

Lecture 17

PPAD-completeness of Nash equilibria (Part I)

Instructor: Prof. Gabriele Farina

As promised, we close this course with a discussion of problems that, from the point of view of its computational complexity, are equivalent to that of computing Nash equilibrium. We have already seen in Lecture 2 that the one can prove existence of Nash equilibrium in general games directly from applying Brouwer’s fixed-point theorem on a continuous function from a compact and convex set to itself. This immediately shows that

Computing Nash equilibrium is not harder than computing Brouwer fixed points.

Yet, no other proof of existence of Nash equilibrium is known that does not rely on Brouwer’s fixed point theorem. Why is that the case? One of the main reasons is the following (very much non-trivial) fact.

The opposite is also true: from a computational point of view, the computation of a Brouwer fixed point is equivalent to computing a Nash equilibrium of some appropriate game of polynomial size.

In other words, Nash equilibria are the *prototypical* example of a Brouwer fixed points!

1 Sperner’s lemma

One of the important merits of complexity theory is that of connecting problems that, despite their disparate looks, are united by their ability to encode one same hard primitive. For example, all NP-complete problem are united by the fact that they are all equivalent methods for checking whether satisfying assignment to a Boolean formula exists.

What is the primitive that unites Brouwer’s fixed points and Nash equilibria? To shed light in this direction, let us consider another problem that, at first glance, has nothing to do with neither fixed points nor equilibria: Sperner’s lemma.




Figure 1: Sample Sperner coloring.

Upon further inspection, it will turn out to be another problem in the same “family” as the previous two.

Sperner’s lemma has to do with colored $N \times N$ grids of points. The rules are simple—each point is colored with one of three colors: red, blue, or yellow. The coloring must however satisfy the following *boundary* conditions:

*These notes are class material that has not undergone formal peer review. The TA and I are grateful for any reports of typos. Some of the content of the lecture was adapted from material from Costis Daskalakis.

- i. The left column cannot contain any blue;
- ii. The bottom row cannot contain any red;
- iii. The right column and top row cannot contain any yellow.

Any coloring that satisfies these conditions is called a *Sperner coloring*. An example of a coloring satisfying these rules is shown in Figure 1.


Given any Sperner coloring, we are interested in finding a trichromatic triangle, that is, a cell whose vertices are colored red, blue, and yellow. Sperner’s lemma guarantees that such a request is always possible to satisfy.

Theorem 1.1 (Sperner’s lemma). Any Sperner coloring must have at least *one* trichromatic triangle. In fact, it must have an *odd* number of trichromatic triangles.

We illustrate the previous theorem on the grid of Figure 1.


Example 1.1. In the grid of Figure 1, there are a total of *five* trichromatic triangles, as highlighted in green on the right.

Indeed, five is an odd number, validating the prediction of Theorem 1.1.



1.1 The connection between Brouwer and Sperner

What does Sperner’s lemma have to do with Brouwer’s fixed point theorem? The connection is not immediate, but upon second thought, several glimpses of connections emerge. For one, both results are nonconstructive existence results. Furthermore, both results are trivially false if the boundary conditions are not satisfied. In the case of Brouwer’s fixed point theorem, the boundary conditions are that the continuous function must map the compact set to itself. In the case of Sperner’s lemma, they are the coloring rules on the boundary. As it turns out, Sperner’s lemma is fundamentally a *discretized version of Brouwer’s fixed point theorem* for continuous functions on $[0, 1]^2$. The colors correspond to a discretization of the direction in which $f(z) - z$ points, according to the coloring rule shown on the left.

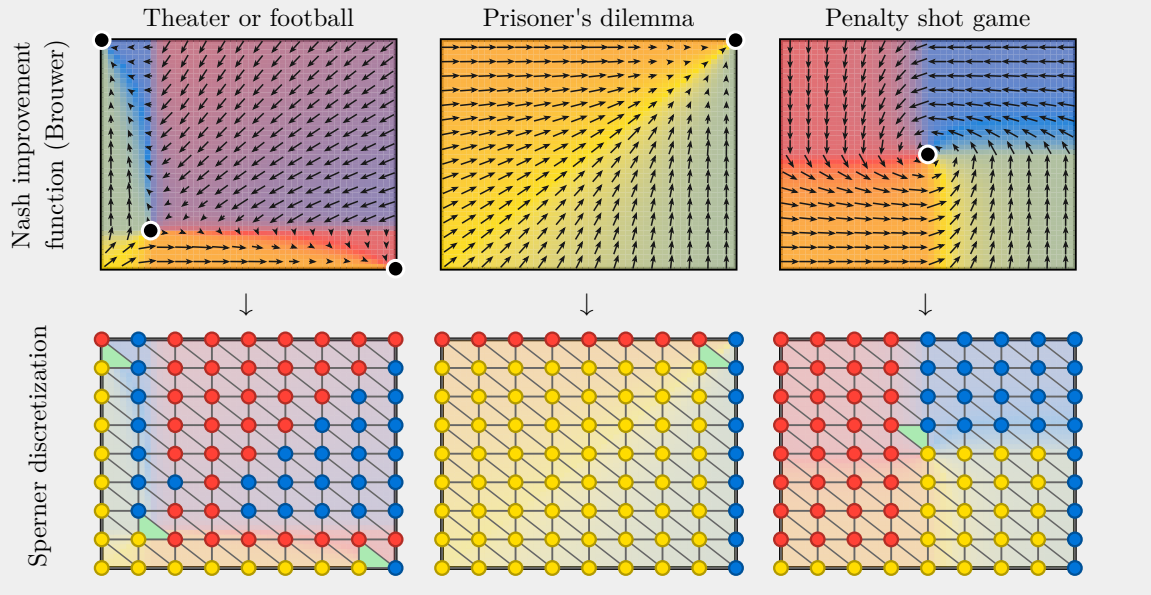


(This was the same coloring we used in Lecture 2, with a bit of foreboding).

Following this connection, it should then be immediate why trichromatic triangles have value: they correspond to cells in which the function points in three different directions, at the three corners of the triangle, which gives hope that—due to continuity—the function must have a point somewhere in the interior of the triangle in which all directions coexist and cancel out leading to a fixed point.

If we had access to an algorithm to find trichromatic triangles, we could use it to find approximate Brouwer fixed points of continuous functions $[0, 1]^2 \rightarrow [0, 1]^2$ by discretizing the space and coloring the points according to the direction in which the function points. The discretization parameter would then determine the precision of the fixed point. We illustrate this with an example.

Example 1.2. The following plots illustrate the Sperner discretization of the Nash improvement function in three games we used in Lecture 2.



1.2 Formalizing the connection

To make the argument formal, we need to establish a formal connection between a trichromatic triangle and an approximate Brouwer fixed point, and connect that to a choice of discretization parameter.

By the Heine-Cantor theorem, any f continuous on a compact set is *uniformly* continuous, meaning that

$$\forall \epsilon > 0, \exists \delta(\epsilon) > 0 : \|z - w\|_\infty \leq \delta(\epsilon) \implies \|f(z) - f(w)\|_\infty \leq \epsilon.$$

Consider now a Sperner discretization in which the diameter of every cell is at most $\delta \leq \delta(\epsilon)$. Then, the following approximation bound can be established.

Theorem 1.2. If z_Y is the yellow corner of a trichromatic triangle in a Sperner discretization with cell diameter $\delta \leq \delta(\epsilon)$, then

$$\|f(z_Y) - z_Y\|_\infty < \epsilon + \delta.$$

Proof. Let z_R, z_B , and z_Y be the red, blue, and yellow vertices of the trichromatic triangle. The key observation is that, by the coloring rule,

$$\begin{aligned} (f(z_Y) - z_Y)_x & \text{ and } (f(z_B) - z_B)_x & \text{ have opposite signs, and} \\ (f(z_Y) - z_Y)_y & \text{ and } (f(z_R) - z_R)_y & \text{ have opposite signs.} \end{aligned}$$

Thus, we can write

$$\begin{aligned} |(f(z_Y) - z_Y)_x| & \leq |(f(z_Y) - z_Y)_x - (f(z_B) - z_B)_x| \\ & \leq |(f(z_Y) - f(z_B))_x - (z_Y - z_B)_x| \\ & \leq \|f(z_Y) - f(z_B)\|_\infty + \|z_Y - z_B\|_\infty < \epsilon + \delta. \end{aligned}$$

and similarly

$$\begin{aligned}
|(f(z_Y) - z_Y)_y| &\leq |(f(z_Y) - z_Y)_y - (f(z_R) - z_R)_y| \\
&\leq |(f(z_Y) - f(z_R))_y - (z_Y - z_R)_y| \\
&\leq \|f(z_Y) - f(z_R)\|_\infty + \|z_Y - z_R\|_\infty < \epsilon + \delta.
\end{aligned}$$

From here, we can just use the definition of infinity norm:

$$\|f(z_Y) - z_Y\|_\infty = \max\{|(f(z_Y) - z_Y)_x|, |(f(z_Y) - z_Y)_y|\} < \epsilon + \delta. \quad \square$$

Corollary 1.1. Let $\epsilon > 0$ be an arbitrary precision. The Sperner discretization of the function with parameter $\delta := \min\{\delta(\epsilon), \epsilon\}$ is a 2ϵ -approximate Brouwer fixed point.

While this is not necessary for our discussion of the complexity of equilibrium computation, we remark that Corollary 1.1 immediately implies, using a standard compactness argument, Brouwer's fixed point theorem for the two-dimensional (2D) case.

Corollary 1.2 (Brouwer's fixed point theorem, 2D case). Any continuous function $[0, 1]^2 \rightarrow [0, 1]^2$ has a fixed point.

Proof. Consider the sequence of approximations $\epsilon_i := 2^{-i}$ for $i \in \mathbb{N}_{\geq 1}$, and the corresponding discretization parameters $\delta_i := \min\{\delta(\epsilon_i), \epsilon_i\}$. For each i , we can isolate a yellow vertex $z_{Y,i}$. Since $z_{Y,i} \in [0, 1]^2$, and $[0, 1]^2$ is a compact set, there exists a convergent subsequence $z_{Y,j}$; let z_Y denote the limit of such a subsequence. By the continuity of f , the function $d(z) := \|f(z) - z\|_\infty$ is also continuous. Hence,

$$d(z_Y) = \lim d(z_{Y,j}).$$


Since $d(z_{Y,j}) \in [0, 2 \cdot 2^{-j}]$ by Corollary 1.1, we conclude $d(z_Y) = 0$, which is equivalent to $f(z_Y) = z_Y$. This proves that a fixed point exists. \square


2 Proof of Sperner's lemma

How can one prove Sperner's lemma? As it turns out, the lemma can be restated as a pretty simple property of graphs. By following this process, we will be able to shed more light into the nonconstructive nature of the proof of existence.

Before jumping into the proof, let's operate a simplification. Without loss of generality, we will assume that the boundary of the Sperner coloring is as in the figure on the right: red on the left (except for the bottom-left corner), yellow on the bottom (except for the bottom-right corner), and blue everywhere else. We will call any such instance a *standard* Sperner coloring.

The previous assumption is without loss of generality. Indeed, if the boundary of the instance does not respect the assumption (e.g., Figure 1), we can always pad the grid with a boundary that satisfies the assumption, and embed the non-respecting grid in the inside. This will not change the number of positions of the trichromatic triangles.





The proof of Sperner's lemma is based on a graph-theoretic argument. We can convert any standard Sperner coloring into a *Sperner graph* as follows:

- Each cell is converted into a node;
- Two neighboring cells u, v (nodes) are connected by a directed edge $u \rightarrow v$ if to go from cell u to cell v one passes through a *red-yellow* portal, *i.e.*, a segment that has a red color on the left and a yellow color on the right.

For the standard Sperner coloring of Figure 1, the corresponding graph is shown just above on the left.

2.1 Properties of the Sperner graph

As you might have guessed from the picture, the following key properties hold.

Theorem 2.1. In any Sperner graph, the following properties hold:

- (1) every node has outdegree and indegree at most 1;
- (2) any node with indegree 1 and outdegree 0 is a trichromatic triangle (marked green in the figure above);
- (3) any node with outdegree 1 and indegree 0 is a trichromatic triangle (marked green), with the only exception of the bottomleft node (marked purple).

Proof. Property (1) follows from case analysis.

We prove property (2) by contradiction. Take any node with indegree 1 and outdegree 0, and assume for contradiction that it is not a trichromatic triangle. Since the indegree is 1, one of the sides of the cell corresponding to the node is a red-yellow gate. Let's consider now the third vertex of the cell. Since by assumption the cell is not trichromatic, the third vertex is either red or yellow. Either case produces another red-yellow gate, so the cell must be on the boundary of the Sperner coloring. A simple case analysis reveals that it is impossible that such a cell exists given the boundary conditions.

A similar argument can be made for property (3). □

At this point, the proof of Sperner's lemma is immediate. A graph in which each node has indegree at most one and outdegree at most one is composed of connected components that can only be singleton nodes, lines, or simple cycles. Only lines have nodes with outdegree 1 and indegree 0, or outdegree 0 and indegree 1; each has exactly one of each. Furthermore, the bottomleft node is part of a line, and is the only node with outdegree 1 and indegree 0 that is not a trichromatic triangle. Hence, there are an odd number of trichromatic triangles in any Sperner coloring.

3 The PPAD complexity class

The proof of Sperner's lemma seen above distilled computing approximate Brouwer fixed points into computing any end of a path in the Sperner graph. This is the essence of the PPAD complexity class: a problem is in PPAD if it can be reduced to finding the end of a line in a directed graph in which every node has indegree and outdegree at most 1, and a given start of a line is given.

As just stated, the problem might seem kind of trivial: after all, why can't we just follow the line until the end? The problem is in the size of Sperner discretization that we need. Assuming the function for which we need an ϵ -approximate Brouwer fixed point is 1-Lipschitz continuous, the size N of the Sperner coloring

we need to consider is roughly of order $1/\epsilon$, where ϵ is the desired precision. This quantity is *exponentially large* in the representation of ϵ specified in the input, which only takes $\log(1/\epsilon)$ bits to store.

To capture the above, we can define the following problem. Suppose that a directed graph with vertex set $\{0, 1\}^n$ is described by a pair of circuits P and N , which both have n -bit inputs and n -bit outputs, as follows.

- P takes as input a node $v \in \{0, 1\}^n$ and outputs another node $u \in \{0, 1\}^n$ that is a “possible previous” node for v .
- N takes as input a node $v \in \{0, 1\}^n$ and outputs a node $u \in \{0, 1\}^n$ that is a “possible next” node for v .
- The graph defined by P and N has an edge from node $v_1 \in \{0, 1\}^n$ to node $v_2 \in \{0, 1\}^n$ whenever

$$P(v_2) = v_1 \wedge N(v_1) = v_2.$$

That is, an edge is established between the nodes if they agree on their next/previous relationship.

No matter what P and N are, the graph defined by them has, every node has in-degree and out-degree at most one. Hence, the graph comprises directed paths, directed cycles, and isolated vertices just like in the proof of Sperner’s lemma above. We define the following search problem on such graphs:

Definition 3.1 (END-OF-THE-LINE). Given P and N , if 0^n is an unbalanced node, *i.e.*, if it has different in- and out-degrees, find another unbalanced node. Otherwise output 0^n .

The PPAD complexity class is then the class of all search problems that can be reduced efficiently to END-OF-THE-LINE. To give mathematical meaning to this statement, we formalize a few concepts.

3.1 The complexity class FNP

The Functional NP (FNP) complexity class, which is the class of all search problems with short, efficiently verifiable solutions. In particular, PPAD is a subclass of TFNP, which is the subset of FNP problems that are *total*, that is, always admit a solution.

Definition 3.2. A *search problem* L is defined by a relation $R_L \subseteq \{0, 1\}^* \times \{0, 1\}^*$ with the convention that $(x, y) \in R_L$ iff y is a solution to x . A search problem is called *total* iff for all x there exists y such that $(x, y) \in R_L$.

Definition 3.3 (FNP class). A search problem L belongs to the class FNP if there exists a polynomial-time algorithm A_L (the “verifier”) and a polynomial function p_L such that:

- The verifier is correct, that is, whenever it says that y is a solution to x , y must truly be a solution to x , and vice versa. In symbols, for all $x, y : A_L(x, y) = 1 \iff (x, y) \in R_L$; and
- If there is a solution to x , there must be a “short” solution to x , meaning polynomially long in $|x|$. In symbols, for all x : if $\exists y$ s.t. $A_L(x, y) = 1$ then $\exists z$ with $|z| < p_L(|x|)$ such that $A_L(x, z) = 1$, where $|x|, |z|$ represent respectively the lengths of x and z as bit strings.

Clearly END-OF-THE-LINE \in FNP, as we can check efficiently if a node is an end of a line by verifying its indegree and outdegree.

Next, recall the concept of polynomial-time reductions between problems. Intuitively, when a problem L is polynomial-time reducible to problem L' , this means that L' is “at least as hard as” L since, given an efficient algorithm for L' , composing this algorithm with the reduction gives an efficient algorithm for L . In particular, given an input x to L , we can efficiently compute the input $f(x)$ to L' . Then we can use the

efficient algorithm for L' to find a solution y to $f(x)$. Finally, we can efficiently compute $g(y)$, the solution to x for problem L .

SAT, the problem of finding satisfying truth assignments for Boolean formulas, is an example of an FNP-complete problem. With this notion of completeness in mind, we can now formally define the PPAD complexity class.

Definition 3.4. PPAD is the class of search problems in FNP that are polynomial-time reducible to END-OF-THE-LINE.


3.2 Relationship between PPAD and other classes

So far we have polynomial-time reductions from Nash equilibria to END-OF-THE-LINE. We also know that END-OF-THE-LINE \in FNP. Given the absence of efficient algorithms for Nash equilibria, it is natural to ask whether END-OF-THE-LINE is FNP-complete. With our current notion of reduction the answer is *no*, because END-OF-THE-LINE always has a solution whereas a SAT instance may not.

To circumvent this, suppose that we tried updating our notion of reduction to be less stringent as follows. For a reduction from SAT to END-OF-THE-LINE suppose that we required the END-OF-THE-LINE solution to inform us whether the SAT instance is satisfiable or not, and provide us with a satisfying assignment in the “yes” case.

This too is problematic because, if such a reduction existed, it could be turned into a non-deterministic algorithm for the problem NON-SATISFIABILITY: given a Boolean formula we would find the END-OF-THE-LINE instance it reduces to; this would inform us about whether the Boolean formula is satisfiable or not. Having a non-deterministic algorithm for NON-SATISFIABILITY would imply $NP = co-NP$, which is believed to be false.

The believed location of PPAD is shown on the right.



MIT OpenCourseWare
<https://ocw.mit.edu>

6.S890 Topics in Multiagent Learning
Fall 2024

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>