

A PSPACE-complete graph nim

KYLE BURKE AND OLIVIA C. GEORGE

We build off G. Stockman's game NIMG to create a version named NEIGHBORING NIM. By reducing from GEOGRAPHY, we show that this game is PSPACE-hard. The games created by the reduction share strong similarities with UNDIRECTED (VERTEX) GEOGRAPHY and regular NIM, though these are both solvable in polynomial-time. This application of graphs can be used as a form of game sum with any rulesets, not only NIM.

1. Background

1.1. Algorithmic combinatorial game theory. Most of the results here revolve around the computational complexity of determining which player has a winning strategy from a given game position. There exist faster algorithms to solve this problem for some rulesets than for others. For each ruleset, we consider the computational problem that could be solved by such an algorithm. We will refer to both the ruleset and problem by the same name.

We strongly encourage readers unfamiliar with these topics to refer to [1].

1.2. Terminology. A small amount of nonstandard terminology is used:

- We use the word *sticks* to refer to the objects in nim heaps. Thus, a nim heap of size six contains six sticks.
- An *optimal sequence set* is a set of sequences of plays for both players such that any move deviating from all of the sequences results in an \mathcal{N} -position (meaning, the \mathcal{N} ext player has a winning move). No move in that sequence should be nonoptimal for either player. Thus, if a player does not know whether they have a winning strategy, adhering to an optimal sequence is at least as good as any other move.

1.3. NIM. NIM is an impartial game played on a collection of heaps, each with a nonnegative number of sticks. On a player's turn, they choose a nonempty pile

MSC2010: primary 91A46; secondary 68Q17.

Keywords: nim, graph nim, neighboring nim, combinatorial game theory, impartial games.

and remove as many sticks as desired (at least one) from that pile. A player loses when they cannot remove sticks (all piles are empty).

NIM is a classic impartial game, being the basis of Nimbers and Sprague–Grundy theory [8; 5]. NIM has lots of nice properties, from easy evaluation of games to obvious composition of two NIM games (the sum is just a new NIM game).

1.4. NIMG. NIM has been extended to incorporate graphs so that nim heaps are assigned to either edges or vertices. There are three different versions of the game named NIMG. In all three versions, a turn consists of both traversing an edge of the graph and removing sticks from a visited element.

1.4.1. Edge-heap NimG. Fukuyama describes NIMG where nim heaps are embedded into the edges of the graph [4]. On each turn, the current player chooses an edge to traverse (which has at least 1 stick on it) and removes any number of sticks from that edge. The next player then starts on the vertex on the other end of that edge and must choose an adjacent edge for their move. When there are no more edges with sticks adjacent to the current vertex, the current player loses. Many results for this game are known on complete graphs [2].

1.4.2. Vertex-heap NimG. In VERTEX NIMG, players similarly move from one vertex to another, but heaps are connected to the vertices instead of edges [9]. The two variants can be easily described here as: “remove sticks, then move” and “move, then remove sticks”. In both cases, a player loses if they cannot complete their turn. The main topic of this paper is a variant of “move, then remove”.

1.5. Geography. We will use GEOGRAPHY to show the PSPACE-hardness of NEIGHBORING NIM. There are many flavors of GEOGRAPHY; we use the term to refer to DIRECTED VERTEX GEOGRAPHY. This impartial game is played on a directed graph; each turn begins with a vertex already chosen. The current player’s turn consists of selecting an arc leaving the chosen vertex that leads to a vertex that hasn’t yet been visited during the game. The next player then starts their turn with the resulting vertex selected. We formally describe the ruleset as follows:

Definition 1.1 (GEOGRAPHY — DIRECTED VERTEX). Geography positions are described by $G = (V, E)$ and $v \in V$. Move options for (G, v) are all (G', v') where

- $(v, v') \in E$,
- $V' = V \setminus \{v\}$,
- E' is the subset of E induced by V' , and
- $G' = (V', E')$.

GEOGRAPHY is known to be PSPACE-complete [6; 7].

2. NEIGHBORING NIM

We define the ruleset NEIGHBORING NIM to be similar to the “move, then remove” version of NIMG, but also allow players to choose to play on the same vertex as the last move as though each vertex has a self-loop. Note that standard NIM is equivalent to a game of NEIGHBORING NIM on a complete graph with each heap on a separate vertex. A more formal definition follows.

Definition 2.1 (NEIGHBORING NIM). NEIGHBORING NIM positions are described by $G = (V, E)$, $w : V \rightarrow \mathbb{N}$, and $x \in V$. The options for (G, w, x) are all (G, w', x') where $w' : V \rightarrow \mathbb{N}$ and

- $x' = x$ or $\{x, x'\} \in E$,
- $w'(x') < w(x')$, and
- $\forall v \in V \setminus \{x'\} : w'(v) = w(v)$.

Our main result for this paper is that NEIGHBORING NIM is PSPACE-hard. Since our analysis uses graphs with a small number of sticks on each vertex, we define a version of the game with a bounded number of sticks per vertex.

Definition 2.2 (k -NEIGHBORING NIM). k -NEIGHBORING NIM is the same ruleset as NEIGHBORING NIM, except that the weight function w has bounded range $[0, k]$.

We are able to show that 2-NEIGHBORING NIM is PSPACE-complete, and thus c -NEIGHBORING NIM is also PSPACE-complete for any constant $c \geq 2$. The case for 1-NEIGHBORING NIM is solvable in polynomial time, since this game is equivalent to UNDIRECTED (VERTEX) GEOGRAPHY [3]. Thus, if $P \neq PSPACE$, allowing a second stick on vertex-heaps is enough to increase the computational hardness of determining the winning player!

3. Computational complexity of NEIGHBORING NIM

3.1. PSPACE-hardness. The following is the main result of this paper:

Theorem 3.1 (hardness). NEIGHBORING NIM is PSPACE-hard.

We will show the hardness of this problem by reducing from the game GEOGRAPHY, which is PSPACE-hard [6].

Proof. Given any GEOGRAPHY position, we will give an algorithm to construct an equivalent NEIGHBORING NIM state, meaning that there is a win in the GEOGRAPHY position exactly when there is a win in corresponding NEIGHBORING NIM position. First we will describe the method for generating these positions, then prove their equivalence.




Figure 1. Our main gadget: reduce each directed edge from y to z to the undirected weighted graph shown here.

Let GG be a GEOGRAPHY position on the directed and unweighted graph $G = (V, E)$. We define a new undirected graph, $G' = (V', E')$ with weights on the vertices $w : V' \rightarrow \mathbb{N}$ in the following way: $\forall v \in V$: let $X_v \in V'$ and set $w(X) = 1$. Also, $\forall (y, z) \in E$: (edge directed from y to z) let $a_{y,z}, b_{y,z}, c_{y,z}, d_{y,z}, e_{y,z}, f_{y,z}, g_{y,z} \in V'$ where, ignoring the (y, z) -subscripts,

- $w(a) = w(b) = w(c) = w(e) = w(f) = w(g) = 1$,
- $w(d) = 2$, and
- $\{(X_y, a), (a, b), (b, c), (c, d), (b, e), (e, f), (d, f), (d, g), (f, g), (g, X_z)\} \subset E'$.

See Figure 1 for a visual description of this gadget.

We soon show that the resulting G' is the graph for our NEIGHBORING NIM position equivalent to GG . The only final step in the reduction is to declare that if GG has a starting vertex v then $X_v \in V'$ is the starting vertex (where the previous play had been made) in our game and $w(X_v)$ is set to 0 instead of 1.

To complete the reduction, we must show that the structure in Figure 1 “acts” like a directed edge in GEOGRAPHY. Thus, we must prove:

- Moving “backwards” is a losing play. If the previous play was at X_z , then a backwards play would be to remove the only stick at $g_{(y,z)}$. A backwards play results in an \mathcal{N} -position.
- The same player moving into the gadget should also move out. If a player moves from X_y to $a_{(y,z)}$, then in an optimal sequence of plays, the same player will move from $g_{(y,z)}$ to X_z .

We prove the former in Lemma 3.2 and the latter in Lemma 3.3. The result is that each of these gadgets (as in Figure 1) in the NEIGHBORING NIM position works just like a (directed) edge in GEOGRAPHY. Trying to go backwards will result in losing and, if players play optimally, they both might as well continue normally through each gadget.




Figure 2. This sequence of moves is “safe” for both players to traverse the gadget. The two gray arrows indicate the potential additional moves in the second sequence. Each move assumes that exactly one object is taken from a vertex.

Lemma 3.2 (don’t go backwards). *Any play from (X_z) to $g_{(y,z)}$ (for all y) results in an \mathcal{N} -position.*

(See Appendix A for a proof of this claim.) This implies that our gadgets are directed: if a player tries to go “backwards” from an X -vertex to an i -vertex, the opponent will have a winning strategy.

To finish showing that our gadget acts like a directed edge, we must prove that “nothing can go wrong” during a regular forward traversal of the structure. To this end, we find two sequences that constitute an optimal sequence set through the gadget, thus showing that neither player benefits from deviating from the sequence. In order to get from one end of the gadget (as in Figure 1) to the other, the following sequence of moves suffices (let Alice and Bob be our two players; we will again ignore subscripts): Alice “takes” a , Bob takes b , Alice takes e , Bob takes f , Alice decrements d by 1, Bob takes g , Alice takes X_z . Note that the same player (in this example, Alice) who chooses to take a also moves to X_z . The other sequence is where Bob takes c instead of g ; here Alice will take the remaining object at d and Bob will be forced to take g , rejoining with the first sequence. See Figure 2 for a visual description of the safe sequences. We must prove that neither player benefits from deviating from these sequences. To do this, we show that any deviation is a losing move.

Lemma 3.3 (stick to the script). *Let the notation $k(p)$ denote taking p objects from vertex k in a turn. Then, after the plays $(\dots, X_y(1), a(1))$, any play deviating from the following sequences is a losing move:*

$$(b(1), e(1), f(1), d(1), g(1), X_z(1)),$$

$$(b(1), e(1), f(1), d(1), c(1), d(1), g(1), X_z(1)).$$

(See Appendix B for the proof of this claim.) This implies that once a player makes an appropriate move onto the gadget (playing on an a -node) any “safe”

sequence of moves in the gadget results in that same player making the play at the opposite X node. The two above claims combined show that our gadget correctly models a directed edge in a graph just between the X nodes.

Thus, for any edge (y, z) in our GEOGRAPHY position GG , the move to $a_{(y,z)}$ will result in the same player moving to X_z as desired. Also, since we proved players shouldn't go backwards, this game is equivalent to GG ; the first player has a winning strategy in GG exactly when the first player has a winning strategy in this NEIGHBORING NIM position.

Thus, NEIGHBORING NIM is PSPACE-hard. \square

The hardness of VERTEX NIMG follows directly.

Corollary 3.4 (VERTEX NIMG hardness). *VERTEX NIMG is PSPACE-hard.*

Proof. Neighboring Nim is a special case of VERTEX NIMG where all vertices have self-loops. Thus, VERTEX NIMG is also PSPACE-hard. \square

3.2. Speculation on completeness. Unfortunately, NEIGHBORING NIM is not automatically PSPACE-complete as games could take a number of moves exponential in the size of the description of the game. For example, a vertex can have a number of sticks exponential in the amount of bits needed to express that number and the rest of the graph. We leave this unsolved as Open Problem 6.1. There are good arguments to conjecture either way.

On one hand, it seems to not be inside PSPACE. Games can last an exponential number of turns, so the game trees are extremely tall. A straight-forward brute-force traversal can't be performed in polynomial space.

On the other hand, it might be inside PSPACE. Although there are many EXPTIME-hard rulesets, the authors know only of loopy examples. This means they can have positions that repeat during the course of a game, which cannot occur in NEIGHBORING NIM. Additionally, Nim heaps are well-understood; perhaps increasing the size of the heaps doesn't greatly increase the difficulty of finding strategies. It may also be that if there are only m different heap sizes on vertices, you can substitute them with a set of m significantly smaller sizes.

3.3. PSPACE-complete versions. We can sidestep this problem a bit by using our bounded-heap-size version of the game.

Corollary 3.5 (2-NEIGHBORING NIM completeness). *k -NEIGHBORING NIM is PSPACE-complete for any $k \geq 2$.*

Proof. The result of the reduction from Theorem 3.1 is always a 2-NEIGHBORING-NIM position. Thus, the PSPACE-hardness holds for this subset of positions as well. The positions are in PSPACE because k bounds the maximum number of moves per vertex. \square

4. Generalization

This graph-embedding technique works with games other than NIM. Given a graph, assign different game states to the vertices, and use similar rules: players may make one move legal in the game in any vertex neighboring the last play. We define this formally.

Definition 4.1 (NEIGHBORING- R). Given any ruleset R , NEIGHBORING- R has positions of the form $G = (V, E)$, $w : V \rightarrow \text{positions}(R)$, and $x \in V$. The left options for (G, w, x) are (G, w', x') where $w' : V \rightarrow \text{positions}(R)$ and

- $x' = x$ or $\{x, x'\} \in E$,
- $w'(x')$ is a left option of $w(x')$, and
- $\forall v \in V \setminus \{x'\} : w'(v) = w(v)$.

The right options are defined analogously: $w'(x')$ must be a right option of $w(x')$.

4.1. Inequivalent positions. This new definition allows a NEIGHBORING NIM vertex to contain multiple heaps instead of only a single heap. Although each nim position is equivalent to a single heap, that equivalence doesn't carry over in the neighboring situation. Consider the two NEIGHBORING NIM games in Figure 3. (The previous move was made on the leftmost vertex in both cases.) The values of the games embedded in the left vertices are both 0, the values on the middle vertices are both 0, and the values of the rightmost games are both $*$. However, the overall value of the positions are not equivalent.

In the top game, there are no move options, so the value is 0. In the bottom position, the next player can move to the middle vertex, even though the value of the nim game there is also zero. After that move there are exactly two moves remaining. Thus, the initial game has exactly three moves remaining and has a value of $*$.

4.2. Generalized hardness. The next result allows us to say something about the hardness of graph-embedded versions of many impartial games.




Figure 3. Two NEIGHBORING NIM positions. In both, the last move was made on the dashed vertex. The value of the top game is zero, the value of the bottom game is $*$.

Theorem 4.2 (neighboring- R hardness). *For any ruleset R which has positions identical to 0, $*$, and $*2$, NEIGHBORING- R is PSPACE-hard.*

Proof. Two positions are *identical* if they have isomorphic game trees. Replacing the nim heaps of size 1 and 2 with $*$ and $*2$, respectively, in the reduction of Theorem 3.1 doesn't change the winnability of the resulting games. Thus, the reduction applies to R . \square

5. Conclusions

Building on algorithmic work analyzing different versions on NIMG, we present NEIGHBORING NIM, a new PSPACE-hard game.

An interesting aspect of the hardness of NEIGHBORING NIM is the juxtaposition with VERTEX GEOGRAPHY. 1-NEIGHBORING NIM is the same ruleset as UNDIRECTED VERTEX GEOGRAPHY, which is solvable efficiently [3]. 2-NEIGHBORING NIM, however, is PSPACE-hard.

Furthermore, we can replace NIM and apply the graph-embedding concept to any other ruleset R to create NEIGHBORING- R .

6. Future work

There are many extensions to the work described here. The most prominent is certainly the unknown completeness of NEIGHBORING NIM with any number of sticks.

Open Problem 6.1. NEIGHBORING NIM \in PSPACE?

Additionally, the computational hardness of other flavors of NIMG remains unsolved.

Open Problem 6.2. What is the computational complexity of EDGE NIMG?

Open Problem 6.3. What is the computational complexity of VERTEX NIMG on graphs without self-loops?

Other explorable problems include the hardness of other versions of NEIGHBORING- R .

Open Problem 6.4. Is NEIGHBORING- R PSPACE-hard if R includes any positions *equivalent* to $*$ and $*2$?

(Note that Open Problem 6.4 is a stronger statement than shown here because equivalent does not necessarily mean identical.)

Open Problem 6.5. For which other computationally easy rulesets R is NEIGHBORING- R hard?

Open Problem 6.6. Are there strictly partisan positions of a ruleset R that can be used to show NEIGHBORING- R is hard? How small can the game trees be to get a hard game?

7. Acknowledgments

The authors would like to thank those who played NEIGHBORING NIM with them, including three excellent Wittenberg University students: Deanna Fink, Dang Mai and Ernie Heyder, as well as Professor Doug Andrews who defeated the first author over and over again. We would also like to thank Professor Adam Parker for listening to initial versions of the hardness proof and proofreading this paper.

Appendix A: Proof of Lemma 3.2

Lemma 3.2 (don't go backwards). *Any play from (X_z) to $g_{(y,z)}$ (for all y) is suboptimal.*

We will refer to the player who moves from X_z to g (we will leave out the subscript for the internal vertices) as the “foe” while the other player is the “hero”. We will show that the hero has a winning strategy after a backwards move. We can now look at two cases, each depending on the state of the game outside the gadget.

The first is the case where the move from a to X_y would be a winning play. In this case, the hero can next move from g to d and take both of the objects there. The foe has two options, both of which, we show, allow the hero to win:

- (1) *The foe moves to c .* In this case the hero must choose to go to b . The foe can now either choose to move to a — in which case the hero will gladly move to X_y and win as we assumed — or to e . Then the hero simply takes the object at f and, as there are no more moves, the hero has won.
- (2) *The foe moves to f .* The hero must then take e and the foe must take b . The hero can then move to c and win the game.

The second major case assumes that the move from a to X_y is a losing play. Here, the hero can still move to d (from g) but will take only one of the objects. Now the foe has three options: taking the other object at d , moving to c or moving to f . We show all to be losses:

- (1) *Foe moves to c .* Now the hero should take the remaining object at d . The following sequence must occur: foe must take f , hero at e , foe at b , hero at a , followed by the foe at X_y , a losing move by our assumption.
- (2) *Foe takes the remaining object at d .* The hero will choose to take c , so the foe must take b . The hero can then take a , forcing the foe to take X_y , a losing move by our assumption.

- (3) *Foe takes f .* The hero should then take e so the foe must take b . Again, the hero can take a , so the foe must move to X_y , a losing move by our assumption.

Thus, it is a losing play to move from an X -vertex to a g -vertex.

Appendix B: Proof of Lemma 3.3

Lemma 3.3 (stick to the script). *Let the notation $k(p)$ denote taking p objects from vertex k in a turn. Then, after the plays $(\dots, X_y(1), a(1))$ any play deviating from the following sequences is a losing move:*

$$(b(1), e(1), f(1), d(1), g(1), X_z(1)),$$

$$(b(1), e(1), f(1), d(1), c(1), d(1), g(1), X_z(1)).$$

We continue by analyzing all possible deviations from these sequences and show that they result in a loss. In this claim, we will refer to the deviating player as the foe and the other player as the hero. We will show that the foe loses in each case. It may be helpful to refer to Figure 2 during these case descriptions.

- (1) *$c(1)$ instead of $e(1)$.* Here we have two subcases: either moving from g to X_z is a winning (result is a \mathcal{P} -position) or losing (an \mathcal{N} -position) move. If it's in \mathcal{N} , then the hero can respond to $c(1)$ with $d(1)$. If the foe then chooses $g(1)$, the hero can take the remaining stick in d with $d(1)$; $f(1)$ and $e(1)$ must follow with the hero winning. If the foe instead chooses $f(1)$, the hero can win instantly by choosing $e(1)$. For the foe's last chance, they could select $d(1)$, removing the other stick from d . The hero should respond with $g(1)$. The foe will lose by selecting $f(1)$, because the hero will win at $e(1)$, but the foe will also lose with $X_z(1)$, an \mathcal{N} -position, as assumed.
If $X_z(1)$ is instead leaves the board in \mathcal{P} , the hero should respond to $c(1)$ with $d(2)$. The foe could choose $f(1)$, but the hero can then win with $e(1)$. Instead, the foe can choose $g(1)$ in which case the hero can choose $X_z(1)$ and win, as assumed.
- (2) *$d(2)$ instead of (the first) $d(1)$.* Here the hero has a simple move to win. By taking $c(1)$ there are no further moves and the foe has lost.
- (3) *$g(1)$ instead of (the first) $d(1)$.* The hero can respond with $d(1)$. This leaves two different adjacent vertices with 1 object apiece and no other adjacent nonempty vertices. Either move by the foe results in one remaining move and a win for the hero.
- (4) *$d(1)$ instead of $c(1)$.* The hero can respond with $c(1)$ and win.
- (5) *$d(1)$ instead of $X_z(1)$.* This cannot happen in the second sequence, but if it happens in the first, the hero can respond with $c(1)$ and win.

Thus, any deviation from the two sequences specified in the claim puts the game in an \mathcal{N} -position.

References

- [1] Erik D. Demaine and Robert A. Hearn. Playing games with algorithms: Algorithmic combinatorial game theory. In Michael H. Albert and Richard J. Nowakowski, editors, *Games of No Chance 3*, volume 56 of *Mathematical Sciences Research Institute Publications*, pages 3–56. Cambridge University Press, 2009.
- [2] Lindsay Erickson. Nim on complete graphs. *Ars Combin.*, 119:117–127, 2015.
- [3] Aviezri S. Fraenkel, Edward R. Scheinerman, and Daniel Ullman. Undirected edge geography. *Theor. Comput. Sci.*, 112(2):371–381, 1993.
- [4] Masahiko Fukuyama. A nim game played on graphs. *Theor. Comput. Sci.*, 1-3(304):387–399, 2003.
- [5] P. M. Grundy. Mathematics and games. *Eureka*, 2:198–211, 1939.
- [6] David Lichtenstein and Michael Sipser. Go is polynomial-space hard. *J. ACM*, 27(2):393–401, 1980.
- [7] C. H. Papadimitriou. *Computational Complexity*. Addison Wesley, Reading, Massachusetts, 1994.
- [8] R. P. Sprague. Über mathematische Kampfspiele. *Tôhoku Math. J.*, 41:438–444, 1935.
- [9] Gwendolyn Stockman. Presentation: The game of nim on graphs: NimG, 2004. Available at http://www.aladdin.cs.cmu.edu/reu/mini_probes/papers/final_stockman.ppt.

kwburke@plymouth.edu

*Department of Computer Science and Technology,
Plymouth State University, Plymouth, NH, United States*

