On the Computational Power of Players in Two-Person Strategic Games

Advisor: Prof. Yuh-Dauh Lyuu

Ching-Lueh Chang
Department of Computer Science and Information Engineering
National Taiwan University

Contents

Abstract

We consider families of two-person strategic games parameterized by a positive integer n. We assume that each of the players, the row player and the column player, has 2^n strategies to choose from and can take mixed strategies. We also assume that the games are not too "risky" in that the payoffs are at most polynomial in n in absolute value. The row player is said to guarantee an expected payoff of $t \in \mathbb{R}$ against all column players of a certain class when his expected payoff is at least t against all column players of that class. This thesis studies the expected payoff the row player could guarantee against all column players of a certain computational power. In our main result we consider the case where the row player is informed of how the column player chooses his action but is not allowed to see the internal coin tosses of the column player. Roughly speaking, we show that when the computational power of the column player shrinks polynomially, the row player can have his number of pure strategies shrunk exponentially without harming his guaranteed expected payoff too much. We obtain several corollaries regarding the computational power needed by the row player to guarantee a good expected payoff against computationally bounded column players in situations where the row player is aware of the output strategy of the column player, the mixed strategy of the column player, or only $O(\log n)$ bits of information about the column player.

Chapter 1

Introduction

1.1 Two-Person Strategic Games

In a two-person strategic game there are two players, the row player and the column player. Each player has several pure strategies to choose from. There is a payoff function mapping each combination of the players' pure strategies to their payoffs. Each player can also adopt a mixed strategy, that is, a probability distribution over pure strategies. In a two-person game, a pair of mixed strategies (α_1, α_2) is a mixed strategy Nash Equilibrium if no player i has a strategy yielding an expected payoff higher than when he chooses α_i, given that the other player j chooses α_j. The expectation is over the mixed strategies adopted by both players. If the payoffs of the two players sum to 0, the game is said to be zero-sum. For a zero-sum game, the expected payoff for the row player when a mixed strategy Nash equilibrium is played is called a game value under mixed strategies. The game value thus defined is unique [OR94]. For a precise definition of these and some other terms please refer to Section 2.1.

1.2 Our Results

We consider families of two-person strategic games parameterized by a positive integer n. In the game parameterized by n, each player is given input 1^n and oracle access to the payoff function mapping each combination of the players' pure strategies to their payoffs. The payoffs are assumed to be bounded by n^k in absolute value for some constant k. For technical reasons we will sometimes need to assume that the payoffs can be represented in binary using polynomially many bits (in n). Each player has $\{0,1\}^n$ as his set of pure strategies and can take mixed strategies. The row player is said to guarantee an expected payoff of $t \in \mathbb{R}$ against all column players of a certain class if his expected payoff is at least t against all column players of that class. This thesis considers the case where the row player is informed of how the column player chooses his action but is not allowed to see the internal coin tosses of the column player. In this case, therefore, the row player is said to guarantee an expected payoff of t against all column players of a certain class if, against each column player of that class, the row player can take some strategy to force an expected payoff of at least t. In particular, the row player is interested in the expected payoff he could guarantee against all column players being randomized circuits (with oracle gates for the players' payoff functions) of a certain size s. This thesis shows that when s shrinks polynomially (in n), the row player could have his number of pure strategies shrunk exponentially (in n) without harming his guaranteed expected payoff too much. We obtain several corollaries regarding the computational power needed by the row player to guarantee a good expected payoff against randomized circuits (acting as the column player) of a certain size, in situations where the row player is informed of the action taken by the column player, the mixed strategy of the column player, or only some $O(\log n)$ bits of information about the column player.

1.3 Related Works

The computational power of players in games is a topic being actively researched. We briefly describe several results in this area. For the definitions of standard complexity classes please refer to the complexity zoo [AK]. In this section when a game is given to an algorithm, it is given by explicitly specifying the payoffs for both players given each combination of pure strategies being adopted by the players, unless otherwise specified. It is well-known that for a two-person zero-sum game, a mixed strategy Nash equilibrium, and hence the game value, can be found in polynomial time by linear programming [Kha79, Owe82]. It is also known that given a two-person zero-sum game and a number v, computing whether the game value equals v is **P**-hard [FIKU05]. Given a polynomial-size circuit that computes the payoff function of a two-person zero-sum game and a number v, computing whether the game value equals v is known to be **EXP**-complete [FKS95]. A mixed strategy α in a two-person zero-sum game is said to be ϵ-optimal if, when the corresponding player uses α, his least possible expected payoff (against arbitrarily malicious players) is at most ϵ lower than that when a mixed strategy Nash equilibrium is played. Newman has shown that in any two-person zero-sum game with payoffs in $[\,0, 1\,]$, each player has an ϵ-optimal strategy which chooses uniformly from a multiset of $O(\log N/\epsilon^2)$ pure strategies, where N is the number of pure strategies of the other player [New91]. Finding ϵ-optimal strategies can be done efficiently in parallel [GK92, GK95, LN93, PST95], as well as sequentially in sublinear time by a randomized algorithm [GK95]. Given a polynomial-size circuit that computes the payoff function of a two-person zero-sum game, approximating the game value to within an additive factor is complete for the class promise-$\mathbf{S_2^p}$ [FIKU05], and a pair of ϵ-optimal strategies can be constructed by a $\mathbf{ZPP^{NP}}$ algorithm [FIKU05]. After a series of exciting developments initiated by Daskalakis, Goldberg, and Papadimitriou [DGP05], Cheng and Deng [CD06] recently show that computing a mixed strategy Nash equilibrium is **PPAD**-complete for two-person games.

The rest of the thesis is organized as follows. Chapter 2 presents the necessary background knowledge and some conventions. Chapter 3 shows how to extract a small number of strategies good against players of a certain computational power using an idea similar to that in [FPS03]. Chapter 4 applies the result to analyze the computational power needed by a player to guarantee a certain expected payoff in several cases.

Chapter 2

Preliminaries

2.1 Basic Terms in Game Theory

A finite N-person game G in strategic form consists of players $1, \ldots, N$, N finite sets $A_1, A_2, \ldots, A_N$, and for $i = 1, \ldots, N$, a payoff function $M_i : A_1 \times A_2 \times \cdots \times A_N \to \mathbb{R}$. For $i = 1, \ldots, N$, A_i is the finite set of pure strategies player i can take, and M_i maps the players' pure strategies to the payoff for player i.

Player i can adopt a probability distribution over A_i as his strategy. In this case we say that player i takes a mixed strategy. A mixed strategy α_i of player i is also used alternatively to represent a random variable whose distribution is specified by α_i as the ambiguity will not arise. A pure strategy is obviously a degenerate mixed strategy.

Now suppose α_i is the mixed strategy of player i and $\alpha_1, \ldots, \alpha_N$ are independent. We say that α_i is a best response to $\times_{1 \leq k \leq N, k \neq i} \alpha_k$ if for every mixed strategy ${\alpha_i}'$ of player i,

$$\begin{aligned} & E[M_i(\alpha_1, \ldots, \alpha_{i-1}, {\alpha_i}', \alpha_{i+1}, , \ldots, \alpha_N)] \\ \leq \; & E[M_i(\alpha_1, \ldots, \alpha_{i-1}, \alpha_i, \alpha_{i+1}, \ldots, \alpha_N)]. \end{aligned}$$

The randomness in the first expectation comes from $\alpha_1, \ldots, \alpha_{i-1}, {\alpha_i}', \alpha_{i+1}, \ldots, \alpha_N$, that of the second expectation comes from $\alpha_1, \ldots, \alpha_N$. We say that $(\alpha_1, \ldots, \alpha_N)$ is a mixed strategy Nash equilibrium for G if for each i, α_i is a best response to $\times_{1 \leq k \leq N, k \neq i} \alpha_k$. A celebrated theorem due to Nash states that every finite N-person game in strategic form has a mixed strategy Nash equilibrium [Nas51]. The following fact is well-known.

Fact 1. *([OR94]) Let $\alpha_1, \ldots, \alpha_N$ be mixed strategies for players $1, \ldots, N$ in a finite N-person game in strategic form. For $1 \leq i \leq N$, there is a pure strategy for the ith player which is a best response to $\times_{1 \leq k \leq N, k \neq i} \alpha_k$.*

When there are only two players, player 1 is called the row player and player 2 the column player. The corresponding payoff functions will be denoted $M_{\text{row}} : A_{\text{row}} \times A_{\text{col}} \to \mathbb{R}$ and $M_{\text{col}} : A_{\text{row}} \times A_{\text{col}} \to \mathbb{R}$, respectively. For convenience, we will sometimes combine M_{row} and M_{col} into a single function $M(\cdot) = (M_{\text{row}}(\cdot), M_{\text{col}}(\cdot))$. M_{row}, M_{col}, and M will be referred to as the payoff function for the row player, the payoff function for the column player, and simply the payoff function, respectively. When a two-person game satisfies $M_{\text{row}} = -M_{\text{col}}$, the game is said to be zero-sum. Whenever $(\alpha_{\text{row}}, \alpha_{\text{col}})$ is a mixed strategy Nash equilibrium in a two-person zero-sum game, we say that $E[M_{\text{row}}(\alpha_{\text{row}}, \alpha_{\text{col}})]$ is a value for that game under mixed strategies. The game value thus defined is known to be unique [OR94]. The set of all distributions over the pure strategies of the row player, or equivalently, the set of all mixed strategies of the row player, will be denoted $\mathcal{P}_{\text{row}}$. $\mathcal{P}_{\text{col}}$ is defined similarly.

2.2 Randomized Circuits

Let $o : \{0,1\}^l \to \{0,1\}^m$ be a function. A randomized o-oracle circuit $\mathcal{C}$ is a collection of gates with ordered input and output pins, and directed wires going from output pins to input pins, together with a specification of several distinct ordered output pins $p_1, \ldots, p_n$ as the output bits of the whole circuit. Cycles are not permitted. That is, we cannot start from some gate, keep traveling along outgoing wires and finally reach the gate we started with. The set of gates is $G_{\text{inp}} \cup G_{\text{rand}} \cup G_{\text{std}} \cup G_{\text{oracle}}$, where $G_{\text{inp}} = \{\omega_1, \ldots, \omega_t\}$ is the set of input gates, G_{rand} the set of random gates, G_{std} the set of standard gates, and G_{oracle} the set of oracle gates. Each non-oracle gate has exactly one output pin. Each input pin should receive exactly one incoming wire. The indegree of a gate is the number of its incoming wires. Each input or random gate has indegree zero, each standard gate has indegree zero, one, or two, and each oracle gate has indegree l. The output pin of each standard gate of indegree zero is labeled either by the constant 0 or by the constant 1. The output pin of each standard gate of indegree one is labeled by the negation function $\text{NOT} : \{0,1\} \to \{0,1\}$, which maps 0 to 1 and 1 to 0. The output pin of each standard gate of indegree two is labeled either by the conjuction $\text{AND} : \{0,1\}^2 \to \{0,1\}$ or by the disjunction $\text{OR} : \{0,1\}^2 \to \{0,1\}$. The AND function outputs a 1 when its inputs are both 1 and a 0 otherwise. The OR function outputs a 0 when its inputs are both 0 and a 1 otherwise. Each oracle gate has m output pins labeled with functions $\text{pin}_1 : \{0,1\}^l \to \{0,1\}, \ldots, \text{pin}_m : \{0,1\}^l \to \{0,1\}$. The function pin_i is the ith output bit of o.

The wires in $\mathcal{C}$ carry Boolean values 0 and 1. For any assignment of Boolean values to the input gates, we compute the outputs of $\mathcal{C}$ by first labeling independently and equiprobably a constant 0 or 1 for each output pin of a random gate. We then proceed by propagating values along the wires and computing the functions labeled with the respective pins until the output bits of $\mathcal{C}$ are obtained. The size of a randomized o-oracle circuit is defined to be the number of its gates. A randomized circuit is a randomized o-oracle circuit without oracle gates. A circuit is deterministic if it has no random gates.

A well-known counting technique due to Shannon can be adapted to bound the number of randomized o-oracle circuits of size s with n output bits.

Lemma 1. *([Sha49]) Fix a Boolean function $o : \{0,1\}^l \to \{0,1\}^m$. For each $s \in \mathbb{N}^+$, there are fewer than $2^{O(ls\log(sm)+n\log(sm))}$ randomized o-oracle circuits of size s with n output bits.*

Proof. The input gates are $\omega_1, \ldots, \omega_t$, for some $1 \leq t \leq s$. Each gate, except the input gates, is one of seven kinds. Each gate receives at most $\max(l,2)$ incoming wires, each coming from some output pin of another gate. Finally, there are at most $(sm)^n$ ways to select the output bits in order. To sum up, the number of such circuits is at most $s7^s((sm)^{\max(l,2)})^s(sm)^n$. A direct computation completes the proof. □

We will need a well-known circuit called a multiplexer.

Lemma 2. *([KB04]) There is an $O((m+n)2^m)$-size deterministic circuit that, on Boolean input values $b_{i,j}$, $1 \leq i \leq 2^m, 1 \leq j \leq n$ and an integer $1 \leq g \leq 2^m$ in binary, outputs $b_{g,j}, 1 \leq j \leq n$.*

Proof. For $1 \leq i \leq 2^m$, an $O(m)$-size circuit computes whether $g = i$. Having computed whether $g = i$ for $1 \leq i \leq 2^m$, each output bit $b_{g,j}$ can be computed by an $O(2^m)$-size deterministic circuit. □

2.3 Tail Inequalities

We will need two results regarding tail probabilities. The first is the famous Chernoff bound.

Fact 2. *([Che52]) Let $X_1, X_2, \ldots, X_n$ be independent* 0-1 *random variables such that for $1 \leq i \leq n$, $\Pr[X_i = 1] = p_i$. Then, for any $\delta > 0$,*

$$\Pr\left[\frac{\Sigma_{i=1}^n X_i}{n} > (1+\delta)\frac{\Sigma_{i=1}^n p_i}{n}\right] < \left(\frac{e^\delta}{(1+\delta)^{1+\delta}}\right)^{\Sigma_{i=1}^n p_i}.$$

□

The next famous bound is Hoeffding's inequality.

Fact 3. *([Hoe63]) Let $X_1, X_2, \ldots, X_n$ be n independent random variables with the same probability distribution, each ranging over the (real) interval $[\,a, b\,]$, and let μ denote the expected value of each of these variables. Then, for every $\epsilon > 0$,*

$$\Pr\left[\left|\frac{\Sigma_{i=1}^n X_i}{n} - \mu\right| > \epsilon\right] < 2e^{-\frac{2\epsilon^2 n}{(b-a)^2}}.$$

□

2.4 Conventions and Assumptions

We consider families of two-person games in strategic form, parameterized by a positive integer n. In the game parameterized by n, each player is given a parameter 1^n and has $\{0,1\}^n$ as his set of pure strategies. For each player, we use either $\{0,1\}^n$ or numbers 0 to $2^n - 1$ to represent his pure strategies. The payoff functions M_{row}, M_{col}, and M for the game parameterized by n will be written as $M_{\text{row}}^{(n)}$, $M_{\text{col}}^{(n)}$, and $M^{(n)}$, respectively. We assume $\|M_{\text{row}}^{(n)}\|_\infty \leq n^k$ and $\|M_{\text{col}}^{(n)}\|_\infty \leq n^k$ for some positive constant k, where for every real-valued, continuous, bounded function f, its supremum norm $\|f\|_\infty$ is the supremum of the image of $|\,f\,|$ [Rud76]. Whenever $M^{(n)}$ is used as an oracle, both its output components, $M^{(n)}_{\text{row}}$ and $M^{(n)}_{\text{col}}$, are in binary.

The players may be either computationally unbounded, time-bounded Turing machines, or polynomial-size (in n) randomized $M^{(n)}$-oracle circuits. Using the tableau method, randomized polynomial-size $M^{(n)}$-oracle circuits can be simulated by randomized polynomial-time Turing machines with polynomial advice and an $M^{(n)}$-oracle, and vice versa [Sip05]. The set of randomized $M^{(n)}$-oracle circuits of size s with n output bits is denoted $\text{SIZE}^{M^{(n)}}_{s,n}$, or simply SIZE^{M}_{s} when the value of n is clear from the text. Similarly, the set of randomized circuits (without oracle gates) of size s with n output bits is $\text{SIZE}_{s,n}$, or SIZE_s when the value of n is clear from the text. Throughout this thesis, we say that a (deterministic or nondeterministic) Turing machine runs in polynomial time (resp. logarithmic space) if its time (resp. space) complexity is polynomial (resp. logarithmic) in its input length, which is n in the game parameterized by n.

We abuse our notation slightly by using $M^{(n)}_{\text{row}}(R, C)$ to denote the payoff of the row player when the row player is R and the column player is C, and $E[M^{(n)}_{\text{row}}(R, C)]$ its expected value. The expectation is over the random coin flips of R and C. If C always uses pure strategy j, we also denote $M^{(n)}_{\text{row}}(R, C)$ by $M^{(n)}_{\text{row}}(R, j)$. The same convention is adopted for $M^{(n)}_{\text{col}}$. Similarly, $M^{(n)}(R, C)$ denotes $(M^{(n)}_{\text{row}}(R, C), M^{(n)}_{\text{col}}(R, C))$, and $M^{(n)}(R, j)$ denotes $(M^{(n)}_{\text{row}}(R, j), M^{(n)}_{\text{col}}(R, j))$.

Chapter 3

Extraction of Small Numbers of Good Strategies

In this chapter, we will often encounter a function $\epsilon : \mathbb{N}^+ \to (0,1)$. For convenience, we will write $\epsilon(n)$ simply as ϵ when the value of n is clear from the context. We first prove the following theorem.

Theorem 1. *Let $k > 0$, $d \geq 1$, and $c > d + k + 1$ be constants. Consider a family of two-person strategic games parameterized by $n \in \mathbb{N}^+$. Let $M^{(n)} : \{0,1\}^n \times \{0,1\}^n \to [-n^k, n^k]^2$ be the payoff function for the game parameterized by n. Let $\epsilon : \mathbb{N}^+ \to (0,1)$ be a function. We assume that the binary representation of each number in the range of $M^{(n)}_{row}$ and $M^{(n)}_{col}$ is of length polynomial in n. For each n, there is a set $S \subseteq \{0,1\}^n$ of size $O((n^{k+d+1} \log n)/\epsilon)$ such that*

$$\min_{C \in \mathit{SIZE}^M_{n^d}} \max_{i \in S} E[M^{(n)}_{row}(i,C)] > \min_{C \in \mathit{SIZE}^M_{n^c \log(1/\epsilon)/\epsilon}} \max_{i \in \{0,1\}^n} E[M^{(n)}_{row}(i,C)] - \epsilon$$

The symmetric statement, with the roles of players exchanged, holds.

Here is the interpretation of Theorem 1. Suppose the row player is informed of the mixed strategy of the column player. For each column player

C, the row player chooses a best response and obtains an expected payoff of

$$\max_{\alpha \in \mathcal{P}_{\text{row}}} E[M_{\text{row}}^{(n)}(\alpha, C)] = \max_{i \in \{0,1\}^n} E[M_{\text{row}}^{(n)}(i, C)].$$

The equality is from Fact 1. The first expectation is over α and the coin flips of C and the second is over the coin flips of C. The expected payoff the row player could guarantee against all randomized $M^{(n)}$-oracle circuits of size $n^c \log(1/\epsilon)/\epsilon$ is therefore

$$\min_{C \in \text{SIZE}^M_{n^c \log(1/\epsilon)/\epsilon}} \max_{i \in \{0,1\}^n} E[M_{\text{row}}^{(n)}(i, C)].$$

Similarly,

$$\min_{C \in \text{SIZE}^M_{n^d}} \max_{i \in S} E[M_{\text{row}}^{(n)}(i, C)]$$

is the expected payoff the row player could guarantee by using only strategies in S against all randomized $M^{(n)}$-oracle column players of size n^d, provided that the row player is informed of the mixed strategy of the column player. For ϵ not too small, Theorem 1 shows that when the circuit size of the column player shrinks polynomially (from $n^c \log(1/\epsilon)/\epsilon$ to n^d), the set of strategies of the row player could shrink exponentially (from 2^n to $O((n^{k+d+1}) \log n)/\epsilon)$ without affecting the guaranteed expected payoff of the row player too much.

Proof of Theorem 1. We will always assume n is sufficiently large whenever needed since the requirement on the size of S contains a big-O and the theorem is therefore immediately true for small values of n.

Denote

$$t = \min_{C \in \text{SIZE}^M_{n^c \log(1/\epsilon)/\epsilon}} \max_{i \in \{0,1\}^n} E[M_{\text{row}}^{(n)}(i, C)]$$

for convenience. We say that a pure row strategy i is good against a randomized n^d-size $M^{(n)}$-oracle circuit C if $E[M_{\text{row}}^{(n)}(i, C)] > t - \epsilon$.

S will be formed in stages. We keep a set of survivors which is initially the set of all randomized n^d-size $M^{(n)}$-oracle circuits with n output bits. In each stage we put one pure row strategy into S and kill the survivors against which this strategy is good. We do so until there are no survivors left. If the number of stages is $O((n^{k+d+1}\log n)/\epsilon)$, then we are done.

Let Survivors_i denote the set of circuits that have not been killed after stage i. Initially, Survivors_0 consists of all randomized n^d-size $M^{(n)}$-oracle circuits with n output bits. According to Shannon's counting argument (Lemma 1 with n^d assigned to s, $2n$ assigned to l and $\text{poly}(n)$ assigned to m),

$$|\,\text{Survivors}_0\,| \leq 2^{O(n^{d+1}\log n)}. \tag{3.1}$$

Given Survivors_i, we now show how to obtain Survivors_{i+1}. Let T be the smallest power of 2 not less than $(n^{k+1}\log n)/\epsilon$. Hence $(n^{k+1}\log n)/\epsilon \leq T \leq 2(n^{k+1}\log n)/\epsilon$. Consider any collection of T circuits in Survivors_i, possibly with repetitions: $C_1, C_2, \ldots, C_T$. We construct a randomized $M^{(n)}$-oracle circuit C that feeds independent random inputs to $C_1, C_2, \ldots, C_T$ and chooses equiprobably one of the n-bit outputs of $C_1, C_2, \ldots, C_T$ as the output of C. Summing up the sizes of $C_1, C_2, \ldots, C_T$ and including a multiplexer (Lemma 2 with T assigned to 2^m), we see that C is of size

$$\begin{aligned} & O\left((n^{k+1}\log n)/\epsilon \cdot n^d + (n^{k+1}\log n)/\epsilon \cdot (\log((n^{k+1}\log n)/\epsilon) + n)\right) \\ < \; & n^c \log(1/\epsilon)/\epsilon \end{aligned} \tag{3.2}$$

for sufficiently large n. Inequality (3.2) and the definition of t show that there is an i^* with $E[M^{(n)}_{\text{row}}(i^*, C)] \geq t$. This is equivalent to saying

$$\frac{E[M^{(n)}_{\text{row}}(i^*, C_1)] + \cdots + E[M^{(n)}_{\text{row}}(i^*, C_T)]}{T} \geq t. \tag{3.3}$$

Let f be the fraction of values above $t - \epsilon$ among

$$E[M_{\text{row}}^{(n)}(i^*, C_1)], E[M_{\text{row}}^{(n)}(i^*, C_2)], \ldots, E[M_{\text{row}}^{(n)}(i^*, C_T)].$$

Since $\|M_{\text{row}}^{(n)}\|_\infty \leq n^k$, it is clear that

$$\begin{aligned} & \frac{E[M_{\text{row}}^{(n)}(i^*, C_1)] + \cdots + E[M_{\text{row}}^{(n)}(i^*, C_T)]}{T} \\ \leq \; & f n^k + (1-f)(t-\epsilon). \end{aligned} \tag{3.4}$$

Inequalities (3.3) and (3.4) imply that $t \leq fn^k + (1-f)(t-\epsilon)$, or that $f \geq \epsilon/(n^k - t + \epsilon)$. This and the fact that $\epsilon \in (0,1)$ and $|\,t\,| \leq \|M_{\text{row}}^n\|_\infty \leq n^k$ result in

$$f > \epsilon/(3n^k)$$

for sufficiently large n. That is, i^* is good against more than a $\epsilon/(3n^k)$ fraction of players among $C_1, C_2, \ldots, C_T$ for sufficiently large n.

Next suppose we actually pick each of $C_1, C_2, \ldots, C_T$ independently and uniformly from Survivors_i. Fix arbitrarily, if any, a pure row strategy i' good against less than an $\epsilon/(4n^k)$ fraction of Survivors_i. Let $f_{i'}\,\epsilon/(3n^k)$ be the fraction of players in Survivors_i against which i' is good. By the choice of i', we have $f_{i'} < 3/4$. The Chernoff bound (Fact 2 with $X_j = 1$ if i' is good against C_j and 0 otherwise, $p_j = f_{i'}\,\epsilon/(3n^k)$, T assigned to n and $-1 + 1/f_{i'}$ assigned to δ) gives

$$\begin{aligned} & \Pr\left[i' \text{ is good against more than an } \epsilon/(3n^k) \text{ fraction of } C_1, \ldots, C_T\right] \\ < \; & (f_{i'} e^{1-f_{i'}})^{T\epsilon/(3n^k)} \leq e^{-\Omega(n \log n)}. \end{aligned}$$

The probability is over the picking of $C_1, \ldots, C_T$. The last inequality is true because the function xe^{1-x} has positive derivative on $(0,1)$ and $(3/4)e^{1-3/4} < 1$. By summing over $i' \in \{0,1\}^n$, the probability is at most $2^n e^{-\Omega(n \log n)}$

that some pure row strategy good against less than an $\epsilon/(4n^k)$ fraction of Survivors$_i$ is good against more than an $\epsilon/(3n^k)$ fraction of randomly picked $C_1, C_2, \ldots, C_T$. Since $2^n e^{-\Omega(n \log n)} < 1$ for sufficiently large n, a probabilistic argument shows that there is a choice of $C_1, C_2, \ldots, C_T$ such that *every* pure row strategy good against more than an $\epsilon/(3n^k)$ fraction of $C_1, C_2, \ldots, C_T$ must be good against at least an $\epsilon/(4n^k)$ fraction of Survivors$_i$. We have seen in the last paragraph that for this (in fact, every) choice of $C_1, C_2, \ldots, C_T$, there is an i^* good against more than an $\epsilon/(3n^k)$ fraction of $C_1, C_2, \ldots, C_T$. This i^* must therefore be good against at least an $\epsilon/(4n^k)$ fraction of Survivors$_i$. We add this i^* to S and obtain

$$|\,\text{Survivors}_{i+1}\,| \leq (1 - \epsilon/(4n^k))|\,\text{Survivors}_i\,|.$$

From this and inequality (3.1), after $O((n^{k+d+1} \log n)/\epsilon)$ stages the number of survivors will be less than one. At that time, there must be no survivors left. □

The same theorem holds for circuits without oracle gates.

Corollary 1. *Let $k > 0$, $d \geq 1$ and $c > d + k + 1$ be constants. Consider a family of two-person strategic games parameterized by $n \in \mathbb{N}^+$. Let $M^{(n)} : \{0,1\}^n \times \{0,1\}^n \to [-n^k, n^k]^2$ be the payoff function for the game. Let $\epsilon : \mathbb{N}^+ \to (0,1)$ be a function. We assume that the binary representation of each number in the range of $M^{(n)}_{row}$ and $M^{(n)}_{col}$ is of length polynomial in n. For each n, there is a set $S \subseteq \{0,1\}^n$ of size $O((n^{k+d+1} \log n)/\epsilon)$ such that*

$$\begin{aligned} &\min_{C \in SIZE_{n^d}} \max_{i \in S} E[M^{(n)}_{row}(i, C)] \\ &> \min_{C \in SIZE_{n^c \log(1/\epsilon)/\epsilon}} \max_{i \in \{0,1\}^n} E[M^{(n)}_{row}(i, C)] - \epsilon \end{aligned}$$

The symmetric statement, with the roles of players exchanged, holds.

Proof. The proof is the same as that of Theorem 1 but with each occurrence of "randomized $M^{(n)}$-oracle circuit" replaced with "randomized circuit." □

As have been pointed out in [FKS95], Newman's result implies the following, which is relevant to our work [New91].

Lemma 3. *([New91]) Consider a family of two-person zero-sum strategic games parameterized by $n \in \mathbb{N}^+$. Let $M_{row}^{(n)} : \{0,1\}^n \times \{0,1\}^n \to [-n^k, n^k]$ be the payoff function for the row player. Let v be the game value under mixed strategies for the game parameterized by n, and $\epsilon : \mathbb{N}^+ \to (0,1)$ be a function. For every $T \geq 2n^{2k+1}/\epsilon^2$ there is a multiset of T pure row strategies such that if the row player selects equiprobably one of these strategies, his expected payoff is at least $v - \epsilon$ against any mixed strategy in $\mathcal{P}_{col}$ adopted by the column player. Similarly, there is a multiset of T pure column strategies such that if the column player selects equiprobably one of these strategies, his expected payoff is at least $-v - \epsilon$ against any mixed strategy in $\mathcal{P}_{row}$ adopted by the row player.*

Proof. Let $T \geq 2n^{2k+1}/\epsilon^2$. Let each of the independent random variables α_1, α_2, ..., α_T be distributed identically as the mixed strategy of the row player in a mixed strategy Nash equilibrium. For an arbitrary pure column strategy j, we have $E[M_{\text{row}}^{(n)}(\alpha_i, j)] \geq v$, $i = 1, \ldots, T$, by the definition of the game value under mixed strategies. From Hoeffding's inequality (Fact 3 with $[a, b] = [-n^k, n^k]$, $M_{\text{row}}^{(n)}(\alpha_i, j)$ assigned to X_i and T assigned to n), we have

$$\Pr\left[\frac{M_{\text{row}}^{(n)}(\alpha_1, j) + \cdots + M_{\text{row}}^{(n)}(\alpha_T, j)}{T} < v - \epsilon\right] < 2e^{-n}.$$

So with probability at most $2^n 2e^{-n} < 1$, there exists a j with

$$\frac{M_{\text{row}}^{(n)}(\alpha_1, j) + \cdots + M_{\text{row}}^{(n)}(\alpha_T, j)}{T} < v - \epsilon.$$

Hence there exist pure strategies $i_1, i_2, \ldots, i_T$ for the row player with

$$\frac{M_{\mathrm{row}}^{(n)}(i_1, j) + \cdots + M_{\mathrm{row}}^{(n)}(i_T, j)}{T} \geq v - \epsilon, \forall j.$$

By selecting equiprobably from among i_1, i_2, $\ldots$, i_T, the row player can guarantee an expected payoff of at least $v - \epsilon$ against any mixed strategy of the column player. The second part of this corollary can be proved by observing that by exchanging the roles of the two players, the game value becomes $-v$. □

We now consider the effects of Lemma 3 on Theorem 1.

Corollary 2. *Let $k > 0$ and $c > 2k + 2$ be constants. Consider a family of two-person strategic games parameterized by $n \in \mathbb{N}^+$. Let $M^{(n)} : \{0,1\}^n \times \{0,1\}^n \to [-n^k, n^k]^2$ be the payoff function for the game G_n. Let $\epsilon : \mathbb{N}^+ \to (0,1)$ be a function. There is a multiset $S \subseteq \{0,1\}^n$ of size $O(n^{2k+1}/\epsilon^2)$ such that the row player R_S who plays each strategy in S equiprobably guarantees*

$$\min_{\beta \in \mathcal{P}_{col}} E[M_{row}^{(n)}(R_S, \beta)] > \min_{C \in SIZE^M_{n^c \log(1/\epsilon)/\epsilon^2}} \max_{i \in \{0,1\}^n} E[M_{row}^{(n)}(i, C)] - \epsilon$$

for sufficiently large n.

Proof. We may define a zero-sum game G'_n with the payoff function $-M_{\mathrm{row}}^{(n)}$ for the column player and $M_{\mathrm{row}}^{(n)}$ for the row player. Let v' be the value of G'_n under mixed strategies. Let T be the smallest power of 2 not less than $2n^{2k+1}/\epsilon^2$. Using Lemma 3 on G'_n, we see that there is a multiset S' of size T such that, for the column player $C_{S'}$ choosing equiprobably a strategy in S',

$$\forall \alpha \in \mathcal{P}_{\mathrm{row}}, E[M_{\mathrm{row}}^{(n)}(\alpha, C_{S'})] \leq v' + \epsilon/3. \tag{3.5}$$

The same holds in the original game G_n because the games G_n and G'_n adopt the same payoff function for the row player and $C_{S'}$ makes no queries. $C_{S'}$

could be implemented as a circuit by hardwiring S', adding random gates, and including a multiplexer. From Lemma 2 with T assigned to 2^m, such a circuit is of size

$$\begin{aligned} & O\left(n^{2k+1}/\epsilon^2 \cdot n + n^{2k+1}/\epsilon^2 \cdot (\log(n^{2k+1}/\epsilon^2) + n)\right) \\ < \;& n^c \log(1/\epsilon)/\epsilon^2, \end{aligned} \tag{3.6}$$

for sufficiently large n. Denote

$$t = \min_{C \in \mathrm{SIZE}^M_{n^c \log(1/\epsilon)/\epsilon^2}} \max_{i \in \{0,1\}^n} E[M^{(n)}_{\mathrm{row}}(i, C)].$$

Inequalities (3.5), (3.6), and our definition of t gives $t \leq v' + \epsilon/3$ for sufficiently large n, or, equivalently,

$$v' \geq t - \epsilon/3, \tag{3.7}$$

for sufficiently large n. Applying Lemma 3 to G'_n, we see that by selecting equiprobably from among some T strategies S, the row player can guarantee an expected payoff of at least $v' - \epsilon/3$ against every mixed strategy of the column player. Again, this must also be true in G_n. Inequality (3.7) implies $v' - \epsilon/3 \geq t - 2\epsilon/3 > t - \epsilon$ and completes the proof. □

Since

$$\forall \beta \in \mathcal{P}_{\mathrm{col}}, E[M^{(n)}_{\mathrm{row}}(R_S, \beta)] \leq \max_{i \in S} E[M^{(n)}_{\mathrm{row}}(i, \beta)],$$

a direct computation on the sizes of S and the circuit sizes in Theorem 1 and Corollary 2 shows that Theorem 1 with $d = k + 3$ and $\epsilon > 1/n$ is implied by Corollary 2 , and hence by Lemma 3. It is not known if Theorem 1 is implied by Lemma 3, however.

Chapter 4

Implications on the Power of Players

We use Theorem 1 to investigate the computational resource needed for a player to guarantee a good expected payoff against computationally-bounded players. A naive player performs exponentially (in n) many queries to $M^{(n)}$ to determine an optimal strategy, pure or mixed, but we show how a player could run in less than exponential time without degrading his payoff too much.

Corollary 3. *Let $k > 0$, $d \geq 1$, $c > k + d + 1$, and $0 < \epsilon < 1$ be constants. Consider a family of two-person strategic games parameterized by a positive integer n. Let $M^{(n)} : \{0,1\}^n \times \{0,1\}^n \to [-n^k, n^k]^2$ be the payoff function. We assume that the binary representation of each number in the range of $M^{(n)}_{row}$ and $M^{(n)}_{col}$ is of length polynomial in n. Denote $t = \min_{C \in SIZE^M_{n^c \log(1/\epsilon)/\epsilon}} \max_{i \in \{0,1\}^n} E[M^{(n)}_{row}(i, C)]$. The following statements are true.*

(i) If R is informed of the circuit computing C, then he needs only polynomial time with polynomial advice (in n), private fair coin flips, and

an $M^{(n)}$-oracle to guarantee that $E[M^{(n)}_{row}(R,C)] > t-\epsilon$ for every randomized n^d-size $M^{(n)}$-oracle C.

(ii) *Assume R is allowed to produce his output after he sees the output strategy of C. If the language $\{i,j,k,n|$ the kth bit in the binary representation of $M^{(n)}_{row}(i,j)$ is a 1 $\}$ is in **NL**/poly, then there is a nondeterministic logarithmic-space row player R with polynomial advice that always has a unique accepting branch, and $E[M^{(n)}_{row}(R,C)] > t-\epsilon$ holds for every randomized n^d-size $M^{(n)}$-oracle C. Here the output of R is defined to be the string it leaves on its output tape on the unique accepting branch.*

(iii) *If R is required to determine his output strategy simultaneously with C (which is unknown to R), and if he obtains an additional $O(\log n)$ bits of information about C, then he needs only deterministic polynomial time with polynomial advice (in n) to guarantee that $E[M^{(n)}_{row}(R,C)] > t-\epsilon$ holds for every randomized n^d-size $M^{(n)}$-oracle circuit C.*

Here is the interpretation of Corollary 3. Assuming that R is informed of the mixed strategy (or the circuit) of the column player, the value t is the expected payoff R could guarantee against all randomized $M^{(n)}$-oracle column players of size $n^c \log(1/\epsilon)/\epsilon$. Item (i) states that with polynomial advice and oracle access to $M^{(n)}$, a polynomial-time R suffices to guarantee a $t-\epsilon$ payoff against all randomized n^d-size $M^{(n)}$-oracle column players. Item (ii) further assumes that R chooses his output strategy after the column player does, and $M^{(n)}$ is computed by a nondeterministic logarithmic-space Turing machine with polynomial advice. Under these assumptions, the polynomial-time Turing machine R in item (i) can be further reduced to an unambiguous logarithmic-space one. The guaranteed expected payoff is still at least $t-\epsilon$

against randomized n^d-size $M^{(n)}$-oracle column players. Item (iii) suggests that, even if R knows only some $O(\log n)$ bits of information about the column player (of size n^d), his guaranteed expected payoff is almost as if he is informed of the whole circuit of the column player (of size $n^c \log(1/\epsilon)/\epsilon$).

Proof of Corollary 3. We have seen in Theorem 1 that there is a set S of size polynomial in n such that, for every randomized n^d-size $M^{(n)}$-oracle column player C, there is a strategy $i \in S$ with $E[M^{(n)}_{\text{row}}(i, C)] > t - \epsilon/2$. We give S as the advice to R in all the following cases.

Proof of (i). Let C be an arbitrary randomized $M^{(n)}$-oracle circuit of size n^d with n output bits. Being informed of the circuit computing C, R could simulate C independently $50n^{2k+1}/\epsilon^2$ times and obtains outputs $O_1, \ldots, O_{50n^{2k+1}/\epsilon^2}$. For each $i \in S$, he computes $M^{(n)}_{\text{row}}(i, O_j)$, $1 \le j \le O_{50n^{2k+1}/\epsilon^2}$. Hoeffding's inequality (Fact 3 with $M^{(n)}_{\text{row}}(i, O_j)$ assigned to X_j, $50n^{2k+1}/\epsilon^2$ assigned to n, $\epsilon/5$ assigned to ϵ and $[a, b] = [-n^k, n^k]$) yields

$$\begin{aligned} & \Pr\left[\left|\frac{\Sigma_{j=1}^{50n^{2k+1}/\epsilon^2} M^{(n)}_{\text{row}}(i, O_j)}{50n^{2k+1}/\epsilon^2} - E[M^{(n)}_{\text{row}}(i, C)]\right| > \epsilon/5\right] \\ < \quad & 2e^{-n}, \end{aligned} \tag{4.1}$$

where the probability is over the coin tosses of R. For each $i \in S$, R estimates $E[M^{(n)}_{\text{row}}(i, C)]$ as

$$\begin{aligned} & \tilde{E}[M^{(n)}_{\text{row}}(i, C)] \\ \equiv \quad & \frac{M^{(n)}_{\text{row}}(i, O_1) + \cdots + M^{(n)}_{\text{row}}(i, O_{50n^{2k+1}/\epsilon^2})}{50n^{2k+1}/\epsilon^2}. \end{aligned}$$

Note that the random variable $\tilde{E}[M^{(n)}_{\text{row}}(i, C)]$ depends solely on the random coin tosses of R, and not those of C. Using inequality (4.1), we see that

$$\Pr\left[\exists i \in S, \left|\tilde{E}[M^{(n)}_{\text{row}}(i, C)] - E[M^{(n)}_{\text{row}}(i, C)]\right| > \epsilon/5\right] < |S|2e^{-n}, \tag{4.2}$$

where the probability is over the random coin tosses of R. The selection of S guarantees

$$\max_{i \in S} E[M_{\text{row}}^{(n)}(i, C)] > t - \epsilon/2. \tag{4.3}$$

Consider the row player R outputting an $i \in S$ with the largest $\tilde{E}[M_{\text{row}}^{(n)}(i, C)]$. Let i^* be the output of R. From inequality (4.2), with probability more than $1 - |S| 2e^{-n}$

$$\begin{aligned} & E[M_{\text{row}}^{(n)}(i^*, C)] \\ \geq\ & \tilde{E}[M_{\text{row}}^{(n)}(i^*, C)] - \epsilon/5 \\ =\ & \max_{i \in S} \tilde{E}[M_{\text{row}}^{(n)}(i, C)] - \epsilon/5 \\ \geq\ & \max_{i \in S} E[M_{\text{row}}^{(n)}(i, C)] - 2\epsilon/5 \\ >\ & t - 9\epsilon/10. \end{aligned} \tag{4.4}$$

The last inequality is from inequality (4.3). Since R and C use independent coin tosses, and since $\|M_{\text{row}}^{(n)}\| \leq n^k$,

$$\begin{aligned} & E[M_{\text{row}}^{(n)}(R, C)] \\ >\ & (1 - |S|2e^{-n})(t - 9\epsilon/10) + |S|2e^{-n}(-n^k) \\ >\ & t - \epsilon \end{aligned}$$

for sufficiently large n, as required. For smaller values of n so that the last inequality does not hold, item (i) is immediately true because a polynomial time Turing machine could use every strategy in $\{0, 1\}^n$.

Proof of (ii). In this case R himself evaluates $M_{\text{row}}^{(n)}$ on C's output strategy and each strategy in S. He chooses the best strategy in S for output.

Since the Immerman-Szelepcsenyi theorem [Imm88] can be directly extended to say that **NL**/poly = **coNL**/poly, our assumption implies that asking if the kth bit of $M^{(n)}_{\text{row}}(i,j)$ is 0 is also in **NL**/poly. We will also use the fact that **NL**/poly = **UL**/poly from [RA97].

Let C's output strategy be j and denote $S = \{i_1, \ldots, i_{|S|}\}$. The row player computes $M^{(n)}_{\text{row}}(i_t,j)$ for $t = 1, \ldots, |S|$. When computing $M^{(n)}_{\text{row}}(i_t,j)$, he guesses its first bit using a nondeterministic branch and verifies it in unambiguous logarithmic space. This can be done since **NL**/poly = **coNL**/poly = **UL**/poly. He then does the same for the second, third bit and so on. It is clear that only one branch guesses completely correctly and survives. Others are rejected. In this manner R proceeds by computing $M^{(n)}_{\text{row}}(i_t,j)$, $t = 1, \ldots, |S|$, one by one. Instead of saving all these values, which takes space polynomial in n, he needs only store the best strategy he has evaluated so far. The corresponding $M^{(n)}_{\text{row}}$-value can be computed on the fly whenever it is needed. These observations yield that R runs in unambiguous logarithmic space (in n).

For an arbitrary randomized n^d-size $M^{(n)}$-oracle player C, let $i \in S$ be such that $E[M^{(n)}_{\text{row}}(i,C)] > t - \epsilon$. Since whatever strategy C takes, R always chooses a strategy no worse than i, we must have $E[M^{(n)}_{\text{row}}(R,C)] > t - \epsilon$.

Proof of (iii). In this case R just needs to know which $i \in S$ makes $E[M^{(n)}_{\text{row}}(i,C)] > t - \epsilon$. This information takes $O(\log n)$ bits.

□

Appendix A

Complexity Classes

A brief definition of several complexity classes are given. For detailed definitions, please refer to [AK].

1. **P** is the class of languages decidable in polynomial time.
2. **P**-hard is the class of languages logarithmic-space reducible from all languages in **P**.
3. **EXP** is the class of languages decidable in exponential time.
4. **EXP**-hard is the class of languages polynomial-time reducible from all languages in EXP.
5. **EXP**-complete = **EXP** $\cap$ **EXP**-hard.
6. promise-$\mathbf{S_2^P}$ is the class of languages L such that there are disjoint sets Π^+, Π^- with $\Pi^+ \subseteq L$, $\Pi^- \cap L = \emptyset$ and there is a polynomial-time computable predicate $R(x, y, z)$ for $|y| = |z| = poly(|x|)$ satisfying the following: $\forall x \in \Pi^+$, $\exists y \forall z R(x, y, z) = 1$ and $\forall x \in \Pi^-$, $\exists z \forall y R(x, y, z) = 0$.

7. **RP** is the class of languages L such that there is a nondeterministic polynomial time Turing machine which on input $x \in L$ accepts on at least 1/2 of its computation paths and on input $x \notin L$ rejects on every computation path.

8. co**RP** is the class of languages whose complement is in **RP**.

9. $\mathbf{ZPP} = \mathbf{RP} \cap \mathrm{co}\mathbf{RP}$

10. **NP** is the class of languages L such that there is a nondeterministic polynomial-time Turing machine which on input x accepts on at least one computation path if and only if $x \in L$.

11. **PPAD** is the class of function problems of the following form. Given a polynomial time algorithm P that, on any input x, implicitly defines a directed graph $G(x)$ with nodes $\Sigma^{p(|x|)}$ by outputting for each $y \in \Sigma^{p(|x|)}$ the vertices pointing to or from y, where p is a polynomial and Σ is some constant size alphabet. The graph $G(x)$ is restricted to be one in which each vertex has indegree and outdegree at most one. Given a source of $G(x)$ (i.e., one with indegree zero), the problem is to find another source or sink.

12. **PPAD**-complete is the class of function problems in **PPAD** reducible from every function problem in **PPAD**.

13. **NL**/poly is the class of languages L such that there is a nondeterministic logarithmic-space Turing machine and a sequence of polynomially-long (in n) advices $\{a_n\}_{n \in \mathbb{N}}$ such that on inputs x and $a_{|x|}$, the Turing machine has an accepting computation path if and only if $x \in L$.

14. **UL**/poly is the same as **NL**/poly except the logarithmic-space Turing

machine must have exactly one computation path whatever input it is given.

15. co**NL**/poly is the class of languages whose complement is in **NL**/poly.

Bibliography

[AK] S. Aaronson and G. Kuperberg, http://www.complexityzoo.com/.

[CD06] X. Chen and X. Deng, *2D-SPERNER is **PPAD**-complete*, Submitted to STOC (2006).

[Che52] H. Chernoff, *A measure of the asymptotic efficiency of tests of a hypothesis based on the sum of observations*, Annals of Mathmematical Statistics **23** (1952), 493–507.

[DGP05] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou, *The complexity of computing a nash equilibrium*, Tech. Report TR05-115, Electronic Colloquium on Computational Complexity, 2005.

[FIKU05] L. Fortnow, R. Impagliazzo, V. Kabanets, and C. Umans, *On the complexity of succinct zero-sum games*, Proceedings of the 20th IEEE Conference on Computational Complexity, 2005, pp. 323–332.

[FKS95] J. Feigenbaum, D. Koller, and P. Shor, *A game-theoretic classification of interactive complexity classes*, Proceedings of the 10th Annual IEEE Conference on Computational Complexity, 1995, pp. 227–237.

[FPS03] L. Fortnow, A. Pavan, and S. Sengupta, *Proving sat does not have small circuits with an application to the two queries problem*, Proceedings of the 18th Annual IEEE Conference on Computational Complexity, 2003, pp. 347–357.

[GK92] M. Grigoriadis and L. Khachiyan, *Approximating solution of matrix games in parallel*, Advances in Optimization and Parallel Computing (Amsterdam) (P. Pardalos, ed.), Elsevier, 1992, pp. 129–136.

[GK95] ———, *A sublinear-time randomized approximation algorithm for matrix games*, Operations Research Letters **18** (1995), no. 2, 53–58.

[Hoe63] W. Hoeffding, *Probability inequalities for sums of bounded random variables*, Journal of the American Statistical Association **58** (1963), no. 301, 13–30.

[Imm88] N. Immerman, *Nondeterministic space is closed under complementation*, SIAM Journal on Computing (1988), 935–938.

[KB04] R. H. Katz and G. Borriello, *Contemporary logic design*, 2nd ed., Prentice Hall, 2004.

[Kha79] L. G. Khachiyan, *A polynomial algorithm in linear programming*, Soviet Mathematics Doklady **20** (1979), 191–194.

[LN93] M. Luby and N. Nisan, *A parallel approximation algorithm for positive linear programming*, Proceedings of the 25th Annual ACM Symposium on Theory of Computing, 1993, pp. 448–457.

[Nas51] J. Nash, *Noncooperative games*, Annals of Mathematics **54** (1951), 289–295.

[New91] I. Newman, *Private vs. common random bits in communication complexity*, Information Processing Letters **39** (1991), 67–71.

[OR94] M. J. Osborne and A. Rubinstein, *A course in game theory*, MIT Press, 1994.

[Owe82] G. Owen, *Game theory*, Academic Press, 1982.

[PST95] S. Plotkin, D. Shmoys, and E. Tardos, *Fast approximation algorithms for fractional packing and covering problems*, Mathematics of Operations Research **20** (1995), no. 2, 257–301.

[RA97] Klaus Reinhardt and Eric Allender, *Making nondeterminism unambiguous*, Proceedings of the 38th Annual IEEE Symposium on Foundations of Computer Science, 1997, pp. 244–253.

[Rud76] W. Rudin, *Principles of mathematical analysis*, 3rd ed., McGraw-Hill, 1976.

[Sha49] C. E. Shannon, *Communication in the presence of noise*, IRE **37** (1949), 10–21.

[Sip05] M. Sipser, *Introduction to the theory of computation*, 2nd ed., Course Technology, 2005.