Game Theory and Algorithms*
Lecture 7: PPAD and Fixed-Point Theorems

March 17, 2011

Summary: The ultimate goal of this lecture is to finally prove Nash's theorem. First, we introduce and prove *Sperner's lemma*. On face value it looks totally unrelated to Nash equilibria, but we discuss how they are equivalent, via PPAD-completeness. Then, we mention another PPAD-complete problem, which is Brouwer's fixed point theorem: it was used in the original proof of Nash's theorem. We focus on a slightly easier proof via Kakutani's fixed point theorem. This also leads to other existence theorems for equilibria.

1 Sperner's Lemma

Sperner's lemma (1928) is a proof that in a certain kind of 3-coloured triangulation, a triangle with all 3 colours exists. The detailed definition is as follows.

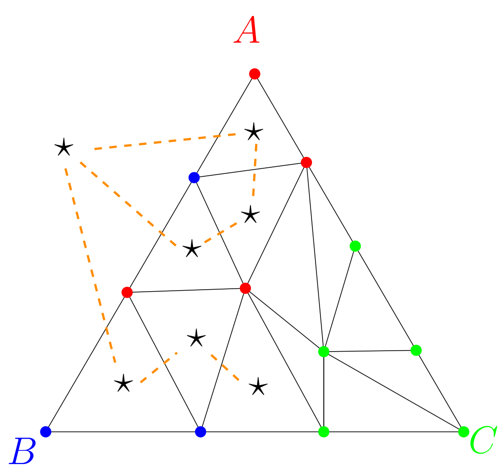

The triangulation. We are given a large triangle ABC; its sides are subdivided, and there are additional vertices in the interior of ABC. We draw edges (black in the diagram) connecting these vertices such that every face of the resulting graph is a small triangle. (No black edges cross).

The colouring. Name our colours a, b, c. The colouring assigns one colour to each vertex of the triangulation, under the following constraints: vertex A gets colour a (and similarly for b, c); and each vertex subdividing AB gets either colour a or colour b (and similarly for the other AC, BC). Interior vertices can get any colour.

Now that we have defined the colouring, we can give Sperner's Lemma and its proof.

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

Lemma 1. *For any triangulation and colouring of the form above, there is a small triangle such that all 3 colours appear on one of its vertices. (A* tricoloured triangle.*)*

Proof. We will define a new graph and apply the following fact to the graph:

Handshake Lemma: the sum of the degrees of the vertices is two times the number of edges

Draw a star inside of every small triangle with at least one a-coloured-vertex and at least one b-coloured-vertex. Also, draw a star on the outside of ABC. These stars form the nodes of a new graph. Whenever an edge of the original graph has one endpoint coloured a and one coloured b, (an $\{a, b\}$*-edge*) note that both faces on the side of this edge have stars: now connect these stars by an edge in the new graph.

The rest of the proof proceeds to analyze the degrees of the vertices in this graph.

- First check the star on the outside of ABC. It has edges crossing the subdivision edges along AB. Moreover, the degree of this star must be odd, since the degree counts the number of changes between colour a and colour b as we walk from A to B along the exterior of the subdivision, and an even degree would imply that A and B get the same colour, which we know is false.

- For stars inside the triangulation there are two cases. First, any small tricoloured triangle has exactly one $\{a, b\}$-edge, so the star inside this triangle has degree 1.

- Next, any small triangle with a star that is *not* tricoloured has all of its vertices coloured a and b, two of one and one of the other. In either case, we see that this star has degree 2.

The handshake lemma implies that the number of stars with odd degree is even. So, the number of tricoloured triangles is odd. In particular, it is nonzero, which finishes the proof. □

You would probably agree that on face value, Sperner's lemma and the problem of computing Nash equilibria don't seem to be related at all. However, they are indeed related. The first hint of this is that the argument in the *proof* of Sperner's lemma is a lot like the proof of termination for the Lemke-Howson algorithm: they rely on an indirect argument having to do with the parity of degree sums in a graph. In fact, more is true: there is an efficient algorithm to find NE of two-player games if and only if there is an efficient algorithm to find tricoloured triangles in exponentially-large implicit graphs. We'll make this more precise as we discuss several topics:

- The definition of the class PPAD
- Formalization of finding Nash equilibria as a PPAD problem
- Formalization of Sperner's lemma as a PPAD problem

Then we will briefly discuss Brouwer's fixed-point theorem. It is also PPAD complete, but for us the more important fact is that we can use it to give a proof of Nash's theorem (for any number of players)!

2 PPAD

Like P and NP, PPAD is a complexity class of computational problems. Its name stands for *polynomial parity argument, directed* and it was invented in 1994 by Papadimitriou. His original paper proved several problems were PPAD-complete[1] (e.g., 3-dimensional versions of Sperner, Brouwer, and Kakutani's theorems). Some other problems (including 2-dimensional versions of the aforementioned, and 2D NE) were observed to be in PPAD, and shown to be complete later on (in 2006 papers by X. Chen and X. Deng).

2.1 Definition

The prototypical problem in PPAD (called "end of the line") is the following. You are given an *implicit directed graph* of size 2^n in the following sense: the vertices are all n-bit binary strings, and you have a poly(n)-time algorithm (or Turing machine) whose input is a vertex and whose output lists all its out-neighbours and in-neighbours. We require that every node of the graph so defined has in-degree at most 1, and out-degree at most 1. Moreover, the input to an instance of this problem must include the label of a *source vertex* which has in-degree 0 and out-degree 1. (So the input is n, the Turing machine, and the source.) The correct output for an instance of this problem is any *other* vertex whose degree (in-degree + out-degree) equals 1. Finally, a problem is in PPAD if it can be solved by a polynomial-time algorithm which makes a single call to a solver for "end of the line." (This is a "Karp reduction.")

Notice that every problem in PPAD has a trivial exponential-time algorithm: just follow the path starting at the source; eventually you have to hit a sink. At the same time, it is hard to imagine that any faster algorithm is possible. (If the graph were arbitrary indeed no faster algorithm would be possible, but in principle the fact that we can look at the "source code" for the polynomial-time Turing machine might help. Whether PPAD = P is an open problem, much like "P vs NP.")

2.2 2-Player NE is in PPAD

The definition of PPAD is very similar to the argument we used to prove correctness and termination of the Lemke-Howson algorithm. Indeed, it would show that the problem of computing mixed NE in 2-player games is in PPAD (e.g. we encode each vertex of P_1 as a $|A_1| + |A_2|$-bit binary string according to the tight inequalities), but there is one obstacle: the Lemke-Howson argument used an *undirected* graph. In fact, there exists an algebraic "orientation rule" which can be used to make the graph directed.

[1] A problem is PPAD-complete if (i) it is in PPAD, and (ii) you can reduce any PPAD problem to it in polynomial time.

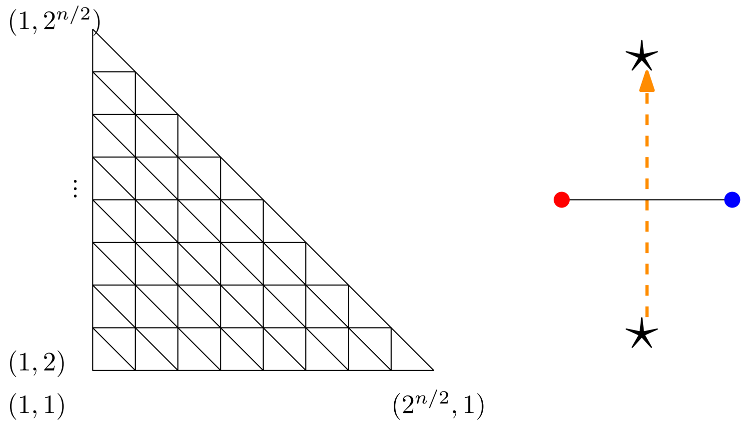

Figure 1: Left: a standard exponential-size triangulation. Right: the orientation of the arcs used to show that the 2D Sperner problem lies in PPAD.

2.3 2D Sperner is in PPAD

We encode 2D Sperner in the following way. An input to 2D Sperner is an *implicit colouring* of a certain standard exponential-size triangulation (see Figure 1), given by a polynomial(n)-time Turing machine which takes the coordinates of a point as input, and outputs the colour of that point. (The colouring must meet the same boundary conditions as specified earlier.) The output for this problem is the set of coordinates of a tricoloured triangle.

Again, aside from an orientation issue, it is not hard to see how to modify the proof we gave of Sperner's Lemma to show that 2D Sperner lies in PPAD. This time the solution to the orientation issue is easier to describe: orient each edge so that, in passing from the tail to the head, the a-coloured vertex is on the left and the b-coloured vertex is on the right. (See Figure 1).

Remark. Why do we go to the trouble of using this complicated Turing machine-based colouring scheme? The reason is the following: if the input is an explicit list of all colours for all vertices then the resulting problem is easily polynomial-time solvable (just check all triangles to find a tricoloured one). In contrast, this implicit-Sperner problem has a very interesting property, which we explain next.

2.4 Equivalence of 2-Player NE and 2D Sperner

Theorem 2 (Chen & Deng, 2006)**.** *There is a polynomial-time algorithm to find Nash equilibria of 2-player games if and only if there is a polynomial-time algorithm to solve the 2D Sperner problem.*

Chen & Deng showed this by showing (in separate papers) that both problems are PPAD-complete. This implies that also the "End of the Line" problem can be added to the list of problems in the theorem statement. The proofs of PPAD-completeness are too technical for our course but we mention some history and related work. The earliest NE-type completeness result was that "ϵ-approximate 4-player Nash equilibria" is PPAD-complete, followed by the 3-player version. Brouwer's fixed-point theorem (see the next section) is also PPAD-complete in any fixed dimension $k \geq 2$.

To define an *ϵ-approximate Nash equilibrium*, assume that all payoffs are between 0 and 1; then a mixed action profile α is an ϵ-approximate Nash equilibrium if, for every α'_i, we have $u_i(\alpha'_i, \alpha_{-i}) \leq u_i(\alpha) + \epsilon$. In other words, there is no very profitable deviation for any player (but small profitable deviations might exist). The PPAD-hardness results hold when ϵ is inverse polynomial or inverse exponential in ϵ; computing approximate equilibria for 2 players has the same hardness. (Aside: the best known algorithms for approximate equilibria are somewhat faster than the best known algorithms for exact equilibria, but still have exponential time complexity when ϵ is inverse polynomial. For $\epsilon = 1/2$ and two players, a linear-time algorithm is known.)

Here is one reason that approximate equilibria arise in our studies: we know that 2-player games have rational equilibria (if all $u_i(a)$ are rational), but k-player games for $k \geq 3$ can have irrational equilibria, and it is computationally difficult just to represent a general equilibrium for the latter game. So the chain of reductions from 3-player NE to 2-player NE naturally requires abandoning exact computation. In addition, many "gadgets" in the reductions are built around the notion of ϵ-approximate equilibria. (Aside: Bubelis in 1979 proved that you can *exactly* reduce any k-player game to a 3-player game, showing that the qualitative distinction only arises between 2-player and 3-player games.)

3 Fixed Points and Proof of Nash's Theorem

Nash's original proof of the existence of mixed equilibria in finite strategic games used the following topological theorem.

Theorem 3 (Brouwer, 1910). *If $S \subset \mathbb{R}^d$ is compact, convex, and nonempty, and $f : S \to S$ is continuous, then there exists $x \in S$ such that $f(x) = x$.*

We call x a *fixed point* of f.

We review the definitions involved in the statement of the theorem:

- S is *compact* means (i) *boundedness*, i.e. that there is some ball of finite radius containing all of S, or equivalently that it does not go off to infinity, and (ii) *closedness*, i.e. if $\{x_i\}_i$ is a convergent sequence of points in S, then $\lim_{i\to\infty} x$ is also in S

- S is *convex* means that whenever points x and y lie in S, then so does the line segment connecting them, i.e. $\{\lambda x + (1-\lambda)y \mid 0 \leq \lambda \leq 1\} \subset S$. Here are two convex shapes: ■, ●; and a star $\star$ is non-convex since the line segment connecting two adjacent tips of the star does not lie entirely inside the star.

Exercise. Prove Brouwer's theorem for $d = 1$.

The general proof of Brouwer's theorem takes us too far into topology, so we will skip it. But, Sperner's lemma (and it's higher-dimensional generalization) is one way of giving a proof, by taking the limit of successively finer and finer discretizations of f.

Exercise. In this exercise we show that the conditions of Brouwer's theorem cannot be weakened. Give a counterexample (S, f) when:

- All conditions are satisfied, except that S is empty.
- All conditions are satisfied, except that S is not bounded.
- All conditions are satisfied, except that S is not closed.
- All conditions are satisfied, except that S is not convex.
- All conditions are satisfied, except that f is not continuous.

We will discuss Nash's original proof later; first we give a simpler proof based on a stronger fixed-point theorem. A *set-function* $f : S \to \mathbf{2}^S$ means a function f whose input is a single point $x \in S$, and whose output is any subset of S. (You can think of it as a function with multiple outputs.) The proper generalization of continuity is that f has a *closed graph* if, whenever convergent sequences $\{x_i\}_i$ and $\{y_i\}_i$ satisfy $x_i \in f(y_i)$ for all i, then $\lim_{i\to\infty} x_i \in f(\lim_{i\to\infty} y_i)$.

Theorem 4 (Kakutani, 1941)**.** *Suppose $S \subset \mathbb{R}^d$ is compact, convex, and nonempty, and $f : S \to \mathbf{2}^S$ has a closed graph. Additionally, assume that for every $x \in S$, we have that $f(x)$ is nonempty and convex. Then there exists $x \in S$ such that $x \in f(x)$.*

Notice that in the special case $|f(x)| = 1$ for all x, this is exactly equivalent to Brouwer's fixed point theorem.

When we use Kakutani's theorem to prove Nash's theorem, we introduce the following notation.

Definition 5. Let T be a finite set. The *simplex* $\triangle^T$ is the polyhedron

$$\{x \mid \forall t \in T : x_t \geq 0; \sum_{t\in T} x_t = 1\}.$$

So we may think of $\triangle^T$ as having $|T|$ variables. We also can view $\triangle^T$ as a subset of $\triangle^U$ whenever $T \subset U$; in this case for all $x \in \triangle^T$ we take $x_u = 0$ for all $u \in U \backslash T$.

Proof of Nash's Theorem. We are given a finite strategic game with its A_i and u_i; we want to find a mixed Nash equilibrium α.

Let $S = \prod_{i=1}^n \triangle^{A_i}$ be the set of all mixed strategy profiles. Define the set-valued function $f : S \to \mathbf{2}^S$ by

$$f(\alpha) := \prod_{i=1}^n \triangle^{B_i(\alpha_{-i})}.$$

In other words, $\beta \in f(\alpha)$ if for all i, $\mathtt{supp}(\beta_i) \subseteq B_i(\alpha_{-i})$. If we can show f satisfies the hypotheses of Kakutani's theorem, then we will be done: the fixed point α is a Nash equilibrium because of the Support Characterization.

It is easy to see $f(\alpha)$ is convex and nonempty for all α. To see that it has a closed graph is technical but straightforward. Consider a sequence $\alpha^{(k)}$ for $k = 1, 2, \dots$ of mixed strategy profiles which converge to α^*, and corresponding $\beta^{(k)} \in f(\alpha^{(k)})$ with limit β^*. Although the support can get smaller in the limit (the limit of positive numbers can be zero), the reverse can never happen. So for all i and sufficiently large k, $\mathtt{supp}(\beta_i^*) \subseteq \mathtt{supp}(\beta_i^{(k)})$. It is relatively straightforward to verify that best responses are preserved in the limit, so for all i and sufficiently large k, $B_i(\alpha_{-i}^{(k)}) = B_i(\alpha_{-i}^*)$. So for all i and large enough k

$$\mathtt{supp}(\beta_i^*) \subseteq \mathtt{supp}(\beta_i^{(k)}) \subseteq B_i(\alpha_{-i}^{(k)}) = B_i(\alpha_{-i}^*),$$

where the middle containment holds since $\beta^{(k)} \in f(\alpha^{(k)})$. Taking the above containments together, we get $\beta^* \in f(\alpha^*)$ as needed. □

Exercise (A class of games with pure equilibria)**.** By mimicking the proof above, prove the following theorem of Debreu-Fan-Glicksburg (1952). The setting is a game where each A_i is a closed convex nonempty subset of $\mathbb{R}^{d_i}$. We require that each u_i is a continuous *quasi-concave* function, meaning that the *level sets* $\{a \mid u_i(a) \geq C\}$ are convex for all i and all $C \in \mathbb{R}$. Prove that this game has a *pure* Nash equilibrium.

Note the Debreu-Fan-Glicksburg theorem immediately implies that the Cournot duopoly game has a pure Nash equilibrium.

Exercise. Prove that the condition of quasi-concavity cannot be removed in the Debreu-Fan-Glicksburg theorem: consider the two-player zero-sum game with $A_1 = A_2 = [-1, 1]$ and $u_1 = a_1 a_2 + a_1^2 - a_2^2$; show it has no pure Nash equilibrium.

We should remark that in the absence of quasi-concavity, such games still admit a *mixed* Nash equilibrium (this is a 1952 theorem of Glicksburg), but a stronger fixed-point theorem is needed (since we take a probability distribution over infinite sets. If we also remove the requirement that u_i is continuous, there may even fail to be a mixed Nash equilibrium (Sion & Wolfe 1957).

Exercise (Symmetric games have symmetric equilibria)**.** A game is *symmetric* if

- $A_i = A_j$ for all players i, j;
- $u_i(a_i, a_{-i}) = u_i(a_i, a'_{-i})$ whenever a_i is a permutation of a'_i; and
- $u_i(a_i, a_{-i}) = u_j(b_j, b_{-j})$ whenever $a_i = b_j$ and a_{-i} is a permutation of b_{-j}.

Show that every such game with $|A_i|$ finite has a mixed Nash equilibrium α with $\alpha_i = \alpha_j$ for all i, j.

3.1 Proof via Brouwer's Theorem

The original proof of Nash's theorem was in his thesis, where he used Brouwer's theorem. We give a similar but somewhat simpler version here. (The simpler proof via Kakutani's theorem was given in the 1950 journal version; Nash credits the observation that this shortcut is possible to David Gale.)

To rewrite the proof in such a way that Brouwer's theorem is used rather than Kakutani's, we need to alter the definition of f in such a way that it has just a *single* output. We want that $f(\alpha) = \alpha$ iff α is a Nash equilibrium. The general idea is that f should push α "towards" best responses, leaving α_i alone if it already only uses best responses. We also require that f is continuous.

To accomplish this, define

$$\alpha_i'(a_i) := \alpha_i(a_i) + \max\{0, u_i(a_i, \alpha_{-i}) - u_i(\alpha)\}$$

so the probability of anything which is a better response than α_i is raised. Then define

$$f(\alpha) := (\texttt{nrml}(\alpha_1'), \texttt{nrml}(\alpha_2'), \ldots, \texttt{nrml}(\alpha_n')).$$

On the one hand, if α_i is a best response to α_{-i}, then $\alpha_i' = \alpha_i$. Otherwise, by the Support Characterization α_i assigns positive probability to a non-best response. We now consider two cases and leave the details as an exercise: one case is that α_i assigns zero probability to some best response; the other case is that $\texttt{supp}(\alpha_i) \supsetneq B_i(\alpha_{-i})$. In either case we can show $\texttt{nrml}(\alpha_i') \neq \alpha_i$ and thus we are done.