

The Design of a Reconfigurable Maze

Craig S. Kaplan
Cheriton School of Computer Science
University of Waterloo
csk@uwaterloo.ca

Abstract

A reconfigurable maze is a small set of modular components, called tiles, that can be assembled into a variety of distinct mazes. Although the idea of a reconfigurable maze is simple enough to articulate, there are some difficult mathematical challenges to overcome in order to make them a reality. I introduce the concept of a reconfigurable maze, and work through the design of a first prototype based on a 2×2 grid of square tiles. The prototype is a partial solution to the original problem, leaving open many opportunities for future innovation.

1 Introduction

The construction of a maze is a process that blends ideas from art, design, mathematics, and computation. A professional maze designer may draw the walls of a maze based purely on intuition, but that intuition must still encompass an implicit understanding of the topological properties of the result (for example, the requirement that all the passages be connected).

Conversely, many mathematical and algorithmic tools may be brought to bear on the problem of maze design. At the simplest level, any spanning tree algorithm can be used to break walls in a square grid until a maze emerges in which every two cells are connected by exactly one path. Such mazes are referred to as *perfect*. By using Kruskal's algorithm, with the attendant union-find set data structure [1, Section 23.2], maze construction becomes an instructive programming exercise in an introductory algorithms course. Contemporary research in computer graphics and mathematical art offers many techniques for automatic or semi-automatic creation of mazes based on images [4, 6, 7], mazes designed to be especially challenging as puzzles [8], and labyrinths with guaranteed symmetry properties [5].

In 2013 I was contacted by a manufacturer attempting to develop a new maze-based toy. The toy was to consist of a set of moulded plastic square "tiles", each one decorated with an extruded maze fragment. To be clear, there would be a small number of plastic tiles, each of which would be inscribed with a substantial portion of a maze; the reader may wish to look ahead to the examples in Figure 6 for a demonstration. This design stands in contrast to the work of Bosch et al. [5, Section 2.7], in which the authors exhibit a large set of small 3D printed square tiles that can be rearranged to produce all possible labyrinths.

When our tiles are laid out in a 2×2 grid, the fragments would snap together to create a single large maze. The novelty lay in the fact that the tiles could be placed in any configuration, and rotated by any multiple of 90° , and the result would still be a valid puzzle. The tiles could be very simple fragments based on grids, but they would conspire to produce a large number of distinct mazes. I was offered the challenge of designing what we came to call a *reconfigurable maze*. (I later learned of the "3-D Magic Maze Cube", a plastic toy in which a ball bearing can roll around mazes inscribed in the faces of a cube; in that toy, the faces can be rearranged inside a clear plastic box.)

In this paper, I describe the mathematical and algorithmic ideas that I applied to the development of a prototype reconfigurable maze. I begin by articulating the requirements for a successful design. Ultimately some of these requirements give way under the burden of the mathematical constraints implied by reconfigurability. The result is partially successful; to view the toy as a puzzle, we must stretch the traditional rules of mazes a bit.

2 Goals and requirements

Although the basic idea of a reconfigurable maze is straightforward enough, it is important to state our design goals first. At the outset, we do not know whether such a puzzle can exist at all. A clear statement of goals will help to make the design problem more concrete, and offer some sense of which goals we might choose to renege on if the problem turns out to be overconstrained. After some consideration of the problem, I arrived at a short list of requirements:

1. Ideally every configuration of tiles should produce a distinct perfect maze: for every two points in the maze, there should be a unique path connecting them. We can reduce the goal of perfection to two requirements: first, that the maze not have any cycles; and second, that no part of the maze form an isolated set of passages cut off from the rest.
2. Once a maze is assembled, there should be a clear objective. Ideally, the assembled maze would have a single obvious entrance, and a single obvious exit. Failing that, it should at least be possible to state the objective in a way that can be understood by the intended audience of children.
3. Solutions should not be “short”—we should never find that the only legal solution to a given configuration enters one tile and immediately leaves from the same tile. We would like to maximize the importance of the tiles by forcing any valid solution to use more than one of them.
4. When in doubt we should prioritize reconfigurability over other goals. A reconfigurable maze will be less novel if it can assume only two configurations. With four tiles, we would prefer to have access to hundreds or even thousands of configurations. Furthermore, if there are restrictions that make certain configurations legal and others illegal, the distinction should be obvious; that is, under normal use a child should not unwittingly be able to place the tiles into an illegal configuration.

In the rest of this paper, we will see that we must weaken many of these goals because of the combinatorial realities of the design problem. However, we will eventually arrive at a prototype that can function as a compelling puzzle.

3 Tile entrances

In order to produce a single large maze from a set of tiles, the overall configuration must include some passages between neighbouring tiles. We therefore expect that each tile will have some number of openings along its boundary. We refer to these openings as *entrances*, though of course they also function as exits! In our set of four tiles, we must locate entrances on tile boundaries so that they can be brought into coincidence with each other when the tiles are placed in a 2×2 grid.

In a first attempt at reconfigurable maze design, the most natural arrangement of entrances is for each square tile to have four, one at the centre of each of its four sides. With entrances at edge centres, neighbouring tiles will always be interconnected regardless of their orientations (including reflections). If we were to introduce an off-centre entrance, we would either limit the possible orientations of neighbouring tiles to force their entrances to coincide, or else we would need to add the entrance’s symmetric twin across the bisector of the tile’s edge, greatly increasing the total number of entrances under consideration. These possibilities are probably still worth exploring, and I discuss them further in Section 6.

With an entrance on every side of every tile, a 2×2 grid of tiles will inevitably have eight entrances around its border. This proliferation of entrances interferes with Goal 2: where are the start and end of the puzzle? We must offer the solver an explicit objective, rather than having it be obvious from the structure of the maze.

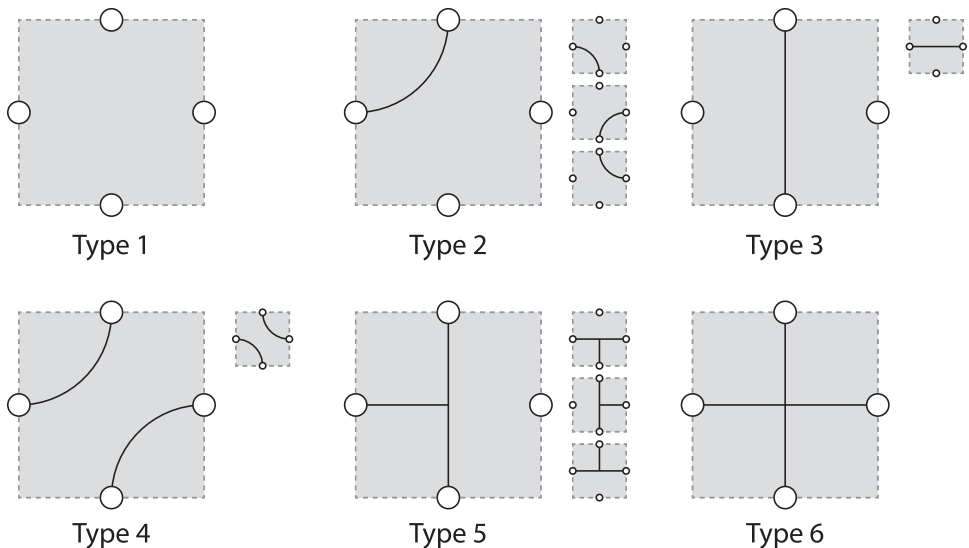


Figure 1 : The six routing types for a maze tile with four entrances. A single large representative is given for each type, together with the rest of its family of rotational equivalents.

4 Tile routing

At the outset, it is clear that the precise contents of each tile do not matter for the analysis of reconfigurability. More important is the abstract “topology” of the tile—the manner in which its entrances are interconnected. Viewed as a binary relation on pairs of entrances, interconnectedness is certainly an equivalence relation. It may therefore be regarded as a means of partitioning the tile’s entrances into classes, where two entrances are connected by a path if and only if they belong to the same class.

Let a tile’s n entrances be labelled in order around its boundary as E_1, \dots, E_n . Although every valid maze routing is certainly a partition of the E_i , not every partition represents a valid routing, a consequence of the planarity of a maze. In particular, when $n = 4$ there is no way to route E_1 to E_3 and E_2 to E_4 without simply connecting all four entrances.

The concept of a *noncrossing partition* [3] precisely captures valid maze routings. If n points are spaced evenly around a circle, a noncrossing partition may be defined as one for which the convex hulls of the partitioned subsets of points are pairwise disjoint. We discover that of the 15 possible partitions of four entrances, only $\{E_1, E_3\}, \{E_2, E_4\}$ mentioned above must contain a crossing, leaving a set of 14 noncrossing partitions. (Interestingly, the number of noncrossing partitions of n is given by the n th Catalan number.) Many of these are equivalent under rotations or reflections. Because we will permit our tiles to be rotated freely, we group them by these equivalences into six types, as shown in Figure 1. These types are easily discovered through experimentation; but an understanding of their connection to noncrossing partitions could be useful future work, particularly in generalizations to tiles with more entrances.

For the purposes of maze design, I identify two important subsets of these six types. First, note that Types 4 and 6 are the only two for which every entrance is connected to at least one other entrance. In every other type, there is at least one disconnected entrance (such an entrance will not be blocked off altogether, but will lead to a dead end somewhere inside the tile). I will refer to Types 4 and 6 as the *complete* types. Second, it will be important to distinguish Types 2, 4, 5, and 6 as those for which there is at least one connection between two adjacent entrances. I refer to these as *corner* types because they include a path that skirts one corner of a tile.

With four tiles, each having fixed entrances at the centres of its four sides, we are left with the task of choosing




Figure 2: Examples of configurations that must be avoided. In any type multiset with four corner tiles, as in $\{2,4,5,6\}$ on the left, the tiles can be placed in a configuration that contains a cycle (shown in bold). In any multiset with two non-corner tiles, as in $\{3,4,5,6\}$ on the right, it is possible to create a set of isolated passages (in this case, containing the joined entrance marked in bold).

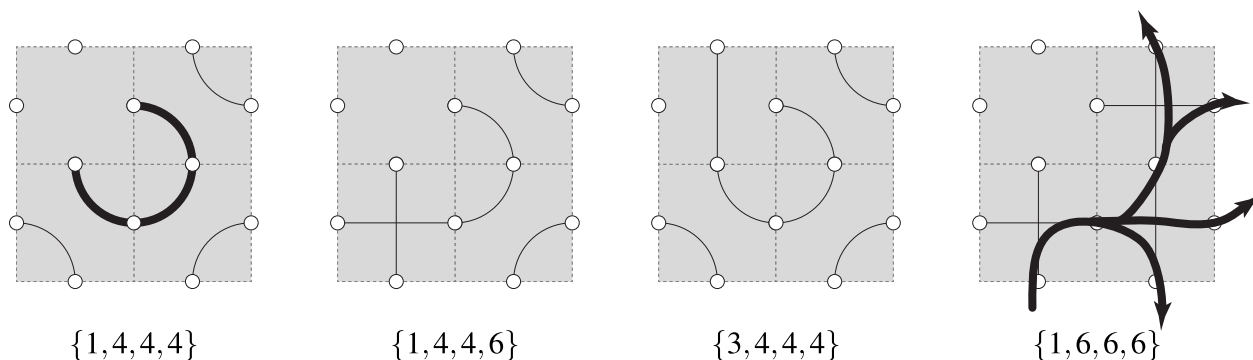


Figure 3: Examples of specific configurations for four sets of tile types that demonstrate why those sets might be undesirable. In $\{1,4,4,4\}$, tiles can be arranged to produce an isolated portion of the maze, highlighted in bold. In $\{1,4,4,6\}$ and $\{3,4,4,4\}$, there are arrangements with no solutions that cross tile boundaries. In $\{1,6,6,6\}$ (and $\{3,6,6,6\}$, not shown), the abundance of Type 6 tiles allows too many distinct solutions.

four types from the six above. The types can be chosen with replacement, and the choice should be unordered, so that the type sequences $(3, 1, 6, 3)$ and $(1, 3, 3, 6)$ should be considered equivalent. Our choice, then, amounts to selecting a 4-multiset combination, a gentle modification of the usual theory of combinations in which repetitions are permitted. Using standard notation for multisets, the number of k -multiset combinations from a set of size n is given by

$$\binom{n}{k} = \binom{k+n-1}{k},$$

yielding 126 ways to choose a multiset of four tile types from the full set of six.

We might immediately proceed to analyze each of these 126 possibilities to determine which ones are most effective as puzzles, but we can greatly reduce the number of candidates by virtue of two simple observations based on the constraints in Section 2.

First, note that if all four tiles are corner tiles, then there will always exist a configuration of those four tiles in which the distinguished corners are gathered at the centre of the 2×2 grid. In that case, the four corners will collectively define a cycle in the centre of the maze, a feature that we set out to avoid. An example is shown on the left in Figure 2. We can conclude that any choice of tile types must include at least one non-corner tile, i.e., a tile of Type 1 or 3.

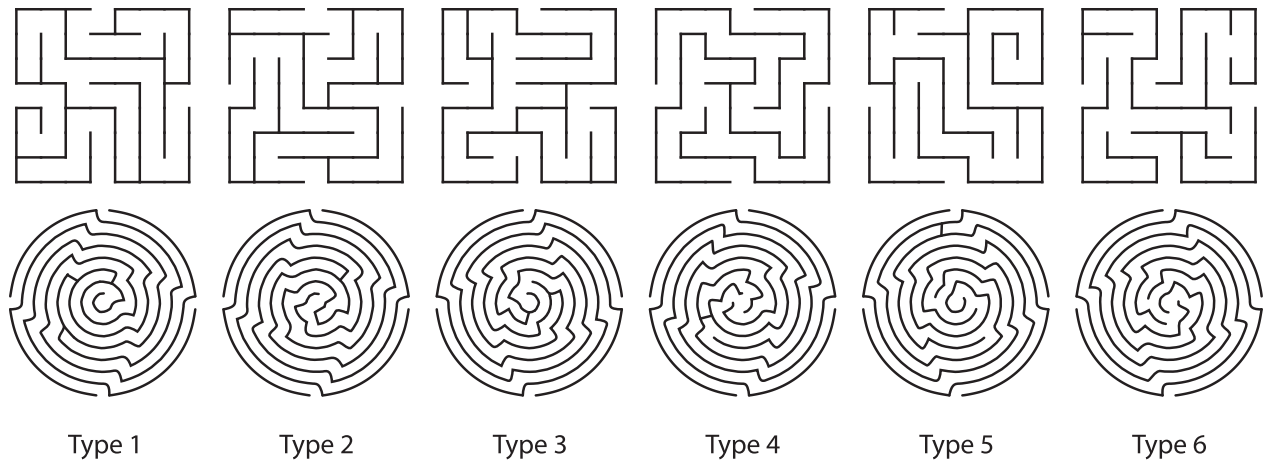


Figure 4: Examples of maze tiles that belong to each of the six tile types of Section 4. The top row shows manually constructed 7×7 grid maze tiles; the bottom row shows vortices.

Second, if our set includes two tiles that are not complete, then there will be a configuration in which two isolated entrances will be brought into coincidence, creating an undesirable pocket of passages isolated from the rest of the maze. See the diagram on the right in Figure 2. We can conclude that a satisfactory design cannot include more than one incomplete tile.

Given these two observations, the only remaining possibilities are multisets containing one tile of Type 1 or 3, and three tiles in a mix of Types 4 and 6. There are eight such multisets: $\{1, 4, 4, 4\}$, $\{1, 4, 4, 6\}$, $\{1, 4, 6, 6\}$, $\{1, 6, 6, 6\}$, $\{3, 4, 4, 4\}$, $\{3, 4, 4, 6\}$, $\{3, 4, 6, 6\}$, and $\{3, 6, 6, 6\}$. Some of these we rule out by consideration of individual configurations that we deem undesirable, as shown in Figure 3. Certainly $\{1, 4, 4, 4\}$ is unacceptable because it can produce a longer-range isolated pocket. In the cases of $\{1, 4, 4, 6\}$ and $\{3, 4, 4, 4\}$, we discover that there are configurations in which the only solutions are short: all paths through the maze enter through one exposed edge of a tile and exit immediately through the other exposed edge of the same tile. I further rule out $\{1, 6, 6, 6\}$ and $\{3, 6, 6, 6\}$ on the grounds that they are too “easy”: there are always too many long routes through the resulting maze. We are left with three type multisets that offer a balanced amount of complexity: $\{1, 4, 6, 6\}$, $\{3, 4, 4, 6\}$ and $\{3, 4, 6, 6\}$.

5 Puzzle construction

Once the routing types have been decided, all that remains is to construct tiles that exhibit those routings. There are several ways to accomplish this task. If the tiles are to be based on $n \times n$ square grids (a simple, popular style of maze), then they can easily be constructed by hand. A bit of practice may be required to create non-trivial paths within tiles that nevertheless achieve the desired routings. Note also that n must be odd, so that each tile edge can include an entrance in its centre that passes into a whole grid cell. The top row of Figure 4 shows examples of grid maze tiles for each of the six routings of Section 4. Of course, this process can be automated using a generalization of Kruskal’s algorithm that permits non-interacting routes between multiple entrances [7].

There are many other ways to construct maze-like tiles that connect four entrances, opening up a wide range of aesthetic possibilities. One possible approach to explore is the use of vortices [8]. A vortex is a standard device in maze construction that effectively obfuscates the relationships between its entrances. Examples of vortices for each of the tile types appear in the bottom row of Figure 4.

In the original design project that launched this investigation, the goal was to create tiles based on simple square

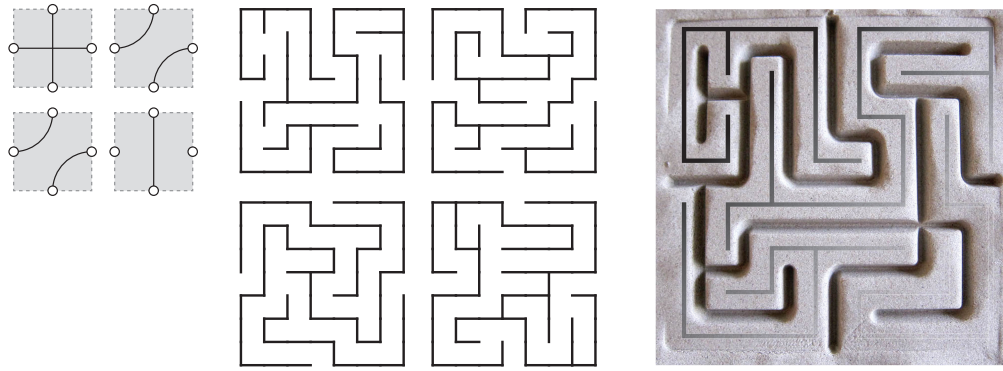


Figure 5: *The final set of tiles created as a first reconfigurable maze design. The four routing diagrams on the left reveal the type multiset to be $\{3, 4, 4, 6\}$. The four central maze fragments show the final tile designs. The photograph on the right shows the impression of a plastic model of the top-left tile in sand. The plastic part renders thin recessed passages and thick walls. A faded copy of the underlying tile is superimposed to show the equivalence of the structure.*




Figure 6: *Example configurations of the four tiles shown in Figure 5. The first three examples use only rotations of the original tiles; in the final example, two of the tiles are also reflected. In every case, the challenge is to find a path through the maze that enters through one tile and exits through some other tile.*

grid mazes. Through discussion with the manufacturer who commissioned the tiles, we arrived at a set of four tiles with type multiset $\{3, 4, 4, 6\}$, as shown in Figure 5. These were then manufactured in injection-moulded plastic, as a children’s toy. The plastic tiles could be pressed into sand, leaving a composite maze as an impression to be traced with a finger. After experimenting with moulds, the manufacturer settled on “inverting” the usual structure of a maze, rendering the maze’s passages as the thin protruding elements of the toy, rather than the maze’s walls. The impression in sand is still easily interpreted as a maze, as shown in the photograph on the right of Figure 5.

Of course, the the mazes produced from these tiles will be ambiguous as puzzles, violating Goal 2: unlike ordinary mazes, each configuration will have a total of eight entrances around its border. Based on the way in which the original design goals influenced the available choices of tile types, it now becomes necessary to offer instructions on solving these puzzles. In this puzzle, the solver should seek a path from an entrance on one tile to an entrance of a different tile (thereby satisfying Goal 3). The articulation of clear instructions is complicated by the fact that the tile divisions may be unclear after they are assembled together, and the fact that the puzzle is intended for children. Note also that some configurations will still have isolated sets of passages, for example when a Type 4 tile has one of its corner-skirting entrance pairs pointing to an outer corner of the maze. I accept this deficiency as a necessary evil, having at least suppressed all isolated passages in the *interior* of the maze.

How reconfigurable is the resulting puzzle? The maze fragment used for each tile has no internal symmetries,

meaning that every orientation of the tile is distinct, regardless of its routing type. I will first assume that tiles cannot be reflected (i.e., flipped over); perhaps each tile must have a handle on one side so that children can manipulate them easily. In that case, there are $4! = 24$ arrangements of the tiles in a 2×2 grid, and four orientations for each tile, for a total of $4! \cdot 4^4 = 6144$ mazes. We might argue that each resulting *maze* occurs in four orientations in that enumeration, and that there are therefore only 1536 distinct mazes. I consider both answers acceptable, depending on context; a child might not immediately recognize the rotational equivalence. Similarly, if we allow tiles to be flipped over, then there are $4! \cdot 8^4 = 98304$ configurations, of which 12288 are distinct under rotations and reflections of the whole maze.

6 Conclusions

Beginning with the initial challenge of creating a reconfigurable maze, I developed a suite of mathematical ideas needed to specify and understand the problem (drawn mostly from combinatorics). I produced an initial design of four square tiles that can be placed into a number of maze-like configurations. These configurations do not yield maze puzzles in the usual sense; every puzzle has eight entrances, meaning that there are no obvious start or end points. With the addition of instructions requiring the solver to start at one tile and end at another, we arrive at a puzzle with relatively few solutions in every configuration.

There are many possible generalizations worth considering of the simple 2×2 grid of square tiles. We might consider other tile shapes, such as rectangles, equilateral triangles, regular hexagons, and so on. Tile shapes with different symmetries imply different amount of reconfigurability; also, for non-rectangular tiles we may need to use a maze construction technique that does not rely on a simple square grid.

With any tile shape, we might experiment with different numbers and locations of entrances on tile edges. For example, perhaps we could achieve more balanced complexity by working with square tiles with *two* entrances on each edge. Certainly we would then have access to a much richer vocabulary of routings, but perhaps that richness would be counterbalanced by the greater number of ways that passages in the maze might be unacceptable, as in the analysis of Section 4. It would be valuable to develop some simple heuristics that rule out large classes of type multisets over these eight entrances, or perhaps to switch to more of a brute-force computer search.

In the limit, perhaps tiles could be developed that are completely open on every edge. I find this idea to be very appealing aesthetically, but I expect that the restrictions on the way those open passages might interconnect would make it impossible to construct a non-trivial reconfigurable maze.

It would also be interesting to investigate generalizations to larger numbers of tiles. Unfortunately, even a 3×3 grid of square tiles is impossible given the restrictions of Section 4: there is no 9-multiset of the six tile types that avoids cycles and isolated passages. Perhaps with a larger grid, we can restore reconfigurability by limiting legal orientations of individual tiles, for example, permitting only translations. This option affords an interesting opportunity for toy design: a sliding 15-puzzle printed with maze fragments, producing up to $15!$ distinct mazes under translation alone (assuming the empty cell must remain in the lower-right corner of the puzzle). At the mathematical end of this continuum, we might examine the limit of reconfigurable mazes in which each tile is a single cell in a grid maze, as in the work of Bosch et al. [5]. In that case, we would need to add considerable intelligence to tiles to force them to form legal mazes when assembled. We might look towards augmenting tiles with matching conditions in the style of Wang tiles [2, Chapter 10], though it is hard to see how to use Wang tiles to avoid long-range problems like cycles.

Acknowledgments

Thanks to Gary Fallowes of NewMetro Design for approaching me with the problem of reconfigurable maze design, and to Gary Greenfield for telling me about the 3-D Magic Maze Cube after seeing an early version of this work.

References

- [1] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, third edition, 2009.
- [2] Branko Grünbaum and G. C. Shephard. *Tilings and Patterns*. W. H. Freeman, 1987.
- [3] G. Kreweras. Sur les partitions non croisées d'un cycle. *Discrete Mathematics*, 1(4):333–350, 1972.
- [4] Yoshio Okamoto and Ryuhei Uehara. How to make a picturesque maze. In *Proceedings of the 21st Canadian Conference on Computational Geometry (CCCG2009)*, pages 137–140, 2009.
- [5] Mäneka Puligandla Robert Bosch, Sarah Fries and Karen Ressler. From path-segment tiles to loops and labyrinths. In George W. Hart and Reza Sarhangi, editors, *Proceedings of Bridges 2013: Mathematics, Music, Art, Architecture, Culture*, pages 119–126, Phoenix, Arizona, 2013. Tessellations Publishing. Available online at <http://archive.bridgesmathart.org/2013/bridges2013-119.html>.
- [6] Liang Wan, Xiaopei Liu, Tien-Tsin Wong, and Chi-Sing Leung. Evolving mazes from images. *IEEE Transactions on Visualization and Computer Graphics*, 16(2):287–297, March/April 2010.
- [7] Jie Xu and Craig S. Kaplan. Image-guided maze construction. *ACM Trans. Graph.*, 26(3):29, July 2007.
- [8] Jie Xu and Craig S. Kaplan. Vortex maze construction. *Journal of Mathematics and the Arts*, 1(1):7–20, March 2007.