

# ZHED is NP-complete

Sagnik Saha\*

Erik D. Demaine†

## Abstract

We prove that the 2017 puzzle game ZHED is NP-complete, even with just tiles numbered 1. Such a puzzle is defined by a set of unit-square tiles in a square grid, and a target square of the grid. A move consists of selecting a previously unselected tile and then filling the next unfilled square in a chosen direction from that tile (similar to Tipover and Cross Purposes). We prove the NP-completeness of deciding whether the target square can be filled, by a reduction from rectilinear planar monotone 3SAT.

## 1 Introduction

ZHED [Gro17] is a puzzle game available for Android, iOS, Switch, and Steam, first released in 2017. An instance of this puzzle is played on a  $n \times n$  square grid; refer to Figure 1 (left). One of the squares is designated as the *target*. Several of the other squares are *filled* by *tiles*. Each initial tile has an integer *number* between 1 and  $n - 1$  written on it.

In each move, the player selects one of the remaining numbered tiles and a direction (up, down, left, or right); refer to Figure 1. If the tile was numbered  $k$ , the move replaces the tile with a blank tile (keeping the square filled, but removing the number) and fills the  $k$  closest unfilled squares in the specified direction with blank tiles. The objective of the game is to fill the target square.




Figure 1: A typical move in ZHED. Tiles are drawn in dark blue. The target square has a small square in the middle.

The number of moves in such a puzzle is bounded by the number of numbered tiles, so the puzzle is in

NP. In this paper, we prove that the puzzle is in fact NP-complete, even when all tiles are numbered 1.

The mechanic of ZHED moves is similar to the idea of a tower of specified height that falls over in a specified direction. This “falling tower” mechanic is present in the puzzle board game Tipover, which is NP-complete [Hea06], and the 2-player board game Cross Purposes, which is PSPACE-complete [Hea05]. The key distinction is that both Tipover and Cross Purposes prevent a tower from falling on top of occupied/filled squares.

## 2 Rectilinear Planar Monotone 3SAT

Our NP-hardness reduction is from the known NP-complete problem Rectilinear Planar Monotone 3SAT (henceforth called *RPM-3SAT*) [BK12]; refer to Figure 2. In this problem, we are given a Boolean formula  $\mathcal{F}$  over  $n$  Boolean variables  $x_1, x_2, \dots, x_n$ . Formula  $\mathcal{F}$  is in conjunctive normal form, so the formula  $\mathcal{F}$  is the logical AND ( $\wedge$ ) of  $m$  *clauses*. Each clause in  $\mathcal{F}$  is a logical OR ( $\vee$ ) of at most three *literals*, and the literals in each clause are either all positive, consisting of non-negated variables of the form  $x_i$ , or all negative, consisting of negated variables of the form  $\bar{x}_i$ . (The last is the “monotonicity” property.) We are also given a planar embedding of the variable–clause incidence graph connecting each variable to the clauses containing them, which satisfies the following five conditions:

1. All variables and clauses are grid-aligned rectangles of height 1.
2. All of the variables lie along a single horizontal line.
3. All edges lie along vertical lines.
4. All positive clauses lie above the line of the variables.
5. All negative clauses lie below the line of the variables.

The goal in the RPM-3SAT problem is to determine whether there exists a value assignment to the  $n$  Boolean variables that satisfies  $\mathcal{F}$ . It is known to be NP-complete [BK12].

## 3 Reduction: Puzzle Construction

We now show how to build a ZHED puzzle that represents an arbitrary instance of RPM-3SAT, in such a

\*Work done while at MIT.

†MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar St., Cambridge, MA 02139, USA, edemaine@mit.edu

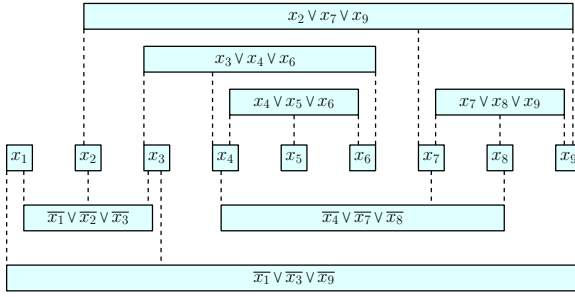


Figure 2: A sample instance of Rectilinear Planar Monotone 3SAT (RPM-3SAT). This planar embedding represents the conjunctive-normal-form formula  $(x_2 \vee x_7 \vee x_9) \wedge (x_3 \vee x_4 \vee x_6) \wedge (x_4 \vee x_5 \vee x_6) \wedge (x_7 \vee x_8 \vee x_9) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3) \wedge (\bar{x}_4 \vee \bar{x}_7 \vee \bar{x}_8) \wedge (\bar{x}_1 \vee \bar{x}_3 \vee \bar{x}_9)$ .

way that the puzzle we construct will be solvable if and only if the original problem is satisfiable (as proved in Section 4). We start with a description of the various gadgets in our construction, and then describe how they are combined together. In our construction, we only need to use tiles with the number 1, so in the figures we omit the number on tiles.

### 3.1 Threshold Gadget: Wires, AND, OR

Our main workhorse is the *threshold gadget*, which consists of a line of tiles on alternate squares, as shown in Figure 3. The empty squares between the tiles each function as a *source*. The threshold gadget is parameterized by a nonnegative integer  $k$ , and we call the square at a distance of  $k + 1$  from the last tile in the gadget (on the same line in the direction the gadget is meant to be expanded) the *target*.




Figure 3: Threshold gadget. We assume that at most 5 of the source squares can be filled by outside tiles, and this gadget activates as long as at least 3 of those squares are filled before the tiles in this gadget are expanded rightwards.

**Threshold property.** This gadget has the property that the target square can be filled if and only if at least  $k$  of the source squares were already filled by tiles from other gadgets. The intended activation is to expand

each tile in the gadget toward the target, in decreasing order of distance. Without any sources already filled, this will reach only one square from the last tile in the gadget. If, however,  $j$  sources were already filled, then it will reach distance  $j + 1$ . Thus, the target will get filled if and only if  $j \geq k$ .

**Wire, turn, AND, OR.** This gadget is versatile, and forms the basis of several other gadgets. If we connect only one source to other gadgets and set  $k = 1$ , then we get simple *wire gadget* which propagate a signal (represented by a square being filled) from that source to the target. We can *turn* the wire by connecting the target of one wire at a source of another orthogonal wire. On the other hand, if  $m > 1$  sources are attached to other gadgets (overlapping each source with the target of the other gadget, representing  $m$  input signals), then we obtain an *OR gadget* by setting  $k = 1$ , and we obtain an *AND gadget* by setting  $k = m$ .

**Bounding region.** Because of the spacing between the tiles, a threshold gadget can only possibly fill squares on that line within a distance of  $b + 1$  from one end of the gadget, where  $b$  is the maximum number of sources that can be filled by other gadgets. Because each tile has the number 1, the gadget can also affect at most the two adjacent lines of squares. Therefore, the dotted rectangle in Figure 3 serves as a *bounding region* of the gadget. No sequence of moves in the overall puzzle can lead to a tile in the threshold gadget directly filling a square outside its bounding region.

**Chaining threshold gadgets.** As alluded to above, we can chain threshold gadgets by making the target of one a source of the next, as shown in Figure 4. This allows us to aggregate signals and turn wires, as chaining necessarily introduces a right-angle turn. During activation, we simply expand the first gadget’s tiles first, which (maybe) fills its target square, and then expand the tiles of the second gadget.

This interaction is one of only two ways different gadgets are allowed to interact with each other in our construction. The first gadget’s bounding region grows by 1 unit in the direction it propagates, because the second gadget may be expanded before it. The composite bounding region is simply the union of the two gadgets’ individual bounding regions.

### 3.2 Shift Gadget

To resolve parity issues in our construction, we sometimes need to shift threshold gadgets by 1 square. To do so, we simply add an extra tile at the beginning of the gadget, adjacent to the old starting tile, as in Figure 5. To activate this gadget, we expand all tiles toward the

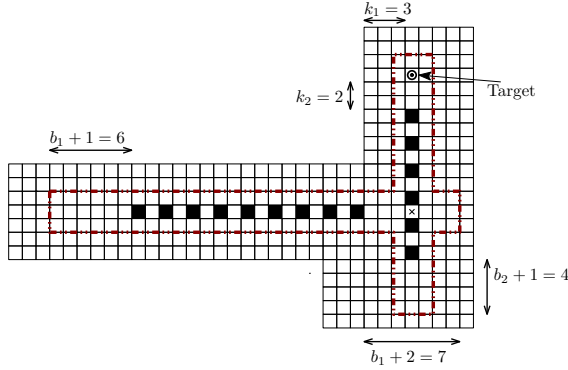


Figure 4: Chaining together two threshold gadgets (with parameters  $k_1$  and  $k_2$  respectively). The target of the horizontal threshold gadget (marked by  $\times$ ) acts as a source for the vertical threshold gadget.

target, in decreasing order by distance (as in the threshold gadget).




Figure 5: Shift gadget. The threshold gadget from Figure 3 is shifted by one square to the right.

**Bounding region.** As a result of the extra tile, the bounding region of the original threshold gadget expands by 1 square in the forward direction, and 2 squares in the backward direction (counting the new tile itself, as well as the extra square it can potentially fill). We move the target square one square forward, and everything else remains unchanged.

### 3.3 Variable Gadget

A *variable gadget* consists of a large even number  $L$  of tiles consecutive on a horizontal row; refer to Figure 6. These tiles are all meant to be expanded in the same direction, either left or right. In our construction, expanding all tiles left corresponds to setting the variable to *false*, and expanding all tiles to the right corresponds to setting the variable to *true*.

Activating a variable gadget as intended fills  $L$  squares in one direction. We attach vertical wires that propagate the filling of certain squares to the left or right of the tiles (to clause gadgets), thus “reading” the value of the variable. All such wires are at a distance




Figure 6: Variable gadget with  $L = 8$  tiles, along with four threshold gadgets “reading” the value of this variable. The two threshold gadgets right of the variable can get one of their source squares filled if the variable tiles are all expanded right, setting it to true. The two gadgets on the left can be activated when the variable is set to false.

between  $L/2 + 1$  and  $L$  from the variable gadget. This choice ensures that an unintended activation of the gadget, which expands some tiles left and other tiles right, cannot activate wires on both sides of the variable. For example, if  $x$  of tiles are expanded left and  $y$  are expanded right, then at least one of  $x$  and  $y$  will be less than  $L/2 + 1$ , so no wire in that direction will be activated.

We choose the value of  $L$  based on the number of vertical wires needed. In the rest of the construction described below, we specify the number of wires needed for each clause gadget, and the gaps needed between consecutive wire. Any value of  $L$  large enough so that all of those wires fit within  $L/2 - 1$  columns suffices.

**Bounding region.** The bounding region of the variable gadget is again only 3 rows high because all the tiles have number 1. There are less than  $L/2$  vertical wires on either side within the reach of the tiles in this gadget, so we obtain an upper bound on the bounding region of the variable gadget by supposing that all of those wires were already expanded before any tile in this gadget is expanded. Thus the bounding region of a variable gadget spans at most  $L + L/2 = 3L/2$  columns outside the tiles in each direction, as shown in Figure 6.

### 3.4 Clause Gadget

The threshold gadget enables us to build a *clause gadget*; refer to Figure 7. Because our 3SAT instance is monotone, each clause uses only positive or only negative literals. For positive clauses, we create a horizontal threshold gadget going rightwards above the variables; and for negative clauses, we create the threshold gadget below the variables. Then we connect the clause to corresponding variables (on the left side of the variable for positive clauses, and on the right side of the variable for negative clauses) via vertical “thick wires”. A *thick wire* is a group of  $g$  parallel wires connecting the same

gadgets, each separated 4 units apart from each other; thus, a single input actually advances the reach of the gadget by  $g$ . For the threshold gadget to function as an OR of the literals, we set its  $k$  to  $g$  (instead of 1). We refer to the combination of the thick wires and the horizontal OR gadget as the clause gadget. Thus the target square for the clause gadget is  $g+1$  distance to the right of the rightmost tile in the gadget. The thickness  $g$  is a parameter that may vary across different clauses and will be decided later.




Figure 7: Clause gadget for  $(\bar{x} \vee \bar{y})$ . We use thick wires, each made of 2 individual wires, to connect the clause with the variables. The vertical wire chained from the OR gadget (with the orange bounding region) is a clause propagator that uses the target square of this clause gadget as a source, and propagates the signal elsewhere.

**Bounding regions.** The clause gadget consists mainly of a threshold gadget which can have  $3g$  filled sources, so it already has a defined bounding region. By our choice of separation, the bounding regions of the individual wires comprising each thick wire do not overlap.

### 3.5 Crossover Gadget

We use a *crossover gadget* when we need a vertical threshold gadget to intersect the horizontal threshold gadget belonging to a clause. As shown in Figure 8, we simply position the two threshold gadgets so that the intersection square is blank (with four tiles surrounding that intersection square like a plus sign). Next, we increase the  $k$  of the vertical threshold gadget by 1. Essentially, we assume that the horizontal gadget will be activated before the vertical one, and we compensate for the extra filled source square for the vertical gadget by moving its target square one unit farther. If the player violates this assumption, the horizontal gadget will effectively have its  $k$  off by 1, an error we will tolerate using thick wires (see below).

**Bounding region.** The bounding region of the crossover gadget is easy to calculate because there are only four tiles involved. At the point of intersection, a  $3 \times 3$  square centered at the intersection point suffices. The bounding regions of both gadgets involved get longer in their long direction by 2 units (1 unit on each end) because of the potential extra source square that may be filled.




Figure 8: Crossover gadget. The horizontal threshold gadget is meant to be expanded first, so its target remains unchanged. The vertical threshold gadget is to be activated second, when the intersection square is already filled in. This pushes the target by one square.

### 3.6 Putting It All Together

To construct the overall puzzle, we transform the planar embedding of the provided RPM-3SAT instance; refer to Figure 9. All the variable gadgets go on the central horizontal row, spaced far enough apart that none of their bounding regions overlap. We then draw the clause gadgets according to the provided clause rectangles (above and below the central row for positive and negative clauses respectively). We space out the clauses more than the original drawing in order to leave room for the thick-wire connections between variables and clauses. Because we started with a planar monotone drawing, we obtain a noncrossing drawing.

Next we add vertical wires (henceforth referred to as *clause propagators*) to transport the output signals from all of the target squares of the positive clauses up. We place an AND gadget horizontally on a row high enough so that its bounding region doesn't collide with the bounding region of any clause gadget, and ensure that the target of all positive clause propagators coincide with sources of the AND gadget. We perform a symmetric construction below the negative clauses, resulting in an upper and a lower AND gadget.

Each clause propagator may intersect with the horizontal threshold gadgets from some of the other clauses (nested above this one). We use a crossover gadget for each such point of intersection. For each clause, we count the number  $x$  of crossovers in its OR gadget, and set the thickness  $g$  of the thick wires attached to the clause gadget to be strictly larger:  $g > x$ . This choice ensures that all the propagators intersecting a clause combined have less influence on the clause threshold gadget than a single thick wire, so no clause target can be satisfied by the propagator intersections alone.

We design the upper and lower AND gadgets to both go rightwards and have their targets on the same col-

umn, which is farther right than the bounding region of every other gadget. We combine these two signals using a vertical AND gadget, with a target square near the bottom right of the board. This square is the actual *target* for the entire puzzle in our construction.

The parity issue in the above construction is that the threshold gadget has a repeating unit of size 2, and hence we need to account for parity in lining everything up. This issue may cause problems when we want multiple threshold gadget targets to be on the same line for a chained gadget. Parity issues may also become significant for crossover gadgets, which require the intersection square to be empty. We fix these issues using shift gadgets. This approach allows us to adjust the rows of the clause gadgets and the columns of the clause propagators without parity restrictions. If we put all the clause gadgets on rows of the same parity, we can ensure that all the crossover gadgets line up. We use the same trick to ensure proper chaining of the clause propagators into the upper and lower AND gadgets, and the chaining of those two AND gadgets into the final vertical AND gadget. See Figure 9 for an example.

**Puzzle size.** During this construction, we space the gadgets out sufficiently so that no two bounding regions collide except as part of an intended interaction (through crossover gadgets or chaining threshold gadgets). The thickness  $g$  of each clause gadget’s thick wires are determined at this stage, and the size  $L$  of each variable gadget then depends on the total number of individual wires using it. We’ll show that because all of our gadgets have bounding-box dimensions of polynomial size, the overall board remains polynomial sized in the input parameters.

Note that there are a total of  $m$  clause propagators corresponding to the  $m$  clause gadgets. Therefore, any given clause gadget can have a maximum of  $m - 1$  crossovers in its OR gadget, and we don’t need any thick wire with thickness  $g$  exceeding  $m$ . Each variable can be connected to at most  $m$  clauses, and may therefore need  $m$  thick wires *reading* it. So we need to accommodate a maximum of  $m^2$  vertical wires, each spaced 4 units apart within  $L/2 - 1$  columns. This implies that the variables can all have sizes  $L \leq 8m^2 + 2$ .

We can draw the clause gadgets’ horizontal wires just long enough to intersect all the thick wires comprising them, and therefore the bounding region of a clause gadget should end no more than  $3g + 2 \leq 3m^2 + 2$  units to the right of the rightmost thick wire. To account for at most  $m$  crossover gadgets intersecting the clause, it can extend a further  $m$  units.

The horizontal dimension of the puzzle only needs to be big enough to fit the bounding regions of all the variable and clause gadgets, followed by the AND gadgets. The total length of the bounding region of a variable gadget

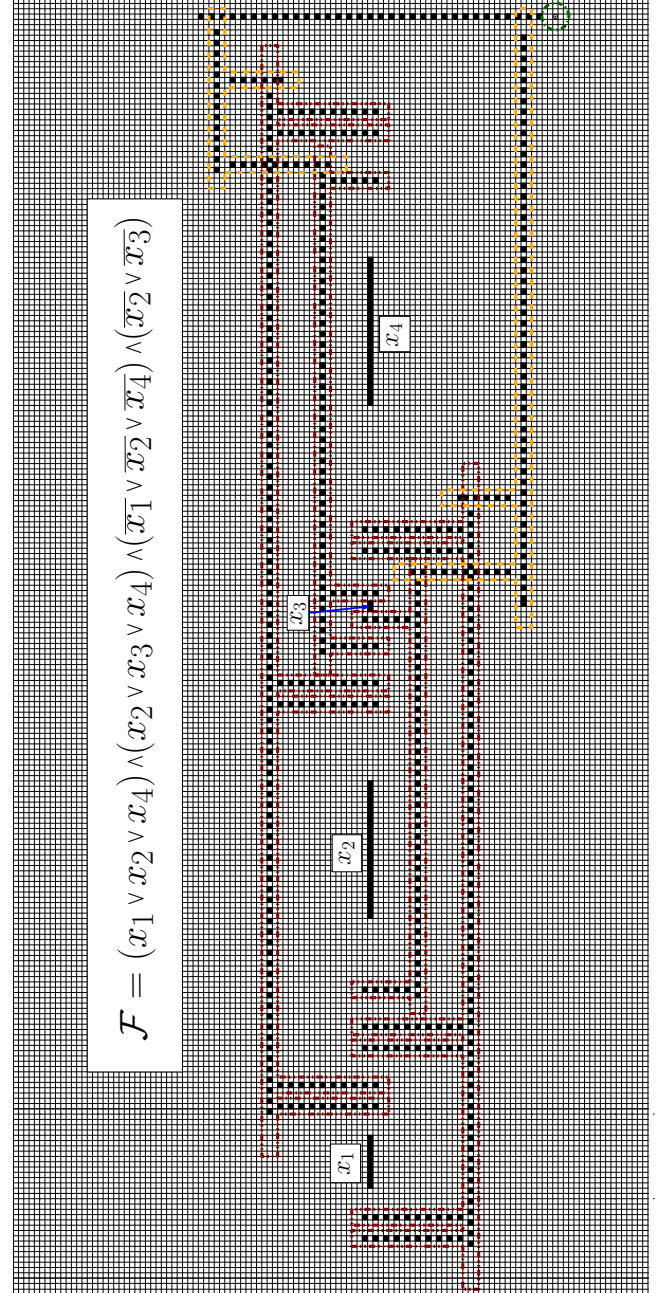


Figure 9: The entire puzzle board for a sample formula. The bounding regions for the clause gadgets are marked in red. The bounding regions of the clause propagators together with the upper and lower AND gadgets are colored orange. The final target square is at the lower-right of the board, marked by a green circle. Note that different variables have different sizes  $L$  depending on how many wires connect them to clause gadgets:  $v_3$  uses only 2 tiles while  $v_4$  uses 28. Similarly, the first and third clauses use thick wires while the other two use single wires to connect with variables. We use shift gadgets for the lower AND gadget and for a couple of clause propagators.

is  $7L$ , and the  $n$  variables account for  $7n(8m^2 + 2)$ . The clause gadgets can add another  $3m^2 + m + 2$ . The upper and lower AND gadgets have up to  $m$  clause propagators, and the vertical AND gadget therefore needs to be another  $m$  squares to the right of the rightmost clause gadget's bounding region. This shows that the horizontal dimension of our generated puzzle is  $\mathcal{O}(nm^2)$ .

The vertical size of the puzzle is much easier to upper bound. The bounding region of a basic horizontal threshold gadget is only 3 squares tall. Our generated puzzle grows vertically because we need to draw multiple clause gadgets on top of each other without their bounding regions colliding. All the vertical wires can only be as long as necessary to intersect with the variable gadgets and the horizontal OR gadgets of the clauses. Therefore, the total height of our puzzle is  $\mathcal{O}(m)$ .

## 4 Reduction Correctness

Finally we show that the construction of Section 3 is a valid reduction, i.e., an instance of RPM-3SAT is satisfiable if and only if the obtained ZHED puzzle is solvable.

### 4.1 Satisfiable Formula $\Rightarrow$ Puzzle Solution

Suppose there exists an assignment of the Boolean variables  $v_i$ s that makes the formula evaluate to true. Then we solve the constructed puzzle as follows.

First we expand each of the variable tiles. If  $v_i = \text{true}$  in the solution, we expand the corresponding  $L$  tiles rightward; otherwise, we expand the tiles leftward. Then we expand all the vertical wires connecting variables to clauses, toward the clauses. Next we expand the threshold gadgets in the clause wires rightward. Because each clause is satisfied in the formula, each of the clause gadgets will be activated by at least one thick wire, and hence each of the target squares for the clause gadgets will be filled after this stage.

Now we expand all of the clause propagator wires. These propagate the filled source corresponding to their clause's target square to the extreme upper and lower horizontal AND gadgets, filling all of their sources. Next, we expand the upper and lower AND gadgets rightward, which fills the two sources of the extreme-right vertical AND gadget. Finally, we expand the last vertical AND gadget downwards, which fills the puzzle target as desired.

### 4.2 Puzzle Solution $\Rightarrow$ Satisfiable Formula

Suppose we have a solution to the ZHED puzzle. We argue about the different gadgets in the reverse order from the intended activation sequence of Section 4.1, and use the fact that a square can only be filled by gadgets whose bounding regions contain it, to obtain a sat-

isfying Boolean assignment. We call a threshold gadget *successfully activated* if it fills its target square.

- The puzzle's target square only belongs to the bounding region of the extreme-right vertical AND gadget. Therefore that AND gadget was successfully activated, with at least  $k = 2$  of its sources filled by other gadgets.
- There are only 2 sources of the extreme-right vertical AND gadget that might be filled by other gadgets, so both of them must have been. This implies that both the extreme upper and lower horizontal AND gadgets were successfully activated.
- For the extreme horizontal AND gadgets to be successfully activated, all clause propagators had to have successfully activated their target squares.
- This implies that all clause gadgets successfully activated their target squares.
- If none of the thick wires activated for any particular clause gadget, then there will not be enough sources in that clause's OR gadget that belong to other gadget's bounding regions (specifically, crossovers) to reach its target square. Therefore, for every clause gadget, there is at least one thick wire with one or more component wires successfully activated.
- For any individual wire connecting a variable to a clause to be successfully activated, over half of the tiles in the corresponding variable gadget must be expanded in the direction of that clause (right if the clause is positive and left otherwise). We set each variable to false or true based on which direction (among left or right) the majority of the tiles of its corresponding variable gadget expanded; only clauses in this direction will be satisfied by the variable gadget. We set the value arbitrarily if there is a tie, in which case the variable did not satisfy any clauses.
- This assignment must be a satisfying assignment, because all clauses must have at least one variable satisfying them.

### 4.3 Main Theorem

We just established that determining whether a ZHED puzzle is solvable is NP-hard. It is also easy to prove membership in NP: a solution can be described by a sequence of moves, and if there are  $n$  tiles, each move takes  $\mathcal{O}(\log n)$  bits to describe. The number of moves cannot exceed the number  $n$  of tiles, so the total size of a proposed solution is polynomial in  $n$ . We can verify the solution in polynomial time by simply maintaining

the state of the board after each move and simulating the moves. Together, these assertions prove our main result:

**Theorem 1** *It is NP-complete to decide if a ZHED puzzle is solvable, even when all tile numbers are 1.*

## References

- [Hea05] Robert A. Hearn. “Amazons, Konane, and Cross Purposes are PSPACE-complete”. In: *Games of No Chance III: Proceedings of the BIRS Workshop on Combinatorial Games*. 2005, pp. 287–306. DOI: 10 . 1017 / CB09780511807251.015.
- [Hea06] Robert A. Hearn. “Tipover is NP-complete”. In: *The Mathematical Intelligencer* 28 (June 2006), pp. 10–14. DOI: 10.1007/BF02986878.
- [BK12] Mark de Berg and Amirali Khosravi. “Optimal Binary Space Partitions for Segments in the Plane”. In: *International Journal of Computational Geometry & Applications* 22.3 (2012), pp. 187–206. DOI: 10 . 1142 / S0218195912500045.
- [Gro17] Ground Control Games. *ZHED*. <https://playzhed.com/>. 2017.