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 5B

INTRODUCTION TO BACKTRACKING

DONALD E. KNUTH *Stanford University*

ADDISON–WESLEY

September 11, 2019

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.

See also `http://www-cs-faculty.stanford.edu/~knuth/programs.html` for various experimental programs that I wrote while writing this material (and some data files).

Zeroth printing (revision 2), 11 September 2019

CARROLL
Internet

PREFACE

> *Begin at the* beginning, *and do not allow yourself to gratify a mere idle curiosity by dipping into the book, here and there. This would very likely lead to your throwing it aside, with the remark "This is* much *too hard for me!," and thus losing the chance of adding a very* large *item to your stock of mental delights.*
>
> — LEWIS CARROLL, in *Symbolic Logic* (1896)

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 the opening pages of Section 7.2.2 of *The Art of Computer Programming*, entitled "Backtrack programming." The preceding section, 7.2.1, was about "Generating basic combinatorial patterns" — namely tuples, permutations, combinations, partitions, and trees. Now it's time to consider the non-basic patterns, the ones that have a much less uniform structure. For these we generally need to make tentative choices and then we need to back up when those choices need revision. Several subsections (7.2.2.1, 7.2.2.2, etc.) will follow this introductory material.

* * *

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.

stamping
Knuth

Please look, for example, at the exercises that I've classed as research problems (rated with difficulty level 46 or higher), namely exercises 15, . . . ; I've also implicitly mentioned or posed additional unsolved questions in the answers to exercises 7, 9, 46, 48, ... 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 29, 32, 36(b), 38, 47, 52, 53, 61, 62, 63, 67, 75–78, ... Furthermore I've credited exercises . . . to unpublished work of . . . . Have any of those results ever appeared in print, to your knowledge?

I've got a historical question too: Have you any idea who originated the idea of "stamping" in data structures? (See 7.2.2–(26). This concept is quite different from the so-called time stamps in persistent data structures, and quite different from the so-called time stamps in depth-first search algorithms, and quite different from the so-called time stamps in cryptology, although many programmers do use the name "time stamp" for those kinds of stamp.) It's a technique that I've seen often, in programs that have come to my attention during recent decades, but I wonder if it ever appeared in a book or paper that was published before, say, 1980.

* * *

Special thanks are due to . . . for their detailed comments on my early attempts at exposition, as well as to 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:–)

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 D. E. K.
19 June 2019

MPR
English words
Internet

Part of the Preface to Volume 4B

During the years that I've been preparing Volume 4, I've often run across basic techniques of probability theory that I would have put into Section 1.2 of Volume 1 if I'd been clairvoyant enough to anticipate them in the 1960s. Finally I realized that I ought to collect most of them together in one place, near the beginning of Volume 4B, because the story of these developments is too interesting to be broken up into little pieces scattered here and there.

Therefore this volume begins with a special section entitled "Mathematical Preliminaries Redux," and future sections use the abbreviation 'MPR' to refer to its equations and its exercises.

* * *

Several exercises involve the lists of English words that I've used in preparing examples. You'll need the data from

`http://www-cs-faculty.stanford.edu/~knuth/wordlists.tgz`

if you have the courage to work those exercises.

MATHEMATICAL PRELIMINARIES REDUX

MANY PARTS of this book deal with *discrete probabilities*, namely with a finite or countably infinite set Ω of atomic events ω, each of which has a given probability $\Pr(\omega)$, where

$$0 \le \Pr(\omega) \le 1 \qquad \text{and} \qquad \sum_{\omega \in \Omega} \Pr(\omega) = 1. \tag{1}$$

...

For the complete text of the special MPR section, please see Pre-Fascicle 5a. Incidentally, Section 7.2.2 intentionally begins on a left-hand page, and its illustrations are numbered beginning with Fig. 68, because Section 7.2.1 ended on a right-hand page and its final illustration was Fig. 67. The editor has decided to treat Chapter 7 as a single unit, even though it will be split across several physical volumes.

Nowhere to go but out,
Nowhere to come but back.
— BEN KING, in *The Sum of Life* (c. 1893)

Lewis back-tracked the original route up the Missouri.
— LEWIS R. FREEMAN, in *National Geographic Magazine* (1928)

When you come to one legal road that's blocked,
you back up and try another.
— PERRY MASON, in *The Case of the Black-Eyed Blonde* (1944)

KING
Lewis
FREEMAN
MASON
Gardner ES
backtrack
Walker
domain
cutoff
properties: logical propositions (relations)
$P_0()$
lexicographically

7.2.2. Backtrack Programming

Now that we know how to generate simple combinatorial patterns such as tuples, permutations, combinations, partitions, and trees, we're ready to tackle more exotic patterns that have subtler and less uniform structure. Instances of almost *any* desired pattern can be generated systematically, at least in principle, if we organize the search carefully. Such a method was christened "backtrack" by R. J. Walker in the 1950s, because it is basically a way to examine all fruitful possibilities while exiting gracefully from situations that have been fully explored.

Most of the patterns we shall deal with can be cast in a simple, general framework: We seek all sequences $x_1x_2\ldots x_n$ for which some property $P_n(x_1,x_2,\ldots,x_n)$ holds, where each item x_k belongs to some given domain D_k of integers. The backtrack method, in its most elementary form, involves the invention of intermediate "cutoff" properties $P_l(x_1,\ldots,x_l)$ for $1\le l<n$, such that

$$P_l(x_1,\ldots,x_l) \text{ is true whenever } P_{l+1}(x_1,\ldots,x_{l+1}) \text{ is true;} \qquad (1)$$

$$P_l(x_1,\ldots,x_l) \text{ is fairly easy to test, if } P_{l-1}(x_1,\ldots,x_{l-1}) \text{ holds.} \qquad (2)$$

(We assume that $P_0()$ is always true. Exercise 1 shows that all of the basic patterns studied in Section 7.2.1 can easily be formulated in terms of domains D_k and cutoff properties P_l.) Then we can proceed lexicographically as follows:

Algorithm B (*Basic backtrack*). Given domains D_k and properties P_l as above, this algorithm visits all sequences $x_1x_2\ldots x_n$ that satisfy $P_n(x_1,x_2,\ldots,x_n)$.

B1. [Initialize.] Set $l\leftarrow 1$, and initialize the data structures needed later.

B2. [Enter level l.] (Now $P_{l-1}(x_1,\ldots,x_{l-1})$ holds.) If $l>n$, visit $x_1x_2\ldots x_n$ and go to B5. Otherwise set $x_l\leftarrow \min D_l$, the smallest element of D_l.

B3. [Try x_l.] If $P_l(x_1,\ldots,x_l)$ holds, update the data structures to facilitate testing P_{l+1}, set $l\leftarrow l+1$, and go to B2.

B4. [Try again.] If $x_l\ne \max D_l$, set x_l to the next larger element of D_l and return to B3.

B5. [Backtrack.] Set $l\leftarrow l-1$. If $l>0$, downdate the data structures by undoing the changes recently made in step B3, and return to B4. (Otherwise stop.) ▮

The main point is that if $P_l(x_1,\ldots,x_l)$ is false in step B3, we needn't waste time trying to append any further values $x_{l+1}\ldots x_n$. Thus we can often rule out huge regions of the space of all potential solutions. A second important point is that very little memory is needed, although there may be many, many solutions.

n queens–
diagonal
backtrack tree
profile

For example, let's consider the classic *problem of n queens*: In how many ways can n queens be placed on an $n \times n$ board so that no two are in the same row, column, or diagonal? We can suppose that one queen is in each row, and that the queen in row k is in column x_k, for $1 \le k \le n$. Then each domain D_k is $\{1, 2, \ldots, n\}$; and $P_n(x_1, \ldots, x_n)$ is the condition that

$$x_j \ne x_k \quad \text{and} \quad |x_k - x_j| \ne k - j, \qquad \text{for } 1 \le j < k \le n. \tag{3}$$

(If $x_j = x_k$ and $j < k$, two queens are in the same column; if $|x_k - x_j| = k - j$, they're in the same diagonal.)

This problem is easy to set up for Algorithm B, because we can let property $P_l(x_1, \ldots, x_l)$ be the same as (3) but restricted to $1 \le j < k \le l$. Condition (1) is clear; and so is condition (2), because P_l requires testing (3) only for $k = l$ when P_{l-1} is known. Notice that $P_1(x_1)$ is always true in this example.

One of the best ways to learn about backtracking is to execute Algorithm B by hand in the special case $n = 4$ of the n queens problem: First we set $x_1 \leftarrow 1$. Then when $l = 2$ we find $P_2(1, 1)$ and $P_2(1, 2)$ false; hence we don't get to $l = 3$ until trying $x_2 \leftarrow 3$. Then, however, we're stuck, because $P_3(1, 3, x)$ is false for $1 \le x \le 4$. Backtracking to level 2, we now try $x_2 \leftarrow 4$; and this allows us to set $x_3 \leftarrow 2$. However, we're stuck again, at level 4; and this time we must back up all the way to level 1, because there are no further valid choices at levels 3 and 2. The next choice $x_1 \leftarrow 2$ does, happily, lead to a solution without much further ado, namely $x_1x_2x_3x_4 = 2413$. And one more solution (3142) turns up before the algorithm terminates.

The behavior of Algorithm B is nicely visualized as a tree structure, called a search tree or *backtrack tree*. For example, the backtrack tree for the four queens problem has just 17 nodes,

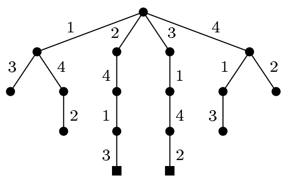

, (4)

corresponding to the 17 times step B2 is performed. Here x_l is shown as the label of an edge from level $l - 1$ to level l of the tree. (Level l of the algorithm actually corresponds to the tree's level $l - 1$, because we've chosen to represent patterns using subscripts from 1 to n instead of from 0 to $n-1$ in this discussion.) The *profile* $(p_0, p_1, \ldots, p_n)$ of this particular tree — the number of nodes at each level — is $(1, 4, 6, 4, 2)$; and we see that the number of solutions, $p_n = p_4$, is 2.

Figure 68 shows the corresponding tree when $n = 8$. This tree has 2057 nodes, distributed according to the profile $(1, 8, 42, 140, 344, 568, 550, 312, 92)$. Thus the early cutoffs facilitated by backtracking have allowed us to find all 92 solutions by examining only 0.01% of the $8^8 = 16{,}777{,}216$ possible sequences $x_1 \ldots x_8$. (And 8^8 is only 0.38% of the $\binom{64}{8} = 4{,}426{,}165{,}368$ ways to put eight queens on the board.)

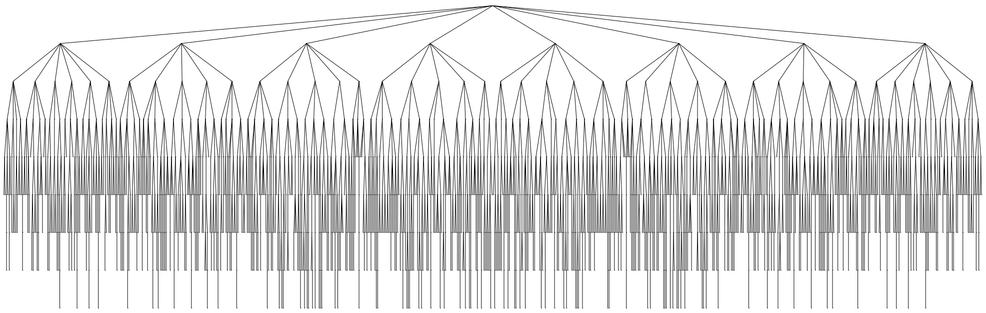

Fig. 68. The problem of placing eight nonattacking queens has this backtrack tree.

data structures–
mems
downdating vs updating+
undoes

Notice that, in this case, Algorithm B spends most of its time in the vicinity of level 5. Such behavior is typical: The backtrack tree for $n = 16$ queens has 1,141,190,303 nodes, and its profile is (1, 16, 210, 2236, 19688, 141812, 838816, 3998456, 15324708, 46358876, 108478966, 193892860, 260303408, 253897632, 171158018, 72002088, 14772512), concentrated near level 12.

Data structures. Backtrack programming is often used when a *huge* tree of possibilities needs to be examined. Thus we want to be able to test property P_l as quickly as possible in step B3.

One way to implement Algorithm B for the n queens problem is to avoid auxiliary data structures and simply to make a bunch of sequential comparisons in that step: "Is $x_l - x_j \in \{j-l, 0, l-j\}$ for some $j < l$?" Assuming that we must access memory whenever referring to x_j, given a trial value x_l in a register, such an implementation performs approximately 112 billion memory accesses when $n = 16$; that's about 98 mems per node.

We can do better by introducing three simple arrays. Property P_l in (3) says essentially that the numbers x_k are distinct, and so are the numbers $x_k + k$, and so are the numbers $x_k - k$. Therefore we can use auxiliary Boolean arrays $a_1 \ldots a_n$, $b_1 \ldots b_{2n-1}$, and $c_1 \ldots c_{2n-1}$, where a_j means 'some $x_k = j$', b_j means 'some $x_k + k - 1 = j$', and c_j means 'some $x_k - k + n = j$'. Those arrays are readily updated and downdated if we customize Algorithm B as follows:

B1*. [Initialize.] Set $a_1 \ldots a_n \leftarrow 0 \ldots 0$, $b_1 \ldots b_{2n-1} \leftarrow 0 \ldots 0$, $c_1 \ldots c_{2n-1} \leftarrow 0 \ldots 0$, and $l \leftarrow 1$.

B2*. [Enter level l.] (Now $P_{l-1}(x_1, \ldots, x_{l-1})$ holds.) If $l > n$, visit $x_1 x_2 \ldots x_n$ and go to B5*. Otherwise set $t \leftarrow 1$.

B3*. [Try t.] If $a_t = 1$ or $b_{t+l-1} = 1$ or $c_{t-l+n} = 1$, go to B4*. Otherwise set $a_t \leftarrow 1$, $b_{t+l-1} \leftarrow 1$, $c_{t-l+n} \leftarrow 1$, $x_l \leftarrow t$, $l \leftarrow l + 1$, and go to B2*.

B4*. [Try again.] If $t < n$, set $t \leftarrow t + 1$ and return to B3*.

B5*. [Backtrack.] Set $l \leftarrow l - 1$. If $l > 0$, set $t \leftarrow x_l$, $c_{t-l+n} \leftarrow 0$, $b_{t+l-1} \leftarrow 0$, $a_t \leftarrow 0$, and return to B4*. (Otherwise stop.) ▮

Notice how step B5* neatly undoes the updates that step B3* had made, in the reverse order. Reverse order for downdating is typical of backtrack algorithms,

registers
Walker
cutoff properties
visits
bitwise operations
historical notes+
Sprague

although there is some flexibility; we could, for example, have restored a_t before b_{t+l-1} and c_{t-l+n}, because those arrays are independent.

The auxiliary arrays a, b, c make it easy to test property P_l at the beginning of step B3*, but we must also access memory when we update them and downdate them. Does that cost us more than it saves? Fortunately, no: The running time for $n = 16$ goes down to about 34 billion mems, roughly 30 mems per node.

Furthermore we could keep the bit vectors a, b, c entirely in registers, on a machine with 64-bit registers, assuming that $n \le 32$. Then there would be just two memory accesses per node, namely to store $x_l \leftarrow t$ and later to fetch $t \leftarrow x_l$. However, quite a lot of in-register computation would become necessary.

Walker's method. The 1950s-era programs of R. J. Walker organized backtracking in a somewhat different way. Instead of letting x_l run through all elements of D_l, he calculated and stored the set

$$S_l \leftarrow \{x \in D_l \mid P_l(x_1, \ldots, x_{l-1}, x) \text{ holds}\} \tag{5}$$

upon entry to each node at level l. This computation can often be done efficiently all at once, instead of piecemeal, because some cutoff properties make it possible to combine steps that would otherwise have to be repeated for each $x \in D_l$. In essence, he used the following variant of Algorithm B:

Algorithm W (*Walker's backtrack*). Given domains D_k and cutoffs P_l as above, this algorithm visits all sequences $x_1 x_2 \ldots x_n$ that satisfy $P_n(x_1, x_2, \ldots, x_n)$.

W1. [Initialize.] Set $l \leftarrow 1$, and initialize the data structures needed later.

W2. [Enter level l.] (Now $P_{l-1}(x_1, \ldots, x_{l-1})$ holds.) If $l > n$, visit $x_1 x_2 \ldots x_n$ and go to W4. Otherwise determine the set S_l as in (5).

W3. [Try to advance.] If S_l is nonempty, set $x_l \leftarrow \min S_l$, update the data structures to facilitate computing S_{l+1}, set $l \leftarrow l + 1$, and go to W2.

W4. [Backtrack.] Set $l \leftarrow l - 1$. If $l > 0$, downdate the data structures by undoing changes made in step W3, set $S_l \leftarrow S_l \setminus x_l$, and retreat to W3. ▮

Walker applied this method to the n queens problem by computing $S_l = U \setminus A_l \setminus B_l \setminus C_l$, where $U = D_l = \{1, \ldots, n\}$ and

$$A_l = \{x_j \mid 1 \le j < l\},\ B_l = \{x_j + j - l \mid 1 \le j < l\},\ C_l = \{x_j - j + l \mid 1 \le j < l\}. \tag{6}$$

He represented these auxiliary sets by bit vectors a, b, c, analogous to (but different from) the bit vectors of Algorithm B* above. Exercise 10 shows that the updating in step W3 is easy, using bitwise operations on n-bit numbers; furthermore, no downdating is needed in step W4. The corresponding run time when $n = 16$ turns out to be just 9.1 gigamems, or 8 mems per node.

Let $Q(n)$ be the number of solutions to the n queens problem. Then we have

$n =$	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
$Q(n) =$	1	1	0	0	2	10	4	40	92	352	724	2680	14200	73712	365596	2279184	14772512

and the values for $n \le 11$ were computed independently by several people during the nineteenth century. Small cases were relatively easy; but when T. B. Sprague had finished computing $Q(11)$ he remarked that "This was a very heavy piece of work, and occupied most of my leisure time for several months. ... It will, I imagine, be scarcely possible to obtain results for larger boards, unless a number of persons co-operate in the work." [See *Proc. Edinburgh Math. Soc.* **17** (1899), 43–68; Sprague was the leading actuary of his day.] Nevertheless, H. Onnen went on to evaluate $Q(12) = 14{,}200$ — an astonishing feat of hand calculation — in 1910. [See W. Ahrens, *Math. Unterhaltungen und Spiele* **2**, second edition (1918), 344.]

All of these hard-won results were confirmed in 1960 by R. J. Walker, using the SWAC computer at UCLA and the method of exercise 10. Walker also computed $Q(13)$; but he couldn't go any further with the machine available to him at the time. The next step, $Q(14)$, was computed by Michael D. Kennedy at the University of Tennessee in 1963, commandeering an IBM 1620 for 120 hours. S. R. Bunch evaluated $Q(15)$ in 1974 at the University of Illinois, using about two hours on an IBM System 360-75; then J. R. Bitner found $Q(16)$ after about three hours on the same computer, but with an improved method.

Computers and algorithms have continued to get better, of course, and such results are now obtained almost instantly. Hence larger and larger values of n lie at the frontier. The whopping value $Q(27) = 234{,}907{,}967{,}154{,}122{,}528$, found in 2016 by Thomas B. Preußer and Matthias R. Engelhardt, probably won't be exceeded for awhile! [See *J. Signal Processing Systems* **88** (2017), 185–201. This distributed computation occupied a dynamic cluster of diverse FPGA devices for 383 days; those devices provided a total peak of more than 7000 custom-designed hardware solvers to handle 2,024,110,796 independent subproblems.]

Permutations and Langford pairs. Every solution $x_1 \ldots x_n$ to the n queens problem is a permutation of $\{1, \ldots, n\}$, and many other problems are permutation-based. Indeed, we've already seen Algorithm 7.2.1.2X, which is an elegant backtrack procedure specifically designed for special kinds of permutations. When that algorithm begins to choose the value of x_l, it makes all of the appropriate elements $\{1, 2, \ldots, n\} \setminus \{x_1, \ldots, x_{l-1}\}$ conveniently accessible in a linked list.

We can get further insight into such data structures by returning to the problem of Langford pairs, which was discussed at the very beginning of Chapter 7. That problem can be reformulated as the task of finding all permutations of $\{1, 2, \ldots, n\} \cup \{-1, -2, \ldots, -n\}$ with the property that

$$x_j = k \quad \text{implies} \quad x_{j+k+1} = -k, \qquad \text{for } 1 \le j \le 2n \text{ and } 1 \le k \le n. \tag{7}$$

For example, when $n = 4$ there are two solutions, namely $234\bar{2}1\bar{3}\bar{1}\bar{4}$ and $413\bar{1}2\bar{4}\bar{3}\bar{2}$. (As usual we find it convenient to write $\bar{1}$ for -1, $\bar{2}$ for -2, etc.) Notice that if $x = x_1 x_2 \ldots x_{2n}$ is a solution, so is its "dual" $-x^R = (-x_{2n}) \ldots (-x_2)(-x_1)$.

Here's a Langford-inspired adaptation of Algorithm 7.2.1.2X, with the former notation modified slightly to match Algorithms B and W: We want to maintain pointers $p_0 p_1 \ldots p_n$ such that, if the positive integers not already present in $x_1 \ldots x_{l-1}$ are $k_1 < k_2 < \cdots < k_t$ when we're choosing x_l, we have the linked list

$$p_0 = k_1,\ p_{k_1} = k_2,\ \ldots,\ p_{k_{t-1}} = k_t,\ p_{k_t} = 0. \tag{8}$$

Such a condition turns out to be easy to maintain.

Sprague
Onnen
Ahrens
Walker
SWAC computer
UCLA
Kennedy
University of Tennessee
IBM 1620
Bunch
University of Illinois
IBM System 360-75
Bitner
Preußer
Engelhardt
distributed computation
FPGA
permutation
data structures
Langford pairs
dual
linked list

undoes
deletion operation
dancing links
search tree
cutoff principle

Algorithm L (*Langford pairs*). This algorithm visits all solutions $x_1 \ldots x_{2n}$ to (7) in lexicographic order, using pointers $p_0 p_1 \ldots p_n$ that satisfy (8), and also using an auxiliary array $y_1 \ldots y_{2n}$ for backtracking.

L1. [Initialize.] Set $x_1 \ldots x_{2n} \leftarrow 0 \ldots 0$, $p_k \leftarrow k+1$ for $0 \le k < n$, $p_n \leftarrow 0$, $l \leftarrow 1$.

L2. [Enter level l.] Set $k \leftarrow p_0$. If $k = 0$, visit $x_1 x_2 \ldots x_{2n}$ and go to L5. Otherwise set $j \leftarrow 0$, and while $x_l < 0$ set $l \leftarrow l+1$.

L3. [Try $x_l = k$.] (At this point we have $k = p_j$.) If $l + k + 1 > 2n$, go to L5. Otherwise, if $x_{l+k+1} = 0$, set $x_l \leftarrow k$, $x_{l+k+1} \leftarrow -k$, $y_l \leftarrow j$, $p_j \leftarrow p_k$, $l \leftarrow l+1$, and return to L2.

L4. [Try again.] (We've found all solutions that begin with $x_1 \ldots x_{l-1} k$ or something smaller.) Set $j \leftarrow k$ and $k \leftarrow p_j$, then go to L3 if $k \ne 0$.

L5. [Backtrack.] Set $l \leftarrow l - 1$. If $l > 0$ do the following: While $x_l < 0$, set $l \leftarrow l - 1$. Then set $k \leftarrow x_l$, $x_l \leftarrow 0$, $x_{l+k+1} \leftarrow 0$, $j \leftarrow y_l$, $p_j \leftarrow k$, and go back to L4. Otherwise terminate the algorithm. ▮

Careful study of these steps will reveal how everything fits together nicely. Notice that, for example, step L3 removes k from the linked list (8) by simply setting $p_j \leftarrow p_k$. That step also sets $x_{l+k+1} \leftarrow -k$, in accordance with (7), so that we can skip over position $l+k+1$ when we encounter it later in step L2.

The main point of Algorithm L is the somewhat subtle way in which step L5 undoes the deletion operation by setting $p_j \leftarrow k$. The pointer p_k still retains the appropriate link to the next element in the list, *because p_k has not been changed by any of the intervening updates.* (Think about it.) This is the germ of an idea called "dancing links" that we will explore in Section 7.2.2.1.

To draw the search tree corresponding to a run of Algorithm L, we can label the edges with the positive choices of x_l as we did in (4), while labeling the nodes with any previously set negative values that are passed over in step L2. For instance the tree for $n = 4$ is

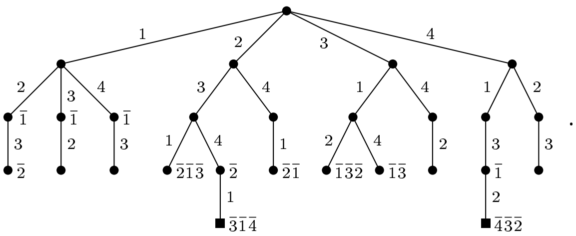

. (9)

Solutions appear at depth n in this tree, even though they involve $2n$ values $x_1 x_2 \ldots x_{2n}$.

Algorithm L sometimes makes false starts and doesn't realize the problem until probing further than necessary. Notice that the value $x_l = k$ can appear only when $l + k + 1 \le 2n$; hence if we haven't seen k by the time l reaches $2n - k - 1$, we're *forced* to choose $x_l = k$. For example, the branch $12\bar{1}$ in (9) needn't be pursued, because 4 must appear in $\{x_1, x_2, x_3\}$. Exercise 20 explains how to incorporate this cutoff principle into Algorithm L. When $n = 17$, it reduces the number of nodes in the search tree from 1.29 trillion to 330 billion, and reduces the running time from 25.0 teramems to 8.1 teramems. (The amount of work has gone up from 19.4 mems per node to 24.4 mems per node, because of the extra tests for cutoffs, yet there's a significant overall reduction.)

Furthermore, we can "break the symmetry" by ensuring that we don't consider both a solution and its dual. This idea, exploited in exercise 21, reduces the search tree to just 160 billion nodes and costs just 3.94 teramems — that's 24.6 mems per node.

Word rectangles. Let's look next at a problem where the search domains D_l are much larger. An $m \times n$ *word rectangle* is an array of n-letter words* whose columns are m-letter words. For example,

```
S T A T U S
L O W E S T
U T O P I A
M A K I N G
S L E D G E
```
(10)

is a 5×6 word rectangle whose columns all belong to `WORDS`(5757), the collection of 5-letter words in the Stanford GraphBase. To find such patterns, we can suppose that column l contains the x_lth most common 5-letter word, where $1 \le x_l \le 5757$ for $1 \le l \le 6$; hence there are $5757^6 = 36{,}406{,}369{,}848{,}837{,}732{,}146{,}649$ ways to choose the columns. In (10) we have $x_1 \ldots x_6 = 1446\ 185\ 1021\ 2537\ 66\ 255$. Of course very few of those choices will yield suitable rows; but backtracking will hopefully help us to find all solutions in a reasonable amount of time.

We can set this problem up for Algorithm B by storing the n-letter words in a trie (see Section 6.3), with one trie node of size 26 for each l-letter prefix of a legitimate word, $0 \le l \le n$.

For example, such a trie for $n = 6$ represents 15727 words with 23667 nodes. The prefix `ST` corresponds to node number 260, whose 26 entries are

$$(484,0,0,0,1589,0,0,0,2609,0,0,0,0,0,1280,0,0,251,0,0,563,0,0,0,1621,0); \quad (11)$$

this means that `STA` is node 484, `STE` is node 1589, ..., `STY` is node 1621, and there are no 6-letter words beginning with `STB`, `STC`, ..., `STX`, `STZ`. A slightly different convention is used for prefixes of length $n - 1$; for example, the entries for node 580, '`CORNE`', are

$$(3879,0,0,3878,0,0,0,0,0,0,0,9602,0,0,0,0,0,171,0,5013,0,0,0,0,0,0), \quad (12)$$

meaning that `CORNEA`, `CORNED`, `CORNEL`, `CORNER`, and `CORNET` are ranked 3879, 3878, 9602, 171, and 5013 in the list of 6-letter words.

* Whenever five-letter words are used in the examples of this book, they're taken from the 5757 Stanford GraphBase words as explained at the beginning of Chapter 7. Words of other lengths are taken from *The Official SCRABBLE® Players Dictionary*, fourth edition (Hasbro, 2005), because those words have been incorporated into many widely available computer games. Such words have been ranked according to the British National Corpus of 2007 — where '`the`' occurs 5,405,633 times and the next-most common word, '`of`', occurs roughly half as often (3,021,525). The OSPD4 list includes respectively (101, 1004, 4002, 8887, 15727, 23958, 29718, 29130, 22314, 16161, 11412) words of lengths (2, 3, ..., 12), of which (97, 771, 2451, 4474, 6910, 8852, 9205, 8225, 6626, 4642, 3061) occur at least six times in the British National Corpus.

break the symmetry
dual
word rectangles–
footnote doesn't show up here
n-letter words
WORDS
5-letter words
Stanford GraphBase
trie

Suppose x_1 and x_2 specify the 5-letter column-words SLUMS and TOTAL as in (10). Then the trie tells us that the next column-word x_3 must have the form $c_1c_2c_3c_4c_5$ where $c_1 \in \{\mathtt{A},\mathtt{E},\mathtt{I},\mathtt{O},\mathtt{R},\mathtt{U},\mathtt{Y}\}$, $c_2 \notin \{\mathtt{E},\mathtt{H},\mathtt{J},\mathtt{K},\mathtt{Y},\mathtt{Z}\}$, $c_3 \in \{\mathtt{E},\mathtt{M},\mathtt{O},\mathtt{T}\}$, $c_4 \notin \{\mathtt{A},\mathtt{B},\mathtt{O}\}$, and $c_5 \in \{\mathtt{A},\mathtt{E},\mathtt{I},\mathtt{O},\mathtt{U},\mathtt{Y}\}$. (There are 221 such words.)

Let $a_{l1} \ldots a_{lm}$ be the trie nodes corresponding to the prefixes of the first l columns of a partial solution to the word rectangle problem. This auxiliary array enables Algorithm B to find all solutions, as explained in exercise 24. It turns out that there are exactly 625,415 valid 5×6 word rectangles, according to our conventions; and the method of exercise 24 needs about 19 teramems of computation to find them all. In fact, the profile of the search tree is

$$(1,\ 5757,\ 2458830,\ 360728099,\ 579940198,\ 29621728,\ 625415), \qquad (13)$$

indicating for example that just 360,728,099 of the $5757^3 = 190{,}804{,}533{,}093$ choices for $x_1x_2x_3$ will lead to valid prefixes of 6-letter words.

With care, exercise 24's running time can be significantly decreased, once we realize that every node of the search tree for $1 \le l \le n$ requires testing 5757 possibilities for x_l in step B3. If we build a more elaborate data structure for the 5-letter words, so that it becomes easy to run though all words that have a specific letter in a specific position, we can refine the algorithm so that the average number of possibilities per level that need to be investigated becomes only

$$(5757.0,\ 1697.9,\ 844.1,\ 273.5,\ 153.5,\ 100.8); \qquad (14)$$

the total running time then drops to 1.15 teramems. Exercise 25 has the details. And exercise 28 discusses a method that's faster yet.

Commafree codes. Our next example deals entirely with *four*-letter words. But it's not obscene; it's an intriguing question of coding theory. The problem is to find a set of four-letter words that can be decoded even if we don't put spaces or other delimiters between them. If we take any message that's formed from words of the set by simply concatenating them together, `likethis`, and if we look at any seven consecutive letters $\ldots x_1x_2x_3x_4x_5x_6x_7\ldots$, exactly one of the four-letter substrings $x_1x_2x_3x_4$, $x_2x_3x_4x_5$, $x_3x_4x_5x_6$, $x_4x_5x_6x_7$ will be a codeword. Equivalently, if $x_1x_2x_3x_4$ and $x_5x_6x_7x_8$ are codewords, then $x_2x_3x_4x_5$ and $x_3x_4x_5x_6$ and $x_4x_5x_6x_7$ aren't. (For example, `iket` isn't.) Such a set is called a "commafree code" or a "self-synchronizing block code" of length four.

Commafree codes were introduced by F. H. C. Crick, J. S. Griffith, and L. E. Orgel [*Proc. National Acad. Sci.* **43** (1957), 416–421], and studied further by S. W. Golomb, B. Gordon, and L. R. Welch [*Canadian Journal of Mathematics* **10** (1958), 202–209], who considered the general case of m-letter alphabets and n-letter words. They constructed optimum commafree codes for all m when $n = 2$, 3, 5, 7, 9, 11, 13, and 15; and optimum codes for all m were subsequently found also for $n = 17$, 19, 21, … (see exercise 37). We will focus our attention on the four-letter case here ($n = 4$), partly because that case is still very far from being resolved, but mostly because the task of finding such codes is especially instructive. Indeed, our discussion will lead us naturally to an understanding of several significant techniques that are important for backtrack programming in general.

teramems
profile
search tree
data structure
commafree codes–
4-letter codewords
four-letter words
coding theory
concatenating
codewords
self-synchronizing block code
block code
Crick
Griffith
Orgel
Golomb
Gordon
Welch

periodic
aperiodic
cyclic shifts
prime strings
lookahead
dynamic ordering–
search rearrangement, see dynamic ordering
cutoff properties

To begin, we can see immediately that a commafree codeword cannot be "periodic," like `dodo` or `gaga`. Such a word already appears within two adjacent copies of itself. Thus we're restricted to *aperiodic* words like `item`, of which there are $m^4 - m^2$. Notice further that if `item` has been chosen, we aren't allowed to include any of its cyclic shifts `temi`, `emit`, or `mite`, because they all appear within `itemitem`. Hence the maximum number of codewords in our commafree code cannot exceed $(m^4 - m^2)/4$.

For example, consider the binary case, $m = 2$, when this maximum is 3. Can we choose three four-bit "words," one from each of the cyclic classes

$$\begin{aligned} [0001] &= \{0001, 0010, 0100, 1000\}, \\ [0011] &= \{0011, 0110, 1100, 1001\}, \\ [0111] &= \{0111, 1110, 1101, 1011\}, \end{aligned} \qquad (15)$$

so that the resulting code is commafree? Yes: One solution in this case is simply to choose the smallest word in each class, namely 0001, 0011, and 0111. (Alert readers will recall that we studied the smallest word in the cyclic class of *any* aperiodic string in Section 7.2.1.1, where such words were called *prime strings* and where some of the remarkable properties of prime strings were proved.)

That trick doesn't work when $m = 3$, however, when there are $(81 - 9)/4 = 18$ cyclic classes. Then we cannot include 1112 after we've chosen 0001 and 0011. Indeed, a code that contains 0001 and 1112 can't contain either 0011 or 0111.

We could systematically backtrack through 18 levels, choosing x_1 in [0001] and x_2 in [0011], etc., and rejecting each x_l as in Algorithm B whenever we discover that $\{x_1, x_2, \ldots, x_l\}$ isn't commafree. For example, if $x_1 = 0010$ and we try $x_2 = 1001$, this approach would backtrack because x_1 occurs inside $x_2 x_1$.

But a naïve strategy of that kind, which recognizes failure only after a bad choice has been made, can be vastly improved. If we had been clever enough, we could have looked a little bit ahead, and never even considered the choice $x_2 = 1001$ in the first place. Indeed, after choosing $x_1 = 0010$, we can automatically exclude *all* further words of the form $*001$, such as 2001 when $m \ge 3$ and 3001 when $m \ge 4$.

Even better pruning occurs if, for example, we've chosen $x_1 = 0001$ and $x_2 = 0011$. Then we can immediately rule out all words of the forms $1{*}{*}{*}$ or ${*}{*}{*}0$, because $x_1 1{*}{*}{*}$ includes x_2 and ${*}{*}{*}0x_2$ includes x_1. Already we could then deduce, in the case $m \ge 3$, that classes [0002], [0021], [0111], [0211], and [1112] *must* be represented by 0002, 0021, 0111, 0211, and 2111, respectively; each of the other three possibilities in those classes has been wiped out!

Thus we see the desirability of a lookahead mechanism.

Dynamic ordering of choices. Furthermore, we can see from this example that it's not always good to choose x_1, then x_2, then x_3, and so on when trying to satisfy a general property $P_n(x_1, x_2, \ldots, x_n)$ in the setting of Algorithm B. Maybe the search tree will be much smaller if we first choose x_5, say, and then turn next to some other x_j, depending on the particular value of x_5 that was selected. Some orderings might have much better cutoff properties than others, and every branch of the tree is free to choose its variables in any desired order.

frame
move
sequential lists–
downdating vs updating
cache-friendly
unordered sequential list
cells of memory
MEM, an array of "cells"–

Indeed, our commafree coding problem for ternary 4-tuples doesn't dictate any particular ordering of the 18 classes that would be likely to keep the search tree small. Therefore, instead of calling those choices x_1, x_2, ..., x_{18}, it's better to identify them by the various class names, namely x_{0001}, x_{0002}, x_{0011}, x_{0012}, x_{0021}, x_{0022}, x_{0102}, x_{0111}, x_{0112}, x_{0121}, x_{0122}, x_{0211}, x_{0212}, x_{0221}, x_{0222}, x_{1112}, x_{1122}, x_{1222}. (Algorithm 7.2.1.1F is a good way to generate those names.) At every node of the search tree we then can choose a convenient variable on which to branch, based on previous choices. After beginning with $x_{0001} \leftarrow 0001$ at level 1 we might decide to try $x_{0011} \leftarrow 0011$ at level 2; and then, as we've seen, the choices $x_{0002} \leftarrow 0002$, $x_{0021} \leftarrow 0021$, $x_{0111} \leftarrow 0111$, $x_{0211} \leftarrow 0211$, and $x_{1112} \leftarrow 2111$ are forced, so we should make them at levels 3 through 7.

Furthermore, after those forced moves are made, it turns out that they don't force any others. But only two choices for x_{0012} will remain, while x_{0122} will have three. Therefore it will probably be wiser to branch on x_{0012} rather than on x_{0122} at level 8. (Incidentally, it also turns out that there *is* no commafree code with $x_{0001} = 0001$ and $x_{0011} = 0011$, *except* when $m = 2$.)

It's easy to adapt Algorithms B and W to allow dynamic ordering. Every node of the search tree can be given a "frame" in which we record the variable being set and the choice that was made. This choice of variable and value can be called a "move" made by the backtrack procedure.

Dynamic ordering can be helpful also after backtracking has taken place. If we continue the example above, where $x_{0001} = 0001$ and we've explored all cases in which $x_{0011} = 0011$, we aren't obliged to continue by trying another value for x_{0011}. We do want to remember that 0011 should no longer be considered legal, until x_{0001} changes; but we could decide to explore next a case such as $x_{0002} \leftarrow 2000$ at level 2. In fact, $x_{0002} = 2000$ is quickly seen to be impossible in the presence of 0001 (see exercise 39). An even more efficient choice at level 2 turns out to be $x_{0012} \leftarrow 0012$, because that branch immediately forces $x_{0002} \leftarrow 0002$, $x_{0022} \leftarrow 0022$, $x_{0122} \leftarrow 0122$, $x_{0222} \leftarrow 0222$, $x_{1222} \leftarrow 1222$, and $x_{0011} \leftarrow 1001$.

Sequential allocation redux. The choice of a variable and value on which to branch is a delicate tradeoff. We don't want to devote more time to planning than we'll save by having a good plan.

If we're going to benefit from dynamic ordering, we'll need efficient data structures that will lead to good decisions without much deliberation. On the other hand, elaborate data structures need to be updated whenever we branch to a new level, and they need to be downdated whenever we return from that level. Algorithm L illustrates an efficient mechanism based on linked lists; but sequentially allocated lists are often even more appealing, because they are cache-friendly and they involve fewer accesses to memory.

Assume then that we wish to represent a set of items as an unordered sequential list. The list begins in a cell of memory pointed to by `HEAD`, and `TAIL` points just beyond the end of the list. For example,

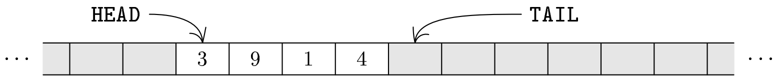

(16)

is one way to represent the set $\{1,3,4,9\}$. The number of items currently in the set is $\mathtt{TAIL}-\mathtt{HEAD}$; thus $\mathtt{TAIL}=\mathtt{HEAD}$ if and only if the list is empty. If we wish to insert a new item x, knowing that x isn't already present, we simply set

$$\mathtt{MEM[TAIL]} \leftarrow x, \quad \mathtt{TAIL} \leftarrow \mathtt{TAIL}+1. \tag{17}$$

Conversely, if $\mathtt{HEAD} \le \mathtt{P} < \mathtt{TAIL}$, we can easily delete $\mathtt{MEM[P]}$:

$$\mathtt{TAIL} \leftarrow \mathtt{TAIL}-1; \quad \text{if } \mathtt{P} \ne \mathtt{TAIL}, \text{ set } \mathtt{MEM[P]} \leftarrow \mathtt{MEM[TAIL]}. \tag{18}$$

(We've tacitly assumed in (17) that MEM[TAIL] is available for use whenever a new item is inserted. Otherwise we would have had to test for memory overflow.)

We can't delete an item from a list without knowing its MEM location. Thus we will often want to maintain an "inverse list," assuming that all items x lie in the range $0 \le x < M$. For example, (16) becomes the following, if $M = 10$:

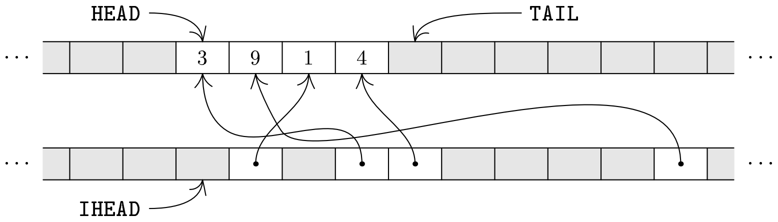

(19)

(Shaded cells have undefined contents.) With this setup, insertion (17) becomes

$$\mathtt{MEM[TAIL]} \leftarrow x, \quad \mathtt{MEM[IHEAD}+x\mathtt{]} \leftarrow \mathtt{TAIL}, \quad \mathtt{TAIL} \leftarrow \mathtt{TAIL}+1, \tag{20}$$

and TAIL will never exceed $\mathtt{HEAD}+M$. Similarly, deletion of x becomes

$$\begin{aligned}&\mathtt{P} \leftarrow \mathtt{MEM[IHEAD}+x\mathtt{]}, \quad \mathtt{TAIL} \leftarrow \mathtt{TAIL}-1;\\ &\quad \text{if } \mathtt{P} \ne \mathtt{TAIL}, \text{ set } y \leftarrow \mathtt{MEM[TAIL]},\ \mathtt{MEM[P]} \leftarrow y,\ \mathtt{MEM[IHEAD}+y\mathtt{]} \leftarrow \mathtt{P}.\end{aligned} \tag{21}$$

For example, after deleting '9' from (19) we would obtain this:

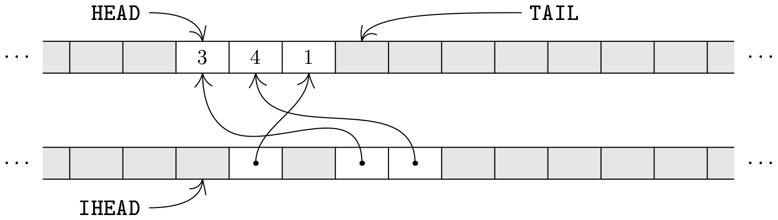

(22)

In more elaborate situations we also want to test whether or not a given item x is present. If so, we can keep more information in the inverse list. A particularly useful variation arises when the list that begins at IHEAD contains a *complete* permutation of the values $\{\mathtt{HEAD}, \mathtt{HEAD}+1, \ldots, \mathtt{HEAD}+M-1\}$, and the memory cells beginning at HEAD contain the inverse permutation — although only the first $\mathtt{TAIL}-\mathtt{HEAD}$ elements of that list are considered to be "active."

For example, in our commafree code problem with $m = 3$, we can begin by putting items representing the $M = 18$ cycle classes [0001], [0002], …, [1222] into memory cells HEAD through $\mathtt{HEAD}+17$. Initially they're all active, with

empty
insert
overflow
inverse list
inverse permutation
active

delete
data structures
periodic
radix m

TAIL $=$ HEAD $+ 18$ and MEM[IHEAD $+ c$] $=$ HEAD $+ c$ for $0 \le c < 18$. Then whenever we decide to choose a codeword for class c, we delete c from the active list by using a souped-up version of (21) that maintains full permutations:

$$\begin{aligned}&\mathtt{P} \leftarrow \mathtt{MEM[IHEAD} + c\mathtt{]}, \quad \mathtt{TAIL} \leftarrow \mathtt{TAIL} - 1;\\ &\quad \text{if } \mathtt{P} \ne \mathtt{TAIL}, \text{ set } y \leftarrow \mathtt{MEM[TAIL]}, \quad \mathtt{MEM[TAIL]} \leftarrow c, \quad \mathtt{MEM[P]} \leftarrow y,\\ &\qquad\qquad \mathtt{MEM[IHEAD} + c\mathtt{]} \leftarrow \mathtt{TAIL}, \quad \mathtt{MEM[IHEAD} + y\mathtt{]} \leftarrow \mathtt{P}.\end{aligned} \tag{23}$$

Later on, after backtracking to a state where we once again want c to be considered active, we simply set TAIL $\leftarrow$ TAIL $+ 1$, because c will already be in place!

Lists for the commafree problem. The task of finding all four-letter commafree codes is not difficult when $m = 3$ and only 18 cycle classes are involved. But it already becomes challenging when $m = 4$, because we must then deal with $(4^4 - 4^2)/4 = 60$ classes. Therefore we'll want to give it some careful thought as we try to set it up for backtracking.

The example scenarios for $m = 3$ considered above suggest that we'll repeatedly want to know the answers to questions such as, "How many words of the form 02** are still available for selection as codewords?" Redundant data structures, oriented to queries of that kind, appear to be needed. Fortunately, we shall see that there's a nice way to provide them, using sequential lists as in (19)–(23).

In Algorithm C below, each of the m^4 four-letter words is given one of three possible states during the search for commafree codes. A word is *green* if it's part of the current set of tentative codewords. It is *red* if it's not currently a candidate for such status, either because it is incompatible with the existing green words or because the algorithm has already examined all scenarios in which it is green in their presence. Every other word is *blue*, and sort of in limbo; the algorithm might or might not decide to make it red or green. All words are initially blue — except for the m^2 periodic words, which are permanently red.

We'll use the Greek letter α to stand for the integer value of a four-letter word x in radix m. For example, if $m = 3$ and if x is the word 0102, then $\alpha = (0102)_3 = 11$. The current state of word x is kept in MEM[α], using one of the arbitrary internal codes 2 (GREEN), 0 (RED), or 1 (BLUE).

The most important feature of the algorithm is that *every blue word $x = x_1x_2x_3x_4$ is potentially present in seven different lists*, called P1(x), P2(x), P3(x), S1(x), S2(x), S3(x), and CL(x), where

- P1(x), P2(x), P3(x) are the blue words matching $x_1{*}{*}{*}$, $x_1x_2{*}{*}$, $x_1x_2x_3{*}$;
- S1(x), S2(x), S3(x) are the blue words matching ${*}{*}{*}x_4$, ${*}{*}x_3x_4$, ${*}x_2x_3x_4$;
- CL(x) hosts the blue words in $\{x_1x_2x_3x_4, x_2x_3x_4x_1, x_3x_4x_1x_2, x_4x_1x_2x_3\}$.

These seven lists begin respectively in MEM locations P1OFF $+ p_1(\alpha)$, P2OFF $+ p_2(\alpha)$, P3OFF $+ p_3(\alpha)$, S1OFF $+ s_1(\alpha)$, S2OFF $+ s_2(\alpha)$, S3OFF $+ s_3(\alpha)$, and CLOFF $+ 4cl(\alpha)$; here (P1OFF, P2OFF, P3OFF, S1OFF, S2OFF, S3OFF, CLOFF) are respectively $(2m^4,\ 5m^4,\ 8m^4,\ 11m^4,\ 14m^4,\ 17m^4,\ 20m^4)$. We define $p_1((x_1x_2x_3x_4)_m) = (x_1000)_m$, $p_2((x_1x_2x_3x_4)_m) = (x_1x_200)_m$, $p_3((x_1x_2x_3x_4)_m) = (x_1x_2x_30)_m$; and $s_1((x_1x_2x_3x_4)_m) = (x_4000)_m$, $s_2((x_1x_2x_3x_4)_m) = (x_3x_400)_m$, $s_3((x_1x_2x_3x_4)_m) = (x_2x_3x_40)_m$; and finally $cl((x_1x_2x_3x_4)_m)$ is an internal number, between 0 and

reflection
symmetry breaking
closed

Table 1

LISTS USED BY ALGORITHM C (m = 2), ENTERING LEVEL 1

	0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f	
0	RED	BLUE	BLUE	BLUE	RED	RED	BLUE	BLUE	RED	BLUE	RED	BLUE	BLUE	BLUE	BLUE	RED	
10		20	21	22			23	24		29		2c	28	2b	2a		
20	0001	0010	0011	0110	0111				1100	1001	1110	1101	1011				P1
30	25								2d								
40		50	51	52			54	55		58		59	5c	5e	5d		
50	0001	0010	0011		0110	0111			1001	1011			1100	1110	1101		P2
60	53				56				5a				5f				
70		80	82	83			86	87		88		8a	8c	8d	8e		
80	0001		0010	0011			0110	0111	1001		1011		1100	1101	1110		P3
90	81		84		84		88		89		8b		8e		8f		
a0		b8	b0	b9			b1	bb		ba		bd	b2	bc	b3		
b0	0010	0110	1100	1110					0001	0011	1001	0111	1101	1011			S1
c0	b4								be								
d0		e4	e8	ec			e9	ed		e5		ee	e0	e6	ea		
e0	1100				0001	1001	1101		0010	0110	1110		0011	0111	1011		S2
f0	e1				e7				eb				ef				
100		112	114	116			11c	11e		113		117	118	11a	11d		
110			0001	1001	0010		0011	1011	1100		1101		0110	1110	0111		S3
120	110		114		115		118		119		11b		11e		11f		
130		140	141	144			145	148		147		14b	146	14a	149		
140	0001	0010			0011	0110	1100	1001	0111	1110	1101	1011					CL
150	142				148				14c								

This table shows MEM locations 0000 through 150f, using hexadecimal notation. (For example, MEM[40d] =5e; see exercise 41.) Blank entries are unused by the algorithm.

$(m^4-m^2)/4-1$, assigned to each class. The seven MEM locations where x appears in these seven lists are respectively kept in inverse lists that begin in MEM locations P1OFF $-\, m^4+\alpha$, P2OFF $-\, m^4+\alpha$, …, CLOFF $-\, m^4+\alpha$. And the TAIL pointers, which indicate the current list sizes as in (19)–(23), are respectively kept in MEM locations P1OFF $+\, m^4+p_1(\alpha)$, P2OFF $+\, m^4+p_2(\alpha)$, …, CLOFF $+\, m^4+cl(\alpha)$. (Whew; got that?)

This vast apparatus, which occupies $22m^4$ cells of MEM, is illustrated in Table 1, at the beginning of the computation for the case $m=2$. Fortunately it's not really as complicated as it may seem at first. Nor is it especially vast: After all, $22m^4$ is only 13,750 when $m=5$.

(A close inspection of Table 1 reveals incidentally that the words 0100 and 1000 have been colored red, not blue. That's because we can assume without loss of generality that class [0001] is represented either by 0001 or by 0010. The other two cases are covered by left-right reflection of all codewords.)

Algorithm C finds these lists invaluable when it is deciding where next to branch. But it has no further use for a list in which one of the items has become green. Therefore it declares such lists "closed"; and it saves most of the work of list maintenance by updating *only* the lists that remain open. A closed list is represented internally by setting its TAIL pointer to HEAD $-$ 1.

For example, Table 2 shows how the lists in MEM will have changed just after $x=0010$ has been chosen to be a tentative codeword. The elements $\{0001,0010,0011,0110,0111\}$ of P1(x) are effectively hidden, because the tail

undoing–
Floyd
compiler

Table 2

LISTS USED BY ALGORITHM C ($m = 2$), ENTERING LEVEL 2

	0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f	
0	RED	RED	GREEN	BLUE	RED	RED	BLUE	BLUE	RED	RED	RED	BLUE	BLUE	BLUE	BLUE	RED	
10												29	28	2b	2a		
20									1100	1011	1110	1101					P1
30	1f								2c								
40							54	55				58	5c	5e	5d		
50					0110	0111			1011				1100	1110	1101		P2
60	4f				56				59				5f				
70							86	87				8a	8c	8d	8e		
80							0110	0111			1011		1100	1101	1110		P3
90	80		81		84		88		88		8b		8e		8f		
a0				b9				bb				b8		ba			
b0									1011	0011	1101	0111					S1
c0	af								bc								
d0				ec				ed				ee	e0	e4			
e0	1100				1101								0011	0111	1011		S2
f0	e1				e5				e7				ef				
100				116			11c	11e				117	118	11a	11d		
110							0011	1011	1100		1101		0110	1110	0111		S3
120	110		112		113		118		119		11b		11e		11f		
130				144			145	148				14b	146	14a	149		
140					0011	0110	1100		0111	1110	1101	1011					CL
150	13f				147				14c								

The word 0010 has become green, thus closing its seven lists and making 0001 red. The logic of Algorithm C has also made 1001 red. Hence 0001 and 1001 have been deleted from the open lists in which they formerly appeared (see exercise 42).

pointer `MEM[30]` $=$ `1f` $= 20-1$ marks that list as closed. (Those list elements actually do still appear in `MEM` locations 200 through 204, just as they did in Table 1. But there's no need to look at that list while any word of the form $0{*}{*}{*}$ is green.)

A general mechanism for doing and undoing. We're almost ready to finalize the details of Algorithm C and to get on with the search for commafree codes, but a big problem still remains: The state of computation at every level of the search involves all of the marvelous lists that we've just specified, and those lists aren't tiny. They occupy more than 5000 cells of `MEM` when $m = 4$, and they can change substantially from level to level.

We could make a new copy of the entire state, whenever we advance to a new node of the search tree. But that's a bad idea, because we don't want to perform thousands of memory accesses per node. A much better strategy would be to stick with a single instance of `MEM`, and to update and downdate the lists as the search progresses, if we could only think of a simple way to do that.

And we're in luck: There *is* such a way, first formulated by R. W. Floyd in his classic paper "Nondeterministic algorithms" [*JACM* **14** (1967), 636–644]. Floyd's original idea, which required a special compiler to generate forward and backward versions of every program step, can in fact be greatly simplified when all of the changes in state are confined to a single `MEM` array. All we need to do is to replace every assignment operation of the form '`MEM[`a`]` $\leftarrow v$' by the slightly more cumbersome operation

$$\text{store}(a, v): \quad \text{Set } \texttt{UNDO[}u\texttt{]} \leftarrow \bigl(a, \texttt{MEM[}a\texttt{]}\bigr),\ \texttt{MEM[}a\texttt{]} \leftarrow v, \text{ and } u \leftarrow u + 1. \qquad (24)$$

Here UNDO is a sequential stack that holds (address, value) pairs; in our application we could say '$\texttt{UNDO[}u\texttt{]} \leftarrow (a \ll 16) + \texttt{MEM[}a\texttt{]}$', because the cell addresses and values never exceed 16 bits. Of course we'll also need to check that the stack pointer u doesn't get too large, if the number of assignments has no a priori limit.

Later on, when we want to undo all changes to MEM that were made after the time when u had reached a particular value u_0, we simply do this:

$$\text{unstore}(u_0): \quad \text{While } u > u_0,\ \text{ set } u \leftarrow u - 1,$$
$$(a, v) \leftarrow \texttt{UNDO[}u\texttt{]},\ \text{ and } \texttt{MEM[}a\texttt{]} \leftarrow v. \qquad (25)$$

In our application the unstacking operation '$(a, v) \leftarrow \texttt{UNDO[}u\texttt{]}$' here could be implemented by saying '$a \leftarrow \texttt{UNDO[}u\texttt{]} \gg 16,\ v \leftarrow \texttt{UNDO[}u\texttt{]}\ \&\ {}^{\#}\texttt{ffff}$'.

A useful refinement of this reversible-memory technique is often advantageous, based on the idea of "stamping" that is part of the folklore of programming. It puts only one item on the UNDO stack when the same memory address is updated more than once in the same round.

$$\text{store}(a, v): \quad \text{If } \texttt{STAMP[}a\texttt{]} \neq \sigma, \text{ set } \texttt{STAMP[}a\texttt{]} \leftarrow \sigma,$$
$$\texttt{UNDO[}u\texttt{]} \leftarrow \bigl(a, \texttt{MEM[}a\texttt{]}\bigr), \text{ and } u \leftarrow u + 1.$$
$$\text{Then set } \texttt{MEM[}a\texttt{]} \leftarrow v. \qquad (26)$$

Here STAMP is an array with one entry for each address in MEM. It's initially all zero, and σ is initially 1. Whenever we come to a fallback point, where the current stack pointer will be remembered as the value u_0 for some future undoing, we "bump" the current stamp by setting $\sigma \leftarrow \sigma + 1$. Then (26) will continue to do the right thing. (In programs that run for a long time, we must be careful when integer overflow causes σ to be bumped to zero; see exercise 43.)

Notice that the combination of (24) and (25) will perform five memory accesses for each assignment and its undoing. The combination of (26) and (25) will cost seven mems for the first assignment to MEM[a], but only two mems for every subsequent assignment to the same address. So (26) wins, if multiple assignments exceed one-time-only assignments.

Backtracking through commafree codes. OK, we're now equipped with enough basic knowhow to write a pretty good backtrack program for the problem of generating all commafree four-letter codes.

Algorithm C below incorporates one more key idea, which is a lookahead mechanism that is specific to commafree backtracking; we'll call it the "poison list." Every item on the poison list is a pair, consisting of a suffix and a prefix that the commafree rule forbids from occurring together. Every green word $x_1x_2x_3x_4$ — that is, every word that will be a final codeword in the current branch of our backtrack search — contributes three items to the poison list, namely

$$(*x_1x_2x_3, x_4{*}{*}{*}), \quad ({*}{*}x_1x_2, x_3x_4{*}{*}), \quad \text{and} \quad ({*}{*}{*}x_1, x_2x_3x_4*). \qquad (27)$$

UNDO
stack
reversible-memory
stamping
fallback point
bump
overflow
lookahead
poison list

If there's a green word on both sides of a poison list entry, we're dead: The commafree condition fails, and we must backtrack. If there's a green word on one side but not the other, we can kill off all blue words on the other side by making them red. And if either side of a poison list entry corresponds to an *empty* list, we can remove this entry from the poison list because it will never affect the outcome. (Blue words become red or green, but red words stay red.)

inverse list
Laxdal
Jiggs
stamping++

For example, consider the transition from Table 1 to Table 2. When word 0010 becomes green, the poison list receives its first three items:

$$(*001, 0***), \quad (**00, 10**), \quad (***0, 010*).$$

The first of these kills off the $*001$ list, because $0***$ contains the green word 0010. That makes 1001 red. The last of these, similarly, kills off the $010*$ list; but that list is empty when $m = 2$. The poison list now reduces to a single item, $(**00, 10**)$, which remains poisonous because list $**00$ contains the blue word 1100 and $10**$ contains the blue word 1011.

We'll maintain the poison list at the end of MEM, following the CL lists. It obviously will contain at most $3(m^4 - m^2)/4$ entries, and in fact it usually turns out to be quite small. No inverse list is required; so we shall adopt the simple method of (17) and (18), but with two cells per entry so that TAIL will change by ± 2 instead of by ± 1. The value of TAIL will be stored in MEM at key times so that temporary changes to it can be undone.

The case $m = 4$, in which each codeword consists of four *quaternary* digits $\{0, 1, 2, 3\}$, is particularly interesting, because an early backtrack program by Lee Laxdal found that no such commafree code can make use of all 60 of the cycle classes [0001], [0002], ..., [2333]. [See B. H. Jiggs, *Canadian Journal of Math.* **15** (1963), 178–187.] Laxdal's program also reportedly showed that at least three of those classes must be omitted; and it found several valid 57-word sets. Further details were never published, because the proof that 58 codewords are impossible depended on what Jiggs called a "quite time-consuming" computation.

Because size 60 is impossible, our algorithm cannot simply assume that a move such as 1001 is forced when the other words 0011, 0110, 1100 of its class have been ruled out. We must also consider the possibility that class [0011] is entirely absent from the code. Such considerations add an interesting further twist to the problem, and Algorithm C describes one way to cope with it.

Algorithm C (*Four-letter commafree codes*). Given an alphabet size $m \le 7$ and a goal g in the range $L - m(m-1) \le g \le L$, where $L = (m^4 - m^2)/4$, this algorithm finds all sets of g four-letter words that are commafree and include either 0001 or 0010. It uses an array MEM of $M = \lfloor 23.5m^4 \rfloor$ 16-bit numbers, as well as several more auxiliary arrays: ALF of size 16^3m; STAMP of size M; X, C, S, and U of size $L + 1$; FREE and IFREE of size L; and a sufficiently large array called UNDO whose maximum size is difficult to guess.

C1. [Initialize.] Set ALF$[(abcd)_{16}] \leftarrow (abcd)_m$ for $0 \le a, b, c, d < m$. Set STAMP$[k] \leftarrow 0$ for $0 \le k < M$ and $\sigma \leftarrow 0$. Put the initial prefix, suffix, and class lists into MEM, as in Table 1. Also create an empty poison list by setting MEM[PP] $\leftarrow$ POISON, where POISON $= 22m^4$ and PP $=$ POISON $- 1$. Set FREE[k] $\leftarrow$ IFREE[k] $\leftarrow k$ for $0 \le k < L$. Then set $l \leftarrow 1$, $x \leftarrow$ #0001, $c \leftarrow 0$, $s \leftarrow L - g$, $f \leftarrow L$, $u \leftarrow 0$, and go to step C3. (Variable l is the level, x is a trial word, c is its class, s is the "slack," f is the number of free classes, and u is the size of the UNDO stack.)

C2. [Enter level l.] If $l > L$, visit the solution $x_1 \ldots x_L$ and go to C6. Otherwise choose a candidate word x and class c as described in exercise 44.

C3. [Try the candidate.] Set U[l] $\leftarrow u$ and $\sigma \leftarrow \sigma + 1$. If $x < 0$, go to C6 if $s = 0$ or $l = 1$, otherwise set $s \leftarrow s - 1$. If $x \ge 0$, update the data structures to make x green, as described in exercise 45, escaping to C5 if trouble arises.

C4. [Make the move.] Set X[l] $\leftarrow x$, C[l] $\leftarrow c$, S[l] $\leftarrow s$, $p \leftarrow$ IFREE[c], $f \leftarrow f - 1$. If $p \ne f$, set $y \leftarrow$ FREE[f], FREE[p] $\leftarrow y$, IFREE[y] $\leftarrow p$, FREE[f] $\leftarrow c$, IFREE[c] $\leftarrow f$. (This is (23).) Then set $l \leftarrow l + 1$ and go to C2.

C5. [Try again.] While $u >$ U[l], set $u \leftarrow u - 1$ and MEM[UNDO[u] $\gg$ 16] $\leftarrow$ UNDO[u] & #ffff. (Those operations restore the previous state, as in (25).) Then $\sigma \leftarrow \sigma + 1$ and redden x (see exercise 45). Go to C2.

C6. [Backtrack.] Set $l \leftarrow l - 1$, and terminate if $l = 0$. Otherwise set $x \leftarrow$ X[l], $c \leftarrow$ C[l], $f \leftarrow f - 1$. If $x < 0$, repeat this step (class c was omitted from the code). Otherwise set $s \leftarrow$ S[l] and go back to C5. ▮

Exercises 44 and 45 provide the instructive details that flesh out this skeleton.

Algorithm C needs just 13, 177, and 2380 megamems to prove that no solutions exist for $m = 4$ when g is 60, 59, and 58. It needs about 22800 megamems to find the 1152 solutions for $g = 57$; see exercise 47. There are roughly (14, 240, 3700, 38000) thousand nodes in the respective search trees, with most of the activity taking place on levels 30 ± 10. The height of the UNDO stack never exceeds 2804, and the poison list never contains more than 12 entries at a time.

Running time estimates. Backtrack programs are full of surprises. Sometimes they produce instant answers to a supposedly difficult problem. But sometimes they spin their wheels endlessly, trying to traverse an astronomically large search tree. And sometimes they deliver results just about as fast as we might expect.

Fortunately, we needn't sit in the dark. There's a simple Monte Carlo algorithm by which we can often tell in advance whether or not a given backtrack strategy will be feasible. This method, based on random sampling, can actually be worked out by hand *before* writing a program, in order to help decide whether to invest further time while following a particular approach. In fact, the very act of carrying out this pleasant pencil-and-paper method often suggests useful cutoff strategies and/or data structures that will be valuable later when a program is being written. For example, the author developed Algorithm C above after first doing some armchair experiments with random choices of potential commafree codewords; these dry runs revealed that a family of lists such as those in Tables 1 and 2 would be quite helpful when making further choices.

To illustrate the method, let's consider the n queens problem again, as represented in Algorithm B* above. When $n = 8$, we can obtain a decent "ballpark

Running time
estimates of run time–
Monte Carlo algorithm
random sampling
pencil-and-paper method
cutoff strategies
data structures
author

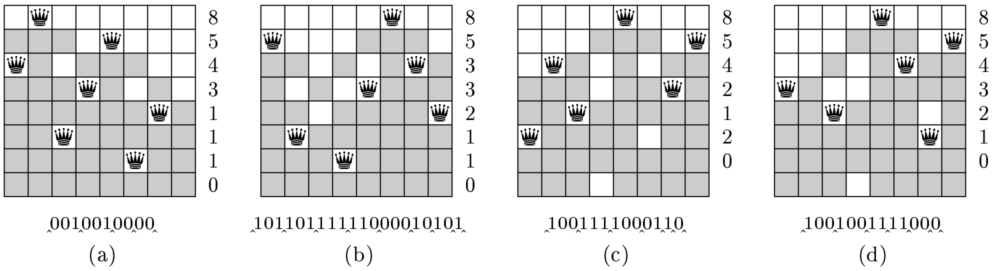

Fig. 69. Four random attempts to solve the 8 queens problem. Such experiments help to estimate the size of the backtrack tree in Fig. 68. The branching degrees are shown at the right of each diagram, while the random bits used for sampling appear below. Cells have been shaded in gray if they are attacked by one or more queens in earlier rows.

search tree
root node
degree
pi as source+
rejection method
e, as source
phi, as source
gamma, as source
random variables
cost function

estimate" of the size of Fig. 68 by examining only a few random paths in that search tree. We start by writing down the number $D_1 \leftarrow 8$, because there are eight ways to place the queen in row 1. (In other words, the root node of the search tree has degree 8.) Then we use a source of random numbers — say the binary digits of $\pi \bmod 1 = (.001001000011\ldots)_2$ — to select one of those placements. Eight choices are possible, so we look at three of those bits; we shall set $X_1 \leftarrow 2$, because 001 is the second of the eight possibilities (000, 001, ..., 111).

Given $X_1 = 2$, the queen in row 2 can't go into columns 1, 2, or 3. Hence five possibilities remain for X_2, and we write down $D_2 \leftarrow 5$. The next three bits of π lead us to set $X_2 \leftarrow 5$, since 5 is the second of the available columns (4, 5, 6, 7, 8) and 001 is the second value of (000, 001, ..., 100). If π had continued with 101 or 110 or 111 instead of 001, we would incidentally have used the "rejection method" of Section 3.4.1 and moved to the next three bits; see exercise 49.

Continuing in this way leads to $D_3 \leftarrow 4$, $X_3 \leftarrow 1$; then $D_4 \leftarrow 3$, $X_4 \leftarrow 4$. (Here we used the two bits 00 to select X_3, and the next two bits 00 to select X_4.) The remaining branches are forced: $D_5 \leftarrow 1$, $X_5 \leftarrow 7$; $D_6 \leftarrow 1$, $X_6 \leftarrow 3$; $D_7 \leftarrow 1$, $X_7 \leftarrow 6$; and we're stuck when we reach level 8 and find $D_8 \leftarrow 0$.

These sequential random choices are depicted in Fig. 69(a), where we've used them to place each queen successively into an unshaded cell. Parts (b), (c), and (d) of Fig. 69 correspond in the same way to choices based on the binary digits of $e \bmod 1$, $\phi \bmod 1$, and $\gamma \bmod 1$. Exactly 10 bits of π, 20 bits of e, 13 bits of ϕ, and 13 bits of γ were used to generate these examples.

In this discussion the notation D_k stands for a branching degree, not for a domain of values. We've used uppercase letters for the numbers D_1, X_1, D_2, etc., because those quantities are random variables. Once we've reached $D_l = 0$ at some level, we're ready to estimate the overall cost, by implicitly assuming that the path we've taken is representative of *all* root-to-leaf paths in the tree.

The cost of a backtrack program can be assessed by summing the individual amounts of time spent at each node of the search tree. Notice that every node on level l of that tree can be labeled uniquely by a sequence $x_1 \ldots x_{l-1}$, which defines the path from the root to that node. Thus our goal is to estimate the sum of all $c(x_1 \ldots x_{l-1})$, where $c(x_1 \ldots x_{l-1})$ is the cost associated with node $x_1 \ldots x_{l-1}$.

subtree
size of the tree
standard deviation

For example, the four queens problem is represented by the search tree (4), and its cost is the sum of 17 individual costs

$$c() + c(1) + c(13) + c(14) + c(142) + c(2) + c(24) + \cdots + c(413) + c(42). \quad (28)$$

If $C(x_1 \ldots x_l)$ denotes the total cost of the subtree rooted at $x_1 \ldots x_l$, then

$$C(x_1 \ldots x_l) = c(x_1 \ldots x_l) + C(x_1 \ldots x_l x_{l+1}^{(1)}) + \cdots + C(x_1 \ldots x_l x_{l+1}^{(d)}) \quad (29)$$

when the choices for x_{l+1} at node $x_1 \ldots x_l$ are $\{x_{l+1}^{(1)}, \ldots, x_{l+1}^{(d)}\}$. For instance in (4) we have $C(1) = c(1) + C(13) + C(14)$; $C(13) = c(13)$; and $C() = c() + C(1) + C(2) + C(3) + C(4)$ is the overall cost (28).

In these terms a Monte Carlo estimate for $C()$ is extremely easy to compute:

Theorem E. *Given D_1, X_1, D_2, X_2, … as above, the cost of backtracking is*

$$C() = \mathrm{E}\bigl(c() + D_1(c(X_1) + D_2(c(X_1X_2) + D_3(c(X_1X_2X_3) + \cdots))))\bigr). \quad (30)$$

Proof. Node $x_1 \ldots x_l$, with branch degrees d_1, …, d_l above it, is reached with probability $1/d_1 \ldots d_l$; so it contributes $d_1 \ldots d_l c(x_1 \ldots x_l)/d_1 \ldots d_l = c(x_1 \ldots x_l)$ to the expected value in this formula. ▮

For example, the tree (4) has six root-to-leaf paths, and they occur with respective probabilities 1/8, 1/8, 1/4, 1/4, 1/8, 1/8. The first one contributes 1/8 times $c() + 4(c(1) + 2(c(13)))$, namely $c()/8 + c(1)/2 + c(13)$, to the expected value. The second contributes $c()/8 + c(1)/2 + c(14) + c(142)$; and so on.

A special case of Theorem E, with all $c(x_1 \ldots x_l) = 1$, tells us how to estimate the total size of the tree, which is often a crucial quantity:

Corollary E. *The number of nodes in the search tree, given D_1, D_2, …, is*

$$\mathrm{E}(1 + D_1 + D_1D_2 + \cdots) = \mathrm{E}\bigl(1 + D_1\bigl(1 + D_2(1 + D_3(1 + \cdots))\bigr)\bigr). \quad (31)$$

For example, Fig. 69 gives us four estimates for the size of the tree in Fig. 68, using the numbers D_j at the right of each 8×8 diagram. The estimate from Fig. 69(a) is $1+8\bigl(1+5\bigl(1+4(1+3(1+1(1+1(1+1))))\bigr)\bigr) = 2129$; and the other three are respectively 2689, 1489, 2609. None of them is extremely far from the true number, 2057, although we can't expect to be so lucky all the time.

The detailed study in exercise 53 shows that the estimate (31) in the case of 8 queens turns out to be quite well behaved:

$$\bigl(\text{min } 489, \quad \text{ave } 2057, \quad \text{max } 7409, \quad \text{dev } \sqrt{1146640} \approx 1071\bigr). \quad (32)$$

The analogous problem for 16 queens has a much less homogeneous search tree:

$$\bigl(\text{min } 2597105, \text{ ave } 1141190303, \text{ max } 131048318769, \text{ dev} \approx 1234000000\bigr). \quad (33)$$

Still, this standard deviation is roughly the same as the mean, so we'll usually guess the correct order of magnitude. (For example, ten independent experiments predicted .632, .866, .237, 1.027, 4.006, .982, .143, .140, 3.402, and .510 billion nodes, respectively. The mean of these is 1.195.) A thousand trials with $n = 64$ suggest that the problem of 64 queens will have about 3×10^{65} nodes in its tree.

binomial tree
estimating solutions

Let's formulate this estimation procedure precisely, so that it can be performed conveniently by machine as well as by hand:

Algorithm E (*Estimated cost of backtrack*). Given domains D_k and properties P_l as in Algorithm B, together with node costs $c(x_1 \ldots x_l)$ as above, this algorithm computes the quantity S whose expected value is the total cost $C()$ in (30). It uses an auxiliary array $y_1 y_2 \ldots$ whose size should be $\geq \max(|D_1|, \ldots, |D_n|)$.

E1. [Initialize.] Set $l \leftarrow D \leftarrow 1$, $S \leftarrow 0$, and initialize any data structures needed.

E2. [Enter level l.] (At this point $P_{l-1}(X_1, \ldots, X_{l-1})$ holds.) Set $S \leftarrow S + D \cdot c(X_1 \ldots X_{l-1})$. If $l > n$, terminate the algorithm. Otherwise set $d \leftarrow 0$ and set $x \leftarrow \min D_l$, the smallest element of D_l.

E3. [Test x.] If $P_l(X_1, \ldots, X_{l-1}, x)$ holds, set $y_d \leftarrow x$ and $d \leftarrow d + 1$.

E4. [Try again.] If $x \neq \max D_l$, set x to the next larger element of D_l and return to step E3.

E5. [Choose and try.] If $d = 0$, terminate. Otherwise set $D \leftarrow D \cdot d$ and $X_l \leftarrow y_I$, where I is a uniformly random integer in $\{0, \ldots, d-1\}$. Update the data structures to facilitate testing P_{l+1}, set $l \leftarrow l + 1$, and go back to E2. ▮

Although Algorithm E looks rather like Algorithm B, it never backtracks.

Of course we can't expect this algorithm to give decent estimates in cases where the backtrack tree is wildly erratic. The *expected* value of S, namely $\mathrm{E}\, S$, is indeed the true cost; but the *probable* values of S might be quite different.

An extreme example of bad behavior occurs if property P_l is the simple condition '$x_1 > \cdots > x_l$', and all domains are $\{1, \ldots, n\}$. Then there's only one solution, $x_1 \ldots x_n = n \ldots 1$; and backtracking is a particularly stupid way to find it!

The search tree for this somewhat ridiculous problem is, nevertheless, quite interesting. It is none other than the binomial tree T_n of Eq. 7.2.1.3–(21), which has $\binom{n}{l}$ nodes on level $l + 1$ and 2^n nodes in total. If we set all costs to 1, the expected value of S is therefore $2^n = e^{n \ln 2}$. But exercise 52 proves that S will almost always be much smaller, less than $e^{(\ln n)^2 \ln \ln n}$. Furthermore the average value of l when Algorithm E terminates with respect to T_n is only $H_n + 1$. When $n = 100$, for example, the probability that $l \geq 20$ on termination is only 0.0000000027, while the vast majority of the nodes are near level 51.

Many refinements of Algorithm E are possible. For example, exercise 54 shows that the choices in step E5 need not be uniform. We shall discuss improved estimation techniques in Section 7.2.2.9, after having seen numerous examples of backtracking in practice.

***Estimating the number of solutions.** Sometimes we know that a problem has more solutions than we could ever hope to generate, yet we still want to know roughly how many there are. Algorithm E will tell us the approximate number, in cases where the backtrack process never reaches a dead end — that is, if it never terminates with $d = 0$ in step E5. There may be another criterion for successful termination in step E2 even though l might still be $\leq n$. The expected final value of D is exactly the total number of solutions, because every solution $X_1 \ldots X_l$ constructed by the algorithm is obtained with probability $1/D$.

king
chessboard
simple paths
paths, simple
π for guidance
statistics
sample variance
variance
error bars
discarded data
Diaconis
Chatterjee

For example, suppose we want to know the number of different paths by which a king can go from one corner of a chessboard to the opposite corner, without revisiting any square. One such path, chosen at random using the bits of π for guidance as we did in Fig. 69(a), is shown here. Starting in the upper left corner, we have 3 choices for the first move. Then, after moving to the right, there are 4 choices for the second move. And so on. We never make a move that would disconnect us from the goal; in particular, two of the moves are actually forced. (Exercise 58 explains one way to avoid fatal mistakes.)

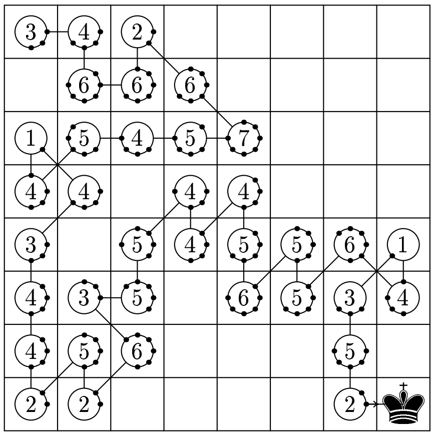

The probability of obtaining this particular path is exactly $\frac{1}{3}\frac{1}{4}\frac{1}{6}\frac{1}{6}\frac{1}{2}\frac{1}{6}\frac{1}{7}\cdots\frac{1}{2} = 1/D$, where $D = 3\times4\times6\times6\times2\times6\times7\times\cdots\times2 = 1^2\cdot2^4\cdot3^4\cdot4^{10}\cdot5^9\cdot6^6\cdot7^1 \approx 8.7\times10^{20}$. Thus we can reasonably guess, at least tentatively, that there are 10^{21} such paths, more or less.

Of course that guess, based on a single random sample, rests on very shaky grounds. But we know that the average value $M_N = (D^{(1)}+\cdots+D^{(N)})/N$ of N guesses, in N independent experiments, will almost surely approach the correct number.

How large should N be, before we can have any confidence in the results? The actual values of D obtained from random king paths tend to vary all over the map. Figure 70 plots typical results, as N varies from 1 to 10000. For each value of N we can follow the advice of statistics textbooks and calculate the sample variance $V_N = S_N/(N-1)$ as in Eq. 4.2.2–(16); then $M_N \pm \sqrt{V_N/N}$ is the textbook estimate. The top diagram in Fig. 70 shows these "error bars" in gray, surrounding black dots for M_N. This sequence M_N does appear to settle down after N reaches 3000 or so, and to approach a value near 5×10^{25}. That's much higher than our first guess, but it has lots of evidence to back it up.

On the other hand, the bottom chart in Fig. 70 shows the distribution of the *logarithms* of the 10000 values of D that were used to make the top chart. Almost half of those values were totally negligible — less than 10^{20}. About 75% of them were less than 10^{24}. But some of them* exceeded 10^{28}. Can we really rely on a result that's based on such chaotic behavior? Is it really right to throw away most of our data and to trust almost entirely on observations that were obtained from comparatively few rare events?

Yes, we're okay! Some of the justification appears in exercise MPR–124, which is based on theoretical work by P. Diaconis and S. Chatterjee. In the paper cited with that exercise, they defend a simple measure of quality,

$$Q_N \;=\; \max(D^{(1)},\ldots,D^{(N)})/(NM_N) \;=\; \frac{\max(D^{(1)},\ldots,D^{(N)})}{D^{(1)}+\cdots+D^{(N)}}, \qquad (34)$$

* Four of the actual values that led to Fig. 70 were larger than 10^{28}; the largest, $\approx 2.1\times10^{28}$, came from a path of length 57. The smallest estimate, 19361664, came from a path of length 10.

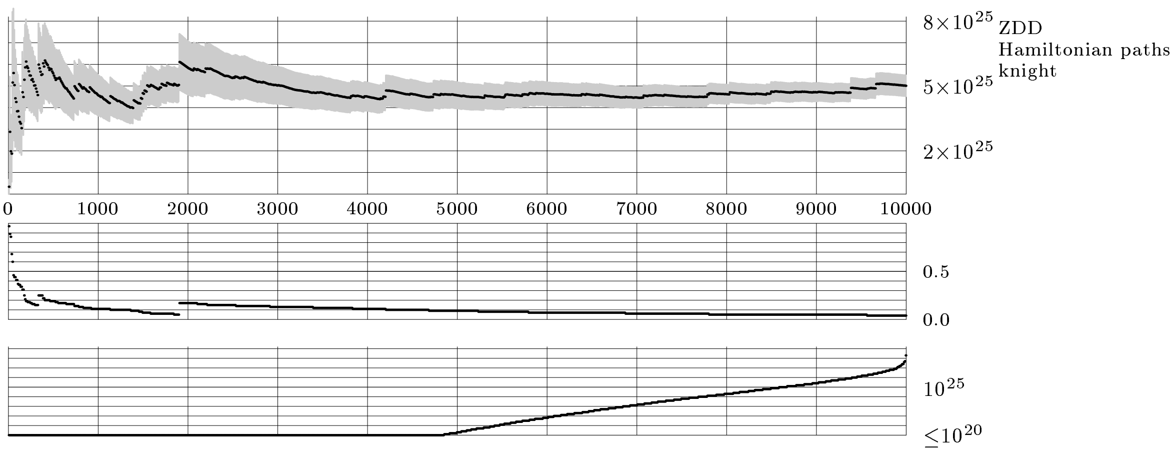

Fig. 70. Estimates of the number of king paths, based on up to 10000 random trials. The middle graph shows the corresponding quality measures of Eq. (34). The lower graph shows the *logarithms* of the individual estimates $D^{(k)}$, after they've been sorted.

arguing that a reasonable policy in most experiments such as these is to stop sampling when Q_N gets small. (Values of this statistic Q_N have been plotted in the middle of Fig. 70.)

Furthermore we can estimate other properties of the solutions to a backtrack problem, instead of merely counting those solutions. For example, the expected value of lD on termination of the random king's path algorithm is the total *length* of such paths. The data underlying Fig. 70 suggests that this total is $(2.66 \pm .14) \times 10^{27}$; hence the average path length appears to be about 53. The samples also indicate that about 34% of the paths pass through the center; about 46% touch the upper right corner; about 22% touch both corners; and about 7% pass through the center and both corners.

For this particular problem we don't actually need to rely on estimates, because the ZDD technology of Section 7.1.4 allows us to compute the *true* values. (See exercise 59.) The total number of simple corner-to-corner king paths on a chessboard is exactly 50,819,542,770,311,581,606,906,543; this value lies almost within the error bars of Fig. 70 for all $N \geq 250$, except for a brief interval near $N = 1400$. And the total length of all these paths turns out to be exactly 2,700,911,171,651,251,701,712,099,831, which is a little higher than our estimate. The true average length is therefore ≈ 53.15. The true probabilities of hitting the center, a given corner, both corners, and all three of those spots are respectively about 38.96%, 50.32%, 25.32%, and 9.86%.

The total number of corner-to-corner king paths of the maximum length, 63, is 2,811,002,302,704,446,996,926. This is a number that can *not* be estimated well by a method such as Algorithm E without additional heuristics.

The analogous problem for corner-to-corner *knight* paths, of any length, lies a bit beyond ZDD technology because many more ZDD nodes are needed. Using Algorithm E we can estimate that there are about $(8.6 \pm 1.2) \times 10^{19}$ such paths.

Factoring the problem. Imagine an instance of backtracking that is equivalent to solving two *independent* subproblems. For example, we might be looking for all sequences $x = x_1x_2 \ldots x_n$ that satisfy $P_n(x_1, x_2, \ldots, x_n) = F(x_1, x_2, \ldots, x_n)$, where

$$F(x_1, x_2, \ldots, x_n) \;=\; G(x_1, \ldots, x_k) \wedge H(x_{k+1}, \ldots, x_n). \tag{35}$$

Then the size of the backtrack tree is essentially the *product* of the tree sizes for G and for H, even if we use dynamic ordering. Hence it's obviously foolish to apply the general setup of (1) and (2). We can do much better by finding all solutions to G first, then finding all solutions to H, thereby reducing the amount of computation to the *sum* of the tree sizes. Again we've divided and conquered, by factoring the compound problem (35) into separate subproblems.

We discussed a less obvious application of problem factorization near the beginning of Chapter 7, in connection with latin squares: Recall that E. T. Parker sped up the solution of 7–(6) by more than a dozen orders of magnitude, when he discovered 7–(7) by essentially factoring 7–(6) into ten subproblems whose solutions could readily be combined.

In general, each solution x to some problem F often implies the existence of solutions $x^{(p)} = \phi_p(x)$ to various simpler problems F_p that are "homomorphic images" of F. And if we're lucky, the solutions to those simpler problems can be combined and "lifted" to a solution of the overall problem. Thus it pays to be on the lookout for such simplifications.

Let's look at another example. F. A. Schossow invented a tantalizing puzzle [*U.S. Patent 646463* (3 April 1900)] that went viral in 1967 when a marketing genius decided to rename it "Instant Insanity." The problem is to take four cubes such as

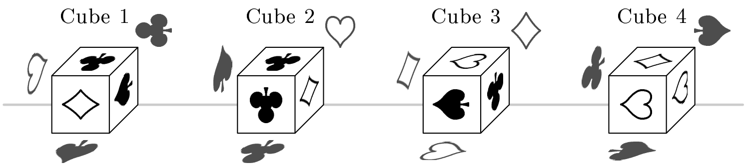

(36)

where each face has been marked in one of four ways, and to arrange them in a row so that all four markings appear on the top, bottom, front, and back sides. The placement in (36) is incorrect, because there are two $\clubsuit$s (and no $\spadesuit$) on top. But we get a solution if we rotate each cube by 90°.

There are 24 ways to place each cube, because any of the six faces can be on top and we can rotate four ways while keeping the top unchanged. So the total number of placements is $24^4 = 331776$. But this problem can be factored in an ingenious way, so that all solutions can be found quickly by hand! [See F. de Carteblanche, *Eureka* **9** (1947), 9–11.] The idea is that any solution to the puzzle gives us two each of $\{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\}$, if we look only at the top and bottom or only at the front and back. That's a much easier problem to solve.

For this purpose a cube can be characterized by its three pairs of markings on opposite faces; in (36) these face-pairs are respectively

$$\{\clubsuit\spadesuit, \clubsuit\diamondsuit, \spadesuit\heartsuit\}, \quad \{\clubsuit\clubsuit, \clubsuit\heartsuit, \spadesuit\diamondsuit\}, \quad \{\heartsuit\heartsuit, \spadesuit\diamondsuit, \clubsuit\diamondsuit\}, \quad \{\spadesuit\diamondsuit, \spadesuit\heartsuit, \clubsuit\heartsuit\}. \tag{37}$$

factorization of problems–
independent subproblems
backtrack tree
tree sizes
dynamic ordering
divide and conquer paradigm
latin squares
Parker
homomorphic images
lifted
Schossow
Patent
Tantalizer, see Instant InsanityTM
Armbruster
Instant Insanity
Carteblanche

Which of the $3^4 = 81$ ways to choose one face-pair from each cube will give us $\{\clubsuit, \clubsuit, \diamondsuit, \diamondsuit, \heartsuit, \heartsuit, \spadesuit, \spadesuit\}$? They can all be discovered in a minute or two, by listing the nine possibilities for cubes $(1,2)$ and the nine for $(3,4)$. We get just three,

$$(\clubsuit\diamondsuit, \clubsuit\heartsuit, \spadesuit\diamondsuit, \spadesuit\heartsuit), \quad (\spadesuit\heartsuit, \clubsuit\heartsuit, \clubsuit\diamondsuit, \spadesuit\diamondsuit), \quad (\spadesuit\heartsuit, \spadesuit\diamondsuit, \clubsuit\diamondsuit, \clubsuit\heartsuit). \tag{38}$$

Notice furthermore that each solution can be "halved" so that one each of $\{\clubsuit, \diamondsuit, \heartsuit, \spadesuit\}$ appears on both sides, by swapping face-pairs; we can change (38) to

$$(\diamondsuit\clubsuit, \clubsuit\heartsuit, \spadesuit\diamondsuit, \heartsuit\spadesuit), \quad (\heartsuit\spadesuit, \clubsuit\heartsuit, \diamondsuit\clubsuit, \spadesuit\diamondsuit), \quad (\heartsuit\spadesuit, \spadesuit\diamondsuit, \diamondsuit\clubsuit, \clubsuit\heartsuit). \tag{39}$$

Each of these solutions to the opposite-face subproblem can be regarded as a 2-regular graph, because every vertex of the multigraph whose edges are (say) $\diamondsuit - \clubsuit$, $\clubsuit - \heartsuit$, $\spadesuit - \diamondsuit$, $\heartsuit - \spadesuit$ has exactly two neighbors.

A solution to "Instant Insanity" will give us *two* such 2-regular factors, one for top-and-bottom and one for front-and-back. Furthermore those two factors will have disjoint edges: We can't use the same face-pair in both. Therefore problem (36) can be solved only by using the first and third factor in (39).

Conversely, whenever we have two disjoint 2-regular graphs, we can always use them to position the cubes as desired, thus "lifting" the factors to a solution of the full problem.

Exercise 67 illustrates another kind of problem factorization. We can conveniently think of each subproblem as a "relaxation" of constraints.

Historical notes. The origins of backtrack programming are obscure. Equivalent ideas must have occurred to many people, yet there was hardly any reason to write them down until computers existed. We can be reasonably sure that James Bernoulli used such principles in the 17th century, when he successfully solved the "Tot tibi sunt dotes" problem that had eluded so many others (see Section 7.2.1.7), because traces of the method exist in his exhaustive list of solutions.

Backtrack programs typically traverse the tree of possibilities by using what is now called depth-first search, a general graph exploration procedure that Édouard Lucas credited to a student named Trémaux [*Récréations Mathématiques* **1** (Paris: Gauthier-Villars, 1882), 47–50].

The eight queens problem was first proposed by Max Bezzel [*Schachzeitung* **3** (1848), 363; **4** (1849), 40] and by Franz Nauck [*Illustrirte Zeitung* **14**, 361 (1 June 1850), 352; **15**, 377 (21 September 1850), 182], perhaps independently. C. F. Gauss saw the latter publication, and wrote several letters about it to his friend H. C. Schumacher. Gauss's letter of 27 September 1850 is especially interesting, because it explained how to find all the solutions by backtracking—which he called 'Tatonniren', from a French term meaning "to feel one's way." He also listed the lexicographically first solutions of each equivalence class under reflection and rotation: 15863724, 16837425, 24683175, 25713864, 25741863, 26174835, 26831475, 27368514, 27581463, 35281746, 35841726, and 36258174.

Computers arrived a hundred years later, and people began to use them for combinatorial problems. The time was therefore ripe for backtracking to be described as a general technique, and Robert J. Walker rose to the occasion [*Proc. Symposia in Applied Math.* **10** (1960), 91–94]. His brief note introduced

2-regular graph
disjoint
lifting
relaxation
Bernoulli
Tot tibi
depth-first search
Lucas
Trémaux
eight queens problem
Bezzel
Nauck
Gauss
Schumacher
lexicographically
Walker

Algorithm W in machine-oriented form, and mentioned that the procedure could readily be extended to find variable-length patterns $x_1 \ldots x_n$ where n is not fixed.

The next milestone was a paper by Solomon W. Golomb and Leonard D. Baumert [*JACM* **12** (1965), 516–524], who formulated the general problem carefully and presented a variety of examples. In particular, they discussed the search for maximum commafree codes, and noted that backtracking can be used to find successively better and better solutions to combinatorial optimization problems. They introduced certain kinds of lookahead, as well as the important idea of dynamic ordering by branching on variables with the fewest remaining choices.

Backtrack methods allow special cutoffs when applied to integer programming problems [see E. Balas, *Operations Research* **13** (1965), 517–546]. A. M. Geoffrion simplified and extended that work, calling it "implicit enumeration" because many cases aren't enumerated explicitly [*SIAM Rev.* **9** (1967), 178–190].

Other noteworthy early discussions of backtrack programming appear in Mark Wells's book *Elements of Combinatorial Computing* (1971), Chapter 4; in a survey by J. R. Bitner and E. M. Reingold, *CACM* **18** (1975), 651–656; and in the Ph.D. thesis of John Gaschnig [Report CMU-CS-79-124 (Carnegie Mellon University, 1979), Chapter 4]. Gaschnig introduced techniques of "backmarking" and "backjumping" that we shall discuss later.

Monte Carlo estimates of the cost of backtracking were first described briefly by M. Hall, Jr., and D. E. Knuth in *Computers and Computing, AMM* **72**, 2, part 2, Slaught Memorial Papers No. 10 (February 1965), 21–28. Knuth gave a much more detailed exposition a decade later, in *Math. Comp.* **29** (1975), 121–136. Such methods can be considered as special cases of so-called "importance sampling"; see J. M. Hammersley and D. C. Handscomb, *Monte Carlo Methods* (London: Methuen, 1964), 57–59. Studies of random self-avoiding walks such as the king paths discussed above were inaugurated by M. N. Rosenbluth and A. W. Rosenbluth, *J. Chemical Physics* **23** (1955), 356–359.

Backtrack applications are nicely adaptable to parallel programming, because different parts of the search tree are often completely independent of each other; thus disjoint subtrees can be explored on different machines, with a minimum of interprocess communication. Already in 1964, D. H. Lehmer explained how to subdivide a problem so that two computers of different speeds could work on it simultaneously and finish at the same time. The problem that he considered had a search tree of known shape (see Theorem 7.2.1.3L); but we can do essentially similar load balancing even in much more complicated situations, by using Monte Carlo estimates of the subtree sizes. Although many ideas for parallelizing combinatorial searches have been developed over the years, such techniques are beyond the scope of this book. Readers can find a nice introduction to a fairly general approach in the paper by R. Finkel and U. Manber, *ACM Transactions on Programming Languages and Systems* **9** (1987), 235–256.

M. Alekhnovich, A. Borodin, J. Buresh-Oppenheim, R. Impagliazzo, A. Magen, and T. Pitassi have defined *priority branching trees*, a general model of computation with which they were able to prove rigorous bounds on what backtrack programs can do, in *Computational Complexity* **20** (2011), 679–740.

Golomb
Baumert
commafree codes
optimization
lookahead
dynamic ordering
minimum remaining values
integer programming problems
Balas
Geoffrion
implicit enumeration
Wells
Bitner
Reingold
Gaschnig
backmarking
backjumping
Monte Carlo estimates
Hall
Knuth
importance sampling
Hammersley
Handscomb
self-avoiding walks
king paths
Rosenbluth
Rosenbluth
parallel programming
search tree
Lehmer
load balancing
Finkel
Manber
Alekhnovich
Borodin
Buresh-Oppenheim
Impagliazzo
Magen
Pitassi
priority branching trees

EXERCISES

▶ **1.** [*22*] Explain how the tasks of generating (i) n-tuples, (ii) permutations of distinct items, (iii) combinations, (iv) integer partitions, (v) set partitions, and (vi) nested parentheses can all be regarded as special cases of backtrack programming, by presenting suitable domains D_k and cutoff properties $P_l(x_1, \ldots, x_l)$ that satisfy (1) and (2).

2. [*10*] True or false: We can choose D_1 so that $P_1(x_1)$ is always true.

3. [*20*] Let T be any tree. Is it possible to define domains D_k and cutoff properties $P_l(x_1, \ldots, x_l)$ so that T is the backtrack tree traversed by Algorithm B?

4. [*16*] Using a chessboard and eight coins to represent queens, one can follow the steps of Algorithm B and essentially traverse the tree of Fig. 68 by hand in about three hours. Invent a trick to save half of the work.

▶ **5.** [*20*] Reformulate Algorithm B as a *recursive* procedure called *try*$\,(l)$, having global variables n and $x_1 \ldots x_n$, to be invoked by saying '*try*$\,(1)$'. Can you imagine why the author of this book decided *not* to present the algorithm in such a recursive form?

6. [*20*] Given r, with $1 \le r \le 8$, in how many ways can 7 nonattacking queens be placed on an 8×8 chessboard, if no queen is placed in row r?

7. [*20*] (T. B. Sprague, 1890.) Are there any values $n > 5$ for which the n queens problem has a "framed" solution with $x_1 = 2$, $x_2 = n$, $x_{n-1} = 1$, and $x_n = n - 1$?

8. [*20*] Are there two 8-queen placements with the same $x_1x_2x_3x_4x_5x_6$?

9. [*21*] Can a $4m$-queen placement have $3m$ queens on "white" squares?

▶ **10.** [*22*] Adapt Algorithm W to the n queens problem, using bitwise operations on n-bit numbers as suggested in the text.

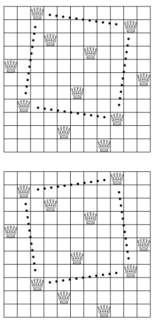

11. [*M25*] (W. Ahrens, 1910.) Both solutions of the n queens problem when $n = 4$ have *quarterturn symmetry*: Rotation by 90° leaves them unchanged, but reflection doesn't.

a) Can the n queens problem have a solution with reflection symmetry?

b) Show that quarterturn symmetry is impossible if $n \bmod 4 \in \{2, 3\}$.

c) Sometimes the solution to an n queens problem contains four queens that form the corners of a tilted square, as shown here. Prove that we can always get another solution by tilting the square the other way (but leaving the other $n - 4$ queens in place).

d) Let C_n be the number of solutions with 90° symmetry, and suppose c_n of them have $x_k > k$ for $1 \le k \le n/2$. Prove that $C_n = 2^{\lfloor n/4 \rfloor} c_n$.

12. [*M28*] (*Wraparound queens.*) Replace (3) by the stronger conditions '$x_j \ne x_k$, $(x_k - x_j) \bmod n \ne k - j$, $(x_j - x_k) \bmod n \ne k - j$'. (The $n \times n$ grid becomes a torus.) Prove that the resulting problem is solvable if and only if n is not divisible by 2 or 3.

13. [*M30*] For which $n \ge 0$ does the n queens problem have at least one solution?

14. [*M25*] If exercise 12 has $T(n)$ toroidal solutions, show that $Q(mn) \ge Q(m)^n T(n)$.

15. [*HM47*] Does $(\ln Q(n))/(n \ln n)$ approach a positive constant as $n \to \infty$?

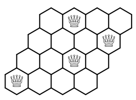

16. [*21*] Let $H(n)$ be the number of ways that n queen bees can occupy an $n \times n$ honeycomb so that no two are in the same line. (For example, one of the $H(4) = 7$ ways is shown here.) Compute $H(n)$ for small n.

17. [*15*] J. H. Quick (a student) noticed that the loop in step L2 of Algorithm L can be changed from 'while $x_l < 0$' to 'while $x_l \ne 0$', because x_l cannot be positive at

n-tuples
tuples
permutations
combinations
integer partitions
partitions
set partitions
nested parentheses
parentheses
domains
cutoff
properties
backtrack tree
recursive
global variables
author
recursion versus iteration
Sprague
bitwise operations
Ahrens
quarterturn symmetry
Rotation by 90°
Wraparound queens
broken diagonal, see wraparound
torus
n queens problem
queen bees
bees
honeycomb
hexagons
Quick

that point of the algorithm. So he decided to eliminate the minus signs and just set $x_{l+k+1} \leftarrow k$ in step L3. Was it a good idea?

18. [*17*] Suppose that $n = 4$ and Algorithm L has reached step L2 with $l = 4$ and $x_1x_2x_3 = 241$. What are the current values of $x_4x_5x_6x_7x_8$, $p_0p_1p_2p_3p_4$, and $y_1y_2y_3$?

19. [*M19*] What are the domains D_l in Langford's problem (7)?

▸ **20.** [*21*] Extend Algorithm L so that it forces $x_l \leftarrow k$ whenever $k \notin \{x_1, \ldots, x_{l-1}\}$ and $l \geq 2n - k - 1$.

▸ **21.** [*M25*] If $x = x_1x_2 \ldots x_{2n}$, let $x^D = (-x_{2n}) \ldots (-x_2)(-x_1) = -x^R$ be its dual.

a) Show that if n is odd and x solves Langford's problem (7), we have $x_k = n$ for some $k \leq \lfloor n/2 \rfloor$ if and only if $x_k^D = n$ for some $k > \lfloor n/2 \rfloor$.

b) Find a similar rule that distinguishes x from x^D when n is even.

c) Consequently the algorithm of exercise 20 can be modified so that exactly one of each dual pair of solutions $\{x, x^D\}$ is visited.

22. [*M26*] Explore "loose Langford pairs": Replace '$j + k + 1$' in (7) by '$j + \lfloor 3k/2 \rfloor$'.

23. [*17*] We can often obtain one word rectangle from another by changing only a letter or two. Can you think of any 5×6 word rectangles that almost match (10)?

24. [*20*] Customize Algorithm B so that it will find all 5×6 word rectangles.

▸ **25.** [*25*] Explain how to use *orthogonal lists*, as in Fig. 13 of Section 2.2.6, so that it's easy to visit all 5-letter words whose kth character is c, given $1 \leq k \leq 5$ and $\mathtt{a} \leq c \leq \mathtt{z}$. Use those sublists to speed up the algorithm of exercise 24.

26. [*21*] Can you find nice word rectangles of sizes 5×7, 5×8, 5×9, 5×10?

27. [*22*] What profile and average node costs replace (13) and (14) when we ask the algorithm of exercise 25 for 6×5 word rectangles instead of 5×6?

▸ **28.** [*23*] The method of exercises 24 and 25 does n levels of backtracking to fill the cells of an $m \times n$ rectangle one column at a time, using a trie to detect illegal prefixes in the rows. Devise a method that does mn levels of backtracking and fills just *one* cell per level, using tries for *both* rows and columns.

29. [*20*] Do any 5×6 word rectangles contain fewer than 11 different words?

30. [*22*] *Symmetric* word squares, whose columns are the same as their rows, were popular in England during the 1850s. For example, A. De Morgan praised the square

```
L E A V E
E L L E N
A L O N E
V E N O M
E N E M Y
```

because it actually is "meaningful"! Determine the total number of symmetric 5×5 word squares, by adapting the method of exercise 28. How many belong to `WORDS`(500)?

31. [*20*] (Charles Babbage, 1864.) Do any of the symmetric 5×5 word squares also have valid words on both diagonals?

32. [*22*] How many symmetric word squares of sizes 2×2, 3×3, ..., are supported by *The Official* SCRABBLE® *Players Dictionary* (fourth edition, 2005)?

33. [*21*] Puzzlers who tried to construct word squares by hand found long ago that it was easiest to work from bottom to top. Therefore they used "reverse dictionaries," whose words appear in colex order. Does this idea speed up computer experiments?

Langford pairs+
domains
dual
loose Langford pairs
word rectangle
orthogonal lists
trie
word squares
De Morgan
Babbage
OSPD4
reverse dictionaries
colex order

34. [*15*] What's the largest commafree subset of the following words?

```
aced babe bade bead beef cafe cede dada dead deaf face fade feed
```

▸ **35.** [*22*] Let $w_1, w_2, \ldots, w_n$ be four-letter words on an m-letter alphabet. Design an algorithm that accepts or rejects each w_j, according as w_j is commafree or not with respect to the accepted words of $\{w_1, \ldots, w_{j-1}\}$.

36. [*M22*] A two-letter block code on an m-letter alphabet can be represented as a digraph D on m vertices, with $a \to b$ if and only if ab is a codeword.

a) Prove that the code is commafree $\iff$ D has no oriented paths of length 3.
b) How many arcs can be in an m-vertex digraph with no oriented paths of length r?

▸ **37.** [*M30*] (W. L. Eastman, 1965.) The following elegant construction yields a commafree code of maximum size for any *odd* block length n, over any alphabet. Given a sequence $x = x_0x_1 \ldots x_{n-1}$ of nonnegative integers, where x differs from each of its other cyclic shifts $x_k \ldots x_{n-1}x_0 \ldots x_{k-1}$ for $0 < k < n$, the procedure outputs a cyclic shift σx with the property that the set of all such σx is commafree.

We regard x as an infinite periodic sequence $\langle x_n \rangle$ with $x_k = x_{k-n}$ for all $k \ge n$. Each cyclic shift then has the form $x_k x_{k+1} \ldots x_{k+n-1}$. The simplest nontrivial example occurs when $n = 3$, where $x = x_0x_1x_2x_0x_1x_2x_0 \ldots$ and we don't have $x_0 = x_1 = x_2$. In this case the algorithm outputs $x_k x_{k+1} x_{k+2}$ where $x_k > x_{k+1} \le x_{k+2}$; and the set of all such triples clearly satisfies the commafree condition.

One key idea is to think of x as partitioned into t substrings by boundary markers b_j, where $0 \le b_0 < b_1 < \cdots < b_{t-1} < n$ and $b_j = b_{j-t} + n$ for $j \ge t$. Then substring y_j is $x_{b_j} x_{b_j+1} \ldots x_{b_{j+1}-1}$. The number t of substrings is always odd. Initially $t = n$ and $b_j = j$ for all j; ultimately $t = 1$, and $\sigma x = y_0$ is the desired output.

Eastman's algorithm is based on comparison of adjacent substrings y_{j-1} and y_j. If those substrings have the same length, we use lexicographic comparison; otherwise we declare that the longer substring is bigger.

The second key idea is the notion of "dips," which are substrings of the form $z = z_1 \ldots z_k$ where $k \ge 2$ and $z_1 \ge \cdots \ge z_{k-1} < z_k$. It's easy to see that any string $y = y_0y_1 \ldots$ in which we have $y_i < y_{i+1}$ for infinitely many i can be factored into a sequence of dips, $y = z^{(0)}z^{(1)} \ldots$, and this factorization is unique. For example,

$$3141592653589793238462643383 \ldots = 314\ 15\ 926\ 535\ 89\ 79\ 323\ 846\ 26\ 4338\ 3 \ldots .$$

Furthermore, if y is a periodic sequence, its factorization into dips is also ultimately periodic, although some of the initial factors may not occur in the period. For example,

$$123443550123443550123443550 \ldots = 12\ 34\ 435\ 501\ 23\ 4435\ 501\ 23\ 4435\ \ldots .$$

Given a periodic, nonconstant sequence y described by boundary markers as above, where the period length t is odd, its periodic factorization will contain an odd number of odd-length dips. Each round of Eastman's algorithm simply retains the boundary points at the left of those odd-length dips. Then t is reset to the number of retained boundary points, and another round begins if $t > 1$.

a) Play through the algorithm by hand when $n = 19$ and $x = 3141592653589793238$.
b) Show that the number of rounds is at most $\lfloor \log_3 n \rfloor$.
c) Exhibit a binary x that achieves this worst-case bound when $n = 3^e$.
d) Implement the algorithm with full details. (It's surprisingly short!)
e) Explain why the algorithm yields a commafree code.

commafree
two-letter block code
digraph
commafree
Eastman
aperiodic words
periodic sequence
substrings
boundary markers
lexicographic comparison
dips
pi, as "random" data
worst-case bound
analysis of algs

38. [*HM28*] What is the probability that Eastman's algorithm finishes in one round? (Assume that x is a random m-ary string of odd length $n > 1$, unequal to any of its other cyclic shifts. Use a generating function to express the answer.)

39. [*18*] Why can't a commafree code of length $(m^4 - m^2)/4$ contain 0001 and 2000?

▸ **40.** [*15*] Why do you think sequential data structures such as (16)–(23) weren't featured in Section 2.2.2 of this series of books (entitled "Sequential Allocation")?

41. [*17*] What's the significance of (a) `MEM[40d]`=`5e` and (b) `MEM[904]`=`84` in Table 1?

42. [*18*] Why is (a) `MEM[f8]` = `e7` and (b) `MEM[a0d]` = `ba` in Table 2?

43. [*20*] Suppose you're using the undoing scheme (26) and the operation $\sigma \leftarrow \sigma + 1$ has just bumped the current stamp σ to zero. What should you do?

▸ **44.** [*25*] Spell out the low-level implementation details of the candidate selection process in step C2 of Algorithm C. Use the routine store(a, v) of (26) whenever changing the contents of `MEM`, and use the following selection strategy:

a) Find a class c with the least number r of blue words.
b) If $r = 0$, set $x \leftarrow -1$; otherwise set x to a word in class c.
c) If $r > 1$, use the poison list to find an x that maximizes the number of blue words that could be killed on the other side of the prefix or suffix list that contains x.

▸ **45.** [*28*] Continuing exercise 44, spell out the details of step C3 when $x \geq 0$.

a) What updates should be done to `MEM` when a blue word x becomes red?
b) What updates should be done to `MEM` when a blue word x becomes green?
c) Step C3 begins its job by making x green as in part (b). Explain how it should finish its job by updating the poison list.

46. [*M30*] Is there a *binary* ($m = 2$) commafree code with one codeword in each of the $(\sum_{d\backslash n} \mu(d)2^{n/d})/n$ cycle classes, for every word length n?

47. [*HM29*] A commafree code on m letters is equivalent to at most $2m!$ such codes if we permute the letters and/or replace each codeword by its left-right reflection.

Determine all of the nonisomorphic commafree codes of length 4 on m letters when m is (a) 2 (b) 3 (c) 4 and there are (a) 3 (b) 18 (c) 57 codewords.

48. [*M42*] Find a maximum-size commafree code of length 4 on $m = 5$ letters.

49. [*20*] Explain how the choices in Fig. 69 were determined from the "random" bits that are displayed. For instance, why was X_2 set to 1 in Fig. 69(b)?

50. [*M15*] Interpret the value $\mathrm{E}(D_1 \ldots D_l)$, in the text's Monte Carlo algorithm.

51. [*M22*] What's a simple martingale that corresponds to Theorem E?

▸ **52.** [*HM25*] Elmo uses Algorithm E with $D_k = \{1, \ldots, n\}$, $P_l = [x_1 > \cdots > x_l]$, $c = 1$.

a) Alice flips n coins independently, where coin k yields "heads" with probability $1/k$. True or false: She obtains exactly l heads with probability ${n \brack l}/n!$.
b) Let $Y_1, Y_2, \ldots, Y_l$ be the numbers on the coins that come up heads. (Thus $Y_1 = 1$, and $Y_2 = 2$ with probability 1/2.) Show that Pr(Alice obtains $Y_1, Y_2, \ldots, Y_l$) = Pr(Elmo obtains $X_1 = Y_l$, $X_2 = Y_{l-1}, \ldots, X_l = Y_1$).
c) Prove that Alice q.s. obtains at most $(\ln n)(\ln \ln n)$ heads.
d) Consequently Elmo's S is q.s. less than $\exp((\ln n)^2(\ln \ln n))$.

▸ **53.** [*M30*] Extend Algorithm B so that it also computes the minimum, maximum, mean, and variance of the Monte Carlo estimates S produced by Algorithm E.

Eastman
generating function
analysis of algs
data structures
Sequential Allocation
undoing scheme
bumped
stamp
poison list
nonisomorphic
symmetries
"random" bits
Monte Carlo
martingale
coins
Stirling cycle numbers
q.s.
variance

54. [*M21*] Instead of choosing each y_i in step E5 with probability $1/d$, we could use a biased distribution where $\Pr(I = i \mid X_1, \ldots, X_{l-1}) = p_{X_1 \ldots X_{l-1}}(y_i) > 0$. How should the estimate S be modified so that its expected value in this general scheme is still $C()$?

55. [*M20*] If all costs $c(x_1, \ldots, x_l)$ are positive, show that the biased probabilities of exercise 54 can be chosen in such a way that the estimate S is always exact.

▸ **56.** [*M25*] The commafree code search procedure in Algorithm C doesn't actually fit the mold of Algorithm E, because it incorporates lookahead, dynamic ordering, reversible memory, and other enhancements to the basic backtrack paradigms. How could its running time be reliably estimated with Monte Carlo methods?

57. [*HM21*] Algorithm E can potentially follow M different paths $X_1 \ldots X_{l-1}$ before it terminates, where M is the number of leaves of the backtrack tree. Suppose the final values of D at those leaves are $D^{(1)}, \ldots, D^{(M)}$. Prove that $\bigl(D^{(1)} \ldots D^{(M)}\bigr)^{1/M} \ge M$.

58. [*27*] The text's king path problem is a special case of the general problem of counting simple paths from vertex s to vertex t in a given graph.

We can generate such paths by random walks from s that don't get stuck, if we maintain a table of values `DIST(`v`)` for all vertices v not yet in the path, representing the shortest distance from v to t through unused vertices. For with such a table we can simply move at each step to a vertex for which `DIST(`v`)` $< \infty$.

Devise a way to update the `DIST` table dynamically without unnecessary work.

59. [*26*] A ZDD with 3,174,197 nodes can be constructed for the family of all simple corner-to-corner king paths on a chessboard, using the method of exercise 7.1.4–225. Explain how to use this ZDD to compute (a) the total length of all paths; (b) the number of paths that touch any given subset of the center and/or corner points.

▸ **60.** [*20*] Experiment with biased random walks (see exercise 54), weighting each non-dead-end king move to a new vertex v by $1 + \texttt{DIST}(v)^2$ instead of choosing every such move with the same probability. Does this strategy improve on Fig. 70?

61. [*HM26*] Let P_n be the number of integer sequences $x_1 \ldots x_n$ such that $x_1 = 1$ and $1 \le x_{k+1} \le 2x_k$ for $1 \le k < n$. (The first few values are 1, 2, 6, 26, 166, 1626, ...; this sequence was introduced by A. Cayley in *Philosophical Magazine* (4) **13** (1857), 245–248, who showed that P_n enumerates the partitions of $2^n - 1$ into powers of 2.)

a) Show that P_n is the number of different profiles that are possible for a binary tree of height n.

b) Find an efficient way to compute P_n for large n. *Hint:* Consider the more general sequence $P_n^{(m)}$, defined similarly but with $x_1 = m$.

c) Use the estimation procedure of Theorem E to prove that $P_n \ge 2^{\binom{n}{2}}/(n-1)!$.

▸ **62.** [*22*] When the faces of four cubes are colored randomly with four colors, estimate the probability that the corresponding "Instant Insanity" puzzle has a unique solution. How many 2-regular graphs tend to appear during the "factored" solution process?

63. [*20*] Find *five* cubes, each of whose faces has one of *five* colors, and where every color occurs at least five times, such that the corresponding puzzle has a unique solution.

64. [*24*] Assemble five cubes with uppercase letters on each face, using the patterns

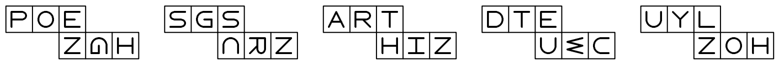

By extending the principles of "Instant Insanity," show that these cubes can be placed in a row so that four 5-letter words are visible. (Each word's letters should have a consistent orientation. The letters C and U, H and I, N and Z are related by 90° rotation.)

biased distribution
commafree code
lookahead
dynamic ordering
reversible memory
simple paths
random walks
dynamic shortest distances
shortest distances, dynamic
ZDD
corner-to-corner
king paths
chessboard
biased random walks
Cayley
binary partitions
partitions
profiles
height n
Instant Insanity
2-regular graphs
factored
speedy schizophrenia
cubes
uppercase letters
LEXI-CUBES
alphabet blocks
Instant Insanity
5-letter words

65. [*25*] Show that the generalized "Instant Insanity" problem, with n cubes and n colors on their faces, is NP-complete, even though cases with small n are fairly easy.

▸ **66.** [*23*] (*The Fool's Disk.*) "Rotate the four disks of the left-hand illustration below so that the four numbers on each ray sum to 12." (The current sums are $4+3+2+4=13$, etc.) Show that this problem factors nicely, so that it can be solved readily by hand.

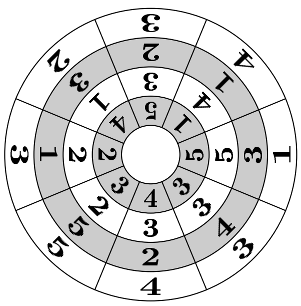

The Fool's Disk

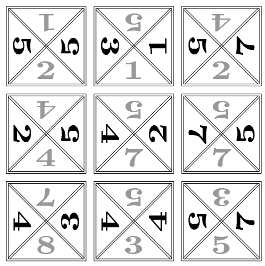

The Royal Aquarium Thirteen Puzzle

▸ **67.** [*26*] (*The Royal Aquarium Thirteen Puzzle.*) "Rearrange the nine cards of the right-hand illustration above, optionally rotating some of them by 180°, so that the six horizontal sums of gray letters and the six vertical sums of black letters all equal 13." (The current sums are $1+5+4=10$, ..., $7+5+7=19$.) The author of *Hoffmann's Puzzles Old and New* (1893) stated that "There is no royal road to the solution. The proper order must be arrived at by successive transpositions until the conditions are fulfilled." Prove that he was wrong: "Factor" this problem and solve it by hand.

▸ **68.** [*28*] (Johan de Ruiter, 14 March 2018.) Put a digit into each empty box, in such a way that every box names the exact number of distinct digits that it points to.

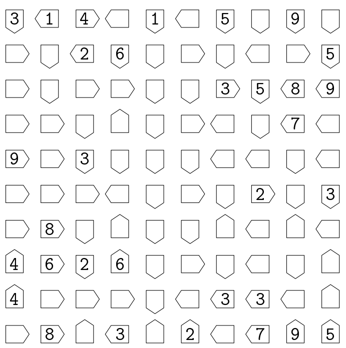

69. [*40*] Is there a puzzle like exercise 68 whose clues contain *more* than 32 digits of π?

70. [*HM40*] (M. Bousquet-Mélou.) Consider self-avoiding paths from the upper left corner of an $m \times n$ grid to the lower right, where each step is either up, down, or to the right. If we generate such paths at random, making either 1 or 2 or 3 choices at each step as in Algorithm E, the expected value $\mathrm{E}\,D_{mn}$ is the total number of such paths, m^{n-1}. But the variance is considerably larger: Construct polynomials $P_m(z)$ and $Q_m(z)$ such that we have $G_m(z) = \sum_{n=1}^{\infty}(\mathrm{E}\,D_{mn}^2)z^n = zP_m(z)/Q_m(z)$ for $m \geq 2$. For example, $G_3(z) = (z+z^2)/(1-9z-6z^2) = z+10z^2+96z^3+924z^4+8892z^5+\cdots$. Prove furthermore that $\mathrm{E}\,D_{mn}^2 = \Theta(\rho_m^n)$, where $\rho_m = 2^m + O(1)$.

NP-complete
Fool's Disk
Royal Aquarium Thirteen Puzzle
Le Nombre Treize, see Royal Aquarium Thirteen
Hoffmann
Factor
de Ruiter
pi day
self-reference
π
Bousquet-Mélou
self-avoiding paths
variance

Woods
questionnaire
paradoxical

Table 666

TWENTY QUESTIONS (SEE EXERCISE 71)

---

1. The first question whose answer is A is:
(A) 1 (B) 2 (C) 3 (D) 4 (E) 5

2. The next question with the same answer as this one is:
(A) 4 (B) 6 (C) 8 (D) 10 (E) 12

3. The only two consecutive questions with identical answers are questions:
(A) 15 and 16 (B) 16 and 17 (C) 17 and 18 (D) 18 and 19 (E) 19 and 20

4. The answer to this question is the same as the answers to questions:
(A) 10 and 13 (B) 14 and 16 (C) 7 and 20 (D) 1 and 15 (E) 8 and 12

5. The answer to question 14 is:
(A) B (B) E (C) C (D) A (E) D

6. The answer to this question is:
(A) A (B) B (C) C (D) D (E) none of those

7. An answer that appears most often is:
(A) A (B) B (C) C (D) D (E) E

8. Ignoring answers that appear equally often, the least common answer is:
(A) A (B) B (C) C (D) D (E) E

9. The sum of all question numbers whose answers are correct and the same as this one is:
(A) $\in [59\,..\,62]$ (B) $\in [52\,..\,55]$ (C) $\in [44\,..\,49]$ (D) $\in [61\,..\,67]$ (E) $\in [44\,..\,53]$

10. The answer to question 17 is:
(A) D (B) B (C) A (D) E (E) wrong

11. The number of questions whose answer is D is:
(A) 2 (B) 3 (C) 4 (D) 5 (E) 6

12. The number of *other* questions with the same answer as this one is the same as the number of questions with answer:
(A) B (B) C (C) D (D) E (E) none of those

13. The number of questions whose answer is E is:
(A) 5 (B) 4 (C) 3 (D) 2 (E) 1

14. No answer appears exactly this many times:
(A) 2 (B) 3 (C) 4 (D) 5 (E) none of those

15. The set of odd-numbered questions with answer A is:
(A) $\{7\}$ (B) $\{9\}$ (C) not $\{11\}$ (D) $\{13\}$ (E) $\{15\}$

16. The answer to question 8 is the same as the answer to question:
(A) 3 (B) 2 (C) 13 (D) 18 (E) 20

17. The answer to question 10 is:
(A) C (B) D (C) B (D) A (E) correct

18. The number of prime-numbered questions whose answers are vowels is:
(A) prime (B) square (C) odd (D) even (E) zero

19. The last question whose answer is B is:
(A) 14 (B) 15 (C) 16 (D) 17 (E) 18

20. The maximum score that can be achieved on this test is:
(A) 18 (B) 19 (C) 20 (D) indeterminate
(E) achievable only by getting this question wrong

▸ **71.** [*M29*] (Donald R. Woods, 2000.) Find all ways to maximize the number of correct answers to the questionnaire in Table 666. Each question must be answered with a letter from A to E. *Hint:* Begin by clarifying the exact meaning of this exercise. What answers are best for the following two-question, two-letter "warmup problem"?

1. (A) Answer 2 is B. (B) Answer 1 is A.

2. (A) Answer 1 is correct. (B) Either answer 2 is wrong or answer 1 is A, but not both.

72. [*HM28*] Show that exercise 71 has a surprising, somewhat paradoxical answer if two changes are made to Table 666: 9(E) becomes '$\in [39\,..\,43]$'; 15(C) becomes '$\{11\}$'.

▸ **73.** [*30*] (*A clueless anacrostic.*) The letters of 29 five-letter words

$\underline{\ }_{1}\underline{\ }_{2}\underline{\ }_{3}\underline{\ }_{4}\underline{\ }_{5}$, $\underline{\ }_{6}\underline{\ }_{7}\underline{\ }_{8}\underline{\ }_{9}\underline{\ }_{10}$, $\underline{\ }_{11}\underline{\ }_{12}\underline{\ }_{13}\underline{\ }_{14}\underline{\ }_{15}$, $\underline{\ }_{16}\underline{\ }_{17}\underline{\ }_{18}\underline{\ }_{19}\underline{\ }_{20}$, …, $\underline{\ }_{141}\underline{\ }_{142}\underline{\ }_{143}\underline{\ }_{144}\underline{\ }_{145}$,

all belonging to `WORDS`(1000), have been shuffled to form the following mystery text:

$\underline{\ }_{30}\underline{\ }_{29}\underline{\ }_{9}$ $\underline{\ }_{140}\underline{\ }_{12}\underline{\ }_{13}\underline{\ }_{145}\underline{\ }_{90}\underline{\ }_{45}\underline{\ }_{99}$ $\underline{\ }_{26}\underline{\ }_{107}$ $\underline{\ }_{47}\underline{\ }_{84}\underline{\ }_{53}\underline{\ }_{51}\underline{\ }_{27}\underline{\ }_{133}\underline{\ }_{39}$ $\underline{\ }_{137}\underline{\ }_{139}$ $\underline{\ }_{66}\underline{\ }_{112}\underline{\ }_{69}\underline{\ }_{14}\underline{\ }_{8}\underline{\ }_{20}\underline{\ }_{91}\underline{\ }_{129}\underline{\ }_{70}$
$\underline{\ }_{16}\underline{\ }_{7}\underline{\ }_{93}\underline{\ }_{19}\underline{\ }_{85}$ $\underline{\ }_{101}\underline{\ }_{76}\underline{\ }_{78}\underline{\ }_{44}\underline{\ }_{10}\underline{\ }_{106}\underline{\ }_{60}$ $\underline{\ }_{118}\underline{\ }_{119}$ $\underline{\ }_{24}\underline{\ }_{25}\underline{\ }_{100}$ $\underline{\ }_{1}\underline{\ }_{5}\underline{\ }_{64}\underline{\ }_{11}\underline{\ }_{71}$ $\underline{\ }_{42}\underline{\ }_{122}\underline{\ }_{123}$
$\underline{\ }_{103}\underline{\ }_{104}\underline{\ }_{63}\underline{\ }_{49}\underline{\ }_{31}\underline{\ }_{121}\underline{\ }_{98}\underline{\ }_{79}\underline{\ }_{80}$ $\underline{\ }_{46}\underline{\ }_{48}$ $\underline{\ }_{134}\underline{\ }_{135}\underline{\ }_{131}$ $\underline{\ }_{143}\underline{\ }_{96}\underline{\ }_{142}\underline{\ }_{120}\underline{\ }_{50}\underline{\ }_{132}\underline{\ }_{33}\underline{\ }_{43}\underline{\ }_{34}\underline{\ }_{40}$. …
$\underline{\ }_{111}\underline{\ }_{97}\underline{\ }_{113}\underline{\ }_{105}\underline{\ }_{38}\underline{\ }_{102}\underline{\ }_{62}\underline{\ }_{65}\underline{\ }_{114}$ $\underline{\ }_{74}\underline{\ }_{82}\underline{\ }_{81}\underline{\ }_{83}\underline{\ }_{136}\underline{\ }_{37}\underline{\ }_{21}\underline{\ }_{61}\underline{\ }_{88}\underline{\ }_{86}\underline{\ }_{55}$ ($\underline{\ }_{32}\underline{\ }_{35}$ $\underline{\ }_{117}\underline{\ }_{116}\underline{\ }_{23}\underline{\ }_{52}$
$\underline{\ }_{56}\underline{\ }_{17}\underline{\ }_{18}\underline{\ }_{94}\underline{\ }_{67}$ $\underline{\ }_{128}\underline{\ }_{15}\underline{\ }_{57}\underline{\ }_{58}\underline{\ }_{89}$ $\underline{\ }_{87}\underline{\ }_{109}$ $\underline{\ }_{2}\underline{\ }_{4}\underline{\ }_{6}\underline{\ }_{28}\underline{\ }_{95}\underline{\ }_{3}\underline{\ }_{126}\underline{\ }_{77}\underline{\ }_{144}\underline{\ }_{54}\underline{\ }_{41}$) $\underline{\ }_{68}\underline{\ }_{115}$
$\underline{\ }_{75}\underline{\ }_{138}\underline{\ }_{73}\underline{\ }_{124}\underline{\ }_{36}\underline{\ }_{130}\underline{\ }_{127}\underline{\ }_{141}$ $\underline{\ }_{22}\underline{\ }_{92}$ $\underline{\ }_{72}\underline{\ }_{59}$ $\underline{\ }_{108}\underline{\ }_{125}\underline{\ }_{110}$.

Furthermore, their initial letters $\underline{\ }_{1}$, $\underline{\ }_{6}$, $\underline{\ }_{11}$, $\underline{\ }_{16}$, …, $\underline{\ }_{141}$ identify the source of that quotation, which consists entirely of common English words. What does it say?

74. [*21*] The fifteenth mystery word in exercise 73 is '$\underline{\ }_{134}\underline{\ }_{135}\underline{\ }_{131}$'. Why does its special form lead to a partial *factorization* of that problem?

▸ **75.** [*30*] (*Connected subsets.*) Let v be a vertex of some graph G, and let H be a connected subset of G that contains v. The vertices of H can be listed in a canonical way by starting with $v_0 \leftarrow v$ and then letting v_1, v_2, … be the neighbors of v_0 that lie in H, followed by the neighbors of v_1 that haven't yet been listed, and so on. (We assume that the neighbors of each vertex are listed in some fixed order.)

For example, if G is the 3×3 grid $P_3 \,\square\, P_3$, exactly 21 of its connected five-element subsets contain the upper left corner element v. Their canonical orderings are

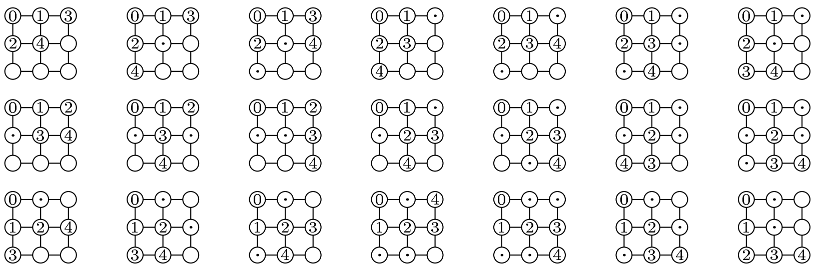

if we order the vertices top-to-bottom, left-to-right when listing their neighbors. (Vertices labeled ⓪, ①, ②, ③, ④ indicate v_0, v_1, v_2, v_3, v_4. Other vertices are not in H.)

Design a backtrack algorithm to generate all of the n-element connected subsets that contain a specified vertex v, given a graph that is represented in SGB format (which has `ARCS`, `TIP`, and `NEXT` fields, as described near the beginning of Chapter 7).

76. [*23*] Use the algorithm of exercise 75 to generate *all* of the connected n-element subsets of a given graph G. How many such subsets does $P_n \,\square\, P_n$ have, for $1 \le n \le 9$?

77. [*M22*] A v-reachable subset of a directed graph G is a nonempty set of vertices H with the property that every $u \in H$ can be reached from v by at least one oriented path in $G\,|\,H$. (In particular, v itself must be in H.)

a) The digraph $P_3^{\rightarrow} \,\square\, P_3^{\rightarrow}$ is like $P_3 \,\square\, P_3$, except that all arcs between vertices are directed downward or to the right. Which of the 21 connected subsets in exercise 75 are also v-reachable from the upper left corner element v of $P_3^{\rightarrow} \,\square\, P_3^{\rightarrow}$?

clueless
anacrostic
five-letter words
WORDS(n)
mystery text
English words
factorization
Connected subsets
canonical
grid
pentominoes
SGB format
ARCS
TIP
NEXT
polyominoes
v-reachable subset
reachable subsets
grid, oriented
oriented grid

b) True or false: H is v-reachable if and only if $G\,|\,H$ contains a dual oriented spanning tree rooted at v. (An oriented tree has arcs $u \longrightarrow p_u$, where p_u is the parent of the nonroot node u; in a *dual* oriented tree, the arcs are reversed: $p_u \longrightarrow u$.)

c) True or false: If G is undirected, so that $w \longrightarrow u$ whenever $u \longrightarrow w$, its v-reachable subsets are the same as the connected subsets that contain v.

d) Modify the algorithm of exercise 75 so that it generates all of the n-element v-reachable subsets of a digraph G, given n, v, and G.

78. [*22*] Extend the algorithm of exercise 77 to weighted graphs, in which every vertex has a nonnegative weight: Generate all of the connected induced subgraphs whose total weight w satisfies $L \le w < U$.

▸ **79.** [*M30*] The author and his wife own a pipe organ that contains 812 pipes, each of which is either playing or silent. Therefore 2^{812} different sounds (including silence) can potentially be created. However, the pipes are controlled by a conventional organ console, which has only $56+56+32 = 144$ keys and pedals that can be played by hands and feet, together with 20 on-off switches to define the connections between keys and pipes. Therefore at most 2^{164} different sounds are actually playable! The purpose of this exercise is to determine the exact number of n-pipe playable sounds, for small n.

The keys are binary vectors $s = s_0 s_1 \ldots s_{55}$ and $g = g_0 g_1 \ldots g_{55}$; the pedals are $p = p_0 p_1 \ldots p_{31}$; the console control switches are $c = c_0 c_1 \ldots c_{19}$; and the pipes are $r_{i,j}$ for $0 \le i < 16$ and $0 \le j < 56$. Here are the precise rules that define the pipe activity $r_{i,j}$ in terms of the input vectors s, g, p, and c that are governed by the organist:

$$r_{i,j} = \begin{cases} c_i p_j \lor c_{i+15} p_{j-12}, & i \in \{0,1\}; \\ c_i p_j, & i \in \{2\}; \end{cases}$$
$$r_{i,j} = \begin{cases} (c_i \lor c_{i+1}[j<12])\, s_j^*, & i \in \{3\}; \\ c_i[j \ge 12]\, s_j^*, & i \in \{4,8\}; \\ c_i s_j^*, & i \in \{5,6,7\}; \end{cases}$$
$$r_{i,j} = \begin{cases} (c_i \lor c_{i+1}[j<12])\, g_j^*, & i \in \{9\}; \\ c_i[j \ge 12]\, g_j^*, & i \in \{10\}; \\ c_i g_j^*, & i \in \{11,12\}; \\ (c_{13} \lor c_{14})\, g_j^*, & i \in \{13\}; \\ c_{14} g_j^*. & i \in \{14,15\}. \end{cases}$$

Here $p_j = 0$ if $j < 0$ or $j \ge 32$; $s_j^* = s_j \lor c_{17} g_j \lor c_{18} p_j$; $g_j^* = g_j \lor c_{19} p_j$. [In organ jargon, the array of pipes has 16 "ranks"; ranks $\{0,1,2\}$, $\{3,\ldots,8\}$, $\{9,\ldots,15\}$ constitute the Pedal, Swell, and Great divisions. Ranks 3 and 4 share their lower 12 pipes, as do ranks 9 and 10. Ranks 13, 14, and 15 form a "mixture," c_{14}. Unit ranks c_{15} and c_{16} extend ranks 0 and 1, twelve notes higher. Console switches c_{17}, c_{18}, c_{19} are "couplers" Swell → Great, Swell → Pedal, Great → Pedal, which explain the formulas for s_j^* and g_j^*.]

A *playable sound* S is a set of pairs (i,j) such that we have $r_{i,j} = [(i,j) \in S]$ for at least one choice of the input vectors s, g, p, c. For example, the first chord of Bach's *Toccata in D minor* is the 8-pipe sound $\{(3,33), (3,45), (4,33), (4,45), (5,33), (5,45), (6,33), (6,45)\}$, which is achievable when $s_{33} = s_{45} = c_3 = c_4 = c_5 = c_6 = 1$ and all other inputs are 0. We want to find the number Q_n of playable sounds with $\|S\| = n$.

a) There are 16×56 variables $r_{i,j}$ but only 812 actual pipes, because some of the ranks are incomplete. For which pairs (i,j) is $r_{i,j}$ always false?

b) True or false: If $s \subseteq s'$, $g \subseteq g'$, $p \subseteq p'$, and $c \subseteq c'$, then $r \subseteq r'$.

c) Show that every playable sound is achievable with $c_{17} = c_{18} = c_{19} = 0$.

d) Find a 5-pipe playable sound in which just five of the s_j, g_j, p_j, c_j are nonzero.

e) For which i and i' are the 2-pipe sounds $\{(i,40), (i',50)\}$ playable?

f) Determine Q_1 by hand, and explain why it is less than 812.

g) Determine Q_{811} by hand.

h) Determine $Q_2, \ldots, Q_{10}$ by computer, and compare them to $\binom{812}{2}, \ldots, \binom{812}{10}$.

dual oriented spanning tree
oriented tree
parent
directed graph versus undirected
undirected graph versus directed
weighted graphs
author
wife
pipe organ
royalties, use of
music
organ
playable sounds
playable sound
Bach

SECTION 7.2.2

restricted growth strings
Dewey decimal notation
reflection
Schumacher
Gauss
symmetric
inner loop
backtracking, efficient
iteration versus recursion
Preußer
Engelhardt
de Jaenisch
mask

1. Although many formulations are possible, the following may be the nicest: (i) D_k is arbitrary (but hopefully finite), and P_l is always true. (ii) $D_k = \{1, 2, \ldots, n\}$ and P_l = '$x_j \neq x_k$ for $1 \le j < k \le l$'. (iii) For combinations of n things from N, $D_k = \{1, \ldots, N+1-k\}$ and P_l = '$x_1 > \cdots > x_l$'. (iv) $D_k = \{0, 1, \ldots, \lfloor n/k \rfloor\}$; P_l = '$x_1 \ge \cdots \ge x_l$ and $n-(n-l)x_l \le x_1 + \cdots + x_l \le n$'. (v) For restricted growth strings, $D_k = \{0, \ldots, k-1\}$ and P_l = '$x_{j+1} \le 1+\max(x_1, \ldots, x_j)$ for $1 \le j < l$'. (vi) For indices of left parentheses (see 7.2.1.6–(8)), $D_k = \{1, \ldots, 2k-1\}$ and P_l = '$x_1 < \cdots < x_l$'.

2. True. (If not, set $D_1 \leftarrow D_1 \cap \{x \mid P_1(x)\}$.)

3. Let $D_k = \{1, \ldots, \text{max degree on level } k-1\}$, and let $P_l(x_1, \ldots, x_l)$ = '$x_1.\cdots.x_l$ is a label in T's Dewey decimal notation' (see Section 2.3).

4. We can restrict D_1 to $\{1, 2, 3, 4\}$, because the reflection $(9-x_1)\ldots(9-x_8)$ of every solution $x_1 \ldots x_8$ is also a solution. (H. C. Schumacher made this observation in a letter to C. F. Gauss, 24 September 1850.) Notice that Fig. 68 is left-right symmetric.

5. $try(l)$ = "If $l > n$, visit $x_1 \ldots x_n$. Otherwise, for $x_l \leftarrow \min D_l$, $\min D_l + 1$, $\ldots$, $\max D_l$, if $P_l(x_1, \ldots, x_l)$ call $try(l+1)$."

This formulation is elegant, and fine for simple problems. But it doesn't give any clue about why the method is called "backtrack"! Nor does it yield efficient code for important problems whose inner loop is performed billions of times. We will see that the key to efficient backtracking is to provide good ways to update and downdate the data structures that speed up the testing of property P_l. The overhead of recursion can get in the way, and the actual iterative structure of Algorithm B isn't difficult to grasp.

6. Excluding cases with $j = r$ or $k = r$ from (3) yields respectively (312, 396, 430, 458, 458, 430, 396, 312) solutions. (With column r also omitted there are just (40, 46, 42, 80, 80, 42, 46, 40).)

7. Yes, almost surely for all $n > 16$. One such is $x_1x_2\ldots x_{17}$ = 2 17 12 10 7 14 3 5 9 13 15 4 11 8 6 1 16. [See *Proc. Edinburgh Math. Soc.* **8** (1890), 43 and Fig. 52.] Preußer and Engelhardt found 34,651,355,392 solutions when $n = 27$.

8. Yes: (42736815, 42736851); also therefore (57263148, 57263184).

9. Yes, at least when $m = 4$; e.g., $x_1 \ldots x_{16}$ = 5 8 13 16 3 7 15 11 6 2 10 14 1 4 9 12. There are no solutions when $m = 5$, but 7 10 13 20 17 24 3 6 23 11 16 21 4 9 14 2 19 22 1 8 5 12 15 18 works for $m = 6$. (Are there solutions for all even $m \ge 4$? C. F. de Jaenisch, *Traité des applications de l'analyse mathématique au jeu des échecs* **2** (1862), 132–133, noted that all 8-queen solutions have four of each color. He proved that the number of white queens must be even, because $\sum_{k=1}^{4m}(x_k + k)$ is even.)

10. Let bit vectors a_l, b_l, c_l represent the "useful" elements of the sets in (6), with $a_l = \sum\{2^{x-1} \mid x \in A_l\}$, $b_l = \sum\{2^{x-1} \mid x \in B_l \cap [1\,..\,n]\}$, $c_l = \sum\{2^{x-1} \mid x \in C_l \cap [1\,..\,n]\}$. Then step W2 sets bit vector $s_l \leftarrow \mu \,\&\, \bar{a}_l \,\&\, \bar{b}_l \,\&\, \bar{c}_l$, where μ is the mask $2^n - 1$.

In step W3 we can set $t \leftarrow s_l \,\&\, (-s_l)$, $a_l \leftarrow a_{l-1} + t$, $b_l \leftarrow (b_{l-1} + t) \gg 1$, $c_l \leftarrow ((c_{l-1} + t) \ll 1) \,\&\, \mu$; and it's also convenient to set $s_l \leftarrow s_l - t$ at this time, instead of deferring this change to step W4.

(There's no need to store x_l in memory, or even to compute x_l in step W3 as an integer in $[1\,..\,n]$, because x_l can be deduced from $a_l - a_{l-1}$ when a solution is found.)

11. (a) Only when $n = 1$, because reflected queens can capture each other.

(b) Queens not in the center must appear in groups of four.

(c) The four queens occupy the same rows, columns, and diagonals in both cases.

(d) In each solution counted by c_n we can independently tilt (or not) each of the $\lfloor n/4 \rfloor$ groups of four. [*Mathematische Unterhaltungen und Spiele* **1**, second edition (Leipzig: Teubner, 1910), 249–258.]

12. With distinct x_k, $\sum_{k=1}^{n}(x_k+k) = 2\binom{n+1}{2} \equiv 0$ (modulo n). If the $(x_k+k) \bmod n$ are also distinct, we have also $\sum_{k=1}^{n} k \equiv \binom{n+1}{2}$. But that's impossible when n is even.

Now suppose further that the numbers $(x_k-k) \bmod n$ are distinct. Then we have $\sum_{k=1}^{n}(x_k+k)^2 \equiv \sum_{k=1}^{n}(x_k-k)^2 \equiv \sum_{k=1}^{n} k^2 = n(n+1)(2n+1)/6$. And we also have $\sum_{k=1}^{n}(x_k+k)^2 + \sum_{k=1}^{n}(x_k-k)^2 = 4n(n+1)(2n+1)/6 \equiv 2n/3$, which is impossible when n is a multiple of 3. [See W. Ahrens, *Mathematische Unterhaltungen und Spiele* **2**, second edition (1918), 364–366, where G. Pólya cites a more general result of A. Hurwitz that applies to wraparound diagonals of other slopes.]

Conversely, if n isn't divisible by 2 or 3, we can let $x_n = n$ and $x_k = (2k) \bmod n$ for $1 \le k < n$. [The rule $x_k = (3k) \bmod n$ also works. See Édouard Lucas, *Récréations Mathématiques* **1** (1882), 84–86.]

13. The $(n+1)$ queens problem clearly has a solution with a queen in a corner if and only if the n queens problem has a solution with a queen-free main diagonal. Hence by the previous answer there's always a solution when $n \bmod 6 \in \{0,1,4,5\}$.

Another nice solution was found by J. Franel [L'*Intermédiaire des Mathématiciens* **1** (1894), 140–141] when $n \bmod 6 \in \{2,4\}$: Let $x_k = \bigl(n/2 + 2k - 3[2k \le n]\bigr) \bmod n + 1$, for $1 \le k \le n$. With this setup we find that $x_k - x_j = \pm(k-j)$ and $1 \le j < k \le n$ implies $(1 \text{ or } 3)(k-j) + (0 \text{ or } 3) \equiv 0$ (modulo n); hence $k-j = n - (1 \text{ or } 3)$. But the values of x_1, x_2, x_3, x_{n-2}, x_{n-1}, x_n give no attacking queens except when $n=2$.

Franel's solution has empty diagonals, so it provides solutions also for $n \bmod 6 \in \{3,5\}$. We conclude that only $n=2$ and $n=3$ are impossible.

[A more complicated construction for all $n>3$ had been given earlier by E. Pauls, in *Deutsche Schachzeitung* **29** (1874), 129–134, 257–267. Pauls also explained how to find all solutions, in principle, by building the tree level by level (*not* backtracking).]

14. For $1 \le j \le n$, let $x_1^{(j)} \ldots x_m^{(j)}$ be a solution for m queens, and let $y_1 \ldots y_n$ be a solution for n toroidal queens. Then $X_{(i-1)n+j} = (x_i^{(j)} - 1)n + y_j$ (for $1 \le i \le m$ and $1 \le j \le n$) is a solution for mn queens. [I. Rivin, I. Vardi, and P. Zimmermann, *AMM* **101** (1994), 629–639, Theorem 2.]

15. [Rivin, Vardi, and Zimmermann, in the paper just cited, observe that in fact the sequence $(\ln Q(n))/(n \ln n)$ appears to be *increasing*.]

16. Let the queen in row k be in cell k. Then we have a "relaxation" of the n queens problem, with $|x_k - x_j|$ becoming just $x_k - x_j$ in (3); so we can ignore the b vector in Algorithm B* or in exercise 10. We get

n	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14
$H(n)$	1	1	1	3	7	23	83	405	2113	12657	82297	596483	4698655	40071743	367854835

[N. J. Cavenagh and I. M. Wanless, *Discr. Appl. Math.* **158** (2010), 136–146, Table 2.]

17. It fails spectacularly in step L5. The minus signs, which mark decisions that were previously forced, are crucial tags for backtracking.

18. $x_4 \ldots x_8 = \bar{2}\bar{1}0\bar{4}0$, $p_0 \ldots p_4 = 33300$, and $y_1y_2y_3 = 130$. (If $x_i \le 0$ the algorithm will never look at y_i; hence the current state of $y_4 \ldots y_8$ is irrelevant. But y_4y_5 happens to be 20, because of past history; y_6, y_7, and y_8 haven't yet been touched.)

19. We could say D_l is $\{-n, \ldots, -2, -1, 1, 2, \ldots, n\}$, or $\{k \mid k \ne 0 \text{ and } 2-l \le k \le 2n-l-1\}$, or anything in between. (But this observation isn't very useful.)

Ahrens
Pólya
Hurwitz
Lucas
Franel
Pauls
tree
breadth first search
Rivin
Vardi
Zimmermann
semi-queens
Cavenagh
Wanless

parity argument
trie
signature
trie
bitwise AND

20. First we add a Boolean array $a_1 \ldots a_n$, where a_k means "k has appeared," as in Algorithm B*. It's $0 \ldots 0$ in step L1; we set $a_k \leftarrow 1$ in step L3, $a_k \leftarrow 0$ in step L5.

The loop in step L2 becomes "while $x_l < 0$, go to L5 if $l \geq n-1$ and $a_{2n-l-1} = 0$, otherwise set $l \leftarrow l+1$." After finding $l+k+1 \leq 2n$ in L3, and before testing x_{l+k+1} for 0, insert this: "If $l \geq n-1$ and $a_{2n-l-1} = 0$, while $l+k+1 \neq 2n$ set $j \leftarrow k$, $k \leftarrow p_k$."

21. (a) In any solution $x_k = n \iff x_{k+n+1} = -n \iff x^D_{n-k} = n$.

(b) $x_k = n-1$ for some $k \leq n/2$ if and only if $x^D_k = n-1$ for some $k > n/2$.

(c) Let $n' = n - [n \text{ is even}]$. Change '$l \geq n-1$ and $a_{2n-l-1} = 0$' in the modified step L2 to '($l = \lfloor n/2 \rfloor$ and $a_{n'} = 0$) or ($l \geq n-1$ and $a_{2n-l-1} = 0$)'. Insert the following before the other insertion into step L3: "If $l = \lfloor n/2 \rfloor$ and $a_{n'} = 0$, while $k \neq n'$ set $j \leftarrow k$, $k \leftarrow p_k$." And in step L5 — this subtle detail is needed when n is even — go to L5 instead of L4 if $l = \lfloor n/2 \rfloor$ and $k = n'$.

22. The solutions $1\bar{1}$ and $21\bar{1}\bar{2}$ for $n = 1$ and $n = 2$ are self-dual; the solutions for $n = 4$ and $n = 5$ are $431\bar{1}2\bar{3}\bar{4}\bar{2}$, $245\bar{2}31\bar{1}\bar{4}\bar{3}\bar{5}$, $451\bar{1}23\bar{4}\bar{2}\bar{5}\bar{3}$, and their duals. The total number of solutions for $n = 1, 2, \ldots$ is 1, 1, 0, 2, 4, 20, 0, 156, 516, 2008, 0, 52536, 297800, 1767792, 0, 75678864, …; there are none when $n \bmod 4 = 3$, by a parity argument.

Algorithm L needs only obvious changes. To compute solutions by a streamlined method like exercise 21, use $n' = n - (0,1,2,0)$ and substitute '$l = \lfloor n/4 \rfloor + (0,1,2,1)$' for '$l = \lfloor n/2 \rfloor$', when $n \bmod 4 = (0,1,2,3)$; also replace '$l \geq n-1$ and $a_{2n-l-1} = 0$' by '$l \geq \lceil n/2 \rceil$ and $a_{\lfloor (4n+2-2l)/3 \rfloor} = 0$'. The case $n = 15$ is proved impossible with 397 million nodes and 9.93 gigamems.

23. `slums` $\to$ `sluff`, `slump`, `slurs`, `slurp`, or `sluts`; (`slums`, `total`) $\to$ (`slams`, `tonal`).

24. Build the list of 5-letter words and the trie of 6-letter words in step B1; also set $a_{01}a_{02}a_{03}a_{04}a_{05} \leftarrow 00000$. Use $\min D_l = 1$ in step B2 and $\max D_l = 5757$ in step B4. To test P_l in step B3, if word x_3 is $c_1c_2c_3c_4c_5$, form $a_{l1} \ldots a_{l5}$, where $a_{lk} = \mathit{trie}[a_{(l-1)k}, c_k]$ for $1 \leq k \leq 5$; but jump to B4 if any a_{lk} is zero.

25. There are 5×26 singly linked lists, accessed from pointers h_{kc}, all initially zero. The xth word $c_{x1}c_{x2}c_{x3}c_{x4}c_{x5}$, for $1 \leq x \leq 5757$, belongs to 5 lists and has five pointers $l_{x1}l_{x2}l_{x3}l_{x4}l_{x5}$. To insert it, set $l_{xk} \leftarrow h_{kc_{xk}}$, $h_{kc_{xk}} \leftarrow x$, and $s_{kc_{xk}} \leftarrow s_{kc_{xk}} + 1$, for $1 \leq k \leq 5$. (Thus s_{kc} will be the length of the list accessed from h_{kc}.)

We can store a "signature" $\sum_{c=1}^{26} 2^{c-1}[\mathit{trie}[a,c] \neq 0]$ with each node a of the trie. For example, the signature for node 260 is $2^0+2^4+2^8+2^{14}+2^{17}+2^{20}+2^{24} = \texttt{\#1124111}$, according to (11); here $\texttt{A} \leftrightarrow 1$, …, $\texttt{Z} \leftrightarrow 26$.

The process of running through all x that match a given signature y with respect to position z, as needed in steps B2 and B4, now takes the following form: (i) Set $i \leftarrow 0$. (ii) While $2^i \mathbin{\&} y = 0$, set $i \leftarrow i+1$. (iii) Set $x \leftarrow h_{z(i+1)}$; go to (vi) if $x = 0$. (iv) Visit x. (v) Set $x \leftarrow l_{xz}$; go to (iv) if $x \neq 0$. (vi) Set $i \leftarrow i+1$; go to (ii) if $2^i \leq y$.

Let $\mathit{trie}[a,0]$ be the signature of node a. We choose z and $y = \mathit{trie}[a_{(l-1)z}, 0]$ in step B2 so that the number of nodes to visit, $\sum_{c=1}^{26} s_{zc}[2^{c-1} \mathbin{\&} y \neq 0]$, is minimum for $1 \leq z \leq 5$. For example, when $l = 3$, $x_1 = 1446$, and $x_2 = 185$ as in (10), that sum for $z = 1$ is $s_{11}+s_{15}+s_{19}+s_{1(15)}+s_{1(18)}+s_{1(21)}+s_{1(25)} = 296+129+74+108+268+75+47 = 997$; and the sums for $z = 2, 3, 4, 5$ are 4722, 1370, 5057, and 1646. Hence we choose $z = 1$ and $y = \texttt{\#1124111}$; only 997 words, not 5757, need be tested for x_3.

The values y_l and z_l are maintained for use in backtracking. (In practice we keep x, y, and z in registers during most of the computation. Then we set $x_l \leftarrow x$, $y_l \leftarrow y$, $z_l \leftarrow z$ before increasing $l \leftarrow l+1$ in step B3; and we set $x \leftarrow x_l$, $y \leftarrow y_l$, $z \leftarrow z_l$ in

author
McIlroy
compressed trie
trie, compressed
poetic
palindromes

step B5. We also keep i in a register, while traversing the sublists as above; this value is restored in step B5 by setting it to the zth letter of word x, decreased by `'A'`.)

26. Here are the author's favorite 5×7 and 5×8, and the *only* 5×9's:

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

No 5×10 word rectangles exist, according to our ground rules.

27. (1, 15727, 8072679, 630967290, 90962081, 625415) and (15727.0, 4321.6, 1749.7, 450.4, 286.0). Total time ≈ 18.3 teramems.

28. Build a separate trie for the m-letter words; but instead of having trie nodes of size 26 as in (11), it's better to convert this trie into a *compressed* representation that omits the zeros. For example, the compressed representation of the node for prefix '`CORNE`' in (12) consists of five consecutively stored pairs of entries ('`T`', 5013), ('`R`', 171), ('`L`', 9602), ('`D`', 3878), ('`A`', 3879), followed by $(0,0)$. Similarly, each shorter prefix with c descendants is represented by c consecutive pairs (character, link), followed by $(0,0)$ to mark the end of the node. Steps B3 and B4 are now very convenient.

Level l corresponds to row $i_l = 1 + (l-1) \bmod m$ and column $j_l = 1 + \lfloor (l-1)/m \rfloor$. For backtracking we store the n-trie pointer a_{i_l,j_l} as before, together with an index x_l into the compressed m-trie.

This method was used by M. D. McIlroy in 1975 (see answer 32). It finds all 5×6 word rectangles in just 400 gigamems; and its running time for "transposed" 6×5 rectangles turns out to be slightly less (380 gigamems). Notice that only one mem is needed to access each (character, link) pair in the compressed trie.

29. Yes, exactly 1618 of the 625415 solutions have repeated words. For example:

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

30. The use of a single compressed trie both horizontally and vertically leads to a very pretty algorithm, which needs only 120 Mμ to find all 541,968 solutions. De Morgan's example isn't among them, because the proper name '`ELLEN`' doesn't qualify as a word by our conventions. But some of the squares might be "meaningful," at least poetically:

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

Just six of the solutions belong to the restricted vocabulary `WORDS`(500); three of them actually belong to `WORDS`(372), namely `**ASS|*IGHT|AGREE|SHEEP|STEPS`, where `***` is either `CLL` or `GLL` or `GRR`. (And `***` = `GRL` gives an *unsymmetric* 5×5 in `WORDS`(372). There are $(1787056 - 541968)/2 = 622544$ unsymmetric squares in `WORDS`(5757).)

31. Yes, 27 of them. The search is greatly facilitated by noting that the NE-to-SW diagonal word must be one of the 18 palindromes in `WORDS`(5757). '`SCABS|CANAL|ANGLE|BALED|SLEDS`', which belongs to `WORDS`(3025), has the most common words. [See the end of Chapter 18 in Babbage's *Passages from The Life of a Philosopher* (London: 1864).]

32. There are (717, 120386, 2784632, 6571160, 1117161, 13077, 6) of sizes 2×2, ..., 8×8, and none larger than this. Each of these runs needed fewer than 6 gigamems of computation. Example solutions with words as common as possible are

```
                                          CURTAIL  NEREIDES
                             HEART ESTATE UTERINE  ETERNISE
              ITS  AWAY      ERROR SLAVES REVERTS  RELOCATE
         TO   THE  WERE      ARGUE TALENT TREBLES  EROTIZED
         OF   SEE  AREA      ROUTE AVENUE AIRLINE  INCITERS
                   YEAR      TREES TENURE INTENSE  DIAZEPAM
                                   ESTEEM LESSEES  ESTERASE
                                                   SEEDSMEN
```

with the following numeric ranks of "minimax rarity" within their lists: `TO` = 2, `SEE` = 25, `AREA` = 86, `ERROR` = 438, `ESTEEM` = 1607, `TREBLES` = 5696, `ETERNISE` = 23623.

[Word squares go back thousands of years; '`SATOR|AREPO|TENET|OPERA|ROTAS`', a famous 5×5 example that is found in many places including the ruins of Pompeii, actually has fourfold symmetry. But 6×6 squares appear to have been unknown until William Winthrop, the U.S. consul in Malta(!), published '`CIRCLE|ICARUS|RAREST|CREATE|LUSTRE|ESTEEM`' in *Notes & Queries* (2) **8** (2 July 1859), page 8, claiming to have thereby "squared the circle." (If he had been told not to use a proper name like Icarus, he could have said '`CIRCLE|INURES|RUDEST|CREASE|LESSER|ESTERS`'.)]

> *The conclusion to be drawn about exercises of this kind is that four letters are nothing at all; that five letters are so easy that nothing is worth notice unless the combination have meaning; that six letters, done in any way, are respectable; and that seven letters would be a triumph.*
>
> — AUGUSTUS DE MORGAN, in *Notes & Queries* (3 September 1859)

Henry Dudeney constructed several 7×7 examples and used them in clever puzzles, beginning with '`PALATED|ANEMONE|LEVANTS|AMASSES|TONSIRE|ENTERER|DESSERT`' [*The Weekly Dispatch* (25 October and 8 November 1896)] and '`BOASTER|OBSCENE|ASSERTS|SCEPTRE|TERTIAN|ENTRANT|RESENTS`' [*The Weekly Dispatch* (21 November and 5 December 1897)]. Years later he was particularly pleased to have found '`NESTLES|ENTRANT|STRANGE|TRAITOR|LANTERN|ENGORGE|STERNER`' [*Strand* **55** (1918), 488; **56** (1919), 74; *The World's Best Word Puzzles* (1925), Puzzles 142 and 145]. M. Douglas McIlroy was the first to apply computers to this task [*Word Ways* **8** (1975), 195–197], discovering 52 examples such as '`WRESTLE|RENEWAL|ENPLANE|SELFDOM|TWADDLE|LANOLIN|ELEMENT`'. Then he turned to the more difficult problem of *double word squares*, which are unsymmetric and contain $2n$ *distinct* words: He presented 117 double squares, such as '`REPAST|AVESTA|CIRCUS|INSECT|SCONCE|MENTOR`', in *Word Ways* **9** (1976), 80–84. (His experiments allowed proper names, but avoided plurals and other derived word forms.)

For an excellent history of word squares and word cubes, chronicling the subsequent computer developments as well as extensive searches for 10×10 examples using vast dictionaries, see Ross Eckler, *Making the Alphabet Dance* (New York: St. Martin's Griffin, 1997), 188–203; *Tribute to a Mathemagician* (A. K. Peters, 2005), 85–91.

33. Working from bottom to top and right to left is equivalent to working from top to bottom and left to right on the word *reversals*. This idea does make the tries smaller; but unfortunately it makes the programs slower. For example, the 6×5 computation of answer 28 involves a 6347-node trie for the 6-letter words and a 63060-node compressed trie for the 5-letter words. Those sizes go down to 5188 and 56064, respectively, when we reverse the words; but the running time goes up from 380 $G\mu$ to 825 $G\mu$.

minimax
fourfold symmetry
palindromes
Winthrop
DE MORGAN
Dudeney
McIlroy
double word squares
word square, double
history of word squares
word cubes
word squares, history
Eckler

34. Leave out `face` and (of course) `dada`; the remaining eleven are fine.

35. Keep tables p_i, p'_{ij}, p''_{ijk}, s_i, s'_{ij}, s''_{ijk}, for $0 \le i, j, k < m$, each capable of storing a ternary digit, and initially zero. Also keep a table x_0, x_1, … of tentatively accepted words. Begin with $g \leftarrow 0$. Then for each input $w_j = abcd$, where $0 \le a, b, c, d < m$, set $x_g \leftarrow abcd$ and also do the following: Set $p_a \leftarrow p_a \dot{+} 1$, $p'_{ab} \leftarrow p'_{ab} \dot{+} 1$, $p''_{abc} \leftarrow p''_{abc} \dot{+} 1$, $s_d \leftarrow s_d \dot{+} 1$, $s'_{cd} \leftarrow s'_{cd} \dot{+} 1$, $s''_{bcd} \leftarrow s''_{bcd} \dot{+} 1$, where $x \dot{+} y = \min(2, x+y)$ denotes saturating ternary addition. Then if $s_{a'}p''_{b'c'd'} + s'_{a'b'}p'_{c'd'} + s''_{a'b'c'}p_{d'} = 0$ for all $x_k = a'b'c'd'$, where $0 \le k \le g$, set $g \leftarrow g+1$. Otherwise reject w_j and set $p_a \leftarrow p_a - 1$, $p'_{ab} \leftarrow p'_{ab} - 1$, $p''_{abc} \leftarrow p''_{abc} - 1$, $s_d \leftarrow s_d - 1$, $s'_{cd} \leftarrow s'_{cd} - 1$, $s''_{bcd} \leftarrow s''_{bcd} - 1$.

36. (a) The word bc appears in message $abcd$ if and only if $a \to b$, $b \to c$, and $c \to d$.

(b) For $0 \le k < r$, put vertex v into class k if the longest path from v has length k. Given any such partition, we can include all arcs from class k to class $j < k$ without increasing the path lengths. So it's a question of finding the maximum of $\sum_{0 \le j < k < r} p_j p_k$ subject to $p_0 + p_1 + \cdots + p_{r-1} = m$. The values $p_j = \lfloor (m+j)/r \rfloor$ achieve this (see exercise 7.2.1.4–68(a)). When $r = 3$ the maximum simplifies to $\lfloor m^2/3 \rfloor$.

37. (a) The factors of the period, 15 926 535 89 79 323 8314, begin at the respective boundary points 3, 5, 8, 11, 13, 15, 18 (and then $3 + 19 = 22$, etc.). Thus round 1 retains boundaries 5, 8, and 15. The second-round substrings $y_0 = 926$, $y_1 = 5358979$, $y_2 = 323831415$ have different lengths, so lexicographic comparison is unnecessary; the answer is $y_2 y_0 y_1 = x_{15} \ldots x_{33}$.

(b) Each substring consists of at least three substrings of the previous round.

(c) Let $a_0 = 0$, $b_0 = 1$, $a_{e+1} = a_e a_e b_e$, $b_{e+1} = a_e b_e b_e$; use a_e or b_e when $n = 3^e$.

(d) We use an auxiliary subroutine 'less(i)', which returns $[y_{i-1} < y_i]$, given $i > 0$: If $b_i - b_{i-1} \ne b_{i+1} - b_i$, return $[b_i - b_{i-1} < b_{i+1} - b_i]$. Otherwise, for $j = 0, 1, \ldots$, while $b_i + j < b_{i+1}$, if $x_{b_{i-1}+j} \ne x_{b_i+j}$ return $[x_{b_{i-1}+j} < x_{b_i+j}]$. Otherwise return 0.

The tricky part of the algorithm is to discard initial factors that aren't periodic. The secret is to let i_0 be the smallest index such that $y_{i-3} \ge y_{i-2} < y_{i-1}$; then we can be sure that a factor begins with y_i.

O1. [Initialize.] Set $x_j \leftarrow x_{j-n}$ for $n \le j < 2n$, $b_j \leftarrow j$ for $0 \le j < 2n$, and $t \leftarrow n$.

O2. [Begin a round.] Set $t' \leftarrow 0$. Find the smallest $i > 0$ such that less(i) = 0. Then find the smallest $j \ge i + 2$ such that less($j - 1$) = 1 and $j \le t + 2$. (If no such j exists, report an error: The input x was equal to one of its cyclic shifts.) Set $i \leftarrow i_0 \leftarrow j \bmod t$. (Now a dip of the period begins at i_0.)

O3. [Find the next factor.] Find the smallest $j \ge i + 2$ such that less ($j - 1$) = 1. If $j - i$ is even, go to O5.

O4. [Retain a boundary.] If $j < t$, set $b'_{t'} \leftarrow b_j$; otherwise set $b'_k \leftarrow b'_{k-1}$ for $t' \ge k > 0$ and $b'_0 \leftarrow b_{j-t}$. Finally set $t' \leftarrow t' + 1$.

O5. [Done with round?] If $j < i_0 + t$, set $i \leftarrow j$ and return to O3. Otherwise, if $t' = 1$, terminate; σx begins at item $x_{b'_0}$. Otherwise set $t \leftarrow t'$, $b_k \leftarrow b'_k$ for $0 \le k < t$, and $b_k \leftarrow b_{k-t} + n$ for $k \ge t$ while $b_{k-t} < n$. Return to O2. ▮

(e) Say that a "superdip" is a dip of odd length followed by zero or more dips of even length. Any infinite sequence y that begins with an odd-length dip has a unique factorization into superdips. Those superdips can, in turn, be regarded as atomic elements of a higher-level string that can be factored into dips. The result σx of Algorithm O is an infinite periodic sequence that allows repeated factorization into infinite periodic sequences of superdips at higher and higher levels, until becoming constant.

Notice that the first dip of σx ends at position i_0 in the algorithm, because its length isn't 2. Therefore we can prove the commafree property by observing that, if codeword $\sigma x''$ appears within the concatenation $\sigma x\,\sigma x'$ of two codewords, its superdip factors are also superdip factors of those codewords. This yields a contradiction if any of σx, $\sigma x'$, or $\sigma x''$ is a superdip. Otherwise the same observation applies to the superdip factors at the next level. [Eastman's original algorithm was essentially the same, but presented in a more complicated way; see *IEEE Trans.* **IT-11** (1965), 263–267. R. A. Scholtz subsequently discovered an interesting and totally different way to define the set of codewords produced by Algorithm O, in *IEEE Trans.* **IT-15** (1969), 300–306.]

Eastman
Scholtz
aperiodic
stacks
queues
memory constraints, historic
unordered sets
Linked lists
deletion

38. Let $f_k(m)$ be the number of dips of length k for which $m > z_1$ and $z_k < m$. The number of such sequences with $z_2 = j$ is $(m-j-1)\binom{m-j+k-3}{k-2} = (k-1)\binom{m-j+k-3}{k-1}$; summing for $0 \le j < m$ gives $f_k(m) = (k-1)\binom{m+k-2}{k}$. Thus $F_m(z) = \sum_{k=0}^{\infty} f_k(m)z^k = (mz-1)/(1-z)^m$. (The fact that $f_0(m) = -1$ in these formulas turns out to be useful!)

Algorithm O finishes in one round if and only if some cyclic shift of x is a superdip. The number of aperiodic x that finish in one round is therefore $n[z^n]\,G_m(z)$, where

$$G_m(z) = \frac{F_m(-z) - F_m(z)}{F_m(-z) + F_m(z)} = \frac{(1+mz)(1-z)^m - (1-mz)(1+z)^m}{(1+mz)(1-z)^m + (1-mz)(1+z)^m}.$$

To get the stated probability, divide by $\sum_{d\backslash n} \mu(d) m^{n/d}$, the number of aperiodic x. (See Eq. 7.2.1.1–(60). For $n = 3, 5, 7, 9$ these probabilities are 1, 1, 1, and $1-3/\binom{m^3-1}{3}$.)

39. If so, it couldn't have 0011, 0110, 1100, or 1001.

40. That section considered such representations of stacks and queues, but not of unordered sets, because large blocks of sequential memory were either nonexistent or ultra-expensive in olden days. Linked lists were the only decent option for families of variable-size sets, because they could more readily fit in a limited high-speed memory.

41. (a) The blue word x with $\alpha = \texttt{d}$ (namely 1101) appears in its `P2` list at location `5e`.
(b) The `P3` list for words of the form 010∗ is empty. (Both 0100 and 0101 are red.)

42. (a) The `S2` list of 0010 has become closed (hence 0110 and 1110 are hidden).
(b) Word 1101 moved to the former position of 1001 in its `S1` list, when 1001 became red. (Previously 1011 had moved to the former position of 0001.)

43. In this case, which of course happens rarely, it's safe to set all elements of `STAMP` to zero and set $\sigma \leftarrow 1$. (Do *not* be tempted to save one line of code by setting all `STAMP` elements to -1 and leaving $\sigma = 0$. That might fail when σ reaches the value -1!)

44. (a) Set $r \leftarrow m+1$. Then for $k \leftarrow 0, 1, \ldots, f-1$, set $t \leftarrow \texttt{FREE[}k\texttt{]}$, $j \leftarrow \texttt{MEM[CLOFF+}4t + m^4\texttt{]} - (\texttt{CLOFF} + 4t)$, and if $j < r$ set $r \leftarrow j$, $c \leftarrow t$; break out of the loop if $r = 0$.
(b) If $r > 0$ set $x \leftarrow \texttt{MEM[CLOFF} + 4cl(\texttt{ALF[}x\texttt{]})\texttt{]}$.
(c) If $r > 1$ set $q \leftarrow 0$, $p' \leftarrow \texttt{MEM[PP]}$, and $p \leftarrow \texttt{POISON}$. While $p < p'$ do the following steps: Set $y \leftarrow \texttt{MEM[}p\texttt{]}$, $z \leftarrow \texttt{MEM[}p+1\texttt{]}$, $y' \leftarrow \texttt{MEM[}y+m^4\texttt{]}$, and $z' \leftarrow \texttt{MEM[}z+m^4\texttt{]}$. (Here y and z point to the heads of prefix or suffix lists; y' and z' point to the tails.) If $y = y'$ or $z = z'$, delete entry p from the poison list; this means, as in (18), to set $p' \leftarrow p'-2$, and if $p \ne p'$ to store$(p, \texttt{MEM[}p'\texttt{]})$ and store$(p+1, \texttt{MEM[}p'+1\texttt{]})$. Otherwise set $p \leftarrow p+2$; if $y'-y \ge z'-z$ and $y'-y > q$, set $q \leftarrow y'-y$ and $x \leftarrow \texttt{MEM[}z\texttt{]}$; if $y'-y < z'-z$ and $z'-z > q$, set $q \leftarrow z'-z$ and $x \leftarrow \texttt{MEM[}y\texttt{]}$. Finally, after p has become equal to p', store$(\texttt{PP}, p')$ and set $c \leftarrow cl(\texttt{ALF[}x\texttt{]})$. (Experiments show that this "max kill" strategy for $r > 1$ slightly outperforms a selection strategy based on r alone.)

45. (a) First there's a routine rem(α, δ, o) that removes an item from a list, following the protocol (21): Set $p \leftarrow \delta + o$ and $q \leftarrow \texttt{MEM[}p+m^4\texttt{]} - 1$. If $q \ge p$ (meaning that list p isn't closed or being killed), store$(p + m^4, q)$, set $t \leftarrow$ MEM[$\alpha + o - m^4$]; and if $t \neq q$ also set $y \leftarrow$ MEM[q], store(t, y), and store(ALF[y] $+ o - m^4, t)$.

Now, to redden x we set $\alpha \leftarrow$ ALF[x], store$(\alpha,$ RED$)$; then rem$(\alpha, p_1(\alpha),$ P1OFF$)$, rem$(\alpha, p_2(\alpha),$ P2OFF$)$, ..., rem$(\alpha, s_3(\alpha),$ S3OFF$)$, and rem$(\alpha, 4cl(\alpha),$ CLOFF$)$.

(b) A simple routine close(δ, o) closes list $\delta{+}o$: Set $p \leftarrow \delta{+}o$ and $q \leftarrow$ MEM[$p + m^4$]; if $q \neq p - 1$, store$(p + m^4, p - 1)$.

Now, to green x we set $\alpha \leftarrow$ ALF[x], store$(\alpha,$ GREEN$)$; then close$(p_1(\alpha),$ P1OFF$)$, close$(p_2(\alpha),$ P2OFF$)$, ..., close$(s_3(\alpha),$ S3OFF$)$, and close$(4cl(\alpha),$ CLOFF$)$. Finally, for $p \leq r < q$ (using the p and q that were just set within 'close'), if MEM[r] $\neq x$ redden MEM[r].

(c) First set $p' \leftarrow$ MEM[PP] $+6$, and store$(p'-6, p_1(\alpha)+$S1OFF$)$, store$(p'-5, s_3(\alpha)+$ P3OFF$)$, store$(p'-4, p_2(\alpha) +$ S2OFF$)$, store$(p'-3, s_2(\alpha) +$ P2OFF$)$, store$(p'-2, p_3(\alpha) +$ S3OFF$)$, store$(p'-1, s_1(\alpha) +$ P1OFF$)$; this adds the three poison items (27).

Then set $p \leftarrow$ POISON and do the following while $p < p'$: Set y, z, y', z' as in answer 44(c), and delete poison entry p if $y = y'$ or $z = z'$. Otherwise if $y' < y$ and $z' < z$, go to C6 (a poisoned suffix-prefix pair is present). Otherwise if $y' > y$ and $z' > z$, set $p \leftarrow p + 2$. Otherwise if $y' < y$ and $z' > z$, store$(z + m^4, z)$, redden MEM[r] for $z \leq r < z'$, and delete poison entry p. Otherwise (namely if $y' > y$ and $z' < z$), store$(y + m^4, y)$, redden MEM[r] for $y \leq r < y'$, and delete poison entry p.

Finally, after p has become equal to p', store(PP$, p')$.

46. Exercise 37 exhibits such codes explicitly for all odd n. The earliest papers on the subject gave solutions for $n = 2$, 4, 6, 8. Yoji Niho subsequently found a code for $n = 10$ but was unable to resolve the case $n = 12$ [*IEEE Trans.* **IT-19** (1973), 580–581].

This problem can readily be encoded in CNF and given to a SAT solver. The case $n = 10$ involves 990 variables and 8.6 million clauses, and is solved by Algorithm 7.2.2.2C in 10.5 gigamems. The case $n = 12$ involves 4020 variables and 175 million clauses. After being split into seven independent subproblems (by appending mutually exclusive unit clauses), it was proved *unsatisfiable* by that algorithm after about 86 teramems of computation.

So the answer is "No." The maximum-size code for $n = 12$ remains unknown.

47. (a) There are 28 commafree binary codes of size 3 and length 4; Algorithm C produces half of them, because it assumes that cycle class [0001] is represented by 0001 or 0010. They form eight equivalence classes, two of which are symmetric under the operation of complementation-and-reflection; representatives are $\{0001, 0011, 0111\}$ and $\{0010, 0011, 1011\}$. The other six are represented by $\{0001, 0110, 0111$ or $1110\}$, $\{0001, 1001, 1011$ or $1101\}$, $\{0001, 1100, 1101\}$, $\{0010, 0011, 1101\}$.

(b) Algorithm C produces half of the 144 solutions, which form twelve equivalence classes. Eight are represented by $\{0001$, 0002, 1001, 1002, 2201, 2001, 2002, 2011, 2012, 2102, 2112, 2122 or $2212\}$ and ($\{0102, 1011, 1012\}$ or $\{2010, 2101, 2101\}$) and ($\{1202, 2202, 2111\}$ or $\{2021, 2022, 1112\}$); four are represented by $\{0001$, 0020, 0021, 0022, 1001, 1020, 1021, 1022, 1121 or 1211, 1201, 1202, 1221, 2001, 2201, $2202\}$ and ($\{1011, 1012, 2221\}$ or $\{1101, 2101, 1222\}$).

(c) Algorithm C yields half of the 2304 solutions, which form 48 equivalence classes. Twelve classes have unique representatives that omit cycle classes [0123], [0103], [1213], one such being the code $\{0010$, 0020, 0030, 0110, 0112, 0113, 0120, 0121, 0122, 0130, 0131, 0132, 0133, 0210, 0212, 0213, 0220, 0222, 0230, 0310, 0312, 0313, 0320, 0322, 0330, 0332, 0333, 1110, 1112, 1113, 2010, 2030, 2110, 2112, 2113, 2210, 2212, 2213, 2230, 2310, 2312, 2313, 2320, 2322, 2330, 2332, 2333, 3110, 3112, 3113, 3210, 3212, 3213, 3230, 3310, 3312, $3313\}$. The others each have two representatives that omit

Niho
CNF
SAT solver
unit clauses

classes [0123], [0103], [0121], one such pair being the code {0001, 0002, 0003, 0201, 0203, 1001, 1002, 1003, 1011, 1013, 1021, 1022, 1023, 1031, 1032, 1033, 1201, 1203, 1211, 1213, 1221, 1223, 1231, 1232, 1233, 1311, 1321, 1323, 1331, 2001, 2002, 2003, 2021, 2022, 2023, 2201, 2203, 2221, 2223, 3001, 3002, 3003, 3011, 3013, 3021, 3022, 3023, 3031, 3032, 3033, 3201, 3203, 3221, 3223, 3321, 3323, 3331} and its isomorphic image under reflection and (01)(23).

48. (The maximum size of such a code is currently unknown. Algorithm C isn't fast enough to solve this problem on a single computer, but a sufficiently large cluster of machines and/or an improved algorithm should be able to discover the answer. The case $m = 3$ and $n = 6$ is also currently unsolved; a SAT solver shows quickly that a full set of $(3^6 - 3^3 - 3^2 + 3^1)/6 = 116$ codewords cannot be achieved.)

49. The 3-bit sequences 101, 111, 110 were rejected before seeing 000. In general, to make a uniformly random choice from q possibilities, the text suggests looking at the next $t = \lceil \lg q \rceil$ bits $b_1 \ldots b_t$. If $(b_1 \ldots b_t)_2 < q$, we use choice $(b_1 \ldots b_t)_2 + 1$; otherwise we reject $b_1 \ldots b_t$ and try again. [This simple method is optimum when $q \le 4$, and the best possible running time for other values of q uses more than half as many bits. But a better scheme is available for $q = 5$, using only $3\frac{1}{3}$ bits per choice instead of $4\frac{4}{5}$; and for $q = 6$, one random bit reduces to the case $q = 3$. See D. E. Knuth and A. C. Yao, *Algorithms and Complexity*, edited by J. F. Traub (Academic Press, 1976), 357–428, §2.]

50. It's the number of nodes on level $l+1$ (depth l) of the search tree. (Hence we can estimate the profile. Notice that $D = D_1 \ldots D_{l-1}$ in step E2 of Algorithm E.)

51. $Z_0 = C()$, $Z_{l+1} = c() + D_1 c(X_1) + D_1 D_2 c(X_1 X_2) + \cdots + D_1 \ldots D_l c(X_1 \ldots X_l) + D_1 \ldots D_{l+1} C(X_1 \ldots X_{l+1})$.

52. (a) True: The generating function is $z(z+1)\ldots(z+n-1)/n!$; see Eq. 1.2.10–(9).

(b) For instance, suppose $Y_1 Y_2 \ldots Y_l = 1457$ and $n = 9$. Alice's probability is $\frac{1}{1}\frac{1}{2}\frac{2}{3}\frac{1}{4}\frac{1}{5}\frac{5}{6}\frac{1}{7}\frac{7}{8}\frac{8}{9} = \frac{1}{3}\frac{1}{4}\frac{1}{6}\frac{1}{9}$. Elmo obtains $X_1 X_2 \ldots X_l = 7541$ with probability $\frac{1}{9}\frac{1}{6}\frac{1}{4}\frac{1}{3}$.

(c) The upper tail inequality (see exercise 1.2.10–22 with $\mu = H_n$) tells us that $\Pr(l \ge (\ln n)(\ln\ln n)) \le \exp(-(\ln n)(\ln\ln n)(\ln\ln\ln n) + O(\ln n)(\ln\ln n))$.

(d) If $k \le n/3$ we have $\sum_{j=0}^{k}\binom{n}{j} \le 2\binom{n}{k}$. By exercise 1.2.6–67, the number of nodes on the first $(\ln n)(\ln\ln n)$ levels is therefore at most $2(ne/((\ln n)(\ln\ln n)))^{(\ln n)(\ln\ln n)}$.

53. The key idea is to introduce recursive formulas analogous to (29):

$$m(x_1 \ldots x_l) = c(x_1 \ldots x_l) + \min\bigl(m(x_1 \ldots x_l x_{l+1}^{(1)})d, \ldots, m(x_1 \ldots x_l x_{l+1}^{(d)})d\bigr);$$
$$M(x_1 \ldots x_l) = c(x_1 \ldots x_l) + \max\bigl(M(x_1 \ldots x_l x_{l+1}^{(1)})d, \ldots, M(x_1 \ldots x_l x_{l+1}^{(d)})d\bigr);$$
$$\widehat{C}(x_1 \ldots x_l) = c(x_1 \ldots x_l)^2 + \sum_{i=1}^{d}\bigl(\widehat{C}(x_1 \ldots x_l x_{l+1}^{(i)})d + 2c(x_1 \ldots x_l)C(x_1 \ldots x_l x_{l+1}^{(i)})\bigr).$$

They can be computed via auxiliary arrays `MIN`, `MAX`, `KIDS`, `COST`, and `CHAT` as follows:

At the beginning of step B2, set `MIN[l]` $\leftarrow \infty$, `MAX[l]` $\leftarrow$ `KIDS[l]` $\leftarrow$ `COST[l]` $\leftarrow$ `CHAT[l]` $\leftarrow 0$. Set `KIDS[l]` $\leftarrow$ `KIDS[l]` $+ 1$ just before $l \leftarrow l + 1$ in step B3.

At the beginning of step B5, set $m \leftarrow c(x_1 \ldots x_{l-1}) +$ `KIDS[l]` $\times$ `MIN[l]`, $M \leftarrow c(x_1 \ldots x_{l-1}) +$ `KIDS[l]` $\times$ `MAX[l]`, $C \leftarrow c(x_1 \ldots x_{l-1}) +$ `COST[l]`, $\widehat{C} \leftarrow c(x_1 \ldots x_{l-1})^2 +$ `KIDS[l]` $\times$ `CHAT[l]` $+ 2 \times$ `COST[l]`. Then, after $l \leftarrow l - 1$ is positive, set `MIN[l]` $\leftarrow \min(m,$ `MIN[l]`$)$, `MAX[l]` $\leftarrow \max(M,$ `MAX[l]`$)$, `COST[l]` $\leftarrow$ `COST[l]` $+ C$, `CHAT[l]` $\leftarrow$ `CHAT[l]` $+ \widehat{C}$. But when l reaches zero in step B5, return the values m, M, C, $\widehat{C} - C^2$.

uniformly random
reject
Knuth
Yao
complexity of calculation
computational complexity
Traub
profile
generating function
tail inequality
cumulative binomial distribution
recursive formulas

54. Let $p(i) = p_{X_1 \ldots X_{l-1}}(y_i)$, and simply set $D \leftarrow D/p(I)$ instead of $D \leftarrow Dd$. Then node $x_1 \ldots x_l$ is reached with probability $\Pi(x_1 \ldots x_l) = p(x_1)p_{x_1}(x_2) \ldots p_{x_1 \ldots x_l}(x_l)$, and $c(x_1 \ldots x_l)$ has weight $1/\Pi(x_1 \ldots x_l)$ in S; the proof of Theorem E goes through as before. Notice that $p(I)$ is the *a posteriori* probability of having taken branch I.

(The formulas of answer 53 should now use '$/p(i)$' instead of 'd'; and that algorithm should be modified appropriately, no longer needing the `KIDS` array.)

55. Let $p_{X_1 \ldots X_{l-1}}(y_i) = C(x_1 \ldots x_{l-1}y_i)/(C(x_1 \ldots x_{l-1}) - c(x_1 \ldots x_{l-1}))$. (Of course we generally need to know the cost of the tree before we know the exact values of these ideal probabilities, so we cannot achieve zero variance in practice. But the form of this solution shows what kinds of bias are likely to reduce the variance.)

56. The effects of lookahead, dynamic ordering, and reversible memory are all captured easily by a well-designed cost function at each node. But there's a fundamental difference in step C2, because different codeword classes can be selected for branching at the same node (that is, with the same ancestors $x_1 \ldots x_{l-1}$) after C5 has undone the effects of a prior choice. The level l never surpasses $L + 1$, but in fact the search tree involves hidden levels of branching that are implicitly combined into single nodes.

Thus it's best to view Algorithm C's search tree as a sequence of *binary* branches: Should x be one of the codewords or not? (At least this is true when the "max kill" strategy of answer 44 has selected the branching variable x. But if $r > 1$ and the poison list is empty, an r-way branch is reasonable (or an $(r + 1)$-way branch when the slack is positive), because r will be reduced by 1 and the same class c will be chosen after x has been explored.)

If x has been selected because it kills many other potential codewords, we probably should bias the branch probability as in exercise 54, giving smaller weight to the "yes" branch because the branch that includes x is less likely to lead to a large subtree.

57. Let $p_k = 1/D^{(k)}$ be the probability that Algorithm E terminates at the kth leaf. Then $\sum_{k=1}^{M}(1/M)\lg(1/(Mp_k))$ is the Kullback–Leibler divergence $D(q\|p)$, where q is the uniform distribution (see exercise MPR–121). Hence $\frac{1}{M}\sum_{k=1}^{M}\lg D^{(k)} \geq \lg M$. (The result of this exercise is essentially true in *any* probability distribution.)

58. Let ∞ be any convenient value $\geq n$. When vertex v becomes part of the path we will perform a two-phase algorithm. The first phase identifies all "tarnished" vertices, whose `DIST` must change; these are the vertices u from which every path to t passes through v. It also forms a queue of "resource" vertices, which are untarnished but adjacent to tarnished ones. The second phase updates the `DIST`s of all tarnished vertices that are still connected to t. Each vertex has `LINK` and `STAMP` fields in addition to `DIST`.

For the first phase, set $d \leftarrow$ `DIST`(v), `DIST`$(v) \leftarrow \infty + 1$, `R` $\leftarrow \Lambda$, `T` $\leftarrow v$, `LINK`$(v) \leftarrow \Lambda$, then do the following while `T` $\neq \Lambda$: $(*)$ Set $u \leftarrow$ `T`, `T` $\leftarrow$ `S` $\leftarrow \Lambda$. For each w — u, if `DIST`$(w) < d$ do nothing (this happens only when $u = v$); if `DIST`$(w) \geq \infty$ do nothing (w is gone or already known to be tarnished); if `DIST`$(w) = d$, make w a resource (see below); otherwise `DIST`$(w) = d + 1$. If w has no neighbor at distance d, w is tarnished: Set `LINK`$(w) \leftarrow$ `T`, `DIST`$(w) \leftarrow \infty$, `T` $\leftarrow w$. Otherwise make w a resource (see below). Then set $u \leftarrow$ `LINK`(u), and return to $(*)$ if $u \neq \Lambda$.

The queue of resources will start at `R`. We will stamp each resource with v so that nothing is added twice to that queue. To make w a resource when `DIST`$(w) = d$, do the following (unless $u = v$ or `STAMP`$(w) = v$): Set `STAMP`$(w) \leftarrow v$; if `R` $= \Lambda$, set `R` $\leftarrow$ `RT` $\leftarrow w$; otherwise set `LINK(RT)` $\leftarrow w$ and `RT` $\leftarrow w$. To make w a resource when `DIST`$(w) = d + 1$ and $u \neq v$ and `STAMP`$(w) \neq v$, put it first on stack `S` as follows: Set `STAMP`$(w) \leftarrow v$; if `S` $= \Lambda$, set `S` $\leftarrow$ `SB` $\leftarrow w$; otherwise set `LINK`$(w) \leftarrow$ `S`, `S` $\leftarrow w$.

a posteriori
cost function
search tree
poison list
slack
bias
Kullback
Leibler
divergence
queue

Finally, when $u = \Lambda$, we append `S` to `R`: Nothing needs to be done if $\mathtt{S} = \Lambda$. Otherwise, if $\mathtt{R} = \Lambda$, set $\mathtt{R} \leftarrow \mathtt{S}$ and $\mathtt{RT} \leftarrow \mathtt{SB}$; but if $\mathtt{R} \neq \Lambda$, set $\mathtt{LINK(RT)} \leftarrow \mathtt{S}$ and $\mathtt{RT} \leftarrow \mathtt{SB}$. (These shenanigans keep the resource queue in order by `DIST`.)

Phase 2 operates as follows: Nothing needs to be done if $\mathtt{R} = \Lambda$. Otherwise we set $\mathtt{LINK(RT)} \leftarrow \Lambda$, $\mathtt{S} \leftarrow \Lambda$, and do the following while $\mathtt{R} \neq \Lambda$ or $\mathtt{S} \neq \Lambda$: (i) If $\mathtt{S} = \Lambda$, set $d \leftarrow \mathtt{DIST(R)}$. Otherwise set $u \leftarrow \mathtt{S}$, $d \leftarrow \mathtt{DIST}(u)$, $\mathtt{S} \leftarrow \Lambda$; while $u \neq \Lambda$, update the neighbors of u and set $u \leftarrow \mathtt{LINK}(u)$. (ii) While $\mathtt{R} \neq \Lambda$ and $\mathtt{DIST(R)} = d$, set $u \leftarrow \mathtt{R}$, $\mathtt{R} \leftarrow \mathtt{LINK}(u)$, and update the neighbors of u. In both cases "update the neighbors of u" means to look at all $w \longrightarrow u$, and if $\mathtt{DIST}(w) = \infty$ to set $\mathtt{DIST}(w) \leftarrow d+1$, $\mathtt{STAMP}(w) \leftarrow v$, $\mathtt{LINK}(w) \leftarrow \mathtt{S}$, and $\mathtt{S} \leftarrow w$. (It works!)

generating function
Venn diagram
recurrence

59. (a) Compute the generating function $g(z)$ (see exercise 7.1.4–209) and then $g'(1)$.

(b) Let (A, B, C) denote paths that touch (center, NE corner, SW corner). Recursively compute eight counts $(c_0, \ldots, c_7)$ at each node, where c_j counts paths π with $j = 4[\pi \in A] + 2[\pi \in B] + [\pi \in C]$. At the sink node $\boxed{\top}$ we have $c_0 = 1$, $c_1 = \cdots = c_7 = 0$. Other nodes have the form $x = (\bar{e}?\ x_l\colon x_h)$ where e is an edge. Two edges go across the center and affect A; three edges affect each of B and C. Say that those edges have types 4, 2, 1, respectively; other edges have type 0. Suppose the counts for x_l and x_h are $(c'_0, \ldots, c'_7)$ and $(c''_0, \ldots, c''_7)$, and e has type t. Then count c_j for node x is $c'_j + [t=0]c''_j + [t \mathbin{\&} j \neq 0](c''_j + c''_{j-t})$.

(This procedure yields the following exact "Venn diagram" set counts at the root: $c_0 = |\overline{A} \cap \overline{B} \cap \overline{C}| = 76536853848890196480916\,04$; $c_1 = c_2 = |\overline{A} \cap \overline{B} \cap C| = |\overline{A} \cap B \cap \overline{C}| = 7755019053779199171839134$; $c_3 = |\overline{A} \cap B \cap C| = 7857706970503366819944024$; $c_4 = |A \cap \overline{B} \cap \overline{C}| = 4888524166534573765995071$; $c_5 = c_6 = |A \cap \overline{B} \cap C| = |A \cap B \cap \overline{C}| = 4949318991771252110605148$; $c_7 = |A \cap B \cap C| = 5010950157283718807987280$.)

60. Yes, the paths are less chaotic and the estimates are better:

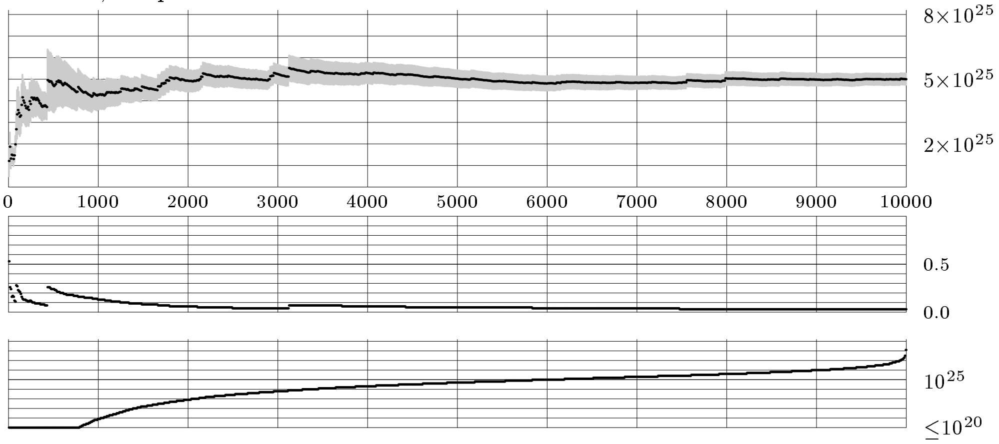

61. (a) Let x_k be the number of nodes at distance $k-1$ from the root.

(b) Let $Q_n^{(m)} = P_n^{(1)} + \cdots + P_n^{(m)}$. Then we have the joint recurrence $P_1^{(m)} = 1$, $P_{n+1}^{(m)} = Q_n^{(2m)}$; in particular, $Q_1^{(m)} = m$. And for $n \geq 2$, we have $Q_n^{(m)} = \sum_{k=1}^{n} a_{nk}\binom{m}{k}$ for certain constants a_{nk} that can be computed as follows: Set $t_k \leftarrow P_n^{(k)}$ for $1 \leq k \leq n$. Then for $k = 2, \ldots, n$ set $t_n \leftarrow t_n - t_{n-1}, \ldots, t_k \leftarrow t_k - t_{k-1}$. Finally $a_{nk} \leftarrow t_k$ for $1 \leq k \leq n$. For example, $a_{21} = a_{22} = 2$; $a_{31} = 6$, $a_{32} = 14$, $a_{33} = 8$. The numbers $P_n^{(m)}$ have $O(n^2 + n \log m)$ bits, so this method needs $O(n^5)$ bit operations to compute P_n.

Cook
Kleber
Mahler
unique solution
Gridgeman
Robertson
Munro
exact cover problem
invariant

(c) $P_n^{(m)}$ corresponds to random paths with $X_1 = m$, $D_k = 2X_k$, $X_{k+1} = \lceil 2U_k X_k \rceil$, where each U_k is an independent uniform deviate. Therefore $P_n^{(m)} = \mathrm{E}(D_1 \ldots D_{n-1})$ is the number of nodes on level n of an infinite tree. We have $X_{k+1} \geq 2^k U_1 \ldots U_k m$, by induction; hence $P_n^{(m)} \geq \mathrm{E}\bigl(2^{\binom{n}{2}} U_1^{n-2} U_2^{n-3} \ldots U_{n-2}^1 m^{n-1}\bigr) = 2^{\binom{n}{2}} m^{n-1}/(n-1)!$.

[M. Cook and M. Kleber have discussed similar sequences in *Electronic Journal of Combinatorics* **7** (2000), #R44. See also K. Mahler's asymptotic formula for binary partitions, in *J. London Math. Society* **15** (1940), 115–123, which shows that $\lg P_n = \binom{n}{2} - \lg(n-1)! + \binom{\lg n}{2} + O(1)$.]

62. Random trials indicate that the expected number of 2-regular graphs is ≈ 3.115, and that the number of disjoint pairs is (0, 1, …, 9, and $\geq$10) approximately (74.4, 4.3, 8.7, 1.3, 6.2, 0.2, 1.5, 0.1, 2.0, 0.0, and 12.2) percent of the time. If the cubes are restricted to cases where each color occurs at least five times, these numbers change to ≈ 4.89 and (37.3, 6.6, 17.5, 4.1, 16.3, 0.9, 5.3, 0.3, 6.7, 0.2, 5.0).

However, the concept of "unique solution" is tricky, because a 2-regular graph with k cycles yields 2^k ways to position the cubes. Let's say that a set of cubes has a *strongly unique* solution if (i) it has a unique disjoint pair of 2-regular graphs, and furthermore (ii) both elements of that pair are n-cycles. Such sets occur with probability only about 0.3% in the first case, and 0.4% in the second.

[N. T. Gridgeman, in *Mathematics Magazine* **44** (1971), 243–252, showed that puzzles with four cubes and four colors have exactly 434 "types" of solutions.]

63. It's easy to find such examples at random, as in the second part of the previous answer, since strongly unique sets occur about 0.5% of the time (and weakly unique sets occur with probability $\approx 8.4\%$). For example, the pairs of opposite faces might be $(12, 13, 34)$, $(02, 03, 14)$, $(01, 14, 24)$, $(04, 13, 23)$, $(01, 12, 34)$.

(Incidentally, if we require each color to occur exactly *six* times, every set of cubes that has at least one solution will have at least *three* solutions, because the "hidden" pairs can be chosen in three ways.)

64. Each of these cubes can be placed in 16 different ways that contribute legitimate letters to all four of the visible words. (A cube whose faces contain only letters in {C, H, I, N, O, U, X, Z} can be placed in 24 ways. A cube with a pattern like $\boxed{\mathsf{B}\,|\,\rhd\,|\,\mathsf{D}}$ cannot be placed at all.) We can restrict the first cube to just two placements; thus there are $2 \cdot 16 \cdot 16 \cdot 16 \cdot 16 = 131072$ ways to place those cubes without changing their order. Of these, only 6144 are "compatible," in the sense that no right-side-up-only letter appears together with an upside-down-only letter in the same word.

The 6144 compatible placements can then each be reordered in $5! = 120$ ways. One of them, whose words before reordering are GRHTI, NCICY, O⅁ꓤMN, INN∩O, leads to the unique solution. (There's a partial solution with three words out of four. There also are 39 ways to get two valid words, including one that has UNTIL adjacent to HOURS, and several with SYRUP opposite ECHOS.)

65. E. Robertson and I. Munro, in *Utilitas Mathematica* **13** (1978), 99–116, have reduced the exact cover problem to this problem.

66. Call the rays N, NE, E, SE, S, SW, W, NW; call the disks 1, 2, 3, 4 from inside to outside. We can keep disk 1 fixed. The sum of rays N, S, E, W must be 48. It is 16 (on disk 1) plus 13 or 10 (on disk 2) plus 8 or 13 (on disk 3) plus 11 or 14. So it is attained either as shown, or after rotating disks 2 and 4 clockwise by 45°. (Or we could rotate any disk by a multiple of 90°, since that keeps the desired sum invariant.)

Ernst
Goldstein
Slocum
Botermans
Washington Monument Puzzle, see Fool's Disk
Rotating Century Puzzle, see Fool's Disk
Safe Combination Puzzle, see Fool's Disk
digraph
stable
$d^+(v)$, the out-degree of v
bitwise tricks
$\nu(n)$, the sideways sum
sideways addition
Japanese arrow puzzles
Natsuhara

Next, with optional 90° rotations, we must make the sum of rays N + S equal to 24. In the first solution above it is 9 plus (6 or 7) plus (4 or 4) plus (7 or 4), hence never 24. But in the other solution it's 9 plus (4 or 6) plus (4 or 4) plus (5 or 9); hence we must rotate disk 2 clockwise by 90°, and possibly also disk 3. However, 90° rotation of disk 3 would make the NE + SW sum equal to 25, so we musn't move it.

Finally, to get NE's sum to be 12, via optional rotations by 180°, we have 1 plus (2 or 5) plus (1 or 5) plus (3 or 4); we must shift disks 3 and 4. Hurrah: That makes all eight rays correct. Factoring twice has reduced 8^3 trials to $2^3 + 2^3 + 2^3$.

[See George W. Ernst and Michael M. Goldstein, *JACM* **29** (1982), 1–23. Such puzzles go back to the 1800s; three early examples are illustrated on pages 28 of Slocum and Botermans's *New Book of Puzzles* (1992). One of them, with six rings and six rays, factors from 6^5 trials to $2^5 + 3^5$. A five-ray puzzle would have defeated factorization.]

67. Call the cards **1525**, **5113**, ..., **3755**. The key observation is that all 12 sums must be odd, so we can first solve the problem mod 2. For this purpose we may call the cards **1101**, **1111**, ..., **1111**; only three cards now change under rotation, namely **1101**, **0100**, and **1100** (which are the mod 2 images of **1525**, **4542**, and **7384**).

A second observation is that each solution gives $6 \times 6 \times 2$ others, by permuting rows and/or columns and/or by rotating all nine cards. Hence we can assume that the upper left card is **0011** (**8473**). Then **0100** (**4542**) must be in the first column, possibly rotated to **0001** (**4245**), to preserve parity in the left two black sums. We can assume that it's in row 2. In fact, after retreating from 13 mod 2 to 13, we see that it *must* be rotated. Hence the bottom left card must be either **4725**, **7755**, or **3755**.

Similarly we see that **1101** (**1525**) must be in the first row, possibly rotated to **0111** (**2515**); we can put it in column 2. It must be rotated, and the top right card must be **3454** or **3755**. This leaves just six scenarios to consider, and we soon obtain the solution: **8473**, **2515**, **3454**; **4245**, **2547**, **7452**; **7755**, **1351**, **5537**.

68. In general, let's say that a vertex labeling of a digraph is *stable* if v's label is the number of distinct labels among $\{w \mid v \longrightarrow w\}$, for all v. We wish to find all stable labelings that extend a given partial labeling. We may assume that no vertex is a sink.

Let $\Lambda(v)$ be a set of digits that includes every label that v could possibly have, in a solution to this extension problem. Initially, $\Lambda(v) = \{d\}$ if v's label is supposed to be d; otherwise $\Lambda(v) = \{1, \ldots, d^+(v)\}$. These sets are conveniently represented as the binary numbers $L(v) = \sum\{2^{k-1} \mid k \in \Lambda(v)\}$. Our goal is to reduce each $L(v)$ to a 1-bit number. A nice backtrack routine called "refine(v)" proves to be helpful in this regard.

Let $v_0 = v$ and let $v_1, \ldots, v_n$ be v's successors. Let $a_j = L(v_j)$. Following the outline of Algorithm B, we let $x_l \subseteq a_l$ be a 1-bit number, accepted in step B3 only if $2^{\nu s_l - 1} \subseteq g_l$, where $s_l = x_1 \mid \cdots \mid x_l$ and where the goal sets g_l are defined by $g_n = a_0$, $g_l = (g_{l+1} \mid g_{l+1} \gg 1) \,\&\, (2^l - 1)$. We start with all $b_j \leftarrow 0$; then when visiting a solution $x_1 \ldots x_n$, we set $b_j \leftarrow b_j \mid x_j$ for $1 \le j \le n$, and $b_0 \leftarrow b_0 \mid 2^{\nu s_n - 1}$. After finding all solutions we'll have $b_j \subseteq a_j$ for all j; and whenever $b_j \ne a_j$ we can reduce $L(v_j) \leftarrow b_j$.

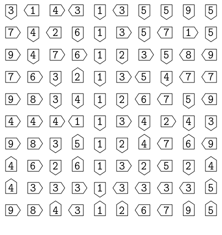

Operate in rounds, where all vertices are refined in round 1; subsequent rounds refine only the vertices whose parameters a_j have changed. In each round we first refine the vertices with smallest product $(\nu a_1) \ldots (\nu a_n)$, because they have the fewest potential solutions $x_1 \ldots x_n$. This method isn't guaranteed to succeed; but fortunately it does solve the stated problem, after 301 refinements in 6 rounds. [Such "Japanese arrow puzzles" were introduced by Masanori Natsuhara on page 75 of *Puzuraa* **128** (July 1992).]

69. (The 33rd boxed clue will, of course, have to point *outside* the 10×10 array. Maybe there's even a puzzle whose empty boxes are symmetrical, as in exercise 68.)

70. An extremely instructive analysis [*Combinatorics, Probability and Computing* **23** (2014), 725–748] leads to the recurrences $P_m = (5+9z)P_{m-2} - 4P_{m-4}$, $Q_m = (5+9z)Q_{m-2} - 4Q_{m-4}$, for $m \geq 6$, where the initial values are $(P_2, P_3, P_4, P_5) = (1, 1+z, 1+3z, 1+10z+9z^2)$; $(Q_2, Q_3, Q_4, Q_5) = (1-4z, 1-9z-6z^2, 1-19z-18z^2, 1-36z-99z^2-54z^3)$. The denominator $Q_m(z)$ has all real roots, exactly one of which is positive, namely $1/\rho_m$.

71. Suppose there are n questions, whose answers each lie in a given set S. A *student* supplies an answer list $\alpha = a_1 \ldots a_n$, with each $a_j \in S$; a *grader* supplies a Boolean vector $\beta = x_1 \ldots x_n$. There is a Boolean function $f_{js}(\alpha, \beta)$ for each $j \in \{1, \ldots, n\}$ and each $s \in S$. A graded answer list (α, β) is *valid* if and only if $F(\alpha, \beta)$ is true, where

$$F(\alpha,\beta) = F(a_1 \ldots a_n, x_1 \ldots x_n) = \bigwedge_{j=1}^{n} \bigwedge_{s \in S} \big([a_j = s] \;\Rightarrow\; x_j \equiv f_{js}(\alpha,\beta)\big).$$

The *maximum score* is the largest value of $x_1 + \cdots + x_n$ over all graded answer lists (α, β) that are valid. A *perfect score* is achieved if and only if $F(\alpha, 1 \ldots 1)$ holds.

Thus, in the warmup problem we have $n = 2$, $S = \{\mathrm{A}, \mathrm{B}\}$; $f_{1\mathrm{A}} = [a_2 = \mathrm{B}]$; $f_{1\mathrm{B}} = [a_1 = \mathrm{A}]$; $f_{2\mathrm{A}} = x_1$; $f_{2\mathrm{B}} = \bar{x}_2 \oplus [a_1 = \mathrm{A}]$. The four possible answer lists are:

$$\begin{array}{ll}
\mathrm{AA:} & F = \big(x_1 \equiv [\mathrm{A}=\mathrm{B}]\big) \wedge \big(x_2 \equiv x_1\big) \\
\mathrm{AB:} & F = \big(x_1 \equiv [\mathrm{B}=\mathrm{B}]\big) \wedge \big(x_2 \equiv \bar{x}_2 \oplus [\mathrm{A}=\mathrm{A}]\big) \\
\mathrm{BA:} & F = \big(x_1 \equiv [\mathrm{B}=\mathrm{A}]\big) \wedge \big(x_2 \equiv x_1\big) \\
\mathrm{BB:} & F = \big(x_1 \equiv [\mathrm{B}=\mathrm{A}]\big) \wedge \big(x_2 \equiv \bar{x}_2 \oplus [\mathrm{B}=\mathrm{A}]\big)
\end{array}$$

Thus AA and BA must be graded 00; AB can be graded either 10 or 11; and BB has no valid grading. Only AB can achieve the maximum score, 2; but 2 isn't guaranteed.

In Table 666 we have, for example, $f_{1\mathrm{C}} = [a_2 \neq \mathrm{A}] \wedge [a_3 = \mathrm{A}]$; $f_{4\mathrm{D}} = [a_1 = \mathrm{D}] \wedge [a_{15} = \mathrm{D}]$; $f_{12\mathrm{A}} = [\Sigma_\mathrm{A} - 1 = \Sigma_\mathrm{B}]$, where $\Sigma_s = \sum_{1 \leq j \leq 20} [a_j = s]$. It's amusing to note that $f_{14\mathrm{E}} = [\{\Sigma_\mathrm{A}, \ldots, \Sigma_\mathrm{E}\} = \{2,3,4,5,6\}]$.

The other cases are similar (although often more complicated) Boolean functions — except for 20D and 20E, which are discussed further in exercise 72.

Notice that an answer list that contains both 10E and 17E must be discarded: It can't be graded, because 10E says '$x_{10} \equiv \bar{x}_{17}$' while 17E says '$x_{17} \equiv x_{10}$'.

By suitable backtrack programming, we can prove first that no perfect score is possible. Indeed, if we consider the answers in the order (3, 15, 20, 19, 2, 1, 17, 10, 5, 4, 16, 11, 13, 14, 7, 18, 6, 8, 12, 9), many cases can quickly be ruled out. For example, suppose $a_3 = \mathrm{C}$. Then we must have $a_1 \neq a_2 \neq \cdots \neq a_{16} \neq a_{17} = a_{18} \neq a_{19} \neq a_{20}$, and early cutoffs are often possible. (We might reach a node where the remaining choices for answers 5, 6, 7, 8, 9 are respectively $\{\mathrm{C},\mathrm{D}\}$, $\{\mathrm{A},\mathrm{C}\}$, $\{\mathrm{B},\mathrm{D}\}$, $\{\mathrm{A},\mathrm{B},\mathrm{E}\}$, $\{\mathrm{B},\mathrm{C},\mathrm{D}\}$, say. Then if answer 8 is forced to be B, answer 7 can only be D; hence answer 6 is also forced to be A. Also answer 9 can no longer be B.) An instructive little propagation algorithm will make such deductions nicely at every node of the search tree. On the other hand, difficult questions like 7, 8, 9, are best *not* handled with complicated mechanisms; it's better just to wait until all twenty answers have been tentatively selected, and to check such hard cases only when the checking is easy and fast. In this way the author's program showed the impossibility of a perfect score by exploring just 52859 nodes, after only 3.4 megamems of computation.

student
grader
valid
dynamic ordering
propagation algorithm
search tree
author

recursively
quantified
fixed point
indeterminate

The next task was to try for score 19 by asserting that only x_j is false. This turned out to be impossible for $1 \le j \le 18$, based on very little computation whatsoever (especially, of course, when $j = 6$). The hardest case, $j = 15$, needed just 56 nodes and fewer than 5 kilomems. But then, ta da, three solutions were found: One for $j = 19$ (185 kilonodes, 11 megamems) and two for $j = 20$ (131 kilonodes, 8 megamems), namely

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	
D	C	E	A	B	E	B	C	E	A	B	E	A	E	D	B	D	A	b	B	(i)
A	E	D	C	A	B	C	D	C	A	C	E	D	B	C	A	D	A	A	c	(ii)
D	C	E	A	B	A	D	C	D	A	E	D	A	E	D	B	D	B	E	e	(iii)

(The incorrect answers are shown here as lowercase letters. The first two solutions establish the truth of 20B and the falsity of 20E.)

72. Now there's only *one* list of answers with score ≥ 19, namely (iii). But that is paradoxical — because it claims 20E is false; hence the maximum score *cannot* be 19!

Paradoxical situations are indeed possible when the global function F of answer 71 is used recursively within one or more of the local functions f_{js}. Let's explore a bit of recursive territory by considering the following two-question, two-letter example:

1. (A) Answer 1 is incorrect. (B) Answer 2 is incorrect.
2. (A) Some answers can't be graded consistently. (B) No answers achieve a perfect score.

Here we have $f_{1\mathrm{A}} = \bar{x}_1$; $f_{1\mathrm{B}} = \bar{x}_2$; $f_{2\mathrm{A}} = \exists a_1 \exists a_2 \forall x_1 \forall x_2 \neg F(a_1a_2, x_1x_2)$; $f_{2\mathrm{B}} = \forall a_1 \forall a_2 \neg F(a_1a_2, 11)$. (Formulas quantified by $\exists a$ or $\forall a$ expand into $|S|$ terms, while $\exists x$ or $\forall x$ expand into two; for example, $\exists a \forall x\, g(a,x) = \bigl(g(\mathrm{A},0) \wedge g(\mathrm{A},1)\bigr) \vee \bigl(g(\mathrm{B},0) \wedge g(\mathrm{B},1)\bigr)$ when $S = \{\mathrm{A},\mathrm{B}\}$.) Sometimes the expansion is undefined, because it has more than one "fixed point"; but in this case there's no problem because $f_{2\mathrm{A}}$ is true: Answer AA can't be graded, since 1A implies $x_1 \equiv \bar{x}_1$. Also $f_{2\mathrm{B}}$ is true, because both BA and BB imply $x_1 \equiv \bar{x}_2$. Thus we get the maximum score 1 with either BA or BB and grades 01.

On the other hand the simple one-question, one-letter questionnaire '**1.** (A) The maximum score is 1' has an *indeterminate* maximum score. For in this case $f_{1\mathrm{A}} = F(\mathrm{A}, 1)$. We find that if $F(\mathrm{A}, 1) = 0$, only $(\mathrm{A}, 0)$ is a valid grading, so the only possible score is 0; similarly, if $F(\mathrm{A}, 1) = 1$, the only possible score is 1.

OK, suppose that the maximum score for the modified Table 666 is m. We know that $m < 19$; hence (iii) isn't a valid grading. It follows that 20E is true, which means that *every valid graded list of score m has x_{20} false.* And we can conclude that $m = 18$, because of the following two solutions (which are the only possibilities with 20C false):

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
B	A	d	A	B	E	D	C	D	A	E	D	A	E	D	E	D	B	E	c
A	E	D	C	A	B	C	D	C	A	C	E	D	B	a	C	D	A	A	c

But wait: If $m = 18$, we can score 18 with 20A true and two errors, using (say)

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
D	e	D	A	B	E	D	e	C	A	E	D	A	E	D	B	D	C	C	A

or 47 other answer lists. This *contradicts* $m = 18$, because x_{20} is true.

End of story? No. This argument has implicitly been predicated on the assumption that 20D is false. What if m is indeterminate? Then a new solution arises

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
D	C	E	A	B	E	D	C	E	A	E	B	A	E	D	B	D	A	d	D

of score 19. With (iii) it yields $m = 19$! If m is determinate, we've shown that m cannot actually be defined consistently; but if m is indeterminate, it's definitely 19.

Question 20 was designed to create difficulties. [:-)]
— DONALD R. WOODS (2001)

73. The 29 words `spark`, `often`, `lucky`, `other`, `month`, `ought`, `names`, `water`, `games`, `offer`, `lying`, `opens`, `magic`, `brick`, `lamps`, `empty`, `organ`, `noise`, `after`, `raise`, `drink`, `draft`, `backs`, `among`, `under`, `match`, `earth`, `roots`, `topic` yield this: "The success or failure of backtrack often depends on the skill and ingenuity of the programmer. ... Backtrack programming (as many other types of programming) is somewhat of an art." — Solomon W. Golomb, Leonard D. Baumert.

That solution can be found interactively, using inspired guesses based on a knowledge of English and its common two-letter and three-letter words. But could a computer that knows common English words discover it without understanding their meanings?

We can formulate that question as follows: Let $w_1, \ldots, w_{29}$ be the unknown words from WORDS(1000), and let $q_1, \ldots, q_{29}$ be the unknown words of the quotation. (By coincidence there happen to be just 29 of each.) We can restrict the q's to words that appear, say, 32 times or more in the British National Corpus. That gives respectively (85, 562, 1863, 3199, 4650, 5631, 5417, 4724, 3657, 2448) choices for words of (2, 3, ..., 11) letters; in particular, we allow 3199 possibilities for the five-letter words q_7, q_{11}, q_{21}, q_{22}, because they aren't required to lie in WORDS(1000). Is there a unique combination of words w_i and q_j that meets the given anacrostic constraints?

This is a challenging problem, whose answer turns out (surprisingly?) to be *no*. In fact, here is the first solution found by the author's machine(!): "The success or failure of backtrack often depends on roe skill and ingenuity at the programmer. ... Backtrack programming (as lacy offal types of programming) as somewhat al an art." (The OSPD4 includes 'al' as the name of the Indian mulberry tree; the BNC has 'al' 3515 times, mostly in unsuitable contexts, but that corpus is a blunt instrument.) Altogether 720 solutions satisfy the stated constraints; they differ from the "truth" only in words of at most five letters.

Anacrostic puzzles, which are also known by other names such as double-crostics, were invented in 1933 by E. S. Kingsley. See E. S. Spiegelthal, *Proceedings of the Eastern Joint Computer Conference* **18** (1960), 39–56, for an interesting early attempt to solve them — *without* backtracking — on an IBM 704 computer.

74. Instead of considering 1000 possibilities for $\underset{131}{\underline{\ \ \ }}\underset{132}{\underline{\ \ \ }}\underset{133}{\underline{\ \ \ }}\underset{134}{\underline{\ \ \ }}\underset{135}{\underline{\ \ \ }}$, it suffices to consider the 43 pairs xy such that $cxyab$ is in WORDS(1000) and abc is a common three-letter word. (Of these pairs `ab`, `ag`, ..., `ve`, only `ar` leads to a solution. And indeed, the 720 solutions factor into three sets of 240, corresponding to choosing `earth`, `harsh`, or `large` as the keyword for $\underset{131}{\underline{\ \ \ }}\underset{132}{\underline{\ \ \ }}\underset{133}{\underline{\ \ \ }}\underset{134}{\underline{\ \ \ }}\underset{135}{\underline{\ \ \ }}$.) Similar reductions, but not so dramatic, occur with respect to $\underset{137}{\underline{\ \ \ }}\underset{139}{\underline{\ \ \ }}$, $\underset{118}{\underline{\ \ \ }}\underset{119}{\underline{\ \ \ }}$, $\underset{46}{\underline{\ \ \ }}\underset{48}{\underline{\ \ \ }}$, and $\underset{32}{\underline{\ \ \ }}\underset{35}{\underline{\ \ \ }}$.

75. The following algorithm uses an integer utility field TAG(u) in the representation of each vertex u, representing the number of times u has been "tagged." The operations "tag u" and "untag u" stand respectively for TAG(u) $\leftarrow$ TAG(u) $+ 1$ and TAG(u) $\leftarrow$ TAG(u) -1. Vertices shown as '$\odot$' in the 21 examples have a nonzero TAG field, indicating that the algorithm has decided *not* to include them in this particular H.

State variables v_l (a vertex), i_l (an index), and a_l (an arc) are used at level l for $0 \le l < n$. We assume that $n > 1$.

WOODS
Golomb
Baumert
two-letter
three-letter
British National Corpus
constraints
author
OSPD4
double-crostics
Kingsley
Spiegelthal
IBM 704
utility field
tagged

R1. [Initialize.] Set $\mathtt{TAG}(u) \leftarrow 0$ for all vertices u. Then set $v_0 \leftarrow v$, $i \leftarrow i_0 \leftarrow 0$, $a \leftarrow a_0 \leftarrow \mathtt{ARCS}(v)$, $\mathtt{TAG}(v) \leftarrow 1$, $l \leftarrow 1$, and go to R4.

R2. [Enter level l.] (At this point $i = i_{l-1}$, $v = v_i$, and $a = a_{l-1}$ is an arc from v to v_{l-1}.) If $l = n$, visit the solution $v_0v_1 \ldots v_{n-1}$ and set $l \leftarrow n-1$.

R3. [Advance a.] Set $a \leftarrow \mathtt{NEXT}(a)$, the next neighbor of v.

R4. [Done with level?] If $a \neq \Lambda$, go to R5. Otherwise if $i = l-1$, go to R6. Otherwise set $i \leftarrow i+1$, $v \leftarrow v_i$, $a \leftarrow \mathtt{ARCS}(v)$.

R5. [Try a.] Set $u \leftarrow \mathtt{TIP}(a)$ and tag u. If $\mathtt{TAG}(u) > 1$, return to R3. Otherwise set $i_l \leftarrow i$, $a_l \leftarrow a$, $v_l \leftarrow u$, $l \leftarrow l+1$, and go to R2.

R6. [Backtrack.] Set $l \leftarrow l-1$, and stop if $l = 0$. Otherwise set $i \leftarrow i_l$. Untag all neighbors of v_k, for $l \geq k > i$. Then set $a \leftarrow \mathtt{NEXT}(a_l)$; while $a \neq \Lambda$, untag $\mathtt{TIP}(a)$ and set $a \leftarrow \mathtt{NEXT}(a)$. Finally set $a \leftarrow a_l$ and return to R3. ▌

Λ the null pointer
backtracking, variant structure
n-omino placement
monotone
factored
generating functions
convolution

This instructive algorithm differs subtly from the conventional structure of Algorithm B. Notice in particular that $\mathtt{TIP}(a_l)$ is not untagged in step R6; that vertex won't be untagged and chosen again until some previous decision has been reconsidered.

76. Let G have N vertices. For $1 \leq k \leq N$, perform Algorithm R on the kth vertex v of G, except that step R1 should tag the first $k-1$ vertices so that they are excluded. (A tricky shortcut can be used: If we untag all neighbors of $v = v_0$ after Algorithm R stops, the net effect will be to tag only v.)

The n-omino placement counts 1, 4, 22, 113, 571, 2816, 13616, 64678, 302574 are computed almost instantly, for small n. (Larger n are discussed in Section 7.2.3.)

77. (a) All but the 13th and 18th, which require an upward or leftward step.

(b) True. If $u \in H$ and $u \neq v$, let p_u be any node of H that's one step closer to v.

(c) Again true: The oriented spanning trees are also ordinary spanning trees.

(d) The same algorithm works, except that step R4 must return to itself after setting $a \leftarrow \mathtt{ARCS}(v)$. (We can no longer be sure that $\mathtt{ARCS}(v) \neq \Lambda$.)

78. Extend Algorithm R to terminate immediately if $\mathtt{WT}(v) \geq U$, otherwise to visit the singleton solution v. Also set $w \leftarrow \mathtt{WT}(v)$ in step R1. Replace steps R2 and R5 by

R2′. [Enter level l]. If $p \geq L$, visit the solution $v_0v_1 \ldots v_{l-1}$.

R5′. [Try a.] Set $u \leftarrow \mathtt{TIP}(a)$ and tag u. If $\mathtt{TAG}(u) > 1$ or $w + \mathtt{WT}(u) \geq U$, return to R3. Otherwise set $i_l \leftarrow i$, $a_l \leftarrow a$, $v_l \leftarrow u$, $w \leftarrow w + \mathtt{WT}(u)$, $l \leftarrow l+1$, and go to R2.

In step R6, set $w \leftarrow w - \mathtt{WT}(v_l)$ just before setting $i \leftarrow i_l$.

79. (a) $(0,j)$ and $(1,j)$ for $j \geq 44$; $(2,j)$ for $j \geq 32$; $(4,j)$, $(8,j)$, $(10,j)$ for $j < 12$.

(b) True, each of the Boolean functions $r_{i,j}$ is clearly monotone.

(c) The "couplers" can be simulated by playing s_j^* and g_j^* instead of s_j and g_j (as if the organist had assistants). Therefore the problem can be *factored* into independent subproblems for the Pedal, Swell, and Great separately: Let there be P_n, S_n, G_n playable sounds on the Pedal, Swell, and Great, and define $P(z) = \sum_n P_n z^n$, $S(z) = \sum_n S_n z^n$, $G(z) = \sum_n G_n z^n$; then $Q(z) = \sum_n Q_n z^n$ is the convolution $P(z)S(z)G(z)$.

(d) $p_0 = p_{12} = c_0 = c_1 = c_{15} = 1$ gives $(0,0)$, $(0,12)$, $(0,24)$, $(1,0)$, $(1,12)$; $s_0 = s_{19} = s_{28} = c_3 = c_4 = 1$ gives (the beautiful) $(3,0)$, $(3,19)$, $(3,28)$, $(4,19)$, $(4,28)$; etc.

(e) It's unplayable if and only if $i \in \{2,14,15\}$ or $i' \in \{0,1,2,14,15\}$ or ($i \neq i'$ and either $3 \leq i,i' \leq 8$ or $9 \leq i,i' \leq 15$).

(f) $Q_1 = 812 - 112 = 700$, because we can't have $(14,j)$ or $(15,j)$ without $(13,j)$.

generating functions

(g) $Q_{811} = 12$ sounds lack only one pipe: With all inputs 1 except p_j, for $12 \le j < 24$, only $r_{2,j}$ is 0. (Thankfully there isn't enough wind pressure to actually play this.)

(h) Brute-force backtrack programs can be written, using the monotonicity property (b) for cutoffs, in order to check small values and to list the actual sounds. But the best way to compute P_n, S_n, G_n, and Q_n is to use generating functions.

For example, let $G(z) = G_0(z) + G_1(z) + \cdots + G_{63}(z)$, where $G_k(z)$ for $k = (c_{14}c_{13}c_{12}c_{11}c_{10}c_9)_2$ enumerates the sounds for a given setting of console switches, excluding sounds already enumerated by $G_j(z)$ for $j < k$. Then $G_0(z) = 1$; $G_k(z) = 0$ if $c_{13}c_{14} = 1$; otherwise $G_k(z) = f(c_9 + c_{11} + c_{12} + c_{13} + 3c_{14})$ when $c_{10} = 0$, and $G_k(z) = g(c_9 + 1 + c_{11} + c_{12} + c_{13} + 3c_{14}, 1 + c_{11} + c_{12} + c_{13} + 3c_{14})$ when $c_{10} = 1$, where

$$f(n) = (1 + z^n)^{56} - 1, \qquad g(m, n) = (1 + z^n)^{12}\bigl((1 + z^m)^{44} - 1\bigr).$$

Thus $G(z) = 1 + 268z + 8146z^2 + 139452z^3 + \cdots + 178087336020z^{10} + \cdots + 12z^{374} + z^{380}$.

Similarly, with $S(z) = \sum_{k=0}^{63} S_k(z)$ and $k = (c_8c_7c_6c_5c_4c_3)_2$, we have $S_0(z) = 1$; $S_{32}(z) = (1 + z)^{44} - 1$; otherwise $S_k(z) = f(c_3 + c_5 + c_6 + c_7)$ when $c_4 = c_8 = 0$, $S_k(z) = g\bigl(c_3 + c_4 + c_5 + c_6 + c_7 + c_8, \max(c_3, c_4) + c_5 + c_6 + c_7\bigr)$ when $c_4 + c_8 > 0$. Thus $S(z) = 1 + 312z + 9312z^2 + 155720z^3 + \cdots + 180657383126z^{10} + \cdots + 12z^{308} + z^{312}$. [Curiously we have $S_n > G_n$ for $1 \le n \le 107$.]

The generating functions for $P(z) = \sum_{k=0}^{31} P_k(z)$, with $k = (c_{16}c_{15}c_2c_1c_0)_2$, are trickier. Let $h(w, z) = (1 + 3wz^2 + 2w^2z^3 + z^2z^4 + w^3z^4)^8\bigl((1 + 2wz^2 + w^2z^3)^4 - 1\bigr)$. Then $P_{31}(z) = h(z, z^2)$, and there are three main cases when $0 < k < 31$: If $c_0c_{15} = c_1c_{16} = 0$, then $P_k(z) = (1 + z^{c_{15}+c_{16}})^{32} - (1 + z^{c_{15}+c_{16}})^{20}$ if $c_0 + c_1 + c_2 = 0$, otherwise $P_k(z) = (1 + z^{c_0+c_1+c_2+c_{15}+c_{16}})^{32} - 1$. If $c_0 = c_{15}$, $c_1 = c_{16}$, $c_2 = 0$, then $P_k(z) = q(z^{c_0+c_1})$,

$$q(z) = (1 + 3z^2 + 2z^3 + z^4)^8(1 + 2z^2 + z^3)^4 - 2(1 + 2z^2 + z^3)^8(1 + z^2)^4 + (1 + z^2)^8.$$

Otherwise we have $P_k(z) = h(z^{c_0+c_1+c_2+c_{15}+c_{16}-2}, z)$. Thus $P(z) = 1 + 120z + 2336z^2 + 22848z^3 + \cdots + 324113168z^{10} + \cdots + 8z^{119} + z^{120}$. And $Q(z) = 1 + 700z + 173010z^2 + 18838948z^3 + 1054376915z^4 + 38386611728z^5 + 1039287557076z^6 + 22560539157160z^7 + 410723052356833z^8 + 6457608682396156z^9 + 89490036797524716z^{10} + \cdots + 12z^{811} + z^{812}$. So $(Q_2/\binom{812}{2}, \ldots, Q_{10}/\binom{812}{10}) \approx (.5, .2, .06, .01, .003, .0005, .00009, .00002, .000003)$.

INDEX AND GLOSSARY

He writes indexes to perfection.
— OLIVER GOLDSMITH, *Citizen of the World* (1762)

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.