Geography

Start with a chip sitting on a vertex v of a graph or digraph G.

A move consists of moving the chip to a neighbouring vertex. In edge geography, moving the chip from x to y deletes the edge (x, y). In vertex geography, moving the chip from x to y deletes the vertex x.

The problem is given a position (G, v), to determine whether this is a P or N position.

Complexity Both edge and vertex geography are Pspace-hard on digraphs. Edge geography is Pspace-hard on an undirected graph. Only vertex geography on a graph is polynomial time solvable.

1 Undirected Vertex Geography – UVG

Theorem 1. *(G, v) is an N-position in UVG iff every maximum matching of G covers v.*

Proof (i) Suppose that M is a maximum matching of G which covers v. Player 1's strategy is now: Move along M-edge that contains current vertex.

If Player 1 were to lose, then there would exist a sequence of edges $e_1, f_1, \ldots, e_k, f_k$ such that $v \in e_1$, $e_1, e_2, \ldots, e_k \in M$, $f_1, f_2, \ldots, f_k \notin M$ and $f_k = (x, y)$ where y is the current vertex for Player 1 and y is not covered by M. But then if $A = \{e_1, e_2, \ldots, e_k\}$ and $B = \{f_1, f_2, \ldots, f_k\}$ then $(M \setminus A) \cup B$ is a maximum matching (same size as M) which does not cover v, contradiction.

(ii) Suppose now that there is some maximum matching M which does not cover v. Then if (v, w) is Player 1's move, w must be covered by M, else M is not a maximum matching. Player 2's strategy is now: Move along M-edge that contains current vertex. If Player 2 were to lose then there exists $e_1 = (v, w), f_1, \ldots, e_k, f_k, e_{k+1} = (x, y)$ where y is the current vertex for Player 2 and y is not covered by M. But then we have defined an augmenting path from v to y and so M is not a maximum matching, contradiction. □

Note that we can determine whether or not v is covered by all maximum matchings as follows: Find the size σ of the maximum matching G. This can be done in $O(n^3)$ time on an n-vertex graph. Then find the size σ' of a maximum matching in $G - v$. Then v is covered by all maximum matchings of G iff $\sigma \neq \sigma'$.

2 Undirected Edge Geography – UEG on a bipartite graph

An *even kernel* of G is a non-empty set $S \subseteq V$ such that (i) S is an independent set and (ii) $v \notin S$ implies that $deg_S(v)$ is even, (possibly zero). ($deg_S(v)$ is the number of neighbours of v in S.)

Lemma 1. *If S is an even kernel and $v \in S$ then (G, v) is a P-position in UEG.*

Proof Any move at a vertex in S takes the chip outside S and then Player 2 can immediately put the chip back in S. After a move from $x \in S$ to $y \notin S$, $deg_S(y)$ will become odd and so there is an edge back to S. making this move, makes $deg_S(y)$ even again. Eventually, there will be no $S : \bar{S}$ edges and Player 1 will be stuck in S. □

We now discuss Bipartite UEG i.e. we assume that G is bipartite, G has bipartion consisting of a copy of $[m]$ and a disjoint copy of $[n]$ and edges set E. Now consider the $m \times n$ 0-1 matrix A with $A(i, j) = 1$ iff $(i, j) \in E$.

We can play our game on this matrix: We are either positioned at row i or we are positioned at column j. If say, we are positioned at row i, then we choose a j such that $A(i,j) = 1$ and (i) make $A(i,j) = 0$ and (ii) move the position to column j. An analogous move is taken when we positioned at column j.

Lemma 2. *Suppose the current position is row i. This is a P-position iff row i is in the span of the remaining rows (is the sum (mod 2) of a subset of the other rows) or row i is a zero row. A similar statement can be made if the position is column j.*

Proof If row i is a zero row then vertex i is isolated and this is clearly a P-position. Otherwise, assume the position is row 1 and there exists $I \subseteq [m]$ such that $1 \in I$ and

$$r_1 = \sum_{i \in I \setminus \{1\}} r_i (mod\ 2) \text{ or } \sum_{i \in I} r_i = 0 (mod\ 2) \tag{1}$$

where r_i denotes row i.

I is an even kernel: If $x \notin I$ then either (i) x corresponds to a row and there are no x, I edges or (ii) x corresponds to a column and then $\sum_{i \in I} A(i,x) = 0(mod\ 2)$ from (1) and then x has an even number of neighbours in I.

Now suppose that (1) does not hold for any I. We show that there exists a ℓ such that $A(1,\ell) = 1$ and putting $A(1,\ell) = 0$ makes column ℓ dependent on the remaining columns. Then we will be in a P-position, by the first part.

Let e_1 be the m-vector with a 1 in row 1 and a 0 everywhere else. Let A^* be obtained by adding e_1 to A as an $(n+1)$th column. Now the row-rank of A^* is the same as the row-rank of A (here we are doing all arithmetic modulo 2). Suppose not, then if r_i^* is the ith row of A^* then there exists a set J such that

$$\sum_{i \in J} r_i = 0(mod\ 2) \neq \sum_{i \in J} r_i^* (mod\ 2).$$

Now $1 \notin J$ because r_1 is independent of the remaining rows of A, but then $\sum_{i \in J} r_i = 0(mod\ 2)$ implies $\sum_{i \in J} r_i^* = 0(mod\ 2)$ since the last column has all zeros, except in row 1.

Thus rank A^* = rank A and so there exists $K \subseteq [n]$ such that

$$e_1 = \sum_{k \in K} c_k (mod\ 2) \text{ or } e_1 + \sum_{k \in K} c_k = 0(mod\ 2) \tag{2}$$

where c_k denotes column k of A. Thus there exists $\ell \in K$ such that $A(1,\ell) = 1$. Now let $c_j' = c_j$ for $j \neq \ell$ and c_ℓ' be obtained from c_ℓ by putting $A(1,\ell) = 0$ i.e. $c_\ell' = c_\ell + e_1$. But then (2) implies that $\sum_{k \in K} c_k' = 0(mod\ 2)$ ($K = \{k\}$ is a possibility here).. □

Tic Tac Toe and extensions

We consider the following multi-dimensional version of Tic Tac Toe (Noughts and Crosses to the English). The *board* consists of $[n]^d$. A point on the board is therefore a vector $(x_1, x_2, \ldots, x_d)$ where $1 \leq x_i \leq n$ for $1 \leq i \leq d$.

A *line* is a set points $(x_j^{(1)}, x_j^{(2)}, \ldots, x_j^{(d)})$, $j = 1, 2, \ldots, n$ where each sequence $x^{(i)}$ is either (i) of the form $k, k, \ldots, k$ for some $k \in [n]$ or is (ii) $1, 2, \ldots, n$ or is (iii) $n, n-1, \ldots, 1$. Finally, we cannot have Case (i) for all i.

Thus in the (familiar) 3×3 case, the top row is defined by $x^{(1)} = 1, 1, 1$ and $x^{(2)} = 1, 2, 3$ and the diagonal from the bottom left to the top right is defined by $x^{(1)} = 3, 2, 1$ and $x^{(2)} = 1, 2, 3$

Lemma 3. *The number of winning lines in the (n, d) game is $\frac{(n+2)^d - n^d}{2}$.*

Proof In the definition of a line there are n choices for k in (i) and then (ii), (iii) make it up to $n+2$. There are d independent choices for each i making $(n+2)^d$. Now delete n^d choices where only Case (i) is used. Then divide by 2 because replacing (ii) by (iii) and vice-versa whenever Case (i) does not hold produces the same set of points (traversing the line in the other direction). □

The game is played by 2 players. The Red player (X player) goes first and colours a point red. Then the Blue player (0 player) colours a different point blue and so on. A player wins if there is a line, all of whose points are that players colour. If neither player wins then the game is a draw. The second player does not have a wnning strategy:

Lemma 4. *Player 1 can always get at least a draw.*

Proof We prove this by considering *strategy stealing.* Suppose that Player 2 did have a winning strategy. Then Player 1 can make an arbitrary first move x_1. Player 2 will then move with y_1. Player 1 will now win playing the winning strategy for Player 2 against a first move of y_1. This can be carried out until the strategy calls for move x_1 (if at all). But then Player 1 can make an arbitrary move and continue, since x_1 has already been made. □

2.1 Pairing Strategy

$$\begin{bmatrix} 11 & 1 & 8 & 1 & 12 \\ 6 & 2 & 2 & 9 & 10 \\ 3 & 7 & * & 9 & 3 \\ 6 & 7 & 4 & 4 & 10 \\ 12 & 5 & 8 & 5 & 11 \end{bmatrix}$$

The above array gives a strategy for Player 2 the 5×5 game ($d = 2, n = 5$). For each of the 12 lines there is an associated pair of positions. If Player 1 chooses a position with a number i, then Player 2 responds by choosing the other cell with the number i. This ensures that Player 1 cannot take line i. If Player 1 chooses the * then Player 2 can choose any cell with an unused number. So, later in the game if Player 1 chooses a cell with j and Player 2 already has the other j, then Player 1 can choose an arbitrary cell. Player 2's strategy is to ensure that after all cells have been chosen, he/she will have chosen one of the numbered cells asociatded with each line. This prevents Player 1 from taking a whole line. This is called a *pairing* strategy.

We now generalise the game to the following: We have a family $\mathcal{F} = A_1, A_2, \ldots, A_N \subseteq A$. A move consists of one player, taking an uncoloured member of A and giving it his colour. A player wins if one of the sets A_i is completely coloured with his colour.

A pairing strategy is a collection of distinct elements $X = \{x_1, x_2, \ldots, x_{2N-1}, x_{2N}\}$ such that $x_{2i-1}, x_{2i} \in A_i$ for $i \geq 1$. This is called a *draw forcing pairing.* Player 2 responds to Player 1's choice of $x_{2i+\delta}, \delta = 0, 1$ by choosing $x_{2i+3-\delta}$. If Player 1 does not choose from X, then Player 2 can choose any uncoloured element of X. In this way, Player 2 avoids defeat, because at the end of the game Player 2 will have coloured at least one of each of the pairs x_{2i-1}, x_{2i} and so Player 1 cannot have completely coloured A_i for $i = 1, 2, \ldots, N$.

Theorem 2. *If*

$$\left| \bigcup_{A \in \mathcal{G}} A \right| \geq 2|\mathcal{G}| \qquad \forall \mathcal{G} \subseteq \mathcal{F} \tag{3}$$

then there is a draw forcing pairing.

Proof We define a bipartite graph Γ. A will be one side of the bipartition and $B = \{b_1, b_2, \ldots, b_{2N}\}$. Here b_{2i-1} and b_{2i} both represent A_i in the sense that if $a \in A_i$ then there is an edge (a, b_{2i-1}) and an edge (a, b_{2i}). A draw forcing pairing corresponds to a complete matching of B into A and the condition (3) implies that Hall's condition is satisfied. $\square$

Corollary 3. *If $|A_i| \geq n$ for $i = 1, 2, \ldots, n$ and every $x \in A$ is contained in at most $n/2$ sets of $\mathcal{F}$ then there is a draw forcing pairing.*

Proof The degree of $a \in A$ is at most $2(n/2)$ in Γ and the degree of each $b \in B$ is at least n. This implies (via Hall's condition) that there is a complete matching of B into A. $\square$

Consider Tic tac Toe when case $d = 2$. If n is even then every array element is in at most 3 lines (one row, one column and at most one diagonal) and if n is odd then every array element is in at most 4 lines (one row, one column and at most two diagonals). Thus there is a draw forcing pairing if $n \geq 6$, n even and if $n \geq 9$, n odd. (The cases $n = 4, 7$ have been settled as draws. $n = 7$ required the use of a computer to examine all possible strategies.

In general we have

Lemma 5. *If $n \geq 3^d - 1$ and n is odd or if $n \geq 2^d - 1$ and n is even, then there is a draw forcing pairing of (n, d) Tic tac Toe.*

Proof We only have to estimate the number of lines through a fixed point $\mathbf{c} = (c_1, c_2, \ldots, c_d)$. If n is odd then to choose a line L through $\mathbf{c}$ we specify, for each index i whether L is (i) constant on i, (ii) increasing on i or (iii) decreasing on i. This gives 3^d choices. Subtract 1 to avoid the all constant case and divide by 2 because each line gets counted twice this way.

When n is even, we observe that once we have chosen in which positions L is constant, L is determined. Suppose $c_1 = x$ and 1 is not a fixed position. Then every other non-fixed position is x or $n - x + 1$. Asuning w.l.o.g. that $x \leq n/2$ we see that $x < n - x = 1$ and the positions with x increase together at the same time as the positions with $n - x + 1$ decrease together. Thus the number of lines through $\mathbf{c}$ in this case is bounded by $\sum_{i=0}^{d-1} \binom{d}{i} = 2^d - 1$. $\square$

2.2 Quasi-probabilistic method

We now prove a theorem of Erdős and Selfridge.

Theorem 4. *If $|A_i| \geq n$ for $i \in [N]$ and $N < 2^{n-1}$, then Player 2 can get a draw in the game defined by $\mathcal{F}$.*

Proof At any point in the game, let C_j denote the set of elements in A which have been coloured with Player j's colour, $j = 1, 2$ and $U = A \setminus C_1 \cup C_2$. Let

$$\Phi = \sum_{i : A_i \cap C_2 = \emptyset} 2^{-|A_i \cap U|}.$$

Suppose that the players choices are $x_1, y_1, x_2, y_2, \ldots,$. Then we observe that immediately after Player 1's first move, $\Phi < N 2^{-(n-1)} < 1$.

We will show that Player 2 can keep $\Phi < 1$ through out. Then at the end, when $U = \emptyset$, $\Phi = \sum_{i : A_i \cap C_2 = \emptyset} 1 < 1$ implies that $A_i \cap C_2 \neq \emptyset$ for all $i \in [N]$.

So, now let Φ_j be the value of Φ after the choice of $x_1, y_1, \ldots, x_j$. then if U, C_1, C_2 are defined at precisely this time,

$$\begin{aligned}
\Phi_{j+1} - \Phi_j &= -\sum_{\substack{i:A_i\cap C_2=\emptyset \\ y_j\in A_i}} 2^{-|A_i\cap U|} + \sum_{\substack{i:A_i\cap C_2=\emptyset \\ y_j\notin A_i, x_{j+1}\in A_i}} 2^{-|A_i\cap U|} \\
&\le -\sum_{\substack{i:A_i\cap C_2=\emptyset \\ y_j\in A_i}} 2^{-|A_i\cap U|} + \sum_{\substack{i:A_i\cap C_2=\emptyset \\ x_{j+1}\in A_i}} 2^{-|A_i\cap U|}
\end{aligned}$$

We deduce that $\Phi_{j+1} - \Phi_j \le 0$ if Player 2 chooses y_j to maximise over y, $\sum_{\substack{i:A_i\cap C_2=\emptyset \\ y\in A_i}} 2^{-|A_i\cap U|}$.

In this way, Player 2 keeps $\Phi < 1$ and obtains a draw. $\square$

In the case of (n, d) Tic Tac Toe, we see that Player 2 can force a draw if (see Lemma 3)

$$\frac{(n+2)^d - n^d}{2} < 2^{n-1}$$

which is implied, for n large, by

$$n \ge (1+\epsilon)d\log_2 d$$

where $\epsilon > 0$ is a small positive constsnt.

Shannon Switching Game Start with a connected multi-graph $G = (V, E)$.
Two players: Player A goes first and deletes edges and player B fortifies edges making them invulnerable to deletion by B. Player B wins iff the fortified edges contain a spanning tree of G.

Theorem 5. *Player B wins iff G contains two edge disjoint spanning trees.*

Proof (a) Here we assume that G has two edge disjoint spanning trees T_1, T_2 We prove this by induction on $|V|$. If $|V| = 2$ then G must contain at least two parallel edges joining the two vertices and so B can win. Suppose next that $|V| > 2$. Suppose that A deletes an edge $e = (x, y)$ of T_2 red. This breaks T_2 into two sub-trees T_2', T_2''. B will choose an edge $f = (u, v) \in T_1$ with one end in $V(T_2')$ and the other end in $V(T_2'')$. Now contract the edge f. In the new graph G^*, both T_1 and T_2 become spanning trees T_1^* and T_2^* and they are edge disjoint. It follows by induction that B can win the game on G^* and then wins the game on G by uncontracting the edge f. Of course f is chosen first of all still!

If A chooses an edge x in neither of the trees then B can choose an arbitrary edge f of T_1. Now let e be any edge of the unique cycle contained in $T_2 + e$. B can continue playing on $G - x$ as though e was the deleted edge. We can contract f as before and apply the above inductive argument.

(b) For this part we use a Theorem due to Nash-Williams:

Theorem 6. *Let k be a positive integer. Then G contains k edge disjoint spanning trees iff for every partition $\mathcal{P} = (V_1, V_2, \ldots, V_\ell)$ of V we have*

$$e(\mathcal{P}) = |E(\mathcal{P})| = \sum_{1 \le i < j \le \ell} e(V_i, V_j) \ge k(\ell - 1). \tag{4}$$

Here $E(\mathcal{P})$ is the set of edges joining different parts of the partition and $e(V_i, V_j)$ is the number of edges joining V_i and V_j.

Let us apply Theorem 6 with $k = 2$. If G does not contain two edge disjoint spanning trees, then it contains a partition $\mathcal{P} = (V_1, V_2, \ldots, V_\ell)$ with $e(\mathcal{P}) \le 2\ell - 3$. A starts by deleting an edge $e \in E(\mathcal{P})$. B will fortify an edge $f = (u, v)$. If u, v join different sets in the partition $\mathcal{P}$ then we can merge them and consider $\mathcal{P}'$ which has one less part and satisfies $e(\mathcal{P}') \le e(\mathcal{P}) - 2$ (edegs e, f have gone from the count). Otherwise B chooses an edge entirely inside a part of $\mathcal{P}$ and the number of parts does not change, but $e(\mathcal{P})$ goes down by one. Eventually, we come to a point where one part is joined to the rest of the graph by a single edge ($2\ell - 3 = 1$ when $\ell = 2$) and A wins by deleting this edge. □

Sketch of proof of Theorem 6

If $\mathcal{P} = (V_1, V_2, \ldots, V_\ell)$ is a partition and T is a spanning tree then T contains at least $\ell - 1$ edges of $E(\mathcal{P})$ and the only if part is straightforward.

Suppose now that (4) holds for all partitions. Let $\mathcal{F}$ be the set of edge disjoint forests containing the maximum number of edges. If $F = (F_1, F_2, \ldots, F_k) \in \mathcal{F}$ and $e \in E \setminus E[\mathcal{F}]$ then every $F_i + e$ contains a cycle. If e' belongs to this cycle then $F' \in \mathcal{F}$ where $F_j' = F_j$ for $j \ne i$ and $F_i' = F_i + e' - e$. We say that F' is obtained from F by a *replacement.*

Consider now a fixed $F^0 = (F_1^0, F_2^0, \ldots, F_k^0) \in \mathcal{F}$ and let $\mathcal{F}^0$ be the set of k-tples in $\mathcal{F}$ that can be obtained from F^0 by a sequence of replacements. Then let

$$E^0 = \bigcup_{F \in \mathcal{F}^0} (E \setminus E([F]).$$

Claim 1. *For every $e^0 \in E \setminus E([F^0]$ there exists a set $U \subseteq V$ that contains the endpoints of e^0 and induces a connected tree in F_i^0 for $1 \le i \le k$.*

Assume the claim for the moment. Suppose that not every F_i^0 is a spanning tree. Then G contains at least $k(|V|-1)$ edges (from (4) applied to the parttion of V into singletons) and so there exists $e^0 \in E \setminus E[F^0]$. Shrink the vertices of the set U in the claim to a single vertex v_U to obtain a graph G'. Apply induction to G' to get a set of k disjoint spanning trees $T_1', T_2', \ldots, T_k'$ of G'. Now expand v_U back to U. Each T_i' expands to a spanning tree of G. In this way we get k edge-disjoint spanning trees of G.

Proof of Claim 1

Let $G^0 = (V, E^0)$ and let C_0 be the component of G^0 that contains e^0. Let $U = V(C^0)$. First verify that if $F = (F_1, F_2, \ldots, F_k) \in \mathcal{F}^0$ and F' is obtained from F by a replacement and x, y are the ends of a path in $F_i' \cap U$ then x, y are joined by a path $xF_iy \subseteq U$. (Exercise).

We now show that $F_i^0 \cap U$ is connected. Let (x, y) be an edge of C^0. Since C^0 is connected, we only have to show that F_i^0 contains a path from x to y, all of whose vertices belong to U. But this follows by using the exercise and backwards induction starting from some $F \in \mathcal{F}^0$ for which F_i contains the edge (x, y). □