

On Rolling Cube Puzzles

Kevin Buchin* Maike Buchin* Erik D. Demaine†
 Martin L. Demaine† Dania El-Khechen‡ Sándor P. Fekete§
 Christian Knauer* André Schulz* Perouz Taslakian¶

Abstract

We analyze the computational complexity of various rolling cube puzzles.

1 Introduction

Consider the simple rolling cube puzzle in Figure 1(a). The objective is to roll the die over all labeled cells of the board such that the label on the top face of the die is always the same as the label of the cell it lies on. The die may be rolled between neighboring cells by tipping it over along one edge that touches the board. The die may not be rotated within the same cell. In the example of Figure 1(a), it may roll on white cells multiple times but on labeled cells only once.

Rolling cube puzzles were popularized by Martin Gardner. In three of his *Mathematical Games* columns published in *Scientific American*, he discusses rolling cube puzzles. The first problem, called “Heavy Boxes”, was described in [5, 7] and originally posed by Roland Sprague in [17, 18]. It is shown in Figure 2(a). The puzzle poses the following questions. Is it possible to roll the five dice, each labeled “A” on exactly one face, from the upper left of the board to the lower right? Which die in this row started out as the center die of the cross formation?


*Institute of Computer Science, Free University Berlin, {buchin,mbuchin,knauer,schulza}@inf.fu-berlin.de

†Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, {edemaine, mdemaine}@mit.edu


‡Department of Computer Science and Software Engineering, Concordia University, d.elkhec@cs.concordia.ca

§Department of Computer Science, Braunschweig University of Technology, s.fekete@tu-bs.de

¶School of Computer Science, McGill University, perouz@cs.mcgill.ca



(a) by Joseph O’Rourke



(b) by Martin Demaine

Figure 1: Rolling cube puzzles posed at CCCG 2005 [4].




Figure 2: Rolling Cube Puzzles from Martin Gardner’s *Mathematical Games* column.

The next rolling cube problem, called “The red-faced cube”, was described in [6, 8] and is due to John Harris. It is shown in Figure 2(b). The die has one face that is colored red; all other faces are colored white. The board has size 8×8 . The puzzle poses the following questions. Can the die, starting in the upper-left corner and ending in the upper-right corner, roll over each cell of the board exactly once and never show the red face on top except for the first and last cell? Can it be rolled over the whole board visiting each cell exactly once and never show the red face on top?

Later on [9, 10], Martin Gardner discussed another problem by John Harris [12], the *single vacancy rolling cube problem*. This puzzle is shown in Figure 2(c). Here there are eight dice, each colored red on one face and black on the opposite face. Can they be rolled from showing all black faces to showing all red faces in the positions shown in the figure?

More recently, Robert Abbott has posed several rolling cube puzzles in books [1, 2, 3].¹ He also describes how to create rolling cube mazes [2]. Two easy examples of his puzzles are shown in Figure 3. We will illustrate some of the ideas presented in Section 2 using these mazes.

In Abbott’s two mazes, the task is to find a way from the *start* to the *goal* and they are therefore called *rolling cube mazes*. In the first maze, one starts with the die oriented as shown next to the maze with the 6 on top and the 4 front-facing. The die may be tipped onto a square if the number currently shown on top is the same as the number on the new cell. The die can be tipped onto squares with an asterisk no matter which number is facing up. The die can be tipped onto all cells multiple times. In the second maze, one starts with the 6 on top. The die can be tipped onto a white square with any number except 1. Again, the die can be tipped onto all cells multiple times. The die cannot be tipped onto the shaded squares.

Based on rolling cube mazes, Richard Tucker invented *rolling block mazes*¹ where one rolls shapes composed by several dice. Rolling block mazes are also described by van Deventer [20]. Trigg [19] considered puzzles with a tetrahedron, instead of a cube, rolling on a triangular grid.

In the open-problem session at CCCG 2005 [4], Joseph O’Rourke posed the computational complexity of rolling cube puzzles like the one in Figure 1(a). During the discussion, Martin Demaine developed the multiple-dice puzzle shown in Figure 1(b).

In general, a *rolling cube puzzle* consists of one or more dice, a board, a task, and a set of rules. A *die* is a cube with (some) labeled faces. We consider the case of a *standard die*, that is, a die with faces labeled 1 to 6 and with the labels on opposite sides adding up to 7. There are two standard dice. They can be distinguished by how the numbers 1, 2, and 3 are oriented with respect to each other, i.e., either *clockwise* or *counterclockwise*, and are also called either *right-handed* or *left-handed* dice. See Figure 4(a–b) for an illustration.

Here we use a right-handed orientation, which is more common and is also the orientation of the

¹See also Robert Abbott’s website, <http://www.logicmazes.com>.





Figure 3: Rolling Cube Mazes by Robert Abbott. Used with permission.




Figure 4: The two standard dice and the standard casino die.




Figure 5: Solids with suitable grids. From left to right: octahedron, icosahedron, cuboctahedron, truncated tetrahedron, dodecahedron.

standard casino die for which furthermore the orientation of the marks is fixed (see Figure 4(c)). In Section 5 we consider puzzles with a two-colored die where the task is to color as many cells as possible in one color (as in the “The red-faced cube” problem).

Instead of a cube, a platonic solid with triangular faces can be used. Charles W. Trigg [19] studied the case of rolling a tetrahedron. He shows that placing the tetrahedron on a cell of the board fixes the orientation of the tetrahedron on all other cells. Thus, puzzles using a tetrahedron are not so interesting. To the best of our knowledge, rolling an octahedron or an icosahedron over a triangular grid has not been studied. Also other solids and tilings might result in interesting puzzles, e.g., rolling a cuboctahedron or a truncated tetrahedron on suitable Archimedean tilings. Using a tiling with irregular pentagons, one might even consider rolling a dodecahedron. See Figure 5 for an illustration.

The *board* is a grid with (some) labeled cells. Given a board, the objective is to roll the die over the cells of the board to accomplish some task, such as to visit all the labeled cells; sometimes we may also be given a starting position of the die and an ending position. We consider the case where all labeled cells must be visited.

In Section 2 we show that puzzles are easy (for a computer, not for a human) if labeled cells may be visited several times. Thus we later concentrate on puzzles where labeled cells must be visited exactly once.

Cells can be of three types: *labeled*, *blocked*, or *free*. Labeled cells must be visited exactly once with the label appearing on the top face of the die being the same as the label of the cell. Note that this is different from Robert Abbott’s puzzles for which the label on top before rolling is relevant. Blocked cells cannot be visited by the die. Free cells can be visited with any label on the top face of the die and any number of times. We restrict ourselves to puzzles with one die, but puzzles can also involve several dice as in Figure 1(b).

In Section 3 we prove that it is NP-complete to decide whether we can roll a die over the labeled cells of a board that has some free cells. This solves the open problem posed by Joseph O’Rourke. Free cells seem to be essential for the hardness of the problem; thus in Section 4 we present an algorithmic approach to puzzles with no free cells, and show that the solution to a puzzle with labeled (and possibly blocked) cells is not necessarily unique. The computational complexity of such puzzles remains open.

2 Basic Properties

2.1 State Graph

The *state graph* has a vertex for each possible state of the die and an edge for each possible transition between two states. A *state* consists of a board position and the entire orientation of the die. In particular, the state encodes which label is on the top face, but even fixing this top label, there are four possible orientations, defined by the label facing a fixed side of the board.

We will denote these orientations as x^y where x is the label on the top of the die and y the label on the back face (i.e., the side facing the top of the page in this paper). For example, in the first maze in Figure 3, the die starts in the state 6^3 , because 3 is on the opposite side of the front 4.

An edge of the state graph connects vertices corresponding to adjacent cells on the board for which it is possible to roll a die from one cell to the other, respecting the orientations of the two states. (Moves are reversible, so the graph is undirected.)

2.2 Parity Property

An important property of the state graph is that its vertex set naturally falls into two parts of equal size with no edges between the two parts. For a labeled cell, only two of the four orientations need to be considered if we restrict the die to one part of the state graph. This corresponds to the *parity property* of rolling cube puzzles [5, 12, 17, 18] which gives the solution to the puzzle in Figure 2(a).

In the following we look at this property in a different way. We consider how a 2-coloring of the corners of the cube induces a 2-coloring of the corners of the board. We color the corners of the cube alternately black and white as in Figure 6(a). Rolling this cube over the board generates one of two 2-colorings of the corners of the board, i.e., a checkerboard pattern. One possible 2-coloring is shown in Figure 6(c); the other is its complement with black and white exchanged. Consider a cube with a given label on top. Each of the two possible 2-colorings of the board is generated by two of the four orientations of the cube. It is not possible to return to a cell with the cube rotated by 90° .

A different perspective on the coloring argument is the following. If only two opposite corners of the cube are colored, as shown in Figure 6(b), then the sequence of corners of the board touched by the colored corners corresponds to the way a chess knight moves. As on a checkerboard, the color changes with every move. A natural question concerning checkerboard is the following: can




Figure 6: A die and a board with 2-colored corners.




Figure 7: Rolling cube maze using the parity property.

the faces of a cube be colored in black and white in a way that rolling it over a board generates a chess board pattern? It is easy to see that this is not the case.

We designed the rolling cube maze in Figure 7 based on an idea of Robert Abbott² to connect the two parts of the state graph by a *bridge*: rolling along the top or bottom row changes the parity of the die. The rules are the same as for the first maze in Figure 3.

²See also <http://www.logicmazes.com>.

2.3 Local Structure

For the case of a labeled cell, the above argument reduces the number of possible states to two, assuming we know the part of the state graph. Thus, if we consider the state graph at a cell, it has the structure of one of the configurations shown in Figure 8 or a subgraph of one of these. The vertices represent the two possible states of a cell, and the edges represent the possible transitions to states of neighboring cells. Degree-1 edges are shown dotted because they can be ignored when searching for a Hamiltonian cycle.

Different configurations at a cell in the state graph yield different difficulties. The upper-left configuration allows for only one state at a cell but for several possibilities (three if a cell can be visited only once, four otherwise) to roll in this state. The three other configurations in the upper row allow for two states but only one possibility to roll (two if cells can be visited several times). The configurations in the lower row are not that interesting, because they allow only one state and two possibilities to roll if cells can only be visited once.




Figure 8: Possible configurations for the states corresponding to a labeled cell.

Situations where only one state is possible can be easily read off the labels of the neighboring cells. If we want to roll the die from a neighboring cell onto a cell, this determines the state at this cell. Again assuming that cells may only be visited once, a cell can be only visited in the state for which the *majority of the neighbors votes*, i.e., which is compatible with the neighboring cells. A tie is only possible in the two against two situations already seen in Figure 8 (upper row, except for the left most). This simple observation is frequently useful to cut down the number of possible states. For example, in a fully labeled board all boundary cells allow only for one state because there are at most three neighbors.

2.4 Changing Direction Twice

If the die changes direction twice, the label on its top face is either the same as before or the opposite, depending on whether the last turn was back or not. More precisely, suppose the die starts with label x and rolls one cell in one direction, then an arbitrary number of cells in an orthogonal direction, and then one cell in the original or opposite direction. Then the die shows the label x if the turn was a *U-turn*, and $7 - x$ if it was a *Z-turn*; see Figure 9.

Reconsider the second maze in Figure 3. It can be easily seen that the rollable path must be a sequence of U-turns, the last U-turn being possibly incomplete. The best solution (in terms of path length) is a set of eight and a half U-turns.




Figure 9: U-turn and Z-turn.

2.5 Rolling over Labeled Cells Multiple Times

Puzzles where labeled cells may be rolled on several times can be solved by checking the reachable part of the state graph, starting at some initial state (or a set of possible initial states). For instance, the puzzles in Figure 3 can be solved in this way. Finding a rolling path for the die corresponds to finding a sufficiently large connected component in the state graph. The reachable part can for instance be determined by a breadth-first search of the state graph.

Thus we obtain the following result:

Proposition 1 *When the labeled cells of a rolling cube puzzle may be visited arbitrarily often, the puzzle can be solved in polynomial time. The time needed is at most linear in the complexity of the state graph.*

We therefore restrict our attention to the case where every labeled face must be visited exactly once.

3 Boards with Free Cells

In this section we show that solving puzzles on boards with labeled, free, and blocked cells is NP-complete. We then refine the proof to show that the problem remains NP-complete with only labeled and free cells.

3.1 Free and Blocked Cells

We show NP-hardness by a reduction from the Hamiltonian path problem in grid graphs. An (*induced*) *grid graph* is an induced subgraph of the infinite grid graph that has vertices (i, j) , $i, j \in \mathbb{Z}$, and edges between vertices of distance one. Grid graphs are uniquely determined by their vertex set. Detecting a Hamiltonian path or cycle in grid graphs is NP-complete [13]. For the reduction, we construct a board with labeled cells representing the vertices of the grid graph, free cells representing the edges, and with all other cells blocked.

Theorem 2 *It is NP-complete to decide whether a die can roll along a path or cycle over a board with labeled, blocked, and free cells, visiting each labeled cell exactly once.*

Proof: The problem clearly is in NP, it remains to prove the NP-hardness. Let V be the vertex set of an induced grid graph and let n denote its size, i.e., $|V| = n$. We can assume V to lie in an $n \times n$ subgrid, because else it would contain isolated vertices. Let us further assume the subgrid to have its lower left corner in the origin, i.e., $V \subset \{1, 2, \dots, n\} \times \{1, 2, \dots, n\}$.

We will construct a board of size $(2n + 1) \times (2n + 1)$ and label its cells such that there is path or cycle in the grid graph if and only if there is a path or cycle, respectively, visiting the labeled




Figure 10: Reduction of a grid graph to a labeled board with blocked and free cells.

cells of the board that can be rolled by a die. Because the board is part of the infinite grid, this shows the NP-hardness for both finite and infinite grids.

The cells of the board are labeled as follows. With each vertex (i, j) of the grid graph, we associate the cell $(2i, 2j)$ of the board. These cells associated with vertices of the grid graph are labeled alternately with the labels 1 and 6. All cells $(2i, 2j)$ where (i, j) is not a vertex of the grid graph are blocked. Furthermore, all cells (k, l) where both k and l are even, and all cells on the border of the board, are blocked. All remaining cells are left free. See Figure 10 for an example.

More specifically, each cell (k, l) is assigned as follows:

$$(k, l) \rightarrow \begin{cases} \text{label 1} & \text{if } k, l \text{ even} \wedge \left(\frac{k}{2}, \frac{l}{2}\right) \in V \wedge \frac{k+l}{2} \text{ even} \\ \text{label 6} & \text{if } k, l \text{ even} \wedge \left(\frac{k}{2}, \frac{l}{2}\right) \in V \wedge \frac{k+l}{2} \text{ odd} \\ \text{blocked} & \text{if } k, l \text{ odd} \vee \left(k, l \text{ odd} \wedge \left(\frac{k+1}{2}, \frac{l+1}{2}\right) \notin V\right) \\ & \vee k \in \{1, 2n+1\} \vee l \in \{1, 2n+1\} \\ \text{free} & \text{else} \end{cases}$$

We claim that there is a Hamiltonian path in the grid graph if and only if there is a Hamiltonian path labeling the board. To see this, first observe that there are some (free) “dead end” cells in the board, namely all those (free) cells that are bordered on three sides by blocked cells or the boundary of the board. These cells cannot be used by a Hamiltonian path and we can therefore also label them blocked as in Figure 10 (left). But now it is obvious that we have just mapped the grid graph to the board.

Any path in the grid graph is rollable on the board because we labeled the cells using alternating opposite labels in every second cell. Giving each cell corresponding to a vertex in the grid graph a label (other than free or blocked) ensures that any labeling Hamiltonian path in the board is also a Hamiltonian path in the grid graph. \square

Note that the reduction uses only two labels, and in fact one label would suffice on a larger board. Using two labels make the puzzles a special case of two-colored dot mazes [2].

3.2 Free Cells

Now we extend the proof for blocked and free cells to hold also for only free cells.

Theorem 3 *It is NP-complete to decide whether a die can roll along a path or cycle over a board with labeled and free cells, visiting each labeled cell exactly once.*




Figure 11: Reduction of an induced grid graph to a labeled board with free cells.

Proof: Again it suffices to prove the NP-hardness. Let again V be the vertex set of an induced grid graph of size n , i.e., $|V| = n$. Again we can assume that $V \subset \{1, 2, \dots, n\} \times \{1, 2, \dots, n\}$. We enlarge our construction from the previous proof, i.e., we now label a board of size $(4n+3) \times (4n+3)$ as follows. With each vertex (i, j) of the grid graph we associate the cell $(4i, 4j)$ of the board and label these cells with the labels 1. All cells $(4i, 4j)$ where (i, j) is not a vertex of the grid graph but a neighbor of (i, j) is a vertex in the grid graph, are labeled 6.

In between all “vertices” and “nonvertices” we label the 3×3 subboards as shown on the right: the middle vertex is labeled 6 and the five corner vertices of the 3×3 subboard are labeled alternatingly with 1 and 6. We do this on all 3×3 subboards that touch at least one “vertex”.

1		6
	6	
6		1

In this construction we cannot put a blocked border around the board. Instead we have constructed a border consisting of 3×3 subboards and cells labeled 6 for “nonvertices” of the grid graph which neighbor vertices in the grid graph.

There are two important observations concerning this construction:

1. From a cell associated with a vertex of the grid graph it is not possible to roll onto any other labeled cell of the board but the cells associated with vertices of the grid graph. A (Hamiltonian) path or cycle rolling over all cells associated with vertices of the grid graph corresponds (exactly) to a (Hamiltonian) path or cycle in the grid graph.
2. All labeled cells which are not associated with a vertex in the grid graph, i.e., also those which are associated with “nonvertices” of the grid graph, can be rolled by (many different) consecutive paths.

In Figure 11, paths and cycles that roll over the labeled cells are indicated by dotted lines.

Proof of Observation 1: Two cells associated with neighboring vertices in the grid graph are connected in the board by three free cells. Rolling from on cell to the next therefore requires four turns of the die, after which the die will have the same number on top as before. To see that it is not possible to roll on any other labeled cell than these, consider the cell in the very center of Figure 11. Assume it can roll on some other labeled cell. Let C be the first labeled cell it can reach. To C it has to roll directly, because there are not enough free cells around it to allow it

			6				6	
					1		6	
						6		
					6		1	
6				1		6		1
		1		6		1		6
			6		6			
		6		1		6		1
6				1				1

Figure 12: Entry to the cells associated with vertices in the grid graph.

to maneuver. But rolling directly to any of the labeled cells in its vicinity will produce the wrong label.

Proof of Observation 2: We partition the labeled cells not associated to vertices of the grid graph in two: The first set are those labeled 6 which are either associated to nonvertices of the grid graph or which lie in the center of a 3×3 subboard. The second set are those labeled 1 or 6 in the corners of a 3×3 subboard. In Figure 11 these two sets are colored black and green, respectively. By a similar spacing argument as for the cells associated with vertices of the grid graph, one sees that it is possible to roll between these sets of cells. Furthermore it is possible to go from one of these paths to the other (arbitrarily often) by leaving the labeled $(4n + 3) \times (4n + 3)$ board. In the free space outside one can easily maneuver to produce any label on the top of the die. Note that this is not possible from the cells associated with vertices of the grid graph, because from these it is not possible to leave the labeled $(4n + 3) \times (4n + 3)$ board.

So now we have two sets of labeled cells, the ones associated with vertices in the grid graph and all others. We know that the cells associated with vertices in the grid graph can be rolled by a Hamiltonian cycle if and only if this is possible in the grid graph. All other labeled cells can also be rolled by a path or cycle. But currently there is no possibility to go between these two sets of labeled cells. Therefore the final twist of the construction is to construct one “entry” to the cells associated with vertices of the grid graph. We make this “entry” such that it can be used only once, in the case we are searching for a path, or twice, in case we are searching for a cycle. Then it will be possible to first roll all cells but the cells associated with vertices, then roll the cells associated with vertices, if possible, by a Hamiltonian path or cycle in the grid graph and, if we are searching for a cycle on the board, leave the cells associated with vertices of the grid graph again and return to the starting point.

We will place the “entry” at the leftmost top vertex in the grid graph. (i.e., of all top vertices we choose the leftmost). We will call the cell associated with this vertex the *entry vertex*. This vertex does not have a top or left neighbor. Therefore to be included in a Hamiltonian cycle, it has to have a right and lower neighbor to which it has to connect.

To make an entry we simply delete the 3×3 subboard to the upper left of the entry cell and label the cells two to the right and two down of the entry vertex with a 6. Cf. Figure 12. Furthermore when searching for a cycle we also delete the label of the entry cell. Then we can roll

onto the (possibly unlabeled) entry vertex from a cell not associated to a vertex in the grid graph by maneuvering in the white space (e.g., from the black 6 on its left, roll two cells yielding a 1 and from there two U-turns to the left).

The two cells labeled 6 two cells down and right of the entry vertex, ensure that we roll a Hamiltonian cycle on the cells associated with vertices of the grid graph (and not a Hamiltonian path). Note that, in contrast to the previous proof, we reduce to Hamiltonian cycle and not path even when searching for a path in the grid, because we can enter the cycle on the grid vertices at any point. \square

3.3 NP-Hardness and Puzzle Design

We have shown that rolling cube puzzles on boards with free and possibly blocked cells where each labeled cell has to be visited exactly once are NP-complete. Does this help us to design difficult puzzles? It does not immediately help us, because difficult instances of NP-complete problems are usually unsolvable instances. And we are now interested in difficult solvable rolling cube puzzles. In particular, we are not interested in the decision problem but in actually finding a solution.

Similar questions have been asked for the 3-satisfiability (3SAT) problem. Extensive research has been done on solving 3SAT instances. To benchmark 3SAT solvers, researchers have considered empirically difficult instances, in particular difficult positive instances.³ For example, the spin glass approach [14] creates instances with many local maxima (in the number of satisfied clauses) and only one satisfying assignment. These are difficult to solve by local search algorithms.

We could construct rolling cube puzzles by starting with a (difficult) 3SAT formula and following the reduction. To reduce the size of the resulting puzzle, our reduction can be modified to directly reduce from planar Hamiltonian cycle [11] or even planar 3SAT [16]. But the resulting puzzles are still too large.

Empirically difficult 3SAT problems demonstrate that it should be possible to create challenging rolling cube puzzles. However, for us this is still an open problem.

Problem 1 *Design a small but difficult rolling cube puzzle (e.g., with many local maxima) where the task is to roll over every labeled cell exactly once.*

With additional rules it is possible to design smaller puzzles from 3SAT formulas; see Figure 13 for an example. The puzzle corresponds to a 3SAT formula with 4 variables (encoded by the sides of the main square) and 16 clauses (encoded by cells labeled 1 and 6).⁴ The 3SAT formula is the smallest nontrivial case of the spin glass approach. The additional rules could be avoided by taking a planar 3SAT instance and by leaving more space between the paths originating from a variable gadget.

4 Boards without Free Cells

What happens when the board contains only labeled and possibly blocked cells? It remains open whether a polynomial-time algorithm can determine Hamiltonicity of a board, even when all the cells are labeled and when the labeling specifies the orientation of the die. In this section we provide

³A SAT-solver competition is part of the annual *International Conference on Theory and Applications of Satisfiability Testing*.

⁴The key to the solution is the following. Rolling along the main square in clockwise orientation, the four sides of the square are entered by cells labeled 2, 4, 2, and 4. At each of these cells one must choose a direction. The correct directions are down, left, down, and right.




Figure 13: A puzzle corresponding to a 3SAT formula. Start and end at the blue die and roll over every labeled cell exactly once. Additional rules are: It is not allowed to make a turn at a crossing. After rolling onto a cell labeled 1 or 6, one must turn back.

a discussion of the problem, with the hope that our observations will lead to establishing a result. On fully labeled boards the following simple observation holds:

Observation 1 *Hamiltonian cycles cannot be rolled on fully labeled boards with an odd number of cells.*

4.1 Elimination and Cutting

Consider the version of the problem where the labeled cells of the board also specify full die orientation, and let G be the state graph induced by such a board. First observe that, if G contains a vertex of degree at most 2, then a Hamiltonian cycle has only one way of visiting this vertex. Let a *chain* of the state graph be a path where all vertices except the first and the last have degree 2 and at least one vertex has degree 2. Such chains are effectively *forced*: they must appear in any Hamiltonian cycle. On the other hand, if a vertex u has degree more than 2 but is connected to two vertices of degree 2, then a Hamiltonian cycle will visit u using the two edges connecting u to its neighbors of degree 2.

Based on these two observations, we define two operations on state graphs: (1) *Elimination*: if a vertex is incident to two forced edges, remove all other incident edges. (2) *Cutting*: if the two endpoints of a chain are also connected by an edge, remove the edge. We apply these two operations exhaustively on G to get a subgraph G' with possibly fewer edges, and such that G' has a Hamiltonian cycle if and only if G does. Observe that it is possible that G' is not a grid graph. The main question is whether G' has properties that enable us to determine its Hamiltonicity in polynomial time.

One of the properties that differentiates state graphs from grid graphs is that state graphs have forbidden configurations. One example of a forbidden configuration is a cycle of length 4 on a 2×2 grid. Other forbidden configurations are the maximum cycles on 3×2 and 4×2 grids. In fact, the shortest cycle in the state graph has length 8 and is the maximum cycle on a 3×3 grid. The cycle with the next shortest length has 10 vertices on a 5×2 grid.




Figure 14: The shortest possible cycle in a state graph.

Fact 1 *The shortest possible cycle in a state graph is the length-8 cycle shown in Figure 14. All other cycles have length at least 10. This fact remains true even if the cycle starts and ends with different die orientations but matching top faces.*


It turns out that these forbidden configurations disallow a large number of vertices of degree more than 2 to be packed closely together in the state graph. Let a *blob* of the state graph be a maximal set of connected vertices each having degree at least 3. The *depth* of a vertex v of a blob is the length of the shortest path from v to a vertex that is adjacent to at least one vertex of degree 2. The *depth of a blob* is the depth of the vertex with maximum depth among all vertices of the blob.

As a step towards computational tractability, we show that depth of any blob is at most 2, and at most 1 after elimination. First we need a preliminary lemma about depth-2 blobs:

Lemma 4 *If a vertex u in a blob has depth 2, then its distance-3 neighborhood is subgraph of a rotation and/or reflection of Figure 15(c).*


Proof: We consider successive neighborhoods of u in the state graph for u to have depth at least 2 in its blob.

Distance 1: Because u is in a blob, u has at least 3 vertices v_1, v_2, v_3 at distance 1 in the state graph; see Figure 15(a). Up to reflection and rotation, we can assume that u has right, up, and left neighbors; label them v_1, v_2 , and v_3 , respectively. The number of empty grid points adjacent to at least one of v_1, v_2, v_3 is equal to 7. We will call this the *availability*; hence the availability at distance 1 is at most 7.



(a) Distance 1 from u in a blob.

(b) Distance 2 from u of depth at least 1.



(c) Distance 3 from u of depth at least 2.

Figure 15: The distance-1, -2, and -3 neighborhoods around a vertex u of sufficiently high depth in a blob. Empty squares represent possibly empty grid points within the neighborhood; empty circles represent possible neighbor locations at the next larger distance.

Distance 2: Because u has depth at least 1 in the blob, each of v_1, v_2, v_3 must have degree at least 3. No v_i and v_j can have a common neighbor w other than u , because otherwise u, v_i, w, v_j, u would be a cycle of length 4, while Fact 1 implies that any cycle in a state graph has length at least 8. Thus, each of v_1, v_2, v_3 must be adjacent to at least two distinct vertices each adjacent to exactly one of v_1, v_2, v_3 . Hence, we need to create at least $3 \cdot 2 = 6$ new vertices w_1, w_2, \dots, w_6 , each at distance at most 2 from u in the lattice; see Figure 15(b). Applying an appropriate reflection, v_2 has a left neighbor w_2 . This neighbor forces v_3 to have a left and bottom neighbor, while the position of the remaining vertices w_1, w_5, w_6 remains flexible (provided the grid point above v_1 and to the right of v_2 is not shared by more than one of v_1, v_2, v_3). Thus we have five different embeddings of the state graph, modulo rotation and reflection. In all except one of these four embeddings, the availability is 11, while for the embedding of the graph in Figure 15(b) the availability is 12.

Distance 3: Because u has depth at least 2, each of the 6 vertices w_1, w_2, \dots, w_6 must have degree at least 3. Again each w_i is adjacent to two distinct vertices each adjacent to exactly one w_i , because otherwise we would form a cycle of length 6. Thus, we need to create at least $6 \cdot 2 = 12$ new vertices x_1, x_2, \dots, x_{12} , each at distance at most 3 from u in the lattice; see Figure 15(c). Thus we need to use the availability-12 embedding of the distance-2 neighborhood state graph in Figure 15(b). Moreover, we claim that the positions of the vertices x_1, x_2, \dots, x_{12} are then forced. Vertex w_2 can only have a top and left neighbor. As a result, w_3 must have a left and bottom neighbor; w_4 must have a bottom and right neighbor; w_5 must have a bottom and right neighbor; w_6 must have a right and top neighbor; and w_1 must have a right and top neighbor. Therefore, up to rotation and reflection, Figure 15(c) is the only way of embedding (this subgraph of) the distance-3 neighborhood on a grid. \square

Lemma 5 *The depth of every blob in a state graph is at most 2.*

Proof: Suppose for contradiction that vertex u of a blob has depth greater than 2. By Lemma 4, the distance-3 neighborhood includes as a subgraph the configuration in Figure 15(c) (up to rotation and reflection). Now we proceed to the next level.

Distance 4: Because u has depth at least 3, each of the 12 vertices x_1, x_2, \dots, x_{12} must have degree at least 3. Thus, we need to have at least two new neighbors for each x_i . However, this is impossible. Consider the right neighbor x_8 of w_4 in Figure 15(c). We cannot give it two more neighbors without creating a cycle of length less than 8, which contradicts Fact 1. Hence, not all vertices at distance 3 from u have degree at least 3. Therefore, the maximum possible depth of any vertex in a blob is at most 2. \square

For the next lemma, we need to avoid degree-1 vertices in the state graph. This condition is certainly necessary if we wish to find a Hamiltonian cycle. For Hamiltonian paths, there are at most two such vertices, forcing endpoints of the path, and they can be dealt with separately.

Lemma 6 *After applying the elimination operation to every vertex of a state graph having no vertices of degree 1, the depth of every blob is at most 1.*

Proof: By Lemma 5, the depth of every blob is at most 2. Consider a vertex u of depth 2, and the required configuration of a subgraph of its distance-3 neighborhood from Lemma 4. If, after applying elimination to all the vertices of the graph, vertex u still has depth 2, then every w_i still has degree at least 3, which in turn implies that at most one of the vertices adjacent to each w_i has degree 2. Thus, in order for u to have depth 2 after elimination, each w_i must be adjacent to at most one vertex of degree 2. Together with the assumption that we do not have vertices of degree

1 in our state graph, we can conclude that we must add at least 18 ($= 12 + 6$) new edges from x_1, x_2, \dots, x_{12} to the graph.

As in Lemma 5, vertex x_8 can have only one new neighbor, on the bottom, because otherwise we would create a cycle of length less than 8. This neighbor and the degree of x_8 forces x_7 to have a bottom and left neighbor. Then x_6 can have only a left neighbor, so x_5 must have a left and top neighbor. Note that, if x_7 shared its left neighbor with x_6 or a right neighbor with x_8 , then this would create a cycle of length less than 8, contradicting Fact 1. On the other hand, if x_4 shared a left neighbor with x_5 , this would create a cycle of length 8, but not of the form required by Fact 1, a contradiction. As a result, both x_4 and x_3 must have degree 2, because no more grid positions are available to create more neighbors; moreover, they cannot share a neighboring vertex because that would create a cycle of length 4. Therefore, w_2 has two neighbors of degree 2, causing the elimination operation to disconnect w_2 from v_2 and thus decreasing u 's depth to at most 1. \square

Whether the bound on the depth of blobs in state graphs is enough to determine its Hamiltonicity in polynomial time remains unclear. We conjecture that it is:

Conjecture 1 *On boards with labeled and possibly blocked cells, rollable Hamiltonian cycles can be determined in polynomial time.*

4.2 Uniqueness of Cycles?

At first we conjectured that, if a state graph contains a Hamiltonian cycle, then this cycle is unique. If true, this property would possibly make it easier to determine Hamiltonicity in polynomial time, as it would increase the restrictions on the Hamiltonian cycle if one exists. The conjecture, however, is false for boards that contain blocked cells:

Observation 2 *There are boards with labeled and blocked cells in which rollable Hamiltonian cycles are not unique.*

The state graph in the example of Figure 16 is composed of two cycles of length 8, where each cell of the first cycle is connected by a rollable path to a cell of the second. By carefully connecting copies of the state graph induced by the labeling of the board in Figure 16, we can generate a state graph containing multiple Hamiltonian cycles. It is still unclear whether this counterexample leads to a hardness proof, or whether it can be conquered by dynamic programming. We believe the latter is the case.

It remains open whether nonuniqueness of Hamiltonian cycles holds for fully labeled boards. For such boards we pursued the following two approaches: enumeration for small boards and constructing cycles from corners. On small fully labeled boards we can test uniqueness of Hamiltonian cycles by enumerating all possible solutions by hand or by computer. To achieve this, we first enumerate all Hamiltonian cycles on the board; we then check whether the cycles are *rollable*, that is, whether a die can be rolled along the cycle, generating a consistent labeling, and starting and ending at the same state. For all rollable cycles, we check whether they are *uniquely rollable*, that is, given a labeling obtained by rolling a die along the cycle, we check whether this labeling can also be obtained by a different cycle.

Consider the examples in Figure 17. For $n = 2$, one rollable Hamiltonian path exists which is shown in the figure. This path is uniquely rollable (up to change of direction): by the U-turn argument, the first and last cell need to have the same label and the other cells need to have different labels. Thus only one U is possible. As mentioned earlier, the only Hamiltonian cycle on the board is not rollable: any two neighboring cells would form a U -turn and would therefore




Figure 16: A labeling of the board containing more than one Hamiltonian cycle. The dark cells are blocked.




Figure 17: Boards of small size: $n = 2, 4$.

need to have the same label. That is, all cells would need to have the same label, which is not rollable. For $n = 4$, several unrollable Hamiltonian cycles exist, but (up to rotation) there is only one rollable Hamiltonian cycle; see Figure 17.

By enumerating Hamiltonian cycles on boards up to side lengths 8 by computer, we obtain the following observation.

Observation 3 *All rollable Hamiltonian cycles in fully labeled boards with side lengths at most 8 are unique.*

We explored how *border obsessiveness* (a term used to describe a strategy for solving jigsaw puzzles) can be applied to rolling cube puzzles. For jigsaw puzzles, people basically use one of two strategies: a person is either an “opportunistic” or “border obsessive” jigsaw puzzler [15]. In the following we explore how border obsessiveness helps for our problem of rolling a die.

The majority votes in Section 2.3 directly yield that we know the orientation of the die for all border squares. For the four-corner squares we even know the neighboring squares on the cycle, because there are only two neighboring squares. A feasible approach is therefore to start at a

corner and to construct the cycle from there, for instance by inductively deducing the edges of the cycle diagonal by diagonal. We show that this is possible at least for the first squares. Let the k -neighborhood of a cell be all cells that can be reached in k steps ignoring the labels of the cells. The following proposition can be proved by a case distinction. It shows that this strategy works at least to some extent.

Proposition 7 *If a labeling of an $n \times n$ board admits a Hamiltonian cycle, then the path of the die within the 3-neighborhoods of the corners is unique (up to direction). For $n = 4$, the rolling pattern at a corner is determined by the 2-neighborhood of this corner. For $n > 4$, the 5-neighborhood determines the pattern.*

Proof: For $n \leq 3$, there is no possibility to roll a cycle. For $n = 4$, there are two symmetric rollable Hamiltonian cycles, the “H” and the rotated “H”. These two cycles can be distinguished by looking at a 2-neighborhood of a corner.

The interesting case is when $n > 4$. Figure 18 shows all possible 3-neighborhoods of the upper-left corner. For each possibility the rolling pattern is shown and below that the labeling we get if we fix the label and orientation of the corner to 1^2 . The cases $A'-D'$ are symmetric to $A-D$. In the figure we included the orientations of squares which can be deduced using the majority votes (Section 2.3). Assume a set of labels allows for two different rolling patterns in the upper-left corner. By overlaying the labelings in Figure 18 we can directly distinguish all cases by looking at the 3-neighborhood, except possibly the following pairs: (A, B) , (A, E) , (B, D) and (C, E) , and the symmetric cases (A', B') , (A', E) , (B', D') and (C', E) . It suffices to consider the first four pairs.

Assume A and B are not distinguishable. Then we are in the situation of Figure 19(a). By a majority vote we actually know that the 5 without orientation is actually a 5^6 . Thus, for B to be possible the neighbors of this square must be as in Figure 19(b). Now we know that the 6 is a 6^1 because it cannot connect to its neighboring 6^2 . Also, it must be able to connect to the square below it because the cycle when going through the 6 has only one further possibility. Therefore, the square below is a 3^1 and to allow B the neighbors of the 3^1 must be as in Figure 19(b). Now finally, it is clear that the 6^2 must connect to the 5^6 and to its neighbor below the 3^6 . But this means that the 3^6 cannot connect to this square and therefore must connect to 5^6 . With this we are done because with 6^2 and 3^6 connected to 5^6 it must be pattern B :

The cases A and E are distinguishable because the 4^6 in case E would require a 6^3 below it but for case A the 6^5 to the left of this square would require a 4^5 on this square. The cases B and D can be handled in the same way: the 3^2 would require a 2^4 below it, but the 2^1 requires a 3^1 on the right of it.

The cases C and E are again slightly more involved. The overlay of the labels are shown in Figure 20(a). The 5^1 needs to connect as shown in Figure 20(b). The 4^5 needs a 1^5 to the right of it as shown in Figure 20(c) and then the neighbors of 1^5 must be as in Figure 20(d). Then the 4^1 must connect to 5^1 and below, the 1^2 must connect to 5^1 and below. But then we must be in case C .

Thus all patterns are distinguishable, and for this we used at most the 5-neighborhood of a corner. \square

4.3 Finding Long Rollable Cycles

Suppose we are given a fully labeled board and we ask for a rollable cycle that visits the maximum number of cells, without necessarily visiting them all. We show that finding a maximum cycle on a fully labeled board is NP-complete even when we are given a starting position and orientation.






Figure 18: Possible configurations of a corner's 3-neighborhood.




Figure 19: Distinguishing configurations A and B .




Figure 20: Distinguishing configurations A and E .

Theorem 8 *Deciding whether a die can roll along a cycle of length K over a board with labeled cells is NP-complete, even given starting position and orientation.*

Proof: Again, we reduce from Hamiltonicity of grid graphs; see Figure 21. Consider a grid graph $G = (V, E)$, described by a set of n grid points, inducing m edges. Scale all distances by a factor of 4, resulting in a different drawing of G ; place this on and parallel to the orthogonal subdivision, such that grid vertices lie inside of squares. This drawing uses a subset S of $n + 3m$ squares, labeled on blue in the figure; the n squares with a vertex of G gets the label 1, marked with red circles in the figure. Furthermore, any edge (shown in red) uses an additional set of three adjacent squares between two vertices. Any triple carrying a horizontal edge get the labels (west to east) 3, 6, 4, while a triple carrying a vertical edge gets the labels (north to south) 5, 6, 2. Finally, all other, “unused” squares of the plane get label 6. (In the figure, these are shown in black.)

Now place the die on a vertex square, 1 facing up, 2 facing south, 3 facing west. We claim that there is a closed tour of length $4n$ if and only if the original grid graph G has a Hamiltonian cycle.

It is easy to see that a cycle allows a long closed tour; we argue the converse. Observe that no feasible tour or path can contain an L-shaped portion for which both end points are identical or opposite numbers. Similarly, no feasible tour or path can contain a straight subpath with three vertices for which both end points contain identical numbers. Finally, 6–6 is also illegal.

This implies that no tour can contain an unused square: clearly, the only possible neighbors are used squares; only an unused square diagonally adjacent to a vertex square with a 1 has two used neighbors. Thus consider using both of these neighbors; avoiding an illegal I-shape forces the next




Figure 21: Finding long cycles is NP-complete.

neighbor to be another used square, either a 1 or a 6, forcing an illegal L-shape. Now the claim follows easily.

To see the same for paths, we only have to deal with unused squares near the path ends; but the resulting marginal difference in cost does not change the fact that the total length of the path is large only if there is a Hamiltonian cycle in the original grid graph. \square

5 Coloring a Board with a Black–White Cube

We now come to a different class of puzzles, namely, a generalized version of the “The red-faced cube” shown in Figure 2(b). Related to these puzzles is also the second puzzle in Figure 3.

We consider a cube with black and white faces and a quadratic board. When we roll the cube over a cell of the board, we color this cell with the color on the bottom of the cube. Every of the n board cells can be visited only once and we want to color all cells of the board. In the end, we have a number of black cells n_b and a number of white cells. We ask the following question:

What is the highest percentage of black cells n_b/n that can be achieved with such a rolling?

In other words we ask for the value of

$$q_b := \limsup_{n \rightarrow \infty} \max_{\text{ham.paths}} n_b/n = ?$$

Of course the results depend on the cube we are using.

We will identify all different black-white cubes by the following system. First we count the number of black faces f_b . If f_b is one of 0, 1, 5, 6, then the cube is uniquely described. If f_b is 2, we have two variants of the cube: the symmetric variant (two opposite faces are black) and the asymmetric variant (two incident faces are black). For $f_b = 4$, we classify the cube in the same way by considering the white faces. Finally, if we have three black faces, we call the situation *symmetric* where all of the faces share a vertex and *asymmetric* otherwise. To distinguish between the different cube variants, we abbreviate every cube with $C_{f_b}^x$, where $x \in \{\emptyset, s, a\}$ denotes the symmetric or asymmetric variant if necessary. The problem to find the maximal percentage of black cell is called the $C_{f_b}^x$ *blackening problem*.

In the following, we discuss two cases of the problem in more detail. An interesting case is the C_3^s blackening problem, because it is the perfect equally distributed case and will not give the black faces any advantage over the white ones. Another natural question is the C_5 blackening problem, which is related to the “red-faced cube” problem introduced by John Harris [6, 8].

5.1 C_5 Blackening Problem

We consider a C_5 cube with only one white face. We want to roll the cube over the board using the white face as little as possible.

Theorem 9 *The solution for the C_5 blackening problem is $q_b = 1$.*

Proof: We will construct the Hamiltonian path by using the building block shown in Figure 22. The block has a height of 4 cells. The width is not determined because we can enlarge the middle zig-zag path. The width is $6 + 2i$ where i is the number of iterations of the zig-zag path. We will




Figure 22: A block which can be transversed without using the white face.

make the block as wide as the board length allows us to. Then we pile up copies of the block until we reach the board's top boundary. Every block will be traversed by the Hamiltonian path drawn in Figure 22. Because the starting and end points of the paths of several blocks meet, we come up with a large Hamiltonian path for most of our board. In Figure 22, one can trace the position of the white face along the path. The position is abbreviated with L (left), T (top), B (back), and F (front). Notice that the orientations of the different blocks fit together.

While sweeping through the blocks we do not use the white face at all. What is left is to color the rest of the board (if there is any). This leads to $q_b = 1$ in the limit. \square

5.2 C_3^s Blackening Problem

Theorem 10 *The solution for the C_3^s blackening problem is $q_b = 0.75$.*

Proof: *Upper bound:* First we show that every Hamiltonian path produces at most $3/4$ black cells. We do not need the Hamiltonicity here because even for every general path this bound holds. Throughout the proof we assume that we are using a regular die, where the faces 1, 2, 3 are colored black and the faces 4, 5, 6 are colored white. This represents the situation for the C_3^s cube but allows us to identify the orientation more easily.

Without loss of generality our path starts with a 1 and the 2 is on the right side. We observe that there is only one way to go four consecutive steps while using only black faces (down, right, up). Now we have to use a white face. We have two possibilities, going up or going right.

Case 1: We go right. The only way to use a black face is now to go two down and one left. Now we have to consume a white face again. Going left is not a good idea, because this will cost us two white cells, therefore we have to go down. Then we can go two right and one up. This leaves us in the situation at the beginning of case 1. Figure 23(a) shows the best possible path.

Case 2: We go up. Then our only choice is two left and one down. It turns out that the best possible path is the rotated variant of Case 1; see Figure 23(b) for the complete path.




Figure 23: The best possible paths of length 7 in the two cases.

In both cases the best sequence we can color is first four black cells and then one white cell and three black ones iterated. Clearly this leads to a bound of $3/4$ of black cells in the limit.

Lower bound: We prove the lower bound by showing a way how to color a board with asymptotically $3/4$ black cells. We need six different types of building blocks for the construction. The most essential block is the *central block*. It will be responsible for the high number of black colored cells. The Hamiltonian path we are using crosses the block four times. See Figure 24 for the detailed path. Passing through the block (on any of its four paths) will not change the orientation of the cube, only its position. When we roll the block like shown in Figure 24 we will cover $3/4$ of the cells black. One can observe that central blocks fit together and the described paths looks like a diagonal zig-zag-path across the board.

It remains to complete the boundary of the board. We have to take care that the turns we make keep the orientation of the cube when we enter a central block the next time. The blocks which are necessary to do this are called *h block* at the top and bottom and *side block* on the left and right side of the board (see Figure 25). We could finish the paths arbitrarily, but we can even do better and finish with a Hamiltonian cycle. We need three other blocks for this (and their rotated counterparts). All three are shown in Figure 26. Finally, we can plug everything together and will




Figure 26: Blocks for the Hamiltonian cycle.




Figure 27: An example with 4 central blocks.

C_2^a blackening problem: The upper bound for this case can be constructed in the same way as was done in Theorem 10. We have used a computer program to backtrack all possibilities. It turns out that we cannot get better than $q_b = 1/2$. Notice that this bound is tight. We know that q_b for the C_3^s case is $3/4$. We now take a board colored by an optimal path for the C_3^s cube. We had associated the numbers 1, 2, 3 with the black faces. $3/4$ of the cells contain one of these numbers. We take the number which has colored the least cells. The associated face will now become white. The resulting cube has the form C_2^s and colors $1/2$ of the cells black by the pigeon hole principle.

C_3^a blackening problem: It is easy to observe that the path with the best result for q_b simply rolls straight and colors three cells in a row black. This leads to an upper bound of $3/4$. A good example for the lower bound is the spiral walk of Figure 29. We color along the ingoing spiral $3/4$ -th of the cells black and along the outgoing $1/4$. This gives us a lower bound of $5/8$.




Figure 28: An optimal path for the C_1 blackening problem.




Figure 29: An optimal path for the C_2^s blackening problem.

C_4^a blackening problem: We found an upper bound with the help of a computer program by showing that every path can color at most $11/12$ of the cells black. This was done by backtracking. For the lower bound we use the construction given in Theorem 10. When we take a closer look at the central block we discover that the numbers 1, 2, 3, 5 color $7/8$ of the area black. Considering these numbers as black faces of the C_4^a cube gives us a lower bound of $7/8$.

C_4^s blackening problem: We show that $q_b = 1$ by the following construction. We start at the boundary and construct an empty area in the middle which looks like a rectangle minus the 4 corner cells. Then we begin to eat up the free cells along a spiral path (see Figure 30). This leads to a completely black colored area, except for the boundary and the diagonals.

Table 1 summarizes our results for the different cubes.




Figure 30: An optimal path for the C_4^s blackening problem.

cube	lower bound	upper bound	gap
C_1	$1/3$	$1/3$	—
C_2^s	$1/2$	$1/2$	—
C_2^a	$1/2$	$1/2$	—
C_3^s	$3/4$	$3/4$	—
C_3^a	$3/4$	$5/8$	$1/8$
C_4^s	1	1	—
C_4^a	$7/8$	$11/12$	$1/24$
C_5	1	1	—

Table 1: Summary of the results for blackening problems.

Acknowledgments

We would like to thank Joseph O’Rourke for suggesting and discussing the problem with us. We would also like to thank Robert Abbott, Frank Hoffmann, Klaus Kriegel, Günter Rote, and the participants of the 4th GWOP (Group Emo Workshop on Open Problems) for helpful discussions. Furthermore, we would like to thank Robert Abbott for allowing us to include his mazes.

References

- [1] Robert Abbott. *Mad Mazes: Intriguing Mind Twisters for Puzzle Buffs, Game Nuts and Other Smart People*. Adams, 1990.
- [2] Robert Abbott. *SuperMazes: Mind Twisters for Puzzle Buffs, Game Nuts, and Other Smart People*. Prima Publishing, 1997.

- [3] Robert Abbott. A maze with rules. In Martin Gardner, Elwyn Berlekamp, and Tom Rodgers, editors, *The Mathemagician and Pied Puzzler: A Collection in Tribute to Martin Gardner*, pages 27–28. AK Peters, Ltd, 1999.
- [4] Erik D. Demaine and Joseph O’Rourke. Open problems from CCCG 2005. In *Proc. 18th Canadian Conference on Computational Geometry (CCCG’06)*, pages 75–80, 2006.
- [5] Martin Gardner. Mathematical games column. *Scientific American*, 209(6):144, December 1963. Solution in January 1964 column.
- [6] Martin Gardner. Mathematical games column. *Scientific American*, 213(5):120–123, November 1965. Problem 9. Solution in December 1965 column.
- [7] Martin Gardner. *Martin Gardner’s Sixth Book of Mathematical Diversions from Scientific American*. W H Freeman, San Francisco, 1971. Chapter 8.
- [8] Martin Gardner. *Mathematical Carnival*. Borzoi / Alfred A Knopf, New York, 1975. Chapter 9, Problem 1: The red-faced cube.
- [9] Martin Gardner. Mathematical games column. *Scientific American*, 232(3):112–116, March 1975. Solution in April 1975 column.
- [10] Martin Gardner. *Time Travel and Other Mathematical Bewilderments*. W H Freeman, New York, 1988. Chapter 9, Problem 8: Rolling cubes.
- [11] M. R. Garey, D. S. Johnson, and R. Endre Tarjan. The planar Hamiltonian circuit problem is NP-complete. *SIAM Journal on Computing*, 5(4):704–714, 1976.
- [12] John Harris. Single vacancy rolling cube problems. *J. Recreational Mathematics*, 7(3):220–223, 1974.
- [13] Alon Itai, Christos H. Papadimitriou, and Jayme Luiz Szwarcfiter. Hamilton paths in grid graphs. *SIAM J. Comput.*, 11(4):676–686, 1982.
- [14] Haixia Jia, Cristopher Moore, and Bart Selman. From spin glasses to hard satisfiable formulas. In *Theory and Applications of Satisfiability Testing, 7th Intl. Conf. (SAT 2004)*, volume 3542 of *Lecture Notes in Computer Science*, pages 199–210. Springer, 2005.
- [15] Hilary Johnson and Joanne Hyde. Towards modeling individual and collaborative construction of jigsaws using task knowledge structures (tks). *ACM Trans. Comput.-Hum. Interact.*, 10(4):339–387, 2003.
- [16] David Lichtenstein. Planar formulae and their uses. *SIAM Journal on Computing*, 11(2):329–343, 1982.
- [17] Ronald Sprague. *Unterhaltsame Mathematik*. Vieweg & Sohn, Braunschweig, 1961. Problem 3: Schwere Kisten.
- [18] Ronald Sprague. *Recreations in Mathematics*. Blackie, Glasgow, 1963. Problem 3: Heavy Boxes. Translation by T. H. O’Beirne.
- [19] Charles W. Trigg. Tetrahedron rolled onto a plane. *J. Recreational Mathematics*, 3(2):82–87, 1970.

- [20] M. Oskar van Deventer. Rolling block mazes. In Barry Cipra, Erik D. Demaine, Martin L. Demaine, and Tom Rodgers, editors, *Tribute to a Mathemagician*, pages 241–250. AK Peters, Ltd, 2004.