

Strings-and-Coins and Nimstring are PSPACE-complete

Erik D. Demaine*

Yevhenii Diomidov*

Abstract

We prove that Strings-and-Coins — the combinatorial two-player game generalizing the dual of Dots-and-Boxes — is strongly PSPACE-complete on multigraphs. This result improves the best previous result, NP-hardness, argued in *Winning Ways*. Our result also applies to the Nimstring variant, where the winner is determined by normal play; indeed, one step in our reduction is the standard reduction (also from *Winning Ways*) from Nimstring to Strings-and-Coins.

In memoriam Elwyn Berlekamp (1940–2019)

1 Introduction

Elwyn Berlekamp loved Dots and Boxes. He wrote an entire book, *The Dots and Boxes Game: Sophisticated Child’s Play* [Ber00] devoted to explaining the mathematical underpinnings of the game, after they were first revealed in Berlekamp, Conway, and Guy’s classic book *Winning Ways* exploring many such combinatorial games [BCG03, ch. 16]. At book signings for both books,¹ and after talks he gave about these topics [AMS03], Elwyn routinely played simultaneous exhibitions of Dots and Boxes — him against dozens of players, in the style of Chess masters.

As many children will tell you, Dots-and-Boxes is a simple pencil-and-paper game taking place on an $m \times n$ grid of dots. Two players alternate drawing edges of the grid, with one special rule: when a player completes the fourth edge of one or two 1×1 boxes, that player gains one or two points, respectively, and *must* immediately draw another edge (a “free move” that is often a blessing and a curse). The game ends when all grid edges have been drawn; then the player with the most points wins. (Draws are possible on boards with an even number of squares.)

An equivalent way to think about Dots-and-Boxes is in the dual of the grid graph. Think of each 1×1 square as a dual vertex or *coin* worth one point, “tied down” by four incident *strings* or dual edges. Interior strings connect two coins, while boundary strings connect a coin to the *ground* (not worth any points). (Equivalently, boundary edges have only one endpoint.) Now players alternate cutting (removing) strings, and when a player *frees* one or two coins (removing the last strings attached to them), that player gains the corresponding number of points and must move again. The game ends when all strings have been cut; then the player with the most points wins.

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

¹The first author had the honor of playing such a game against Elwyn at a book signing on April 13, 2004, at Quantum Books in Cambridge, Massachusetts. Elwyn won.

Strings-and-Coins [BCG03, pp. 550–551], [Ber00, ch. 2] is the generalization of this game to arbitrary graphs, where vertices represent coins and edges represent strings which can connect up to two coins (the other endpoints being considered “ground”). *Nimstring* [BCG03, pp. 552–554], [Ber00, ch. 6] is the closely related game where we modify the win condition to *normal play*: the last player to make a move wins. Nimstring is known to be a special case of Strings-and-Coins, a fact we use in our results; see Lemma 2.2 below.

Related work. Dots-and-Boxes, Strings-and-Coins, and Nimstring are surprisingly intricate games with intricate strategy [BCG03, Ber00]. On the mathematical side, even $1 \times n$ Dots-and-Boxes is largely unsolved [GN02, CDDL15].

To formalize this difficulty, *Winning Ways* [BCG03] argued in 1984 that deciding the winner of a Strings-and-Coins position is *NP-hard* by a reduction from vertex-disjoint cycle packing. Around 2000, Eppstein [Epp] pointed out that this reduction can be adapted to apply to Dots-and-Boxes as well by restricting to maximum-degree-3 planar graphs and because exact path lengths are insignificant; see [DH09].

This work left some natural open problems, first explicit posed in 2001 [DH09]: are Dots-and-Boxes, Strings-and-Coins, and Nimstring NP-complete or do they extend into a harder complexity class? Being bounded two-player games, all three naturally lie within PSPACE; are they PSPACE-complete?

Our results. In this paper, we settle two out of three of these 20-year-old open problems by proving that Strings-and-Coins and Nimstring are PSPACE-complete. This is the first improvement beyond NP-hardness since the original *Winning Ways* result from 1984. Our reductions from Game SAT are relatively simple but subtle. Along the way, we prove PSPACE-completeness of a new Strings-and-Coins variant called *Coins-Are-Lava*, where the first person to free a coin loses.

Our constructed game positions rely on *multigraphs* with multiple copies of some edges/strings, a feature not present in instances corresponding to Dots-and-Boxes. Thus our results do not apply to Dots-and-Boxes. A generalization of Dots-and-Boxes that we might be able to target is *weighted* Dots-and-Boxes, where each grid edge has a specified number of times it must be drawn before it is “complete” and thus can form the boundary of a 1×1 box. This game corresponds to Strings-and-Coins on planar multigraphs whose vertices can be embedded at grid vertices such that edges have unit length. However, our multigraphs are neither planar nor maximum-degree-4, so they cannot be drawn on a square grid, and the complexity of weighted Dots-and-Boxes remains open.

2 Nimstring

We begin with more formal definitions of the games of interest, and some known lemmas about them:

Definition 2.1 (STRINGS-AND-COINS(G) and NIMSTRING(G)). A *multigraph* G consists of vertices, also called *coins*, and edges, also called *strings*, where each edge $e \in E$ is a subset of at most two vertices in V . STRINGS-AND-COINS(G) and NIMSTRING(G) represent the games played on G where two players alternate removing edges and, if a player removes the last edge incident to one or two coins (*freeing* the coin), then that player gains a point and must move again. The games end when there are no more strings; in STRINGS-AND-COINS(G), the player with the most points wins, while in NIMSTRING(G), the last player to move wins.

Next we prove the standard result that Nimstring is a special case of Strings-and-Coins, and thus hardness of the former implies hardness of the latter:

Lemma 2.2. [BCG03, p. 552] *For every graph G , there exists an efficiently computable graph H such that the winner of $\text{NIMSTRING}(G)$ is the same as the winner of $\text{STRINGS-AND-COINS}(H)$.*

Proof. Let $H = G \cup C_n$ where C_n is a cycle on $n > |V(G)|$ vertices. If a player cuts any string in this cycle, then the opponent can claim $n > \frac{|V(H)|}{2}$ coins in a single turn, winning the game. Therefore the players will try to only cut edges in G , and the player who cannot do so loses. This goal is equivalent to just playing $\text{NIMSTRING}(G)$. \square

A final known result we will need is about “loony” positions in Nimstring:

Lemma 2.3. [BCG03, p. 557] *If G has a degree-2 vertex adjacent to exactly one degree-1 vertex, then the first player can always win in $\text{NIMSTRING}(G)$. Such positions are known as **loony** positions.*



(a) String b connects to a vertex of degree at least 2.

(b) String b connects to the ground.

Figure 1: Two loony positions

Proof. Let a be the string between the two coins, b be the other string connected to a degree-2 coin, and G' be the rest of the graph (Figure 1). One of the players has a winning strategy in $\text{NIMSTRING}(G')$.

- If the first player has a winning strategy in $\text{NIMSTRING}(G')$, then we cut strings a and b in this order. We get exactly G' and it is still our turn. By assumption we can win.
- If the second player has a winning strategy in $\text{NIMSTRING}(G')$, then we just cut string b . We get graph G' (plus an extra edge that does not affect the game), and it is our opponent's turn. By assumption opponent cannot win. \square

3 Coins-are-Lava

We introduce a variant game played on strings and coins that we find easier to analyze, called *Coins-are-Lava*:²

²For a “practical” motivation for this game, consider the 1933 Double Eagle U.S. coin: until 2002, possession of this coin could result in imprisonment [USM02].

Definition 3.1 ($\text{COINS-ARE-LAVA}(G)$). Given a multigraph G , $\text{COINS-ARE-LAVA}(G)$ represents the game played on G where two players alternate removing edges and, if a player frees a coin, that player loses. Equivalently, players are forbidden from removing an edge that would free a coin, and the winner is determined according to normal play.

Now we show that Coins-are-Lava is a special case of Nimstring. Thus, its hardness will imply the hardness of both Nimstring and (by Lemma 2.2) Strings-and-Coins.

Lemma 3.2. *For every graph G , there exists an efficiently computable graph H such that the winner of $\text{COINS-ARE-LAVA}(G)$ is the same as the winner of $\text{NIMSTRING}(H)$.*




Figure 2: (a) The graph H ; (b) freeing a coin in G results in a loony position; (c–f) cutting a string outside G results in a loony position.

Proof. Let H be a graph obtained from G by connecting every coin to the ground with a long chain (length ≥ 5); see Figure 2a.

If a player cuts a string in one of these chains, or cuts all strings in G attached to the same coin, this creates a loony position and ends their turn; see Figure 2. By Lemma 2.3, their opponent can then win.

Therefore the players will try to avoid cutting strings outside G or freeing a coin in G . The first player to fail to do so loses. This goal is equivalent to $\text{COINS-ARE-LAVA}(G)$. \square

4 PSPACE-Hardness

It remains to prove that $\text{COINS-ARE-LAVA}(G)$ is PSPACE-complete. Our reduction is from the following known PSPACE-complete problem.

Definition 4.1 ($\text{GAME-SAT}(\mathcal{F})$). Given a positive DNF formula \mathcal{F} (an OR of ANDs of variables without negation), $\text{GAME-SAT}(\mathcal{F})$ is the following game played by two players, Trudy and Fallon. Initially each variable is *unset*. In each turn, the player may *set* a variable to TRUE or FALSE, or the player may *skip* their turn (do nothing). The game ends when all variables are set; then Trudy wins if formula \mathcal{F} is true, while Fallon wins if formula \mathcal{F} is false.

We allow players to skip turns and to set variables to the “wrong” value (Trudy to false or Fallon to true). The player with a winning strategy can always avoid such moves, however, replacing them with dominating “good” moves that do not skip and play the “right” value (Trudy to true or Fallon to false), as such moves never hurt the winning player’s final goal.

Schaefer [Sch78] proved that this game is PSPACE-complete, under the name $G_{pos}(\text{POS DNF})$.

Theorem 4.2. *Coins-are-Lava is PSPACE-complete.*

Proof. Let \mathcal{F} be a positive DNF formula with n variables, m clauses, and k_i occurrences of each variable x_i . Without loss of generality, every clause contains at least 2 variables and every variable appears in at least 1 clause. Fix a sufficiently large number $N \gg m^2 n^2$.

First we define several useful gadgets. Many of these gadgets are parametrized by an integer *level*. Intuitively, doing anything to a level- $(\ell + 1)$ gadget requires an order of magnitude more time than doing anything to a level- ℓ gadget. This way we can make sure that players interact with gadgets in the right order. However since each level- ℓ gadget uses $N^{\theta(\ell)}$ strings, we can only use a constant number of levels.




Figure 3: A width-5 rope




Figure 4: Variable gadget

A *rope* (Figure 3) is a collection of strings that share both endpoints. The number of strings in a rope is called its *width*. We say that a rope has been *cut* when all of its strings have been cut. When the game ends, every rope has either been cut completely, or it has only 1 string remaining. (Otherwise, a string in the rope can always be safely cut without freeing any coin.)

A *variable* gadget (Figure 4) consists of an *internal* string and an *output* string. We say that it is *set to false* if the internal string has been cut, *set to true* if the output string has been cut, and *unset* if neither string has been cut. A variable implicitly has level 0.

A level- ℓ *wire* gadget (Figure 5) consists of a width- $N^{2\ell-1}$ *input* rope and a width- $N^{2\ell}$ *output* rope connected in series. We say that it is *disabled* if the input rope has been cut, *activated* if the output rope has been cut. The *HP* (Hit Points) of the wire is the number of strings remaining in




Figure 5: Wire gadget




Figure 6: Clause gadget




Figure 7: Graph G for formula $(x_1 \wedge x_2 \wedge x_3) \vee (x_2 \wedge x_3) \vee (x_3 \wedge x_4)$. Blue labeled clauses are real; green clauses are singleton; and the red clause is empty. Dotted lines indicate that there are supposed to be $k_i - 1$ wires there, but $k_i - 1 = 0$.

the input rope. Note that activating a wire typically takes a lot longer than disabling it. Intuitively this means that you can only activate a wire if disabling it would free a coin.

A level- ℓ **clause** gadget (Figure 6) consists of a single width- $N^{2\ell-1}$ **input** rope connected to the ground. We say that it is **disabled** if the input string has been cut. The **HP** of the clause is the number of strings remaining in the input rope.

The winner is determined solely by the parity of the number of removed strings. We can easily flip this parity, for example by adding an extra ground-to-ground string. So without loss of generality, Fallon wins if (but not only if) every variable and wire has one string remaining and all m clauses have no strings. Then Trudy wins if every variable and wire has one string remaining, $m - 1$ clauses have no strings, and the final clause has one string. In fact, we will show that the game has to end in one of these two specific ways.

Let \mathcal{F}' be a new formula with the following clauses:

- all clauses from \mathcal{F} , which we call **real** clauses;
- for every variable, a **singleton** clause containing just that variable; and
- one additional **empty** clause that contains no variables and is always satisfied.

We construct a multigraph G by connecting the gadgets as follows; refer to Figure 7:

- a variable gadget for each variable;
- a clause gadget for each clause;

- a level-1 wire from each variable x_i to each of k_i real clauses that contain that variable;
- $k_i - 1$ level-1 wires from each variable to the corresponding singleton clause;
- a single vertex called the *root coin*;
- a level-2 wire from the root coin to every real clause and every singleton clause;
- and
- $n + m - 1$ level-2 wires from the root coin to the empty clause.

First we describe how typical gameplay in G should look (without proofs) to give some intuition for why this construction makes sense, and then we prove that it works more formally.

1. First Trudy and Fallon set variable gadgets to true and false respectively.
2. Then the players disable all wires from false variables, and disable all but one wire from each true variable (disabling all wires from a true variable would free a coin). Then they activate the level-1 wires that have not been disabled (one from each true variable). Note that almost half the wires from each variable go to singleton clauses. If all real clauses are false, then those wires form the majority, and Fallon can ensure that one of them gets activated. But if even one real clause is true, the wires to true clauses (real or singleton) now form a majority, and Trudy can ensure that one of them gets activated.
3. Then the players disable all but one level-2 wire and activate the remaining level-2 wire (disabling all of them would free a coin). Almost half of these wires go to the empty clause. If all real clauses are false, then they form a majority, and Fallon can ensure that one of them gets activated. But if even one real clause is true, then it together with the empty clause forms a majority, and Trudy can ensure that one of them gets activated.
4. Finally, the players disable the clause gadgets. A clause can be disabled unless all wires pointing at it got activated (in that case, disabling it would free a coin). If the formula is not satisfied, then all clauses get disabled and Fallon wins. If the formula is satisfied, then exactly one clause remains and Trudy wins.

We want to show that the winner of $\text{COINS-ARE-LAVA}(G)$ is the same as the winner of $\text{GAME-SAT}(\mathcal{F})$. We do a case split on the winner of $\text{GAME-SAT}(\mathcal{F})$, and in each case provide a winning $\text{COINS-ARE-LAVA}(G)$ strategy for that player.

If Fallon can win $\text{GAME-SAT}(\mathcal{F})$, then they can win $\text{COINS-ARE-LAVA}(G)$ using the following strategy:

1. There is a natural mapping f from states of $\text{COINS-ARE-LAVA}(G)$ to states of $\text{GAME-SAT}(\mathcal{F})$: a variable x_i in $\text{GAME-SAT}(\mathcal{F})$ is set to true if the corresponding variable gadget is set to true, set to false if the gadget is set to false, and unset if the gadget is unset. Every move in $\text{COINS-ARE-LAVA}(G)$ maps to a valid move in $\text{GAME-SAT}(\mathcal{F})$, where moves outside of variable gadgets map to skip moves. Also, if we played $\text{COINS-ARE-LAVA}(G)$ for less than $2n$ moves, then we can perform any move that is valid in the corresponding $\text{GAME-SAT}(\mathcal{F})$ state. This does not free a coin, because the relevant coin has degree $\Omega(N) \gg 2n$. So we can transfer the strategy from $\text{GAME-SAT}(\mathcal{F})$ to $\text{COINS-ARE-LAVA}(G)$: for every opponent's move in $\text{COINS-ARE-LAVA}(G)$, map it to $\text{GAME-SAT}(\mathcal{F})$, find the best response, and map it back to $\text{COINS-ARE-LAVA}(G)$. We play like this until we have set all variable gadgets to some assignment that does not satisfy \mathcal{F} , as guaranteed by Fallon's winning strategy.

2. Call a level-1 wire from a true variable x_i **good** if it points at a real clause and **bad** if it points at a singleton clause. Wires from false variables are **neutral**. Each true variable x_i has k_i good wires and $k_i - 1$ bad ones. For each true variable x_i , the total HP of bad wires is still at most $(k_i - 1)N^1$ and total HP of all good wires is at least $k_i N^1 - O(n) > (k_i - 1)N^1$ (the opponent could cut up to $O(n)$ strings here while we were setting variables).
 - (a) Disable all bad wires, Specifically, if the opponent reduced HP of a good wire connected to some true variable x_i , we respond by reducing HP of a bad wire connected to the same x_i ; if the opponent did something else or x_i has no bad wires left, we reduce HP of a bad wire connected an arbitrary variable x_j . This maintains the invariant that for each true variable x_i , HP of x_i 's good wires is higher than HP of x_i 's bad wires. The opponent cannot activate any bad wires because that would take $\Theta(N^2) \gg \sum_i k_i N^1$ turns.
 - (b) Disable good and neutral level-1 wires until there is only one good wire remaining per true variable. Once again, the opponent cannot activate these wires because that would take too long.
 - (c) Activate the remaining good wires. Opponent cannot disable these wires, because that would free a coin.
 - (d) We have activated exactly one good wire per true variable. There are no activated level-1 wires pointing at satisfied clauses, because real clauses are unsatisfied and singleton clauses are bad.
3. Call a level-2 wire **good** if it points to a real or singleton clause and **bad** if it points to the empty clause. The total HP of bad wires is at most $(n + m - 1)N^3$ and the total HP of good wires is at still least $(n + m)N^3 - O(nmN^2) > (n + m - 1)N^3$.
 - (a) Disable all bad wires. The opponent cannot disable all good wires before we disable the bad ones because good wires have more HP.
 - (b) Disable all but one good wire. The opponent cannot activate any of these wires because it would take $\Theta(N^4)$ turns.
 - (c) Activate the last good wire. The opponent cannot disable it because that would free the root coin.
4. Disable all clause gadgets. This will not free a coin, because every clause has at least one disabled wire: real clauses are unsatisfied so there is a wire from a false variable that we broke in Step 2(b), singleton clauses have bad level-1 wires that we broke in Step 2(a), and the empty clause has a bad level-2 wire that we broke in Step 3(a). We win because there are no clause gadgets remaining.

If Trudy can win $\text{GAME-SAT}(\mathcal{F})$, then they can win $\text{COINS-ARE-LAVA}(G)$ using the following strategy:

1. Set the variable gadgets to some assignment that satisfies \mathcal{F} . Let $C \in \mathcal{F}$ be a satisfied clause.
2. Call a level-1 wire from a true variable $x_i \in C$ **good** if it points at a singleton clause or C and **bad** otherwise. Wires from variables not in C are **neutral**. Each variable $x_i \in C$ has $k_i - 1$ bad wires (k_i to real clauses, but one of them is C) and k_i good ones ($k_i - 1$ to the singleton clause and one to C). Disable all bad wires, then disable all neutral wires and all-but-one good wire per variable in C , and then activate the remaining good wires. Each activated wire

points to a singleton clause or to C . Then either all of them point to C , or at least one of them points to a singleton clause. Either way, we have some satisfied real or singleton clause C' with only activated level-1 wires.

3. Call a level-2 wire **good** if it points to C' or to the empty clause. There are $n + m - 1$ bad wires ($n + m$ to real clauses, but one of them is C') and $n + m$ bad ones ($n + m - 1$ to the empty clause plus one to C'). Disable all bad wires and activate exactly one good wire. Let C'' be the clause pointed by the activated wire (either C' or the empty clause).
4. Disable all clause gadgets other than C'' . This will not free a coin, because every clause other than C'' has a disabled level-2 wire. But C'' cannot be disabled, because all of the wires pointing at it have been activated. We win because there is exactly one clause gadget remaining. \square

Corollary 4.2.1. *Nimstring is PSPACE-complete.*

Proof. This follows immediately from Theorem 4.2 and Lemma 3.2. \square

Corollary 4.2.2. *Strings-and-Coins is PSPACE-complete.*

Proof. This follows immediately from Theorem 4.2.1 and Lemma 2.2. \square

References

- [AMS03] American Mathematical Society. Elwyn Berlekamp gives Arnold Ross lecture. <http://www.ams.org/programs/students/ar12004>, 2003.
- [BCG03] Elwyn R. Berlekamp, John H. Conway, and Richard K. Guy. *Winning Ways for Your Mathematical Plays*, volume 3. A K Peters, Wellesley, MA, 2nd edition, 2003. First edition published in 1982.
- [Ber00] Elwyn Berlekamp. *The Dots and Boxes Game: Sophisticated Child's Play*. A K Peters, Wellesley, MA, 2000.
- [CDDL15] Sébastien Collette, Erik D. Demaine, Martin L. Demaine, and Stefan Langerman. Narrow misère dots-and-boxes. In Richard J. Nowakowski, editor, *Games of No Chance 4*, pages 57–64. Cambridge University Press, 2015.
- [DH09] Erik D. Demaine and Robert A. Hearn. Playing games with algorithms: Algorithmic combinatorial game theory. In *Games of No Chance 3*, pages 3–56. Cambridge University Press, 2009. arXiv:cc.CC/0106019. First version appeared at MFCS 2001.
- [Epp] David Eppstein. Computational complexity of games and puzzles. <http://www.ics.uci.edu/~eppstein/cgt/hard.html>.
- [GN02] Richard K. Guy and Richard J. Nowakowski. Unsolved problems in combinatorial games. In R. J. Nowakowski, editor, *More Games of No Chance*, pages 457–473. Cambridge University Press, 2002.
- [Sch78] Thomas J. Schaefer. On the complexity of some two-person perfect-information games. *Journal of Computer and System Sciences*, 16(2):185–225, 1978.
- [USM02] United States Mint. The United States government to sell the famed 1933 double eagle, the most valuable gold coin in the world. <https://www.usmint.gov/news/press-releases/20020207-the-united-states-government-to-sell-the-fame> February 2002.