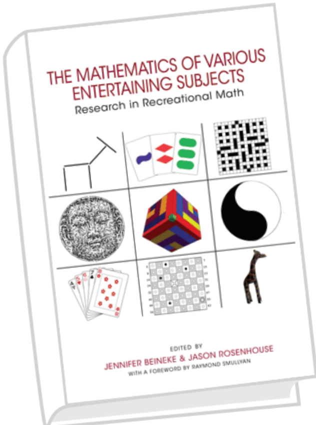

The Mathematics of Various Entertaining Subjects

A Review by Brian Hayes

The Mathematics of Various Entertaining Subjects: Research in Recreational Math
Jennifer Beineke and Jason Rosenhouse, Editors
Foreword by Raymond Smullyan
Princeton University Press, 2016
US$75.00, 272 pages
ISBN: 978-0-691-16403-8

Recreational mathematics is a gateway drug. Give the novices a taste of logic puzzles, magic tricks, and perplexing geometric patterns, then, once they're hooked, move them along to the harder stuff—complex analysis, algebraic geometry, interuniversal Teichmüller theory. As a recruiting tool this strategy has been highly successful. Many a mathematician credits Martin Gardner and his "Mathematical Games" column for inspiring a lifetime of mathematical inquiry.

puzzles and problems meant to engage the attention of research mathematicians

The Mathematics of Various Entertaining Subjects (*MOVES*) is recreational math for a different audience. The puzzles and problems may be the same, but in this presentation they are meant to engage the attention of research mathematicians. But that's certainly not to say there's no fun here. The volume consists of selected papers from a symposium of the same name held in 2013 at the National Museum of Mathematics in New York.

Of course recreational problems have been inspiring "real" math for a long time. Jennifer Beineke and Jason Rosenhouse, the editors of this volume, point out in their preface that probability theory began with games of chance and that graph theory has roots in puzzles such as Euler's tour of the Königsberg bridges and Hamilton's Icosian game. Going farther back, "Among the oldest mathematical documents to have survived to the present is the Egyptian Rhind papyrus, which is largely a collection of ancient brainteasers. The isoperimetric problem is discussed by Virgil in the *Aeneid*."

The book's seventeen chapters survey quite a broad spectrum of problems and puzzles. Robert Bosch, Tim Chartier, and Michael Rowan describe schemes for converting any grayscale image into a maze by solving an optimization problem, such as the traveling salesman problem. Classic coin-weighing puzzles ("Find the counterfeit in no more than k weighings") are revisited and updated by Tanya Khovanova, who in recent years has shown that this thoroughly mined genre still has a multitude of subtle variations. Julie Beier and Carolyn Yackel find fresh results strewn along another well-trodden path—the study of flexagons, which were the subject of Martin Gardner's first proto-column in *Scientific American* sixty years ago. Maureen T. Carroll and Steven T. Dougherty try playing tic-tac-toe on a finite affine plane, where a finite set of points define the positions of Xs and Os, and a finite set of lines (not necessarily straight) mark the potential winning paths.

A review can't do justice to all these varied topics and the mathematics that lies behind them. I am therefore going to focus on just three selected chapters.

Random Towers of Hanoi

Max A. Alekseyev and Toby Berger consider a variant of the Towers of Hanoi puzzle (introduced in 1883 by Édouard Lucas, a significant figure in both serious and recreational math). The original problem is challenging enough: 64 perforated disks are stacked on a peg, with the largest on the bottom and the diameter diminishing step by step up to the top. Generations of monks labor to transfer the entire heap to another peg, moving one disk at a time. They use a third peg as a temporary holding

Brian Hayes is contributing editor at American Scientist *magazine, associate in computer science at Harvard, and Y Combinator Research Fellow. His e-mail address is* brian@bit-player.org.

For permission to reprint this article, please contact: reprint-permission@ams.org.
DOI: http://dx.doi.org/10.1090/noti1478

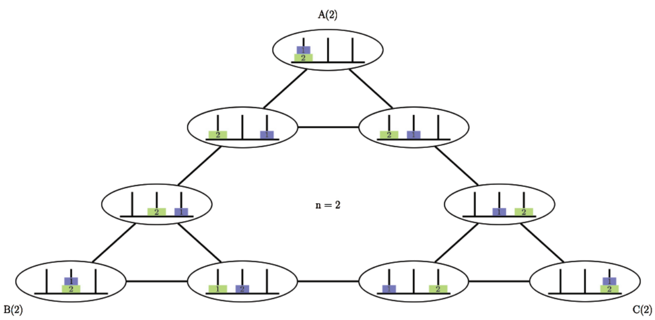

Figure 1. The state-transition diagram for a two-disk Towers of Hanoi problem resembles the Sierpiński gasket. Each vertex of the graph defines a state of the system: the disposition of the disks on the three available pegs. An edge connects two vertices if a legal move connects the corresponding states. Max A. Alekseyev and Toby Berger study the solution of the problem by taking random walks on the graph.

area, but at no stage in the process do they ever place a larger disk above a smaller one on any of the pegs. When the task is complete, according to Lucas, the world ends. The minimum number of moves to complete the transfer is $2^{64} - 1$, or about 10^{19}. Alekseyev and Berger ask: Suppose that instead of applying the most efficient algorithm, the monks shuffle the disks at random—though always choosing only legal moves. What then is the expected length of the solution?

It turns out the graph of the Towers problem is a slight modification of the Sierpiński gasket, a fractal mesh of triangles within triangles. Each node of this graph represents a state of the puzzle—a list of which disks occupy which pegs—and each edge of the graph represents a legal move shifting one disk between pegs. The nodes at the three corners of the gasket represent end states where all the disks are on a single peg. Hence the most efficient strategy is to follow the sequence of states along a side of the gasket or, in other words, the shortest path between two corners.

The random-move strategy corresponds to a random walk on the graph. If the walk begins at the corner node with all disks on peg 1, is it guaranteed to someday reach the corner where all disks are on peg 3? Yes, with probability 1. But how long does it take on average? Alekseyev and Berger prove that for a tower of n rings, the answer is not $2^n - 1$ but rather:

$$E_{1\to 3}(n) = \frac{(3^n - 1)(5^n - 3^n)}{2 \cdot 3^{n-1}}.$$

In the case of $n = 64$, this works out to about $5^{64} \sim 10^{45}$ moves. Should we be surprised by how large this number is or by how small it is? Alekseyev and Berger remark:

> [R]eplacing the minimum-moves strategy with a random walk forestalls the end of the world by a factor of roughly $(\frac{5}{2})^{64} > 2.9 \times 10^{25}$ on average. Although this is reassuring, it would provide further comfort to know that the coefficient of variation of the random number of steps…is small (i.e., that its standard deviation is many times smaller than its mean). Exact determination of said coefficient of variation is an open problem that we may address in future research.

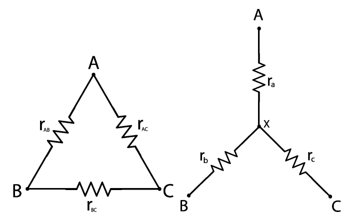

Figure 2. The delta-wye transform, a concept borrowed from electrical engineering, can simplify the analysis of the random Towers of Hanoi problem. Suppose each edge in the Sierpiński-like graph is replaced by a 1 ohm electrical resistor. When a voltage is applied across two vertices such as A and C, the current through each resistor gives the frequency with which a random walker would traverse the corresponding edge of the graph. In a complex graph of many nested triangles, calculating those currents is challenging; it becomes easier when the graph is replaced by an equivalent Y-shaped network. If the resistors in the delta network are each 1, those in the Y-network are each 1/3. In either case the total resistance from A to C is 2/3.

an adventurous transposition of the problem into quite a different realm

Alekseyev and Berger give two proofs of the mean random walk length. One proof is a conventional, recursive calculation of probabilities. The other is an adventurous transposition of the problem into quite a different realm. They model the graph of the Sierpiński gasket as a network of electrical resistors, an idea imported into mathematics from electrical engineering by Peter G. Doyle and J. Laurie Snell in the 1980s. Suppose every edge of the Sierpiński graph is a 1 ohm resistor and you apply a voltage to two corners of the graph, node 1 and node 3. If you then calculate or measure the current flowing through each resistor, the result indicates the relative frequency with which a random walker will pass along that edge of the graph. In itself, this strategy is no improvement over the more straightforward calculus of probabilities. But a rearrangement of the graph known to electrical engineers as the delta-wye transformation greatly simplifies the problem. As the name suggests, the transformation changes a delta, or triangle Δ, into a Y-shaped network. When the delta-wye transformation is applied recursively, the entire network is reduced to one big Y, and calculating the resistance between two corners is a simple summation.

Heartless Poker

In poker, hands are valued in inverse order of expected frequency. Roughly 42 percent of all possible 5-card

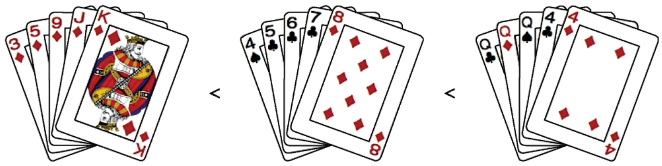

Figure 3. “Heartless poker,” played with a 39-card deck that has just three suits, alters the ranking of certain hands. Whereas in normal poker a flush (5 cards of one suit) beats a straight (5 cards of consecutive ranks), in the heartless variant this ordering is inverted. Dominic Lanphier and Laura Taalman show that every possible ordering of the flush, the straight, and the full house (three of a kind plus a pair) can be achieved by some variation on the standard card deck.

hands feature a single pair (i.e., 2 cards of the same rank and nothing else of value). Fewer than 5 percent have two pairs, and about 2 percent have three of a kind. Based on this ranking, even novice players quickly learn that three of a kind beats two pair, which in turn beats one pair. The ordering of some higher-ranking hands is less obvious. For one thing, those hands are so rare that most players have little chance to develop much intuition about them. Dominic Lanphier and Laura Taalman point out another reason for occasional confusion about the ranking of the straight, the flush, and the full house: their frequencies are not dramatically different. A straight (5 cards in sequence) is less than twice as common as a flush (5 cards of the same suit). A flush is only 1.4 times as common as a full house (three of a kind plus a pair).

Remove all the hearts from the deck for a game of “heartless poker.”

Lanphier and Taalman ask whether changes to the game—and in particular changes to the composition of the deck of cards—might alter the rankings of certain hands. A standard 52-card deck has four suits and thirteen ranks. Suppose we remove all the hearts from the deck and play a game of “heartless poker” with just 39 cards. Eliminating a suit makes it much easier to draw a flush, enough so that a straight becomes rarer and more valuable. Going in the other direction, a “fat pack” of cards is double the size of a normal deck, with eight suits and the usual thirteen ranks. Playing with that 104-card deck makes a flush rarer and more valuable than a full house. Lanphier and Taalman go on to show that all six possible orderings of those three hands can be achieved with some deck that has between three and fifteen suits and between thirteen and thirty-three ranks.

Then comes the question that takes us slightly beyond the usual turf of recreational mathematics: Is there any deck for which two of these hands (or all three) have the same probability and thus should be valued equally? Taking the number of ranks r and the number of suits s to be continuous variables, they define three curves on the rs plane where the straight, the flush, and the full house are pairwise equal in frequency. Furthermore, they show that the three curves have a unique point of intersection. In other words, there exists a pair of r and s values for which the straight, flush, and full house are all equally probable. However, those values of r and s are not integers, so we can’t actually create a physical deck of cards with this property. Further Diophantine analysis shows that there are no integer numbers of ranks and suits where any two of the three hands are of equal value.

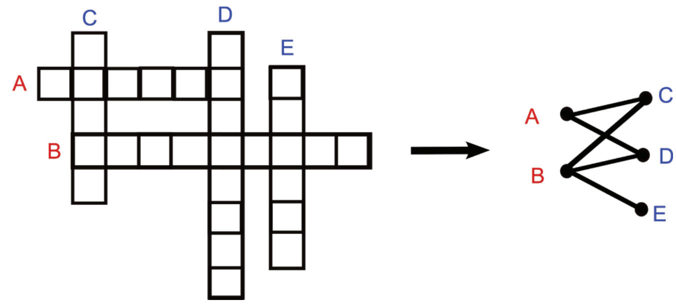

Figure 4. A crossword puzzle can be modeled as a bipartite graph. The words to be entered as answers are the vertices of the graph; two vertices are connected by an edge if the answers have a letter in common. The graph is bipartite because only Across and Down answers can intersect.

Critical Phenomena in Crossword Puzzles

Crossword puzzles are a pleasant diversion for many of us, but they seem to draw mainly on linguistic skills and general knowledge, without much mathematical content. John K. McSweeney demonstrates that even if mathematical analysis won’t help you solve the crossword in Sunday’s newspaper, it can illuminate the inner structure of the puzzle, help measure its difficulty, and explain why completing it is satisfying (or not). McSweeney writes:

> What distinguishes a crossword puzzle from a simple list of trivia questions is that the answers are entered into a grid in crossing fashion, and therefore each correct answer obtained provides partial information about others…Indeed, even if there are only a few easy answers that can be found immediately, these may trigger further answers, and, in such a cascading fashion, many or all of the answers in the puzzle may be found.

McSweeney represents the puzzle as a bipartite graph. The nodes of the graph are the answers to the puzzle clues. Two nodes are connected by an edge if the corresponding answers cross within the puzzle diagram and hence have a letter in common. The graph is bipartite because edges can only connect Across and Down answers; two Across answers or two Down answers never intersect. (McSweeney focuses mainly on American-style puzzles, and in particular on the Sunday *New York Times* crossword, where every letter entered into the grid forms part of both an Across and a Down answer; some British puzzles allow semiisolated letters.)

Knowing some of the letters in an answer should generally make it easier to fill in the rest. This fact leads to a nonlinearity in the puzzle-solving process that McSweeney investigates in detail. He models the solver's task as a purely probabilistic process. Each clue x is assigned a difficulty threshold $\varphi_x \in \mathcal{R}$. The answer to clue x is revealed only if the proportion of letters already known from cross-clues is at least φ_x. If $\varphi_x \le 0$, the answer is known immediately. If $\varphi_x = \frac{1}{2}$, the answer will become clear when half the letters have been revealed. When $\varphi_x \ge 1$, the answer is a stumper that you can't understand even when all the letters are filled in.

Between the easy and the impossible are puzzles exhibiting random fluctuations on a large scale.

In most of his experiments and analyses McSweeney assumes that the difficulty thresholds are drawn from a Gaussian distribution, with a mean μ greater than zero and a standard deviation σ wide enough that at least a few clues can be answered without help from orthogonal answers. The simplest and most symmetric case is a square grid without any black squares so that all Across and Down answers intersect one another. In this type of puzzle (almost never seen in the wild), there are parameter ranges for μ and σ where most puzzles are boringly easy or frustratingly hard, but there's also a more interesting region with a bimodal distribution: The solving algorithm may get stuck early and make very little progress, but if it completes a certain fraction of the grid, it will almost surely go on to solve the entire puzzle. I suspect this bimodal pattern has a lot to do with the pleasure of working crossword puzzles: The solver wants a challenge but also wants the satisfaction of completing the task.

McSweeney also runs simulations based on actual puzzles from the Sunday *Times*. In these grids, patterns of black squares break the puzzle into loosely coupled regions, so that only subsets of Across and Down answers intersect. Because of these barriers, it's not uncommon to solve most of a puzzle but be stymied in a few corners where neither Across nor Down clues yield up their secrets. The bimodal distribution of outcomes is still in evidence, but it takes a somewhat different form. There are easy puzzles and impossible puzzles, but those in between exhibit random fluctuations on a large scale, with big blocks of the puzzle solved but others left blank. The patterns resemble those seen in fluids or magnetic materials near a critical point. In crosswords, the critical parameter values seem to be near $\mu = 0.3, \sigma = 0.23$.

Conclusion

Although *MOVES* addresses an audience of mathematical sophisticates, almost all the chapters are readily accessible to students and amateurs. Where proofs are given, they are explained in detail. Some knowledge of group theory is helpful in the chapter on flexagons; graph theory and probability turn up in many contexts. In general, though, all that's needed to appreciate this work is a little facility in mathematical reasoning, and a dose of enthusiasm.

At a time when the public is once again debating the utility of mathematics—Do plumbers need to know algebra? Will calculus get you a better job?—it's a relief to open this window on this less fretful cosmos, where mathematics is a source of understanding and wonder.

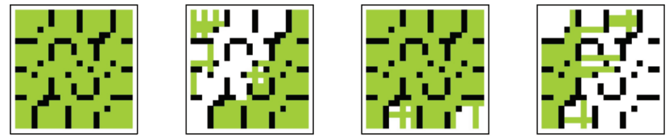

Figure 5. Critical behavior is observed in a random model of the solution process for *New York Times* Sunday crossword puzzles. Each clue is assigned a difficulty threshold drawn from a Gaussian distribution with mean μ and standard deviation σ. The answer to the clue is known whenever the proportion of letters revealed by known intersecting answers exceeds this threshold. In these diagrams green cells have been filled in successfully and white cells remain blank. At $\mu = 0.3$ and $\sigma = 0.23$, the solutions exhibit large-scale fluctuations, similar to those seen in physical systems near a critical point. In some instances the entire puzzle is solved; in others large areas are left blank.

Image and Photo Credits

Figures 1 and 2 are images reproduced from "Solving the Tower of Hanoi with Random Moves," by Max A. Alekseyev and Toby Berger, in *The Mathematics of Various Entertaining Subjects: Research in Recreational Math*, edited by Jennifer Beineke and Jason Rosenhouse, Princeton University Press, 2015.

Figure 3 is courtesy of Dominic Lanphier and Laura Taalman.

Figures 4 and 5 are courtesy of John K. McSweeney.

Photo of Brian Hayes is courtesy of Harvard University.

Brian Hayes

ABOUT THE REVIEWER

Brian Hayes's next book, *Foolproof, and Other Mathematical Meditations*, will be published next year by MIT Press.