

Note to readers:
Please ignore these
sidenotes; they're just
hints to myself for
preparing the index,
and they're often flaky!

KNUTH

THE ART OF COMPUTER PROGRAMMING

VOLUME 4 PRE-FASCICLE 5C

DANCING LINKS

DONALD E. KNUTH *Stanford University*

ADDISON-WESLEY



January 29, 2017

Internet
Stanford GraphBase
MMIX

Internet page <http://www-cs-faculty.stanford.edu/~knuth/taocp.html> contains current information about this book and related books.

See also <http://www-cs-faculty.stanford.edu/~knuth/sgb.html> for information about *The Stanford GraphBase*, including downloadable software for dealing with the graphs used in many of the examples in Chapter 7.

See also <http://www-cs-faculty.stanford.edu/~knuth/mmixware.html> for downloadable software to simulate the MMIX computer.

Copyright © 2016 by Addison–Wesley

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, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher, except that the official electronic file may be used to print single copies for personal (not commercial) use.

Zeroth printing (revision -80), 29 January 2017

January 29, 2017

PREFACE

*With this issue we have terminated the section "Short Notes."
... It has never been "crystal clear" why a Contribution cannot be short,
just as it has occasionally been verified in these pages
that a Short Note might be long.*

— ROBERT A. SHORT, *IEEE Transactions on Computers* (1973)

THIS BOOKLET contains draft material that I'm circulating to experts in the field, in hopes that they can help remove its most egregious errors before too many other people see it. I am also, however, posting it on the Internet for courageous and/or random readers who don't mind the risk of reading a few pages that have not yet reached a very mature state. *Beware:* This material has not yet been proofread as thoroughly as the manuscripts of Volumes 1, 2, 3, and 4A were at the time of their first printings. And alas, those carefully-checked volumes were subsequently found to contain thousands of mistakes.

Given this caveat, I hope that my errors this time will not be so numerous and/or obtrusive that you will be discouraged from reading the material carefully. I did try to make the text both interesting and authoritative, as far as it goes. But the field is vast; I cannot hope to have surrounded it enough to corral it completely. So I beg you to let me know about any deficiencies that you discover.

To put the material in context, this portion of fascicle 5 previews Section 7.2.2.1 of *The Art of Computer Programming*, entitled "Dancing links." It develops an important data structure technique that is suitable for *backtrack programming*, which is the main focus of Section 7.2.2. Several subsections (7.2.2.2, 7.2.2.3, etc.) will follow.

* * *

The explosion of research in combinatorial algorithms since the 1970s has meant that I cannot hope to be aware of all the important ideas in this field. I've tried my best to get the story right, yet I fear that in many respects I'm woefully ignorant. So I beg expert readers to steer me in appropriate directions.

Please look, for example, at the exercises that I've classed as research problems (rated with difficulty level 46 or higher), namely exercises 182, ...; I've also implicitly mentioned or posed additional unsolved questions in the answers to exercises 82, 210, Are those problems still open? Please inform me if you know of a solution to any of these intriguing questions. And of course if no

solution is known today but you do make progress on any of them in the future, I hope you'll let me know.

I urgently need your help also with respect to some exercises that I made up as I was preparing this material. I certainly don't like to receive credit for things that have already been published by others, and most of these results are quite natural "fruits" that were just waiting to be "plucked." Therefore please tell me if you know who deserves to be credited, with respect to the ideas found in exercises 5, 6, 20, 21, 31, 40, 70, 158, 163, 177, 198(d), 206, 207, 208, 210, 218, 222, Furthermore I've credited exercises . . . to unpublished work of Have any of those results ever appeared in print, to your knowledge?

Jellis
Huang
Sicherman
FGbook
Knuth

* * *

Special thanks are due to George Jellis for answering dozens of historical queries, as well as to Wei-Hwa Huang, George Sicherman, and . . . for their detailed comments on my early attempts at exposition. And I want to thank numerous other correspondents who have contributed crucial corrections.

* * *

I happily offer a "finder's fee" of \$2.56 for each error in this draft when it is first reported to me, whether that error be typographical, technical, or historical. The same reward holds for items that I forgot to put in the index. And valuable suggestions for improvements to the text are worth 32¢ each. (Furthermore, if you find a better solution to an exercise, I'll actually do my best to give you immortal glory, by publishing your name in the eventual book:—)

In the preface to Volume 4B I plan to introduce the abbreviation *FGbook* for my book *Selected Papers on Fun and Games* (Stanford: CSLI Publications, 2011), because I will be making frequent reference to it in connection with recreational problems.

Cross references to yet-unwritten material sometimes appear as '00'; this impossible value is a placeholder for the actual numbers to be supplied later.

Happy reading!

Stanford, California
99 Umbruary 2016

D. E. K.

*What a dance
do they do
Lordy, how I'm tellin' you!*

— HARRY BARRIS, *Mississippi Mud* (1927).

BARRIS
FIELDS
undoing
exact covering-
0s and 1s

*Don't lose your confidence if you slip,
Be grateful for a pleasant trip,
And pick yourself up, dust yourself off, start all over again.*

— DOROTHY FIELDS, *Pick Yourself Up* (1936)

7.2.2.1. Dancing links. One of the chief characteristics of backtrack algorithms is the fact that they usually need to *undo* everything that they *do* to their data structures. Blah blah de blah blah blah.

* * *

Exact cover problems. We will be seeing many examples where links dance happily and efficiently, as we study more and more examples of backtracking. The beauty of the idea can perhaps be seen most naturally in an important class of problems known as *exact covering*: We're given an $m \times n$ matrix A of 0s and 1s, and the problem is to find a subset of rows whose sum is exactly 1 in every column. For example, consider the 6×7 matrix

$$A = \begin{pmatrix} 0 & 0 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 & 1 \end{pmatrix}. \quad (20)$$

Each row of A corresponds to a subset of a 7-element universe. A moment's thought shows that there's only one way to cover all seven of these columns with disjoint rows, namely by choosing rows 1, 4, and 5. We want to teach a computer how to solve such problems, when there are many, many rows and many columns.

DUDENEY
 CLARKE
 GOLOMB
 Golomb
 Conway

*If mounted on cardboard, [these pieces]
 will form a source of perpetual amusement in the home.*

— HENRY E. DUDENEY, *The Canterbury Puzzles* (1907)


*Very gently, he replaced the titanite cross
 in its setting between the F, N, U, and V pentominoes.*

— ARTHUR C. CLARKE, *Imperial Earth* (1976)


*Which English nouns ending in -o pluralize with -s and which with -es?
 If the word is still felt as somewhat alien, it takes -s,
 while if it has been fully naturalized into English, it takes -es.
 Thus, echoes, potatoes, tomatoes, dingoes, embargoes, etc.,
 whereas Italian musical terms are altos, bassos, cantos, pianos, solos, etc.,
 and there are Spanish words like tangos, armadillos, etc.
 I once held a trademark on 'Pentomino(-es)', but I now prefer
 to let these words be my contribution to the language as public domain.*

— SOLOMON W. GOLOMB, letter to Donald Knuth (16 February 1994)

Everybody agrees that seven of the pentominoes should be named after seven consecutive letters of the alphabet:



But two different systems of nomenclature have been proposed for the other five:



where Golomb likes to think of the word 'Filipino' while Conway prefers to map the twelve pentominoes onto the twelve consecutive letters. Conway's scheme tends to work better in computer programs.

A minimum number of blocks of simple form are employed. . . . Experiments and calculations have shown that from the set of seven blocks it is possible to construct approximately the same number of geometrical figures as could be constructed from twenty-seven separate cubes.

— PIET HEIN, *United Kingdom Patent Specification 420,349* (1934)

HEIN
cuboids
parallelepipeds
Hein
Soma
Gardner
Parker Brothers
pentominoes
canonical
factoring

* * *

The simplest polycubes are *cuboids*—also called rectangular parallelepipeds by people who like long names. But things get even more interesting when we consider noncuboidal shapes. Piet Hein noticed in 1933 that the seven smallest shapes of that kind, namely



1: bent 2: ell 3: tee 4: skew 5: L-twist 6: R-twist 7: claw

can be put together to form a $3 \times 3 \times 3$ cube, and he liked the pieces so much that he called them *Soma*. Notice that the first four pieces are essentially planar, while the other three are inherently three-dimensional. Moreover, the two twists are mirror images: We can't change one into the other without entering the *fourth* dimension. Martin Gardner wrote about the joys of Soma in *Scientific American* **199, 3** (September 1958), 182–188, and it soon became wildly popular: More than two million SOMA[®] cubes were sold in America alone, after Parker Brothers began to market a well-made set with an instruction booklet written by Hein.

The task of packing these seven pieces into a cube is easy to formulate as an exact cover problem, just as we did when packing pentominoes. This time we have 24 3D-rotations of the pieces to consider, instead of 8 2D-rotations and/or 3D-reflections; so exercise 200 is used instead of exercise 140 to generate the rows of the problem. It turns out that there are 688 rows, involving 34 columns that we can call 1, 2, . . . , 7, 111, 112, . . . , 333. For example, the first row

1 111 121 211

characterizes one of the potential ways to place the “bent” piece 1.

Algorithm D needs just 407 megamems to find all 11,520 solutions to this problem. Furthermore, we can save most of that time by taking advantage of symmetry: Every solution can be rotated into a unique “canonical” solution in which the “ell” piece 2 has not been rotated; hence we can restrict that piece to only six placements, namely (111, 121, 131, 211), (112, 122, 132, 212), . . . , (213, 223, 233, 313)—all shifts of each other. This removes 138 rows, and the algorithm now finds the 480 canonical solutions in just 20 megamems. (These canonical solutions form 240 mirror-image pairs.)

Factoring an exact cover problem. In fact, we can simplify the Soma cube problem much further, so that all of its solutions can actually be found by hand in a reasonable time, by *factoring* the problem in a clever way. . . .

Color-controlled covering. *Take a break!* Before reading any further, please spend a minute or two solving the “word search” puzzle in Fig. 71; comparatively mindless puzzles like this one provide a low-stress way to sharpen your word-recognition skills. It can be solved easily—for instance, by making eight passes over the array—and the solution appears in Fig. 72.

color-controlled-
word search
color codes

Fig. 71. Find the mathematicians*:

Put ovals around the following names where they appear in the 15×15 array shown here, reading either forward or backward or upward or downward, or diagonally in any direction. After you’ve finished, the leftover letters will form a hidden message. (The solution appears on the next page.)

ABEL	HENSEL	MELLIN
BERTRAND	HERMITE	MINKOWSKI
BOREL	HILBERT	NETTO
CANTOR	HURWITZ	PERRON
CATALAN	JENSEN	RUNGE
FROBENIUS	KIRCHHOFF	STERN
GLAISHER	KNOPP	STIELTJES
GRAM	LANDAU	SYLVESTER
HADAMARD	MARKOFF	WEIERSTRASS

O	T	H	E	S	C	A	T	A	L	A	N	D	A	U
T	S	E	A	P	U	S	T	H	O	R	S	R	O	F
T	L	S	E	E	A	Y	R	R	L	Y	H	A	P	A
E	P	E	A	R	E	L	R	G	O	U	E	M	S	I
N	N	A	R	R	C	V	L	T	R	T	A	A	M	A
I	T	H	U	O	T	E	K	W	I	A	N	D	E	M
L	A	N	T	N	B	S	I	M	I	C	M	A	A	W
L	G	D	N	A	R	T	R	E	B	L	I	H	C	E
E	R	E	C	I	Z	E	C	E	P	T	N	E	D	Y
M	E	A	R	S	H	R	H	L	I	P	K	A	T	H
E	J	E	N	S	E	N	H	R	I	E	O	N	E	T
H	S	U	I	N	E	B	O	R	F	E	W	N	A	R
T	M	A	R	K	O	F	F	O	F	C	S	O	K	M
P	L	U	T	E	R	P	F	R	O	E	K	G	R	A
G	M	M	I	N	S	E	J	T	L	E	I	T	S	G

Our goal in this section is not to discuss how to *solve* such puzzles; instead, we shall consider how to *create* them. It’s by no means easy to pack those 27 names into the box in such a way that their 184 characters occupy only 135 cells, with eight directions well mixed. How can that be done with reasonable efficiency?

For this purpose we shall extend the idea of exact covering by introducing “color codes.” . . .

* The journal *Acta Mathematica* celebrated its 21st birthday by publishing a special *Table Générale des Tomes 1–35*, edited by Marcel Riesz (Uppsala: 1913), 179 pp. It contained a complete list of all papers published so far in that journal, together with portraits and brief biographies of all the authors. The 27 mathematicians mentioned in Fig. 71 are those who were subsequently mentioned in Volumes 1, 2, or 3 of *The Art of Computer Programming*—except for people like MITTAG-LEFFLER or POINCARÉ, whose names contain special characters.


Fig. 72. Solution to the puzzle of the hidden mathematicians (Fig. 71). Notice that the central letter R actually participates in six different names:

- BERTRAND
- GLAISHER
- HERMITE
- HILBERT
- KIRCHHOFF
- WEIERSTRASS

The T to its left participates in five.

Here's what the leftover letters say:

These authors of early papers in *Acta Mathematica* were cited years later in *The Art of Computer Programming*.



EXERCISES — First Set

- **5.** [26] Let T be any tree. Construct an unsolvable exact cover problem for which T is the backtrack tree traversed by Algorithm D; a *unique* column should have the minimum size whenever step D? is encountered. Illustrate your construction when $T = \wedge \wedge \wedge$.
- 6.** [25] Continuing exercise 5, let T be a tree in which certain leaves have been distinguished from the others and designated as “solutions.”
- Show that some such trees never match the behavior of Algorithm D.
 - Characterize all such trees that *do* arise, having solutions where indicated.
- 10.** [M21] The solution to an exact cover problem with matrix A can be regarded as a binary vector x such that $xA = 11\dots 1$. The *distance* between two solutions x and x' can then be defined as the Hamming distance $d(x, x') = \nu(x \oplus x')$, the number of places where x and x' differ. The *diversity* of A is the minimum distance between two of its solutions. (If A has at most one solution, its diversity is infinite.)
- Is it possible to have diversity 1?
 - Is it possible to have diversity 2?
 - Is it possible to have diversity 3?
 - Prove that if A represents a *uniform* exact cover problem, the distance between solutions is always even.
 - Most of the exact cover problems that arise in applications are at least *quasi-uniform*, in the sense that they have a nonempty subset C of primary columns such that $A \upharpoonright C$ has the same number of 1s in every row. (For example, every polyomino or polycube packing problem is quasi-uniform, because every row of the matrix specifies exactly one piece name.) Can such problems have odd distances?
- 19.** [M16] Given an exact cover problem A , construct an exact cover problem A' that has exactly one more solution than A does. [Consequently it is NP-hard to determine whether an exact cover problem with at least one solution has more than one solution.] Assume that A contains no all-zero rows.
- 20.** [M25] Given an exact cover problem A , construct an exact cover problem A' such that (i) A' has at most three 1s in every column; (ii) A' and A have exactly the same number of solutions.
- 21.** [M21] Continuing exercise 20, construct A' having *exactly* three 1s per column.
- **24.** [30] Given an $m \times n$ exact cover problem A with exactly three 1s per column, construct a generalized “instant insanity” problem with $N = O(n)$ cubes and N colors that is solvable if and only if A is solvable. (See 7.2.2-(36).)
- **26.** [M24] A *grope* is a set G together with a binary operation \circ , in which the identity $x \circ (y \circ x) = y$ is satisfied for all $x \in G$ and $y \in G$.
- Prove that the identity $(x \circ y) \circ x = y$ also holds, in every grope.
 - Which of the following “multiplication tables” define a grope on $\{0, 1, 2, 3\}$?

0123	0321	0132	0231	0312
1032	3210	1023	3102	2130
2301	2103	3210	1320	3021
3210	1032	2301	2013	1203

(In the first example, $x \circ y = x \oplus y$; in the second, $x \circ y = (-x - y) \bmod 4$. The last two have $x \circ y = x \oplus f(x \oplus y)$ for certain functions f .)

- For all n , construct a grope whose elements are $\{0, 1, \dots, n-1\}$.
- Consider the exact cover problem that has n^2 columns (x, y) for $0 \leq x, y < n$ and the following $n + (n^3 - n)/3$ rows:
















backtrack tree
distance
Hamming distance
diversity
uniform
quasi-uniform
NP-hard
unique solution
instant insanity
grope
binary operation
multiplication tables


- i) $\{(x, x)\}$, for $0 \leq x < n$;
- ii) $\{(x, x), (x, y), (y, x)\}$, for $0 \leq x < y < n$;
- iii) $\{(x, y), (y, z), (z, x)\}$, for $0 \leq x < y, z < n$.

Show that its solutions are in one-to-one correspondence with the multiplication tables of gropes on the elements $\{0, 1, \dots, n-1\}$.

- e) Element x of a grope is *idempotent* if $x \circ x = x$. If k elements are idempotent and $n-k$ are not, prove that $k \equiv n^2 \pmod{3}$.

27. [21] Modify the exact cover problem of exercise 26(d) in order to find the multiplication tables of (a) all idempotent gropes—grope such that $x \circ x = x$ for all x ; (b) all commutative gropes—grope such that $x \circ y = y \circ x$ for all x and y ; (c) all gropes with an identity element—grope such that $x \circ 0 = 0 \circ x = x$ for all x .

30. [21] *Dominosa* is a solitaire game in which you “shuffle” the 28 pieces , , , , , , , , , , , , , ,  of double-six dominoes and place them at random into a 7×8 frame. Then you write down the number of spots in each cell, put the dominoes away, and try to reconstruct their positions based only on that 7×8 array of numbers. For example,



yields the array

$$\begin{pmatrix} 0 & 0 & 5 & 2 & 1 & 4 & 1 & 2 \\ 1 & 4 & 5 & 3 & 5 & 3 & 5 & 6 \\ 1 & 1 & 5 & 6 & 0 & 0 & 4 & 4 \\ 4 & 4 & 5 & 6 & 2 & 2 & 2 & 3 \\ 0 & 0 & 5 & 6 & 1 & 3 & 3 & 6 \\ 6 & 6 & 2 & 0 & 3 & 2 & 5 & 1 \\ 1 & 5 & 0 & 4 & 4 & 0 & 3 & 2 \end{pmatrix}.$$

- a) Show that *another* placement of dominoes also yields the same matrix of numbers.
- b) What domino placement yields the array

$$\begin{pmatrix} 3 & 3 & 6 & 5 & 1 & 5 & 1 & 5 \\ 6 & 5 & 6 & 1 & 2 & 3 & 2 & 4 \\ 2 & 4 & 3 & 3 & 3 & 6 & 2 & 0 \\ 4 & 1 & 6 & 1 & 4 & 4 & 6 & 0 \\ 3 & 0 & 3 & 0 & 1 & 1 & 4 & 4 \\ 2 & 6 & 2 & 5 & 0 & 5 & 0 & 0 \\ 2 & 5 & 0 & 5 & 4 & 2 & 1 & 6 \end{pmatrix}?$$

► **31.** [20] Show that *Dominosa* reconstruction is a special case of 3D MATCHING.

32. [M22] Generate random instances of *Dominosa*, and estimate the probability of obtaining a 7×8 matrix with a unique solution. Use two models of randomness: (i) Each matrix whose elements are permutations of the multiset $\{8 \times 0, 8 \times 1, \dots, 8 \times 6\}$ is equally likely; (ii) each matrix obtained from a random shuffle of the dominoes is equally likely.

39. [20] By setting up an exact cover problem and solving it with Algorithm D, show that the queen graph Q_8 (exercise 7.1.4–241) cannot be colored with eight colors.

40. [21] In how many ways can Q_8 be colored in a “balanced” fashion, using eight queens of color 0 and seven each of colors 1 to 8?

► **50.** [21] If we merely want to count the number of solutions to an exact cover problem, without actually constructing them, a completely different approach based on bitwise manipulation instead of list processing is sometimes useful.

The following naive algorithm illustrates the idea: We’re given an $m \times n$ matrix of 0s and 1s, represented as n -bit vectors r_1, \dots, r_m . The algorithm works with a

idempotent
commutative
identity element
Dominosa
solitaire
game
Pijowski solitaire, see *Dominosa*
dominoes
3D MATCHING
permutations of the multiset
queen graph
colored
exact cover problem
bitwise manipulation
breadth-first
0s and 1s

(potentially huge) database of pairs (s_j, c_j) , where s_j is an n -bit number representing a set of columns, and c_j is a positive integer representing the number of ways to cover that set exactly. Let p be the n -bit mask that represents the primary columns.

N1. [Initialize.] Set $N \leftarrow 1$, $s_1 \leftarrow 0$, $c_1 \leftarrow 1$, $k \leftarrow 1$.

N2. [Done?] If $k > m$, terminate; the answer is $\sum_{j=1}^N c_j [s_j \& p = p]$.

N3. [Append r_k where possible.] Set $t \leftarrow r_k$. For $N \geq j \geq 1$, if $s_j \& t = 0$, insert $(s_j + t, c_j)$ into the database (see below).

N4. [Loop on k .] Set $k \leftarrow k + 1$ and return to N2. ■

To insert (s, c) there are two cases: If $s = s_i$ for some (s_i, c_i) already present, we simply set $c_i \leftarrow c_i + c$. Otherwise we set $N \leftarrow N + 1$, $s_N \leftarrow s$, $c_N \leftarrow c$.

Show that this algorithm can be significantly improved by using the following trick: Set $u_k \leftarrow r_k \& \bar{f}_k$, where $f_k = r_{k+1} | \dots | r_m$ is the bitwise OR of all future rows. If $u_k \neq 0$, we can remove any item from the database for which s_j does not contain $u_k \& p$. We can also exploit the nonprimary columns of u_k to compress the database further.

51. [25] Implement the improved algorithm of the previous exercise, and compare its running time to that of Algorithm D when applied to the n queens problem.

52. [M21] Explain how the method of exercise 50 could be extended to give representations of all solutions, instead of simply counting them.

70. [25] Prove that the exact cover problem with color controls is NP-complete, even if every row of the matrix has only two entries.

80. [22] Using the “word search puzzle” conventions of Figs. 71 and 72, show that the words ONE, TWO, THREE, FOUR, FIVE, SIX, SEVEN, EIGHT, NINE, TEN, ELEVEN, and TWELVE can all be packed into a 6×6 square, leaving one cell untouched.

81. [22] Also pack *two* copies of ONE, TWO, THREE, FOUR, FIVE into a 5×5 square.

► **82.** [32] The first 44 presidents of the U.S.A. had 38 distinct surnames: ADAMS, ARTHUR, BUCHANAN, BUSH, CARTER, CLEVELAND, CLINTON, COOLIDGE, EISENHOWER, FILLMORE, FORD, GARFIELD, GRANT, HARDING, HARRISON, HAYES, HOOVER, JACKSON, JEFFERSON, JOHNSON, KENNEDY, LINCOLN, MADISON, MCKINLEY, MONROE, NIXON, OBAMA, PIERCE, POLK, REAGAN, ROOSEVELT, TAFT, TAYLOR, TRUMAN, TYLER, VANBUREN, WASHINGTON, WILSON.


a) What's the smallest square into which all of these names can be packed, using word search conventions, and requiring all words to be *connected* via overlaps?

b) What's the smallest *rectangle*, under the same conditions?

► **83.** [25] Pack as many of the following words as possible into a 9×9 array, simultaneously satisfying the rules of *both* word search *and* sudoku:

ACRE	COMPARE	CORPORATE	MACRO	MOTET	ROAM
ART	COMPUTER	CROP	META	PARAMETER	TAME

90. [24] Find the unique solutions to the following examples of polyomino sudoku:



100. [M25] Consider a weighted exact cover problem in which we must choose 2 of 4 rows to cover column 1, and 5 of 7 rows to cover column 2; the rows don't interact.

primary columns
bitwise AND
bitwise OR
nonprimary columns
 n queens problem
color controls
NP-complete
2D matching
word search puzzle
presidents
I'm not sure
how many of
these names
should go in
the index
connected
word search
sudoku
polyomino sudoku
sudoku
weighted exact cover problem

- a) What's the size of the search tree if we branch first on column 1, then on column 2? Would it be better to branch first on column 2, then on column 1?
- b) Generalize part (a) to the case when column 1 needs p of $p + d$ rows, while column 2 needs q of $q + d$ rows, where $q > p$ and $d > 0$.

EXERCISES — Second Set

Hundreds of fascinating recreational problems have been based on polyominoes and their cousins (the polycubes, polyiamonds, polyhexes, polysticks, ...). The following exercises explore “the cream of the crop” of such classic puzzles, as well as a few gems that were not discovered until recently.

In most cases the idea is to find a good way to discover all solutions, usually by setting up an appropriate exact cover problem that can be solved without taking an enormous amount of time.

Conway
five-letter words
pentominoes
nonstraight
symmetry
pentominoes
Clarke
tetrominoes
tetrominoes
three-colorable
graph coloring

- **140.** [25] Sketch the design of a utility program that will create sets of rows by which an exact cover solver will fill a given shape with a given set of polyominoes.

148. [18] Using Conway’s piece names, pack five pentominoes into the shape so that they spell a common English word when read from left to right.



- **150.** [21] There are 1010 ways to pack the twelve pentominoes into a 5×12 box, not counting reflections. What’s a good way to find them all, using Algorithm D?

151. [21] How many of those 1010 packings decompose into $5 \times k$ and $5 \times (12 - k)$?


152. [21] In how many ways can the eleven nonstraight pentominoes be packed into a 5×11 box, not counting reflections? (Reduce symmetry cleverly.)

154. [20] There are 2339 ways to pack the twelve pentominoes into a 6×10 box, not counting reflections. What’s a good way to find them all, using Algorithm D?

155. [23] Continuing exercise 154, explain how to find special kinds of packings:

- Those that decompose into $6 \times k$ and $6 \times (10 - k)$.
- Those that have all twelve pentominoes touching the outer boundary.
- Those with all pentominoes touching that boundary *except* for V, which doesn’t.
- Same as (c), with each of the other eleven pentominoes in place of V.
- Those with the *minimum* number of pentominoes touching the outer boundary.
- Those that are characterized by Arthur C. Clarke’s description, as quoted in the text. (That is, the X should touch only the F, N, U, and V—no others.)

157. [21] There are five different *tetrominoes*, namely




In how many essentially different ways can each of them be packed into an 8×8 square together with the twelve pentominoes?

158. [21] If an 8×8 checkerboard is cut up into thirteen pieces, representing the twelve pentominoes together with one of the tetrominoes, some of the pentominoes will have more black cells than white. Is it possible to do this in such a way that U, V, W, X, Y, Z have a black majority while the others do not?

159. [18] Design a nice, simple tiling pattern that’s based on the five tetrominoes.

160. [25] How many of the 6×10 pentomino packings are *strongly three-colorable*, in the sense that each individual piece could be colored red, white, or blue in such a way that no pentominoes of the same color touch each other — not even at corner points?

- **162.** [20] The black cells of a square $n \times n$ checkerboard form an interesting graph called the *Aztec diamond* of order $n/2$. For example, the cases $n = 11$ and 13 are



where (ii) has a “hole” showing the case $n = 3$. Thus (i) has 61 cells, and (ii) has 80.

- Find all ways to pack (i) with the twelve pentominoes and one monomino.
- Find all ways to pack (ii) with the 12 + 5 pentominoes and tetrominoes.

Speed up the process by not producing solutions that are symmetric to each other.

- **163.** [M26] Arrange the twelve pentominoes into a Möbius strip of width 4. The pattern should be “fault-free”: Every straight line must intersect some piece.

164. [40] (H. D. Benjamin, 1948.) Show that the twelve pentominoes can be wrapped around a cube of size $\sqrt{10} \times \sqrt{10} \times \sqrt{10}$. For example, here are front and back views of such a cube, made from twelve colorful fabrics by the author's wife in 1993:



(Photos by Hector Garcia)

What is the best way to do this, minimizing undesirable distortions at the corners?

- **165.** [22] (Craig S. Kaplan.) A polyomino can sometimes be surrounded by non-overlapping copies of itself that form a *fence*: Every cell that touches the polyomino — even at a corner — is part of the fence; conversely, every piece of the fence touches the inner polyomino. Furthermore, the pieces must not enclose any unoccupied “holes.”

Find the (a) smallest and (b) largest fences for each of the twelve pentominoes. (Some of these patterns are unique, and quite pretty.)

166. [22] Solve exercise 165 for fences that satisfy the *tatami* condition of exercise 7.1.4–215: No four edges of the tiles should come together at any “crossroads.”

168. [21] (T. H. O’Beirne, 1961.) The *one-sided pentominoes* are the eighteen distinct 5-cell pieces that can arise if we aren’t allowed to flip pieces over:



Notice that there now are two versions of F, L, P, N, Y, and Z.


In how many ways can all eighteen of them be packed into rectangles?

169. [21] Suppose you want to pack the twelve pentominoes into a 6×10 box, *without* turning any pieces over. Then 2^6 different problems arise, depending on which sides of the one-sided pieces are present. Which of those 64 problems has (a) the fewest (b) the most solutions?

checkerboard
Aztec diamond
symmetric
Möbius strip
fault-free
Benjamin
cube, wrapped
Knuth, Jill
Garcia, Hector
Kaplan
fence
holes
tatami
crossroads
O’Beirne
one-sided pentominoes


170. [21] When tetrominoes are both checkered and one-sided (see exercises 158 and 168), ten possible pieces arise. In how many ways can all ten of them fill a rectangle?

175. [20] There are 35 *hexominoes*, first enumerated in 1934 by the master puzzlist H. D. Benjamin. At Christmastime that year, he offered ten shillings to the first person who could pack them into a 14×15 rectangle — although he wasn't sure whether or not it could be done. The prize was won by F. Kadner, who proved that the hexominoes actually *can't* be packed into *any* rectangle. Nevertheless, Benjamin continued to play with them, eventually discovering that they fit nicely into the triangle shown here.



Prove Kadner's theorem. *Hint:* See exercise 158.

176. [24] (Frans Hansson, 1947.) The fact that $35 = 1^2 + 3^2 + 5^2$ suggests that we might be able to pack the hexominoes into three boxes that represent a *single* hexomino shape at three levels of magnification, such as




For which hexominoes can this be done?

► **177.** [30] Show that the 35 hexominoes can be packed into five “castles”:




In how many ways can this be done?

178. [41] For which values of m can the hexominoes be packed into a box like this?



179. [41] Perhaps the best hexomino packing uses a 5×45 rectangle with 15 holes



proposed by W. Stead in 1954. In how many ways can the 35 hexominoes fill it?

► **181.** [22] In how many ways can the twelve pentominoes be placed into an 8×10 rectangle, leaving holes in the shapes of the five tetrominoes? (The holes should not touch the boundary, nor should they touch each other, even at corners; one example is shown at the right.) Explain how to encode this puzzle as an exact cover problem with color controls.



182. [46] If possible, solve the analog of exercise 181 for the case of 35 *hexominoes* in a 5×54 rectangle, leaving holes in the shapes of the twelve *pentominoes*.

► **198.** [HM35] A *parallelogram polyomino*, or “parallomino” for short, is a polyomino whose boundary consists of two paths that each travel only north and/or east. (Equivalently, it is a “skew Young tableau” or a “skew Ferrers board,” the difference between

tetrominoes
checkered
one-sided
checkerboard dissections
hexominoes
Benjamin
Kadner
Hansson
magnification
triplication
castles
Stead
pentominoes
tetrominoes
color controls
hexominoes
parallelogram polyomino
parallomino
skew Young tableau
Young tableaux
skew Ferrers board
Ferrers diagrams

the diagrams of two tableaux or partitions; see Sections 5.1.4 and 7.2.1.4.) For example, there are five parallominoes whose boundary paths have length 4:




- a) Find a one-to-one correspondence between the set of ordered trees with m leaves and n nodes and the set of parallominoes with width m and height $n - m$. The area of each parallomino should be the path length of its corresponding tree.
- b) Study the generating function $G(w, x, y) = \sum_{\text{parallominoes}} w^{\text{area}} x^{\text{width}} y^{\text{height}}$.
- c) Prove that the parallominoes whose width-plus-height is n have total area 4^{n-2} .
- d) Part (c) suggests that we might be able to pack all of those parallominoes into a $2^{n-2} \times 2^{n-2}$ square, *without* rotating them or flipping them over. Such a packing is clearly impossible when $n = 3$ or $n = 4$; but is it possible when $n = 5$ or $n = 6$?

200. [20] Extend exercise 140 to three dimensions. How many base placements do each of the seven Soma pieces have?


- ▶ **202.** [22] The *Somap* is the graph whose vertices are the 240 distinct solutions to the Soma cube problem, with $u - v$ if and only if u can be obtained from v by changing the positions of at most three pieces. (Using the terminology of exercise 10(d), adjacent vertices correspond to solutions of *semidistance* ≤ 3 .) The *strong Somap* is similar, but it has $u - v$ only when a change of just *two* pieces gets from one to the other.
 - a) What are the degree sequences of these graphs?
 - b) How many connected components do they have? How many bicomponents?

- ▶ **204.** [M25] Use factorization to prove that Fig. 80's W-wall cannot be built.
- 205.** [24] Figure 80(a) shows some of the many "low-rise" (2-level) shapes that can be built from the seven Soma pieces. Which of them is hardest (has the fewest solutions)? Which is easiest? Answer these questions also for the 3-level prism shapes in Fig. 80(b).

- ▶ **206.** [M23] Generalizing the first four examples of Fig. 80, study the set of *all* shapes obtainable by deleting three cubies from a $3 \times 5 \times 2$ box. (Two examples are shown here.) How many essentially different shapes are possible? Which shape is easiest? Which shape is hardest?




- 207.** [22] Similarly, consider (a) all shapes that consist of a $3 \times 4 \times 3$ box with just three cubies in the top level; (b) all 3-level prisms that fit into a $3 \times 4 \times 3$ box.



208. [25] How many of the 1285 *nonominoes* define a prism that can be realized by the Soma pieces? Do any of those packing problems have a unique solution?

210. [M40] Make empirical tests of Piet Hein's belief that the number of shapes achievable with seven Soma pieces is approximately the number of 27-cubie polycubes.

212. [20] (B. L. Schwartz, 1969.) Show that the Soma pieces can make shapes that appear to have more than 27 cubies, because of holes hidden inside or at the bottom:



staircase

penthouse

pyramid

In how many ways can these three shapes be constructed?

tableaux
partitions
trees
path length
generating function
base placements
Somap
Soma cube
semidistance
degree sequences
connected components
bicomponents
factorization
W-wall
Soma pieces
nonominoes
Hein
Schwartz

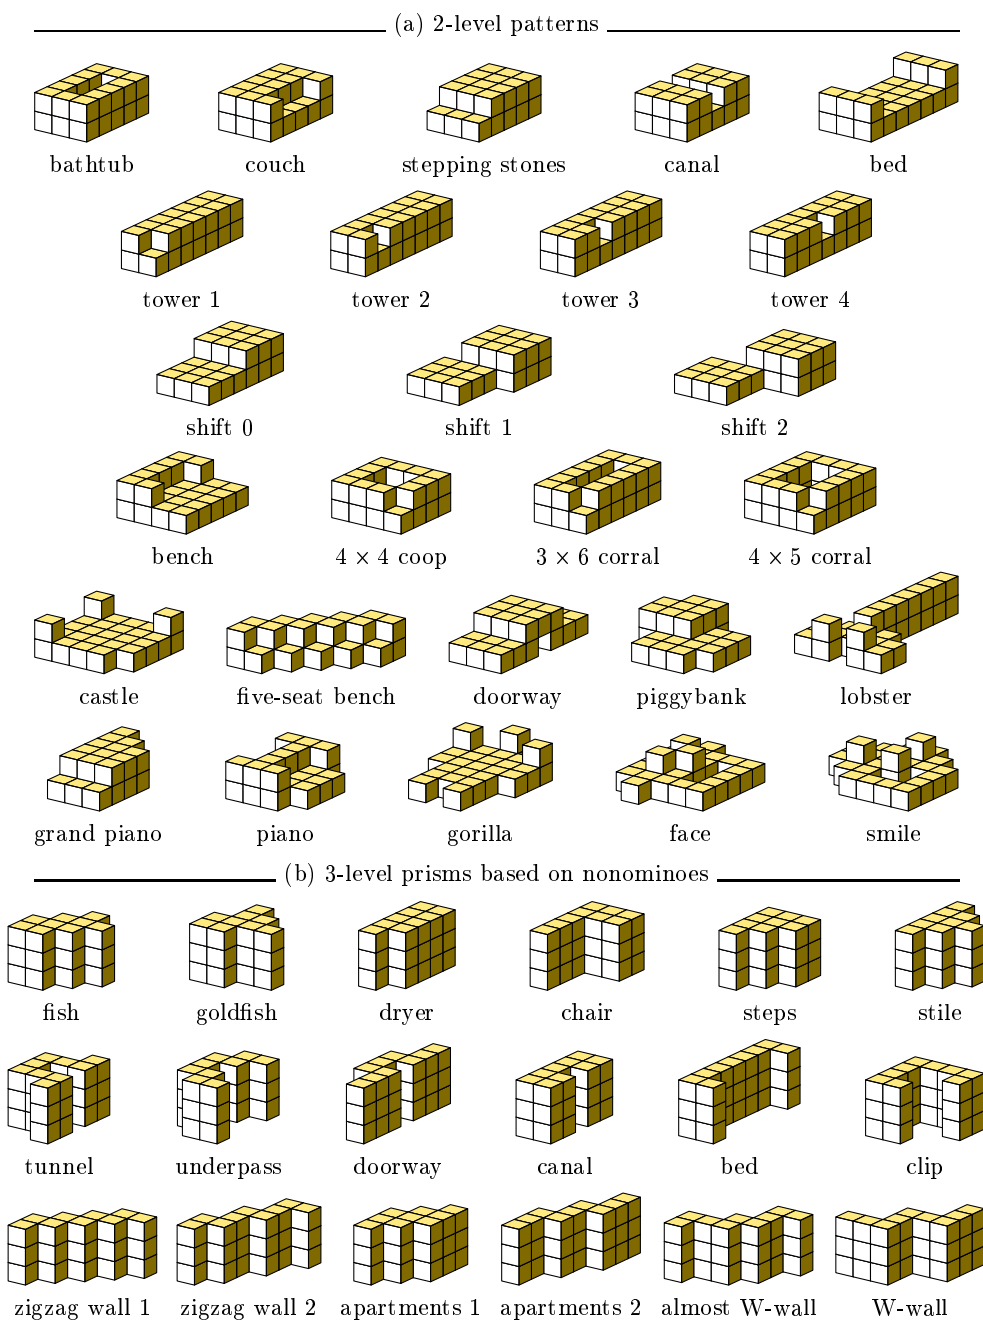




Fig. 80. Gallery of noteworthy polycubes that contain 27 cubies. All of them can be built from the seven Soma pieces, except for the W-wall. Many constructions are also stable when tipped on edge and/or when turned upside down. (See exercises 204–214.)

213. [22] Show that the seven Soma pieces can also make structures such as



which are “self-supporting” via gravity. (You may need to place a small book on top.)

► 214. [M32] Impossible structures *can* be built, if we insist only that they look genuine when viewed from the front (like façades in Hollywood movies)! Find all solutions to




that are visually correct. (In order to solve this exercise, you need to know that the illustrations here use the non-isometric projection $(x, y, z) \mapsto (30x - 42y, 14x + 10y + 45z)u$ from three dimensions to two, where u is a scale factor.) All seven pieces must be used.

215. [30] The earliest known example of a polycube puzzle is the “Cube Diabolique,” manufactured in late nineteenth century France by Charles Watilliaux; it contains six flat pieces of sizes 2, 3, . . . , 7:



- a) In how many ways do these pieces make a $3 \times 3 \times 3$ cube?
- b) Are there six polycubes, of sizes 2, 3, . . . , 7, that make a cube in just *one* way?


217. [22] Show that there are exactly eight different *tetracubes* — polycubes of size 4. Which of the following shapes can they make, respecting gravity? How many solutions are possible?



218. [25] How many of the 369 *octominoes* define a 4-level prism that can be realized by the tetracubes? Do any of those packing problems have a unique solution?

self-supporting
gravity
façades
movies
isometric
projection
three dimensions
Cube Diabolique
Diabolical Cube
Watilliaux
tetracubes
gravity
octominoes

220. [30] There are 29 *pentacubes*, conveniently identified with one-letter codes:




pentacubes
 solid pentacubes
 flat pentacubes
 mirror images
 pentominoes
 $5 \times 5 \times 5$ cube
 Dowler's Box
 chiral
 mirror

Pieces o through z are called, not surprisingly, the *solid pentominoes* or *flat pentacubes*.


- a) What are the mirror images of a, b, c, d, e, f, A, B, C, D, E, F, j, k, l, . . . , z?
- b) In how many ways can the solid pentominoes be packed into an $a \times b \times c$ cuboid?
- c) What “natural” set of 25 pentacubes is able to fill the $5 \times 5 \times 5$ cube?

► **221.** [25] The full set of 29 pentacubes can build an enormous variety of elegant structures, including a particularly stunning example called “Dowler’s Box.” This $7 \times 7 \times 5$ container, first considered by R. W. M. Dowler in 1979, is constructed from five flat slabs. Yet only 12 of the pentacubes lie flat; the other 17 must somehow be worked into the edges and corners.




Despite these difficulties, Dowler’s Box has so many solutions that we can actually impose many further conditions on its construction:

- a) Build Dowler’s Box in such a way that the chiral pieces a, b, c, d, e, f and their images A, B, C, D, E, F all appear in horizontally mirror-symmetric positions.




horizontally symmetric c and C



diagonally symmetric c and C

- b) Alternatively, build it so that those pairs are *diagonally* mirror-symmetric.
- c) Alternatively, place piece x in the center, and build the remaining structure from four congruent pieces that have seven pentacubes each.

222. [25] The 29 pentacubes can also be used to make the shape shown here, exploiting the curious fact that $3^4 + 4^3 = 29 \cdot 5$. But Algorithm D will take a long, long time before telling us how to construct it, unless we’re lucky, because the space of possibilities is huge. How can we find a solution quickly?



999. [M00] this is a temporary exercise (for dummies)

Dr Pell was wont to say, that in the Resolution of Questiones, the main matter is the well stating them: which requires a good mother-witt & Logick: as well as Algebra: for let the Question be but well-stated, and it will worke of it selfe: ... By this way, an man cannot intangle his notions, & make a false Steppe.

— JOHN AUBREY, *An Idea of Education of Young Gentlemen* (c. 1684)

Pell
AUBREY
semidistance
Matsui
Matsui
NP-complete
minimum remaining values heuristic

SECTION 7.2.2.1

5. If T has only a root node, let there be one column, no rows. Otherwise let T have $d \geq 1$ subtrees T_1, \dots, T_d , and assume that we've constructed problems with rows R_j and columns C_j for each T_j . Let $C = C_1 \cup \dots \cup C_d \cup \{1, \dots, d\}$. The problem for T is obtained by appending $d+1$ new columns $\{0, 1, \dots, d\}$ and the following new rows:
(i) '0 and all columns of $C \setminus C_j$ ', for $1 \leq j \leq d$; (ii) 'all columns of $C \setminus j$ ', for $1 \leq j \leq d$. This construction works except when $d = 1$ and T_1 is a leaf; in that case we can use columns $\{0, 1, 2, 3\}$, rows '0 1 2', '1 3', '2 3'. The matrix for the example tree has 17 columns and 16 rows.

```
01111110000000000
10111110000000000
11011110000000000
11100110000000000
11101010000000000
11100110000000000
0000000111111000
0000000101111000
0000000110111000
00000001110011000
00000001110101000
00000001110110000
00000001111111111
11111110000000111
111111111111111001
1111111111111010
```

6. (a) If a solution isn't at the root, its parent must have exactly one child. (Alternatively, if duplicate rows are permitted, all siblings of a solution must be solutions.)

(b) Use the previous construction; a solution node corresponds to column 0, row '0'.

10. (a) No. Otherwise A would have a row that's zero in all primary columns.

(b) Yes, but only if A has two rows that are identical in all primary columns.

(c) Yes, but only if A has two rows whose sum is also a row, when restricted to primary columns.

(d) The number of places, j , where $x = 1$ and $x' = 0$ must be the same as the number where $x = 0$ and $x' = 1$. For if A has exactly k primary 1s in every row, exactly jk primary columns are being covered in different ways.

(e) Again the distances must be even, because every solution to A is also a solution to the uniform problem $A|C$. (Therefore it makes sense to speak of the *semidistance* $d(x, x')/2$ between solutions of quasi-uniform exact covering problem. The semidistance in a polyform packing problem is the number of pieces that are packed differently.)

19. (Solution by T. Matsui.) Add one new column at the left of A , all 0s. Then add two rows of length $n+1$ at the bottom: $10\dots 0$ and $11\dots 1$. This $(m+2) \times (n+1)$ matrix A' has one solution that chooses only the last row. All other solutions choose the second-to-last row, together with rows that solve A .

20. (Solution by T. Matsui.) Assume that all 1s in column 1 appear in the first t rows, where $t > 3$. Add two new columns at the left, and two new rows $1100\dots 0$, $1010\dots 0$ of length $n+2$ at the bottom. For $1 \leq k \leq t$, if row k was $1\alpha_k$, replace it by $010\alpha_k$ if $k \leq t/2$, $011\alpha_k$ if $k > t/2$. Insert 00 at the left of the remaining rows $t+1$ through m .

This construction can be repeated (with suitable row and column permutations) until no column sum exceeds 3. If the original column sums were (c_1, \dots, c_n) , the new A' has $2T$ more rows and $2T$ more columns than A did, where $T = \sum_{j=1}^n (c_j \div 3)$.

One consequence is that the exact cover problem is NP-complete even when restricted to cases where all row and column sums are at most 3.

Notice, however, that this construction is *not* useful in practice, because it disguises the structure of A : It essentially *destroys* the minimum remaining values heuristic, because all columns whose sum is 2 look equally good to the solver!

21. Take a matrix with column sums (c_1, \dots, c_n) , all ≤ 3 , and extend it with three columns of 0s at the right. Then add the following four rows: $(x_1, \dots, x_n, 0, 1, 1)$, $(y_1, \dots, y_n, 1, 0, 1)$, $(z_1, \dots, z_n, 1, 1, 0)$, and $(0, \dots, 0, 1, 1, 1)$, where $x_j = [c_j < 3]$, $y_j = [c_j < 2]$, $z_j = [c_j < 1]$. The bottom row must be chosen in any solution.

24. Consider a set of cubes and colors called $\{*, 0, 1, 2, 3, 4, \dots\}$, where (i) all faces of cube $*$ are colored $*$; (ii) colors 1, 2, 3, 4 occur only on cubes 0, 1, 2, 3, 4; (iii) the opposite face-pairs of those five cubes are respectively $(00, 12, **)$, $(11, 12, 34)$, $(22, 34, \alpha)$, $(33, 12, \beta)$, $(44, 34, \gamma)$, where α, β, γ are pairs of colors $\notin \{1, 2, 3, 4\}$. Any solution to the cube problem has disjoint 2-regular graphs X and Y containing two faces of each color. Since X and Y both contain $**$ from cube $*$, we can assume that X contains 00 and Y contains 12 from cube 0. Hence Y can't contain 11 or 22; it must contain 12 from cube 1 or cube 3. If X doesn't contain 11 or 22, it must contain 12 from cube 1 and cube 3. Hence X contains 11, 22, 33, and 44. We're left with only three possibilities for Y from cubes 1, 2, 3, 4, namely $(34, \alpha, 12, 34)$, $(12, 34, \beta, 34)$, $(34, 34, 12, \gamma)$.

Now let a_{j1}, a_{j2}, a_{j3} denote the 1s in column j of A . We construct $N = 8n + 1$ cubes and colors called $*, a_{jk}, b_{jl}$, where $1 \leq j \leq n, 1 \leq k \leq 3, 0 \leq l \leq 4$. The opposite face-pairs of $*$ are $(**, **, **)$. Those of a_{jk} are $(a_{jk}a_{jk}, a_{jk}a_{jk}, a_{jk}b_{j'0})$, where j' is the column of a_{jk} 's cyclic successor to the right in its row. Those of $b_{j0}, b_{j1}, b_{j2}, b_{j3}, b_{j4}$ are respectively $(b_{j0}b_{j0}, b_{j1}b_{j2}, **)$, $(b_{j1}b_{j1}, b_{j1}b_{j2}, b_{j3}b_{j4})$, $(b_{j2}b_{j2}, b_{j3}b_{j4}, b_{j0}a_{j1})$, $(b_{j3}b_{j3}, b_{j1}b_{j2}, b_{j0}a_{j2})$, $(b_{j4}b_{j4}, b_{j3}b_{j4}, b_{j0}a_{j3})$. By the previous paragraph, solutions to the cube problem correspond to 2-regular graphs X and Y such that, for each j , X or Y contains all the pairs $b_{jl}b_{jl}$ and the other "selects" one of the three pairs $b_{j0}a_{jk}$. The face-pairs of each selected a_{jk} ensure that a_{jk} 's cyclic successor is also selected.

[See E. Robertson and I. Munro, *Utilitas Mathematica* **13** (1978), 99–116.]

26. (a) $(x \circ y) \circ x = (x \circ y) \circ (y \circ (x \circ y)) = y$.

(b) All five are legitimate. (The last two are gropes because $f(t + f(t)) = t$ for $0 \leq t < 4$ in each case. They are isomorphic if we interchange any two elements. The third is isomorphic to the second if we interchange $1 \leftrightarrow 2$. There are 18 grope tables of order 4, of which $(4, 12, 2)$ are isomorphic to the first, third, and last tables shown here.)

(c) For example, let $x \circ y = (-x - y) \pmod n$. (More generally, if G is any group and if $\alpha \in G$ satisfies $\alpha^2 = 1$, we can let $x \circ y = \alpha x^{-1} \alpha y^{-1} \alpha$. If G is commutative and $\alpha \in G$ is arbitrary, we can let $x \circ y = x^{-1} y^{-1} \alpha$.)

(d) For each row of type (i) in an exact covering, define $x \circ x = x$; for each row of type (ii), define $x \circ x = y, x \circ y = y \circ x = x$; for each row of type (iii), define $x \circ y = z, y \circ z = x, z \circ x = y$. Conversely, every grope table yields an exact covering in this way.

(e) Such a grope covers n^2 columns with k rows of size 1, all other rows of size 3. [F. E. Bennett proved, in *Discrete Mathematics* **24** (1978), 139–146, that such gropes exist for all k with $0 \leq k \leq n$ and $k \equiv n^2 \pmod 3$, except when $k = n = 6$.]

Notes: The identity $x \circ (y \circ x) = y$ seems to have first been considered by E. Schröder in *Math. Annalen* **10** (1876), 289–317 [see '(C₀)' on page 306], but he didn't do much with it. In a class for sophomore mathematics majors at Caltech in 1968, the author defined gropes and asked the students to discover and prove as many theorems about them as they could, by analogy with the theory of groups. The idea was to "grope for results." The official modern term for a grope is a real jawbreaker: *semisymmetric quasigroup*.

27. (a) Eliminate the n columns for (x, x) ; use only the $2 \binom{n}{3}$ rows of type (iii) for which $y \neq z$. (Idempotent gropes are equivalent to "Mendelsohn triples," which are families of $n(n-1)/3$ 3-cycles (xyz) that include every ordered pair of distinct elements. N. S.


2-regular graphs
Robertson
Munro
isomorphic
isomorphic
Bennett
Caltech
author
groups
quasigroup
semisymmetric quasigroup
Mendelsohn triples

Mendelsohn proved [*Computers in Number Theory* (New York: Academic Press, 1971), 323–338] that such systems exist for all $n \not\equiv 2 \pmod{3}$, except when $n = 6$.)

(b) Use only the $\binom{n+1}{2}$ columns (x, y) for $0 \leq x \leq y < n$; replace rows of type (ii) by $\{(x, x), (x, y)\}$ and $\{(x, y), (y, y)\}$ for $0 \leq x < y < n$; replace those of type (iii) by $\{(x, y), (x, z), (y, z)\}$ for $0 \leq x < y < z < n$. (Such systems, Schröder's ' (C_1) ' and ' (C_2) ', are called totally symmetric quasigroups; see S. K. Stein, *Trans. Amer. Math. Soc.* **85** (1957), 228–256, §8. If idempotent, they're equivalent to Steiner triple systems.)

(c) Omit columns for which $x = 0$ or $y = 0$. Use only the $2\binom{n-1}{3}$ rows of type (iii) for $1 \leq x < y, z < n$ and $y \neq z$. (Indeed, such systems are equivalent to idempotent gropes on the elements $\{1, \dots, n-1\}$.)

30. In (a), four pieces change; in (b) the solution is unique:



Notice that the spot patterns $\begin{smallmatrix} \bullet & \bullet \\ \bullet & \bullet \end{smallmatrix}$, $\begin{smallmatrix} \bullet & \bullet \\ \bullet & \bullet \end{smallmatrix}$, and $\begin{smallmatrix} \bullet & \bullet \\ \bullet & \bullet \end{smallmatrix}$ are rotated when a domino is placed vertically; these visual clues, which would disambiguate (a), don't show up in the matrix.

[Dominosa was invented in Germany by O. S. Adler [Reichs Patent #71539 (1893); see his booklet written with F. Jahn, *Sperr-Domino und Dominosa* (1912), 23–64. Similar problems of “quadrilles” had been studied earlier by E. Lucas and H. Delannoy; see Lucas's [*Récréations Mathématiques 2* (Paris: Gauthier-Villars, 1883), 52–63].

31. Define 28 vertices Dxy for $0 \leq x \leq y \leq 6$; 28 vertices ij for $0 \leq i < 7, 0 \leq j < 8$, and $i + j$ even; and 28 similar vertices ij with $i + j$ odd. The matching problem has 49 triples of the form $\{Dxy, ij, i(j+1)\}$ for $0 \leq i, j < 7$, as well as 48 of the form $\{Dxy, ij, (i+1)j\}$ for $0 \leq i < 6$ and $0 \leq j < 8$, corresponding to potential horizontal or vertical placements. For example, the triples for exercise 30(a) are $\{D00, 00, 01\}$, $\{D05, 01, 02\}$, \dots , $\{D23, 66, 67\}$; $\{D01, 00, 10\}$, $\{D04, 01, 11\}$, \dots , $\{D12, 57, 67\}$.

32. Model (i) has $M = 56!/8!^7 \approx 4.10 \times 10^{42}$ equally likely possibilities; model (ii) has $N = 1292697 \cdot 28! \cdot 2^{21} \approx 8.27 \times 10^{41}$, because there are 1292697 ways to pack 28 dominoes in a 7×8 frame. (Algorithm D will quickly list them all.) The expected number of solutions per trial in model (i) is therefore $N/M \approx 0.201$.

Ten thousand random trials with model (i) gave 216 cases with at least one solution, including 26 where the solution was unique. The total number $\sum x$ of solutions was 2256; and $\sum x^2 = 95918$ indicated a heavy-tailed distribution whose empirical standard deviation is ≈ 3.1 . The total running time was about $250 M\mu$.

Ten thousand random trials with model (ii), using random choices from a precomputed list of 1292687 packings, gave 106 cases with a unique solution; one case had 2652 of them! Here $\sum x = 508506$ and $\sum x^2 = 144119964$ indicated an empirical mean of ≈ 51 solutions per trial, with standard deviation ≈ 109 . Total time was about $650 M\mu$.

Schröder
totally symmetric quasigroups
Stein
Steiner triple systems
Adler
Jahn
quadrilles
Lucas
Delannoy
dimer tilings
heavy-tailed distribution
empirical standard deviation

39. Each of the 92 solutions to the eight queens problem (see Fig. 68) occupies eight of the 64 cells; so we must find eight disjoint solutions. Only 1897 updates of Algorithm D are needed to show that such a mission is impossible. [In fact no *seven* solutions can be disjoint, because each solution touches at least three of the twenty cells 13, 14, 15, 16, 22, 27, 31, 38, 41, 48, 51, 58, 61, 68, 72, 77, 83, 84, 85, 86. See Thorold Gosset, *Messenger of Mathematics* **44** (1914), 48. Henry E. Dudeney found the illustrated way to occupy all but two cells, in *Tit-Bits* **32** (11 September 1897), 439; **33** (2 October 1897), 3.]

40. This is an exact cover problem with $92 + 312 + 396 + \dots + 312 = 3284$ rows (see exercise 7.2.2-5). Algorithm D needs about 2 million updates to find the solution shown, and about 83 billion to find all 11,092 of them.

50. Set $f_m \leftarrow 0$ and $f_{k-1} \leftarrow f_k \mid r_k$ for $m \geq k > 1$. The bits of u_k represent columns that are being changed for the last time.

Let $u_k = u' + u''$, where $u' = u_k \& p$. If $u_k \neq 0$ at the beginning of step N4, we compress the database as follows: For $N \geq j \geq 1$, if $s_j \& u' \neq u'$, delete (s_j, c_j) ; otherwise if $s_j \& u'' \neq 0$, delete (s_j, c_j) and insert $((s_j \& \bar{u}_k) \mid u', c_j)$.

To delete (s_j, c_j) , set $(s_j, c_j) \leftarrow (s_N, c_N)$ and $N \leftarrow N - 1$.

When this improved algorithm terminates in step N2, we always have $N \leq 1$. Furthermore, if we let $p_k = r_1 \mid \dots \mid r_{k-1}$, the size of N never exceeds 2^{ν_k} , where $\nu_k = \nu(p_k r_k f_k)$ is the size of the “frontier” (see exercise 7.1.4-55).

[In the special case of n queens, represented as the exact cover problem in $(\star\star)$, this algorithm is due to I. Rivin, R. Zabih, and J. Lamping, *Inf. Proc. Letters* **41** (1992), 253-256. They proved that the frontier for n queens never has more than $3n$ columns.]

51. The author has had reasonably good results using a triply linked binary search tree for the database, with randomized search keys. (Beware: The swapping algorithm used for deletion was difficult to get right.) This implementation was, however, limited to exact cover problems whose matrix has at most 64 columns; hence it could do n queens via $(\star\star)$ only when $n < 12$. When $n = 11$ its database reached a maximum size of 75,009, and its running time was about 25 megamems. But Algorithm D was a lot better: It needed only about 780K updates to find all $Q(11) = 2680$ solutions.

In theory, this method will need only about 2^{3n} steps as $n \rightarrow \infty$, times a small polynomial function of n . A backtracking algorithm such as Algorithm D, which enumerates each solution explicitly, will probably run asymptotically slower (see exercise 7.2.2-14). But in practice, a breadth-first approach needs too much space.


On the other hand, this method did beat Algorithm D on the n queen bees problem of exercise 7.2.2-15: When $n = 11$ its database grew to 364,864 items; it computed $H(11) = 596,483$ in just 30 M μ , while Algorithm D needed 27 mega-updates.

52. The set of solutions for s_j can be represented as a regular expression α_j instead of by its size, c_j . Instead of inserting $(s_j + t, c_j)$ in step N3, insert $\alpha_j k$. If inserting (s, α) , when (s_i, α_i) is already present with $s_i = s$, change $\alpha_i \leftarrow \alpha_i \cup \alpha$. [Alternatively, if only one solution is desired, we could attach a single solution to each s_j in the database.]

70. Given a 3SAT problem with clauses $(l_{i1} \vee l_{i2} \vee l_{i3})$ for $1 \leq i \leq m$, with each $l_{ij} \in \{x_1, \bar{x}_1, \dots, x_n, \bar{x}_n\}$, construct an exact cover problem with $3m$ primary columns ij ($1 \leq i \leq m, 1 \leq j \leq 3$) and n secondary columns x_k ($1 \leq k \leq n$), having the following rows: (i) $'l_{i1} l_{i2}', 'l_{i2} l_{i3}', 'l_{i3} l_{i1}'$; (ii) $'l_{ij} x_k:1'$ if $l_{ij} = x_k$, $'l_{ij} x_k:0'$ if $l_{ij} = \bar{x}_k$. That problem has a solution if and only if the given clauses are satisfiable.

eight queens problem
Gosset
Dudeney
frontier
 n queens
Rivin
Zabih
Lamping
author
triply linked
binary search tree
backtracking algorithm
asymptotically
theory vs practice
practice vs theory
 n queen bees
regular expression
3SAT

80. There are just five solutions; the latter two are flawed by being disconnected:




disconnected
Gibat
author
interactive method
Gordon
Eckler
branch, choice of
choice of column to cover
best column
Huang
Snyder

Historical note: Word search puzzles were invented by Norman E. Gibat in 1968.

81. When Algorithm C is generalized to allow non-unit column sums as in Algorithm ?, it needs just 24 megamems to prove that there are exactly eight solutions; hence the solution shown here is essentially unique.



82. (a, b) The author's best solutions, thought to be minimal (but there is no proof), are below. In both cases, and in Fig. 71, an interactive method was used: After the longest words were placed strategically by hand, Algorithm C packed the others nicely.



P	I	E	R	C	(E)	I	S	E	N	H	O	(W)	(R)	U	H	T	(R)	(A)	(H)	(A)	(R)	(D)	(I)	(N)	(G)	(A)	(R)	(F)	(I)	(E)	(L)	(D)	(N)	(A)	(L)	(E)	(V)	(E)	(L)	(C)	(T)	
O	(B)	(A)	(M)	(A)	(D)	(I)	(S)	(O)	(N)	(O)	(S)	(L)	(I)	(W)	(A)	(S)	(H)	(I)	(N)	(G)	(T)	(O)	(N)	(O)	(S)	(I)	(R)	(R)	(A)	(H)	(O)	(O)	(V)	(E)	(R)	(E)	(A)	(G)	(A)	(N)	(A)	(P)
L	(I)	(N)	(C)	(O)	(L)	(N)	(O)	(S)	(K)	(C)	(A)	(J)	(E)	(F)	(F)	(E)	(R)	(S)	(O)	(N)	(E)	(R)	(U)	(B)	(N)	(A)	(V)	(A)	(D)	(A)	(M)	(S)	(E)	(Y)	(A)	(H)	(S)	(U)	(B)	(F)	(O)	
K	(E)	(N)	(N)	(E)	(D)	(Y)	(E)	(L)	(N)	(I)	(K)	(C)	(M)	(O)	(N)	(R)	(O)	(E)	(J)	(O)	(H)	(N)	(S)	(O)	(N)	(O)	(X)	(I)	(N)	(A)	(N)	(A)	(H)	(C)	(U)	(B)	(R)	(E)	(T)	(R)	(A)	(C)
F	(I)	(L)	(L)	(M)	(O)	(R)	(E)	(L)	(Y)	(T)	(A)	(Y)	(L)	(O)	(R)	(O)	(O)	(S)	(E)	(V)	(E)	(L)	(T)	(R)	(U)	(M)	(A)	(N)	(O)	(T)	(N)	(I)	(L)	(C)	(O)	(O)	(L)	(I)	(D)	(G)	(E)	


[Solution (b) applies an idea by which Leonard Gordon was able to pack the names of presidents 1–42 with one less column. See A. Ross Eckler, *Word Ways* 27 (1994), 147.]

83. To pack w given words, use primary columns $\{P_{ij}, Ric, Cic, Bic, \#k \mid 1 \leq i, j \leq 9, 1 \leq k \leq w, c \in \{A, C, E, M, O, P, R, T, U\}\}$ and secondary columns $\{ij \mid 1 \leq i, j \leq 9\}$. There are 729 rows ' $P_{ij} Ric Cjc Bbc ij:c'$ ', where $b = 3[(i - 1)/3] + [j/3]$, together with a row ' $\#k i_1j_1:c_1 \dots i_lj_l:c_l$ ' for each placement of an l -letter word $c_1 \dots c_l$ into cells $(i_1, j_1), \dots, (i_l, j_l)$. Furthermore, it's important to *modify* step ?? of the algorithm so that the "best column" always has the form $\#k$, unless it has length ≤ 1 .

A brief run then establishes that COMPUTER and CORPORATE cannot both be packed. But all of the words *except* CORPORATE do fit together; the (unique) solution shown is found after only 7.3 megamems, most of which are needed simply to input the problem. [This exercise was inspired by a puzzle in *Sudoku Masterpieces* (2010) by Huang and Snyder.]



90. (The author designed these puzzles with the aid of exercises ??-??.)



author
 profile
 log-concave
 author
 UNIX
 extended hexadecimal digits
 hexadecimal notation, extended
 base placements
 sorted
 packed integers
 straight tromino
 secondary
 break symmetry
 author
 color controls
 Mercer
 Gardner

100. (a) To cover 2 of 4, we have 3 choices at the root, then 3 or 2 or 1 at the next level, hence (1, 3, 6) nodes at levels (0, 1, 2). To cover 5 of 7, there are (1, 3, 6, 10, 15, 21) nodes at levels (0, 1, ..., 5). Thus the profile with column 1 first is (1, 3, 6, 6 · 3, 6 · 6, 6 · 10, 6 · 15, 6 · 21). The other way is better: (1, 3, 6, 10, 15, 21, 21 · 3, 21 · 6).

(b) With column 1 first the profile is $(a_0, a_1, \dots, a_p, a_p a_1, \dots, a_p a_q)$, where $a_j = \binom{j+d}{d}$. We should branch on column 2 first because $a_{p+1} < a_p a_1, a_{p+2} < a_p a_2, \dots, a_q < a_p a_{q-p}, a_q a_1 < a_p a_{q-p+1}, \dots, a_q a_{p-1} < a_p a_{q-1}$. (These inequalities follow because the sequence $\langle a_j \rangle$ is strongly log-concave: It satisfies the condition $a_j^2 > a_{j-1} a_{j+1}$ for all $j \geq 1$. See exercise MPR-125.)

140. Let the given shape be specified as a set of integer pairs (x, y) . These pairs might simply be listed one by one in the input; but it's much more convenient to accept a more compact specification. For example, the utility program with which the author prepared the examples of this book was designed to accept UNIX-like specifications such as '[14-7]2 5[0-3]' for the seven pairs $\{(1, 2), (4, 2), (5, 2), (6, 2), (7, 2), (5, 0), (5, 1), (5, 3)\}$. The range $0 \leq x, y < 62$ has proved to be sufficient in almost all instances, with such integers encoded as single "extended hexadecimal digits" 0, 1, ..., 9, a, b, ..., z, A, B, ..., Z. The specification '[1-3][1-k]' is one way to define a 3×20 rectangle.

Similarly, each of the given polyominoes is specified by stating its piece name and a set T of typical positions that it might occupy. Such positions (x, y) are specified using the same conventions that were used for the shape; they needn't lie within that shape.

The program computes *base placements* by rotating and/or reflecting the elements of that set T . The first base placement is the shifted set $T_0 = T - (x_{\min}, y_{\min})$, whose coordinates are nonnegative and as small as possible. Then it repeatedly applies an elementary transformation, either $(x, y) \mapsto (y, x_{\max} - x)$ or $(x, y) \mapsto (y, x)$, to every existing base placement, until no further placements arise. (That process becomes easy when each base placement is represented as a sorted list of packed integers $(x \ll 16) + y$.) For example, the typical positions of the straight tromino might be specified as '1[1-3]'; it will have two base placements, $\{(0, 0), (0, 1), (0, 2)\}$ and $\{(0, 0), (1, 0), (2, 0)\}$.

After digesting the input specifications, the program defines the columns of the exact problem, which are the piece names together with the cells xy of the given shape.

Finally, it defines the rows: For each piece p and for each base placement T' of p , and for each offset (δ_x, δ_y) such that $T' + (\delta_x, \delta_y)$ lies fully within the given shape, there's a row that names the columns $\{p\} \cup \{(x + \delta_x, y + \delta_y) \mid (x, y) \in T'\}$.

(The output of this program is often edited by hand, to take account of special circumstances. For example, some columns may change from primary to secondary; some rows may be eliminated in order to break symmetry. The author's implementation also allows the specification of secondary columns with color controls, along with base placements that include such controls.)

148. RUSTY. [Leigh Mercer posed a similar question to Martin Gardner in 1960.]

150. As in the 3×20 example considered in the text, we can set up an exact cover problem with $12 + 60$ columns, and with rows for every potential placement of each

piece. This gives respectively (52, 292, 232, 240, 232, 120, 146, 120, 120, 30, 232, 120) rows for pieces (O, P, . . . , Z) in Conway's nomenclature, thus 1936 rows in all.

To reduce symmetry, we can insist that the X occurs in the upper left corner; then it contributes just 10 rows instead of 30. But some solutions are still counted twice, when X is centered in the middle row. To prevent this we can add a *secondary column* 's', and append 's' to the five rows that correspond to those centered appearances; we also append 's' to the 60 rows that correspond to placements where the Z is flipped over.

Without those changes, Algorithm D would use 9.76 $G\mu$ to find 4040 solutions; with them, it needs just 2.86 $G\mu$ to find 1010.

This approach to symmetry breaking in pentomino problems is due to Dana Scott [Technical Report No. 1 (Princeton University Dept. of Electrical Engineering, 10 June 1958)]. Another way to break symmetry would be to allow X anywhere, but to restrict the W to its 30 *unrotated* placements. That works almost as well: 2.87 $G\mu$.

151. There's a unique way to pack P, Q, R, U, X into a 5×5 square, and to pack the other seven into a 5×7 . (See below.) With independent reflections, together with rotation of the square, we obtain 16 of the 1010. There's also a unique way to pack P, R, U into a 5×3 and the others into a 5×9 (noticed by R. A. Fairbairn in 1967), yielding 8 more. And there's a unique way to pack O, Q, T, W, Y, Z into a 5×6 , plus two ways to pack the others, yielding another 16. (These paired 5×6 patterns were apparently first noticed by J. Pestiau; see answer 169.) Finally, the packings in the next exercise give us 264 decomposable 5×12 s altogether.

[Similarly, C. J. Bouwkamp discovered that S, V, T, Y pack uniquely into a 4×5 , while the other eight can be put into an 4×10 in five ways, thus accounting for 40 of the 368 distinct 4×15 s. See *Journal of Recreational Mathematics* **3** (1970), 125.]



152. Without symmetry reduction, 448 solutions are found in 1.21 $G\mu$. But we can restrict X to the upper left corner, flagging its placements with 's' when centered in the middle row or middle column (but not both). Again the 's' is appended to flipped Z's. Finally, when X is placed in dead center, we append *another* secondary column 'c', and append 'c' to the 90 rotated placements of W. This yields 112 solutions, after 0.34 $G\mu$.

Or we could leave X unhindered but curtail W to 1/4 of its placements. That's easier to do (although not *quite* as clever) and it finds those 112 in 0.42 $G\mu$.

Incidentally, there *aren't* actually any solutions with X in dead center.

154. The exact cover problem analogous to that in exercise 150 has $12 + 60$ columns and (56, 304, 248, 256, 248, 128, 1152, 128, 128, 32, 248, 128) rows. It finds 9356 solutions after 15.93 $G\mu$ of computation, without symmetry reduction. But if we insist that X be centered in the upper left quarter, by removing all but 8 of its placements, we get 2339 solutions after just 3.93 $G\mu$. (The alternative of restricting W's rotations is *not* as effective in this case: 5.43 $G\mu$.) These solutions were first enumerated by C. B. and Jenifer Haselgrove [Eureka: *The Archimedeans' Journal* **23** (1960), 16–18].

155. (a) Obviously only $k = 5$ is feasible. All such packings can be obtained by omitting all rows of the cover problem that straddle the "cut." That leaves 1507 of the original 2032 rows, and yields 16 solutions after 104 $M\mu$. (Those 16 boil down to just the two 5×6 decompositions that we already saw in answer 151.)

Conway
secondary column
Scott
break symmetry
Fairbairn
Pestiau
Bouwkamp
Haselgrove, Colin
Haselgrove, Jenifer

(b) Now we remove the 763 rows for placements that don't touch the boundary, and obtain just the two solutions below, after 100 $M\mu$. (This result was first noticed by Tony Potts, who posted it to Martin Gardner on 9 February 1960.)


(c) Now there are 1237 placements/rows; the *unique* solution is found after 83 $M\mu$.

(d) There are respectively (0, 9, 3, 47, 16, 8, 3, 1, 30, 22, 5, 11) solutions for pentominoes (O, P, Q, . . . , Z). (The I/O pentomino can be “framed” by the others in 11 ways; but all of those packings also have at least one other interior pentomino.)


(e) Despite many ways to cover all boundary cells with just seven pentominoes, none of them lead to an overall solution. Thus the minimum is eight; 207 of the 2339 solutions attain it. To find them we might as well generate and examine all 2339.

(f) The question is ambiguous: If we're willing to allow the X to touch unnamed pieces at a corner, but not at an edge, there are 25 solutions (8 of which happen to be answers to part (a)). In each of these solutions, X also touches the outer boundary. (The cover and frontispiece of Clarke's book show a packing in which X doesn't touch the boundary, but it *doesn't* solve this problem: There's an edge where X meets I, and there's a point where X meets P.) There also are two packings in which the edges of X touch only F, N, U, and the boundary, but not V.

On the other hand, there are just 6 solutions if we allow only F, N, U, V to touch X's corner points. One of them, shown below, has X touching the short side and seems to match the quotation best. These 6 solutions can be found in just 47 $M\mu$, by introducing 60 secondary columns as sort of an “upper level” to the board: All placements of X occupy the normal five lower-level cells, plus up to 16 upper-level cells that touch them; all placements of F, N, U, V are unchanged; all placements of the other seven pieces occupy both the lower and the upper level. This nicely forbids them from touching X.




157. Restrict X to five essentially different positions; if X is on the diagonal, also keep Z unflipped by using the second column 's' as in answer 152. There are respectively (16146, 24600, 23619, 60608, 25943) solutions, found in (19.8, 35.4, 27.3, 66.6, 34.5) $G\mu$.







In each case the tetromino can be placed anywhere that doesn't immediately cut off a region of one or two squares. [The twelve pentominoes first appeared in print when H. E. Dudeney published *The Canterbury Puzzles* in 1907. His puzzle #74, “The Broken Chessboard,” presented the first solution shown above, with pieces checkered in black and white. That parity restriction, with the further condition that no piece is turned over, would reduce the number of solutions to only 4, findable in 120 $M\mu$.]

The 60-element subsets of the chessboard that *can't* be packed with the pentominoes has been characterized by M. Reid in *J. Recreational Math.* **26** (1994), 153–154.




158. Yes, in seven essentially different ways. To remove symmetry, we can make the I vertical and put the X in the right half. (The pentominoes will have a total of $6 \times 2 + 5 \times 3 + 4 = 31$ black squares; therefore the tetromino *must* be )



159. These shapes can't be packed in a rectangle. But we can use the “supertile”  to make an infinite strip \cdots  \cdots . We can also tile the plane with a supertile like , or even use a generalized torus such as  (see exercise 7–137). That supertile was used in 2009 by George Sicherman to make tetromino wallpaper.


160. The 2339 solutions contain 563 that satisfy the “tatami” condition: No four pieces meet at any one point. Each of those 563 leads to a simple 12-vertex graph coloring problem; for example, the SAT methods of Section 7.2.2.2 typically need at most two or three kilomems to decide each case.

It turns out that exactly 94 are three-colorable, including the second solution to exercise 155(b). Here are the three for which W, X, Y, Z all have the same color:



162. Both shapes have 8-fold symmetry, so we can save a factor of nearly 8 by placing the X in (say) the north-northwest octant. If X thereby falls on the diagonal, or in the middle column, we can insist that the Z is not flipped, by introducing a secondary column ‘s’ as in answer 152. Furthermore, if X occurs in dead center — this is possible only for shape (i) — we use ‘c’ as in that answer to prohibit also any rotation of the W.

Thus find (a) 10 packings, in $3.5 G\mu$; (b) 7302 packings, in $353 G\mu$; for instance




It turns out that the monomino must appear in or next to a corner, as shown. [The first solution to shape (i) with monomino in the corner was sent to Martin Gardner by H. Hawkins in 1958. The first solution of the other type was published by J. A. Lindon in *Recreational Mathematics Magazine* #6 (December 1961), 22. Shape (ii) was introduced and solved much earlier, by G. Fuhendorf in *The Problemist: Fairy Chess Supplement* 2, 17 and 18 (April and June, 1936), problem 2410.]

163. (Notice that width 3 would be impossible, because every fault-free placement of the V needs width 4 or more.) We can set up an exact cover problem for a 4×19 rectangle in the usual way; but then we make cell $(x, y + 15)$ identical to $(3 - x, y)$ for $0 \leq x < 4$ and $0 \leq y < 5$, essentially making a half-twist when the pattern begins to wrap around. There are 60 symmetries, and care is needed to remove them properly. The easiest way is to put X into a fixed position, and allow W to rotate at most 90° .

symmetry
torus
torus, generalized
Sicherman
wallpaper
tatami
SAT
secondary column
Gardner
Hawkins
Lindon
Fuhendorf
symmetries

This exact cover problem has 850 solutions, 502 of which are fault-free. Here's one of the 29 strongly three-colorable ones, shown before and after its ends are joined:


three-colorable
color controls




164. It's also possible to wrap *two* cubes of size $\sqrt{5} \times \sqrt{5} \times \sqrt{5}$, as shown by F. Hansson; see *Fairy Chess Review* **6** (1947–1948), problems 7124 and 7591. A full discussion appears in *FGbook*, pages 685–689.



165. It's easy to set up an exact cover problem in which the cells touching the polyomino are primary columns, while other cells are secondary, and with rows restricted to placements that contain at least one primary column. Postprocessing can then remove spurious solutions that contain holes. Typical answers for (a) are




representing respectively (9, 2153, 37, 2, 17, 28, 18, 10, 9, 2, 4, 1) cases. For (b) they're



representing (16, 642, 1, 469, 551, 18, 24, 6, 4, 2, 162, 1). The total number of fences is respectively (3120, 1015033, 8660380, 284697, 1623023, 486, 150, 2914, 15707, 2, 456676, 2074), after weeding out respectively (0, 0, 16387236, 398495, 2503512, 665, 600, 11456, 0, 0, 449139, 5379) cases with holes. (See *MAA Focus* **36**, 3 (June/July 2016), 26; **36**, 4 (August/September 2016), 33.) Of course we can also make fences for one shape by using *other* shapes; for example, there's a beautiful way to fence a Z with 12 Ws, and a unique way to fence one pentomino with only *three* copies of another.

166. The small fences of answer 165(a) already meet this condition—except for the X, which has *no* tatami fence. The large fences for T and U in 165(b) are also good. But the other nine fences can no longer be as large:




[The tatami condition can be incorporated into the exact cover problem by using color controls: Introduce a secondary column for every potential edge between tiles, with values **t** and **f**. Also introduce a primary column *p* for every corner point; *p* will appear only in four rows '*p e:f*', one for each edge *e* that touches *p*. In every row for the

placement of a piece, include the columns ‘e:f’ for every edge *internal* to that piece, and ‘e:t’ for every edge at the *boundary* of that piece. Then every point will be next to a nonedge. However, for this exercise it’s best simply to apply the tatami condition directly to each ordinary solution, before postprocessing for hole-removal.]

168. This exercise, with 3×30 , 5×18 , 6×15 , and 9×10 rectangles, yields four increasingly difficult benchmarks for the exact cover problem, having respectively (46, 686628, 2562928, 10440433) solutions. Symmetry can be broken as in exercise 152. The 3×30 case was first resolved by J. Haselgrove; the 9×10 packings were first enumerated by A. Wassermann and P. Östergård, independently. [See *New Scientist* **12** (1962), 260–261; J. Meeus, *J. Recreational Math.* **6** (1973), 215–220; and *FGbook* pages 455, 468–469.] Algorithm D needs (.006, 5.234, 15.576, 63.386) teramems to find them. (I plan to give statistics for improved versions too; please stay tuned.)


169. Two solutions are now equivalent only when related by 180° rotation. Thus there are $2 \cdot 2339/64 = 73.09375$ solutions per problem, on average. The minimum (42) and maximum (136) solution counts occur for the cases



[In *U.S. Patent 2900190* (1959, filed 1956), J. Pestiau remarked that these 64 problems would give his pentomino puzzle “unlimited life and utility.”]

170. There are no ways to fill 2×20 ; 4×66 ways to fill 4×10 ; 4×84 ways to fill 5×8 . None of the solutions are symmetrical.


[See R. K. Guy, *Nabla* **7** (1960), 99–101.]



175. Most of the hexominoes will have three black cells and three white cells, in any “checkering” of the board. However, eleven of them (shown as darker gray in the illustration) will have a two-to-four split. Thus the total number of black cells will always be an even number between 94 and 116, inclusive. But a 210-cell rectangle always contains exactly 105 black cells. [See *The Problemist: Fairy Chess Supplement* **2**, 9–10 (1934–1935), 92, 104–105; *Fairy Chess Review* **3**, 4–5 (1937), problem 2622.]

Benjamin’s triangular shape, on the other hand, has $1+3+5+\dots+19 = 10^2 = 100$ cells of one parity and $\binom{20}{2} - 10^2 = 110$ of the other. It can be packed with the 35 hexominoes in a huge number of ways, probably not feasible to count exactly.

176. The parity considerations in answer 175 tell us that this is possible only for the “unbalanced” hexominoes, such as the one shown. And in fact, Algorithm D readily finds solutions for all eleven of those, too numerous to count. Here’s an example:



[See *Fairy Chess Review* **6** (April 1947) through **7** (June 1949), problems 7252, 7326, 7388, 7460, 7592, 7728, 7794, 7865, 7940, 7995, 8080. See also the similar problem 7092.]

177. Each castle must contain an odd number of the eleven unbalanced hexominoes (see answer 175). Thus we can begin by finding all sets of seven hexominoes that can be packed into a castle: This amounts to solving $\binom{11}{1} + \binom{11}{3} + \binom{11}{5} + \binom{11}{7} = 968$ exact cover problems, one for each potential choice of unbalanced elements. Each of those



gadget
benchmarks
Haselgrove
Wassermann
Östergård
Meeus
 180° rotation
central symmetry
Patent
Pestiau
Guy
checkering
parity
parity
exact cover


problems is fairly easy; the 24 balanced hexominoes provide secondary columns, while the castle cells and the chosen unbalanced elements are primary. In this way we obtain 39411 suitable sets of seven hexominoes, with only a moderate amount of computation.

That gives us *another* exact cover problem, having 35 columns and 39411 rows. This secondary problem turns out to have exactly 1201 solutions (found in just 115 G μ), each of which leads to at least one of the desired overall packings. Here's one:




In this example, two of the hexominoes in the rightmost castle can be flipped vertically; and of course the entire contents of each castle can independently be flipped horizontally. Thus we get 64 packings from this particular partition of the hexominoes (or maybe $64 \times 5!$, by permuting the castles), but only two of them are “really” distinct. Taking multiplicities into account, there are 1803 “really” distinct packings altogether.

[Frans Hansson found the first way to pack the hexominoes into five equal shapes, using  as the container; see *Fairy Chess Review* 8 (1952–1953), problem 9442. His container admits 123189 suitable sets of seven, and 9298602 partitions into five suitable sets instead of only 1201. Even more packings are possible with the container , which has 202289 suitable sets and 3767481163 partitions!]


In 1965, M. J. Povah packed all of the hexominoes into containers of shape , using *seven* sets of *five*; see *The Games and Puzzles Journal* 2 (1996), 206.

178. By exercise 175, m must be odd, and less than 35. F. Hansson posed this question in *Fairy Chess Review* 7 (1950), problem 8556. He gave a solution for $m = 19$,




and claimed without proof that 19 is optimum. The 13 dark gray hexominoes in this diagram cannot be placed in either “arm”; so they must go in the center. (Medium gray indicates pieces that have parity restrictions in the arms.) Thus we cannot have $m \geq 25$.

When $m = 23$, there are 39 ways to place all of the hard hexominoes, such as




However, none of these is completable with the other 22; hence $m \leq 21$.

When $m = 21$, the hard hexominoes can be placed in 791792 ways, without creating a region whose size isn't a multiple of 6 and without creating more than one region that matches a particular hexomino. Those 791792 ways have 69507 essentially distinct “footprints” of occupied cells, and the vast majority of those footprints appear to be impossible to fill. But in 2016, George Sicherman found the remarkable packing



which not only solves $m = 21$, it yields solutions for $m = 19, 17, 15, 11, 9, 7, 5$, and 3 by simple modifications. Sicherman also found separate solutions for $m = 13$ and $m = 1$.

179. Stead's original solution makes a very pleasant three-colored design:



[See *Fairy Chess Review* **9** (1954), 2–4; also *FGbook*, pages 659–662.]

This problem is best solved via the techniques of dynamic programming (Section 7.7), *not* with Algorithm D, because numerous subproblems are equivalent.

181. Make rows for the pentominoes in cells xy for $0 \leq x < 8$, $0 \leq y < 10$ as in exercise 140, and also for the tetrominoes in cells xy for $1 \leq x < 7$, $1 \leq y < 9$. In the latter rows include also columns $xy':0$ for all cells xy in the tetromino, as well as $xy':1$ for all other cells xy touching the tetromino, where the columns xy' for $0 \leq x < 8$ and $0 \leq y < 10$ are secondary. We can also assume that the center of the X pentomino lies in the upper left corner. There are 168 solutions, found after 1.5 Tμ of computation. (Another way to keep the tetrominoes from touching would be to introduce secondary columns for the *vertices* of the grid. Such columns are more difficult to implement, however, because they behave differently under the rotations of answer 140.)

[Many problems that involve placing the tetrominoes and pentominoes together in a rectangle were explored by H. D. Benjamin and others in the *Fairy Chess Review*, beginning already with its predecessor *The Problemist: Fairy Chess Supplement* (1936), problem 2171. But this particular question seems to have been raised first by Michael Keller in *World Game Review* **9**, (1989), xx.]

182. At present, not a single solution to this puzzle is known, although intuition suggests that enormously many of them ought to be possible. P. J. Torbijn and J. Meeus [*J. Recreational Mathematics* **32** (2003), 78–79] have exhibited solutions for rectangles of sizes 6×45 , 9×30 , 10×27 , and 15×18 .

198. (a) Represent the tree as a sequence $a_0 a_1 \dots a_{2n+1}$ of nested parentheses; then $a_1 \dots a_{2n}$ will represent the corresponding root-deleted forest, as in Algorithm 7.2.1.6P. The left boundary of the corresponding paralomino is obtained by mapping each '(' into N or E, according as it is immediately followed by '(' or ')'. The right boundary, similarly, maps each ')' into N or E according as it is immediately preceded by ')' or '(' . For example, the paralomino for forest 7.2.1.6–(2) is shown below with part (d).

(b) This series $wxy + w^2(xy^2 + x^2y) + w^3(xy^3 + 2x^2y^2 + x^3y) + \dots$ can be written $wxyH(w, wx, wy)$, where $H(w, x, y) = 1/(1 - x - y - G(w, x, y))$ generates a sequence of “atoms” corresponding to places x, y, G where the juxtaposed boundary paths have the respective forms $\begin{smallmatrix} E \\ E \end{smallmatrix}, \begin{smallmatrix} N \\ N \end{smallmatrix}$, or $\begin{smallmatrix} N \\ E \end{smallmatrix}(\text{inner})\begin{smallmatrix} E \\ N \end{smallmatrix}$. The area is thereby computed by diagonals between corresponding boundary points. (In the example from (a), the area is $1+1+1+1+2+2+2+2+2+2+2+2+2+2+1+1$; there's an “outer” G , whose H is $wxyGy$, and an “inner” G , whose H is $wxyGy$.) Thus we can write G as a continued fraction,

$G(w, x, y) = wxy / (1 - x - y - wxy / (1 - wx - wy - w^3xy / (1 - w^2x - w^2y - w^5xy / (\dots))))$.

[A completely different form is also possible, namely $G(w, x, y) = x \frac{J_1(w, x, y)}{J_0(w, x, y)}$, where

$$J_0(w, x, y) = \sum_{n=0}^{\infty} \frac{(-1)^n y^n w^{n(n+1)/2}}{(1-w)(1-w^2) \dots (1-w^n)(1-xw)(1-xw^2) \dots (1-xw^n)};$$

$$J_1(w, x, y) = \sum_{n=1}^{\infty} \frac{(-1)^{n-1} y^n w^{n(n+1)/2}}{(1-w)(1-w^2) \dots (1-w^{n-1})(1-xw)(1-xw^2) \dots (1-xw^n)}.$$


This form, derived via *horizontal* slices, disguises the symmetry between x and y .]

strongly three-colorable
dynamic programming
Benjamin
Keller
Torbijn
Meeus
nested parentheses
forest
continued fraction
Bessel functions, gen'lized

(c) Let $G(w, z) = G(w, z, z)$. We want $[z^n] G'(1, z)$, where differentiation is with respect to the first parameter. From the formulas in (b) we know that $G(1, z) = z(C(z) - 1)$, where $C(z) = (1 - \sqrt{1-4z})/(2z)$ generates the Catalan numbers. Partial derivatives $\partial/\partial w$ and $\partial/\partial z$ then give $G'(1, z) = z^2/(1-4z)$ and $G_l(1, z) = 1/\sqrt{1-4z} - 1$.

(d) This problem has four symmetries, because we can reflect about either diagonal. When $n = 5$, Algorithm D finds 4×801 solutions, of which 4×129 satisfy the tatami condition, and 4×16 are strongly three-colorable. (The tatami condition is easily enforced via secondary columns in this case, because we need only stipulate that the upper right corner of one parallomino doesn't match the lower left corner of another.) When $n = 6$ there are oodles and oodles of solutions. All of the trees/parallominoes thereby appear together in an attractive compact pattern.

Catalan numbers
tatami
strongly three-colorable
secondary columns
geek art
Klarner
Rivest
Bender
Goulden
Jackson
Delest
Viennot
Woan
Shapiro
Rogers
Flajolet
Sedgewick
pendant vertex: of degree 1
diameter




[References: D. A. Klarner and R. L. Rivest, *Discrete Math.* **8** (1974), 31–40; E. A. Bender, *Discrete Math.* **8** (1974), 219–226; I. P. Goulden and D. M. Jackson, *Combinatorial Enumeration* (New York: Wiley, 1983), exercise 5.5.2; M.-P. Delest and G. Viennot *Theoretical Comp. Sci.* **34** (1984), 169–206; W.-J. Woan, L. Shapiro, and D. G. Rogers, *AMM* **104** (1997), 926–931; P. Flajolet and R. Sedgewick, *Analytic Combinatorics* (Cambridge Univ. Press, 2009), 660–662.]

200. The same ideas apply, but with three coordinates instead of two, and with the elementary transformations $(x, y, z) \mapsto (y, x_{\max} - x, z)$, $(x, y, z) \mapsto (y, z, x)$.

Pieces (1, 2, . . . , 7) have respectively (12, 24, 12, 12, 12, 12, 8) base placements, leading to $144 + 144 + 72 + 72 + 96 + 96 + 64$ rows for the $3 \times 3 \times 3$ problem.

202. It's tempting, but wrong, to try to compute the Somap by considering only the 240 solutions that have the tee in a fixed position and the claw restricted; the pairwise semidistances between these special solutions will miss many of the actual adjacencies. To decide if $u \sim v$, one must compare u to the 48 solutions equivalent to v .

(a) The strong Somap has vertex degrees $7^1 6^7 5^{19} 4^{31} 3^{59} 2^{63} 1^{45} 0^{15}$; so an “average” solution has $(1 \cdot 7 + 7 \cdot 6 + \dots + 15 \cdot 0)/240 \approx 2.57$ strong neighbors. (The unique vertex of degree 7 has the level-by-level structure $\frac{333}{462} \frac{114}{652} \frac{473}{662}$ from bottom to top.)

The full Somap has vertex degrees $21^2 18^1 16^9 15^{13} 14^{10} 13^{16} 12^{17} 11^{12} 10^{16} 9^{28} 8^{26} 7^{25} 6^{26} 5^{16} 4^{17} 3^3 2^1 1^1 0^1$, giving an average degree ≈ 9.14 . (Its unique isolated vertex is $\frac{333}{462} \frac{465}{762} \frac{715}{772}$, and its only pendant vertex is $\frac{333}{222} \frac{755}{462} \frac{771}{466}$. Two other noteworthy solutions, $\frac{333}{466} \frac{412}{762} \frac{115}{772}$ and $\frac{333}{466} \frac{419}{765} \frac{112}{772}$, are the only ones that contain the two-piece substructure .)

(b) The Somap has just two components, namely the isolated vertex and the 239 others. The latter has just three bicomponents, namely the pendant vertex, its neighbor, and the 237 others. Its diameter is 8 (or 21, if we use the edge lengths 2 and 3).

The strong Somap has a much sparser and more intricate structure. Besides the 15 isolated vertices, there are 25 components of sizes $\{8 \times 2, 6 \times 3, 4, 3 \times 5, 2 \times 6, 7, 8, 11, 16, 118\}$. Using the algorithm of Section 7.4.1, the large component breaks down into nine bicomponents (one of size 2, seven of size 1, the other of size 109); the 16-vertex component breaks into seven; and so on, totalling 58 bicomponents altogether.

[The Somap was first constructed by R. K. Guy, J. H. Conway, and M. J. T. Guy, without computer help. It appears on pages 910–913 of Berlekamp, Conway, and Guy’s *Winning Ways*, where all of the strong links are shown, and where enough other links are given to establish near-connectedness. Each vertex in that illustration has been given a code name; for example, the five special solutions mentioned in part (a) have code names B5f, R7d, LR7g, YR3a, and R3c, respectively.]

204. Let the cubic coordinates be $51z$, $41z$, $31z$, $32z$, $33z$, $23z$, $13z$, $14z$, $15z$, for $z \in \{1, 2, 3\}$. Replace matrix A of the exact cover problem by a simplified matrix A' having only columns $(1, 2, 3, 4, 5, 6, 7, S)$, where S is the sum of all columns xyz of A where xyz is odd. Any solution to A yields a solution to A' with column sums $(1, 1, 1, 1, 1, 1, 1, 10)$. But that’s impossible, because the rows of A' all have the forms ‘1 [S]’, ‘2 [S] [S]’, ‘3 [S] [S]’, ‘4 [S]’, ‘5 [S]’, ‘6 [S]’, ‘7 [S]’. [See the Martin Gardner reference in answer 213.]

205. (a) The solution counts, ignoring symmetry reduction, are: 4×5 corral (2), gorilla (2), smile (2), 3×6 corral (4), face (4), lobster (4), castle (6), bench (16), bed (24), doorway (28), piggybank (80), five-seat bench (104), piano (128), shift 2 (132), 4×4 coop (266), shift 1 (284), bathtub (316), shift 0 (408), grand piano (526), tower 4 (552), tower 3 (924), canal (1176), tower 2 (1266), couch (1438), tower 1 (1520), stepping stones (2718). So the 4×5 corral, gorilla, and smile are tied for hardest, while stepping stones are the easiest. (The bathtub, canal, bed, and doorway each have four symmetries; the couch, stepping stones, tower 4, shift 0, bench, 4×4 coop, castle, five-seat bench, piggybank, lobster, piano, gorilla, face, and smile each have two. To get the number of *essentially distinct* solutions, divide by the number of symmetries.)

(b) Notice that the canal, bed, and doorway appear also in (a), as does the dryer (which is the same as “stepping stones”). The solution counts are: W-wall (0), almost W-wall (12), bed (24), apartments 2 (28), doorway (28), clip (40), tunnel (52), zigzag wall 2 (52), zigzag wall 1 (92), underpass (132), chair (260), stile (328), fish (332), apartments 1 (488), goldfish (608), canal (1176), steps (2346), dryer (2718); hence “almost W-wall” is the hardest of the possible shapes. Notice that the dryer, chair, steps, and zigzag wall 2 each have two symmetries, while the others in Fig. 80(b) all have four. The $3 \times 3 \times 3$ cube, with its 48 symmetries, probably is the easiest possible shape to make from the Soma pieces.

[Piet Hein himself published the tower 1, shift 2, stile, and zigzag wall 1 in his original patent; he also included the bathtub, bed, canal, castle, chair, steps, stile, stepping stones, shift 1, five-seat bench, tunnel, W-wall, and both apartments in his booklet for Parker Brothers. Parker Brothers distributed four issues of *The SOMA*[®] *Addict* in 1970 and 1971, giving credit for new constructions to Noble Carlson (fish, lobster), Mrs. C. L. Hall (clip, underpass), Gerald Hill (towers 2–4), Craig Kenworthy (goldfish), John W. M. Morgan (cot, face, gorilla, smile), Rick Murray (grand piano), and Dan Smiley (doorway, zigzag wall 2). Sivy Farhi published a booklet called *Somacubes* in 1977, containing the solutions to more than one hundred Soma cube problems including the bench, the couch, and the piggybank.]

206. By eliminating symmetries, there are (a) 421 distinct cases with cubies omitted on both layers, and (b) 129 with cubies omitted on only one layer. All are possible, except in the one case where the omitted cubies disconnect a corner cell. The easiest of type (a) omits (111, 112, 311) and has 3599 solutions; the hardest omits (211, 222, 231) and has 45×2 solutions. The easiest of type (b) omits (111, 151, 311) and has 3050 solutions; the hardest omits (211, 221, 251) and has 45×2 solutions. (The two examples illustrated have 821×2 and 68×4 solutions. Early Soma solvers seem to have overlooked them!)


Guy
Conway
Guy
Berlekamp
Conway
Guy
Gardner
symmetries
Hein
Parker Brothers
Carlson
Hall
Hill
Kenworthy
Morgan
Murray
Smiley
Farhi
symmetries

207. (a) The 60 distinct cases are all quite easy. The easiest has 3497 solutions and uses (113, 123, 213) on the top level; the hardest has 268 solutions and uses (113, 223, 313).

(b) Sixteen of the 60 possibilities are disconnected. Three of the others are also impossible — namely those that omit $(12z, 24z, 32z)$ or $(21z, 22z, 23z)$ or $(21z, 22z, 24z)$. The easiest has 3554 solutions and omits $(11z, 12z, 34z)$; the hardest of the possibles has only 8 solutions and omits $(11z, 23z, 24z)$.

(The two examples illustrated have 132×2 and 270×2 solutions.)

208. All but 216 are realizable. Five cases have unique (1×2) solutions:



210. Every polycube has a minimum enclosing box for which it touches all six faces. If those box dimensions $a \times b \times c$ aren't too large, we can generate such polycubes uniformly at random in a simple way: First choose 27 of the abc possible cubies; try again if that choice doesn't touch all faces; otherwise try again if that choice isn't connected.

For example, when $a = b = c = 4$, about 99.98% of all choices will touch all faces, and about 0.1% of those will be connected. This means that about $.001 \binom{64}{27} \approx 8 \times 10^{14}$ of the 27-cubie polycubes have a $4 \times 4 \times 4$ bounding box. Of these, about 5.8% can be built with the seven Soma pieces.

But most of the relevant polycubes have a larger bounding box; and in such cases the chance of solvability goes down. For example, $\approx 6.2 \times 10^{18}$ cases have bounding box $4 \times 5 \times 5$; $\approx 3.3 \times 10^{18}$ cases have bounding box $3 \times 5 \times 7$; $\approx 1.5 \times 10^{17}$ cases have bounding box $2 \times 7 \times 7$; and only 1% or so of those cases are solvable.

Section 7.2.3 will discuss the enumeration of polycubes by their size.

212. Each interior position of the penthouse and pyramid that might or might not be occupied can be treated as a secondary column in the corresponding exact cover problem. We obtain 10×2 solutions for the staircase; $(223, 286) \times 8$ solutions for the penthouse with hole at the (bottom, middle); and 32×2 solutions for the pyramid, of which 2×2 have all three holes on the diagonal and 3×2 have no adjacent holes.

213. A full simulation of gravity would be quite complex, because pieces can be prevented from tipping with the help of their neighbors above and/or at their side. If we assume a reasonable coefficient of friction and an auxiliary weight at the top, it suffices to define stability by saying that a piece is stable if and only if at least one of its cubies is immediately above either the floor or a stable piece.

The given shapes can be packed in respectively 202×2 , 21×2 , 270×2 , 223×8 , and 122×2 ways, of which 202×2 , 8×2 , 53×2 , 1×8 , and 6×2 are stable. Going from the bottom level to the top, the layers $\begin{smallmatrix} 1 & 1 & 1 & 1 & 1 \\ 3 & 6 & 3 & 6 & 3 \\ 3 & 1 & 3 & 1 & 3 \end{smallmatrix}$ give a decently stable cot; a fragile vulture comes from $\begin{smallmatrix} 2 & 2 & 2 & 2 & 2 \\ 3 & 5 & 3 & 5 & 3 \\ 3 & 5 & 3 & 5 & 3 \end{smallmatrix}$; a delicate mushroom comes from $\begin{smallmatrix} 7 & 7 & 7 & 7 & 7 \\ 4 & 6 & 4 & 6 & 4 \\ 4 & 6 & 4 & 6 & 4 \end{smallmatrix}$; and a delicate cantilever from $\begin{smallmatrix} 2 & 2 & 2 & 2 & 2 \\ 2 & 2 & 2 & 2 & 2 \\ 5 & 7 & 5 & 7 & 5 \end{smallmatrix}$. The author's cherished set of Skjøde Skjern Soma pieces, made of rosewood and purchased in 1967, includes a small square base that nicely stabilizes both mushroom and cantilever. The vulture needs a book on top.


[The casserole and cot are due respectively to W. A. Kustes and J. W. M. Morgan. The mushroom, which is hollow, is the same as B. L. Schwartz's "penthouse," but turned upside down; John Conway noticed that it then has a unique stable solution. See Martin Gardner, *Knotted Doughnuts* (1986), Chapter 3.]

secondary column
author
Skjøde Skjern
Knutsen, see Skjøde Skjern
Kustes
Morgan
Schwartz
Conway
Gardner

214. Infinitely many cubies lie behind a wall; but it suffices to consider only the hidden ones whose distance is at most $27 - v$ from the v visible ones. For example, if the W-wall has coordinates as in answer 204, we have $v = 25$ and the two invisible cubies are $\{332, 331\}$. We're allowed to use any of $\{241, 242, 251, 252, 331, 332, 421, 422, 521, 522\}$ at distance 1, and $\{341, 342, 351, 352, 431, 432, 531, 532, 621, 622\}$ at distance 2. (The stated projection doesn't have left-right symmetry.) The X-wall is similar, but it has $v = 19$ and potentially $(9, 7, 6, 3, 3, 2, 1)$ hidden cubies at distances 1 to 7 (omitting cases like 450, which is invisible at distance 2 but "below ground").

gravitationally stable
 Francillon
 Hoffmann
 Mikusinski's Cube
 Steinhaus
 pentacubes
 Reid
 Sicherman
 Holy Grail
 Shindo
 Neo Diabolical Cube

Using secondary columns for the optional cubies, we must examine each solution to the exact cover problem and reject those that are disconnected or violate the gravity constraint of exercise 213. Those ground rules yield 282 solutions for the W-wall, 612 for the X-wall, and a whopping 1,130,634 for the cube itself. (These solutions fill respectively 33, 275, and 13842 different sets of cubies.) Here are examples of some of the more exotic shapes that are possible, as seen from behind and below:



There also are ten surprising ways to make the cube façade if we allow hidden "underground" cubies: The remarkable construction $\begin{matrix} \dots & 446 & 744 & 333 \\ \dots & 56 & 209 & 276 \\ \dots & 2 & 9 & 211 \end{matrix}$ raises the entire cube one level *above* the floor, and is gravitationally stable, by exercise 213's criteria! Unfortunately, though, it falls apart — even with a heavy book on top.

[The false-front idea was pioneered by Jean Paul Francillon, whose construction of a fake W-wall was announced in *The SOMA® Addict* 2, 1 (spring 1971).]

215. (a) Each of 13 solutions occurs in 48 equivalent arrangements. To remove the symmetry, place piece 7 horizontally, either (i) at the bottom or (ii) in the middle. In case (ii), add a secondary 's' column as in answer 150, and append 's' also to all placements of piece 6 that touch the bottom more than the top. Run time: 400 Kμ.

[This puzzle was number 39 in *Hoffmann's Puzzles Old and New* (1893). Another $3 \times 3 \times 3$ polycube dissection of historical importance, "Mikusinski's Cube," was described by Hugo Steinhaus in the 2nd edition of his *Mathematical Snapshots* (1950). That one consists of the ell and the two twist pieces of the Soma cube, plus the pentacubes B, C, and f of exercise 220; it has 24 symmetries and just two solutions.]

(b) Yes: Michael Reid, circa 1995, found the remarkable set





which also makes $9 \times 3 \times 1$ uniquely(!). George Sicherman carried out an exhaustive analysis of all relevant flat polyominoes in 2016, finding exactly 320 sets that are unique for $3 \times 3 \times 3$, of which 19 are unique also for $9 \times 3 \times 1$. In fact, one of those 19,



is the long-sought "Holy Grail" of $3 \times 3 \times 3$ cube decompositions: Its pieces not only have flatness and double uniqueness, they are nested (!!). There's also Yoshiya Shindo's



known as the "Neo Diabolical Cube" (1995); notice that it has 24 symmetries, not 48.


217. The straight tetracube  and the square tetracube , together with the size-4 Soma pieces in (30), make a complete set.

We can fix the tee's position in the twin towers, saving a factor of 32; and each of the resulting 40 solutions has just one twist with the tee. Hence there are five inequivalent solutions, and 5×256 altogether.

The double claw has 63×6 solutions. But the cannon, with 1×4 solutions, can be formed in essentially only one way. (*Hint:* Both twists are in the barrel.)

There are no solutions to 'up 3'. But 'up 4' and 'up 5' each have 218×8 solutions (related by turning them upside down). Gravitationally, four of those 218 are stable for 'up 5'; the stable solution for 'up 4' is unique, and unrelated to those four.

References: Jean Meeus, *Journal of Recreational Mathematics* **6** (1973), 257–265; Nob Yoshigahara, *Puzzle World No. 1* (San Jose: Ishi Press International, 1992), 36–38.

218. All but 48 are realizable. The unique “hardest” realizable case, , has 2×2 solutions. The “easiest” case is the $2 \times 4 \times 4$ cuboid, with $11120 = 695 \times 16$ solutions.

220. (a) A, B, C, D, E, F, a, b, c, d, e, f, j, k, l, . . . , z. (It's a little hard to see why reflection doesn't change piece 'l'. In fact, S. S. Besley once patented the pentacubes under the impression that there were 30 different kinds! See *U.S. Patent 3065970* (1962), where Figs. 22 and 23 illustrate the same piece in slight disguise.)

Historical notes: R. J. French, in *Fairy Chess Review* **4** (1940), problem 3930, was first to show that there are 23 different pentacube shapes, if mirror images are considered to be identical. The full count of 29 was established somewhat later by F. Hansson and others [*Fairy Chess Review* **6** (1948), 141–142]; Hansson also counted the $35 + 77 = 112$ mirror-inequivalent hexacubes. Complete counts of hexacubes (166) and heptacubes (1023) were first established soon afterwards by J. Niemann, A. W. Baillie, and R. J. French [*Fairy Chess Review* **7** (1948), 8, 16, 48].

(b) The cuboids $1 \times 3 \times 20$, $1 \times 4 \times 15$, $1 \times 5 \times 12$, and $1 \times 6 \times 10$ have of course already been considered. The $2 \times 3 \times 10$ and $2 \times 5 \times 6$ cuboids can be handled by restricting X to the bottom upper left, and sometimes also restricting Z, as in answers 150 and 152; we obtain 12 solutions (in $350 M\mu$) and 264 solutions (in $2.5 G\mu$), respectively.

The $3 \times 4 \times 5$ cuboid is more difficult. Without symmetry-breaking, we obtain 3940×8 solutions in about $200 G\mu$. To do better, notice that X can appear in 11 essentially different positions: $(1+1^*)(1+1^*)$ in a 4×5 plane, 2^*+2^{**} in a 3×5 plane, and 2^*+1^{**} in a 3×4 plane, where '*' denotes a case where symmetry needs to be broken down further because X is fixed by some symmetry. With 11 separate runs we can find $(923 + 558/2 + 402/2 + 376/4) + (1268/2 + 656/2 + 420/4 + 752/4) + (1480/2 + 720/2 + 352/4) = 3940$ solutions, in $4.9 + 3.3 + 3.1 + 2.4 + \dots + 2.1 \approx 50 G\mu$.

[The fact that solid pentominoes will fill these cuboids was first demonstrated by D. Nixon and F. Hansson, *Fairy Chess Review* **6** (1948), problem 7560 and page 142. Exact enumeration was first performed by C. J. Bouwkamp in 1967; see *J. Combinatorial Theory* **7** (1969), 278–280, and *Indagationes Math.* **81** (1978), 177–186.]

(c) Almost *any* subset of 25 pentacubes can probably do the job. But a particularly nice one is obtained if we simply omit o, q, s, and y, namely those that don't fit in a $3 \times 3 \times 3$ box. R. K. Guy proposed this subset in *Nabla* **7** (1960), 150, although he wasn't able to pack a $5 \times 5 \times 5$ at that time. The same idea occurred independently to J. E. Dorie, who trademarked the name “Dorian cube” [*U.S. Trademark 1,041,392* (1976)].

An amusing way to form such a cube is to make 5-level prisms in the shapes of the P, Q, R, U, and X pentominoes, using pieces {a, e, j, m, w}, {f, k, l, p, r}, {A, d, D, E, n},

Meeus
Yoshigahara
Besley
Patent
French
Hansson
hexacubes
heptacubes
Niemann
Baillie
symmetry-breaking
Nixon
Hansson
Bouwkamp
Guy
Dorie
Dorian cube
pentominoes

{c, C, F, u, v}, {b, B, t, x, z}; then use the packing in answer 151(!). This solution can be found with six very short runs of Algorithm D, taking only 300 megamems overall.

Sillke
partition

Another nice way, due to Torsten Sillke, is more symmetrical: There are 70,486 ways to partition the pieces into five sets of five that allow us to build an X-prism in the center (with piece x on top), surrounded by four P-prisms.

One can also assemble a Dorian cube from five cuboids, using one $1 \times 3 \times 5$, one $2 \times 2 \times 5$, and three $2 \times 3 \times 5$ s. Indeed, there are zillions more ways, too many to count.

221. (a) Make an exact cover problem in which a and A, b and B, . . . , f and F are required to be in symmetrical position; there are respectively (86, 112, 172, 112, 52, 26) placements for such 10-cubic “super-pieces.” Furthermore, the author decided to force piece m to be in the middle of the top wall. Solutions were found immediately! So piece x was placed in the exact center, as an additional desirable constraint. Then there were exactly 20 solutions; the one below has also n, o, and u in mirror-symmetrical locations.

author
geek art
Sillke
Künzell
Farhi

(b) The super-pieces now have (59, 84, 120, 82, 42, 20) placements; the author also optimistically forced j, k, and m to be symmetrical about the diagonal, with m in the northwest corner. A long and apparently fruitless computation (34.3 teramems) ensued; but — hurrah — two closely related solutions were discovered at the last minute.

(c) This computation, due to Torsten Sillke [see *Cubism For Fun* 27 (1991), 15], goes much faster: The quarter-of-a-box shown here can be packed with seven non-x pentacubes in 55356 ways, found in 1.3 Gμ. As in answer 177, this yields a new exact cover problem, with 33412 different rows. Then 11.8 Gμ more computation discovers seven suitable partitions into four sets of seven, one of which is illustrated here.



l	l	l	q	q	q	q
l	o	o	o	o	o	q
f	f	u	u	u	F	F
D	f	u	m	u	F	d

m	o	o	o	o	o	s
m	m	x	q	q	q	q
m	x	x	x	b	b	b
r	n	x	e	e	b	a

v	E	z	z	t	A	A
E	E	z	s	t	t	t
E	z	z	s	t	F	F
f	f	s	s	a	F	k

l	l	f	D	D	D	m	m	d	d	d	F	q	q	
l	f	f	C	C	D	D	m	d	d	c	C	F	F	r
v	v	v	B	C	C	C	x	c	c	c	b	r	r	r
v	w	B	B	B	x	x	x	b	b	b	b	r	r	z
v	w	w	A	A	A	A	x	a	a	a	a	z	z	z
k	k	w	E	E	A	n	a	e	e	y	z	j	j	j
k	k	s	s	s	E	E	n	e	e	y	y	y	j	j
s	E	n	n	n	e	e	y	y	y	j	j	j	j	j
s	s	p	t	t	t	t	y	y	y	j	j	j	j	j
k	s	p	p	t	j	j	j	j	j	j	j	j	j	j
k	k	p	p	t	j	j	j	j	j	j	j	j	j	j

m	m	m	r	r	n	e	e	e	a	a	a	a	b	q	s
t	m	r	r	r	n	n	n	e	a	D	D	D	b	q	s
t	t	t	r	E	p	p	p	v	F	C	C	D	D	s	s
t	w	w	E	E	p	p	v	F	C	C	C	s	s	z	z
w	w	B	E	A	A	v	v	v	F	F	F	z	z	z	z
w	l	B	B	A	d	c	f	f	k	k	F	z	j	j	j
l	l	B	B	A	d	c	f	k	k	k	u	u	j	j	j
A	d	d	c	f	f	k	k	u	u	j	j	j	j	j	j
B	B	d	c	u	u	u	u	u	u	u	u	u	u	u	u
l	l	l	l	l	l	l	l	l	l	l	l	l	l	l	l
l	y	y	y	y	j	j	j	j	j	j	j	j	j	j	j

v	E	E	f	B	f	s	w	a	k	k	k	F	t	A
v	E	f	f	B	B	w	w	a	k	k	D	F	F	A
v	v	v	B	B	w	w	x	a	a	D	D	y	y	A
u	u	u	j	j	j	x	x	x	D	D	y	y	y	y
u	p	u	j	m	j	C	x	c	d	o	o	o	o	o
p	p	e	m	m	C	b	c	d	d	n	n	n	n	l
p	p	e	e	m	C	C	b	c	c	d	d	n	n	l
e	C	q	b	b	c	d	d	n	n	l	l	l	l	l
e	e	q	b	r	n	n	n	n	n	n	n	n	n	n
p	e	q	r	r	r	r	r	r	r	r	r	r	r	r
p	q	q	r	r	r	r	r	r	r	r	r	r	r	r

(a)
(b)
(c)

222. As in previous exercises, the key is to reduce the search space drastically, by asking for solutions of a special form. (Such solutions are likely to exist, because pentacubes are so versatile.) Here we can break the given shape into four pieces: Three modules of size $3^3 + 2^3$ to be packed with seven pentacubes, and one of size $4^3 - 3 \cdot 2^3$ to be packed with eight pentacubes. The smaller problem has 13,587,963 solutions, found with 2.5 Tμ of computation; these reduce to 737,695 distinct sets of seven pentacubes. The larger problem has 13,800 solutions, found with 400 Mμ and reduced to 2075 sets of eight. Exactly covering those sets yields 1,132,127,589 suitable partitions; the first one found, {a, A, b, c, j, q, t, y}, {B, C, d, D, e, k, o}, {E, f, l, n, r, v, x}, {F, m, p, s, u, w, z}, works fine. (We need only one partition, so we needn't have computed more than a thousand or so solutions to the smaller problem.)



Pentacubes galore: Since the early 1970s, Ekkehard Künzell and Sivy Farhi have independently published booklets that contain hundreds of solved pentacube problems.

999. . . .

INDEX AND GLOSSARY

Pope
Homer
WHEATLEY

There is a curious poetical index to the Iliad in Pope's Homer, referring to all the places in which similes are used.

— HENRY B. WHEATLEY, *What is an Index?* (1878)

When an index entry refers to a page containing a relevant exercise, see also the *answer* to that exercise for further information. An answer page is not indexed here unless it refers to a topic not included in the statement of the exercise.

Barris, Harry, 1.

DIMACS: DIMACS Series in Discrete Mathematics and Theoretical Computer Science, inaugurated in 1990.

Fields, Dorothy, 1.

MPR: Mathematical Preliminaries Redux, v.

Short, Robert Allen, iii.

Nothing else is indexed yet (sorry).

Preliminary notes for indexing appear in the upper right corner of most pages.

If I've mentioned somebody's name and forgotten to make such an index note, it's an error (worth \$2.56).