Mathematical Games

Undirected edge geography

Aviezri S. Fraenkel *

Department of Applied Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel

Edward R. Scheinerman †

Department of Mathematical Sciences, The Johns Hopkins University, Baltimore, MD 21218-2689, USA

Daniel Ullman ‡

Department of Mathematics, The George Washington University, Washington, DC 20052, USA

Abstract

Fraenkel, A.S., E.R. Scheinerman and D. Ullman, Undirected edge geography, Theoretical Computer Science 112 (1993) 371–381.

The game of *edge geography* is played by two players who alternately move a token on a graph from one vertex to an adjacent vertex, erasing the edge in between. The player who first has no legal move is the loser. We show that the decision problem of determining whether a position in this game is a win for the first player is PSPACE-complete. Further, the problem remains PSPACE-complete when restricted to planar graphs with maximum degree 3. However, if the underlying graph is bipartite we provide (1) a linear algebraic characterization of the P- and N-positions, yielding (2) a polynomial time algorithm for deciding whether any given position is P or N, and also (3) a polynomial time algorithm to find winning moves.

Correspondence to: A.S. Fraenkel, Department of Applied Mathematics and Computer Science, Weizmann Institute of Science, Rehovot 76100, Israel. Email addresses of the authors: fraenkel@wisdom.weizmann.ac.il, ers@cs.jhu.edu and dullman@math.gwu.edu.

1. Introduction

The word "geography" is used to describe a collection of games played by two players who alternately move a token on a graph until one of them can no longer play legally and loses. The game acquires its name from the familiar word game in which two players alternately name geographical places subject to the restriction that the first letter of every place matches the last letter of the previously named place and that no place may be named twice. In this paper, we discuss generalizations of this word game to the category of finite graphs.

The graph game has several variations. First, the underlying graph may be directed (D) or undirected (U); we refer to the two options naturally as *directed* geography and *undirected* geography. Second, we consider the case where no vertex may be repeated (V) as well as the variation where vertices may be repeated but no edge may be repeated (E); we refer to the two options as *vertex* geography and *edge* geography. In each case, a position in the game is a rooted graph (directed or undirected) and the critical issue is to classify which rooted graphs are N-positions (wins for the *N*ext or first player) or P-positions (wins for the *P*revious or second player).

There are numerous other variations of the rules that we do not consider in this paper. For example, another natural way to begin a game of geography is to allow the first player to choose any vertex of an unrooted graph. Other options include multiple tokens which either block or annihilate one another, or tokens of two colors in which players may only move the tokens of their own color. See [1,3–6].

The directed case was first explored in [9], where it was shown that the question whether the first player can win when starting from a distinguished vertex in directed edge geography (DEG) is PSPACE-complete. The same result holds for the vertex case (DVG). In fact, both versions are PSPACE-complete even for bipartite, planar graphs with in/out degrees at most 2 and total degree at most 3. See [7,8].

In this note, then, we focus on geography played on undirected graphs. We can dismiss the vertex case (UVG), which turns out to be governed by Theorem 1.1, which is closely related to an exercise in [2, p. 71, Problem 4.1.4].

First some definitions. A *matching* in a graph is a set of independent edges. A matching is *maximum* if it is a matching of maximum cardinality. A matching M *saturates* a vertex v if v is incident to some edge in M.

Theorem 1.1. *Let G be a graph and v a distinguished vertex of G. Then the game* UVG *starting at position (G, v) is a first player win if and only if every maximum matching of G saturates v.*

Proof. Suppose that some maximum matching M does not saturate v. Then a

winning strategy for the second player is always to move the token along an edge in M. Such a move is always available by the maximality of M. Hence the second player need not lose.

Conversely, suppose that every maximum matching saturates v. Then a winning strategy for the first player is to choose any maximum matching M and always to move the token along an edge in M. Should such a move be unavailable, then there would be another matching of the same cardinality as M which does not saturate v, contradicting our assumption. Hence the first player need not lose. $\square$

It follows that in *unrooted* UVG (where the first player chooses the starting vertex) the second player has a winning strategy if and only if the graph has a perfect matching; this is the above-mentioned exercise in [2].

There are polynomial-time algorithms for determining the size of the maximum matching of a graph. By comparing the size of the maximum matching of G to the size of the maximum matching of $G - v$, one can check the condition of Theorem 1.1 in polynomial time. Hence there is a polynomial-time algorithm for determining whether the first or second player has the advantage in a game of UVG.

In Section 2, we show that the corresponding decision problem for undirected edge geography (UEG) is PSPACE-complete. In Section 3, we prove a characterization theorem for the same question in the case that the underlying graph is bipartite; this yields a polynomial-time algorithm for the decision problem. In Section 4, we apply our characterization theorem to some familiar classes of bipartite graphs. We close in Section 5 with some open problems.

2. General graphs

We are concerned with the following decision problem:

UNDIRECTED EDGE GEOGRAPHY
Instance: A rooted graph (G, v).
Question: Is (G, v) a P-position in UEG?

We show that this decision problem is in the complexity class PSPACE, meaning that it can be solved on a computer using only a polynomial amount of memory (though possibly exponential time). In fact, we show that it is PSPACE-complete, roughly meaning that it is no more easily computed than any other problem in PSPACE. This is very strong evidence of intractability, at least as strong as NP-completeness.

**Theorem 2.1.** *The* UEG *decision problem is PSPACE-complete.*

Fig. 1. A pseudoarc A composed of edges.

Proof. One can see directly that the problem is in PSPACE, since the game can last only $|E|$ moves. We show that UEG is PSPACE-complete by providing a polynomial transformation from DEG.

Given an arbitrary rooted directed graph (D, w) (i.e., an instance of DEG), we construct a rooted graph (G, v) (i.e., an instance of UEG) by replacing every arc (x, y) in D with a copy of the graph A pictured in Fig. 1. (We call the graph A a *pseudoarc* because, as we discuss below, it functions in UEG as an arc from x to y.) Putting $w = v$ completes the description of the rooted graph. This construction is clearly polynomial.

One checks that a player of UEG who moves into the pseudoarc A from x will arrive 4 moves later at y, assuming optimal play by both players. (It is curious that any move straying from this path is immediate suicide.) Moreover, no optimal player will ever play from y into A (unless forced), since the opponent has a (local) forced victory by responding to z. Hence (G, v) plays as a UEG position exactly as (D, w) plays as a DEG position. In other words, (D, w) is a P-position of DEG if and only if (G, v) is a P-position of UEG. Since the DEG decision problem is known to be PSPACE-complete, the same can be said for UEG.   □

The DEG decision problem remains PSPACE-complete even for planar directed graphs in which no vertex has degree greater than 3; see [7]. Our local replacement construction preserves this property, hence the UEG decision problem is PSPACE-complete even for planar graphs with maximum degree 3.

Although deciding whether (G, v) is a P-position is in general intractable, there is a simple sufficient condition that settles some special cases. An *even kernel* for G is a nonempty set S of vertices such that (1) no two elements of S are adjacent and (2) every vertex not in S is adjacent to an even number (possibly 0) of vertices in S.

Proposition 2.2. *If* S *is an even kernel for* G *and* $v \in S$, *then* (G, v) *is a P-position in* UEG.

Proof. The second player wins with the following simple strategy: Always move to a vertex in S (it does not matter which one). Since S is a stable (i.e., an

Fig. 2. A P-position of UEG in which the root u is *not* in any even kernel.

independent) set of vertices, the first player always moves to a vertex not in S. Since each vertex not in S has an even number of neighbors in S, the second player is never at a loss for a move. □

Corollary 2.3. *Suppose that nonadjacent vertices v and w in a graph G have the same neighbors. Then (G, v) is a P-position in* UEG.

Proof. The set $S = \{v, w\}$ is an even kernel for G. □

Corollary 2.4. *For $n \geq 2$, (K_n, v) is an N-position of* UEG.

Proof. Suppose that the first player moves from v to w. Then v and w have the same neighbors in $K_n - vw$, so, by the previous corollary, $(K_n - vw, w)$ is a P-position. Thus this (and every) first move by the first player is a winning move. Hence (K_n, v) is an N-position. □

Even kernels only provide an even answer to the UEG decision problem. There are P-positions whose root is not in any even kernel, such as vertex u in Fig. 2; and, of course, there are graphs without an even kernel, such as K_n (with $n \geq 1$). However, as we show in the next section, even kernels give a complete description of how to play UEG on bipartite graphs.

3. Bipartite graphs

As mentioned above, both DEG and DVG remain PSPACE-complete even when the underlying graph is restricted to be bipartite. In contrast, we show here that the bipartite case of UEG is polynomial. In order to describe the proof, we make use of an alternate description of bipartite UEG. Associated with any bipartite graph G with parts X and Y is an $|X|$-by-$|Y|$ matrix M called the *bipartite adjacency matrix* whose i, j entry is 1 if $x_i y_j$ is an edge of G and 0 otherwise. By a *line* of a matrix we mean either a row or a column. We say that two lines are *perpendicular* if one is a row and one is a column; otherwise, they are *parallel*. Since the vertices of G are in one-to-one correspondence with the lines of M, we can play UEG on M and forget about G. A legal move consists of choosing a 1 in the distinguished line of M, changing it to a 0, and identifying the perpendicular line through this entry to be the new

distinguished line. In other words, the players alternately change 1's to 0's subject to the restriction that the row player must change a 1 in the same row as the previously changed 1 while the column player must change a 1 in the same column as the previously changed 1.

The advantage of this new model is that we can make use of linear algebra over the 2-element field GF(2). The following theorem reveals the key to bipartite UEG.

Theorem 3.1. *Given a binary matrix M and a distinguished line l, (M, l) is a P-position in* UEG *if and only if l is in the span over* GF(2) *of the other lines parallel to it.*

Proof. Suppose, without loss of generality, that l is a row, say the first row r_1. Suppose also that r_1 is in the span of the remaining rows, i.e., that

$$0 = \sum_{i \in I} r_i,$$

where I is some subset of the rows with $1 \in I$. (Note: all arithmetic in this proof is modulo 2.) Then observe that the set S of vertices of G associated with the rows r_i with $i \in I$ is an even kernel for G. Thus by Proposition 2.2, (M, l) is a P-position.

Suppose, conversely, that r_1 is linearly independent of the remaining rows. Then we shall show that there is a 1 in row r_1, say $m_{1j} = 1$, which when changed to a 0 makes column c_j dependent on the other columns. By the first part of this proof, such a move is a winning move for the first player, since the opponent now faces a P-position. Hence the position is an N-position.

Let M^* be the matrix obtained from M by adding the column $e_1 = (1, 0, 0, ..., 0)^T$. Since r_1 was independent of the other rows, M^* has the same row rank as M. Hence M^* has the same column rank as M, which is to say that one may express e_1 as a linear combination of other columns:

$$e_1 = \sum_j b_j c_j \quad \text{(each } b_j = 0 \text{ or } 1). \tag{1}$$

Choose an index k so that $b_k = 1$ and the first entry of c_k is 1. (Such a k must exist, since the first entry of e_1 is one.) Then add c_k to both sides of (1) to obtain

$$c_k + e_1 = \sum_{j \neq k} b_j c_j. \tag{2}$$

The left-hand side of (2) is the column c_k with first bit switched from 1 to 0. The right-hand side of (2) shows that this can be expressed as a linear combination of remaining columns. $\square$

The proof shows that finding an even kernel is key to playing UEG in bipartite graphs. Thus we have the following strengthening of Proposition 2.2.

Corollary 3.2. *Let v be a vertex of a bipartite graph G. Then (G,v) is a P-position of* UEG *if and only if v is in an even kernel of G.*

The conditions in Theorem 3.1 can be readily checked in polynomial time using Gaussian elimination.

Corollary 3.3. *There is a polynomial-time algorithm for* UEG *in case the rooted graph is bipartite.*

Not only can we efficiently decide whether a rooted graph is a P- or an N-position in bipartite UEG, but we can efficiently (e.g., using Gaussian elimination) compute a winning move (if one exists) at each turn.

Corollary 3.4. *If G is a bipartite Eulerian graph and v is any vertex, then (G,v) is a P-position in* UEG.

Proof. All vertices in the two parts of G have even degree, so each part is an even kernel. □

In fact, we may strengthen the statement of this corollary to the following.

Corollary 3.5. *If G is a bipartite graph with parts X and Y, v is a vertex in X, and every vertex in Y has even degree, then (G,v) is a P-position.*

A variation of Theorem 3.1 can be found for *unrooted* bipartite UEG. This game is the same as UEG except that at the beginning, there is no root; player 1 chooses freely from among all vertices. That vertex is now the root and player 2 continues as in UEG.

Corollary 3.6. *Let G be a bipartite graph with bipartite adjacency matrix M. The second player has a winning strategy in unrooted* UEG *if and only if M is invertible (i.e., square and nonsingular) over* $GF(2)$.

4. Cubes and grids

In this section we apply Theorem 3.1 to the n-cube and to the $m \times n$-grid graph in order to classify in a combinatorial way which bipartite UEG positions are P-positions. The n-cube Q_n is the graph whose vertex set is the set of 0, 1 sequences of length n with an edge between two such sequences if they differ

in exactly one position. The $m \times n$-grid $G_{m \times n}$ is the graph whose vertex set is $\{1, 2, \ldots, m\} \times \{1, 2, \ldots, n\}$ with an edge between (i, j) and (k, l) if and only if $|i - k| + |j - l| = 1$.

Theorem 4.1. *Let v be any vertex in the n-cube, Q_n. Then (Q_n, v) is a P-position if and only if n is even.*

Proof. The n-cube, Q_n, is bipartite. When n is even, it is also Eulerian; hence Corollary 3.4 tells us that (Q_n, v) is a P-position.

Suppose n is odd. Without loss of generality, let $v = (0, 0, 0, \ldots, 0)$. Let $w = (1, 0, 0, \ldots, 0)$ and let S be the set of $0, 1$ sequences of length n with precisely one 1. Then S is an even kernel for $(Q_n - vw, w)$, so this is a P-position. Hence (Q_n, v) is an N-position, and the winning strategy for the first player is to always play into S. □

The situation with grid graphs is a bit more complicated.

Theorem 4.2. *Take $m, n \geqslant 2$ and let $v = (1, 1)$. Then $(G_{m \times n}, v)$ is a P-position if and only if $\gcd(m + 1, n + 1) \neq 1$.*

Proof. Suppose that $d > 1$ is a common divisor of $m + 1$ and $n + 1$. Let

$$S = \{(i, j) : d \nmid i, d \nmid j, \text{ and either } 2d \mid (i - j) \text{ or } 2d \mid (i + j)\}.$$

The set may be described in words as follows: divide the rectangle with vertices $(0, 0), (m + 1, 0), (0, n + 1), (m + 1, n + 1)$ into d-by-d squares. Color these square in checkerboard fashion. The set S is the set of points on the interiors of the diagonals of these squares, taking only one diagonal of each square, choosing this diagonal according to the color of the square. Fig. 3 shows the set S for $G_{14 \times 9}$.

One checks that S is an even kernel for $G_{m \times n}$, and so $(G_{m \times n}, (1, 1))$ is a P-position.

Now suppose that $\gcd(m + 1, n + 1) = 1$. One of m and n must be even; without loss of generality we assume m is even. In order to show that $(G_{m \times n}, (1, 1))$ is an N-position, we show that at least one of the two opening moves for the first player leads to a P-position.

Imagine a billiard ball rolling on an $(m + 1)$-by-$(n + 1)$ billiard table (with corners coordinatized by $(0, 0), (m + 1, 0), (0, n + 1)$, and $(m + 1, n + 1)$). Our grid graph is drawn on this table by placing vertex (i, j) at the point (i, j).

The billiard ball begins its journey at corner $(m + 1, 0)$, always maintaining a 45° angle with the sides of the table. As it travels, the ball bounces off the

Fig. 3. The even kernel S for $G_{14 \times 9}$.

Fig. 4. Lower left corner of the billiard table.

edges of the table. After traveling through $\mathrm{lcm}(m+1, n+1) = (m+1)(n+1)$ steps, it is "absorbed" in another corner.

The ball's trajectory can take it through any given vertex of the grid graph at most twice (once with slope $+1$ and once with slope -1). Further, it can only travel through *half* the vertices: those vertices (i, j) with $i + j \equiv 1 \pmod 2$—one of the partite sets of $G_{m \times n}$. Thus one checks that, in fact, the orbit visits every vertex (i, j) (with $i + j \equiv 1 \pmod 2$) exactly twice.

Let us examine this trajectory near $(1, 2)$ and $(2, 1)$. See Fig. 4. Four portions of the trajectory are labeled by a, b, c and d. Two of these four portions lead to corners of the table. It is not the case that *both* the a and d portions lead to a corner (for otherwise we would only visit $(1, 2)$ and $(2, 1)$ once); likewise at most one of b and c lead to a corner. Therefore, exactly one of portions b or c lead to a corner; let us say it is portion b.

Finally, let S be the set of all vertices which are traversed an odd number of times on portion b. See Fig. 5. Observe that every vertex not in S (an

Fig. 5. An even kernel in $G_{12 \times 9} - [(1,1)(1,2)]$.

independent set) is adjacent to an even number of members of S *except* vertex $(1,1)$. Thus a move from $(1,1)$ to $(1,2)$ results in a P-position. It follows that this is a winning move for the first player, i.e., $(G_{m \times n}, (1,1))$ is an N-position. $\square$

5. Open problems

We close with some questions raised by our investigation of UEG.

(1) **Even Kernel Computational Complexity.** Which graphs have even kernels? In particular, what is the computational complexity of the following decision problem:

> EVEN KERNEL
> *Instance*: A graph G.
> *Question*: Does there exist a nonempty, independent set of vertices S so that every vertex in G is adjacent to an even number (possibly 0) of elements of S?

Corollaries 3.2 and 3.3 imply that this question is polynomially decidable for bipartite graphs. For example (Theorem 4.1), the n-cube has a even kernel if and only if n is even.

A related question is: Which graphs G have the property that (G, v) is a P-position of UEG if and only if v is in an even kernel? Complete graphs and bipartite graphs share this property. Which others?

(2) **Unrooted Undirected Edge Geography.** What is the computational complexity of the *unrooted* UEG decision problem? In particular, given a graph G,

how do we decide whether (G, v) is an N-position in UEG for *every* vertex $v \in V(G)$?

(3) **Random graphs.** Let $G = G_{n,p}$ be a random graph on n labeled vertices $\{1, \ldots, n\}$ with fixed edge probability p. Is it the case that for almost all such G, $(G, 1)$ is an N-position in UEG? We suspect that the probability that $(G_{n,p}, 1)$ is an N-position goes to 1 as $n \to \infty$.

(4) **Grids.** For which values of m, n, i and j is $[G_{m \times n}, (i, j)]$ a P-position in UEG? (We showed that for $(i, j) = (1, 1)$ we have a P-position if and only if $m + 1$ and $n + 1$ have a nontrivial common factor.)

One can also consider higher-dimensional grid graphs.

(5) **Partizan Versions.** In *partizan* geography there are two tokens, one assigned to each player. Players may only move their own token along edges/arcs without repeating vertices/edges. There are now four decision problems: Under optimal play, can player 1 always force a win, regardless of which of the two tokens is assigned to that player? Can player 2? Can the player with the red token (regardless of who moves first)? Can the player with the blue token (regardless of who moves first)?

The question as to whether player 2 can win is NP-hard for rooted partizan UVG; there is a fairly straightforward reduction from HAMILTONIAN PATH (the idea is as in Theorem 2.3 of [4]). Is it actually PSPACE-complete? What is the computational complexity of this question for rooted UEG?

References

[1] H.L. Bodlaender, Complexity of path-forming games, *Theoret. Comput. Sci.* **110** (1993) 215–245.

[2] A. Bondy and U.S.R. Murty, *Graph Theory with Applications* (North-Holland, Amsterdam, 1976).

[3] A.S. Fraenkel and E. Goldschmidt, PSPACE-hardness of some combinatorial games, *J. Combin. Theory (A)* **46** (1987) 21–38.

[4] A.S. Fraenkel and S. Simonson, Geography, *Theoret. Comput. Sci.* **110** (1993) 197–214.

[5] A.S. Fraenkel and Y. Yesha, Complexity of problems in games, graphs and algebraic equations, *Discr. Appl. Math.* **1** (1979) 15–30.

[6] A.S. Fraenkel and Y. Yesha, Theory of annihilation games I, *J. Combin. Theory (B)* **33** (1982) 60–86.

[7] M.R. Garey and D.S. Johnson, *Computers and Intractability* (Freeman, San Francisco, 1979).

[8] D. Lichtenstein and M. Sipser, Go is polynomial-space hard, *J. ACM* **27** (1980) 393–401.

[9] T.J. Schaefer, On the complexity of some two-person perfect-information games, *J. Comput. System Sci.* **16** (1978) 185–225.