CS 354: Unfulfilled Algorithmic Fantasies	10 Apr 2019

Lecture 4: Communication Complexity of ϵ-Nash Equilibrium

Lecturer: Aviad Rubinstein *Scribe: Brian Zhang*

1 Introduction and Review

As a refresher, we will first review some definitions.

Definition 1.1. Suppose Alice and Bob play a two-player finite game. If $A, B \in [-1,1]^{N\times N}$ are the payoff matrices of Alice and Bob respectively, and $x, y \in \Delta(N)$ are probability distributions over the N actions available to the players (which should be interpreted as mixed strategies for Alice and Bob respectively), then we say that (x, y) is an *ϵ-Nash equilibrium* of the game if, for every pure strategy (unit vector) $e \in \Delta(N)$,

(1) $x^T Ay \geq e^T Ay - \epsilon$; and

(2) $x^T By \geq x^T Ae - \epsilon$.

Since $x^T Ay$ and $x^T By$ are the payoffs for Alice and Bob respectively if they play mixed strategies x and y, this condition states that neither Alice nor Bob can increase their payoffs by more than ϵ by switching to any pure strategy.

Definition 1.2. For $\epsilon > 0$, the *ϵ-Nash equilibrium (ϵ-NE) problem* is as follows: given two matrices $A, B \in [-1,1]^{N\times N}$ representing payoff matrices for Alice and Bob respectively on some game, output a ϵ-Nash equilibrium of the game. In the query version of the problem, A and B are not given directly to the algorithm, but the algorithm is given query access to the entries of A and B. In the communication version, Alice is given A, and Bob is given B.

Last time, we stated the following theorem:

Theorem 1.3 (Theorem 1 in [**?**]). *There is a constant $\epsilon > 0$ for which the communication complexity of ϵ-NE is $\Omega(N^{2-o(1)})$.*

Today's lecture will focus on a proof sketch of the following weakened version of the theorem. We will take many liberties (i.e. cheat) throughout the lecture, which we will attempt to point out as they come along.

Theorem 1.4 (Main Theorem of [1]). *There are constants $\epsilon, \delta > 0$ for which the communication complexity of ϵ-NE is $\Omega(N^\delta)$.*

We first take some time to review the end-of-line problem and Brouwer fixed point theorem, and related query and communication complexity results.

Definition 1.5. The *end-of-a-line* problem is the following: given a base graph $G = (V, E)$, a successor function $S : V \to V$, and a predecessor function $P : V \to V$ (both of which obey the edges of the base graph), and a start vertex v_0 such that $P(v_0) = v_0$, find the v^* at the end of a line; i.e. a vertex v^* such that $S(v^*) = v^*$. In the query version, the algorithm is given the base graph G and vertex v_0; and query access to the functions S and P.

Lemma 1.6. *The end-of-a-line problem takes $\Omega(|V|)$ queries in the worst case.*

A brief argument for this lemma was given in Lecture 2.

Definition 1.7. The *ϵ-Brouwer fixed point problem* is the following: given a function $f : [0,1]^n \to [0,1]^n$, find an $x \in [0,1]^n$ such that $\|f(x) - x\|_2^2 \le \epsilon n$. In the query version, the algorithm is given query access to f.

The Brouwer fixed point theorem guarantees that this problem is solvable; in particular it guarantees the existence of some x such that $f(x) = x$.

Theorem 1.8 ([3, 4])**.** *There is an $\epsilon > 0$ for which the query complexity of the ϵ-Brouwer fixed point problem is $2^{\Omega(n)}$.*

A proof of this theorem was given informally in Lecture 2. The proof used a reduction from end-of-a-line to Brouwer, showing that we could embed end-of-a-line problems with $2^{\Omega(n)}$ nodes in a Brouwer fixed point problem in n dimensions.

2 Useful Tools

In this section, we will discuss two tools that will be useful to us later on in the lecture.

2.1 The Lifting (Simulation) Theorem

Theorem 2.1 (Theorem 1 in [2], rephrased)**.** *Suppose we have a query problem $\mathcal{A}$ in which an algorithm is given query access to a function $f : \{0,1\}^n \to \{0,1\}$ that outputs bits. Then there is a communication problem $\mathcal{B}$ with the following characteristics:*

(1) Alice and Bob are given, as input, respectively strings $\alpha_i, \beta_i \in \mathbb{Z}_2^{\log n}$ for each possible query $x_i \in \{0,1\}^n$. The strings are such that $\alpha_i \cdot \beta_i = f(x_i)$ for all i.

(2) The goal of Alice and Bob is to solve the same problem as $\mathcal{A}$.

(3) The communication complexity of $\mathcal{B}$ is the same as the query complexity of $\mathcal{A}$, up to a log factor.

This theorem is somewhat remarkable: it allows us to turn any query lower bound into a communication lower bound in a black-box fashion. The proof of this theorem is beyond the scope of this course; the overachieving reader should read [2].

2.2 An Attempted Query Complexity Reduction to Approximate Nash

In this section, we will give an attempted reduction in the query complexity model from the Brouwer fixed-point problem to the query version of the ϵ-NE problem. Though this attempt will ultimately fail, it will inform our later attempts to prove the main theorem of this lecture (Theorem 1.4).

Let $f : [0,1]^N \to [0,1]^N$ be continuous. We will define a game such that, if ϵ-NE is solvable with few queries, we can find a fixed point of f with few queries as well. Define the *Brouwer imitation game* as follows: Alice and Bob play respectively some vectors $x, y \in [0,1]^N$. Then Alice gets payoff $A(x,y) = -\|x - y\|_2^2$, and Bob gets payoff $B(x,y) = -\|f(x) - y\|_2^2$.

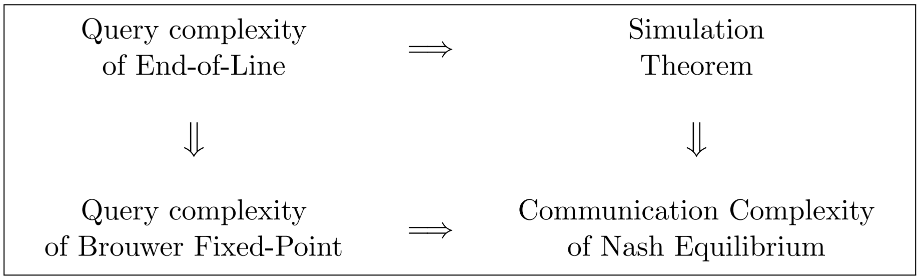

Figure 1: The four main ingredients of our reduction

Observe that, if Bob plays a mixed strategy given by a random variable y, then Alice's unique best response is to play $x = \mathbb{E}\, y$. This follows by observing that $\mathbb{E}\, A(x, y) = -\|x - \mathbb{E}\, y\|_2^2 - \operatorname{var} y$, and Alice cannot control $\operatorname{var} y$. Similarly, if Alice plays mixed strategy given by variable x, Bob's unique best response is to play $y = \mathbb{E}\, f(x)$. Thus, the unique Nash equilibrium is when $y = x = f(y)$; i.e. they both play at a fixed point of f. This result holds roughly with relaxation as well.

This construction is informal: indeed, we have created a game with infinitely many actions, which our prior definitions did not allow, and which doesn't say anything immediately about what happens with finitely many actions! We will call this a cheat and move on.

There is also another problem with this construction: since Bob can compute the function f exactly from his utility function, he doesn't need any communication at all—he can simply compute the Nash equilibrium himself and send it to Alice, solving the problem far more efficiently than we would like. This problem is more fundamental, as it means this reduction is not actually a reduction. However, we will see later that this construction is still useful to us, as it will form part of a larger construction that will allow us to prove Theorem 1.4.

Figure 1 gives the four major ingredients of what will eventually be our proof: (1) the query complexity of end-of-line; (2) the reduction from end-of-line to Brouwer fixed-point; (3) the simulation theorem, which we will use to go from query complexity to communication complexity; and (4) the Brouwer imitation game, which we will use to go from Brouwer to Nash. In the next section, we will see in more detail how to properly combine these ideas to arrive at our desired result.

3 Proving a Communication Complexity Lower Bound on ϵ-NE

Given the tools of the previous sections, we now have several different ideas for attacking the main result of this lecture (Theorem 1.4).

(1) We can try to apply the lifting theorem to the Brouwer imitation game. Sadly, this fails because the resulting vectors α_i, β_i that are received by Alice and Bob cannot reasonably be interpreted as a game.

(2) We can try to apply the lifting theorem to the end-of-line problem. Sadly, this fails too for a similar reason: the resulting communication problem has no apparent connection to anything else that is useful.

(3) We can try to apply the lifting theorem to the Brouwer fixed point problem. Intuitively, the main advantage of understanding the complexity of Brouwer's fixed point is that it is a

continuous problem, like (mixed) Nash equilibrium. The downside with this approach is that the lifting theorem gives us a *discrete* object (binary vectors), so we lose everything that was good about our reduction from end-of-line to Brouwer.

We will now present a working construction. The general idea of the construction is to modify the end-of-line-to-Brouwer-to-ϵ-NE reduction in such a way that it incorporates the lifting result in a black-box fashion, which will result in a game with high communication complexity. Let $G = (V, E)$ be a base graph with $N = 2^n$ vertices, and consider the end-of-line problem on G. We can embed G into $[0,1]^n$. Define a function $u : [0,1]^n \to V$ so that $u(x)$ is the vertex in V that is embedded nearest to x.[1] Further, apply the lifting theorem to the end-of-line problem to obtain sets of vectors α_i, β_i for each $i \in V$.[2] We will call our game the *lifting imitation game*, and define it as follows.

(1) Alice plays $x \in [0,1]^n$, some $\hat{a} \in \mathbb{Z}_2^n$, and some $\hat{r} \in \mathbb{Z}_2$.

(2) Bob plays $y \in [0,1]^n$, $\hat{b} \in \mathbb{Z}_2^n$, and some vertex $u^* \in V$.

(3) Alice's payoff is $A = -\|x - y\|_2^2 + [\hat{\alpha} = \alpha_{u(y)}] + [\hat{r} = \hat{\alpha} \cdot \hat{\beta}]$, where $[\cdot]$ denotes an indicator function.

(4) Bob's payoff is $-\|f_{\hat{r}}(x) - y\|_2^2 + [u^* = u(x)] + [\hat{\beta} = \beta_{u(x)}]$.

We can now make the following claims about this game:

(1) In NE, Alice must play $\mathbb{E}\, y$ and Bob must play $\mathbb{E}\, f(x)$. Thus again $x = y = f(y)$. This follows from the same argument as in the Brouwer imitation game.

(2) In NE, Bob must play u^* such that $u^* = u(x)$, and $\hat{\beta}$ such that $\hat{\beta} = \beta_{u(x)}$ (since he knows Alice will play $x = y$). Otherwise, he loses 1 utility. Similarly, Alice must play $\hat{\alpha} = \alpha_{u(y)}$ and $\hat{r} = \hat{\alpha} \cdot \hat{\beta}$.

(3) We can fix the cheat from before regarding Bob having to know the Brouwer function f: Since $\alpha_{u(x)}$ and $\beta_{u(x)}$ are played by Alice and Bob respectively, Bob doesn't need to know f everywhere: he can locally compute the value of f at x using the information $\hat{r} = r_{u(x)}$ about the vertex $u(x)$ in the end-of-line instance; we denote this value $f_{\hat{r}}(x)$. Indeed, if we recall the end-of-line-to-Brouwer reduction, knowing what vertices are near x (and what f maps those vertices to) is sufficient to completely understand the behavior of f.

Again, these claims still hold (approximately) upon relaxing to an ϵ-NE. We thus notice that the unique Nash equilibrium of this game is a Brouwer fixed point, and that computing this Nash equilibrium requires Alice and Bob to share $\hat{\alpha}$ and $\hat{\beta}$ in order to compute each evaluation of the function f; thus, this is actually a proper reduction from the query complexity of the Brouwer fixed point problem to the communication complexity of approximate Nash equilibrium!

[1] Notice that the nearest vertex may not be unique, but this is a problem that we will again overlook and call it a cheat.

[2] This again is a cheat: it assumes that the end-of-line problem in the query model has oracles that return only $O(1)$ bits per vertex.

4 Fixing the Cheats

Finally, we will discuss briefly some of the cheats that we have used throughout this class, and discuss possible ways of fixing them.

1. *The imitation games have infinitely large action spaces.* This can be fixed by an adequate discretization.
2. *The end-of-line problem may require $\omega(1)$ bits of information to be sent for each vertex.* Not so if the base graph G has constant degree (e.g. is a grid)!
3. *The vertex embedded nearest to x may not be unique.* This is actually more or less okay: the only utility of the function u is to let Bob correctly compute $f(y)$ via the inner product gadget from the lifting theorem. However, if x is far from all the vertices, recall from Lecture 2 that the reduction defines $f(x)$ to simply point downwards; thus, no query is necessary to evaluate $f(x)$, and similarly Alice and Bob do not need to communicate to evaluate $f(x)$.

References

[1] Yakov Babichenko and Aviad Rubinstein. Communication complexity of approximate nash equilibria. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 878–889, 2017.

[2] Mika Göös, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for BPP. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 132–143, 2017.

[3] Michael D. Hirsch, Christos H. Papadimitriou, and Stephen A. Vavasis. Exponential lower bounds for finding brouwer fix points. *J. Complexity*, 5(4):379–416, 1989.

[4] Aviad Rubinstein. Settling the complexity of computing approximate two-player nash equilibria. In *FOCS*, pages 258–265. IEEE Computer Society, 2016.