COMS E6998: Algorithmic Game Theory and Economics — Feb 17,2010

Lecture 5 – Approximation of Nash Equilibria

Instructor: *Xi Chen* — Scribes: *Yundi Zhang*

1 Brouwer's fixed point theorem

We are going to show Brouwer's fixed point theorem for triangles:

Theorem 1. *Let $\triangle$ be a triangle in $\mathbb{R}^2$. Then every continuous function f from $\triangle$ to itself has a fixed point $\mathbf{x} \in \triangle$ such that $f(\mathbf{x}) = \mathbf{x}$.*

Proof. For convenience, we will prove the theorem for the following particular triangle $\triangle$:

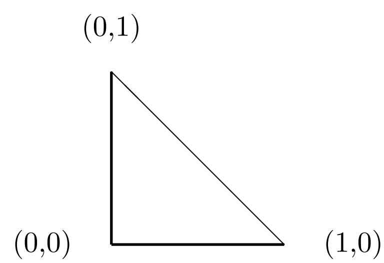

Without loss of generality, we assume the three vertices of $\triangle$, $(0,0), (1,0)$ and $(0,1)$, are not fixed point; otherwise we are already done.

The plan of the proof is as follows. For every $i > 1$, we let S_i denote the standard i-th triangulation of $\triangle$. Given f, we will define carefully a 3-coloring C_i over the vertices of S_i for every $i > 1$. We will then show that C_i is a proper coloring for all i and thus, by Sperner's lemma, has at least one trichromatic triangle, denoted by $\triangle_i$. As i goes up, this sequence of trichromatic triangles $\{\triangle_i\}$ becomes smaller and smaller. We use $\mathbf{u}_i, \mathbf{v}_i, \mathbf{w}_i \in \mathbb{R}^2$ to denote the red, blue, green vertex of $\triangle_i$, respectively. Because $\{\mathbf{u}_i\}$ is an infinite sequence of points in $\triangle$, it must have a converging subsequence $\{\mathbf{u}_{i_j}\}_{j\geq 1}$. We finish the proof by showing that the limit of $\{\mathbf{u}_{i_j}\}_{j\geq 1}$ must be a fixed point of f.

Definition of C_i over S_i:

Let $\mathbf{v} = (v_1, v_2)$ be any vertex of S_i, where v_1 is the x-coordinate of $\mathbf{v}$ and v_2 is the y-coordinate of $\mathbf{v}$. We use $f(\mathbf{v}) = (f_1(\mathbf{v}), f_2(\mathbf{v}))$ to color $\mathbf{v}$ as follows:

1. If $f_2(\mathbf{v}) < v_2$, set $C_i(\mathbf{v})$ to be red;
2. If $f_2(\mathbf{v}) \geq v_2$ and $f_1(\mathbf{v}) < v_1$, set $C_i(\mathbf{v})$ to be blue;
3. Otherwise (when $f_2(\mathbf{v}) \geq v_2$ and $f_1(\mathbf{v}) \geq v_1$), set $C_i(\mathbf{v})$ to be green.

By following these rules (and using the assumption that the three vertices of $\triangle$ are not fixed point of f), it is easy to check that the color of $(0,1)$ must be red; the color of $(1,0)$ must be blue; the color of $(0,0)$ must be green; and moreover, C_i is a proper 3-coloring over S_i.

Then by Sperner's lemma, we know that C_i has at least one trichromatic triangle $\triangle_i$, and we denote its red, blue and green vertex by $\mathbf{u}_i, \mathbf{v}_i$ and $\mathbf{w}_i \in \triangle$, respectively. In this way, we get an infinite sequence of points $\{\mathbf{u}_i\}$ in $\triangle$. Because $\triangle$ is clearly bounded and closed, $\{\mathbf{u}_i\}$ must have a converging subsequence $\{\mathbf{u}_{i_j}\}_{j\geq 1}$, with indices $i_1 < i_2 < \cdots$. We use $\mathbf{x} \in \triangle$ to denote its limit.

x is a fixed point of f:

First of all, since the size of $\triangle_i$ decreases strictly as i goes up. It is easy to show that both sequences $\{\mathbf{v}_{i_j}\}_{j\geq 1}$ and $\{\mathbf{w}_{i_j}\}_{j\geq 1}$ converge to $\mathbf{x}$.

Because every point $\mathbf{u}$ in the sequence $\{\mathbf{u}_{i_j}\}$ is red, by the coloring rule, we have

$$f_2(\mathbf{u}) < u_2, \qquad \text{for all } \mathbf{u} \text{ in the sequence.}$$

Since $\mathbf{x}$ is the limit of the sequence $\{\mathbf{u}_{i_j}\}$ and f is continuous, we have $f_2(\mathbf{x}) \leq x_2$.

Similarly, because every point $\mathbf{v}$ in the sequence $\{\mathbf{v}_{i_j}\}$ is blue, by the coloring rule, we have

$$f_1(\mathbf{v}) < v_1, \qquad \text{for all } \mathbf{v} \text{ in the sequence.}$$

Since $\mathbf{x}$ is the limit of the sequence $\{\mathbf{v}_{i_j}\}$ and f is continuous, we have $f_1(\mathbf{x}) \leq x_1$.

Finally, because every point $\mathbf{w}$ in the sequence $\{\mathbf{w}_{i_j}\}$ is green, by the coloring rule, we have

$$f_1(\mathbf{w}) \geq w_1 \quad \text{and} \quad f_2(\mathbf{w}) \geq w_2, \qquad \text{for all } \mathbf{w} \text{ in the sequence.}$$

Since $\mathbf{x}$ is the limit of the sequence $\{\mathbf{w}_{i_j}\}$ and f is continuous, we have $f_1(\mathbf{x}) \geq x_1$ and $f_2(\mathbf{x}) \geq x_2$.

Combining all four inequalities, we conclude that $\mathbf{x}$ is a fixed point of f. □

2 Existence of Rational Nash Equilibria in Two-Player Games

Given a two-player game $G = (\mathbf{A}, \mathbf{B})$, where every entry of $\mathbf{A}$ and $\mathbf{B}$ is rational, does it always have a rational Nash equilibrium $(\mathbf{x}, \mathbf{y})$ in which all the entries are rational numbers? In Nash's paper, he gave a very simple three-player rational game with no rational Nash equilibrium. However, we will prove the following theorem:

Theorem 2. *Every rational two-player game* $G = (\mathbf{A}, \mathbf{B})$ *has a rational Nash Equilibrium* $(\mathbf{x}, \mathbf{y})$*. Moreover, the number of bits needed to describe* $(\mathbf{x}, \mathbf{y})$ *is polynomial in the input size of* G*.*

Proof. We will use the following property: If a rational linear program has a solution, then it always has a rational solution. Furthermore, the number of bits needed to describe it is polynomial in the input size of the linear program.

We use the idea of **Support Enumeration**. Given vectors $\mathbf{x} \in \Delta_m$ and $\mathbf{y} \in \Delta_n$, we let

$$\text{Supp}(\mathbf{x}) = \{i | x_i > 0\} \subseteq [m] \quad \text{and} \quad \text{Supp}(\mathbf{y}) = \{j | y_j > 0\} \subseteq [n].$$

Now for any pair of nonempty subsets $S \subseteq [m]$ and $T \subseteq [n]$, we let $\text{LP}(S, T)$ denote the following linear program with variables $x_1, \ldots, x_m, y_1, \ldots, y_n$:

$$
\begin{aligned}
&\mathrm{LP}(S,T):\\
&\sum_{i\in[m]} x_i = 1\\
&x_i > 0 \quad \text{for all } i \in S\\
&x_i = 0 \quad \text{for all } i \notin S\\
&\sum_{j\in[n]} y_j = 1\\
&y_j > 0 \quad \text{for all } j \in T\\
&y_j = 0 \quad \text{for all } j \notin T\\
&\mathbf{A}_{i,*}\mathbf{y} \geq \mathbf{A}_{j,*}\mathbf{y} \quad \text{for all } i \in S \text{ and } j \in [m]\\
&\mathbf{x}^T\mathbf{B}_{*,i} \geq \mathbf{x}^T\mathbf{B}_{*,j} \quad \text{for all } i \in T \text{ and } j \in [n]
\end{aligned}
$$

By the definition of Nash equilibria, it is easy to prove the following property:

Property 3. *The linear program $LP(S,T)$ has a solution $(\mathbf{x},\mathbf{y})$ if and only if $(\mathbf{x},\mathbf{y})$ is a Nash equilibrium of $G = (\mathbf{A},\mathbf{B})$ such that*

$$\mathrm{Supp}(\mathbf{x}) = S \quad \textit{and} \quad \mathrm{Supp}(\mathbf{y}) = T.$$

Now by Nash's theorem, we know $G = (\mathbf{A},\mathbf{B})$ has a Nash equilibrium $(\mathbf{x}^*,\mathbf{y}^*)$ (even though it may not be rational). Let $S = \mathrm{Supp}(\mathbf{x}^*)$ and $T = \mathrm{Supp}(\mathbf{y}^*)$, then by the property above, $(\mathbf{x}^*,\mathbf{y}^*)$ is a solution to the linear program $\mathrm{LP}(S,T)$. As a result, we know that $\mathrm{LP}(S,T)$ has a rational solution which we denote by $(\mathbf{x},\mathbf{y})$, and the number of bits needed to describe $(\mathbf{x},\mathbf{y})$ is polynomial in the input size of G. Using the property above again, $(\mathbf{x},\mathbf{y})$ must also be an equilibrium of G, and the theorem follows. $\square$

3 Approximation of Nash Equilibrium

We combine support enumeration with the probabilistic method to give an approximation algorithm for Nash equilibria. In this section, we always assume that the entries of $\mathbf{A}$ and $\mathbf{B}$ are between 0 and 1. For simplicity, we also assume that both matrices are n-by-n.

Definition 4. *Given $G = (\mathbf{A},\mathbf{B})$ with $A_{i,j}, B_{i,j} \in [0,1]$ for all $i,j \in [n]$, we say $(\mathbf{x},\mathbf{y})$ is an ϵ-approximate Nash equilibrium for some $\epsilon > 0$ if*

$$
\begin{aligned}
\mathbf{x}^T\mathbf{A}\mathbf{y} &\geq (\mathbf{x}')^T\mathbf{A}\mathbf{y} - \epsilon, \qquad \textit{for all } \mathbf{x}' \in \Delta_n;\\
\mathbf{x}^T\mathbf{B}\mathbf{y} &\geq \mathbf{x}^T\mathbf{B}\mathbf{y}' - \epsilon, \qquad \textit{for all } \mathbf{y}' \in \Delta_n.
\end{aligned}
$$

We say a probability distribution $\mathbf{x}$ is k-uniform, for some $k \geq 1$, if every entry x_i of $\mathbf{x}$ is a multiple of $1/k$. One way to interpret k-uniform distributions is to imagine that there are k balls numbered from 1 to k; and there are n bins numbered from 1 to n. The balls are then tossed arbitrarily into the bins.

Every possible result gives a k-uniform distribution:

$$x_i = \frac{\text{\# of balls in bin } i}{k}, \qquad \text{for all } i \in [n].$$

We will use the probability method to prove the following theorem:

Theorem 5. *Given any two-player game $G = (\mathbf{A}, \mathbf{B})$ with $A_{i,j}, B_{i,j} \in [0,1]$ for all $i, j \in [n]$, it has an ϵ-approximate Nash Equilibrium $(\mathbf{x}, \mathbf{y})$ in which both $\mathbf{x}$ and $\mathbf{y}$ are k-uniform distributions with*

$$k = \frac{100 \cdot \ln n}{\epsilon^2}.$$

This theorem gives us the following algorithm to compute an ϵ-approximate Nash equilibrium:

1. Enumerate all pairs of k-uniform probability distributions $(\mathbf{x}, \mathbf{y})$;
2. Output $(\mathbf{x}, \mathbf{y})$ if it is an ϵ-approximate Nash equilibrium.

The correctness of this (support enumeration) algorithm (that it always outputs an ϵ-approximate Nash equilibrium of G) follows from the theorem above. It is also easy to see that its time complexity is

$$n^{O(k)} = n^{O(\ln n/\epsilon^2)}.$$

Proof Sketch of Theorem 5. First of all, by Nash's theorem, G has an equilibrium $(\mathbf{x}, \mathbf{y})$ with $\mathbf{x}, \mathbf{y} \in \Delta_n$.

We use the probabilistic method. To this end, we randomly pick a k-uniform probability distribution $\widehat{\mathbf{x}}$ as follows:

1. For each $i \in [k]$, independently put ball i into bin j, $j \in [n]$, with probability x_j;
2. Then set $\widehat{x}_j = \text{\# of balls in bin } j \big/ k$ for all $j \in [n]$.

We also randomly and independently pick a k-uniform distribution $\widehat{\mathbf{y}} \in \Delta_n$ using $\mathbf{y}$ in the same way.

To prove there exists an ϵ-approximate Nash equilibrium in which both distributions are k-uniform, it suffices to show that

$$\Pr\Big[(\widehat{\mathbf{x}}, \widehat{\mathbf{y}}) \text{ is an } \epsilon\text{-approximate Nash equilibrium of } G = (\mathbf{A}, \mathbf{B})\Big] > 0.$$

By the definition of ϵ-approximate equilibria, it suffices to show that

$$\Pr\begin{pmatrix} \forall i,\ \mathbf{A}_{i,*}\widehat{\mathbf{y}} \le \widehat{\mathbf{x}}^T\mathbf{A}\widehat{\mathbf{y}} + \epsilon \\ \forall j,\ \widehat{\mathbf{x}}^T\mathbf{B}_{*,j} \le \widehat{\mathbf{x}}^T\mathbf{B}\widehat{\mathbf{y}} + \epsilon \end{pmatrix} > 0.$$

This will follow directly if we can prove that: For every $i \in [n]$,

$$\Pr\Big[\mathbf{A}_{i,*}\widehat{\mathbf{y}} - \widehat{\mathbf{x}}^T\mathbf{A}\widehat{\mathbf{y}} > \epsilon\Big] < \frac{1}{n^{20}} \tag{1}$$

and for every $j \in [n]$,

$$\Pr\Big[\widehat{\mathbf{x}}^T\mathbf{B}_{*,j} - \widehat{\mathbf{x}}^T\mathbf{B}\widehat{\mathbf{y}} > \epsilon\Big] < \frac{1}{n^{20}}. \tag{2}$$

To prove (1), we rewrite it as

$$\Pr\left[(\mathbf{A}_{i,*}\widehat{\mathbf{y}} - \mathbf{A}_{i,*}\mathbf{y}) + (\mathbf{A}_{i,*}\mathbf{y} - \mathbf{x}^T\mathbf{A}\mathbf{y}) + (\mathbf{x}^T\mathbf{A}\mathbf{y} - \widehat{\mathbf{x}}^T\mathbf{A}\mathbf{y}) + (\widehat{\mathbf{x}}^T\mathbf{A}\mathbf{y} - \widehat{\mathbf{x}}^T\mathbf{A}\widehat{\mathbf{y}}) > \epsilon\right]. \tag{3}$$

Because $(\mathbf{x}, \mathbf{y})$ is a Nash equilibrium, it follows from the definition that $(\mathbf{A}_{i,*}\mathbf{y} - \mathbf{x}^T\mathbf{A}\mathbf{y}) \le 0$. As a result, the probability in (3) is upper bounded by

$$\begin{aligned}
&\Pr\left[\mathbf{A}_{i,*}\widehat{\mathbf{y}} - \mathbf{A}_{i,*}\mathbf{y} > \epsilon/3 \;\text{ or }\; \mathbf{x}^T\mathbf{A}\mathbf{y} - \widehat{\mathbf{x}}^T\mathbf{A}\mathbf{y} > \epsilon/3 \;\text{ or }\; \widehat{\mathbf{x}}^T\mathbf{A}\mathbf{y} - \widehat{\mathbf{x}}^T\mathbf{A}\widehat{\mathbf{y}} > \epsilon/3\right] \\
&\le \Pr\left[\mathbf{A}_{i,*}\widehat{\mathbf{y}} - \mathbf{A}_{i,*}\mathbf{y} > \epsilon/3\right] + \Pr\left[\mathbf{x}^T\mathbf{A}\mathbf{y} - \widehat{\mathbf{x}}^T\mathbf{A}\mathbf{y} > \epsilon/3\right] + \Pr\left[\widehat{\mathbf{x}}^T\mathbf{A}\mathbf{y} - \widehat{\mathbf{x}}^T\mathbf{A}\widehat{\mathbf{y}} > \epsilon/3\right].
\end{aligned}$$

We bound the first term using the Hoeffding inequality. For each $\ell \in [k]$, let z_ℓ denote the following random variable: $z_\ell = A_{i,j}$, if ball ℓ is tossed into bin j (when generating $\widehat{\mathbf{y}}$ randomly from $\mathbf{y}$), which happens with probability y_j. As a result,

$$E(z_\ell) = \sum_{j\in[n]} A_{i,j} \cdot y_j = \mathbf{A}_{i,*}\mathbf{y} \quad \text{and} \quad E(z_1 + \cdots + z_k) = k \cdot \mathbf{A}_{i,*}\mathbf{y}.$$

We also have

$$z_1 + \cdots + z_k = \sum_{j\in[n]} A_{i,j} \cdot \#\text{ of balls in bin } j = k \cdot \mathbf{A}_{i,*}\widehat{\mathbf{y}}.$$

By using the Hoeffding inequality, we have

$$\Pr\left[\mathbf{A}_{i,*}\widehat{\mathbf{y}} - \mathbf{A}_{i,*}\mathbf{y} > \epsilon/3\right] = \Pr\left[(z_1 + \cdots + z_k) - E(z_1 + \cdots + z_k) > k \cdot \epsilon/3\right] \le e^{-\frac{2k\epsilon^2}{9}} = e^{-\frac{200\ln n}{9}} \ll \frac{1}{n^{20}}.$$

We can see that the probability of event $\left[\mathbf{A}_{i,*}\widehat{\mathbf{y}} - \mathbf{A}_{i,*}\mathbf{y} > \epsilon/3\right]$ is pretty small. Similarly, one can use Hoeffding inequality to prove that both

$$\Pr\left[\mathbf{x}^T\mathbf{A}\mathbf{y} - \widehat{\mathbf{x}}^T\mathbf{A}\mathbf{y} > \epsilon/3\right] \quad \text{and} \quad \Pr\left[\widehat{\mathbf{x}}^T\mathbf{A}\mathbf{y} - \widehat{\mathbf{x}}^T\mathbf{A}\widehat{\mathbf{y}} > \epsilon/3\right]$$

are small. This proves (1). The same argument can be used to prove (2), and the theorem follows. $\square$