These lecture notes are based on a manuscript of a book on matrix rigidity by Chi-Ning Chou and Sasha Golovnev.

Chapter 1

Introduction

1.1 Definitions and examples

Lecture 1

One of the main questions in computational complexity is that of proving lower bounds on the size of Boolean circuits computing explicitly given functions. While most Boolean functions of n inputs require circuits of size $2^n/n$ [Sha49a, Lup59a], we can only prove small linear lower bounds for explicitly defined functions [LR01, IM02, Blu83, DK11, FGHK16]. [1]

The same question remains open for *linear* circuits computing *linear* Boolean functions. Since any linear function with one output can be computed by a circuit of size at most n, we study linear functions with n inputs and n outputs. A random linear map with n outputs requires circuits of size $n^2/\log n$ [Lup56], but the best known lower bound for an explicit linear map is only $3n - o(n)$ [Cha94a].

The notion of matrix rigidity was introduced by Valiant [Val77] as a tool for proving lower bounds against linear circuits. (A related notion of separability was introduced by Grigoriev [Gri76].)

We will use the following notation. A matrix A is called s-sparse, if the number of non-zero entries in A is at most s. We will use $I_n, 0_n$ and J_n to denote the identity matrix, zero matrix, and all-ones matrix of size $n \times n$. For a matrix $A \in \mathbb{F}^{n\times n}$, by $\|A\|_0$ we denote the number of non-zero entries in A.

Definition 1.1 (Rigidity)**.** *Let $\mathbb{F}$ be a field, $A \in \mathbb{F}^{n\times n}$ be a matrix, and $0 \le r \le n$. The rigidity of A over $\mathbb{F}$, denoted by $\mathcal{R}_A^{\mathbb{F}}(r)$, is the Hamming distance between A and the set of matrices of rank at most r. Formally,*

$$\mathcal{R}_A^{\mathbb{F}}(r) := \min_{\operatorname{rank}(A+C)\le r} \|C\|_0 \ .$$

In other words, a matrix A has rigidity $\mathcal{R}_A^{\mathbb{F}}(r) \ge s$ if and only if $A \in \mathbb{F}^{n\times n}$ *cannot* be written as a sum

$$A = S + L \ ,$$

where $S \in \mathbb{F}^{n\times n}$ is $(s-1)$-sparse matrix, and $L \in \mathbb{F}^{n\times n}$ is low-rank: $\operatorname{rank}(L) \le r$.

Valiant [Val77] proved that any linear map $A \in \mathbb{F}^{n\times n}$ computed by a linear circuit (over a field $\mathbb{F}$) of depth $O(\log n)$ and size $o(n \log\log n)$ has rigidity at most $\mathcal{R}_A^{\mathbb{F}}(\varepsilon n) \le n^{1+\delta}$ for every constant $\varepsilon, \delta > 0$. Therefore, an explicit matrix of higher rigidity would give us a super-linear lower bound against linear circuits of logarithmic depth. Despite more than 40 years of research, the problem of proving super-linear lower bounds for such circuits remains open.

Let us now see the rigidity of a few specific matrices.

- If $A \in \mathbb{F}^{n\times n}$ has rank $\operatorname{rank}(A) = k$ over the field $\mathbb{F}$, then $\mathcal{R}_A^{\mathbb{F}}(r) = 0$ for every $r \ge k$. Indeed, A can be written as a sum of A and 0_n, where $\operatorname{rank}(A) \le r$ and 0_n is 0-sparse. Similarly, an s-sparse matrix $A \in \mathbb{F}^{n\times n}$ has rigidity $\mathcal{R}_A^{\mathbb{F}}(r) \le s$ for any value of r.

[1] Here by explicit functions we mean functions computable in time polynomial in n. We will later discuss the notion of explicitness in greater detail.

- For any $0 \le r \le n$, $\mathcal{R}^{\mathbb{F}}_{I_n}(r) = n - r$. Indeed, if we change $n-r$ ones of I_n to zeros, then the resulting matrix has rank r, which implies that $\mathcal{R}^{\mathbb{F}}_{I_n}(r) \le n-r$. On the other hand, for any $(n-r)$-sparse matrix B, from subadditivity of rank,

$$\mathrm{rank}(I_n + B) \ge \mathrm{rank}(I_n) - \mathrm{rank}(B) \ge n - (n-r) = r \ ,$$

which gives us that $\mathcal{R}^{\mathbb{F}}_{I_n}(r) \ge n - r$.

- Let n be a multiple of $2r$, and let $M_n \in \mathbb{F}^{n \times n}$ be a matrix consisting of matrices I_{2r} stacked together side by side:

$$M_n = \begin{pmatrix} I_{2r} & \cdots & I_{2r} \\ \vdots & \ddots & \vdots \\ I_{2r} & \cdots & I_{2r} \end{pmatrix} \ .$$

We will show that this matrix has rigidity $\mathcal{R}^{\mathbb{F}}_{A}(r) = \frac{n^2}{4r}$.

Theorem 1.2 ([Mid05]). *For any field $\mathbb{F}$, and any n divisible by $1 \le 2r \le n$,*

$$\mathcal{R}^{\mathbb{F}}_{M_n}(r) = \frac{n^2}{4r} \ .$$

Proof of Theorem 1.2. M_n consists of $\frac{n^2}{4r^2}$ copies of the identity matrix I_{2r}. In order to drop the rank of A to r, the rank of each copy of I_{2r} must be dropped to r. From the previous example we know that in order to decrease the rank of I_{2r} to r, one needs to change at least r elements. Thus, $\frac{n^2}{4r^2} \cdot r = \frac{n^2}{4r}$ entries of M_n must be changed. Note that this bounds is tight, *i.e.*, $\mathcal{R}^{\mathbb{F}}_{M_n}(r) = \frac{n^2}{4r}$. □

The bound of Theorem 1.2 easily generalizes to all values $r \le n/2$ with a loss of a multiplicative factor of 2. This theorem was proven by Midrijānis [Mid05], and it gives a simple matrix with rigidity $\mathcal{R}^{\mathbb{F}}_{M_n}(r) \ge \frac{n^2}{8r}$.

We will see later that there exist matrices with much higher rigidity $\tilde{\Omega}\left((n-r)^2\right)$. Embarrassingly, the best known lower bound for an *explicit* matrix improves on the $\frac{n^2}{8r}$ bound only by a logarithmic factor.

1.2 Circuit Complexity

A circuit corresponds to a simple straight line program where every instruction performs a binary operation on two operands, each of which is either an input or the result of a previous instruction. The structure of this program is extremely simple: no loops, no conditional statements. Still, we know no functions in P (or even NP, or even E^{NP}) that requires even $3.1n$ binary instructions ("size") to compute on inputs of length n. This is in sharp contrast with the fact that it is easy to *non-constructively* find such functions: simple counting arguments show a random function on n variables has circuit size $\Omega(2^n/n)$ with probability $1 - o(1)$ [Sha49b, Lup59b].

For small-depth circuits we know several strong lower bounds. (Note that when working with circuits of constant depth, we do not pose bounds on the fan-ins of the gates.) Depth-2 circuits (after a simple normalization) are just CNFs or DNFs. It is easy to see that the parity function $\oplus_n$ of n inputs requires CNFs and DNFs of size $\Omega(2^n)$. For depth-d circuits, we know a lower bound of $2^{\Omega(n^{(1/(d-1))})}$ [Hås86, HJP93, PPZ97, Bop97, PPSZ05, MW17]. Thus, for depth $d = o(\log n / \log\log n)$ we have non-trivial lower bounds even if the fan-ins of the gates are unbounded. For circuits with fan-in 2, we known functions which cannot be computed by circuits of depth $1.99 \log n$ [Nec66]. Thus, a problem on the frontier is

Problem 1.3. *Prove a lower bound of $10n$ against circuits of depth $10 \log n$.*

More generally, a lower bound of $\omega(n)$ against circuits of depth $O(\log n)$.

Super-linear lower bounds are not known even for linear circuits, i.e., circuits consisting of only gates computing linear combinations of their two inputs. Note that every linear function with one output has a circuit of size $n-1$ (and depth $\log n$). For linear circuits, we consider *linear transformations*, multi-output functions of the form $f(x) = Ax$ where $A \in \mathbb{F}^{n\times n}$. For a random matrix $A \in \{0,1\}^{n\times n}$, the size of the smallest linear circuit computing Ax is $\Theta(n^2/\log n)$ [Lup56] with probability $1 - o(1)$, but for explicitly-constructed matrices the strongest known lower bound is $3n - o(n)$ [Cha94b]. This leads us to another problem on the frontier:

Problem 1.4. *Prove a lower bound of $\omega(n)$ against* linear *circuits of depth $O(\log n)$.*

Formally, Problem 1.3 and Problem 1.4 are incomparable, as in the linear case we study a weaker computational model (which makes it easier to prove lower bounds), but are limited to proving lower bounds for a smaller class of problem (which makes it harder to prove lower bounds).

1.3 Circuits and Rigidity

In this section, we will present a seminal result of Valiant [Val77] showing that rigid matrices require log-depth circuits of super-linear size. We start with the definition of linear circuits.

Definition 1.5 (Linear circuits)**.** *Let $\mathbb{F}$ be a field and $n \in \mathbb{N}$. A circuit C with n inputs and n outputs is a directed acyclic graph where n vertices have fan-in zero and are labeled by the inputs, all other vertices have fan-in two and are labeled with* affine functions *(over $\mathbb{F}$) of their two inputs, n of these vertices are labeled as outputs. For every fixed input, the value at each node is computed by applying the corresponding functions. Such a circuit C naturally defines a linear map $f\colon \mathbb{F}^n \to \mathbb{F}^n$, and the corresponding matrix $A \in \mathbb{F}^{n\times n}$ such that $f(x) = Ax$.*

The depth $d(C)$ of a circuit C is the length of the longest path in the circuit. The size $s(C)$ of C is defined as the number of vertices in C.

The following theorem shows a connection between lower bounds for linear circuits and matrix rigidity.

Theorem 1.6. *Let $\mathbb{F}$ be a field, and $A \in \mathbb{F}^{n\times n}$ be a family of matrices for $n \in \mathbb{N}$. If $\mathcal{R}^{\mathbb{F}}_A(\varepsilon n) > n^{1+\delta}$ for constant $\varepsilon, \delta > 0$, then any $O(\log n)$-depth linear circuit computing $x \to Ax$ must be of size $\Omega(n \cdot \log\log n)$.*

The proof of Theorem 1.6 repeatedly uses the following beautiful graph theoretic lemma due to Erdös, Graham, and Szemerédi [EGS76]: If G is a directed acyclic graph with s edges and of depth d, then there is a set of $s/\log d$ edges whose removal decreases the depth of G by a factor of two. We will follow the proof of this lemma from [Vio09].

Lemma 1.7 ([EGS76])**.** *Let G be an acyclic digraph with s edges and of depth $d = 2^k$. There exists a set of $s/\log d$ edges in G such that after their removal, the longest path in G has length at most $d/2$.*

Proof of Lemma 1.7. For ease of exposition, we follow [Vio09] and define a *depth function*. Let $G = (V, E)$ be an acyclic digraph. We say that $D : V \to \{0, 1, \ldots, d\}$ is a depth function for G if for any $(a, b) \in E$, $D(a) < D(b)$. It is not difficult to see that G has depth at most d if and only if there exists a depth function $D : V \to \{0, 1, \ldots, d-1\}$ for G.

We start with G of depth at most $d = 2^k$, and its depth function $D : V \to \{0, 1, \ldots, 2^k\}$. Now, consider the following partition of E using the depth function D. For each $i \in [k]$, define

$$E_i = \{(a, b) \in E : \text{ the most significant bit where } D(a), D(b) \text{ differ is the } i^{\text{th}} \text{ bit}\}.$$

As $\{E_i\}_{i\in[k]}$ is a partition of E, by the averaging argument, there exists $i^* \in [k]$ such that

$$|E_{i^*}| \le \frac{|E|}{k} \le \frac{|E|}{\log d}\,.$$

Now, it suffices to show that the depth of $G' = (V, E')$, where $E' = E \backslash E_{i^*}$, is at most 2^{k-1}. This can be shown by exhibiting a depth function $D' : V \to \{0, 1, \ldots, 2^{k-1} - 1\}$ for G'. The following shows that we can take $D'(v)$ to be $D(v)$ without the $i^{*\text{th}}$ bit.

Consider an edge $(a, b) \in E'$. Since $(a, b) \in E$, $D(a) < D(b)$. In particular, there exists $i \in [k]$ such that the most significant bit where $D(a)$ and $D(b)$ differ is i. Since $(a, b) \in E'$, the edge (a, b) was not removed, so $i \ne i^*$. Therefore, after removing the bit i^*, this bit i is still the most significant bit where $D'(a)$ and $D'(b)$ differ. This implies that $D'(a) < D'(b)$, and that $D' : V \to \{0, 1, \ldots, 2^{k-1} - 1\}$ is a depth function for G'. □

Bibliography

[Blu83] Norbert Blum. A Boolean function requiring $3n$ network size. *Theoretical Computer Science*, 28(3):337–345, 1983.

[Bop97] Ravi B. Boppana. The average sensitivity of bounded-depth circuits. *Inf. Process. Lett.*, 63(5):257–261, 1997.

[Cha94a] Aleksandr V. Chashkin. On the complexity of Boolean matrices, graphs and their corresponding Boolean functions. *Diskretnaya matematika*, 4(3):229–257, 1994.

[Cha94b] Aleksandr V. Chashkin. On the complexity of Boolean matrices, graphs and their corresponding Boolean functions. *Discrete Math. and Appl.*, 4(3):229–257, 1994.

[DK11] Evgeny Demenkov and Alexander S Kulikov. An elementary proof of a $3n - o(n)$ lower bound on the circuit complexity of affine dispersers. In *Proceedings of the 36th Internationl Symposium on the Mathematical Foundations of Computer Science (MFCS 2011)*, pages 256–265. Springer, 2011.

[EGS76] Paul Erdös, Ronald L. Graham, and Endre Szemerédi. On sparse graphs with dense long paths. *Computers and Mathematics with Applications*, pages 365–369, 1976.

[FGHK16] Magnus Find, Alexander Golovnev, Edward A. Hirsch, and Alexander S. Kulikov. A better-than-$3n$ lower bound for the circuit complexity of an explicit function. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2016)*, pages 89–98. IEEE, 2016.

[Gri76] Dmitrii Yu. Grigoriev. Application of separability and independence notions for proving lower bounds of circuit complexity. *Zapiski Nauchnykh Seminarov POMI*, 60:38–48, 1976.

[Hås86] Johan Håstad. Almost optimal lower bounds for small depth circuits. In *STOC 1986*, pages 6–20, 1986.

[HJP93] Johan Håstad, Stasys Jukna, and Pavel Pudlák. Top-down lower bounds for depth 3 circuits. In *FOCS 1993*, pages 124–129, 1993.

[IM02] Kazuo Iwama and Hiroki Morizumi. An explicit lower bound of $5n - o(n)$ for Boolean circuits. In *Proceedings of the 27th Internationl Symposium on the Mathematical Foundations of Computer Science (MFCS 2002)*, pages 353–364. Springer, 2002.

[LR01] Oded Lachish and Ran Raz. Explicit lower bound of $4.5n - o(n)$ for Boolean circuits. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)*, pages 399–408. ACM, 2001.

[Lup56] Oleg B. Lupanov. On rectifier and switching-and-rectifier schemes. In *Dokl. Akad. Nauk SSSR*, volume 111, pages 1171–1174, 1956. In Russian.

[Lup59a] Oleg B. Lupanov. A method of circuit synthesis. *Izv. VUZov, Radiofizika*, 1:120–140, 1959. In Russian.

[Lup59b] Oleg B. Lupanov. A method of circuit synthesis. *Izv. VUZov, Radiofizika*, 1:120–140, 1959. In Russian.

[Mid05] Gatis Midrijānis. Three lines proof of the lower bound for the matrix rigidity. *arXiv preprint cs/0506081*, 2005.

[MW17] Or Meir and Avi Wigderson. Prediction from partial information and hindsight, with application to circuit lower bounds. In *ECCC*, volume 24, 2017.

[Nec66] Edward I. Nechiporuk. On a Boolean function. *Dokl. Akad. Nauk SSSR*, 169(4):765–766, 1966.

[PPSZ05] Ramamohan Paturi, Pavel Pudlák, Michael E Saks, and Francis Zane. An improved exponential-time algorithm for k-SAT. *J. ACM*, 52(3):337–364, 2005.

[PPZ97] Ramamohan Paturi, Pavel Pudlák, and Francis Zane. Satisfiability coding lemma. In *FOCS 1997*, pages 566–574, 1997.

[Sha49a] Claude E. Shannon. The synthesis of two-terminal switching circuits. *Bell Syst. Tech. J.*, 28:59–98, 1949.

[Sha49b] Claude E. Shannon. The synthesis of two-terminal switching circuits. *Bell Syst. Tech. J.*, 28:59–98, 1949.

[Val77] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In *Proceedings of the 2nd Internationl Symposium on the Mathematical Foundations of Computer Science (MFCS 1977)*, pages 121–127. Springer, 1977.

[Vio09] Emanuele Viola. On the power of small-depth computation. *Foundations and Trends in Theoretical Computer Science*, 5(1):1–72, 2009.