall that a weakly inaccessible cardinal is an uncountable regular limit cardinal. We say that a cardinal κ is a *strong limit* if for every $\lambda < \kappa$ we have $|\mathcal{P}(\lambda)| < \kappa$. An uncountable regular strong limit cardinal is called *strongly inaccessible* (or just inaccessible). Every strongly inaccessible cardinal is weakly inaccessible. If we assume the generalized continuum hypothesis (GCH), then every weakly inaccessible cardinal is also strongly inaccessible. (See Exercise 6.) The hypothesis GCH is independent of ZFC, and has inspired a great deal of interesting work [173].

We say that κ is *weakly compact* if κ is uncountable and $\kappa \rightarrow (\kappa)_2^2$. This arrow notation is the same used in Section 3.2, so $\kappa \rightarrow (\kappa)_2^2$ means that if we color the edges of a complete graph with κ vertices using two colors, then it must contain a monochromatic complete subgraph with κ vertices. These cardinals reappear in the next section.

We will also look at some results concerning subtle cardinals, another type of large cardinal defined in terms of colorings of unordered n -tuples. Suppose that κ is a cardinal and let $[\kappa]^n$ denote the set of n -element subsets of κ . We say that a function $S : [\kappa]^n \rightarrow \mathcal{P}(\kappa)$ is an (n, κ) -sequence if for each element of $[\kappa]^n$ of the form $\alpha_1 < \alpha_2 < \dots < \alpha_n < \kappa$, we have $S(\{\alpha_1, \alpha_2, \dots, \alpha_n\}) \subset \alpha_1$. A subset $C \subset \kappa$ is closed if the limit of each sequence of elements in C is either κ or in C . The subset $C \subset \kappa$ is unbounded if for each $\alpha \in \kappa$, there is a $\beta \in C$ such that $\alpha < \beta$. We abbreviate closed and unbounded by writing *club*. The cardinal κ is n -subtle if for every (n, κ) -sequence S and every club set $C \subset \kappa$, there exist elements $\beta_1, \beta_2, \dots, \beta_{n+1} \in C$ such that

$$S(\beta_1, \beta_2, \dots, \beta_n) = \beta_1 \cap S(\beta_2, \beta_3, \dots, \beta_{n+1}).$$

The basic idea is that given a coloring S and a large set C , the large set must contain some elements that are monochromatic. The n -subtle cardinals are closely related to n -ineffable cardinals. More information on both these types of cardinals can be found in [18], [19], [80], and especially [157].

One variation on coloring n -tuples for fixed values of n is to color n -tuples for all $n \in \omega$ simultaneously. Let $[\kappa]^{<\omega}$ denote the set of all finite subsets of κ . We say κ is a *Ramsey cardinal* and write $\kappa \rightarrow (\kappa)_2^{<\omega}$ if for every function $f : [\kappa]^{<\omega} \rightarrow 2$ there is a set X of size κ such that for each n , f is constant on $[X]^n$. Note that the same X works for all n , though when $j \neq k$ the j -tuples may not be the same color as the k -tuples.

One type of cardinal often mentioned in the literature is bigger than anything we have listed so far. We say that κ is a *measurable cardinal* if there is a κ -additive two-valued measure on κ . Roughly, this means that there is a way of assigning a value $\mu(X)$ to each $X \subset \kappa$ so that μ acts a lot like the measures that appear in analysis.

One way to organize all these cardinals is by comparing the sizes of the least example of each type of cardinal. Suppose we assign letters as follows:

- W : Least weakly inaccessible cardinal,
- I : Least strongly inaccessible cardinal,
- C : Least weakly compact cardinal,
- $S1$: Least 1-subtle cardinal,
- $S2$: Least 2-subtle cardinal,
- \vdots
- S_n : Least n -subtle cardinal,
- \vdots
- R : Least Ramsey cardinal,
- M : Least measurable cardinal.

Then we have the following relationships:

$$W \leq I < C < S1 < S2 < \dots < S_n < \dots < R < M.$$

The proofs of these relationships are frequently nontrivial. Good references include [80] and [168].

Exercises

1. For each $k \in \omega$ show that $|L_{k+1}| = 2^{|L_k|}$.
2. Using Exercise 1, prove that L_ω is countable.
3. Prove that for each ordinal α , L_α is transitive. (Hint: L_α is transitive if $x \in y \in L_\alpha$ implies $x \in L_\alpha$. Use induction on the ordinals.)
4. Prove that if $\alpha < \beta$, then $L_\alpha \subset L_\beta$.
5. Prove that if κ is a limit cardinal and $x \in L_\beta$ for some $\beta < \kappa$, then we also have $\cup x \in L_\kappa$.
6. The generalized continuum hypothesis (GCH) asserts that for every ordinal α , $|\mathcal{P}(\aleph_\alpha)| = \aleph_{\alpha+1}$. Assuming GCH, prove that every limit cardinal is a strong limit cardinal. As a corollary, show that GCH implies that every weakly inaccessible cardinal is strongly inaccessible.
7. Construct a 2-coloring f of $[\omega]^{<\omega}$ such that f is constant on $[\omega]^n$ for each n , but no pair has the same color as any triple.

3.7 Weakly Compact Cardinals

Watch out for that tree!

— *George of the Jungle* theme song

Theorem 3.3 says that if we 2-color a complete graph G with \aleph_0 vertices, then it must contain a monochromatic subgraph with \aleph_0 vertices. In arrow notation, this is written as $\aleph_0 \rightarrow (\aleph_0)_2^2$, where the first \aleph_0 is the size of G and the second \aleph_0 is the size of the desired monochromatic subgraph. It would be nice to know what other cardinals κ satisfy $\kappa \rightarrow (\kappa)_2^2$. We call an uncountable cardinal with this property *weakly compact*. Judging from the next theorem, our list of weakly compact cardinals may be very short.

Theorem 3.36. $|\mathbb{R}| \not\rightarrow (\aleph_1)_2^2$ and consequently, $\aleph_1 \not\rightarrow (\aleph_1)_2^2$.

Proof. Let $\kappa = |\mathbb{R}|$ and let $g : \kappa \rightarrow \mathbb{R}$ be a matching between the ordinals less than κ and the reals. Let G be a complete graph with κ vertices. We can think of each vertex of G as having two labels, an ordinal $\alpha < \kappa$ and a real number $g(\alpha)$. Color the edges of G using the scheme

$$\chi(\alpha\beta) = \begin{cases} \text{red} & \text{if } \alpha < \beta \leftrightarrow g(\alpha) < g(\beta), \\ \text{blue} & \text{if } \alpha < \beta \leftrightarrow g(\beta) < g(\alpha). \end{cases}$$

Informally, χ colors the edge $\alpha\beta$ red if the order on the ordinal labels agrees with the order on the real number labels, and colors the edge blue if the orders disagree.

Suppose that S is a subgraph of G and $|S| = \aleph_1$. We will show that S is not monochromatic. Since the ordinal labels for the vertices of S are a well-ordered subset of κ , we can list them in increasing order as $\langle \alpha_\gamma \mid \gamma < \aleph_1 \rangle$. We consider two cases.

First, suppose that S is red. Then the ordering on the real labels of the vertices of S agrees with the ordering on the ordinal labels. This gives us an uncountable well-ordered increasing sequence of reals, $\langle g(\alpha_\gamma) \mid \gamma < \aleph_1 \rangle$. Using the fact that \mathbb{Q} is dense in \mathbb{R} , for each $\gamma \in \aleph_1$, choose a rational q_γ such that $g(\alpha_\gamma) < q_\gamma < g(\alpha_{\gamma+1})$. Then $\langle q_\gamma \mid \gamma < \aleph_1 \rangle$ is an uncountable sequence of distinct rationals, contradicting the countability of \mathbb{Q} . (See Exercise 5 in Section 3.5.) Thus, S is not red.

Second, suppose that S is blue. This yields an uncountable decreasing sequence of reals. By choosing q_γ such that $g(\alpha_\gamma) > q_\gamma > g(\alpha_{\gamma+1})$ we obtain another contradiction in the same fashion as in the preceding case.

Summarizing, S is neither red nor blue. Thus G contains no monochromatic subgraph of size \aleph_1 .

To prove the last statement in the theorem, we note that by Theorem 3.18 and Corollary 3.23, $\aleph < \mathbb{R}$. Thus $\aleph_0 < |\mathbb{R}|$, and so $|\mathbb{R}| \geq \aleph_1$. Since there is a way to color a graph with $|\mathbb{R}|$ vertices so that no \aleph_1 -sized monochromatic subgraphs exist, we can certainly do the same for a (possibly smaller) graph with a mere \aleph_1 vertices. □

Summarizing, we know that $\aleph_0 \rightarrow (\aleph_0)_2^2$, but $\aleph_1 \not\rightarrow (\aleph_1)_2^2$. We can generalize Theorem 3.36 to show that $\aleph_{\alpha+1} \not\rightarrow (\aleph_{\alpha+1})_2^2$, eliminating all the successor cardinals from our hunt for weakly compact cardinals.

Theorem 3.37. $|\mathcal{P}(\aleph_\alpha)| \not\rightarrow (\aleph_{\alpha+1})_2^2$.

Proof. Imitate the preceding proof using $\mathcal{P}(\aleph_\alpha)$ in the role of \mathbb{R} . To do this, prove and use the fact that when $\mathcal{P}(\aleph_\alpha)$ is ordered by the relation

$$X < Y \quad \text{if and only if} \quad \min((X - Y) \cup (Y - X)) \in Y,$$

it contains no increasing or decreasing sequences of size $\aleph_{\alpha+1}$. □

At this point we know that any weakly compact cardinal must be an uncountable limit cardinal. In the following theorem, we emulate Erdős and Tarski [97] in showing that any weakly compact cardinal is inaccessible, and therefore large. Their studies of weakly compact cardinals were motivated by questions in infinite combinatorics stated in their 1943 paper [96].

Theorem 3.38. *If κ is weakly compact, then κ is strongly inaccessible.*

Proof. Suppose that $\kappa \rightarrow (\kappa)_2^2$. We need to show that κ is regular and a strong limit cardinal.

First suppose that κ is not regular, so for some λ there is a function $f : \lambda \rightarrow \kappa$ such that $\cup_{\alpha < \lambda} f(\alpha) = \kappa$. We may assume that f is increasing. We will use f to construct a 2-coloring of a complete graph with κ vertices. For each $\alpha < \beta < \kappa$, color the edge $\alpha\beta$ using the scheme

$$\chi(\alpha\beta) = \begin{cases} \text{red} & \text{if } \exists \gamma(\alpha < f(\gamma) \leq \beta), \\ \text{blue} & \text{otherwise.} \end{cases}$$

Informally, f chops κ into λ intervals. An edge is red if it connects two intervals. Thus, no red subgraph can be larger than size λ . Also, each interval chopped out by f is smaller than κ , so there is no blue subgraph of size κ . This contradicts $\kappa \rightarrow (\kappa)_2^2$, proving that no function like f exists and that κ is regular.

Now suppose that $\aleph_\alpha < \kappa$. If $\kappa \leq |\mathcal{P}(\aleph_\alpha)|$, then by Theorem 3.37, $\kappa \not\rightarrow (\aleph_{\alpha+1})_2^2$. But $\aleph_{\alpha+1} \leq \kappa$, so this implies that $\kappa \not\rightarrow (\kappa)_2^2$, contradicting weak compactness. Thus, if $\aleph_\alpha < \kappa$, then $|\mathcal{P}(\aleph_\alpha)| < \kappa$, proving that κ is a strong limit cardinal. □

We have seen that if κ is uncountable and $\kappa \rightarrow (\kappa)_2^2$, then κ is a large cardinal. By the results in Section 3.6, we know that ZFC cannot prove that such cardinals exist. Interestingly enough, increasing the number of colors (to any value less than κ) or increasing the size of the n -tuples (from 2 to any $n < \omega$) does not lead to larger cardinals. That is, if $\lambda < \kappa$, $n \in \omega$, and κ is weakly compact, then $\kappa \rightarrow (\kappa)_\lambda^n$. (Details are left as exercises.) We could color n -tuples for all $n \in \omega$, and look for a single set that is monochromatic for n -tuples for each n . This leads to the Ramsey cardinals mentioned in Section 3.6, which are considerably larger than the weakly compact cardinals.

In Section 3.1 we saw a close relationship between Ramsey’s Theorem and König’s Lemma. This relationship reappears at higher cardinalities. We say that a cardinal κ has the *tree property* if whenever T is a tree with κ many nodes and every level of T has size less than κ , then T must have a path of size κ . König’s Lemma (Theorem 3.2) says that \aleph_0 has the tree property. The next theorem shows the connection between weakly compact cardinals and cardinals with the tree property.

Theorem 3.39. *κ is weakly compact if and only if κ is strongly inaccessible and has the tree property.*

Pointers: Excellent proofs of this result can be found in the books of Drake (Theorems 3.5 and 3.7 in Chapter 7 of [80]), Jech (Lemma 29.6 in Chapter 5 of [168]), and Roitman (Theorem 36 in Chapter 7 of [241]). □

Theorem 3.39 does not characterize the smaller cardinals with the tree property. As we have already noted, \aleph_0 has the tree property. However, ZFC proves that \aleph_1 does not. There is a tree T such that $|T| = \aleph_1$, the cardinality of each level of T is less than \aleph_1 , and T has no paths of size \aleph_1 . Such a tree is called an \aleph_1 -Aronszajn tree, and stands as a counterexample to \aleph_1 having the tree property.

The existence of an \aleph_2 -Aronszajn tree is deducible from GCH in ZFC. On the other hand, Silver has shown that from the existence of a weakly compact cardinal we can prove the consistency of ZFC and “ \aleph_2 has the tree property.” Mitchell and Silver have a number of other results pertaining to the tree property. Finally, using the ordered list of cardinals that appears in Section 3.6, if κ is the least strongly inaccessible cardinal, then κ is strictly less than the least weakly compact cardinal, so by Theorem 3.39, κ cannot have the tree property. Thus if κ is the least strongly inaccessible cardinal, then there is a κ -Aronszajn tree.

Exercises

1. Prove that if $\kappa \rightarrow (\kappa)_2^2$, then $\kappa \rightarrow (\kappa)_\lambda^2$ for every $\lambda < \kappa$.
2. Prove that if $\kappa \rightarrow (\kappa)_\lambda^2$ for each $\lambda < \kappa$, then $\kappa \rightarrow (\kappa)_2^3$.
3. Assuming GCH, show that if κ is the least weakly inaccessible cardinal, then there is a κ -Aronszajn tree.
4. Find a proof that there is an \aleph_1 -Aronszajn tree. (Hint: A library is a good place to look for proofs.)

3.8 Infinite Marriage Problems

Infinite matching theory may seem rather mature and complete as it stands, but there are still fascinating unsolved problems. . .

— Reinhard Diestel [75]

We have considered the problem of matchmaking in the guise of graph matchings in Section 1.7, systems of distinct representatives in Section 1.7.2 and Section 3.4, and stable marriages in Section 2.9. Now we will study formulations of infinite marriage problems, expressed in some anthropocentric terminology.

Suppose M is a set of men. For each man $m \in M$, let $W(m)$ denote the women on his list of potential wives. For a set $S \subset M$, we will write $W(S) = \cup_{m \in S} W(m)$ for the combined lists of all the men in S . We will call the ordered pair (M, W) a *society*. A society is *espousable* if there is a one-to-one function $f : M \rightarrow W(M)$ such that for every $m \in M$ we have $f(m) \in W(m)$. We are requiring that f is an injection, so polygamy is disallowed. Every man must marry someone on his list, but some women may be left unmarried. Implicitly we assume that the collections of men and women are disjoint, so $M \cap W(M) = \emptyset$.

Some readers may argue that this terminology is quaint or sexist. We use it in deference to earlier authors (e.g. [2]) and because it makes the concepts very concrete and clear. Technically any matching application could be substituted, such as callers and circuits, readers and books, or pigeons and single occupancy pigeon holes.

Our goal is to determine exactly which societies are espousable. We can address certain situations by applying the theorems of various Halls.

3.8.1 Hall and Hall

She made it perfectly plain that she was his . . .

— Hall and Oates

If we have some group of seven men whose combined lists contain only five women, we have no hope of finding a wife for every man. In general, given any society (M, W) with a subpopulation $S \subset M$ such that $|S| > |W(S)|$, we know that the society is not espousable. When M is finite, this condition is the only possible barrier to solving a marriage problem, as shown by the following theorem of Philip Hall [147]. (An alternate form of this theorem was published earlier by D. König [180].)

Theorem 3.40. *If (M, W) is a society and $|M| < \aleph_0$, then the following are equivalent:*

1. *For every $S \subset M$, $|S| \leq |W(S)|$.*
2. *(M, W) is espousable.*

Proof. Given a society (M, W) , construct the graph G consisting of a vertex for each person and an edge from each man to each woman on his list. Thus the vertex set is $M \cup W(M)$ and the set of edges is $\{(m, W(m)) \mid m \in M\}$. This graph is bipartite; separate the vertices by gender. If the cardinality condition in item 1 holds, then by Theorem 1.51 there is a matching of the men into the women, so (M, W) is espousable. To prove the converse, suppose that (M, W) is espousable and that f is the injection of M into $W(M)$ matching each man to his wife. Then for each $S \subset M$, f restricted to S is an injection of S into $W(S)$, so by the definition of cardinality we must have $|S| \leq |W(S)|$. \square

The preceding theorem is usually viewed as a result about finite societies, but the only requirement is that the population of men is finite. The number of women could be anything (e.g. a woman for every real number) and the theorem still holds. Thus, as long as the number of men is finite, this theorem completely settles the question of which societies are espousable. If we allow an infinite number of men, problems may arise. By restricting ourselves to the case where each man has a finite list of women, we can use a theorem of Marshall Hall, Jr. [145] to settle the problem.

Theorem 3.41. *Suppose (M, W) is a society and for every $m \in M$ we have $|W(m)| < \aleph_0$. Then the following are equivalent:*

1. *For every $S \subset M$, $|S| \leq |W(S)|$.*
2. *(M, W) is espousable.*

Proof. First, we will prove the theorem for countable sets of men using a result from a previous section. Given a society (M, W) with $|M| \leq \aleph_0$, consider the family of sets defined by the formula $T = \{W(m) \mid m \in M\}$. Since $|M| \leq \aleph_0$,

we could index the members of T with natural numbers. Also, since each $W(m)$ is finite, we have that T is a countable family of finite sets. By Theorem 3.16, T has a system of distinct representatives (SDR) if and only if the condition in item 1 holds. From the definition of an SDR, T has an SDR if and only if (M, W) is spousesable. Thus item 1 holds if and only if item 2 holds.

Now we will prove the theorem for arbitrarily large sets of men. Our proof that item 1 implies item 2 will use the following compactness principle: “If T is a set of formulas and every finite subset of T has a model, then T has a model.” Just as in Section 3.6.2, if U is a model of T , then every object mentioned in formulas of T will appear in U , and every formula of T will be true of the objects in U . Suppose that (M, W) is a society satisfying item 1 in the theorem. For each $m \in M$, if $W(m) = \{w_1, \dots, w_n\}$ is the set of women on m 's list, add the formula $f(m) = w_1 \vee \dots \vee f(m) = w_n$ to T . For every pair $m_1, m_2 \in M$ with $m_1 \neq m_2$, add the formula $f(m_1) \neq f(m_2)$ to T . By the condition in item 1, every finite subset of T has a model. By the compactness principle, T has a model. Let f be the function in this model. Then f is a matching of the men into the women, as desired. The proof that item 2 implies item 1 is identical to that for Theorem 3.40. \square

From the proof, it appears that Theorem 3.41 for countable M is just Theorem 3.16 with some of the terminology changed. Indeed, Theorem 3.41 for countable M could be added to the list of results in Theorem 3.17. Thus, in ZF we can prove that this marriage theorem for countable M is equivalent to König's Lemma and also to the countable axiom of choice for finite sets (CACF). Even in axiom systems much weaker than ZF, a countable version of this marriage theorem can be shown to be equivalent to a weak version of König's Lemma [162].

Hall's [145] original proof of Theorem 3.41 for arbitrarily large sets of men uses Zorn's Lemma, a statement equivalent to the full axiom of choice (AC). Our proof uses the compactness principle from logic, which is equivalent to the Prime Ideal Theorem (PIT). Since PIT is weaker than AC (see [167]), we can also conclude that working in ZF it is not possible to deduce AC from Theorem 3.41. It is not known whether or not PIT can be deduced from Theorem 3.41 in ZF.

It may seem odd that we restrict the length of the lists of potential wives in Theorem 3.41. If some man has an infinite list, then he must not be very picky. On the surface, it seems like it should be easy to find him a wife. However, our intuition here is based on finite societies, and infinite societies (like infinite sets) are very peculiar.

Consider the following situation. Let $M = \{m_0, m_1, \dots\}$ be the men in our society and let $Y = \{w_0, w_1, \dots\}$ be the women. Let $W(m_0) = Y$ and for $j > 0$, let $W(m_j) = \{w_{j-1}\}$. If $S \subset M$, then the structure of $W(S)$ falls into two nice cases. If $m_0 \in S$, then $W(S) \supset W(m_0) \supset Y$, so $W(S) = Y$. If $m_0 \notin S$, then $W(S) = \{w_j \mid m_{j+1} \in S\}$. In either case, $|S| \leq |W(S)|$. Thus, the society (M, W) satisfies item 1 in Theorem 3.41 and satisfies all the hypotheses of Theorem 3.41 except that $|W(m_0)| = \aleph_0$. However, this society is not spousesable. By its construction, m_1 must marry w_0 , m_2 must marry w_1 ,

and so on. If m_0 marries some w_j , then m_{j+1} will be deprived of a wife. Thus, because m_0 has an infinite list, item 1 of Theorem 3.41 is no longer sufficient to guarantee that the society is espousable.

Exercises

1. Prove that if (M, W) is any espousable society then for every $S \subset M$, $|S| \leq |W(S)|$. (Hint: Since the society may have infinitely men or a man with an infinite list, neither Theorem 3.40 nor Theorem 3.41 applies directly. However, you could adapt part of a proof.)
2. Without using Theorem 3.17, prove in ZF that Theorem 3.41 for countable M implies CACF.
3. Without using Theorem 3.17, prove in ZF that Theorem 3.41 for countable M implies König's Lemma.
4. Prove in ZF that Theorem 3.41 implies the axiom of choice for finite sets (ACF).
5. Consider the society containing men $M = \{m_0, m_1, \dots\}$ and women $Y = \{w_0, w_1, \dots\}$, where $W(m_0) = Y$ and $W(m_j) = w_j$ for $j > 0$. Is (M, W) espousable?
6. Construct a society such that all the following hold:
 - (a) $W(m_0)$ is infinite.
 - (b) There are infinitely many women who are not on m_0 's list.
 - (c) For every finite $S \subset M$, $|S| \leq |W(S)|$.
 - (d) (M, W) is *not* espousable.
7. Suppose (M, W) is a society in which exactly one man has an infinite list and for all finite sets $S \subset M$, $|S| < |W(S)|$. (Note the strict inequality.) Is (M, W) espousable?

3.8.2 Countably Many Men

The squad may count off in a line or column formation.

– Drill sergeant study guide

In this subsection, we will examine four marriage theorems for the situation where the population of men is countable. We begin with the theorem of Damerell and Milner [66]. The statement of this theorem incorporates a margin function μ_{ω_1} , inspired by a conjecture of Nash-Williams [211]. Although the exact definition of μ_{ω_1} is somewhat technical, the underlying notion is easy. Suppose (M, W) is a society and let $Y \subset W(M)$ be some subset of the women. Build the set of all the men whose entire list lies within Y , so $S = \{m \in M \mid W(m) \subset Y\}$. If Y

and S are finite, then $|Y| - |S|$ indicates how many extra women Y contains. If $|Y| - |S|$ is ever negative, then the society is not espousable. The margin function μ_{ω_1} generalizes this notion of measuring the extra women in a set. We will define it precisely after we see how it is used in the following theorem of Damerell and Milner [66].

Theorem 3.42. *If (M, W) is a society and $|M| \leq \aleph_0$ then the following are equivalent:*

1. *For every $Y \subset W(M)$ we have $\mu_{\omega_1}(Y) \geq 0$.*
2. *(M, W) is espousable.*

In a moment, we will define μ_{ω_1} and make some comments on the proof of the theorem. Before we do that, it is worth taking a minute to note how this theorem compares with those in the preceding and following subsections. First, in this theorem, each man may have any number of women on his list. This is a change from M. Hall's Theorem 3.41 where each man was restricted to a finite number of women. On the other hand, Theorem 3.41 places no restriction on $|M|$ and this theorem requires the set of men to be countable. This is not an idle restriction. In [66], Damerell and Milner provide an example of a society with an uncountable collection of men that satisfies item 1 of the theorem, but fails to be espousable. This example and a theorem for uncountable collections of men is given in Section 3.8.3.

To clarify the statement of the theorem we must define μ_{ω_1} . To do this we need some ancillary definitions. Let (M, W) be a society with $|M| \leq \aleph_0$. For a set $Y \subset W(M)$, a *tower* on Y is defined to be an infinite nested sequence of sets $T = \langle T_n \mid n < \omega \rangle$ such that $T_0 \subset T_1 \subset T_2 \subset \dots$ and $Y = \cup_{n < \omega} T_n$. Given a tower T on Y , let $d(T)$ be the number of men whose lists are subsets of Y but not subsets of any element in the tower. That is,

$$d(T) = |\{m \in M \mid W(m) \subset Y \wedge \forall n W(m) \not\subset T_n\}|.$$

Let \mathbb{Z}^+ denote the extended integers, including the usual integers plus symbols for infinities, so $\mathbb{Z}^+ = \{\dots, -2, -1, 0, 1, 2, \dots\} \cup \{-\infty, \infty\}$. Now we can define μ_α for each ordinal α . In the case where $\alpha = 0$, for $Y \subset W(M)$ define

$$\mu_0(Y) = \begin{cases} |Y| - |\{m \in M \mid W(m) \subset Y\}| & \text{if } |Y| < \aleph_0 \\ \infty & \text{if } |Y| \geq \aleph_0. \end{cases}$$

The definition of the margin function for successor ordinals uses a special class of towers. Suppose that μ_α has been defined and that it maps subsets of $W(M)$ into \mathbb{Z}^+ . For each $Y \subset W(M)$ let $A_\alpha(Y)$ denote the collection of all towers $T = \langle T_n \mid n \in \omega \rangle$ on Y satisfying both $\mu_\alpha(T_0) < \infty$ and for all $n < \omega$, $\mu_\alpha(T_n) = \mu_\alpha(T_0)$. Thus, $A_\alpha(Y)$ consists of all towers on Y such that μ_α is constant and finite on the elements of the tower.

With the formulation of $A_\alpha(Y)$ in hand, we can complete the definition of the margin function. Suppose $\alpha = \beta + 1$ and μ_β and $A_\beta(Y)$ have been defined. Then for each $Y \subset W(M)$ define μ_α by

$$\mu_\alpha(Y) = \begin{cases} \inf_{T \in A_\beta(Y)} (\mu_\beta(T_0) - d(T)) & \text{if } A_\beta(Y) \neq \emptyset \\ \infty & \text{if } A_\beta(Y) = \emptyset. \end{cases}$$

If α is a limit ordinal, and μ_β is defined for all $\beta < \alpha$, then define

$$\mu_\alpha(Y) = \inf_{\beta < \alpha} \mu_\beta(Y).$$

In particular, $\mu_{\omega_1}(Y) = \inf_{\beta < \omega_1} \mu_\beta(Y)$ where ω_1 is the smallest uncountable ordinal.

In light of the complexity of the construction of μ_{ω_1} , it is not too surprising that the proof of Theorem 3.42 is beyond the scope of this book. However, we will make a few comments on the general structure of the proof and verify the theorem for a couple of societies. Theorem 3.42 states that a society with a countable number of men is espousable if and only if $\forall Y \subset W(M) \mu_{\omega_1}(Y) \geq 0$. The proof that the society is espousable actually constructs the matching one man at a time. The condition on μ_{ω_1} insures that the construction can proceed at each stage. The countability of M insures that the men can be counted off as m_0, m_1, m_2, \dots so that the construction does not involve a limit stage. To prove the converse, the authors assume the existence of a matching and show that for each $\alpha < \omega_1$ and each $Y \subset W(M)$ the inequality $\mu_\alpha(Y) \geq 0$ holds. This argument is carried out in a number of lemmas, largely proved by transfinite induction. These induction arguments have base cases, successor cases, and limit cases, mirroring the construction of μ_{ω_1} . For full details, see [66].

As an aside, we should elaborate on this use of constructions and induction proofs indexed by ordinals. These proof techniques are referred to as transfinite recursion and transfinite induction. Both principles can be proved in ZFC. In the case of recursion, ZFC proves that given any formula $\theta(x, y)$ and any ordinal α , if for every x there is a unique y such that $\theta(x, y)$ holds, then there is a function f with domain $\text{dom}(f) = \alpha$ such that for every $\beta < \alpha$ we have $\theta(f \upharpoonright_\beta, f(\beta))$. (The notation $f \upharpoonright_\beta$ denotes the function that is the same as f for inputs less than β and is undefined elsewhere. We call this the *restriction of f to β* .) For a proof of this transfinite recursion theorem, see [88]. In terms of the construction of our margin function, the transfinite recursion theorem asserts that there is a function f with domain $\omega_1 + 1$ such that for each $\alpha \leq \omega_1$, $f(\alpha) = \mu_\alpha$. Thus, not only does μ_{ω_1} exist, but a function encoding the entire construction of μ_{ω_1} exists.

The transfinite induction principle states that given an ordinal α and a formula $\theta(x)$, if both $\theta(0)$ and $\forall \beta \leq \alpha ((\forall \gamma < \beta \theta(\gamma)) \rightarrow \theta(\beta))$, then $\forall \beta \leq \alpha \theta(\beta)$. Just as in regular induction, we can view $\theta(0)$ as a base case and the formula $\forall \beta \leq \alpha ((\forall \gamma < \beta \theta(\gamma)) \rightarrow \theta(\beta))$ as an induction step. The proof of the induction step is often broken into cases where β is a successor ordinal and where β is a limit ordinal. Like the standard induction scheme, transfinite induction can be viewed

as a consequence of a least element principle. If $\theta(\beta)$ fails for some $\beta \leq \alpha$, then the set $\{\beta \leq \alpha \mid \neg\theta(\beta)\}$ exists. This set is well ordered, and thus has a least element β_0 . If $\beta_0 = 0$ then the base case fails, while if $\beta_0 > 0$ then the induction step fails. In this fashion, the existence of a least β_0 proves the contrapositive of the transfinite induction principle. This concludes our aside on transfinite recursion and induction; we return to our discussion of Theorem 3.42.

We will now verify Theorem 3.42 for two societies. We begin with the society from Subsection 3.8.1 that was not espousable.

Example (M, W) . Let (M, W) be the society in which

$$\begin{aligned} M &= \{m_0, m_1, m_2, \dots\}, \\ W(M) &= X = \{w_0, w_1, w_2, \dots\}, \\ W(m_0) &= X, \text{ and} \\ W(m_j) &= \{w_{j-1}\} \text{ for each } j > 0. \end{aligned}$$

As shown in Subsection 3.8.1, the society in Example (M, W) is not espousable. Consequently, we should be able to find a set of women $Y \subset W(M)$ such that $\mu_{\omega_1}(Y) < 0$. To do this, we will need to calculate some values of μ_0 and μ_1 , using the definitions following the statement of Theorem 3.42.

If $Y = \{w_{i_1}, \dots, w_{i_j}\}$ is a finite set of women, then the collection of all men m such that $W(m) \subset Y$ is exactly $\{m_{i_1+1}, \dots, m_{i_j+1}\}$. Because Y is finite and $W(m_0)$ is infinite, m_0 is never included in such a list. Since $\mu_0(Y)$ is the cardinality of Y less the number of these men, $\mu_0(Y) = 0$. On the other hand, if Y is infinite, then $\mu_0(Y) = \infty$. Summarizing,

$$\mu_0(Y) = \begin{cases} 0 & \text{if } |Y| < \aleph_0 \\ \infty & \text{if } |Y| \geq \aleph_0. \end{cases}$$

Given $\mu_0(Y)$, we can calculate $A_0(Y)$ for each set of women Y . When Y is finite, μ_0 is constantly 0 on every subset of Y , so $A_0(Y)$ will consist of all infinite nested sequences of subsets of Y that include Y . (Note that elements in the tower can be repeated.) If Y is infinite, $A_0(Y)$ consists of all infinite nested sequences of finite subsets of Y that eventually contain every element of Y .

With μ_0 and $A_0(Y)$ in hand, we are ready to calculate a value of μ_1 . Consider the tower $T = \langle T_i \mid i < \omega \rangle$ defined by $T_i = \{w_j \mid j \leq i\}$. We know that $T \in A_0(W(M))$. For each $j > 0$, $W(m_j) \subset T_j$, so m_0 is the only man satisfying both $W(m_0) \subset W(M)$ and $\forall n W(m_0) \not\subset T_n$. From the definition of $d(T)$, we have $d(T) = 1$. Now

$$\begin{aligned} \mu_1(W(M)) &= \inf_{T' \in A_0(W(M))} (\mu_0(T'_0) - d(T')) \\ &\leq \mu_0(T_0) - d(T) \\ &= 0 - 1 \\ &= -1. \end{aligned}$$

Thus $\mu_1(W(M)) < 0$. It can be shown that whenever $\alpha < \beta$, $\mu_\beta(Y) \leq \mu_\alpha(Y)$. (See the exercises.) Consequently, $\mu_{\omega_1}(W(M)) \leq \mu_1(W(M)) < 0$, and the first item of Theorem 3.42 fails, verifying the theorem.

To distinguish our next example from the work we have just completed, we will use boys and girls.

Example (B, G) . Let (B, G) be the society in which

$$\begin{aligned} B &= \{b_0, b_1, b_2, \dots\}, \\ G(B) &= \{g_0, g_1, g_2, \dots\}, \\ G(b_0) &= \{g_0, g_1\}, \text{ and} \\ G(b_j) &= \{g_j\} \text{ for each } j > 0. \end{aligned}$$

It is easy to see that the map $f(b_i) = g(i)$ is a solution to this marriage problem. We should be able to show that $\mu_{\omega_1}(Y) \geq 0$ for every possible set of girls Y .

We begin by finding $\mu_0(Y)$. If we choose a finite set of girls Y that contains g_0 but omits g_1 , then $\{m \in M \mid W(m) \subset Y\}$ will consist of exactly those boys whose indices match the indices of the non- g_0 girls of Y . In this case, g_0 is “extra,” and $\mu_0(Y) = 1$. For other finite Y , $\mu_0(Y)$ will always be 0. Summarizing, we have

$$\mu_0(Y) = \begin{cases} 0 & \text{if } |Y| < \aleph_0 \wedge (g_0 \notin Y \vee g_1 \in Y) \\ 1 & \text{if } |Y| < \aleph_0 \wedge g_0 \in Y \wedge g_1 \notin Y \\ \infty & \text{if } |Y| \geq \aleph_0. \end{cases}$$

Any sequence included in $A_0(Y)$ must be an infinite nested sequence of finite subsets of Y whose union is Y . Additionally, to insure that for every n we have $\mu_0(T_n) = \mu_0(T_0)$, either we must have both $g_0 \in T_0$ and $g_1 \notin Y$, or we must have that $g_0 \in T_k$ implies $g_1 \in T_k$ for all k . In the first case $\mu(T_n) = 1$ for all n , and in the second case $\mu(T_n) = 0$ for all n . If $W(m_0) = \{g_0, g_1\} \subset Y$, then $W(m_0) \subset T_n$ for some n . This implies that $d(T) = 0$ for all $T \in A_0(Y)$.

Now we are ready to calculate μ_1 . If $g_0 \in Y$ and $g_1 \notin Y$, then $\mu_0(T_0) = 1$ for all $T \in A_0(Y)$. Since $d(T) = 0$, we have $\mu_1(Y) = \mu_0(T_0) - d(T) = 1$ in this case. On the other hand, if $g_0 \notin Y$ or $g_1 \in Y$ then for every $T \in A_0(Y)$, $g_0 \in T_k$ implies $g_1 \in T_k$. In this case $\mu_0(T_0) = 0$, yielding $\mu_1(Y) = \mu_0(T_0) - d(T) = 0 - 0 = 0$. Summarizing, we have

$$\mu_1(Y) = \begin{cases} 0 & \text{if } g_0 \notin Y \vee g_1 \in Y \\ 1 & \text{if } g_0 \in Y \wedge g_1 \notin Y. \end{cases}$$

We have shown that for all $Y \subset G(B)$, $\mu_1(Y) \geq 0$. Furthermore, it is not too hard to see that for every Y , $A_0(Y) = A_1(Y)$. Since $d(T)$ is always 0, it immediately follows that $\mu_2(Y) = \mu_1(Y)$. An easy transfinite induction proves that $\mu_\alpha(Y) = \mu_1(Y)$ for all α and Y . In particular, when $\alpha = \omega_1$ we have $\mu_{\omega_1}(Y) = \mu_1(Y) \geq 0$ for all Y , so the first item of Theorem 3.42 holds, verifying the theorem for this espousable society.

Of course, our two verifications do not constitute a valid proof, nor do they even provide a significant body of empirical evidence. However, they do reveal the structure of the margin functions, and the interaction between the margin functions and the matchings (or absence of matchings) in societies.

After Damerell and Milner, other mathematicians devised necessary and sufficient conditions for the espousability of societies with countable sets of men. We will collect three statements of this sort below in Theorem 3.43. To unify their presentation, we will introduce the terminology of strings and admissibility as found in Wojciechowski [286].

The simplest form of admissibility is due to Podewski and Steffens [223]. The marriage problem (M, W) is *c-admissible* if there is no subset $J \subset M$ for which there is an element $m \in M - J$ such that $W(m) \subset W(J)$ and the society (J, W) has a unique solution. Note that if such a J and $m \in M - J$ exist, then the society $(J \cup \{m\}, W)$ is not espousable. Thus a society is *c-admissible* if it avoids these egregiously solutionless subsocieties.

The other two forms of admissibility use the following notion. A *string* in (M, W) is a one-to-one function from an ordinal into $M \cup W(M)$. We can have strings of men, strings of women, or mixed gender strings. A string f is *saturated* if no man precedes any woman on his list. That is, f is saturated if whenever $f(\beta) = m$, then $W(m) \subset \{f(\alpha) \mid \alpha < \beta\}$.

Saturated strings are used in the definition of μ -admissibility introduced by Wojciechowski [286]. Define the function μ on strings f as follows. Suppose α is $\text{dom}(f)$ (the domain of f). If $\alpha = 0$, then let $\mu(f) = 0$. If $\alpha = \beta + 1$, then define $\mu(f)$ by

$$\mu(f) = \begin{cases} \mu(f \upharpoonright_\beta) + 1 & \text{if } f(\beta) \in W(M) \\ \mu(f \upharpoonright_\beta) - 1 & \text{if } f(\beta) \in M. \end{cases}$$

If α is a limit, then define $\mu(f) = \liminf_{\beta \rightarrow \alpha} \mu(f_\beta)$. We say that (M, W) is μ -admissible if $\mu(f) \geq 0$ for every saturated string of (M, W) . Informally, a saturated string is like a line of people. The margin function $\mu(f)$ indicates how many spare women are in the line defined by f . If (M, W) is μ -admissible, then the only way a line can have a negative number of spare women is if some man jumps in front of a woman on his list.

Our third form of admissibility is due to Nash-Williams [212]. Suppose f is a string of women. Define $q(f)$ as follows. Suppose $\alpha = \text{dom}(f)$. If $\alpha = 0$, then define $q(f) = -|\{m \in M \mid W(m) = \emptyset\}|$. In the following let $\text{ran}(f)$ denote the range of f and let $f \upharpoonright_\beta$ denote the restriction of f to β . If $\alpha = \beta + 1$, then define

$$q(f) = q(f \upharpoonright_\beta) - |\{m \in M \mid W(m) \subset \text{ran}(f) \wedge W(m) \not\subset \text{ran}(f \upharpoonright_\beta)\}|.$$

When α is a limit, define

$$q(f) = \liminf_{\beta \rightarrow \alpha} q(f \upharpoonright_\beta) - |\{m \in M \mid W(m) \subset \text{ran}(f) \wedge \forall \beta < \alpha \ W(m) \not\subset \text{ran}(f \upharpoonright_\beta)\}|.$$

We say that (M, W) is *q-admissible* if $q(f) \geq 0$ for every string f of women. The construction of q parallels the definition of μ_{ω_1} , but substitutes strings for towers. This is Nash-Williams' simplification of the concepts in Damerell and Milner's work on his conjecture.

With all this terminology, we can state Theorem 3.43, a sort of historical tour of marriage theorems for societies with countable populations of men.

Theorem 3.43. *If (M, W) is a society and $|M| \leq \aleph_0$, then the following are equivalent:*

1. (M, W) is *espousable*.
2. (M, W) is *c-admissible*.
3. (M, W) is *q-admissible*.
4. (M, W) is *μ -admissible*.

Proofs of the equivalences above can be found in the papers of Podewski and Steffens [223] (for item 2), Nash-Williams [212] (for item 3), and Wojciechowski [286] (for item 4). Verifications of these marriage theorems for the previously analyzed examples are left as exercises.

Exercises

1. Prove that for all Y , $\mu_{\alpha+1}(Y) \leq \mu_\alpha(Y)$.
2. Use transfinite induction to prove that for every pair of ordinals α and β , if $\alpha \leq \beta$ then $\mu_\beta(Y) \leq \mu_\alpha(Y)$ for all Y . (Hint: Fix α , let $\beta = \alpha + \gamma$ and use induction on γ . Exercise 1 is the induction step.)
3. Prove directly that the society in Example (M, W) is not *c-admissible*.
4. Prove directly that the society in Example (M, W) is not *q-admissible*.
5. Prove directly that the society in Example (M, W) is not *μ -admissible*.
6. Prove directly that the society in Example (B, G) is *c-admissible*.
7. Prove directly that the society in Example (B, G) is *q-admissible*.
8. Prove directly that the society in Example (B, G) is *μ -admissible*.

3.8.3 Uncountably Many Men

*If you can count the leaves of the trees,
Or the foaming waves of the untamed seas,
Then I will entrust to you alone
To reckon the amours I have known.*

— Anacreontea, Ode 14 (translated by J. F. Davidson [67])

The marriage theorems of the previous section allow each man an arbitrarily long list of prospective wives, but require that the set of men is countable. This is not an idle restriction. The theorems all fail for certain societies with uncountably many men.

Consider the following example, taken from [66]. We use ω_1 to denote the first uncountable ordinal and define the society (M, W) by

$$M = \{m_\alpha \mid \omega \leq \alpha < \omega_1\},$$

$$W(M) = \{y_\alpha \mid \alpha < \omega_1\}, \text{ and}$$

$$\forall m_\alpha \in M \ W(m_\alpha) = \{y_\beta \mid \beta < \alpha\}.$$

Thus we have a man for every infinite ordinal less than ω_1 , a woman for every ordinal (infinite or finite) less than ω_1 , and each man's list consists of those women whose indices are strictly less than his own. Since m_ω is the man of lowest index (poor guy), he has the shortest list. Even he has an infinitely long list of potential wives. In fact, every man's list contains exactly \aleph_0 women.

We will show that (M, W) is not espousable. By way of contradiction, suppose that $f : M \rightarrow W(M)$ is a one-to-one function. Beginning with any α such that $\omega \leq \alpha < \omega_1$, build a sequence of men as follows. Start with m_α . If $f(m_\alpha) = y_\beta$ and $\beta \geq \omega$, add m_β to the left of m_α . If $f(m_\gamma) = y_\alpha$, add m_γ to the right of m_α . Since f maps men to women with smaller indices, we must have $\beta < \alpha < \gamma$. Continue in this fashion, adding men to the right of the right end of the sequence and to the left of the left end of the sequence for as long as possible. Since the sequence to the left gives a descending sequence of ordinals, the sequence must have a left termination point. The only way for the sequence to terminate is if the last man is mapped to a woman with an index less than ω . Thus every sequence constructed in this fashion terminates on the left with a man married to some woman y_k with $k < \omega$. Because f is one-to-one, no two sequences can terminate with men married to the same woman, so each sequence terminates on the left with a man married to a *unique* woman y_k with $k < \omega$. This shows that there are countably many sequences. The sequence for any particular man m_α could extend infinitely to the right. However, since there is no limit stage in the construction, the sequence to the right will have at most \aleph_0 elements. Thus we have countably many sequences, each containing countably many men. Every man is in one of these sequences, so there are countably many men. However, we said that the set of men is $M = \{m_\alpha \mid \omega \leq \alpha < \omega_1\}$, which is uncountable. This contradiction shows that the society (M, W) is not espousable.

If we could apply Theorem 3.43 here, then (M, W) would not be c -admissible. However, (M, W) is c -admissible, as the following argument shows. Suppose $J \subset M$ is any subset of men such that the subsociety (J, W) is espousable. Let $f : J \rightarrow W(J)$ be a one-to-one function espousing (J, W) . On the one hand, suppose that for some $m_\alpha \in J$, $f(m_\alpha) = y_\beta$ and $\beta + 1 < \alpha$. We can modify the solution f by setting $f(m_\alpha) = y_{\beta+1}$ and, if necessary, marrying the former husband of $y_{\beta+1}$ to y_β . In this case, f is not a unique solution to the marriage

problem. On the other hand, if $f(m_\alpha) = y_\beta$ always implies $\beta + 1 \geq \alpha$, then nobody is married to y_0 . Thus, we can modify f by marrying someone to y_0 . In either case, the solution of (J, W) is not unique, so (M, W) is c -admissible. Similar proofs could be used to show that (M, W) is q -admissible, μ -admissible, and satisfies item 1 of Theorem 3.42. Thus none of the results of the previous section remain true when the hypothesis of the countability of M is removed from their statements.

In [2] and [1], Aharoni, Nash-Williams, and Shelah present a marriage theorem that does allow uncountable populations of men. Unlike Theorem 3.41, there is no restriction on the lengths of the lists of the men. Thus the hypothesis for the statement is particularly brief. Here is the theorem.

Theorem 3.44. *Suppose (M, W) is a society. The following are equivalent:*

1. (M, W) is espousable.
2. For every cardinal κ such that $0 < \kappa \leq \aleph_0$ or κ is regular, there is no κ -obstruction in (M, W) .

Since the proof of this theorem in [1] is twenty-five pages long, we will not replicate it here. However, to understand the statement of the theorem, we should unravel the definition of a κ -obstruction. This will require application of some of our knowledge of cardinal numbers. Item 2 above mentions finite cardinals, which are just natural numbers, and regular cardinals, which are defined in Section 3.5.3. Good examples of regular cardinals include \aleph_0 and \aleph_1 . By contrast, \aleph_ω is good example of a singular (i.e. not regular) cardinal.

Every κ -obstruction turns out to be a subsociety of (M, W) , and some new notation will prove helpful. In this setting, it is easiest to identify (M, W) with the bipartite graph with vertices in M and $W(M)$. Given some subset Δ of the vertices (people) in (M, W) , we can easily construct the subsociety (which we will also denote with Δ) corresponding to the subgraph of (M, W) with vertices in Δ . We can also form $(M, W) - \Delta$, the society whose graph is the subgraph of (M, W) on the vertices that do not appear in Δ . We say that Δ is a *saturated* subsociety of (M, W) if whenever $m \in \Delta$, $W(m) \subset \Delta$.

Our definition of κ -obstructions is a transfinite recursion with infinitely many base cases. Here are the base cases. Suppose that $0 < \kappa < \aleph_0$ or $\kappa = \aleph_0$. We say that Δ is a κ -obstruction of (M, W) if Δ is a saturated subsociety and there is a set $L \subset M$ with $|L| = \kappa$ such that $\Delta - L$ is espousable, but for every woman $w \in \Delta$, the society $\Delta - (L \cup \{w\})$ is not espousable. Informally, if Δ is a κ -obstruction (for $\kappa \leq \aleph_0$), then L is a set of excess men preventing Δ from having an espousal. The requirement that Δ is saturated transfers the unespousability of Δ to the society (M, W) . Thus if (M, W) has a κ -obstruction for $\kappa \leq \aleph_0$, then (M, W) is not espousable.

To define κ -obstructions for uncountable regular cardinals κ , we will need to review and introduce some set theoretic terminology. In Section 3.6.3 we said that a set $C \subset \kappa$ is *club* (shorthand for **closed and unbounded**) if

- the limit of each sequence of elements in C is either κ or in C , and
- $\forall \alpha \in \kappa \exists \beta \in C (\alpha < \beta)$.

We say that a set $S \subset \kappa$ is *stationary* if for every club $C \subset \kappa$, we have $S \cap C \neq \emptyset$. To use a topological metaphor, stationary sets meet every club set.

Now we are ready for the inductive step in the definition of κ -obstructions of (M, W) . Suppose that κ is regular, and we have defined μ -obstructions for every $\mu < \kappa$ such that $0 < \mu \leq \aleph_0$ or μ is uncountable and regular. We say that a subsociety Δ of (M, W) is a κ -obstruction of (M, W) if there is a transfinite sequence $\langle \Delta_\alpha \rangle_{\alpha < \kappa}$ of disjoint subsocieties such that

- $\Delta = \cup_{\alpha < \kappa} \Delta_\alpha$,
- for each $\alpha < \kappa$, either Δ_α is a single woman, or for some $\beta < \kappa$, the subsociety Δ_α is a β -obstruction in $(M, W) - \cup_{\theta < \alpha} \Delta_\theta$, and
- the set of $\alpha < \kappa$ for which Δ_α is not a single woman is stationary.

This completes the definition of κ -obstructions for all those cardinals κ such that $0 < \kappa \leq \aleph_0$ or κ is regular.

At the beginning of this section we showed that the society (M, W) that has $M = \{m_\alpha \mid \omega \leq \alpha < \omega_1\}$ and $W(m_\alpha) = \{y_\beta \mid \beta < \alpha\}$ for all infinite $\alpha < \omega_1$ is not espousable. In light of Theorem 3.44, this society must have a κ -obstruction for some κ that either satisfies $0 < \kappa \leq \aleph_0$ or is a regular cardinal. In fact, we can find an \aleph_1 -obstruction for (M, W) .

We define the \aleph_1 -obstruction for (M, W) by applying transfinite recursion. Let $\Delta_0 = \{y_0\}$. For each successor ordinal $\alpha + 1 < \aleph_1$, let $\Delta_{\alpha+1} = \{y_{\alpha+1}\}$. Finally, if $\lambda < \aleph_1$ is a limit ordinal, let $\Delta_\lambda = \{m_\lambda, m_{\lambda+1}, y_\lambda\}$. This defines a sequence $\langle \Delta_\alpha \rangle_{\alpha < \aleph_1}$ of disjoint subsocieties of (M, W) . Let $\Delta = \cup_{\alpha < \aleph_1} \Delta_\alpha$. We must use the definition of a κ -obstruction to prove that Δ is an \aleph_1 -obstruction. For each $\alpha < \aleph_1$, either the subsociety Δ_α is a single woman or α is a limit. If α is a limit, then $\Delta_\alpha = \{m_\alpha, m_{\alpha+1}, y_\alpha\}$. Since y_α is the only woman in $(M, W) - \cup_{\theta < \alpha} \Delta_\theta$ with an index less than $\alpha + 1$, Δ_α is a saturated subsociety of $(M, W) - \cup_{\theta < \alpha} \Delta_\theta$. Furthermore, the subsociety $\Delta_\alpha - \{m_\alpha\}$ is espousable, but $\Delta_\alpha - \{m_\alpha, y_\alpha\}$ is not. Thus if α is a limit, then Δ_α is a 1-obstruction of $(M, W) - \cup_{\theta < \alpha} \Delta_\theta$. Finally, the set of $\alpha < \aleph_1$ such that Δ_α is not a single woman is exactly the set $S = \{\lambda < \aleph_1 \mid \lambda \text{ is a limit}\}$, which is stationary. (This is an exercise.) This completes the verification that Δ is an \aleph_1 -obstruction of (M, W) .

If we think of a society as a bipartite graph, then every κ -obstruction is a subgraph of the society. With this viewpoint, Theorem 3.44 states that the graph for a society has a matching of the men into the women if and only if it contains no subgraphs of a particular form. Thus, Theorem 3.44 is a “forbidden subgraph” characterization of all the espousable societies.

Exercises

1. Prove that if κ is an uncountable regular cardinal, then the set of ordinals $\{\lambda < \kappa \mid \lambda \text{ is a limit}\}$ is stationary.
2. Prove that if a society (M, W) has an \aleph_0 -obstruction, then it also has a 1-obstruction.
3. Suppose that (M, W) is the society with

$$M = \{m_\alpha \mid \alpha < \omega\},$$

$$W(m_{\alpha+1}) = \{w_\alpha\} \text{ for each } \alpha < \omega, \text{ and}$$

$$W(m_0) = \{w_\alpha \mid \alpha < \omega\}.$$

Complete the following.

- (a) Find a 1-obstruction for (M, W) .
 - (b) Prove that (M, W) has no 2-obstructions.
4. Suppose that (M, W) is the society with

$$M = \{m_\alpha \mid \omega \leq \alpha < \omega_1\},$$

$$W(M) = \{y_\alpha \mid \alpha < \omega_1\}, \text{ and}$$

$$\forall m_\alpha \in M \ W(m_\alpha) = \{y_\beta \mid \beta < \alpha\}.$$

Complete the following.

- (a) Prove that (M, W) has no 1-obstruction.
- (b) Prove that (M, W) has no \aleph_0 -obstruction. (Hint: You could do and apply Exercise 2.)
- (c) Prove that (M, W) is μ -admissible.
- (d) If you survived part 4c, then prove that (M, W) is p -admissible.

3.8.4 Espousable Cardinals

It is easy to deduce from this that weakly compact cardinals are very large...

— E. C. Milner [206]

For the past two sections we have been modifying our marriage theorems in order to make them hold for larger and larger cardinals. This process has been pretty successful, resulting in Theorem 3.44 that can handle arbitrarily large societies. However, the statement of Theorem 3.44 is substantially more complicated than that of Marshall Hall Jr.'s Theorem 3.41. An alternative approach is to concoct a natural simple generalization of Theorem 3.41 and ask which cardinals satisfy it. This is the motivation for the following definition.

Definition. An infinite cardinal κ is *espousable* if every society (M, W) that satisfies the following criteria is espousable.

1. $|M| = \kappa$,
2. $\forall m \in M |W(m)| < \kappa$, and
3. if $M_0 \subset M$ and $|M_0| < \kappa$, then (M_0, W) is espousable.

Informally, “ κ is espousable” means that a result very like Theorem 3.41 holds for all societies of size κ . There is some variation from the theorem. The third clause requires espousability of all small subsocieties where Theorem 3.41 just specifies that there is a sufficiently large number of women. This modification is necessary, since the unspousable uncountable society from the beginning of Section 3.8.3 satisfies the cardinality requirements from Theorem 3.41 and could be embedded in any uncountable society.

From Theorem 3.41, we know that \aleph_0 is espousable. The unspousable and uncountable society from the beginning of Section 3.8.3 shows that \aleph_1 is not espousable. (You may wish to check the details of these claims as exercises.) Thus the espousable cardinals form a proper subclass of the class of all cardinals. The following two theorems of Shelah [252] give additional properties of espousable cardinals.

Theorem 3.45. *If κ is an uncountable espousable cardinal, then κ is a limit cardinal.*

Proof. It is easiest to show that if κ is an uncountable successor cardinal, then κ is not espousable. We will show that a particular uncountable successor cardinal, \aleph_2 , is not espousable, and leave the generalization of the argument as an exercise.

Our goal is to show that \aleph_2 is not espousable. We modify the example from the beginning of Section 3.8.3 as follows. Let

$$M = \{m_\alpha \mid \aleph_1 \leq \alpha < \aleph_2\},$$

$$W(M) = \{y_\alpha \mid \alpha < \aleph_2\}, \text{ and}$$

$$\text{for each } m_\alpha \in M, W(m_\alpha) = \{y_\beta \mid \beta < \alpha\}.$$

We must verify that (M, W) satisfies all the criteria in the definition of espousable cardinal. The cardinality of M is \aleph_2 and for each man m_α , $|W(m_\alpha)| = \aleph_1 < \aleph_2$. Furthermore, if $M_0 \subset M$ and $|M_0| < \aleph_2$, then $|M_0| = \aleph_1$, so any bijection between M_0 and $\{y_\alpha \mid \alpha < \aleph_1\}$ is an espousal of the subsociety (M_0, W) . Summarizing, (M, W) satisfies all the criteria listed in the definition of espousable cardinals. However, the argument on page 337 can be adapted to show that (M, W) is not espousable. Thus \aleph_2 is not an espousable cardinal. Generalization of this argument to other successor cardinals is left as an exercise. \square

Theorem 3.46. *If κ is an uncountable espousable cardinal, then κ is regular.*

Proof. As in the preceding theorem’s proof, it is easiest to show that if κ is both uncountable and not regular then κ is not espousable. Again, we will prove the theorem for a particular case and leave the generalization as an exercise.

Suppose that κ is \aleph_ω . We know that \aleph_ω is not regular because $\cup_{\alpha < \omega} \aleph_\alpha = \aleph_\omega$. We will show that \aleph_ω is not espousable by gluing together versions of the society from the end of Section 3.8.1. Let (M, W) be the society defined by:

$$\begin{aligned} M &= \{p\} \cup \{m_\alpha \mid \alpha < \aleph_\omega\}, \\ W(M) &= \{w_\alpha \mid \alpha < \aleph_\omega\}, \\ W(m_\alpha) &= \{w_\alpha\} \text{ if } \alpha \notin \{\aleph_n \mid n \in \omega\}, \\ W(m_{\aleph_n}) &= \{w_\beta \mid \aleph_n \leq \beta < \aleph_{n+1}\}, \text{ and} \\ W(p) &= \{w_{\aleph_n} \mid n \in \omega\}. \end{aligned}$$

It is easy to verify that (M, W) satisfies the first two criteria in the definition of an espousable cardinal. Suppose that $M_0 \subset M$ and $|M_0| < \aleph_\omega$. If $p \notin M_0$ then matching each man to the woman with his index yields an espousal. If $p \in M_0$, we must work a little harder. Find the least j with $M_0 \not\supseteq \{m_\alpha \mid \aleph_j \leq \alpha < \aleph_{j+1}\}$, and let α_0 be some ordinal in $[\aleph_j, \aleph_{j+1})$ such that $m_{\alpha_0} \notin M_0$. If we let $f(p) = w_{\aleph_j}$, $f(m_{\aleph_j}) = w_{\alpha_0}$, and $f(m_\beta) = w_\beta$ for all other $\beta \in M_0$, then F is an espousal of the subsociety for M_0 . If \aleph_ω was espousable, then (M, W) would have an espousal. However, any assignment of a wife to man p leads to a contradiction. Thus \aleph_ω is not an espousable cardinal. To generalize this argument to other non-regular cardinals, just replace $\aleph_0, \aleph_1, \dots$ with any cofinal sequence witnessing that the cardinal under consideration is singular. \square

On the basis of the preceding two theorems, we know that if κ is an uncountable espousable cardinal, then κ is an uncountable regular limit cardinal, and so κ is weakly inaccessible. Applying Theorem 3.35, ZFC cannot prove the existence of an uncountable espousable cardinal. However, the following theorem shows that espousable cardinals are not too high in our hierarchy of large cardinals. (Weakly compact cardinals are very large, but not too large.)

Theorem 3.47. *Every weakly compact cardinal is espousable.*

Sketch of proof. Suppose that κ is weakly compact and (M, W) is a society that satisfies the criteria in the definition of espousable cardinals. Without any loss of generality, let $M = \{m_\alpha \mid \alpha < \kappa\}$. Build a tree of height κ in which the nodes at level α of T are labeled with the names of women in $W(m_\alpha)$, and each path through the tree is an initial segment of an espousal of (M, W) . Since κ is weakly compact, it is a regular strong limit cardinal. This can be used to prove that each level of T has size less than κ . Applying Theorem 3.39, use the tree property for κ to find a path through T of length κ . This path is the desired espousal of (M, W) . \square

We now have both upper and lower bounds on the size of the least uncountable espousable cardinal, but there are some obvious questions. Is every uncountable espousable cardinal a strong limit cardinal? Does every uncountable espousable cardinal have the tree property? Is every uncountable cardinal with the tree property espousable? At this writing, these questions appear to be open.

Exercises

1. Show that \aleph_0 is espousable.
2. Show that \aleph_1 is not espousable.
3. Complete the proof of Theorem 3.45 by generalizing the argument to other successor cardinals.
4. Complete the proof of Theorem 3.46 by generalizing the argument to other singular cardinals.
5. Fill in the details of the proof of Theorem 3.47.

3.8.5 Perfect Matchings

I actually found someone for me.

— Whitney Houston

A society is called *symmetrically espousable* if it has an espousal that maps the men *onto* the women. What a difference one word makes! Requiring the espousal to be onto insures that every person in the society will have a spouse, without regard to their gender. If we view societies as bipartite graphs, an onto espousal gives a perfect matching between the two sets of vertices.

A bit more notation makes it easy to formulate symmetric marriage theorems. A *symmetric society* is a quadruple $(M, W, \overline{W}, \overline{M})$ such that (M, W) is a society consisting of the men and their lists, $(\overline{W}, \overline{M})$ is a society consisting of the women and their lists, and the lists match appropriately. More precisely, we require that a woman is on a man's list if and only if he is on her list, so $w \in W(m)$ if and only if $m \in \overline{M}(w)$. We can now formulate and prove a symmetric version of Theorem 3.41.

Theorem 3.48. *Suppose $(M, W, \overline{W}, \overline{M})$ is a symmetric society in which each person has a finite list. More formally, for every $m \in M$ and $w \in \overline{W}$, we have $|W(m)| < \aleph_0$ and $|\overline{M}(w)| < \aleph_0$. Then the following are equivalent:*

1. *For every $S \subset M$, $|S| \leq |W(S)|$ and for every $T \subset \overline{W}$, $|T| \leq |\overline{M}(T)|$.*
2. *$(M, W, \overline{W}, \overline{M})$ is symmetrically espousable.*

Proof. To prove item 2 implies item 1, it suffices to note that if $(M, W, \overline{W}, \overline{M})$ is symmetrically espousable, then (M, W) and $(\overline{W}, \overline{M})$ are espousable. Applying Theorem 3.41 twice yields item 1.

To prove the converse assuming item 1, Theorem 3.41 says that the societies (M, W) and $(\overline{W}, \overline{M})$ have espousals. Thus we have one-to-one functions from the men into the women and from the women into the men. The statement of the Cantor–Bernstein Theorem (Theorem 3.20) says that we must have a bijection between the men and the women. This does not quite finish the proof, since an arbitrary bijection might not map a person to someone on their list. However, the *proof* of Theorem 3.20 shows that this bijection can be constructed using only values from the initial injections. In our setting, suppose $f : M \rightarrow \overline{W}$ is an espousal of the men, $g : \overline{W} \rightarrow M$ is an espousal of the women, and $h : M \rightarrow \overline{W}$ is the bijection constructed as in the proof of Theorem 3.20. Then $h(m) = w$ implies that either $f(m) = w$ or $g(w) = m$. Thus, either w is on m 's list or m is on w 's list. By the symmetry of the society, each must be on the other's list. Thus h is a one-to-one and onto map between the men and the women that matches each person with someone on their list. \square

The preceding theorem and proof can be adapted to formulate and prove new symmetric versions of all of our marriage theorems. We leave these results as exercises.

Exercises

1. Formulate and prove a symmetric version of the finite marriage theorem, Theorem 3.40.
2. Formulate and prove a symmetric version of the theorem of Damerell and Milner, Theorem 3.42.
3. Formulate and prove a symmetric version of the countable omnibus marriage theorem, Theorem 3.43.
4. Formulate and prove a symmetric version of the uncountable marriage theorem, Theorem 3.44.
5. Define symmetrically espousable cardinals. Using your definition, is there a cardinal that is espousable but not symmetrically espousable?

3.9 Finite Combinatorics with Infinite Consequences

Does mathematics need new axioms?

— Solomon Feferman [104]

In this section we will discuss a remarkable result due to H. Friedman [113]. Friedman has concocted a finite combinatorial statement that he calls Proposition B. Under the appropriate consistency assumptions, he shows the following:

1. For each $k \in \omega$, the axiom system consisting of ZFC plus “there is a k -subtle cardinal” *cannot* prove Proposition B.
2. The axiom system consisting of ZFC plus “for every $k \in \omega$ there is a k -subtle cardinal” *can* prove Proposition B.

Thus, the use of large cardinals is *required* in the proof of the finite combinatorial statement of Proposition B.

Proposition B is not simple, but it can be understood with a modest effort. We will start with the full statement, and then gradually define any mysterious terminology.

Proposition B. *For every $k > 0$ and $p > 0$, there is an integer n such that if $\{f_X \mid X \subset [n]^k\}$ is any $\#$ -decreasing function assignment for $[n]^k$, then we can find sets A and E satisfying*

- E is subset of $\{0, 1, 2, \dots, n\}$ with p elements,
- $[E]^k \subset A \subset [n]^k$, and
- f_A has no more than k^k regressive values on $[E]^k$.

In the statement of the proposition, p , n , and k represent natural numbers. Here, the notation $[n]^k$ denotes the set of all ordered k -tuples with coordinates in $\{0, 1, 2, \dots, n - 1\}$.

If x is a k -tuple, then $|x|$ is the maximum coordinate of x and $\min(x)$ is the minimum coordinate. For example, if $x = (1, 3, 0)$, then $|x| = 3$ and $\min(x) = 0$. If f is a function mapping k -tuples to k -tuples, we say that y is a *regressive value* for f on $[E]^k$ if for some $x \in [E]^k$, $y = f(x)$ and $|y| < \min(x)$. For example, if $f(7, 6, 5) = (1, 3, 2)$, then $(1, 3, 2)$ is a regressive value for f because $3 < 5$. A single regressive value may have many witnesses. The function $g : [8]^3 \rightarrow [8]^3$ defined by $g(x, y, z) = (0, 0, 0)$ for all $0 \leq x, y, z < 8$ has only one regressive value, namely $(0, 0, 0)$.

A function assignment on $[n]^k$ assigns one function to each nonempty subset $X \subset [n]^k$. The function assigned to X is required to map X to X , so we always have $f_X : X \rightarrow X$. We say that the function assignment $\{f_X \mid X \subset [n]^k\}$ is $\#$ -decreasing if whenever $A \subset [n]^k$ and $x \in [n]^k$, either

- $f_{A \cup \{x\}}$ and f_A agree on all elements of A , or
- there is a y such that $|y| > |x|$ and $|f_A(y)| > |f_{A \cup \{x\}}(y)|$.

Some sense of the nature of $\#$ -decreasing function assignments can be gained from looking at a very small example. Suppose we set $k = 1$ and consider the possible function assignments for $[2]^1$. Since we are required to set $f_{\{0\}}(0) = 0$ and $f_{\{1\}}(1) = 1$, every function assignment on $[2]^1$ is completely determined by the values of $f_{\{0,1\}}$. We will look at all four possible cases.

Case 01: Suppose $f_{\{0,1\}}(0) = 0$ and $f_{\{0,1\}}(1) = 1$. Since $f_{\{0,1\}}$ extends both $f_{\{0\}}$ and $f_{\{1\}}$, this is a $\#$ -decreasing function assignment.

Case 00: Suppose $f_{\{0,1\}}(0) = 0$ and $f_{\{0,1\}}(1) = 0$. Since $f_{\{0,1\}}$ extends $f_{\{0\}}$ and $f_{\{1\}}(1) = 1 > 0 = f_{\{1\} \cup \{0\}}(1)$, this is also a $\#$ -decreasing function assignment.

Case 10: Suppose $f_{\{0,1\}}(0) = 1$ and $f_{\{0,1\}}(1) = 0$. This is an acceptable function assignment, but it is not $\#$ -decreasing. If we let $A = \{0\}$ and set $x = 1$, then $f_{A \cup \{x\}}(0) = 1 \neq 0 = f_A(0)$, so the first clause of the definition of $\#$ -decreasing fails. Because $x = 1$, there is no $y \in A$ such that $|y| > |x|$, so the second clause fails also.

Case 11: Suppose $f_{\{0,1\}}(0) = 1$ and $f_{\{0,1\}}(1) = 1$. This is an acceptable function assignment. Because $f_{\{0,1\}}(0) = 1$, imitating the argument in the preceding case will show that this function assignment is not $\#$ -decreasing.

In general, the second clause of the definition of $\#$ -decreasing function assignment forces the values of the functions to be pushed down. One might think that this would always lead to a proliferation of regressive values. But Proposition B asserts that if we start with a large enough domain, then even when the function assignment is $\#$ -decreasing there will be a function on a large subset that is not regressive at too many values.

Proposition B is a remarkable statement. Perhaps it is too remarkable to be true. It is important to remember that proving the proposition inherently requires large cardinal assumptions that are well beyond the strength of the axioms of ZFC. Should we automatically tack these large cardinal axioms onto our collection of everyday set theory axioms? Feferman [104] would consider this rash. Does Friedman’s result offer novel insights into the consequences of assuming large cardinal axioms? Absolutely.

Exercises

1. Let I denote the function assignment on $[n]^k$ defined by setting $f_A(x) = x$ for every A and every $x \in A$. Show that I is a $\#$ -decreasing function assignment.
2. Let M denote the function assignment on $[n]^1$ defined by setting $f_A(x) = \min(A)$ for every A and every $x \in A$.
 - (a) Show that M is a $\#$ -decreasing function assignment.
 - (b) If $f_A \in M$, what is the maximum number of regressive values that f_A can have? What is the maximum number of witnesses that a regressive value for f_A can have?
3. Let S denote the function assignment on $[n]^1$ defined by setting

$$f_A(x) = \begin{cases} x & \text{if } x < |A|, \\ \min(A) & \text{if } x = |A|, \end{cases}$$

for every A and every $x \in A$. Is S a $\#$ -decreasing function assignment?

4. Find a function assignment on $[3]^1$ that is not $\#$ -decreasing. If you would like a challenge, find all of them.

3.10 k -critical Linear Orderings

It appears that Feferman is using the word ‘need’ in a sense that requires discussion.

— Harvey Friedman [105]

In [114], H. Friedman reveals the connection between k -subtle cardinals and a very natural Ramsey-style property on linear orderings. These new results are not expressions of finite combinatorics like those of Section 3.9, but they have the advantage of requiring less terminology.

Suppose that $\langle X, \leq \rangle$ is a linear ordering. As you may recall from Section 3.5.2, this means that the relation \leq is antisymmetric, transitive, and satisfies the trichotomy law. We say that X has *no endpoints* if it has neither a maximum nor a minimum element. The ordinal ω is certainly a linear ordering, with 0 as an endpoint. The integers are a familiar example of a linear ordering without endpoints, as are the rational numbers and the real numbers.

Suppose that $\langle X, \leq \rangle$ is a linear ordering without endpoints and that $k > 0$ is a natural number. We say that $f : [X]^k \rightarrow X$ is a *regressive* function if for every increasing k -tuple $x_1 < x_2 < \dots < x_k$, we have $f(x_1, x_2, \dots, x_k) < x_1$. A linear ordering is *k -critical* if for every regressive function $f : [X]^k \rightarrow X$ there is a sequence of elements $b_1 < \dots < b_{k+1}$ such that $f(b_1, \dots, b_k) = f(b_2, \dots, b_{k+1})$.

This is all the machinery that we need to state (some of) Friedman’s results! Here are two of his theorems.

Theorem 3.49. *For each nonzero $k \in \omega$, ZFC proves that κ is the least cardinality of a $(k + 1)$ -critical linear ordering if and only if κ is the least k -subtle cardinal.*

Theorem 3.50. *ZFC proves that there is a k -critical linear ordering for every nonzero $k \in \omega$ if and only if there is a k -subtle cardinal for every nonzero $k \in \omega$.*

For proofs of these theorems and many related results, see [114]. It is worth noting that Theorem 3.49 could be reformulated without using the word “least.” Working in ZFC, suppose there is a $(k + 1)$ -critical linear ordering X . Let $\lambda = |X|$. Applying the axiom of separation we can form the subset S of those cardinals $\alpha \leq \lambda$ such that there is a $(k + 1)$ -critical linear ordering Y with $|Y| = \alpha$. Since S is a set of cardinals, it has a least element; call it κ . Then κ is the least cardinality of a $(k + 1)$ -critical linear ordering. Summarizing, if there is a $(k + 1)$ -critical linear ordering, then there is a least cardinality of a $(k + 1)$ -critical linear ordering. The converse of this implication is obvious. A similar argument based on the definition of subtle cardinals shows that there is a k -subtle cardinal if and only if there is a least k -subtle cardinal. Consequently, Theorem 3.49 could be stated as “there is a $(k + 1)$ -critical linear ordering if and only if there is a k -subtle cardinal.”

Is it reasonable to assert that k -critical linear orderings exist? Familiar linear orderings like the integers, the rationals, and the reals, are much too small to be even 2-critical. (See the exercises.) The smallest 2-critical linear ordering must be larger than some weakly compact cardinal. On the other hand, the least Ramsey cardinal is larger than some k -critical linear ordering for every k . If we believe that k -subtle cardinals exist, then we must also accept the existence of k -critical linear orderings. If we reject the existence of weakly compact cardinals, then we must also reject the existence of k -critical linear orderings. The existence of these objects is independent of ZFC (assuming ZFC is consistent), so we are back to deciding whether or not to add new axioms.

We can be certain about one thing. Friedman's theorems give us a characterization of k -subtle cardinals that does not mention (n, κ) -sequences or club sets. This elegant description of k -subtle cardinals could be a help in discussions about their existence.

Exercises

1. Let $\mathbb{Z} = \{\dots, -2, -1, 0, 1, 2, \dots\}$ denote the set of all integers and define $f : \mathbb{Z} \rightarrow \mathbb{Z}$ by $f(x) = x - 1$. Show that f is regressive.
2. Prove that \mathbb{Z} is not 2-critical. (Hint: Problem 1 is useful.)
3. Let \mathbb{Q} denote the rational numbers. Show that \mathbb{Q} is not 2-critical.
4. Let \mathbb{R} denote the real numbers. Show that \mathbb{R} is not 2-critical.
5. Prove that no countable linear ordering with no endpoints is 2-critical. (Hint: Any such linear ordering has countably many pairs and also has a countable unbounded descending sequence. These can be used to build the needed regressive function.)
6. Suppose that $\kappa \geq \aleph_0$ and κ is regular. Prove that if L is a linear ordering of cardinality κ containing an unbounded descending sequence of size κ , then L is not 2-critical.
7. (This one is very challenging.) Without using Friedman's theorem, try to prove that if κ is a Ramsey cardinal then there is a 2-critical linear ordering of size κ .

3.11 Points of Departure

Where do you want to go today? [®]

— Microsoft

This chapter can be summarized as a study of pigeonhole principles, König's Lemma, and Ramsey's Theorem and their connections to cardinal numbers. Our

slightly obsessive focus on cardinals is not the only approach to studying infinite objects. In this section we take brief peeks at computability, reverse mathematics, and the analytical hierarchy. For each topic area, we outline an application to graph theory and combinatorics, and state an open question. The section closes with a list of Ramsey-style theorems and named trees.

Computability

We say that a function $f : \mathbb{N} \rightarrow \mathbb{N}$ is *computable* if there is a program written in \mathcal{C} that given an input of n always halts with output $f(n)$. Any programming language can be substituted for \mathcal{C} . A set $M \subset \mathbb{N}$ is *computable* if its characteristic function χ_M is a computable function. The function χ_M is defined by setting $\chi_M(n) = 1$ if $n \in M$ and $\chi_M(n) = 0$ if $n \notin M$. Our definition extends in the obvious way to handle functions on \mathbb{N}^k and subsets of \mathbb{N}^k . The computable functions are sometimes called Turing computable functions or general recursive functions. The study of these functions is called computability theory or recursion theory. Good introductory texts include [31] and [259], and [120] is a survey of computability in graph theory and combinatorics.

Every computable function has a program. We can think of this program as an integer; it certainly would be stored in a computer as a string of zeros and ones that could be thought of as a single big integer. We use $\{m\}(n) = k$ as shorthand for saying that when the machine with code m receives input n , it halts with output k . We frequently refer to these codes as *indices* for computable functions.

Not every subset of \mathbb{N} is computable. For example, consider the set

$$H = \{m \in \mathbb{N} \mid \{m\}(m) = 0\}.$$

To show that H is not computable, we consider χ_H and prove the following theorem.

Theorem 3.51. χ_H is not computable.

Proof. Suppose that χ_H is computable. Let m be the code for a program that computes χ_H . Then

$$\begin{aligned} \chi_H(m) = 0 &\leftrightarrow \{m\}(m) = 0 \\ &\leftrightarrow m \in H \\ &\leftrightarrow \chi_H(m) = 1. \end{aligned}$$

Since $0 \neq 1$, χ_H is not computable. □

Because computer programs process in discrete steps, we can examine computations at various stages. We write $\{m\}(n) \downarrow_s$ if the program with index m and input n halts in fewer than s steps. We also write $\{m\}(n) \downarrow$ if the program eventually halts. It is easy to mechanically check whether a program halts in s steps, but impossible to mechanically determine which programs will eventually halt. Indeed, the set H is referred to as the solution to the self-halting problem.

Using our extended shorthand, we can encode noncomputable sets like H in graph theory problems. For example, suppose we define a coloring F on $[\mathbb{N}]^3$ by the following rule. Let $\{i, j, k\}$ be any three-element subset of \mathbb{N} , assume that we have $i < j < k$, and let

$$F(i, j, k) = \begin{cases} \text{red} & (\forall m \leq i) \{m\}(m) \downarrow_j \leftrightarrow \{m\}(m) \downarrow_k, \\ \text{blue} & \text{otherwise.} \end{cases}$$

Thus $F(i, j, k)$ is red when for every machine with a code $m \leq i$, we get the same information about whether or not $\{m\}(m)$ halts by checking j steps as we do by checking k steps. If $F(i, j, k)$ is blue, then for some $m \leq i$ the machine $\{m\}(m)$ halts between step j and step k . By Ramsey's Theorem, F must have an infinite monochromatic set; call it $T = \{t_0, t_1, t_2, \dots\}$ and assume that we have listed the elements in increasing order. No matter how we select three elements from T , applying F should always give us the same color. Since there are only finitely many machines with codes less than t_0 , we cannot have one of them stopping between t_i steps and t_{i+1} steps for every $i > 1$. (Smell any pigeons?) Thus T must be a red monochromatic set. The elements of T are very handy for computing the set H . Pick any machine m . We know that $m \leq t_m$, and if $\{m\}(m)$ ever halts, it must halt by step t_{m+1} . To see this, suppose that $\{m\}(m)$ halts at a later step k . Then $t_{k+1} > k$ and $F(t_m, t_{m+1}, t_{k+1})$ would be blue, contradicting the claim that T is a monochromatic red set. Thus, using χ_T as a subprogram, we could easily write a program that computes χ_H . Since χ_H is not computable, χ_T is not computable either, and we have proved the following theorem.

Theorem 3.52. *There is a computable coloring of $[\mathbb{N}]^3$ that has no infinite computable monochromatic set.*

Many other types of coloring problems lead to noncomputable sets. For example, there is a computable 2-regular 2-colorable graph with no computable 2-coloring. However, bumping up the size of our color palette can lead to computable colorings. Schmerl [248] proved that if G is a computable d -regular graph that is n -colorable, then G has a computable $2n - 1$ coloring. Depending on the value of d , it may or may not be possible to find a computable coloring with fewer than $2n - 1$ colors. For $2 \leq n \leq m \leq 2n - 2$, let $D(n, m)$ be the least degree d such that there is a d -regular n -colorable graph with no recursive m coloring. Beyond a few bounds (see [120] and [248]), little is known about the values of $D(n, m)$. For example, we know that $6 \leq D(5, 5) \leq 7$, but the exact value of $D(5, 5)$ is not known.

Reverse Mathematics

Reverse mathematics is a program of mathematical logic that was founded by H. Friedman (as in Section 3.9) and S. Simpson. The goal of the program is to measure the logical strength of mathematical theorems by proving that each one is equivalent to some statement in a hierarchy of axioms for second order arithmetic.

These proofs are carried out in a weak base system, RCA_0 , which consists of PA with restricted induction plus a comprehension axiom that essentially asserts the existence of computable subsets of \mathbb{N} . We refer to these systems as *second order arithmetic* because the formulas include variables for numbers and variables for sets of numbers. Simpson's book [254] is a comprehensive source for information about reverse mathematics. A number of theorems about infinite graphs have been analyzed, including the following example.

Theorem 3.53. *RCA_0 can prove that the following are equivalent:*

1. *Every 2-regular graph with no cycles of odd length is bipartite.*
2. *Weak König's Lemma: Every infinite tree in which every node is labeled 0 or 1 contains an infinite path.*

Since a graph is bipartite if and only if it is 2-colorable, it is not hard to adapt the solution of Exercise 3 of Section 3.1 to prove that 2 implies 1. The proof that 1 implies 2 is less obvious, and the published proof uses an intermediate equivalent statement [162].

There are many open questions in reverse mathematics that are related to combinatorics and graph theory. For example, can RCA_0 prove that Ramsey's Theorem for pairs and two colors implies Weak König's Lemma? It is known that the converse is independent of RCA_0 and that Ramsey's Theorem for triples and two colors implies a much stronger version of König's Lemma.

The Analytical Hierarchy

Subsets of \mathbb{N} that are definable by formulas of second order arithmetic are called *analytical*, and can be organized by the complexity of their defining formulas. The resulting hierarchy of sets is analogous (but not identical) to the hierarchy of projective sets studied by descriptive set theorists. Both [240] and [161] provide useful background for the study of the analytical hierarchy.

To state any results, we need some terminology. A set is Σ_1^1 -*definable* if it is definable by a formula containing no universal set quantifiers. We say that a set is Σ_1^1 -*complete* if it is Σ_1^1 -definable and every other Σ_1^1 -definable set is 1-reducible to it. Here, B is 1-*reducible* to A if there is a computable one-to-one function f such that for all n , $n \in B$ if and only if $f(n) \in A$. Naïvely, each Σ_1^1 -complete set embodies all the information content of every other Σ_1^1 -definable set. One interesting characteristic of these sets is that any defining formula for a Σ_1^1 -complete set must contain a set quantifier. Thus, no formula containing only quantifiers on numbers can possibly define a Σ_1^1 -complete set.

An infinite graph G has a *Hamiltonian path* if there is a sequence of vertices v_0, v_1, v_2, \dots such that each vertex of G appears exactly once in the list and $v_j v_{j+1}$ is an edge of G for each $j \in \mathbb{N}$. An index for a computable graph is the code for the characteristic function of the edge and vertex sets for the graph. Using all this terminology, we can state the following theorem of Harel [153].

Theorem 3.54. *The set of indices of computable graphs with Hamiltonian paths is Σ_1^1 -complete.*

For a similar problem, suppose we weight the edges of an infinite graph with rational numbers that sum to 1. Such a graph may or may not have a minimal spanning tree. Is the set of indices of computable graphs with minimal spanning trees a Σ_1^1 -complete set? This question was motivated by a related reverse mathematics result for directed graphs [57], and the answer was unknown when the first edition of this book was published. Since that time, Schmerl [249] has shown that the answer is no, the set of indices of computable graphs with minimal spanning trees can be defined by an arithmetical formula.

Lists of Theorems and Trees

There are a number of interesting Ramsey-style theorems that are worthy of exploration. Here is a list, with pointers to some good references.

- Van der Waerden's Theorem (and the related Szemerédi's Theorem) [136]
- Hindman's Theorem (and the related Folkman's Theorem) [136]
- Milliken's Theorem [205]
- Galvin–Prikry Theorem [119]
- Erdős–Rado Theorem [168]

Other good prospects for further study include special and regular Aronszajn trees, Suslin trees, and Kurepa trees, all appearing in [168] and [187].

3.12 References

His was a name to conjure with in certain circles.

— E. Wallace, *The Just Men of Cordova*

In this section we will name a few authors whose books will be helpful to those interested in the further study of infinite sets, graphs, and combinatorics. Many good introductory texts on axiomatic set theory are available. For example, the books of Enderton [88], Moschovakis [209], and Roitman [241] are all accessible and remarkably distinct. For a more technical approach, Drake [80] and Levy [187] are good choices. Jech's encyclopedic text [168] is an invaluable reference (and a good read).

Several books give extended treatments of some specific topics in this chapter. Devlin's book [73] gives a comprehensive coverage of constructible sets. Large cardinals and partition cardinals play a central role in Drake's text [80]. *The Axiom of Choice* is the title and subject of another nice book by Jech [167]. For a detailed treatment of cardinal and ordinal arithmetic, it is hard to beat Sierpiński's old gem

[253]. It is very interesting to compare Sierpiński's development with that of the ante-ZFC papers of Cantor, [47] and [48], both translated in [49].

Much of the historical content of the chapter was gleaned from van Heijenoort's anthology [272] and Kanamori's insightful article [173]. For the final word in many matters of logic, one can consult Kleene's blue bible [176] or the more accessible text of Mendelson [201]. Finally, other treatments of infinite graphs and combinatorics can be found in the books of Ore [218] and Diestel [74].

References

- [1] R. Aharoni, C. St. J. A. Nash-Williams, and S. Shelah, *A general criterion for the existence of transversals*, Proc. London Math. Soc. (3) **47** (1983), no. 1, 43–68.
- [2] ———, *Marriage in infinite societies*, Progress in Graph Theory (Waterloo, ON, 1982) (J. Bondy and U. Murty, eds.), Academic Press, Toronto, ON, 1984, pp. 71–79.
- [3] M. Aigner, *Graph Theory: A Development from the 4-Color Problem*, BCS Associates, Moscow, ID, 1987.
- [4] ———, *Combinatorial Theory*, Springer-Verlag, Berlin, 1997.
- [5] ———, *A Course in Enumeration*, Grad. Texts in Math., vol. 238, Springer-Verlag, New York, 2007.
- [6] M. Aigner and G. M. Ziegler, *Proofs from The Book*, 3rd ed., Springer-Verlag, Berlin, 2004.
- [7] I. Anderson, *Perfect matchings of a graph*, J. Combin. Theory Ser. B **10** (1971), no. 3, 183–186.
- [8] G. E. Andrews, *Euler’s pentagonal number theorem*, Math. Mag. **56** (1983), no. 5, 279–284.
- [9] ———, *The Theory of Partitions*, Cambridge Univ. Press, Cambridge, 1998.
- [10] G. E. Andrews and K. Eriksson, *Integer Partitions*, Cambridge Univ. Press, Cambridge, 2004.
- [11] T. M. Apostol, *Mathematical Analysis*, 2nd ed., Addison-Wesley, Reading, MA, 1974.
- [12] K. Appel and W. Haken, *Every planar map is four colorable*, Illinois J. Math. **21** (1977), no. 3, 429–567.
- [13] D. Babić, D. J. Klein, J. von Knop, and N. Trinajstić, *Combinatorial enumeration in chemistry*, Chemical Modelling: Applications and Theory (A. Hinchliffe, ed.), Vol. 3, Royal Society of Chemistry, London, 2004, pp. 126–170.
- [14] M. Baiou and M. Balinski, *Many-to-many matching: Stable polyandrous polygamy (or polygamous polyandry)*, Discrete Appl. Math. **101** (2000), no. 1, 1–12.
- [15] D. C. Banks, S. A. Linton, and P. Stockmeyer, *Counting cases in substitute algorithms*, IEEE Trans. Vis. Comput. Graphics **10** (2004), no. 4, 371–384.

- [16] D. C. Banks and P. Stockmeyer, *DeBruijn counting for visualization algorithms*, Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration (T. Möller, B. Hamann, and R. Russell, eds.), Math. Vis., Springer-Verlag, New York, 2009. To appear.
- [17] A.-L. Barabási, *Linked*, Perseus Publishing, Cambridge, MA, 2002.
- [18] J. E. Baumgartner, *Ineffability properties of cardinals I*, Infinite and Finite Sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday), Vol. I (A. Hajnal, R. Rado, and V. Sós, eds.), Colloq. Math. Soc. János Bolyai, vol. 10, North-Holland, Amsterdam, 1975, pp. 109–130.
- [19] ———, *Ineffability properties of cardinals II*, Fifth Internat. Congr. Logic, Methodology and Philos. of Sci., Part I (Univ. Western Ontario, London, ON, 1975), Logic, Foundations of Mathematics and Computability Theory (R. Butts and J. Hintikka, eds.), Univ. Western Ontario Ser. Philos. Sci., vol. 9, Reidel, Dordrecht, 1977, pp. 87–106.
- [20] P. Bedrossian, *Forbidden subgraph and minimum degree conditions for Hamiltonicity*, Ph.D. Thesis, Memphis State Univ., 1991.
- [21] L. W. Beineke and R. J. Wilson (eds.), *Graph Connections*, Oxford Lecture Ser. Math. Appl., vol. 5, Oxford Univ. Press, New York, 1997.
- [22] A. T. Benjamin and J. J. Quinn, *Proofs that Really Count: The Art of Combinatorial Proof*, Dolciani Math. Exp., vol. 27, Math. Assoc. America, Washington, DC, 2003.
- [23] C. Berge, *Two theorems in graph theory*, Proc. Natl. Acad. Sci. USA **43** (1957), 842–844.
- [24] ———, *Principles of Combinatorics*, Math. Sci. Engrg., vol. 72, Academic Press, New York, 1971.
- [25] H. Bielak and M. M. Sysło, *Peripheral vertices in graphs*, Studia Sci. Math. Hungar. **18** (1983), no. 2-4, 269–275.
- [26] N. L. Biggs, E. K. Lloyd, and R. J. Wilson, *Graph Theory: 1736–1936*, Oxford Univ. Press, 1976.
- [27] G. D. Birkhoff and D. C. Lewis, *Chromatic polynomials*, Trans. Amer. Math. Soc. **60** (1946), no. 3, 355–451.
- [28] A. Bloch and G. Pólya, *On the roots of certain algebraic equations*, Proc. London Math. Soc. (2) **33** (1932), no. 1, 102–114.
- [29] B. Bollobás, *Modern Graph Theory*, Grad. Texts in Math., vol. 184, Springer-Verlag, New York, 1998.
- [30] W. E. Bonnice, *On convex polygons determined by a finite planar set*, Amer. Math. Monthly **81** (1974), no. 7, 749–752.
- [31] G. S. Boolos and R. C. Jeffrey, *Computability and Logic*, 3rd ed., Cambridge Univ. Press, Cambridge, 1989.
- [32] E. Borel, *Leçons sur la Théorie des Fonctions*, Gauthier-Villars et fils, Paris, 1898.
- [33] P. Borwein, *Sylvester’s problem and Motzkin’s theorem for countable and compact sets*, Proc. Amer. Math. Soc. **90** (1984), no. 4, 580–584.
- [34] P. Borwein and W. O. J. Moser, *A survey of Sylvester’s problem and its generalizations*, Aequationes Math. **40** (1990), 111–135.
- [35] P. Borwein and M. J. Mossinghoff, *Polynomials with height 1 and prescribed vanishing at 1*, Experiment. Math. **9** (2000), no. 3, 425–433.
- [36] ———, *Newman polynomials with prescribed vanishing and integer sets with distinct subset sums*, Math. Comp. **72** (2003), no. 242, 787–800.

- [37] P. Brass, W. Moser, and J. Pach, *Research Problems in Discrete Geometry*, Springer-Verlag, New York, 2005.
- [38] I. Broere, M. Dorfling, J. E. Dunbar, and M. Frick, *A path(ological) partition problem*, *Discuss. Math. Graph Theory* **18** (1998), no. 1, 113–125.
- [39] H. Broersma and H. J. Veldman, *Restrictions on induced subgraphs ensuring Hamiltonicity or pancyclicity of $K_{1,3}$ -free graphs*, *Contemporary Methods in Graph Theory* (R. Bodendiek, ed.), BI Wissenschaftsverlag, Mannheim, 1990, pp. 181–194.
- [40] J. G. Broida and S. G. Williamson, *A Comprehensive Introduction to Linear Algebra*, Addison-Wesley, Redwood City, CA, 1989.
- [41] R. L. Brooks, *On colouring the nodes of a network*, *Proc. Cambridge Philos. Soc.* **37** (1941), 194–197.
- [42] F. Buckley and F. Harary, *Distance in Graphs*, Addison-Wesley, Redwood City, CA, 1990.
- [43] F. Buckley and M. Lewinter, *A Friendly Introduction to Graph Theory*, Pearson Education, Upper Saddle River, NJ, 2003.
- [44] F. Buckley, Z. Miller, and P. J. Slater, *On graphs containing a given graph as center*, *J. Graph Theory* **5** (1981), no. 4, 427–434.
- [45] W. Burnside, *Theory of Groups of Finite Order*, Cambridge Univ. Press, London, 1897. 2nd ed., Cambridge Univ. Press, London, 1911, reprinted by Dover, New York, 1955.
- [46] S. A. Burr, *Generalized Ramsey theory for graphs—A survey*, *Capital Conference on Graph Theory and Combinatorics* (George Washington Univ., Washington, DC, 1973), *Graphs and Combinatorics* (R. A. Bari and F. Harary, eds.), *Lecture Notes in Math.*, vol. 406, Springer-Verlag, Berlin, 1974, pp. 52–75.
- [47] G. Cantor, *Beiträge zur Begründung der transfiniten Mengenlehre*, *Math. Ann.* **46** (1895), no. 4, 481–512; English transl., G. Cantor, *Contributions to the Founding of the Theory of Transfinite Numbers* (1955), 85–136.
- [48] ———, *Beiträge zur Begründung der transfiniten Mengenlehre II*, *Math. Ann.* **49** (1897), no. 2, 207–246; English transl., ———, *Contributions to the Founding of the Theory of Transfinite Numbers* (1955), 137–201.
- [49] ———, *Contributions to the Founding of the Theory of Transfinite Numbers*, Dover, New York, 1955. Translated by P. E. B. Jourdain.
- [50] A. Cayley, *A theorem on trees*, *Quart. J. Math.* **23** (1889), 276–378.
- [51] D. Chakerian and D. Logothetti, *Cube slices, pictorial triangles, and probability*, *Math. Mag.* **64** (1991), no. 4, 219–241.
- [52] G. Chartrand and L. Lesniak, *Graphs & Digraphs*, 4th ed., Chapman & Hall, Boca Raton, FL, 2005.
- [53] V. Chvátal, *Tree-complete graph Ramsey numbers*, *J. Graph Theory* **1** (1977), no. 1, 93.
- [54] V. Chvátal and P. Erdős, *A note on Hamiltonian circuits*, *Discrete Math.* **2** (1972), no. 2, 111–113.
- [55] V. Chvátal and F. Harary, *Generalized Ramsey theory for graphs III: Small off-diagonal numbers*, *Pacific J. Math.* **41** (1972), no. 2, 335–345.
- [56] J. Clark and D. A. Holton, *A First Look at Graph Theory*, World Scientific, Teaneck, NJ, 1991.
- [57] P. G. Clote and J. L. Hirst, *Reverse mathematics of some topics from algorithmic graph theory*, *Fund. Math.* **157** (1998), no. 1, 1–13.
- [58] P. J. Cohen, *The independence of the continuum hypothesis*, *Proc. Natl. Acad. Sci. USA* **50** (1963), 1143–1148.

- [59] ———, *The independence of the continuum hypothesis II*, Proc. Natl. Acad. Sci. USA **51** (1964), 105–110.
- [60] L. Comtet, *Advanced Combinatorics: The Art of Finite and Infinite Expansions*, D. Reidel, Dordrecht, 1974.
- [61] J. H. Conway and R. K. Guy, *The Book of Numbers*, Springer-Verlag, New York, 1996.
- [62] H. S. M. Coxeter, *A problem of collinear points*, Amer. Math. Monthly **55** (1948), no. 1, 26–28.
- [63] G. P. Csicsery (director), *N is a Number: A Portrait of Paul Erdős*, Zala Films, Oakland, CA, 1993. Documentary film.
- [64] J. Csima and E. T. Sawyer, *There exist $6n/13$ ordinary points*, Discrete Comput. Geom. **9** (1993), no. 2, 187–202.
- [65] ———, *The $6n/13$ theorem revisited*, Graph Theory, Combinatorics, and Algorithms: Proceedings of the Seventh Quadrennial International Conference on the Theory and Applications of Graphs (Western Michigan Univ., Kalamazoo, MI, 1992) (Y. Alavi and A. Schwenk, eds.), Wiley-Intersci. Publ., vol. 1, John Wiley & Sons, New York, 1995, pp. 235–249.
- [66] R. M. Damerell and E. C. Milner, *Necessary and sufficient conditions for transversals of countable set systems*, J. Combin. Theory Ser. A **17** (1974), no. 3, 350–374.
- [67] J. F. Davidson, *The Anacreontea & Principal Remains of Anacreon of Teos, in English verse*, E. P. Dutton, New York, 1915.
- [68] N. G. de Bruijn, *Pólya's theory of counting*, Applied Combinatorial Mathematics (E. F. Beckenbach, ed.), John Wiley & Sons, New York, 1964, pp. 144–184.
- [69] ———, *Color patterns that are invariant under a given permutation of the colors*, J. Combin. Theory **2** (1967), 418–421.
- [70] ———, *A survey of generalizations of Pólya's enumeration theorem*, Nieuw Arch. Wisk. (3) **19** (1971), 89–112.
- [71] S. de Wannemacker, T. Laffey, and R. Osburn, *On a conjecture of Wilf*, J. Combin. Theory Ser. A **114** (2007), no. 7, 1332–1349.
- [72] W. Degen, *Pigeonhole and choice principles*, MLQ Math. Log. Q. **46** (2000), no. 3, 313–334.
- [73] K. J. Devlin, *Constructibility*, Perspect. Math. Logic, Springer-Verlag, Berlin, 1984.
- [74] R. Diestel, *Graph Decompositions*, Oxford Sci. Publ., Oxford Univ. Press, New York, 1990.
- [75] ———, *Graph Theory*, 3rd ed., Grad. Texts in Math., vol. 173, Springer-Verlag, Berlin, 2005.
- [76] E. W. Dijkstra, *A computing scientist's approach to a once-deep theorem of Sylvester's* (1988), 14 pp. Manuscript number EWD1016, available at the *E. W. Dijkstra Archive*, <http://www.cs.utexas.edu/users/EWD>.
- [77] G. A. Dirac, *Some theorems on abstract graphs*, Proc. London Math. Soc. (3) **2** (1952), no. 1, 69–81.
- [78] G. A. Dirac and S. Schuster, *A theorem of Kuratowski*, Indag. Math. **16** (1954), 343–348. Corrigendum, Indag. Math. **23** (1961), 360.
- [79] G. Dobiński, *Summirung der Reihe $\sum \frac{n^m}{n!}$ für $m = 1, 2, 3, 4, 5, \dots$* [sic], Grunert Archiv (Arch. Math. Phys.) **61** (1877), 333–336.
- [80] F. Drake, *Set Theory: An Introduction to Large Cardinals*, Stud. Logic Found. Math., vol. 76, Elsevier, Amsterdam, 1974.
- [81] D. Duffus, M. S. Jacobson, and R. J. Gould, *Forbidden subgraphs and the Hamiltonian theme*, The Theory and Applications of Graphs: Fourth International Conference (Western Michigan Univ., Kalamazoo, MI, 1980) (G. Chartrand, Y. Alavi, D. L. Goldsmith, L. Lesniak-Foster, and D. R. Lick, eds.), John Wiley & Sons, New York, 1981, pp. 297–316.

- [82] A. Dumitrescu, *A remark on the Erdős-Szekeres theorem*, Amer. Math. Monthly **112** (2005), no. 10, 921–924.
- [83] J. E. Dunbar and M. Frick, *Path kernels and partitions*, J. Combin. Math. Combin. Comput. **31** (1999), 137–149.
- [84] ———, *The path partition conjecture is true for claw-free graphs*, Discrete Math. **307** (2007), no. 11–12, 1285–1290.
- [85] A. W. F. Edwards, *Pascal's Arithmetical Triangle: The Story of a Mathematical Idea*, Johns Hopkins Univ. Press, Baltimore, 2002.
- [86] J. Eeckhout, *On the uniqueness of stable marriage matchings*, Econom. Lett. **69** (2000), no. 1, 1–8.
- [87] E. Egerváry, *Matrixok Kombinatorikus Tulajdonságairól*, Mat. Fiz. Lapok **38** (1931), 16–28.
- [88] H. B. Enderton, *Elements of Set Theory*, Academic Press, New York, 1977.
- [89] P. Erdős, *Problem 4065*, Amer. Math. Monthly **50** (1943), no. 1, 65.
- [90] P. Erdős, *Some remarks on the theory of graphs*, Bull. Amer. Math. Soc. **53** (1947), 292–294.
- [91] P. Erdős, *Combinatorial problems in geometry and number theory*, Relations Between Combinatorics and Other Parts of Mathematics (Ohio State Univ., Columbus, OH, 1978) (D. K. Ray-Chaudhuri, ed.), Proc. Sympos. Pure Math., vol. 34, Amer. Math. Soc., Providence, R.I., 1979, pp. 149–162.
- [92] ———, *Some of my favorite problems and results*, The Mathematics of Paul Erdős I (R. L. Graham and J. Nešetřil, eds.), Algorithms Combin., vol. 13, Springer-Verlag, Berlin, 1997, pp. 47–67.
- [93] P. Erdős and G. Purdy, *Extremal problems in combinatorial geometry*, Handbook of Combinatorics (R. L. Graham, M. Grötschel, and L. Lovász, eds.), Vol. 1, Elsevier Science, Amsterdam, and MIT Press, Cambridge, 1995, pp. 809–874.
- [94] P. Erdős and G. Szekeres, *A combinatorial problem in geometry*, Compositio Math. **2** (1935), 463–470; reprinted in I. Gessel and G.-C. Rota (eds.), *Classic Problems in Combinatorics*, Birkhäuser, Boston, 1987.
- [95] ———, *On some extremum problems in elementary geometry*, Ann. Univ. Sci. Budapest. Eötvös Sect. Math. **3–4** (1961), 53–62.
- [96] P. Erdős and A. Tarski, *On families of mutually exclusive sets*, Ann. of Math. (2) **44** (1943), no. 2, 315–329.
- [97] ———, *On some problems involving inaccessible cardinals*, Essays on the Foundations of Mathematics (Y. Bar-Hillel, E. Poznansik, M. Rabin, and A. Robinson, eds.), Magnes Press, Hebrew Univ., Jerusalem, 1961, pp. 50–82.
- [98] L. Euler, *Solutio problematis ad geometriam situs pertinentis*, Comment. Academiae Sci. I. Petropolitanae **8** (1736), 128–140.
- [99] ———, *Demonstratio nonnullarum insignium proprietatum quibus solida hedris planis inclusa sunt praedita*, Novi Comm. Acad. Sci. Imp. Petropol **4** (1758), 140–160.
- [100] H. Eves, *In Mathematical Circles*, Prindle, Weber & Schmidt, Boston, 1969.
- [101] I. Fáry, *On straight line representation of planar graphs*, Acta Univ. Szeged. Sect. Sci. Math. **11** (1948), 229–233.
- [102] R. J. Faudree and R. J. Gould, *Characterizing forbidden pairs for Hamiltonian properties*, Discrete Math. **173** (1997), no. 1–3, 45–60.
- [103] T. Feder, *Stable Networks and Product Graphs*, Mem. Amer. Math. Soc., vol. 116, American Math. Soc., Providence, RI, 1995.

- [104] S. Feferman, *Does mathematics need new axioms?*, Amer. Math. Monthly **106** (1999), no. 2, 99–111.
- [105] S. Feferman, H. M. Friedman, P. Maddy, and J. R. Steel, *Does mathematics need new axioms?*, Bull. Symbolic Logic **6** (2000), no. 4, 401–446.
- [106] L. R. Foulds, *Graph Theory Applications*, Universitext, Springer-Verlag, New York, 1992.
- [107] A. Fraenkel, *Über die Zermelosche Begründung der Mengenlehre*, Jahresber. Deutsch. Math.-Verein. **30** (1921), 97–98. 2nd section.
- [108] ———, *Axiomatische Begründung der transfiniten Kardinalzahlen I*, Math. Z. **13** (1922), no. 1, 153–188.
- [109] ———, *Der Begriff “definit” und die Unabhängigkeit des Auswahlaxioms*, Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse (1922), 253–257; English transl., J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931* (1967), 284–289.
- [110] ———, *Zu den Grundlagen der Cantor-Zermeloschen Mengenlehre*, Math. Ann. **86** (1922), no. 3–4, 230–237.
- [111] F. Franklin, *Sur le développement du produit infini $(1 - x)(1 - x^2)(1 - x^3) \dots$* , C. R. Acad. Sci. Paris **82** (1881), 448–450.
- [112] G. Frege, *Über die Begriffsschrift des Herrn Peano und meine eigene*, Berichte über die Verhandlungen der Königlich Sächsischen Gesellschaft der Wissenschaften zu Leipzig, Mathematisch-physikalische Klasse **48** (1896), 361–378; English transl. in J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931* (1967). A translation of the “*Summum bonum*” quotation appears on p. 2.
- [113] H. M. Friedman, *Finite functions and the necessary use of large cardinals*, Ann. of Math. (2) **148** (1998), no. 3, 803–893.
- [114] ———, *Subtle cardinals and linear orderings*, Ann. Pure Appl. Logic **107** (2001), no. 1–3, 1–34.
- [115] F. G. Frobenius, *Über die Congruenz nach einem aus zwei endlichen Gruppen gebildeten Doppelmodul*, J. Reine Angew. Math. **101** (1887), 273–299; reprinted in F. G. Frobenius, *Gesammelte Abhandlungen*, Vol. 2, Springer-Verlag, Berlin, 1968.
- [116] S. Fujita, *Symmetry and Combinatorial Enumeration in Chemistry*, Springer-Verlag, Berlin, 1991.
- [117] D. Gale and L. S. Shapley, *College admissions and the stability of marriage*, Amer. Math. Monthly **69** (1962), no. 1, 9–15.
- [118] T. Gallai, *Problem 4*, Colloquium on Graph Theory (Tihany, Hungary, 1966) (P. Erdős and G. Katona, eds.), Academic Press, New York, 1968, pp. 362.
- [119] F. Galvin and K. Prikry, *Borel sets and Ramsey’s theorem*, J. Symbolic Logic **38** (1973), 193–198.
- [120] W. Gasarch, *A survey of recursive combinatorics*, Handbook of Recursive Mathematics, Vol. 2 (Y. Ershov, S. Goncharov, A. Nerode, and J. Remmel, eds.), Stud. Logic Found. Math., vol. 139, North-Holland, Amsterdam, 1998, pp. 1041–1176.
- [121] T. Gerken, *Empty convex hexagons in planar point sets*, Discrete Comput. Geom. **39** (2008), no. 1–3, 239–272.
- [122] I. M. Gessel, *The Smith College diploma problem*, Amer. Math. Monthly **108** (2001), no. 1, 55–57.
- [123] I. Gessel and G.-C. Rota (eds.), *Classic Papers in Combinatorics*, Birkhäuser, Boston, 1987.

- [124] K. Gödel, *Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I*, *Monatsh. Math. Phys.* **38** (1931), no. 1, 173–198; English transl., K. Gödel, *Collected Works Vol. I* (1986), 144–195.
- [125] ———, *The consistency of the Axiom of Choice and the Generalized Continuum Hypothesis*, *Proc. Natl. Acad. Sci. USA* **24** (1938), 556–557. Reprinted in K. Gödel, *Collected Works Vol. I*, (S. Feferman, ed.), Oxford Univ. Press, New York, 1986, 26–27.
- [126] ———, *Collected Works Vol. I* (S. Feferman, ed.), Oxford Univ. Press, New York, 1986.
- [127] S. Goodman and S. Hedetniemi, *Sufficient conditions for a graph to be Hamiltonian*, *J. Combin. Theory Ser. B* **16** (1974), no. 2, 175–180.
- [128] R. J. Gould, *Graph Theory*, Benjamin Cummings, Menlo Park, CA, 1988.
- [129] ———, *Advances on the Hamiltonian problem—A survey*, *Graphs Combin.* **19** (2003), no. 1, 7–52.
- [130] R. J. Gould and M. S. Jacobson, *Forbidden subgraphs and Hamiltonian properties of graphs*, *Discrete Math.* **42** (1982), no. 2-3, 189–196.
- [131] R. L. Graham, M. Grötschel, and L. Lovász (eds.), *Handbook of Combinatorics*, Vol. 1, Elsevier Science, Amsterdam, and MIT Press, Cambridge, 1995.
- [132] ———(ed.), *Handbook of Combinatorics*, Vol. 2, Elsevier Science, Amsterdam, and MIT Press, Cambridge, 1995.
- [133] R. L. Graham, D. E. Knuth, and O. Patashnik, *Concrete Mathematics*, 2nd ed., Addison-Wesley, Reading, MA, 1994.
- [134] R. L. Graham and J. Nešetřil (eds.), *The Mathematics of Paul Erdős I*, *Algorithms Combin.*, vol. 13, Springer-Verlag, Berlin, 1997.
- [135] ———(ed.), *The Mathematics of Paul Erdős II*, *Algorithms Combin.*, vol. 14, Springer-Verlag, Berlin, 1997.
- [136] R. L. Graham, B. L. Rothschild, and J. H. Spencer, *Ramsey Theory*, 2nd ed., Wiley-Intersci. Ser. Discrete Math. Optim., John Wiley & Sons, New York, 1990.
- [137] J. Grantham, *There are infinitely many Perrin pseudoprimes* (2006), 11 pp. Preprint available at <http://www.pseudoprime.com/pseudo>.
- [138] A. Granville, *Zaphod Beeblebrox's brain and the fifty-ninth row of Pascal's triangle*, *Amer. Math. Monthly* **99** (1992), no. 4, 318–331.
- [139] J. Gross and J. Yellen, *Graph Theory and its Applications*, 2nd ed., Chapman & Hall, Boca Raton, FL, 2006.
- [140] J. Grossman, *The Erdős Number Project*, March 1, 2007. <http://www.oakland.edu/enp>.
- [141] B. Grünbaum, *Vertices missed by longest paths or circuits*, *J. Combin. Theory Ser. A* **17** (1974), no. 1, 31–38.
- [142] J. Guare, *Six Degrees of Separation: A Play*, Random House, New York, 1990.
- [143] D. Gusfield and R. W. Irving, *The Stable Marriage Problem: Structure and Algorithms*, MIT Press, Cambridge, MA, 1989.
- [144] H. Hadwiger and H. Debrunner, *Combinatorial Geometry in the Plane*, Holt, Rinehart & Winston, New York, 1964. Translated from the 1960 German edition by V. Klee.
- [145] M. Hall Jr., *Distinct representatives of subsets*, *Bull. Amer. Math. Soc.* **54** (1948), 922–926.
- [146] ———, *Combinatorial Theory*, 2nd ed., John Wiley & Sons, New York, 1983.
- [147] P. Hall, *On representation of subsets*, *J. London Math. Soc.* **10** (1935), 26–30.
- [148] F. Harary, *Graph Theory*, Addison-Wesley, Reading, MA, 1969.

- [149] F. Harary and E. M. Palmer, *The power group enumeration theorem*, J. Combin. Theory **1** (1966), 157–173.
- [150] ———, *Graphical Enumeration*, Academic Press, New York, 1973.
- [151] H. Harborth, *Konvexe Fünfecke in ebenen Punktmengen*, Elem. Math. **33** (1978), no. 5, 116–118.
- [152] G. H. Hardy and S. Ramanujan, *Asymptotic formulae in combinatory analysis*, Proc. London Math. Soc. (2) **17** (1918), no. 1, 75–115.
- [153] D. Harel, *Hamiltonian paths in infinite graphs*, Israel J. Math. **76** (1991), no. 3, 317–336.
- [154] J. M. Harris and M. J. Mossinghoff, *Traceability in small claw-free graphs*, Util. Math. **70** (2006), 263–271.
- [155] F. Hartogs, *Über das Problem der Wohlordnung*, Math. Ann. **76** (1915), no. 4, 438–443.
- [156] P. J. Heawood, *Map-colour theorem*, Quart. J. Math. **24** (1890), 332–339.
- [157] C. Henrion, *Properties of subtle cardinals*, J. Symbolic Logic **52** (1987), no. 4, 1005–1019.
- [158] J. Herman, R. Kučera, and J. Šimša, *Counting and Configurations: Problems in Combinatorics, Arithmetic, and Geometry*, CMS Books Math./Ouvrages Math. SMC, vol. 12, Springer-Verlag, New York, 2003. Translated from the 1997 Czech edition by K. Dilcher.
- [159] C. Hierholzer, *Ueber die Möglichkeit, einen Linienzug ohne Wiederholung und ohne Unterbrechnung zu umfahren*, Math. Ann. **6** (1873), 30–32.
- [160] D. Hilbert, *Über das Unendliche*, Math. Ann. **95** (1926), no. 1, 161–190; English transl., J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931* (1967), 367–392.
- [161] P. G. Hinman, *Recursion-theoretic Hierarchies*, Perspect. Math. Logic, Springer-Verlag, Berlin, 1978.
- [162] J. L. Hirst, *Marriage theorems and reverse mathematics*, Logic and Computation (Pittsburgh, PA, 1987) (W. Sieg, ed.), Contemp. Math., vol. 106, Amer. Math. Soc., Providence, RI, 1990, pp. 181–196.
- [163] P. Hoffman, *The Man Who Loved Only Numbers*, Hyperion, New York, 1998.
- [164] J. D. Horton, *Sets with no empty convex 7-gons*, Canad. Math. Bull. **26** (1983), no. 4, 482–484.
- [165] *The Internet Movie Database*, IMDb.com, Inc., 2008. <http://www.imdb.com>.
- [166] R. W. Irving, *An efficient algorithm for the “stable roommates” problem*, J. Algorithms **6** (1985), no. 4, 577–595.
- [167] T. J. Jech, *The Axiom of Choice*, Stud. Logic Found. Math., vol. 75, North-Holland, Amsterdam, 1973.
- [168] ———, *Set Theory*, 3rd ed., Springer Monogr. Math., Springer-Verlag, Berlin, 2003.
- [169] S. Johnson, *A new proof of the Erdős-Szekeres convex k -gon result*, J. Combin. Theory Ser. A **42** (1986), no. 2, 318–319.
- [170] C. Jordan, *Sur les assemblages des lignes*, J. Reine Angew. Math. **70** (1869), 185–190.
- [171] J. D. Kalbfleisch, J. G. Kalbfleisch, and R. G. Stanton, *A combinatorial problem on convex n -gons*, Proceedings of the Louisiana Conference on Combinatorics, Graph Theory and Computing (Louisiana State Univ., Baton Rouge, LA, 1970) (R. C. Mullin, K. B. Reid, and D. P. Roselle, eds.), Louisiana State Univ., Baton Rouge, LA, 1970, pp. 180–188.
- [172] J. G. Kalbfleisch and R. G. Stanton, *On the maximum number of coplanar points containing no convex n -gons*, Util. Math. **47** (1995), 235–245.
- [173] A. Kanamori, *The mathematical development of set theory from Cantor to Cohen*, Bull. Symbolic Logic **2** (1996), no. 1, 1–71.

- [174] A. B. Kempe, *On the geographical problem of the four colors*, Amer. J. Math. **2** (1879), no. 3, 193–200.
- [175] G. Kirchhoff, *Über der Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Verteilung galvanischer Ströme geführt wird*, Ann. Phys. Chem. **72** (1847), 497–508.
- [176] S. C. Kleene, *Introduction to Metamathematics*, Van Nostrand, Princeton, NJ, 1952. Reprinted by North-Holland, Amsterdam, 1980.
- [177] D. E. Knuth, *Two notes on notation*, Amer. Math. Monthly **99** (1992), no. 5, 403–422. Addendum, *Stirling numbers*, **102** (1995), no. 6, 562.
- [178] ———, *Stable Marriage and its Relation to other Combinatorial Problems*, CRM Proc. Lec. Notes, vol. 10, Amer. Math. Soc., Providence, RI, 1997.
- [179] D. König, *Über eine Schlussweise aus dem Endlichen ins Unendliche*, Acta Litt. Acad. Sci. Hung. (Szeged) **3** (1927), 121–130.
- [180] D. König, *Graphen und Matrizen*, Math. Fiz. Lapok **38** (1931), 116–119.
- [181] J. König, *Sur la theorie des ensembles*, C. R. Acad. Paris **143** (1906), 110–112.
- [182] V. A. Koshelev, *On the Erdős-Szekeres problem in combinatorial geometry*, Electron. Notes Discrete Math. **29** (2007), 175–177.
- [183] J. B. Kruskal Jr., *On the shortest spanning subtree of a graph and the traveling salesman problem*, Proc. Amer. Math. Soc. **7** (1956), no. 1, 48–50.
- [184] K. Kuratowski, *Sur la notion del l'ordre dans la théorie des ensembles*, Fund. Math. **2** (1921), 161–171.
- [185] ———, *Sur le problème des courbes gauches en topologie*, Fund. Math. **15** (1930), 271–283.
- [186] N. J. Lennes, *On the foundations of the theory of sets*, Bull. Amer. Math. Soc. **28** (1922), no. 6, 300. Abstract from the 222nd reg. meeting of the Amer. Math. Soc. (Univ. Chicago, 1922).
- [187] A. Levy, *Basic Set Theory*, Springer-Verlag, Berlin, 1979. Reprinted by Dover, Mineola, NY, 2002.
- [188] M. Lewin, *A new proof of a theorem of Erdős and Szekeres*, Math. Gaz. **60** (1976), no. 412, 136–138.
- [189] L. Lovász, *On decomposition of graphs*, Studia Sci. Math. Hungar. **1** (1966), 237–238.
- [190] ———, *Three short proofs in graph theory*, J. Combin. Theory Ser. B **19** (1975), no. 3, 269–271.
- [191] ———, *Combinatorial Problems and Exercises*, 2nd ed., North-Holland, Amsterdam, 1993.
- [192] É. Lucas, *Théorie des fonctions numériques simplement périodiques*. Sections 1-5, Amer. J. Math. **1** (1878), no. 2, 184–196.
- [193] ———, *Théorie des fonctions numériques simplement périodiques*. Sections 6-23, Amer. J. Math. **1** (1878), no. 3, 197–240.
- [194] ———, *Théorie des fonctions numériques simplement périodiques*. Sections 24-30, Amer. J. Math. **1** (1878), no. 4, 289–321.
- [195] ———, *Récréations Mathématiques IV*, Gauthiers-Villars et fils, Paris, 1894.
- [196] D. F. Manlove, R. W. Irving, K. Iwama, S. Miyazaki, and Y. Morita, *Hard variants of stable marriage*, Theoret. Comput. Sci. **276** (2002), no. 1-2, 261–279.
- [197] J.-L. Marichal and M. J. Mossinghoff, *Slices, slabs, and sections of the unit hypercube*, Online J. Anal. Comb. **3** (2008), Art. 1, 11 pp.
- [198] M. M. Matthews, *Every 3-connected $K_{1,3}$ -free graph with fewer than 20 vertices is Hamiltonian*, Technical Report 82-004, Dept. Comp. Sci., Univ. South Carolina, 1982.

- [199] M. M. Matthews and D. P. Sumner, *Hamiltonian results in $K_{1,3}$ -free graphs*, J. Graph Theory **8** (1984), no. 1, 139–146.
- [200] J. Matoušek, *Lectures on Discrete Geometry*, Grad. Texts in Math., vol. 212, Springer-Verlag, New York, 2002.
- [201] E. Mendelson, *Introduction to Mathematical Logic*, 4th ed., Chapman & Hall, London, 1997.
- [202] K. Menger, *Zur allgemeinen Kurventheorie*, Fund. Math. **10** (1927), 95–115.
- [203] E. P. Miles Jr., *Generalized Fibonacci numbers and associated matrices*, Amer. Math. Monthly **67** (1960), no. 8, 745–752.
- [204] S. Milgram, *The small-world problem*, Psychology Today **1** (1967), no. 1, 61–67.
- [205] K. R. Milliken, *Ramsey's theorem with sums or unions*, J. Combin. Theory Ser. A **18** (1975), no. 3, 276–290.
- [206] E. C. Milner, *Review of "Weakly compact cardinals: A combinatorial proof" by S. Shelah*, MR **0550384 (81i:04009)**, Mathematical Reviews, Amer. Math. Soc., Providence, RI (1981).
- [207] D. Mirimanoff, *Les antinomies de Russell et de Burali-Forti et le problème fondamental de la théorie des ensembles*, Enseign. Math. **19** (1917), 37–52.
- [208] W. Morris and V. Soltan, *The Erdős-Szekeres problem on points in convex position—A survey*, Bull. Amer. Math. Soc. (N.S.) **37** (2000), no. 4, 437–458.
- [209] Y. Moschovakis, *Notes on Set Theory*, 2nd ed., Undergrad. Texts Math., Springer-Verlag, New York, 2006.
- [210] C. Mulcahy, *An ESPeriment with cards*, Math Horizons (Feb. 2007), 10–12.
- [211] C. St. J. A. Nash-Williams, *Which infinite set-systems have transversals?—A possible approach*, Conference on Combinatorial Mathematics (Oxford Univ., 1972), Combinatorics (D. J. A. Welsh and D. R. Woodall, eds.), Inst. Math. Appl., Southend, 1972, pp. 237–253.
- [212] ———, *Another criterion for marriage in denumerable societies*, Cambridge Combinatorial Conference (Trinity College, 1977), Advances in Graph Theory (B. Bollobás, ed.), Ann. Discrete Math., vol. 3, 1978, pp. 165–179.
- [213] P. M. Neumann, *A lemma that is not Burnside's*, Math. Sci. **4** (1979), no. 2, 133–141.
- [214] J. R. Newman (ed.), *The World of Mathematics*, Vol. I–IV, Simon & Schuster, New York, 1956. Reprinted by Dover, Mineola, NY, 2003.
- [215] C. M. Nicolás, *The empty hexagon theorem*, Discrete Comput. Geom. **38** (2007), no. 2, 389–397.
- [216] A. Nijenhuis and H. S. Wilf, *Combinatorial Algorithms*, Academic Press, New York, 1975.
- [217] O. Ore, *Note on Hamilton circuits*, Amer. Math. Monthly **67** (1960), no. 1, 55.
- [218] ———, *Theory of Graphs*, Amer. Math. Soc. Colloq. Publ., vol. 38, Amer. Math. Soc., Providence, RI, 1962.
- [219] M. Overmars, *Finding sets of points without empty convex 6-gons*, Discrete Comput. Geom. **29** (2003), no. 1, 153–158.
- [220] G. Peano, *Super theoremata de Cantor–Bernstein*, Rend. Circ. Mat. Palermo **21** (1906), 136–143.
- [221] J. Petersen, *Die Theorie der regulären Graphen*, Acta Math. **15** (1891), 193–220.
- [222] M. Petkovšek, H. S. Wilf, and D. Zeilberger, *$A = B$* , A K Peters, Wellesley, MA, 1996.
- [223] K.-P. Podewski and K. Steffens, *Injective choice functions for countable families*, J. Combin. Theory Ser. B **21** (1976), no. 1, 40–46.
- [224] G. Pólya, *Kombinatorische Anzahlbestimmungen für Gruppen, Graphen, und chemische Verbindungen*, Acta Math. **68** (1937), 145–254; English transl. in G. Pólya and R. C. Read, *Combinatorial Enumeration of Groups, Graphs, and Chemical Compounds* (1987).

- [225] ———, *On picture-writing*, Amer. Math. Monthly **63** (1956), 689–697; reprinted in I. Gessel and G.-C. Rota (eds.), *Classic Problems in Combinatorics*, Birkhäuser, Boston, 1987.
- [226] G. Pólya and R. C. Read, *Combinatorial Enumeration of Groups, Graphs, and Chemical Compounds*, Springer-Verlag, New York, 1987.
- [227] G. Pólya, R. E. Tarjan, and D. R. Woods, *Notes on Introductory Combinatorics*, Progr. Comput. Sci., vol. 4, Birkhäuser, Boston, 1983.
- [228] R. C. Prim, *Shortest connection networks and some generalizations*, Bell Syst. Tech. J. **36** (1957), 1389–1401.
- [229] H. Prüfer, *Beweis eines Satzes über Permutationen*, Arch. Math. Phys. **27** (1918), 742–744.
- [230] H. Rademacher, *On the partition function $p(n)$* , Proc. London Math. Soc. (2) **43** (1937), 241–254.
- [231] S. P. Radziszowski, *Small Ramsey numbers*, Electron. J. Combin. (August 1, 2006), 60 pp. Dynamic Survey 1, revision 11.
- [232] F. P. Ramsey, *On a problem of formal logic*, Proc. London Math. Soc. (2) **30** (1930), no. 1, 264–286; reprinted in I. Gessel and G.-C. Rota (eds.), *Classic Problems in Combinatorics*, Birkhäuser, Boston, 1987.
- [233] J. H. Redfield, *The theory of group-reduced distributions*, Amer. J. Math. **49** (1927), no. 3, 433–455.
- [234] P. Reynolds and B. Tjaden, *The Oracle of Bacon at Virginia*, 2008. <http://oracleofbacon.org>.
- [235] J. Riordan, *An Introduction to Combinatorial Analysis*, John Wiley & Sons, New York, 1958. Reprinted by Dover, Mineola, NY, 2002.
- [236] ———, *Combinatorial Identities*, John Wiley & Sons, New York, 1968. Reprinted by Krieger, Huntington, NY, 1979.
- [237] ———, *A budget of rhyme scheme counts*, Second International Conference on Combinatorial Mathematics (New York, 1978) (A. Gewirtz and L. V. Quintas, eds.), Ann. New York Acad. Sci., vol. 319, New York Acad. Sci., New York, 1979, pp. 455–465.
- [238] F. S. Roberts and B. Tesman, *Applied Combinatorics*, 2nd ed., Pearson Education, Upper Saddle River, NJ, 2005.
- [239] N. Robertson, D. Sanders, P. Seymour, and R. Thomas, *The four-colour theorem*, J. Combin. Theory Ser. B **70** (1997), no. 1, 2–44.
- [240] H. Rogers Jr., *Theory of Recursive Functions and Effective Computability*, 2nd ed., MIT Press, Cambridge, MA, 1987.
- [241] J. Roitman, *Introduction to Modern Set Theory*, Pure Appl. Math., John Wiley & Sons, New York, 1990.
- [242] J. Rosser, *Extensions of some theorems of Gödel and Church*, J. Symbolic Logic **1** (1936), no. 3, 87–91.
- [243] G.-C. Rota (ed.), *Studies in Combinatorics*, MAA Stud. Math., vol. 17, Math. Assoc. America, Washington, DC, 1978.
- [244] B. Russell, *On some difficulties in the theory of transfinite numbers and order types*, Proc. London Math. Soc. (2) **4** (1906), 29–53; reprinted in B. Russell, *Essays in Analysis* (D. Lackey, ed.), George Braziller, New York, 1973.
- [245] Z. Ryjáček, *On a closure concept in claw-free graphs*, J. Combin. Theory Ser. B **70** (1997), no. 2, 217–224.
- [246] H. J. Ryser, *Combinatorial Mathematics*, Carus Math. Monogr., vol. 14, Math. Assoc. America, 1963.

- [247] B. Schechter, *My Brain is Open*, Simon & Schuster, New York, 1998.
- [248] J. H. Schmerl, *Recursive colorings of graphs*, *Canad. J. Math.* **32** (1980), no. 4, 821–830.
- [249] ———, *Minimal spanning trees*, *Proc. Amer. Math. Soc.* **132** (2004), no. 2, 333–340.
- [250] W. Schmitz, *Über längste Wege und Kreise in Graphen*, *Rend. Sem. Mat. Univ. Padova* **53** (1975), 97–103.
- [251] I. Schur, *Über die Kongruenz $x^m + y^m \equiv z^m \pmod{p}$* , *Jahresber. Deutsch. Math.-Verein.* **25** (1916), 114–117.
- [252] S. Shelah, *Notes on partition calculus*, *Infinite and Finite Sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)*, Vol. III (J. Hajnal, R. Rado, and V. Sós, eds.), *Colloq. Math. Soc. János Bolyai*, vol. 10, North-Holland, Amsterdam, 1975, pp. 1257–1276.
- [253] W. Sierpiński, *Cardinal and Ordinal Numbers*, 2nd ed., *Monografie Matematyczne*, vol. 34, Państwowe Wydawnictwo Naukowe, Warsaw, 1965.
- [254] S. G. Simpson, *Subsystems of Second Order Arithmetic*, *Perspect. Math. Logic*, Springer-Verlag, Berlin, 1999.
- [255] T. Skolem, *Einige Bemerkungen zur axiomatischen Begründung der Mengenlehre*, *Mathematikerkongressen i Helsingfors den 4–7 Juli 1922, Den femte skandinaviska matematikerkongressen, Redogörelse*, Akademiska-Bokhandeln, Helsinki (1922), 217–232; English transl., J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931* (1967), 290–301.
- [256] ———, *Two remarks on set theory*, *Math. Scand.* **5** (1957), 40–46; reprinted in ———, *Selected works in logic* (J. Fenstad, ed.), Universitetsforlaget, Oslo, 1970.
- [257] N. J. A. Sloane, *The On-Line Encyclopedia of Integer Sequences*, AT&T Knowledge Ventures, 2007. <http://www.research.att.com/~njas/sequences>.
- [258] N. J. A. Sloane and S. Plouffe, *The Encyclopedia of Integer Sequences*, Academic Press, San Diego, CA, 1995.
- [259] R. I. Soare, *Recursively Enumerable Sets and Degrees*, *Perspect. Math. Logic*, Springer-Verlag, Berlin, 1987.
- [260] R. P. Stanley, *Generating functions*, *Studies in Combinatorics* (G.-C. Rota, ed.), *MAA Stud. Math.*, vol. 17, Math. Assoc. America, Washington, DC, 1978, pp. 100–141.
- [261] ———, *Enumerative Combinatorics, Vol. 1*, *Cambridge Stud. Adv. Math.*, vol. 49, Cambridge Univ. Press, Cambridge, 1997.
- [262] ———, *Enumerative Combinatorics, Vol. 2*, *Cambridge Stud. Adv. Math.*, vol. 62, Cambridge Univ. Press, Cambridge, 1999.
- [263] D. Stanton and D. White, *Constructive Combinatorics*, *Undergrad. Texts Math.*, Springer-Verlag, New York, 1986.
- [264] R. Steinberg, *Solution to problem 4065: Three point collinearity*, *Amer. Math. Monthly* **51** (1944), no. 3, 169–171.
- [265] J. J. Sylvester, *Questions for solution, problem 11851*, *Educational Times* **46** (1893), 156.
- [266] G. Szekeres and L. Peters, *Computer solution to the 17-point Erdős-Szekeres problem*, *ANZIAM J.* **48** (2006), no. 2, 151–164.
- [267] R. Thomas, *An update on the four-color theorem*, *Notices Amer. Math. Soc.* **45** (1998), no. 7, 848–859.
- [268] G. Tóth and P. Valtr, *The Erdős-Szekeres theorem: Upper bounds and related results*, *Combinatorial and Computational Geometry* (J. E. Goodman, J. Pach, and E. Welzl, eds.), *Math. Sci. Res. Inst. Publ.*, vol. 52, Cambridge Univ. Press, Cambridge, 2005, pp. 557–568.

- [269] W. T. Tutte, *The factorization of linear graphs*, J. London Math. Soc. **22** (1947), no. 2, 107–111.
- [270] V. R. R. Uppuluri and J. A. Carpenter, *Numbers generated by the function $\exp(1 - e^x)$* , Fibonacci Quart. **7** (1969), no. 4, 437–448.
- [271] P. Valtr, *On empty hexagons*, Surveys on Discrete and Computational Geometry: Twenty Years Later (J. E. Goodman, J. Pach, and R. Pollack, eds.), Contemp. Math., vol. 453, Amer. Math. Soc., Providence, RI, 2008, pp. 433–442.
- [272] J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931*, Harvard Univ. Press, Cambridge, MA, 1967.
- [273] J. H. van Lint and R. M. Wilson, *A Course in Combinatorics*, 2nd ed., Cambridge Univ. Press, Cambridge, 2001.
- [274] O. Veblen, *An application of modular equations in analysis situs*, Ann. of Math. **2** (1912/13), no. 14, 86–94.
- [275] J. Venn, *Symbolic Logic*, Macmillan, London, 1881.
- [276] J. von Neumann, *Eine Axiomatisierung der Mengenlehre*, J. Reine Angew. Math. **154** (1925), 34–56; English transl., J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931* (1967), 393–413.
- [277] K. Wagner, *Bemerkungen zum Vierfarbenproblem*, Jahresber. Deutsch. Math.-Verein. **46** (1936), 21–22.
- [278] H. Walther, *Über die Nichtexistenz eines Knotenpunktes, durch den alle längsten Wege eines Graphen gehen*, J. Combin. Theory **6** (1969), 1–6.
- [279] H. Walther and H. J. Voss, *Über Kreise in Graphen*, VEB Deutscher Verlag der Wissenschaften, Berlin, 1974.
- [280] G. S. Warrington, *Juggling probabilities*, Amer. Math. Monthly **112** (2005), no. 2, 105–118.
- [281] D. B. West, *Introduction to Graph Theory*, 2nd ed., Prentice Hall, Upper Saddle River, NJ, 2000.
- [282] Western Maryland College Problems Group, *Problem 11007*, Amer. Math. Monthly **110** (2003), no. 4, 340. Solution, R. Tauraso, <http://www.mat.uniroma2.it/~tauraso/AMM/amm2003.html>.
- [283] N. Wiener, *A simplification of logic of relations*, Proc. Cambridge Phil. Soc. **17** (1912–1914), 387–390.
- [284] H. S. Wilf, *Generatingfunctionology*, 3rd ed., A K Peters, Wellesley, MA, 2006.
- [285] R. Wilson, *Four Colors Suffice*, Princeton Univ. Press, Princeton, NJ, 2002.
- [286] J. Wojciechowski, *A criterion for the existence of transversals of set systems*, J. London Math. Soc. (2) **56** (1997), no. 3, 491–503.
- [287] J. Worpitzky, *Studien über die Bernoullischen und Eulerischen Zahlen*, J. Reine Angew. Math. **94** (1883), 203–232.
- [288] E. M. Wright, *Burnside's lemma: A historical note*, J. Combin. Theory Ser. B **30** (1981), no. 1, 89–90.
- [289] Y. Yang, *On a multiplicative partition function*, Electron. J. Combin. **8** (2001), no. 1, R19, 14 pp.
- [290] T. Zamfirescu, *On longest paths and circuits in graphs*, Math. Scand. **38** (1976), no. 2, 211–239.
- [291] E. Zermelo, *Beweis, daß jede Menge wohlgeordnet werden kann*, Math. Ann. **59** (1904), no. 4, 514–516; English transl., J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931* (1967), 139–144.
- [292] ———, *Neuer Beweis für die Möglichkeit einer Wohlordnung*, Math. Ann. **65** (1908), no. 1, 107–128; English transl., J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931* (1967), 183–198.

- [293] ———, *Untersuchungen über die Grundlagen der Mengenlehre I*, *Math. Ann.* **65** (1908), no. 2, 261–281; English transl., J. van Heijenoort, *From Frege to Gödel. A Source Book in Mathematical Logic, 1879–1931* (1967), 199–215.

Index

- Acquaintance Graph, 26
- adjacency
 - matrix, 22
 - of vertices, 5
- Aharoni, Ron, 338
- Aigner, Martin, 127, 277, 278
- alephs, 308
 - $\aleph_0, \aleph_1, \dots$, 312
- Alexander the Great, 52
- Alice in Wonderland*, 1
- Alma, Alabama, 144, 148
- Amarillo, Texas, 137
- Anacreontea*, 336
- analytical sets, 351
- Anastasia*, 111
- Anderson, Poul, 202
- Andrews, George E., 225, 278
- Anquetil, Jacques, 200
- Anthony and Cleopatra*, 83
- anthracene, 207
- anvil salesman, 51
- Appel, Kenneth, 95
- approximating irrationals, 153
- Aragorn, 301
- Arizona Republic, The*, 150
- Armstrong, Lance, 200
- arrow notation, 287, 322
- Arthur, King of the Britons, 191, 227
- Atlantic Coast Conference, 4
- average degree, 38, 80, 93
- axiom of choice
 - AC, 298
 - AC2, 298
 - ACF, 302
 - CAC, 302, 303
 - CACF, 302, 329
 - equivalences, 298–303, 317
 - PIT, 329
 - weak versions, 302–303
 - Zorn’s Lemma, 317
- axioms of ZFC
 - choice, 298
 - empty set, 292
 - extensionality, 292
 - infinity, 296
 - pairing, 293
 - power set, 294
 - regularity, 296
 - replacement, 296
 - separation, 294

- union, 293
- Bacon, Kevin
 - game, 27
 - number, 27
 - person, 26, 27
- Baïou, Mourad, 279
- Balinski, Michel L., 279
- Ballad*, 102
- ballots, 136
- Banks, David C., 278
- Barabási, Albert-László, 127
- baseball pitchers, 154
- Bedivere, Sir, 190, 191, 228
- Beineke, Lowell W., 126
- Bell numbers, 233, 237–239
 - complementary, 240, 279
- Benjamin, Arthur T., 278
- benzene, 206
- Berge, Claude, 277
- Berge's Theorem, 104
- Bernstein, Felix, 307
 - Theorem, *see* Cantor–Bernstein
- Bert and Ernie, 4
- Bielak, Halina, 20
- big one, the, 304, 305
- Biggs, Norman L., 126
- bijection, 192
- binary sequence, 184
- binomial coefficients, 133
 - absorption/extraction, 142, 170
 - addition, 138, 170
 - cancellation, 142, 170
 - expansion, 138
 - generalized, 168
 - hexagon identity, 143
 - negating upper index, 169
 - parallel summation, 143, 170
 - summing on upper index, 141
 - symmetry, 138
- binomial theorem, 139
 - for factorial powers, 143
 - generalized, 168
- bipartite graph, 13
 - complete, 13
- Birkhoff, George, 97
- Birkhoff Diamond, 87
- Birkhoff–Lewis Reduction Theorem, 98
- Blakley, George Robert, ix
- Bloch, André, 156
- Bobet, Louision, 200
- Bobo, Mississippi, 150
- Bollobás, Béla, 126
- Bond, James, 101
- boots and socks, 298
- Borwein, Peter B., 267, 278, 279
- bound degree, 76
- Boundin'*, 209
- box principle, 151
- bracelets, 209, 213, 215
- Brass, Peter, 267, 274, 279
- bridge
 - card game, 133
 - in a graph, 8, 10
 - in Königsberg, 1, 52–54
- Bridges, Robert, 30
- Brooks's Theorem, 90
- Browning, Elizabeth Barrett, 248
- Buckley, Fred, 126
- buffet line, 162
- Bug Tussle, Texas, 150
- Burnside, Ambrose E., 196
- Burnside, William, 199
 - lemma of, 199
- burnt orange, 217
- Burr, Stefan, 125
- Butler, Samuel, 2
- C++ variable names, 134
- cafeteria, 177
- Calverly, C. S., 102
- Campbell, Thomas, 17
- Candide*, 292
- Cantor, Georg, 290, 353
- Cantor's Theorem, 305
- Cantor–Bernstein Theorem, 306
 - applications, 307–308, 312, 314, 316, 343
 - proof, 306, 316

- carboxyl group, 207
- cardinal, 311–312
 - color, 217
 - espousable, 341
 - large, 322
 - measurable, 323
 - Ramsey, 323, 326
 - regular, 313, 314
 - Richelieu, 322
 - singular, 313
 - strong limit, 322
 - strongly inaccessible, 322
 - subtle, 323, 345
 - weakly compact, 322, 324–326
 - weakly inaccessible, 320–322
- carnival, 171
- Carpenter, John A., 279
- Carroll, Lewis, 1, 109
- Cartesian product, 295, 300
- Catalan numbers, 185–188
- Cauchy–Binet Formula, 50
- Cayley, Arthur, 43, 44
- Cayley’s Tree Formula, 44
- ceiling function ($\lceil x \rceil$), 153
- center
 - of a graph, 18, 19
 - of a tree, 36
- Chakerian, Gulbank D., 279
- changing money, 171–175
 - in 1875, 175
- characteristic path length, 29
- Chartrand, Gary, 126
- Chava and Hodel, 250
- chemistry, 32, 206–208
- child vertex, 188
- chromatic number, 86
 - bounds, 88–93
- chromatic polynomial, 98, 159, 163
 - properties, 101
 - relation to Four Color Problem, 101
- Chvátal, Václav, 63, 124, 125
- circuit, 6
 - Eulerian, 55
- claw, 66, 69, 72
- clique number, 92
- clustering coefficient
 - of a graph, 29
 - of a vertex, 29
- cofactor, 48
- Coffey, Paul, 242
- college admissions, 249, 263
- color classes, 86
- k -colorability of a graph, 86
- coloring
 - of edges, 116
 - of vertices, 86
- k -coloring of a graph, 86
 - proper, 86
- combinatorial geometry, 264
- commemorative coins, 176, 177
- complement, 11
- complementary Bell numbers,
 - see* Bell numbers
- complete
 - bipartite graph, 13
 - graph, 10
 - multipartite graph, 16
- complete graph, 8
- composition, 226
- computability, 349
- Comtet, Louis, 241, 277
- concert hall, 242
- Conehead, Connie, 304
- connected component, 8
- connected graph, 7
- Connecticut Yankee in King Arthur’s Court*, A, 227
- connectives, 291
- connectivity of a graph, 8
- constructible universe (L), 320
- convolution of two sequences, 186
- Conway, John H., 279
- Cooleemee, North Carolina, 150
- countable sets, 305
- countable union, 314
- cover, 109
- Cowell, Charlie (anvil salesman), 51
- Coxeter, Harold Scott MacDonald, 265

- k -critical graphs, 88
- Csima, József, 267
- cube, 81
- cut
 - in a network, 110
 - capacity, 111
 - set, 8
 - vertex, 8, 10
- cycle
 - Hamiltonian, 60
 - index (of a group), 201
 - notation (for permutations), 192
 - the graph C_n , 12
 - within a graph, 6
- Damerell, R. M., 331
- de Bruijn, Nicolaas Govert, 278
 - enumeration formula, 212
- de Wannemacker, Stefan, 240
- Debrunner, Hans E., 279
- deck of cards, 133
- Dedekind finite, 303
- degree
 - average, 38, 80, 93
 - matrix, 48
 - of a vertex, 6
 - sequence, 6
- DeMorgan, Augustus, 2, 94
- Denver, John, 116
- deranged twins, 163
- derangements, 160–161, 163, 231
- detour
 - order, 67, 93
 - path, 67, 93
- Devlin, Keith J., 352
- diameter of a graph, 18
- dice, six-sided, 200, 202, 206
- Diestel, Reinhard, 126, 327, 353
- difference operator, 137
- digraph, *see* directed graph
- Dijkstra, Edsger Wybe, 265
- Dirac, Gabriel Andrew, 62, 84
- directed graph (digraph), 3
- Dirichlet, Johann Peter Gustav Lejeune, 151, 153
- Dirichlet's approximation theorem, 153
- disjointification, 303
- Dissertatio de Arte Combinatoria*, vii, 129
- distance
 - between vertices, 18
 - matrix, 25
- Dobiński's formula, 239
- dodecahedron, 60, 81
- Drake, Frank R., 352
- Duffus, Dwight, 66
- Dumas, Alexandre, 124
- Dumitrescu, Adrian, 275
- Dunbar, Jean, 72
- eccentricity, 18
- Edberg, Stefan, 231
- edge, 2
 - deletion, 7
 - edge set, 5
- edge cover, 109, *see* cover
- Edmonton Oilers, 242
- Edwards, Anthony William Fairbank, 277
- Eeckhout, Jan, 255
- Ehrenfeucht, Andrzej, 276
- Einstein, Albert, 97, 130
- Ekai, Mumon, 304
- empty
 - graph, 11
 - product, 132
 - set (\emptyset), 290, 299
 - axiom, *see* axioms of ZFC
- end vertices
 - of a walk, 6
 - of an edge, 5
- Enderton, Herbert B., 352
- Epcot Center, 80
- equivalence
 - class, 197
 - relation, 197, 315
- Erdős, Paul, 28, 30, 63, 116, 122, 123, 127, 152, 264, 270, 279, 280, 325

- number, 28
- Eriksson, Kimmo, 278
- Euler, Leonhard
 - and Königsberg bridge problem, 1, 52–54
 - characterization of Eulerian graphs, 55
 - Formula (for planar graphs), 78, 81
 - losing sight, 78
 - Pentagonal Number Theorem, 222
 - φ function, 158, 162
- Eulerian
 - circuit, 55
 - graph, 55
 - numbers, 242–247
 - trail, 55
- Fáry, István, 77
- factorial, 132
- falling factorial power, 132
- Faudree, Ralph, 69
- Feder, Tomás, 279
- Federer, Roger, 231
- Feferman, Solomon, 344, 346
- Ferrers diagram, *see* Young diagram
- Fezzik, 191
- Fibonacci numbers, 177–179
 - generalized, 185
- Fiddler on the Roof*, 250
- Filthy Frank, 4
- finger, 304
- finite set, 282
- First Theorem of Graph Theory, 6
- Five Color Theorem, 95
- flags, 161
- Flatliners*, 26
- Fleury's algorithm, 59
- floor function ($\lfloor x \rfloor$), 153
- flow, 110
- football
 - American, 248
 - International Football Association, 135
 - Laws of the Game, 135
- forbidden subgraphs, 65
- forest, 31
 - number of edges, 35
- Foulds, Leslie R., 126
- Four Color Problem, 2, 93
- Four Color Theorem, 94
- fractional part function ($\{x\}$), 153
- Fraenkel, Abraham A., 290, 294, 296
- Franklin, Benjamin, 277
- Franklin, Fabian, 222
- Frege, Gottlob, 290
- Frick, Marietjie, 72
- Friedman, Harvey M., 344, 347, 350
- Frink, Orrin, 84
- Frobenius, Ferdinand G., 199
- From Russia with Love*, 101
- Frost, Robert, 5, 282
- Fuhr, Grant, 242
- Fujita, Shinsaku, 278
- full house, 134
- function assignment, 345
- fusion, 283
- Gödel, Kurt F., 318
- Gödel's Incompleteness Theorem
 - First, 318–320
 - Second, 319–320
- Galahad, Sir, 191
- Gale, David, 250, 279
- Gale–Shapley algorithm, 250
- Gallai, Tibor, 68, 265
- Gateless Gate, The*, 304
- Gawain, Sir, 191
- Gekko, Gordon, 88
- general position, 270
- generating function, 164
 - exponential, 238
- geometry of position, 54
- George of the Jungle, 324
- Gerken, Tobias, 277
- Gessel, Ira M., 279, 280
- Gibson, William, 308
- Gilbert, Sir William S., 137
- Giving Tree, The*, 34

- Goodman, Seymour, 65
 Gould, Ronald J., viii, 62, 66, 69, 126
 Graham, Ronald L., 127, 171, 277, 278, 280
 Grand Slam matches, 231
 Grant, Cary, 166
 Grant, Ulysses S., 196
 Grantham, Jon, 181
 Granville, Andrew, 277
 graph
 bipartite, 13
 center, 18
 clique number, 92
 clustering coefficient, 29
 complement, 11
 complete, 8, 10
 bipartite, 13
 multipartite, 16
 connected, 7
 connectivity, 8
 critical, 88
 definition, 2
 diameter, 18
 directed, 3
 empty, 11
 Eulerian, 55
 Hamiltonian, 60
 infinite, 4
 isomorphic pairs, 15
 k th power, 21
 line, 16, 64, 66, 67, 70, 93
 matching, 104
 order, 5
 periphery, 18
 planar, 74
 radius, 18
 Ramsey theory, 124
 regular, 11
 self-centered, 21
 size, 5
 traceable, 61
 weighted, 39
Great Gatsby, The, allusion to, 260
 greedy algorithm, 88
 Gretzky, Wayne, 242
 Gross, Jonathon L., 126
 Grossman, Jerry, 28
 group
 abelian, 191
 alternating, 194
 cyclic, 193
 definition, 191
 dihedral, 193
 generated by an element, 193
 symmetric, 192
 Grünwald, Tibor, 265
 Guare, John, 26
 Guinness, Alec, 30
 gurus, 157
 Gusfield, Dan, 279
 Guthrie, Francis, 94
 Guy, Richard K., 279

 Hadwiger, Hugo, 279
 Haken, Wolfgang, 95
 Hall, Daryl and Oates, John, 328
 Hall, Monty, 104, 105
 Hall, Philip, 105, 328
 Hall Jr., Marshall, 277, 328
 Hall's Theorem, 105, 328
 corollary for regular graphs, 112
 halting problem, 349
 hamburgers, 170
 Hamilton, Al, 242
 Hamilton, Sir William Rowan, 61, 94
 Hamiltonian
 cycle, 60
 graph, 60
 path, 60, 351
 handshakes, 190
 Hanoi, Tower of, 181
 Harary, Frank, 124, 126, 278
 Harborth, Heiko, 277
 Hardy, Godfrey H., 224
 harey problem, 177
 harmonic mean, 155
 Harris, Priscilla, viii
 Harris, Sophie, viii

- Harris, Will
 allusion to, vi
Harry Potter and the Prisoner of Azkaban, 27
- Heawood, Percy, 94, 95
- Hedetniemi, Stephen, 19, 65
- Heine–Borel Theorem, 285
- heliotrope, 209, 217
- Herman, Jiří, 279
- Hierholzer, Carl, 55
- Hierholzer’s algorithm, 57
- Higgledy-Piggledy*, 2
- Hilbert, David, 281, 290
- Hinault, Bernard, 200
- Hirst, the other Prof., ix
- Hitchcock, Alfred, 166
- Hoffman, Paul, 127
- Hollywood Graph, 27
- hopscotch, 181
- Horton, Joseph D., 277
- Houston, Whitney, 343
- humuhumunukunukuapua’a, 150
- hungry
 fraternity brother, 235
 math major, 177
- Hunting of the Snark, The*, 109
- hydroxyl group, 206
- hypergraph, 4
- ichthyologists, 150
- icosahedron, 81
- Icosian Game, The, 60
- incidence
 matrix, 48
 of vertex and edge, 5
- independence number, 63, 93
- independent set of vertices, 63
- independent zeros, 111
- induced subgraph, 12
- Indurain, Miguel, 200
- infinite set, 282
- injective function, 192
- Internet Movie Database, 27
- intersecting detour paths, 67
- intersection, 294
- invariant set, 198
- Irish Blessing, An*, 10
- irrational numbers, 153
- Irving, Robert W., 279
- isomer, 206
- isomorphism, 15
- Jacobson, Michael S., 66
- Jech, Thomas J., 352
- Jefferson, Thomas, 176
- jelly beans, 170
- JFK*, 27
- Johnson, Scott, 275
- Just Men of Cordova, The*, 352
- Kalbfleisch, James G., 274
- Kanamori, Akihiro, 353
- Kelly, Leroy M., 265
- Kempe, Alfred, 94, 95
- killer rabbits, 177
- Kirchhoff, Gustav, 47, 48
- Klee Jr., Victor L., 264, 279
- Kleene, Stephen C., 318, 353
- Klein, Esther, 264, 270
- knights of the round table, 227–228
- Knuth, Donald E., ix, 171, 255, 277–279
- König, Dénes, 284, 328
- König, Julius, 284, 307
- König–Egerváry Theorem, 109
- König’s Lemma, 283, 302, 326, 351
- Königsberg Bridge Problem, 1, 52–54
- Kruskal’s algorithm, 40–42
- Kučera, Radan, 279
- Kuratowski, Kazimierz, 83, 84
- Kuratowski’s Theorem, 84
- Kurri, Jari, 242
- L*, constructible universe, 320
- Laffey, Thomas, 240
- Lancelot, Sir, 228
- lapis lazuli, 201, 203, 214
- Last of the Mohicans, The*, allusion to, 260

- lauwiliwilinukunuku'oi'oi, 150
 lavender, 217
 lazy professor, 160, 163
 leaf, 31
 number in tree, 35
 tea, 34
Leaves of Grass, 181
 Lehrer, Tom, 312
 Leibniz, Gottfried Wilhelm, vii, 129
 LeMond, Greg, 200
 length, 6
 Leonardo of Pisa, 177
 Lesniak, Linda, 126
 Levy, Azriel, 352
 Lewinter, Marty, 126
 Lewis and Clark expedition, 176
Liar, Liar, 67
 limerick, 236
 Lincoln, Abraham, 176, 177
 line graph, 16, 64, 66, 67, 70, 93
 linear ordering, 309
 k -critical, 347
 Linton, Stephen A., 278
 Lloyd, E. Keith, 126
 Logothetti, David E., 279
London Snow, 30
 Longfellow, Henry W., 38, 285
Looney Tunes, 265
 lottery
 Florida Fantasy 5, 144
 Florida Lotto, 144
 Lotto Texas, 133, 137, 141
 repetition allowed, 171
 Rhode Island Wild Money, 137
 Texas Two Step, 136
 Virginia Win For Life, 144
 Lovász, László, 72, 90, 277
Love's Labour's Lost, 74
 Lucas, Édouard, 181
 numbers, 180

 Makai, Endre, 274
 Man in black, 297
 Marichal, Jean-Luc, 279
 Mark, gospel of, 171

 maroon, 217
 marriage
 infinite sets, 327–344
 Secrets of a Successful, 218
 stable, *see* stable marriage
 matching, 102
 graph, 104
 M -alternating path, 104
 M -augmenting path, 104
 many-to-many, 264
 many-to-one, 263
 maximal, 102
 maximum, 102
 perfect, 102, 111, 343
 saturated edges, 102
 stable, 248
 optimal, 252
 pessimal, 253
 strongly stable, 259
 super-stable, 259
 weakly stable, 259
 Mathematical Collaboration Graph,
 28
 Matoušek, Jiří, 275, 279
 Matrix Tree Theorem, 48
Matrix, The, 21
 Matthews, Manton, 70
 Matthews and Sumner's Conjecture,
 69
 Max Flow Min Cut Theorem, 111
 maximal planar graph, 80
 maximum degree, 6
 McEnroe, John, 231
 McPeake, Sharon, viii
 Meade, George G., 196
 Meeks, Randy, v
 Mendelson, Elliott, 353
 Menger's Theorem, 110
 Merckx, Eddy, 200
Merry Wives of Windsor, The, 217
 Messier, Mark, 242
 methyl group, 206
 metric, 17
 Mihók, Peter, 72
 Miles Jr., Ernest P., 278

- Milgram, Stanley, 26
 Milner, Eric C., 331, 340
 Milton, John, 320
 minimum degree, 6
 minimum weight spanning trees, 39–42
Moby Dick, allusion to, 260
 model of ZFC, 321
 Moe, Larry, and Curly, 87
Mona Lisa Overdrive, 308
 monotonic subsequences, 152
 Montagues, Capulets, and Hatfields, 73
 Monticello, 176
Monty Python and the Holy Grail, 168, 190
 Morris Jr., Walter D., 275, 279
 Moschovakis, Yiannis, 352
 Moser, William O. J., 267, 274, 279
 Mossinghoff, Alexandra
 allusion to, vi
 Mossinghoff, Amanda
 allusion to, ix
 Mossinghoff, Kristine, ix
 Mossinghoff, Michael J., 278, 279
 Mulcahy, Colm, 278
 multigraph, 3
 multinomial coefficients, 144–149
 analogue of Vandermonde’s convolution, 150
 addition, 146
 expansion, 145
 symmetry, 145
 multinomial theorem, 147, 167
 for factorial powers, 149
 multiset, 147
Music Man, The, 51

N is a Number, 30
 naphthalene, 206
 naphthol, 206
 Nash-Williams, Crispin St. John Al-
 vah, 335, 338
 Nastase, Ilie, 231
 National Basketball Association, 135
 National Resident Matching Program, 249
 necklaces, 191, 198–203
 neighborhood
 closed, of a vertex, 6
 of a set, 6
 open, of a vertex, 5
 Nešetřil, Jaroslav, 280
 neurotic running back, 248
 Nicolás, Carlos M., 277
Night of the Lepus, 177
 Nijenhuis, Albert, 277
 North American Numbering Plan, 131
North by Northwest, 166

 occupancy problems, 217–218
 Oconomowoc, Wisconsin, 150
 octahedron, 81, 217
 ogre and ogress, 255, 258
 Oldman, Gary, 27
 OPPOSITES ATTRACT, 209, 213, 215
 orbit, 198
 order of a graph, 5
 ordinal, 309–312, 317
 ordinary line, 265
 Ore, Oystein, 63, 353
 Osburn, Robert, 240
Othello, 237
 Overmars, Mark, 277

 Pach, Janos, 267, 274, 279
 Padovan sequence, 180
 Palmer, Edgar M., 278
 Pangloss, 292
 paradise
 Cantor’s, 290
 Lost, 320
 tasting, 292
 parent vertex, 188
 partite set, 13
 τ -partitionable graphs, 71
 partitions, 175, 218–225
 conjugate, 221
 distinct parts, 221, 225
 Pascal, Blaise, 80, 139

- pyramid of, 146
- triangle of, 139
- Patashnik, Oren, 171, 277, 278
- path
 - closed, 6
 - Hamiltonian, 60, 351
 - the graph P_n , 12
 - within a graph, 6
- Path Partition Conjecture, 71
- pattern inventory, 203
- Peano Arithmetic (PA), 318
- Pearl, The*, allusion to, 260
- Pensées*, 80
- pentagonal number, 226
- Percival, Sir, 191, 228
- perfect matching, 102, 111
 - in regular graphs, 114
- periphery of a graph, 18, 20
- permutation, 132
 - as function, 191
 - even and odd, 194
- Perrin sequence, 180
- Peters, Lindsay, 274
- Petersen, Julius, 114, 115
- Petersen graph, 64, 87, 115
- Petersen's Theorem, 115
- Petkovšek, Marko, 277
- phenomenology exam, 135
- Phoenix, Arizona, 150, 151
- phone numbers, 131
- pigeonhole principle
 - finite, 118, 150–152, 154, 281, 312
 - infinite, 282, 313
 - ultimate, 313
 - variations, 318
- ping-pong balls, 148, 235
- pipe organ, 242, 244
- Pirates of Penzance, The*, 137
- planar graph, 74
 - maximal, 80
 - straight line representation, 77
- planar representation, 74
- Pleasures of Hope, The*, 17
- Pliny the Younger, 31
- Plouffe, Simon, 188, 279
- Podewski, Klaus-Peter, 335
- poker
 - card game, 133, 136, 162
 - chips, 148, 169, 218, 219
 - multiple decks, 167, 169
 - two decks, 166–167
- Pólya, George, 156, 171, 190, 277, 278, 280
 - enumeration formula, 203
- polyhedra, 80
- Poor Richard's Almanack*, 277
- power set (\mathcal{P}), 290, 299
 - axiom, *see* axioms of ZFC
- Prüfer sequence, 51
- Prim's algorithm, 43
- prime numbers, 159, 163
- Princess Bride, The*, 191, 297
- Princess Fiona, 255
- Princess Leia, 93
- principle of inclusion and exclusion, 158
 - generalization, 163
- product rule, 131
- proverbial alien, 118
- Prüfer, Heinz, 44
- pseudograph, 3
- Purdy, George B., 279
- quantifiers, 291
- Quinn, Jennifer J., 278
- Rademacher, Hans, 224
- radius of a graph, 18
- Radziszowski, Stanisław P., 127
- Ramanujan, Srinivasa, 224
- Ramsey, Frank P., 116, 271, 280, 287
- Ramsey numbers
 - classical, 116
 - known bounds, 123, 273
 - known values, 122
 - graph, 124
- Ramsey's Theorem
 - failure at \aleph_1 , 324
 - finite, 272, 286, 288

- hypergraphs, 272, 275
- infinite, 287
- pairs, 286, 351
- triples, 289, 351
- variations, 352
- ransom note, 176
- Read, Ronald C., 278
- Redfield, J. Howard, 278
- region, 75
- registrar, 162
- regressive value, 345
- regular
 - graph, 11
 - polyhedra, 81
- relatively prime, 158
- relief agency, 208, 216
- restaurant
 - gourmet, 162
 - steakhouse, 162
- restriction, 332
- Return of the King, The*, 301
- reverse mathematics, 350
- Reynolds, Patrick, 27
- rhodonite, 201, 203, 214
- rhyming schemes, 236, 240, 279
- ridgeline, 190
- Riordan, John, 277, 279
- rising factorial power, 132
- Roberts, Fred S., 126
- Robertson, Neil, 95
- Robin Hood*, 126
- Roitman, Judith, viii, 352
- Romeo and Juliet*, 73
- rose quartz, 201, 203, 214
- Rota, Gian-Carlo, 129, 279, 280
- Rothschild, Bruce L., 127
- round tables, 191, 197, 199, 227, 229
- Russell, Bertrand, 295, 298
- Ryjáček, Zdeněk, 70
- Ryser, Herbert J., 277

- Sadie Hawkins dance, 254
- Sanders, Daniel, 95
- Sawyer, Eric T., 267
- Schechter, Bruce, 127

- Schlömilch, Oskar X., 241
- Schmitz, Werner, 69
- Schröder, *see* Cantor–Bernstein
- Schubfachprinzip, 151
- Schur, Issai, 276
- Schuster, Seymour, 84
- Scream 2*, v
- SDR, 107, 301, 327
 - version of Hall's Theorem, 107
- self-centered graph, 21
- separated set, 268
- separating set, 110
- Seuss, Dr., 85
- Seymour, Paul, 95
- Shakespeare, William, 73, 74, 83, 164, 217, 237
 - characters, 257
- Shapley, Lloyd S., 250, 279
- Shelah, Saharon, 338, 341
- Shin, Jae-II, viii
- Shrek 2*, 255
- Sierpiński, Waclaw, 352
- Σ_1^1 -complete sets, 351
- Silverstein, Shel, 34
- Simpson, Homer, 218
- Simpson, Stephen G., 350
- Šimša, Jaromír, 279
- six degrees of separation, 26
- size of a graph, 5
- Skolem, Thoralf, 290, 294, 296
- Sloane, Neil James Alexander, 188, 279
- small world networks, 28
- Smith, Paul, 84
- soccer team, 135
- socks, 161
- Soltan, Valeriu P., 275, 279
- Song of Hiawatha, The*, 285
- sonnet, 236
- Sonnets from the Portuguese*, 248
- Soso, Mississippi, 150
- Sound of Trees, The*, 282
- space cruiser, 200
- spanning tree, 39
 - counting, 43

- of minimum weight, 39–42
- Spencer, Joel H., 127
- stabilizer, 198
- stable enrollment, 264
- stable marriage
 - algorithm, 250
 - main theorem, 252
 - problem, 248–262
 - with indifference, 259–261
 - with sets of different sizes, 261–262
 - with unacceptable partners, 256–259
- stable roommates, 249, 279
- staircase, 189
- Stanley, Richard P., 188, 203, 277, 278
- Stanton, Dennis, 277
- Stanton, Ralph Gordon, 274
- star, 34
- Star Wars*, 93
- stationary set, 339
- Steffens, Karsten, 335
- Stirling
 - cycle numbers, 227–230
 - set numbers, 231–235
- Stockmeyer, Paul, 278
- strikeouts, 154
- stump, 31
- subdivision
 - of a graph, 84, 85
 - of an edge, 84
 - of Verona, 73
- subgraph, 12
 - forbidden, 65
 - induced, 12
- subgroup, 193
- Sullivan, Sir Arthur S., 137
- sum rule, 131
- Sumner, David, 70
- surjective function, 192
- Sweet 16, 32
- Sylvester
 - James Joseph, 265
 - Looney Tunes cat, 265
 - problem of, 265–267
 - system of distinct representatives, *see* SDR
- Syśło, Maciej, 20
- Szekeres, George, 122, 152, 264, 274, 279, 280
- Tarjan, Robert E., 277
- Tarski, Alfred, 325
- Tarsy, Michael, 275
- tennis, 231
- termination argument, 265
- Tesman, Barry, 126
- tetrahedron, 81
- tetramethylnaphthalene, 206
- tetraphenylmethane, 208
- Texas
 - cities, 137, 150, 249, 254
 - lottery, *see* lottery
- thistle, 217
- Thomas, Robin, 95
- Thompson, Emma, 27
- Thornhill, Roger, 166
- Three Musketeers, The*, 322
- Thys, Philippe, 200
- Tolkein, J. R. R., 301
- Tour de France, 200
- trace of a square matrix, 25
- traceable graph, 61
- trail, 6
 - closed, 6
 - Eulerian, 55
- Traité du Triangle Arithmétique*, 139
- transfinite
 - induction, 332
 - recursion, 332
- transitive set, 309, 316
- transposition, 194
- tree, 30, 283
 - Aronszajn, 326
 - as a model, 31–32
 - as a subgraph, 36
 - binary decision, 32
 - characterization, 35, 38, 42
 - definition, 31

- in chemistry, 32
- in probability, 31
- in programming, 32
- labeled, 43
- labels, 283, 284
- named, 352
- number of edges, 34
- palm, 34
- property, 326
- rooted, 188
- small, 31
- spanning, 39, 352
- strictly binary, 188
- triangulation, 189
- tribonacci numbers, 184
- trimethylantracene, 207
- trinomial coefficients, 150
- triphenylamine, 207
- Tristram, Sir, 191, 228
- Tutte's Theorem, 112
- Twain, Mark, 227
- twig, 31
- Typee*, allusion to, 260
- typesetter's comfort, 290
- Tyson, Mike, 150, 151

- Unalaska, Alaska, 150
- union, 290, 299
 - axiom, *see* axioms of ZFC
- United Nations, 135
- universal set, 295
- Uppuluri, V. R. Rao, 279
- Urban Legend*, 18

- van Heijenoort, Jean, 353
- van Lint, Jacobus H., 277
- Vandermonde's convolution, 142, 150
- Veblen, Oswald, 55
- Venn, John, 156
- Venn diagram, 157
- vertex, 2
 - cut set, 8
 - deletion, 7
 - vertex set, 5
- vexillologist, 161

- Village Blacksmith, The*, 38
- Vizzini, 297
- volleyball tournament, 184
- Voyage Round the World, A, 60

- Wagner, Klaus, 77
- walk, 6
- Walla Walla, Washington, 150
- Wallace, Edgar, 352
- Wall Street*, 88
- Walther, Hansjoachim, 68
- Warrington, Gregory S., 279
- Washington, George, 177
- weakly ordered rankings, 259
- weight function, 39
- weighted graph, 39
- well-ordering, 309–311
- West, Douglas B., 126
- White, Dennis, 277
- Whitman, Walt, 181
- Wilf, Herbert S., 240, 277, 278
- Wilson, Richard M., 277
- Wilson, Robin J., 126
- Winter's Tale, The*, 164
- wisteria, 217
- Wojciechowski, Jerzy, 335
- Woods, Donald R., 277
- Worpitzky's identity, 247

- Yang, Yifan, 240
- Yellen, Jay, 126
- Young diagram, 220

- Zamfirescu, Tudor, 68
- Zeilberger, Doron, 277
- Zermelo, Ernst, 290, 311
- ZF and ZFC, 292
- ZFC, 290
 - axioms, *see* axioms of ZFC
 - limitations, 318, 320, 344
- Ziegler, Günter M., 278
- zodiac sign, 170
- Zorn's Lemma, 311, 317