Rolling a Cube can be Tricky

Marzio De Biasi
marziodebiasi [at] gmail [dot] com

May 2012
Version 0.03

Abstract

We settle two open problems related to the rolling cube puzzle: Hamiltonian cycles are not unique even in fully labeled boards and rolling cube puzzle is NP-complete in labeled boards without free cells and with blocked cells.

1 Introduction

In a *rolling cube puzzle* a die must be rolled on a board visiting all its labeled square cells and return to its starting location; at every step the number on the top face of the die must match the number of the underlying *labeled cell*. The die can be rolled between two adjacent cells by tipping it over along one of its bottom edges that touches the board; it cannot be rotated within the same cell.

Rolling cube puzzles have been made popular by Marting Gardner who wrote about them in three of his famous *Mathematical Games* columns published in *Scientific American* review.

There are many variants of the game, for example: *Heavy boxes* [3], *Red-Faced Cube* [4], *Single Vacancy Problem* [5], *rolling cube mazes* [1]. In this paper we focus on the variant studied in [2]: the *die* is a *standard right-handed die* (see Figure 1), the *board* is a grid with square cells; cells can be of three types: *labeled* with a number between 1 and 6, *free* or *blocked*. The die must visit all labeled cells exactly once and the number on the top face must match the number of the cell; free cells can be visited any number of times and with any number on the top face; finally the die cannot roll over a blocked cell. The decision problem "Can a die be rolled over a labeled board with some free cells and some blocked cells?" is NP-complete [2].

In Section 2 we define the notation used in the rest of the paper. In Section 3 we give an example of a full labeled board without free cells and without blocked cells with two distinct Hamiltionian cycles. This solves an open problem posed in [2].

In Section 4 we define the *gadgets* that will be used in Section 5 to prove that it is NP-compete to decide whether a die can be rolled over a labeled board without free cells and with some blocked cells. This improves the result achieved in [2].

The complexity of rolling a die over a full labeled board without free cells and without blocked cells remains an open problem.

2 Notation

A standard right-handed die and its unfolded faces are show in Figure 1.

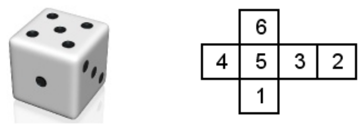

Figure 1: A die and a valid rollable path on a labeled board.

The board is represented using a grid with labeled square cells; blocked cells are represented with empty squares. We will not use free cells.

The *state* of the die can be specified by its board position, its top face and its orientation. Indeeed if we fix the top face, there are still four possible orientations. We will represent the orientation with a small dash placed upward, downward, leftward or rightward near the number of the top face: if the number on the top face of the die is 2, 3, 4, or 5 then the dash indicates the position of the face with number 1; if the number on the top face is 1 or 6 then the dash indicates the position of the face with number 2. For example the state of the die in Figure 1 is $\underline{5}$: a five with the small dash placed downward.

The *state graph* G has a vertex for each possible state of the die and an edge for each possible transition between two states, i.e. an edge represent two adjacent cells for which it is possible to roll a die from one cell to the other respecting the orientations. The moves are bidirectional, so the state graph is undirected.

Valid rolling paths, i.e. the edges of the state graph, are represented with colored blue lines (an example is shown in Figure 2).

The two rules of the puzzle that a labeled cell can be visited only once and that the die must return to its starting location after visiting all the cells, can be used to easily exclude some rolling paths and reduce the state graph.

For example if the state graph contains a vertex of degree 2, then a Hamiltonian cycle has only one way of visiting it. A *chain* of the state graph is a path where all vertices except the first and the last have degree 2; these chains must appear in any Hamiltonian cycle, so we say that the edges of chain are *forced*.

We will use the two reductions *elimination* and *cutting* defined in [2]:

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 can apply these two reductions exhaustively on G to get a simpler subgraph G' but such that G' has a Hamiltonian cycle if and only if the original graph G does.

Valid rolling paths after elimination and cutting are represented with red lines (an example is shown in Figure 3).

In summary:

- **bold number**: label of the cell;
- **small dash**: the small dash placed near the labels shows the allowed die orientation (up, left, right, bottom);
- **blue line**: rollable path(s);
- **red line**: rollable path(s) after *elimination* and *cutting* reductions [2];
- **white cells**: blocked (unrollable) cells.

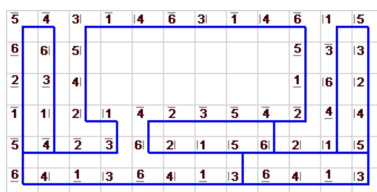

Figure 2: A labeled board with some blocked cells.

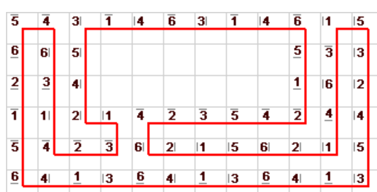

Figure 3: The same board of Figure 2 after elimination and cutting.

3 Hamiltonian cycles are not unique even in fully labeled boards

Nonuniquess of Hamiltonian cycles for boards with labeled and blocked cells has been showed in [2]. But nonuniqueness of Hamiltonian cycles holds also for fully labeled boards without blocked cells.

Proposition 3.1. *There are full labeled boards in which rollable Hamiltonian cycles are not unique.*

Proof. Figure 4 shows a board without free cells and without blocked cells, that has two distinct Hamiltonian cycles.

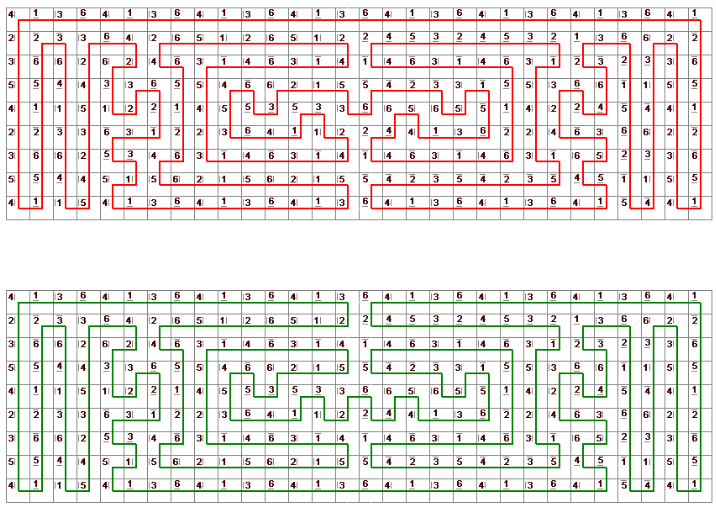

Figure 4: Full labeled board with two distinct Hamiltonian cycles.

□

4 Some gadgets

In the next sections we will use a smaller 16×9 board with some blocked cells that nevertheless has two Hamiltonian cycles (see Figure 5) and we will call it *basic gadget*. We label the cells placed in the middle of the left, top, bottom and right border (E_1, E_2), (F_1, F_2), (G_1, G_2) and (H_1, H_2) respectively.

If we extend the basic gadget by placing two labeled cells above the cells (F_1, F_2), two labeled cells below the cells (G_1, G_2), and two labeled cells beside the cells $(E1, E2)$ we obtain the extended board in Figure 6.

We can observe that the two pairs of cells (F_1, G_1) and (F_2, G_2) *act like a switch.*

- if the die uses one of the top cells F_1 or F_2 to enter or exit the basic gadget (or rolls over them), then it will not be able to use the corresponding bottom cell G_1 or G_2 (or both) to enter or exit the gadget;

- if the die uses one of the bottom cells G_1 or G_2 to enter or exit the basic gadget (or rolls over them), then it will not be able to use the corresponding top cell G_1 or G_2 (or both) to enter or exit the gadget.

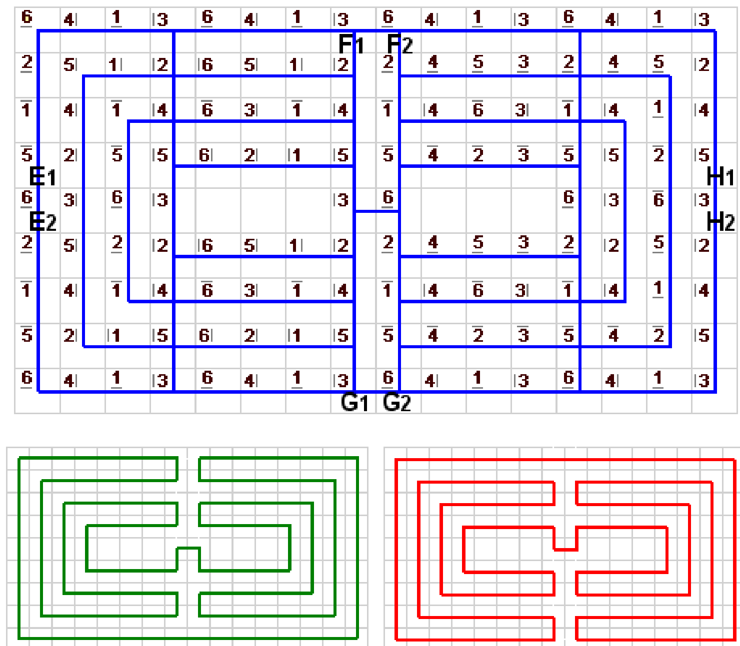

Figure 5: Basic gadget with some blocked cells and its two possible Hamiltonain cycles.

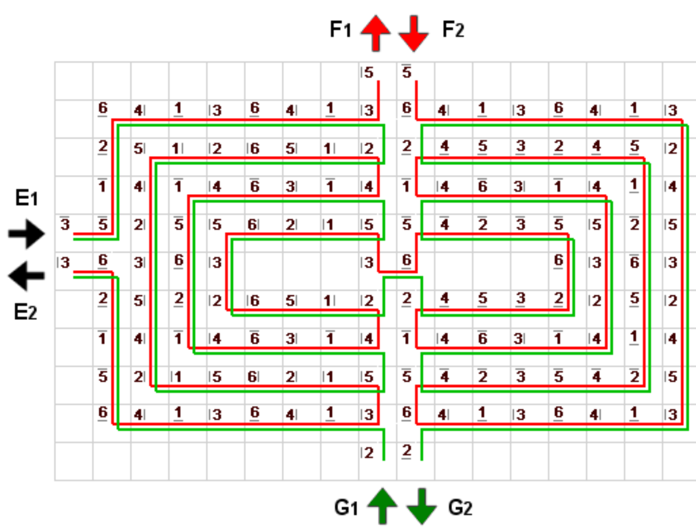

Figure 6: First attempt to extend the basic gadget. Cells (F_1, G_1) and (F_2, G_2) act like a switch.

The proof is done by enumerating the possible rolls of a die that starts outside of the border and enters the basic gadget (see Appendix A.1).

If we apply the elimination and cutting algorithm [2], and remove cells with degree 2, we can represent the choice gadget with the simpler planar graph shown in Figure 7. We will use this simpler graph to represent the more complex gadgets.

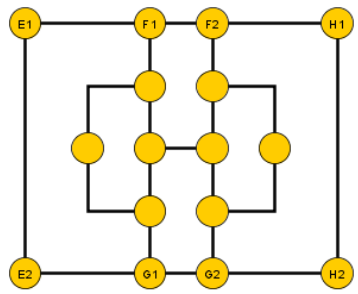

Figure 7: A graph representation of the basic gadget.

4.1 Choice gadget

We will need to stack two basic gadgets vertically without any cells between them, but this is not possible due to the number configuration of the top and bottom borders. So we can slighty modify it adding some cells above and below, but keeping the same rolling properties. We obtain the *choice gadget* shown in Figure 8.

The possible rolls of a die that starts outside of the border and enters the choice gadget are shown in Appendix A.2 and are equivalent to those found for the basic gadget.

Two choice gadgets can be stacked simply starting from cells (G_1, G_2) of a choice gadget and mimic a die roll that "draws" another choice gadget below it. The border between them is showed in Figure 9. If we apply the elimination and cutting algorithm the only way for a die to cross the border is through cells $G_1 \leftrightarrow F_1$ or $G_2 \leftrightarrow F_2$ (it cannot roll along the vertical paths colored in blue). So we can still represent a choice gaget using the graph in Figure 7.

4.2 T-gadgets

If we arrange six choice gadgets in a 3×2 larger grid and connect them horizontally through cells (E_1, E_2), (H_1, H_2), and vertically through cells (F_1, F_2), (G_1, G_2), we obtain the gadget in Figure 10 (the equivalent graph is shown in Figure 11). The lower choice gadgets don't need to be connected, so their bottom border corresponds to the bottom border of the basic gadget.

With this configuration we can exploit the mutual exclusion property of the cells F_1, G_1 and F_2, G_2 of the choice gadgets: in order to completely roll the

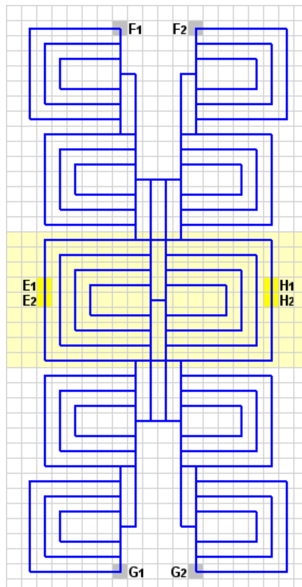

Figure 8: The choice gadget; yellow cells correspond to the original basic gadget.

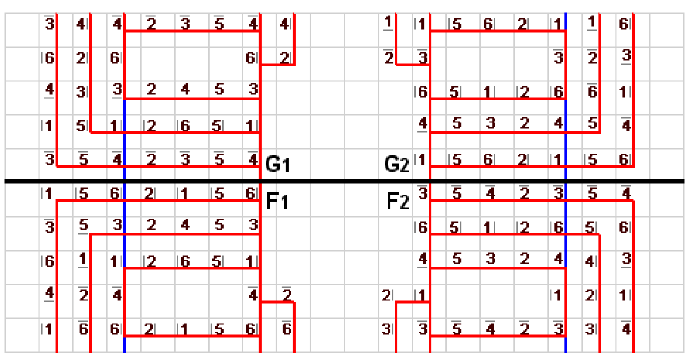

Figure 9: The border between two stacked choice gadgets. After applying the elimination and cutting algorithm, the only way for a die to cross the border is through cells $G_1 \leftrightarrow F_1$ or $G_2 \leftrightarrow F_2$ (blue edges are discarded).

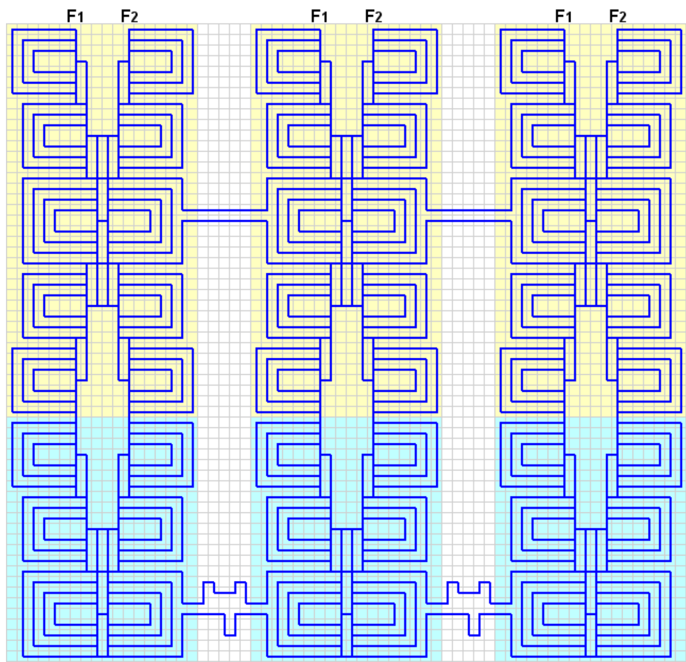

Figure 10: Six choice gadgets arranged in a 3 × 2 larger grid and linked horizontally and vertically.

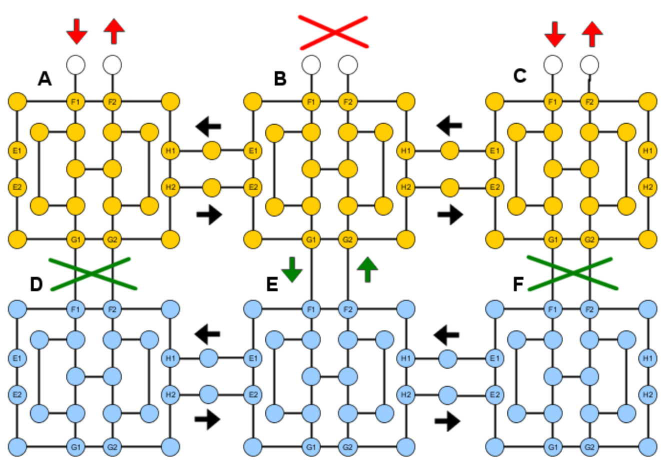

Figure 11: A graph equivalent to the six choice gadgets arranged in a 3 × 2 larger grid.

gadget in Figure 11, the die must enter through one of the top choice gadgets, suppose it enters from A. For the property of the choice gadget, it cannot go down to D and it must roll on the right to B. In order to complete the three lower choice gadgets, the die must choose to go to choice gadget E or F from the top, if it chooses E, then it can completely roll the three lower choice gadgets $(D, E, F,)$ and can return to B from E (if it chooses to go from F to C, then it cannot go up from C). At this point the die cannot go up from B, so it must roll to C and from C it can go up and roll another part of the gadget and then return back to complete C, B and A. The complete sequence is:

```
OUT->A->B->E->D!->E->F!->E!->B->C->OUT->C!->B!->A!->OUT
```

(the exclamation mark indicates when the choice gadget is completely rolled).

By rotating and rearranging the gadget of Figure 10, we can build four *T-gadgets* in which the central choice gadget is oriented respectively upward, downward, rightward, or leftward. The graph representation of the upward T-gadget is showed in Figure 12a. For better clarity, in each T-gadget we will label:

- L_1, L_2 the cells F_1, F_2 of the leftward choice gadget;
- U_1, U_2 the cells F_1, F_2 of the upward choice gadget;
- R_1, R_2 the cells F_1, F_2 of the rightward choice gadget;
- D_1, D_2 the cells F_1, F_2 of the downward choice gadget;

We well call these cells *iterface cells*.

In order to build a T-gadget we need a 138×138 labeled board with blocked cells. The number configurations on the borders are such that the T-gadgets can be combined horizontally and vertically: a die can roll between interface cells $L_1 \leftrightarrow R_1$, $L_2 \leftrightarrow R_2$, $U_1 \leftrightarrow D_1$, $U_2 \leftrightarrow D_2$. See Appendix A.3 for an example and more details.

4.3 Link-gadgets

In addition to the T-gadgets, we will need six more *link-gadgets*; each one is built using two choice gadgets connected together through cells (E_1, E_2) and (H_1, H_2) and oriented respectively: leftward/upward, upward/rightward, rightward/downward, downward/leftward, leftward/rightward, upward/downward (see Figure 13).

Every link-gadget can fit in a 138×138 labeled board with blocked cells. The number configurations on the borders are such that the link-gadgets can be combined horizontally and vertically (and can be combined with T-gadgets as well): a die can roll between cells $L_1 \leftrightarrow R_1$, $L_2 \leftrightarrow R_2$, $U_1 \leftrightarrow D_1$, $U_2 \leftrightarrow D_2$. See Appendix A.4 for more details.

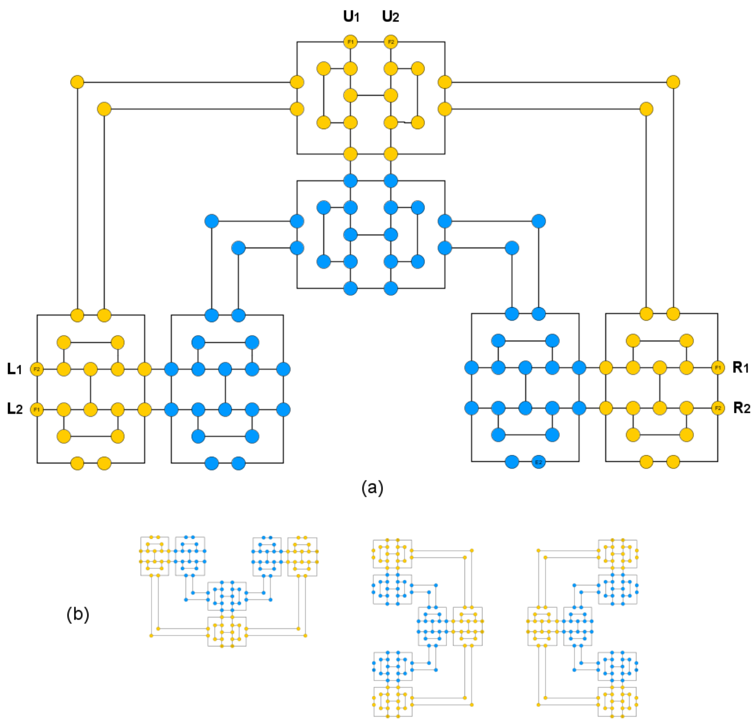

Figure 12: The graph representation of the upward T-gadget (a), and downward, rightward, leftward T-gadgets (b).

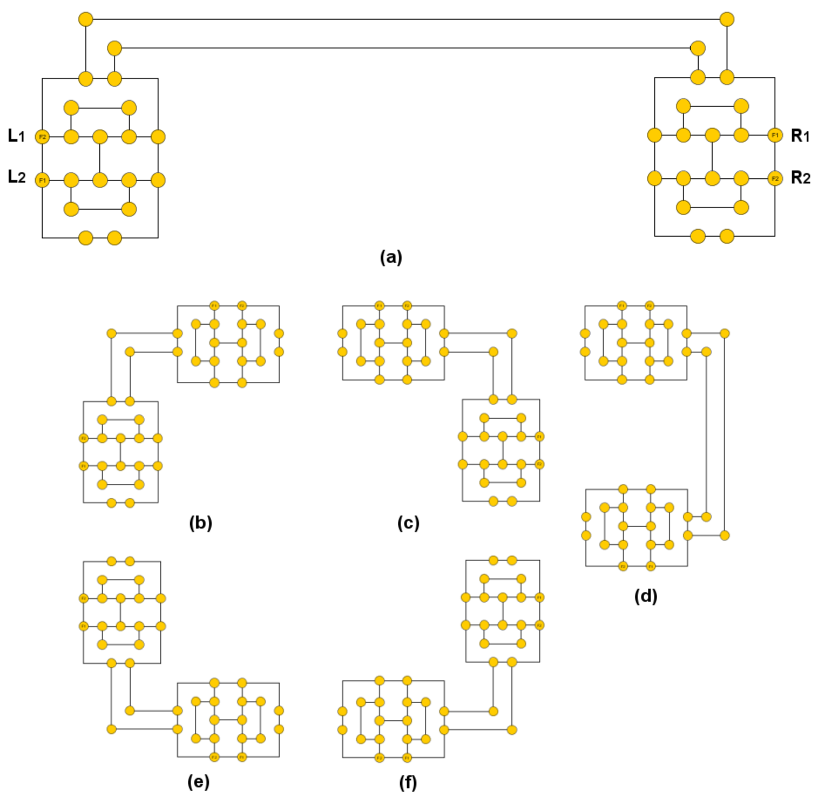

Figure 13: Graph representation of the six link-gadgets: left/right (a), right/down (b), down/left (c), up/down (d), up/right (e) and left/up (f).

4.4 Endpoint-gadgets

We will also use four *endpoint-gadgets*: a few labeled cells placed on one of the four sides of a 138×138 board (see Figure 14). The endpoints will be used to make a direct "rollable shortcut" between the cells L_1 and L_2 (or R_1 and R_2, or U_1 and U_2, or D_1 and D_2) of a larger t-gadget or link-gadget. See Appendix A.5 for more details.

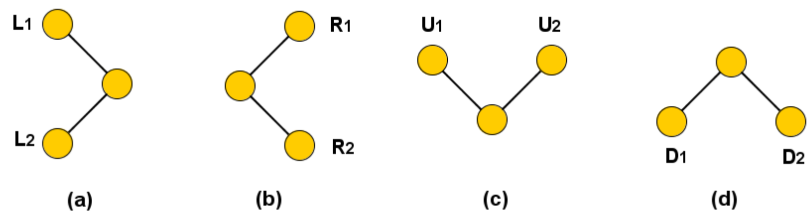

Figure 14: Graph representation of the four endpoint-gadgets: left (a), right (b), up (c), down (d).

5 Main theorem

We use the gadgets to *mimic a grid graph traversal with a rolling die.*

5.1 Grid graphs

The *Hamiltonian Path problem* (HP) is NP-complete for various classes of graphs including grid graphs with degree ≤ 3 [6] (see Figure 15). In particular the problem *"Given a grid graph G with degree ≤ 3 does an Hamiltonian path without specified endpoints exist?"* is NP-complete.

A *grid graph* $G = (V, E)$ is a node-induced finite subgraph of the (infinite) grid, i.e. $V \subset Z \times Z$, and $E = \{\{(x, y), (x', y')\} : (x, y), (x', y') \in V$ and $|x - x'| + |y - y'| = 1\}$.

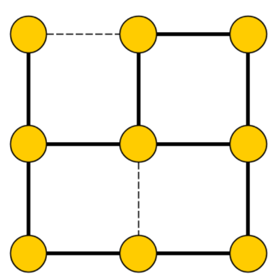

Figure 15: Grid graph with degree ≤ 3.

5.2 Simulating a traversal of a grid graph

Given an $w \times h$ grid graph G in which every node $n_i, 1 \leq i \leq m = w \times h$ has degree $deg(n_i) \leq 3$, we build labeled board B of size $138w \times 138h$. For every node n at coordinates (x, y) $x \in [0..w-1], y \in [0..h-1]$ in the grid G, we put a gadget at coordinate $(138x, 138y)$ in B; in particular we use:

- a **T-gadget** if $deg(n) = 3$ (oriented according to the three edges incident to n);
- a **link-gadget** if $deg(n) = 2$ (oriented according to the two edges incident to n);
- an **endpoint** if $deg(n) = 1$ (oriented according to the single edge incident to n);

For example, the grid graph of Figure 15 is transformed to a labeled board equivalent to the graph in Figure 16.

5.3 NP-completeness

Theorem 5.1. *Given a $w \times h$ grid graph G with degree ≤ 3, we can build in polynomial time a corresponding labeled board B with blocked cells replacing each node $n_i, i \in [1..m]$ with a gadget g_i of type:*

- *T-gadget if $deg(n) = 3$;*
- *link-gadget if $deg(n) = 2$;*
- *endpoint-gadget if $deg(n) = 1$.*

The graph G has an Hamiltonian path if and only if the board B is rollable.

Proof($\rightarrow$). Suppose that G has an Hamiltonian path $s_1 \rightarrow s_2 \rightarrow ... \rightarrow s_m$. We can put the die in a cell of the gadget g_1 corresponding to node s_1, and roll it through the gadgets $g_2, g_3, ..., g_m$ corresponding to the nodes $s_2, s_3, ..., s_m$ following the Hamiltonian path. Then we can roll back through the gadgets in reverse order $g_{m-1}, g_{m-2}, ..., g_1$ and complete the Hamiltonian cycle.

By construction the gadgets are fully rollable, and a die in gadget g_{i-1} can enter an adjacent gadget g_i, partially roll it and continue to gadget g_{i+1} through a pair of interface cells; then it can re-enter gadget g_i through the same pair of interface cells, complete it and return back to gadget g_{i-1}. Each node $s_i \in G$ with $deg(s_i) = 3$ is visited only once (only two edges are used), and this guarantees that the die uses only two pairs of interface cells and can fully roll the corresponding T-gadget g_i (it cannot use all the three pairs).

□

Proof($\leftarrow$). Suppose that the board B is rollable (i.e. has an Hamiltonian cycle that visits all its labeled cells). The die can roll from gadget g_i to the adjacent gadget g_j only through a pair of interface cells and must re-enter through the

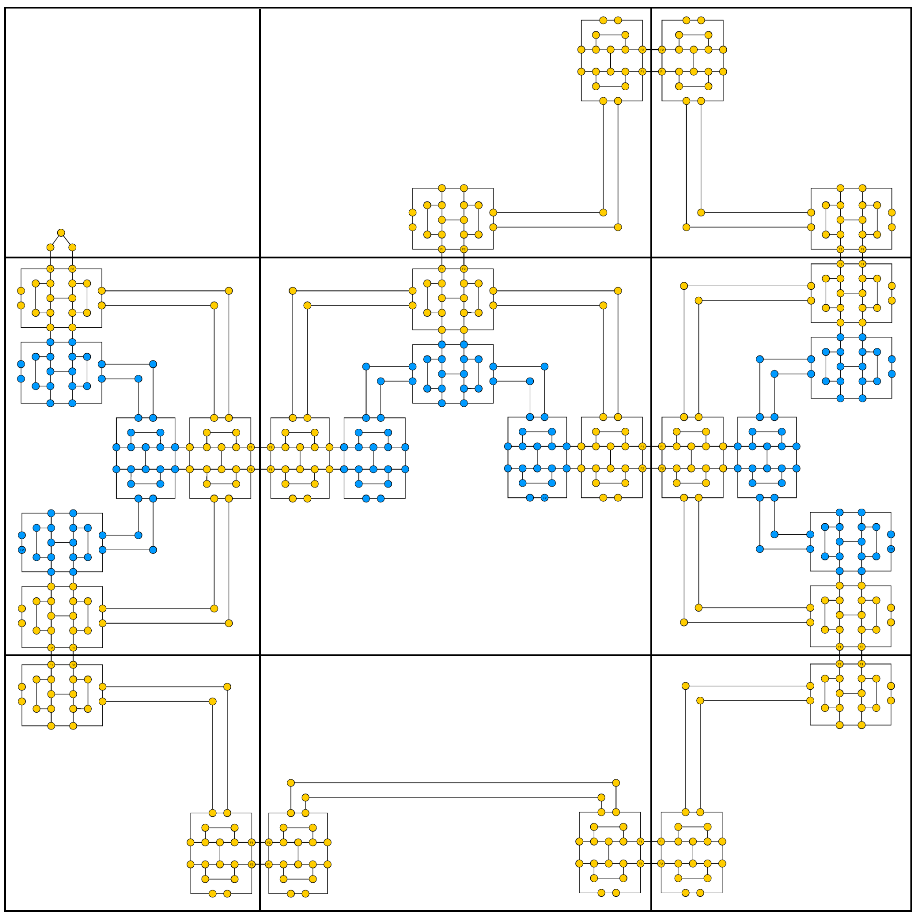

Figure 16: The graph representation of the reduction from the grid graph in Figure 15 to a labeled board.

same pair $((L_1, L_2) \leftrightarrow (R_1, R_2)$ or $(U_1, U_2) \leftrightarrow (D_1, D_2))$ otherwise there is no way way to completely roll g_i. The two pairs correspond to the edge (s_i, s_j) between the two adjacent nodes s_i and s_j of the graph G. Furthermore, in each gadget only two pairs of interface cells can be used (only one pair in the endpoint-gadgets). We can take the set P of edges in G corresponding to the pairs of interface cells crossed by the die in its cycle. We observe that:

- the die must visit all the gadgets, so for every node s_i of G there is at least one edge $(s_i, s_j) \in P$;
- a node s_i cannot be contained in three distinct edges in P otherwise in the corresponding gadget g_i three pairs of interface cells are crossed: $\{s_i, s_j\}, \{s_i, s_k\}, \{s_i, s_h\} \in P \Rightarrow s_j = s_k \vee s_j = s_h \vee s_k = s_h$;
- there is no way for the die to "jump": all the gadgets (and the corresponding nodes of G) are connected.

So the set P forms a path that visits all the nodes of G once.

□

From the NP-completeness of Hamiltonian path on grid graphs with degree ≤ 3 and without specified endpoints we can derive the NP-completeness of the die rolling problem over a labeled board with blocked cells.

Corollary 5.2. *Rolling a die over a labeled grid without free cells and with blocked cells is NP-complete.*

6 Conclusion

We proved that in the rolling cube puzzle Hamiltonian cycles are not unique even in fully labeled boards and deciding if a solution exists in labeled boards without free cells and with blocked cells is NP-complete.

The complexity of the puzzle in labeled boards without free cells and without blocked cells is still an open problem.

A Gadgets details

A.1 Rolls over the first basic gadget extension

If we suppose that the die can enter or exit only through cells $E_1, E_2, F_1, F_2, G_1, G_2$, then the rolls that fully visit the inner basic gadget are (see Figure 17):

```
1) OUT<->E1<->E2<->OUT
2) OUT<->F1<->F2<->OUT
3) OUT<->E1<->F1<->OUT<->F2<->E2<->OUT
4) OUT<->E1<->G2<->OUT<->G1<->E2<->OUT
5) OUT<->E1<->F1<->OUT<->G2<->E2<->OUT
```

```
6) OUT<->E1<->F2<->OUT<->G1<->E2<->OUT
7) OUT<->F2<->G1<->OUT
8) OUT<->G1<->G2<->OUT
9) OUT<->F1<->G2<->OUT
```

If we add the cells H_1, H_2 on the right border, we get nine additional rolls.

A.2 Rolls over the choice gadget

Three rolls over the choice gadget are shown in Figure 18. Other rolls can be obtained by mirroring the paths horizontally and vertically, or adding the cells H_1, H_2 on the right border. In all cases, the pairs F_1, G_1 and F_2, G_2 act like a switch as they do in the basic gadget.

A.3 T-gadgets details

The upward T-gadget is shown in Figure 19.

The rolls over a T-gadget have been tested using **Mathematica 8**. The T-gadget of Figure 10 has been converted to `graphml` format (`http://graphml.graphdrawing.org/`) after applying the elimination and cut algorithm. The file file can be downloaded from:

- `http://www.fractalmuse.org/rollingcubenpc/t-gadget.graphml`

The *id* of the nodes are `n<x>_<y>`, where `x` is the horizontal coordinate [0..137] in the grid, and `y` is the vertical coordinate [0..137] except for the interface cells which have *id*: `L1, L2, U1, U2, D1, D2`. The nodes are labeled with the corresponding number [1..6] except for the *interface* cells which are labeled: `L1, L2, U1, U2, D1, D2`. The T-gadget can be displayed correctly with the free graph editor `yEd` (`http://www.yworks.com/en/products_yed_about.html`).

For example to test that the T-gadget can be fully rolled if the die enters from L_1 and exits from L_2, we can load the `.graphml` file in Mathematica, add a node OUT_1, add two undirected edges $L_1 \leftrightarrow OUT_1$ and $OUT_1 \leftrightarrow L_2$ and use the `FindHamiltonianCycle` function:

```
g = Import["t-gadget.graphml", "Graph"];
g = VertexAdd[g, "OUT1"];
FindHamiltonianCycle[EdgeAdd[g, {"L1" <-> "OUT1", "OUT1" <-> "L2"}]]
Answer: {edgelist}
```

The possible valid rolls found are:

```
L1<->OUT1, OUT1<->L2: YES
U1<->OUT1, OUT1<->U2: YES
R1<->OUT1, OUT1<->R2: YES
L1<->OUT1, OUT1<->L2, U1<->OUT2, OUT2<->U2: YES
```

```
L1<->OUT1, OUT1<->L2, R1<->OUT2, OUT2<->R2:  YES
U1<->OUT1, OUT1<->U2, R1<->OUT2, OUT2<->R2:  YES
L1<->OUT1, OUT1<->U1, L2<->OUT2, OUT2<->U2:  YES
L1<->OUT1, OUT1<->R1, L2<->OUT2, OUT2<->R2:  YES
R1<->OUT1, OUT1<->U1, R2<->OUT2, OUT2<->U2:  YES
```

So, a T-gadget can be fully rolled only if the die uses at most two pairs of interface cells. An example is shown in Figure 20.

A.4 Link-gadgets details

The left/right link-gadget is shown in Figure 21. The rolls that fully visit the left/right link-gadget are:

```
1) OUT<->L1<->L2<->OUT
2) OUT<->L1<->R1<->OUT<->R2<->L2<->OUT
3) OUT<->R1<->R2<->OUT
```

The first two rolls are shown in Figure 22. The possible rolls over the other link-gadgets are similar.

A.5 Endpoint-gadgets details

The four endpoints-gadgets are shown in Figure 21.

References

[1] Robert Abbott and Bob Abbott. *Supermazes: Mind Twisters for Puzzle Buffs, Game Nuts, and Other Smart People.* Prima Pub., 1997.

[2] Kevin Buchin, Maike Buchin, Erik D. Demaine, Martin L. Demaine, Dania El-Khechen, Sándor Fekete, Christian Knauer, André Schulz, and Perouz Taslakian. On rolling cube puzzles. In *Proc. 19th Canadian Conference on Computational Geometry (CCCG)*, pages 141 – 144, 2007. full version in electronic proceedings, 30 pages.

[3] Martin Gardner. *Martin Gardner's Sixth Book of Mathematical Diversions from Scientific American.* University of Chicago Press, 1971.

[4] Martin Gardner. *Mathematical carnival: from penny puzzles, card shuffles and tricks of lightning calculators to roller coaster rides into the fourth dimension.* Knopf, 1975.

[5] John Harris. Single vacancy rolling cube problems. *J. Recreational Mathematics*, 7(3):220–223, 1974.

[6] Christos H Papadimitriou and Umesh V Vazirani. On two geometric problems related to the travelling salesman problem. *Journal of Algorithms*, 5(2):231 – 246, 1984.

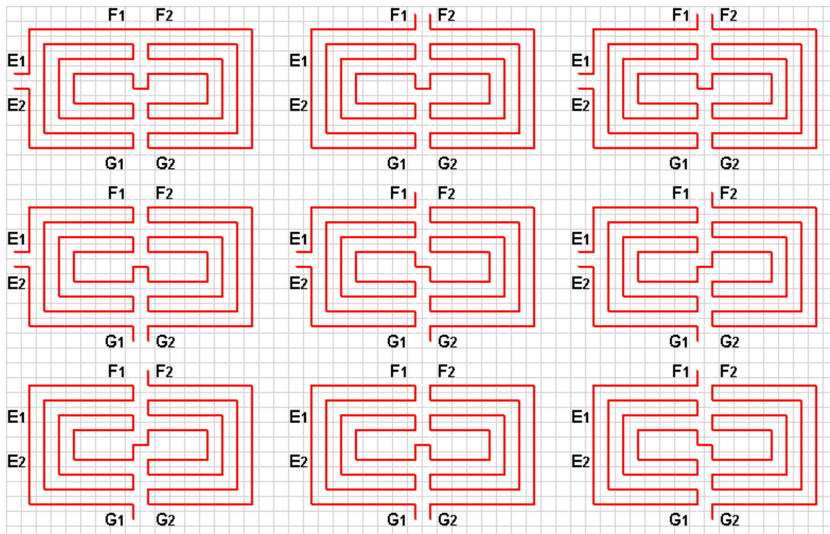

Figure 17: Rolls over basic gadget; cells F_1, G_1 and F_2, G_2 act like a switch.

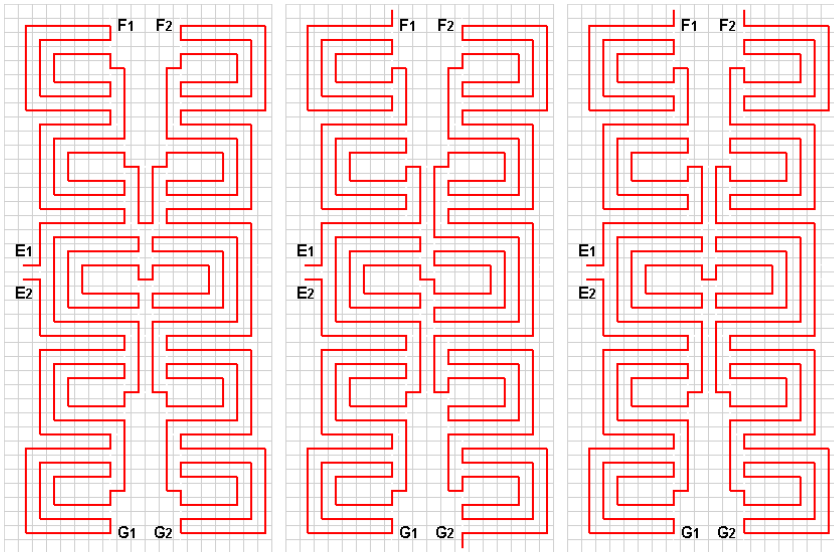

Figure 18: Rolls over the choice gadget; other rolls can be obtained by mirroring the paths horizontally and vertically.

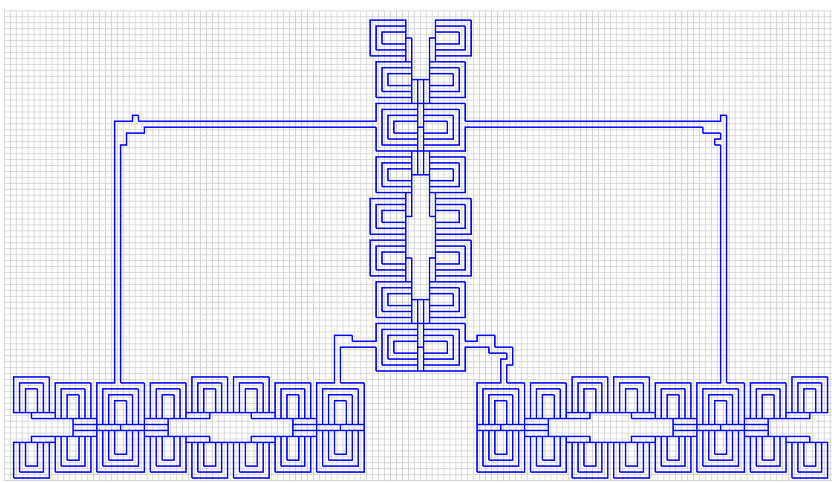

Figure 19: Upward T-gadget that fits in a 138×138 labeled board with blocked cells.

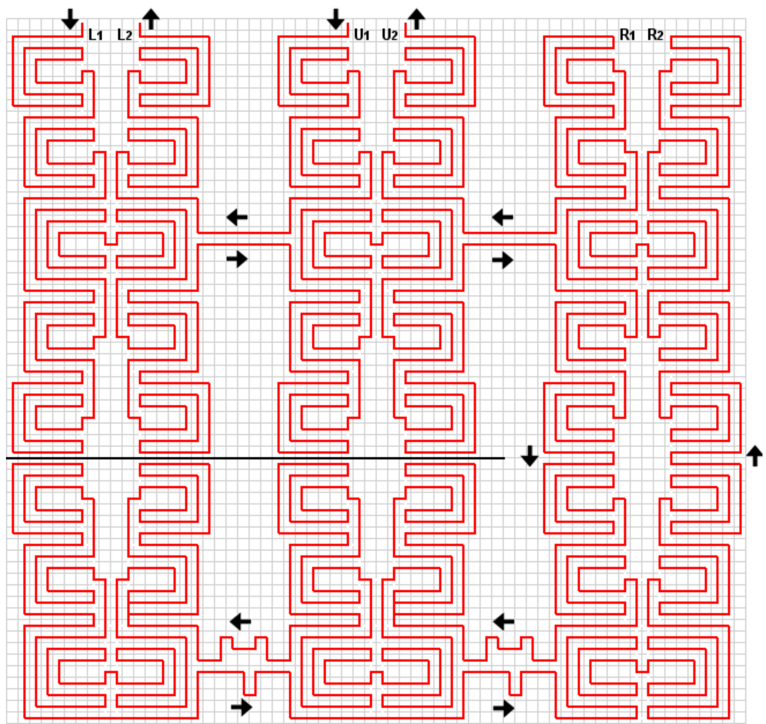

Figure 20: An example of a complete roll over the T-gadget `OUT1->L1->U2->OUT->U1->L2->OUT`; the interface cells R_1 and R_2 cannot be crossed.

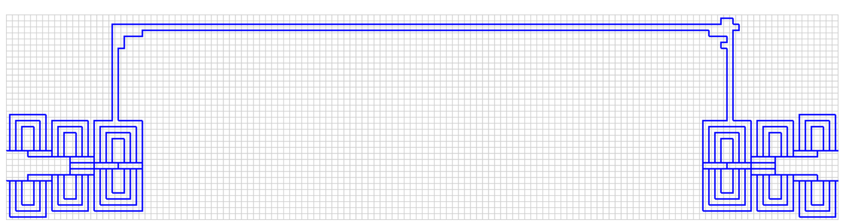

Figure 21: Left/right link-gadget that fits in a 138×138 labeled board with blocked cells.

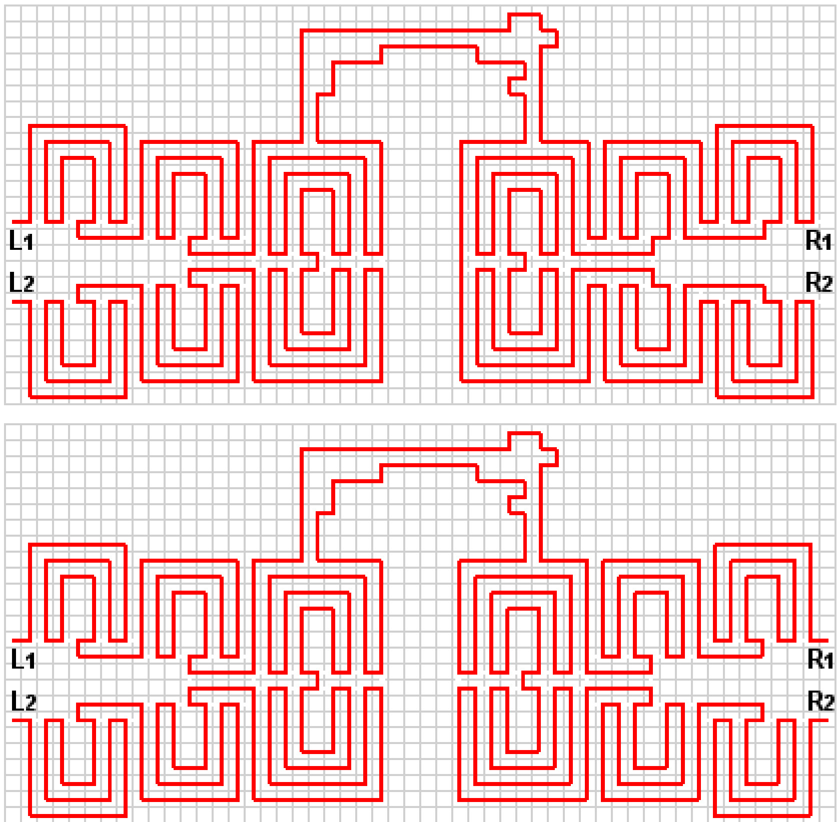

Figure 22: Two rolls over the left/right link-gadget.

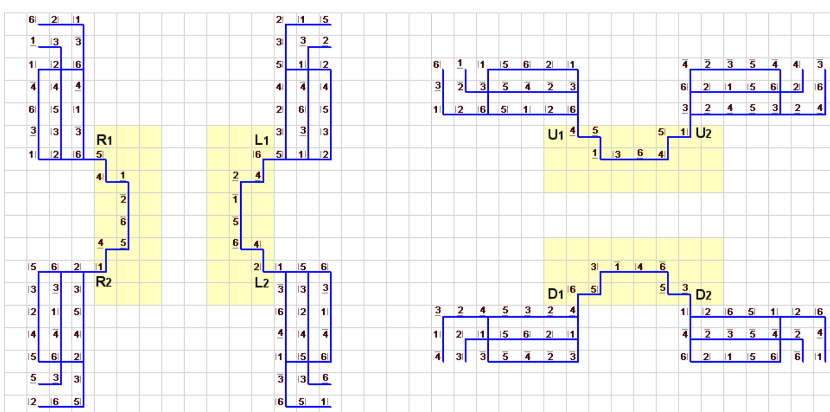

Figure 23: The four endpoints-gadgets (yellow area) and how they are connected to the T-gadgets or link-gadgets.