

OXFORD LECTURE SERIES IN MATHEMATICS
AND ITS APPLICATIONS · 28

Graphs and Homomorphisms



Pavol Hell and
Jaroslav Nešetřil

Oxford Lecture Series in
Mathematics and its Applications 28

Series Editors

John Ball Dominic Welsh

OXFORD LECTURE SERIES
IN MATHEMATICS AND ITS APPLICATIONS

1. J. C. Baez (ed.): *Knots and quantum gravity*
2. I. Fonseca and W. Gangbo: *Degree theory in analysis and applications*
3. P. L. Lions: *Mathematical topics in fluid mechanics, Vol. 1: Incompressible models*
4. J. E. Beasley (ed.): *Advances in linear and integer programming*
5. L. W. Beineke and R. J. Wilson (eds): *Graph connections: Relationships between graph theory and other areas of mathematics*
6. I. Anderson: *Combinatorial designs and tournaments*
7. G. David and S. W. Semmes: *Fractured fractals and broken dreams*
8. Oliver Pretzel: *Codes and algebraic curves*
9. M. Karpinski and W. Rytter: *Fast parallel algorithms for graph matching problems*
10. P. L. Lions: *Mathematical topics in fluid mechanics, Vol. 2: Compressible models*
11. W. T. Tutte: *Graph theory as I have known it*
12. Andrea Braides and Anneliese Defranceschi: *Homogenization of multiple integrals*
13. Thierry Cazenave and Alain Haraux: *An introduction to semilinear evolution equations*
14. J. Y. Chemin: *Perfect incompressible fluids*
15. Giuseppe Buttazzo, Mariano Giaquinta and Stefan Hildebrandt: *One-dimensional variational problems: an introduction*
16. Alexander I. Bobenko and Ruedi Seiler: *Discrete integrable geometry and physics*
17. Doina Cioranescu and Patrizia Donato: *An introduction to homogenization*
18. E. J. Janse van Rensburg: *The statistical mechanics of interacting walks, polygons, animals and vesicles*
19. S. Kuksin: *Hamiltonian partial differential equations*
20. Alberto Bressan: *Hyperbolic systems of conservation laws: the one-dimensional Cauchy problem*
21. B. Perthame: *Kinetic formulation of conservation laws*
22. Andrea Braides: *Γ -convergence for beginners*
23. Robert Leese and Stephen Hurley (eds): *Methods and algorithms for radio channel assignment*
24. Charles Semple and Mike Steel: *Phylogenetics*
25. Luigi Ambrosio and Paolo Tilli: *Topics on analysis in metric spaces*
26. Eduard Feireisl: *Dynamics of viscous compressible fluids*
27. Anotnín Novotný and Ivan Straškraba: *Introduction to mathematical theory of compressible flow*
28. Pavol Hell and Jaroslav Nešetřil: *Graphs and homomorphisms*

Graphs and Homomorphisms

Pavol Hell

Simon Fraser University, Burnaby, B.C., Canada

and

Jaroslav Nešetřil

Charles University, Prague, The Czech Republic

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Bangkok Buenos Aires Cape Town Chennai
Dares Salaam Delhi Hong Kong Istanbul Karachi Kolkata
Kuala Lumpur Madrid Melbourne Mexico City Mumbai Nairobi
São Paulo Shanghai Taipei Tokyo Toronto

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© Oxford University Press 2004

The moral rights of the author have been asserted
Database right Oxford University Press (maker)

First published 2004

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose this same condition on any acquirer

A catalogue record for this title is available from the British Library

Library of Congress Cataloging in Publication Data

(Data available)

ISBN 0 19 852817 5 (Hbk)

10 9 8 7 6 5 4 3 2 1

Typeset by the authors using L^AT_EX

Printed in Great Britain
on acid-free paper by
Biddles Ltd, King's Lynn

PREFACE

This is a book about graph homomorphisms. While graph theory is now an established discipline (within the field of combinatorics), the study of graph homomorphisms is only beginning to gain wide acceptance. As little as a few years ago, most graph theorists, while passively aware of a few classical results on graph homomorphisms, would not include homomorphisms among the topics of central interest in graph theory. We believe that this perception is changing, principally because of the usefulness of the homomorphism perspective in areas such as graph reconstruction, products, and fractional and circular colourings, and applications in complexity theory, artificial intelligence, telecommunication, and, most recently, statistical physics. At the same time, the homomorphism framework strengthens the link between graph theory and other parts of mathematics, making graph theory more attractive, and understandable, to other mathematicians.

We feel that the time is ripe to introduce this exciting topic to a wider audience. It was our intention to bring together what we see as the highlights of the theory and its many applications. We hope that our book will be seen as a *sampler* of this rich theory, of its most interesting results, techniques, and applications. Sample additional results have been included in the Exercises and referred to in the Remarks. We hope that the reader will be motivated to further explore the rich literature. We have tried not to set our focus too narrowly; thus the techniques and points of view vary, from algebraic and algorithmic to applied, extremal, and randomized. The resulting lack of homogeneity means that we have had to occasionally make certain compromises on continuity, level of exposition, terminology, or organization. We hope the reader will be understanding.

One challenge we faced was the intermingling of the various versions of graphs, digraphs, and more general systems. It is typical of the area to freely jump from graphs to digraphs, allowing or disallowing loops, as is dictated by the context. We have tried to be clear at each point what is the correct context, but the reader may find it useful to keep in mind the main possibilities, illustrated in Fig. 1.1. In general, our most basic context is that of digraphs, i.e., sets with one binary relation; graphs are viewed as a subclass of digraphs. Occasionally, we shall consider more general relational systems, i.e., sets with several relations of various arities. Homomorphisms are defined the same way in all these contexts—they simply have to preserve all the relations. Most homomorphism-related concepts transfer between these contexts without difficulty—this is precisely what makes jumping between the contexts possible. Many of the results we shall discuss apply in the most general context of relational systems; however, if the generalization does not bring a new perspective, we usually just stick to

the context of digraphs, or even just graphs. One generality we have, for the most part, completely avoided is infinite graphs. By definition, all our graphs, digraphs, and relational systems are *finite*. This includes, in the case of general relational systems, the number of relations and their arities. However, it does not include sets of graphs and associated notions such as category and partial order; these sets (categories, orders) can be infinite.

If the book is to be a sampler, we have written Chapter 1 as a mini-sampler. In it, we introduce motivational examples and applications, which are usually taken up in more detail in later chapters. It gives the flavour of algorithmic aspects, to be taken up again in Chapter 5, retractions, to be further discussed in Chapter 2, duality, investigated in Chapters 3 and 5, constraint satisfaction problems, discussed in Chapter 5, as well as structural properties of homomorphism composition, to which we devote Chapter 4. The highlights of Chapter 1 include a simple proof of the Colouring Interpolation Theorem, a generalization of the No-Homomorphism Lemma, the construction of a triangle-free graph to which all cubic triangle-free graphs are homomorphic, a case of the Edge Reconstruction Conjecture, and a generalization of a theorem of Frucht on graphs with prescribed automorphism groups.

Chapter 2 focuses on certain basic constructions that occur frequently in the rest of the book, emphasizing the product and the retract, but also considering other constructs. These include the shift graph, the exponential graph, and the Lovász vector—each of which plays an auxiliary role in Chapter 2, and is a useful concept in its own right. Other basic constructions are discussed in later chapters; we mention here the replacement operation alluded to in Chapter 1 and explored in greater detail in Chapter 4, the indicator and subindicator constructions described in Chapter 5, the Kneser graphs and the rational complete graphs studied in Chapter 6, and so on. Taken together, such constructions are the common thread and the leitmotiv of this book. The highlight of the sections on products include a linear algebra based lower bound on the dimension of a graph, a stronger version of the edge reconstruction result from Chapter 1, a discussion of cancellation and unique factorization properties, a construction of graphs with arbitrarily high odd girth and chromatic number, an exploration of the Product Conjecture, and an elementary proof of the multiplicativity of oriented cycles of prime power length. In the sections on retracts, we prove that an isometric tree, and a shortest cycle, is always a retract of a reflexive graph; we prove a similar result for shortest odd cycles in irreflexive nearly bipartite graphs. We characterize absolute reflexive retracts in several different ways, including characterizations in terms of majority functions, in terms of the variety of paths, and in terms of the Helly property (or the absence of holes). We prove that a reflexive graph admits a winning strategy for the cop, in the game of cops and robbers, if and only if it is dismantlable, and relate dismantlable graphs to absolute reflexive retracts. Finally, we introduce median graphs and relate them to retracts of hypercubes; we also discuss an application of median graphs and retractions in a resource location problem.

In Chapter 3 we consider the order homomorphisms induce on the set of all cores. This order is rich enough to represent all countable partial orders. We consider antichains in the homomorphism order, i.e., collections of incomparable graphs (graphs without homomorphisms between any two of them). Of particular interest are finite maximal antichains, and their structure turns out to be surprisingly revealing. Graphs only have trivial finite maximal antichains, while digraphs have many such antichains, of all possible sizes, arising from duality relationships. This chapter also contains the (probabilistic) proof of the Sparse Incomparability Lemma, of the fact that asymptotically almost all graphs on n vertices are cores, and of the fact that the number of incomparable graphs on n vertices differs little (asymptotically) from the total number of non-isomorphic graphs on n vertices. The density of the homomorphism order is related to duality, revealing an unexpected connection between these two seemingly unrelated concepts. Finally, we show that one can gain interesting insights into many traditional graph topics, such as, say, Hadwiger’s conjecture, when interpreting them as statements about the homomorphism order.

In Chapter 4, we explore the *structure*, as opposed to just the existence, of the family of homomorphisms among a set of graphs. The difference is noticeable with even just one graph—consider, for instance, a graph having only the identity homomorphisms to itself. Such graphs are called rigid and they are the building blocks of many constructions. We construct many useful examples of rigid graphs, prove that asymptotically almost all graphs are rigid, and construct infinite rigid graphs with arbitrary cardinality. The homomorphisms among a set of graphs impose the algebraic structure of a category. We show that every finite category is represented by a set of graphs. This is the generalization of the theorem of Frucht alluded to above. Also, as in the case studied by Frucht, we show that the representing graphs can be required to have any of a number of graph theoretic properties. However, we prove that these properties cannot include having bounded degrees—somewhat surprisingly, since Frucht proved that cubic graphs represent all finite groups.

In Chapter 5, we explore algorithmic aspects of graph homomorphisms and of similar partition problems. The highlights include the dichotomy classification of graph homomorphisms to a fixed target graph H , a proof of the fact that dichotomy for digraph homomorphisms would imply dichotomy for all constraint satisfaction problems, a presentation of consistency-based algorithms, and associated dualities, that seem to be applicable to all known polynomial cases of the digraph homomorphism problem. We also discuss the use of polymorphisms for the design of polynomial algorithms, and prove that graphs with the same set of polymorphisms define polynomially equivalent problems. We explain how the polymorphism known as the majority function can be used to construct a polynomial time algorithm. We prove the dichotomy classification of list homomorphism problems for reflexive graphs. We present list matrix partition problems in the language of trigraph homomorphisms, and illustrate the richness of the associated algorithms on the case of clique cutsets and generalized split graphs.

Finally, Chapter 6 sets out certain particular classes of homomorphism problems that have been investigated as variants of colourings. The homomorphism perspective unifies these concepts and offers new questions. We include a discussion of the circular chromatic number, the fractional chromatic number, the T -span, and the oriented chromatic number. The highlights include a number of equivalent definitions of the circular chromatic number, in terms of H -colourability, in terms of a geometric representation, in terms of orientations, implying, say, Minty's result on chromatic numbers, and in terms of schedule concurrency. For fractional chromatic numbers we also give equivalent formulations, in terms of Kneser graphs, integer linear programs, and zero-sum games, and relate them in several ways to the circular chromatic numbers. We discuss homomorphisms amongst Kneser graphs and a proof of Kneser's conjecture. We prove that the span, for any set T , of the cliques K_n has a limit, closely related to the fractional chromatic number of an associated graph. We also give bounds on the oriented chromatic numbers of planar and outerplanar graphs, and relate the oriented chromatic numbers to acyclic chromatic numbers.

Our book can be used as a textbook for a second course in graph theory, at the level of a beginning graduate student. (In fact, we have used it for just such a course at Simon Fraser University, Vancouver, Charles University, Prague, Eidgenössische Technische Hochschule Zürich, Universidade Federal do Rio de Janeiro, and Universitat Politècnica de Catalunya.) Because of the relative independence of the chapters, the book can also be used as a supplementary text for a more varied course (at the same graduate or even undergraduate level). One can, for instance, just present Chapter 1, our mini-sampler. In addition, Chapter 1 can be supplemented by a sequence of combinatorial topics from Chapters 2 and 6. If time permits, a more intensive sequence could complement Chapter 1 with a selection of algebraic topics from Chapters 2, 3, and 4, or of algorithmic topics from Chapters 2 and 5.

The exercises vary in difficulty. The first few are usually intended to give the reader an opportunity to practice the concepts introduced in the chapter; the later ones explore related concepts or even introduce new ones. For the harder exercises we usually give a hint or a reference.

We thank our students, friends, and collaborators for checking some of the details in this book. Special thanks to M. Bálek, T. Feder, J. Foniok, J. Fiala, J. Huang, A. Kazda, J. Kratochvíl, L. Lovász, J. Matoušek, R. Naserasr, A. Raspaud, V. Rödl, R. Šámal, I. Švejdarová, and U. Wagner, who have made numerous suggestions. We are particularly grateful to C. Tardif and X. Zhu for their valuable input; with their participation, we are currently writing a more comprehensive follow-up book. Last but not least, we express our deep gratitude to our teachers, and pioneers of the area, Zdeněk Hedrlín, Aleš Pultr, and Gert Sabidussi.

We are the only ones to blame for any remaining errors and inconsistencies. The book was written over an extended period of time, and we can only hope that we have managed to make all the parts fit together.

We will maintain a webpage at

www.cs.sfu.ca/~pavol/hombook.html

to record any corrections found after printing, and provide other useful information and links. The reader's input would be appreciated.

Finally, we dedicate this book to all who have encouraged and inspired us in this endeavor, especially Marion Hellová, Helena Nešetřilová, Heather Mitchell, Catherine and Julia Taylor-Hell, Jakub Nešetřil, Sam and Erin Hogg, and Jan and Lenka Hřebejk.

Pavol Hell, Jaroslav Nešetřil

Vancouver, Prague, Christmas 2003.

This page intentionally left blank

CONTENTS

1	Introduction	1
1.1	Graphs, digraphs, and homomorphisms	1
1.2	Homomorphisms preserve adjacency	3
1.3	Homomorphisms generalize colourings	6
1.4	The existence of homomorphisms	10
1.5	Homomorphisms generalize isomorphisms	16
1.6	Homomorphic equivalence	18
1.7	The composition of homomorphisms	20
1.8	Homomorphisms model assignments and schedules	27
1.9	Remarks	33
1.10	Exercises	34
2	Products and retracts	37
2.1	The product	37
2.2	Dimension	40
2.3	The Lovász vector and the Reconstruction Conjecture	43
2.4	Exponential digraphs	46
2.5	Shift graphs	47
2.6	The Product Conjecture and graph multiplicativity	50
2.7	Projective digraphs and polymorphisms	57
2.8	The retract	58
2.9	Isometric trees and cycles	60
2.10	Reflexive absolute retracts	64
2.11	Reflexive dismantlable graphs	68
2.12	Median graphs	72
2.13	Remarks	76
2.14	Exercises	78
3	The partial order of graphs and homomorphisms	81
3.1	The partial orders \mathcal{C} and \mathcal{C}_S	81
3.2	Representing ordered sets	82
3.3	Incomparable graphs and maximal antichains in \mathcal{C}_S	85
3.4	Sparse graphs with specified homomorphisms	89
3.5	Incomparable graphs with additional properties	93
3.6	Incomparable graphs on n vertices	94
3.7	Density	96
3.8	Duality and gaps	98
3.9	Maximal antichains in \mathcal{C}	101
3.10	Bounds	102

3.11	Remarks	106
3.12	Exercises	107
4	The structure of composition	109
4.1	Introduction	109
4.2	Rigid digraphs	109
4.3	An excursion to infinity	115
4.4	The replacement operation	117
4.5	Categories	122
4.6	Representation	128
4.7	A combinatorial obstacle to representation	132
4.8	Some categories are not rich enough	135
4.9	Remarks	138
4.10	Exercises	139
5	Testing for the existence of homomorphisms	142
5.1	The H -colouring problem	142
5.2	Dichotomy for graphs	143
5.3	Digraph homomorphisms and CSPs	151
5.4	Duality and consistency	161
5.5	Pair consistency and majority functions	166
5.6	List homomorphisms and retractions	170
5.7	Trigraph homomorphisms	178
5.8	Generalized split graphs	183
5.9	Remarks	186
5.10	Exercises	187
6	Colouring—variations on a theme	192
6.1	Circular colourings	192
6.2	Fractional colourings	200
6.3	T -colourings	210
6.4	Oriented and acyclic colourings	212
6.5	Remarks	218
6.6	Exercises	219
	References	222
	Index	239

INTRODUCTION

In this introductory chapter we shall describe some basic views of homomorphisms, and mention a few typical examples, which should help the reader develop intuition about homomorphisms, and which illustrate some of their appeal. We will return to many of these topics in subsequent chapters.

1.1 Graphs, digraphs, and homomorphisms

Our most basic objects are digraphs. A *digraph* G is a finite set $V = V(G)$ of *vertices*, together with a binary relation $E = E(G)$ on V . The elements (u, v) of E are called the *arcs* of G . A digraph is *symmetric*, or *reflexive*, or *irreflexive*, etc., if the relation E is symmetric, or reflexive, or irreflexive, etc., respectively. Later on we shall introduce more complex systems consisting of a finite set with several relations of various arities, for instance a set with two binary relations and one ternary relation. Much of the theory extends to such systems, and sometimes they offer interesting new insights.

Symmetric digraphs are more conveniently viewed as (undirected) graphs. Formally, a *graph* G is a set $V = V(G)$ of vertices together with a set $E = E(G)$ of edges, each of which is a two-element set of vertices. If we allow *loops*, i.e., edges that only consist of one vertex, we have a *graph with loops allowed*. Finally, if every vertex has a loop, we have a *reflexive graph*.

Suppose G is a graph with loops allowed. The *corresponding symmetric digraph* (of G) is obtained from G by replacing each edge $\{u, v\}$ with the two arcs $(u, v), (v, u)$, and each loop $\{w\}$ with the arc (w, w) . This correspondence allows us to view each graph as a digraph. Specifically, a graph with loops allowed corresponds to a symmetric digraph (and conversely), a graph corresponds to an irreflexive symmetric digraph (and conversely), and a reflexive graph corresponds to a reflexive symmetric digraph (and conversely). We prefer this slightly cumbersome terminology so that we can have the basic term ‘graph’ reserved for the object most commonly investigated, i.e., an irreflexive symmetric digraph.

It is important to bear in mind that even though we speak of graphs as a different kind of objects from digraphs, we view the class of graphs, with loops allowed, as a subclass of the class of digraphs, via their corresponding symmetric digraphs. In Fig. 1.1, we indicate the relationship amongst the various graph classes, and illustrate them by examples.

There are additional natural transformations between graphs and digraphs.

Given a graph G , we may replace each edge $\{u, v\}$ with exactly one of the arcs $(u, v), (v, u)$, obtaining an *orientation* of G . (If G is a graph with loops allowed,




FIG. 1.1. Digraphs, graphs, reflexive graphs, and graphs with loops allowed.

we replace each loop $\{u\}$ with (u, u) .) An *oriented graph* is an orientation of some graph; clearly a digraph is an oriented graph if and only if it has no *symmetric pair of arcs*, i.e., no pair $(u, v), (v, u)$ for some vertices u, v .

Given a digraph G , the *underlying graph* of G is the graph with the same vertices as G , in which $\{u, v\}$ is an edge whenever at least one of $(u, v), (v, u)$ is an arc of G . Finally, the *symmetric part* of G is the graph with the same vertices as G , in which u and v are adjacent whenever both (u, v) and (v, u) are edges of G .

We shall use the usual simplified notation for arcs and edges, in which uv represents the arc (u, v) , or the edge $\{u, v\}$, depending on the context. A loop at u is written as uu . If $uv \in E(G)$, we say that u and v are *adjacent*. If G is a graph, we have $uv = vu$. If uv is an arc in a digraph, we say that u and v are adjacent *in the direction from u to v* , or that u is an *inneighbour* of v and v is an *outneighbour* of u . In any case, u and v are adjacent in a digraph as long as at least one of uv, vu is an arc; in that case we also say that u and v are *neighbours*. The number of neighbours of v (other than v) is called the *degree* of v ; the number of inneighbours respectively outneighbours of v is called the *indegree* respectively *outdegree* of v . We say that G is a *subgraph* of H , and H a *supergraph* of G , if $V(G) \subseteq V(H)$ and $E(G) \subseteq E(H)$. Also, G is an *induced subgraph* of H if it is a subgraph of H and contains all the arcs (edges) of H

amongst the vertices in G . A *clique* in a graph G is a complete subgraph of G . All other standard notions (*bipartite graph*, *planar graph*, *regular graph*, etc.), will be used in their usual sense, as in, say, [36, 339].

These definitions mean that all our graphs and digraphs, and the more general systems to be introduced later, are finite. Occasionally, we will allow infinite graphs or digraphs, in which case we make a point of explicitly mentioning this fact. They also mean that all our graphs and digraphs are *simple*, in the sense that parallel edges, or arcs in the same direction, are not allowed. Note that the definitions allow the possibility that the vertex set is empty; however, we shall normally avoid complicating the statements of theorems by considering digraphs with empty vertex sets.

The fact that we deal with various versions of graphs at various times makes a certain demand on the reader, to keep in mind the right context. This is typical of the area, and the ease with which most notions and techniques transfer amongst the variants, while others offer substantial differences, is one of the area's characteristic features.

Here is our most basic definition.

Let G and H be any digraphs. A *homomorphism* of G to H , written as $f : G \rightarrow H$ is a mapping $f : V(G) \rightarrow V(H)$ such that $f(u)f(v) \in E(H)$ whenever $uv \in E(G)$. A homomorphism of G to H is also called an *H -colouring of G* (Proposition 1.7 suggests why). If there exists a homomorphism $f : G \rightarrow H$ we shall write $G \rightarrow H$, and $G \not\rightarrow H$ means there is no such homomorphism. If $G \rightarrow H$ we shall say that G is *homomorphic to H* , or that G is *H -colourable*. It is easy to see that the composition $f \circ g$ of homomorphisms $g : G \rightarrow H$ and $f : H \rightarrow K$ is a homomorphism of G to K (cf. Section 1.7). Thus the binary relation ‘is homomorphic to’ on the set of digraphs is transitive. We denote by $\text{HOM}(G, H)$ the set of all homomorphisms $f : G \rightarrow H$, and let $\text{hom}(G, H)$ denote the number of elements in $\text{HOM}(G, H)$.

If G and H are graphs, we can apply the above definition of homomorphism to the corresponding symmetric digraphs of G and H . Clearly, this is equivalent to reading the same definition of $f : G \rightarrow H$ as a mapping $f : V(G) \rightarrow V(H)$ such that $f(u)f(v) \in E(H)$ whenever $uv \in E(G)$, with $f(u)f(v)$ and uv being edges rather than arcs. Hence homomorphisms of graphs preserve adjacency, while homomorphisms of digraphs also preserve the directions of the arcs. Therefore, a homomorphism of digraphs $G \rightarrow H$ is also a homomorphism of the underlying graphs, but not conversely.

Note that for graphs (or, more generally, for irreflexive digraphs) $f(u)f(v) \in E(H)$ implies that $f(u) \neq f(v)$, since each edge of H consists of two distinct elements.

1.2 Homomorphisms preserve adjacency

For simplicity, we begin by focusing on graphs.

The fact that homomorphisms are mappings of the vertices that preserve adjacency has interesting implications, for instance, for homomorphisms of paths

and cycles. A *walk* in a graph G is a sequence of vertices v_0, v_1, \dots, v_k of G such that v_{i-1} and v_i are adjacent, for each $i = 1, 2, \dots, k$. A walk is *closed* if $v_0 = v_k$. A *path* in G is a walk in which all the vertices are distinct. The integer k is called the *length* of the walk, respectively path.

The graph with vertices $0, 1, \dots, k$ and edges $01, 12, \dots, (k-1)k$ is called *the path* P_k . Note that P_k has $k+1$ vertices and k edges (Fig. 1.3).

Proposition 1.1 *A mapping $f : V(P_k) \rightarrow V(G)$ is a homomorphism of P_k to G if and only if the sequence $f(0), f(1), \dots, f(k)$ is a walk in G . \square*

In particular, homomorphisms of G to H map paths in G to walks in H , and hence do not increase distances. If we denote by $d_G(u, v)$ the distance (length of a shortest path) from u to v in G , then we have the following fact.

Corollary 1.2 *If $f : G \rightarrow H$ is a homomorphism, then $d_H(f(u), f(v)) \leq d_G(u, v)$, for any two vertices u, v of G .*

Proof If $u = v_0, v_1, \dots, v_k = v$ is a path in G , then $f(u) = f(v_0), f(v_1), \dots, f(v_k) = f(v)$ is a walk of the same length k in H . Since every walk from $f(u)$ to $f(v)$ contains a path from $f(u)$ to $f(v)$, we must have $d_H(f(u), f(v)) \leq k$. \square

A *cycle* in a graph G is a sequence of distinct vertices v_1, v_2, \dots, v_k of G such that each $v_i, i = 2, 3, \dots, k$, is adjacent to v_{i-1} , and v_1 is adjacent to v_k . Note that a cycle is a closed walk, and thus the definition of length is still applicable.

The graph with vertices $0, 1, \dots, k-1$ and edges $i(i+1)$ for $i = 0, 1, \dots, k-1$ (with addition modulo k) is called *the cycle* C_k . Note that C_k has k vertices and k edges.


Proposition 1.3 *A mapping $f : V(C_k) \rightarrow V(G)$ is a homomorphism of C_k to G if and only if $f(0), f(1), \dots, f(k-1)$ is a closed walk in G . \square*

Corollary 1.4 $C_{2k+1} \rightarrow C_{2l+1}$ if and only if $l \leq k$.

Proof An odd cycle has no closed odd walk shorter than its length, and has a closed walk of any odd length greater than or equal to its length. \square

Figure 1.2 illustrates a homomorphism $f : C_7 \rightarrow C_5$; the images $f(v), v \in V(C_7)$ are shown in C_5 . (Hence we see the closed walk $f(0), f(1), \dots, f(6), f(0)$ in C_5 .)

A homomorphism $f : G \rightarrow H$ is a mapping of $V(G)$ to $V(H)$, but since it preserves adjacency it also naturally defines a mapping $f^\#$ of $E(G)$ to $E(H)$ by setting $f^\#(uv) = f(u)f(v)$ for all $uv \in E(G)$. We shall call a homomorphism $f : G \rightarrow H$ *vertex-injective*, or *vertex-surjective*, or *vertex-bijective*, if the mapping $f : V(G) \rightarrow V(H)$ is injective, or surjective, or bijective respectively; and call it *edge-injective*, or *edge-surjective*, or *edge-bijective*, if the mapping $f^\# : E(G) \rightarrow E(H)$ is injective, or surjective, or bijective, respectively. Finally, a homomorphism f is an *injective homomorphism*, or a *surjective homomorphism*,

FIG. 1.2. A homomorphism $f : C_7 \rightarrow C_5$.

or a *bijective homomorphism*, if it is both vertex- and edge- injective, or surjective, or bijective, respectively. Note that a homomorphism that is vertex-injective is also edge-injective (but not conversely), and, as long as H has no isolated vertices, a homomorphism that is edge-surjective is also vertex-surjective (but not conversely). In other words, injective homomorphisms are the same as vertex-injective homomorphisms, while surjective homomorphisms are, in the absence of isolated vertices, the same as edge-surjective homomorphisms.

We denote by $\text{INJ}(G, H)$ the set of all injective homomorphisms of G to H , and let $\text{inj}(G, H)$ denote the number of elements in $\text{INJ}(G, H)$. The sets $\text{SUR}(G, H)$, $\text{BIJ}(G, H)$, and numbers $\text{sur}(G, H)$, $\text{bij}(G, H)$, are defined analogously.

An *Euler trail* in a graph G is a walk v_0, v_1, \dots, v_k of G such that $v_{i-1}v_i$, $i = 1, 2, \dots, k$, contains every edge of G exactly once. An *Euler circuit* is a closed Euler trail.

In the following proposition m denotes the number of edges of G .

Proposition 1.5 *Let G be a graph. Then*

- *an Euler trail is an edge-bijective homomorphism of $P_m \rightarrow G$, and*
- *an Euler circuit is an edge-bijective homomorphism $C_m \rightarrow G$.*

Proof Propositions 1.1 and 1.3 explain how such homomorphisms can be viewed as the corresponding trails. \square

We now turn to digraphs. The definitions of injective, surjective, and bijective homomorphisms apply verbatim. Many other definitions also extend to digraphs if we interpret, as defined above, the word *adjacent* to mean adjacent in at least one direction. Hence we may apply the definitions of walks, paths, cycles (and their lengths), to digraphs, since we have stated them in terms of adjacency. Note that these walks, paths, and cycles, do not require the edges to be oriented in the same direction; to underline this fact, we usually refer to them as *oriented* walks (respectively paths or cycles).

Specifically, an *oriented walk* in a digraph G is a sequence of vertices v_0, v_1, \dots, v_k of G such that v_{i-1} and v_i are adjacent in G , for each $i = 1, 2, \dots, k$. An arc $v_{i-1}v_i \in E(G)$ is called a *forward arc* of the walk, and an arc $v_i v_{i-1}$ is called a *backward arc* of the walk. The *net length* of a walk is the difference between the number of forward arcs and the number of backward arcs, in the walk. Note that the net length may be negative. A *directed walk* in a digraph G is an oriented walk in which all arcs are forward arcs. Oriented and directed paths and cycles are defined correspondingly. Oriented and directed walks are used to define the connectivity of digraphs. A digraph is *connected* if any two vertices are joined by an oriented walk, and it is *strongly connected* if any two vertices are joined by a directed walk (in each of the two directions).

Directed paths and cycles \vec{P}_k and \vec{C}_k are defined exactly as the graphs P_k and C_k , only this time $i(i+1)$ are arcs and not edges (see Fig. 1.3).



FIG. 1.3. The paths P_2 and \vec{P}_2 .

Much of what we have discussed for graphs also applies, with obvious changes, to digraphs. In particular, homomorphisms from \vec{P}_k are *directed walks*, since homomorphisms not only preserve adjacency but also directions of the arcs. Euler trails are directed walks that contain every arc exactly once, and they correspond to edge-bijective homomorphisms of \vec{P}_m , and similarly for closed Euler trails and edge-bijective homomorphisms of \vec{C}_m , exactly as in Proposition 1.5.

On the other hand, the notion of net length is specific to digraphs, and we have the following observation.

Proposition 1.6 *Let G and H be digraphs, and $f : G \rightarrow H$ a homomorphism. If v_0, v_1, \dots, v_k is a walk in G , then $f(v_0), f(v_1), \dots, f(v_k)$ is a walk in H , of the same net length.* \square

1.3 Homomorphisms generalize colourings

A good way to develop intuition about homomorphisms is to relate them to a notion familiar to all students of graph theory, namely vertex colourings. A k -*colouring* of a graph G is an assignment of k colours to the vertices of G , in which adjacent vertices have different colours. Denote by K_k the complete graph on vertices $1, 2, \dots, k$, and suppose that these integers are also used as the ‘colours’ in k -colourings. Then a k -colouring of G may be viewed as a mapping $f : V(G) \rightarrow \{1, 2, \dots, k\}$; the requirement that adjacent vertices have distinct colours means that $f(u) \neq f(v)$ whenever $uv \in E(G)$. It now only remains to observe that the condition $f(u) \neq f(v)$ is equivalent to the condition $f(u)f(v) \in E(K_k)$, and we may conclude the following fact.

Proposition 1.7 *Homomorphisms $f : G \rightarrow K_k$ are precisely the k -colourings of G .* \square

This observation allows us, in particular, to derive the following corollary concerning chromatic numbers. Recall that the *chromatic number*, $\chi(G)$, of a graph G , is defined as the smallest k such that G admits a k -colouring.

Corollary 1.8 *If $G \rightarrow H$ then $\chi(G) \leq \chi(H)$.*

Proof Let $h : G \rightarrow H$ be a homomorphism. Whenever H has a k -colouring $f : H \rightarrow K_k$ then $h \circ f$ is a k -colouring of G (cf. Section 1.7). Thus G has a $\chi(H)$ -colouring, and $\chi(G) \leq \chi(H)$. \square

We have a similar result concerning odd girth, Exercise 2: *if $G \rightarrow H$, then the odd girth of G is at least as large as the odd girth of H* . The *odd girth* of a nonbipartite graph G is the minimum length of an odd cycle in G ; the *girth* of a graph with cycles is the minimum length of a cycle in G .

Corollary 1.8 and Exercise 2 allow us to obtain graphs G and H that are *incomparable*, in the sense that $G \not\rightarrow H$ and $H \not\rightarrow G$. Suppose that G and H are graphs such that the chromatic number of G is greater than the chromatic number of H , and the odd girth of G is greater than the odd girth of H . Then $G \not\rightarrow H$ by Corollary 1.8 and $H \not\rightarrow G$ by Exercise 2. This is a construction repeatedly used in this book. (A typical example is Proposition 3.4.) However, it depends on the existence of graphs with arbitrarily high chromatic numbers and odd girths. The following landmark result of Erdős guarantees that there is a graph with arbitrarily high chromatic number and arbitrarily high girth (and hence odd girth).

Theorem 1.9 *For any positive integers k, ℓ there exists a graph of chromatic number k , and girth at least ℓ .*

We will prove this fact in Chapter 3, as Corollary 3.13. That construction uses random graphs; however, in Section 2.5 we will give an explicit construction of graphs with arbitrarily chromatic numbers and arbitrarily high odd girths.

The more traditional ‘static’ view defines a k -colouring of a graph G as a partition of $V(G)$ into k independent sets. (A set of vertices is *independent* in G if it contains no pair of adjacent vertices.) If f is a k -colouring as defined above, i.e., as an assignment of colours mapping $V(G)$ to $\{1, 2, \dots, k\}$, then we *associate with f* the partition θ_f of $V(G)$ into the independent sets $f^{-1}(1), f^{-1}(2), \dots, f^{-1}(k)$. Conversely, to each partition θ of $V(G)$ into independent sets S_1, S_2, \dots, S_k , we associate the mapping f that colours each vertex of the set S_i with the colour i .

This is just the standard association between mappings and partitions, and it is equally useful for dealing with homomorphisms. If G, H are digraphs, and $f : G \rightarrow H$ a homomorphism, the associated partition θ_f consists of the preimages of f , i.e., the sets $f^{-1}(x), x \in V(H)$. If there is no loop at vertex x , then the set S_x must be independent, as before. (Hence if H is a graph, or an irreflexive digraph, all parts of the associated partition are independent.) The structure of

If f is any homomorphism of G to H , then the digraph with vertices $f(v), v \in V(G)$, and arcs $f(v)f(w), vw \in E(G)$, is called the *homomorphic image of G under f* , and denoted $f(G)$. Note that $f(G)$ is a subgraph of H , and that $f : G \rightarrow f(G)$ is a surjective homomorphism. Conversely, if $f : G \rightarrow H$ is a surjective homomorphism, then $H = f(G)$. In fact, there is a close connection between quotients and homomorphic images, made explicit in the following corollary. (Isomorphisms are formally introduced in Section 1.5).

Corollary 1.11 *Every quotient of G is a homomorphic image of G , and conversely, every homomorphic image of G is isomorphic to a quotient of G .*

Proof Suppose θ is a partition of $V(G)$ into nonempty parts $S_i, i \in I$. Then the canonical mapping $f : V(G) \rightarrow I$ defined above is a surjective homomorphism of G to its quotient G/θ , and hence G/θ is the homomorphic image of G under f . Conversely, the partition θ_f associated to a surjective homomorphism $f : G \rightarrow H$ defines the quotient G/θ_f , which is isomorphic to H , via the isomorphism that assigns to each part $f^{-1}(x)$ of θ_f the vertex x of H . \square

For digraphs, the situation is analogous, and we give an example partition θ and its associated quotient (homomorphic image) in Fig. 1.5.




FIG. 1.5. A partition and its quotient, for a digraph.

A surjective homomorphism of a graph G to the complete graph K_k is called a *complete k -colouring* of G . Thus each complete k -colouring of G is associated with a partition of $V(G)$ into k nonempty independent sets, any two of which are joined by at least one edge. Note that any k -colouring of a graph G with $\chi(G) = k$ must be complete.

The following result, called the *Colouring Interpolation Theorem* is most naturally proved in the framework of associated partitions.

Corollary 1.12 *If a graph G admits a complete k -colouring and a complete ℓ -colouring then it admits a complete i -colouring for all i between k and ℓ .*

Proof Let A_1, A_2, \dots, A_k be a partition of $V(G)$ into k independent sets, and let B_1, B_2, \dots, B_ℓ be a partition of $V(G)$ into ℓ independent sets, with $k < \ell$.

Clearly, it will suffice to construct a complete $(k + 1)$ -colouring of G . For each $i = 0, 1, 2, \dots, \ell$, let $C_i = \cup_{1 \leq j \leq i} B_j$, and consider the partition θ_i of $V(G)$ into those sets from $B_1, B_2, \dots, B_i, A_1 - C_i, A_2 - C_i, \dots, A_k - C_i$, which are nonempty. The partition θ_0 has parts A_1, A_2, \dots, A_k ; the partition θ_ℓ has parts B_1, B_2, \dots, B_ℓ (since the other parts are empty). Therefore, G/θ_ℓ is isomorphic to K_ℓ , and there must exist a first subscript i , $i = 0, 1, 2, \dots, \ell$, such that G/θ_i is not k -colourable. Note that the minimality of i implies that G/θ_i is $(k + 1)$ -colourable—just colour B_i with the $(k + 1)$ -st colour. Then G/θ_i is a homomorphic image of G that admits a complete $(k + 1)$ -colouring, and hence so does G . \square

The largest k such that the graph G admits a complete k -colouring is called the *achromatic number* of G . Thus we conclude from the Colouring Interpolation Theorem, that G admits a complete k -colouring for any k between its chromatic and achromatic number.

As a last remark on quotients, we note that the concept of a minor also easily fits into this framework. Suppose G is a graph and θ a partition of $V(G)$ in which each part induces a connected subgraph of G . Then the graph G' obtained from the quotient G/θ by deleting all loops is called a *contraction* of G . A *minor* of G is a contraction of any subgraph of G .

1.4 The existence of homomorphisms

Graph colourability is not the only property that can be nicely expressed in the framework of homomorphisms. Consider the following two natural concepts for digraphs.

We say that a digraph G is *acyclic*, if it does not contain a directed cycle, and that G is *balanced*, if every cycle in G has net length zero.

The *transitive tournament* \vec{T}_k has vertices $1, 2, \dots, k$ and arcs ij for all $i < j$. (Occasionally, we may take a different set of k vertices.)

Proposition 1.13 *A digraph G with n vertices is*

- *acyclic if and only if $G \rightarrow \vec{T}_n$, and*
- *balanced if and only if $G \rightarrow \vec{P}_{n-1}$.*

Proof If $G \rightarrow \vec{T}_n$, then G cannot have a directed cycle since a homomorphism must take a directed cycle to a directed closed walk, and \vec{T}_n has no such walks. On the other hand, if G is acyclic then we label each vertex v of G by the integer $F(v) = 1 + f(v)$, where $f(v)$ is the maximum number of arcs in a directed walk in G , ending at v . The absence of directed cycles in G implies that $f(v)$ is well defined, and at most equal to $n - 1$. It is clear that if $vw \in E(G)$ then $f(v) < f(w)$, thus F is a homomorphism of G to \vec{T}_n . The partition associated with f is sometimes called a *topological sort* of G .

Similarly, it is easy to see that \vec{P}_{n-1} is balanced, and hence so is any G with $G \rightarrow \vec{P}_{n-1}$. On the other hand, if G is balanced, we may label its vertices by integers as follows. In each component of G , pick a vertex and label it 0. Once a vertex has been labeled by the integer i , label all of its outneighbours by $i + 1$

and all of its inneighbours by $i - 1$. It is easy to see that this labeling will produce unique labels (after the initial choices of one vertex per component with label 0). Indeed, if a vertex should obtain two different labels starting from the same initial vertex, then by tracing back the way the two labels propagated, we obtain two paths of different net lengths between the same pair of vertices—creating a closed walk of net length not equal to zero. It is easy to see that such a closed walk would contain a cycle whose net length is not zero. Since the labels we obtain must be consecutive, we can shift them (add the same positive integer to all labels) so that the smallest label is 0 (and the largest at most $n - 1$). Now the labels define a homomorphism $G \rightarrow \vec{P}_{n-1}$. \square

In Exercise 14, we ask the reader to consider when $G \rightarrow \vec{P}_k$ for $0 \leq k < n - 1$.

Let G be a connected balanced digraph. The procedure explained in the above proof assigns nonnegative integer labels to the vertices of G , so that if $uv \in E(G)$, then the labels of u and v are i and $i + 1$, for some nonnegative integer i . Since G is connected, and the smallest label is zero, the labels are unique. We call the label of u the *level of u* , and we call the maximum level of a vertex the *height of G* . It follows from the above proof that any walk from a vertex u to a vertex v has a net length equal to the difference between the height of u and the height of v .

Proposition 1.14 *If G and H are two balanced digraphs of the same height, then any homomorphism of G to H preserves the levels of vertices.*

Proof Suppose a vertex u of level i is mapped to a vertex $f(u)$ of level j , with $j \neq i$, under some homomorphism $f : G \rightarrow H$. If $j < i$, then consider an oriented walk in G from any vertex of level 0 to the vertex u . It is easy to see that such a walk must have net length i . Therefore, Proposition 1.6 implies that the image of this walk, under f , also has a net length i . However, in H there is no walk of net length i ending in $f(u)$, since the level of $f(u)$ is $j < i$, and levels are nonnegative. A similar consideration of oriented walks starting in u and $f(u)$, and the fact that G and H have the same height, leads to a contradiction in the case when $j > i$. \square

Next we discuss when $G \rightarrow \vec{C}_k$. (Recall that \vec{C}_k is the directed cycle with vertices $0, 1, \dots, k-1$.) In terms of the associated partition, we have the following condition. *A given digraph G satisfies $G \rightarrow \vec{C}_k$ if and only if the vertices of G can be partitioned into k independent sets S_0, S_1, \dots, S_{k-1} so that each arc of G goes from S_i to S_{i+1} for some $i = 0, 1, \dots, k-1$ (with addition modulo k) (Fig. 1.6).*

We will develop a criterion to decide when such a partition exists, i.e., when a digraph G satisfies $G \rightarrow \vec{C}_k$. The criterion turns out to be intimately related to a (polynomial time) algorithm to find such a partition. This brings us to the first discussion of algorithmic aspects of homomorphisms, featured prominently in Chapter 5.




FIG. 1.6. A partition whose quotient is the directed cycle \vec{C}_3 .

We shall again try to construct a homomorphism $f : G \rightarrow \vec{C}_k$ by labeling each vertex v with the label $f(v)$. The algorithm proceeds as above, by choosing a starting vertex in each component of G , to be labeled 0, and then whenever a vertex has been labeled i , all its outneighbours are labeled $i + 1$, and all its inneighbours are labeled $i - 1$. This time, however, all calculations are done modulo k ; hence the labels are automatically vertices of \vec{C}_k .

We wish to show that our algorithm is correct, i.e., the labels define a homomorphism $f : G \rightarrow \vec{C}_k$ if and only if such a homomorphism exists. This is best accomplished by simultaneously providing the following good characterization of graphs G with $G \rightarrow \vec{C}_k$.

Proposition 1.15 *The following statements are equivalent.*

1. *The algorithm succeeds (all vertices obtain unique labels)*
2. $G \rightarrow \vec{C}_k$
3. *the net length of every closed walk in G is divisible by k .*

Proof It should be clear that 1 implies 2 (the labels define a homomorphism of G to \vec{C}_k), and that 2 implies 3 (homomorphisms do not change the net length of a closed walk, cf. Proposition 1.6). The proof that 3 implies 1 is analogous to the proof of Proposition 1.13. The only way that our algorithm can fail is by trying to label a vertex by two different labels. Hence there must exist between two fixed vertices two walks of net lengths not congruent modulo k . Then the concatenation of the first walk with the reversal of the second walk is a closed walk of net length not divisible by k . \square

From the equivalence of 1 and 2 we obtain the following.

Corollary 1.16 *The algorithm is correct, i.e., it will find a homomorphism of G to \vec{C}_k if and only if one exists.*

From the equivalence of 2 and 3 we obtain the following.

Corollary 1.17 *A digraph G satisfies $G \rightarrow \vec{C}_k$ if and only if the net length of every closed walk in G is divisible by k .*

It is worth noting how the proof of the correctness of the algorithm and the proof of the characterization theorem are interconnected, and each acts as a tool for proving the other.

Recalling that a closed walk is a homomorphic image of a cycle, we can reformulate the last result as follows.

Corollary 1.18 *A digraph G satisfies $G \not\rightarrow \vec{C}_k$ if and only if $C \rightarrow G$, for some oriented cycle C of net length not divisible by k .*

Thus the nonexistence of one homomorphism is equivalent to the existence of certain other homomorphisms. We refer to this kind of good characterization as *homomorphism duality*, and we will return to it on several occasions, in Chapters 3 and 5. Exercise 8 suggests why one may want to use the term ‘duality’ here. A prototype example of this sort of homomorphism duality is the well-known theorem of König, characterizing bipartite graphs as graphs without odd cycles. In the above formalism it can be expressed as follows.

Corollary 1.19 *A graph G satisfies $G \not\rightarrow K_2$ if and only if $C_\ell \rightarrow G$ for some odd integer $\ell \geq 3$. \square*

We have stated it as a Corollary, since it follows from Corollary 1.18 for $k = 2$ by considering the associated digraphs.

There are many additional homomorphism duality results on digraphs. For instance, a more careful analysis of the proof of Proposition 1.13 implies the following duality of homomorphisms to transitive tournaments \vec{T}_k and homomorphisms from directed paths \vec{P}_k .

Proposition 1.20 *A digraph G satisfies $G \not\rightarrow \vec{T}_k$ if and only if $\vec{P}_k \rightarrow G$.*

Proof It is clear that $\vec{P}_k \not\rightarrow \vec{T}_k$ since \vec{P}_k has $k + 1$ vertices and \vec{T}_k has only k vertices and no directed closed walks. Hence if $\vec{P}_k \rightarrow G$ then $G \not\rightarrow \vec{T}_k$, otherwise the composition of such two homomorphisms would imply that $\vec{P}_k \rightarrow \vec{T}_k$. On the other hand, if $\vec{P}_k \not\rightarrow G$ then G is acyclic and the labeling F defined in the proof of Proposition 1.13 is a homomorphism of G to \vec{T}_k , since there does not exist in G a directed walk with k arcs. \square

Proposition 1.20 implies the following well-known fact.

Corollary 1.21 *A graph G is k -colourable if and only if there exists an acyclic orientation of G which does not contain the directed path \vec{P}_k .*

Proof If G is k -colourable, then orienting all edges from lower numbered colours to higher colours produces an acyclic orientation of G which does not contain \vec{P}_k . On the other hand, if G is not k -colourable, then any orientation \vec{G} of G satisfies $\vec{G} \not\rightarrow \vec{T}_k$, and hence $\vec{P}_k \rightarrow \vec{G}$ by Proposition 1.20. If \vec{G} is acyclic, then it must contain \vec{P}_k . \square

The corollary in fact holds for all orientations of G , i.e., we can remove the word ‘acyclic’. Indeed, for an arbitrary orientation \vec{G} of G which does not contain \vec{P}_k we can apply the above statement to a maximal acyclic subgraph G' of \vec{G} , resulting in a colouring of the underlying graph of G' . It is the easy to argue that this is a colouring of the entire graph G . While it is not difficult to prove Corollary 1.21 directly, the above interpretation sheds some additional light on it, since Proposition 1.20 replaces a global assertion about all orientations by a local assertion about each orientation.

Homomorphism duality allows us to certify the *nonexistence* of homomorphisms by the *existence* of certain other kinds of homomorphisms. This is one of very few tools available to certify the nonexistence of homomorphisms. One other tool is Corollary 1.8 (and similar results are proved in Sections 6.1 and 6.2). Another useful tool is explained in the following proposition. Let X, Y be arbitrary graphs; we denote by $n(X)$ the number of vertices of the graph X , and by $n(X, Y)$ the maximum number of vertices in an induced subgraph of X that is homomorphic to Y . A graph H is said to be *vertex-transitive* if for any vertices u, v of H some automorphism of H (i.e., a bijective homomorphism of H to H , cf. Section 1.5) takes u to v .

Proposition 1.22 *Suppose G, H, K are graphs, where H is vertex-transitive. If $G \rightarrow H$ then*

$$\frac{n(G, K)}{n(G)} \geq \frac{n(H, K)}{n(H)}.$$

Proof Suppose H_1, H_2, \dots, H_q are all the induced subgraphs of H homomorphic to K with $n(H, K)$ vertices. Since H is vertex-transitive, each vertex of H belongs to the same number, say p , of the subgraphs H_i . Thus $qn(H, K) = pn(H)$. We now fix a homomorphism $f : G \rightarrow H$ and denote by G_i the subgraph of G induced by the vertex set $f^{-1}(V(H_i))$. Since each G_i is homomorphic to K , and each vertex of G belongs to p subgraphs G_i , we must have $qn(G, K) \geq pn(G)$. The result follows. \square

In particular, when $K = K_1$, the value of $n(X, K)/n(X)$ is known as the *independence ratio* of X , and will be denoted by $i(X)$. We denote, as usual, by $\alpha(X)$ the *independence number* of the graph X , i.e., the maximum size of an independent set of vertices of X . Thus $i(X)$ is the ratio of $\alpha(X)$ to the total number of vertices of X . The following special case of the proposition is known as the *No-Homomorphism Lemma*.

Corollary 1.23 *Suppose H is a vertex-transitive graph. If $G \rightarrow H$ then $i(G) \geq i(H)$.* \square

For certain classes of graphs, homomorphisms always exists. For instance, the Four Colour Theorem asserts that all planar graphs are homomorphic to K_4 . Similarly, the theorem of Grötzsch says that all triangle-free planar graphs are homomorphic to K_3 . Consider also the following example; the construction will be generalized in Chapter 3. (A graph is *cubic* if it is regular of degree three, and *triangle-free* if it doesn't have a subgraph isomorphic to K_3 .)

Proposition 1.24 *There exists a triangle-free graph to which all cubic triangle-free graphs are homomorphic.*

Proof Let X be a set with at least 22 elements. We define a graph $U(X)$ as follows. The vertices of $U(X)$ are ordered pairs (x, T) , where $x \in X$ and T is a set of three distinct elements (a triple) of X such that $x \notin T$. Two vertices $(x, T), (x', T')$ are adjacent in $U(X)$ if T and T' are disjoint and $x \in T', x' \in T$. It is clear that $U(X)$ is triangle-free, as any vertex (x'', T'') adjacent to both adjacent vertices $(x, T), (x', T')$ would have to have $x'' \in T \cap T'$ contradicting the fact the T, T' are disjoint.

Consider now a cubic triangle-free graph G . We construct a homomorphism $f : G \rightarrow U(X)$ as follows. The graph G' obtained from G by joining any two vertices of distance at most three in G (the ‘third power’ of G) is easily seen to have all degrees at most 21. (In G a vertex has three vertices at distance one, at most 6 vertices at distance two, and at most 12 vertices at distance three.) Therefore, there exists a 22-colouring of G' , i.e., a colouring c of the vertices of G by the elements of X such that $c(u) \neq c(v)$ whenever u, v have distance at most three in G . We now define each $f(u)$ to be the pair $(c(u), c(N(u)))$, where $c(N(u))$ is the triple of colours $c(u_1), c(u_2), c(u_3)$ of the three neighbours u_1, u_2, u_3 of u . It follows that each $f(u)$ is a vertex of $U(X)$, since all of $c(u), c(u_1), c(u_2), c(u_3)$ are distinct. Moreover, if u and v are adjacent, then the colours of the neighbours u_1, u_2, u_3 of u and the colours of the neighbours v_1, v_2, v_3 of v are all distinct, since their mutual distances in G are at most three. Since u is one of v_1, v_2, v_3 , and v is one of u_1, u_2, u_3 , we have $f(u)$ adjacent to $f(v)$ in $U(X)$, and f is indeed a homomorphism. \square

In Chapter 3 we will interpret Proposition 1.24 in the context of the partial order of graph homomorphisms, and extend it in various ways; cf. Exercise 8 in Chapter 3.

Let U be any $U(X)$ with X having at least 22 elements. We can interpret Proposition 1.24 as the following ‘single-graph’ duality theorem for cubic graphs.

Corollary 1.25 *A cubic graph G satisfies $G \not\rightarrow U$ if and only if $K_3 \rightarrow G$.*

\square

This duality is simpler than those mentioned in Corollaries 1.18 and 1.19, in that the nonexistence of a homomorphism to U is equivalent to the existence of a homomorphism from one single graph, K_3 . Without degree restrictions we do not have nontrivial single-graph dualities of this sort, but we shall discuss them for digraphs in Section 3.8.

There is, however, one trivial single-graph duality that is useful to keep in mind:

$$G \not\rightarrow K_1 \text{ if and only if } K_2 \rightarrow G.$$

1.5 Homomorphisms generalize isomorphisms

Further intuition can be derived from comparing homomorphisms to isomorphisms. For digraphs G, H , an *isomorphism of G to H* is a homomorphism $f : G \rightarrow H$ which is a bijective (i.e., both vertex- and edge- bijective). In fact, it is easy to see (by counting) that it suffices to require that f be a vertex-injective homomorphism and $|V(G)| = |V(H)|, |E(G)| = |E(H)|$. (Recall that digraphs are by definition finite.) We write $G \simeq H$ if the digraphs G, H are isomorphic, i.e., if there exists an isomorphism $G \rightarrow H$. An isomorphism $f : G \rightarrow G$ is called an *automorphism* of G , and the set of all automorphisms of G is denoted by $\text{AUT}(G)$.

Recall that a digraph G is a subgraph of a digraph H provided $V(G) \subseteq V(H)$ and $E(G) \subseteq E(H)$. If G is a subgraph of H , then the *inclusion* mapping $i : G \rightarrow H$, defined by $i(u) = u$, for all $u \in V(G)$, is an injective homomorphism.

It is easy to see that every homomorphism is a composition of an injective homomorphism and a surjective homomorphism.

Proposition 1.26 *Let G, H be any digraphs. Every homomorphism $f : G \rightarrow H$, can be written as $f = i \circ s$ where s is a surjective homomorphism and i is an injective homomorphism.*

Proof Note that $f(G)$ is a subgraph of H and a quotient of G with respect to the partition θ_f . Thus we can take s to be canonical surjective homomorphism of G onto $G/\theta_f = f(G)$, and i the inclusion homomorphism of $f(G)$ in H . \square

We may refine this argument to derive a useful identity involving the number of homomorphisms $f : G \rightarrow H$. Recall that we use $\text{hom}(G, H)$ to denote the number of homomorphisms of G to H , and $\text{inj}(G, H)$ to denote the number of injective homomorphisms of G to H . We denote by Θ the set of all partitions of $V(G)$.

Corollary 1.27 *For any digraphs G and H ,*

$$\text{hom}(G, H) = \sum_{\theta \in \Theta} \text{inj}(G/\theta, H).$$

Proof We have associated with each homomorphism $f : G \rightarrow H$ the pair (s, i) , where s is the canonical surjective homomorphism onto $G/\theta_f = f(G)$, and i the natural injective homomorphism of $f(G)$ to H . Grouping together all homomorphisms f with the same s , we obtain the identity. \square

A similar identity is obtained by further grouping all homomorphisms f with the same (i.e., isomorphic) graph $f(G)$, cf. Exercise 7.

Counting injective homomorphisms turns out to be crucial in a nice application of homomorphisms to the following *Edge Reconstruction Conjecture*.

Conjecture 1.28 *Graphs G and H are isomorphic if and only if all their edge-removed subgraphs are isomorphic, i.e., if and only if there exists a bijection $\beta : E(G) \rightarrow E(H)$ such that $G - e \simeq H - \beta(e)$ for every $e \in E(G)$.*

This conjecture, and especially a corresponding vertex reconstruction conjecture, where edge-removed subgraphs are replaced by vertex-removed subgraphs, has captured the imagination of many graph theorists. It is customary to speak of a *deck* of G , meaning the collection of (nonisomorphic) subgraphs $G - e$, $e \in E(G)$, and restate the Edge Reconstruction Conjecture as follows.

Conjecture 1.29 *Two graphs are isomorphic if and only if they have the same deck.*

The conjecture fails for a few (four, to be exact) pairs of small graphs, and consequently is understood to be restricted to graphs with at least four edges. Clearly, the condition is necessary; it is the sufficiency that is the essence of the conjecture.

Since we are only making statements about isomorphic graphs, we may assume that both G and H have the same vertex set V . In fact, all graphs in this proof will have the vertex set V , and we shall describe them by simply giving their edge sets. If the desired bijective mapping β between $E(G)$ and $E(H)$ exists, we may extend it to associate *subsets* of $E(G)$ with subsets of $E(H)$ so that for any nonempty set of edges A we have $G - A \simeq H - \beta(A)$. Equivalently, for any set of edges A not equal to the entire set $E(G)$ we have the graph formed by A isomorphic to the graph formed by $\beta(A)$. This is possible, as can easily be seen by counting. Assume that A consists of $k < |E(G)| = |E(H)| = m$ edges, and let $N(G)$ denote the number of copies of A in G , and $N_e(G)$ the number of copies of A in $G - e$. Then counting in two ways we obtain

$$(m - k)N(G) = \sum_e N_e(G) = \sum_e N_e(H) = (m - k)N(H),$$

and hence the number of copies of any particular A is the same in G and H .

Later on in the book we shall prove one of the strongest known results supporting the validity of the conjecture (Theorem 2.16). In this section we illustrate the technique by proving a simpler version. The crux of both the simpler version and the full version is that the Edge Reconstruction Conjecture holds for graphs with sufficiently many edges, and both are proved by counting homomorphisms.

Theorem 1.30 *Let G, H be graphs with n vertices and $m > \frac{1}{2} \binom{n}{2}$ edges. If there exists a bijective mapping $\beta : E(G) \rightarrow E(H)$ such that $G - e \simeq H - \beta(e)$ for every $e \in E(G)$, then $G \simeq H$.*

Proof We shall apply the inclusion-exclusion principle to count the number of injective homomorphisms of G to the complement \overline{H} of H . Let $e = vw$ be an edge of G ; we say that an injective mapping $f : V \rightarrow V$ has the property \mathcal{P}_e , if $f(v)f(w)$ is an edge of H . Thus an injective mapping is an injective homomorphism of G to \overline{H} if and only if it has none of the properties \mathcal{P}_e , $e \in E(G)$. The number of all injective mappings V to V is equal to $\text{inj}(\emptyset, H)$ (recall that we give a graph by describing its set of edges, so \emptyset denotes the graph S with

$V(S) = V, E(S) = \emptyset$). Similarly, the number of injective mappings with property \mathcal{P}_e (and perhaps others) is equal to $\text{inj}(A, H)$, where $A = \{e\}$, and so on. Letting A denote an arbitrary subset of $E(G)$ (and hence an arbitrary subgraph of G), we obtain

$$\begin{aligned} \text{inj}(G, \overline{H}) &= \text{inj}(\emptyset, H) - \sum_{|A|=1} \text{inj}(A, H) + \sum_{|A|=2} \text{inj}(A, H) \\ &\quad - \sum_{|A|=3} \text{inj}(A, H) + \cdots + (-1)^m \text{inj}(G, H). \end{aligned}$$

In particular, the number of injective homomorphisms of H to \overline{H} can be calculated as follows (the sums are over subsets A of $E(\overline{H})$).

$$\begin{aligned} \text{inj}(H, \overline{H}) &= \text{inj}(\emptyset, H) - \sum_{|A|=1} \text{inj}(A, H) + \sum_{|A|=2} \text{inj}(A, H) \\ &\quad - \sum_{|A|=3} \text{inj}(A, H) + \cdots + (-1)^m \text{inj}(H, H). \end{aligned}$$

Note that, except for the last term, the subgraphs induced by particular A 's in G and in H are in one-to-one correspondence (according to our assumption that $A \simeq \beta(A)$). Hence most of the terms on the right-hand side of these expressions are pairwise the same, and we obtain

$$\text{inj}(G, \overline{H}) - \text{inj}(H, \overline{H}) = (-1)^m (\text{inj}(G, H) - \text{inj}(H, H)).$$

The left-hand side of this expression is zero as both G and H have too many edges. On the other hand, $\text{inj}(H, H) \neq 0$, as the identity mapping on H is an injective homomorphism, and hence $\text{inj}(G, H) \neq 0$ also. Thus there exists an injective homomorphism $G \rightarrow H$, and, since G, H have the same number of vertices and edges, we conclude that they must be isomorphic. \square

We close this section by pointing out that counting homomorphisms has a long tradition in graph theory. Indeed, for a fixed graph G , the function $\text{hom}(G, K_k)$ turns out to be a polynomial in the variable k , and is called the *chromatic polynomial* of G . Much effort has been devoted to understanding the properties of the chromatic polynomial.

1.6 Homomorphic equivalence

Two digraphs such that each is homomorphic to the other are called *homomorphically equivalent*. It is easy to verify that homomorphic equivalence is indeed an equivalence relation. It follows from Corollary 1.8 that homomorphically equivalent graphs have the same chromatic number.

Suppose that the digraph H is a subgraph of the digraph G . A *retraction* of G to H is a homomorphism $r : G \rightarrow H$ such that $r(x) = x$ for all $x \in V(H)$. If there

is a retraction of G to H we say that G *retracts to* H , and that H is a *retract* of G . When H is a retract of G , we have the retraction homomorphism of G to H and the inclusion homomorphism of H to G , thus G and H are homomorphically equivalent.

Figure 1.7 shows a retraction of the whole graph to the subgraph induced by the black vertices (each black vertex is mapped to itself). Clearly, K_n is always a retract of any n -colourable graph that contains it. On the other hand, K_n does not have a retract other than itself. Similarly, an odd cycle C_ℓ or any directed cycle \vec{C}_k does not have a retract other than itself.




FIG. 1.7. An example retraction.

The study of retractions has a long tradition, especially in topology, and we shall return to it in the next chapter. For the time being, we shall use them to define the following useful construct, and give an example application in a certain graph pursuit game.

A *core* is a digraph which does not retract to a proper subgraph. As we observed above, any K_n is a core, as is C_ℓ for odd ℓ , or any \vec{C}_k .

Proposition 1.31 *A digraph is a core if and only if it is not homomorphic to a proper subgraph.*

Proof If G retracts to a proper subgraph, then it is homomorphic to it. Conversely, if G is homomorphic to a proper subgraph, then let H be a proper subgraph of G with the fewest vertices to which G is homomorphic. Then H must not be homomorphic to a proper subgraph of itself, and hence any homomorphism of H to H must be an automorphism of H . Let f be a homomorphism of G to H . Thus the restriction a of f to H is an automorphism of H , and hence it has an inverse automorphism a^{-1} . Then clearly $a^{-1} \circ f$ is a retraction of G to H . \square

Note that if both subgraphs H and H' are retracts of G then there exist homomorphisms $f : H \rightarrow H'$ and $g : H' \rightarrow H$. If both H and H' are also cores, then $f \circ g$ and $g \circ f$ must be automorphisms, hence H and H' must be isomorphic. Therefore, each digraph has a unique (up to isomorphism) retract that is a core, i.e., has no proper retracts itself. We say that H is *the core* of G .

Corollary 1.32 *Every digraph is homomorphically equivalent to a unique core.* \square

If $r : G \rightarrow H$ is a retraction, then each part $r^{-1}(x)$ of the associated partition θ_r contains the corresponding vertex x ; therefore, if any vertices u, v in the parts $r^{-1}(x), r^{-1}(y)$ (respectively) have $uv \in E(G)$, then also $xy \in E(H)$. (In particular, a retraction is a surjective homomorphism.) This observation has an interesting application in the game of *cops and robbers*. The game, played on an arbitrary graph G , proceeds as follows. First, the ‘cop’ occupies a vertex v and the ‘robber’ occupies another vertex w . Thereafter, they continue taking turns, moving from their current vertex to an adjacent vertex. The whole graph, as well as the current position of both the cop and the robber is known to each player. The objective of the cop is to capture the robber, i.e., to be in the vertex occupied by the robber. The objective of the robber is to continue evading the cop. It is clear that when G is a path, then the cop has a winning strategy (consisting of moving towards the robber); in fact, the cop will win in any tree G . On the other hand, if G is a cycle of length greater than three, then the robber has an obvious winning strategy (keeping as far away from the cop as possible).

Proposition 1.33 *Suppose that H is a retract of G . If the robber has a winning strategy in H then he has a winning strategy in G .*

Proof Let r be a retraction of G to H . The strategy of the (male) robber in G will simulate his strategy in H as follows. Suppose the (female) cop begins by occupying a vertex $v \in r^{-1}(x)$, and suppose that in the game on H if she first occupied x then the robber’s response was to occupy the vertex w . Then his response in G is the same vertex w . Throughout the game, the robber will stay on the vertices of H , which allow him the greatest possible freedom to move. Regardless of which vertex of $r^{-1}(x)$ the cop is at, she is able to make moves only to a set $r^{-1}(y)$ with y adjacent to x . Thus her moves in G can be translated to moves in H . On the other hand, the replies of the robber in H can be viewed as moves in G , since H is a subgraph of G . \square

1.7 The composition of homomorphisms

Homomorphisms compose—this simple fact is an important aspect of homomorphisms, which they share with ‘morphisms’ of other structures, such as, say, group or ring homomorphisms, continuous maps of topological spaces, or monotone mappings of ordered sets. The composition of two homomorphisms, if possible (i.e., if the codomain of the first homomorphism is the domain of the second one), is also a homomorphism. Specifically, if, say, $g : G \rightarrow H, f : H \rightarrow K$ are homomorphisms, then $f \circ g : G \rightarrow K$ is also a homomorphism. This composition operation endows any set of digraphs with a structure of algebraic flavour. Formally, such a structure is best viewed from the perspective of a category, and we shall sketch how to do that in Chapter 4. But one can easily illustrate the ideas without being too formal. The point we wish to make is that this (categorical) setting offers new perspectives and motivates new questions.

We have noted that homomorphisms generalize isomorphisms, and therefore also automorphisms. In the case of automorphisms, the structure of composition has been much studied, and is best viewed using the concept of a group. Indeed, the set $\text{AUT}(G)$ of all automorphisms of a digraph G forms a group under composition, and the group embodies the ‘structure’ of the composition. What would be an analogous notion for homomorphisms?

While an automorphism is related to one single digraph, a homomorphism relates in general a pair of digraphs, and so to study the composition structure of several homomorphisms we need to take into account several digraphs.




FIG. 1.8. Three digraphs with all their homomorphisms.

Figure 1.8 depicts three very simple digraphs, A, B, C , and all the eight homomorphisms a, b, c, d, e, f, g, h that exist amongst them. In part (a) of the figure, the homomorphisms are depicted (by the labeled dotted lines) as they map the individual vertices; part (b) of the figure depicts the homomorphisms without reference to the vertices of A, B, C , revealing just the domain and codomain of each homomorphism. The homomorphisms a, d , and h are identities on the respective digraphs, A, B , and C , and will play a special role.

The composition of these homomorphisms is summarized in a ‘composition table’, which has a row and column for each homomorphism, and the entry in row x and column y is the composition $x \circ y$. While the composition of automorphisms is always defined, homomorphisms can be composed if and only if the codomain of the first is the domain of the second; thus composition is only a partial operation (whence the blank entries in Table 1.1).

Note that the table contains the information about which homomorphisms can be composed (by the presence or absence of an entry), and hence also about their domains and codomains.

Indeed, suppose we are looking at a table such as Table 1.1 that records

	a	b	c	d	e	f	g	h
a	a							
b	b							
c	c							
d		b	c	d	e			
e		c	b	e	d			
f		g	g	f	f			
g	g							
h						f	g	h

Table 1.1 *The composition table for Fig. 1.8.*

the compositions of all homomorphisms of a set of digraphs, but that we don't have the information about what digraphs and what mappings they represent. What can we recover from the table? The key is to begin with the identity homomorphisms. Whatever were the digraphs that produced the table, each of them admitted an identity homomorphism to itself. Thus we can recover the digraphs involved—by associating one digraph to each identity in the table. In our table there are three identities— a , d and h —and we associate three digraphs— A , B and C —with them. The other elements of the table must correspond to other homomorphisms of these digraphs. We can easily detect their domains and codomains: the domain of an element x is the digraph associated with the identity i for which $x \circ i$ is defined, and the codomain of x is the digraph associated with the identity j for which $j \circ x$ is defined. For instance, b can compose with a on the right and d on the left—hence b has domain A and codomain B .

To capture the structure of the composition of automorphisms, we abstract away the structure of the digraphs and view the automorphisms as elements of an abstract group—satisfying just the axioms of a group. Thus automorphisms were taken as an abstract notion, without reference to a mapping, as was the composition on them; the axioms of a group was all that was required. In a sense, we have encoded the structure of composition in the multiplication table of an abstract group. The group *is* the multiplication table (satisfying the usual group axioms of associativity, having an identity, and each element having an inverse). We can proceed similarly for homomorphisms—noting that each homomorphism has a 'domain' and a 'codomain', which are some kind of 'objects', and that there is a partial binary operation, 'composition', which applies if and only if the domain of the second morphism is the codomain of the first one. We can then list the axioms that the morphisms have to fulfill—axioms that are, of course, motivated by the properties of homomorphisms, such as the associativity of composition and the existence of identities. We shall do just this in Chapter 4. This will amount to axiomatizing what a table of composition (such as Table 1.1) will have to satisfy in order to correspond to the homomorphisms of a set

	i	f	l	r	c	c'
i	i	f	l	r	c	c'
f	f	i	r	l	c'	c
l	l	l	l	l	c	c
r	r	r	r	r	c'	c'
c	c	c	c	c	c	c
c'	c'	c'	c'	c'	c'	c'

Table 1.2 A full composition table is a monoid.

of digraphs.

To show the flavour of that definition, and the questions that these notions motivate, we shall look in this introduction at the special case of homomorphisms of *one digraph* G . In this case, there are only the homomorphisms of G to G —which we call *endomorphisms of G* . Thus the composition is a *full* binary operation—defined for any two endomorphisms of G —and the essential features of the composition structure simplify to just an associative binary operation with an identity.

We define a *monoid* to be a finite set M with a special element $1_M \in M$ (called the *identity* of M), and an abstract binary operation \circ that associates to each $m, n \in M$ the element $m \circ n \in M$, such that for any $m, n, p \in M$, we have

$$(m \circ n) \circ p = m \circ (n \circ p) \text{ and } 1_M \circ m = m \circ 1_M = m.$$

(Thus a monoid is just a semigroup with identity.) Note that, continuing with our focus on finite structures, we defined a monoid to be finite. A *group* is a monoid in which each element m has an inverse m^{-1} with $m \circ m^{-1} = m^{-1} \circ m = 1_M$. (In particular, we also consider only finite groups.)

We note that it follows from the definition that the identity in a monoid is unique. (If i and j are identities, then $i = i \circ j = j$.)

The \circ -operation table of a monoid has no blank entries; in fact, it looks very much like the multiplication in a group, except some elements may lack inverses. For instance, we give below an example monoid on six elements, i, f, l, r, c, c' ; the element i is the identity.

We see in the example that i is the identity of the monoid, and that i and f are invertible, while l, r, c, c' are not.

Let $\text{END}(G)$ denote the set of all endomorphisms of G , for an arbitrary digraph G . It is clear that $\text{END}(G)$ is a monoid under composition. It is called the *endomorphism monoid* of G .

Two monoids M, M' are *isomorphic* if there exists a bijective mapping $b : M \rightarrow M'$ such that $b(m \circ n) = b(m) \circ b(n)$ for all $m, n \in M$. Note that it follows that $b(1_M) = 1_{M'}$.

Figure 1.9 shows a digraph G for which $\text{END}(G)$ is isomorphic to the monoid given in Table 1.2. The endomorphisms corresponding to i, f, l, r are depicted

by the labeled dashed lines. The endomorphisms corresponding to c, c' are the constants taking all vertices to the left respectively right vertex with a loop. Note that the automorphisms of G are precisely the invertible elements—in this example i and f . In other words, a monoid M is a group if and only if each element m is invertible ($m \circ n = n \circ m = 1_M$ for some n).




FIG. 1.9. A digraph with the given endomorphism monoid.

It is a basic result of R. Frucht that *every group is isomorphic to the automorphism group of a graph*. In other words, we can ‘represent’ every abstract group—a multiplication table—by the group of automorphisms of some graph. Theorem 1.34 below extends Frucht’s result to *endomorphisms of digraphs*. In Chapter 4 we shall further extend this result in two ways: we shall replace digraphs by *graphs*, properly generalizing Frucht’s theorem, and also from monoids, which are one-object categories, to all finite categories. We shall, moreover, show that the graphs can be required to have some typical graph theoretic properties.

Theorem 1.34 *Every monoid is isomorphic to the endomorphism monoid of a suitable digraph G .*

The proof resembles the usual proof of Frucht’s theorem for automorphism groups in several aspects.

We first establish the result for a more flexible structure. A *binary relational system* S is a finite set $V = V(S)$, together with a finite family of binary relations $R_i(S), i \in I$, on V . To specifically mention the finite set of indices I , we sometimes call such an S a *binary I -system*. If I has one element, S is a digraph; otherwise we can visualize S as a digraph with arcs of various ‘colours’, as illustrated in Fig. 1.10. Note that a pair of vertices may be joined by several arcs of different colours (belonging to different relations). Elements of V are called *vertices* and elements of R_i *arcs of colour i* .

For binary I -systems S, T , we define a homomorphism $f : S \rightarrow T$ as a mapping $f : V(S) \rightarrow V(T)$ such that $(f(u), f(v)) \in R_i(T)$ whenever $(u, v) \in R_i(S)$. (We emphasize that S, T must have the same index set I .) Note that the homomorphisms of digraphs viewed as binary systems are precisely the homomorphisms of digraphs as previously defined. Now many of our other definitions

apply easily to binary systems, including the concept of endomorphism monoid $\text{END}(S)$ of a binary system S .

We now prove the following first step towards Theorem 1.34.

Theorem 1.35 *Every monoid M is isomorphic to the endomorphism monoid of a suitable binary relational system S .*

Proof The proof has two stages. In the first stage we represent M by a set of mappings on some set V , i.e., find mappings $L_m, m \in M$, on V , so that the set $L_m, m \in M$ is a monoid under composition, which is isomorphic to M . This is accomplished by taking $V = M$, and letting L_m be the ‘left multiplication’ by m (taking x to $m \circ x$ for each $x \in V$). It is easy to see that $L_{m \circ n} = L_m \circ L_n$ (they have values $(m \circ n)(x) = m(n(x))$ on an arbitrary $x \in V$), for any $m, n \in M$. Thus the representation $m \rightarrow L_m$ is an isomorphism.

In the second stage we add a finite set of binary relations $R_i, i \in I$, on the set V , whose role will be to ‘select’ as homomorphisms exactly the mappings $L_m, m \in M$, chosen to represent the elements of the monoid M . In other words, we wish the endomorphisms of the resulting relational system S to be precisely the mappings $L_m, m \in M$. We again let $I = M$, and define each relation R_i to be the ‘right multiplication’ by i (consisting of all ordered pairs $(x, x \circ i), x \in V$). Now each left multiplication L_m preserves all relations R_i , since $(a, b) \in R_i$ means $b = a \circ i$ and hence $m \circ b = m \circ (a \circ i) = (m \circ a) \circ i$ which in turn means $L_m(a)L_m(b) \in R_i$. Furthermore, no other mapping preserves all edge sets $R_i, i \in M$. If π is a mapping of V to itself which preserves each R_i , then $(\pi(1_M), \pi(i)) \in R_i$ (as $(1_M, i) \in R_i$ by definition), and so $\pi(i) = \pi(1_M)i$, for all $i \in V$, i.e., $\pi = L_{\pi(1_M)}$. Consequently, the endomorphisms of the system are precisely the mappings L_m , which we have shown to form a monoid isomorphic to M . \square

The role of the binary relational systems is auxiliary—once we have a binary relational system S with the desired endomorphism monoid, we can transform it into a digraph G . This is the third and final stage of constructing a digraph with a prescribed monoid of endomorphisms. The idea is to choose for each relation R_i a suitable digraph J_i , and to replace each arc of R_i by a different copy of J_i . Each J_i comes with two specified vertices j_i, j'_i , and we replace an arc of relation R_i by a copy of J_i so that j_i is identified with the beginning of the arc, and j'_i with the end of the arc. It is easy to see that each endomorphism of the binary system gives rise to a corresponding endomorphism of the resulting digraph G , in which copies of the *replacement digraphs* map to each other according to the mapping of the arcs of the various colours. We want to be careful when choosing the replacement digraphs J_i , so that there are no other endomorphisms of the final digraph G . In other words, every endomorphism of G *arises* (in the above way) *from an endomorphism of S* . For this purpose, we shall choose the digraphs J_i in such a way that each homomorphism of J_i to G takes it identically onto a copy of the same J_i . We devote Section 4.4 to analysing this *replacement operation*, and we view it as one of our basic tools. We can, for instance, make the resulting

digraph G have certain desired properties by carefully choosing the replacement digraphs J_i . This will allow us to claim results such as *every monoid (respectively group) is isomorphic to the endomorphism monoid (respectively automorphism group) of a three-colourable graph; or of a balanced digraph; etc.* For the time being we shall give an example family of replacement digraphs that suits our purpose.




FIG. 1.10. A binary relational system S with two relations, two replacement digraphs, and the resulting digraph G .

The replacement digraph J_i has a directed path v_1, v_2, \dots, v_i , all vertices of which are outneighbours of a vertex u . Moreover, the vertex j_1 has as outneighbours u, v_1 , and the vertex j'_1 has as outneighbours vertices u, v_i . In Fig. 1.10, we illustrate the replacement construction. The figure shows a binary relational system S with two relations, and for each of the relations R_i (one represented by solid lines and the other by dotted lines) the replacement digraph J_i . After the replacement operation is performed, we obtain the last digraph in the figure. Note in particular the effect of identifying the vertices j_i and j'_i of J_i when the

replacement is applied to a loop. We shall denote by J_i^* the homomorphic image of J_i when j_i and j'_i are identified. (In other words, the quotient of J_i with respect to the partition in which $\{j_i, j'_i\}$ is the only nonsingleton part.) These replacement digraphs J_i have the following property.

Lemma 1.36 *Let S be any relational system with two binary relations, and let G be obtained from S by the replacement operation, using the above replacement digraphs J_1, J_2 . Then every endomorphism of G arises from an endomorphism of S .*

Proof First we notice that each J_i , as well as each J_i^* , is very highly connected—any two vertices can be joined by an oriented path x_1, x_2, \dots, x_k in which $x_i x_{i+2}$ or $x_{i+2} x_i$ is an arc for each $i = 1, 2, \dots, k - 2$. (In other words, we can join any two vertices by a sequence of oriented triangles, with consecutive triangles sharing an arc.) This high connectivity must be preserved under homomorphisms, hence any homomorphism of a J_i or a J_i^* to G must take it to a subgraph of some J_k or J_k^* .

We claim that a homomorphism $J_i \rightarrow J_k$ is only possible when $i = k$, and in that case it must be the identity. Similarly, we claim that $J_i^* \rightarrow J_k^*$ implies that $i = k$ and the mapping is the identity. Finally, we claim that there is no homomorphism of $J_i^* \rightarrow J_k$ for any i, k , and that $J_i \rightarrow J_k^*$ implies that $i = k$ and the homomorphism is the canonical mapping onto the quotient. It will follow from these claims that each endomorphism of G faithfully mimicks an endomorphism of S .

To prove the claims we note that each J_i, J_i^* is irreflexive and acyclic. Hence we only need to consider possible irreflexive and acyclic quotients of the J_i, J_i^* . Hence no two vertices joined by a directed path can be identified by a homomorphism of some J_i or J_i^* to some J_j or J_j^* . Since J_i^* has a Hamiltonian directed path, it does not admit any proper irreflexive acyclic quotient. Similarly, J_i has a directed path joining any two given vertices, except for j_i, j'_i . Thus all homomorphisms amongst the digraphs J_i, J_i^* , are either injective, or are the canonical homomorphisms $J_i \rightarrow J_i^*$. It is easy to check that no J_i or J_i^* is a subgraph of another J_k or J_k^* , and that the only automorphism of each J_i and J_i^* is the identity. \square

The lemma ensures that the endomorphism monoid of G is isomorphic to the endomorphism monoid of S , and hence Theorem 1.34 follows from Theorem 1.35. \square

A digraph is *rigid*, if it has no endomorphism other than the identity. We claimed in the above proof that the replacement digraphs J_i are rigid. Rigid digraphs and graphs are discussed in detail in Chapter 4.

1.8 Homomorphisms model assignments and schedules

A homomorphism of G to H is an assignment of vertices of H to the vertices of G , satisfying the adjacency constraints. Thus it is a useful model for situations

where we need to assign one kind of objects to another kind of objects, while preserving some sort of adjacency.

The simplest assignment problems can often be modeled just with graph colourings. For instance, in the *exam scheduling problem* one is assigning exam periods (colours) to courses, so that courses with common students obtain different colours. By defining the *course graph* G to have the courses as vertices and two courses adjacent just if they have a student in common, we model the exam scheduling problem as a problem of colouring the course graph G . (Typically, one wants to minimize the number of colours, i.e., the number of required exam periods.)

In more complex situations, the constraints cannot be modeled just by graph colouring. There is a general paradigm for modeling problems where constraints are to be satisfied—the constraint satisfaction problem (CSP). Most assignment and scheduling problems can be modeled as CSPs. In our terminology we can formulate the constraint satisfaction problem as follows.

A *general relational system* S is a finite set $V = V(S)$ (the *vertices of* S), together with a finite set of relations (the *relations of* S) $R_i(S), i \in I$, where $R_i(S)$ is a k_i -ary relation on V . The finite set I and the integers $k_i, i \in I$, describing the arities of the relations form the *pattern* (or *type*) P of the general system S . (Binary relational systems are general relational systems in which all $k_i = 2$.) For general relational systems S, T with the same pattern, we define a *homomorphism* $f : S \rightarrow T$ to be a mapping $f : V(S) \rightarrow V(T)$ such that $(f(v_1), f(v_2), \dots, f(v_{k_i})) \in R_i(T)$ whenever $(v_1, v_2, \dots, v_{k_i}) \in R_i(S)$. Note that the homomorphisms between relational systems which are binary coincide with the homomorphisms of binary relational systems as defined in Section 1.7 (and thus the same holds for systems which are digraphs or graphs).

A constraint satisfaction problem (or CSP) is simply the problem of finding a homomorphism between two general relational systems S, T .

For instance, consider a crossword puzzle in which the words to be entered are given, but without specifying their exact locations. This is in effect an assignment problem—we wish to assign to each ‘block’ of the puzzle (a horizontal or vertical line of consecutive blank cells) one of a given set of words. We construct two general systems. In the *block system* S the vertices are the blocks of the puzzle. There is a unary (arity one) relation R_i consisting of all blocks of length i (for all possible i). There is also, for each possible a and b , a binary relation $R_{a,b}$ consisting of all pairs of blocks u, v such that the a -th cell of block u coincides with the b -th cell of block v (this happens only if one of the blocks a, b is horizontal and the other vertical). Finally, there is another binary relation D consisting of all pairs of different blocks. In the *word system* T the vertices are the given words. The unary relation R'_i consists of all words of length i . The binary relation $R'_{a,b}$ consists of all pairs of words u, v such that the a -th letter of word u is the same as the b -th letter of word v . Finally, the binary relation D' consists of all pairs of different words. A homomorphism of S to T must take different blocks to different words, observe the length, and match the common letters of intersecting words.

Therefore, *the solutions of the crossword puzzle are precisely the homomorphisms of S to T .*

In this book, we shall always consider constraint satisfaction problems in which the second relational system (T) is fixed. Thus let T be a fixed general relational system with the pattern P , called the *template*. The *constraint satisfaction problem, with template T* , written as T -CSP, is the decision problem in which we ask whether or not an input general relational system S with the same pattern P admits a homomorphism to T .

Consider, for instance, the fixed system N with vertices 0, 1, and one ternary relation $R = \{(1, 0, 0), (0, 1, 0), (0, 0, 1), (1, 1, 0), (1, 0, 1), (0, 1, 1)\}$. We can view an input system S (with the same pattern, i.e., one ternary relation) as follows. Each vertex of S corresponds to a Boolean variable; value 0 corresponds to mapping the vertex to 0 in N and value 1 to mapping it to vertex 1 of N . In this sense, the triples in the ternary relation of S correspond to clauses, where each clause is to have a true and a false variable. In other words, N -CSP is precisely the problem NOT-ALL-EQUAL 3-SAT without negated variables.

To model this way the better known *three-satisfiability problem* 3-SAT, we need several ternary relations, since in this problem the disjunctive clauses may contain negations of the variables. Each clause has precisely three literals. A literal is either a variable or a negation of a variable. We define the relational system T with vertices 0, 1 and eight ternary relations $R'_{000}, R'_{001}, \dots, R'_{111}$, where

$$\begin{aligned} R'_{000} &= \{(i, j, k) : (i, j, k) \neq (0, 0, 0)\} \\ R'_{001} &= \{(i, j, k) : (i, j, k) \neq (0, 0, 1)\} \\ &\vdots \\ R'_{111} &= \{(i, j, k) : (i, j, k) \neq (1, 1, 1)\}. \end{aligned}$$

Now, each instance of 3-SAT, i.e., a set of disjunctive clauses of three literals each, corresponds to a relational system S whose vertices are the variables, and whose relations R_{abc} are defined as follows.

$$\begin{aligned} R_{000} &\text{ consists of all ordered triples } (x, y, z) \text{ for which } x \vee y \vee z \text{ is a clause;} \\ R_{001} &\text{ consists of all ordered triples } (x, y, z) \text{ for which } x \vee y \vee \bar{z} \text{ is a clause;} \\ &\vdots \\ R_{111} &\text{ consists of all ordered triples } (x, y, z) \text{ for which } \bar{x} \vee \bar{y} \vee \bar{z} \text{ is a clause.} \end{aligned}$$

It is easy to see that satisfying truth assignments are precisely the homomorphisms of S to T . Note that the system T is fixed for the entire problem 3-SAT, while each system S corresponds to one instance of 3-SAT. In this way, we have modeled 3-SAT as T -CSP.

Sometimes the problem T -CSP is also called the *homomorphism problem* with respect to the template T . In particular, when the template T has just one binary relation, i.e., T is a digraph (or a graph if the relation is symmetric), we always refer to the problem as the homomorphism problem for digraphs (or graphs). In the special case when T is the complete graph K_n , the homomorphism problem is precisely the well-known problem of *graph n -colourability*. By analogy, we sometimes denote the problem T -CSP also as T -COL. This is always the notation we use when T is a digraph (or graph); in that case it is customary to use H in place of T .

We have seen that graph colourings can model some assignment and scheduling problems, but for the most general problems we need constraint satisfaction problems (i.e., homomorphisms of general relational systems). Graph homomorphisms are a robust subclass of CSPs. It can be argued that graph homomorphisms are more applicable and flexible than graph colourings, yet retain much of their combinatorial appeal and elegance. On the other hand, graph homomorphism problems, while not as flexible as the general constraint satisfaction problems, exhibit much of CSP's complexity and modeling power.

One of the important open problems in theoretical computer science asks whether or not, for each template T , the problem T -CSP is polynomial time solvable or NP -complete. Such *dichotomy* is only known for a few subclasses of CSPs, including the homomorphism problem for undirected graphs (the template T has one symmetric binary relation), which we shall discuss in Section 5.2. The *Dichotomy Conjecture* (Conjecture 5.13) states that dichotomy holds for all CSPs. The important role homomorphisms play is underscored by the following seminal result of T. Feder and M. Vardi, which will follow from Theorem 5.14 proved in Chapter 5.

Theorem 1.37 *If there is dichotomy for the homomorphism problem for digraphs then there is dichotomy for all CSPs.*

Let us now consider some assignment problems which are naturally modeled by graph homomorphisms.

A typical example is the problem of assigning frequencies to transmitters for wireless communication. Since the range of usable frequencies is small (say 3kHz–300GHz), the International Telecommunication Union, and its national homologues, apportion bands of frequencies to various applications and operators. (For instance, the range of 300MHz–3GHz is reserved for cellular phones, global positioning systems, and UHF television.) An operator acquires a band of frequencies and divides them into channels; thus different channels ensure a certain minimum separation of frequencies. Let us focus on the example of a GSM cellular phone network in a typical European country; in today's technology, it may have on the order of 100 channels available. These channels are assigned to the so called transmitter/receiver units (TRXs). A typical cell of the system (corresponding to an antenna of a base transceiver station) may have up to about a dozen TRXs, and the whole network may have on the order of

10,000 TRXs. Hence the operator must find a way to reuse the channels, in a way which minimizes interference. This is known as the *frequency assignment problem* or *channel assignment problem*.

Interference will surely happen on two TRXs on the same antenna; these must receive different channels. Interference may also happen, to a smaller degree, on two TRXs on different antennas but in the same base transceiver station, or located in a particular way with respect to a geographical feature, etc. Furthermore, there may be interference even if two such TRXs operate on adjacent, or nearby, channels, or on a channel and its harmonic. There may also be restrictions on which channels a TRX may use—for instance at the edge of the operator's territory there may be constraints imposed by a neighbouring country or operator.

Trying to extract a simplified model of all these requirements is not easy. However, the following simplification (obtained, basically, just by ignoring the degree of interference) typifies the problems handled in practice. Let G be the graph of the transmitters, where the vertices are the TRXs and two TRXs are adjacent if they are located so that there could be interference. Let H be the graph of the frequencies, where the vertices are the channels and two channels are adjacent if they are separated so they cannot interfere with each other. Now *a suitable assignment of frequencies to transmitters is precisely a homomorphism of G to H .*

The graphs H that arise in practice are often very simple; for instance the vertices (channels) may be $1, 2, \dots, N$ and the edges may be all ij with $|i - j|$ not from a specified sets of forbidden integers. Homomorphisms to graphs H of this type are called T -colourings; we discuss them in Section 6.3.

Suppose G, H are graphs and each vertex $v \in V(G)$ has a list (set) $L(v) \subseteq V(H)$. A *list homomorphism* of G to H , with respect to the lists L , is a homomorphism $f : G \rightarrow H$ such that for each $v \in V(G)$ we have $f(v) \in L(v)$. If each TRX has a (possibly) restricted list of allowed channels, we are looking for a list homomorphism of the graph of transmitters to the graph of channels.

List homomorphisms naturally arise in a number of other assignment problems. A very simple example of this is as follows. Given a set of tasks T and a set of processors P , find an assignment of the tasks to the processors with the property that the tasks that need to communicate frequently are executed on processors that are directly connected. Such an assignment is clearly just a homomorphism of the appropriate graphs. If there are further constraints, requiring each task $t \in T$ to be assigned to one of a 'list' $L(t) \subseteq P$ of allowed processors, we need again a list homomorphism. Figure 1.11 shows an example set of tasks, processors, and allowed lists, as well as a possible assignment, indicated by the black dots.

We note that in practice, there may be other constraints, such as having roughly the same number of tasks assigned to each processor, etc.

The task graph T has the tasks as vertices and two tasks are adjacent just if the tasks need to communicate frequently. The processor graph P has processors




FIG. 1.11. Assigning tasks to processors.

as vertices and two processors are adjacent just if they are directly connected. The processor graphs are reflexive by convention. A *suitable assignment* is a *homomorphism of the task graph T to the processor graph P* . (Note how reflexivity is a natural assumption here—tasks that need frequent communication can certainly be scheduled on the same processor.) If there are list constraints on which tasks can be executed by which processors, we obtain a list homomorphism. In Section 5.6, we shall investigate algorithms for finding homomorphisms and list homomorphisms. It will turn out, for example, that this problem can be solved in linear time, if the processor graph is an *interval graph*, i.e., the intersection graph of a family of intervals on the real line. (Note that we take an interval graph to be reflexive, since each interval intersects itself. It is easy to see that the graph in Fig. 1.11 is an interval graph.)

Recently, graph homomorphisms have been found useful to model configurations in statistical physics. A configuration is an assignment of ‘spins’ or ‘states’ to the sites of a (usually very symmetric, crystal-like) structure, such that certain constraints are satisfied. For instance, in the *Widom–Rowlinson gas model* with three kinds of particles, there is a regular grid of sites each possibly containing one of the particles of types a, b, c . A configuration in this model must not have two particles of different types in adjacent sites. Of course, there may be sites that are not occupied by any particle. Thus a configuration is an assignment of labels a, b, c , or ‘blank’ to the sites. (The label ‘blank’ represents an unoccupied site.) In Fig. 1.12 we depict a (reflexive) graph H which has one vertex for each kind of particle, and a central vertex adjacent to all other vertices, which is unlabeled. An assignment of labels to a given graph of sites which satisfies the adjacency constraint, i.e., a *configuration in the three-particle Widom–Rowlinson gas model*, is precisely a homomorphism of the graph of sites to H .

In statistical physics much effort has gone into counting the number of configurations, finding its generating functions, or at least obtaining good estimates. A number of individual cases, like the Widom–Rowlinson configurations, have been studied more or less independently under various names. Graph homomorphisms offer a unifying perspective. Nevertheless, these problems are inherently




FIG. 1.12. The graph H for the three-particle Widom–Rowlinson gas model, and an example configuration.

very difficult, even for the simplest of models, cf. Exercise 9. Finally, we note, as we observed at the end of Section 1.5, that counting colourings has a long tradition in graph theory, in the theory of chromatic polynomials. All these problems can be viewed as special cases of counting homomorphisms.

1.9 Remarks

The study of graph homomorphisms is over forty years old. It was pioneered by G. Sabidussi [308], and by Z. Hedrlín and A. Pultr [146]. (We are proud to count all three as our teachers.) Since digraphs are just binary relations, there were earlier developments for general mathematical structures [125]; however, by focusing on such a restricted world, interesting new problems arose. There was also an earlier paper in graph theory by K. Čulík [70], who set himself the same goal of investigating properties of general structures restricted to graphs; however, he chose a different definition of homomorphisms. (It is explored in Exercise 10.) The term homomorphism was initially also used for minors [78]. More background on history is given in Section 4.9.

The Colouring Interpolation Theorem, Corollary 1.12, is from [138]; it seems to find its proper setting in the context of universal algebra, [241, 326], cf. also [34]. For achromatic numbers see [94, 184] and Exercise 16. Corollary 1.21, for which we presented an unusual proof, is a theorem of T. Gallai [113] and B. Roy [304], which was in fact proved earlier by L. Vitaver [333] and M. Hasse [139]. The No-Homomorphism Lemma, Corollary 1.23, is due to M. Albertson and K. Collins [2]; its extension in Proposition 1.22 is from [35]. In Chapter 3 we shall view such results (including Exercise 2) as describing monotonicity of the various graph parameters with respect to the so called homomorphism order.

Proposition 1.24 is from [133]; its meaning in the homomorphism order, together with extensions from [81], will be discussed in Chapter 3. (See also Exercise 8 in Chapter 3.) The algorithmic implications of Corollary 1.25 are explored in Exercise 16 in Chapter 5. The Edge Reconstruction result in Theorem 1.30 is due to L. Lovász [227]; in Chapter 2 we discuss an extension from [250], which more substantially uses graph homomorphisms. The game of cops and robbers, and its connection to retractions in reflexive graphs, was independently investigated by R. Nowakowski and P. Winkler [286] and by A. Quilliot [298]. The importance of cores in the study of graph homomorphisms was independently discovered by several people (who gave them different names such as ‘minimal graphs’ and ‘unretractive graphs’) [108, 169, 200, 287]; our approach follows [169]. (We remark that for infinite graphs there are complications; even Proposition 1.31 no longer holds [31].) The composition structure of graph homomorphisms was systematically investigated by the so called Prague School of Category Theory, see Section 4.9. Theorem 1.35 is from [87], Theorem 1.34 is due to Z. Hedrlín and A. Pultr [147], and the corresponding result of R. Frucht on automorphism groups, which it generalizes, is from [112]. Graphs with given automorphism groups and prescribed graph properties were studied by G. Sabidussi [305], cf. also [189]. Constraint satisfaction problems have been identified as a useful modeling tool since [248], and their formulation as homomorphism problems for general relational systems is due to T. Feder and M. Vardi [105, 106]. The Dichotomy Conjecture and Theorem 1.37 is also due to T. Feder and M. Vardi [105, 106]. It has been verified, for instance, for all T -CSPs where T has only two or three vertices [48, 313]. For frequency assignments, we followed [88], and for a graph theoretic perspective on the applications in statistical physics, we suggest [45, 46]. Much general information on graph homomorphisms, often complementing the results discussed here, can be found in the survey [135], and the book [122]. For general reference we recommend the books [36, 195, 236, 339].

Exercise 9 is from [45]; it corresponds to the so called *hardcore gas model*, and illustrates how difficult counting homomorphisms can be. (Counting independent sets is an intractable problem.) Exercise 15 is from [162]; it is shown there that there are only a few noncomplete graphs H with this property. (For complete graphs, see Exercise 16, and also Exercise 27 in Chapter 5). Exercise 18 is also discussed in [135]; see also [335].

1.10 Exercises

1. Prove that for cycles we have $C_k \rightarrow C_\ell$ if and only if k is even, or ℓ is odd and $\ell \leq k$.
2. Prove that $G \rightarrow H$ implies that the odd girth of G is at least as large as the odd girth of H . Does such a statement hold for girth?
3. Prove that a graph G has an odd cycle if and only if $C_k \rightarrow G$ for an odd k . Does such a statement hold for even cycles?

4. Prove that for directed cycles we have $\vec{C}_k \rightarrow \vec{C}_\ell$ if and only if ℓ divides k . Deduce that the set of directed cycles \vec{C}_p , where p is a prime, is *incomparable*, i.e., $\vec{C}_p \rightarrow \vec{C}_q$ if and only if $p = q$.
5. Prove that the core of a connected graph is connected. Prove that the core of a vertex-transitive graph is vertex-transitive.
6. Suppose G is a vertex-transitive graph in which the number of vertices and the maximum size of an independent set are relatively prime. Prove that G is a core.
Suppose G is a *vertex-critical graph*, i.e., that the chromatic number of G is greater than the chromatic number of any subgraph $G - u$. Prove that G is a core.
7. Show that if $f : G \rightarrow H$ is an injective homomorphism, then G is isomorphic to the subgraph $f(G)$ of H .
Prove that $\text{hom}(G, H) = \sum_Q \text{sur}(G, Q) \cdot \text{inj}(Q, H)$, where the summation is over all (nonisomorphic) graphs Q such that G admits a surjective homomorphism to Q .
8. Show that a graph G satisfies the duality statement

$$G \not\rightarrow K_n \text{ if and only if } K_{n+1} \rightarrow G$$

if and only if the chromatic number of G is equal to the size of a maximum clique of G . Deduce that such a duality holds for any class \mathcal{C} of perfect graphs G . (A graph G is *perfect* if G and all its induced subgraphs have their chromatic number equal to the size of a largest clique.)

9. Show that the number of homomorphisms of a graph G to the graph H in Fig. 1.13 is precisely the number of independent sets of G .



FIG. 1.13. A special graph H for counting independent sets.

10. Let G, H be graphs. A *full homomorphism* $f : G \rightarrow H$ is a homomorphism which satisfies $f(u)f(v) \in E(H)$ if and only if $uv \in E(G)$. Prove the following statements.
 - (a) If G is an induced subgraph of H then the inclusion mapping $i : G \rightarrow H$ is a full injective homomorphism.
 - (b) If $f : G \rightarrow H$ is a full injective homomorphism then G is isomorphic to an induced subgraph of H .
 - (c) There exists a full homomorphism $f : G \rightarrow H$ if and only if the vertices of G can be partitioned into independent sets $S_x, x \in V(H)$, such that if $xy \notin E(H)$ then no edge of G joins the set S_x to the set S_u , and if $xy \in E(H)$ then all vertices of S_x are joined in G to all vertices of S_y .

11. Prove that if G and G_1, G_2, \dots, G_k are graphs such that

$$E(G) = E(G_1) \cup E(G_2) \cup \dots \cup E(G_k),$$

then

$$\chi(G) \leq \chi(G_1) \cdot \chi(G_2) \cdot \dots \cdot \chi(G_k).$$

12. Show that each oriented graph G , in which the binary relation $E(G)$ is transitive, retracts to some \vec{T}_k .
13. Suppose G_1, G_2 are graphs such that $V(G_1) = V(G_2)$ and every component of each G_i is a complete graph. (Each relation $E(G_i)$ is an equivalence.) Let G be the graph with $V(G) = V(G_1) = V(G_2)$ and $E(G) = E(G_1) \cup E(G_2)$. Prove that the core of G is a complete graph.
14. Give a polynomial time algorithm to decide whether or not an input digraph G admits a homomorphism to the directed path \vec{P}_k , k fixed. Deduce from the algorithm that $G \not\rightarrow \vec{P}_k$ if and only if $P \rightarrow G$ for some oriented path of net length $k + 1$.
15. A graph G is *point determining* if distinct vertices have distinct sets of neighbours. Prove that for $H = K_{1,3}$ and $H = C_4$ there exist only finitely many point determining bipartite graphs which do not have H as a homomorphic image.
16. [161] Let G be a graph, and let θ be the equivalence relation on $V(G)$ in which two vertices are equivalent if and only if they have exactly the same set of neighbours. Let $t(G)$ denote the number of equivalence classes of θ . (Note that a point-determining graph G has $t(G) = |V(G)|$.) Prove that the achromatic number of G is bounded both below and above by a function of $t(G)$. Deduce that the number of nonisomorphic point-determining graphs with a given achromatic number is finite.
17. [144] Prove that if a graph G admits a surjective homomorphism to both K_n and K_{n+1} , then it also admits a surjective homomorphism to some graph X with $n + 2$ vertices such that X admits a surjective homomorphism to both K_n and K_{n+1} .
18. [309] Prove that every vertex-transitive graph is a retract of a Cayley graph. (The *Cayley graph* of a group Γ and a set S of generators closed under taking inverses, is defined to have vertices $v \in \Gamma$, and edges $v(sv)$ for all $v \in \Gamma, s \in S$.)

PRODUCTS AND RETRACTS

In this chapter, we shall introduce some basic algebraic constructs, emphasizing the product and the retract. One can frequently aid the analysis of a complex digraph by writing it as a product of simpler digraphs. Even if this is not possible, we may gain some insight by embedding the digraph in a product of simpler digraphs. This turns out to be most useful when the given digraph is a retract of the product.

2.1 The product

The product most relevant to homomorphisms is defined as follows.

Let G and H be digraphs. The *product* of G and H is the digraph $G \times H$ with the vertex set $V(G \times H) = V(G) \times V(H)$, in which $(u, v)(u', v') \in E(G \times H)$ whenever $uu' \in E(G)$ and $vv' \in E(H)$.

The definition is stated in terms of digraphs, but applies to graphs (even when loops are allowed) via their corresponding symmetric digraphs. This amounts to applying the definition as written, with uu', vv' , and $(u, v)(u', v')$ being edges (and loops) rather than arcs. In Fig. 2.1, we illustrate the product of two (a) digraphs, (b) graphs with loops allowed, (c) graphs, and (d) reflexive graphs.

The relevance of this product to graph homomorphisms is based on the following simple facts (Fig. 2.2).

Proposition 2.1 *For digraphs G and H ,*

- $G \times H \rightarrow G$ and $G \times H \rightarrow H$
- if $X \rightarrow G$ and $X \rightarrow H$ then $X \rightarrow G \times H$.

Proof Indeed, consider the two *projections* π, ρ defined by $\pi(u, v) = u$ and $\rho(u, v) = v$ for all $(u, v) \in V(G \times H)$. It follows from the definition of $G \times H$ that the projections are homomorphisms $\pi : G \times H \rightarrow G$ and $\rho : G \times H \rightarrow H$. Moreover, if X is a digraph and if $\pi' : X \rightarrow G, \rho' : X \rightarrow H$ are two homomorphisms, then the mapping $f : X \rightarrow G \times H$ defined by $f(x) = (\pi'(x), \rho'(x))$ is also a homomorphism. \square

We can strengthen this observation by noting that the homomorphism f in the above proof satisfies $\pi \circ f = \pi'$ and $\rho \circ f = \rho'$, and is the unique mapping with this property. In fact, these properties characterize products (and projections), and we have the following theorem.

Theorem 2.2 *For any digraphs G, H , there exists a unique (up to isomorphism) digraph P and homomorphisms $p : P \rightarrow G, r : P \rightarrow H$ such that*




FIG. 2.1. Example products.

- for every digraph X with homomorphisms $p' : X \rightarrow G, r' : X \rightarrow H$ there is a unique homomorphism $f : X \rightarrow P$ with $p \circ f = p'$ and $r \circ f = r'$.




FIG. 2.2. Commuting diagram.

Proof We have observed above that $P = G \times H$ and $p = \pi, r = \rho$ have the required property. Suppose now that P_1 with homomorphisms p_1, r_1 , and P_2 with p_2, r_2 both have the required property. In particular, there exist homomorphisms $f : P_1 \rightarrow P_2, f' : P_2 \rightarrow P_1$ with $p_2 \circ f = p_1, r_2 \circ f = r_1$, and $p_1 \circ f' = p_2, r_1 \circ f' = r_2$. Consider now the composition $f' \circ f : P_1 \rightarrow P_1$: since $p_1 \circ (f' \circ f) = p_1, r_1 \circ (f' \circ f) = r_1$ we must have $f' \circ f$ equal to the identity mapping on P_1 , according to the uniqueness property. Similarly, $f \circ f'$ must be the identity mapping on P_2 , whence $f' = f^{-1}$ is an isomorphism of P_2 and P_1 . \square

We interpret the theorem to mean that we didn't have any choice in defining the 'correct' product for digraphs. If we want the product to satisfy these basic homomorphism properties, we *must* define it this way. Because of the theorem, one can also *define* the product of digraphs to be the unique P from the theorem. Such a 'vertex-free' definition has its disadvantages—we do not learn directly what are the vertices and edges of the product. But it also has some advantages—for instance it applies to general relational systems. In Exercise 1 we ask the reader to give a *direct* definition of the product of general relational systems, which satisfies the properties in Theorem 2.2. In Exercise 2 we explore a similar 'vertex-free' definition of the disjoint union $G + H$ of graphs.

We also conclude that the *numbers* of homomorphisms to the product can be computed from the numbers of homomorphisms to the factors.

Corollary 2.3 *For any digraph X ,*

$$\text{hom}(X, G \times H) = \text{hom}(X, G) \cdot \text{hom}(X, H).$$

Proof Every homomorphism $f : X \rightarrow G \times H$ is uniquely determined by the pair of homomorphisms $(\pi \circ f, \rho \circ f)$. \square

It is easy to see (directly from the definitions or by using Theorem 2.2) that the product is commutative and associative, and that it distributes over the disjoint union $+$.

Proposition 2.4 *For any digraphs G, H, K , we have*

- $G \times H \simeq H \times G$,
- $G \times (H \times K) \simeq (G \times H) \times K$, and
- $G \times (H + K) \simeq G \times H + G \times K$.

\square

Because of the associativity we may speak of the product $\prod_{i \in I} G_i$ of a finite family of digraphs G_i . We may also define $\prod_{i \in I} G_i$ directly. Given digraphs $G_i, i \in I$ we let $\prod_{i \in I} G_i$ be the digraph on the vertex set $\prod_{i \in I} V(G_i)$, with $(u_i)_{i \in I} (u'_i)_{i \in I} \in E(\prod_{i \in I} G_i)$ whenever $u_i u'_i \in E(G_i)$ for each $i \in I$. We also define, for each $j \in I$, the j -th projection $\pi_j : \prod_{i \in I} G_i \rightarrow G_j$ by $\pi_j(u_i)_{i \in I} = u_j$. Note that this definition allows infinite index sets I , cf. Exercise 6.

Let L denote the one-vertex digraph with a loop. Clearly each digraph G is isomorphic to $G \times L$. A digraph which is isomorphic to a product of two smaller digraphs will be called *composite*, and a digraph which is not composite is called *prime digraph*. It is clear that a digraph is (isomorphic to) the product of prime digraphs—either it is prime itself, or is isomorphic to a product of smaller digraphs, each of which in turn is either prime, or a product of smaller digraphs, and so on. Since our digraphs are finite, we must eventually end with a product of prime factors—called a *prime factorization* of G , in analogy with prime factorization of integers. Unlike the case of integers, prime factorization need not be unique, even for graphs, as we illustrate in Fig. 2.3. (On the other hand, see Theorem 2.14.)

For instance, for graphs with loops allowed, we can just notice that $2L \times K_2 \simeq 2K_2 \simeq K_2 \times K_2$, and both $2L$ and K_2 are prime, as they have a prime number of vertices. (In fact, $2L \times G \simeq 2G \simeq K_2 \times G$, for any bipartite graph G .) In Fig. 2.3, we illustrate a similar situation for graphs (without loops). The heavier edges are intended to be a visual aid in seeing an isomorphism between $C_6 \times K_2$ and $2C_3 \times K_2$.




FIG. 2.3. $C_6 \times K_2 \simeq 2C_6 \simeq 2C_3 \times K_2$.

The construction of the product is deceptively simple, yet, as we hope to demonstrate, very useful. We explore this theme in the next few sections, and return to it repeatedly.

2.2 Dimension

The product construction allows for very complex structures to be built from extremely simple building blocks. Even starting with complete graphs yields a surprisingly rich family of products. In this section we focus on graphs.

Theorem 2.5 *Every graph is isomorphic to an induced subgraph of a product of complete graphs.*

Proof Let G be any graph, and K a complete graph on the same vertex set $V(G)$. Let $I = \text{HOM}(G, K)$ and take the product $\prod_{h \in I} F_h$ where each $F_h = K$. For each $h \in I$ (a homomorphism $G \rightarrow K$), we can define a homomorphism $f_h : G \rightarrow F_h = K$ by setting $f_h(x) = h(x)$. We now claim that the mapping f defined by $f(u) = (f_h(u))_{h \in I}$ is an isomorphism of G to an induced subgraph of $\prod_{h \in I} F_h$. It is easy to check that f is a homomorphism. It is injective, since one of the $h \in I$ is the identity mapping on $V(G)$ and hence $u \neq v$ implies

$f_h(u) \neq f_h(v)$. Finally, it is an isomorphism onto an induced subgraph of the product, since if u, v are not adjacent in G then some $h \in I$ has $h(u) = h(v)$ and so $f_h(u), f_h(v)$, and hence also $f(u), f(v)$ are not adjacent in the product. \square

An isomorphism of a graph G to an induced subgraph of a product of complete graphs will be called an *encoding of G* . An encoding of G associates to each vertex v of G a vector, which is a vertex in the product of complete graphs; we call this vector *the code of v* , in the given encoding of G . In Fig. 2.4 we write the codes as strings.




FIG. 2.4. An encoding of P_3 in $K_2 \times K_3$.

The last proof is essentially identical to the well-known proof showing that every partial order is a suborder of a direct product of chains. (That result may be formulated in our context as follows: *every reflexive, transitive, and antisymmetric digraph is isomorphic to an induced subgraph of a product of reflexive transitive tournaments.*) As for ordered sets, we wish to measure the complexity of a graph by how many factors are needed in such a product.

The *dimension* of G , denoted $\dim G$, is the smallest integer d such that G is isomorphic to an induced subgraph of a product of d complete graphs. Theorem 2.5 ensures that $\dim G$ is well defined, but its proof only yields the upper bound $\dim G \leq |I| \leq n^n$. In order to obtain better bounds we consider the following reformulation.

Proposition 2.6 *Let G be a graph. Then $\dim G \leq d$ if and only if there exist equivalence relations R_1, R_2, \dots, R_d on $V(G)$ such that*

1. *two adjacent vertices of G are nonequivalent in all R_i ,*
2. *two nonadjacent vertices of G are equivalent in some R_i , and*
3. *two distinct vertices of G are nonequivalent in some R_i .*

Proof Condition 3 says that the equivalences ‘separate’ vertices, while conditions 1 and 2 say that the union of the equivalences is the complement graph \overline{G} . If G is encoded in the product of d complete graphs, we set R_i to be the equivalence relation in which two vertices are equivalent just if they have the same i -th coordinate. Conditions 1, 2, 3 are now easy to check. Conversely, if R_1, R_2, \dots, R_d are equivalence relations on $V(G)$ satisfying 1, 2, 3, we can encode G in a product of d complete graphs, by setting the i -th coordinate of a vertex $v \in V(G)$ to be the name of the equivalence class of R_i containing v . \square

Corollary 2.7 *If G is a graph with n vertices, then $\dim G \leq n$.*

Proof Since the maximum degree in \overline{G} is at most $n - 1$, Vizing’s theorem implies the edges of \overline{G} can be coloured with n colours, so that incident edges

obtain different colours. Let E_i denote the set of edges of colour $i = 1, 2, \dots, n$. Even though some E_i may be empty, we can define equivalence relations $R_i, i = 1, 2, \dots, n$, as follows. Each equivalence relation R_i corresponds to the partition of $V(G)$ in which the classes are the two-vertex subsets corresponding to the edges of E_i (if any), and the one-vertex subsets corresponding to the vertices not belonging to any edge of E_i . (Recall that each vertex can belong to at most one edge of any E_i .) This collection of equivalence relations clearly satisfies 1–3. \square

In fact, it can be shown that with a small number of exceptions, $\dim G \leq n-1$. (In this sense the situation resembles the theorem of Brooks.)

The *chromatic index* of a graph is the minimum number of colours needed to colour the *edges* of G so that incident edges have different colours. The preceding proof implies that the dimension of G is at most the chromatic index of the complement \overline{G} , as long as this number is at least two. (When \overline{G} has chromatic index one, we need an extra equivalence to ensure condition 3.) If, in addition, the complement \overline{G} is triangle-free, then we need to consider *only* equivalences with classes of size one or two (as above).

Corollary 2.8 *Let G be a graph. If the complement \overline{G} is triangle-free and has chromatic index $k \geq 2$, then $\dim G = k$.* \square

This corollary allows us to find graphs of arbitrarily high dimension, by taking triangle-free graphs with high minimum degree, and hence high chromatic index.

Another very useful technique from linear algebra relies on matchings. A *matching* M of a graph G is a set of disjoint edges $x_i y_i, i = 1, 2, \dots, k$, of G . A matching $M = \{x_1 y_1, x_2 y_2, \dots, x_k y_k\}$ of G is *strict* if $x_i y_j \notin E(G)$ whenever $i < j$.

Theorem 2.9 *If G has a strict matching with k edges, then $\dim G \geq \log_2 k$.*

Proof Suppose G is encoded in a product of m complete graphs. We may assume that the code of each vertex v of G is an m -vector of positive integers, $\vec{c}(v) = (c(v)_1, c(v)_2, \dots, c(v)_m)$. Let $L = \{1, 2, \dots, m\}$. We further associate to v two 2^m -vectors $\vec{a}(v) = (a(v)_S)_{S \subseteq L}$ and $\vec{b}(v) = (b(v)_S)_{S \subseteq L}$, where $a(v)_S = \prod_{i \in S} c(v)_i$, and $b(v)_S = \prod_{i \notin S} c(v)_i$. This definition implies that for two vertices x, y of G , the dot product of vectors $\vec{a}(x)$ and $\vec{b}(y)$ is $\prod_{i=1}^m (c(x)_i - c(y)_i)$, which is zero if and only if x and y are nonadjacent in G .

Let now $M = \{x_1 y_1, x_2 y_2, \dots, x_k y_k\}$ be a strict matching in G . It then easily follows from what we have just observed that the matrix of dot products $(\vec{a}(x_i) \cdot \vec{b}(y_j))_{i,j}$ is nonsingular. Therefore, the vectors $\vec{a}(x_1), \vec{a}(x_2), \dots, \vec{a}(x_k)$ are linearly independent. Indeed, $\sum_{i=1}^m \lambda_i \vec{a}(x_i) = 0$ implies $\sum_{i=1}^m \lambda_i (\vec{a}(x_i) \cdot \vec{b}(y_j)) = 0$, for all $j = 1, 2, \dots, k$. This is a system of homogeneous linear equations with the above matrix we have just proved nonsingular. Hence all λ_i are zero. Finally, we observe that since there exist k linearly independent 2^m -vectors, we must have $k \leq 2^m$. \square

It can be deduced that the dimension of a product of m nontrivial complete graphs is exactly m , i.e., cannot be also expressed as a subgraph of a product of fewer than m complete graphs. Consider the product of m copies of K_2 . It is readily seen to be the union of 2^m disjoint copies of K_2 , and hence admits a strict matching with 2^m edges. Thus we obtain another family of graphs with arbitrarily high dimension.

Corollary 2.10 $\dim(K_2)^m = m$. □

Of course, a similar result holds about $\dim(K_n)^m$.

2.3 The Lovász vector and the Reconstruction Conjecture

It turns out that the numbers of homomorphisms uniquely describe a digraph. Two digraphs are isomorphic if and only if they have exactly the same *number* of homomorphisms from every other digraph.

Theorem 2.11 *Let G, H be digraphs. Then G and H are isomorphic if and only if for every digraph X we have*

$$\text{hom}(X, G) = \text{hom}(X, H).$$

Let X_1, X_2, \dots be a fixed enumeration of all nonisomorphic digraphs. (In other words, these digraphs are pairwise nonisomorphic and each digraph is isomorphic to one of them.) Then the *Lovász vector* of a digraph G is the countable sequence (n_1, n_2, \dots) , where each $n_i = \text{hom}(X_i, G)$.

The theorem states that two digraphs are isomorphic if and only if their Lovász vectors are equal. Recall the deck introduced for the Edge Reconstruction Conjecture (Conjecture 1.29). While we don't know that the deck is sufficient to reconstruct G , the *Lovász vector is sufficient*.

Proof Clearly, the cardinality condition is necessary. To prove its sufficiency, note that it is enough to prove that it implies, for every digraph X ,

$$\text{inj}(X, G) = \text{inj}(X, H),$$

since taking $X = G$ and $X = H$ implies the existence of both an injective homomorphism of G to H and an injective homomorphism of H to G —and thus of an isomorphism between G and H . We verify the above equality by induction on the number of vertices of X . When X has only one vertex, then every homomorphism of X is injective, so the equality holds by assumption. Let us assume we have already proved the equality for all digraphs which have fewer vertices than X . We have observed in Lemma 1.27 that

$$\text{hom}(X, Y) = \sum_{\theta} \text{inj}(X/\theta, Y),$$

where the sum is taken over all partitions θ of $V(X)$. Let ι denote the trivial partition in which each class is a singleton. Then

$$\text{hom}(X, Y) = \text{inj}(X, Y) + \sum_{\theta \neq \iota} \text{inj}(X/\theta, Y).$$

Since each X/θ on the right-hand side has fewer vertices than X , we have $\text{inj}(X/\theta, G) = \text{inj}(X/\theta, H)$ by the induction hypothesis, and hence we conclude that $\text{inj}(X, G) = \text{inj}(X, H)$. \square

We now return to the question of factorizations, which motivated Theorem 2.11. We have seen that the same digraph may have two different prime factorizations. In fact, we had seen examples where $G \times K \simeq H \times K$ without $G \simeq H$ —an equation from which it is impossible to cancel K . Theorem 2.11 allows us to prove the following cancellation law for digraphs.

Corollary 2.12 *If G, H, K are digraphs, where K contains a loop, then $G \times K \simeq H \times K$ if and only if $G \simeq H$.*

Proof Clearly the condition $G \simeq H$ is sufficient. For its necessity, take any digraph X and note that using Corollary 1.27, we obtain

$$\begin{aligned} \text{hom}(X, G) \cdot \text{hom}(X, K) &= \text{hom}(X, G \times K) \\ &= \text{hom}(X, H \times K) = \text{hom}(X, H) \cdot \text{hom}(X, K). \end{aligned}$$

Since K has at least one loop, the factor $\text{hom}(X, K)$ is not zero and hence $\text{hom}(X, G) = \text{hom}(X, H)$. \square

For graphs, it turns out that any odd cycle in K (not just a loop, i.e., a one-cycle) ensures that K can be cancelled.

Theorem 2.13 [224, 225] *If K is a nonbipartite graph, then $G \times K \simeq H \times K$ if and only if $G \simeq H$.*

On the other hand, for connected reflexive and nonbipartite graphs much more can be proved.

Theorem 2.14 [240] *The class of connected nonbipartite reflexive graphs has unique prime factorization.*

Another interesting result that follows (in a similar fashion) from Theorem 2.11 is the following corollary dealing with numeric powers $G^k = G \times G \times \cdots \times G$ of digraphs.

Corollary 2.15 *$G^k \simeq H^k$ if and only if $G \simeq H$.* \square

We now return to the original Edge Reconstruction Conjecture. Recall that we have proved the conjecture, in Theorem 1.30, if the graphs G and H have n vertices and more than $\frac{1}{2}\binom{n}{2}$ edges. The approach based on counting homomorphisms can be extended to prove a stronger result. For simplicity we shall again stay with graphs, even though the technique applies equally to digraphs (and more general structures).

Theorem 2.16 *Let G, H be graphs with n vertices and $m > n(\log_2 n - 1)$ edges. If there exists a bijective mapping $\beta : E(G) \rightarrow E(H)$ such that $G - e \simeq H - \beta(e)$ for every $e \in E(G)$, then $G \simeq H$.*

Proof As in the proof of Theorem 1.30, we shall apply inclusion–exclusion, but this time in a somewhat more sophisticated way. Let $e = vw$ be an edge of G ; recall that we say that an injective mapping $f : V \rightarrow V$ has the property \mathcal{P}_e , if $f(v)f(w)$ is an edge of H . We say that f is an injective homomorphism of G to \overline{H} with *defect* D ($D \subseteq E(G)$), if it has precisely the properties $\mathcal{P}_e, e \in D$. Let $\text{inj}_D(G, \overline{H})$ denote the number of injective homomorphisms of G to \overline{H} with defect D .

We have previously shown that $\text{inj}(G, \overline{H}) = \text{inj}(\emptyset, H) - \sum_{|A|=1} \text{inj}(A, H) + \sum_{|A|=2} \text{inj}(A, H) - \sum_{|A|=3} \text{inj}(A, H) + \cdots + (-1)^m \text{inj}(G, H)$. A similar equation can be derived, by inclusion–exclusion, for $\text{inj}_D(G, \overline{H})$.

$$\begin{aligned} \text{inj}_D(G, \overline{H}) &= \text{inj}(D, H) - \sum_{|A|=1} \text{inj}(D \cup A, H) + \sum_{|A|=2} \text{inj}(D \cup A, H) \\ &\quad - \sum_{|A|=3} \text{inj}(D \cup A, H) + \cdots + (-1)^{m-|D|} \text{inj}(G, H). \end{aligned}$$

The sets A are assumed to be subsets of $E(G) - D$. For $0 \leq d \leq m$, we let $\text{inj}_d(G, \overline{H})$ denote the sum

$$\sum_{|D|=d} \text{inj}_D(G, H);$$

then we can rewrite the above equations as follows, with all sets $B \subseteq E(G)$:

$$\begin{aligned} \text{inj}_d(G, \overline{H}) &= \sum_{|B|=d} \text{inj}(B, H) - \sum_{|B|=d+1} \binom{d+1}{d} \text{inj}(B, H) \\ &\quad + \sum_{|B|=d+2} \binom{d+2}{d} \text{inj}(B, H) - \sum_{|B|=d+3} \binom{d+3}{d} \text{inj}(B, H) \\ &\quad + \cdots + (-1)^{m-d} \binom{m}{d} \text{inj}(G, H). \end{aligned}$$

(Note that taking $d = 0$ yields the formula derived in the proof of Theorem 1.30.) If we again let $H = G$, we obtain the equation (with all sets $B \subseteq E(H)$)

$$\begin{aligned} \text{inj}_d(H, \overline{H}) &= \sum_{|B|=d} \text{inj}(B, H) - \sum_{|B|=d+1} \binom{d+1}{d} \text{inj}(B, H) \\ &\quad + \sum_{|B|=d+2} \binom{d+2}{d} \text{inj}(B, H) - \sum_{|B|=d+3} \binom{d+3}{d} \text{inj}(B, H) \\ &\quad + \cdots + (-1)^{m-d} \binom{m}{d} \text{inj}(H, H). \end{aligned}$$

Except for the last term, the subgraphs corresponding to any particular B 's in G and in H are in one-to-one correspondence, as before. Thus most of the terms on the right-hand side of these expressions are pairwise the same, and we obtain

$$\text{inj}_d(G, \overline{H}) - \text{inj}_d(H, \overline{H}) = (-1)^{m-d} \binom{m}{d} (\text{inj}(G, H) - \text{inj}(H, H)).$$

Assume, for a proof by contradiction, that G and H are not isomorphic, i.e., that $\text{inj}(G, H) = 0$. Summing up the above equations over all $0 \leq d \leq m$, we obtain

$$2^m \leq 2^m \text{inj}(H, H) = \sum_{d=0}^m |\text{inj}_d(G, \overline{H}) - \text{inj}_d(H, \overline{H})| \leq 2n!$$

since $\sum_d \text{inj}_d(G, \overline{H}) = n!$. From $n! < (n/2)^n$ we conclude that $m < 1 + n \cdot (\log_2 n - 1)$, contrary to our assumption. \square

2.4 Exponential digraphs

Let G, H be digraphs. The *exponential digraph* H^G has as its vertices all *mappings* $V(G) \rightarrow V(H)$ and as its edges all ff' such that $f(v)f'(v') \in E(H)$ for every $vv' \in E(G)$. The exponential digraph H^G is also called (the G -th) *power* of H .

If G, H are graphs, the exponential graph H^G is, in general, a graph with loops allowed. In fact, f is a loop of G^H if and only if $f : G \rightarrow H$ is a homomorphism.

In Fig. 2.5, we illustrate the exponential digraph $\vec{C}_3^{\vec{C}_2}$ for directed cycles \vec{C}_2, \vec{C}_3 . Note that there are no loops in this power, since $\vec{C}_2 \not\rightarrow \vec{C}_3$.




FIG. 2.5. The exponential digraph $\vec{C}_3^{\vec{C}_2}$ for directed cycles \vec{C}_2, \vec{C}_3

The exponentiation of digraphs has properties that are very similar to the usual laws of exponentiation of real numbers. We explore some of these properties in the next proposition, and a few more are taken up in the exercises (Exercise 8).

Proposition 2.17 *For any digraphs F, G , and H*

1. $H^{G+F} \simeq H^G \times H^F$
2. $H^{G \times F} \simeq (H^G)^F$.

Proof We prove 2 and leave 1 as an exercise. To each mapping m of $V(F)$ to $V(H^G)$ we assign the mapping $\Phi(m)$ of $V(G \times F)$ to $V(H)$ which maps the vertex (u, v) of $G \times F$ to the vertex $(m(v))(u)$ of H . (Recall that $m(v)$ is a mapping from $V(G)$ to $V(H)$.) It is clear that Φ is a bijective mapping between $V((G^H)^F)$ and $V(H^{G \times F})$. It is also easy to check that Φ is a homomorphism from $(G^H)^F$ to $H^{G \times F}$. Indeed, if $mm' \in E((G^H)^F)$, then $\Phi(m)\Phi(m') \in E(H^{G \times F})$, since $(\Phi(m))(u, v) = (m(v))(u)$ is adjacent to $(\Phi(m'))(u', v') = (m'(v'))(u')$ whenever (u, v) is adjacent to (u', v') in $G \times F$. \square

Corollary 2.18 *$G \times F \rightarrow H$ if and only if $F \rightarrow H^G$.*

Proof In fact, it is clear that $\text{hom}(G \times F, H) = \text{hom}(F, H^G)$, as both sides count the number of loops in $(H^G)^F$. \square

We make two other observations.

Proposition 2.19 *For any digraphs G and H*

1. H is isomorphic to an induced subgraph of H^G
2. $G \times H^G \rightarrow H$

Proof The constant mapping $f_h : V(G) \rightarrow V(H)$, which takes all vertices of G to the vertex $h \in V(H)$, is a vertex of H^G . It is easy to check that the association $h \mapsto f_h$ is an isomorphism of H onto an induced subgraph of H^G , proving 1. To prove 2, define η to be the *evaluation* mapping on $V(G) \times V(H)^{V(G)}$, i.e., let $\eta(u, m) = m(u)$, where $m \in V(H)^{V(G)}$, $u \in V(G)$. It is again a simple exercise to show that η is a homomorphism $G \times H^G \rightarrow H$. \square

2.5 Shift graphs

Theorem 1.9 states that there exist graphs of arbitrarily high chromatic number and girth. The constructive proofs of this result are quite difficult. In contrast, there are simple constructions of graphs with high *odd* girth and chromatic number. Recall that the odd girth of a nonbipartite graph G is the length of a shortest odd cycle in G , and that the odd girth is more relevant than girth from the perspective of graph homomorphisms (cf. Exercise 2 in Chapter 1).

General shift graphs have as vertices fixed-length strings over a fixed alphabet, and adjacency is defined as follows. If $\ell \geq 2$, the string $a_1 a_2 \cdots a_\ell$ is adjacent to $b_1 b_2 \cdots b_\ell$ just if $b_i = a_{i+1}$ for all $i = 1, 2, \dots, \ell - 1$. The generic shift graph (known as the *de Bruijn graph*) $dB(n, \ell)$, has as vertices all strings of length ℓ over the alphabet $0, 1, 2, \dots, n-1$. For our purposes we shall take an induced subgraph of the de Bruijn graph. The *shift graph* $R(n, \ell)$ is the subgraph of $dB(n, \ell)$ induced by all monotone strings, i.e., strings $a_1 a_2 \cdots a_\ell$ with $a_1 < a_2 < \cdots < a_\ell$.

We observe that shift graphs have an inherent orientation—reading the definition (‘is adjacent to’) in the context of digraphs results in a *shift digraph* $\vec{R}(n, \ell)$.

This digraph has arcs from a vertex $a_1a_2\cdots a_\ell$ to all vertices $a_2a_3\cdots a_\ell a$, for all characters $a, a > a_\ell$. In this way the arc from $a_1a_2\cdots a_\ell$ to $a_2a_3\cdots a_\ell a_{\ell+1}$ can be associated with a string of length $\ell + 1$, namely $a_1a_2\cdots a_\ell, a_{\ell+1}$. Thus there is a natural connection to line digraphs. The *line digraph* $L(G)$ of a digraph G has as its vertices the arcs of G , and has an arc from $ab \in E(G)$ to $cd \in E(G)$, just if $b = c$. It follows from the above observation that shift digraphs satisfy $\vec{R}(n, \ell + 1) = L(\vec{R}(n, \ell))$. From the definition of $\vec{R}(n, 2)$ we see that $\vec{R}(n, 2) = L(\vec{T}_n)$, therefore, it is convenient to define $\vec{R}(n, 1)$ as \vec{T}_n , the transitive tournament, this time taken on the vertices $0, 1, \dots, n - 1$.

Lemma 2.20 *Let n be a positive integer. Then*

- $\vec{R}(n, 1) = \vec{T}_n$, and
- $\vec{R}(n, \ell + 1) = L(\vec{R}(n, \ell))$, for all $\ell \geq 1$. □

Note that all shift digraphs are irreflexive.

In the following, we take the chromatic number of an irreflexive digraph to be the chromatic number of the underlying graph. The chromatic number of the line digraph of G is bounded on both sides by a (logarithmic) function of the chromatic number of G .

Lemma 2.21 *For any digraph G ,*

$$\log_2 \chi(G) \leq \chi(L(G)) \leq \min \left\{ k : \chi(G) \leq \binom{k}{\lfloor k/2 \rfloor} \right\}.$$

Proof In this section we only need the lower bound, which we prove first. A *proper colouring* of $L(G)$ is a colouring of the arcs of G in which no two arcs ab and bc have the same colour. Thus consider such a proper colouring of $L(G)$ with t colours, and form the spanning subgraphs G_1, G_2, \dots, G_t of G , where $E(G_i)$ consists of all the arcs that obtained colour i . Since each vertex of G_i has either outdegree zero or indegree zero, the underlying graph of G_i is bipartite, for every i . It is easy to check (cf. Exercise 11 in Chapter 1) that if each c_i is a two-colouring of G_i , then c defined by $c(v) = (c_1(v), c_2(v), \dots, c_t(v))$ is a proper colouring of G , and so $\chi(G) \leq 2^t$.

To prove the upper bound, we show that $\chi(G) \leq \binom{k}{\lfloor k/2 \rfloor}$ implies that $\chi(L(G)) \leq k$. Indeed, any colouring c of G with $\binom{k}{\lfloor k/2 \rfloor}$ colours can be viewed as an assignment of $\lfloor k/2 \rfloor$ -subsets of $\{1, 2, \dots, k\}$ to the vertices of G , in which adjacent vertices obtain disjoint subsets. Define a mapping f of $V(L(G))$ to $\{1, 2, \dots, k\}$ in which the arc xy of G (being a vertex of $L(G)$) receives the value $f(xy) = u$ where u is any element of $f(x) - f(y)$. It is easy to see that f is indeed a k -colouring of $L(G)$, i.e., that arcs xy and yz cannot obtain the same colour. Note that the upper bound is roughly of the order $\log_4 \chi(G)$; it will be used in Section 2.6. □

Lemma 2.22 *The graph $R(n, \ell)$ has odd girth at least $2\ell + 1$.*

Proof We proceed by induction on ℓ . Indeed, the statement is trivial when $\ell = 1$. Hence consider an odd cycle C in $R(n, \ell)$. It consists of vertices r_1, r_2, \dots, r_t (t odd) of $R(n, \ell)$. Thus the strings r_1, r_2, \dots, r_t are arcs of $R(n, \ell - 1)$, and for each consecutive pair r_i, r_{i+1} of arcs, the head of one equals the tail of the other. This defines a closed walk of the same length t in $R(n, \ell - 1)$. Of course, a closed odd walk must contain a closed odd cycle, so we only have to observe that the closed walk r_1, r_2, \dots, r_t is itself not a cycle (hence the odd cycle is shorter and we can apply the induction hypothesis). This follows from the fact that the arc that is the lexicographically smallest string has both its neighbours in C have tails equal to its head, and the arc that is the lexicographically largest string has both its neighbours in C have heads equal to its tail (Fig. 2.6). \square

The situation is illustrated in Fig. 2.6; for convenience we show parts of the graph $R(n, 2)$ and the digraph $\vec{R}(n, 1)$.




FIG. 2.6. A five-cycle in $R(n, 2)$ corresponds to a closed walk of length five in $\vec{R}(n, 1) = \vec{T}_n$.

Since the chromatic number of $R(n, \ell)$ is controlled by both n and ℓ , while the odd girth only by ℓ , we can now construct graphs of arbitrarily high odd girth and chromatic number.

Theorem 2.23 *Let g, k be positive integers, $g \geq 3$ odd. Then there exists a graph $S(g, k)$ with odd girth at least g and chromatic number at least k .*

Proof It is enough to set $S(g, k) = R(n, \ell)$ with $2\ell + 1 = g$ and $n = 2^{2^{\dots^k}}$, where the tower of powers has length ℓ . (In other words, $n = n_\ell$ is defined recursively as $n_1 = 2^k$, and $n_{\ell+1} = 2^{n_\ell}$.) \square

It may appear that it will be much harder to construct graphs of arbitrarily high odd girth g and chromatic number k which are *uniquely k -colourable*, but the product construction makes this easy.

Proposition 2.24 *If G is a connected graph which is not k -colourable, then $G \times K_k$ is uniquely k -colourable.*

Proof The second projection $\rho : G \times K_k \rightarrow K_k$ (taking each (u, x) to x) is a k -colouring of $G \times K_k$. We shall show that every other k -colouring is obtained from

ρ by renaming colours, i.e., is the projection ρ followed by an automorphism of K_k . Thus let $c : G \times K_k \rightarrow K_k$ be a homomorphism, i.e., a k -colouring of $G \times K_k$. Suppose first that for some vertex $u \in V(G)$ the values $c(u, x) = x, x \in V(K_k)$, are all distinct. In this case, for any v adjacent to u , and any $x \in V(K_k)$, the vertex (v, x) is adjacent to all $(u, y), y \neq x$, and so $c(v, x) = c(u, x)$. In particular, all $c(v, x), x \in V(K_k)$, are again distinct. By the connectivity of G we conclude that all values $c(u, x)$ with the same x are equal, i.e., that the colouring c differs from ρ only in renaming the colours.

We prove the existence of such a vertex u by contradiction. Suppose that for each $u \in V(G)$ there are two vertices x_u, y_u of K_k with $c(u, x_u) = c(u, y_u)$, and denote by $c(u)$ this common value $c(u, x_u) = c(u, y_u)$. We claim that c is a k -colouring of G . Consider an edge $uv \in E(G)$. Note that for every $x \in V(K_k)$ at least one of $(u, x_u), (u, y_u)$ is adjacent to (v, x) , and so $c(v, x) \neq c(u)$; in particular, $c(v) \neq c(u)$. \square

Corollary 2.25 *Let g, k be positive integers, g odd. Then there exists a uniquely k -colourable graph of odd girth at least g .*

(Note that a uniquely k -colourable graph must have chromatic number k .)

Proof Let G be a graph of chromatic number at least $k + 1$ from the theorem (which clearly may be chosen to be connected). Then the proposition implies that $G \times K_k$ is uniquely k -colourable. It only remains to observe that $G \times K_k \rightarrow G$ and hence must not have odd girth smaller than G . \square

The existence of uniquely k -colourable graphs with high *girth* is discussed in Corollary 3.17.

The proof of Proposition 2.24 has a nice interpretation in the power $K_k^{K_k}$. According to the isomorphism of $K_k^{K_k \times G}$ and $(K_k^{K_k})^G$ from Proposition 2.17, 2, each homomorphism $c : G \times K_k \rightarrow K_k$ corresponds to a set of vertices of $K_k^{K_k}$ —one for each vertex $u \in V(G)$ (being the mapping that takes x to $c(u, x)$). Since this mapping is a homomorphism $K_k \rightarrow K_k$ if and only if all values $c(u, x) = x$, for $x \in V(K_k)$, are distinct, the two parts of the proof of Proposition 2.24 correspond respectively to the following two assertions.

Corollary 2.26 *Consider the power $K_k^{K_k}$.*

- *The subgraph induced by the vertices with loops has no other edges.*
- *The subgraph induced by the vertices without loops is k -colourable.* \square

2.6 The Product Conjecture and graph multiplicativity

One of the most challenging open problems in this area is the following *Product Conjecture*.

Conjecture 2.27 *If G, H are graphs, then*

$$\chi(G \times H) = \min(\chi(G), \chi(H)).$$

Since $G \times H \rightarrow G$ and $G \times H \rightarrow H$, we have $\chi(G \times H) \leq \min(\chi(G), \chi(H))$, whence the Product Conjecture amounts to really just proving the opposite inequality, i.e., that the product of two graphs which are *not* k -colourable is also not k -colourable. For convenience we also give the conjecture in the positive form.

Conjecture 2.28 *If $G \times H$ is k -colourable, then G or H is also k -colourable.*

For $k = 1$, this just says that the product of two graphs that are not edgeless is a graph that is also not edgeless—clearly true. (A graph G is *edgeless* if $E(G) = \emptyset$.) For $k = 2$, we are saying that the product of two nonbipartite graphs is also nonbipartite. If the odd girth of G is a and the odd girth of H is b with $b \geq a$, then $C_b \rightarrow G$ and $C_b \rightarrow H$, and hence $C_b \rightarrow G \times H$, so that $G \times H$ cannot be bipartite (cf. also Exercise 4). In both cases, we have used duality—for $k = 1$, the trivial simple duality that says $G \not\rightarrow K_1$ if and only if $K_2 \rightarrow G$, and for $k = 2$, the duality from Corollary 1.19 which says that $G \not\rightarrow K_2$ if and only if $C_n \rightarrow G$ for some odd n (the theorem of König).

There is again a natural connection to the exponential digraph. To see the connection more clearly, we define a digraph K to be *multiplicative* if $G \times H \rightarrow K$ implies that $G \rightarrow K$ or $H \rightarrow K$. (Again, the condition is often used in the contrapositive form, $G \not\rightarrow K, H \not\rightarrow K$ imply that $G \times H \not\rightarrow K$.) Thus the Product Conjecture asserts that each complete graph K_k is multiplicative.

Lemma 2.29 *A digraph K is multiplicative if and only if $G \not\rightarrow K$ implies $K^G \rightarrow K$.*

Proof Suppose first that K is multiplicative, and $G \not\rightarrow K$. If also $K^G \not\rightarrow K$, then $G \times K^G \not\rightarrow K$, contrary to Proposition 2.19. On the other hand, assume that $G \not\rightarrow K$ implies $K^G \rightarrow K$, and take any G, H with $G \not\rightarrow K, H \not\rightarrow K$. If we had $G \times H \rightarrow K$, then by Corollary 2.18 $H \rightarrow K^G$, contradicting $H \not\rightarrow K$ (since $K^G \rightarrow K$). \square

We shall say that a digraph G is *K -persistent* if $G \not\rightarrow K$ implies $K^G \rightarrow K$. (Note that any digraph G with $G \rightarrow K$ is trivially K -persistent.) Therefore, K is multiplicative if and only if each digraph is K -persistent. It follows from the previous proof that G is K -persistent if and only if $G \rightarrow K$, or $G \times H \not\rightarrow K$ whenever $H \not\rightarrow K$. Thus the following two propositions describe situations in which the Product Conjecture holds for a graph G and *all* H with the same chromatic number.

Proposition 2.30 *Any graph G in which each vertex belongs to a K_k is K_k -persistent.*

Proof Let each vertex of G belongs to a K_k , and $G \not\rightarrow K_k$. Suppose $H \not\rightarrow K_k$, but there is a homomorphism $f : G \times H \rightarrow K_k$. Without loss of generality we may assume that H is connected. Since each vertex of G belongs to a K_k , and $K_k \times H$ is uniquely colourable by Proposition 2.24, we must have $f(u, v) = f(u, v')$ for

all $u \in V(G)$, $v, v' \in V(H)$. It follows that $c(u) = f(u, w)$, where w is any fixed vertex of H , defines a homomorphism of G to K_k , contrary to assumption. \square

Proposition 2.31 *Any graph G without induced $2K_2$ is K_k -persistent for all k .*

Proof Let G contain no induced $2K_2$, and suppose $G \not\rightarrow K_k$. We define a homomorphism $c : K_k^G \rightarrow K_k$ as follows. Each mapping $m \in V(K_k^G)$ must have $m(a) = m(b)$ on some adjacent pair of vertices a, b of G —since $G \not\rightarrow K_k$. Let $c(m)$ denote such a value $m(a) = m(b)$. We claim that c is indeed a homomorphism $K_k^G \rightarrow K_k$. Suppose that m and m' are adjacent in K_k^G , and $c(m) = c(m')$. Let $ab, a'b' \in E(G)$ have $m(a) = m(b) = c(m)$, $m'(a') = m'(b') = c(m') (= c(m))$. Since G contains no induced $2K_2$, one of the vertices a, b must be adjacent or equal to one of the vertices a', b' . In each case, we obtain adjacent vertices x, x' of G with $m(x) = m'(x')$, contradicting the fact that $mm' \in E(K_k^G)$. \square

We have noted that the Product Conjecture holds for $k = 1, 2$. It has also been proved for $k = 3$, using Lemma 2.29.

Theorem 2.32 [89] *If $G \times H$ is three-colourable, then so is G or H .*

In other words, the undirected cycle C_3 is multiplicative. Of course, every even cycle C_{2k} is multiplicative, since it is homomorphically equivalent to K_2 , which is multiplicative. In fact, it can be shown that all odd cycles are also multiplicative.

Theorem 2.33 [132] *Every undirected cycle C_n is multiplicative.*

By contrast, we don't know whether or not the Product Conjecture itself holds for any $k > 3$. In fact, the first multiplicative graphs K with $\chi(K) > 3$ have been discovered only recently [324].




FIG. 2.7. For directed cycles we have $\vec{C}_2 \times \vec{C}_3 \simeq \vec{C}_6$.

In the case of digraphs, we must focus on irreflexive digraphs. Figure 2.7 illustrates the fact that $\vec{C}_2 \times \vec{C}_3 \simeq \vec{C}_6$. Since $\vec{C}_2 \not\rightarrow \vec{C}_6$ and $\vec{C}_3 \not\rightarrow \vec{C}_6$, it follows that \vec{C}_6 is not multiplicative.

Multiplicativity of some digraphs is easy to establish using duality.

Theorem 2.34 *Each transitive tournament \vec{T}_n is multiplicative.*

Proof Suppose G, H are digraphs with $G \not\rightarrow \vec{T}_n, H \not\rightarrow \vec{T}_n$. According to Proposition 1.20 this means that there exist homomorphisms $f : \vec{P}_n \rightarrow G, g : \vec{P}_n \rightarrow H$; therefore $F : \vec{P}_n \rightarrow G \times H$ defined by $F(x) = (f(x), g(x))$ is a homomorphism. This means that $G \times H \not\rightarrow \vec{T}_n$ by another application of Proposition 1.20. \square

Duality can also be used in the following cases.

Theorem 2.35 *Each directed path \vec{P}_n is multiplicative.*

Proof Suppose G, H are digraphs with $G \not\rightarrow \vec{P}_n, H \not\rightarrow \vec{P}_n$. According to Exercise 14 in Chapter 1 this means that there exist homomorphisms $f : P \rightarrow G, g : Q \rightarrow H$ for some oriented paths P, Q of net length $n + 1$. The proof will be concluded as before, if we prove the following lemma. \square

Lemma 2.36 *If P and Q are oriented paths of net length ℓ , then $P \times Q$ contains an oriented path of net length ℓ .*

If W is the walk v_0, v_1, \dots, v_k , then a *subwalk* of W is any walk v_i, v_{i+1}, \dots, v_j with $0 \leq i \leq j \leq k$, and a *prefix* of W is a walk v_0, v_1, \dots, v_j with $0 \leq j \leq k$.

Proof Assume that P has consecutive vertices p_0, p_1, \dots, p_a and Q consecutive vertices q_0, q_1, \dots, q_b . By taking subpaths of P and Q , we can make sure that the net lengths of all prefixes of P and Q are between 0 and ℓ . (These net lengths are precisely the levels as defined after Proposition 1.13.)

Note that a path of net length ℓ exists if and only if a walk of net length ℓ exists, since P, Q , and hence also $P \times Q$ are balanced, and so any closed walk has net length zero. We shall show the existence of a walk of net length ℓ between specific vertices.

Claim $P \times Q$ contains a walk of net length ℓ from (p_0, q_0) to (p_a, q_b) .

Figure 2.8 illustrates the product of two oriented paths. Depending on the orientations of the arcs incident to $u \in P - p_0 - p_a, v \in Q - q_0 - q_b$, the product $P \times Q$ has the vertex (u, v) either with no neighbours, or with four inneighbours, or with four outneighbours, or with two inneighbours or outneighbours. In any case, the number of neighbours of such (u, v) is always even. Moreover, when $v = q_0$, a vertex (u, v) still has even degree when u has level 0 in Q , because of the assumption that no vertex has negative level in Q . For similar reasons, the degree of (u, v) is even when $v = q_b$ and u has level ℓ in Q ; when $u = p_0$ and v has level 0 in Q ; and when $u = p_a$ and v has level ℓ in Q . (The levels are marked in the illustration.)

Consider the following traversal of $P \times Q$: starting at (p_0, q_0) , continue taking arcs in either direction, as long as possible while never taking an arc already taken. This traversal results in a walk W without repeated arcs. Suppose W' is a subwalk of W , from (p_0, q_0) to some (u, q_0) . Since the projection of W' onto Q has net length 0, so must the projection of W' onto P , i.e., the level of u in P must be 0. It follows from this observation (and the similar cases of subwalks to $(u, q_b), (p_0, v), (p_a, v)$), that W must end in (p_a, q_b) , since all other vertices encountered have an even degree. \square




FIG. 2.8. The product of two oriented paths of net length two.

Theorem 2.37 *The directed cycle \vec{C}_n is multiplicative if and only if n is a prime power.*

Proof If n is not a prime power, it can be written as $n = pq$ where p and q are relatively prime. It is then easy to see, as in Fig. 2.7, that $\vec{C}_p \times \vec{C}_q \simeq \vec{C}_n$. Since $\vec{C}_p \not\rightarrow \vec{C}_n, \vec{C}_q \not\rightarrow \vec{C}_n$, the digraph \vec{C}_n is not multiplicative.

Suppose now that n is a power of a prime. We shall again use duality—this time the duality of Corollary 1.18, asserting that $G \not\rightarrow \vec{C}_n$ if and only if some oriented cycle C of net length not divisible by n has $C \rightarrow G$.

Suppose $G \not\rightarrow \vec{C}_n$ because $C \rightarrow G$ and $H \not\rightarrow \vec{C}_n$ because $D \rightarrow H$, where C is an oriented cycle of net length $m_1 > 0$ not divisible by n , and D is an oriented cycle of net length $m_2 > 0$ not divisible by n . Let $m = k_1 m_1 = k_2 m_2$ be the least common multiple of m_1 and m_2 . Since n is a prime power, m is also not divisible by n . We shall construct an oriented cycle X of net length m such that $X \rightarrow C \times D$, whereby proving the theorem. Instead of arguing about X , we shall look at its image in $C \times D$, which will be a closed walk of net length m in $C \times D$.

We now assume that C and D are given by consecutive vertices $c_0, c_1, \dots, c_a = c_0$ respectively $d_0, d_1, \dots, d_b = d_0$ such that the net lengths of all prefixes c_0, c_1, \dots, c_x , respectively d_0, d_1, \dots, d_y , are nonnegative. (This is easily accomplished by making the right choices for c_0 and d_0 . Note however, that these net lengths may exceed m_1 or m_2 .) Let P be a path obtained by going around C exactly k_1 times, and Q a path obtained by going around D exactly k_2 times. To be precise, P has distinct vertices $c_0^1, c_1^1, \dots, c_a^1 = c_0^2, \dots, c_a^2, \dots, c_0^{k_1}, \dots, c_a^{k_1}$ and $c_j^i c_{j+1}^i$ being a forward or backward arc exactly as $c_j c_{j+1}$ in C ; and similarly for Q . Both these paths have net length m and all their prefixes have nonnegative

net lengths. Suppose M is an integer strictly greater than the net length of all prefixes of P and Q . Consider now extending P by going around C again until the first time the net length reaches M ; call this path P' , and let P'' be the extension, i.e., the subpath of P' obtained by removing the prefix P . We obtain Q' and Q'' in the same way from Q and D . Let c and d be the last vertices of P' and Q' respectively.

In the paths P' and Q' all prefixes have net lengths between 0 and M , hence the Claim from the proof of Lemma 2.36 ensures that $P' \times Q'$ has a walk W' of net length M from (c_0^1, d_0^1) to (c, d) . In the paths P'' and Q'' all prefixes have net lengths between 0 and $M - m$, hence the same Claim ensures that $P'' \times Q''$ (which is a subgraph of $P' \times Q'$) contains a walk W'' from $(c_0^{k_1}, d_0^{k_2})$ to (c, d) of net length $M - m$. Concatenating W' with the reverse of W'' yields a walk in $P' \times Q'$ of net length m , which corresponds in $C \times D$ to a closed walk, since $c_0^{k_1}$ and c_0 (and similarly $d_0^{k_2}$ and d_0) correspond to the same vertex. \square

Proposition 2.38 *Let K be the symmetric digraph corresponding to the graph K_k . If $k > 2$, then K is nonmultiplicative.*

Proof Let G be the transitive tournament \vec{T}_{k+1} with vertices $1, 2, \dots, k, k+1$ and arcs ij for all $i < j$, and let H be obtained from G by reversing the arc $1(k+1)$. Then $G \not\rightarrow K$ and $H \not\rightarrow K$, because the underlying graphs of G and H are not k -colourable. However, the digraph $G \times H$ has special structure: all vertices (i, j) with $i \neq j$ are isolated, except for $(1, k+1)$ and $(k+1, 1)$, which are adjacent. The only other component of $G \times H$ is on vertices (i, i) , $i = 1, 2, \dots, k+1$, and is isomorphic to \vec{T}_{k+1} with the arc $1(k+1)$ removed. Thus a homomorphism $f : G \times H \rightarrow K$ is obtained by identifying vertices $(1, 1)$ and $(k+1, k+1)$. \square

Proposition 2.38 implies that the Product Conjecture should not be applied to digraphs—for each $k > 2$ there are irreflexive digraphs G and H which are not k -colourable but such that $G \times H$ is k -colourable. (Recall that we take the chromatic number of a digraph to be the chromatic number of its underlying undirected graph.) Therefore, for *digraphs* we have no shortage of highly chromatic multiplicative graphs.

It is certainly discouraging that not much more is known to support the Product Conjecture. One might expect to be able to prove that the chromatic number $\chi(G \times H)$ is bounded from below at least by some function of the minimum of $\chi(G), \chi(H)$, but no such results are known. It seems possible that there are graphs G, H of arbitrarily high chromatic number, yet having $\chi(G \times H) \leq 10$, say. (In fact, this leads some to conclude the conjecture may well be false.) Somewhat surprisingly, it turns out that if one can ensure that $\chi(G \times H)$ reaches ten by requiring G, H to have a high chromatic number, then the same will hold for any constant c in place of ten.

Theorem 2.39 *As $k \rightarrow \infty$, the minimum chromatic number of the product of two graphs with chromatic number k , either tends to ∞ , or is always smaller than ten.*

Let $u(k)$ be the minimum chromatic number of any product $G \times H$ of graphs G, H with $\chi(G) = \chi(H) = k$, and let, similarly, $d(k)$ be the minimum chromatic number of any product $G \times H$ of digraphs G, H with $\chi(G) = \chi(H) = k$. The Product Conjecture claims that $u(k) = k$ for all k , while Proposition 2.38 implies that $d(k) < k$ for all $k > 2$. We aim to prove that if $u(k) \geq 10$ for some k then there is no upper bound to the values $u(k)$. Our first step is to prove a similar result for $d(k)$.

Theorem 2.40 *If $d(k) \geq 4$ for some k , then $\lim_{k \rightarrow \infty} d(k) = \infty$.*

Proof Assume that $\lim_{k \rightarrow \infty} d(k) = c < \infty$. Since $d(k)$ is a nondecreasing function of k we must have $d(k) = c$ for all sufficiently large k , say all $k \geq k_0$. Let $k_1 = 2^{k_0}$, and let G, H be digraphs of chromatic number k_1 with $\chi(G \times H) = c$. Then according to Lemma 2.21 $\chi(L(G)) \geq k_0, \chi(L(H)) \geq k_0$, and hence $\chi(L(G) \times L(H)) \geq c$. However, it is easy to check that $L(G) \times L(H) \simeq L(G \times H)$, and hence the Lemma 2.21 also implies that $\chi(L(G) \times L(H)) \geq c$. This means, again according to Lemma 2.21, that $c = \chi(G \times H) > \binom{c-1}{\lfloor (c-1)/2 \rfloor}$. Therefore, we must have $c \leq 4$, i.e., all $d(k) \leq 4$.

It takes only a small additional effort to reduce the bound to $d(k) \leq 3$, proving the theorem by contradiction. Thus suppose that some $d(k) = 4$, i.e., that $d(k) = 4$ for all $k \geq k_0$. Let again $k_1 = 2^{k_0}$, and also let $k_2 = 2^{k_1}$. Take digraphs G, H with $\chi(G) = \chi(H) = k_2, \chi(G \times H) = 4$. It follows, as above, that the chromatic number of $L(L(G \times H))$ is still at least four. On the other hand, the digraph $G \times H$ is four-colourable, i.e., homomorphic to the complete symmetric digraph K_4 . It follows that $L(L(G \times H)) \rightarrow L(L(K_4))$, and in Exercise 18 we ask the reader to verify that $L(L(K_4))$ is three-colourable, a contradiction. \square

Given a digraph G , we denote by G^{-1} the digraph obtained from G by reversing all arcs. It is straightforward to see that the above proof actually applies to the following modified function $d'(k)$. Let $d'(G, H) = \max\{\chi(G \times H), \chi(G \times H^{-1})\}$, and define $d'(k)$ to be the minimum value of $d'(G, H)$ over all digraphs G, H with $\chi(G) = \chi(H) = k$. Then we have the following.

Corollary 2.41 *If $d'(k) \geq 4$ for some k , then $\lim_{k \rightarrow \infty} d'(k) = \infty$.* \square

We are now in a position to prove our Theorem 2.39, by showing that $u(k)$ is always less than 10, or goes to infinity.

Corollary 2.42 *If $u(k) \geq 10$ for some k , then $\lim_{k \rightarrow \infty} d(k) = \infty$.*

Proof Let G, H be digraphs with the underlying graphs G', H' respectively. Then the underlying graph of $G \times H$ is a subgraph of $G' \times H'$, whence $\chi(G \times H) \leq \chi(G' \times H')$ and hence $d'(k) \leq u(k)$. On the other hand, if G, H are graphs and G', H' the corresponding symmetric digraphs (obtained by replacing

each undirected edge by the two opposite arcs), then the edge set of $G' \times H'$ is partitioned into $E(G \times H)$ and $E(G \times H^{-1})$, and thus

$$\chi(G' \times H') \leq \chi(G \times H) \cdot \chi(G \times H^{-1})$$

(cf. Exercise 11 in Chapter 1), implying that $u(k) \leq (d'(k))^2$. Corollary 2.41 implies that $d'(k)$ is either bounded by three or is unbounded, and we may conclude that $u(k)$ is either bounded by nine or is unbounded. \square

2.7 Projective digraphs and polymorphisms

Recall that the numeric power G^t of a digraph G denotes the product of t copies of G , i.e., the digraph $G^t = G \times G \times \cdots \times G$. Recall also that an endomorphism of G is a homomorphism of G to G . A homomorphism of G^t to G , for any integer $t \geq 2$, is called a *polymorphism*, and the integer t is called the *order* of the polymorphism. It is clear that any projection π_j of G^t to G is a polymorphism of G . In this section we focus on polymorphisms with an additional property. We say that a polymorphism $f : G^t \rightarrow G$ is *idempotent* if $f(u, u, \dots, u) = u$, for all $u \in V(G)$. Clearly all projections are idempotent. We shall encounter other important kinds of polymorphisms later in this chapter, and in Chapter 5.

We say that a digraph G is *projective* if the only idempotent polymorphisms $f : G^t \rightarrow G$ are the projections. More specifically, we also say that a digraph is t -projective if the only polymorphisms of order t are the projections; thus G is projective if and only if it is t -projective for all $t \geq 2$.

Theorem 2.43 *Each complete graph K_n is two-projective.*

Proof Let $G = K_n$. If $f : G \times G \rightarrow G$ is a homomorphism, then $f(u, v) = f(u', v')$ implies that $u = u'$ or $v = v'$, as otherwise (u, v) and (u', v') are adjacent in $G \times G$. Now assume that f is idempotent. First, we note that this means that each $f(u, v)$ is equal to either u or v : if $f(u, v) = i$ then $f(u, v) = f(i, i)$, whence $u = i$ or $v = i$. Finally, we note that if $f(u, v) = u$ for $u \neq v$, then for all (u', v') we must have $f(u', v') = u'$. Indeed, suppose that $f(u', v') = v'$ for some $u' \neq v'$, and consider $f(v', u)$. If $f(v', u) = v'$ then $f(v', u) = f(u', v')$, and if $f(v', u) = u$ then $f(v', u) = f(u, v)$, both of which are impossible unless $u = v'$. However, $u = v'$ is also impossible, as $f(u, v) = f(u', v')$ would imply that $u = u' = v'$ or $v = v' = u$. Hence f is the first projection. \square

It turns out that a two-projective graph is always projective, and that asymptotically almost all graphs are projective (cf. Exercise 12 in Chapter 3). We will only need the following easier fact.

Corollary 2.44 *Each complete graph K_n with $n \geq 3$ is projective.*

Proof We shall prove that the only idempotent homomorphisms of K_n^t to K_n are the projections. The theorem implies this for $t = 2$, and we proceed by induction on t . Thus let f be an idempotent homomorphism of K_n^{t+1} to K_n . For any i, j with $i < j$, we let $g_{i,j}$ be the following mapping of $V(K_n^t)$ to $V(K_n)$:

$g_{i,j}(x_1, x_2, \dots, x_t) = f(y_1, y_2, \dots, y_{t+1})$, where $y_a = x_a$ for $a = 1, 2, \dots, j-1$, then $y_j = x_i$, and finally $y_a = x_{a-1}$ for $a = j+1, j+2, \dots, t+1$. It is easy to check that $g_{i,j}$ is an idempotent homomorphism of K_n^t to K_n , and hence a projection, by the induction hypothesis. Suppose first that $g_{i,j}$ is the a -th projection, where $a < j$ and $a \neq i$. We claim that this implies that f is also the a -th projection. Indeed, let $v = f(y_1, y_2, \dots, y_{t+1})$, and consider the following vertex $(z_1, z_2, \dots, z_{t+1})$ of K_n^{t+1} . For $b \neq i, j, a$, we let z_b be any vertex different from y_b ; for $b = i, j$, we let $z_i = z_j$ be any vertex different from both y_i and y_j (recall that $n \geq 3$). Finally, for $b = a$, consider all possible vertices $z_a \neq y_a$. Since $g_{i,j}$ is the a -th projection, we must have $f(z_1, z_2, \dots, z_{t+1}) = z_a$, and since $(z_1, z_2, \dots, z_{t+1})$ is adjacent to $(y_1, y_2, \dots, y_{t+1})$ in $V(K_n^t)$, we also have $f(z_1, z_2, \dots, z_{t+1}) \neq v$. Therefore, $v \neq z_a$, for all $z_a \neq y_a$, i.e., $v = y_a$. In other words, f is also the a -th projection. A similar calculation shows that if $a \geq j$ then f is the $(a+1)$ -st projection. This argument will apply unless for any pair i, j with $i < j$, the homomorphism $g_{i,j}$ is the i -th projection. In particular, this means that $f(x_1, x_2, \dots, x_{t+1}) = v$ as soon as at least two of the arguments x_i are equal to v . This is clearly impossible if $t \geq 4$. In the case $t = 3$, it is also impossible if there exist three distinct vertices a, b, c in K_n , as $f(a, b, c)$ would have to be adjacent to $f(b, a, a) = a$, $f(b, c, b) = b$, and $f(c, c, a) = c$, and thus equal to some d different from a, b, c . However, in that case, $f(d, d, d) = d$ would have to be adjacent to $f(a, b, c) = d$ as well, a final contradiction. \square

Corollary 2.45 *If $n \geq 3$ then any homomorphism $f : K_n^t \rightarrow K_n$ is the composition $f = \phi \circ \pi_i$ where ϕ is an automorphism of K_n (a permutation of vertices) and π_i is the i -th projection.*

Proof The mapping $\psi(x) = f(x, x, \dots, x)$ is an automorphism of K_n , and $\psi \circ f$ is an idempotent homomorphism of K_n^t to K_n , hence some projection π_i . Letting $\phi = \psi^{-1}$ completes the proof. \square

2.8 The retract

A *retraction* is a homomorphism r of a digraph G onto a subgraph H which satisfies $r(x) = x$ for all vertices x of H ; if H admits a retraction from G it is called a *retract* of G . Retractions are at the heart of the problem of extending homomorphisms.

Proposition 2.46 *Suppose that H is a subgraph of G , and $f : H \rightarrow K$ is a homomorphism. Let G' be the digraph obtained from the disjoint union $G + K$ by identifying each $x \in V(H)$ with $f(x) \in V(K)$.*

Then there exists a homomorphism $F : G \rightarrow K$ with $F(x) = f(x)$ for all $x \in V(H)$ if and only if K is a retract of G' .

Proof Note that the digraph G' contains K as a subgraph, but G' might not contain G or H , since distinct vertices of H may be identified (if they have the same image under f). Any retraction $r : G' \rightarrow K$ defines a homomorphism $F : G \rightarrow K$ where these identified vertices are assigned the same image. Clearly

this homomorphism F extends f . Conversely, any homomorphism $F : G \rightarrow K$ extending f can be used to define a retraction $r : G' \rightarrow K$ since those vertices of G that have been identified already have the same image under f . It is easy to see that r is a retraction. \square

More on the relation between retracts and homomorphism extensions can be found in Exercise 16.

Suppose H is a retract of G . We then have a retraction homomorphism $r : G \rightarrow H$ and the inclusion homomorphism $i : H \rightarrow G$ (and $r \circ i$ is the identity endomorphism on H). Thus a retract of a digraph is homomorphically equivalent to it, and hence inherits many basic properties of it—e.g., for graphs these include the chromatic number, the odd girth, and so on (Corollary 1.8, Exercise 2 in Chapter 1).

For complete graphs, we can also nicely formulate multiplicativity as a property related to retracts.

Proposition 2.47 *K_n is multiplicative if and only if K_n is a retract of graphs G or H , whenever it is a retract of $G \times H$.*

Proof It is easy to verify from the definitions that (up to isomorphism)

$$(G + K_n) \times (H + K_n) = (G \times H) + (G \times K_n) + (K_n \times H) + (K_n \times K_n).$$

Suppose that whenever K_n is a retract of a product of two graphs, it is a retract of one of the factors, and let $G \times H \rightarrow K_n$. According to the above equation, we must have $(G + K_n) \times (H + K_n) \rightarrow K_n$. Since $(G + K_n) \times (H + K_n)$ contains K_n , and K_n is a core, it follows that K_n is a retract of $(G + K_n) \times (H + K_n)$. Therefore, K_n is a retract of $G + K_n$ or of $H + K_n$, and hence $G \rightarrow K_n$ or $H \rightarrow K_n$, i.e., K_n is multiplicative.

Conversely, suppose that K_n is multiplicative, and assume it is a retract of the graph $G \times H$. Then $G \times H \rightarrow K_n$ and hence $G \rightarrow K_n$ or $H \rightarrow K_n$. Since K_n is a core, this means that K_n is a retract of G or of H . \square

Surprisingly the following weaker version of the Product Conjecture is implicit in what we already proved.

Theorem 2.48 *K_n is a retract of graphs G or H whenever it is a retract of $G \times H$ and both G and H are connected.*

Proof Suppose G and H are connected, and K_n is a retract (and thus a subgraph) of $G \times H$. Since the image of K_n under each projection is isomorphic to K_n , we may assume that $K_n \times K_n$ is a subgraph of $G \times H$, which contains the retract K_n as the subgraph induced by the vertices (i, i) , $i = 1, 2, \dots, n$ of $K_n \times K_n$. Consider now a retraction $r : G \times H \rightarrow K_n$. The restriction of r to $K_n \times K_n$ is an idempotent homomorphism to K_n , and hence by Theorem 2.43 a projection. Without loss of generality, assume that r restricted to $K_n \times K_n$ is the projection π to the first coordinate, i.e., $r(i, j) = i$ for all $i, j = 1, 2, \dots, n$. Note that if $r(u, j') = r(u, j'') = k$ for some $u \in V(G)$, $j' \neq j''$, then $r(v, j) \neq k$

for all v adjacent to u in G and all $j = 1, 2, \dots, n$, since each such (v, j) is adjacent to (u, j') or (u, j'') . Since G is connected, it easily follows that for each $u \in V(G)$ there exist two vertices x_u, y_u in K_n with $r(u, x_u) = r(u, y_u)$. If we denote by $f(u)$ this common value $r(u, x_u) = r(u, y_u)$, we easily see (cf. the proof of Proposition 2.24) that f is a retraction of G to the K_n induced by $1, 2, \dots, n$. \square

2.9 Isometric trees and cycles

In this section we focus on graphs with loops allowed. There are some natural necessary conditions for H to be a retract of G . For instance, as we have observed above, each retract of G has the same chromatic number as G . (The chromatic number of a graph which has a loop is taken to be infinity.) It is also clear that a retract H of a graph G is necessarily an induced subgraph of G . In fact a stronger property holds, defined as follows. We say that a subgraph H of a graph G is *isometric* if $d_H(x, y) = d_G(x, y)$ for any two vertices x, y of H . (Recall that $d_H(a, b)$ denotes the distance, i.e., the number of edges in a shortest path, from a to b in G ; loops do not make a difference here.)

Proposition 2.49 *Every retract of G is an isometric subgraph of G .*

Proof Suppose $r : G \rightarrow H$ is a retraction. If P is a path of length k joining x and y , then $r(P)$ is a walk of length k joining $r(x) = x$ and $r(y) = y$. \square

In some cases the isometry condition is also sufficient.

Proposition 2.50 *Let H be a reflexive path. If H is an isometric subgraph of G , then H is a retract of G .*

Proof Suppose the vertices of the path H are consecutively named $0, 1, \dots, k$, and assume H is an isometric subgraph of a graph G with loops allowed. We can define a retraction $r : G \rightarrow H$ by $r(v) = \min(k, d_G(0, v))$. (We can take $d_G(0, v)$ infinite if v and 0 belong to different components of G .) It is clear that $r(i) = i$ for vertices i of the path H , by isometry. If u and v are adjacent vertices of G , then their distances from 0 in G are either equal, or differ by one. Since H is reflexive, it is easy to check that this implies that $r(u)$ and $r(v)$ are adjacent. \square

In the next section we generalize the result to any reflexive *tree*.

For reflexive cycles, the isometry condition is not sufficient, as we shall see in Fig. 2.13. However, if a reflexive cycle H is not only isometric, but actually a shortest nontrivial (meaning other than a loop) cycle in G , then we still have a retraction.

Proposition 2.51 *Let H be a reflexive cycle. If H is a shortest nontrivial cycle in G , then H is a retract of G .*

Proof It is easy to see that H satisfies the following property. Suppose u and v are two vertices of H , and P_1, P_2 the two simple paths joining u and v in H .

If P is any path in G different from the paths P_1, P_2 , then the length of P is at least the length of the longer of P_1, P_2 . This follows from the fact that H is the shortest cycle in G , since the union of P and either P_i contains a cycle. (We note that the isometry of H only implies that the length of P is at least the length of the shorter of P_1, P_2 .) This means that when we remove an arbitrary edge of H we obtain a reflexive path H' isometric in G' , and hence a retraction of G' to H' is implied by the previous proposition. It is easy to see that the same mapping is also a retraction of G to H . \square

For graphs (without loops), similar results hold when H is an isometric path or shortest cycle in G , and G is bipartite (Exercise 26). However, when G is not bipartite, a tree or an even cycle will not be a retract of G , if the chromatic number of G is greater than three, by Corollary 1.8. A shortest odd cycle in G may also fail to be a retract of G , even if G has chromatic number three. The graph in Fig. 2.9 does not retract to a five-cycle, in fact it does not have any proper retracts, in other words, it is a core. (Any homomorphic image in which two vertices of degree three are identified must contain a loop or a triangle, and hence is not a subgraph of the pictured graph.) Nevertheless, there is a natural class of graphs, with chromatic number three, in which a shortest odd cycle always is a retract. We say that a graph G is *nearly bipartite* if it admits an orientation in which the net length of each cycle is at most one. Observe that a graph is bipartite if and only if it admits an orientation in which each cycle has net length zero, i.e., a balanced orientation. The following theorem will be derived as Corollary 6.14.




FIG. 2.9. A three-colourable core.

Theorem 2.52 *If H is a shortest odd cycle in a nearly bipartite graph G , then H is a retract of G .*

It follows that each nearly bipartite graph is three-colourable. There is in fact an elegant characterization of the class of nearly bipartite graphs.

Theorem 2.53 [120] *A graph G is nearly bipartite if and only if it does not contain as an induced subgraph any of the graphs in Fig. 2.10.*




FIG. 2.10. The forbidden subgraphs.

The dashed lines and the dotted lines are pairwise disjoint paths, with the dashed lines representing paths with at least one edge, while the dotted lines may have length zero. The lengths of the paths are such that the marked cycles have odd length. In general, a *subdivision* of a graph G is a graph G' obtained from G by replacing the edges by pairwise disjoint paths. (Thus, for instance, A is a subdivision of K_4 .) If all the replacing paths have even length, we say that G' is an *odd subdivision* of G .

It is easy to construct examples of graphs containing one of the forbidden subgraphs which nevertheless retracts to a shortest odd cycle. It suffices, for instance, to take A in which abc is a triangle, and at least one of the paths from d to a, b, c has even length. Nevertheless, we have the following characterization.

Corollary 2.54 *A graph G is nearly bipartite if and only if any odd subdivision G' of G retracts to each shortest odd cycle in G' .*

Proof If G is nearly bipartite and G' is an odd subdivision of G , then clearly G' is also nearly bipartite—it suffices to orient the odd paths alternately forward and backward, starting and ending with the orientation of the edge they replaced. If G is not nearly bipartite, then it contains an induced subgraph from Fig. 2.10. We shall construct an odd subdivision of that subgraph which does not retract to a shortest odd cycle.

Consider first the graph B in Fig. 2.10. For each of the vertex pairs a, b , and c, d , and e, f , one of the paths joining them must be odd and the other even (so that the cycle formed by these two paths is odd). Suppose that k is a positive integer greater than the lengths of all three dotted paths b, c and d, e and f, a . A suitable odd subdivision B' of B has all the paths of the lengths indicated in Fig. 2.11, where $l > 2k$. The three shortest odd cycles in B' have length $4l + 1$, and we focus on the cycle C formed by the two paths joining a and b . Suppose that $r : B' \rightarrow C$ is a retraction. Then $r(f)$ is a vertex of C of distance (in C) less

than k from a . It is easy to see that this implies that $r(e)$ is of distance at most k from b . A similar argument shows that $r(c)$ is a vertex of distance less than k from b , and hence $r(d)$ a vertex of distance at most k from a . However, this is impossible, as $r(d)$ and $r(e)$ have distance less than k and $2l > 3k$. Therefore B' does not retract to C .




FIG. 2.11. A subdivision of B that does not retract to a shortest odd cycle.

A similar proof applies to the graph A in Fig. 2.10. First, we note that the cycle including a, b and d but not c must also be odd, since it is modulo two sum of the cycles formed by a, b, c and b, c, d and a, c, d . If one of these four odd cycles has all three sides odd, we have either all indicated paths of odd length, or three of them starting from one vertex are even (Fig. 2.12(a), (b)). If each of the odd cycles has two even sides, we must have the situation depicted in Fig. 2.12(c).




FIG. 2.12. Possible parities of subdivisions of A .

A suitable odd subdivision A' of A then has all even lengths equal to $2l$ and

all odd lengths equal to $2l - 1$. In all three cases, the cycle C formed by a, b, c is a shortest odd cycle. Suppose r is a retraction of A' to C . In the case (a), the path from a to d and then b can only retract to the path from a to c and then to b , as that is the only even walk in C from a to b . However, that would mean that $r(d) = b$, and hence the path from d to b would have to be mapped to a closed walk of length $2l - 1$ in C , which is impossible. The same argument shows that A' does not retract to C in case (c) either. In case (b) we similarly conclude that $r(d)$ is a neighbour of b on C , but then the path from d to b of length $2l$ would have to map to an even walk from b to a neighbour, which is also impossible. Thus in all cases there is no retraction of A' to a shortest odd cycle. \square

Note that each bipartite graph is nearly bipartite; the corollary applies to a bipartite graph vacuously (there are no odd cycles).

In the next few sections we shall focus on reflexive graphs. Recall that $G \rightarrow H$ if H is reflexive, since a constant mapping to a vertex with a loop is always a homomorphism. However, it is not clear when a subgraph H of a reflexive graph G is a retract of G . The cop and robber example from Chapter 1, illustrates a context where it is useful to assume that the graph is reflexive. The loops allow the cop and the robber to take a turn in which no move is made. If there are no loops the cop and the robber have to move at each turn; the robber can win by placing himself in the vertex chosen by the cop, and then duplicating all of her moves. We shall focus on reflexive graphs throughout the rest of this chapter, even though much of the theory can be extended to graphs with loops allowed.

Perhaps the main reason that reflexive graphs are often most natural in the context of retractions is due to the fact that reflexivity makes graphs with the usual distance function more naturally into (integer-valued) metric spaces. More specifically, it is clear that each connected graph G defines a metric space $S(G)$ on $V(G)$, via the integer-valued distance function d_G . Each graph homomorphism $f : G \rightarrow H$ can be viewed as a *nonexpansive mapping* $S(G) \rightarrow S(H)$, i.e., it satisfies $d_H(f(x), f(y)) \leq d_G(x, y)$ for any $x, y \in V(G)$. If G, H are reflexive graphs, then the converse also holds, i.e., every nonexpansive mapping $S(G) \rightarrow S(H)$ is a homomorphism $G \rightarrow H$. Therefore, graph retractions of reflexive graphs behave much like metric space retractions.

2.10 Reflexive absolute retracts

In this section all graphs are assumed to be reflexive. We have noted that for a reflexive path, isometry suffices to ensure the existence of a retraction. In this section, we shall characterize graphs with this property. A connected graph H is called an *absolute retract* if it is a retract of any graph G that contains H as an isometric subgraph. Thus Proposition 2.50 asserts that every reflexive path is an absolute retract. In Fig. 2.13, we illustrate the fact that the hexagon C_6 is not an absolute retract. (Note that this remains true regardless of any loops it may or may not have.)

The problem is that the hexagon contains three vertices a, b, c of which any two have a common neighbour, but all three do not have a common neighbour.




FIG. 2.13. The outer hexagon, with filled vertices, is an isometric subgraph, but not a retract, of G .

Such a ‘hole’ in H can be exploited as in the figure to produce a graph G which contains H as an isometric subgraph but not as a retract. Formally, we define a *hole* of H to consist of vertices a_1, a_2, \dots, a_k and positive integers x_1, x_2, \dots, x_k , with $k \geq 3$ (k is called the *size of the hole*), such that

- H contains no vertex v with $d_H(v, a_i) \leq x_i$ for all $i = 1, 2, \dots, k$, and
- for every $i, j = 1, 2, \dots, k$, we have $d_H(a_i, a_j) \leq x_i + x_j$.

The first condition states the absence of a vertex satisfying all the distance requirements; the second condition, ensures that there is for each $i, j = 1, 2, \dots, k$ a vertex $v \in V(H)$ with $d_H(v, a_i) \leq x_i$ and $d_H(v, a_j) \leq x_j$, i.e., that the distance conditions can be satisfied at least pairwise.

Theorem 2.55 *A graph H is an absolute retract if and only if it has no holes.*

Proof Suppose the vertices a_1, a_2, \dots, a_k , and integers x_1, x_2, \dots, x_k , form a hole of H . Let G be obtained from H by introducing a new vertex v and disjoint paths from v to each a_i , of length x_i . It is easy to verify that the second condition of a hole guarantees that H is an isometric subgraph of G , and that there is no retraction $G \rightarrow H$ since there is no possible image for v . (The example G in Fig. 2.13 is constructed in this way from the hexagon, and the hole formed by a, b, c and $1, 1, 1$.)

On the other hand, suppose that H has no holes. If H is not an absolute retract, then there exists a graph G such that

- G contains H as an isometric subgraph,
- H is not a retract of G , and
- G has the smallest number of vertices.

We may assume that G has more vertices than H . Let v be a vertex of G that is not in H ; let a_1, a_2, \dots, a_n be an enumeration of the vertices of H and let $x_i = d_G(v, a_i)$. Since H has no holes, there must exist a vertex v' in H that satisfies $d_H(v', a_i) \leq x_i, i = 1, 2, \dots, n$. Let G' be the quotient of G obtained by identifying v and v' .

We first observe that G' still contains H as a subgraph, in fact as an isometric subgraph. Indeed, any path of length less than $d_H(a, b)$ joining a, b in G' could not be a path in G , and hence would have to consist of a path in G joining v and one of a, b and a path in G joining v' and the other one of a, b . Say, $d_{G'}(a, b) = d_G(v, a) + d_G(v', b) < d_H(a, b)$. However, this is impossible, as the condition $d_H(v', a) \leq d_G(v, a)$ implies that

$$d_H(a, b) \leq d_H(a, v') + d_H(v', b) \leq d_G(a, v) + d_G(v', b).$$

By the minimality of G , we know that there exists a retraction $r : G' \rightarrow H$. Clearly, r induces a retraction $G \rightarrow H$, in which the image of v is v' . (This retraction is the composition of r and the canonical homomorphism of G onto its quotient G' .) This contradicts the second item above, establishing the theorem. \square

Absolute retracts turn out to be an important class of graphs. We illustrate this with the following results.

A *graph variety* is a class \mathcal{V} of graphs which is closed under taking products and retracts. If \mathcal{F} is a family of graphs, then *the variety of \mathcal{F}* is the smallest variety containing \mathcal{F} . Varieties of graphs (and ordered sets) are studied as a means of developing a structure theory and a classification scheme [82].

Corollary 2.56 *The class of all absolute retracts is equal to the variety of reflexive paths.*

Proof We have noted that each reflexive path is an absolute retract. Thus we will have proved that the class of absolute retracts contains the variety of reflexive paths, if we can show that the product of absolute retracts is an absolute retract and a retract of an absolute retract is also an absolute retract.

This is routine to prove directly, but it also follows from our theorem, as it is easy to see that products and retracts of graphs without holes can not themselves have holes.

It remains to show that each absolute retract is in the variety of reflexive paths. It turns out that there is a natural mapping—which we call the standard isometric representation—of each graph H onto an isomorphic copy of H which is an isometric subgraph of a product of reflexive paths. When H is an absolute retract this copy is necessarily a retract of the product.

Let H be a connected graph with vertices a_1, a_2, \dots, a_n . The *standard isometric representation* ϕ of a connected graph H is defined as follows. Let Π be the product of reflexive paths P^1, P^2, \dots, P^n , where P^i is of length equal to the *eccentricity* of a_i , i.e., the maximum of $d_H(a_i, x)$ over all vertices x of H . Then we set

$$\phi(x) = (d_H(a_1, x), d_H(a_2, x) \cdots, d_H(a_n, x)).$$

It is easy to see that $\phi(H)$ is isomorphic to H . It is also straightforward to check that $\phi(H)$ is an isometric subgraph of Π . \square

We now introduce an important class of polymorphisms of order three. A *majority function* on a reflexive graph H is a homomorphism $m : H \times H \times H \rightarrow H$ such that $m(a, b, c) = x$ if at least two of the arguments a, b, c equal x . We shall also encounter majority functions again in Chapter 5 as a tool for designing polynomial time algorithms for the homomorphism and the retraction problems.

Corollary 2.57 *A reflexive graph H has a majority function if and only if it is an absolute retract.*

Proof First, suppose that H admits a majority function m . If H contains a hole $a_1, a_2, \dots, a_n, x_1, x_2, \dots, x_n$, then we may assume that $n \geq 3$ and there is no hole of size smaller than n . This implies that H contains a vertex u such that $d_H(u, a_i) \leq x_i$ for all $i \neq 1$, a vertex v with $d_H(v, a_i) \leq x_i$ for all $i \neq 2$, and a vertex w with $d_H(w, a_i) \leq x_i$ for all $i \neq 3$. We now derive the conclusion that $z = m(u, v, w)$ satisfies all the distance conditions, i.e., $d_H(z, a_i) \leq x_i, i = 1, 2, \dots, n$. This is a contradiction, proving that H cannot contain a hole, and hence is an absolute retract.

Clearly $m(u, a_1, a_1) = a_1$. Since v and w are both at most distance x_1 from a_1 , we conclude that

$$d_{H \times H \times H}((u, v, w), (u, a_1, a_1)) \leq x_1.$$

This implies

$$d_H(m(u, v, w), m(u, a_1, a_1)) = d_H(z, a_1) \leq x_1,$$

since m is a homomorphism. Similarly, z is at most distance x_2 from a_2 and x_3 from a_3 . For each $a_i, i = 4, 5, \dots, n$ we have $d_H(u, a_i) \leq x_i, d_H(v, a_i) \leq x_i, d_H(w, a_i) \leq x_i$, and hence

$$d_{H \times H \times H}((u, v, w), (a_i, a_i, a_i)) \leq x_i,$$

and we also conclude that

$$d_H(m(u, v, w), m(a_i, a_i, a_i)) = d_H(z, a_i) \leq x_i.$$

We now show that the class of graphs that admit a majority function forms a variety which contains all reflexive paths, and hence it also contains all absolute retracts. Let H and H' be two graphs with majority functions m_H and $m_{H'}$. Consider the graph $H \times H'$. It is easy to check that $m((u_1, u'_1), (u_2, u'_2), (u_3, u'_3)) = (m_H(u_1, u_2, u_3), m_{H'}(u'_1, u'_2, u'_3))$ is a majority function for $H \times H'$. Now suppose H is a retract of G , with $r : G \rightarrow H$ a retraction, and let m be a majority function on G . Then $r \circ (m|_{H \times H \times H})$ is easily seen to be a majority function on H . Hence the class of graphs that admit majority functions forms a variety.

To see that each reflexive path has a majority function, let (a, b, c) be a triple of vertices of some reflexive path P . Define $m(a, b, c)$ to be the 'middle vertex' of the triple. (This middle vertex is well defined for any three, possibly equal, vertices of a path.) It is easy to check this is indeed a majority function, since the paths are reflexive. \square

Corollary 2.58 *Every reflexive tree is an absolute retract.*

Proof Every reflexive tree T has the following natural majority function m . If u, v, w lie on a path P in T then $m(u, v, w)$ is defined as on P , i.e., it is the middle vertex. Otherwise $m(u, v, w)$ is the unique vertex m whose deletion from T puts all three vertices u, v, w in different components of $T - t$. It is not difficult to verify, by a case analysis, that m is indeed a majority function. \square

2.11 Reflexive dismantlable graphs

All graphs in this section shall also be assumed to be reflexive. (It is again the case that nearly everything can be extended to graphs with loops allowed, but we choose to keep our focus on reflexive graphs.) Let H be a graph, and u, v two vertices of H such that every neighbour of u (including u itself) is also a neighbour of v . Then there is a natural retraction of H to $H - v$, taking v to u (and fixing everything else). We call such retraction a *fold*. A graph is *dismantlable* if it can be reduced, by a sequence of folds, to one vertex. Dismantlable graphs are used to model physical processes such as phase transitions. They also turn out to be precisely those graphs on which the cop from Section 1.6 has a winning strategy.

Proposition 2.59 *A graph H is dismantlable if and only if the cop has a winning strategy.*

Proof Suppose first that H admits a winning strategy for the cop. Consider the last couple of moves in a particular game in which the cop won following her strategy even when the robber did his best to keep escaping. Say just before the last move of the robber, the cop was in a vertex v' and the robber in a vertex v . Wherever the robber moved next (including remaining at v), the cop would capture him. Thus every neighbour of v (including v) is also a neighbour of v' , i.e., H can be folded onto $H - v$. According to Proposition 1.33, the cop has a winning strategy in the retract $H - v$ of H . It follows by induction (on the number of vertices of H) that H is dismantlable.

On the other hand, suppose the graph H is dismantlable. Clearly, the one vertex graph (with a loop) admits a winning strategy for the cop. We will be finished by induction again, if we can prove that a graph in which the robber wins cannot be folded onto a graph in which the cop wins. (This is false for retracts in general !) Thus suppose v was folded onto v' , and assume that $H - v$ admits a winning strategy for the cop. Then she can use the same strategy on H , and if the robber is in v , the cop can play as though the robber was in v' . When the game ends in $H - v$, with the cop capturing the robber, either the cop actually caught the robber in H , or the cop is in v' and the robber is in v , whereupon she captures him on the next move—since v and v' are adjacent. \square

A graph that is not dismantlable may admit a sequence of folds, retracting to a subgraph with more than one vertex. Any such sequence of folds (not necessarily a maximal sequence) is called a *dismantling sequence*, and any retraction

that is the composition of a sequence of folds is called a *dismantling retraction*. Graphs that do not admit further folds are termed *stiff*. Dismantling sequences are an interesting graph theoretic example of a combinatorial structure called a *greedoid* [205]. Each maximal dismantling sequence in H leads to a stiff subgraph of H . This subgraph is a retract of H , but it may fail to be the core of H . (For instance the reflexive four-cycle C is itself stiff but its core is L , the one-vertex graph with a loop; in this case C retracts to L but does not fold to L .) Nevertheless, the following result highlights a property of stiff retracts that is similar to the property of cores.

Proposition 2.60 *For any graph H there is a unique (up to isomorphism) stiff graph G to which H retracts by a dismantling retraction.*

Proof Suppose X and Y are stiff subgraphs of H , and $f : H \rightarrow X, g : H \rightarrow Y$ are dismantling retractions. Consider the associated partition θ_f of the homomorphism f ; it has the parts $f^{-1}(x), x \in V(X)$, and each $f^{-1}(x)$ contains the vertex x . Let $A_x, x \in V(X)$, be a collection of sets of vertices of H obtained recursively as follows. Initially, each $A_x = \{x\}$. Then,

- whenever $s \in V(H)$ is adjacent to a vertex in every A_z such that z is adjacent to x in X , then we place s in the set A_x .

We now claim that each $A_x \subseteq f^{-1}(x)$. Indeed, this is true initially, and we can proceed recursively. Suppose this was true up to the point when s was placed into A_x , and let $x_z \in A_z \subseteq f^{-1}(z)$ be a vertex adjacent to s , for every z adjacent to x . Since this means that each $f(x_z) = z$, we must have $f(s)$ adjacent to all neighbours of x , and since X is stiff, we must have $f(s) = x$, i.e., $s \in f^{-1}(x)$.

Next we show that the dismantling retraction $g : H \rightarrow Y$ maintains an isomorphic copy of X , with the vertex corresponding to any x belonging to the set A_x . This is certainly true before the first fold of g , so suppose it was true up to a certain fold, and consider the next fold (in the sequence making up g). Say, that fold takes the vertex $x' \in A_x$ corresponding to x , to some vertex v . Since the neighbours of x' are all also neighbours of v , we must have $v \in A_x$. We claim that we can replace x' by v as the vertex corresponding to x . We have already observed that v will be adjacent to all vertices corresponding to the neighbours of x in X , since x' had this property. On the other hand, if u is not a neighbour of x in X , then the sets A_u and A_x have no edges joining them (being subsets of $f^{-1}(u)$ and $f^{-1}(x)$ respectively, cf. property 2 of Proposition 1.10), and so v is not adjacent to any vertex of A_u , including the vertex representing u . \square

Recall that $\text{HOM}(G, H)$ denotes the set of all homomorphisms of G to H . We now extend this notation to mean a *graph* whose vertices are the homomorphisms of G to H , with two homomorphisms f and f' adjacent if $f(v) = f'(v)$ for all vertices v of G , except possibly for one vertex. (Thus $\text{HOM}(G, H)$ is a reflexive graph.)

Proposition 2.61 *A graph H is dismantlable if and only if the graph $\text{HOM}(G, H)$ is connected, for each graph G .*

Proof Suppose first that H is a stiff graph with more than one vertex. Consider the graph $G = H \times K$, where K is the complete reflexive graph with two vertices, 1, 2. In other words, G has vertices (x, i) where $x \in V(H)$ and $i = 1, 2$, and $(x, i), (x', i')$ are adjacent if and only if x, x' are adjacent in H . (Since the graphs are reflexive, G contains all loops and all edges $(x, 1)(x, 2)$.) The first projection π of G to H , which takes each (x, i) to x , is a homomorphism, which is not adjacent (in $\text{HOM}(G, H)$) to any other homomorphism f of G to H . Otherwise π and f would differ on a unique vertex (x, i) of G . Hence they would agree on (x, i') where $i' \neq i$, i.e., $f(x, i') = x$; since (x, i) is adjacent to all neighbours of (x, i') in G , this would mean that $f(x, i) \neq x$ is adjacent to all neighbours of x in H , contradicting the fact that H is stiff. Since all constant maps are vertices of $\text{HOM}(G, H)$, we see that $\text{HOM}(G, H)$ has more than one vertex and hence is not connected.

Suppose next that H admits a dismantling retraction to a stiff graph with more than one vertex. We claim that some $\text{HOM}(G, H)$ is disconnected. Clearly, it suffices to prove that if H folds to $H - x$ and $\text{HOM}(G, H - x)$ is disconnected, then $\text{HOM}(G, H)$ is also disconnected. Let $f : H \rightarrow H - x$ be the fold, and let G be a graph such that $\text{HOM}(G, H - x)$ is disconnected. Then $\text{HOM}(G, H)$ must also be disconnected, otherwise any $\phi, \psi \in \text{HOM}(G, H - x)$ could be joined by the path $\alpha, \beta, \dots \in \text{HOM}(G, H)$ and hence the walk $f \circ \alpha, f \circ \beta, \dots$ would join them in $\text{HOM}(G, H - x)$.

On the other hand, if H is dismantlable, then consider a dismantling sequence of folds starting from $H = H_0, H_1, \dots$, and ending in $H_k = L$ where L is the one-vertex reflexive graph on $v \in V(H)$. Let f_i be the fold from H_i to H_{i+1} . For any homomorphism $\phi : G \rightarrow H$ we define $\phi_0 = \phi$, and $\phi_{i+1} = f_i \circ \phi_i$. It is easy to see that ϕ_k is the constant mapping onto v . Moreover, each $\phi_i \phi_{i+1}$ is an edge of $\text{HOM}(G, H)$. Since every vertex of $\text{HOM}(G, H)$ has a path to one fixed vertex, the graph $\text{HOM}(G, H)$ is connected. \square

While absolute retracts have nice properties concerning retractions from supergraphs, dismantlable graphs have nice properties with respect to retractions to subgraphs. We have the following connections between these two interesting classes of graphs.

Proposition 2.62 *Every absolute retract is dismantlable.*

Proof Let H be an absolute retract, and d the diameter of H . The result is obvious if $d = 1$, since then H is a complete graph. Thus assume that $d \geq 2$, and let a, b be two vertices of distance d in H . Form the graph G by adding to H a new vertex z , adjacent to a and all neighbours of a , and a path P of length $d - 1$ from z to b . It is clear that H is an isometric subgraph of G , and hence a retract of G . Under any retraction, the image a' of the vertex z is a vertex different from a , since a does not have a path of length $d - 1$ to b . But a and each neighbour of a is also a neighbour of a' . Hence H folds to $H - a$. Since $H - a$ is a retract of H it is also an absolute retract, and we can continue to dismantle H . \square

Not every dismantlable graph is an absolute retract, cf. Fig. 2.14. In the example, H is dismantlable in the order a, d, f, b, c, e , but has a hole with the vertices a, d, f and values 1, 1, 1, and hence is not an absolute retract. Note that this H contains three maximal cliques, the triangles abc, bde, cef , which are pairwise intersecting but do not all intersect. We call a graph *clique-Helly* if its maximal cliques have the *Helly property*, i.e., if every family of pairwise intersecting maximal cliques has a nonempty intersection.




FIG. 2.14. A dismantlable H that is not an absolute retract.

Proposition 2.63 *Every absolute retract is clique-Helly.*

Proof Suppose \mathcal{C} is a family of pairwise intersecting maximal cliques of an absolute retract H , and let w_1, w_2, \dots, w_k be the set of all vertices of all the cliques in \mathcal{C} . Any two vertices in the same member of \mathcal{C} are adjacent, and any two vertices in different members of \mathcal{C} have distance at most two, since the members intersect. Consider the vertices w_1, w_2, \dots, w_k with the values $1, 1, \dots, 1$. Since the absolute retract H cannot have a hole, there must exist a vertex $x \in V(H)$ adjacent to all w_i , $i = 1, 2, \dots, k$.

We claim that x lies in each maximal clique $C \in \mathcal{C}$. Otherwise C would not be maximal, since x is adjacent to all of its vertices. \square

It can be shown that these two conditions are also sufficient.

Theorem 2.64 [22] *A graph is an absolute retract if and only if it is dismantlable and clique-Helly.*

We now turn our attention to fixed point properties. An endomorphism of G (a homomorphism of G to itself) *fixes* a set $U \subseteq V(G)$, and the set U is *fixed* by the endomorphism f , if $f(U) = U$. It is easy to see that a nontrivial graph always admits an endomorphism that fixes no vertex. If H has no edge other than the loops, then choose two distinct vertices u and v . Otherwise, choose two adjacent vertices u and v . The mapping that sends u to v and all other vertices to u is an endomorphism of H without a fixed vertex. However, there are graphs H in which every endomorphism fixes an edge. We will characterize such graphs H below. We begin with a more general positive result.

Theorem 2.65 *Any endomorphism of a dismantlable graph fixes some clique.*

Proof We proceed by induction on the number n of vertices of the dismantlable graph H . The result certainly holds for dismantlable graphs with one or two vertices. (Note that a dismantlable graph must be connected.) Thus let H be a dismantlable graph with $n + 1$ vertices and let a dismantling sequence begin with the first fold ϕ of a vertex x to a vertex y (whose neighbourhood contains the neighbourhood of x). Hence ϕ is a retraction of H to $H' = H - x$. Consider an endomorphism f of H , and define a corresponding endomorphism f' of H' by restricting $\phi \circ f$ to $V(H')$. By the induction hypothesis, f' has a fixed clique C . If C is not a fixed subgraph of f , then some $z \in C$ has $f(z) = x$, $f'(z) = y$, while all other $w \in C$ have $f(w) = f'(w) \in C$. We first observe that x is adjacent to all vertices of C . It is adjacent to y by definition, and any other w has a unique v with $f(v) = f'(v) = w$; since $vz \in E(H)$ we must have $f(v)f(z) = wx \in E(H)$. A similar argument establishes that $f(x)$ is adjacent to x (since $xz \in E(G)$ and $f(x)f(z) = f(x)x$), to y (since the neighbourhood of y contains the neighbourhood of x), and all other vertices of C (as before). In fact, the arguments apply to any $f^i(x) = f^{i+1}(z)$, which must be adjacent to all vertices of C , as well as to $f(z), f^2(z), \dots, f^i(z)$. Since the sequence $f^j(z)$ must be periodic (the graph H is finite), some set $f^a(z), f^{a+1}(z), \dots, f^{a+b}(z)$ forms a clique that is fixed by f . \square

Corollary 2.66 *Every endomorphism of a reflexive graph H has a fixed edge if and only if H is a reflexive tree.*

Proof If H is a reflexive tree, then H is an absolute retract and hence is dismantlable, so the result follows from our theorem, since a clique in a tree is either an edge or a vertex. In the first case there is a fixed edge in the second case there is a fixed loop, which we also view as an edge.

On the other hand, if a reflexive graph H is not a tree, then it is disconnected or it contains nontrivial cycles. If H is not connected then we can pick two vertices u, v in different components of H and define an endomorphism of H where all vertices are mapped to u , except the vertices of the component containing u , which are mapped to v . This endomorphism has no fixed edge. If H contains a nontrivial cycle, then let r be a retraction of H to its shortest nontrivial cycle C ensured by Proposition 2.51, and let a be an automorphism of C that does not fix any edge of C (cycles can be rotated). Then $a \circ r$ is an endomorphism of H without a fixed edge. \square

2.12 Median graphs

In this section we again consider only reflexive graphs. A connected graph H is a *median graph* if for every three vertices a, b, c , there exists a unique vertex v that lies simultaneously on a shortest a, b -path, a shortest b, c -path, and a shortest a, c -path. The vertex v is called the *median* of a, b, c .

An odd cycle is not a median graph, since two adjacent vertices a, b , and the vertex c opposite to the edge ab have no median. The same argument applies in

any nonbipartite graph and a shortest nontrivial odd cycle. Thus median graphs are bipartite.

Each tree T is a median graph, since the three vertices a, b, c either lie so that one is on the path between the other two, and then it is the median, or there is a unique vertex x such that a, b, c belong to different components of $T - x$, and then x is the median of a, b, c .

The *hypercube* of dimension k , denoted by Q_k , is the reflexive graph whose vertices are all binary strings of length k , with $a_1a_2 \cdots a_k$ adjacent to $b_1b_2 \cdots b_k$ whenever at least $k - 1$ equalities $a_i = b_i$ hold.

Proposition 2.67 *Each hypercube is a median graph.*

Proof Suppose $a_1a_2 \cdots a_k, b_1b_2 \cdots b_k, c_1c_2 \cdots c_k$ are vertices of Q_k , and let x_i be the majority value of a_i, b_i, c_i . Then $x_1x_2 \cdots x_k$ is on the shortest path between any two of the vertices, and is the unique vertex of Q_k with the property. (If, say, $a_i = b_i$, then any shortest path from $a_1a_2 \cdots a_k$ to $b_1b_2 \cdots b_k$ only contains vertices $y_1y_2 \cdots y_k$ with $y_i = a_i = b_i$.) \square

Proposition 2.68 *A retract of a median graph is a median graph.*

Proof Suppose G is a median graph, and $r : G \rightarrow H$ a retraction. If $a, b, c \in V(H)$ have median x in G , then x is on a shortest path between any two of the vertices in G . Since H is an isometric subgraph of G , the distances amongst a, b, c , are the same in G and in H . It follows that $r(x)$ is on a shortest path between any two vertices from a, b, c , and is the unique vertex with the property. \square

We have proved that each retract of a hypercube is a median graph. It turns out that the converse also holds.

Theorem 2.69 [16] *Median graphs are precisely the retracts of hypercubes.*

\square

In the *dynamic location problem* for a connected graph G one vertex s_0 contains a resource, which may be needed at other vertices of G . When the resource is at vertex s_{i-1} and is requested at vertex r_i , we may service the request at the cost $d_G(s_{i-1}, r_i)$, or move it to an intermediate location s_i , at the cost of $d_G(s_{i-1}, s_i)$, and service the request from there, at the additional cost of $d_G(s_i, r_i)$. The latter option may make future requests less costly to service. If r_1, r_2, \dots, r_n is the sequence of requests, and $s_0, s_1, s_2, \dots, s_n$ the corresponding sequence of locations of the resource (with s_0 being the given initial location), then the total cost is the sum of the moving costs and the servicing costs

$$\sum_{i=1}^n (d_G(s_{i-1}, s_i) + d_G(s_i, r_i)).$$

Letting $r_0 = s_0$, we then have the given sequence $r_0, r_1, r_2, \dots, r_n$ and we are interested in finding the sequence of locations s_1, s_2, \dots, s_n that results in the

minimum total cost [62]. When the entire sequence of requests is known ahead of time, the minimum sequence of locations is easily computed by dynamic programming.

The *window index* of G is the minimum integer ('the time window') k such that some algorithm can find the optimal sequence of locations ('schedule') for every sequence r_0, r_1, \dots, r_n of requests on G , by choosing s_i only from the knowledge of $r_0, r_1, \dots, r_{i+k-1}$, i.e., without knowing r_{i+k}, \dots, r_n .

It is easy to see that a nontrivial graph cannot have window index 1, i.e., that it is impossible to optimally predict where to move the resource only on the basis of knowing the next request. For instance, if a and b are adjacent vertices of G and the resource is initially on a , then the request sequences b, a, a, \dots, a and b, b, b, \dots, b appear the same, yet have different optimal move.

On the other hand, it turns out that a graph has window index two if and only if it is a median graph. In preparation for this result, and an extension to any finite window index, we prove the following simple facts.

Proposition 2.70 *If H is a retract of G then the window index of H does not exceed the window index of G .*

Proof Suppose $f : G \rightarrow H$ is a retraction. If algorithm A solves the dynamic location problem on G with window k , then we can solve the problem on H , also with window k , as follows. Whenever a sequence r_0, r_1, \dots, r_n of requests in H is given, we use algorithm A to find, with time window k , an optimal sequence s_1, s_2, \dots, s_n in the whole graph G . The sequence $f(s_1), f(s_1), \dots, f(s_n)$, now consisting of vertices of H , and computed also with time window k , is an optimal solution in H , since its cost, $\sum_{i=1}^n (d_G(f(s_{i-1}), f(s_i)) + d_G(f(s_i), r_i))$, is bounded by the original cost $\sum_{i=1}^n (d_G(s_{i-1}, s_i) + d_G(s_i, r_i))$. (Recall Proposition 2.49.) \square

The *cartesian product* of graphs G and H has the vertex set $V(G) \square V(H)$ and $(u, v)(u', v')$ is an edge whenever $u = u'$ and $vv' \in E(H)$ or $v = v'$ and $uu' \in E(G)$.

Proposition 2.71 *If both G and H have window index at most k then so does their cartesian product.*

Proof If $(u_0, v_0), (u_1, v_1), \dots, (u_n, v_n)$ is the request sequence in the cartesian product of G and H , then the cost of a schedule $(a_1, b_1), (a_2, b_2), \dots, (a_n, b_n)$ is equal to the sum of the costs of the schedules a_1, a_2, \dots, a_n servicing the requests u_0, u_1, \dots, u_n , and b_1, b_2, \dots, b_n servicing the requests v_0, v_1, \dots, v_n . This is a consequence of the simple fact that the distance between (u, v) and (u', v') in the cartesian product is the sum of the distances between u and u' in G and between v and v' in H . Thus $(a_1, b_1), (a_2, b_2), \dots, (a_n, b_n)$ is optimal if and only if both a_1, a_2, \dots, a_n and b_1, b_2, \dots, b_n are optimal, and if both can be obtained with time window k then so can $(a_1, b_1), (a_2, b_2), \dots, (a_n, b_n)$. \square

Both the definition and the proposition extend to cartesian products of a finite number of graphs, in the obvious way.

Proposition 2.72 *The window index of a complete graph with n vertices is equal to n .*

Proof With window n , we can schedule as follows. If the resource is in vertex s_i and the next n requests are r_{i+1}, \dots, r_{i+n} , then set s_{i+1} equal to r_{i+1} if r_{i+1} is the first vertex repeated in the sequence $s_i, r_{i+1}, \dots, r_{i+n}$; otherwise set s_{i+1} equal to s_i .

We claim that this strategy will produce an optimal schedule. We shall show by induction on j that some optimal schedule agrees with the schedule proposed by our strategy in the first j moves. This is trivially true when $j = 0$, so assume it holds for some $j - 1$, and let a_1, a_2, \dots, a_n be an optimal schedule that agrees with our schedule s_1, s_2, \dots, s_n on the first $j - 1$ moves, i.e., with $a_i = s_i$ for $i = 1, 2, \dots, j - 1$. We now show that some optimal schedule also agrees with our schedule on the first j moves.

If $a_j = s_j$, we are done. If a_j is neither a_{j-1} nor r_j , then it contributes at least two to the cost function, via the sum $d_G(a_{j-1}, a_j) + d_G(a_j, r_j) + d_G(a_j, a_{j+1})$. Hence we can replace a_j by s_j without increasing the cost of the schedule, since s_j is a_{j-1} or r_j , and hence it contributes at most two to the same sum. (Recall that the graph is complete, so all distances are at most one.)

If $a_j = a_{j-1} = s_{j-1}$, then $a_j = s_j$ again, or $s_j = r_j$, which means that r_j was the first repeated vertex in the sequence $s_{j-1}, r_j, r_{j+1}, \dots, r_{j+n-1}$, say $r_j = r_{j+\ell}$. Let $a_{j+1} = a_{j+2} = \dots = a_{j+k} = s_{j-1}$ but $a_{j+k+1} \neq s_{j-1}$ (possibly $k = 0$), and let $k' = \min(k, \ell)$. It is easy to see that replacing all $a_i (= s_{j-1})$ with $i = j, j+1, \dots, j+k'$ by $a_i = r_j$ will not increase the cost. The resulting optimal schedule agrees with our schedule for the first j moves.

Finally, if $a_j = r_j$, then we have $a_j = s_j$ again, or $s_j = a_{j-1}$, whence the first repeated vertex was not r_j . Let $r_{j+\ell}$ be the second occurrence of the first repeated vertex, and let k be the greatest integer such that $a_j = a_{j+1} = \dots = a_{j+k} = r_j$. Then we can again replace all $a_i (= r_j)$ with $i = j, j+1, \dots, j+k'$ by $a_i = s_{j-1}$ without increasing the cost of the schedule, and obtain an optimal schedule that agrees with our schedule on the first j moves. (The definition of k' is exactly as in the previous case.)

It remains to show that we cannot optimally schedule with window $n - 1$. Suppose the resource is in s_0 and the first $n - 1$ requests r_1, r_2, \dots, r_{n-1} are the remaining $n - 1$ vertices of the complete graph. If the n -th request is s_0 , then an optimal strategy requires the resource to stay at s_0 for the first n moves since this strategy only has cost $n - 1$ while all other strategies have cost n . On the other hand, if the n -th request is r_1 , then an optimal strategy must have $s_1 = r_1$ for a similar reason. Therefore, no strategy can yield an optimal schedule. \square

Corollary 2.73 *A graph G has window index two if and only if it is a median graph.*

Proof A median graph is a retract of a hypercube, and each hypercube is a cartesian product of copies of the complete graph on two vertices. Thus a median

graph has window index at most two, and therefore two. We omit the proof of the converse [62]. \square

Similarly, one has the following result.

Theorem 2.74 [62] *A graph G has finite window index if and only if it is the retract of a cartesian product of complete graphs.*

2.13 Remarks

The product we focused on is the most natural product from the homomorphism point of view, because of the properties in Theorem 2.2. Since it is the ‘right’ product for binary relations, it has long been considered from an algebraic perspective, and results such as Theorem 2.2 are folklore of the algebraic (and categorical) perspective. Its investigation in graph theory was pioneered by G. Sabidussi [307]. Many other products are possible, on the vertex set $V(G) \times V(H)$; systematic catalogues of all products satisfying certain natural axioms have been compiled in [185, 294]. The recent book by W. Imrich and S. Klavzar [185] is devoted entirely to the study of graph products; it contains a nice historical account. The dimension of graphs is introduced in [266, 271], and Theorem 2.9 is from [228]. (See also [3, 86].) The Lovász vector and the exponential graph were both introduced in graph theory in [224], although similar exponential objects were studied earlier in algebra and category theory [295]. The application of the exponential graph to the Product Conjecture was pioneered by M. El-Zahar and N. Sauer [89], who used it to prove Theorem 2.32. It was also used in [132] to prove Theorem 2.33, as in Lemma 2.29. The same technique has in fact been previously used by K. Vesztergombi to bound the chromatic number [329]. (Exercise 12, and Exercise 21 in Chapter 6, are also from [329].) The Lovász vector and Theorem 2.11 were used in [224] to prove the cancellation properties Corollary 2.12, Theorem 2.13, and Corollary 2.15. The nonuniqueness of prime factorization was first noted by D. J. Miller [245]. The unique prime factorization theorem for reflexive connected, nonbipartite graphs, Theorem 2.14, is due to R. McKenzie [240]. The ingenious counting of graph homomorphisms proving the edge reconstruction result in Theorem 2.16 is due to V. Müller [250]. Shift graphs are related to shift register sequences [71]; more recently they have been found useful because they are large relative to their diameter and maximum degree, cf. [73]. The construction of $S(g, k)$ in Theorem 2.23 is from [270]. The Product Conjecture was first explicitly stated in [141], and popularized in [55]. It is perhaps the most important conjecture related to graph products and homomorphisms, and there have been many papers dealing with the subject, including the surveys [312, 350], as well as the book [122], cf. for instance [132, 199, 213, 267, 323]. Persistent graphs were introduced in [83], and the Proposition 2.31 is from [328]. Multiplicativity was studied in [132, 267], and our proof of Theorem 2.35 is from [267]. Theorem 2.37 was first proved in [132] using the Lipschitz lemma from algebraic topology; the combinatorial

proof given here is from [344]. Connections between multiplicativity and duality have been explored in [132, 341, 342]. Other results on multiplicativity of oriented paths, cycles, and other digraphs can be found in [75, 174, 324]; in particular, C. Tardif [324] has recently shown that all rational complete graphs $K_{p/q}$ with $p/q < 4$ are multiplicative. Our treatment of Theorem 2.39 is combined from [291, 292]. (Similar results were obtained independently by J. H. Schmerl and X. Zhu.) The projectivity of K_n is obtained by generalizing [126], and appears in [283], cf. also [214, 215, 229, 250]. Projectivity also plays an important role both in Chapter 3 and Chapter 5. (See Exercise 12 in Chapter 3.) Projectivity of *reflexive* graphs has been previously studied in [140].

Retracts in graph theory were first studied in [150, 153, 156], although they were popular in topology and algebra since the time of K. Borsuk [38]. The connection to multiplicativity in Proposition 2.47 is from [213]. Proposition 2.51 is due to G. Sabidussi [310], and was the starting point for the investigations in [150]. The proof given here, as well as that of Corollary 2.66, is adapted from [287]. Theorem 2.52 was first proved in [119] and the simpler proof given in Chapter 6 is from [120]. Theorem 2.53 is also proved in [120]; its proof depends on Tutte's characterization of regular matroids. Absolute retracts in reflexive graphs are similar to absolute retracts in bipartite graphs; direct translations between the characterizing theorems are given in [19]. The results given here, in terms of the variety of paths, Corollary 2.56, in terms of the absence of holes, Theorem 2.55, and in terms of the existence of majority functions, Corollary 2.57, are from [190, 191, 284, 285, 296–298] in the context of reflexive graphs, and their bipartite analogues can be traced back to [18, 150, 153, 154]. Our treatment is based on [173]. The absence of holes is often expressed as the Helly property of disks, as in Exercise 24. Related retract concepts in metric spaces have been studied in [5, 80, 187]. Absolute retracts in graphs, digraphs, and more general structures are considered in [18, 21, 151, 190, 191], and especially in the book by E. Pesch [289]. Absolute retracts can also be defined in terms of sufficiency of conditions other than just isometry; often the resulting class of graphs is still a variety [173, 223]; we explore one example in Exercise 31 in Chapter 5. The graphs $\text{HOM}(G, H)$ are used to model physical systems with ‘hard constraints’, cf. Section 1.8; Proposition 2.61 is from [46]. The connection of dismantlability to the game of cops and robbers, Proposition 2.59, is from [286, 296]. Theorem 2.64, attributed to H. -J. Bandelt and E. Prisner [22], also depends on a deep result on Bandelt and Pesch [20]. Median graphs, segments, and algebras have been the subject of many papers [7, 15, 188, 249, 255, 315]; the application of retracts in the dynamic location problem is from [62], cf. also [289] for earlier connections between the two concepts.

Exercise 3 is from [180], which also contains a nice application of products and homomorphisms in artificial intelligence. Exercise 6 is from [245]. Exercise 8 is based on [132]. Exercise 5 is from [336]. Exercise 10 is from [154]. It is related to subdirect representations of bipartite graphs; subdirect representations of graphs in general are investigated in [310]. Exercise 11 is from [245]. Exercise 21 suggests

studying graphs G which have $G \square G \rightarrow G$; it turns out that they are precisely Cayley graphs of normal groups [216]. Exercise 22 can be strengthened to prove that each graph that does not contain an induced A from Fig. 2.10 is three-colourable [60, 119].

2.14 Exercises

1. Using Theorem 2.2 as a guide, explicitly define the product $S \times S'$ for relational systems S, S' with one ternary relation. Apply your definition to describe the numeric powers N^k of the system N introduced in Section 1.8.
2. Let $G + H$ denote the disjoint union of graphs G, H . Prove that
 - (a) $G \rightarrow G + H$ and $H \rightarrow G + H$.
 - (b) $G \rightarrow X$ and $H \rightarrow X$ imply $G + H \rightarrow X$.
 - (c) For any graphs G, H , there exists a unique (up to isomorphism) graph S and homomorphisms $p : G \rightarrow S, r : H \rightarrow S$ such that for every graph X with homomorphisms $p' : G \rightarrow X, r' : H \rightarrow X$ there is a unique homomorphism $f : S \rightarrow X$ with $f \circ p = p'$ and $f \circ r = r'$.
3. Let G and H be digraphs. Prove that $G \times H$ has a directed cycle if and only if both G and H have a directed cycle.
4. Prove that if $a \leq b$ are both odd integers, then the graph $C_a \times C_b$ contains C_b .
5. Let G and H be graphs. Prove that (u, v) and (u', v') are in the same component of $G \times H$ if and only if there exists a u, u' -walk in G and an v, v' -walk in H whose lengths have the same parity. Deduce that $G \times H$ is connected if and only if both G and H are connected and at least one is nonbipartite.
6. Prove that the infinite product of odd cycles $\prod_n C_{2n+1}$ is bipartite.
7. For graphs G, H define $G \cdot H$ to be the graph with vertex set $V(G) \times V(H)$ and edges $(u, v)(u', v')$ where $uu' \notin E(G)$ or $vv' \in E(H)$. Prove that $G \rightarrow H$ if and only if $G \cdot H$ has a clique with $|V(G)|$ vertices.
8. Prove that
 - (a) G is isomorphic to a subgraph of $H^{(H^G)}$
 - (b) if $G \rightarrow H$ then $K^H \rightarrow K^G$ and $G^K \rightarrow H^K$.
 - (c) $G^L \simeq G$
 - (d) $H \rightarrow H^G$
 - (e) $G \rightarrow H$ if and only if $L \rightarrow H^G$.

(L is the digraph with one vertex and one loop.)
9. Prove that $\dim G \leq 2$ if and only if \overline{G} is the line graph of a bipartite graph.
10. Prove that every bipartite graph is an isometric subgraph of a product of paths.

11. Prove that for connected nontrivial graphs G and H , we have $G \times H \simeq G \square H$ if and only if $G \simeq H \simeq C_{2k+1}$.
12. The *strong product* $G \boxtimes H$ of graphs G, H has vertex set $V(G) \times V(H)$ and edges $(u, v)(u', v')$ where $u = u'$ and $vv' \in E(H)$, or $v = v'$ and $uu' \in E(G)$, or $uu' \in E(G)$ and $vv' \in E(H)$. (As the notation suggests, the strong product is the union of the product and the cartesian product.) Prove that the strong product satisfies the properties of the product from Theorem 2.2, in the context of reflexive graphs.
Prove that for a nonbipartite G we have $\chi(G \boxtimes C_5) = 5$ if and only if $G \rightarrow C_5$.
13. Prove that for graphs G, H , the distance in $G \boxtimes H$ between (u, v) and (u', v') is equal to $\max(d_G(u, u'), d_H(v, v'))$.
14. The *lexicographic product* $G[H]$ of graphs G and H has vertex set $V(G) \times V(H)$ and edges $(u, v)(u', v')$ where $uu' \in E(G)$, or $u = u'$ and $vv' \in E(H)$. Let G, H be connected graphs with cores G', H' , respectively, and let G be triangle-free. Prove that the core of $G[H]$ is $G'[H']$.
15. Prove that for graphs G and H with $\chi(H) = k$, we have $\chi(G[H]) = \chi(G[K_k])$.
16. Suppose H is a subgraph of G . Prove that H is a retract of G if and only if any homomorphism $H \rightarrow X$ can be extended to a homomorphism $G \rightarrow X$.
17. Prove that each graph of girth at least five is projective.
18. Prove that $L(L(K_4))$ is three-colourable. (Hint: Colour the triple abc by b if $b \neq 4$ and by $d \neq a, b, c$ if $b = 4$.)
19. Prove that the class of reflexive dismantlable graphs is a variety.
20. Prove that the reflexive 6-cycle does not have a finite window index.
21. Prove that
 - (a) $K_n \square K_n \rightarrow K_n$ for every n .
 - (b) $\chi(G \square H) = \max(\chi(G), \chi(H))$.
 - (c) Each cycle C_k satisfies $C_k \square C_k \rightarrow C_k$.
 - (d) The Petersen graph P has $P \square P \not\rightarrow P$.
22. Prove that each nearly bipartite graph is three-colourable.
23. Define the reflexive graph G^+ to be obtained from G by subdividing each edge into a path of length two and adding a new vertex adjacent to all vertices. Show that G is triangle-free if and only if G^+ is a median graph.
24. Prove that a reflexive graph H is an absolute retract if and only if its family of disks $(D_k(x) = \{y : d_H(x, y) \leq k\}, k \in N, x \in V(H))$ has the Helly property. (A family of sets has the *Helly property* if each subfamily in which any two members intersect has a nonempty intersection.)
25. Suppose H is a reflexive graph without cycles of length greater than five. Show that H is an absolute retract.
26. Show that a graph H is a retract of a graph G , whenever H is an (isometric path or shortest cycle) in G , and G is bipartite.

27. [203] Prove that $\prod_{i \in I} K_{n(i)} \simeq \prod_{j \in J} K_{m(j)}$ if and only if there is a bijection ϕ of I to J with each $n(i) = m(\phi(j))$.
28. [41] Prove that every homomorphic image of a graph G is isomorphic to a retract of G if and only if G does not contain an induced P_3 or $2K_2$. Prove that the same property in the class of reflexive graphs is equivalent to being a *threshold graph*, i.e., not containing an induced (reflexive) P_3, C_4 , or $2K_2$.
29. [119] Prove that a graph G such that all odd cycles contain one fixed vertex is nearly bipartite.
30. [119] A *cobipartite graph* is a graph whose complement is bipartite. Prove the following max–min relation: if G is nearly bipartite, then the odd girth of G is equal to the maximum number of pairwise edge-disjoint cobipartite subgraphs of G .

THE PARTIAL ORDER OF GRAPHS AND HOMOMORPHISMS

In this chapter we examine the partial order determined by the *existence* of homomorphisms. This point of view motivates many nice questions, and sheds new light on some classical problems. In the next chapter we shall return to the study of the structure of sets of homomorphisms and the existence of homomorphisms will again be the focus of Chapter 5. A *partial order* is a set (not necessarily finite) with a reflexive, transitive, and anti-symmetric relation, usually denoted by \leq . (We do not require a partial order to be a finite set, because we shall consider the partial order of all graphs, which is infinite; however, for the same reason, in practice we will only consider countable partial orders.)

3.1 The partial orders \mathcal{C} and \mathcal{C}_S

Let \mathcal{G} denote the set of all digraphs, and write $G \leq H$ for $G \rightarrow H$ (with $G, H \in \mathcal{G}$). Since ‘homomorphisms compose’ (Section 1.7), \leq is a transitive relation on \mathcal{G} . It is also clear that \leq is a reflexive relation (each identity map is a homomorphism). However \leq is in general not antisymmetric; in fact, $G \leq H$ and $H \leq G$ just means that G and H are homomorphically equivalent. A binary relation that is reflexive and transitive, but not necessarily antisymmetric is called a *quasiorder*; thus \leq defines a quasiorder on \mathcal{G} .

There are standard ways to transform a quasiorder into a partial order—by identifying equivalent objects, or by choosing a particular representative for each equivalence class. In our context, we will do the latter, choosing the representatives to be the cores. Recall from Proposition 1.31 that a core is a graph that is not homomorphic to a proper subgraph of itself. We have shown in Corollary 1.32 that every graph is homomorphically equivalent to a unique core (up to isomorphism). Therefore, if we denote by \mathcal{C} the set of all nonisomorphic cores, we cannot have two distinct elements G, H of \mathcal{C} satisfy both $G \leq H$ and $H \leq G$.

Theorem 3.1 *The set \mathcal{C} is a partial order under \leq .* □

It is interesting to interpret statements such as Corollary 1.8, Corollary 1.23, and Exercise 2 in Chapter 1, in terms of this partial order. They claim the monotonicity of the corresponding graph parameters. For instance, Corollary 1.8 claims that the chromatic number is a monotone parameter in the sense that $G \leq H$ implies $\chi(G) \leq \chi(H)$. Corollary 1.23 claims the same kind of monotonicity for the independence ratio, in the special case when the graph H is vertex-transitive.

We note the following property of \mathcal{C} .

Proposition 3.2 *The partial order \mathcal{C} is a lattice.*

Proof Any two elements G, H of \mathcal{C} have a greatest lower bound in \mathcal{C} —namely the core of the product $G \times H$. Indeed, Proposition 2.1 implies that $G \times H \leq G$ and $G \times H \leq H$, as well as that $X \leq G, X \leq H$ yields $X \leq G \times H$. Note that $G \times H$ may not be a core, but the core of $G \times H$ satisfies exactly the same set of inequalities in \mathcal{C} . Similarly, using Exercise 2 from Chapter 2, we prove that the core of $G + H$ is the least upper bound of G and H in \mathcal{C} . \square

The fact that $G \times H$ is (homomorphically equivalent to) the greatest lower bound of G and H in \mathcal{C} often has interesting implications. For instance in certain models in learning theory [180], a clause C corresponds to a digraph G_C in such a way that a clause C generalizes a clause C' if and only if $G_C \rightarrow G_{C'}$. Suppose we are given a set of (background knowledge) clauses describing a certain, otherwise unknown predicate. These clauses are instances, both positive (examples satisfying the predicate) and negative (examples not satisfying it). We may want to predict, based on the background knowledge, which other clauses will be satisfied by it. In such a case, one may want to compute the least general common generalization of the given background clauses. The proposition implies that the least general common generalization of clauses $C_i, i \in I$, corresponds to the digraph $\prod_{i \in I} G_{C_i}$.

The partial order \mathcal{C} is unusual in the sense that it is in general difficult to decide if $G \leq H$. For instance, just deciding if a given graph satisfies $G \leq K_3$ is the well-known NP -complete problem of graph three-colourability. (More on this topic in Chapter 5). In fact, \mathcal{C} is a particularly rich and complex partial order, as we shall discover in this chapter. For instance, we have the following theorem.

Theorem 3.3 [295] *Any countable partial order is isomorphic to a suborder of \mathcal{C} .*

We will paraphrase this kind of result by saying that any countable partial order admits a *representation* in \mathcal{C} .

Recall that we view graphs as a subclass of digraphs. It is clear from the proof of Proposition 3.2 that graphs form a sublattice (and a suborder) of \mathcal{C} . We will denote this suborder by \mathcal{C}_S . Theorem 3.3 can be strengthened to say that *any countable partial order is isomorphic to a suborder of \mathcal{C}_S* (cf. also Proposition 3.8).

We will not be proving (either version of) this deep result, but in the following sections we will derive several of its consequences.

3.2 Representing ordered sets

In this section we shall consider some special cases of Theorem 3.3. As a warm-up we consider the following two extreme cases. A *chain* in a partial order is a set of pairwise comparable elements. The countable chain C_ω is the usual numeric order of positive integers. An *antichain* in a partial order is a set of

pairwise incomparable elements. The countable antichain A_ω is the equality order of positive integers (in which any $i \neq j$ are incomparable).

Proposition 3.4 *The countable chain C_ω is isomorphic to a suborder of \mathcal{C}_S . The countable antichain A_ω is isomorphic to a suborder of \mathcal{C}_S .*

Proof For the first assertion, we can use the complete graphs $K_i, i \in N$. We know that $K_i \leq K_j$ if and only if $i \leq j$. For the second assertion, we can use the graphs $S(i, i), i$ odd, from Theorem 2.23. Because of Corollary 1.8 and Exercise 2 in Chapter 1, we know that $S(i, i)$ and $S(j, j)$ with $i \neq j$ are incomparable. \square

In the case of *finite* ordered sets, this is sufficient to obtain the desired result.

Theorem 3.5 *Every finite partial order is isomorphic to a suborder of \mathcal{C}_S .*

Proof It is well known that each partial order P is isomorphic to the inclusion order of a suitable collection of sets. One can, for instance, represent each element $p \in P$ by the set $L_p = \{x : x \leq p \text{ in } P\}$. It is easy to check that $p \leq q$ in P if and only if $L_p \subseteq L_q$. Thus we may assume that P is the inclusion partial order of sets $M \in \mathcal{M}$. Let X be the union of all sets M in \mathcal{M} . Without loss of generality, we may assume that X is a set of positive integers. Let $S(i, i), i \in X$, be the graphs used above in representing the antichain A_ω . We then assign, to each set $M \in \mathcal{M}$, the graph G_M which is the disjoint union of $S(i, i), i \in M$. It is easy to see that there is a homomorphism G_M to $G_{M'}$ if and only if $M \subseteq M'$. \square

Of course, there are many different ways to represent a fixed ordered set in \mathcal{C}_S or in \mathcal{C} . For instance, the chain C_ω can also be represented by directed paths $P_i, i \in N$, and the antichain A_ω by the collection of directed cycles of prime lengths (Exercise 4 in Chapter 1).

The following family of oriented paths is often useful. The *zig-zag of order k* , denoted by Z_k , is the oriented path of height three given in Fig. 3.1. We consider a to be the *initial vertex* of Z_k , and d to be the *terminal vertex* of Z_k . (In other words, these oriented paths are understood to have a beginning and an end.)




FIG. 3.1. The zig-zag Z_k (for $k \geq 1$).

Proposition 3.6 $Z_k \leq Z_\ell$ if and only if $k \geq \ell$.

Proof Use Proposition 1.14. □

This means that associating to each Z_k the negative integer $(-k)$ is an isomorphism between the suborder of \mathcal{C} induced by the zig-zags, and the natural order of negative integers. The *concatenation* of oriented paths P and Q , denoted $P \cdot Q$, is the oriented path obtained from P and Q by identifying the terminal vertex of P with the initial vertex of Q , leaving the initial vertex of P as the initial vertex of the concatenation, and the terminal vertex of Q as the terminal vertex of the concatenation. (In Fig. 3.2 we show the concatenation $Z_2 \cdot Z_2$). Let $Z(i)$ denote the i -fold concatenation of Z_i , i.e., the digraph $Z_i \cdot Z_i \cdots Z_i$.

Proposition 3.7 The digraphs $Z(i)$, $i \in \mathbb{N}$, are pairwise incomparable in \mathcal{C} .

Proof This easily follows from Proposition 1.14 and Proposition 3.6. □




FIG. 3.2. The concatenation $Z(2) = Z_2 \cdot Z_2$.

We now prove a variant of Theorem 3.5, namely that *every finite partial order is isomorphic to a suborder of \mathcal{C}* . It is well known that each partial order P is the intersection of chains. Hence we may assume that the elements of P are k -tuples (x_1, x_2, \dots, x_k) , where each x_i is an integer between 0 and $n-1$, with the order in P being the coordinate-wise order of the k -tuples, i.e., $(x_1, x_2, \dots, x_k) \leq (y_1, y_2, \dots, y_k)$ in P if and only if each $x_i \leq y_i$ as integers.

We assign, to each element $x = (x_1, x_2, \dots, x_k)$ of P the digraph $G_x = P_n \cdot Z_{n-x_1} \cdot Z_{n-x_2} \cdots Z_{n-x_k} \cdot P_n$, where P_n is the directed path with n vertices. It is again easy to check, from Propositions 1.14 and 3.6, that $x \leq y$ in P if and only if $G_x \rightarrow G_y$. □

This construction is interesting in that the representing family consists entirely of oriented paths. Let \mathcal{C}_P denote the suborder of \mathcal{C} consisting of all oriented paths.

Proposition 3.8 Every finite partial order is isomorphic to a suborder of \mathcal{C}_P . □

This result has been recently extended to all countable partial orders in [182]. Representability by oriented paths is also discussed at the beginning of Section 4.8.

Let us write $G < H$ to mean $G \leq H$ and $H \not\leq G$. The order \mathcal{C}_S is dense in the usual sense, that between any two graphs there exists another graph. This is discussed in Section 3.7, and there is an obvious exception, noted there. For the time being, we only note that between the graphs K_2 and K_3 (which satisfy $K_2 < K_3$), there is an infinite chain of graphs

$$K_2 < \cdots < C_7 < C_5 < C_3 = K_3.$$

This suggests that the graphs $C_{2\ell+1}$ act as fractional complete graphs, say $C_{2\ell+1} = K_{2+(1/\ell)}$, resulting in

$$K_2 < \cdots < C_7 = K_{2+\frac{1}{3}} < C_5 = K_{2+\frac{1}{2}} < C_3 = K_3.$$

In Section 6.1, we shall formally extend the definition of complete graphs K_n , with integer subscripts n , to cores K_r , where r is any rational number $r \geq 2$, such that $K_r \rightarrow K_{r'}$ if and only if $r \leq r'$ (Theorem 6.3). In other words, if we denote by \mathcal{Q} the set of positive rational numbers under the usual numeric \leq , we have the following representation result.

Proposition 3.9 *\mathcal{Q} is isomorphic to a suborder of \mathcal{C}_S .* □

3.3 Incomparable graphs and maximal antichains in \mathcal{C}_S

We have already seen infinite antichains in \mathcal{C} , i.e., infinite sets of incomparable graphs. However, not every antichain in \mathcal{C} can be extended to an infinite antichain. In this section we focus on finite antichains that are *maximal* with respect to inclusion.

We begin by considering graphs and the order \mathcal{C}_S . The graphs K_1 and K_2 play a special role in \mathcal{C}_S . They are the only possible cores of bipartite graphs. The graph K_1 is homomorphically equivalent to any edgeless graph (graph with no edges). The graph K_2 is homomorphically equivalent to any bipartite graph that has edges. Moreover, both $\{K_1\}$ and $\{K_2\}$ are maximal antichains in \mathcal{C}_S —since $K_1 \leq H$ for any H in \mathcal{C}_S , and $K_2 \leq H$ or $H \leq K_2$ for any H in \mathcal{C}_S (the former case for all graphs H with an edge, the latter case for all edgeless graphs H).

We note that while, say, K_2 cannot be extended by an incomparable *graph* in \mathcal{C}_S , the directed three-cycle \vec{C}_3 is incomparable with it in \mathcal{C} . Thus $\{K_2\}$ is *not* a maximal antichain in \mathcal{C} .

We now proceed to prove that $\{K_1\}$ and $\{K_2\}$ are the *only* finite maximal antichains in \mathcal{C}_S . The essence of the argument is the following extension property of graphs.

Theorem 3.10 *For any set G_1, G_2, \dots, G_t of nonbipartite graphs, there exists a graph G_{t+1} incomparable with all G_i for $i = 1, 2, \dots, t$.*

Proof Let g be an odd integer greater than the odd girth of all the graphs $G_i, i = 1, 2, \dots, t$, and let k be an integer greater than the chromatic number of all the graphs $G_i, i = 1, 2, \dots, t$. Then the graph $G_{t+1} = S(g, k)$ from Theorem

2.23 satisfies the claim—it is not homomorphic to any $G_i, i = 1, 2, \dots, t$ because of the chromatic number, and no G_i with $i = 1, 2, \dots, t$ is homomorphic to it because of the odd girth. \square

Corollary 3.11 *The order \mathcal{C}_S has exactly two finite maximal antichains, namely $\{K_1\}$ and $\{K_2\}$.*

Proof We have established above that cores other than K_1 and K_2 are not bipartite; thus Theorem 3.10 shows that any other finite antichain can be extended. \square

Theorem 3.10 implies that any set of $t \geq 2$ incomparable graphs can be extended to a set of $t + 1$ incomparable graphs. This situation in \mathcal{C}_S is in sharp contrast with \mathcal{C} . Indeed, we shall show in Section 3.9 that there exist in \mathcal{C} maximal antichains of arbitrary finite size. (An extension to countably infinite graphs can be found in [276].)

We shall now proceed to strengthen this extension property in several ways. The techniques used for the proofs are useful in other situations; they involve random graphs. There is a deep connection with classical results typical of the probabilistic method, such as Theorem 1.9, derived below as Corollary 3.13.

The following general theorem is traditionally called the *Sparse Incomparability Lemma*.

Recall that we write $G < H$ for $G \leq H$ and $H \not\leq G$. The Sparse Incomparability Lemma asserts that when $G < H$ there is a sparse (meaning high-girth) graph G' incomparable with G which still has $G' < H$. The precise formulation is as follows.

Theorem 3.12 *Let ℓ be a positive integer. For any nonbipartite graphs G, H with $G < H$, there exists a graph G' such that*

- G' is incomparable with G ,
- $G' < H$, and
- G' has girth at least ℓ .

Before proving the Sparse Incomparability Lemma, we illustrate some of its many applications. To begin with, it proves Theorem 1.9.

Corollary 3.13 *For any positive integers k, ℓ there exists a graph of chromatic number k , and girth at least ℓ .*

Proof For $k \leq 3$ this is obvious. Otherwise we can apply the Sparse Incomparability Lemma, Theorem 3.12, with $G = K_{k-1}$ and $H = K_k$. \square

We also derive the following fact.

Corollary 3.14 *For every nonbipartite graph G and every integer ℓ there exists a graph of girth at least ℓ , incomparable with G .*

Proof Put $H = K_n$, where $n = |V(G)|$. \square

The last corollary implies that *for every nonbipartite graph G there exists a sparse (high-girth) graph G' that is not homomorphic to G* . This is perhaps the most natural extension of Erdős's result on the existence of sparse graphs with high chromatic number (Corollary 3.13).

Proof of Theorem 3.12 The existence of G' will follow from the fact that a certain random graph almost certainly has a strengthened version of the required properties.

Assume that G has t vertices, and H has s vertices, say $1, 2, \dots, s$. Note that both s and t are fixed. Consider a large positive integer n .

Consider the random graph $\mathbf{G}(s, n, p)$ defined as follows.

The vertex set V of $\mathbf{G}(s, n, p)$ will be the union of s disjoint sets V_1, V_2, \dots, V_s , each of size n . The sets V_i are in one-to-one correspondence with the vertices of H ; we will say that two sets V_i, V_j are *adjacent* if the corresponding vertices i, j are adjacent in H . The graph $\mathbf{G}(s, n, p)$ will have no edges within any of the sets V_i , or between any two nonadjacent sets V_i, V_j , and each edge between two adjacent sets V_i, V_j will be chosen to be present independently, at random, with probability $p = n^{\delta-1}$, where $0 < \delta < 1/\ell$.

A set $A \subset V$ will be called *large* if it has at least n/t elements in each of a pair of adjacent sets V_i, V_j , i.e., if there is an edge ij of H such that $|A \cap V_i| \geq n/t$ and $|A \cap V_j| \geq n/t$. Any edge ij of H certifying in this way that A is large will be called a *good edge* for A .

For a large set A , let $m(A)$ be a random variable denoting the smallest number of edges $\mathbf{G}(s, n, p)$ has between any $A \cap V_i$ and $A \cap V_j$, over all good edges ij for A .

We first estimate the probability

$$\alpha = \Pr(A \text{ large implies } m(A) \geq n).$$

We have

$$1 - \alpha \leq \sum_{A \text{ large}} \Pr(m(A) < n).$$

We shall estimate this sum as follows. The number of large sets is at most 2^{sn} (the total number of subsets of V). The fact that $m(A) < n$ means that for some good edge ij , out of the possible ab edges between the $a \geq n/t$ vertices of $A \cap V_i$ and the $b \geq n/t$ vertices of $A \cap V_j$, at most $n - 1$ are present, i.e., all the $ab - (n - 1) \geq (n^2/t^2) - n$ edges are absent. Summing over all good edges we estimate the probability that a particular large A has $m(A) < n$ to be at most

$$\binom{sn}{n} \cdot (1 - p)^{(n^2/t^2) - n}.$$

Indeed, the number of ways the at most $n - 1$ edges are chosen in the various good pairs is smaller than all possible ways of choosing $n - 1$ (and hence smaller than n) edges on V . Now we roughly estimate, for large n ,

$$\binom{\binom{sn}{2}}{n} \leq \binom{s^2 n^2}{n} \leq s^{2n} n^{2n} < e^{cn \log_2 n}$$

and

$$(1-p)^{(n^2/t^2)-n} \leq e^{-p((n^2/t^2)-n)},$$

using the fact that $\ln(1+x) \leq x$.

Putting all this together, we see that

$$1-\alpha \leq \sum_{A \text{ large}} \Pr(m(A) < n) \leq 2^{sn} \binom{\binom{sn}{2}}{n} \cdot (1-p)^{(n^2/t^2)-n} < e^{c'n \log_2 n - c''n^{1+\delta}},$$

for some positive constants c' and c'' which are independent of n .

We conclude that $\Pr(A \text{ large implies } m(A) \geq n) = 1-o(1)$. Thus a large set A asymptotically almost always has at least n edges between the sets $A \cap V_i, A \cap V_j$, for all good edges ij for A . At the same time, we shall show that $\mathbf{G}(s, n, p)$ does not have too many cycles of length less than ℓ . Let q be the random variable denoting the number of such cycles. Then the expected value of q is at most

$$3! \binom{sn}{3} p^3 + 4! \binom{sn}{4} p^4 + \cdots + (\ell-1)! \binom{sn}{\ell-1} p^{\ell-1} \leq \ell \cdot \frac{s^\ell n^\ell n^{\delta\ell}}{n^\ell} < C n^{\delta\ell} = o(n),$$

(for a constant C), according to our choice of δ .

It now follows that there exists a graph G'' on V , with no edges within the sets V_i , in which every large set A has $m(A) \geq n$, and which has at most $n-1$ cycles of length less than ℓ . Let G' be the graph obtained from G'' by deleting an edge in each cycle of length less than ℓ . We claim that G' has girth at least ℓ and is not homomorphic to G . The first claim follows from the definition of G' . Suppose that $f: G' \rightarrow G$ is a homomorphism. Then each V_i contains a set W_i of at least n/t vertices v with the same image $f(v)$ in G . Let $j(i)$ be the image in G of all the vertices of W_i . The union A of all these sets W_i is a large set for which all edges ij of H are good. Hence A contains at least n edges of G'' , and therefore at least one edge of G' , between any adjacent sets V_i, V_j . This allows us to define a homomorphism $g: H \rightarrow G$ by letting $g(i) = j(i)$, contrary to our assumption that $G < H$. We have still not shown that G and H are not homomorphic to G' , but this is easily ensured by choosing ℓ greater than the girth of both G and H . \square

Let us say that graphs G, G' are k -equivalent (for a positive integer k) provided $G \leq K$ if and only if $G' \leq K$ for all graphs K with at most k vertices. It is easy to see that this is an equivalence relation on \mathcal{C} . We now observe that the above proof actually implies a stronger statement, which can be nicely formulated in terms of k -equivalent graphs. Note, in particular, that the new statement no longer requires that the graphs be restricted to be nonbipartite.

Theorem 3.15 *Let k, ℓ be any integers. Every graph G has a k -equivalent graph G' of girth at least ℓ .*

Proof Let G be a graph. In the above probabilistic proof of the existence of G' , let G play the role of the graph H , and let k play the role of t . (The graph called G in that proof will not figure in this argument.) Then the graph G' obtained satisfies $G' < G$; it follows from the last part of the proof that for any K with at most k vertices we have $G \leq K$ if and only if $G' \leq K$. \square

Note that this is a generalization of the Sparse Incomparability Lemma. Suppose that G, H are nonbipartite graphs with $G < H$. Let k be the maximum of $|V(G)|, |V(H)|$, and assume, without loss of generality, that $\ell \geq k$. We now apply the theorem to k, ℓ , and the graph H . There exists a graph G' of girth at least ℓ , which is k -equivalent to H . It follows that $G' \leq H$ and $G' \not\leq G$. On the other hand, the girth of G' ensures that $H \not\leq G', G \not\leq G'$, since G, H are not bipartite. Thus G' is incomparable to G and satisfies $G' < H$.

3.4 Sparse graphs with specified homomorphisms

When G and G' are k -equivalent, we may seek a bijection between the sets of homomorphisms $\text{HOM}(G, H)$ and $\text{HOM}(G', H)$, via the composition with a fixed homomorphism of G' to G . It turns out that such a bijection may not exist, as we explain below. However, if the graph H has a special property, such a bijection can always be found. We say that a graph H is G -pointed, if there do not exist two homomorphisms $f_1, f_2 : G \rightarrow H$ such that $f_1(x) \neq f_2(x)$ holds for exactly one vertex x of G . (Hence H is G -pointed if and only if the reflexive graph $\text{HOM}(G, H)$ has no edges other than loops.)

We shall describe several applications of the following result, which we offer without proof.

Theorem 3.16 [283] *Let k, ℓ be integers. For every graph G there exists a k -equivalent graph G' of girth at least ℓ , and a surjective homomorphism $c : G' \rightarrow G$, such that for any G -pointed graph H with at most k vertices, and any homomorphism $g : G' \rightarrow H$, there exists a unique homomorphism $f : G \rightarrow H$ with $g = f \circ c$ (Fig. 3.3).*




FIG. 3.3. A commuting diagram.

We note that the last statement may fail to hold without some kind of assumption on H . Suppose, for instance, that there exist two homomorphisms

$f_1, f_2 : G \rightarrow H$ with $f_1(x) \neq f_2(x)$, which agree on all other vertices of G , and suppose that there exist two distinct vertices y, z of G' with $c(y) = c(z) = x$. We define $g : G' \rightarrow H$ by setting $g(v) = f_1(c(v))$ for all $v \neq z$ and $g(z) = f_2(c(z))$. Since G is irreflexive, there is no edge between y and z , and hence g is indeed a homomorphism. However, g cannot be written as $g = f \circ c$ for any homomorphism $f : G \rightarrow H$, because $g(y) \neq g(z)$ while $c(y) = c(z)$.

We will not prove Theorem 3.16; however a version in which girth is replaced by odd girth is proved below as Theorem 3.20.

Our first application of Theorem 3.16 is to uniquely H -colourable graphs. We say a graph G is *uniquely H -colourable* if there is a surjective homomorphism $c : G \rightarrow H$, such that every homomorphism of G to H is the composition $\sigma \circ c$ of c and some automorphism σ of H . It is easy to see that if H is not a core then no graph can be uniquely H -colourable. Moreover, when H is a core, then H itself is H -pointed, as any homomorphism of H to itself is an automorphism of H , and two automorphisms f_1, f_2 which differ on x also differ on $f_1^{-1}(f_2(x))$.

Corollary 3.17 *Let ℓ be a positive integer. For every core H there exists a uniquely H -colourable graph G of girth at least ℓ .*

Proof We apply Theorem 3.16 to the integers ℓ and $k = |V(H)|$, and the graph H . The theorem ensures that there exists a graph G' of girth at least ℓ , such that for every homomorphism $g : G' \rightarrow H$ there is a unique homomorphism $f : H \rightarrow H$ with $g = f \circ c$. Since H is a core, f is an automorphism of H . Thus G' is uniquely H -colourable. \square

We note that this corollary implies that there exist uniquely n -colourable graphs of arbitrarily high girth—compare with Corollary 2.25.

Our second application of Theorem 3.16 is a different strengthening of the same fact.

Corollary 3.18 *Let k, ℓ, t be positive integers, $k > 2$.*

Let $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_t$ be distinct partitions of a finite set A , each with k (possibly empty) parts. Then there exists a graph G of girth at least ℓ and chromatic number k , such that

- A is a subset of $V(G)$,
- G has precisely t k -colourings c_1, c_2, \dots, c_t , and
- the partition associated with the k -colouring c_i , restricted to the set A , is precisely \mathcal{A}_i .

In Fig. 3.4 we illustrate the statement of the lemma. Two partitions of A into k parts are shown in (a), one indicated by solid lines, the other by dashed lines. In (b), A is shown included in the vertex set $V(G)$ of some (sparse) graph G , and both partitions are extended to the whole of $V(G)$. The corollary states that each of these partitions of $V(G)$ is a k -colouring of G , and that there are no other k -colourings of G .

We shall actually derive the corollary in the following, more general, form. The fact that it implies Corollary 3.18 is due to Corollary 2.45, which depends




FIG. 3.4. Extending partitions.

on the fact that complete graphs are projective. We also replaced each partition \mathcal{A}_i by the associated mapping f_i .

Corollary 3.19 *Let H be a projective core with k vertices, ℓ a positive integer. Let A be a finite set and let f_1, f_2, \dots, f_t be distinct mappings $A \rightarrow V(H)$. Then there exists a graph G of girth at least ℓ such that*

- A is a subset of $V(G)$,
- for every $i = 1, 2, \dots, t$ there exists a unique homomorphism $g_i : G \rightarrow H$ such that g_i restricted to the set A coincides with the mapping f_i , and
- for every homomorphism $f : G \rightarrow H$ there exists an $i, 1 \leq i \leq t$, and an automorphism h of H , such that $h \circ f_i = f$.

Proof Assume that A has n elements. Put $N = 1 + \max(n, k^{kt})$. Let f_0 be an injective mapping of A into $\{1, 2, \dots, N\} = V(K_N)$. Let K be the product $K_N \times H \times H \times \dots \times H$, with t copies of H . We can view A as a subset of $V(K)$, by identifying $a \in A$ with $\phi(a) = (f_0(a), f_1(a), f_2(a), \dots, f_t(a)) \in V(K)$. (Since f_0 is injective, ϕ is an injective mapping of A to $\phi(A)$.) We now apply Theorem 3.16 to K in place of G . The theorem yields a graph G' of girth at least ℓ and a surjective homomorphism $c : G' \rightarrow K$ such that any homomorphism $g : G' \rightarrow H$ is the composition of c and some homomorphism $f : K \rightarrow H$. Each homomorphism $g : K \rightarrow H$ induces a homomorphism of $H \times H \times \dots \times H$ to H . Indeed, define, for each $v \in V(K_N)$, a mapping f_v of $H \times H \times \dots \times H$ to H , where $f_v(x_1, \dots, x_t) = g(v, x_1, \dots, x_t)$. Since $N > k^{kt}$, some $f_v = f_w$. It is then easy

to check that $f = f_v = f_w$ is a homomorphism. (The following proof contains a similar statement verified in detail). As H is projective and a core, that the only homomorphisms of K to H are projections followed by an automorphism of H (Exercise 1). We can now view A as a subset of $V(G')$ by identifying each $\phi(a)$ with one arbitrary element of $c^{-1}(\phi(a))$. It follows that all homomorphisms of G' to H extend to the mappings f_i , and hence extend the partitions \mathcal{A}_i . We let G be the graph G' . \square

Since all these applications use the unproved Theorem 3.16, we now give a constructive proof of the corresponding result where girth is replaced by odd girth. (Exercise 2 in Chapter 1 argues that the odd girth is a more relevant parameter, from the homomorphism perspective.)

Theorem 3.20 *Let k, ℓ be integers. For every graph G there exists a k -equivalent graph G' of odd girth at least ℓ , and a surjective homomorphism $c : G' \rightarrow G$, such that for any G -pointed graph H with at most k vertices, and any homomorphism $g : G' \rightarrow H$, there exists a unique homomorphism $f : G \rightarrow H$ with $g = f \circ c$.*

In fact, we can choose $G' = G \times G_0$, where G_0 is any connected graph with odd girth at least ℓ and chromatic number greater than $k^{|V(G)|}$, and let c be the projection onto G (Fig. 3.5).




FIG. 3.5. Another commuting diagram.

Proof A graph G_0 satisfying the required properties exists by Theorem 2.23. Clearly, the odd girth of $G \times G_0$ is also at least ℓ .

We first show that $G \times G_0$ is k -equivalent to G . Let H be a graph with at most k vertices. If $G \rightarrow H$, then $G \times G_0 \rightarrow H$ by composition with the second projection. Thus let $g : G \times G_0 \rightarrow H$ be a homomorphism; we show that $G \rightarrow H$, by the same device used in the previous proof. Specifically, for every $y \in V(G_0)$ define the mapping $f_y : V(G) \rightarrow V(H)$ by $f_y(x) = g(x, y)$. As $\chi(G_0) > k^{|V(G)|}$ there must exist an edge yy' in G_0 with $f_y = f_{y'}$. In this case, the mapping f_y is a homomorphism of G to H : for an edge xx' of G we have the edge $(x, y)(x', y')$ in $G \times G_0$ and hence $g(x, y)g(x', y') = f_y(x)f_{y'}(x') = f_y(x)f_y(x')$ is an edge of H .

Suppose next that H is a G -pointed graph with at most k vertices, and $g : G \times G_0 \rightarrow H$ is a homomorphism.

Claim Suppose y is a vertex of G_0 such that the above mapping f_y is a homomorphism of G to H , and assume y' is a vertex of G_0 adjacent to y . Then $f_{y'} = f_y$.

Otherwise $f_{y'}(x_0) \neq f_y(x_0)$ for some vertex $x_0 \in V(G)$. Define a mapping $\phi : V(G) \rightarrow V(H)$ as follows: let $\phi(x) = f_y(x)$ for $x \neq x_0$, and let $\phi(x_0) = f_{y'}(x_0)$. Then ϕ is a homomorphism $G \rightarrow H$. It suffices to check that the edges of G incident with x_0 are preserved. If $xx_0 \in E(G)$ then $(x, y)(x_0, y') \in E(G \times G_0)$, and therefore $f_y(x)f_{y'}(x_0) = \phi(x)\phi(x_0) \in E(H)$. However, ϕ differs from f_y on exactly one vertex, contradicting the fact that H is G -pointed.

Since G_0 is connected, this means that $g(x, y) = f_y(x)$, i.e., $g = f_y \circ c$. \square

3.5 Incomparable graphs with additional properties

Here we give another proof of the fact that, for nonbipartite graphs, any finite set of pairwise incomparable graphs can be extended by another graph incomparable to all of them. This proof applies to many restricted classes of graphs.

We define an *ideal class* to be a class of graphs \mathcal{K} satisfying the following conditions:

- each induced subgraph of a graph $K \in \mathcal{K}$ is in \mathcal{K} ,
- for each graph $K \in \mathcal{K}$ and any graph G , the product $K \times G$ is in \mathcal{K} , and
- for any two graphs $K, K' \in \mathcal{K}$, the disjoint union $K + K'$ is in \mathcal{K} .

Two typical ideal classes of graphs are:

- the set of all graphs G homomorphic to a fixed H , and
- the set of all graphs G to which a fixed connected graph F is *not* homomorphic.

It is clear from the definition that the intersection of a family of ideal classes is an ideal class.

Note that if a core F is not connected, then it can be written as $F = A + B$, and the set of graphs to which F is not homomorphic is not closed under disjoint unions. (F is homomorphic to neither A nor B , but it is homomorphic to their union.)

Let \mathcal{F} be a family of connected graphs, and let $\text{Forb}\mathcal{F}$ denote the set of graphs to which no $F \in \mathcal{F}$ is homomorphic. Then $\text{Forb}\mathcal{F}$ is an intersection of ideal classes, and hence an ideal class. The converse of this statement also holds.

Proposition 3.21 \mathcal{K} is an ideal class if and only if $\mathcal{K} = \text{Forb}\mathcal{F}$ for some family \mathcal{F} of connected graphs.

Proof Indeed, suppose \mathcal{K} is an ideal class, and let \mathcal{F} consist of all those connected graphs that are not induced subgraphs of any $K \in \mathcal{K}$. We claim that $\mathcal{K} = \text{Forb}\mathcal{F}$. We will first show that \mathcal{F} is closed under taking homomorphisms. Assume that $f : F \rightarrow F'$ is a homomorphism, where $F \in \mathcal{F}$, $F' \notin \mathcal{F}$. This means that F' is an induced subgraph of some $K \in \mathcal{K}$, and therefore $F' \in \mathcal{K}$. Therefore, $F' \times F \in \mathcal{K}$ as well, and since $F \rightarrow F'$, it is easy to see that F is an induced subgraph of $F' \times F$. This would mean that $F \in \mathcal{K}$ and thus not in \mathcal{F} , a

contradiction. It now follows that $\text{Forb}\mathcal{F}$ is the class of graphs that do not have any graph from \mathcal{F} as an induced subgraph, i.e., \mathcal{K} . \square

Theorem 3.22 *Let \mathcal{K} be an ideal class.*

1. *For any nonbipartite $A \in \mathcal{K}$ which is not the maximum element of \mathcal{K} there exists a $B \in \mathcal{K}$ incomparable with A .*
2. *For any family of (pairwise) incomparable graphs $A_1, A_2, \dots, A_t, (t > 1)$ in \mathcal{K} , there exists a graph $A_{t+1} \in \mathcal{K}$, incomparable with all $A_i, i = 1, 2, \dots, t$.*

Proof To prove 1, assume that a nonbipartite A is not the maximum element of \mathcal{K} ; thus there exists some $A' \in \mathcal{K}$ not homomorphic to A . Let G be a graph with odd girth greater than the odd girth of A , and chromatic number greater than n^n , where n is the number of vertices of A . We claim that $B = A' \times G \in \mathcal{K}$ is incomparable with A . Indeed, B has odd girth greater than the odd girth of A , hence $A \not\rightarrow B$; on the other hand, $B \not\rightarrow A$ by the same argument as in the proof of Theorem 3.20.

To prove 2, we proceed analogously. Assume $A_1, A_2, \dots, A_t, t > 1$, are incomparable (and hence nonbipartite) graphs in \mathcal{K} . There exists an $A' \in \mathcal{K}$ to which all A_i are homomorphic, since \mathcal{K} is an ideal class; on the other hand, this A' cannot be homomorphic to any A_i , since $t > 1$ and the graphs A_i are pairwise incomparable. Now the proof proceeds exactly as above. \square

Corollary 3.23 *Every ideal class contains an infinite set of incomparable graphs.* \square

3.6 Incomparable graphs on n vertices

In this section we shall estimate the maximum number of incomparable graphs on n vertices. In this context we shall use the random graph $\mathbf{G}(n, p)$, defined in the usual way, on the vertex set $1, 2, \dots, n$, to have each edge present with probability p . In our application we shall only take $p = \frac{1}{2}$. The properties of $\mathbf{G}(n, p)$ have been thoroughly investigated [4, 93, 196]. In particular, routine large deviation techniques can be used to prove the following results.

Theorem 3.24 *Let ϵ be an arbitrary positive number. The random graph $\mathbf{G}(n, p)$ has, with probability $1 - o(c^{-\log n})$ (where $c > 1$ is a constant), the following properties:*

1. *all degrees are at least $\frac{n}{2}(1 - \epsilon)$ and at most $\frac{n}{2}(1 + \epsilon)$,*
2. *the number of common neighbours of any two vertices is at least $\frac{n}{4}(1 - \epsilon)$ and at most $\frac{n}{4}(1 + \epsilon)$,*
3. *both the largest clique and the largest independent set have fewer than $2 \log_2(1 + \epsilon)n$ vertices,*
4. *each set of $m > 30 \log_2 n$ vertices induces a subgraph with at most $\frac{3}{4} \binom{m}{2}$ edges, and*
5. *there is a set of $m > 60 \ln n$ disjoint pairs of vertices $a_i b_i, i = 1, 2, \dots, m$, such that when each a_i is identified with the corresponding b_i the resulting*

graph has more than $\frac{3}{4}\binom{m}{2}$ edges amongst the a_i 's. (In other words, there are more than $\frac{3}{4}\binom{m}{2}$ pairs ij such that one of $a_i a_j, a_i b_j, b_i a_j, b_i b_j$ is an edge.)

Proof We shall give a brief sketch of the proof.

For any $\epsilon > 0$, the probability that $\mathbf{G}(n, 1/2)$ has property 1 tends to one as n goes to infinity.

We use the linearity of expected values, and Chernoff inequality. (More elementary proofs are possible, cf. [204] and the comments below.) If X_{ij} is the random variable denoting whether ij is an edge of \mathcal{G}_n ($X_{ij} = 1$) or not ($X_{ij} = -1$), then clearly the expected value of each X_{ij} is 0, and hence so is the expected value of each sum S_i , ($i = 1, 2, \dots, n$), of X_{ij} over all j . This situation (sum of random variables with expected value zero) is the context of the Chernoff inequality [4], which implies that the probability of S_i being greater than ϵ is less than $e^{-\epsilon^2/2n}$, for any positive ϵ . We can conclude from this that the probability that the degree of vertex i differs from $n/2$ by more than $\epsilon n/2$ is less than $2e^{-\epsilon^2 n/2}$. (The difference between the degree of i and $n/2$ is one half of the sum of X_{ij} , since edges contribute 1 and nonedges -1 .) We now claim that the probability that all degrees are between $\frac{n}{2}(1 - \epsilon)$ and $\frac{n}{2}(1 + \epsilon)$ tends to 1 as $n \rightarrow \infty$. Indeed, the probability that there exists a vertex i with degree greater than $\frac{n}{2}(1 + \epsilon)$ or smaller than $\frac{n}{2}(1 - \epsilon)$ is at most n times the probability that a particular vertex i has such a degree, which we have bounded by $e^{-\epsilon^2 n/2}$. Since $ne^{-\epsilon^2 n/2}$ tends to 0 as n tends to infinity, we have proved property 1.

Similar arguments prove property 2, and, with some additional refinements, also the other properties. Property 3 is a reformulation of the classical bound of Erdős for Ramsey numbers, cf. [124]. \square

We are now ready to state and prove the main theorem.

Theorem 3.25 *Suppose G and H are graphs on $1, 2, \dots, n$ satisfying properties 1–5. Then every homomorphism $f : G \rightarrow H$ is injective.*

Proof Suppose that f is a homomorphism of G to H which has $f(x) = f(y) = z \in V(H)$, for some $x \neq y \in V(G)$. Thus the set A of vertices adjacent to x or y in G must be mapped by f to the set B of vertices adjacent to z in H . Since A contains both the neighbourhood of x and the neighbourhood of y , which have at most $\frac{n}{4}(1 + \epsilon)$ vertices in common, A must have at least $n(1 - \epsilon) - \frac{n}{4}(1 + \epsilon) = \frac{3}{4}n - \frac{5}{4}n\epsilon$ vertices. On the other hand, B has at most $\frac{n}{2}(1 + \epsilon)$ vertices. Moreover, at most $2\log_2(1 + \epsilon)n$ vertices of A can be mapped by f to any vertex of B , according to property 3. It easily follows from this, that some $m > 30\log_2 n$ vertices v_1, v_2, \dots, v_m of B have the property that $f^{-1}(v_i)$ has at least two elements. This means, according to property 5, that there are at least $\frac{3}{4}\binom{m}{2}$ edges among the classes $f^{-1}(v_i)$, and hence among the vertices v_1, v_2, \dots, v_m , contrary to property 4. \square

We derive the following two corollaries.

Corollary 3.26 *Every graph satisfying properties 1–5 is a core* \square

Corollary 3.27 *Two graphs satisfying properties 1–5, and having the same number of vertices and edges, are either isomorphic or incomparable* \square

Since Theorem 3.24 implies that asymptotically almost all graphs satisfy properties 1–5, we apply the above two corollaries as follows.

Corollary 3.28 *Asymptotically almost all graphs are cores.* \square

Corollary 3.29 *When n is large, there exist*

$$\frac{1}{n!} \binom{\binom{n}{2}}{\lfloor \frac{1}{2} \binom{n}{2} \rfloor} (1 - \epsilon)$$

mutually incomparable graphs on n vertices. \square

(Recall that ϵ was an arbitrary positive number.)

We observe that this asymptotically the best lower bound on the possible number of mutually incomparable graphs, since the maximum number of graphs on a fixed set of n vertices, which are just inclusion-free (a much weaker property than mutual incomparability) is

$$\binom{\binom{n}{2}}{\lfloor \frac{1}{2} \binom{n}{2} \rfloor},$$

by Sperner's theorem. We shall show in Theorem 4.7 that asymptotically almost every graph has no automorphisms other than identity, and thus gives rise to $n!$ different graphs on the same vertex set.

3.7 Density

A partial order is *dense* if for any $a < b$ there exists an element c with $a < c < b$. We clearly cannot hope for the order \mathcal{C}_S to be dense, as no graph X satisfies $K_1 < X < K_2$. Indeed, we have $K_2 \leq X$ or $X \leq K_1$ for any X (depending on whether X has edges or not); in fact the same argument shows that no digraph X has $K_1 < X < K_2$.

However, we shall prove that this is the only exception, and the order \mathcal{C}_S is otherwise dense.

Theorem 3.30 *Let graphs G, H be cores such that $G < H$ and $G \neq K_1$ or $H \neq K_2$. Then there exists a graph X such that $G < X < H$.*

Proof If $G = K_1$ and $\chi(H) > 2$ then we can take $X = K_2$. On the other hand $G \neq K_1, H = K_2$ and $G < H$ is impossible. Thus we may assume that $\chi(G) > 1$ and $\chi(H) > 2$.

Let Z be a graph whose odd girth exceeds the odd girth of H , and whose chromatic number exceeds the chromatic number of the power graph G^H . Such a graph exists by Theorem 2.23. We now let $X = G + (H \times Z)$. (Recall that $+$ denotes the disjoint union). It is clear that G is homomorphic to X , via a

mapping that takes it identically to the copy of G in X , and also that $X \rightarrow H$, as $G \rightarrow H$ by assumption and $H \times Z \rightarrow H$ via a projection.

Since $H \not\rightarrow G$, some component C of H must have $C \not\rightarrow G$. If $H \rightarrow X$ then $C \rightarrow H \times Z \rightarrow Z$. Since $\chi(G) > 1$ the component C is not bipartite, and we have a contradiction since the odd girth of C is smaller than that of Z .

It remains to show that $X \not\rightarrow G$. Otherwise $H \times Z \rightarrow G$, and so $Z \rightarrow G^H$, by Corollary 2.18. This contradicts our assumption on the chromatic number of Z . \square

Another proof of the Theorem can be based on the Sparse Incomparability Lemma, Theorem 3.12 as follows. As above, we may assume that $G < H$ and H is nonbipartite. According to the Sparse Incomparability Lemma, there exists a graph G' with $G' < H$ and girth greater than $|V(H)|$ that is incomparable with G . Let $X = G + G'$ and note that $G < X \leq H$. If H was homomorphic to X , then each component of H would be homomorphic to G —as the girth of G' ensures there is no homomorphism of a nonbipartite component of H to G' . This contradicts the fact that $G < H$. Thus $G < X < H$. \square

Let us say that the pair $[G, H]$ with $G, H \in \mathcal{C}_S, G < H$, is a *gap* in \mathcal{C}_S , if no $X \in \mathcal{C}_S$ satisfies $G < X < H$. Therefore, we may reformulate Theorem 3.30 as follows.

Corollary 3.31 *There is only one gap in \mathcal{C}_S , namely $[K_1, K_2]$.* \square

By contrast, the partial order \mathcal{C} has many gaps. For instance, for zig-zags, each $[Z_{n+1}, Z_n]$ is a gap. It is known that for each oriented tree T there exists a digraph G such that $[G, T]$ is a gap, see below. It turns out, however, that if $G < H$ in \mathcal{C} , and if H is a connected core which is not an oriented tree, then $[G, H]$ is not a gap.

Theorem 3.32 *Suppose $G < H$ are digraphs, where H is a connected core which contains an oriented cycle. Then there exists a digraph X with $G < X < H$.*

Note that this theorem implies Theorem 3.30, by specializing to graphs.

Proof Let ab be an arc of H which belongs to a cycle. Let I be the digraph obtained from H by deleting the arc ab , adding a new vertex a' and the arc $a'b$ (Fig. 3.6(a)). Let n be an integer greater than the number of vertices of G . Finally, let G' be the digraph obtained from the transitive tournament \vec{T}_n by replacing each arc xy with a copy of I , identifying x with a and y with a' (Fig. 3.6(b)). (More about this replacement operation in the next chapter; cf. also Fig. 1.10.) Note that all vertices of \vec{T}_n are included in G' . We claim that $X = G' + G$ satisfies $G < X < H$ as claimed. The first inequality follows from the fact that any homomorphism $G' \rightarrow G$ would have to identify two vertices of \vec{T}_n and hence contain a homomorphic image of H in G , contrary to our assumption that $G < H$.

To prove the second inequality, assume that there exists a homomorphism $g: H \rightarrow G'$. Let $f: G' \rightarrow H$ be the homomorphism which maps all the vertices




FIG. 3.6. (a) The replacement digraph I ; (b) a local view of G' .

of \vec{T}_n to a and all other vertices of G' identically to their corresponding vertices of H . Since H is a core, the composition $f \circ g : H \rightarrow H$ is an automorphism of H , and hence there exists a homomorphism $h : H \rightarrow H$ (the automorphism of H which is the inverse of $f \circ g$) such that $f \circ g \circ h$ is the identity on H . Consider now the composition $g \circ h : H \rightarrow G'$; it takes a to some vertex u in G' such that $f(u) = a$. This means that u is a vertex of \vec{T}_n . (None of the other vertices of G' are taken to a by f .) Note that $g \circ h$ cannot take a vertex of H different from a to a vertex of \vec{T}_n , since $f \circ g \circ h$ is the identity on H . Therefore, the image of the connected digraph H under $g \circ h$ is included in union of the set of copies of I meeting at u (Fig. 3.6(b)). Now consider $u' = g(h(b))$. Since $f(u') = b$, the vertex u' is an outneighbour of u , and must in fact be one of the vertices marked by an asterisk in Fig. 3.6, i.e., must correspond to b in one of the copies of I which have their a' identified with u . We now observe that uu' does not belong to a cycle in $g(h(H))$; this is a contradiction, as $g(h(H))$ must be isomorphic to H ($f \circ g \circ h$ is the identity on H), and ab does belong to a cycle in H . \square

3.8 Duality and gaps

We relate the results of the previous section to the notion of homomorphism duality introduced in Chapter 1 (and further discussed in Section 5.4). Let \mathcal{C}' be a suborder of \mathcal{C} . (Typically \mathcal{C}' will be either \mathcal{C} or \mathcal{C}_S .) The simplest duality statements involve just two digraphs. We say that a pair of digraphs (F, H) is a *simple duality pair* (or a *single-graph duality*) if

$$G \not\rightarrow H \text{ if and only if } F \rightarrow G,$$

for all digraphs $G \in \mathcal{C}'$.

(In particular, we must have $F \not\rightarrow H$ in a simple duality pair.) In this case, we also say that H is a *dual* of F in \mathcal{C}' . In a simple duality pair, the nonexistence

of a homomorphism *to* H is certified by the existence of a homomorphism *from* F . We may assume that both F and H are cores. It is easy to see that this implies that F is connected. (If $F = F_1 + F_2$, then $F \not\rightarrow F_i, i = 1, 2$, since F is a core; thus each $F_i \rightarrow H$, and hence $F \rightarrow H$, a contradiction.) It follows from the definition that a dual is unique, up to homomorphic equivalence, so we refer to the core H as *the* dual of F .

For instance, the obvious equivalence

$$G \not\rightarrow K_1 \text{ if and only if } K_2 \rightarrow G$$

for graphs G verifies that (K_2, K_1) is a simple duality pair in \mathcal{C}_S .

Theorem 3.33 *(K_2, K_1) is the only simple duality pair in \mathcal{C}_S .*

Proof Consider a simple duality pair (F, H) , where F and H are cores, F not K_2 . Then F is not bipartite, and hence contains an odd cycle, say C_{2k+1} . Let G be a graph with odd girth greater than $2k + 1$ and chromatic number greater than the chromatic number of H (cf. Theorem 2.23). Then $F \not\rightarrow G$ because of the girth, and $G \not\rightarrow H$ because of the chromatic number. Thus $F = K_2$; then $K_2 \rightarrow G$ if and only if G has an edge, so $G \not\rightarrow H$ is only possible for $H = K_1$. \square

The reader may have noticed that the only gap in \mathcal{C}_S seems related to the only simple duality pair in \mathcal{C}_S . This is not a coincidence. Indeed, gaps and simple duality pairs are in a one-to-one correspondence.

Theorem 3.34 *In the partial orders \mathcal{C} or \mathcal{C}_S , we have the following relationship between gaps and simple duality pairs.*

- If cores (F, H) form a simple duality pair, then $[F \times H, F]$ is a gap.
- If cores $[A, B]$ form a gap and B is connected, then (B, A^B) is a simple duality pair.

The theorem applies more generally in any structures with products and exponentiation, such as general relational systems. We have stated the proof in a way that makes this clear. (This is one of the advantages of the homomorphism approach.)

Proof Suppose first that (F, H) is a simple duality pair, with F, H cores. We claim that there does not exist a K with $F \times H < K < F$. Indeed $F \not\rightarrow K$ means $K \rightarrow H$, and since $K \rightarrow F$ we must also have $K \rightarrow F \times H$, a contradiction.

On the other hand, suppose $[A, B]$ is a gap, A and B are cores, and B is connected. We claim that $G \not\rightarrow A^B$ if and only if $B \rightarrow G$. If $B \rightarrow G \rightarrow A^B$, then $B \rightarrow B \times A^B$ and since $B \times A^B \rightarrow A$ according to Corollary 2.18, we would have $B \rightarrow A$, a contradiction. If $G \not\rightarrow A^B$ and $B \not\rightarrow G$, then $C = A + G \times B$ would satisfy $A < C < B$ contrary to our assumption that $[A, B]$ was a gap. To verify $C \not\rightarrow A$, note that $G \times B \rightarrow A$ is equivalent to $G \rightarrow A^B$, again according to Corollary 2.18. To verify that $B \not\rightarrow C$, note that otherwise we would have $B \rightarrow G \times B$, contradicting $B \not\rightarrow G$. \square

In Proposition 1.20 we have found the simple duality pair (\vec{P}_k, \vec{T}_k) . There are other examples in \mathcal{C} . For instance, Fig. 3.7 shows graphs F and H forming a duality pair. The reader is invited to verify that $G \not\rightarrow H$ if and only if $F \rightarrow G$. It follows from Theorems 3.32 and 3.34 that in any simple duality pair (F, H) , the core of the digraph F is an oriented tree. We now prove a converse. (Note that it applies to any F whose core is a tree.)




FIG. 3.7. A simple duality pair.

Theorem 3.35 *Every oriented tree F has a dual H in \mathcal{C} .*

Proof Let F be an oriented tree. We define $V(H)$ to consist of all mappings $f : V(F) \rightarrow V(F)$ for which each $f(u)$ is a neighbour (in- or out- neighbour) of u . Two such mappings form an arc fg in H , if for all arcs uv of F $f(u) \neq v$ or $g(v) \neq u$ (in other words, it is not the case that $f(u)g(v) = vu$).

We now show that H is a dual of F , i.e., that (F, H) is a simple duality pair. We first check that $F \not\rightarrow H$ (and hence $G \not\rightarrow H$ for any G with $F \rightarrow G$). Suppose instead that $\phi : F \rightarrow H$ is a homomorphism. Let u_0 be any vertex of F , and define the sequence $u_i, i \in \mathbb{N}$, recursively by $u_{i+1} = \phi(u_i)(u_i)$. Then u_{i+1} is a neighbour of u_i , for all nonnegative integers i . Since F is finite, some $u_m = u_n$; since F is an oriented tree, the smallest n which has such an $m < n$ must be $m + 2$. This means that $\phi(u_{m+1})(u_{m+1}) = u_m$, which together with $\phi(u_m)(u_m) = u_{m+1}$ (by definition) implies that $\phi(u_m)$ and $\phi(u_{m+1})$ are not neighbours in H , contrary to the fact that ϕ is a homomorphism.

It remains to show that if $F \not\rightarrow G$, then $G \rightarrow H$. For this purpose, we label all arcs of F as follows: all arcs incident with the leaves of F receive the label 1. Having labeled some arcs by $1, 2, \dots, i-1$, so that the remaining arcs form a subtree F' of F , we label by i all arcs incident with the leaves of F' . (Note that after this step, the unlabeled arcs still form a subtree of F .) Suppose u and v are neighbours in F . When the arc joining u and v is deleted from F , we obtain two trees— F_u , containing u , and F_v , containing v . We denote by $F_{u,v}$ the subtree consisting of F_v , together with u and the arc joining u and v .

We now define a homomorphism $\phi : G \rightarrow H$, assuming that $F \not\rightarrow G$. Consider an arbitrary vertex x of G . Since $F \not\rightarrow G$, each vertex u of F must have a

neighbour v for which no homomorphism of $F_{u,v}$ to G takes u to x . Let $F_x(u)$ be the set of arcs uv such that no homomorphism of $F_{u,v}$ to G takes u to x . Let $f_x(u)$ be a neighbour v for which $uv \in F_x(u)$ and the label of the arc uv is maximal. We let $\phi(x)$ be the mapping $f_x : V(F) \rightarrow V(F)$. Suppose xy is an arc of G . We shall show that for each arc uv of F we have $f_x(u) \neq v$ or $f_y(v) \neq u$, i.e., that $f_x f_y$ is an arc of H . Indeed, if the labels of all arcs incident to u other than uv are smaller than the label of uv , then $F_{v,u}$ admits a homomorphism to G taking v to y , and hence $f_y(v) \neq u$. Alternately, the labels of all arcs incident to v are smaller than the label of uv , and then $F_{u,v}$ admits a homomorphism taking u to x , and hence $f_x(u) \neq v$. \square

The last two theorems imply the following corollary.

Corollary 3.36 *For each oriented tree T there exists a unique digraph G such that $[G, T]$ is a gap in \mathcal{C} . No other gaps $[A, B]$ with connected B exist in \mathcal{C} .*

\square

More generally, we say that digraphs F_1, F_2, \dots, F_t and H have finite duality, or that H is the dual of F_1, F_2, \dots, F_t , in a suborder \mathcal{C}' of \mathcal{C} , if

$$G \not\rightarrow H \text{ if and only if } F_i \rightarrow G \text{ for some } i = 1, 2, \dots, t,$$

for all digraphs $G \in \mathcal{C}'$.

Theorem 3.37 *Any set $F_1, F_2, \dots, F_t, t \geq 1$, of oriented trees has a dual H in \mathcal{C} .*

Proof According to Theorem 3.35, there exist digraphs H_1, H_2, \dots, H_t , such that $G \not\rightarrow H_i$ holds if and only if $F_i \rightarrow G$ holds, for each $i = 1, 2, \dots, t$. Let H be the product $\prod_{i=1}^t H_i$.

If $G \not\rightarrow H$, then for some i we have $G \not\rightarrow H_i$, and hence $F_i \rightarrow G$. On the other hand, if $G \rightarrow H$ then $G \rightarrow H_i$ and hence $F_i \not\rightarrow G$ for every i . \square

It can also be shown [277] that if H is the dual in \mathcal{C} of any family F_1, F_2, \dots, F_t , then each F_i must be an oriented tree.

3.9 Maximal antichains in \mathcal{C}

We continue with our investigation of antichains maximal with respect to inclusion. We first note that there are singleton maximal antichains in \mathcal{C} . Clearly, the one-vertex digraph $\{\vec{P}_0\}$ (without arcs) forms a maximal antichain, since $\vec{P}_0 \leq H$ for any H in \mathcal{C} . Similarly, $\{\vec{P}_1\}$ is a maximal antichain, since $\vec{P}_1 \leq H$ for any H that has arcs, and $H \leq \vec{P}_1$ for any H that has no arcs. The antichain $\{\vec{P}_2\}$ is also maximal, since a digraph H for which $\vec{P}_2 \not\leq H$ cannot contain a vertex with both positive in-degree and positive out-degree. This means that H is homomorphic to \vec{P}_1 , and hence \vec{P}_2 . (All vertices of H with zero in-degree can be mapped to the vertex of \vec{P}_1 which has zero in-degree, all other vertices of H can be mapped to the other vertex of \vec{P}_1 .)

Proposition 3.38 *The partial order \mathcal{C} has exactly three maximal antichains of size one— $\{\vec{P}_0\}, \{\vec{P}_1\}, \{\vec{P}_2\}$.*

Proof Recall that a digraph G is balanced if each cycle has the same number of arcs going forward and backward. Assume first that G is not balanced. Let r be the number of arcs in an unbalanced cycle of G and let s be the chromatic number of the underlying graph of G . According to Corollary 3.13, there exists a graph H with chromatic number greater than s and girth greater than r . Then G and any orientation H' of H are incomparable. We have $G \not\preceq H'$ because the unbalanced cycle with r arcs in G must map to an unbalanced subgraph of H' , but because of the girth of H all subgraphs of H' of size at most r are balanced. We also have $H' \not\preceq G$ because the chromatic number of H exceeds that of G .

For balanced digraphs, we proceed differently. Assume G is a balanced core different from P_1, P_2, P_3 . It is easy to see that this implies that the height h of G is at least three; assume also that G has n vertices. Then $Z_n \cdot Z_n \cdots Z_n$, with sufficiently many factors, is not homomorphic to G because of height, and G is not homomorphic to it because of the lengths of the zig-zags. \square

Consider a simple duality pair (F, H) in \mathcal{C} , such as the two digraphs in Fig. 3.7. It follows from the definition of duality, that F, H is a maximal antichain of size two—since every other digraph admits either a homomorphism to H or a homomorphism from F . Similarly, we have the following general observation.

Proposition 3.39 *If F_1, F_2, \dots, F_t and H have finite duality in \mathcal{C}' , then F_1, F_2, \dots, F_t, H is a maximal antichain in \mathcal{C}' .* \square

Therefore, duality is also useful for finding maximal antichains of any size.

Corollary 3.40 *Let k be a positive integer. There exists in \mathcal{C} a maximal antichain of size k .*

Proof For $k = 1$, apply Proposition 3.38. For $k \geq 2$, take any set of incomparable oriented trees F_1, F_2, \dots, F_{k-1} (for instance the oriented paths from Proposition 3.7), together with their dual digraph H from Theorem 3.37. \square

This result stands out in contrast to Corollary 3.11, which implies that in \mathcal{C}_S there are no maximal antichains of any finite size $t > 1$.

In conclusion, in \mathcal{C} we know all the maximal antichains of size one, and a generous supply of antichains of size $t \geq 2$, arising from finitary dualities, Theorem 3.37. It is not known whether or not there are other maximal antichains of size $t \geq 2$ in \mathcal{C} ; however, in more general relational structures there are known to be others [277].

3.10 Bounds

An (*upper*) *bound* of a suborder \mathcal{C}' of \mathcal{C} is a digraph H such that $G \leq H$ for all $G \in \mathcal{C}'$. A bound of \mathcal{C}' which belongs to \mathcal{C}' is called *the maximum* of \mathcal{C}' . Recall that \mathcal{C} consists of cores only, and hence we treat homomorphically equivalent

graphs as being the same. In particular, if a maximum exists, it is unique. A bound H of \mathcal{C}' is a *supremum* of \mathcal{C}' if no bound H' of \mathcal{C}' has $H' < H$. Of course, a supremum of \mathcal{C}' belongs to \mathcal{C}' if and only if it is the maximum of \mathcal{C}' . A supremum of \mathcal{C}' which does not belong to \mathcal{C}' is called a *proper supremum*.

Many classical results about chromatic numbers and homomorphisms can be expressed in the language of bounds. Consider the following two examples. The Four Colour Theorem (*'every planar graph is four-colourable'*) says that K_4 is a bound for the class \mathcal{P} of planar graphs (or cores). In this case, $K_4 \in \mathcal{P}$; in other words, K_4 is actually the maximum of the suborder \mathcal{P} . In particular, there is no better bound H ($H < K_4$) for \mathcal{P} .

The theorem of Grötzsch (*'each triangle-free planar graph is three-colourable'*) says that K_3 is a bound for the class \mathcal{T} of triangle-free planar graphs. In this case, the bound K_3 is not a member of \mathcal{T} , and it turns out that better bounds H ($H < K_3$) can be found, Exercise 19.

Another result that can be nicely expressed in this language is the Sparse Incomparability Lemma, Theorem 3.12. Let H be a nonbipartite graph, and let $\mathcal{C}' \subseteq \mathcal{C}$ be the class of all G , $G < H$, with girth at least ℓ . Then H is the supremum of \mathcal{C}' .

There are well-known results concerning the chromatic number of bounded degree graphs. Consider for instance the class \mathcal{K} of all cubic graphs. (The logic is similar for the class of graphs with all degrees at most three, or at most any $d \geq 3$). It is easy to see that K_4 is a bound of \mathcal{K} , but the theorem of Brooks asserts that for graphs in \mathcal{K} without K_4 , the graph K_3 is a better bound. Moreover, when K_3 is forbidden, then the graphs in \mathcal{K} have an even better bound, as explained in Proposition 1.24. The same situation keeps repeating itself: the more graphs we forbid (in the sense explained below), the better bound we will obtain.

Recall that $\text{Forb}\mathcal{F}$ denotes the set of all graphs G such that no $F \in \mathcal{F}$ has $F \leq G$. Recall also, that the set $\text{Forb}\mathcal{F}$ is an ideal class. We denote by $\text{Forb}_d\mathcal{F}$ the set of all graphs in $\text{Forb}\mathcal{F}$ which have all vertices of degree at most d . While $\text{Forb}_d\mathcal{F}$ is no longer an ideal class in general, it is still closed under taking subgraphs and disjoint unions.

We shall use the following extension of Proposition 1.24.

Proposition 3.41 *Let $d \geq 3$ be an integer, let H be a graph, and let \mathcal{F} be a finite set of connected graphs.*

Denote by \mathcal{S} the set of all H -colourable graphs in $\text{Forb}_d\mathcal{F}$. Then there exists an H -colourable graph $H' \in \text{Forb}\mathcal{F}$ which is a bound of \mathcal{S} .

Note that when $d = 3$, $H = K_3$ and $\mathcal{F} = \{K_3\}$, the proposition implies that there exists a three-colourable triangle-free graph H' to which all cubic (three-colourable) triangle-free graphs are homomorphic. Note that a triangle-free cubic graph is automatically three-colourable, by Brooks' theorem; see also Exercise 8.

Proof We only sketch the proof, the details are checked in analogy with the proof of Proposition 1.24. Let a be the maximum number of vertices of a graph

in \mathcal{F} , and let X be a fixed set with $b = 1 + d^{2a}$ elements. The vertices of H' are triples (A, x, v) , where A is a connected graph in $\text{Forb}_d\mathcal{F}$ with vertices from X , x is a vertex of A , and $v \in V(H)$. Two vertices $(A, x, v), (A', x', v')$ are adjacent in H' if vv' is an edge of H , xx' is an edge of both A and A' , each vertex of A of distance at most a from x and each vertex of A' of distance at most a from x' is a vertex of both A and A' , and, moreover, $E(A)$ and $E(A')$ agree on $V(A) \cap V(A')$. \square

Theorem 3.42 *Let $d \geq 3$, and let \mathcal{F} be a finite set of connected graphs.*

If H is a bound for $\text{Forb}_d\mathcal{F}$, then either H is itself in $\text{Forb}_d\mathcal{F}$, or there is a better bound $H' < H$ for $\text{Forb}_d\mathcal{F}$.

In other words, the set $\text{Forb}_d\mathcal{F}$ does not have a proper supremum.

Proof Assume H is a bound for $\text{Forb}_d\mathcal{F}$ and $H \notin \text{Forb}_d\mathcal{F}$. Thus a component K of H is not a member of $\text{Forb}_d\mathcal{F}$. We let \mathcal{F}' denote the set $\mathcal{F} \cup \{K\}$. (Our new family \mathcal{F}' still consists of connected graphs only.)

We now claim that $\text{Forb}_d\mathcal{F}'$ has a bound H' with $H' < H$. Indeed, applying the previous proposition to \mathcal{F}' , we conclude that $\text{Forb}_d\mathcal{F}'$ has a bound $H' \in \text{Forb}_d\mathcal{F}'$ with $H' \leq H$. Since K is a component of H that is not homomorphic to H' , we must have $H' < H$.

It remains to prove that H' is a bound for $\text{Forb}_d\mathcal{F}$. In fact, it turns out that $\text{Forb}_d\mathcal{F} = \text{Forb}_d\mathcal{F}'$. The inclusion $\text{Forb}_d\mathcal{F}' \subseteq \text{Forb}_d\mathcal{F}$ is obvious, since the family \mathcal{F}' contains the family \mathcal{F} . Suppose, for a proof by contradiction, that some $G \in \text{Forb}_d\mathcal{F}$ was not a member of $\text{Forb}_d\mathcal{F}'$, i.e., that $K \leq G$. Let L be a component of G with $K \leq L$. We have $G \leq H$, since $G \in \text{Forb}_d\mathcal{F}$. This is only possible if $L \leq K$, since H is a core. For the same reason, $K = L$. We have reached the desired contradiction, since G , and hence L , belongs to $\text{Forb}_d\mathcal{F}$, contradicting our assumption on K . \square

We close this chapter by expressing, in the framework of the homomorphism order, the well-known conjecture of Hadwiger. Recall that we have defined minors and contractions at the end of Section 1.3. As a generalization of the Four Colour Conjecture (now Theorem), Hadwiger conjectured that any graph G with chromatic number n contains K_n as a minor. A class \mathcal{K} of graphs is *minor-closed* if each minor of a member of \mathcal{K} is also in \mathcal{K} ; and \mathcal{K} is called *proper* if it does not contain all graphs. It is known that a proper minor-closed class of graphs must have bounded chromatic number. (For instance, Mader's Theorem states that any graph with high enough minimum degree contains as a minor any prescribed graph G [234].) Therefore, each graph in a proper minor-closed class has a vertex of degree smaller than some d , and hence is d -colourable.

Proposition 3.43 *Hadwiger's Conjecture is equivalent to the following statement: every proper minor-closed class \mathcal{K} has a maximum.*

Proof Note that the class of all graphs is minor-closed, but has no maximum. Hence the assumption that \mathcal{K} is proper is necessary. Assuming Hadwiger's conjecture holds, consider a proper minor-closed class \mathcal{K} . Since the chromatic number

of the graphs in \mathcal{K} is bounded, let G be a graph in \mathcal{K} with the greatest chromatic number, $\chi(G) = n$. Then \mathcal{K} is bounded by K_n . Since Hadwiger's conjecture holds, G contains K_n as a minor, and hence $K_n \in \mathcal{K}$. Thus K_n is the maximum element of \mathcal{K} .

Conversely, assume that each proper minor-closed class has a maximum.

We first prove that each minor-closed class of graphs with a maximum must have the maximum (homomorphically equivalent to) a complete graph. Otherwise, let \mathcal{C} be a minor-closed class with maximum H , not homomorphically equivalent to a complete graph. We may assume that \mathcal{C} has been chosen to be minimal with respect to inclusion. This means, in particular, that \mathcal{C} consist of H and all its minors. Let K_k be the largest clique in \mathcal{C} . It is clear that $k \geq 3$ —if H had no triangle minor, it would be a forest, and hence homomorphically equivalent to K_1 or K_2 . Since $K_k \rightarrow H$, the graph H contains a clique K with k vertices. We now distinguish two cases. First assume that K contains a vertex x only adjacent to the other $k - 1$ vertices of K . Note that any $H - x$ must be k -colourable. Indeed, the proper subclass \mathcal{C}' of \mathcal{C} , consisting of $H - x$ and all its minors, which has a maximum according to our global assumption, would have to have a maximum which is a clique, by the assumption of minimality of \mathcal{C} . Obviously, such a clique cannot have more than k vertices, by the definition of k . However, this would mean that H is also k -colourable, since x has only $k - 1$ neighbours. This is a contradiction, since H would be homomorphically equivalent to K_k . In the remaining second case, every vertex of K has a neighbour outside of K . We claim that $H - K$ is connected. Otherwise, suppose induced subgraphs A and B of $H - K$ have no edges joining them and contain all vertices of H not in K . Then the subgraph $A + K$, induced by A and K , is k -colourable, by the same argument as above, and similarly for $B + K$. It is easy to see that H itself has a k -colouring, a contradiction. If we now contract all vertices of $H - K$, we obtain a K_{k+1} -minor, contradicting the choice of k .

If G is any graph with chromatic number n , then consider the class \mathcal{K} consisting of G and all its minors. As a proper minor-closed class of graphs, \mathcal{K} must have a maximum element H . From above we know we may take $H = K_m$ for some m ; whence from $G \leq K_m$ we deduce that $n \leq m$. Therefore, G also contains K_n as a minor, implying the validity of Hadwiger's conjecture. \square

Thus we obtain Hadwiger's conjecture in the following form.

Conjecture 3.44 *Every proper minor-closed class of graphs has a maximum.*

Note that the above proof actually shows that it is enough to prove that every class consisting of one graph and all its minors has a maximum. Therefore, we can further simplify Hadwiger's conjecture to state that *if both G and H are minors of a graph W , then there exists a minor W' of W to which both G and H are homomorphic.*

If we employ a dashed arrow from A to B to indicate that A is a minor of B , we can depict the conjecture by the diagram in Fig. 3.8.

Exercise 11 is discussed in [4]. Exercise 13 is from [151]. Exercise 15 is from [178]. Exercise 17 is based on [279]. Exercises 18, 19 are based on [252]. The first bounds $G < K_3$ were found in [262, 263].

3.12 Exercises

1. Prove that a core H is projective if and only if any polymorphism $f : H^t \rightarrow H$ is a projection followed by an automorphism of H .
2. Let H be a graph. Prove that if there exists a uniquely H -colourable graph then H must be a core.
3. Prove that a nonbipartite graph H is a supremum of the class of graphs G with $G < H$.
4. Let ℓ be an odd integer, $\ell \geq 3$. Prove that $\text{Forb}C_\ell$ is the class of all graphs with odd girth at least ℓ .
5. Prove that any two duals H_1, H_2 of a family of digraphs F_1, F_2, \dots, F_t are homomorphically equivalent.
6. Prove that the digraph H is the dual of $\mathcal{F} = \{F_1, F_2, \dots, F_t\}$ in \mathcal{C} if and only if H is the maximum of $\text{Forb}\mathcal{F}$.
Deduce that $\text{Forb}\mathcal{F}$ for a finite family \mathcal{F} has a maximum in \mathcal{C} , if all members of \mathcal{F} are oriented trees.
7. Let \mathcal{K} be an ideal class containing a nonbipartite graph. Describe all maximal antichains in \mathcal{K} of size one and prove that there are no other finite maximal antichains in \mathcal{K} .
8. Prove directly from Proposition 1.24 that there exists a *three-colourable* triangle-free graph to which all cubic triangle-free graphs are homomorphic. (Hint: Use products.)
9. Prove that asymptotically almost all graphs are G -pointed.
10. Show, using Proposition 3.8, and the techniques of Section 4.4, (Fig. 4.7), that every finite partial order is isomorphic to a suborder of \mathcal{C}_S restricted to planar graphs with maximum degree three.
11. Let k be a fixed positive integer. Show that asymptotically almost all graphs G have the following *k-extension property*: if A, B are disjoint sets of vertices of G with $|A \cup B| = k$, then some vertex x of G is adjacent to all $a \in A$ and nonadjacent to all $b \in B$.
(Hint: The probability that the random graph $\mathbf{G}(n, 1/2)$ does not have such a vertex x for a particular pair of sets A, B is $(1 - 2^k)^{n-k}$ [229].)
12. Using the previous exercise and the following two claims, show that asymptotically almost all graphs are projective.
Assume G is a graph (with at least three vertices) which has the three-extension property, and $f : G \times G \rightarrow G$ is an idempotent homomorphism.
Claim 1 For any $u, v \in V(G)$ we have $f(u, v) = u$ or $f(u, v) = v$.
Claim 2 If $f(u, v) = u$ for some u, v , then $f(x, y) = x$ for all vertices x, y of G .

(Hint: Enough to show asymptotically almost all graphs are two-projective [214]; for the proofs of the claims, cf. Theorem 2.43.)

- 13. Prove that the Four Colour Theorem is a consequence of the assertion that the class of planar graphs has a maximum.
- 14. Prove that each $\{K_n\}$ with $n = 1, 2, 3, 4$ is a maximal antichain in the class of planar cores. (It is not known whether or not there exist other maximal antichains.)
- 15. Find a countable set of incomparable tournaments.
- 16. Let G be a graph. Prove that a graph H is G -pointed if and only if the graph $\text{HOM}(G, H)$ (defined in Section 2.11) has no edges. Deduce that asymptotically almost all graphs are G -pointed. (Hint: Use Corollary 3.28.)
- 17. Prove that the (underlying graph of the) dual of an oriented tree of height h is h -colourable.

Prove that the digraph H in Fig. 3.9 is the dual of the oriented tree T . Note that H is three-colourable even though T has height (net length) four.




FIG. 3.9. A duality pair for Exercise 17.

- 18. Let G be the following Cayley graph: the vertices of G are elements of the group Z_2^5 , i.e., five-tuples of 0's and 1's, with $a_1a_2a_3a_4a_5$ adjacent to $b_1b_2b_3b_4b_5$ if and only if their difference in Z_2^5 is 11000 or 01100 or 00110 or 00011 or 10001. Prove that $G \times K_3 < K_3$.
- 19. [252] Also prove that $G \times K_3$ (from the previous exercise) is a bound for the class \mathcal{T} of triangle-free planar graphs.

THE STRUCTURE OF COMPOSITION

4.1 Introduction

In this chapter we examine the richness of the structure of homomorphisms with respect to their composition. Recall that homomorphisms generalize automorphisms, and that we have proposed in Section 1.7 to extend, to the context of homomorphisms, results representing any group as the automorphism group of a suitable graph (possibly enjoying some prescribed graph properties, such as given connectivity, chromatic number, etc.) In fact, we have already partially extended these results from automorphisms to endomorphisms: we have shown, Theorem 1.35, that every monoid can be represented as the endomorphism monoid of a binary relational system. It will follow from the results of this chapter that the binary system can be replaced by a graph (and, moreover, the graph can be chosen to have some typical graph properties). As promised in the first chapter, we shall do this in the more general context of categories. We prove that this is the case for any finite category, and explore obstacles for such representations of infinite categories.

4.2 Rigid digraphs

Let us begin by representing the *trivial monoid*, consisting of just the identity element 1. The representation produced by Theorem 1.35 is a binary relational system with one vertex and one relation consisting of a loop. Of course, this can easily be interpreted as a graph—and the loop is not even necessary, since the trivial graph K_1 obviously represents the trivial monoid. However, as suggested in Fig. 1.10, we shall need nontrivial graphs with a trivial endomorphism monoid. This is what we shall do here.

We say that a digraph is *asymmetric* if it has no automorphisms other than the identity, and *rigid* if it has no endomorphisms other than the identity. Note that *a digraph is rigid if and only if it is asymmetric and a core*. As observed above, the trivial digraph K_1 is rigid, but it is neither interesting nor useful, and we shall from now consider only nontrivial rigid digraphs.

Rigid digraphs are plentiful. For instance, Proposition 1.14 implies that each directed path \vec{P}_n is rigid. In fact, we have the following easy observation.

Proposition 4.1 *Any digraph which is acyclic and has a directed Hamilton path is rigid.*

Proof No endomorphism can identify two vertices, as this would create a cycle. Moreover, the directed Hamilton path is unique (otherwise there has to be a directed cycle), and so there is no automorphism other than the identity. \square

It follows from the above proof, that digraphs obtained by adding some simple structure to \vec{P}_n remain rigid. Let $\vec{P}_n(i)$ denote the digraph obtained from \vec{P}_n by the addition of the arc $1i$ (Fig. 4.1). Thus each $\vec{P}_n(i)$ is rigid. In fact, the above proof also shows that a homomorphism $f : P_n(i) \rightarrow P_n(j)$ is only possible when $i = j$ and in that case f must be the identity. We call a family of digraphs *incomparable*, if any two members X and Y of the family are incomparable, i.e., have $X \not\rightarrow Y$ and $Y \not\rightarrow X$.




FIG. 4.1. The rigid digraph $\vec{P}_n(i)$.

Corollary 4.2 *The digraphs $\vec{P}_n(i)$, $i = 3, 4, \dots, n-1$, form an incomparable family of rigid digraphs.*

Recall the zig-zags Z_k introduced in Fig. 3.1. We can construct other incomparable families of rigid digraphs using these zig-zags. Concatenation $P \cdot R$ of oriented paths P, R is defined in Chapter 3 (cf. Fig. 3.2). Concatenating Z_2 with directed paths yields the following family of incomparable rigid oriented paths, illustrated in Fig. 4.2.

Proposition 4.3 *The oriented paths $\vec{P}_i Z_2 \vec{P}_{n-i}$, $i = 0, 1, \dots, n$ form an incomparable family of rigid digraphs.*

Proof This follows directly from Proposition 1.14. \square




FIG. 4.2. Incomparable rigid oriented paths $\vec{P}_0 Z_2 \vec{P}_1, \vec{P}_1 Z_2 \vec{P}_0$.

If P is an oriented path with initial vertex u and terminal vertex v , we denote by P^{-1} the *reversal* of P , i.e., the same path but with initial vertex v and terminal vertex u . We now define a family of oriented paths $Q(a, b) = Z(a) \cdot Z(b) \cdot (Z(a))^{-1} \cdot (Z(b))^{-1}$ (Fig. 4.3). The oriented paths $Q(a, b)$ have net length zero, i.e., the initial and terminal vertices are on the same level, while all other vertices are at a level between zero and the height of $Q(a, b)$, which is six. (An easy modification of $Q(a, b)$ changes the height to four.)




FIG. 4.3. The oriented path $Q(2, 3)$.

Proposition 4.4 *The oriented paths $Q(i, 2n + 1 - i), i = 1, 2, \dots, n$, form an incomparable family of rigid digraphs.*

Proof This follows from Propositions 1.14 and 3.6. □

It is much harder to find rigid *graphs*. Even asymmetric graphs are not trivial to construct. In Figs. 4.4 and 4.5, we depict the smallest rigid graph (with eight vertices and fourteen edges), the smallest asymmetric graph (with six vertices and six edges), and the smallest asymmetric tree (with seven vertices) (see Exercise 17).




FIG. 4.4. The smallest rigid graph.




FIG. 4.5. The smallest asymmetric graph and the smallest asymmetric tree.

We now discuss some basic techniques for constructing rigid graphs. A graph G is *vertex-critical* if $\chi(G) > \chi(G')$ for all subgraphs G' with fewer vertices. It is easy to check that each vertex-critical graph is a core (Exercise 6 in Chapter 1). Therefore, one way to obtain rigid graphs is by constructing asymmetric vertex-critical graphs. We apply this way of thinking to the graphs H_k introduced in Fig. 4.6.




FIG. 4.6. The graphs H_k (for $k \geq 1$).

Proposition 4.5 *Each graph H_k is rigid.*

Proof We first note that, for each $k \geq 1$, we have $\chi(H_k) = 4$. Indeed, if a three-colouring of H_k were possible, then without loss of generality we would have vertices 1, 2, 3 coloured by 1, 2, 3, and hence each $3\ell + i$ would be coloured by i , resulting in two adjacent vertices 1 and $3k + 7$ both of colour 1. Moreover, $\chi(H') \leq 3$ for all proper induced subgraphs H' of H_k , since whenever any one vertex is deleted, the remaining vertices can be consecutively three-coloured 1, 2, 3, 1, 2, 3, 1, \dots , 3. Thus each H_k is a core

To see that H_k is asymmetric, we argue as follows. Any automorphism of H_k must preserve the degrees, and hence take vertex 6 to itself (it is the only vertex of degree five). Then the edge 16 must map to itself (and hence 1 must map to

1), since it is the only edge incident to 6 and not in a triangle. Then 2 must map to itself, as the only vertex of degree three not adjacent to another vertex of degree three, and hence 3 must also map to itself, since it is the only vertex adjacent to both 1 and 2. It is now easy to continue like this, showing that the automorphism must be the identity. \square

We also note for future reference that each graph H_k is *triangle-connected*, i.e., any two of its vertices are joined by a path v_1, v_2, \dots, v_p in which each v_i is adjacent to v_{i+2} , as long as $i + 2 \leq p$. (In other words, any two vertices can be joined by a sequence of triangles, with consecutive triangles sharing an edge.) Triangle-connectivity is preserved under homomorphisms, whence it follows that each homomorphic image of H_k is triangle-connected, and in particular is two-connected (has no cutpoints). This property will be helpful in the sequel, to limit the possible homomorphisms of H_k to other graphs.

Another useful construction of rigid graphs is introduced in Fig. 4.7. We present a family of graphs G_k , for all nonnegative integers k . The labeled dashed lines represent internally disjoint paths of the indicated lengths.




FIG. 4.7. The graphs G_k (for $k \geq 0$).

Proposition 4.6 *Each graph G_k is rigid.*

Proof Assume that k is a nonnegative integer. We first show that G_k is a core. Since G_k contains odd cycles, its core cannot be bipartite. Consider proper subgraphs of G_k . Deleting one vertex v from G_k results in a graph $G_k - v$ which is either an odd cycle with three attached paths, or two odd cycles sharing a common path, and two attached paths. Each such graph has a core consisting of one odd cycle. Therefore, if G_k is not a core, its core C would have to be one of the three odd cycles of length $2k + 5$. That would mean that G retracts to C . We shall now argue that this is not possible. It is easy to check that each cycle C of length $2k + 5$ contains a vertex $u \in C$ adjacent to a vertex $v \notin C$, such that for the two neighbours u', u'' of u on C we have both a cycle of length $2k + 5$ containing v, u, u' , and a cycle of length $2k + 5$ containing v, u, u'' . Since G_k does

not contain a cycle of length $2k + 3$, this implies that v cannot be retracted to C .

Finally, we argue that G_k is asymmetric. It is enough to focus on the vertices of degree three. Only the central vertex has the distances 1, 2, 3 to the other vertices of degree three. Hence every automorphism of G_k fixes the central vertex, and hence also its neighbouring vertex of degree three. It easily follows that the only automorphism of G_k is the identity. \square

We observe that the girth of G_k is $2k + 5$. Moreover, each vertex of G_k belongs to an (odd) cycle of length $2k + 5$. This limits the possible homomorphisms to other graphs, since any homomorphic image of G_k must have each vertex in an odd cycle of length at most $2k + 5$. We also note that $G_k - \{a, b, c\}$ is bipartite (in fact, a tree), hence G_k admits a three-colouring in which a, b, c are the only vertices of colour 1; in particular, it admits a three-colouring in which a and b have the same colour.

Even though rigid graphs may appear to be very special, it turns out that nearly every large graph is rigid (and hence asymmetric).

Theorem 4.7 *Asymptotically almost all graphs are rigid.*

Proof Recall that we have shown that the standard random graph $\mathbf{G}(n, p)$, with $p = \frac{1}{2}$, is a core, Corollary 3.28. Thus asymptotically almost all graphs on n vertices are cores, and it will be sufficient to show that asymptotically almost all graphs are asymmetric.

First we estimate the number of graphs on V , with a given group acting on V . Define *pair-orbits* of V to be the orbits of the group acting on the set of unordered pairs of vertices in V . Consider now a particular graph G with the given group as its group of automorphisms. For each pair-orbit, either all pairs are edges of G or no pair is an edge of G . Thus the number of such graphs G is at most 2^k , where k is the number of pair-orbits. Let us now estimate the number of graphs on V that admit an automorphism taking x to y . We claim that all such graphs have an automorphism group with at most $\binom{n}{2} - cn$ pair-orbits, for some positive constant c . Indeed, asymptotically, in almost every graph the degrees are all between $\frac{n}{2}(1 - \epsilon)$ and $\frac{n}{2}(1 + \epsilon)$ (for any positive ϵ), according to Theorem 3.24. Since x is taken to y , all the at least $\frac{n}{2}(1 - \epsilon)$ edges incident on x must map to the edges incident on y , i.e., all but possibly one (the edge xy , if present, could be its own orbit) must belong to pair-orbits of size at least two. Therefore, there are at most $\binom{n}{2} - \frac{n}{2}(1 + \epsilon) + 1$ pair-orbits. Consequently, there are at most $2^{\binom{n}{2} - cn}$ graphs that admit an automorphism taking x to y . Now the number of graphs that admit any nontrivial automorphism is at most the sum, over all pairs x, y , of the number of graphs that admit an automorphism taking x to y , hence at most $n^2 \cdot 2^{\binom{n}{2} - cn}$. Therefore the proportion of such graphs tends to zero:

$$\frac{n^2 \cdot 2^{\binom{n}{2} - cn}}{2^{\binom{n}{2}}} = \frac{n^2}{2^{cn}} \rightarrow 0.$$

□

We note that it follows that we have the same result even when counting nonisomorphic graphs.

Corollary 4.8 *Asymptotically almost all nonisomorphic graphs are rigid.*

Proof Let A and a denote respectively the number of asymmetric graphs on n vertices and the number of nonisomorphic asymmetric graphs on n vertices. It is easy to see that $A = a \cdot n!$, and it follows from the proof of Theorem 4.7 that $A = 2^{\binom{n}{2}}(1 - (n^2/2^{cn}))$. Hence

$$a = \frac{2^{\binom{n}{2}}(1 - (n^2/2^{cn}))}{n!},$$

while the total number of nonisomorphic graphs on n vertices is at most $2^{\binom{n}{2}}/(n!)$. □

4.3 An excursion to infinity

Although we usually consider only finite graphs, at this point we would like to make a detour and mention an important result of interest for infinite digraphs. We have observed above that each directed path is rigid; therefore, there are rigid digraphs with any finite number of vertices. Since rigid graphs and digraphs are the building blocks of many constructions (as we shall see in the next section), it is also important to obtain rigid digraphs (and graphs) of any infinite cardinality.

Theorem 4.9 *There is a rigid digraph on any infinite set.*

Proof Let X be an infinite set. We shall construct a digraph whose vertex set is in a bijective correspondence with the set X . (The correspondence can then be used to transfer the digraph onto X in the obvious way.) We may assume that X is infinite, and that X is an ordinal α , i.e., that the set X consists of all ordinals (starting from $0, 1, 2, \dots$) smaller than α , $X = \{\beta : \beta < \alpha\}$. In this terminology, the ordering $<$ of ordinals corresponds to the usual inclusion, and each set of ordinals has a supremum (corresponding to their union). For additional properties of ordinals the reader may consult any standard book on set theory, such as [194]. Let X' be a disjoint copy of X , with vertices denoted by primed ordinals, i.e., $\beta' \in X'$ corresponding to $\beta \in X$. Further let $\{a, b, c, a', b', c'\}$ be six additional vertices disjoint from X and X' .

We begin constructing a rigid digraph on $X \cup X' \cup \{a, b, c, a', b', c'\}$ as follows: we take the arcs $\beta\gamma$ and $\beta'\gamma'$, for all $\beta < \gamma < \alpha$, the arcs $\beta\beta'$, for all $\beta < \alpha$, and the arcs $0a, ab, bc, c0, b0$ and $0'a', a'b', b'c', c'0', a'c'$.

Finally, for every ordinal $\beta < \alpha$ that is the supremum of a countable sequence of smaller ordinals, we choose one particular increasing sequence $\{\beta_n\}$ such that $\sup \beta_n = \beta$, and we add a disjoint directed path of length $2n + 3$ from β to β'_n , for each n (Fig. 4.8).




FIG. 4.8. An infinite rigid digraph.

The vertex set of the resulting graph G is a union of fewer than α sets of cardinality smaller than α and thus $V(G)$ and X are in a bijective correspondence. We now show that G is indeed rigid. Let $f : G \rightarrow G$ be a homomorphism. Observe that 0 and $0'$ are the only vertices that belong both to a directed cycle and to an infinite subtournament of G . It is easy to see that $f(0) = 0'$ or $f(0') = 0$ are impossible, hence we may assume that $f(0) = 0, f(0') = 0'$. It follows easily that f restricted to the set $\{a, a', b, b', c, c'\}$ is the identity, and that f maps X to X and X' to X' . Since the arcs $\beta\beta'$ are the only arcs between X and X' , we have $(f(\beta))' = f((\beta)')$ for every $\beta < \alpha$. Note also that f restricted to X (or X') is a monotone mapping, with respect to $<$. If f is not the identity mapping then there is a first β with $f(\beta) \neq \beta$; necessarily $f(\beta) > \beta$, since $f(\beta) < \beta$ would imply $f(f(\beta)) < f(\beta)$ by monotonicity, thus contradicting the minimality of β . Therefore, we have $\beta < f(\beta) < f(f(\beta)) < \dots$. In other words, letting $\gamma_0 = \beta, \gamma_{n+1} = f(\gamma_n)$, we obtain a sequence $\gamma_0 < \gamma_1 < \gamma_2 < \dots$; denote by δ the supremum of the sequence. Recall that we have fixed a particular increasing sequence with supremum δ , say δ_n . Because of the paths joining δ to the vertices δ'_n , it is easy to see that the increasing sequence chosen for $f(\delta)$ must be $f(\delta_n)$, and, in particular, the supremum of the latter sequence must be $f(\delta)$. Note that the sequence $f(\gamma_n)$ is the sequence γ_n , with the first term removed. Since the sequences γ_n, δ_n are interlacing, (after each γ_n there is some δ_m and conversely), so are the sequences $f(\gamma_n), f(\delta_n)$, and hence also the sequences $\gamma_n, f(\delta_n)$. Thus the supremum of $f(\delta_n)$ is also δ , and we have $f(\delta) = \delta$. Considering once more the paths joining $f(\delta)$ to $f(\delta'_n)$, we see that each $f(\delta_n) = \delta_n$. This is impossible, since there exist m and n so that $\gamma_m < \delta_n < \gamma_{m+1}$, which means that $f(\gamma_m) = \gamma_{m+1} > \delta_n = f(\delta_n)$, contrary to the monotonicity of f . \square

Corollary 4.10 *There is a rigid graph on every infinite set.*

Proof This will follow from the technique introduced in the next section, and illustrated in Fig. 1.10. In this case, we take a rigid digraph G on a set of equal cardinality, and construct the graph $G * J$, where J is any finite connected rigid strong replacement graph, such as, say, H_k . Note that the resulting graph is connected. \square

In fact, we have the following stronger result.

Corollary 4.11 *On a set of infinite cardinality α there exists an incomparable family of 2^α rigid graphs.*

Proof Let G be a rigid graph with vertex set of cardinality α from Corollary 4.10. Since G is connected, the cardinality of its edge set is also α . Since G is rigid, no two of the 2^α possible orientations of G are isomorphic, and in fact these 2^α digraphs are an incomparable family of rigid digraphs. Now we can apply the replacement operation once more, to obtain graphs instead of digraphs. \square

4.4 The replacement operation

We now formally introduce the technique of replacing each arc of a digraph G by a copy of a fixed graph J . Various versions of this operation occur frequently in our constructions, and we see it as basic to the study of graph homomorphisms. We have illustrated the technique to prove Theorem 1.34 (Fig. 1.10). If the replacement graph J is chosen carefully, the resulting graph will have ‘the same’ homomorphisms as G . Therefore, we will be able to, for instance, construct a rigid graph (or a family of incomparable rigid graphs), from a rigid digraph (respectively a family of incomparable rigid digraphs). More generally, the technique can be applied to a binary relational system, replacing each arc of a particular colour by a particular graph (or digraph, as in Fig. 1.10), with the same goal.

Let J be a fixed graph, and j, j' two nonadjacent vertices of J . We call J a *replacement graph* with respect to the *connector vertices* j, j' . The connector vertices can be either clear from the context, or are arbitrary, and we often omit mentioning them.

For any digraph G , we denote by $G * J$ the graph obtained from G by replacing each arc $xy \in E(G)$ by an isomorphic copy J_{xy} of J , identifying x with j and y with j' ; it is assumed that all the copies J_{xy} are pairwise vertex disjoint. The graph $G * J$ is said to arise from G by the *replacement operation* with respect to J, j, j' . The replacement operation is illustrated in Fig. 4.9. Note that $G * J$ consists of the vertices of G , which we shall call the *branch vertices* of $G * J$ (shown in the figure by the larger circles), and of the vertices inside the individual copies of the replacement graph J ; we call these latter vertices the *inner vertices* of $G * J$ (shown by the smaller circles).

Suppose G, H are digraphs, and J a replacement graph (with respect to the connector vertices j, j'). If $f : G \rightarrow H$ is a homomorphism, we define a




FIG. 4.9. The replacement operation.

homomorphism $f * J$ as follows. If v is a branch vertex of $G * J$, then its image is the same as under f , i.e., $(f * J)(v) = f(v)$. If v is an inner vertex, say in the copy J_{xy} of J , then its image $(f * J)(v)$ is the corresponding vertex in the copy $J_{f(x)f(y)}$. Thus $f * J$ maps adjacent vertices of $G * J$ to adjacent vertices of $H * J$, and $f * J$ is a homomorphism $G * J \rightarrow H * J$. We define the replacement digraph J (with respect to j, j') to be *strong* if J satisfies the following condition:

- for any irreflexive digraph G and any homomorphism $f : J \rightarrow G * J$, the homomorphic image $f(J)$ is contained in some copy J_{xy} .

(Note that the definition applies to graphs J as well, and in fact we shall mostly use it for graphs.)

If we wish to faithfully represent the homomorphisms of G to H by the homomorphisms $G * J \rightarrow H * J$, we must take care to choose J , and its connector vertices j, j' , so that it is both rigid and strong.

Proposition 4.12 *If J is both rigid and strong, then for any two irreflexive digraphs G, H without isolated vertices, each homomorphism $G * J \rightarrow H * J$ is equal to $f * J$ for some homomorphism $f : G \rightarrow H$.*

Proof Consider a homomorphism $h : G * J \rightarrow H * J$. Since G has no loops and no isolated vertices, each vertex of $G * J$ belongs to some copy J_{xy} of J . (This copy is unique for all inner vertices of $G * J$.) Since J is strong, each copy J_{xy} of G maps to a copy J_{uv} of H . Since J is rigid, this mapping takes corresponding vertices identically to each other. Hence restricting h to the connector vertices defines a mapping f of the vertices of G to the vertices of H . Clearly, this mapping is a homomorphism $G \rightarrow H$, and $h = f * J$. \square

Since a rigid digraph, or any digraph in a family of incomparable digraphs, must be irreflexive and without isolated vertices, we obtain the following fact.

Proposition 4.13 *Let J be a strong replacement graph.*

- *If G is a rigid digraph, then $G * J$ is a rigid graph.*
- *If $G_i, i \in I$, is a family of incomparable digraphs, then $G_i * J, i \in I$, is a family of incomparable graphs.* □

A prototype rigid digraph that is not strong is the directed path \vec{P}_k , with respect to the connector vertices $0, k$. Indeed, when $G = \vec{P}_2$, there are many homomorphisms $\vec{P}_k \rightarrow G * \vec{P}_k$ which map \vec{P}_k to neither copy of $J = \vec{P}_k$.

However, many of the rigid digraphs in Section 4.2 are, in fact, strong. Consider, for instance, the digraph $Q(2, 3)$ from Proposition 4.4 (Fig. 4.3). Take $J = Q(2, 3)$ as a replacement digraph, with respect to the first and last vertex, and assume that H is an irreflexive digraph. Then $H * J$ is a balanced digraph, since the connectors vertices have the same level 0 in J . Moreover, $H * J$ has height four, the same as J . Thus J is strong by Proposition 1.14.

Proposition 4.14 *Each replacement digraph $Q(i, k)$ is strong, with respect to the first and last vertex.* □

We next consider the graphs $J = H_k$ from Proposition 4.5 (Fig. 4.6). Since these graphs are all triangle-connected, they are covered by the following proposition.

Proposition 4.15 *If the replacement graph J is rigid and triangle-connected, then it is strong (with respect to any pair of nonadjacent connector vertices).*

Proof Consider a homomorphism $h : J \rightarrow H * J$, for any irreflexive H . Since J is triangle-connected, so is $h(J_{xy})$. We claim that the maximal triangle-connected subgraphs of $H * J$ are the copies J_{ab} of J . Indeed, suppose a path v_1, v_2, \dots, v_p in which each v_i is adjacent to v_{i+2} (for $i \leq p-2$) was joining v_1 in J_{xy} to v_p not in J_{xy} , and suppose v_k was the last vertex on the path which is in J_{xy} . Then v_k is a branch vertex, and hence the inner vertices v_{k-1} and v_{k+1} in different copies of J cannot be adjacent. Therefore, $h(J)$ is included in some copy of J in $H * J$. □

Now we consider the graphs G_k from Proposition 4.6 (Fig. 4.7).

Proposition 4.16 *Each replacement graph G_k is strong, with respect to the connector vertices a, b .*

Proof Let $J = G_k$, and consider a homomorphism $h : J \rightarrow H * J$ for any irreflexive H . Since each vertex of J belongs to an odd cycle of length at most $2k + 5$, the same must hold for $h(J)$. However, the connector vertices a, b have distance $2k$ in J , hence the only odd cycles of length at most $2k + 5$ in $H * J$ are included in copies of J . It follows that $h(J)$ is included in some J_{xy} . □

Since each G_k or H_k is strong, Proposition 4.13 allows us to construct many rigid or incomparable graphs, from the rigid or incomparable digraphs discussed earlier.

In fact, a simple variant of the replacement operation allows us to construct rigid and incomparable graphs that are themselves triangle-connected, and hence strong.

Suppose J is a fixed graph, and $j_1j_2, j'_1j'_2$ two fixed nonadjacent edges of J . (Neither of j_1, j_2 is adjacent to neither of j'_1, j'_2 .) We call J the *edge-based replacement graph* with respect to the *connector edges* $j_1j_2, j'_1j'_2$. (We again often do not explicitly refer to the connector edges.) Suppose J is an edge-based replacement graph (with some connector edges). Given a digraph G , we again denote by $G*J$ the graph obtained from G as follows: first we replace each vertex v of G by two vertices v_1, v_2 ; then we replace each arc uv of G by a copy of J , identifying u_i with j_i and v_i with $j'_i, i = 1, 2$ (Fig. 4.10).




FIG. 4.10. The edge-based replacement operation with respect to connector edges.

Proposition 4.17 *Suppose J is a triangle-connected rigid edge-based replacement graph. Suppose further that each connector edge jj' of J has the following property \mathcal{P} . For every vertex x of J there is a path $x = v_1, v_2, \dots, v_{p-1} = j, v_p = j'$ or a path $x = v_1, v_2, \dots, v_{p-1} = j', v_p = j$, in which each v_i is adjacent to v_{i+2} (as long as $i + 2 \leq p$).*

1. *If G is a connected digraph, then $G * J$ is a triangle-connected graph,*
2. *if G is a rigid digraph, then $G * J$ is a rigid graph,*
3. *if $G_k, k \in K$, is an incomparable family of digraphs, then $G_k * J, k \in K$ is an incomparable family of graphs.* \square

Note that Propositions 4.17 and 4.15, together with Corollary 4.2, or Proposition 4.3, or Propositions 4.4 and 4.5, imply the following fact.

Corollary 4.18 *For any positive integer n , there exists an incomparable family of n rigid strong replacement graphs.* \square

This is illustrated in Fig. 4.10. We took the two incomparable rigid graphs from Fig. 4.2 (or Proposition 4.3), and our triangle-connected rigid edge-based replacement graph $J = H_1$, in which each connector edge satisfies property \mathcal{P} . The result is a pair of incomparable triangle-connected (and hence strong) rigid graphs, which can serve as replacement graphs in other constructions.

The replacement operation applies equally in the more general context of binary relational systems. If G is a binary I -system, and $\mathcal{J} = (J_i), i \in I$, a family of replacement graphs (each with respect to its own connector vertices j_i, j'_i), then the replacement operation applied to G and \mathcal{J} results in the graph $G * \mathcal{J}$ obtained from G by replacing each arc uv of colour i with a copy of the graph J_i , identifying u with j_i and v with j'_i . An example of this construction, for replacement *digraphs*, is given in Fig. 1.10. The definition of $f * \mathcal{J}$ and of strong replacement graphs for binary I -systems is analogous to the definition for digraphs. In particular we have the following result.

Theorem 4.19 *Suppose $\mathcal{J} = (J_i, i \in I)$ is an incomparable family of rigid strong replacement graphs.*

*If G, H are binary I -systems without loops (all relations are irreflexive) or isolated vertices, then each homomorphism $G * \mathcal{J} \rightarrow H * \mathcal{J}$ is equal to some $f * \mathcal{J}$, where f is a homomorphism $G \rightarrow H$.*

Proof This is proved in the same way as Proposition 4.12, noting that a copy of J_i can only be mapped identically to a copy of the same J_i , thus mimicking the action of a homomorphism $G \rightarrow H$. \square

Therefore, assertions about rigidity and incomparability of binary relational systems (such as those in Proposition 4.17) translate directly into the corresponding assertions about graphs.

Finally, we note that the replacement operations can also be defined for replacement *digraphs*, or replacement *binary relational systems*.

Figure 4.11 shows the result $G * J$ of the replacement operation on a digraph G using a replacement digraph $J = Q(2, 3)$ from Fig. 4.3 (or Proposition 4.4). We note, in particular, that $G * J$ is a balanced digraph for any digraph G .




FIG. 4.11. A digraph replacement operation.

Figure 4.12 shows the result $G * \mathcal{J}$ of the replacement operation on the binary relational system G with respect to a family of replacement binary systems J_i , each consisting of just two arcs of different colours. Here we note that the result is always a binary relational system in which all relations are irreflexive.

It is easy to see that in both these cases the replacement relational systems are rigid and strong, so that homomorphisms are preserved under the replacement operation. In other words, if G, H are binary I -systems, then any homomorphism $G * \mathcal{J} \rightarrow H * \mathcal{J}$ is equal to some $f * \mathcal{J}$, where f is a homomorphism $G \rightarrow H$.

4.5 Categories

We have seen in Chapter 1 how the structure of composition can be analysed for endomorphisms. To do a similar analysis for homomorphisms in general, we need to address several problems.

- What abstract structure (analogue of an abstract monoid) does the set of all homomorphisms amongst a set of graphs define?
- How to associate elements of such an abstract structure with concrete sets and mappings?




FIG. 4.12. A replacement operation using systems J_i .

- How to add relations to the associated sets, that ‘select’ the right mappings as homomorphisms?
- How to transform the resulting relational systems into graphs without changing the composition structure of the homomorphisms?

This section focuses on the first problem—we want to define an abstract structure which generalizes a monoid and which allows some compositions to be left undefined. (Two endomorphisms can always be composed, but two homomorphisms can only be composed if the codomain of the first one is the domain of the second one.) Such a structure is called a category. Recall the composition table of the set of homomorphisms amongst the three graphs in Fig. 1.8, in Table 1.1. One can consider such a table without knowing which homomorphisms of which graphs the entries represent (Table 4.1). After all, this is what we do when we give an abstract group, or an abstract monoid, cf. Table 1.2. It makes sense to define a category to be just such a table (satisfying the right set of axioms).

In Table 4.1, we give an example category. The reader should compare this table to the operation table of a group, or a monoid (such as in Table 1.2). A group is an associative total operation with an identity element, with respect to which every element has an inverse. A monoid is an associative total operation with an identity element. A category is an associative *partial* operation that also satisfies certain axioms. We say that a partial operation \circ is associative if $x \circ (y \circ z) = (x \circ y) \circ z$ whenever at least one side is defined. An *identity* in a partial operation \circ is an element u such that $u \circ x = x, y \circ u = y$ whenever these operations are defined. Let us write $x \in M(u, v)$ if u, v are identities and $x \circ u, v \circ x$ are defined. The remaining axioms for a category defined as a partial operation require that each x belongs to a unique $M(u, v)$ and that $x \in M(v, w), y \in M(u, v)$ imply $x \circ y \in M(u, w)$. As explained below, we will usually assume that each set $M(u, v)$ is finite.

However, it is more common, and will be more convenient for our purposes, to define a category differently. The above view emphasizes the ‘morphisms’ and

	a	b	c	d	e	f	g	h	i
a	a		c	d				h	i
b		b			e	f	g		
c	c		c	d				h	i
d	d		d	c				i	h
e		e			b	g	f		
f	f		f	g				b	e
g	g		g	f				e	b
h		h			i	c	d		
i		i			h	d	c		

Table 4.1 *An example category; a and b are the identities.*

their composition, subsuming the actual objects (associated with the identities). We shall define a category in terms of both ‘objects’ and ‘morphisms’. In the above table we have two identities— a and b , and we associate with them two objects, say A and B . The other elements of the table must then correspond to other morphisms of these objects. The domain of an element x is the object associated with the identity u for which $x \circ u$ is defined, and the codomain of x is the object associated with the identity v for which $v \circ x$ is defined. For instance, in Table 4.1, the element f can compose with b on the left and with a on the right (in other words, both $b \circ f$ and $f \circ a$ are defined in the table), thus f has domain A and codomain B . If the table corresponds to a real example, the composition of $m \circ n$ is defined if and only if the codomain of n is the domain of m . Moreover, the composition must be associative, i.e., $(m \circ n) \circ p = m \circ (n \circ p)$ whenever the domains and codomains are appropriate.

Figure 4.13 depicts the situation as a labeled digraph, where an arrow represents a morphism (with the beginning object indicating the domain and the end object indicating the codomain); such a picture tells us in a more transparent way the domains and codomains of the morphisms, and hence indicates which morphisms can be composed. Thus we can also view a category as consisting of abstract objects and morphisms satisfying certain axioms.




FIG. 4.13. The homomorphisms a, \dots, i among A, B .

A *category* \mathcal{C} consists of *objects* $A \in \mathcal{O}$, and *morphisms* $m \in \mathcal{M}$. Each morphism m is associated with a pair of objects—its *domain* $D(m)$, and its *codomain* $C(m)$. In general, both \mathcal{O} and \mathcal{M} can be arbitrary sets, but typically \mathcal{O} consists of sets (with some additional structure) and \mathcal{M} consists of mappings (compatible with the structure). Even if this is not the case (examples appear below), it is helpful to think of a morphism m with domain A and codomain B as a labeled ‘arrow’

$$A \xrightarrow{m} B.$$

We shall denote by $\mathcal{M}(A, B)$ (or sometimes just $\mathcal{C}(A, B)$) the set of all morphisms with domain A and codomain B . The set \mathcal{M} is partitioned into the sets $\mathcal{M}(A, B)$, $A, B \in \mathcal{O}$. Unless explicitly stated otherwise, *we shall assume that \mathcal{O} is a countable set and that all $\mathcal{M}(X, Y)$, $X, Y \in \mathcal{O}$, are finite sets.* (These assumptions reflect our interest in finite objects; they are not standard in category theory which typically deals with infinite objects as well as finite ones.) Note that, in general, the sets \mathcal{O} and \mathcal{M} can be infinite. If, however, \mathcal{O} and \mathcal{M} are finite sets, we say that \mathcal{C} is a *finite category*.

Additionally, there is a partial binary operation *composition* on \mathcal{M} , denoted by \circ . The composition $m \circ n$ is defined just if $C(n) = D(m)$. (It is again not necessarily the case that \circ is a composition of mappings in the usual sense, but it helpful to have this example in mind.) Therefore, for each $A, B, C \in \mathcal{O}$, the composition is a mapping $\mathcal{M}(B, C) \times \mathcal{M}(A, B) \rightarrow \mathcal{M}(A, C)$, taking the morphisms $A \xrightarrow{n} B$ and $B \xrightarrow{m} C$ to the morphism $A \xrightarrow{m \circ n} C$.

There are just two axioms that a category has to satisfy.

- The *associativity axiom* requires that

$$(m \circ n) \circ p = m \circ (n \circ p),$$

whenever at least one side is defined (i.e., for all $A \xrightarrow{p} B, B \xrightarrow{n} C, C \xrightarrow{m} D$.)

- The *identity axiom* requires that for each object A there be a particular morphism $1_A \in \mathcal{M}(A, A)$, such that

$$m \circ 1_A = m, 1_A \circ n = n,$$

whenever the compositions are defined (i.e., whenever $D(m) = C(n) = A$).

Our first example of a category is the motivating example. If \mathcal{G} is any set of graphs, then the *homomorphism category* of \mathcal{G} is the category with objects $\mathcal{O} = \mathcal{G}$, and morphisms $\mathcal{M}(G, H)$ consisting of all homomorphisms of G to H . Since we have motivated the abstract definition of a category with the properties of homomorphisms, it should be clear that the axioms are satisfied.

If \mathcal{G} is the set of *all* graphs, the resulting category—whose objects are all graphs and morphisms are all homomorphisms amongst them—will be denoted by *Graph*. There are some technicalities to be taken care of in order to keep the category *Graph* within our constraints, i.e., to make sure that the number of objects is countable. (Even though our graphs are finite, we need to make

sure that we don't take too many isomorphic copies of each graph.) Usually we can safely ignore these technical issues, but if needed, we shall assume that each graph has a vertex set which is a finite subset of the natural numbers. Then there are only countably many objects possible. It is easy to see that this implies, in particular, that we may also assume, if needed, that different graphs have different vertex sets.

Two similar categories also receive a special mention:

- *IRel* denotes the category of all binary relational *I*-systems and their homomorphisms, and
- *Digraph* denotes the category of all digraphs and their homomorphisms.

If needed, we may assume, as above, that each relational system, or digraph, has a vertex set which is a subset of natural numbers, and that different relational systems, or digraphs, have different vertex sets. Note also that *Digraph* is the same category as *IRel* where $|I| = 1$.

We could, of course, also obtain a category by taking a set of groups and all their group-homomorphisms, or a set of finite topological spaces and all their continuous maps, or a set of finite partial orders and all their monotone mappings. The language of categories is common to all these theories, and it offers a unifying perspective.

One more category shall play an important role:

- *Set* denotes the category of all finite sets and their mappings.

If needed, we may assume, as above, that the objects of *Set* are subsets of natural numbers.

In all the above categories, the objects were sets (perhaps with additional structure), and morphisms were mappings (somehow preserving the structure). There are simple categories where the morphisms are not mappings.

We give three standard examples. First, suppose $S = \{A, B, \dots\}$ is a countable set, partially ordered by \leq . We can define the *thin category* \mathcal{C}_S associated with S by setting $\mathcal{O} = S$, and setting \mathcal{M} to consist of all pairs of objects (A, B) such that $A \leq B$. The domain of (A, B) is A and the codomain is B . The composition of two morphisms is defined by $(B, C) \circ (A, B) = (A, C)$, and the identities are $1_A = (A, A)$. It is easy to see that the axioms are satisfied. Note that all the sets $\mathcal{M}(A, B)$ in a thin category have at most one element.

For our second example, consider a group Γ . It can be viewed as a category with one object, say O , by considering each element $g \in \Gamma$ to be a morphism $O \xrightarrow{g} O$. The identity element of Γ is 1_O , and the composition is the group operation.

Our third example generalizes the second, by replacing the group with a monoid. (Monoids were defined in Chapter 1; it is easy to see that monoids are precisely one-object categories.)

However, a category such as in the above examples may still be isomorphic to a category in which the objects are sets and the morphisms are mappings. It is of course well known that *every group is isomorphic to a group of permutations*

on a set, and we have seen in the introduction how to generalize this to monoids (every monoid is isomorphic to the monoid of mappings—the ‘left translations’—on a set). Even in the above example of the thin category \mathcal{C}_S associated with a partial order S , we can replace each object $A \in S$ by the set $S_A = \{X : X \leq A\}$, and each morphism (A, B) (with $A \leq B$) by the mapping $m_{(A,B)} : S_A \rightarrow S_B$ taking each $X \in S_A$ identically to itself (as $A \leq B$ implies that $S_A \subseteq S_B$). It is easy to see that the compositions and identities in this category of sets and mappings are in a one-to-one correspondence with those in the category \mathcal{C}_S .

Formally, two categories \mathcal{C}, \mathcal{D} are *isomorphic* if there exists a mapping F which bijectively assigns

- to each object A of \mathcal{C} an object $F(A)$ of \mathcal{D} , and
- to each morphism $m \in \mathcal{C}(A, B)$ a morphism $F(m) \in \mathcal{D}(F(A), F(B))$,

so that $F(m \circ n) = F(m) \circ F(n)$ whenever either side is defined.

(It follows from the definition that we have $F(1_A) = 1_{F(A)}$ for every object A .)

Let \mathcal{C}, \mathcal{D} be categories. We say that \mathcal{C} is a *subcategory* of \mathcal{D} if each object or morphism of \mathcal{C} is also an object or morphism of \mathcal{D} (with the same domains and codomains of morphisms), and the composition and the identities in \mathcal{C} are the same as in \mathcal{D} . A subcategory is called *induced* (by the set of objects of \mathcal{C}), if $\mathcal{C}(A, B) = \mathcal{D}(A, B)$, for all $A, B \in \mathcal{O}$. (Induced subcategories are sometimes called *full* subcategories.)

We shall say that a category \mathcal{C} is *represented in* a category \mathcal{D} , if \mathcal{C} is isomorphic to an induced subcategory of \mathcal{D} . Any isomorphism of \mathcal{C} to an induced subcategory of \mathcal{D} will be called a *representation* of \mathcal{C} in \mathcal{D} . It is easy to see that a composition of a representation of \mathcal{C} in \mathcal{D} and a representation of \mathcal{D} in \mathcal{E} is a representation of \mathcal{C} in \mathcal{E} .

Consider the thin category \mathcal{C} associated with a partial order S . To represent \mathcal{C} in *Graph* means to find a set \mathcal{G} of graphs and a bijective mapping associating to each $A \in S$ a graph $G_A \in \mathcal{G}$ such that $G_A \rightarrow G_B$ if and only if $A \leq B$. Thus this use of the term ‘represent’ is consistent with that in Section 3.2.

Similarly, for the one-object category corresponding to a monoid M , respectively a group Γ , to represent M or Γ in *Graph* means to find a graph G with the endomorphism monoid isomorphic to M , respectively the automorphism group isomorphic to Γ . This is also consistent with our informal use of the term ‘represent’ earlier in this chapter as well as in Chapter 1.

We have seen in the introduction how to represent each monoid in a suitable category *IRel*. In the next section we shall generalize this result, and represent each finite category in some *IRel*. This will of course include a representation of each (thin category associated with a) finite partial order (Theorem 3.5).

In the next section, we shall also show how to represent each category *IRel* (for a finite I) in the category *Graph*. Hence we will be able to conclude that each finite category (including all monoids and groups, and all finite partial orders) can be represented in the category *Graph*. We shall also prove that *Graph* can be represented in some of its induced subcategories such as that induced by three-

colourable, or k -connected, subgraphs, and hence obtain similar representation theorems for such graphs.

4.6 Representation

We now proceed with our original goal—to represent a finite category \mathcal{C} (such as the one given by Table 4.1) as the homomorphism category of a family of graphs. The process is very similar to the technique we used to represent each monoid as the endomorphism monoid of a graph (in Section 1.7). We begin by finding an isomorphism of \mathcal{C} to a subcategory of \mathbf{Set} , i.e., by associating each object of \mathcal{C} with a finite set and each morphism of \mathcal{C} with a mapping. In the case of categories one uses the following terminology.

A category \mathcal{C} is called *concrete* if it is isomorphic to a subcategory of \mathbf{Set} . (Note that we do not require that \mathcal{C} be represented in \mathbf{Set} , i.e., that it be isomorphic to an *induced* subcategory of \mathbf{Set} .)


All of our standard example categories in which objects are sets with some additional structure and morphisms are mappings preserving the structure, are concrete—one can simply ‘forget’ the additional structure, to obtain objects that are sets and morphisms that are mappings. For example \mathbf{Graph} admits an isomorphism F to a subcategory of \mathbf{Set} , by setting $F(G) = V(G)$ for every graph G , and $F(f) = f$ for every homomorphism f . (Here we make use of our assumption that different graphs have different vertex sets, so the above mapping is injective.) Moreover, we have seen three example categories where the objects are not sets and morphisms are not mappings—a group, a monoid, and a finite partial order—which were also shown to be concrete.

Thus our first step amounts to proving the following result.

Proposition 4.20 *Every finite category is concrete.*


Proof Suppose \mathcal{C} is a category with finite sets \mathcal{O}, \mathcal{M} of objects and morphisms, respectively. We shall construct an isomorphism F to a subcategory of \mathbf{Set} as follows. For each $A \in \mathcal{O}$ we define the associated set $F(A) = \cup_{X \in \mathcal{O}} \mathcal{M}(X, A)$; for each $m \in \mathcal{M}(A, B)$ we define the associated mapping $F(m)$ to be the ‘left translation’ $F(A) \rightarrow F(B)$ which takes each $p \in F(A)$ (thus $p \in \mathcal{M}(X, A)$ for some object X) to $m \circ p \in F(B)$ (which is well defined since $C(p) = D(m) = A$) (Fig. 4.14). Since \mathcal{O}, \mathcal{M} are finite, each $F(A)$ is finite. Clearly $F(1_A) = 1_{F(A)}$, and it is easy to check that $F(m \circ n) = F(m) \circ F(n)$. Finally, if $A \neq B$ then $F(A) \neq F(B)$ (for instance $1_A \notin F(B)$), and if $m \neq n$ then $F(m) \neq F(n)$. \square

Next we shall add binary relations to the sets $F(A), A \in \mathcal{O}$, from the preceding proof, which will ensure that each mapping $F(m), m \in \mathcal{M}(A, B)$, is a homomorphism of $F(A)$ to $F(B)$, and which also ensure that no other homomorphisms of $F(A)$ to $F(B)$ exist. (The added relations will ‘select’ the correct mappings.) This is again a generalization of what we have done in Section 1.7 for monoids.


 FIG. 4.14. How $F(A)$ maps to $F(B)$.

Theorem 4.21 *Every finite category can be represented in some $I\mathcal{R}el$ of binary relational systems.*

Proof Consider the mapping F defined in the preceding proof, which associates to each object A the set $F(A) = \cup_{X \in \mathcal{O}} \mathcal{M}(X, A)$, and to each morphism $m \in \mathcal{M}(A, B)$ the left translation $F(m) : F(A) \rightarrow F(B)$. Let $I = \mathcal{M}$ and let $\Phi(A)$ be the binary I -system with vertex set $F(A)$ and relations $R_m, m \in I$, where R_m consists of all pairs (a, b) , where $b = a \circ m$ (Fig. 4.15). (Note that if $m \in \mathcal{M}(X, Y)$ then $a \in \mathcal{M}(Y, A) \subseteq F(A)$ and $b \in \mathcal{M}(X, A) \subseteq F(A)$.)


 FIG. 4.15. The pair $(a, b) \in R_m$ on $F(A)$.

To complete the definition of the isomorphism Φ , we let $\Phi(m) = F(m)$, for any morphism m . It remains to verify that each mapping $F(m), m \in \mathcal{M}(A, B)$, is a homomorphism of $\Phi(A)$ to $\Phi(B)$, and that there are no other homomorphisms $\Phi(A) \rightarrow \Phi(B)$. This is easily checked, along the lines of the proof of Theorem 1.35. For instance, to see that there are no other homomorphisms, consider an arbitrary homomorphism $\pi : \Phi(A) \rightarrow \Phi(B)$. for each $m \in \mathcal{M}(X, A), X \in \mathcal{O}$,

we have $(\pi(1_A), \pi(m)) \in R_m$ (as $(1_A, m) \in R_m$ by definition), and so $\pi(m) = \pi(1_A) \circ m$ for all m . Therefore, π is $\Phi(f)$ where $f = \pi(1_A)$. \square

To prove that every finite category can be represented in *Graph*, it remains to find a representation of each *IRel* in *Graph*.

Theorem 4.22 *For any finite set I , the category $IRel$ can be represented in *Graph*.*

A corresponding result for categories of general relational systems with a fixed pattern in discussed in Exercise 10.

Proof The idea is to use the replacement operation. However, we need to be careful, as to apply Theorem 4.19 we need to have binary relational systems that are irreflexive and without isolated vertices.

We first let $I' = I \cup 0$ (assuming $0 \notin I$), and represent *IRel* in $I'Rel$ as follows. Each I -system S will be represented by the I' -system S' obtained from S by adding a new vertex v_S and all arcs $v_S v$ ($v \in V(S)$) in the added relation $R_0(S')$. Each homomorphism of an I -system S_1 to an I -system I_2 extends to a homomorphism of the corresponding I' -system S'_1 to S'_2 by sending the special vertex v_{S_1} to v_{S_2} . It is easy to check that this defines a representation of *IRel* in $I'Rel$. Moreover, it is a representation in the subcategory of $I'Rel$ induced by all connected relational systems (i.e., systems whose underlying graph is connected). It follows that each *IRel* can be represented in the subcategory of some $I'Rel$ induced by relational systems without isolated vertices.

Next we observe that each category *IRel* can be represented in the subcategory some $I'Rel$ induced by irreflexive relational systems. Indeed, it suffices to let $I' = I \times \{0, 1\}$, and use the replacement relational systems from Fig. 4.12. (Each arc of relation R_i is replaced by a path consisting of an arc of relation $R_{(i,0)}$ followed by an arc of relation $R_{(i,1)}$.)

In conclusion, we have represented each *IRel* in the subcategory of some $I'Rel$, induced by irreflexive relational systems without isolated vertices.

Suppose that I' has n elements. Corollary 4.18 ensures that there exists a set of n incomparable rigid strong replacement graphs, and Theorem 4.19 implies that the homomorphism category of the resulting graphs is isomorphic to the homomorphism category of the binary I -systems. \square

Thus we finally obtain our main representation theorem, asserting that each finite category is isomorphic to the homomorphism category of some set of graphs.

Corollary 4.23 *Any finite category can be represented in *Graph*.* \square

This result implies that any group, monoid, or finite partial order, can be represented by the automorphism group, endomorphism monoid, or homomorphism order of a graph or family of graphs.

Many other categories can be represented in *Graph*, including such basic categories as the category of all finite topological spaces and their continuous

maps, the category of all groups and their group homomorphisms, etc. In fact, the following deep result has been shown to hold.

Theorem 4.24 [295] *Every concrete category can be represented in \mathcal{Graph} .*

Therefore, concrete categories that are not finite also admit a representation in \mathcal{Graph} ; this includes, say, the category of finite topological spaces and their continuous maps, or the category of groups and their group homomorphisms. It is important, however, to remember here that we have assumed that \mathcal{O} is a countable set and that all $\mathcal{M}(X, Y)$, $X, Y \in \mathcal{O}$, are finite sets. These assumptions are crucial for this result (otherwise the situation is complicated by a set theoretic assumption). Thus, we also obtain Theorem 3.3 as a corollary (in fact, strengthened to a representation in \mathcal{C}_S , rather than just \mathcal{C}). The proof of Theorem 4.24 is beyond the scope of this book and we direct the interested reader to [295] where a wealth of other results on representations can also be found.

We now turn to representing \mathcal{Graph} in other categories. Observe that if \mathcal{Graph} can be represented in \mathcal{C} , then all the representation results we had about \mathcal{Graph} also apply to \mathcal{C} , including the representability of any group, monoid, countable partial order, or category.

Proposition 4.25 *The category \mathcal{Graph} can be represented in each of the following induced subcategories of $\mathcal{Digraph}$ ($\ell \geq 1, g \geq 3$, are any integers).*

1. *The category of all three-colourable graphs.*
2. *The category of all ℓ -connected graphs.*
3. *The category of all graphs of girth at least g .*
4. *The category of all balanced digraphs.*

Proof Let $\mathcal{Digraph}^*$ denote the subcategory of $\mathcal{Digraph}$ induced by connected irreflexive digraphs, and let \mathcal{Graph}^* denote the subcategory of \mathcal{Graph} induced by connected graphs. The proof of Theorem 4.22 actually shows that each \mathcal{IRel} can be represented in \mathcal{Graph}^* . This includes the category $\mathcal{Digraph}$, and hence also its induced subcategory \mathcal{Graph} . Thus we have a representation of \mathcal{Graph} in \mathcal{Graph}^* . Of course, \mathcal{Graph}^* is an induced subcategory of $\mathcal{Digraph}^*$, and hence we also have a representation of \mathcal{Graph} in $\mathcal{Digraph}^*$. It remains to find representations of $\mathcal{Digraph}^*$ in the categories 1–4.

The category $\mathcal{Digraph}^*$ is represented in the category of graphs with girth at least $2k + 5$, by using the replacement operation with $J = G_k$, the graph in Fig. 4.7, with connector vertices a and b . Proposition 4.6 ensures that each G_k is rigid; Proposition 4.16 ensures that it is strong. The girth of each replaced graph is at least $2k + 5$ because the girth of G_k is $2k + 5$, and the distance in G_k between a and b is $2k + 2$. This proves 3.

In fact, it also proves 1, if we recall our observation, following Proposition 4.6, that $J = G_k$ admits a three-colouring in which a, b have the same colour. This observation implies that each $H * J$ is three-colourable. To prove 2 for $\ell = 2$, we can use the edge-based replacement graphs H_k as in Fig. 4.10. For higher connectivity, we define a more general K_k -based replacement graph $H(k, \ell)$ analogous

to H_k , cf. Exercise 6. (This construction also yields graphs $H * J$ with arbitrary chromatic number $k > 3$.)

Finally, for 4 we use the strong rigid replacement graph $J = Q(2, 3)$, with connector vertices being the first and the last vertex of the path, from Proposition 4.4, (as in Fig. 4.3), cf. Proposition 4.14. \square

Corollary 4.26 *Any finite category (and hence any group, monoid, or finite partial order) can be represented in any of the categories 1–4.* \square

The corollary yields the following facts.

- *Each group is isomorphic to the automorphism group of a three-colourable (respectively k -connected, etc.) graph G .*
- *Each monoid is isomorphic to the endomorphism monoid of a three-colourable (respectively k -connected, etc.) graph G .*
- *Each finite partial order is isomorphic to the homomorphism order of a set of three-colourable (respectively k -connected, etc.) graphs.*

We close with one more example illustrating the power of the replacement operation.

Proposition 4.27 *Let G be a graph, and let \mathcal{C}_G be the subcategory of \mathbf{Graph} induced by all graphs containing an induced subgraph isomorphic to G . Then \mathbf{Graph} can be represented in \mathcal{C}_G .*

Corollary 4.28 *Every graph is an induced subgraph of a rigid graph.*

Proof We first prove the corollary. Given a graph G with n vertices, we take a family of incomparable triangle-connected rigid graphs $J_i, i = 1, 2, \dots, n$, with the property that the chromatic number of each J_i is greater than $\chi(G)$. Such families can be constructed by generalizing the graphs H_k from Fig. 4.6, as outlined in Exercise 6. For each vertex v_i of G , we identify one vertex of J_i with the vertex v_i . It is easy to see that the resulting graph G^* is rigid. (Note that no J_i can be mapped entirely into G since it has greater chromatic number.) It is easy to see that G^* is a strong replacement graph as well, whence the proposition also follows. \square

4.7 A combinatorial obstacle to representation

We have seen that many categories admit a representation in \mathbf{Graph} . One may wonder whether there are *any* restrictions on which categories can be represented in \mathbf{Graph} (or a similarly rich category such as \mathbf{IRel}).

It turns out that there is a nontrivial necessary condition for such a representation to exist. The obstacle turns out to be the first step of our representation process, namely finding an isomorphic subcategory of \mathbf{Set} . Recall that a category is concrete (Section 4.6) if it is isomorphic to a subcategory (not necessarily induced) of \mathbf{Set} . Since the composition of an isomorphism to an induced subcategory and an isomorphism to a subcategory must be an isomorphism to a subcategory, we have the following observation.

Proposition 4.29 *If a category \mathcal{C} is isomorphic to an induced subcategory of $\mathcal{G}raph$ then \mathcal{C} is concrete.* \square

In view of Theorem 4.24, we may ask *is every category concrete?* We shall see that this is an interesting question with a nontrivial answer.

For every two objects A, B , of \mathcal{C} , denote by $F(A, B)$ the set of all pairs (m, n) of morphisms of \mathcal{C} which have $D(m) = D(n)$ and $C(m) = A, C(n) = B$. Each $(m, n) \in F(A, B)$ is called a *fork* over A, B .




FIG. 4.16. Equivalent forks.

On each set $F(A, B)$ we define an equivalence \sim as follows.

Two forks $(m, n), (m', n')$ over A, B are equivalent, $(m, n) \sim (m', n')$, provided $f \circ m = g \circ n$ if and only if $f \circ m' = g \circ n'$, for any two morphisms f, g with $D(f) = A, D(g) = B$ and $C(f) = C(g)$ (Fig. 4.16).

It is easy to see that \sim is an equivalence on the set of forks over A, B .

We have the following necessary condition for a category to be concrete. Let us say that a category \mathcal{C} satisfies *Isbell's condition*, if for each pair A, B of objects, the equivalence \sim has finite index, i.e., has only finitely many equivalence classes.

Theorem 4.30 *A category \mathcal{C} is concrete if and only if it satisfies Isbell's condition.*

Proof Assume first that \mathcal{C} is concrete; we may assume without loss of generality (by identifying the corresponding objects and morphisms under an isomorphism of \mathcal{C} to a subcategory of $\mathcal{S}et$) that \mathcal{C} is a category of (some) sets and of (some) mapping between them. Thus a fork with respect to objects (sets) A, B , is a pair of morphisms (mappings) (m, n) with $D(m) = D(n)$ and $C(m) = A, C(n) = B$. Let $D(m) = X$; since both m and n are mappings of X , we may define a relation $S(m, n) \subseteq A \times B$ by

$$S(m, n) = \{(m(x), n(x)) : x \in X\}.$$

We now claim that the number of inequivalent forks with respect to A, B is bounded. Indeed, if two forks $(m, n), (m', n')$ with respect to A, B define the same relation S , i.e., if $S(m, n) = S(m', n')$, then they must be equivalent, $(m, n) \sim (m', n')$. (Note that $f(m(x)) = g(n(x))$ if and only if $f(m'(x')) = g(n'(x'))$)

for some x' satisfying $(m(x), n(x)) = (m'(x'), n'(x'))$.) Hence the number of equivalence classes of \sim on $F(A, B)$ is at most $2^{|A| \cdot |B|}$.

Assume next that \mathcal{C} satisfies Isbell's condition, i.e., that on each set of forks $F(A, B)$, the equivalence \sim has finite index. This means that we can choose, for each pair of objects A, B a finite set $R(A, B) \subseteq F(A, B)$ of 'representative forks', such that each fork $(m, n) \in F(A, B)$ is equivalent to exactly one fork $r(m, n) = (m', n') \in R(A, B)$. (Thus two forks are equivalent if and only if their representatives are equal.) We may want to try and imitate the proof of Proposition 4.20, and define an isomorphism to a subcategory of \mathcal{Set} by associating with each object A the union of $R(A, X)$ over all objects X . The first problem to overcome is that even though each set $R(A, X)$ is finite, we may have infinitely many objects X .

At this point, we invoke the fact that the set of objects of \mathcal{C} is countable; hence the objects can be enumerated as A_1, A_2, \dots , and we shall write $A < B$ if A precedes B in this enumeration, i.e., if $A = A_i, B = A_j, i < j$.

We say that a fork $(m, n) \in F(A, X)$ is *minimal* if there is no object $Y, Y < X$, such that $m = p \circ q$, for $D(p) = C(q) = Y$. In other words, in a minimal fork (m, n) , the *first morphism*, m , cannot be factored into two morphisms going through an object that precedes the *second object*, X , in the enumeration. Let $M(A, X)$ denote the set of all minimal forks in $F(A, X)$. We have made some progress, as there are now only finitely many objects X with $M(A, X) \neq \emptyset$, since $M(A, X) = \emptyset$ when $X > A$ (any morphism $\xrightarrow{m} A$ factors into $1_A \circ m$ through $A, A < X$). However, we cannot just take the usual representatives of $M(A, X)$, as minimality may fail to be invariant under \sim .

Therefore, we define a new equivalence between forks. We say that forks $(m, n), (m', n')$ over A, X in \mathcal{C} are *strongly equivalent*, $(m, n) \simeq (m', n')$, provided $(m, n) \sim (m', n')$, and moreover for each fork (m, q) over A, Y with $Y < X$, there is a fork (m', q') over A, Y , and for each fork (m', q') over A, Y with $Y < X$, there is a fork (m, q) over A, Y , such that $(m, q) \sim (m', q')$. It is easy to verify that strong equivalence \simeq is indeed an equivalence relation, and is a refinement of the equivalence \sim , on any $F(A, X)$.

Claim *If $(m, n) \simeq (m', n')$ and $(m, n) \in M(A, X)$ then also $(m', n') \in M(A, X)$. If (m', n') is not minimal, then $m' = p' \circ q'$, where $C(q') = D(p') = Y < X$. Since (m', q') is a fork over A, Y , there is a fork (m, q) over A, Y such that $(m, q) \sim (m', q')$. Since $1_A \circ m' = p' \circ q'$, we must have $1_A \circ m = p' \circ q$, i.e., $m = p' \circ q$ also factors through $Y < X$, contradicting the minimality of (m, n) . This proves the Claim.*

Suppose (m, n) is a fork over A, X , and $X = A_k$. We define the k -tuple of sets $R(m, n) = (R_1, R_2, \dots, R_k)$, where $R_k = \{r(m, n)\}$ and each $R_j, j < k$, is the set of all $r(m, q)$ where (m, q) is a fork over A, A_j . Isbell's condition ensures that each R_j is a finite set. Moreover, it is easy to see that $(m, n) \simeq (m', n')$ if and only if $R(m, n) = R(m', n')$.

Hence the above Claim implies that it is meaningful to speak of k -tuples

$R(m, n)$ where (m, n) is a minimal fork over a pair of sets (since it is a property of the k -tuple $R(m, n)$ and does not depend on which (m, n) was used to define $R(m, n)$).

Now we can define an isomorphism F of \mathcal{C} to a subcategory of \mathcal{Set} . To each object A of \mathcal{C} we assign the following finite set $F(A)$:

$$F(A) = \bigcup_{X \in \mathcal{O}} \{R(m, n) : (m, n) \in M(A, X)\} \cup \{O_A\}.$$

(The element O_A is different for every object A , and plays an auxiliary role.) To each morphism $A \xrightarrow{p} B$ of \mathcal{C} we assign the following mapping $F(p) = f : F(A) \rightarrow F(B)$:

- $f(R(m, n)) = R(p \circ m, n)$, if (m, n) is a minimal fork over A, X , as long as $(p \circ m, n)$ is also a *minimal* fork over B, X ; and
- $f(R(m, n)) = 0_B$ if $(p \circ m, n)$ is not minimal; and
- $f(O_A) = O_B$.

It remains to verify that $F(p)$ is well defined (independent of the choice of (m, n)), and that F is an isomorphism of \mathcal{C} to a subcategory of \mathcal{Set} . These are routine verifications and we leave them as an exercise. \square

It is not difficult to construct examples of categories that do not satisfy Isbell's Condition, and hence are not concrete, Exercise 4.

4.8 Some categories are not rich enough

The conclusions of the previous two sections can be summarized as follows. Any concrete category can be represented in a sufficiently rich category, such as the category of graphs, or some of its subcategories including the category of three-colourable graphs, k -connected graphs, and so on. (Here, and also below, k is any positive integer).

Here are some similar looking categories that are *not* rich enough. Not every concrete category can be represented by

- the category of bipartite graphs
- the category of graphs with degrees at most k
- the category of planar graphs.

In these cases, one cannot even represent a one-object category, i.e., a monoid. However, it can be shown that in the first two categories every group M can be represented, i.e., that every group is isomorphic to the automorphism group of a graph that is bipartite, or has degrees bounded by a constant k , with $k \geq 3$ [165, 305]. This is not, however, the case for planar graphs [10]. On the other hand, *every finite partial order is isomorphic to the homomorphism order of a set of planar graphs with degrees at most three* (Exercise 10 in Chapter 3).

The first bullet is easy to see.

Proposition 4.31 *Let M be a group with more than one element. Then there is no bipartite graph G such that $\text{END}(G)$ is isomorphic to M .*

Proof The group M cannot be represented as the endomorphism monoid of a bipartite graph, since a nontrivial bipartite graph always admits a noninvertible endomorphism (for instance, a retraction to any edge). \square

The second bullet is a consequence of the following theorem.

Theorem 4.32 *For every positive integer k there exists a monoid M_k such that each graph G with $\text{END}(G)$ isomorphic to M_k contains a vertex of degree at least k .*

In the proof we shall use the concept of *extension*.

Let X be a set and Π a group of permutations on X . If $\iota : X \rightarrow X'$ is a bijective mapping, we denote by $\iota(\Pi)$ the set of permutations on X' corresponding to Π , i.e., the set of all permutations $\iota \circ \pi \circ \iota^{-1}$ where $\pi \in \Pi$.

Let G be a graph such that $X \subseteq V(G)$. We say that $\text{AUT}(G)$ *extends* Π if the operation of restricting elements of $\text{AUT}(G)$ to X is a bijection between $\text{AUT}(G)$ and Π . (In other words, the restriction of each automorphism of G to the set X is a permutation in Π , and each permutation in Π is the restriction of a unique automorphism of G .) More generally, we say that the automorphism group of an arbitrary graph G extends Π via a bijection ι of X to $X' \subseteq V(G)$, if $\text{AUT}(G)$ extends the group $\iota(\Pi)$ acting on X' . Note that this means that for each $\alpha \in \Pi$ there exists a unique $\beta \in \text{AUT}(G)$ with $\iota \circ \alpha = \beta \circ \iota$, and for each $\beta \in \text{AUT}(G)$ there exists a unique $\alpha \in \Pi$ with the same property.

Given a group Π of permutations on a set X , one may ask whether or not there exists a graph G with $V(G) = X$ such that $\text{AUT}(G) = \Pi$. This is a stronger requirement than just asking for a graph whose automorphism group is isomorphic to the group Π , and the answer is, in general, negative. However, we shall show that any group of permutations can be extended, in the following very strong sense.

Lemma 4.33 *Let Π be a group of permutations on a set X . There exists a monoid M such that for any graph G with $\text{END}(G) \sim M$ and at least three vertices, the automorphism group $\text{AUT}(G)$ extends Π .*

Proof Suppose Π is a group of permutations on X , and let M consist of Π together with all mappings of $f : X \rightarrow X$ for which $|f(X)| \leq 2$. Note that M contains the constant maps $c_x : X \rightarrow \{x\}$, for all $x \in X$. We claim that this monoid M (viewed abstractly, by its composition table) has the properties claimed in the lemma. Indeed, let G be any graph with at least three vertices and $\text{END}(G) \sim M$, and let ϕ be an isomorphism of M to $\text{END}(G)$. Each permutation $\pi \in \Pi$ yields an invertible $\phi(\pi) \in \text{END}(G)$, i.e., an automorphism of G . Suppose $x_0 \neq x_1$ are two fixed distinct elements of X . Then $c_{x_0} \neq c_{x_1}$ in M , and hence $\phi(x_0) \neq \phi(x_1)$ in $\text{END}(G)$. Let v be a vertex of G such that the endomorphisms $\phi(x_0)$ and $\phi(x_1)$ differ on v . We now define a bijection ι between X and $X' \subseteq V(G)$: Let $\iota(x) = \phi(c_x)(v)$.

We claim that ι is an injective mapping of X to $V(G)$. Indeed, assume that $x \neq y$ are elements of X , and consider the mapping $\alpha : X \rightarrow X$ with $\alpha(x) = x_0$ and $\alpha(t) = x_1$ for all $t \neq x$. (The vertices x_0 and x_1 are specified above.) Note that $\alpha \in M$, and

$$\begin{aligned}\phi(\alpha)(\iota(x)) &= \phi(\alpha)(\phi(c_x)(v)) = (\phi(\alpha) \circ \phi(c_x))(v) = \phi(\alpha \circ c_x)(v) = \phi(c_{x_0})(v) \\ &\neq \phi(c_{x_1})(v) = \phi(\alpha \circ c_y)(v) = \cdots = \phi(\alpha)(\iota(y)).\end{aligned}$$

Therefore we must have $\iota(x) \neq \iota(y)$.

To complete the proof of the lemma, we note that, for any $\beta \in M$, we have

$$\phi(\beta)(\iota(x)) = \phi(\beta)(\phi(c_x)(v)) = \phi(\beta \circ c_x)(v) = \phi(c_{\beta(x)})(v) = \iota(\beta(x)).$$

Since G has at least three vertices, each automorphism of G is equal to some $\phi(\beta)$ where $\beta \in \Pi$. \square

We now return to the proof of Theorem 4.32; we will apply the lemma to a particular group Π_p of permutations. Let p be a prime, and let Z_p denote the cyclic group on $0, 1, 2, \dots, p-1$. We take the set $X = Z_p \times \{1, 2, 3\}$ and define on X the permutations $\pi_{i,j}$, for $i, j \in Z_p$, as follows:

- $\pi_{i,j}(k, 1) = (k + i, 1)$
- $\pi_{i,j}(k, 2) = (k + j, 2)$
- $\pi_{i,j}(k, 3) = (k + i + j, 3)$.

We let X_i denote $Z_p \times \{i\}$; thus $X = X_0 \cup X_1 \cup X_2$.

It is easy to see that the permutations $\pi_{i,j}$ form a group, Π_p , which acts on X , and is isomorphic to the abstract group Z_p^2 . We also observe that, for each element $x \in X$, the stabilizer of x has size p , i.e.,

$$|\{\pi_{i,j} \in \Pi_p : \pi_{i,j}(x) = x\}| = p.$$

Let M be the monoid from the lemma, applied to this group of permutations Π_p . Without loss of generality, we may assume that $V(G)$ contains X , i.e., that the bijection (from the definition of extension) is the identity, and hence that the restrictions of the automorphisms of G to X are precisely the permutations $\pi_{i,j}$. It follows that $|\text{AUT}(G)| = p^2$.

We partition the vertices of G into three sets V_0, V_1, V_2 according to the size of their stabilizer. Specifically, for $i = 0, 1, 2$, let V_i consist of all the vertices v of G for which

$$|\{f \in \text{AUT}(G) : f(v) = v\}| = p^i.$$

(The size of the stabilizer must divide the size of the group; also recall that p is a prime.) The vertices in V_2 are fixed by every automorphism of G . The vertices in V_0 are only fixed by the identity automorphism of G . Each vertex in V_1 is fixed by precisely p automorphisms of G ; according to the above observation, this set includes all vertices of X .

If there is in G an edge joining a vertex v of V_2 to a vertex u of V_1 , then the vertex v must have degree at least p , since there are p^2 automorphisms of G that fix v and exactly p of them map u to any fixed neighbour of v (including u). A similar argument establishes that v has degree at least p if a vertex in V_i is joined to a vertex in V_j with $i \neq j$.

Let a and b be distinct vertices of V_1 , and let A_a, A_b respectively be the stabilizers of a and b . We first observe then the sets A_a, A_b are either equal or intersect in the identity permutation. Indeed, the number of automorphisms of G fixing both a and b must divide p ; hence either all automorphisms fixing one also fix the other, or only the identity fixes both. It now follows that G has a vertex of degree at least p if there exist two adjacent vertices $a, b \in V_1$ with $A_a \neq A_b$.

Thus it remains to consider the case when G has components $C_i, i \in I$, each C_i lying entirely in V_0, V_1 , or V_2 . Moreover, we may assume that no C_i contains two vertices from distinct sets X_j , as such vertices have distinct stabilizers and hence there would need to be a vertex of degree at least p by the same argument as above. Since each component of G can be mapped independently by an automorphism, there would have to exist an automorphism of G which fixes X_1 and X_2 but not X_3 , contrary to the fact that $\text{AUT}(G)$ extends Π . Therefore, this case cannot occur. This proves the theorem. \square

It turns out that, more generally, for every graph H there exists a monoid M_H such that each graph G with $\text{END}(G)$ isomorphic to M_H contains a subdivision of H . Thus the category of graphs that do not contain a subdivision of H is also not rich enough. This excludes the category of planar graphs (our third bullet), or graphs embeddable on any fixed surface.

4.9 Remarks

The study of graphs in the context of category theory began in the early sixties, with pioneering work of Z. Hedrlín, A. Pultr, and, independently, G. Sabidussi. It was part of an intense attempt to develop a theory of ‘general mathematical structures’ in the framework of algebra and category theory [233, 247]. The combinatorial side of this development came to be called the Prague School of Category Theory, and is credited with the solution of several basic problems, including the proof, by L. Kučera and Z. Hedrlín, of Theorem 4.24. Prague School remains active to this day; a nice account of the main developments up to 1980 can be found in the book of A. Pultr and V. Trnková [295]. Other useful sources include the surveys [155, 157, 256, 257, 259]. Our discussion here is much simplified by focusing on finite objects, i.e., by working in finite set theory. Many of the techniques and constructions in this case mirror the general case; however, in some cases, set-theoretic complications do arise, which we have been able to avoid. We did illustrate some of the elegance of the infinite case by our one ‘excursion to infinity’, Section 4.3. Theorem 4.9 was first proved in [334]; we followed a simpler proof from [260]. Corollary 4.11 answered a question asked by S. Ulam, who seems to have been the first to consider families of incomparable graphs,

i.e., antichains in the partial order \mathcal{C}_S . Isbell's condition also first appeared in the context of infinite set theory in [186], where its necessity was discovered. Its sufficiency was first proved by P. Freyd [111]; the simpler proof presented here is due to P. Vínárek [331], and is the first proof that also applies to finite set theory. The importance of rigid graphs was first recognized by Z. Hedrlín and A. Pultr in connection with representations of categories [147, 148]; our examples use some of their constructions, supplemented from [164, 63, 149]. The smallest rigid graph is from [149], cf. also [11, 12, 145, 152, 155, 157, 164, 165, 242, 272] for other constructions of rigid graphs and digraphs. The replacement operation, named the *šíp product* (or *arrow construction*) by E. Mendelsohn [242], goes back to [112] and [72]. Theorem 4.7 is an unpublished result of P. Erdős; our proof is taken from [204]. (A similar result on asymmetric graphs appeared in [92].) Corollary 4.28 is from [63]; Theorem 4.32 is based on a result of L. Babai and A. Pultr [13], who proved a more general result on representability by proper minor-closed families of graphs.

The example in Exercise 4 is from [186]. Exercise 8 is due to Z. Hedrlín. Exercise 10 is discussed in [295] (Theorem 5.3). Exercise 11 is based on [275]; it has been conjectured that there are only finitely many minimal asymmetric graphs, and no nontrivial critical rigid oriented graphs [275]. Exercise 20 deals with graphs that have been investigated by I. Rosenberg [303].

4.10 Exercises

1. Prove that in any representation of the category from Table 4.1 in *Digraph*, the homomorphism c must be a retraction, e must be an automorphism, h and i must be injective, and f and g surjective homomorphisms.
2. Prove that asymptotically almost every graph G is not only rigid, but also has the property that any G' obtained from G by changing (i.e., adding or deleting) at most k edges, is also rigid (k is fixed).
3. Suppose G is a rigid graph with m edges.
 - (a) Prove that each orientation \vec{G} of G is a rigid digraph.
 - (b) Prove the set of all 2^m different orientations of G is an incomparable family of digraphs.
 - (c) Construct from this family a large incomparable family of rigid graphs.
4. Prove that the following category \mathcal{C} is not concrete. \mathcal{C} has two special objects X and Y , and two families of objects A_n, B_n with n ranging over all positive integers. \mathcal{C} has identity morphisms for all objects, and there are morphisms α_n with domain A_n and codomain X and γ_n with domain A_n and codomain Y , as well as morphisms β_n with domain X and codomain B_n and δ_n with domain Y and codomain B_n . For each positive integer n , there is a unique morphism with domain A_n and codomain B_n that is simultaneously the composition of α_n, β_n and of γ_n, δ_n . However, for each positive integers $m \neq n$ there are two morphisms with domain A_m and

codomain B_n , one being the composition of α_m, β_n and the other being the composition of γ_m and δ_n .

5. Prove that *Digraph* is isomorphic to an induced subcategory of *Graph* and conversely, that *Graph* is isomorphic to an induced subcategory of *Digraph*.
6. Let $1 \leq \ell \leq k$ and consider the following graph $H(k, \ell)$: the vertices are $1, 2, \dots, 3k + 7$, and the edges are $1(3k + 7), 1(5 + \ell)$ and all ij with $1 \leq |i - j| \leq k + 1$. Prove that
 - $H(k, 1)$ is isomorphic to H_k .
 - Each $H(k, \ell)$ is rigid.
 - There is no homomorphism $H(k, \ell) \rightarrow H(k, \ell')$ with $\ell \neq \ell'$.
 - Each $H(k, \ell)$ is $(2 + k)$ -clique-connected, i.e., any two of its vertices are joined by a path v_1, v_2, \dots, v_p in which each v_i is adjacent to v_{i+k+1} (as long as $i + k + 1 \leq p$).
 - Each homomorphic image of every $H(k, \ell)$ is two-connected.
7. Construct a rigid series-parallel graph. (A graph is *series-parallel* if it has no K_4 minor.)
8. Given sets A, B and mappings $f_1, f_2, \dots, f_q : A \rightarrow B$, construct graphs G, H with $A \subseteq V(G), B \subseteq V(H)$ such that there are exactly q homomorphisms $G \rightarrow H$, namely F_1, F_2, \dots, F_q , where each F_i extends the corresponding f_i . (Hint: Construct a two-object category to include all f_i and all constant mappings on A and B ; then modify the representation from the proof of Theorem 4.21 to focus on the maps f_i .)
9. (Infinite graphs) Construct an incomparable family of 2^ω countable graphs with maximum degree three.
10. Let \mathcal{P} be a fixed pattern (as defined in Section 1.8), consisting of a finite set I and positive integers $k_i, i \in I$. The *category of general relational systems* $\mathcal{P}Rel$ has as objects all general relational systems with the pattern \mathcal{P} and as morphisms all homomorphisms amongst such systems. Prove that for any pattern \mathcal{P} , the category $\mathcal{P}Rel$ can be represented in *Digraph* (and hence also in *Graph*).
11. [275] A *minimal asymmetric graph* is an asymmetric graph such that each induced subgraph admits a nontrivial automorphism. A *critical rigid digraph* is a rigid digraph G such that each vertex-removed subgraph $G - v$ has a nontrivial endomorphism. Prove that there are only finitely many minimal asymmetric trees, and no nontrivial critical rigid oriented trees.
12. [166] Let G be a graph and H an orientation of G . Show that the automorphism group of H is a subgroup of the automorphism group of G . (Easy.) Show that for every group Γ and subgroup Θ there exists a graph G and an orientation H of G such that the automorphism group of G is isomorphic to Γ and the automorphism group of H is isomorphic to Θ .
13. [64] Let G be a graph and c a colouring (not necessarily proper) of its vertices. Show that the automorphisms of G that preserve c (in the sense

that the image of any vertex has the same colour as the vertex itself) form a normal subgroup of the automorphism group of G . (Easy.) Prove that for every group Γ and normal subgroup Θ there exists a graph G and a vertex colouring c of G such that the group of all automorphisms of G is isomorphic to Γ and the group of all colour-preserving automorphisms of G is isomorphic to Θ .

14. [145] Prove that for any two monoids M and M' there exists a graph G with a subgraph G' such that the endomorphism monoid of G is isomorphic to M and the endomorphism monoid of G' is isomorphic to M' .
Show that this is no longer true if we insist that G' is obtained from G by deleting one vertex.
15. [1] Prove that any connected nonbipartite graph is a homomorphic image of a rigid graph.
16. [155] If H is a subgraph of G , we may view homomorphisms of G to H as having domain G and codomain G . This way, we may compose any two homomorphisms of G to H .
Prove that in this sense $\text{HOM}(G, H)$ forms a semigroup under composition. Prove that for any semigroup S there exists a graph G and a vertex g such that the semigroup $\text{HOM}(G, G - g)$ is isomorphic to S .
17. [148, 149] Prove that the graph in Fig. 4.4 is rigid, and that each nontrivial rigid graph has at least as many vertices and at least as many edges. Do the same for the asymmetric graph and asymmetric tree in Fig. 4.5.
18. [164] A k -uniform hypergraph can be viewed as a relational system with one k -ary relation R which has the property that for each $(a_1, a_2, \dots, a_k) \in R$ and any permutation p , the k -tuple $(p(a_1), p(a_2), \dots, p(a_k))$ is also in R . (Homomorphisms of hypergraphs preserve hyperedges.)
Prove that, for any $k \geq 2$, every finite monoid is isomorphic to the endomorphism monoid of a k -uniform hypergraph.
19. [243] A hypergraph is *intersecting* if any two edges have a common vertex. Prove that there are only finitely many nonisomorphic intersecting three-uniform hypergraphs that are cores.
20. [303] Show that asymptotically almost all graphs are both projective and rigid.

TESTING FOR THE EXISTENCE OF HOMOMORPHISMS

In this chapter, we return to the problem of existence of homomorphisms, and investigate it from an algorithmic perspective. We also expand our focus somewhat, and include list homomorphisms and other homomorphism-like partitions. The examples in Section 1.8 illustrate that we frequently encounter situations where we seek homomorphisms to one fixed target graph H , from many different input graphs G . This is the context for the present chapter.

5.1 The H -colouring problem

Let H be a fixed digraph. The *homomorphism problem for H* asks whether or not an input digraph G admits a homomorphism to H . Recall that a homomorphism of G to H is also called an H -colouring of G . Thus the homomorphism problem for H will also be called the *H -colouring problem*, and denoted H -COL. Proposition 1.10 states that G admits an H -colouring if and only if there is a partition of $V(G)$ into sets $S_x, x \in V(H)$, such that S_x is independent when xx is not a loop of H , and there are no arcs from S_x to S_y when $xy \notin E(H)$. Hence the H -colouring problem asks whether or not such a partition $S_x, x \in V(H)$, exists. The H -colouring problem can be analogously stated for any general relational system H , as we have done in Section 1.8. Recall that in that case the problem is also known as the *constraint satisfaction problem (or CSP) with a template H* , and denoted either as H -CSP (if we wish to emphasize that H is a general relational system), or simply as H -COL.

Recall that two digraphs are homomorphically equivalent if each admits a homomorphism to the other. Clearly two homomorphically equivalent digraphs result in the same homomorphism problem. Specifically, if H, H' are homomorphically equivalent, then a digraph G is H -colourable if and only if it is H' -colourable. Also recall from Corollary 1.32 that every digraph is homomorphically equivalent to a unique smallest subgraph called the core. Thus it will suffice to consider the H -colouring problems when the fixed digraph H is a core. Of course, the same remarks apply to any fixed relational system H .

Some H -colouring problems are easy to solve. Before Proposition 1.15, we gave a simple algorithm for solving the H -colouring problem when H is the directed cycle \vec{C}_k . We begin by restating this algorithm in the format we shall use for presenting algorithms.

Algorithm 1**Input:** A connected digraph G .**Task:** Find a \vec{C}_k -colouring of G , if one exists.**Action:** Choose a vertex $v \in V(G)$ and set $f(v) = 0$. Then whenever a vertex x has been assigned an image $f(x) = i$, set $f(y) = i - 1$ for all vertices y that are inneighbours of x , and set $f(z) = i + 1$ for all vertices z that are outneighbours of x . (Addition is taken modulo k .)

The correctness of the algorithm was shown in Corollary 1.16. We have described similar algorithms for homomorphisms to directed paths \vec{P}_k (Exercise 14 in Chapter 1), and transitive tournaments \vec{T}_k (Proposition 1.20). These simple algorithms were associated with duality theorems. For instance, we have proved (Corollary 1.18) that

- $G \not\rightarrow \vec{C}_k$ if and only if $C \rightarrow G$ for some oriented cycle C of net length not divisible by k .

In the case of *graphs*, there are few polynomial time algorithms and duality results. There is the trivial problem of K_1 -colourability, where the duality states $G \not\rightarrow K_1$ if and only if $K_2 \rightarrow G$; the corresponding algorithm solves the K_1 -colouring problem by checking whether or not G has any edges. There is also the easy problem of K_2 -colourability, which is the usual problem of two-colourability. As noted after Corollary 1.19, this case can be viewed as a restriction of \vec{C}_2 -colourability to undirected graphs, and hence solved by Algorithm 1. Therefore, we have polynomial time algorithms (and duality results) for all graphs homomorphically equivalent to K_1 or K_2 , i.e., all bipartite graphs.

For graphs with loops allowed, the situation is even simpler—all graphs which have a loop have core L , the graph with one vertex and a loop. Clearly, the problem of L -colourability is trivial, since every graph with loops allowed is homomorphic to L . To summarize these observations, we have shown the following facts.

Proposition 5.1 *Let H be a graph with loops allowed. If H is bipartite or contains a loop, then the H -colouring problem has a polynomial time algorithm.*

□

5.2 Dichotomy for graphs

Recall that the usual k -colouring problem is polynomial time solvable when $k \leq 2$ and is NP -complete for $k > 2$. Such a classification result implies that we have a *dichotomy* of possibilities—each k -colouring problem is polynomial time solvable or NP -complete. Of course, if $P = NP$, then each k -colouring problem is polynomial time solvable, and this dichotomy is not interesting. But assuming $P \neq NP$ (as is generally believed), we can say that each k -colouring problem is *either* polynomial time solvable *or* NP -complete, and this dichotomy is an interesting fact, since it is known that (under our assumption) there are computational problems in NP that are neither polynomial time solvable nor NP -complete [212]. In other

words, the dichotomy result says that k -colouring problems do not contain such pathological examples.

It turns out that this result can be viewed as a special case of the following deeper dichotomy result, classifying the complexity of all homomorphism problems.

Theorem 5.2 *Let H be a graph with loops allowed.*

- *If H is bipartite or contains a loop, then the H -colouring problem has a polynomial time algorithm.*
- *Otherwise the H -colouring problem is NP-complete.*

The first statement is proved in Proposition 5.1. As for the second statement (the H -colouring problem for nonbipartite graphs H is NP-complete), we shall accept it without proof for a very small (cf. Exercise 12) and special class of graphs, and then derive it for all other graphs.

We say that a graph H is *extra-ordinary* (Fig. 5.1), if for every vertex v

- (P1) the neighbours of v induce a subgraph of H which is a union of disjoint edges $a_1b_1, a_2b_2, \dots, a_kb_k$, and
- (P2) each nonneighbour of v lies in a four-cycle $C_{i,j}$, whose consecutive vertices are adjacent to a_i, a_j , to a_i, b_j , to b_i, b_j , and to b_i, a_j , for some $i \neq j$.

(There could be additional edges joining the unlabeled vertices to each other and to some a_i or b_j ; the four-cycles $C_{i,j}$ could also intersect.)




FIG. 5.1. An extra-ordinary graph.

A nonbipartite graph which is not extra-ordinary will be called *ordinary*. We shall use without proof the following fact.

Proposition 5.3 [168] *Let H be an extra-ordinary graph. Then the H -colouring problem is NP-complete.* \square

The original proof of Proposition 5.3 in [168] is long and technical. There is now a new algebraic proof [50], based on an analysis of polymorphisms of an extra-ordinary graph (and of its subgraphs); we shall briefly discuss the technique, related to projectivity, at the end of this section.

Accepting this fact, we now complete, in the remainder of this section, the **proof** of Theorem 5.2. It is worth noting that the properties (P1) and (P2) (from the definition of an extra-ordinary graph) will play a role only at the end of the proof, just above Corollary 5.9.

We begin by illustrating our first tool.

Proposition 5.4 *The C_5 -colouring problem is NP-complete.*

Proof Note that the C_3 -colouring problem is just the usual three-colouring problem, well-known to be NP-complete. We now describe a reduction from the usual five-colouring problem to the C_5 -colouring problem. Given an instance G of the five-colouring problem, we replace each edge of G with a path of length three, calling the resulting graph $*G$. (Using the terminology of Section 4.4, $*G = G * P_3$ where the replacement graph P_3 uses the first and last vertex of the path as its connector vertices.) We now claim that G has a five-colouring if and only if $*G$ has a C_5 -colouring. Indeed, a five-colouring of G , with colours 0, 1, 2, 3, 4, provides us with images for all the branch vertices of $*G$, i.e., those vertices that were present in G . The inner vertices of $*G$, i.e., those lying on the added paths, can easily be mapped so that edges are preserved, since for any pair of *distinct* vertices of C_5 there is a homomorphism from the path of length three to C_5 , which maps the endpoints to these prescribed vertices. (If a, b are consecutive vertices of C_5 , the path can map to a, b, a, b ; otherwise, we may assume that, say, $a = 0, b = 2$, and the path can map to $0, 4, 3, 2$.) Conversely, any homomorphism of $*G$ to C_5 yields a mapping of the branch vertices in which two branch vertices joined by one of the added paths (and thus adjacent in G) have distinct images. This is so, since C_5 contains no triangle, and hence there is no homomorphism of P_3 to C_5 in which the endpoints have the same image. \square

This technique can be generalized as follows. Let I be a fixed graph with two specified vertices i, j such that some automorphism of I exchanges i and j . The *indicator construction* (with respect to the *indicator* I, i, j) transforms an arbitrary graph H into the graph H^* , with the same set of vertices as H , and with adjacency defined by the following rule: xy is an edge of H^* if and only if there exists a homomorphism of I to H that maps i to x and j to y . (The same construction also applies to digraphs H , except in this case do not require that I has an automorphism exchanging i and j .)

Lemma 5.5 *If the H^* -colouring problem is NP-complete, then so is the H -colouring problem.*

Proof There is a polynomial time reduction from the H^* -colouring problem to the H -colouring problem. Let $*G = G * I$, i.e., let $*G$ be obtained from G by

replacing each edge uu' of G with a copy of I , identifying i with u and j with u' . Then we again easily verify that $*G \rightarrow H$ if and only if $G \rightarrow H^*$. \square

Note that the indicator construction with respect to $I = P_3$, with i, j being the first and last vertex of the path, transforms $H = C_5$ into $H^* = K_5$. Thus Proposition 5.4 is a special case of Lemma 5.5.

In general, the result of the indicator construction, i.e., the graph H^* , can have loops—if I admits a homomorphism to H in which the images of i and j are equal. However, we will want to choose the indicator I, i, j so that H^* -COL is NP -complete, and, therefore, so that H^* is a graph (i.e., has no loops).

Corollary 5.6 *The C_k -colouring problem is NP -complete, for any odd positive integer $k > 1$.*

Proof Let $H = C_k$ and let I be a path of length $k - 2$ with endpoints i, j . Then $H^* = K_k$, and so H^* -COL is NP -complete. The conclusion now follows from Lemma 5.5. \square

In addition to the indicator construction, the proof Theorem 5.2 uses just one other technique—the so called *sub-indicator construction*, illustrated in the following example.

Proposition 5.7 *Let S be the penny graph from Fig. 5.2. Then S -COL is NP -complete.*




FIG. 5.2. The penny graph S .

Proof The graph S is called the penny graph because it can be realized by seven points in the plane with two points adjacent if and only if they have Euclidean distance one. (It is a subgraph of the *unit distance graph*, cf. Exercise 18 in Chapter 6. It is also known as the *Moser spindle*.)

We first remark that S is a core. Indeed, it is easily checked that it has chromatic number four, but every proper induced subgraph is three-colourable. (See Exercise 6 in Chapter 1.) We shall give a polynomial time reduction from the problem of three-colourability to S -COL. Let $+G$ be the following graph

(constructible from G in polynomial time). We start with the disjoint union of G and S , and adjoin, for every vertex u of G , a new vertex u' adjacent only to u and to the two vertices v and w of S .

We now claim that G is three-colourable if and only if ${}^+G \rightarrow S$. If G is three-colourable, then colour it with colours a, b, c . Each vertex u' for which the corresponding vertex u has been coloured b or c can be coloured y , and each vertex u' for which u has been coloured a can be coloured z . Taking S to itself identically completes a homomorphism of ${}^+G$ to S .

Hence suppose ${}^+G \rightarrow S$. Since S is a core, we may assume (by composition with a suitable automorphism of S), that some homomorphism $f : {}^+G \rightarrow S$ takes S to itself identically. In that case, all of the vertices u' , for u in G , will be mapped to y or z , and hence all the vertices of G will be mapped to a, b, c, v , or w . If we now focus only on G , then each vertex coloured v can be re-coloured by b and each vertex coloured w by c . Thus G is three-colourable. \square

This example is generalized as follows. Let J be a fixed graph with specified vertices k_1, k_2, \dots, k_t and j . The *sub-indicator construction* (with respect to the *sub-indicator* $J, k_1, k_2, \dots, k_t, j$) transforms an arbitrary graph H with specified vertices x_1, x_2, \dots, x_t into its subgraph H^+ defined as follows. Let W be the graph obtained from the disjoint union of H and J by identifying each k_i with the corresponding $x_i, i = 1, 2, \dots, t$. Then H^+ is the subgraph of H induced by those vertices x for which some retraction of W maps j to x .

In general, we have the following fact.

Lemma 5.8 *Let H be a graph which is a core. If the H^+ -colouring problem is NP-complete, then so is the H -colouring problem.*

Proof We give again a polynomial time reduction from the H^+ -colouring problem to the H -colouring problem. Given a graph G , we construct a graph ${}^+G$ such that $G \rightarrow H^+$ if and only if ${}^+G \rightarrow H$. The construction of ${}^+G$ is as above—we adjoin to the disjoint union of G and H , one copy of J for each vertex u of G , identifying the vertex j with u , and identifying each vertex k_i with the corresponding vertex x_i , for $i = 1, 2, \dots, t$.

It is again easy to see that $G \rightarrow H^+$ implies ${}^+G \rightarrow H$. When ${}^+G \rightarrow H$ then the copy of H that is a subgraph of ${}^+G$ must map onto H , since H is a core. It is easy to deduce, as above, that this implies that $G \rightarrow H^+$. \square

Note that when J is the graph $K_{1,3}$ with the vertices of degree one named j, k_1, k_2 , then the sub-indicator construction on S with specified vertices $x_1 = v, x_2 = w$ results in S^+ being the subgraph of S induced by $\{a, b, c, v, w\}$, which is homomorphically equivalent to K_3 . Thus Proposition 5.7 is a special case of Lemma 5.8.

Suppose, for future reference, that H is any graph and $f : S \rightarrow H$ a homomorphism. Then the same sub-indicator construction applied to $x_1 = f(v), x_2 = f(w)$, reduces the H -colouring problem to the H^+ -colouring problem. In this case, H^+ is a subgraph of H which contains the triangle $f(a), f(b), f(c)$.

If, moreover, H does not contain a K_4 , then H^+ does not contain the vertices $f(y), f(z)$. Hence in many cases we will be able to derive the NP -completeness of the H -colouring problem when $S \rightarrow H$.

We now continue **proving** Theorem 5.2 by contradiction. Suppose it is false, and let H be a nonbipartite graph such that

1. the H -colouring problem is not NP -complete,
2. subject to 1, H has the smallest number of vertices, and
3. subject to 1 and 2, H has the greatest number of edges.

It follows from the definitions that H is a connected graph and a core. We shall now derive some additional properties of the graph H .

Property 1 *The largest clique in H is the triangle.*

Proof First we show that H must contain a triangle. Otherwise, assume that the shortest odd cycle has length $k > 3$. Let I be the path P_3 (of length three) with end vertices i and j , and perform the indicator construction. It is easy to see from the definitions that H^* contains all edges of H . The absence of triangles ensures that H^* is also a graph (has no loops), and the fact that H has a cycle of length $k > 3$ ensures that H^* has more edges than H . Thus H^* is a nonbipartite graph with no more vertices than H and strictly more edges than H . If the H^* -colouring problem was not NP -complete, we would have a contradiction with the assumptions 1–3 on H . Therefore, H^* -colouring is NP -complete, and by Lemma 5.5 so is H -colouring, contrary to assumption 1.

Next we show that H cannot contain a K_4 . Otherwise, let J be an edge with end vertices j and k_1 , and perform the sub-indicator construction on H with x_1 being a vertex in a K_4 . The resulting graph H^+ is nonbipartite, since it contains a triangle amongst the neighbours of x_1 . On the other hand, H^+ does not contain x_1 , hence it is a smaller graph than H . Once more, if H^+ -colouring is not NP -complete we obtain a contradiction with the assumptions 1–2, and if H^+ -colouring is NP -complete we obtain a contradiction with Lemma 5.8 as above. \square

Property 2 *Each edge of H belongs to a unique triangle.*

Proof We begin by showing that each vertex of H belongs to a triangle. Let J be the disconnected sub-indicator consisting of a triangle j, j', j'' , and an isolated vertex k_1 , and perform the sub-indicator construction on H , with x_1 being any vertex. It follows that H^+ is the subgraph of H induced by the vertices that belong to triangles. Hence H^+ is a nonbipartite graph; it cannot have fewer vertices than H by Lemma 5.8. Hence each vertex belongs to a triangle.

Next, we show that each edge of H belongs to a triangle. Consider the sub-indicator J consisting of a path of length two with end vertices j and k_1 . Using this sub-indicator with any vertex x_1 of H , results in a nonbipartite graph H^+ which contains any triangle containing x_1 (and we have proved there exist such triangles) and does not contain any vertex not having a path of length two to

x_1 . Hence $V(H^+) = V(H)$, and any two vertices of H are joined by a path of length two; in particular, each edge belongs to a triangle.




FIG. 5.3. The indicator I .

It remains to show that an edge cannot belong to two triangles, i.e., that the graph K_4^- , obtained from K_4 by removing an edge, is not a subgraph of H . According to the discussion below Lemma 5.8, we may assume that $S \not\rightarrow H$. This allows us to use the indicator I in Fig. 5.3 without creating loops in H^* . Since each vertex of H lies in a triangle, H^* contains all edges of H . Indeed, consider an edge xx' of H . There is a homomorphism of I taking i to x and j to x' using the triangles incident to x and to x' . Now suppose H contains a K_4^- ; since H does not contain a K_4 by Property 1, we may assume that u, v are the nonadjacent vertices of our K_4^- . Moreover, H is a core, so the neighbourhoods of u and v cannot be the same. Thus some vertex w is adjacent to, say, v but not u . Since w belongs to a triangle, there is a homomorphism of I to H taking i to u and j to w . This means that H^* is an nonbipartite graph with the same vertices as H and having more edges than H , with NP -complete H^* -colouring problem by Lemma 5.5, contrary to our assumptions on H . \square

We denote by P the graph in Fig. 5.4.




FIG. 5.4. The graph P .

Property 3 Any two intersecting triangles of H are contained in some induced subgraph P .

Proof Let $abc, ab'c'$ be two triangles of H intersecting in the vertex a . Property 2 implies that the vertices b, b', c, c' are distinct, as two triangles of H cannot intersect in more than one vertex. Consider the sub-indicator J in Fig. 5.5, applied to H with $x_1 = b'$ and $x_2 = c'$.




FIG. 5.5. The sub-indicator J .

The graph H^+ will contain the triangle $ab'c'$, and hence we again arrive at a contradiction, unless b (and c) is also in H^+ . This means that there exists a vertex d adjacent to b' and a vertex f adjacent to c' and both d and f are adjacent to b and to each other. By Property 2, there exist triangles $b'de, c'fg$; it also follows from Property 2 that all the nine vertices $a, b, b', c, c', d, e, f, g$ are distinct. We now prove that ce, cg, eg are edges of H . Consider the indicator U in Fig. 5.6.




FIG. 5.6. The indicator U .

The discussion following Lemma 5.8 ensures that U does not admit a homomorphism to H in which i and j obtain the same image. Therefore, H^* is a graph (has no loops). It is also easy to see that H^* contains all the edges of H . (U admits a homomorphism f to a triangle, with $f(i) \neq f(j)$.) Moreover, U admits a homomorphism f to the named nine vertices of H with $f(i) = c, f(j) = e$, respectively $f(i) = c, f(j) = g$, and $f(i) = e, f(j) = g$. If either one of these

edges were not in H we would obtain a contradiction with Lemma 5.8 as before. \square

We are now ready to conclude the **proof** of Theorem 5.2. Indeed, Properties 1–3 imply that H is an extra-ordinary graph. Property 2 implies that the neighbours of any vertex v induce a subgraph of H consisting of disjoint edges as in (P1), and Property 3 implies that for any two of these edges there exists a four-cycle as in (P2) (Fig. 5.1). Finally, it is easy to see that this accounts for all the vertices of H , as any vertex u has a path of length two to v (see the proof of Property 2), and thus participates in one of these four-cycles. Therefore, Theorem 5.2 now follows from Proposition 5.3. \square

Note that we have the following byproduct of Theorem 5.2.

Corollary 5.9 *Suppose H' is an induced subgraph of a graph H . There is a polynomial time reduction from the H' -colouring problem to the H -colouring problem.*

Proof We may assume that H' is not bipartite, otherwise the problem is polynomial time solvable and we can reduce each instance of H' -colouring to one of two fixed instances of H -colouring (a positive instance and a negative instance). But then H is also not bipartite and thus H -colouring is NP -complete. It is easy to see that all homomorphism problems are in the class NP . According to a theorem of S. Cook [68], all problems in NP can be polynomial time reduced to any NP -complete problem. Hence there exists a reduction from the H' -colouring problem to the H -colouring problem. \square

It is ironic that we derived Corollary 5.9 in such a roundabout way (using Cook's Theorem). After all, it is precisely the kind of statement which is routinely proved *directly*—by actually constructing a reduction. If we could find such a direct construction, we would have a much more satisfying proof of the above theorem—since every nonbipartite graph contains an odd cycle, and the NP -completeness of H -colouring when H is an odd cycle was easily proved in Corollary 5.6. On the other hand, there may be good reasons for the difficulties in trying to prove Corollary 5.9 directly—for instance it fails to hold for digraphs, as we shall see in the next section.

5.3 Digraph homomorphisms and CSPs

In Fig. 5.7 we have a digraph H which is a subgraph of a digraph H' . Yet H -COL is NP -complete and H' -COL is polynomial.

Proposition 5.10 *Let H be the digraph from Fig. 5.7. Then the H -colouring problem is NP -complete.*

Proof We take this opportunity to explain one additional reduction, useful for proving the NP -completeness of the H -colouring problem for many digraphs H . The reduction is from the well-known NP -complete problem ONE-IN-THREE




FIG. 5.7. A hard digraph H and an easy digraph H' .

satisfiability without negated variables [116]: given a set of clauses with three variables each, none negated, is there a truth assignment in which each clause has exactly one true variable?




FIG. 5.8. The gadget D .

Consider the digraph D depicted in Fig. 5.8. In any homomorphism of D to H the directed triangle $x'y'z'$ of D must map to the directed triangle 012 of H . Suppose, without loss of generality, that the vertex x' is taken to 1. Then x must also be taken to 1, while the other two vertices y, z can each be taken to 0 or 2 independently of each other.

Given a set of three-variable clauses (without negated variables), we construct a digraph by taking a vertex for each variable x , and a copy of D for each clause C , and identifying the vertices x, y, z of D with the three variables occurring in the clause C . We now claim that the resulting digraph admits a homomorphism to H if and only if the clauses can be satisfied as required. Assume first that there is a truth assignment which gives each clause exactly one true variable. Mapping each true variable to 1 and each false variable to 0 extends to a homomorphism to H . On the other hand, having a homomorphism to H , we see that each clause has exactly one variable mapped to 1. Thus we can declare those to be true, and the other variables (mapped to 0 or 2) to be false. \square

Proposition 5.11 *Let H' be the digraph from Fig. 5.7. Then the H' -colouring problem has a polynomial time algorithm.*

Proof In a digraph G , no vertex of positive outdegree can be coloured 3, in any H' -colouring of G . On the other hand, if a vertex of outdegree zero obtained a colour different from 3, in some H' -colouring of G , we can re-colour it by 3, and still have an H' -colouring of G . Hence we can test for the H' -colourability of G by first colouring all vertices of outdegree zero by 3 and removing them from G . The original graph G is H' -colourable if and only if the reduced graph is C_3 -colourable, which can be checked by Algorithm 1. \square

In the example in Fig. 5.7, the graph H is not an induced subgraph of H' . In Exercise 5(b), we show how to construct graphs H, H' where H is an induced subgraph of H' , and it is still the case that H -colouring is hard and H' -colouring easy.

The H -colouring problem for digraphs has received much attention, and yet no graph-theoretic classification has been obtained, or conjectured. In fact, returning to our discussion of dichotomy from the beginning of Section 5.2, even the following dichotomy conjecture is still open.

Conjecture 5.12 *For each digraph H , the H -colouring problem is polynomial time solvable or NP-complete.*

Recall that the dichotomy of all constraint satisfaction problems is a long-standing open problem. Let T be a general relational system, with vertex set V and relations $R_i, i \in I$, where R_i is a k_i -ary relation on the set V . (Recall that the family P of integers k_i , indexed by a finite set, $i \in I$, is called the *pattern* of T .) Each such system T defines a *constraint satisfaction problem with respect to the template T* , denoted by T -CSP, in which we are to decide whether or not a given general relational system S , with the same pattern P , is homomorphic to T . When T is a digraph, i.e., has just one binary relation, the problem T -CSP is precisely the problem T -COL; by extension, we also occasionally use T -COL to denote T -CSP. We shall also consider the *retraction problem for T* , denoted by T -RET. In T -RET, the instance is a relational system S with the same pattern P , which contains T as a subsystem, and we are to decide whether or not there is a retraction of S to T , i.e., a homomorphism f of S to T such that $f(x) = x$ for all x in T . It is easy to verify (Exercise 2) that if T is a core, then T -RET is polynomially equivalent to T -COL. This is true for general relational systems T , and in particular, also for digraphs T .

Recall the general *Dichotomy Conjecture* of Feder and Vardi.

Conjecture 5.13 *For each general relational system T , the constraint satisfaction problem T -CSP is polynomial time solvable or NP-complete.*

The role of digraphs is central to dichotomy, as we have the following fact implying that Conjecture 5.12 implies Conjecture 5.13.

Theorem 5.14 *Every constraint satisfaction problem T -CSP is polynomially equivalent to some H -colouring problem for a suitable digraph H .*

Proof The theorem will be proved by constructing, for each relational system T , a digraph $H = H_T$, such that H -COL is polynomially equivalent to T -CSP. Recall that two problems are polynomially equivalent, if each of them admits a polynomial reduction to the other.

As remarked at the beginning of this section, we may assume that T is a core, i.e., that each homomorphism of T to itself is an automorphism of T . Therefore, as noted above, we may consider, instead of T -CSP, the problem T -RET. Thus it suffices to show that for each core relational system T there exists a core digraph $H = H_T$ such that T -RET and H -RET are polynomially equivalent.

Lemma 5.15 *For each general relational system T there exists a bipartite graph $H = H_T$ such that T -RET and H -RET are polynomially equivalent.*

Proof We shall first introduce an auxiliary property of general relational systems; it will help in formulating the polynomial reductions needed. We say that a general relational system T is *transversal* if the vertex set V of T can be partitioned into subsets V_1, V_2, \dots, V_k such that each relation $R_i, i \in I$, is a subset of the cartesian product of k_i *distinct* sets V_j .

We claim that we may assume that the relational system T is transversal. In other words, we claim that for each relational system T there exists a transversal relational system T' , such that the retraction problems for T and for T' are polynomially equivalent. (We note that T' need not have the same pattern as T .) Suppose k is the maximum arity k_i , over all $i \in I$, and let V_1, V_2, \dots, V_k be disjoint copies of the set V . Hence each vertex $v \in V$ corresponds to a different vertex v^j in each $V_j, j = 1, 2, \dots, k$. The system T' will have the vertex set $V' = V_1 \cup V_2 \cup \dots \cup V_k$, and for each relation $R_i, i \in I$, on V , it will have a relation R'_i , of the same arity k_i , consisting of all k_i -tuples (u^1, v^2, w^3, \dots) such that $(u, v, w, \dots) \in R_i$. Finally, T' will also have the additional binary relations $E_{j,j'}$ consisting of all ordered pairs $(v^j, v^{j'})$ with $v \in V$, for all distinct $j, j' = 1, 2, \dots, k$. We say that T' was obtained from T by *stretching*. It is clear that we have successfully transformed (stretched) T , in the sense that T' is now transversal. It remains to show that the retraction problems for T and T' are polynomially equivalent. For each instance S of T -RET we can define a natural instance S' of T' -RET, obtained by the same stretching construction from S . If S retracts to T , then clearly S' retracts to T' . Moreover, if S' retracts to T' , then, since all relations $E_{j,j'}$ are preserved, we also obtain a retraction of S to T . Conversely, for each instance S' of T' -RET, we define an instance S of T -RET as follows. Repeatedly identify in S' any two vertices related by any of the relations $E_{j,j'}$, until no distinct vertices are in any of these relations; call the final system S . Note that this will cause the subsystem T' of S' to become the system T . It is then easy to argue that S retracts to T if and only if S' retracts to T' .

Thus assume that T is a transversal general relational system, with vertex set V partitioned into parts V_1, V_2, \dots, V_k , satisfying the transversal condition. Let R be the disjoint union of all the sets $R_i, i \in I$. The bipartite graph $H = H_T$ will contain as vertices all elements of V and of R , with the tuple $r \in R$ adjacent

to all $v \in V$ contained in it (and no others). There will also be a vertex for each part $V_j, j = 1, 2, \dots, k$, adjacent to all $v \in V_j$ (and no other $v \in V$), and a vertex for each relation $R_i, i \in I$, adjacent to all tuples $r \in R_i$ (and no other $r \in R$). We denote these vertices also by V_j and R_i . Finally, H also contains two special adjacent vertices a^* and b^* , where a^* is adjacent to all $v \in V$, and b^* is adjacent to all $r \in R$ (Fig. 5.9). We have constructed a bipartite graph H with parts A and B , where A consists of a^* , all $r \in R$, and the vertices corresponding to the sets $V_j, j = 1, 2, \dots, k$, and B consists of b^* , all $v \in V$, and the vertices corresponding to the relations $R_i, i \in I$.




FIG. 5.9. The bipartite graph H .

We first describe a polynomial reduction from T -RET to H -RET. Hence, let S be a relational system containing T and with the same pattern as T . Let X be the set of vertices of S not in T . If $x \in X$ is not involved in any constraint, i.e., a tuple of some relation in S , it can be mapped arbitrarily. Recall that S has the same pattern as T , thus there are relations corresponding to each $R_i, i \in I$, and each R_i is a subset of some cartesian product of distinct sets V_j . Suppose now that $x \in X$ is in a position of one tuple which requires it to map to a set V_j , and another position of another tuple that requires it to map to $V_{j'}$ with $j \neq j'$; then clearly no retraction is possible. Therefore, we may assume that each $x \in X$ is associated with a unique j so that x cannot retract to any vertex other than V_j . We now define a bipartite graph $G = G_S$, containing the subgraph H . In addition to H , it will contain, for each vertex $x \in X$ a vertex adjacent to the V_j associated with x . For each $i \in I$, and each k_i -tuple r of the relation in S corresponding to R_i , we create a vertex r , adjacent to the vertex of B corresponding to R_i and to the k_i vertices corresponding to the vertices of V involved in the tuple r , including the new vertices corresponding to $x \in X$. It is now easy to argue that G_S retracts to H_T if and only if S retracts to T .

Finally, we describe a polynomial reduction from H -RET to T -RET. Let G be a bipartite graph containing H as a subgraph. We may assume G to be bipartite, as otherwise no retraction is possible. We may also assume that no vertex of G is adjacent to two distinct vertices corresponding to $V_j, V_{j'}$, or to

$R_i, R_{i'}$, as such a vertex would have no possible image under a retraction to H . On the other hand, we may also assume that each vertex of G (other than a^*, b^*) is adjacent to some V_j or to some R_i since all other vertices of G can safely be mapped to either a^* or b^* (depending on which part of the bipartition of G they belong to). We can therefore construct a corresponding instance S of T -RET by viewing each vertex of G adjacent to some V_j as a vertex of S and each vertex of G adjacent to some $R_i, i \in I$, as a tuple (constraint) from relation corresponding to R_i . \square

We now show that each bipartite retraction problem, H -RET with H bipartite, is polynomially equivalent to a digraph retraction problem H' -RET where H' is a core, and hence polynomially equivalent to H' -HOM by Exercise 2. Let A, B be a bipartition of H .


We first construct oriented paths $P, Q, P_a, a \in A, Q_b, b \in B$, satisfying the following properties.

1. The height of P and each P_a is $|A| + 1$, the height of Q and each Q_b is $|B| + 1$,
2. the last vertex of each P_a is a , and the first vertex of each Q_b is b ,
3. the last vertex of P is called 0, the first vertex of Q is called 1,
4. $P_a \rightarrow P_{a'}$ implies $a = a'$,
5. $Q_b \rightarrow Q_{b'}$ implies $b = b'$,
6. $P \rightarrow P_a$ for all $a \in A$,
7. $Q \rightarrow Q_b$ for all $b \in B$,
8. if X is a digraph and x a vertex of X such that for some $a \neq a'$ there exist homomorphisms $f : X \rightarrow P_a$ and $f' : X \rightarrow P_{a'}$, with $f(x) = a$ and $f'(x) = a'$, then there exists a homomorphism $F : X \rightarrow P$ with $F(x) = 0$, and
9. if Y is a digraph and y a vertex of Y such that for some $b \neq b'$ there exist homomorphisms $f : Y \rightarrow Q_b$ and $f' : Y \rightarrow Q_{b'}$, with $f(y) = b$ and $f'(y) = b'$, then there exists a homomorphism $F : Y \rightarrow Q$ with $F(y) = 1$.

Such a collection of oriented paths is easily constructed from the zig-zags Z_k in Fig. 3.1. For instance, consider the path Z obtained from Z_2 by deleting the last two vertices, c_2, d . (In other words, Z is the oriented path consisting of two forward arcs followed by one backward arc.) Using exponents to denote repeated concatenation, we may take for the path P the path $Z^{h-2} \cdot \vec{P}_2$, and for the family $P_a, a \in A$, the family of paths $Z^i \cdot \vec{P}_1 \cdot Z^{h-2-i} \cdot \vec{P}_2, i = 0, 1, \dots, h-2$. (The integer h is chosen sufficiently large to have $h-1 \geq |A|$.) The path Q and the family $Q_b, b \in B$, is constructed similarly. The properties 1–9 are easily checked.

The construction of the digraph H' begins by taking H and orienting all edges from A to B . Then the paths $P_a, a \in A$, and $Q_b, b \in B$, are added, attached by their common vertices, a, b , to H . The resulting graph H' , depicted in Fig. 5.10 is balanced, of height $|A| + |B| + 3$, and a core (because of the properties 4 and 5, cf. Proposition 1.14).

We claim that H -RET is polynomially equivalent to H' -RET (and hence to H' -HOM). If G is a bipartite graph containing H , we construct a digraph G' containing H' , with parts A', B' containing A, B respectively, by orienting all

FIG. 5.10. The graph H' .

edges from A' to B' , and attaching at each new vertex of A' the last vertex 0 of a separate copy of P , and attaching at each new vertex of B' the first vertex 1 of a separate copy of Q . Because of properties 6 and 7, Proposition 1.14 implies that G' retracts to H' if and only if G retracts to H . Conversely, suppose we have a digraph G' containing H' . We may assume G' is balanced, else no retraction to the balanced H' would be possible. The two levels of G' corresponding to the vertices of H define an orientation of a bipartite graph, which we shall call G . The components of $G' - G$ fall into two types, those whose vertices are on levels lower than the levels of G and those whose vertices are on levels higher than the levels of G . Let C_1, \dots, C_ℓ be the components of the former type. If some C_j contains two distinct vertices of G , then we may as well identify those vertices, otherwise there could be no retraction of G' to H' , since C_j will have to map to some P_a . Thus we may assume that each C_j is attached to a unique vertex of G . By the same argument, we may assume that C_j is homomorphic to some $P_a, a \in A$. If a is unique, we may map C_j to P_a , modifying the graph G' . Hence we may assume that each C_j is homomorphic to two distinct paths P_{a_1}, P_{a_2} , and hence homomorphic to all the paths $P_a, a \in A$, by properties 6 and 8. Then G' retracts to H' if and only if G retracts to H . This completes the proof of Theorem 5.14. \square

Note that Theorem 5.14 implies Theorem 1.37, i.e., that we now know that dichotomy for digraphs would imply dichotomy for all constraint satisfaction problems. As discussed before the theorem, Conjecture 5.13 is a difficult open problem. Hence we may want to lower our sights somewhat—we should be happy with dichotomy classifications of restricted classes of digraphs. Theorem 5.2 is such a result, since the class of graphs is a subclass of the class of digraphs. A few such results are known, cf. Section 5.9, and the Exercises. However, in general, dichotomy results seem hard to come by. It is, for instance, not known whether

or not the class of oriented trees has dichotomy. Dichotomy is not even known for trees H which consist of three oriented paths meeting at one vertex. Such trees are called *triads*; Fig. 5.11 shows a triad H (constructed by concatenating zig-zags as suggested by the enlarged vertices), for which H -colouring has been shown to be NP -complete.




FIG. 5.11. A tree H for which H -colouring is NP -complete.

By contrast, we shall show in the next section that H -colouring has a polynomial time algorithm whenever H is an oriented path.

There is a conjectured dichotomy classification for digraphs without sources and sinks, i.e., those digraphs in which each vertex has positive indegree and outdegree.

Conjecture 5.16 *Suppose H is a connected digraph without sources and sinks. If the core of H is a directed cycle, then H -colouring is polynomial time solvable. Otherwise H -colouring is NP -complete.*

We note that it is easy to check whether or not the core of a given digraph G is a directed cycle, since it would have to be a shortest directed cycle in G . Thus

finding the length k of a shortest directed cycle and testing for \vec{C}_k -colourability, by the algorithm of Section 1.4, will accomplish the check.

The conjecture has been verified for a number of digraph classes. We may view Theorem 5.2 as verifying the conjecture for the class of graphs (subclass of the class of digraphs), since a connected graph H has a core which is a directed cycle if and only if the graph is bipartite or has a loop.

Before leaving the topic of dichotomies, we should remark that there is an algebraic approach to proving NP -completeness of H -colouring problems, which substantially depends on the generality of constraint satisfaction problems. Recall from Section 2.7 that a polymorphism of G is a homomorphism of any numeric power G^k to G .

Proposition 5.17 *Assume $V(G) = V(H) = V$. If each polymorphism of G is also a polymorphism of H , then the H -colouring problem is polynomially reducible to the G -colouring problem.*

Proof To prove the proposition requires a consideration of general CSPs, even if both G and H are digraphs. Specifically, one can show [193] that if each polymorphism of G is also a polymorphism of H , then H can be obtained from G by a sequence of operations of the following kind.

- Adding one new binary relation E consisting of all pairs $(v, v), v \in V$.
- Taking an existing k -ary relation and a permutation of $1, 2, \dots, k$, and adding a new k -ary relation consisting of all permuted k -tuples.
- Taking an existing k -ary relation R and adding a new $(k + 1)$ -ary relation consisting of all $(k + 1)$ -tuples in which the first k elements form a k -tuple from R .
- Taking an existing k -ary relation R and adding a new $(k - 1)$ -ary relation (for $k > 1$) obtained by deleting the k -th element of each k -tuple in R .
- Taking the intersection of any two existing relations.

(Note that the relations that have been introduced may be used in the rest of the sequence.) It is then shown by routine reductions that if a relational system X is transformed to a relational system Y by one step in the sequence, then Y -COL is polynomially reducible to X -COL.

Consider the first kind of operation; let Y be obtained from X by the addition of the relation E ('equality'). Given an input A for Y -COL, we form A' as follows. Whenever two vertices v, v' of A are related in a binary relation R corresponding to E , we replace all occurrences of v by v' , in all tuples of A , and remove the relation R . It is now easy to check that A is Y -colourable if and only if A' is X -colourable.

For the second operation, suppose Y is obtained from X by adding a relation R' obtained by permuting an existing relation R . Given an input A for Y -COL, we form A' as follows. We remove the relation corresponding to R' , and apply the inverse permutation to all tuples corresponding to R . Then we again have that A is Y -colourable if and only if A is X -colourable.

The other operations are handled similarly. □

The proposition implies that having fewer polymorphisms makes H -colouring problems more likely to be NP -complete. Recall, for instance, that for $n \geq 3$, the graph K_n is projective (Corollary 2.44). This means that the only polymorphisms of K_n are obtained from projections by a permutation. Thus K_n has truly few polymorphisms—and, indeed, we know that K_n -COL is NP -complete. Another general relational system with few polymorphisms is the system N introduced in Section 1.8. Recall that N has two vertices 0, 1, and one ternary relation $E(N) = \{(1, 0, 0), (0, 1, 0), (0, 0, 1), (1, 1, 0), (1, 0, 1), (0, 1, 1)\}$. (See Exercise 13.) However, we do not really need to know all polymorphisms of N —we already know that N -COL is NP -complete, since it is the problem of NOT-ALL-EQUAL 3-SAT. The point is that any general relational system S on two elements (say 0 and 1), which has no other polymorphisms than those of N , i.e., no others than projections followed by a permutation of 0 and 1, has the problem S -COL NP -complete by Proposition 5.17. Similarly, any system S on $k \geq 3$ elements which has no other polymorphisms than those of K_k , i.e., no others than projections followed by a permutation of $V(K_k)$, must have an NP -complete problem S -COL.

It appears that all known NP -complete H -colouring problems, for any general relational system H , fit this framework, i.e., the polymorphisms on H can be shown to be included (possibly after a suitable transformation we will not discuss [54]) in the set of polymorphisms of N or of some K_n with $n \geq 3$. In any event, it has been conjectured in [54] that these are all the NP -complete constraint satisfaction problems.

This approach can be used for a probabilistic analysis of dichotomy. How can one consider the asymptotic behaviour of constraint satisfaction problems? There are two extreme cases. One can fix the pattern P and consider relational systems S with pattern P and large $V(S)$, or one can fix the vertex set V and consider relational systems S with $V(S) = V$ and large patterns P . It can be shown, by analysing the polymorphisms of T , that in both cases asymptotically almost all problems T -CSP are NP -complete.

Theorem 5.18 [230] *Let P be a fixed pattern, $k_i, i \in I$, with some $k_i > 1$. Then asymptotically almost all relational systems T with pattern P are projective, and hence have NP -complete problems T -CSP, as the size of $V(S)$ increases to infinity.*

Let V be a fixed set with at least two elements. Then asymptotically almost all relational systems T which contain at least one k -ary relation are projective, and hence have NP -complete problems T -CSP, as k increases to infinity.

In particular, asymptotically, there is dichotomy for constraint satisfaction problems: most are NP -complete. Therefore, the Dichotomy Conjecture (Conjecture 5.13) amounts to simply asserting that the rare problems that are not NP -complete are in fact polynomial time solvable.

5.4 Duality and consistency

In this section we focus on digraphs (and occasionally graphs). Recall that most of our simple polynomial time algorithms for the various H -colouring problems seem to have been accompanied by duality theorems. (Duality has been discussed in Section 1.4 and 3.8.) In particular, we have given polynomial time algorithms for H -colouring when H is \vec{C}_n , K_2 , \vec{P}_n , or \vec{T}_n , and simultaneously we have established the following duality results (cf. Corollary 1.18, Corollary 1.19, Exercise 14 in Chapter 1, and Proposition 1.20).

- Proposition 5.19**
1. $G \not\rightarrow \vec{C}_n$, for a digraph G , if and only if $C \rightarrow G$ for some oriented cycle C of net length not divisible by n .
 2. $G \not\rightarrow K_2$, for a graph G , if and only if $C_k \rightarrow G$ for some odd k .
 3. $G \not\rightarrow \vec{P}_n$, for a digraph G , if and only if $P \rightarrow G$ for some oriented path P of net length greater than n .
 4. $G \not\rightarrow \vec{T}_n$, for a digraph G , if and only if $\vec{P}_n \rightarrow G$. □

These results suggest the following terminology. Let \mathcal{T} denote the class of all oriented trees. We say that the graph H has *tree duality*, if there exists a set $\mathcal{F} \subseteq \mathcal{T}$ such that

- $G \not\rightarrow H$ if and only if $T \rightarrow G$ for some $T \in \mathcal{F}$.

Note that each member T of \mathcal{F} must have $T \not\rightarrow H$, since the condition must apply to $G = T$. A graph H which has tree duality allows us to certify non- H -colourability of G by an obstruction that is an oriented tree T with $T \rightarrow G$ and $T \not\rightarrow H$. Statements 3 and 4 of Proposition 5.19 show that \vec{P}_n and \vec{T}_n have tree duality.

More generally, we consider orientations of graphs with bounded treewidth. A graph is a k -tree if it can be obtained from K_k by a sequence of vertex additions, where each new vertex is adjacent to a clique of size k in the previously generated graph. A graph has *treewidth at most k* if it is a subgraph of a k -tree. Let \mathcal{T}_k denote the class of all connected digraphs which are orientations of graphs of treewidth at most k . Note that a 1-tree is just a tree, and thus $\mathcal{T} = \mathcal{T}_1$. We say that the graph H has *treewidth k duality* if there exists a set $\mathcal{F} \subseteq \mathcal{T}_k$ such that

- $G \not\rightarrow H$ if and only if $T \rightarrow G$ for some $T \in \mathcal{F}$.

A digraph H has *bounded treewidth duality*, if there is an integer k such that H has treewidth k duality.

It is easy to see that cycles have treewidth two, so statements 1 and 2 above show that \vec{C}_n , K_2 have treewidth two duality.

The examples in Proposition 5.19 suggest that polynomial time H -colouring algorithms are related to bounded treewidth duality of the graphs H . We will prove the following result.

Theorem 5.20 *If H has bounded treewidth duality, then the H -colouring problem has a polynomial time algorithm.*

For digraphs, all known polynomial time solvable H -colouring problems fit this framework, i.e., the digraph H has bounded treewidth duality. (In fact, in all these cases, the graph H has tree duality, or admits a majority function, which means it has treewidth two duality, Theorem 5.31).

The algorithm needed for Theorem 5.20 is a procedure called consistency check. We shall focus on the special case of tree duality (that is, treewidth one duality). By verifying the correctness of the algorithm we will prove Theorem 5.20 for this special case; the general proof is along similar lines, using higher order consistency checks (cf. Section 5.5).

Suppose H is a fixed digraph, and G is a digraph in which every vertex $u \in V(G)$ has a list (that is, a set) $L(u) \subseteq V(H)$. The lists L are said to be *consistent* if for any arc $uv \in E(G)$ and any $x \in L(u)$ there exists an $y \in L(v)$ with $xy \in E(H)$, and for any arc $uv \in E(G)$ and any $y \in L(v)$ there exists an $x \in L(u)$ with $xy \in E(H)$.

Suppose the initial lists are all $L(u) = V(H)$. (In Section 5.6, we will remove this restriction.) The following algorithm reduces the given lists L to consistent lists L^* .

Algorithm 2 (Consistency Check)

Input: A graph G , with lists $L(u) = V(H)$, $u \in V(G)$.

Task: Reduce the lists to $L^*(u) \subseteq V(H)$, $u \in V(G)$, that are consistent.

Action: Initially set all lists $L^*(u) = L(u)$, and then, as long as changes occur, process $uv \in E(G)$ repeatedly as follows: remove from $L^*(u)$ any x for which no element $y \in L^*(v)$ has $xy \in E(H)$, and remove from $L^*(v)$ any y for which no $x \in L^*(u)$ has $xy \in E(H)$.

We say that the consistency check *fails* if some list $L^*(u)$ becomes empty; otherwise we say that the consistency check *succeeds*.

Arcs may have to be reprocessed as the changes propagate in the graph G ; however, since the size of H is fixed, the procedure is linear in the size of G (Exercise 11).

Recall how our polynomial time H -colouring algorithms became associated with duality results. Typically, as in Proposition 1.15, it was clear that the success of the proposed algorithm identified an H -colouring; and if the algorithm did not succeed in finding an H -colouring then it identified an obstruction certifying non- H -colourability. The consistency check behaves differently: a failure is a clear certificate of non- H -colourability.

Lemma 5.21 *If G is H -colourable, then the consistency check will succeed.*

Proof If f is a list H -colouring of G , then $f(u)$ will never be removed from $L^*(u)$. \square

On the other hand, if the check succeeds, we cannot be sure an H -colouring of G exists. For instance, suppose $H = K_2$ and $G = K_3$. The initial lists $L(u) = \{1, 2\}$ are consistent, so the consistency check will not modify them; yet no homomorphism of G to H exists. However, if H enjoys tree duality, then we *can*

conclude that $G \rightarrow H$, and hence verify the correctness of the consistency check for solving the H -colouring problem.

Theorem 5.22 *H has tree duality if and only if $G \rightarrow H$ whenever the consistency check applied to G succeeds.*

(The above example shows that K_2 does not have tree duality. Recall that it does have treewidth two duality.)

The theorem is a consequence of the following interpretation of the lists L^* produced by the consistency check.

Lemma 5.23 *$x \in L^*(u)$ if and only if for each oriented tree T and each $t \in V(T)$ the following condition holds:*

- *If there is a G -colouring of T which takes t to u , then there is an H -colouring of T which takes t to x .*

Proof Suppose first that $x \in L^*(u)$. If a homomorphism $f : T \rightarrow G$ satisfies $f(t) = u$, then we can begin to define a homomorphism $F : T \rightarrow H$ by setting $F(t) = x$. This definition has the property that $F(t) \in L^*(f(t))$. Assume we have already defined F on a subtree T' of T so that $F(s) \in L^*(f(s))$ for all s in T' . Suppose s' is an inneighbour (respectively an outneighbour) of a vertex s for which $F(s)$ has already been defined. Since $F(s) \in L^*(f(s))$, some inneighbour (respectively outneighbour) y of $F(s)$ is in $L^*(f(s'))$, and we can set $F(s') = y$, maintaining the desired property of F . This will define a homomorphism F as required.

We prove the converse by contradiction; thus suppose that there exist pairs u, x with $x \notin L^*(u)$ such that if there is a G -colouring of T which takes t to u , then there is an H -colouring of T which takes t to x . Consider now a concrete execution of the consistency check (Algorithm 2). For each pair u, x with the above property, the vertex x was removed from $L^*(u)$ at some point in the execution. We focus on the pair u, x , for which this point occurred first in the order of execution. Without loss of generality, assume that at this point we were processing the arc uu' . If x was removed from $L^*(u)$ because it had no outneighbours in H , then the tree T consisting of one arc tt' admits a G -colouring which takes t to u , but not an H -colouring which takes t to x , giving us a contradiction. Otherwise, x was removed from $L^*(u)$ because all outneighbours x' of x were missing from $L^*(u')$ at this point in the execution. According to the minimality of the pair u, x , that means that for each outneighbour x' of x there exists a ‘certifying tree’ $T_{x'}$ with a vertex $t_{x'}$, such that $T_{x'}$ admits a G -colouring which takes $t_{x'}$ to u' , but not an H -colouring which takes $t_{x'}$ to x' . Let T be the tree obtained from the union of the trees $T_{x'}$, for all outneighbours x' of x , by identifying all the vertices $t_{x'}$ into a common vertex t' , and adding a new vertex t with an arc tt' . It is easy to see that T contradicts our assumption, since it admits a G -colouring taking t to u , but no H -colouring taking t to x . \square

If H enjoys tree duality, then it follows from the lemma that each oriented tree which is homomorphic to G is also homomorphic to H and hence $G \rightarrow H$

by the tree duality of H . (As to actually finding such a homomorphism, see the proof of Proposition 5.26.) Conversely, if H doesn't have tree duality, then some tree T is homomorphic to G but not to H , and hence some lists $L^*(u)$ will be empty. This proves Theorem 5.22 (and hence Theorem 5.20 in the case of tree duality). \square

We illustrate the use of the technique by discussing a polynomial time algorithm for the H -colouring problem where H is any oriented path.

Theorem 5.24 *If H is an oriented path, then H has tree duality.*

In fact, Exercise 25 shows that the set \mathcal{F} from the definition of tree duality (the certificates of noncolourability) can be chosen to consist of oriented *paths*. We will prove the theorem by showing that *if the consistency check succeeds then an H -colouring exists*. (This implies the result via Theorem 5.22).

Proof Suppose that the vertices of the oriented path H are consecutively named $1, 2, \dots, n$. If the consistency check for an input graph G (with all lists equal to $V(H)$) succeeds, then each final list $L^*(u)$ will be nonempty, and will consist of some integers between 1 and n . We now claim that choosing for each $f(u)$ the smallest integer in $L^*(u)$ will define a homomorphism f of G to H . If this were not the case, then some edge $uu' \in E(G)$ would have $i = \min(L^*(u))$, $i' = \min(L^*(u'))$, and $ii' \notin E(H)$. Recall also that the lists L^* are consistent, thus some $j' > i'$ has $ij' \in E(H)$, and some $j > i$ has $ji' \in E(H)$. Finally, observe that the consecutive numbering of the vertices of the path H ensures that if $k\ell$ is an arc of H , then $k - 1 \leq \ell \leq k + 1$. Hence we have

$$j \leq i' + 1 \leq j' \leq i + 1 \leq j,$$

whence $i = i'$ and $j = j'$. This is impossible, since H has no symmetric arcs ij, ji . \square

In particular, we obtained a polynomial time algorithm for the H -colouring problem, when H is any oriented path.

This idea can be obviously extended to any digraph H which admits an ordering of the vertices $1, 2, \dots, n$ with the following property, called the *X -underbar property*.

- If ij', ji' are arcs of H and if $i < j, i' < j'$, then ii' is also an arc of H .

Note that the X -underbar property can be restated as follows: *if ij and $i'j'$ are arcs of H then $\min(i, i')\min(j, j')$ is also an arc of H* .

The ordering of the vertices of H is also called an *X -underbar enumeration* of H . Thus for any fixed digraph H with an X -underbar enumeration we have the following linear time H -colouring algorithm.

Algorithm 3

Input: A digraph G .

Task: Find an H -colouring of G , if one exists.

Action: Apply the consistency check (Algorithm 2). If the check fails, no H -colouring exists. If the check succeeds, let each $f(u)$ be the smallest element of $L^*(u)$.

Corollary 5.25 *If H has an X -underbar enumeration, then H has tree duality.*

There are several equivalent mechanisms for defining digraphs that have tree duality. We focus on the following property. Let $\mathcal{P}(H)$ be the digraph whose vertices are nonempty subsets of $V(H)$, and two subsets X, X' of $V(H)$ are adjacent vertices of $\mathcal{P}(H)$ if for each $x \in X$ there is an $x' \in X'$, and for each $x' \in X'$ an $x \in X$, such that $xx' \in E(H)$.

Proposition 5.26 *H has tree duality if and only if $\mathcal{P}(H) \rightarrow H$.*

Proof Suppose ϕ is a homomorphism of $\mathcal{P}(H)$ to H . We can use ϕ to define a homomorphism f of G to H whenever the consistency check succeeds: for each $u \in V(G)$, let $f(u) = \phi(L^*(u))$. It follows from the definitions that f is a homomorphism. Conversely, suppose H has tree duality. Let $G = \mathcal{P}(H)$, and let each list $L(X) = V(H)$. It is easy to see that after applying the consistency check each $L^*(X)$ will contain the set X (adjacency in $\mathcal{P}(H)$ is defined in such a way that the lists $L'(X) = X$ are consistent, so nothing in these lists will be eliminated during the consistency check). Since the consistency check succeeds, we know from Theorem 5.22 that there is a homomorphism from $G = \mathcal{P}(H)$ to H . \square

Proposition 5.26 is a decision procedure for tree duality. Even though much of the above discussion can be extended to bounded treewidth duality, no decision procedure for treewidth k duality is known apart from the above case of $k = 1$. Moreover, the homomorphism $\phi : \mathcal{P}(H) \rightarrow H$, which can be viewed as a *certificate* that H has tree duality, is useful in actually finding a homomorphism of G to H when the consistency check succeeds, as seen in the proof.

We remark in passing that Theorem 5.20 allows us to ‘predict’ theorems, using the general expectation that $P \neq NP$. For instance, we have seen that nonbipartite graphs H yield NP -complete H -colouring problems. This means, in view of Theorem 5.20, that we should not expect nonbipartite graphs to have bounded treewidth duality. The predicted theorem is the following.

Theorem 5.27 *Let H be a nonbipartite graph. Then H does not have bounded treewidth duality.*

Since we have seen that a bipartite graph has treewidth two duality, this result completely classifies graphs that have bounded treewidth duality. The corresponding problem for digraphs is wide open.

Proof We seek, for any integer k , a graph G_k that is not H -colourable, but which has the property that all G_k -colourable graphs of treewidth k are H -colourable. To ensure that G_k is not H -colourable, it suffices take G_k to be a

graph of chromatic number greater than H . Since H is nonbipartite, it contains an odd cycle, say $C_{2\ell+1}$. It turns out that by taking G_k of high girth, we can ensure that all G_k -colourable graphs of treewidth k are $C_{2\ell+1}$ -colourable, and hence H -colourable. (Recall that there exist graphs of arbitrarily high girth and chromatic number, Corollary 3.13.) The following proposition is the crucial result used to complete the proof, and is of independent interest. \square

Proposition 5.28 [280] *There is a function $g(k, \ell)$ such that if G is a graph of girth greater than $g(k, \ell)$, then any graph of treewidth k that is G -colourable is also $C_{2\ell+1}$ -colourable.*

5.5 Pair consistency and majority functions

Let us now return to bounded treewidth duality and higher order consistency checks. Consistency could be called ‘single consistency’, in the sense that we are testing whether lists of single vertices are consistent over pairs of vertices. (Indeed, only adjacent pairs pose any restrictions, so checking for all pairs amounts to the same thing as checking over arcs.) There is a natural extension of the concept, called *k-tuple consistency*, in which k -tuples of vertices are tested for consistency over $(k+1)$ -tuples. For our purposes, we shall focus on the case $k = 2$, and talk about *pair consistency*. (Everything we discuss extends in an obvious way to higher values of k .)

Assume that each pair of distinct vertices u, u' of G has a pair list $L(u, u') \subseteq V(H) \times V(H)$. We will always assume that if $(x, x') \in L(u, u')$ then the mapping taking u to x and u' to x' is a homomorphism of the subgraph of G induced by $\{u, u'\}$ to H . (In other words, any arcs and loops present amongst u and u' correspond to arcs and loops amongst x and x' .) We say that the pair lists L are *consistent* if for any three vertices $u, u', u'' \in V(G)$ and any $(x, x') \in L(u, u')$ there exists an $x'' \in V(H)$ with $(x, x'') \in L(u, u'')$ and $(x', x'') \in L(u', u'')$. In this case we will take the following initial pair lists $L(u, u')$. Consider the set F of all homomorphisms of the subgraph of G induced by u, u' to H , and let $L(u, u')$ consist of all pairs $f(u)f(u')$ with $f \in F$. Note that, if there are no edges of G on u, u' , then $L(u, u') = V(H) \times V(H)$; if u has a loop in G , then in any $(x, x') \in L(u, u')$, the vertex x has a loop in H ; and so on. The following generalization of the consistency check reduces the pair lists L to consistent pair lists L^* .

Algorithm 4 (Pair Consistency Check)

Input: A digraph G , with initial pair lists $L(u, u')$, as described.

Task: Reduce the pair lists to $L^*(u, u')$, that are consistent.

Action: Initially set all lists $L^*(u, u') = L(u, u')$, and then, as long as changes occur, process triples u, u', u'' of vertices of G as follows: remove from $L^*(u, u')$ any (x, x') for which no element $x'' \in V(H)$ has $(x, x'') \in L^*(u, u'')$ and $(x', x'') \in L^*(u', u'')$.

We say that the pair consistency check *fails* if some list $L^*(u, u')$ is empty; otherwise we say that the consistency check *succeeds*. It is again the case that

the pair consistency check can be performed in polynomial time, and that the check cannot fail if a homomorphism exists. We also have the following result.

Theorem 5.29 [171] *H has treewidth two duality if and only if $G \rightarrow H$ whenever the pair consistency check applied to G succeeds.*

We see that for digraphs H that enjoy treewidth two duality we can efficiently test H -colourability of input digraphs by running the polynomial time pair consistency check. If the check fails, we know the input digraph is not H -colourable. If the check succeeds, we know the input digraph is H -colourable; but how do we find an H -colouring?

At this point we have, for each vertex $u \in V(G)$, a list of candidate images in H . Consider any other vertex $u' \in V(G)$ and let

$$L^*(u) = \{x \in V(H) : (x, x') \in L^*(u, u') \text{ for some } x' \in V(H)\}.$$

(Because the pair lists L^* are consistent, it follows that $L^*(u)$ does not depend on which u' is chosen). Since the pair consistency check succeeded, each set $L^*(u)$ is nonempty. At this point it would be good if we could concisely describe how to choose a member of each $L^*(u)$ to obtain a homomorphism of G to H . (We were able to do that after the single consistency check using the homomorphism $F : \mathcal{P}(H) \rightarrow H$.) Unfortunately, no general technique for this is known, and all algorithms to find an H -colouring are somewhat more involved.

We first describe a class of graphs with treewidth two duality which admit a particularly simple kind of algorithm to find an H -colouring (following a successful pair consistency check).

Recall that we have introduced, in Chapter 2, a majority function for a reflexive graph H , and used it in Corollary 2.57 to characterize absolute retracts. A majority function can be defined for any digraph (or general relational system) H , by the same definition, as a polymorphism of order three, i.e., a homomorphism $\mu : H \times H \times H \rightarrow H$, satisfying the following condition:

- if at least two of the vertices a, b, c are equal to x then $\mu(a, b, c) = x$.

We shall show that the existence of a majority function for H implies that H -colouring is polynomially solvable, and that, in this case, there *is* a canonical way to define a homomorphism following the success of the pair consistency check.

Before we state this result, we note the following fact.

Lemma 5.30 *Let L^* be the pair lists produced by the pair consistency check, on an input graph G , and suppose that μ is a majority function on H .*

Then $(a, a'), (b, b'), (c, c') \in L^(u, u')$ implies that $(\mu(a, b, c), \mu(a', b', c')) \in L^*(u, u')$.*

Proof The initial pair lists L satisfy this property by the definition of μ . We show that the property is maintained by each step of the pair consistency check. Suppose that the current lists L^* satisfy the property, and the pairs $(a, a'), (b, b'), (c, c')$ are in the next version of the list $L^*(u, u')$ (after one step

of the algorithm). Then for any $u'' \in V(G)$, there exist a'', b'', c'' such that $(a, a''), (b, b''), (c, c'') \in L^*(u, u'')$ and $(a', a''), (b', b''), (c', c'') \in L^*(u', u'')$. Since the current lists satisfy the property in the lemma, we must have

$$\mu(a, b, c)\mu(a'', b'', c'') \in L^*(u, u'') \text{ and } \mu(a', b', c')\mu(a'', b'', c'') \in L^*(u', u''),$$

and therefore the pair $(\mu(a, b, c), \mu(a', b', c'))$ also survives to the next version of the list $L^*(u, u')$. \square

Theorem 5.31 *If H admits a majority function then H has treewidth two duality.*

Proof Suppose $L^*(v, v'), v, v' \in V(G), v \neq v'$, are the nonempty pair lists produced by the pair consistency check, on an input graph G . If f is a mapping of some subset U of $V(G)$ to $V(H)$, we say that f is *consistent*, if $f(v)f(v') \in L^*(v, v')$ for all $v, v' \in U, v \neq v'$.

Let E_k denote, for any integer $k \geq 2$, the following assertion.

- For any digraph G , any set U of k vertices of G , any consistent mapping $f : U \rightarrow V(H)$, and any vertex $u' \in V(G) - U$, there exists a consistent mapping f' of $U \cup \{u'\}$ to $V(H)$ which extends the mapping f .

We shall prove the theorem by proving that the existence of a majority function for H is equivalent to the conjunction of all the assertions E_k for $k = 2, 3, \dots, |V(H)| - 1$. This conjunction implies that a successful pair consistency check allows us to define a homomorphism starting from any two suitable values $((x, x') \in L(v, v')$ for any $v, v' \in V(G))$, introducing one new value at a time (cf. Algorithm 5). Note that a consistent mapping f defined on $U = V(G)$ is a homomorphism of G to H .

Suppose first that all assertions E_k hold. Let G be the product $H \times H \times H$ and U the set of all vertices (a, b, c) in which not all a, b, c are distinct. Finally, let $f(a, b, c) = x$ if at least two of a, b, c are equal to x . According to the assertions $E_{|U|}, E_{|U|+1}, \dots, E_{|V(H)|-1}$, there is a homomorphism of $H \times H \times H$ to H extending these values; clearly this homomorphism is a majority function μ on H .

Conversely, suppose that μ is a majority function on H . We shall use μ to prove all assertions E_k , by induction on k . The assertion E_2 follows trivially from the success of the pair consistency check. Thus suppose that $k \geq 3$ and that E_{k-1} has already been proved, and let U be any set of k vertices of a digraph G , $f : U \rightarrow V(H)$ a consistent mapping and u' a vertex of G not in U . Since $k \geq 3$, we can choose distinct vertices $u_1, u_2, u_3 \in U$. Let $U_i = U - u_i, i = 1, 2, 3$. By the induction hypothesis, each restriction $f|_{U_i}$ can be extended to a consistent mapping $f_i : U_i \cup \{u'\} \rightarrow V(H)$. We now set $f(u') = \mu(f_1(u'), f_2(u'), f_3(u'))$, and claim that the extended mapping $f : U \cup \{u'\} \rightarrow V(H)$ is also consistent. To see this, consider two distinct vertices v, v' in $U \cup \{u'\}$. Clearly, $f(v)f(v') \in L^*(v, v')$ if both v and v' are in U , so we may assume that, say, $v' = u'$. If $v \neq u_i$ for any $i = 1, 2, 3$, then $f(v)f(u') \in L^*(v, u')$ because $f(v) = \mu(f_1(v), f_2(v), f_3(v))$

and $f(u') = \mu(f_1(u'), f_2(u'), f_3(u'))$, and each $(f_i(v), f_i(u')) \in L'(v, u')$ (since f_i is consistent), cf. the discussion preceding the theorem. If $v = u_i$ for some $i = 1, 2, 3$, we proceed analogously, replacing f_i with any consistent mapping f'_i on $\{u_i, u'\}$ with $f'_i(u') = f_i(u')$. \square

Let U be a set of vertices of G . A mapping $f : U \rightarrow V(H)$ which cannot be extended to a homomorphism of G to H , but such that the restriction of f to any proper subset of U can be so extended, is called a *conflict on G with respect to H* . It follows from the above arguments that we have the following result.

Corollary 5.32 *A digraph H admits a majority function if and only if every conflict with respect to H has at most two vertices.*

By applying Theorem 5.20, we also have the desired result.

Corollary 5.33 *If H admits a majority function, then H -colouring is polynomial time solvable.* \square

Recall from Proposition 5.17 that having more polymorphisms makes a problem more likely to be polynomial time solvable. Majority functions are just one example of polymorphisms that are sufficient for polynomial solvability; they are particularly useful when H is a graph. For digraphs, and more general systems, there are other useful polymorphisms that guarantee polynomial solvability [53, 54].

In the case when H has a majority function, there is a canonical way to find a homomorphism following the success of the pair consistency check. Based on the above proof, we can use the following algorithm.

Algorithm 5

Input: A graph G with vertices u_1, u_2, \dots, u_n .

Task: Find an H -colouring of G , if one exists.

Action: Perform the pair consistency check (Algorithm 4), obtaining final pair lists $L^*(u, u')$. If the check fails (some pair list is empty), then G is not H -colourable. Otherwise find an H -colouring of G as follows. Define $f(u_1), f(u_2)$ so that the pair $(f(u_1), f(u_2))$ is in $L^*(u_1, u_2)$. Then, having defined f consistently on u_1, u_2, \dots, u_{i-1} , for $3 \leq i \leq n$, find a vertex $x \in V(H)$, so that letting $f(u_i) = x$ extends the definition of f consistently to u_1, u_2, \dots, u_i .

The proof of Theorem 5.31 ensures that this procedure is correct, i.e., that at each stage one can actually find a required vertex x .

We now return to the general case of a graph H with treewidth two duality and give a self-reduction procedure, obtained by modifying Algorithm 5, that will still find an H -colouring of the input graph, if one exists. (However, the modified algorithm is more complex and less efficient.) We may assume that H is a core. The **Action** begins as above, by performing the pair consistency check, and declaring that there is no H -colouring if the check fails. If the check succeeds, we proceed to assign images $f(u_1), f(u_2), \dots$. Say that an assignment

of $f(u_1) = x_1, f(u_2) = x_2, \dots, f(u_{i-1}) = x_{i-1}$ is *admissible*, if there is an H -colouring of the input graph extending this assignment. Having an admissible assignment $f(u_1) = x_1, f(u_2) = x_2, \dots, f(u_{i-1})$, we test all $x \in L^*(u_i)$ until we find an x such that the assignment $f(u_1) = x_1, f(u_2) = x_2, \dots, f(u_{i-1}) = x_{i-1}$, and $f(u_i) = x$, is also admissible. Since we know an H -colouring of the input graph exists, we will always be able to extend the current assignment so that it remains admissible, until an H -colouring of G is found. It remains to explain how to test whether an assignment $f(u_1) = x_1, f(u_2) = x_2, \dots, f(u_i) = x_i$, is admissible. This is where we use the fact that H is a core. Let G' be obtained from the input graph G and a copy of the graph H by identifying each vertex u_j with the corresponding vertex x_j , for all $j = 1, 2, \dots, i$. Then $G' \rightarrow H$ if and only if there is a homomorphism of G to H extending the assignment $f(u_1) = x_1, f(u_2) = x_2, \dots, f(u_i) = x_i$. Indeed, any homomorphism of G to H extending the given assignment is easily extended to G' . Conversely, if there is a homomorphism of G' to H , then there is such a homomorphism that maps the vertices of the copy of H included in G' identically to the vertices of H . (Since H is a core, the restriction to the copy of H in G' is an automorphism a of H , and we can compose the homomorphism with the inverse of a .) Therefore, the restriction of this homomorphism to G extends the given assignment. Of course, we can test the existence of a homomorphism $G' \rightarrow H$ by performing a pair consistency check, since H has treewidth two duality.

5.6 List homomorphisms and retractions

We now return to the study of list homomorphisms. Let H be a fixed digraph. Assume that for each vertex u of the input graph G we are given a list (set) $L(u) \subseteq V(H)$. A *list homomorphism* of G to H , or a *list H -colouring* of G , with respect to the lists L , is a homomorphism f of G to H , such that $f(u) \in L(u)$ for all $u \in V(G)$. The *list H -colouring problem* asks whether or not an input digraph G with lists L admits a list homomorphism to H with respect to L . We note that when $H = K_n$, a list homomorphism of G to H is a list colouring of G , and hence that the list K_n -colouring problem is essentially the usual list colouring problem (with a restricted set of colours). We also note that each H -colouring problem can be viewed as a list H -colouring problem restricted to inputs G with all lists $L(u) = V(H)$.

List H -colourings tend to be more manageable than H -colourings, since they offer a natural way to recurse. Seeking a list H -colouring of G we may choose a value for $u \in V(G)$ (from the list $L(u)$), modify correspondingly the lists of all neighbours of u (they can no longer use colours not adjacent to the value assigned to u) and then remove u from consideration. Similarly, if the list H -colouring problem is NP -complete, then so is any list H' -colouring problem where H is an induced subgraph of H' . (Restrict the inputs G to have their lists contained in $V(H)$.) Also, many natural applications of homomorphisms, such as frequency assignment, scheduling, and so on, tend to have additional constraints expressible by lists. Finally, it turns out that many algorithms for graph homomorphisms

adapt very naturally to lists. This is the case for virtually all the algorithms we have discussed, and is particularly plain for consistency check, which introduces lists even if lists were not originally present.

In what follows we shall focus mostly on reflexive graphs, in some cases expanding our focus to include all graphs with loops allowed (cf. also Exercise 18 dealing with just graphs).

A graph (with loops allowed) H is *chordal* if it does not contain an induced cycle of length greater than three. Equivalently, a graph H is chordal if and only if its vertices admit a *perfect elimination ordering*, i.e., an ordering $V(H) = \{v_1, v_2, \dots, v_n\}$ such that any two neighbours v_j, v_k of v_i with $i < j, i < k$, are adjacent to each other [123].

An *interval graph* is a graph H whose vertices can be represented by intervals on the real line, such that two vertices are adjacent in H if and only if the corresponding intervals intersect. Note that by this definition interval graphs are reflexive, cf. Fig. 1.11. (It is more common to define adjacency by intersection only for distinct vertices, but for our purposes we need the above definition.) It is easy to see that a reflexive cycle of length greater than three is not an interval graph. Since an induced subgraph of an interval graph must be an interval graph, we conclude that each interval graph is chordal. The following theorem specifies which chordal graphs are interval graphs. The forbidden configuration turns out to be an *asteroidal triple*: three pairwise nonadjacent vertices a, b, c , and for each pair of them a path joining them without containing a neighbour of the third vertex.

Theorem 5.34 [218] *Let H be a reflexive chordal graph. Then H is an interval graph if and only if it does not contain an asteroidal triple.*

In particular, a reflexive graph is an interval graph if and only if it does not contain an induced cycle of length greater than three or an asteroidal triple.

We have the following classification of the complexity of list H -colouring problems for reflexive graphs.

Theorem 5.35 *Let H be a reflexive graph.*

- *If H is an interval graph then the list H -colouring problem has a polynomial time algorithm.*
- *Otherwise the problem is NP-complete.*

Proof To show that list H -colouring problems for interval graphs H have polynomial time algorithms, we shall prove that each interval graph H admits a majority function, with an important additional property.

We shall say that a majority function is *conservative* if $\mu(a, b, c) \in \{a, b, c\}$, for all a, b, c . A conservative majority function on H yields a polynomial time list H -colouring algorithm along the lines of Algorithm 5. We only have to change the definition of a mapping f being *consistent* to include, in addition to the earlier condition, that each $f(u) \in L(u)$. With this modification, and starting

the pair consistency check with the initial lists being the given lists $L(u)$, Algorithm 5 applies directly. The crucial observation, in proving the assertions E_k using the existence of a conservative majority function μ on H , is that the definition of $f(u')$ as $\mu(f_1(u'), f_2(u'), f_3(u'))$ ensures that it is equal to one of $f_1(u'), f_2(u'), f_3(u')$, since μ is conservative; this easily implies that the extension is also consistent, in this new sense, and proves the following fact.

Lemma 5.36 *If H admits a conservative majority function then the list H -colouring problem has a polynomial time algorithm.* \square

We now continue with the proof of Theorem 5.35. According to the lemma, it will suffice to show that each interval graph H has a conservative majority function. Consider $x, y, z \in V(H)$, and the corresponding intervals I_x, I_y, I_z . Let I be the interval whose left endpoint is the middle point of the three left endpoints of I_x, I_y, I_z , and whose right endpoint is the middle point of the three right endpoints of I_x, I_y, I_z . One of the three intervals I_x, I_y, I_z , say I_t , must contain I , since two of the three intervals have left endpoint not to the right of the left endpoint of I , and two of the three intervals have right endpoint not to the left of the right endpoint of I . Note that t is one of the three vertices x, y, z , and we define $\mu(x, y, z) = t$. It is clear that if at least two of x, y, z are the same, say equal to a , then $t = a$. It remains to show that if $xx', yy', zz' \in E(H)$, then $tt' \in E(H)$ for $t = \mu(x, y, z)$ and $t' = \mu(x', y', z')$. Suppose to the contrary that I_t is disjoint from $I_{t'}$, say I_t precedes $I_{t'}$. Then at least two of I_x, I_y, I_z precede at least two of $I_{x'}, I_{y'}, I_{z'}$, so either I_x is disjoint from $I_{x'}$, or I_y from $I_{y'}$, or I_z from $I_{z'}$, a contradiction. This proves the first assertion of Theorem 5.35.

To prove the second assertion, consider a graph H which is not an interval graph. Thus H contains an induced cycle of length greater than three or an asteroidal triple. We first illustrate the proof of NP -completeness in the case H contains an induced (reflexive) fourcycle C (Fig. 5.12). In this case, we reduce the problem of graph four-colourability to the list H -colouring problem as follows. Given a graph G , we let G' be the graph $G * J$ where J is the replacement graph shown in the figure, obtained from C by adding one edge, and with the shown connector vertices i, j . In other words, G' is obtained from G by replacing each edge xy of G with a copy J_{xy} of J , identifying x with i and y with j . We then let G'' be obtained from G' by adjoining a copy C_0 , of C and connecting the copy of C in each J_{xy} to this special C_0 by a chain of three other copies of C , as illustrated in the figure. (These copies of C are disjoint for each J_{xy} , except for the C_0 which is common to all of them.) We now define lists on G'' as follows: for each vertex v of C_0 we let $L(v) = \{v\}$; for all other vertices u of G'' we let $L(u) = V(C_0)$. Consider now any list homomorphism f of G'' to H . Because of the lists, all images $f(x)$ are amongst the four vertices of C_0 . The images $f(u)$ of the vertices u in C_0 are fixed. For any chain of cycles C leading from C_0 up to some J_{xy} , each consecutive cycle in the chain must map to the previous cycle either directly, taking each vertex x to the corresponding vertex x , or it must 'rotate' to the right, taking each vertex x to $x + 1 \pmod{4}$. This ensures that the

images $f(i')$ and $f(j)$ have distance two in C , and hence that $f(x)$ and $f(y)$ are different. Therefore, G is four-colourable. On the other hand, each four-colouring of G gives rise to a list homomorphism of G'' to H , since each chain of copies of C has three members, and hence allows i' and j to map to any two vertices of distance two in C_0 by a suitable number of rotations, and so allows i and j to map to any two distinct vertices of C_0 .




FIG. 5.12. The construction when H is a four-cycle.

Returning to the general case, when C is a cycle of length $k = 2\ell$, we proceed as above, taking for J a copy of C with two opposite vertices i' and j , and with i joined to i' by a path of length $\ell - 1$. Each copy of C in a J_{xy} is joined to the special cycle C_0 by a chain of $k - 1$ copies of C . Exactly as above, we reduce k -colourability to the problem of list G'' -colouring. When C is a cycle of length $k = 2\ell + 1$, we proceed similarly, except J is obtained from C by adding a *new* vertex i' adjacent to the two consecutive vertices of C , opposite to j .

Next, suppose that H contains an asteroidal triple $0, 1, 2$ of vertices with paths P_{01}, P_{02}, P_{12} disjoint, respectively, from the neighbourhoods of $2, 1, 0$. We may assume without loss of generality that these paths P_{xy} are shortest paths in H , and in particular that x and y have unique neighbours in P_{xy} (for $x, y \in \{0, 1, 2\}$). We shall reduce NOT-ALL-EQUAL 3-SAT without negated variables to the list H -colouring problem.

Let x, y be distinct elements of $\{0, 1, 2\}$ and let X, Y be subsets of $\{0, 1, 2\}$. A *chooser with inputs x, y and corresponding output sets X, Y* is a path P from

the *startpoint* a of P to the *endpoint* b of P , with lists $L(p) \subseteq V(H)$, for all $p \in V(P)$, such that

- any list homomorphism f of P to H has $f(a) = x$ and $f(b) \in X$ or $f(a) = y$ and $f(b) \in Y$; and
- for any $x' \in X$ and $y' \in Y$ there is a list homomorphism f of P to H with $f(a) = x$ and $f(b) = x'$ and a list homomorphism g of P to H with $g(a) = y$ and $g(b) = y'$.

To show the correspondence between the inputs and their associated output sets, we shall write $X = o(x)$ and $Y = o(y)$.

Suppose $P_i, i = 1, 2, 3$ are choosers with inputs $0, 1$ and output sets

$o(0) = \{0, 1\}, o(1) = \{1, 2\}$ for P_1 ,

$o(0) = \{1, 2\}, o(1) = \{0, 2\}$ for P_2 ,

$o(0) = \{0, 2\}, o(1) = \{0, 1\}$ for P_3 .

Let T be the tree obtained from P_1, P_2, P_3 by identifying the three vertices b and distinguishing the respective startpoints a by calling them a_1, a_2, a_3 (where a_i is the a of P_i). The lists of T are those inherited from the P_i 's. We note that

- no list homomorphism f of T to H has $f(a_1) = f(a_2) = f(a_3)$, and
- for any three *distinct* values $b_1, b_2, b_3 \in \{0, 1\}$, there exists a list homomorphism f of T such that $f(a_i) = b_i$ for $i = 1, 2, 3$.

Using this tree T we will be able to give a polynomial reduction from Not-All-Equal Three-Satisfiability without negated variables to the list H -colouring problem. It suffices to replace each clause with variables x, y, z by a separate copy of T , identifying a_1 with x , a_2 with y , and a_3 with z . Clearly, the resulting graph admits a list H -colouring if and only if the instance is satisfiable.

It remains to construct the three choosers P_i . We first construct some auxiliary choosers.

Consider the following chooser $P(0, 1, 2)$ with inputs $0, 1$ and output sets $o(0) = \{0, 2\}, o(1) = \{1, 2\}$. Let ℓ be the larger of the lengths of P_{02}, P_{12} , and let A be the set of all vertices in P_{01}, P_{12}, P_{02} except those adjacent to 0 or 1 . (Recall that there are at most two vertices adjacent to 0 and at most two vertices adjacent to 1 on these paths.) Let $0^+, 1^+$ be the first vertices on P_{02}, P_{12} after $0, 1$ respectively. We take $P(0, 1, 2)$ to be a path of length ℓ from a startpoint a to an endpoint b ; we also denote by a^+ the first vertex of $P(0, 1, 2)$ after a . We define the lists on $P(0, 1, 2)$ as follows: $L(a) = \{0, 1\}, L(a^+) = \{0^+, 1^+\}, L(b) = \{0, 1, 2\}$, and $L(p) = A$ for all other $p \in V(P(0, 1, 2))$. This is indeed a chooser with $o(0) = \{0, 2\}$ and $o(1) = \{1, 2\}$. Any list homomorphism taking a to 0 must take a^+ to 0^+ since 0 has no neighbours in P_{12} . If the next vertex of $P(1, 2, 3)$ is taken to 0 then all vertices of $P(1, 2, 3)$ have to map to 0 , including b . Otherwise neither 0 nor 1 can be images of any vertex of $P(1, 2, 3)$ because of the choice of A . There are clearly list homomorphisms taking a to 0 and taking b to 0 , and taking b to 2 , since ℓ is large enough. The output set of 1 is analysed similarly.

Analogous constructions provide us with choosers $P(x, y, z)$ with inputs x, y and output sets $o(x) = \{x, z\}, o(y) = \{y, z\}$, for any permutation xyz of 012 .

Next we construct a chooser $P(0, 1)$ with inputs $0, 1$ and output sets $o(0) = \{0\}, o(1) = \{2\}$. Let m be the length of the path P_{12} and let $B = V(P_{12})$. Recall that 0 is not adjacent to any vertex of B . We take $P(0, 1)$ to be a path of length m from a startpoint a to an endpoint b with all lists $L(p) = B$, except for $L(a) = \{0, 1\}$ and $L(b) = \{0, 2\}$. It is easy to check that $P(0, 1)$ is the required chooser.

Analogous constructions provide us with choosers $P(x, y)$ with inputs x, y and output sets $o(x) = x, o(y) = z$ for any permutation xyz of 012 .

We now concatenate $P(0, 1)$ and $P(0, 2, 1)$, by identifying the endpoint of $P(0, 1)$ with the startpoint of $P(0, 2, 1)$; clearly this is a chooser with the same inputs as $P(0, 1)$, i.e., 0 and 1 . To find the output set $o(0)$ we consider the output 0 of $P(0, 1)$ as the input of $P(0, 2, 1)$ and obtain the corresponding output set $\{0, 1\}$; similarly, $o(1) = \{1, 2\}$. We let P_1 be this concatenation.

Similarly, we let P_2 be the concatenation of $P(0, 1), P(2, 0), P(1, 2)$, and $P(0, 1, 2)$, in that order, and we let P_3 be the concatenation of $P(1, 0)$ and $P(1, 2, 0)$. It is easily checked that these concatenations have the required properties. \square

Note that the lemma and the theorem predict the following result. *A reflexive graph H admits a conservative majority function if and only if H is an interval graph*, (Exercise 8). (See our remarks, preceding Theorem 5.27, on using NP -completeness to predict results.)

The above NP -completeness proof for reflexive graphs containing an induced cycle of length greater than three only uses lists that are either the entire cycle C or a single vertex of C . A natural restriction of the list H -colouring problem requires that each list be connected. The *connected list H -colouring problem* is the restriction of the list H -colouring problem to instances G such that each list $L(u), u \in V(G)$, induces a connected subgraph of H . Thus the above proof actually shows that for reflexive graphs the connected list H -colouring problem is NP -complete unless H is a chordal graph. For reflexive chordal graphs, the problem is polynomial time solvable.

Theorem 5.37 *Let H be a reflexive graph.*

- *If H is chordal then the connected list H -colouring problem is polynomial time solvable.*
- *Otherwise, the connected list H -colouring problem is NP -complete.*

Proof As noted above, our previous NP -completeness proof implies the second assertion. We now describe a polynomial time algorithm proving the first assertion. Thus, let H be a fixed reflexive chordal graph; we fix a particular perfect elimination ordering of H , say x_1, x_2, \dots, x_k . (Therefore, if x_i is adjacent to x_j and to x_k , with $i < j, i < k$, then x_j and x_k are adjacent to each other.)

Algorithm 6

Input: A graph G , with connected lists $L(u) \subseteq V(H)$, $u \in V(G)$.

Task: Find a list H -colouring of G (with respect to L).

Action: Initially we set all lists $L^*(u) = L(u)$. We process the vertices of H in the perfect elimination order: when x_i is processed, we remove it from all lists $L^*(u)$ which contain at least one other vertex; for lists $L^*(u) = \{x_i\}$, we remove all vertices nonadjacent to x_i from each list $L^*(u')$ where u' is adjacent to u .

Once again, we say that the algorithm *succeeds* if all final lists are nonempty, and say that it *fails* otherwise. Let us say that the current lists L^* (during the execution of the algorithm) are *feasible* if there exists a list homomorphism of G to H with respect to L^* . We shall show that each step of the algorithm maintains the following invariant.

- The lists L^* are connected, and L^* are feasible if and only if the initial lists L are feasible.

The invariant certainly holds for the initial lists $L^* = L$. Suppose the lists L^* satisfy the invariant, and we make one step during the processing of x_i . Here a step means modifying one list $L^*(u)$ —either by removing x_i from $L^*(u)$ (if $L^*(u)$ contains at least one other vertex), or by removing from $L^*(u)$ all nonneighbours of x_i where $L^*(u') = \{x_i\}$ and $uu' \in E(G)$. Note that in the second case the removed vertices x_j must have $i < j$, as otherwise x_j could only be in the list of u if $L^*(u) = \{x_j\}$, and then x_i would have been removed from $L^*(u')$ during the processing of x_j .

If the step made was removing x_i from $L^*(u)$, then note that the neighbourhood of x_i in $L^*(u)$ is complete, and so the deletion of x_i does not disconnect $L^*(u)$. Clearly, if the modified L^* is feasible, then so is the (larger) unmodified L^* . So suppose the unmodified L^* was feasible, and a list homomorphism f of G to H with respect to the unmodified L^* assigned $f(u) = x_i$. We now define a new list homomorphism f' of G to H with respect to the modified lists. We only change the image $f(u) = x_i$ to $f'(u) = x_k$, for any x_k in $L^*(u)$ adjacent to x_i . First of all, such an x_k must exist since $L^*(u)$ was connected. Second, we claim that f' is a homomorphism. Indeed, each neighbour x_j of x_i with $j > i$ is also a neighbour of x_k (this applies also to $x_j = x_k$ since H is reflexive). Thus any edge uv of G which was taken by f to an edge $x_i x_j$ of H with $j > i$, is now taken by f' to the edge $x_k x_j$ of H . Furthermore, if any edge uv of G was mapped by f to an edge $x_i x_j$ of H with $j < i$, then $f(v) = x_j$ was assigned after x_j has been processed, and hence we must have $L^*(v) = \{x_j\}$; therefore $L^*(u)$ only contains neighbours of x_j at this point, and so x_k is adjacent to x_j . Hence f' maps the edge uv to the edge $x_k x_j$. Hence f' is a homomorphism, and hence a list homomorphism with respect to the modified lists L^* .

Thus assume the step made was removing from $L^*(u)$ all nonneighbours of x_i , where some neighbour u' of u has $L^*(u') = \{x_i\}$. Then the remaining vertices of the modified $L^*(u)$ are all neighbours x_j of x_i with $j > i$ and hence adjacent to each other, so the modified $L^*(u)$ is again connected. The modification results

in feasible lists if and only if the unmodified lists were feasible (and hence if and only if L are feasible) because a nonneighbour of x_i cannot be assigned to u if x_i is assigned to u' .

This completes the proof of the validity of the invariant throughout the algorithm; at the end of the algorithm the lists L^* either include an empty list, in which case the final lists, and hence the initial lists L , are not feasible, or the final lists define a list H -colouring. (Note that if adjacent vertices u, u' of G were assigned to x_i, x_j with $i < j$ respectively, then the fact that x_j was not removed from $L^*(u')$ implies that x_i, x_j are adjacent in H .) \square

A similar algorithm, directly employing the consistency check is given in Exercise 30.

For graphs (without loops), the dichotomy of list homomorphism problems, analogous to Theorem 5.35, involves (the complements of) circular arc graphs, in place of interval graphs (Exercise 18). For graphs with loops allowed, the general dichotomy result is stated below; it involves the following concept. Let C be a fixed circle with two specified points p and q . A graph H with loops allowed is called a *bi-arc graph*, if each vertex x of H corresponds to a pair of arcs (N_x, S_x) on C , with N_x containing p but not q and S_x containing q but not p , so that for any $x, y \in V(H)$, not necessarily distinct, the following hold:

- if x and y are adjacent, then neither N_x intersects S_y nor N_y intersects S_x ;
 - if x and y are not adjacent, then both N_x intersects S_y and N_y intersects S_x .
- (Note that we cannot have $(N_x, S_x), (N_y, S_y)$ such that N_x intersects S_y but S_x does not intersect N_y .)

Theorem 5.38 [99] *Let H be a fixed graph with loops allowed.*

- *If H is a bi-arc graph then the list H -colouring problem has a polynomial time algorithm.*
- *Otherwise the problem is NP-complete.*

We now return to the context of reflexive graphs. Recall the *retraction problem* H -RET. In the context of reflexive graphs, we can state it as follows. Given a reflexive graph G containing H as a subgraph, decide whether or not there is a retraction of G to H . Clearly, a retraction of G to H exists if and only if there is a list homomorphism of G to H , where all lists are $L(u) = V(H)$ except for vertices x of H , which have lists $L(x) = \{x\}$. Note that the NP-completeness construction in the proof of Theorem 5.35 also shows that if C is a reflexive cycle of length greater than three, then C -RET is NP-complete. (Cf. our remark following the proof of Theorem 5.35.) If H is a connected chordal graph, the problem H -RET is a restriction of the connected list H -colouring problem, and hence Algorithm 6 solves H -RET in polynomial time. The problem H -RET for disconnected chordal graphs H is also easily solved, by observing that a component of G cannot retract to two different components of H , and hence we can apply the H -RET algorithm for connected graphs to each component of H separately. Therefore, we have the following corollary.

Corollary 5.39 *Let H be a reflexive graph.*

- *If H is chordal then H -RET is polynomial time solvable.*
- *If H is a cycle of length greater than three then H -RET is NP-complete.*

□

Despite the fact that the the corollary appears close to a classification, it seems that dichotomy for retraction problems for reflexive graphs will be hard to obtain, cf. Exercise 29.

5.7 Trigraph homomorphisms

Recall the kind of vertex partition problems that can be expressed as graph homomorphism problems. Proposition 1.10 describes this in the context of digraphs. In the special case of graphs, we can seek partitions of $V(G)$ into parts S_x , where certain parts S_x are required to be independent, and certain pairs of parts $S_x, S_{x'}$ are required to have no edges joining them. By requiring the complement of the graph G to admit a homomorphism to the complement of H , we may also express the problems of finding vertex partitions into parts S_x , where certain parts S_x are required to be complete, and certain pairs of parts $S_x, S_{x'}$ are required to have all the edges joining S_x to $S_{x'}$. (In the list version of these problems, we equip the input graphs G with lists restricting the parts to which the vertices of G can be placed.)

A natural extension of such homomorphism problems seeks vertex partitions in which some parts may be required to be either independent or complete, and some pairs of parts may be required to have either no edges or all the edges. These partitions extend the partitions corresponding to homomorphisms (as in Proposition 1.10), and also extend the complemented partitions suggested above.

A natural way to model these partitions is to introduce the following generalization of graphs. A *trigraph* H consists of a set $V(H)$ of *vertices*, a set $W(H)$ of *weak edges*, and a disjoint set $S(H)$ of *strong edges*. We shall focus on (undirected) trigraphs with loops allowed, i.e., the (strong and weak) edges are two-element or one-element subsets of $V(H)$. Trigraphs will be pictured by drawing the strong edges thicker than the weak edges, cf. Fig. 5.13.

The *adjacency matrix* of a trigraph H is the symmetric matrix $M(H)$, with a row and a column for each vertex of H , and entries

- $M(x, x') = 0$ if xx' is not an edge of H ,
- $M(x, x') = *$ if xx' is a weak edge of H , and
- $M(x, x') = 1$ if xx' is a strong edge of H .

A *homomorphism* f of a graph G to a trigraph H is a mapping f of $V(G)$ to $V(H)$ such that the following two properties hold for vertices $u \neq u'$ of G :

- $uu' \in E(G)$ implies $f(u)f(u') \in W(H) \cup S(H)$, and
- $uu' \notin E(G)$ implies $f(u)f(u') \notin S(H)$.




FIG. 5.13. Example trigraphs.

(Homomorphisms of trigraphs to trigraphs can also be defined, and satisfy the expected properties, cf. Exercise 14.)

As always, a homomorphism of a graph G to a trigraph H will also be called an H -colouring of G .

In the remainder of this section, we shall often describe a trigraph H by its adjacency matrix $M = M(H)$. Thus, it is useful to formulate the H -colourability of G in terms of the matrix M .

Proposition 5.40 *Let G be a graph, and H a trigraph with adjacency matrix $M = M(H)$.*

Then G admits a homomorphism to H if and only if the vertices of G can be partitioned into sets $S_x, x \in V(H)$, such that the following properties are satisfied.

1. *if $M(x, x) = 0$, then S_x is independent in G ,*
2. *if $M(x, x) = 1$, then S_x is a clique in G ,*
3. *if $M(x, x') = 0$, then $uu' \notin E(G)$ whenever $u \in S_x$ and $u' \in S_{x'}$.*
4. *if $M(x, x') = 1$, then $uu' \in E(G)$ whenever $u \in S_x$ and $u' \in S_{x'}$.* □

A partition of G satisfying the properties 1–4 of Proposition 5.40 is called an M -partition of G .

Note that when the trigraph H has $S(H) = \emptyset$, an H -colouring of a graph G is simply an ordinary homomorphism of G to the graph consisting of the weak edges

of H . Hence homomorphisms to trigraphs extend homomorphisms to graphs, and we shall use the standard terminology introduced for homomorphisms, e.g. surjective homomorphisms, homomorphic images, etc. In particular, we note that the class of H -colouring problems, for fixed trigraphs H , extends the class of H -colouring problems for fixed graphs H . This simple generalization of the graph homomorphism framework has surprisingly strong modeling power. In addition to homomorphism problems, this framework includes, for instance, the following natural concepts.

1. A graph G is *split* if it can be partitioned into an independent set and a clique. Thus a graph is split if and only if it admits an H -colouring where H is the trigraph in Fig. 5.13(a). Split graphs are a well-understood and useful class of chordal graphs [123].
2. A *clique cutset* of a connected graph G is a clique of G whose removal disconnects G . Hence a graph has a clique cutset if and only if it can be partitioned into a clique and two nonempty sets joined by no edges. In other words, G has a clique cutset if and only if it admits a vertex-surjective H -colouring where H is the trigraph in Fig. 5.13(b). Clique cutsets can be recognized in polynomial time and lead to efficient optimization algorithms for graphs decomposable by clique cutsets [325, 340].
3. Surjective homomorphisms to the trigraphs in Fig. 5.13(d) and (e) correspond respectively to the existence of a so called *skew cutset* and *two-join*. These concepts (and their close relatives) play an important role in the proof of the Strong Perfect Graph Theorem [61].
4. Homomorphisms to the trigraph in Fig. 5.13(c) (with some additional restrictions on the size of preimages, explained below), correspond to the existence of a so called *homogeneous set* of vertices. *Modular decompositions* into homogeneous sets are a useful tool for efficient recognition of several graph properties [318].

We note that it is easy to capture additional requirements, that the homomorphisms be vertex-surjective, or edge-surjective, or have at least a certain number of preimages for specified vertices, or have at least some edges joining them, etc., by lists. For instance, the definition of a homogeneous set requires that $f^{-1}(c)$ has at least two vertices and that $f^{-1}(a) \cup f^{-1}(b)$ is nonempty; this can be done by choosing (in every possible way) three vertices u, v, w of the input graph and specifying that the lists of u, v only consist of c and the list of w consists of a, b . Thus the problem of finding a homogeneous set in a graph with n vertices is reduced to $n(n-1)(n-2)$ list H -colouring problems. (A homogeneous set exists if and only if at least one of the $n(n-1)(n-2)$ choices of u, v, w has a desired list H -colouring). Analogously, one can ensure that there is at least one edge between $f^{-1}(x)$ and $f^{-1}(y)$ by restricting (by the choice of lists) two adjacent vertices u, u' to be mapped to x, y respectively, for all possible choices of an edge uu' in the input graph.

Therefore, we shall consider, for each fixed trigraph H , the following *list H -colouring problem*: given a graph G , with lists $L(u) \subseteq V(H)$, $u \in V(G)$, is there a homomorphism f of G to H such that $f(u) \in L(u)$ for all $u \in V(G)$?

While the known polynomial cases of the H -colouring problem for graphs H seem to follow the same few patterns (tree duality, majority functions), a variety of different algorithm paradigms seems to become relevant in the case of trigraphs H . In the remainder of this section, and in the next section, we shall try to illustrate this wealth of algorithmic ideas.

One of the most useful technique for solving list H -colouring problems for small trigraphs H is the two-satisfiability algorithm.

Suppose first that the fixed trigraph H has only two vertices, say a and b . We can solve the list H -colouring problem by introducing a boolean variable x_v for each vertex v of the input graph G ; we think of the value of x_v as encoding whether or not the vertex v is mapped to a ($x_v = 1$ means v is mapped to a , $x_v = 0$ means v is mapped to b). It is then easy to see that all the constraints, and lists, of the list H -colouring problem can be stated by polynomially many clauses with at most two literals each. For instance, if $f^{-1}(a)$ is to be an independent set ($M(a, a) = 0$), we impose the constraint $\bar{x}_u \vee \bar{x}_v$ for each edge uv of G . Similarly, if, say, every edge from $f^{-1}(a)$ to $f^{-1}(b)$ is to be present ($M(a, b) = 1$), we impose the constraints $x_u \vee \bar{x}_v$ and $x_v \vee \bar{x}_u$ for each nonedge uv of G . Finally, if the list of v is, say, b , we impose the constraint \bar{x}_v . Hence the problem can now be solved by any two-satisfiability algorithm such as [6].

For instance, we can solve this way the *pre-colouring split extension problem*, in which an input graph G is to be partitioned into a clique and an independent set, with some vertices pre-coloured, i.e., placed in the clique or independent set to be constructed. The two parts are respectively complete and independent, and we correspondingly denote them by C and I . Each vertex v corresponds to a variable x_v , and $x_v = 1$ means $v \in C$, $x_v = 0$ means $v \in I$. We obtain one constraint for each edge $vw \in E(G)$, namely $x_v \vee x_w$, and one constraint for each nonedge $vw \notin E(G)$, namely $\bar{x}_v \vee \bar{x}_w$. Moreover, if a vertex v has a restricted list, we add the constraint x_v or \bar{x}_v .

The same technique applies whenever we have an instance in which every list has size at most two.

Proposition 5.41 *For any fixed trigraph H , there is a polynomial time algorithm which solves the list H -colouring problem restricted to instances in which the list of every vertex of the input graph has size at most two.* \square

To illustrate the usefulness of Proposition 5.41, we shall give an algorithm to solve the list version of the clique cutset problem. Recall the clique cutset problem from Fig. 5.13(b); in the list version, we have an input graph G with lists (subsets of $\{a, b, c\}$), and we seek a list homomorphism of the graph G to the trigraph in Fig. 5.13(b). In other words, we seek a partition of the vertex set of G into three sets A, B, C where C is a clique of G , and there are no edges in G between the parts A and B . To simplify the notation, we shall assume

the lists are be subsets of $\{A, B, C\}$ and each vertex of G has to be placed in a part belonging to its list. We call this the *list clique partition problem*. We begin by recalling that there is a polynomial time algorithm to find a *minimal chordal extension* H of a graph G [302]. This is a graph H which contains G as a subgraph, with $V(G) = V(H)$, and with $E = E(H) - E(G)$ inclusion minimal, in the sense that $G \cup E'$ is not chordal for any $E' \subset E, E' \neq E$. It is easy to see that for any partition A, B, C of G , the minimal chordal extension H will not contain any edges between the sets A and B . Indeed, H with all these edges removed would still be a chordal extension of G , since any cycle in it that contains both a vertex of A and a vertex of B goes twice through the clique C and hence must have a chord.

Let x_1, x_2, \dots, x_n be a perfect elimination ordering of H . Let F_i be the *forward clique* of H , consisting of x_i and all adjacent $x_j, j > i$. The fact that each F_i is a clique follows from the definition of a perfect elimination ordering. Moreover, each clique A of H (and hence also each clique A of G) is a subset of some F_i —it is enough to let i be the lowest subscript of any $v_i \in A$. In conclusion, the sets F_i have the following property: for any partition A, B, C , there exists a set F_i that contains C and is disjoint from either A or B . (Indeed, F_i , being a clique of H , cannot contain both a vertex of A and a vertex of B .)

We can now state our algorithm.

Algorithm 7

Input: A graph G , with lists $L(u) \subset \{A, B, C\}, u \in V(G)$.

Task: Find a partition $V(G) = A \cup B \cup C$ where C is a clique, there are no edges between A and B , and each u belongs to a part $X \in L(u)$.

Action:

Compute a minimal chordal extension H of G .

For each forward clique F_i of H try to find a list partition A, B, C in which F_i contains C and is disjoint from A , or contains C and is disjoint from B . (The former is tested by deleting A from all lists of vertices in F_i and deleting C from all lists of vertices not in F_i ; the latter is tested by deleting A from all lists of vertices in F_i and deleting C from all lists of vertices not in F_i . In both cases we have lists of size at most two and we can use Proposition 5.41.)

If all tests fail, no list partition exists.

Corollary 5.42 *The list clique partition problem is solvable in polynomial time.* □

When there are no lists but we require each part A, B, C to be nonempty, we can proceed as above, testing for a list clique cutset for $n(n-1)(n-2)$ choices of vertices x, y, z for which we set $L(x) = \{A\}, L(y) = \{B\}, L(z) = \{C\}$ (all other lists being set to $\{A, B, C\}$). We have shown the following fact.

Corollary 5.43 *There is a polynomial time algorithm to test whether or not a graph has a clique cutset.* □

5.8 Generalized split graphs

Consider the following generalization of split graphs. An (a, b) -graph is a graph whose vertices can be partitioned into a independent sets and b cliques. Such a partition will be called an (a, b) -partition of G . Let M be the symmetric $(a+b) \times (a+b)$ matrix with all off-diagonal entries $*$, and having a 0's and b 1's on the diagonal. Note that M is the adjacency matrix of a trigraph H which contains a complete graph on $a+b$ vertices, with all edges being weak edges, and with exactly b strong loops, and no weak loops. Then G is an (a, b) -graph if and only if it admits a homomorphism to H , and if and only if it admits an M -partition.

The following result classifies the complexity of the recognition problem for (a, b) -graphs. We state it in the terms of M -partitions, to make the role of lists clearer.

Proposition 5.44 *If both $a \leq 2$ and $b \leq 2$, then the list M -partition problem is solvable in polynomial time. Otherwise the list M -partition problem is NP-complete.*

Proof First we note that if $a \geq 3$ then the M -partition problem is NP-complete, since we can decide whether or not an input graph G is three-colourable by endowing all its vertices with the list $\{1, 2, 3\}$, and then asking whether or not it has a list M -partition. (If $b \geq 3$ the proof is similar.)

We illustrate one technique for recognizing $(1, 1)$ -, $(2, 1)$ -, $(1, 2)$ -, $(2, 2)$ -graphs. For simplicity we consider only the case of $(2, 1)$ -graphs; the other cases are treated similarly. Note that a $(2, 1)$ -partition of G can be viewed as a partition of the vertices of G into an induced bipartite subgraph and a clique. By focusing on the union of the two independent sets (the bipartite subgraph) we emphasize the fact that we shall treat their vertices in the same way. In fact, we shall repeatedly use the obvious fact that a bipartite graph and a clique can meet in at most two vertices.

We claim that a graph G on n vertices has at most n^4 different $(2, 1)$ -partitions, and all these partitions can be found in time proportional to n^8 .

Let $V(G) = B \cup C$ be a particular $(2, 1)$ -partition. Then any other $(2, 1)$ -partition $V(G) = B' \cup C'$ has $|B' \cap C| \leq 2$ and $|B \cap C'| \leq 2$, so B' is obtained from B by deleting at most two vertices and inserting at most two new vertices. In fact, if we allow ourselves to insert back a vertex that has just been deleted, we can say that we make exactly two deletions and exactly two insertions. Each of these at most four operations can be made in at most n ways. This observation proves the first part of the claim and allows us to find *all* $(2, 1)$ -partitions if one such partition is known. It amounts to a four-local search (the current S is changed in at most four vertices), and can be performed in time n^4 .

It remains to explain how to find the first $(2, 1)$ -partition. The algorithm proceeds in two phases. The first phase attempts to find as large a bipartite subgraph as possible. This is based on the observation that if $V(G) = B \cup C$ is a $(2, 1)$ -partition and B' a bipartite graph smaller than B , then $B' \cap C$ has at most two vertices, and hence, as above, B' can be enlarged by removing some two

vertices and inserting some three new vertices. Thus, starting with any bipartite subgraph of G , we can increase its size by performing a five-local search (making all possible two deletions and three insertions and testing if the result is bipartite) in time n^5 times the time to check whether a graph is bipartite. After performing this operation at most n times, we reach a situation where the current bipartite subgraph can no longer be enlarged in this way. Clearly, at this point our current bipartite graph B' has the same size as the (unknown) B .

The second phase of the algorithm attempts to change B' , without changing its size, until $V(G) - B$ is a clique. This is accomplished by a four-local search, based on a very similar principle—namely, if $V(G) = B \cup C$ is a $(2, 1)$ -partition and $|B| = |B'|$, then B is obtained from B' by a deletion of two vertices and the insertion of two other vertices. Hence we can test all n^4 possible new sets B' for being bipartite and the corresponding $V(G) - B'$ for being a clique, and if no $(2, 1)$ -partition is found we can be sure none exists.

The most time-consuming operation is the first phase of the algorithm, finding one $(2, 1)$ -partition—taking time n^8 . No low-degree algorithms for $(2, 1)$ -partition seem to be known at this time. (Another high degree algorithm is given in [39, 40].) \square

We summarize the steps of the algorithm.

Algorithm 8

Input: A graph G .

Task: Find a $(2, 1)$ -partition of G , if one exists.

Action: Start with any induced bipartite subgraph B of G (say, a single edge). While the size of the current B can be increased by deleting two current vertices and inserting three new vertices, do so.

When the size of B can no longer be increased in this way, test whether the complement of B is a clique, and if not, delete two current elements of B and insert two new elements to B , keeping B bipartite, until the complement of the current B is a clique (or until all possible ways have been tried).

We now restrict our attention to chordal graphs G . There are polynomial time algorithms to test whether or not a chordal graph is a -colourable (i.e., an $(a, 0)$ -graph), or partitionable into b cliques (i.e., a $(0, b)$ -graph) [123], and hence we may suspect that recognizing chordal (a, b) -graphs is possible in polynomial time for all positive integers (a, b) . It is also known that a chordal graph is a -colourable if and only if it does not contain an $(a + 1)$ -clique, and is partitionable into b cliques if and only if it does not contain $b + 1$ independent vertices, and we seek similar forbidden subgraph characterizations of chordal (a, b) -graphs.

Note that a chordal graph G is an (a, b) -graph if and only if the deletion of some b cliques from G leaves an a -colourable graph, i.e., a graph without $(a + 1)$ -cliques. In other words, G is an (a, b) -graph if and only if G contains b cliques that meet (intersect) all $(a + 1)$ -cliques of G . Let $g(G, r)$ denote the minimum number of cliques of G that meet all r -cliques (copies of K_r) of G . We shall design an efficient algorithm to compute $g(G, r)$ for chordal graphs G ; note that

this solves the recognition problem for chordal (a, b) -graphs, as $g(G, a + 1) \leq b$ if and only if G is an (a, b) -graph.

We say that two r -cliques in G are *independent* if no vertex of the first clique is adjacent to a vertex of the second clique. This means, in particular, that independent r cliques must be disjoint. Let $f(G, r)$ denote the maximum number of independent r -cliques in a graph G . It is clear from these definitions that $f(G, r) \leq g(G, r)$, as a clique of G cannot meet two independent r -cliques. In general this inequality can be strict, but for chordal graphs we have an equality.

Theorem 5.45 *If G is a chordal graph, then $f(G, r) = g(G, r)$ for all positive integers r .*

Note that when $r = 1$, this is the well-known fact that chordal graphs have independence number equal to the minimum number of cliques needed to cover all vertices [123].

Proof Since $f(G, r) \leq g(G, r)$ holds in general, it only remains to prove the opposite inequality. We shall describe an algorithm which, for any chordal graph G , identifies k independent r -cliques in G , as well as k cliques that cover all r -cliques of G . This means that both $f(G, r)$ and $g(G, r)$ are equal to k .

Algorithm 9

Input: A chordal graph G .

Task: Find in G independent r -cliques A_1, A_2, \dots, A_k as well as cliques B_1, B_2, \dots, B_k that meet all r -cliques of G .

Action: Compute a perfect elimination ordering u_1, u_2, \dots, u_n of G . With this ordering of the vertices of G , we perform as many of the following iterations as possible.

In the i -th iteration ($i = 1, 2, \dots, k$), we define the r -clique A_i , and the clique B_i , as follows. Let u_f be the first remaining vertex which is the $(r - 1)$ -st remaining neighbour of some remaining vertex u_p . Then A_i consists of u_p together with its first $(r - 1)$ remaining neighbours, up to and including u_f , and B_i consists of u_f together with all its remaining neighbours u_q with $q > f$. Now remove from G all the vertices of A_i and of B_i .

It is clear that at the end of Algorithm 9, we will have k cliques B_i that meet all r -cliques of G , since otherwise the algorithm would have performed another iteration. We also have k r -cliques A_i ; it remains to show that these r -cliques are independent. Thus suppose there is an edge between A_i and A_j , with $i < j$. Specifically, assume that some vertex u_s of A_i is adjacent to some vertex u_t of A_j . Let a_i be the last vertex of A_i , and a_j the last vertex of A_j . Note that a_i precedes a_j , and that a_i is also the first vertex of B_i , which was entirely removed from G before A_j was constructed. Hence we will have reached a contradiction, if we can show that a_i is adjacent to a vertex of A_j which follows a_i . If u_t follows a_i , then u_s is adjacent to both a_i and u_t which follow it, and hence a_i is adjacent to $u_t \in A_j$, and u_t follows a_i , as claimed. If u_s precedes u_t and u_t precedes a_i , then u_t must be adjacent to a_i , since u_s is adjacent to both of them. Therefore

u_t is adjacent to both a_i and a_j , and both follow it, whence a_i is adjacent to $a_j \in A_j$, and a_j follows a_i as claimed. The last possibility is that u_s follows u_t . In this case u_s must be adjacent to a_j , since u_t is adjacent to both of them, and both follow u_t , and finally a_i must be adjacent to a_j , since u_s is adjacent to both of them, and both follow u_s . This completes the proof. \square

Corollary 5.46 *A chordal graph G is an (a, b) -graph if and only if it does not contain $(b + 1)K_{a+1}$ as an induced subgraph.* \square

5.9 Remarks

The dichotomy classification in Theorem 5.2 was first proved in [168], and we followed the proof from there; a new proof in [50] follows the same idea, but uses the algebraic approach mentioned at the end Section 5.3. The difference is most pronounced in the case of extra-ordinary graphs, i.e., in the proof of Proposition 5.3. Digraph H -colouring problems were first studied by W. Gutjahr, E. Welzl, and G. Woeginger [129], in the context of applications to graph grammars and interpretations, cf. [237–239]. The example in Fig. 5.7 (and the proof of Proposition 5.10) is due to W. Gutjahr [130]. Theorem 5.14 is due to T. Feder and M. Vardi [105, 106]. The proof is extracted from [106]. Conjecture 5.16 is first formulated in [24]; an interesting reformulation can be found in [27]. Many special cases of the conjecture have been verified by J. Bang-Jensen, G. McGillivray, and others [24, 25, 27]. (For instance, it has been proved for vertex-transitive graphs [231].) A nice account of this is given in the survey by G. Hahn and G. MacGillivray [136], or the book of J. Bang-Jensen and G. Gutin [23]. Our presentation here is based on the survey [158]. The example in Fig. 5.11 is from [172]; the first (much larger) tree H with NP -complete H -colouring problem was found in [129]. The algebraic approach to the dichotomy problem, and the use of polymorphisms, as explained in Proposition 5.17, was pioneered by P. G. Jeavons. (There have been earlier related developments by R. Dechter and others [74].) The proof of Proposition 5.17 given here is from [193], see also [192] and [53]. We also recommend the book of N. Pippenger [290]. The connection of duality to polynomial time solvability, Theorem 5.20, was established in [171]; cf. also [175, 176]. As noted in the text, all known polynomial cases of digraph H -colouring problems can be explained by bounded treewidth duality. There are other known mechanisms for designing polynomial time H -colouring algorithms, including datalog programs [105, 106], and semidefinite programming [8, 107]. Yet they all seem to define the same class of polynomial time solvable digraph H -colouring problems. The consistency check (Algorithm 2) is a well-known procedure in artificial intelligence [232]. The X -underbar property is from [129], where it was first used for designing polynomial time algorithms. Theorem 5.27 is from [280]. The use of majority functions for designing polynomial time homomorphism algorithms first appeared in [105, 106]. Section 5.6 is based on [96]; the full classification for all graphs with loops allowed, Theorem 5.38, is from [99]. Motivated by these classification results, A. A. Bulatov has recently proved dichotomy for all list constraint satisfaction problems [49] (cf.


also [97]). The complexity of list k -colouring problems has been studied by J. Kratochvíl and Z. Tuza [208] and others cited there. The complexity of counting homomorphisms and list homomorphisms has been studied in [51, 52, 76, 85, 170]. Parametrized complexity of homomorphism and list homomorphism problems has been investigated by J. Diaz, M. Serna, D. M. Thilikos, and others, cf. the survey [77], and also see [293]. Homomorphisms to trigraphs (under the name of matrix partitions) were pioneered in [101, 102]; Sections 5.7 and 5.8 are based on [102, 103, 159, 160]. In [102] all trigraph H -colouring problems for H with up to four vertices have been classified as NP -complete or *quasi-polynomial* (admitting an $n^{O(\log n)}$ algorithm). Most, but not all, of the quasi-polynomial cases were proved polynomial in [57, 110]. There is one remaining case dubbed the ‘stubborn case’ in [57]; a similar but simpler open problem is given in [97]. These problems are not likely to be NP -complete (because they admit a quasi-polynomial algorithm), yet no polynomial algorithm is known for either of them. (It seems dichotomy is less likely to hold for trigraph homomorphisms.) Clique cutsets were first studied by E. Tarjan, and by S. Whitesides, each proving Corollary 5.43. Generalized split graphs were introduced by A. Brandstädt [39, 40]. The discussion of (a, b) -partitions for chordal graphs, in Theorem 5.45 and Corollary 5.46, follows [159, 160]. In [103], it is shown that there are trigraphs H such that list H -colouring is NP -complete even for chordal graphs.

Exercise 5 is due to W. Gutjahr, E. Welzl, and G. Woeginger. Exercise 6 is from [27]; it is conjectured there that the conclusion can be extended to any digraph G which includes some three vertices a, b, c and the arcs ab, ba, ac, cb . Exercise 15 is from [207]. Exercise 16 is from [114]; the effect of degree restrictions on the complexity of H -colouring problems has been much studied [100, 114, 133, 170, 198, 207]. Exercise 17 is based on [163]. Exercise 8 is from [42]; cf. also [17]. Exercise 10 is from [169]. Exercise 20 is from [176]. Dichotomy for all oriented cycles, including those of zero net length, has been proved in [95]. (Nevertheless, no simple classification distinguishing polynomial time solvable and NP -complete H -colouring problems for oriented cycles H is known.) Exercise 21 is from [209]; see also [109, 210] and the papers cited there for similar results on locally injective and locally surjective homomorphisms. Exercise 23 is based on [85]; a full classification of the complexity of counting problems has recently been achieved by A. A. Bulatov [51]. Exercise 27 is from [94]; the second problem remains NP -complete for interval graphs and cographs [33], and for trees [56].

5.10 Exercises

1. Prove that for a fixed graph G the question whether or not $G \rightarrow H$, for an input graph H , can be decided in polynomial time.
2. Prove that if H is a core then the problems H -RET and H -COL are polynomially equivalent. (Your proof should apply for any general relational system H .)
3. Give a polynomial time algorithm for H -RET when H is a reflexive tree.

4. Prove that if both the G -colouring and the H -colouring problems are polynomial time solvable, then so is the $(G \times H)$ -colouring problem.
5. (a) Let H be a component of H' . Prove that there is a polynomial time reduction from the H -colouring problem to the H' -colouring problem.
 (b) Let A be the digraph obtained from \vec{C}_3 by adding one vertex v and all arcs $vi, i \in V(\vec{C}_3)$, and let B be obtained from \vec{C}_3 by adding one vertex w and all arcs $iw, i \in V(\vec{C}_3)$. Let $H' = A \times B$.
 Prove that the H' -colouring problem is polynomial time solvable.
 Prove that H' contains an induced subgraph isomorphic to the digraph H in Fig. 5.14, and that the H -colouring problem is NP -complete.

 H FIG. 5.14. An NP -complete H -colouring problem for Exercise 5.

6. Let H be a digraph which includes some three vertices a, b, c with all arcs ab, ba, ac, ca, bc, cb . Prove that H -COL is NP -complete.
7. Prove that each interval graph admits an X -underbar enumeration.
8. Prove that a reflexive graph has a conservative majority function if and only if it is an interval graph. (This result is predicted by Theorem 5.35.)
9. Let g be a positive integer. Let M be an edge-minimal graph of girth at least g that is not three-colourable. Let ik be an edge of M , and let I be obtained from M by deleting the edge ik , and adding a vertex j adjacent to k . Note that I is three-colourable, but in any three-colouring the colours of i and j have to be different (but are otherwise arbitrary). Use Lemma 5.5 to prove that three-colourability is NP -complete even when restricted to graphs of girth at least g .
 Extend the arguments to prove NP -completeness of the H -colouring problem restricted to graphs of girth at least g , for any fixed nonbipartite graph H .
10. Prove that it is NP -complete to decide whether or not the core of a digraph is an oriented tree. Prove that it is NP -complete to decide whether or not a graph is a core.

11. Show that the consistency check (Algorithm 2) can be implemented to run in linear time.
12. Prove that asymptotically almost all graphs are ordinary. (Hint: An extraordinary graph does not contain K_4 .)
13. Prove that the relational system N , introduced in Section 1.8, is projective, i.e., all idempotent polymorphisms of N are projections. (Hint: Show that N admits no near-unanimity functions, as defined in Exercise 28, and that an idempotent polymorphism which is not a near-unanimity function must be a projection.)
14. For each trigraph H we introduce two symmetric binary relations $E(H)$ and $N(H)$ as follows: $E(H) = W(H) \cup S(H)$ and $N(H) = V(H) \times V(H) - S(H)$. Thus each trigraph H yields a binary relational system $R(H)$ with vertices $V(H)$ and two binary relations $E(H), N(H)$. A homomorphism of trigraph H to trigraph H' is any mapping of $V(H)$ to $V(H')$ that is a homomorphism of the system $R(H)$ to the system $R(H')$. (Hence homomorphisms of trigraphs can be composed, as usual.) Prove that if G is a graph, homomorphisms of G to a trigraph H as defined in the text are precisely homomorphisms of $R(G)$ to $R(H)$.
15. A graph G is *d-degenerate* if its vertices can be ordered as u_1, u_2, \dots, u_n so that each v_i has degree at most d in the subgraph induced by u_{i+1}, \dots, u_n . Prove the following dichotomy classification for d -degenerate graphs: H -COL is polynomial if H is bipartite or contains K_d ; otherwise H -COL is NP -complete.
16. Prove that for graphs with maximum degree three, C_5 -COL is NP -complete. Let U be the graph constructed in Proposition 1.24. Prove that for cubic graphs, U -COL is polynomial time solvable.
17. Prove that for planar graphs, C_5 -COL is NP -complete.
18. [98] For graphs, the list H -colouring problem is polynomial time solvable if H is a bipartite graph whose complement of H is a *circular arc graph*; and is NP -complete otherwise. (A *circular arc graph* is an intersection graph of a set of arcs on a circle [318].)
19. [25] Prove the following dichotomy theorem for tournaments H . If H is acyclic, or has exactly one directed cycle, then H -colouring is polynomial time solvable. If H has at least two directed cycles then H -colouring is NP -complete.
20. [176] Find a polynomial time algorithm for the H -colouring problem when H is any oriented cycle with nonzero net length. Find an oriented cycle H of zero net length for which the H -colouring problem is NP -complete.
21. [209] A homomorphism f of a graph G to a graph H is *locally bijective*, and G is a *cover* of H , if the neighbours of every vertex u in G are taken bijectively to the neighbours of $f(u)$ in H . (In particular, u and $f(u)$ must always have the same degree.) The *H-cover problem* asks whether or not an input graph G is a cover of H .

Prove that

- (a) each C_k -cover problem is polynomial time solvable
- (b) if H consists of ℓ cycles incident with one vertex but otherwise disjoint, then the H -cover problem is polynomial time solvable
- (c) if $k > 3$, then the K_k -cover problem is NP -complete.

22. [8, 77] Let H be a graph with loops allowed. The problem of *equitable H -colouring* asks whether or not an input graph G admits a homomorphism f to H such that the sizes of any two sets $f^{-1}(x), x \in V(H)$, differ by at most one element.

Prove the following dichotomy classification.

- If each component of H is a complete bipartite graph or a complete reflexive graph, then the equitable H -colouring problem has a polynomial time algorithm.
- Otherwise the equitable H -colouring problem is NP -complete.

23. [85] Let H be a graph. Prove the following dichotomy classification.

- If each component of H is a complete bipartite graph or a complete reflexive graph, then the number of H -colourings can be found by a polynomial time algorithm.
- Otherwise the problem of counting H -colourings is $\#P$ -complete.

24. [104] This exercise is restricted to reflexive digraphs; a reflexive digraph is considered acyclic if it has no directed cycles other than the loops. A homomorphism f of G to H is *acyclic* if each set $f^{-1}(x), x \in V(H)$, induces an acyclic subgraph of G . The *acyclic H -colouring problem* asks whether or not an input digraph G admits an acyclic homomorphism to H . Prove the following dichotomy classification.

- If H is acyclic then the acyclic H -colouring problem is polynomial time solvable.
- Otherwise, the acyclic H -colouring problem is NP -complete.

25. [175] Prove that for an oriented path H we have $G \not\rightarrow H$ if and only if there exists an oriented path P such that $P \rightarrow G$ but $P \not\rightarrow H$.

26. [47] Recall the definition of a contraction from the last remarks in Section 1.3. Let H -CON denote the problem of deciding whether an input graph G contracts to the fixed graph H .

Give polynomial time algorithms for K_3 -CON, $K_{1,3}$ -CON, and P_2 -CON; prove that C_4 -CON and P_3 -CON are NP -complete.

27. [94] Prove that for any fixed k one can decide in polynomial time whether or not the achromatic number of an input graph G exceeds k . (See Exercise 16 in Chapter 1.)

If both G and k are part of the input, then show that it is NP -complete to decide if G has a complete k colouring.

28. [42] A *near-unanimity function* on H is a polymorphism f of H such that $f(u_1, u_2, \dots, u_k) = x$ if at most one u_i is not equal to x . The *order* of the near-unanimity function is the order of the polymorphism, i.e., the integer k . (Note that each near-unanimity function is idempotent, and a near-unanimity function of order three is precisely a majority function.) Show that each reflexive chordal graph has a near-unanimity function. (Hint: Use Algorithm 6.)
- Show that a digraph H has a near-unanimity function of order k if and only if every conflict with respect to H has at most $k - 1$ vertices.
29. [96] Show that for every bipartite graph H there exists a reflexive graph H' such that H -RET and H' -RET are polynomially equivalent. Deduce that dichotomy for reflexive retraction problems would imply dichotomy for all constraint satisfaction problems.
30. [330] Let H be a reflexive chordal graph. Prove that if the consistency check (Algorithm 2) succeeds for a graph G with connected lists L , then there is a list homomorphism of G to H (with respect to the lists L). (Hint: Consider, in each final list $L^*(u)$, the first vertex in the perfect elimination ordering of H .)
31. [223] Prove that the class A of reflexive graphs that admit a near-unanimity function is a variety.
- Let B be the class of reflexive graphs G with the following property: whenever G is a subgraph of H such that the *retraction consistency check* succeeds, then G is a retract of H . (The retraction consistency check is the consistency check in which the lists are as follows: the list of each $u \in V(G)$ is $L(u) = \{u\}$; all other lists are $L(x) = V(G)$.)
- Prove that B is a variety.
 - Prove that $A \subseteq B$.

6

COLOURING—VARIATIONS ON A THEME

The usual graph colourings are homomorphisms to complete graphs. They have an extensive literature and constitute a basic part of graph theory. One of the reasons that complete graphs offer a convenient scale is the fact that

$$K_1 < K_2 < K_3 < \cdots .$$

(Recall that $G < H$ stands for $G \rightarrow H$ and $H \not\rightarrow G$.) In this chapter we shall look at a number of variants of colouring that have been proposed, and sometimes investigated without direct connection to homomorphisms, and illustrate the utility of viewing them under the umbrella of the homomorphism framework. (All these variants turn out to be homomorphisms to other special targets.)

In general, one can choose, instead of the family $K_n, n \in N$, any fixed family \mathcal{F} of graphs, and investigate the largest or smallest (in some order) member of \mathcal{F} to which a given graph admits a homomorphism. In the first section, we consider circular colourings, which are homomorphisms to a family of *rational (or circular) complete graphs* $K_{p/q}, p, q \in N$. In the second section, we consider multicolourings, i.e., homomorphisms to the family of Kneser graphs $K(n, k), n, k \in N$. A discussion of T -colourings, oriented colourings, and acyclic colourings completes the chapter. We argue that the uniting thread to all these variants of colouring is the common perspective of graph homomorphisms.

6.1 Circular colourings

As noted above, we can view the family $\mathcal{F} = K_1 < K_2 < \cdots$ as a scale that allows us to associate to each graph G a positive integer—the chromatic number of G —which in some sense measures the homomorphism capabilities of G . Just as the set of natural numbers can be extended to the set of rationals, we can extend \mathcal{F} to a larger family, allowing us to measure the homomorphism properties of G more finely.

Let us begin by refining the scale only for graphs that have chromatic number three. There we have a natural calibrating family—namely the family of odd cycles $C_{2k+1}, k \geq 1$, which satisfies

$$C_3 > C_5 > C_7 > \cdots .$$

Each graph G with $\chi(G) = 3$ admits a homomorphism to C_3 ; on the other hand G contains some odd cycle C_{2k+1} , and hence does not admit a homomorphism to any longer odd cycle. Thus there exists a largest odd cycle C_{2c+1} to which

G admits a homomorphism. The larger the value c , the closer (in a sense) the graph G is to being bipartite, and hence we choose to associate with G a more specific ‘chromatic number’, namely $2 + 1/c$.

To apply this way of thinking to graphs G with an arbitrary chromatic number, we define the following family of calibrating graphs $K_{p/q}$, where p, q are integers with $0 < q \leq p$. The *rational complete graph* $K_{p/q}$ has vertices $\{0, 1, \dots, p-1\}$ and edges $\{ij : q \leq |i-j| \leq p-q\}$ (cf. Fig. 6.1). Note that $K_{p/q}$ does not have any edges unless $p \geq 2q$, that $K_{p/1}$ is isomorphic to K_p , and that $K_{(2k+1)/k}$ is isomorphic to the odd cycle C_{2k+1} . Hence the family of rational complete graphs contains both the calibrating family of ‘integer complete graphs’ $K_1 < K_2 < \dots$, and the calibrating family of odd cycles $C_3 < C_5 < \dots$.

Denote $|x|_p = \min\{|x|, p-|x|\}$. Then $|i-j|_p$ is the *circular distance* between the vertices i and j on the circle with vertices $\{0, 1, \dots, p-1\}$. It is easy to verify that $q \leq |i-j| \leq p-q$ if and only if $|i-j|_p \geq q$. In particular, the existence of an edge ij in $K_{p/q}$ depends only on the circular distance of i and j . Graphs with this property are called *circulants*, since their adjacency matrices are circulant matrices. Our rational complete graphs $K_{p/q}$ are ‘large-step’ circulants, as the circular distances that correspond to edges are precisely the distances which are at least q . We note for future reference that each $K_{p/q}$, as all circulants, is vertex-transitive.




FIG. 6.1. The rational complete graph $K_{5/2}$.

A homomorphism of G to $K_{p/q}$ will be called a (circular) (p/q) -colouring of G . Recall that $K_{p/1} = K_p$, so a $(p/1)$ -colouring is just a p -colouring in the usual sense. Hence the set of all (p/q) -colourings of a graph enriches the set of all its p -colourings. We call a graph (p/q) -colourable if it admits a (p/q) -colouring, and define the *circular chromatic number* of a graph G , written as $\chi_c(G)$, as

$$\chi_c(G) = \inf \left\{ \frac{p}{q} : G \text{ is } (p/q)\text{-colourable} \right\}.$$

Early papers refer to this concept as the *star chromatic number*, as the rational

complete graphs can be nicely drawn in the shape of stars, cf. Fig. 6.1. The vertex-transitivity of $K_{p/q}$ implies that in a (p/q) -colouring we can arbitrarily pick a vertex and colour it 0.

We have the following monotonicity property of the circular chromatic number (as in Corollary 1.8).

Proposition 6.1 *If $G \rightarrow H$, then $\chi_c(G) \leq \chi_c(H)$.*

Proof By definition, $\chi_c(G)$ is the infimum of all fractions p/q such that $G \rightarrow K_{p/q}$. If $G \rightarrow H$, then $H \rightarrow K_{p/q}$ implies $G \rightarrow K_{p/q}$, whence the conclusion. \square

In particular, we see that homomorphically equivalent graphs have the same circular chromatic number.

An isomorphism of $K_{(2k+1)/k}$ to C_{2k+1} can be viewed as a $(2k+1, k)$ -colouring of C_{2k+1} . We illustrate this in Fig. 6.2.




FIG. 6.2. A $(5/2)$ -colouring of C_5 .

Thus for any k , $\chi_c(C_{2k+1}) \leq 2 + 1/k$. We shall soon prove that, in fact, $\chi_c(C_{2k+1}) = 2 + 1/k$, Corollary 6.5. In particular we see that $\chi_c(G)$ could be smaller than $\chi(G)$, and arbitrarily close to $\chi(G) - 1$. (Corollary 6.11 will establish that it cannot be equal to $\chi(G) - 1$ or smaller.)

On the other hand, for complete graphs we have the following equality.

Proposition 6.2 $\chi_c(K_n) = \chi(K_n) = n$.

Proof The fact that $\chi_c(K_n) \leq \chi(K_n) = n$ is obvious, since each n -colouring is an $(n/1)$ -colouring. Hence suppose that K_n admits a (p/q) -colouring, and let $c_1 = 0 < c_2 < \dots < c_n$ be the colours of the vertices of K_n . Since we must have each $c_{i+1} - c_i \geq q$, as well as $c_n - c_1 \leq p - q$, we conclude that $(n-1)q \leq c_n \leq p - q$, whence $p/q \geq n$. \square

We begin by investigating the existence of homomorphisms amongst the graphs $K_{p/q}$. One important property of the family of complete graphs K_n was

that $K_n \rightarrow K_{n'}$ if and only if $n \leq n'$. Fortunately, we have a similar property for the family of rational complete graphs.

Theorem 6.3 *Let $p/q \geq 2$. Then $K_{p/q} \rightarrow K_{p'/q'}$ if and only if $p/q \leq p'/q'$.*

Proof The necessity of the condition will be proved using the No-Homomorphism Lemma, i.e., Corollary 1.23. (Recall that each $K_{p/q}$ is vertex-transitive.) Thus we will need to know the independence number of $K_{p/q}$.

Claim *The independence number of $K_{p/q}$ is q .*

The q vertices $0, 1, \dots, q-1$ are independent in $K_{p/q}$, therefore, the independence number is at least q . On the other hand, any independent set S in $K_{p/q}$ must be included in the set $\{i - q + 1, i - q + 2, \dots, i, i + 1, i + 2, \dots, i + q - 1\}$, for any $i \in S$. Moreover, for any $j = 1, 2, \dots, q - 1$, S can contain at most one of the adjacent vertices $i - q + j, i + j$. Therefore S has at most q elements.

Since the independence ratio of $K_{p/q}$ is q/p , there can be no homomorphism $K_{p/q}$ to $K_{p'/q'}$ if $p/q > p'/q'$ by the No-Homomorphism Lemma, Corollary 1.23. If $p/q \leq p'/q'$, then we define $f(i) = \lfloor iq'/q \rfloor$. It is easy to see that for any $i \leq p - 1$, we have $f(i) \leq p' - 1$, and a straightforward calculation shows that if $q \leq |i - j| \leq p - q$ then $q' \leq |f(i) - f(j)| \leq p' - q'$. Thus f is a homomorphism of $K_{p/q}$ to $K_{p'/q'}$. \square

Of particular interest is the conclusion that $K_{p/q}$ and $K_{p'/q'}$ with $p/q = p'/q'$ are homomorphically equivalent. This allows us to consider only (p/q) -colourings with p, q relatively prime, and speak meaningfully (up to homomorphic equivalence) of rational complete graphs as K_r , where r is any rational $r \geq 1$. More precisely, we shall write K_r to mean $K_{p/q}$ where p, q are relatively prime and $r = p/q$.

We now return to the notion of the circular chromatic number. Since we have shown that K_r does not admit a homomorphism to any $K_{r'}$ with $r' < r$, we have also the following fact.

Corollary 6.4 *For each rational number $r \geq 2$, we have*

$$\chi_c(K_r) = r.$$

Applying this to $r = (2k + 1)/k$, we obtain the following.

Corollary 6.5 *For any odd cycle, $\chi_c(C_{2k+1}) = 2 + \frac{1}{k}$.* \square

A crucial property of χ_c is that the infimum in its definition is always attained, i.e., can be replaced by minimum.

In order to obtain this result, we shall first analyse how deleting a vertex from $K_{p/q}$ affects its circular chromatic number. When $q = 1$, we are deleting a vertex of K_p , and obtain K_{p-1} , lowering the chromatic (and circular chromatic) number by one. For general values of q , the amount by which the circular chromatic number is lowered can be calculated as follows. If integers p and q are relatively

prime, with $q > 1$ and $p \geq 2q$, then there exist unique positive integers p' and q' with $0 < p' < p, 0 < q' < q$ such that $pq' - qp' = 1$. Note that it follows that

$$\frac{p'}{q'} = \frac{p}{q} - \frac{1}{qq'}.$$

It turns out that the amount by which the circular chromatic number of $K_{p/q}$ is lowered after a deletion of a single vertex, is exactly the quantity $1/qq'$. Note that the vertex deleted, which we call x , is arbitrary, since $K_{p/q}$ is vertex-transitive.

Lemma 6.6 *Let p and q be relatively prime, and $p/q \geq 2$. Let p' and q' be the unique integers with $pq' - qp' = 1$, where $0 < p' < p$ and $0 < q' < q$. Then for any vertex x , the graph $K_{p/q} - x$ is homomorphically equivalent to $K_{p'/q'}$. In particular, we have*

$$\chi_c(K_{p/q} - x) = \frac{p'}{q'}.$$

Proof It suffices to prove the first statement, as homomorphically equivalent graphs have the same circular chromatic number, and we have already established that $\chi_c(K_{p'/q'}) = p'/q'$. Consider first the sequence of vertices

$$0, q + 1, 2q + 1, \dots, (p' - 1)q + 1,$$

of $K_{p/q}$, reduced modulo p . When the vertices of $K_{p/q}$ are drawn around a circle, as in our figures, this sequence winds around the circle q' times, since $p'q + 1 = pq'$. On the other hand, the sequence $0, q', 2q', \dots, (p' - 1)q'$ enumerates all vertices of $K_{p'/q'}$, and also winds around a corresponding circle q' times. It is easy to check that assigning the i -th terms of the sequences to each other is an isomorphism between $K_{p'/q'}$ and the subgraph of $K_{p/q}$ induced by the vertex set X consisting of $0, q + 1, 2q + 1, \dots, (p' - 1)q + 1$. Note that $q \notin X$. We shall construct a retraction of $K_{p/q} - q$ to its subgraph induced by X . This will mean that $K_{p/q} - q$, and hence any $K_{p/q} - x$, is homomorphically equivalent to $K_{p'/q'}$. The retraction f assigns to each vertex y of $K_{p/q}$, different from q , the largest integer of X smaller than or equal to y . It is a routine exercise to check that if $y < y'$ have $y + q \leq y' \leq y + p - q$, then $f(y) + q \leq f(y') \leq f(y) + p - q$, so that adjacency is preserved. Since $f(y) = y$ when $y \in X$, we have indeed a retraction. \square

Since $p'/q' < p/q$, we are in position to observe another similarity between colourings by (integer) complete graphs K_n and by rational complete graphs K_r .

Corollary 6.7 *If G admits a (p/q) -colouring that does not use every colour $0, 1, \dots, p - 1$, there exists a (p'/q') -colouring of G with $p'/q' < p/q$, which does use every colour.*

Proof A (p/q) -colouring of G that does not use every colour is a homomorphism $f : G \rightarrow K_{p/q} - x$, for some vertex x . The composition of f with a homomorphism $g : K_{p/q} - x \rightarrow K_{p'/q'}$ is a (p'/q') -colouring of G with $p'/q' < p/q$ and $p' < p$. If this colouring does not use every colour, we can further decrease the ratio p'/q' .

As the integer p' strictly decreases at each step, the process must terminate, with a colouring which does use every colour. \square

By Corollary 6.7, to determine the circular chromatic number of a graph with n vertices, we only need to consider (p/q) -colourings with $q \leq p \leq n$. As there are only finitely many such pairs (p, q) , the infimum is always attained and hence can be replaced by minimum.

Corollary 6.8 *For a graph G on n vertices, we have*

$$\chi_c(G) = \min \left\{ \frac{p}{q} : p \leq n, G \rightarrow K_{p/q} \right\}$$

In particular, $\chi_c(G)$ is a rational number, and its computation is a finite problem.

We can now conclude the following fact from Theorem 6.3.

Corollary 6.9 *If $\chi_c(G) \leq r$, then $G \rightarrow K_r$.* \square

Using Proposition 6.1 we also conclude the following.

Corollary 6.10 *$\chi_c(G) \leq 2 + 1/k$ if and only if $G \rightarrow C_{2k+1}$.* \square

We also obtain the following inequalities.

Corollary 6.11 $\chi(G) - 1 < \chi_c(G) \leq \chi(G)$

Proof The upper bound follows from the definitions. We prove the lower bound by contradiction. Suppose $\chi_c(G) \leq \chi(G) - 1$. By Corollary 6.8, G admits a p/q -colouring with $p/q \leq \chi(G) - 1$. Theorem 6.3 implies that $K_{p/q} \rightarrow K_c$ where $c = \chi(G) - 1$, and by composition we also have $G \rightarrow K_c$, a contradiction. \square

Thus $\chi(G)$ is exactly the ceiling of $\chi_c(G)$.

In particular, two graphs of the same chromatic number may have different circular chromatic numbers, and hence the circular chromatic number provides us with a much ‘finer scale’ of colourability.

We conclude this section by presenting two equivalent definitions of χ_c .

Let C be a circle of length r . A C -colouring of a graph G is an assignment of unit length open arcs of C to the vertices of G , so that adjacent vertices are assigned disjoint arcs.

Theorem 6.12 *Let G be a graph, and let R be the set of all real numbers r such that G admits a C -colouring for a circle C of length r . Then $\chi_c(G) = \inf R$.*

Proof Suppose G is a graph. It suffices to prove that for any rational number $r = p/q$, there exists a C -colouring of G for a circle C of length r if and only if there exists a (p/q) -colouring of G . We may view C as obtained from the interval $[0, r]$ by identifying the two ends. Thus each point of C is associated with a real number t with $0 \leq t < r$.

Suppose G admits a C -colouring c , with a circle C of length r . For each vertex x of G , let $c^-(x)$ be the counterclockwise endpoint of the associated arc

$c(x)$. Then for any edge xy of G we have $1 \leq |c^-(x) - c^-(y)| \leq r - 1$. Let $\phi : V(G) \rightarrow \{0, 1, \dots, p-1\}$ be defined as $\phi(x) = \lfloor c^-(x)q \rfloor$. Then for any edge xy of G , $q \leq |\phi(x) - \phi(y)| \leq p - q$. Hence ϕ is a (p/q) -colouring of G . Conversely, given a (p/q) -colouring ϕ of G , we define a mapping $c^- : V(G) \rightarrow [0, r)$ as $c^-(x) = \phi(x)/q$. Then for any edge xy of G , we have $1 \leq |c^-(x) - c^-(y)| \leq r - 1$. Let each $c(x)$ be the unit length open arc with counterclockwise endpoint $c^-(x)$. It is easy to see that c is a C -colouring of G where C has length r . \square

This reformulation of $\chi_c(G)$ is very appealing, as it can be shown (Exercise 12) that when the circle is replaced by a line segment, and the circular arcs also by line segments, we obtain an equivalent definition of the ordinary chromatic number $\chi(G)$. Thus the circular chromatic number is truly a ‘circular’ version of the usual chromatic number.

Recall that an orientation of a graph G is a digraph, say \vec{G} , obtained from G by assigning a unique direction to each edge of G . The orientation \vec{G} is *acyclic* if it contains no directed cycle.

Let C be a cycle of G , and choose a fixed direction around C . For any orientation \vec{G} of G , we denote by C^+ the set of arcs of \vec{G} that are in C in the forward direction, and by C^- the set of arcs of \vec{G} that are in C in the backward direction.

Theorem 6.13 *When G is a forest, $\chi_c(G) = 2$. Otherwise (G contains a cycle), $\chi_c(G)$ is the minimum, over all acyclic orientations \vec{G} of G , of the maximum, over all cycles C of G , of*

$$1 + \frac{|C^+|}{|C^-|}.$$

Proof Suppose $\chi_c(G) = p/q$ and c is a (p/q) -colouring of G . Let \vec{G} be obtained from G by orienting each edge xy from x to y just if $c(x) < c(y)$. Clearly, \vec{G} is an acyclic orientation of G . Then each edge xy of \vec{G} has $q \leq c(y) - c(x) \leq p - q$. If C is a cycle of G and $e = xy$ an edge of C , we let $f(e) = c(y) - c(x)$. Then $\sum_{e \in C^+} f(e) = \sum_{e \in C^-} f(e)$. As $\sum_{e \in C^+} f(e) \geq |C^+|q$ and $\sum_{e \in C^-} f(e) \leq |C^-|(p - q)$, we conclude that

$$\frac{|C^+|}{|C^-|} \leq \frac{p - q}{q} = \frac{p}{q} - 1.$$

Conversely, suppose \vec{G} is an acyclic orientation of G . We may define an integer mapping c on the vertices of G as follows. For each oriented walk W in \vec{G} , let W^+ and W^- be the sets of forward edges and backward edges of W , respectively, and let $f(W) = q|W^+| - (p - q)|W^-|$. For each vertex x of G , let

$$c(x) = \max\{f(W) : W \text{ is a walk of } G \text{ that ends at } x\}.$$

Since for every cycle C of \vec{G} we have $|C^+|/|C^-| \leq (p - q)/q$, it will suffice to consider only oriented paths W . (Revisiting a vertex results in no greater

value of $f(W)$.) There are only finitely many oriented paths, and hence $c(x)$ is well-defined.

Suppose xy is an arc of \vec{G} and $c(x) = f(W)$, where W is a walk of \vec{G} that ends at x . Then the concatenation $W(xy)$ is a walk that ends at y . Hence $c(y) \geq f(W(xy)) = f(W) + q$, i.e., $c(y) - c(x) \geq q$. If $c(y) = f(A)$, for a walk A that ends at y , then $A(yx)$ is a walk of G that ends at x . Hence $c(x) \geq f(A(yx)) = f(A) - (p - q)$, i.e., $c(y) - c(x) \leq p - q$. Let $c'(x)$ be the value congruent to $c(x)$ modulo p , which lies between 0 and $p - 1$. Then c' is a (p/q) -colouring of G . \square

We can now easily derive Theorem 2.52.

Corollary 6.14 *A nearly bipartite graph always retracts to any shortest odd cycle.*

Proof Recall that a graph G is nearly bipartite if it admits an orientation \vec{G} in which each cycle has net length at most one. If the minimum length of an odd cycle in G is $2k + 1$, then the greatest ratio of $|C^+|/|C^-| + 1$ is $(k + 1)/k + 1 = 2 + 1/k$. Thus $\chi_c(G) \leq 2 + 1/k$ and hence $G \rightarrow C_{2k+1}$. Since C_{2k+1} is a core, a homomorphism of G to C_{2k+1} followed by a suitable isomorphism is a retraction of G to any shortest odd cycle H in G . \square

Note that the expression in the theorem, $1 + |C^+|/|C^-|$, can also be written as $|C|/|C^-|$. We also derive the following well-known property of the usual chromatic number.

Corollary 6.15 *The chromatic number $\chi(G)$ is the minimum, over all acyclic orientations \vec{G} of G , of the maximum, over all cycles C of G , of*

$$1 + \left\lceil \frac{|C^+|}{|C^-|} \right\rceil.$$

Theorem 6.13 allows us to connect the circular chromatic number to a measure of concurrency used in heavily loaded resource-sharing systems. In such a system, a process is represented by a vertex and two vertices are adjacent whenever the two processes share a resource. In this way, we obtain a graph G representing the resource-sharing system. Two adjacent vertices (corresponding to processes sharing a resource) cannot operate concurrently. In a synchronous model of such a system (where a common clock controls all processes), a schedule chooses, at each time period, an independent set of vertices to operate. However, in order to ensure fairness in a heavily loaded system (dense graph), it is assumed that a vertex which has just operated will not operate again until all its neighbours have had their turns. To avoid deadlocks, schedules operate by choosing an acyclic orientation \vec{G} of the graph G , and operating on a set of sinks (vertices of outdegree zero) in \vec{G} .

Suppose X is an acyclic digraph and Y a set of sinks in X . If all arcs of X incident with each $y \in Y$ are reoriented to lead out of y (so that y becomes a source), we say that Y has been *reversed* in X .

Lemma 6.16 *If X is an acyclic digraph, and Y a set of sinks in X , then the digraph X' obtained from X by reversing Y is also acyclic.*

Proof Any directed cycle of X' is not a cycle of X ; thus it contains a vertex $y \in Y$. Note that the neighbours of y on the cycle are not in Y , since Y is an independent set in X . Therefore, y was a sink in X , and hence is a source in X' , which is impossible. \square

Suppose G is a graph, with n vertices, representing a resource sharing system. Formally, a *schedule on G* consists of an infinite sequence of acyclic orientations $\vec{G}_0, \vec{G}_1, \vec{G}_2, \dots$ of G , where each \vec{G}_{i+1} is obtained from \vec{G}_i by reversing a set of sinks, denoted by S_i . A vertex v of G that belongs to the set S_i is said to be *active* in the i -th time interval; let $A(v, k)$ denote the number of times vertex v has been active in the first k time periods. The *concurrency* of a schedule is the upper limit, as k tends to infinity, of the expression

$$\frac{1}{kn} \sum_{v \in V(G)} A(v, k).$$

A schedule that operates by choosing an initial orientation and then reversing one sink at a time does not seem to have a high degree of concurrency. It is in fact easy to see that $\sum_{v \in V(G)} A(v, k) = k$ for all k , thus its concurrency is zero.

The highest concurrency of a schedule starting from a given acyclic orientation \vec{G}_0 can be shown to be achieved when at each stage all the sinks are reversed. Moreover, its concurrency can be shown to be computed as follows.

Theorem 6.17 [28] *The maximum concurrency of a schedule on G , starting with a given acyclic orientation \vec{G}_0 , is equal to the minimum, over all oriented cycles C in \vec{G}_0 , of $|C^-|/|C|$.*

Let $\text{conc}(G)$ denote the highest concurrency of any schedule on G . We have the following consequence of the last two theorems.


Corollary 6.18

$$\text{conc}(G) = \frac{1}{\chi_c(G)}.$$

6.2 Fractional colourings

One of the most natural generalizations of colouring consists of assigning to each vertex, instead of just one colour, a set of k colours, and requiring that adjacent vertices obtain disjoint sets of colours. Such an assignment is called a *k -tuple colouring*, or a *k -tuple n -colouring* if a total of n colours is used. (We shall always assume that $0 < k \leq n$.) Obviously, a 1-tuple n -colouring is just the usual n -colouring. Figure 6.3 illustrates a pairwise (2-tuple) five-colouring of the pentagon C_5 .

It is again easy to see that these multicolourings are homomorphisms to suitable target graphs. Let $K(n, k)$ denote the graph with vertex set consisting

FIG. 6.3. A pairwise five-colouring of C_5 .

of all k -element subsets of $X = \{1, 2, \dots, n\}$ (occasionally we shall choose a different ground set X of n elements), and with two vertices adjacent just if the subsets are disjoint. Then a k -tuple n -colouring of G is exactly a homomorphism of G to $K(n, k)$. These calibrating graphs $K(n, k)$ are called *Kneser graphs*. Kneser graphs contain the corresponding rational complete graphs.

Proposition 6.19 *The rational complete graph $K_{p/q}$ is isomorphic to an induced subgraph of the Kneser graph $K(p, q)$.*

Proof Take for the ground set of $K(p, q)$ the set $X = \{0, 1, \dots, p-1\}$. It follows from the definitions of $K_{p/q}$ and $K(p, q)$, that the mapping that assigns to each vertex $x = 0, 1, \dots, p-1$ of $K_{p/q}$ the q -tuple $\{x, x+1, \dots, x+q-1\}$ of $K(p, q)$ (addition is modulo p), is an isomorphism of $K_{p/q}$ to the subgraph of $K(p, q)$ induced by these ‘circularly consecutive’ q -element subsets. For instance, in Fig. 6.4 we show (by heavier edges) how the graph $K_{5/2}$ can be viewed as a subgraph of $K(5, 2)$. \square

The *fractional chromatic number* of G , denoted $\chi_f(G)$, is the infimum of the fractions n/k such that G admits a k -tuple n -colouring. The above figure illustrates the fact that the fractional chromatic number of the pentagon is at most $5/2$. In fact we shall soon see that it is exactly $5/2$, Corollary 6.25. For a fixed k , the minimum n such that G admits a k -tuple n -colouring is called the *k -tuple chromatic number* of G , and denoted $\chi^k(G)$. Therefore, $\chi_f(G) = \inf \chi^k(G)/k$.

We conclude from Proposition 6.19 the following lower bound on $\chi_c(G)$ (the upper bound is from Corollary 6.11).

Corollary 6.20 *For each graph G , $\chi_f(G) \leq \chi_c(G) \leq \chi(G)$.* \square

Since $\chi_f(G)$ is the infimum of the fractions n/k such that $G \rightarrow K(n, k)$, we again have monotonicity of the fractional chromatic number in the homomorphism order, as in Proposition 6.1 and Corollary 1.8.




FIG. 6.4. The Petersen graph $K(5, 2)$ contains $K_{5/2}$ as an induced subgraph.

Proposition 6.21 *If $G \rightarrow H$, then $\chi_f(G) \leq \chi_f(H)$.* \square

The fractional chromatic number is the solution of a natural linear program. We begin by formulating the ordinary chromatic number as a 0, 1 linear program. Let G be a graph, and for each independent set $I \subset V(G)$ let x_I be a 0, 1 variable. An n -colouring of G is a set of n independent sets I which partition $V(G)$. If we choose values $x_I = 1$ for the chosen sets I and values $x_I = 0$ for the other sets I , we have a solution to the system

$$\sum_I x_I = n$$
$$\sum_{v \in I} x_I = 1, \text{ for all } v \in V(G).$$

The first equation indicates that exactly n independent sets I have been chosen. The second set of equations indicates that each vertex v of G is in *exactly* one chosen independent set. Conversely, any 0, 1 solution to the above equations corresponds to an n -colouring of G .

Thus we obtain the following result.

Theorem 6.22 *The chromatic number $\chi(G)$ of a graph G is equal to the optimum value of the integer linear program*

$$\min \sum_I x_I$$
$$\sum_{v \in I} x_I = 1, \text{ for all } v \in V(G)$$

$$x_I = 0, 1, \text{ for all } I \text{ independent.}$$

□

Consider now the *continuous relaxation* of this integer linear program. What will be the effect of relaxing the constraint $x_I = 0, 1$ to $x_I \geq 0$? Clearly, we should never need to make a value x_I greater than one, thus we can take the value x_I to mean the ‘degree’ to which the independent set I is to be taken. This view is aided by the fact that a linear program with integer coefficients is known to have an optimum with rational values of the variables [66]. Consider then a rational optimum x_I, I independent, of our linear program. (Our program is feasible and bounded, and hence does have an optimum.) We may assume that all fractions x_I have the same denominator k . By multiplying through by k we make sure that all values of x_I are integers. If the sum

$$\sum_I x_I$$

has value n then we have chosen n independent sets I , with repetition allowed, which cover each vertex of G exactly k times, i.e., we have a k -tuple n -colouring of G . Since we multiplied by k , the original fractional solution has value n/k . On the other hand, it is clear that any k -tuple n -colouring of G gives rise to a feasible solution of the linear program: it corresponds to n independent sets I covering each vertex k times, so if each variable x_I is set to $1/k$ we obtain a feasible solution of value n/k . Therefore we have the following conclusion.

Theorem 6.23 *The fractional chromatic number of a graph G is the minimum fraction n/k such that G admits a k -tuple n -colouring. In particular, $\chi_f(G)$ is always a rational number.* □

Before leaving linear programs behind, we observe that we can change our program to

$$\min \sum_I x_I$$

$$\sum_{v \in I} x_I \geq 1, \text{ for all } v \in V(G)$$

$$x_I \geq 0, \text{ for all } I \text{ independent,}$$

without changing its solution. Indeed, any solution of this modified program can be changed to be a solution of the original program as follows. Whenever a vertex v has $\sum_{v \in I} x_I > 1$ we decrease the value of some x_I, I independent, $v \in I$, and increase by the same amount the value of x_{I-v} (subset of an independent set is

again independent). This modified program is in the standard form [66], and its dual can be formulated as follows.

$$\begin{aligned} \max \quad & \sum_v y_v \\ \sum_{v \in I} y_v & \leq 1, \text{ for all } I \text{ independent} \\ y_v & \geq 0, \text{ for all } v \in V(G). \end{aligned}$$

This linear program assigns nonnegative values y_v to vertices v of G , so that the sum of the values over any independent set is at most 1, and tries to maximize the total sum of values. Clearly, if all values y_v are integers, i.e., 0 or 1 (given the constraints), the set of vertices v with $y_v = 1$ will define a maximum clique in G . Thus the solution to this linear program is called the *fractional clique number* of G .

The importance of the fractional clique number is that it allows us to obtain lower bounds on $\chi_f(G)$. Indeed, any solution $y_v, v \in V(G)$, of the dual linear program yields the lower bound $\sum_v y_v \leq \chi_f(G)$. (We are using a well-known property of primal and dual linear program [66].)

Corollary 6.24 *Suppose G is a graph with n vertices and independence number α . Then*

$$\chi_f(G) \geq \frac{n}{\alpha}.$$

□

Corollary 6.25 *For every integer $k \geq 1$*

$$\chi_f(C_{2k+1}) = 2 + 1/k.$$

Proof Since $C_{2k+1} = K_{(2k+1)/k}$, we have $\chi_f(C_{2k+1}) \leq \chi_c(C_{2k+1}) = 2 + 1/k$. (A k -tuple $2k+1$ -colouring of C_{2k+1} is illustrated, in the case $k = 2$, in Fig 6.3.) On the other hand, $\chi_f(C_{2k+1}) \geq (2k+1)/k$ by Corollary 6.24. □

At this point it would be nice to describe explicitly when $K(n, k) \rightarrow K(n', k')$. Unfortunately, this turns out to be a difficult problem. Figure 6.5 illustrates the homomorphisms that are known to exist. To avoid trivialities, we shall only consider Knese graphs $K(n, k)$ with $n \geq 2k$.

Proposition 6.26 *For all integers n, k , with $n \geq 2k$ and $k \geq 1$,*

1. $K(n, k) \rightarrow K(n+1, k)$
2. $K(n, k) \rightarrow K(tn, tk)$, for every positive integer t
3. $K(n, k) \rightarrow K(n-2, k-1)$, for $k > 1$.




FIG. 6.5. Homomorphisms among the Kneser graphs.

Proof Statement 1 is obvious from the definition, as $K(n, k)$ is actually a subgraph of $K(n + 1, k)$, if the ground sets are taken to be $\{1, 2, \dots, n\}$ and $\{1, 2, \dots, n, n + 1\}$. Statement 2 is best seen if the ground set of $K(n, k)$ is as above, but the ground set of $K(tn, tk)$ is taken to be $\{1, 2, \dots, t\} \times \{1, 2, \dots, n\}$. This allows us to map each vertex $\{x_1, x_2, \dots, x_k\}$ of $K(n, k)$ to the vertex $\{1, 2, \dots, t\} \times \{x_1, x_2, \dots, x_k\}$ of $K(tn, tk)$. This is easily seen to be a homomorphism. To prove 3, we shall transform each k -element subset A of $\{1, 2, \dots, n\}$ into a $(k - 1)$ -element subset A' of $\{1, 2, \dots, n - 2\}$, in such a way that disjoint sets remain disjoint after the transformation. The first idea for such a transformation might be to simply eliminate the largest element of A . Unfortunately, this will not eliminate both n and $n - 1$ from those sets A that contain both of them (but will work fine for all other sets A). Hence we define A' as follows.

If A does not contain *both* n and $n - 1$, then $A' = A - \max A$. Otherwise $A' = A - \{n, n - 1\} \cup \{x\}$, where x is the maximum element absent from A .

It only remains to prove that $A_1 \cap A_2 = \emptyset$ implies $A'_1 \cap A'_2 = \emptyset$. If not, it could only have been the process of adding x that created an intersection. Obviously both sets A_1, A_2 could not have been added to (only one can contain $n, n - 1$), so suppose x was added to A_1 , but already lies in A_2 . Then it is easy to see that x must have been the largest element of A_2 and hence does not belong to A'_2 . \square

It seems possible that no other homomorphisms amongst the Kneser graphs exist, apart from the above three kinds and their composition (as illustrated in Fig. 6.5). As always, the nonexistence of homomorphisms is harder to demon-

strate. An easy case is the following.

Proposition 6.27 *If $2 \leq n'/k' < n/k$, then*

$$K(n, k) \not\rightarrow K(n', k').$$

Proof According to the Claim in the proof of Theorem 6.3, the independence ratio of $K_{n/k}$ is k/n . Proposition 6.19 implies that $K_{n/k} \rightarrow K(n, k)$, and hence $i(K(n, k)) \leq i(K_{n/k}) = k/n$, by the No-Homomorphism Lemma (Corollary 1.23). On the other hand, the set S of all k -tuples of $\{1, 2, \dots, n\}$ which contain the element 1 forms an independent set in $K(n, k)$, with $\binom{n-1}{k-1}$ vertices, whence the independence ratio is $i(K(n, k)) = k/n$.

Since each $i(K(n, k)) = k/n$, the result follows by one more application of the No-Homomorphism Lemma. \square

Note that in the proof above we have derived the following fact, known as (a part of) the Erdős–Ko–Rado Theorem [91].

Corollary 6.28 *The independence number of the Kneser graph $K(n, k)$ (with $n \geq 2k$) is $\binom{n-1}{k-1}$.* \square

Proposition 6.1 is a useful tool to demonstrate nonexistence of certain homomorphisms amongst the Kneser graphs. (Equivalently, the No-Homomorphism Lemma, Corollary 1.23, can be used, cf. Exercise 14.)

The strongest result on the nonexistence of homomorphisms amongst the Kneser graphs is the following theorem of Lovász.

Theorem 6.29 *For every $n, k, n \geq 2k$*

$$\chi(K(n, k)) = n - 2k + 2.$$

Proof The fact that $\chi(K(n, k)) \leq n - 2k + 2$ follows from Proposition 6.26, since

$$K(n, k) \rightarrow K(n - 2, k - 1) \rightarrow K(n - 4, k - 2) \cdots \rightarrow K(n - 2k + 2, 1).$$

Thus the difficult part of the theorem is equivalent to showing that $K(n, k) \not\rightarrow K(n - 2k + 1, 1)$. Lovász's proof, and most subsequent proofs of this result used topology (either directly or indirectly), usually in some form of Borsuk's Antipodal Theorem [235]. We shall use the following geometric result, which is an easy consequence of Borsuk's Antipodal Theorem [235].

Lemma 6.30 [79] *Assume that, for each $i = 1, 2, \dots, d$, we have a family \mathcal{C}_i of compact convex sets in d -space which does not contain two disjoint sets. Then some hyperplane intersects all sets of all families \mathcal{C}_i .*

The trick in applying this geometric lemma is to choose for $K(n, k)$ the base set X of n points in general position in d -space. Then each vertex v of $K(n, k)$ corresponds to a finite set S_v in d -space, and the convex hull C_v of the set S_v is a compact convex set.

Suppose now that there is a homomorphism $f : K(n, k) \rightarrow K_d$, and consider the associated partition F_1, F_2, \dots, F_d of $V(K(n, k))$ into independent sets of $K(n, k)$. For each $i = 1, 2, \dots, d$, let \mathcal{C}_i denote the set of all $C_v, v \in F_i$. Since F_i is an independent set, no two such C_v can be disjoint (since no two S_v 's can be disjoint). According to the geometric lemma, some hyperplane P intersects all $C_v, v \in V(K(n, k))$. Since the base set X of n points was chosen in general position in d -space, P can contain at most d of them, and so one of the two halfspaces determined by P contains at least $(n - d)/2$ points. Thus $(n - d)/2 \leq k - 1$, else some k -element subset S_v of X (and hence also the corresponding C_v) would have been disjoint from P . However, this inequality implies that $d \geq n - 2k + 2$, as claimed. \square

Corollary 6.31 *If $n' - 2k' < n - 2k$ then $K(n, k) \not\rightarrow K(n', k')$.*

These results go some way towards proving that there are no homomorphisms amongst the Kneser graphs other than those illustrated in Fig. 6.5. Each row in the figure contains all $K(n, k)$ with a fixed k ; each column contains all $K(n, k)$ with a fixed value of $n - 2k$. The three kinds of homomorphisms from Proposition 6.26 go to the right (corresponding to item 1 of the proposition), or down (corresponding to item 3), and go from $K(n, k)$ to all $K(tn, tk)$ (item 2). The last corollary implies that there can be no homomorphism from a graph in one column to any graph in a column to the left. Proposition 6.27 implies that no homomorphisms can be any ‘steeper’ than those of Proposition 6.26, item 2. We illustrate the situation again in Fig. 6.6.

From $K(n, k)$ there is a homomorphism to $K((t - 1)n, (t - 1)k)$, and hence to all Kneser graphs below and to the right of $K((t - 1)n, (t - 1)k)$; there is also a homomorphism to $K(tn, tk)$, and to all Kneser graphs below and to the right of $K(tn, tk)$. It is also known that there are no homomorphisms of $K(n, k)$ into any Kneser graphs on the hypotenuse of the marked triangle, i.e., that $K(n, k) \not\rightarrow K(n', k')$ when $n/k = n'/k' > 2$ and k does not divide k' , Exercise 13. Since no homomorphisms can go to the left, or any steeper than those to $K((t - 1)n, (t - 1)k)$ and $K(tn, tk)$, this leaves open just the marked triangle. (There is one such triangle for each value of $t \geq 2$.) The general belief is that $K(n, k)$ does not admit homomorphisms into the interior of the marked triangle, for any values of t, n, k with $t \geq 2, n > 2k, k \geq 1$. Clearly, it suffices to prove that $K(n, k)$ is not homomorphic to the Kneser graph in the lower right corner of the triangle, since all others are homomorphic to it. In other words, the following is the most natural conjecture.

Conjecture 6.32

$$K(n, k) \not\rightarrow K(tn - 2k + 1, tk - k + 1).$$

There is another important connection between the rational complete graphs $K_{p/q}$ and the Kneser graphs $K(p, q)$.

We say that two k -tuples of nonnegative integers $x_1 < x_2 < \dots < x_k$ and $y_1 < y_2 < \dots < y_k$ are *interlaced*, if $x_1 < y_1 < x_2 < y_2 < \dots < x_k < y_k$ or




FIG. 6.6. Homomorphisms from the Kneser graph $K(n, k)$.

$y_1 < x_1 < y_2 < x_2 < \cdots < y_k < x_k$. An *interlaced k -tuple n -colouring* of G is an assignment of k -tuples of integers from $0, 1, \dots, n-1$ to the vertices of G such that adjacent vertices obtain interlaced k -tuples. Of course, an interlaced k -tuple n -colouring of G is a homomorphism of G to an obvious target graph: $I(n, k)$ is the *spanning* subgraph of $K(n, k)$ (that is, one containing all the vertices of $K(n, k)$), in which two k -tuples are adjacent just if they are interlaced. The *interlaced chromatic number* of G , denoted $\chi_i(G)$, is the infimum of the fractions n/k such that G admits an interlaced k -tuple n -colouring.

Recall that $K_{p/q}$ is (isomorphic to) an induced subgraph of $K(p, q)$ (Proposition 6.19), i.e., contains *some* vertices of $K(p, q)$ and *all* edges between them. In contrast, $I(p, q)$ contains *all* vertices of $K(p, q)$ and *some* edges between them. Yet, surprisingly, the two graphs turn out to be homomorphically equivalent, and so the interlaced chromatic number and the circular chromatic number are

equal.

Theorem 6.33 *The graphs $K_{p/q}$ and $I(p, q)$ are homomorphically equivalent, for all $1 \leq q \leq p$.*

Proof We first construct a homomorphism $f : K_{p/q} \rightarrow I(p, q)$. For each $x \in \{0, 1, \dots, p-1\}$, we set $f(x)$ to be the q -tuple

$$\left\lfloor \frac{x+0p}{q} \right\rfloor < \left\lfloor \frac{x+1p}{q} \right\rfloor < \dots < \left\lfloor \frac{x+(q-1)p}{q} \right\rfloor.$$

To see that this is a homomorphism, we shall show that if $q \leq |x-y| \leq p-q$, then the two q -tuples $f(x)$ and $f(y)$ are interlaced. Indeed, for any i we have

$$\left\lfloor \frac{y+ip}{q} \right\rfloor < \left\lfloor \frac{y+q+ip}{q} \right\rfloor \leq \left\lfloor \frac{x+ip}{q} \right\rfloor < \left\lfloor \frac{x+q+ip}{q} \right\rfloor \leq \left\lfloor \frac{y+(i+1)p}{q} \right\rfloor.$$

Now we shall construct a homomorphism $g : I(p, q) \rightarrow K_{p/q}$. We shall compare each q -tuple $X = x_1 < x_2 < \dots < x_q$ of $I(p, q)$ to the ‘evenly spread’ rational numbers $p/q < 2p/q < \dots < p$. The image $g(X)$ is defined to be the product iq (reduced modulo p), with i being the subscript where $i(p/q) - x_i$ is minimized (ties are broken arbitrarily). To see that this is a homomorphism, we shall show that if $X = x_1 < x_2 < \dots < x_q$ and $Y = y_1 < y_2 < \dots < y_q$ are interlaced, say $x_1 < y_1 < x_2 < y_2 < \dots < x_k < y_k$, and if i minimizes $i(p/q) - x_i$ and j minimizes $j(p/q) - y_j$, then we must have

$$q \leq |(i-j)q| \leq p-q$$

(modulo p).

Assume first that $i \leq j$. The minimality of i and j implies that

$$i\frac{p}{q} - x_i \leq (j+1)\frac{p}{q} - x_{j+1}$$

and

$$j\frac{p}{q} - y_j \leq i\frac{p}{q} - y_i,$$

i.e.,

$$(j-i)p \leq (y_j - y_i)q$$

and

$$(x_{j+1} - x_i)q \leq (j+1-i)p.$$

Because of the interlacing of X and Y we must have $y_j - y_i + 1 \leq y_j - x_i \leq x_{j+1} - x_i - 1$, and hence

$$(j-i)p + q \leq (y_j - y_i + 1)q \leq (y_j - x_i)q \leq (x_{j+1} - x_i - 1)q \leq (j-i)p + p - q,$$

whence $q \leq (j-i)q \leq p-q$, modulo p , as desired. (When $j = q$ we actually have to make a small modification, which we leave to the reader.)

In case $i > j$, we derive $x_i - x_{j+1} + 1 \leq x_i - y_j \leq y_i - y_j - 1$ from the interleaving of X and Y , and using the reversed inequalities $(i - j - 1)p \leq (x_i - x_{j+1})q$ and $(y_i - y_j)q \leq (i - j)p$ from above, we derive

$$(i - j - 1)p + q \leq (x_i - x_{j+1} + 1)q \leq (x_i - y_j)q \leq (y_i - y_j - 1)q \leq (i - j - 1)p + p - q,$$

implying $q \leq (i - j)q \leq p - q$ modulo p . \square

Corollary 6.34 *For all graphs G , we have $\chi_i(G) = \chi_c(G)$.* \square

6.3 T -colourings

In Section 1.8 we discussed a frequency (or channel) assignment problem. In its simplest form, the problem may be modelled as follows. Let T be a fixed finite set of nonnegative integers including 0. A T -colouring of a graph G is an assignment f of nonnegative integers to the vertices of G such that $|f(u) - f(u')| \notin T$ for all edges uu' of G .

The graph G corresponds to possible interference among transmitters, f assigns frequencies to transmitters, and T consists of disallowed distances between frequencies assigned to transmitters that could interfere. The set T is characteristic of particular frequency assignment problems; for instance the assignment problem for UHF television used to need sets such as $T = \{0, 7, 14, 15\}$ [244].

Given the small spectrum of available frequencies, the typical frequency assignment problem asks for the smallest *span* of an assignment. The span of f is the difference between the maximum and minimum values of $f(u)$ over $u \in V(G)$. The span of G , denoted by $sp_T(G)$, is the minimum span of a T -colouring of G . (Minimizing the span is not the only possible objective; see Exercise 6.) When computing the span of G , we may assume that the minimum value is $f(u) = 0$; hence the span of G is the smallest maximum value of $f(u)$, $u \in V(G)$, over all T -colourings of G .

We now observe that T -colourings are once again homomorphisms to a special family of graphs. Let $G(T)$ be the infinite graph with the vertex set $0, 1, 2, \dots$, in which vertices i and j are adjacent just if $|i - j| \notin T$. Then a T -colouring of G is precisely a homomorphism of G to $G(T)$. Since T contains 0, the graph $G(T)$ has no loops; since T is finite, $G(T)$ contains cliques of arbitrary size, and, in particular, every G is homomorphic to $G(T)$. Denote by $G_p(T)$ the subgraph of $G(T)$ induced on the vertices $0, 1, 2, \dots, p - 1$. We have proved the following fact.

Proposition 6.35 *For any set T ,*

$$sp_T(G) = p - 1,$$

where p is the smallest integer such that $G \rightarrow G_p(T)$. \square

Corollary 6.36 *For any set T , if $G \rightarrow H$ then $sp_T(G) \leq sp_T(H)$.* \square

Thus, homomorphically equivalent graphs have the same span, and, in particular, we may restrict our attention to cores.

Corollary 6.37 *If C is the core of G , then $sp_T(G) = sp_T(C)$.* \square

For graphs whose cores are cliques, this reduces the problem to computing spans of cliques (see also Exercise 3). If T satisfies an additional property, this is so even for graphs whose cores are not cliques.

Proposition 6.38 *Suppose T is a set such that each $G_p(T)$ is homomorphically equivalent to a clique.*

Then

$$sp_T(G) = sp_T(K_{\chi(G)})$$

for all graphs G .

Proof We have $sp_T(G) \leq sp_T(K_{\chi(G)})$ since $G \rightarrow K_{\chi(G)}$. On the other hand, if $G \rightarrow G_p(T)$ and $G_p(T)$ is homomorphically equivalent to K_ω , then $G \rightarrow K_\omega$, whence $\chi(G) \leq \omega$. Therefore, $K_{\chi(G)} \rightarrow K_\omega$ and hence $K_{\chi(G)} \rightarrow G_p(T)$. \square

Unfortunately, even computing the span of cliques turns out to be a difficult problem. However, we have the following asymptotic result.

Theorem 6.39 *For each set T there exists a rational number r_T such that*

$$\lim_{n \rightarrow \infty} \frac{sp_T(K_n)}{n} = r_T.$$

More precisely, r_T will turn out to be the fractional chromatic number of the complement of the graph $G(T)$. For simplicity, let us denote the complement of $G(T)$ by $H(T)$, i.e., the vertices of $H(T)$ are nonnegative integers, and ij is an edge of $H(T)$ just if $i \neq j$ and $|i - j| \in T$. We will show that $r_T = \chi_f(H(T))$ satisfies the assertion in Theorem 6.39. (The definition of the fractional chromatic number applies without difficulty to infinite graphs.)

Proof Let $s_n = sp_T(K_n)$ and $m = \max(T)$. We focus on the ratio s_n/n . For each n , there exists a T -colouring f_n of K_n such that $f_n(1) = 0 < f_n(2) < \dots < f_n(n) = s_n$. We now define an n -tuple colouring of $H(T)$. To each vertex j of $H(T)$, we assign the n -tuple consisting of the integers

$$j - f_n(n) < j - f_n(n-1) < \dots < j - f_n(2) < j - f_n(1) = j,$$

with all differences reduced modulo $s_n + m + 1$. We claim that adjacent vertices of $H(T)$ will obtain disjoint sets of colours. Indeed, let j and j' be adjacent vertices of $H(T)$, say, $j < j'$ and $j' - j \in T$. If $j - f_n(i) \simeq j' - f_n(i')$ modulo $s_n + m + 1$ for some i, i' , then $f_n(i') - f_n(i) \simeq j' - j$ modulo $s_n + m + 1$. Since $0 < j' - j \leq m$ and $-s_n \leq f_n(i') - f_n(i) \leq s_n$, we must have $f_n(i') - f_n(i) = j' - j \in T$, contradicting the definition of f_n . Therefore,

$$\chi_f(H(T)) \leq s_n + m + 1/n,$$

for any integer n .

On the other hand, consider the subgraph H_n of $H(T)$ induced by the vertices $0, 1, 2, \dots, s_n$. It follows from the definition of f_n that the set $f_n(1), f_n(2), \dots, f_n(s_n)$ is a maximum independent set in H_n . Therefore,

$$\chi_f(H(T)) \geq \chi_f(H_n) \geq s_n + \frac{1}{n},$$

the last inequality being a consequence of Corollary 6.24.

In conclusion, we have, for any n ,

$$\chi_f(H(T)) - \frac{m+1}{n} \leq \frac{s_n}{n} \leq \chi_f(H(T)) - \frac{1}{n}.$$

Since both the upper and lower bound have the limit $\chi_f(H(T))$, we have $\lim_{n \rightarrow \infty} sp_T(K_n)/n = \chi_f(H(T))$. \square

Corollary 6.40 *For every set T ,*

$$\lim_{n \rightarrow \infty} \frac{sp_T(K_n)}{n} = \chi_f(H(T)).$$

\square

Corollary 6.41 *Suppose T consists of integers without a common divisor greater than one. Then $r_T = 2$ if and only if all nonzero members of T are odd.* \square

6.4 Oriented and acyclic colourings

In this section we shall consider *oriented* graphs, i.e., digraphs without opposite arcs. Recall that if the graph G is the underlying graph of the oriented graph G' , then G' is called an orientation of G .

For oriented graphs there is a natural analogue of the usual chromatic number. The chromatic number of a graph G can be viewed as the smallest number of vertices in a graph H such that $G \rightarrow H$. Indeed, such a graph H must be complete, else we could find a smaller graph H' with $H \rightarrow H'$ and hence $G \rightarrow H'$ (H' obtained from H by identifying any pair of nonadjacent vertices). If we apply a similar definition to oriented graphs, we arrive at the concept of the *oriented chromatic number* $\chi_o(G)$ of an oriented graph G : $\chi_o(G)$ is the minimum number of vertices in an oriented graph H with $G \rightarrow H$. Note that it is no longer true that H must have all pairs of vertices adjacent - the trick of identifying a pair of nonadjacent vertices may fail because such an identification could produce opposite arcs. For instance, the oriented chromatic number of a directed four-cycle is four, since any two nonadjacent vertices are joined by a directed path of length two and hence cannot be identified. For the same reason, the oriented chromatic number of a directed five-cycle is five. On the other hand, $\chi_o(C_6) = 3$.

The 'partition view' of oriented colourings is particularly useful. An *oriented k -colouring* of an oriented graph G is a partition of $V(G)$ into k independent sets S_1, S_2, \dots, S_k such that any two sets S_i and S_j have all arcs between them (if any) oriented in the same direction, i.e., either all from S_i to S_j or all from S_j

to S_i . Then the oriented chromatic number of G , $\chi_o(G)$ is the minimum integer k such that G admits an oriented k -colouring.

An oriented k -colouring of G may equivalently be viewed as a function f on the vertices of G such that adjacent vertices obtain different colours, and there do not exist two different arcs $xy, zt \in E(G)$ such that $f(x) = f(t)$ and $f(y) = f(z)$.

We can extend the definition to graphs G . The *oriented chromatic number of a graph* G , also denoted by $\chi_o(G)$, is the maximum $\chi_o(G')$ over all orientations G' of G . Therefore, the oriented chromatic number of a graph G is the oriented chromatic number of a ‘worst’ orientation of G .

Consider a graph with n vertices, m edges, and oriented chromatic number at most k . There are 2^m different orientations of the graph, and each partition of the vertices into k classes can be an oriented k -colouring for at most $2^{\binom{k}{2}}$ of these orientations. Since there are at most k^n such partitions, we must have the inequality

$$2^m \leq 2^{\binom{k}{2}} \cdot k^n.$$

The following example is instructive. In Fig. 6.7 we depict a graph obtained from K_n by subdividing each edge. Its chromatic number is two, yet its oriented chromatic number is at least n , since we can orient it so that any pair of the original vertices is joined by a directed path of length two and hence must obtain a different colour.




FIG. 6.7. A bipartite graph with oriented chromatic number greater than four.

Proposition 6.42 *If G is a tree, then $\chi_o(G) \leq 3$.*

Proof Every oriented tree admits a homomorphism to \vec{C}_3 (the directed three-cycle). We can arbitrarily choose the image of one vertex of the tree, and then all other images are determined recursively from the definition of a homomorphism. Specifically, suppose the vertex v is mapped to i . If vw is an arc of the tree G ,

then w will be mapped to $i + 1$ (modulo 3); and if uv is an arc of G , then u will be mapped to $i - 1$ (modulo 3). \square

Note that the same proof shows that every oriented tree admits a homomorphism to the directed four-cycle \vec{C}_4 . Moreover, an oriented tree which is coloured by two colours (so that adjacent vertices have different colours) admits a homomorphism to \vec{C}_4 in which vertices of one colour are mapped to 0, 2 and vertices of the other colour to 1, 3. (Recall that \vec{C}_4 has vertices 0, 1, 2, 3.)

It is in general not true that the oriented chromatic number of G is at most the greatest oriented chromatic number of its components. For instance, if A is the directed three-cycle and B the transitive triple (Fig. 6.8), then each have oriented chromatic number three but their disjoint union has oriented chromatic number four. On the other hand, it follows from the above proposition, and the remarks following it, that any oriented forest has an oriented three-colouring, and any oriented two-coloured forest has a homomorphism to \vec{C}_4 in which all vertices of one colour have even images and all vertices of the other colour have odd images.




FIG. 6.8. A digraph of oriented chromatic number four.

There is an interesting connection between oriented colourings and another variant of colouring. An *acyclic k -colouring* of a graph G is a k -colouring of G in the usual sense, in which every cycle uses at least three colours. (In other words, a colouring of G is an acyclic colouring just if all the bicoloured sets of vertices are forests.) The *acyclic chromatic number*, $\chi_a(G)$, of a graph G is the minimum integer k such that G admits an acyclic k -colouring.

In analogy with the Four Colour Theorem, it has been proved [37] that every planar graph admits an acyclic five-colouring. The proof is not as difficult as the proof of the four-colour theorem (and is done without the aid of a computer), but it *is* difficult—and uses an ‘unavoidable set’ of some 450 ‘reducible’ configurations. It is also known that this is best possible, i.e., there are graphs which cannot be acyclically four-coloured.

Theorem 6.43 *If G admits an acyclic k -colouring, then $\chi_o(G) \leq k2^{k-1}$.*

Corollary 6.44 *Every oriented planar graph admits an oriented 80-colouring.*

\square

Corollary 6.45 *Graphs with bounded degree, or treewidth, or genus, have bounded oriented chromatic number.*

Proof It is known that graphs of bounded degree, or treewidth, or genus, have bounded acyclic chromatic numbers [69, 206]. \square

We now proceed with the **proof** of the theorem. Let c be an acyclic k -colouring of G . Consider any two colours i and j , say $i < j$. The set of vertices v of G with $c(v) = i$ or $c(v) = j$ induces a forest in G , which we shall denote by $F_{i,j}$. Each orientation G' of G yields a corresponding orientation $F'_{i,j}$ of $F_{i,j}$. Note that the vertices of $F'_{i,j}$ are coloured by c with the two colours i and j , so that adjacent vertices have different colours. As observed above, there is a homomorphism $c_{i,j} : F'_{i,j} \rightarrow C_4$ such that all vertices v with $c(v) = i$ map to even vertices of C_4 and all vertices v with $c(v) = j$ to odd vertices of C_4 .

We now define a mapping on the vertices of G' . Each vertex v of G' belongs to $k - 1$ bicoloured forests. Suppose $c(v) = l$; then v lies in the oriented forests

$$F'_{1,l}, F'_{2,l}, \dots, F'_{l-1,l}, F'_{l,l+1}, \dots, F'_{l,k}.$$

We define

$$f(v) = (l; c_{1,l}(v), c_{2,l}(v), \dots, c_{l-1,l}(v), c_{l,l+1}(v), \dots, c_{l,k}(v)).$$

The number of possible values f can take is $k \cdot 2^{k-1}$, since there are k possible values of l , and for each of them every other component of $f(v)$ has only one of two possible values (those for $c_{i,l}$ are 0, 2 and those for $c_{l,j}$, 1, 3).

We claim that f is an oriented colouring of G' . Indeed, it is clear that the colours of adjacent vertices are distinct, since they have different colours under the colouring c . Thus it remains to verify that the homomorphic image $f(G')$ is an *oriented* graph, i.e., that there do not exist arcs $xy, zt \in E(G')$ with $f(x) = f(t), f(y) = f(z)$. Otherwise, assume xy, zt are such arcs. Then $c(x) = c(t)$ and $c(y) = c(z)$. Without loss of generality assume that $c(x) = c(t) = i < j = c(y) = c(z)$. Then x, y, z, t are all vertices of $F'_{i,j}$, and $c_{i,j}$ is a homomorphism of $F'_{i,j}$ to C_4 . This is impossible as it would imply that both $c_{i,j}(x)c_{i,j}(y)$ and $c_{i,j}(z)c_{i,j}(t) = c_{i,j}(y)c_{i,j}(x)$ are arcs of the oriented graph C_4 . \square

On the other hand, it is known that there exist oriented planar graphs with oriented chromatic number at least 16 [317]. The construction is quite involved, so we present instead, in Fig. 6.9, an oriented graph of chromatic number 15. It is easy to check that it has, between any two nonadjacent vertices, a directed path of length two. Hence it does not have a homomorphism to a smaller oriented graph—its oriented chromatic number is equal to its number of vertices, i.e., 15.

Theorem 6.46 *Every oriented outerplanar graph admits an oriented seven-colouring.*

Proof We shall show that every oriented outerplanar graph admits a homomorphism to the *quadratic residue tournament* T on 7 vertices, depicted in Fig.




FIG. 6.9. A planar graph of oriented chromatic number 15.

6.10. (The vertices are $0, 1, 2, 3, 4, 5, 6$, and the arcs are ij such that $j - i$ is congruent to $1, 2$, or 4 modulo 7 —the quadratic residues modulo seven.) Otherwise, suppose that G is the minimum counterexample, an oriented outerplanar graph which does not admit a homomorphism to T and with the smallest number of vertices. We may assume that G is maximal outerplanar. Thus G contains a vertex x adjacent to precisely two other vertices y, z . Let G' denote the oriented graph obtained from G by removing x and adding yz if $yz, zy \notin E(G)$. Then G' is also outerplanar and hence has a homomorphism to T . The images of y, z are two distinct vertices of T . It is a nice property of T that for any two distinct vertices and any choice of directions there is a vertex joined to the chosen vertices in the chosen directions. Hence the homomorphism of G' to T can be extended to G . \square




FIG. 6.10. The quadratic residue tournament T on seven vertices.

Figure 6.11 shows an oriented outerplanar graph G which has $\chi_o(G) = 7$. (Note how the planar graph in Fig. 6.9 is constructed from two copies of this G .) It is a suprising fact, that the oriented and acyclic chromatic numbers are very closely related. Not only is the oriented chromatic number bounded by a function of the acyclic chromatic number, but also conversely, the acyclic chromatic number is bounded in terms of the oriented chromatic number. We first give a bound in terms of the oriented chromatic number, assuming the edges of the graph can be covered by a bounded number of trees.




FIG. 6.11. An outerplanar digraph with oriented chromatic number seven.

Proposition 6.47 *Let G be a graph. Suppose $\chi_o(G) \leq k$, and $E(G)$ can be partitioned into q forests. Then $\chi_a(G) \leq k^{q+1}$.*

Proof Assume T_1, T_2, \dots, T_q are forests whose edge sets partition $E(G)$. Let G_0 be an arbitrary orientation of G , and let G_i be obtained from G_0 by reversing all the arcs in T_i , $i = 1, 2, \dots, q$. Let f_i be an oriented k -colouring of G_i . We can now define an acyclic colouring of G by giving each x the colour $(f_0(x), f_1(x), \dots, f_q(x))$. It is clear that adjacent vertices of G obtain different colours, as each f_i gives them different images. Assume that there is a cycle C in G whose vertices are alternatingly coloured by two colours, say

$$(a_0, a_1, \dots, a_q) \text{ and } (b_0, b_1, \dots, b_q).$$

Since f_0 is an oriented colouring of G_0 , all edges between vertices of G_0 coloured a_0 and those coloured b_0 go in the same direction, thus C is oriented in G_0 without a directed path of length two. The same argument applies for each G_i —however, this cannot be true as there is a forest T_j that contains (and hence reverses) some, but not all, arcs of C . \square

It is a fundamental result of Nash-Williams [254] that the minimum number of forests to cover the edges of a graph G is equal to the maximum of

$$\left\lceil \frac{|E(H)|}{|V(H)| - 1} \right\rceil$$

over all subgraphs H of G .

Theorem 6.48 *Any graph G with $\chi_o(G) = k$ can have its edges covered by at most $\lceil \log k + k/2 \rceil$ forests.*

(The logarithms are base two.)

Proof Let H be a subgraph of G with n vertices and m edges. If $n \leq k$, then $m/(n-1) \leq n/2 \leq k/2$. Thus assume that $n > k$. Since $\chi_o(H) \leq \chi_o(G) \leq k$, we can apply the inequality $2^{\binom{k}{2} k^n} \geq 2^m$ proved just prior to Proposition 6.42, showing that $\log k \geq m/n - \binom{k}{2}/n$. Hence we have

$$\begin{aligned}
\log k &\geq \frac{m}{n} - \frac{k(k-1)}{2n} \\
&> \frac{m}{n-1} - \frac{m}{n(n-1)} - \frac{k-1}{2} \\
&\geq \frac{m}{n-1} - \frac{1}{2} - \frac{k-1}{2} \\
&\geq \frac{m}{n-1} - \frac{k}{2}.
\end{aligned}$$

Therefore $m/(n-1) \leq \log k + k/2$. □

Corollary 6.49 *If G admits an oriented k -colouring, then*

$$\chi_a(G) \leq k^{1+\lceil k/2+\log k \rceil}.$$

□

We know this bound is not optimal—in fact we have ignored some obvious ways of significantly improving the bound. At this point we just note that χ_a is bounded by a function of χ_o .

6.5 Remarks

In writing Section 6.1 we have benefited from an early version of X. Zhu's comprehensive survey [352]. We have tried to offer a different perspective from the published version, which is widely read. The circular chromatic number was introduced by A. Vince in [332] (and called the star chromatic number). The more combinatorial view taken here was pioneered in [35]. Theorem 6.12 is due to X. Zhu [345], and it marked the beginning of the real appreciation of this concept, now generally believed to be an important extension of the chromatic number. A wealth of generalizations of classical results about ordinary colourings have been extended to circular colourings, including, say, Corollary 3.13: *for each rational number $r \geq 2$ and any integer $\ell \geq 3$ there exists a graph G of girth at least ℓ and circular chromatic number r* [282, 346], see also [283]. Theorem 6.13 is due to L. Goddyn, M. Tarsi, and C. Q. Zhang [121]; Corollary 6.15 is due to G. J. Minty [246]. Theorem 6.17 and the application in maximum concurrency are due to [28], which also contains a proof of Corollary 6.34. A more combinatorial proof is given in [29]; Theorem 6.33 and its use for Corollary 6.34 is due to X. Zhu. Fractional colourings were introduced by D. Geller and S. Stahl [118, 319], who proved the basic results in Proposition 6.26, and Kneser graphs arose from [201]. M. Kneser also conjectured Theorem 6.29, first proved by L. Lovász [226]. The proof given here is adapted from V. L. Dolnikov [79]; see [235]. These techniques have been further developed in [14]. Conjecture 6.32 has been formulated independently (in equivalent ways) by several people, starting as early as [320]. Progress has been made [65, 117]; we refer to [122] for a

summary. The best results to date seem to be due to Y. Kirsch and M. Perles; in an unpublished work they have verified the conjecture for all $k \leq 3$. The integer programming formulation of the fractional chromatic number, Theorem 6.22, is from [65], where it is also attributed independently to E. N. Gilbert; a game-theoretic formulation from [67] is given in Exercise 22. For a general overview of fractional colourings, cliques, and so on, we recommend [314]. The interest in T -colourings arose from applications in frequency assignments [137, 301]; the connection to graph homomorphisms was first made in [221], which also contains Proposition 6.38; see also [127, 220] and papers cited there. Theorem 6.39 was first proved in [299]. Acyclic colourings were investigated in the sixties by B. Grünbaum [131] and others [37]. Oriented colourings arose in connection with B. Courcelle's study of monadic second order logic in graphs [69]. Theorem 6.43 is from [300], Theorem 6.48 and Corollary 6.49 are from [206].

Exercise 5 is from [209]; the study of $(2, 1)$ -colourings is motivated by frequency assignment problems [88]. Exercise 16 is based on [117], where it is shown that, more generally, if $P \neq NP$ then there can be no approximation algorithm for the chromatic number that guarantees to colour an n -chromatic graph with at most $qn + d$ colours, for any $q < 2$. Exercise 20 is from [134, 348]. It is not known whether the parameter $I(G)$ is always a rational number. Exercise 24 is from [177]; the question of gaps amongst the possible circular chromatic numbers of series-parallel graphs, and other families of graphs given by forbidden minors, are discussed in [288, 219, 351].

6.6 Exercises

1. Prove that $\chi_{k+\ell}(G) \leq \chi_k(G) + \chi_\ell(G)$, for any graph G and integers k, ℓ . Deduce that $\chi_f(G) = \lim_{k \rightarrow \infty} \chi_k(G)/k$.
2. Prove that if the complement of G is disconnected, then $\chi_c(G) = \chi(G)$.
3. Let G be a graph with maximum clique size ω and chromatic number χ . Show that for any T we have $sp_T(K_\omega) \leq sp_T(G) \leq sp_T(K_\chi)$.
4. Prove that $sp_T(K_{m+n}) \geq sp_T(K_m) + sp_T(K_n)$ with the inequality being strict unless m or n is 0. Deduce that $\lim s_n/n$ exists.
5. A homomorphism f of a graph G to a graph H is *locally injective* if distinct neighbours of any vertex x are mapped to distinct neighbours of $f(x)$. A mapping $f : V(G) \rightarrow V(H)$ is a $(2, 1)$ - H -colouring if adjacent vertices of G are mapped to vertices of H of distance at least two, and vertices of G of distance two still map to different vertices of H .
Prove that $f : G \rightarrow H$ is a $(2, 1)$ - H -colouring if and only if f is a locally injective homomorphism of G to the complement of H .
6. Prove that for $T = \{0\}$ we have $sp_T(G) = \chi(G) - 1$.
Prove that for any T , and all graphs G , the minimum number of distinct values $f(u)$, $u \in V(G)$, over all T -colourings f of G , is equal to $\chi(G)$.
Show that for $G = C_5$ and $T = \{0, 1, 4, 5\}$ there is no T -colouring f that simultaneously minimizes the span of f and the number of distinct values $f(u)$, $u \in V(G)$.

7. Prove the converse of Proposition 6.38.
8. Assume $T = \{0, 1, 2, \dots, r\} \cup S$ where S contains no multiples of $r + 1$. Show that $G(T)$ is homomorphically equivalent to a clique and compute $sp_T(H)$, for any graph H .
9. One can consider k -tuple T -colourings, assigning a k -tuple of nonnegative integers to each vertex, with all differences between integers assigned to adjacent vertices being outside of T . The *span* of a k -tuple T -colouring is the difference between the largest and smallest integers used in any of the k -tuples. For $T = \{0, 1\}$ show that a graph G with $\chi(G) = 3$ has $sp_T^2(G) = 6$ if and only if $G \rightarrow C_5$.
10. Let T consists of $0, a$, and b , where $\gcd(a, b) = 1$ and a, b have different parity. Prove that $r_T = 2(a + b)/a + b - 1$. (Hint: Use Corollary 6.40.)
11. Prove that $\chi_f(G) = \chi(G)$ if $|V(G)| = \chi(G)\alpha(G)$.
12. Prove that $\chi(G)$ is the infimum of the set of rational numbers r such that the vertices of G can be assigned to unit length open subintervals of $(0, r)$, so that adjacent vertices are assigned disjoint intervals.
13. Prove that $K(n, k) \rightarrow K(n', k')$ with $n/k = n'/k' > 2$ if and only if k divides k' .
14. Derive the No-Homomorphism Lemma, Corollary 1.23, from Proposition 6.21 and Corollary 6.24.
15. Recall the lexicographic product from Exercise 14 in Chapter 2: $G[H]$ has the vertex set $V(G) \times V(H)$ and edges $(u, v)(u', v')$ where $uu' \in E(G)$, or $u = u'$ and $vv' \in E(H)$.
Prove that $\chi(G[H]) = \chi_{\chi(H)}(G)$.
16. Prove that $\chi_3(K(7, 3)) = 7$ and $\chi_4(K(7, 3)) = 10$.
Deduce that no polynomial algorithm can guarantee to colour a 7-colourable graph by at most 9 colours, unless $P = NP$.
17. Prove that $\chi_a(G)$ is unbounded for bipartite graphs, but that $\chi_a(G) \leq f(\max(\chi(H)))$, where the maximum is taken over all minors H of G .
18. Let U be the *unit distance graph*, in which the vertices are the points in the plane and two vertices are adjacent just if they have distance one. Prove that the chromatic number of U is 4, 5, or 6.
19. [115] Prove that if $\chi_c(G) = \chi_f(G)$ then for any graph H we have $\chi_c(G[H]) = \chi_c(G)\chi(H)$. (The lexicographic product $G[H]$ is defined in Exercise 15 above.)
20. [134, 348] Recall the independence ratio $i(G)$ from Corollary 1.23 (the No-Homomorphism Lemma).
Show that $i(G \square H) \leq \min(i(G), i(H))$.
Let $G^{[1]} = G$, $G^{[k]} = G \square G^{[k-1]}$ for $k > 1$, and let $I(G) = \lim i(G^{[k]})$ as k goes to infinity.
Show that $I(G) = i(G)$ if $G \square G \rightarrow G$.
Show that $G \rightarrow H$ implies that $I(H) \leq I(G)$.
Show that $1/I(G)$ lies between $\chi_f(G)$ and $\chi_c(G)$.

21. [329] Prove that $K(n, k)$ and $K_n^{K_k}$ are homomorphically equivalent.
22. [67] Prove that $\chi_f(G) = 1/v$ where v is the value of the following two-person zero-sum game on G . Player A chooses an independent set I while player B chooses a vertex x . The payoff to player B is 1 if $x \in I$, and 0 otherwise.
23. [43] The k -th *power* of a directed path \vec{P}_k has vertices $0, 1, \dots, k$ (like \vec{P}_k), and arcs ij such that $0 < j - i \leq k$.
 Prove that the minimum power of a directed path to which a given digraph G is homomorphic is equal to the maximum *imbalance* (ratio of forward and backward arcs) in any cycle of G . Deduce the second statement in Proposition 1.13.
24. [177] Prove that a series-parallel graph G has $\chi_c(G) \leq 8/3$ if it has no triangle, but $\chi_c(K_3) = 3$. (Series-parallel graphs are defined in Exercise 7 in Chapter 4.)

REFERENCES

- [1] M. E. Adams, J. Nešetřil, and J. Sichler, Quotients of rigid graphs, *J. Combin. Theory B* **30** (1981) 351–359.
- [2] M. O. Albertson and K. L. Collins, Homomorphisms of 3-chromatic graphs, *Discrete Math.* **54** (1985) 127–132.
- [3] N. Alon, Covering graphs by the minimum number of equivalence relations, *Combinatorica* **6** (1986) 201–206.
- [4] N. Alon and J. Spencer, *The Probabilistic Method*, Wiley, New York (1992).
- [5] N. Aronszajn and P. Panitchpakdi, Extension of uniformly continuous transformations and hyperconvex metric spaces, *Pacific J. Math.* **6** (1956) 405–439.
- [6] B. Aspvall, F. Plass, and R. E. Tarjan, A linear time algorithm for testing the truth of certain quantified Boolean formulas, *Inf. Proc. Lett.* **8** 1979 121–123.
- [7] S. P. Avann, Metric ternary semi-distributive lattices, *Proc. Amer. Math. Soc.* **12** (1961) 407–414.
- [8] R. Bačák, Graph homomorphisms and semidefinite programming, Ph.D. Thesis, Simon Fraser University (1996).
- [9] R. Bačák, Equitable graph homomorphisms, unpublished manuscript (1997).
- [10] L. Babai, Automorphism groups of graphs and edge contraction, *Discrete Math.* **8** (1974) 13–22.
- [11] L. Babai and J. Nešetřil, High chromatic rigid graphs I., *Combinatorics* (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. I, *Colloq. Math. Soc. János Bolyai* **18** North-Holland, Amsterdam New York (1978) pp. 53–60.
- [12] L. Babai and J. Nešetřil, High chromatic rigid graphs II., *Algebraic and geometric combinatorics*, North-Holland Math. Stud. **65** North-Holland, Amsterdam (1982) 55–61.
- [13] L. Babai and A. Pultr, Endomorphism monoids and topological subgraphs of graphs, *J. Combin. Theory B* **28** (1980) 278–283.
- [14] E. Babson, D. N. Kozlov, Topological obstructions to graph colorings, arXiv:math.CO/0305300v2 (2003).
- [15] H. -J. Bandelt, Characterizing median graphs, unpublished manuscript (1982).
- [16] H. -J. Bandelt, Retracts of hypercubes, *J. Graph Theory* **8** (1984) 501–510.
- [17] H. -J. Bandelt, Graphs with edge-preserving majority functions, *Discrete Math.* **103** **1992** 1–5.
- [18] H. -J. Bandelt, A. Dählmann, and H. Schütte, Absolute retracts of bipartite graphs, *Discrete Appl. Math.* **16** (1987) 191–215.

- [19] H. -J. Bandelt, M. Farber, and P. Hell, Absolute reflexive retracts and absolute bipartite retracts, *Discrete Appl. Math.* **44** (1993) 9–20.
- [20] H. -J. Bandelt and E. Pesch, Dismantling absolute retracts of reflexive graphs, *European J. Combin.* **10** (1989) 211–220.
- [21] H. -J. Bandelt and E. Pesch, Efficient characterizations of n -chromatic absolute retracts, *J. Combin. Theory B* **53** (1991) 5–31.
- [22] H. -J. Bandelt and E. Prisner, Clique graphs and Helly graphs, *J. Combin. Theory B* **51** (1991) 34–45.
- [23] J. Bang-Jensen and G. Gutin, *Digraphs: Theory, Algorithms and Applications*, Springer (2001).
- [24] J. Bang-Jensen and P. Hell, The effect of two cycles on the complexity of colourings by digraphs, *Discrete Appl. Math.* **26** (1990) 1–23.
- [25] J. Bang-Jensen, P. Hell, and G. MacGillivray, The complexity of colouring by semicomplete digraphs, *SIAM J. on Discrete Math.* **1** (1988) 281–298.
- [26] J. Bang-Jensen, P. Hell, and G. MacGillivray, On the complexity of colouring by superdigraphs of bipartite graphs, *Discrete Math.* **109** (1992) 27–44.
- [27] J. Bang-Jensen, P. Hell, and G. MacGillivray, Hereditarily hard H -colouring problems, *Discrete Math.* **138** (1995) 75–92.
- [28] V. C. Barbosa and E. Gafni, Concurrency in heavily loaded neighbourhood-constrained systems, *ACM Transactions on Prog. Lang. and Systems*, **11** (1989) 562–584.
- [29] V. C. Barbosa, The interleaved multichromatic number of a graph, unpublished manuscript (1990).
- [30] R. Bari, Homomorphism polynomials of graphs, *J. Combin. Inf. Syst. Science* **7** (1982) 56–64.
- [31] B. Bauslaugh, Core-like properties of infinite graphs and structures, *Discrete Math.* **138** (1995) 101–111.
- [32] G. Bloom and S. Burr, On unavoidable digraphs in orientations of graphs, *J. Graph Theory* **11** (1987) 453–462.
- [33] H. Bodländer, Achromatic number is NP -complete for cographs and interval graphs, *Inform. Processing Letters* **31** (1989) 135–138.
- [34] M. Boguszak, S. Poljak, and J. Tuma, A note on homomorphism interpolation theorems, *Comment. Math. Univ. Carolinae* **17** (1976) 105–109.
- [35] J. A. Bondy and P. Hell, A note on the star chromatic number, *J. Graph Theory* **14** (1990), 479–482.
- [36] J. Bondy and U.S.R. Murty, *Graph Theory with Applications*, North Holland, New York (1976).
- [37] O.V. Borodin, On acyclic coloring of planar graphs, *Discrete Math.* **25** (1979) 211–236.
- [38] K. Borsuk, Sur les retractes, *Fund. Math.* **17** (1931) 152–170.
- [39] A. Brandstädt, Partitions of graphs into one or two stable sets and cliques, *Discrete Math.* **152** (1996) 47–54.
- [40] A. Brandstädt, Corrigendum, *Discrete Math.* **186** (1998) 295.

- [41] R. C. Brewster and G. MacGillivray, Homomorphically full graphs, *Discrete Applied Math.* **66** (1996) 23–31.
- [42] R. C. Brewster, T. Feder, P. Hell, J. Huang, and G. MacGillivray, Near-unanimity functions and varieties of graphs, manuscript (2002).
- [43] R. C. Brewster and P. Hell, Homomorphisms to powers of digraphs, *Discrete Math.* **244** (2002) 31–41.
- [44] R. C. Brewster, P. Hell, and G. MacGillivray, The complexity of restricted graph homomorphisms, *Discrete Math.* **167/168** (1997) 145–154.
- [45] G. Brightwell and P. Winkler, Graph homomorphisms and phase transitions, *J. Combin. Theory B* **77** (1999) 415–435.
- [46] G. R. Brightwell and P. Winkler, Gibbs measures and dismantlable graphs, *J. Combin. Theory B* **78** (2000) 141–166.
- [47] A. E. Brouwer and H. J. Veldman, Contractibility and *NP*-completeness, *J. Graph Theory* **11** (1987) 71–79.
- [48] A. A. Bulatov, A dichotomy theorem for constraints on a three-element set, 43rd IEEE FOCS (2002) pp. 649–658.
- [49] A. A. Bulatov, Tractable conservative constraint satisfaction problems, 18th IEEE SOLCS (2003) pp. 321–330.
- [50] A. A. Bulatov, *H*-coloring dichotomy revisited, manuscript (2003).
- [51] A. A. Bulatov, The complexity of the counting constraint satisfaction problem, manuscript (2004).
- [52] A. A. Bulatov and V. Dalmau, Towards a dichotomy for the counting constraint satisfaction problem, 44th IEEE FOCS (2003) pp. 562–571.
- [53] A. A. Bulatov and P. G. Jeavons, Algebraic structures in combinatorial problems, Technical Report MATH-AL-4-2001, Technische Universität Dresden (2001).
- [54] A. A. Bulatov, P. G. Jeavons, and A. A. Krokhin, Constraint satisfaction problems and finite algebras, 27th ICALP (2000) pp. 272–282.
- [55] S. Burr, P. Erdős, and L. Lovász, On graphs of Ramsey type, *Ars Combinatoria* **1** (1976) 167–190.
- [56] N. Cairnie and K. Edwards, Some results on the achromatic number, *J. Graph Theory* **26** (1997) 129–136.
- [57] K. Cameron, E. M. Eschen, C. T. Hoàng, and R. Sritharan, The complexity of the list partition problem for graphs, SODA (2004).
- [58] P. Catlin, Homomorphisms as generalizations of graph colouring, *Congressus Num.* **50** (1985) 179–186.
- [59] P. Catlin, Graph homomorphisms onto the five-cycle, *J. Combin. Theory B* **45** (1988) 199–211.
- [60] P. Catlin, Hajós graph coloring conjecture; variations and counterexamples, *J. Combin. Theory B* **26** (1979) 268–274.
- [61] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas, The strong perfect graph theorem, manuscript (2002).
- [62] F. R. K. Chung, R. L. Graham, and M. E. Saks, A dynamic location problem for graphs, *Combinatorica* **9** (1989) 111–131.

- [63] V. Chvátal, P. Hell, L. Kučera, and J. Nešetřil, Every finite graph is a full subgraph of a rigid graph, *J. Combin. Theory* **11** (1971) 284–286.
- [64] V. Chvátal and J. Sichler, Chromatic automorphisms of graphs, *J. Combin. Theory B* **14** (1973) 209–215.
- [65] V. Chvátal, M. R. Garey, and D. S. Johnson, Two results concerning multicolouring, *Annals of Discrete Math.* **2** (1978) 151–154.
- [66] V. Chvátal, *Linear Programming*, W. H. Freeman and Company (1980).
- [67] F. H. Clarke, R. E. Jamison, Multicolorings, measures and games on graphs, *Discrete Math.* **14** (1976) 241–245.
- [68] S. A. Cook, A hierarchy of nondeterministic time complexity, *J. Comput. System Sci.* **7** (1973) 343–353.
- [69] B. Courcelle, The monadic second order logic of graphs VI: On several representations of graphs by relational structures, *J. Combin. Theory B* **54** (1994) 117–149.
- [70] K. Čulík, Zür Theorie der Graphen, *Časopis Pěst. Mat.* **83** (1957) 133–155.
- [71] N. G. de Bruijn, A combinatorial problem, *Indagationes Math.* **8** (1946) 461–467.
- [72] J. de Groot, Groups represented by homomorphism groups, *Math. Ann.* **138** (1959) 80–102.
- [73] J. de Rumeur (pseudonym of J. -C. Bermond, P. Fraigniaud, A. Germa, M. -C. Heydemann, E. Lazard, P. Michallon, A. Raspaud, D. Sotteau, M. Syska, and D. Trystram, *Communication dans les réseaux de processeurs*, Masson, Paris (1994).
- [74] R. Dechter, From local to global consistency, *Artificial Intelligence* **55** (1992) 87–107.
- [75] C. Delhomme, N. Sauer, Homomorphisms of products of graphs into graphs without fourcycles, *Combinatorica* **22** (2002) 35–46.
- [76] J. Diaz, M. Serna, and D. M. Thilikos, Counting H -colourings of partial k -trees, *Theoret. Comput. Sci.* **281** (2002) 291–309.
- [77] J. Diaz, M. Serna, and D. M. Thilikos, The complexity of parametrized H -colourings, a survey, manuscript (2002).
- [78] G. A. Dirac, Homomorphism theorems for graphs, *Math. Ann.* **153** (1964) 69–80.
- [79] V. L. Dolnikov, Transversals of families of sets, in *Studies in the theory of functions of several variables* (Russian), Yaroslav. Gos. Univ., Yaroslavl (1981), pp. 30–36.
- [80] A. W. M. Dress, Trees, tight extensions of metric spaces, and the cohomological dimension of certain groups: a note on combinatorial properties of metric spaces, *Advances Math.* **53** (1984) 321–402.
- [81] P. A. Dryer, Jr., C. Malon, and J. Nešetřil, Universal H -colourable graphs without a given configuration, *Discrete Math.* **250** (2002) 245–252.
- [82] D. Duffus and I. Rival, A structure theory for ordered sets, *Discrete Math.* **35** (1981) 53–118.

- [83] D. Duffus, B. Sands, and R. Woodrow, On the chromatic number of the products of graphs, *J. Graph Theory* **9** (1985) 487–495.
- [84] D. Duffus and N. Sauer, Lattices arising in categorical investigations of Hedetniemi’s conjecture, *Discrete Math.* **152** (1996) 125–139.
- [85] M. Dyer and C. Greenhill, The complexity of counting graph homomorphisms, *Random Structures and Algorithms* **17** (2000) 260–289.
- [86] N. Eaton and V. Rödl, Graphs of small dimension, *Combinatorica* **16** (1996) 59–85.
- [87] S. Eilenberg and S. MacLane, General theory of natural equivalences, *Trans. Amer. Math. Soc.* **58** (1945) 231–294.
- [88] A. Eisenblätter, M. Grötschel, and A. M. C. A. Koster, Frequency planning and ramifications of coloring, ZIB-Report 00-47 (2000).
- [89] M. El-Zahar and N. Sauer, The chromatic number of the product of two 4-chromatic graphs is 4, *Combinatorica* **5** (1985) 121–126.
- [90] P. Erdős, Graph theory and probability, *Canad. J. Math.* **11** (1959) 34–38.
- [91] P. Erdős, C. Ko, and R. Rado, Intersection theorems for systems of finite sets, *Quart. J. Math. Oxford Ser.* **12** (1961) 313–320.
- [92] P. Erdős, A. Rényi, On random graphs I., *Publ. Math. Debrecen* **6** (1959) 290–297.
- [93] P. Erdős and J. Spencer, *Probabilistic Methods in Combinatorics*, Academic Press, New York (1974).
- [94] M. Farber, G. Hahn, P. Hell, and D. J. Miller, Concerning the achromatic number of graphs, *J. Combin. Theory B* **40** (1986) 21–39.
- [95] T. Feder, Classification of homomorphisms to oriented cycles and of k -partite satisfiability, *SIAM J. on Discrete Math.* **14** (2001) 471–480.
- [96] T. Feder and P. Hell, List homomorphisms to reflexive graphs, *J. Combin. Theory B* **72** (1998) 236–250.
- [97] T. Feder and P. Hell, List constraint satisfaction and list matrix partition, manuscript (2003).
- [98] T. Feder, P. Hell, and J. Huang, List homomorphisms and circular arc graphs, *Combinatorica* **19** (1999) 487–505.
- [99] T. Feder, P. Hell, and J. Huang, Bi-arc graphs and the complexity of list homomorphisms, *J. Graph Theory* **42** (2003) 61–80.
- [100] T. Feder, P. Hell, and J. Huang, List homomorphisms of graphs with bounded degrees, manuscript (2004).
- [101] T. Feder, P. Hell, S. Klein, and R. Motwani, Complexity of graph partition problems, 31st Annual ACM STOC (1999) 464–472.
- [102] T. Feder, P. Hell, S. Klein, and R. Motwani, List partitions, *SIAM J. on Discrete Math.* **16** (2003) 449–478.
- [103] T. Feder, P. Hell, S. Klein, L. Tito Nogueira, and F. Protti, List matrix partitions of chordal graphs, *LATIN04*.
- [104] T. Feder, P. Hell, and B. Mohar, Acyclic homomorphisms and circular colorings of digraphs, *SIAM J. on Discrete Math.* **17** (2003) 161–169.

- [105] T. Feder and M. Y. Vardi, Monotone monadic SNP and constraint satisfaction, 25th Annual ACM STOC (1993) 612–622.
- [106] T. Feder and M. Y. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory, *SIAM J. Comput.* **28** (1998) 57–104.
- [107] U. Feige and L. Lovász, Two-prover one-round proof systems: their power and their problems, 24th ACM STOC, 1992, pp. 733–744.
- [108] W. D. Fellner, On minimal graphs, *Theoret. Comput. Sci.* **17** (1982) 103–110.
- [109] J. Fiala, and J. Kratochvíl, Complexity of partial covers of graphs, in *Algorithms and computation* (Christchurch, 2001), *Lecture Notes in Comput. Sci.* **2223** Springer, Berlin, (2001) pp. 537–549.
- [110] C. M. H. de Figueiredo, S. Klein, Y. Kohayakawa, and B. Reed, Finding skew partitions efficiently, *J. Algorithms* **37** (2000) 505–521.
- [111] P. Freyd, Concreteness, *J. Pure Appl. Alg.* **3** (1973) 171–191.
- [112] R. Frucht, Herstellung von Graphen mit vorgegebener abstrakter Gruppe, *Compos. Math.* **6** (1938) 239–250.
- [113] T. Gallai, On directed paths and circuits, in *Theory of Graphs* (Proc. Colloq. Tihany 1966), Academic Press, New York (1968), pp. 115–118.
- [114] A. Galluccio, P. Hell, and J. Nešetřil, The complexity of H -colouring of bounded degree graphs, *Discrete Math.* **222** (2000) 101–109.
- [115] G. Gao and X. Zhu, Star extremal graphs and the lexicographic product, *Discrete Math.* **152** (1996) 147–156.
- [116] M. R. Garey and D. S. Johnson, *Computers and Intractability*, W.H. Freeman and Company, San Francisco (1979).
- [117] M. R. Garey and D. S. Johnson, The complexity of near-optimal graph colouring, *J. Assoc. Comp. Mach.* **23** (1976) 43–49.
- [118] D. Geller, r -tuple colorings of uniquely colorable graphs, *Discrete Math.* **16** (1976) 9–12.
- [119] A. H. M. Gerards, Homomorphisms of graphs into odd cycles, *J. Graph Theory* **12** (1988), 73–83.
- [120] A. H. M. Gerards, An orientation theorem for graphs, Research Report CORR88-23, University of Waterloo, 1988.
- [121] L. A. Goddyn, M. Tarsi, and C. Q. Zhang, On (k, d) -colorings and fractional nowhere-zero flows, *J. Graph Theory* **28** (1998) 155–161.
- [122] C. Godsil and G. Royle, *Algebraic Graph Theory*, Springer-Verlag, New York (2001).
- [123] M. C. Golumbic, *Algorithmic Graph Theory and Perfect Graphs*, Academic Press, New York (1980).
- [124] R. L. Graham, B. L. Rothschild, J. H. Spencer, *Ramsey Theory*, Wiley (1980).
- [125] G. Grätzer, *Universal Algebra*, Princeton (1968).
- [126] D. Greenwell and L. Lovász, Applications of product colouring, *Acta Math. Acad. Sci. Hungar.* **25** (1974) 335–340.

- [127] J. R. Griggs, D. Liu, The channel assignment problem for mutually adjacent sites, *J. Combin. Theory A* **68** (1994) 169–183.
- [128] D. R. Guichard, Acyclic graph coloring and the complexity of the star chromatic number, *J. Graph Theory* **17** (1993) 129–134.
- [129] W. Gutjahr, E. Welzl and G. Woeginger, Polynomial graph-colorings, *Discrete Applied Math.* **35** (1992) 29–45.
- [130] W. Gutjahr, Graph colourings, Ph.D. Thesis, Freie Universität Berlin (1991).
- [131] B. Grünbaum, Acyclic colorings of planar graphs, *Israel J. Math.* **14** (1973) 390–412.
- [132] R. Häggkvist, P. Hell, D. J. Miller, and V. Neumann Lara, On multiplicative graphs and the product conjecture, *Combinatorica* **8** (1988) 71–81.
- [133] R. Häggkvist and P. Hell, Universality of A -mote graphs, *European J. of Combin.* **14** (1993) 23–27.
- [134] G. Hahn, P. Hell, and S. Poljak, On the ultimate independence ratio of a graph, *European J. Combin.* **16** (1995) 253–261.
- [135] G. Hahn and C. Tardif, Graph homomorphisms: structure and symmetry, in *Graph Symmetry, Algebraic Methods and Applications* (G. Hahn and G. Sabidussi eds.), NATO ASI Series C 497, Kluwer 1997, pp. 107–166.
- [136] G. Hahn and G. McGillivray, Graph homomorphisms: computational aspects and infinite graphs, manuscript (2002).
- [137] W. K. Hale, Frequency assignments: theory and applications, *Proc. IEEE* **68** (1980) 1497–1514.
- [138] F. Harary, S. Hedetniemi and G. Prins, An interpolation theorem for graphical homomorphisms, *Portugal. Math.* **26** (1967) 453–462.
- [139] M. Hasse, Zur algebraischen Begründung der Graphentheorie. I, *Math. Nachrichten* **28** (1964/5) 275–290.
- [140] S. Hazan, On triangle-free projective graphs, *Algebra Universalis* **35** (1996) 185–196.
- [141] S. T. Hedetniemi, Homomorphisms and graph automata, University of Michigan Technical Report 03105-44-T (1966).
- [142] Z. Hedrlín, On universal partly oriented sets and classes, *J. Algebra* **11** (1969) 503–509.
- [143] Z. Hedrlín, Extensions of structures and full embeddings of categories, *Actes du Congrès Internat. des Mathématiciens*, Paris (1971) tome 1, pp. 319–322.
- [144] Z. Hedrlín, P. Hell, and C. S. Ko, Homomorphism interpolation and approximation, *Annals of Discrete Math.* **15** (1982) 213–227.
- [145] Z. Hedrlín and E. Mendelsohn, The category of graphs with a given subgraph —with applications to topology and algebra, *Canad. J. Math.* **21** (1969) 1506–1517.
- [146] Z. Hedrlín and A. Pultr, Relations (graphs) with given finitely generated semigroups, *Monatsh. Math.* **68** (1964) 213–217.
- [147] Z. Hedrlín and A. Pultr, Symmetric relations (undirected graphs) with given semigroups, *Monatshefte für Math.* **69** (1965) 318–322.

- [148] Z. Hedrlín and A. Pultr, On rigid undirected graphs, *Canad. J. Math.* **21** (1966) 1237–1242.
- [149] P. Hell, Rigid undirected graphs with given number of vertices, *Comment. Math. Univ. Carolinae* **9** (1968) 51–59.
- [150] P. Hell, Rétractions de graphes, Ph.D. thesis, Université de Montréal (1972).
- [151] P. Hell, Absolute planar retracts and the four color conjecture, *J. Combin. Theory B* **17** (1974) 5–10.
- [152] P. Hell, On some strongly rigid families of graphs and the full embeddings they induce, *Algebra Universalis* **4** (1974) 108–126.
- [153] P. Hell, Absolute retracts in graphs, in *Graphs and Combinatorics* (R.A. Bari, F. Harary, eds.), *Springer-Verlag Lecture Notes in Mathematics* **406** (1974) 291–301.
- [154] P. Hell, Subdirect products of bipartite graphs, in *Infinite and Finite Sets*, (V.T. Sós et al., eds.), *Colloq. Math. Soc. J. Bolyai* **10** (1975) 857–866.
- [155] P. Hell, On some independence results in graph theory, in *Algebraic Aspects of Combinatorics* (D. Corneil and E. Mendelsohn, eds.), *Congressus Numer.* **13** (1975) 89–122.
- [156] P. Hell, Graph retractions, in *Teorie Combinatorie* (B. Segre et al., eds.), *Atti dei convegni Lincei* **17** (1976) 263–268.
- [157] P. Hell, An introduction to the category of graphs, *Ann. N.Y. Acad. Sci.* **328** (1979) 120–136.
- [158] P. Hell, Algorithmic aspects of graph homomorphisms, in *Surveys in Combinatorics 2003* (C.D. Wensley ed.), *London Math. Soc. Lecture Note Series* **307** Cambridge University Press, pp. 239–276.
- [159] P. Hell, S. Klein, L. Tito Nogueira, and F. Protti, Independent K_r 's in chordal graphs, *XI Latin-Iberian American Congress of Operations Research*, CLAIO2002.
- [160] P. Hell, S. Klein, L. Tito Nogueira, and F. Protti, Partitioning chordal graphs into independent sets and cliques, *Discrete Applied Math.* (2004).
- [161] P. Hell and D. J. Miller, Graphs with given achromatic number, *Discrete Math.* **16** (1976) 195–207.
- [162] P. Hell and D. J. Miller, Graphs with forbidden homomorphic images, *Ann. N.Y. Acad. Sci.* **319** (1979) 270–280.
- [163] P. Hell, R. Naserasr, and C. Tardif, Complexity of homomorphisms to planar graphs, manuscript (2004).
- [164] P. Hell and J. Nešetřil, Graphs and k -societies, *Canad. Math. Bull.* **13** (1970) 375–381.
- [165] P. Hell and J. Nešetřil, Groups and monoids of regular graphs (and of graphs with bounded degrees), *Canad. J. Math.* **25** (1973) 239–251.
- [166] P. Hell and J. Nešetřil, Homomorphisms of graphs and their orientations, *Monatshefte für Math.* **85** (1978) 39–48.
- [167] P. Hell and J. Nešetřil, Universality of directed graphs of a given height, *Archivum Math. (Brno)* **25** (1989) 47–54.

- [168] P. Hell and J. Nešetřil, On the complexity of H -colouring, *J. Combin. Theory B* **48** (1990) 92–110.
- [169] P. Hell and J. Nešetřil, The core of a graph, *Discrete Math.* **109** (1992) 117–126.
- [170] P. Hell and J. Nešetřil, Counting list homomorphisms and graphs with bounded degrees, in *Graphs, Morphisms and Statistical Physics* (J. Nešetřil and P. Winkler, eds.), *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* **63** (2004) 105 – 112.
- [171] P. Hell, J. Nešetřil, and X. Zhu, Duality and polynomial testing of tree homomorphisms, *Trans. Amer. Math. Soc.* **348** (1996) 147–156.
- [172] P. Hell, J. Nešetřil, and X. Zhu, Complexity of tree homomorphisms, *Discrete Applied Math.* **70** (1996) 23–36.
- [173] P. Hell and I. Rival, Absolute retracts and varieties of reflexive graphs, *Canad. J. Math.* **39** (1987) 544–567.
- [174] P. Hell, H. Zhou, and X. Zhu, Multiplicativity of oriented cycles, *J. Combin. Theory B* **60** (1994) 239–253.
- [175] P. Hell and X. Zhu, Homomorphisms to oriented paths, *Discrete Math.* **132** (1994) 107–114.
- [176] P. Hell and X. Zhu, The existence of homomorphisms to oriented cycles, *SIAM J. on Discrete Math.* **8** (1995) 208–222.
- [177] P. Hell and X. Zhu, The circular chromatic number of series-parallel graphs, *J. Graph Theory* **33** (2000) 14–24.
- [178] C. W. Henson, Countable homogeneous relational systems and categorical theories, *J. Symb. Logic* **37** (1972) 494–500.
- [179] A. J. W. Hilton, R. Rado, and S. H. Scott, A (< 5)-colour theorem for planar graphs, *Bull. London Math. Soc.* **5** (1973) 302–306.
- [180] T. Horváth and G. Turán, Learning logic programs with structures background knowledge, *Artificial Intelligence* **128** (2001) 31–97.
- [181] J. Huang and G. McGillivray, Homomorphically full reflexive graphs and digraphs, manuscript (2003).
- [182] J. Hubička and J. Nešetřil, Universal partial order represented by means of trees and other simple graphs, ITI Series Technical Report 2003-128, Charles University.
- [183] J. Hubička and J. Nešetřil, Finite paths are universal, ITI Series Technical Report 2003-129, Charles University.
- [184] F. Hughes and G. MacGillivray, The achromatic number of graphs: a survey and some new results, *Bull. Inst. Combin. Appl.* **19** (1997) 27–56.
- [185] W. Imrich and S. Klavžar, *Product Graphs, Structure and Recognition*, Wiley-Interscience, 2000.
- [186] J. R. Isbell, Two set-theoretical theorems in categories, **Fund. Math.** **53** (1963) 43–49.
- [187] J. R. Isbell, Six theorems about injective metric spaces, *Comment. Math. Helv.* **39** (1964) 65–76.
- [188] J. R. Isbell, Median algebra, *Trans. Amer. Math. Soc.* **260** (1980) 319–362.

- [189] H. Izbicki, Reguläre Graphen beliebige Grades mit vorgegebenen Eigenschaften, *Monatsh. Math.* **64** (1960) 15–21.
- [190] E. M. Jawhari, M. Pouzet, and I. Rival, A classification of reflexive graphs: the use of ‘holes’, *Canad. J. Math.* **38** (1986) 1299–1328.
- [191] E. M. Jawhari, D. Misane, and M. Pouzet, Retracts: graphs and ordered sets from the metric point of view, *Contemp. Math.* (Amer. Math. Soc.) **57** (1986) 175–226.
- [192] P. G. Jeavons, D. Cohen, and M. Gyssens, Closure properties of constraints, *Journal of the ACM* **44** (1997) 527–548.
- [193] P. G. Jeavons, On the algebraic structure of combinatorial problems, *Theoret. Comput. Sci.* **200** (1998) 185–204.
- [194] T. Jech, *Set theory*, Springer (2002).
- [195] T. R. Jensen and B. Toft, *Graph Coloring Problems*, Wiley, New York (1995).
- [196] S. Janson, T. Luczak, and A. Rucinski, *Random graphs*, Wiley, New York (2000).
- [197] B. Jónsson, Unique factorization problems for finite relational structures, *Colloq. Math.* **14** (1966) 1–32.
- [198] J. H. Kim and J. Nešetřil, On colourings of bounded degree graphs, manuscript (2002).
- [199] S. Klavžar, Coloring of graph products - a survey, *Discrete Math.*, **155** (1996), 135–145.
- [200] U. Knauer, Unretractive joins and lexicographic products of graphs, *J. Graph Theory* **11** (1987) 429–440.
- [201] M. Kneser, Aufgabe 300, *Jahresber. Deutsch. Math.-Verein* **58** (1955) 27.
- [202] P. Komárek, Some new good characterizations of directed graphs, *Časopis Pěst. Mat.* **51** (1984) 348–354.
- [203] V. Koubek, J. Nešetřil, and V. Rödl, Representing of groups and semi-groups by products in categories of relations, *Algebra Universalis* **4** (1974) 336–341.
- [204] V. Koubek and V. Rödl, On minimum order of graphs with given semi-group, *J. Combin. Theory B* **36** (1984) 135–155.
- [205] B. Korte and L. Lovász, Structural properties of greedoids, *Combinatorica* **3** (1983) 359–374.
- [206] A. Kostochka, E. Sopena, and X. Zhu, Acyclic and oriented chromatic numbers of graphs, *J. Graph Theory* **24** (1997) 331–340.
- [207] A. Kostochka, J. Nešetřil, and P. Smolíková, Colorings and homomorphisms of degenerate and bounded degree graphs, *Discrete Math.* **233** (2001) 257–276.
- [208] J. Kratochvíl and Z. Tuza, Algorithmic complexity of list colorings, *Discrete Applied Math* **50** (1994) 297–302.
- [209] J. Kratochvíl, A. Proskurowski, and J. A. Telle, Covering regular graphs, *J. Combin. Theory B* **71** (1997) 1–16.

- [210] J. Kratochvíl, A. Proskurowski, and J. A. Telle, Complexity of graph covering problems, *Nordic J. Comput.* **5** (1998) 173–195.
- [211] L. Kučera, Every category is a factorization of a concrete one, *J. Pure Appl. Alg.* **1** (1971) 373–376.
- [212] R. E. Ladner, On the structure of polynomial time reducibility, *J. Assoc. Comput. Mach.* **22** (1975) 155–171.
- [213] B. Larose and C. Tardif, Hedetniemi’s conjecture and the retracts of products of graphs, *Combinatorica* **20** (2000) 531–544.
- [214] B. Larose and C. Tardif, Strongly Rigid Graphs and Projectivity, *Multiple-Valued Logic* **7** (2001), 339–361.
- [215] B. Larose and C. Tardif, Projectivity and independent sets in powers of graphs, *J. Graph Theory* **40** (2002) 162–171.
- [216] B. Larose, F. Laviolette, and C. Tardif, On normal Cayley graphs and hom-idempotent graphs, *European J. Combin.* **19** (1998) 867–881.
- [217] F. W. Lawvere and S. H. Schnauel, *Conceptual Mathematics*, Cambridge University Press (1977).
- [218] C. G. Lekkerkerker and J. Ch. Boland, Representation of a finite graph by a set of intervals on the real line, *Fund. Math.* **51** (1962) 45–64.
- [219] S. -C. Liaw, Z. Pan, and X. Zhu, Construction of K_n -minor free graphs with given circular chromatic number, *Discrete Math.* **263** (2003) 191–206.
- [220] D. Liu, T -colorings and chromatic number of distance graphs, *Ars Combin.* **56** (2000) 65–80.
- [221] D. Liu, T -colorings of graphs, *Discrete Math.* **101** (1992) 203–212.
- [222] M. Loeb, J. Nešetřil, and B. Reed, A note on random homomorphisms to $\mathbb{Z}\mathbb{Z}$, *Discrete Math.* **273** (2003) 173–181.
- [223] C. Loten, Retractions and generalizations of chordality, Ph.D. Thesis, Simon Fraser University, 2003.
- [224] L. Lovász, Operations with structures, *Acta Math. Acad. Sci. Hungar.* **18** (1967) 321–328.
- [225] L. Lovász, On the cancellation law among finite relational structures, *Period. Math. Hungar.* **1** (1971) 145–156.
- [226] L. Lovász, Kneser’s conjecture, chromatic number, and homotopy, *J. Combin. Theory A* **25** (1978) 319–324.
- [227] L. Lovász, A note on the line reconstruction problem, *J. Combin. Theory* **13** (1972) 309–310.
- [228] L. Lovász, J. Nešetřil, and A. Pultr, On a product dimension of graphs, *J. Combin. Theory B* **29** (1980) 47–96.
- [229] T. Luczak and J. Nešetřil, A note on projective graphs, ITI Series Technical Report 2003-130, Charles University.
- [230] T. Luczak and J. Nešetřil, Towards probabilistic analysis of the dichotomy problem, KAM-DIMATIA Series 2003-640, Charles University Prague.
- [231] G. MacGillivray, On the complexity of colourings by vertex-transitive and arc-transitive digraphs, *SIAM J. on Discrete Math.* **4** (1991) 297–408.

- [232] A. K. Mackworth, Consistency in networks of relations, *Artificial Intelligence* **8** (1977) 99–118.
- [233] S. MacLane, *Categories for the Working Mathematician*, Springer-Verlag 1971.
- [234] W. Mader, Homomorphiesätze für Graphen, *Math. Annalen* **178** (1968) 154–168.
- [235] J. Matoušek, *Using the Borsuk-Ulam Theorem*, Springer-Verlag, Berlin (2003).
- [236] J. Matoušek and J. Nešetřil, *Invitation to Discrete Mathematics*, Oxford University Press (1998).
- [237] H. A. Maurer, A. Salomaa, and E. Welzl, On the complexity of the general colouring problem, *Inform. and Control* **51** (1981) 128–145.
- [238] H. A. Maurer, A. Salomaa, and D. Wood, Colorings and interpretations: a connection between graphs and grammar forms, *Discrete Appl. Math.* **3** (1981) 119–135.
- [239] H. A. Maurer, J. H. Sudborough, and E. Welzl, On the complexity of the general coloring problem, *Inform. and Control* **51** (1981) 123–145.
- [240] R. McKenzie, Cardinal multiplication of structures with a reflexive relation, *Fund. Math.* **70** (1971) 59–101.
- [241] G. F. McNulty and W. Taylor, Combinatorial interpolation theorems, *Discrete Math.* **12** (1975) 193–200.
- [242] E. Mendelsohn, On a technique for representing semigroups as endomorphism semigroups of graphs with given properties, *Semigroup Forum* **4** (1972) 283–294.
- [243] D. Micciancio, N. Segerlind, Using hypergraph homomorphisms to guess three secrets, manuscript (2004).
- [244] L. C. Middlekamp, UHF taboos - history and development, *IEEE Trans. Consumer Electron.* **CE-24** (1978) 514–519.
- [245] D. J. Miller, The categorical product of graphs, *Canad. J. Math.* **20** (1968) 1511–1521.
- [246] G. J. Minty, A theorem on n -colouring the points of a linear graph, *Amer. Math. Monthly* **69** (1962) 623–264.
- [247] B. Mitchell, *Theory of Categories*, Academic Press (1965).
- [248] U. Montanari, Networks of constraints: Fundamental properties and applications to picture processing, *Inf. Sci.* **7** (1974) 95–132.
- [249] M. Mulder, The structure of median graphs, *Discrete Math.* **24** (1978) 197–204.
- [250] V. Müller, The edge reconstruction hypothesis is true for graphs with more than $n \log_2 n$ edges, *J. Combin. Theory B* **22** (1977) 281–283.
- [251] V. Müller, On coloring of graphs without short cycles, *Discrete Math.* **26** (1979) 165–176.
- [252] R. Naserasr, Homomorphisms and bounds, Ph.D. Thesis, Simon Fraser University, (2003).

- [253] R. Naserasr, Y. Nigussie, On the new reformulation of Hadwiger's conjecture, KAM-DIMATIA Series Technical Report 2004-667, Charles University.
- [254] C. St. J. A. Nash-Williams, Decomposition of finite graphs into forests, *J. London Math. Soc.* **39** (1964).
- [255] L. Nebeský, Median graphs, *Comment. Math. Univ. Carolinae* **12** (1971) 317–325.
- [256] J. Nešetřil, Amalgamation of Graphs and its Applications, *Ann. N.Y. Acad. Sci.* **319** (1979) 415–428.
- [257] J. Nešetřil, Structure of Graph Homomorphisms I, in *Graph Theory Notes of New York*, New York Acad. Sci. (1998), pp. 7–12.
- [258] J. Nešetřil, Structure of Graph Homomorphisms, *Combinatorics, Probability and Computing* **8** (1999) 177–184.
- [259] J. Nešetřil, Aspects of Structural Combinatorics, *Taiwanese J. Math.* **3**, 4 (1999), 381–424.
- [260] J. Nešetřil, A rigid graph for every set, *J. Graph Theory* **39** (2002) 108–110.
- [261] J. Nešetřil, Bounds and extrema for classes of graphs and finite structures, to appear in *More Sets, Graphs and Numbers*, *Bolyai Society Mathematical Studies*, (E. Gyori, G.O.H. Katona and L. Lovasz, eds), Springer (2004).
- [262] J. Nešetřil and P. Ossona de Mendez, Colouring and homomorphisms of minor-closed classes, *Discrete and Computational Geometry* **25** (2003) 651–664.
- [263] J. Nešetřil and P. Ossona de Mendez, Folding, manuscript (2003).
- [264] J. Nešetřil and P. Ossona de Mendez, Cuts and bounds, *Discrete Math.* (2004).
- [265] J. Nešetřil and P. Ossona de Mendez, Tree depth, subgraph colouring and homomorphism bounds, KAM-DIMATIA Series 2004-656, Charles University Prague.
- [266] J. Nešetřil and A. Pultr, Dushnik-Miller type dimension of graphs and its complexity, *Springer Verlag Lecture Notes in Computer Science* **56** (1977) 482–494.
- [267] J. Nešetřil and A. Pultr, On classes of relations and graphs determined by subobjects and factorobjects, *Discrete Math.* **22** (1978) 287–300.
- [268] J. Nešetřil and A. Pultr, A note on homomorphism independent families, *Discrete Math.* **235** (2001) 327–334.
- [269] J. Nešetřil, A. Pultr, and C. Tardif, Gaps and dualities in Heyting categories, ITI Series Technical Report 2003-173, Charles University.
- [270] J. Nešetřil and V. Rödl, Type theory of partition properties of graphs, in *Recent advances in graph theory* (M. Fiedler, ed.), Academia Prague 1975, pp. 405–412.
- [271] J. Nešetřil and V. Rödl, A simple proof of the Galvin-Ramsey property of graphs and a dimension of a graph, *Discrete Math.* **23** (1978) 49–56.
- [272] J. Nešetřil and V. Rödl, Chromatically optimal rigid graphs, *J. Combin. Theory B* **46** (1989) 133–141.

- [273] J. Nešetřil and V. Rödl, Combinatorial Partitions of Finite Posets and Lattices - Ramsey Lattices, *Algebra Universalis* **19** (1994) 106–119.
- [274] J. Nešetřil and V. Rödl, Partition theory and its applications, in *Surveys in Combinatorics*, Cambridge University Press 1979, pp. 96–156.
- [275] J. Nešetřil and G. Sabidussi, Minimal asymmetric graphs of induced length four, *Graphs and Combin.* **8** (1992) 343–359.
- [276] J. Nešetřil and S. Shelah, On the order of countable graphs, *European J. Combin.* **24** (2003) 649–663.
- [277] J. Nešetřil and C. Tardif, Duality theorems for finite structures (characterising gaps and good characterizations), *J. Combin. Theory B* **80** (2000) 80–97.
- [278] J. Nešetřil and C. Tardif, Density via duality, *Theoret. Comput. Sci.* **287** (2002) 585–591.
- [279] J. Nešetřil and C. Tardif, A dualistic approach to bounding the chromatic number of a graph, ITI Series Technical Report 2001-036, Charles University.
- [280] J. Nešetřil and X. Zhu, On bounded treewidth duality of graphs, *J. Graph Theory* **23** (1996) 151–162.
- [281] J. Nešetřil and X. Zhu, Paths Homomorphisms, *Proc. Cambridge Phil. Soc.* **120** (1996) 207–220.
- [282] J. Nešetřil and X. Zhu, Construction of Sparse Graphs with Prescribed Circular Colorings, *Discrete Math.* **235** (2001) 277–291.
- [283] J. Nešetřil and X. Zhu, On sparse graphs with given colorings and homomorphisms, *J. Combin. Theory B* **90** (2004) 161–172.
- [284] R. Nowakowski and I. Rival, On a class of isometric subgraphs of a graph, *Combinatorica* **2** (1982) 79–90.
- [285] R. Nowakowski and I. Rival, The smallest graph variety containing all paths, *Discrete Math.* **43** (1983) 223–234.
- [286] R. Nowakowski and P. Winkler, Vertex-to-vertex pursuit in a graph, *Discrete Math.* **43** (1983) 235–239.
- [287] R. Nowakowski and I. Rival, Fixed-edge theorem for graphs with loops, *J. Graph Theory* **3** (1979) 339–350.
- [288] Z. Pan and X. Zhu, Density of the circular chromatic numbers of series-parallel graphs, *J. Graph Theory* **46** (2004) 57–68.
- [289] E. Pesch, *Retracts of Graphs*, Athenaeum Verlag, Frankfurt (1988).
- [290] N. Pippenger, *Theories of Computability*, Cambridge University Press (1997).
- [291] S. Poljak, Coloring digraphs by iterated antichains, *Comment. Math. Univ. Carolinae* **32** (1992) 209–212.
- [292] S. Poljak and V. Rödl, On the arc-chromatic number of a digraph, *J. Combin. Theory B* **31** (1981) 190–198.
- [293] A. Proskurowski and S. Arnborg, Linear time algorithms for NP-hard problems restricted to partial k-trees, *Discrete Appl. Math.* **23** (1989) 11–24.
- [294] A. Pultr, Tensor products in the category of graphs, *Comment. Math. Univ. Carolinae* **11** (1970) 619–639.

- [295] A. Pultr and V. Trnková, *Combinatorial, Algebraic and Topological Representations of Groups, Semigroups and Categories*, North-Holland, Amsterdam (1980).
- [296] A. Quilliot, On the Helly property working as a compactness criterion on graphs, *J. Combin. Theory A* **40** (1985) 186–193.
- [297] A. Quilliot, A retraction problem in graph theory, *Discrete Math.* **54** (1985) 61–71.
- [298] A. Quilliot, Thèse 3. cycle, Université de Paris VI (1978).
- [299] J. H. Rabinowitz and V. Krňanová Proulx, An asymptotic approach to the channel assignment problem, *SIAM J. Alg. Disc. Methods* **6** (1985) 507–518.
- [300] A. Raspaud and E. Sopena, Good and semi-strong colorings of oriented planar graphs, *Information Processing Letters* **51** (1994) 171–174.
- [301] F. Roberts, T -colorings of graphs: recent results and open problems, *Discrete Math.* **93** (1991) 229–245.
- [302] D. J. Rose, R. E. Tarjan and G. S. Lueker, Algorithmic aspects of vertex elimination on graphs, *SIAM J. Comput.* **5** (1976) 266–283.
- [303] I. G. Rosenberg, Strongly rigid relations, *Rocky Mountain Journal of Math.* **3** (1973) 631–639.
- [304] B. Roy, Nombre chromatique et plus longs chemins d’un graph, *Rev. Française Informat. Recherche Operationelle* **1** (1967) 129–132.
- [305] G. Sabidussi, Graphs with given group and given graph-theoretical properties, *Canad. J. Math.* **9** (1957) 515–525.
- [306] G. Sabidussi, On a class of fixed-point-free graphs, *Proc. Amer. Math. Soc.* **9** (1958) 800–804.
- [307] G. Sabidussi, Graph multiplication, *Math. Z.* **72** (1960) 446–457.
- [308] G. Sabidussi, Graph derivatives, *Math. Z.* **76** (1961) 385–401.
- [309] G. Sabidussi, Vertex-transitive graphs, *Monatsh. Math.* **68** (1964) 426–438.
- [310] G. Sabidussi, Subdirect representations of graphs, in *Infinite and Finite Sets*, (V.T. Sós et al., eds.), *Colloq. Math. Soc. J. Bolyai* **10** (1975) 1199–1226.
- [311] N. Sauer and X. Zhu, An approach to Hedetniemi’s conjecture, *J. Graph Theory* **16** (1992) 423–436.
- [312] N. Sauer, Hedetniemi’s conjecture—a survey, *Discrete Math.* **229** (2001) 261–292.
- [313] T. J. Schaeffer, The complexity of satisfiability problems, 10th Annual ACM STOC (1978) 216–226.
- [314] E. R. Scheinerman, D. H. Ullman, *Fractional Graph Theory*, Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, New York, 1997.
- [315] M. Sholander, Trees, lattices, order and betweenness, *Proc. Amer. Math. Soc.* **3** (1952) 369–382.
- [316] E. Sopena, On the chromatic number of oriented graphs, *J. Graph Theory* **25** (1997) 191–205.

- [317] E. Sopena, There exist oriented planar graphs with oriented chromatic number at least sixteen, *Information Processing Letters* **81** (2002) 309–312.
- [318] J. P. Spinrad, *Efficient graph representations*, Fields Institute Monographs **19** Amer. Math. Soc. (2000).
- [319] S. Stahl, n -tuple colorings and associated graphs, *J. Combin. Theory B* **20** (1976) 185–203.
- [320] S. Stahl, The multichromatic numbers of some Kneser graphs, *Discrete Math.* **185** (1998) 287–291.
- [321] C. Tardif, A fixed box theorem for the Cartesian product of graphs and metric spaces, *Discrete Math.* **171** (1997) 237–248.
- [322] C. Tardif, Fractional multiples of graphs and the density of vertex-transitive graphs. *J. Algebraic Combin.* **10** (1999) 61–68.
- [323] C. Tardif, The chromatic number of the product of two graphs is at least half the minimum of the fractional chromatic numbers of the factors, *Comment. Math. Univ. Carolinae* **42** (2001) 353–355.
- [324] C. Tardif, Multiplicative graphs and semi-lattice endomorphisms in the category of graphs, manuscript (2004).
- [325] R. E. Tarjan, Decomposition by clique separators, *Discrete Math.* **55** (1985) 221–232.
- [326] A. Tarski, On the calculus of relations, *J. Symbolic Logic* **6** (1941) 73–89.
- [327] W. T. Trotter, *Dimension Theory*, John's Hopkins University Press (1992).
- [328] D. Turzík, A note on chromatic number of direct product of graphs, *Comment. Math. Univer. Carolinae* **24** (1983) 461–463.
- [329] K. Vesztegombi, Chromatic number of strong products of graphs, in *Algebraic Methods in Graph Theory* (L. Lovász, V. T. Sós, eds.), Colloq. Math. Soc. János Bolyai **25** (1981) 819–825.
- [330] N. Vikas, Computational complexity of compaction to reflexive cycles. *SIAM J. Comput.* **32** (2002/03) 253–280
- [331] J. Vinárek, A new proof of the Freyd's theorem, *J. Pure and Appl. Algebra* **8** (1976) 1–4.
- [332] A. Vince, Star chromatic number, *J. Graph Theory* **12** (1988) 551–559.
- [333] L. M. Vitaver, Determination of minimal colouring of vertices of a graph by means of Boolean powers of incidence matrix, *Doklady Akad. Nauk SSSR* **147** (1962) 758–759.
- [334] P. Vopěnka, A. Pultr, and Z. Hedrlín, A rigid relation exists on any set, *Comment. Math. Univ. Carolinae* **6** (1965) 149–155.
- [335] W. E. Watkins, Vertex-transitive graphs that are not Cayley graphs, in *Cycles and rays* (G. Hahn et al. eds.) NATO ASI Ser. C, Kluwer Academic Publishers, Dordrecht 1990, pp. 243–256.
- [336] P. M. Weichsel, The Kronecker product of graphs, *Proc. Amer. Math. Soc.* **13** (1962) 47–52.
- [337] E. Welzl, Symmetric graphs and interpretations, *J. Combin. Theory B* **37** (1984) 235–244.

- [338] E. Welzl, Color families are dense, *Theoret. Comput. Sci.* **17** (1982) 29–41.
- [339] D. B. West, *Introduction to graph theory*, Prentice-Hall, N.J. (1996).
- [340] S. Whitesides, A method for solving certain graph recognition and optimization problems, with applications to perfect graphs, *Annals of Discrete Math.* **21** (1984) 281–297.
- [341] H. Zhou, On the non-multiplicativity of oriented cycles, *SIAM J. on Discrete Math.* **5** (1992) 207–218.
- [342] H. Zhou and X. Zhu, Multiplicativity of acyclic local tournaments, *Combinatorica* **17** (1997) 135–145.
- [343] X. Zhu, On the chromatic number of the products of hypergraphs, *Ars Combinatoria* **34** (1992) 25–31.
- [344] X. Zhu, A simple proof of the multiplicativity of directed cycles of prime power length, *Discrete Applied Math.* **36** (1992) 313–315.
- [345] X. Zhu, Star chromatic numbers and products of graphs, *J. Graph Theory* **16** (1992) 557–569.
- [346] X. Zhu, Uniquely H -colourable graphs with large girth, *J. Graph Theory* **23** (1996) 33–41.
- [347] X. Zhu, Construction of uniquely H -colorable graphs, *J. Graph Theory* **30** (1999) 1–6.
- [348] X. Zhu, On the bounds for the ultimate independence ratio of a graph, *Discrete Math.* **156** (1996) 229–236.
- [349] X. Zhu, Uniquely H -colorable graphs with large girth, *J. Graph Theory* **23** (1996) 33–41.
- [350] X. Zhu, A survey on Hedetniemi’s conjecture, *Taiwanese J. Math.* **2** (1998) 1–24.
- [351] X. Zhu, Circular chromatic number and graph minors, *Taiwanese J. Math.* **4** (2000) 643–660.
- [352] X. Zhu, Circular chromatic number: a survey, *Discrete Math.* **229** (2001) 371–410.

INDEX

- $\alpha(G)$ 14
- $\chi(G)$ 7
- $\chi_a(G)$ 214
- $\chi_c(G)$ 193
- $\chi_f(G)$ 201
- $\chi_i(G)$ 208
- $\chi_o(G)$ 212
- $\prod_{i \in I} G_i$ 39
- absolute retract 64–8, 70–1, 79
- achromatic number 10, 33, 36, 190
- adjacent vertices
 - in a graph 2
 - in a digraph 5
- Albertson, M. 33
- antichain 82–3, 139
- maximal 85–6, 101–2, 106, 107
- arc 1
 - backward 6
 - forward 6
- artificial intelligence 77, 186
- associated partition 7–11, 69, 90, 207
- asteroidal triple 171, 173
- $\text{AUT}(G)$ 21
- automorphism 16, 21–2, 139, 140
 - group 24, 26, 132, 140–1
- Babai, L. 139
- Bandelt, H. -J. 77
- Bang-Jensen, J. 186
- $\text{BIJ}(G, H)$ 5
- $\text{bij}(G, H)$ 5
- Borsuk, K. 77
- bound, 102–4
- branch vertex 117, 119, 145
- Brandstädt, A. 187
- Bulatov, A. A. 186, 187
- C_k 4
- \vec{C}_k 6
- \mathcal{C} 81
- C_S 82
- cancellation 44, 76
- category 122–33, 135, 138–40
 - concrete 128, 131–5, 139
 - thin, 126–7
- chain 82–3
- channel assignment problem 31, 210
- check
 - consistency 162–5, 186, 189
 - failure 162, 167
 - pair consistency 167–9
 - success 162, 167
- Chernoff inequality 95
- chooser 174–5
- chromatic index 42
- chromatic number 7, 10, 49, 56, 199, 202
 - k -tuple 201
 - acyclic 214–16
 - circular 193–8, 210, 219
 - fractional 201–4, 211–12, 219
 - interlaced 208
 - of a digraph 48
 - of a graph with loops allowed 60
 - oriented 212–16
 - of a graph 213
 - star 193, 218; *see also* chromatic number, circular
- chromatic polynomial 18, 33
- circular distance 193
- clique 3, 183–5
- code 41
- codomain 21, 125
- cograph 187
- Collins, K. 33
- colouring 6–7, 13, 90, 192
 - (2, 1)- 219
 - (p/q)- 193–9
 - T - 210, 219
 - k -tuple 200
 - acyclic 214–17, 219
 - circular 192–193
 - complete 9
 - oriented 212–15, 218
- composition 20–2, 34, 109, 122, 125, 141
- concatenation 84
- concurrency 199–200
- configuration 32
- conflict 169, 191
- conjecture 106, 139, 153, 158, 160, 186, 187, 207, 218
 - Dichotomy 30, 34, 153, 157, 160
 - Edge Reconstruction 16, 17, 34, 43, 44
 - Hadwiger’s 104–6
 - Kneser’s 218
 - Product 50–2, 55, 56, 59, 76
- connectivity 109
- connector

- edge 120
- vertex 117
- consistent mapping, 168
- constraint satisfaction problem 28–30, 34, 142, 153–7, 186, 191
- continuous relaxation 203
- contraction 10, 191
- cops and robbers 20, 34, 64, 68, 77
- core 19, 34–6, 95–6, 107, 109, 142, 188, 189–90
- Courcelle, B. 219
- cover 190
- CSP 28
- Čulík, K. 33
- cupoint 113
- cutset
 - clique 180–2
 - skew 180
- cycle 4
- datalog 186
- Dechter, R. 186
- deck 17, 43
- defect 45
- degree 2
- Diaz, J. 187
- dichotomy 30, 143, 144, 153, 158–60, 186
- Digraph* 126
- digraph 1
 - acyclic 10, 13, 109, 189, 198
 - asymmetric 109, 111
 - balanced 10–11, 102, 119
 - composite 39
 - connected 6
 - corresponding symmetric 1
 - exponential 46, 76
 - infinite 3
 - irreflexive 3
 - line 48
 - multiplicative 51–4, 76
 - persistent 51, 76
 - prime 39
 - projective 57, 77, 79, 107
 - reflexive 1
 - rigid 27, 109–11
 - strongly connected 6
 - symmetric 1, 55
- $\dim G$ 41
- dimension 41, 76
- directed
 - cycle 6
 - path 6
 - walk 6
- disjoint union 58, 78, 93
- distance 4
- Dolnikov, V. L. 218
- domain 21–2, 125
- dual 98–100, 107–8
- duality, 13–15, 35, 54 106, 143, 161, 186
 - bounded treewidth, 161–2, 165, 168–9
 - finite 101–2
 - pair 98–100, 108
 - tree 161, 162, 164–5
- dynamic location problem 73, 77
- edge 1
- edge-based replacement graph 120
- El-Zahar, M. 76
- encoding 41
- $\text{END}(G)$ 23
- endomorphism 23, 27, 71
 - monoid 23–7, 109, 132, 141
- Erdős, P. 95, 106, 139
- Euler
 - circuit 5
 - trail 5
- $f(G)$ 9
- family of incomparable
 - digraphs 138
 - graphs 94, 119
- Feder, T. 30, 34, 153, 186
- fixed
 - clique 71
 - edge 72
- fold 68
- $\text{Forb}\mathcal{F}$ 93
- $\text{Forb}_d\mathcal{F}$ 103
- fork 133
- forward clique 182
- fractional clique number 204
- frequency assignment problem, *see* channel assignment problem
- Freyd, P. 139
- Frucht, R. 24, 34
- $G \leq H$ 81
- $G * J$ 119
- G/θ 8
- G_k 113
- $[G, H]$ 97
- $\mathbf{G}(n, p)$ 94, 114
- $\mathbf{G}(s, n, p)$ 87
- Gallai, T. 33
- gap 97–9, 101
- Geller, D. 218
- Gilbert, E. N. 219
- girth 7, 47, 86, 88, 90, 91
 - odd, 7, 34, 47, 49, 50, 92
- Goddyn, L. 218

- Graph* 125
- graph 1
 - (a, b) - 183
 - d -degenerate 189
 - asymmetric 111, 141
 - bi-arc 177
 - bipartite 3, 13, 78, 79, 135–6, 143
 - Cayley 36, 78, 108
 - chordal 171, 175, 191
 - circulant 193
 - circular-arc 189
 - clique-Helly 71
 - complete 6
 - cubic 14–15, 103, 107
 - de Bruijn 47
 - dismantlable 68–71, 77
 - extra-ordinary 144–5, 151
 - G -pointed 89, 92, 107
 - infinite 34, 115, 140
 - interval 32, 171, 175, 187, 188
 - Kneser 201, 205–7, 218
 - median 72, 75, 79
 - multiplicative 52, 59
 - nearly bipartite 61–2, 64, 79, 80, 199
 - ordinary 144
 - oriented 2, 212
 - outerplanar 215
 - penny 146
 - perfect 35
 - Petersen 79, 202
 - planar 3, 103, 107, 108, 135, 189
 - point determining 36
 - projective 57, 79, 91, 107, 141, 189
 - random 87, 94, 107
 - rational complete 192–3, 201, 207
 - reflexive 1
 - regular 3
 - rigid 111–21, 139, 141; *see also* digraph, rigid
 - series-parallel 140, 219, 221
 - shift 47, 76
 - sparse 86, 87, 89, 90
 - split 180, 183, 187
 - stiff 69, 70
 - threshold 80
 - triangle-connected 113, 121
 - triangle-free 15, 42, 79, 103, 107, 108
 - underlying 2
 - uniquely H -colourable 90, 107
 - uniquely colourable 49, 50, 90
 - unit distance 146, 220
 - vertex-critical 35, 112
 - vertex-transitive 14, 35, 36, 106, 186
 - with loops allowed 1
- greedoid 69
- group 22, 23, 26, 126
- Grünbaum, B. 219
- Gutin, G. 186
- Gutjahr, W. 186, 187
- H -COL 142
- H -CON 190
- H -RET 153, 177–8, 187, 191
- H -colouring 3
 - acyclic 190
 - equitable 190
 - list 170, 177, 181
 - problem 142–4
- $H(k, \ell)$ 140
- H_k 112
- Hahn, G. 186
- hardcore gas model 34
- Hasse, M. 33
- Hedrlín, Z. 33, 34, 138, 139
- height 11
- Helly property 71, 77, 79
- hole 65–7, 71, 77
- $\text{HOM}(G, H)$ 40, 69
- $\text{hom}(G, H)$ 16, 39, 43
- homogeneous set 180
- homomorphic
 - equivalence 18, 82, 142, 194, 196, 221
 - graph 3
 - image 9, 36, 141
- homomorphism
 - acyclic 190
 - bijective 5
 - category 125
 - edge-bijective 4
 - edge-injective 4
 - edge-surjective 4
 - full 35
 - injective 4, 8, 35, 95, 139
 - list 31, 170
 - locally bijective 189
 - locally injective 187, 219
 - locally surjective 187
 - of digraphs 3
 - of general relational systems 28
 - of graphs 3
 - of trigraphs 179, 189
 - order 34, 106, 132
 - problem 30, 142, 144
 - surjective 4, 8, 9, 89, 92, 139
 - vertex-bijective 4
 - vertex-injective 4
 - vertex-surjective 4
- hypercube 73
- hypergraph 141
 - intersecting 141

- $I(n, k)$ 208
- $IRel$ 126
- $i(G)$ 14
- ideal class 93–4, 107
- idempotent 57
- identity
 - in a category 123–5
 - in a monoid 23, 109, 123
- imbalance 221
- Imrich, W. 76
- inclusion 16, 35
- incomparable
 - family 110–11, 117, 121, 139–40
 - graphs 85, 86, 93, 94, 96
- indegree 2
- independence
 - number 14, 204
 - ratio 14, 81, 195, 206, 220
- independent
 - cliques 185
 - set 7, 35
- indicator 145
 - construction 145–6
- induced subcategory 127
- $INJ(G, H)$ 5
- $inj(G, H)$ 5, 16, 43
- inneighbour 2
- inner vertex 117, 119, 145
- integer linear program 202
- interlaced k -tuples 207
- Isbell's condition 133–5, 139
- isomorphism 16, 43, 78, 80, 130, 132
 - of categories 127
 - of digraphs 16
 - of monoids 23
- Jeavons, P. G. 186
- join 180
- $K(n, k)$ 201
- K_k 6
- $K_{p/q}$ 193
- k -equivalent graphs 88–9, 92
- k -tree 161
- k -tuple consistency 166
- Kirsch, Y. 219
- Klavzar, S. 76
- Kneser, M. 218
- Kratochvíl, J. 187
- Kučera, L. 138
- lattice 82
- Lemma
 - No-Homomorphism 14, 33, 195, 206, 220
 - Sparse Incomparability 86, 97, 106
- length 4
- level 11
- linear program 202–3
- list 31
- list clique partition problem 182
- loop 1
- Lovász vector 43, 76
- Lovász, L. 34, 218
- $\mathcal{M}(A, B)$ 125
- M -partition 179, 183
- majority function 67, 162, 167–9
 - conservative 172, 175, 188
- matching 42
- maximum 94, 102, 104, 105, 107
- McGillivray, G. 186
- McKenzie, R. 76
- median 72
- Mendelsohn, E. 139
- metric space 64, 77
- Miller, D. J. 76
- minimal chordal extension 182
- minor 10, 104, 105, 220
- Minty, G. J. 218
- modular decomposition 180
- monoid 23–6, 122, 126, 138
- monotonicity 33, 81, 194
- morphism 125
- Moser spindle 146
- Müller, V. 76, 106
- multicolouring 200
- N 29, 160, 189
- near-unanimity function 189, 191
- neighbour 2
- net length 6, 11–13, 53, 111
- nonexpansive mapping 64
- Nowakowski, R. 34
- object 125
- orientation 1, 212
- oriented
 - cycle 77, 187
 - path 77, 106, 111
 - tree 97, 100, 107, 158, 213–14
 - walk 6
- outdegree 2
- outneighbour 2
- $\mathcal{P}(H)$ 165
- P_k 4
- \tilde{P}_k 6
- pair consistency 166
- parametrized complexity 187
- partial order 81–4

- dense 85, 96, 106
- path 4
- pattern 28, 153, 160
- perfect elimination ordering 171, 175
- Perles, M. 106, 219
- Pesch, E. 77
- phase transition 68
- Pippenger, N. 186
- polymorphism 57, 67, 107, 159–60, 191
- polynomially equivalent 154
- power 46, 57
- Prague School of Category Theory 34, 138
- prime factorization 39, 44, 76
- Prisner, E. 77
- product 37, 40, 56, 59, 76, 93
 - cartesian 74, 76, 220
 - infinite 78
 - lexicographic 79, 220
 - strong 79
- projection 37, 58
- Pultr, A. 33, 34, 138, 139

- $Q(a, b)$ 111
- quadratic residue tournament 215
- quasiorder 81
- quasi-polynomial 187
- Quilliot, A. 34
- quotient 8

- $R(n, \ell)$ 47
- $\vec{R}(n, \ell)$ 48
- Ramsey number 95
- reflexive tree 68, 72
- relational system
 - binary 24, 26, 27, 126, 129
 - general 28, 39, 142
- replacement
 - digraph 25
 - graph 117
 - operation 25, 117, 139
- representation 139
 - of a category 127–32
 - of a group 24, 109
 - of a monoid 25, 109
 - of a partial order 82, 84
- retract 19, 58–60, 73–4, 77, 79, 199
- retraction 18, 34, 36, 58, 139
 - problem 153–6
- reversal 111
- Rosenberg, I. 139
- Roy, B. 33

- Sabidussi, G. 33, 34, 76, 77, 138
- Sauer, N. 76
- Schmerl, J. H. 77
- semidefinite programming 186
- semigroup 141
- Serna, M. 187
- Set* 126
- $sp_T(G)$ 210
- span 210
- Stahl, S. 218
- standard isometric representation 66
- statistical physics 32, 34
- stretching 154
- strict matching 42
- strong replacement graph 118–19
- sub-indicator construction 146–8
- subcategory 127
- subdirect representation 77
- subdivision 62
 - odd 62
- subgraph 2
 - induced 2
 - isometric 60, 61, 64–6, 78
- supergraph 2
- supremum 103
- $SUR(G, H)$ 5
- $\text{sur}(G, H)$ 5
- surface 138
- symmetric
 - arcs 2
 - part 2

- T -CSP 29
- T -RET 153
- \vec{T}_k 10
- Tardif, C. 106
- Tarjan, E. 187
- Tarsi, M. 218
- template 29, 142, 153
- theorem
 - Borsuk's Antipodal 206
 - Brooks' 42
 - Colouring Interpolation 9, 33
 - Cook's 151
 - Erdős–Ko–Rado 206
 - Four Colour 14, 103, 104, 108
 - Grötzsch's 14, 103
 - König's 13, 51
 - Nash-Williams' 217
 - predicted 165, 175, 188
 - Sperner's 96
 - Strong Perfect Graph 180
 - Tutte's 77
 - Vizing's 41
- Thilikos, D. M. 187
- topological sort 10
- tournament 189
 - transitive 10, 52

- transversal system 154
- trigraph 178, 187
 - adjacency 178
- Trnková, V. 138
- Tuza, Z. 187
- two-satisfiability algorithm 181

- Ulam, S. 138

- Vardi, M. 30, 34, 153, 186
- variety 66–7, 77, 79, 191
- vertex 1
- Vesztergombi, K. 76
- Vinárek, P. 139
- Vince, A. 218
- Vitaver, L. 33

- walk 4
- Welzl, E. 186, 187
- Whitesides, S. 187
- Widom–Rowlinson gas model 32
- window index 74–6, 79
- Winkler, P. 34
- Woeginger, G. 186, 187

- X -underbar
 - enumeration 164–5, 188
 - property 164, 186

- $Z(k)$ 84
- Z_k 83–4
- zero-sum game 221
- Zhang, C. Q. 218
- Zhu, X. 77, 218
- zig-zag 83, 97, 110