Diplomarbeit

Sylvester-Gallai Theorems and Identities over $\mathbb{R}$

(Sylvester-Gallai Theoreme und Identitäten über $\mathbb{R}$)

Angefertigt am
Mathematischen Institut

Vorgelegt der
Mathematisch-Naturwissenschaftlichen Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

August 2010

Von
Nils Frohberg
Aus
Wien

Diplomarbeit vorgelegt von Nils Frohberg, geboren am 26. Januar 1982 in Wien, Matrikelnummer 1582175.
Betreuer: Prof. Dr. Nitin Saxena, Hausdorff Center for Mathematics, Bonn.

Acknowledgements

I would like to thank Prof. Dr. Nitin Saxena for his kind support, unlimited patience, and invaluable discussions on (and off!) topics of this paper. His clean, simple, and well structured approach to new and existing topics is refreshing and inspiring.

I also want to thank my brother, Jens, and Malte Beecken for critically reading this paper and giving good advice on improving it.

Also, I want to thank my friends and colleagues for great times spent together, the good (and not so good) coffee, and for all the support I could wish for. Go Team Red!

Finally, but most importantly, I am deeply obliged to my mother, Helga, who has made uncountably many sacrifices to give me this opportunity. Thank you.

Contents

Zusammenfassung

Polynomielle Identitätstests, oder *polynomial identity testing* (PIT) behandelt das Problem, zwei Polynome $P, Q \in \mathbb{F}[x_1, \ldots, x_n]$ zu vergleichen. Äquivalent ist es, zu überprüfen, ob ein Polynom das Nullpolynom ist. So wird aus dem ursprünglichen Problem einfach der Test, ob das Polynom $(P - Q) \in \mathbb{F}[x_1, \ldots, x_n]$ für jegliche Eingabe immer 0 auswertet. Wir werden in dieser Arbeit immer letztere Darstellung von PIT benutzen.

Ziel dieser Arbeit ist es, eine in sich abgeschlossene Präsentation der derzeitigen Methoden, dieses Problem anzugehen, zu geben. Außerdem werden neue mögliche Ansätze zur weiteren Behandlung von PIT vorgeschlagen.

Polynome lassen sich durch *circuits* („Schaltungen") darstellen. Diese sind azyklische Graphen, die genau eine Wurzel haben, und deren Blätter aus den Variablen des Polynoms bestehen. Es gibt Additions– und Multiplikationsknoten, und die Äste können Konstanten tragen, die mit dem jeweiligen darunterliegenden Knoten multipliziert werden. Wir werden nur solche Graphen betrachten, die aus drei Ebenen bestehen[1]: Die Wurzelebene ist hat genau einen Additionsknoten. Die nächste Ebene besteht aus genau k Multiplikationsknoten, welche jeweils genau d Additionsknoten in der nächsten Ebene besitzen. Wir arbeiten mit n Variablen.

Es gibt verschiedene unterschiedliche Herangehensweisen an PIT. Wir geben einen Überblick über drei Methoden, und eine ausführliche Darstellung der Methode, die als ersten eine vom Grad des Polynoms unabhängige „Rang-Schranke" aufweist.

Algorithmen, um PIT zu bestimmen, lassen sich in *non–blackbox–* und *blackbox–* Algorithmen einteilen. Erstere haben die möglichkeit, zur Laufzeit in die Darstellung vom Polynom „hineinzusehen", während dies den letzteren verboten ist.

Wegen seiner Einfachheit zeigen wir zuerst einen blackbox–Algorithmus, den Schwarz–Zippel Algorithmus [Zip79, Sch80]. Er basiert darauf, daß die Wahrscheinlichkeit, die Nullstelle eines Polynoms „zufällig" zu treffen, sehr gering ist, solange der Grad des Polynoms klein ist gegenüber dem zugrundeliegenden Körper. Danach präsentieren wir einen non–blackbox–Algorithmus, der PIT mit einer Laufzeit von $\text{poly}(n, d^k)$ bestimmt [KS07]. Weiterhin zeigen wir dann eine Schranke für den Rang eines „circuits" C [SS09] von

$$\text{rank}(C) \leq O(k^3 \log d).$$

Durch eine Identität vom Rang $(\log d + 2)$ über $\mathbb{F}$ werden wir zeigen, dass letztere Schranke tatsächlich fast optimal über beliebige Körper ist. Somit wurde Dvir und Shpilkas Vermutung [DS05], die obere Schranke wäre $O(k)$, für allgemeine Körper widerlegt. Allerdings haben Kayal und Saraf [KS09] es geschafft zu zeigen, daß über

[1] einschließlich der Wurzelebene

$\mathbb{R}$ die obere Schranke tatsächlich nicht vom Grad d abhängt, indem sie eine obere Schranke von

$$\text{rank}(C) \leq k^{O(k)}$$

über $\mathbb{R}$ bewiesen haben.

Der Beweis dieser Schranke verbindet Graphen mit Geometrie, und benutzt hauptsächlich geometrische Ansätze, um ans Ziel zu kommen. In dieser Arbeit werden wir zum ersten Mal einen in sich abgeschlossenen und vollständigen Beweis dieser oberen Schranke erbringen. D.h., daß wir sowohl auf die geometrischen Theoreme eingehen werden, sowie auf den Beweis der Schranke selbst, einschließlich der Verknüpfung zwischen den beiden.

Außerdem präsentieren wir weitere neue Ansätze, um mit PIT umzugehen. Auch diese beruhen auf dem Zusammenhang zwischen Geometrie und PIT. Eine Bearbeitung dieser Ansätze könnte neue Einsichten zum Verhalten von PIT mit Graphen mit 3 Ebenen bringen.

Zuletzt sei noch bemerkt, dass Saxena und Seshadhri [SS10] kürzlich eine quadratische obere Schranke

$$\text{rank}(C) \leq O(k^2)$$

angekündigt haben. Wir werden auch hierauf kurz eingehen. Obwohl dies nun schon sehr nah an Dvir und Shpilkas obere Schranke herankommt, kann es aber durchaus noch Sinn machen, weiter an PIT mit 3 Ebenen zu arbeiten, unter anderem, um daraus womöglich auch Erkenntnisse für PIT mit 4 Ebenen zu gewinnen.

Chapter 1. Introduction

In many fields of mathematics we often need to compare two polynomials, or need to check whether a given polynomial outputs zero on all input. We already learn early on that, although two polynomials may have a different appearance, they can "represent" the same function, i.e. they can always compute the same output, e.g. the difference of squares:

$$a^2 - b^2 = (a - b)(a + b), \tag{1.1}$$

or the sum of four identities:

$$\begin{aligned}(a_1^2 + a_2^2 + a_3^2 + a_4^2)\,(b_1^2 + b_2^2 + b_3^2 + b_4^2) = \\ (a_1b_1 - a_2b_2 - a_3b_3 - a_4b_4)^2 + \\ (a_1b_2 + a_2b_1 + a_3b_4 - a_4b_3)^2 + \\ (a_1b_3 - a_2b_4 + a_3b_1 + a_4b_2)^2 + \\ (a_1b_4 + a_2b_3 - a_3b_2 + a_4b_1)^2.\end{aligned}$$

Comparing two polynomials is also required in a lot of higher level algorithms, e.g. for primality testing, or in fundamental structural results, e.g. in Complexity Theory for PCP related theorems. Most importantly, polynomial identity testing is also connected to circuit lower bounds [KI04, Agr05].

In theory, identities can be easily proven by expanding the products and comparing the resulting monomials. But the problem with this method is that the number of monomials "explodes" when increasing the number of variables and the degree. Thus, we need a way of formalizing the problem.

1.1 Basics

In order to study polynomial identities, we will start off by stating a few basic definitions. We are looking for ways to represent polynomials such that we can classify them and perform computations on their basic properties.

Unless explicitly denoted by $\mathbb{R}$ or $\mathbb{C}$, $\mathbb{F}$ will always be an arbitrary field, i.e. either finite or infinite.

Definition 1.1.1 (Polynomial Identity Testing)**.** Given a polynomial

$$P(x_1, \ldots, x_n) \in \mathbb{F}[x_1, \ldots, x_n],$$

polynomial identity testing (PIT) is the problem of deciding whether P is the zero polynomial, i.e. if

$$P(x_1, \ldots, x_n) = 0$$

for all $(x_1, \ldots, x_n) \in \mathbb{F}^n$.

Remark 1.1.2. The problem of deciding whether two polynomials P and Q are equal can be simply transformed into PIT by deciding whether $(P - Q)$ is the zero polynomial.

We will thus always be testing polynomials against $0 \in \mathbb{F}[x_1, \ldots, x_n]$, i.e. the *zero polynomial*.

Let us now introduce a representation of polynomials taken from Complexity Theory: arithmetic circuits. They are the extension of Boolean circuits over arbitrary fields.

Definition 1.1.3 (Arithmetic circuit)**.** An *arithmetic circuit* C on n variables and a field $\mathbb{F}$ is an acyclic graph with only one root and with input variables at the leaves and output at the root. The two types of nodes, called *gates*, are $\sum$ (*addition*) and $\prod$ (*multiplication*) and correspond to their respective operation in the field $\mathbb{F}$. The edges of C, called *wires*, can hold constants from $\mathbb{F}$ which are multiplied to the value at the tail of the edge. The *fanin of a gate* is the number of its inputs. The *fanin of* C is the largest fanin of any gate inside C.

It should be obvious that a circuit with n different input variables has an n-variate polynomial over $\mathbb{F}$ at its root. It is therefore an elegant way for us to describe polynomials.

Remark 1.1.4. By abuse of notation, we will often denote a circuit and its corresponding polynomial by the same symbol, e.g. C.

Example 1.1.5. The circuit corresponding to the polynomial $(2x - y^2)$ can be graphed as follows:

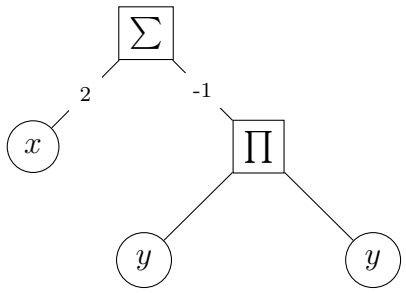

Remark 1.1.6. Just as polynomial representations are not unique, circuit representations are also not unique.

With a small formal trick, we are able to classify circuits. In later chapters, we will restrict ourselves to specific classes of circuits.

Definition 1.1.7. By adding trivial gates of fanin 1, any circuit can be written as alternating levels of $\sum$ and $\prod$ gates. Let C be such a circuit. We call C a $\Sigma\Pi\Sigma\ldots$ or $\Pi\Sigma\Pi\ldots$ circuit, depending on whether it has, respectively, a $\sum$ or $\prod$ root–level gate.

From now on, a *circuit* always has alternating levels, i.e. it is either a $\Sigma\Pi\Sigma\ldots$ or a $\Pi\Sigma\Pi\ldots$ circuit.

Definition 1.1.8. A circuit is called a *depth–h* circuit if it has exactly h levels of gates.

A depth–3 $\Sigma\Pi\Sigma$–circuit with n variables, top–fanin k, and 2^{nd}–level–fanin[1] d (called the *degree* d) will be denoted by $\Sigma\Pi\Sigma(k,d,n)$. We can then also define sets of depth–3 circuits:

$$\begin{aligned}\Sigma\Pi\Sigma &= \{\Sigma\Pi\Sigma(k,d,n) \mid k,d,n \in \mathbb{N}\},\\ \Sigma\Pi\Sigma(k) &= \{\Sigma\Pi\Sigma(k,d,n) \mid d,n \in \mathbb{N}\},\\ \Sigma\Pi\Sigma(k,d) &= \{\Sigma\Pi\Sigma(k,d,n) \mid n \in \mathbb{N}\}.\end{aligned}$$

Remark 1.1.9. It can easily be seen that the degree of a circuit coincides with the degree of its polynomial.

Example 1.1.10. From the difference of squares (1.1), we know that the output of the following $\Sigma\Pi\Sigma\,(3,2,2)$ circuit is always zero:

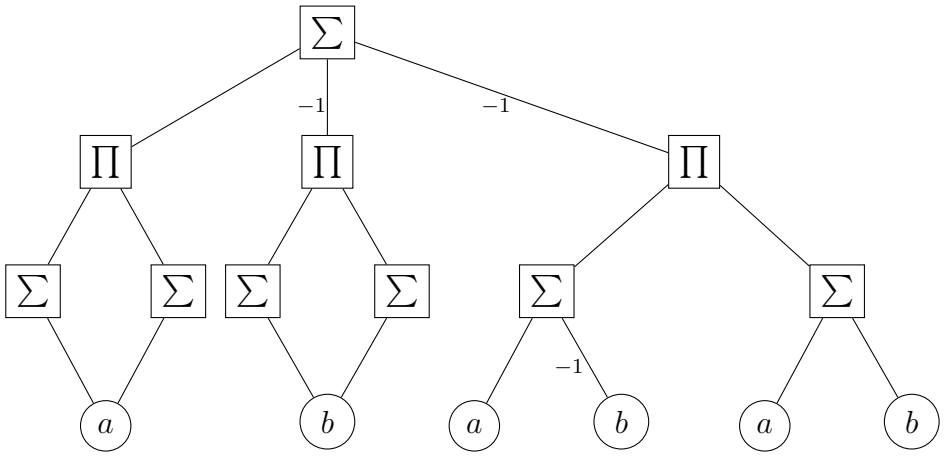

Definition 1.1.11. For $n \in \mathbb{N}$, we denote the set $\{1,2,\ldots,n\}$ by $[n]$.

When working with circuits, we will often need the following properties:

Definition 1.1.12. A depth–3, $\Sigma\Pi\Sigma\,(k,d,n)$, circuit

$$C = T_1 + \ldots + T_k = \prod_{j\in[d]} l_{1j} + \ldots + \prod_{j\in[d]} l_{kj}$$

is called *minimal* if there is no $S \subset [k], |S| > 0$, such that

$$\sum_{j\in S} T_j = 0.$$

[1] i.e., the maximum fanin of any of the $\prod$ gates

The *greatest common denominator* of C, $\gcd(C)$ is defined as

$$\gcd(C) = \gcd_{i \in [k]} T_i.$$

It is called *simple* if

$$\gcd(C) = 1.$$

We define $\mathrm{sim}(C)$ to be

$$\mathrm{sim}(C) := \frac{C}{\gcd(C)}.$$

The *rank of* C is the following: Identify each linear form $l_{ij} = \sum_{\kappa \in [n]} a_{ij\kappa} X_{ij\kappa}$ of C with the vector $(a_{ij1}, \ldots, a_{ijn}) \in \mathbb{R}^n$. Then, set

$$\mathrm{rank}(C) := \dim(\mathrm{span}\{l_{ij} \mid i \in [k], j \in [d]\}).$$

1.2 Dvir and Shpilka Rank Bound Conjecture

In 2005, Dvir and Shpilka conjectured a rank bound of $O(k)$ for simple, minimal depth–3 identities (Chapter 7, [DS05]). I.e., for every fanin size k, there is a constant $c(k)$, such that essentially all $\Sigma\Pi\Sigma(k)$ circuits computing the zero polynomial have rank at most $c(k)$.

Note that the conjecture for $c(k)$ is independent of the degree d of the circuit. All rank bounds prior to Kayal and Saraf's $k^{O(k)}$ rank bound [KS09] also depend on the degree d. We will present these "weaker"[2] results in Chapter 2, and the $k^{O(k)}$ rank bound as part of Chapter 4.

In the meantime, Saxena and Seshadhri could improve the bound[3] to a very promising $O(k^2)$ [SS10]. We briefly pick up this result in Chapter 5.

1.3 Overview

In Chapter 2, we will give an overview of various algorithms and rank bound in order to decide PIT. Here, we will also only be giving restricted versions[4] of their proofs.

Chapter 3 will provide the geometric tools needed by the lower rank bound proofs and for further studying PIT.

The main goal of this paper is to give a self–contained proof of the $k^{O(k)}$–rank bound for PIT. Also, we will suggest new methods for working with matchings, a fundamental tool that proved valuable when studying PIT. Both of these can be found in Chapter 4.

[2]Weaker only over $\mathbb{R}$, since the rank bound of $k^{O(k)}$ only works over $\mathbb{R}$. Therefore, the previous results are still of interest, as they work over arbitrary fields.

[3]Again over $\mathbb{R}$.

[4]I.e., we will be restricting ourselves to special cases, for which we can simplify the proof.

Chapter 2. Depth-3 PITs

While PIT is easily solved for depth–2 circuits (it has a blackbox polynomial time algorithm [KS01]), depth–4 is still quite far away [Sax09]. We will therefore focus on depth–3 $\Sigma\Pi\Sigma$ PIT. (Note that a $\Pi\Sigma\Pi$ circuit is zero iff any factor of the first product is zero, and therefore can be reduced to multiple depth–2 $\Sigma\Pi$ PIT–problems.)

In this chapter, we will only sketch the ideas used in the proofs of the current depth–3 PIT algorithms, mainly following the presentation given in Saxena's survey paper [Sax09].

There are various types of algorithms. In this paper, we will present four different methods[1]. The first three can be found in this chapter, along with a short overview of their proofs[2]. The fourth, the first rank bound $k^{O(k)}$ to be independent of the degree d, will follow in Chapter 4.

First, we will present a randomized algorithm. This dates from 1980, and uses the fact that it is hard to randomly hit a root of a polynomial.[3] Then, we will show an algorithm that can look "inside" the circuit, i.e. a non–blackbox algorithm, which gives a poly(n, d^k)–time algorithm deciding PIT.

One may wish to use algorithms that cannot look "inside" the circuits, i.e. blackbox algorithms[4]. There has been a lot of progress lately for $\Sigma\Pi\Sigma(k, d, n)$ identities with constant top fanin k. We are thus able to show a rank bound of $(k^3 \log d)$ for circuits over any field $\mathbb{F}$ in the third part of this chapter.

2.1 A Randomized PIT

Both Zippel [Zip79] and Schwartz [Sch80] independently found the following randomized polynomial time algorithm, which simply picks a random point in a field, and uses the output of the circuit in that point in order to decide whether the polynomial is the zero polynomial or not:

[1] Actually, there is a short overview of a fifth algorithm at the very end of this paper: Saxena and Seshadhri's newest achievement, a rank bound of $O(k^2)$.

[2] We will try to convey the main ideas used in the proofs, on subsets of $\Sigma\Pi\Sigma(k)$ in order to keep the proofs short and simple.

[3] For polynomials of "small" degree relative to the field size.

[4] The Schwarz–Zippel algorithm is a very obvious example for such a blackbox algorithm, since it only evaluates a circuit C at various random points.

Algorithm 1: Schwartz–Zippel randomized PIT

Input: (P, S): $P \in \mathbb{F}[x_1, \ldots, x_n]$ with $\deg P \leq d$, $S \subseteq \mathbb{F}$ finite

Output: With probability $\left(1 - \frac{d}{|S|}\right)$, outputs TRUE if P is zero on S

begin
  Pick $x \in_{\mathrm{R}} S^n$;
  if $P(x) = 0$ **then**
    output TRUE;
  else
    output FALSE;
  end
end

It is obvious that the algorithm cannot err on outputting FALSE. The probability for TRUE, $\left(1 - \frac{d}{|S|}\right)$, follows from the Schwartz–Zippel Lemma, for which we still introduce the following, simple definition:

Definition 2.1.1. We say that we pick x *randomly from* X, $x \in_{\mathrm{R}} X$, if we choose a purely random x out of a set X using the standard distribution, i.e. each element is picked with a probability of $\frac{1}{|X|}$ for $|X| < \infty$, or with a probability of 0 for $|X| = \infty$.

Lemma 2.1.2 (Schwartz–Zippel)**.** *Let $P \in \mathbb{F}[x_1, \ldots, x_n]$ be a non–zero polynomial of degree $d \geq 0$ over a field $\mathbb{F}$. Let $S \subseteq \mathbb{F}$ be finite.*

Then

$$\operatorname*{Prob}_{x_1,\ldots,x_n \in_R S}[P(x_1, \ldots, x_n) = 0] \leq \frac{d}{|S|}.$$

Proof. We will prove the lemma by induction.

Let $n = 1$, i.e. P is a univariate polynomial. Being non–zero, it can therefore only have at most d roots. So the probability of hitting a root is at most $\frac{d}{|S|}$.

Let the statement hold for polynomials with $(n-1)$ variables. We can view an n-variate polynomial over $\mathbb{F}$ as a univariate polynomial over $\mathbb{F}[x_1, \ldots, x_{n-1}]$ as follows:

$$P(x_1, \ldots, x_n) = \sum_{i=0}^{d} x_n^i P_i(x_1, \ldots, x_{n-1}).$$

Note that, since P is non–zero, at least one P_i must be non–zero. Pick the greatest i such that P_i is non–zero. Then, $\deg(P_i) \leq (d - i)$ holds. From the induction hypothesis, we know that

$$\operatorname*{Prob}_{x_1,\ldots,x_{n-1} \in_{\mathrm{R}} S}[P_i(x_1, \ldots, x_{n-1}) = 0] \leq \frac{d-i}{|S|}.$$

By our choice of i, P is a univariate polynomial over $\mathbb{F}[x_1, \ldots, x_{n-1}]$ of degree i. Thus, by the univariate case, it holds that

$$\operatorname*{Prob}_{x_1,\ldots,x_n \in_{\mathrm{R}} S}[P(x_1, \ldots, x_n) = 0 | P_i(x_1, \ldots, x_{n-1}) \neq 0] \leq \frac{i}{|S|}.$$

We can now see that, for $x_1, \ldots, x_n \in_{\mathrm{R}} S$,

$$\begin{aligned}\text{Prob}\left[P\left(x_1, \ldots, x_n\right) = 0\right] \leq & \text{Prob}\left[P_i\left(x_1, \ldots, x_{n-1}\right) = 0\right] + \\ & \text{Prob}\left[P\left(x_1, \ldots, x_n\right) = 0 | P_i\left(x_1, \ldots, x_{n-1}\right) \neq 0\right] \\ \leq & \frac{d-i}{|S|} + \frac{i}{|S|} \\ = & \frac{d}{|S|}.\end{aligned}$$

□

Corollary 2.1.3. *The probability of randomly picking a root of a polynomial over a field of characteristic 0 is zero.*

2.2 Non–Blackbox PIT

A *blackbox* PIT algorithm is given access to a circuit C, which it can call "for free"[5] with any (valid) input, and uses only the *output* of C in order to decide PIT. A *non–blackbox* algorithm on the other hand can "look into" the circuit C and also make decisions based upon its *structure*, in addition to the output.

In 2005, Kayal and Saxena presented a non–blackbox PIT algorithm that runs in polynomial time on bounded top fanin circuits [KS07]. Although we are mainly interested in blackbox algorithms, we will give a short overview of the ideas involved (for $k \in \{2, 3\}$).

Since we restrict our proof to $k \in \{2, 3\}$, we are able to use a generalization of the ABC Theorem [Sto81, Mas84, dB09] in order to help us find a lower bound on the degree of a polynomial, allowing us to prove the non–blackbox PIT algorithm. The ABC Theorem cannot be applied when $k > 3$, and therefore is not used in [KS07].[6]

Theorem 2.2.1 (ABC Theorem [Sto81, Mas84])**.** *Let $\mathbb{F}$ be an algebraically closed field of characteristic zero. Let $A, B, C \in \mathbb{F}[x]$ be polynomials such that*

$$\gcd(A, B, C) = 1 \tag{2.1}$$

and

$$A + B = C \tag{2.2}$$

hold. Then the number of distinct roots is greater than their largest degree.

Proof. Let $f, g \in \mathbb{F}(x)$ be two rational functions with

$$\begin{aligned}f = & \frac{A}{C}, \\ g = & \frac{B}{C}.\end{aligned}$$

[5] I.e., the cost for running the query are the same as any field operation, both in terms of size and runtime.

[6] By restricting ourselves to $k \in \{2, 3\}$, we can therefore find an alternate, simpler proof than in [KS07].

Then, dividing (2.2) by C gives us

$$f + g = 1.$$

Without loss of generality, let

$$\begin{aligned} A(x) =& c_1 \prod (x - \alpha_i)^{m_i}, \\ B(x) =& c_2 \prod (x - \beta_i)^{n_i}, \\ C(x) =& c_3 \prod (x - \gamma_i)^{r_i}. \end{aligned}$$

Then, by differentiation, we get that

$$f' + g' = \frac{f'}{f} f + \frac{g'}{g} g = 0,$$

and thus it can be determined that

$$\frac{B}{A} = -\frac{\frac{f'}{f}}{\frac{g'}{g}} = -\frac{\sum \frac{m_i}{x-\alpha_i} - \sum \frac{r_i}{x-\gamma_i}}{\sum \frac{n_i}{x-\beta_i} - \sum \frac{r_i}{x-\gamma_i}}.$$

Note that

$$D = \prod (x - \alpha_i) \prod (x - \beta_i) \prod (x - \gamma_i)$$

is a common denominator for $\frac{f'}{f}$ and $\frac{g'}{g}$. Also, the degree of D is the number of distinct roots of A, B, and C. Then, $D\frac{f'}{f}$ and $D\frac{g'}{g}$ have degree less than the number of distinct roots.

The theorem then follows from

$$\frac{B}{A} = -\frac{D\frac{f'}{f}}{D\frac{g'}{g}}$$

and (2.1). □

Theorem 2.2.2 (General ABC Theorem [dB09])**.** *For $r \geq 3$, let*

$$f_1, \ldots, f_r \in \mathbb{C}[x_1, \ldots, x_n]$$

be polynomials, not all constant, such that

$$\gcd_{i,j \in [r], i \neq j} (f_i, f_j) = 1 \tag{2.3}$$

and

$$\sum_{i \in [r]} f_i = 0. \tag{2.4}$$

Then

$$\max_{i \in [r]} \deg (f_i) \leq (r-2) \left(\deg \operatorname{rad} \prod_{i \in [r]} (f_i) - 1 \right).$$

We will refer to [dB09] for the proof of this theorem.

Theorem 2.2.3. *There is a* poly $\left(n, d^k\right)$*-time algorithm that determines whether a* $\Sigma\Pi\Sigma(k, d, n)$ *circuit computes the zero polynomial over* $\mathbb{F}$*.*

Proof. We will only prove the existence for $k \in \{2, 3\}$, for $k \geq 4$ and runtime analysis cf. [KS07, Section 3].

Let the components of $C := \Sigma\Pi\Sigma(k, d, n)$ be denoted as follows:

$$\begin{aligned} C &= \sum_{i\in[k]} T_i \\ &= \sum_{i\in[k]} \prod_{j\in[d]} l_{ij} \\ &= \sum_{i\in[k]} \prod_{j\in[d]} \sum_{\kappa\in[n]} a_{ij\kappa} x_\kappa, \end{aligned}$$

with $a_{ij\kappa} \in \mathbb{F}$. Assume all T_i to be coprime, else replace C with $\mathrm{sim}(C)$.

If $C = T_1 + T_2$, i.e. $k = 2$, we simply have to check whether

$$T_1 = -T_2.$$

Since $\mathbb{F}[x_1, \ldots, x_n]$ is a unique factorization domain, T_1 and $-T_2$ are equal iff the linear forms correspond one-to-one and the monomials match on both sides.

If $C = T_1 + T_2 + T_3$, i.e. $k = 3$: Define $L := \{l_{ij}\}_{i\in[3],j\in[d]}$. Then $C = 0$ iff

$$\forall\, l \in L : C = 0 \mod l. \tag{2.5}$$

Say, without loss of generality, l comes from T_1. Then we have $C = T_2 + T_3 \mod l$, and we can apply the previous case,[7] $k = 2$. Now, if (2.5) holds, it follows that

$$\begin{aligned} & \forall\, l \in L : C = 0 \mod l \\ \Rightarrow\ & C = 0 \mod \prod_{l\in L} l \\ \Rightarrow\ & C = 0 \mod \mathrm{rad}(T_1T_2T_3), \end{aligned}$$

with $\mathrm{rad}(T_1T_2T_3)$ being the radical of the polynomial $(T_1T_2T_3)$. Since we ensured that C is simple, and thus

$$\gcd(T_1, T_2, T_3) = 1,$$

we can apply Theorem 2.2.2, and we get that $\deg(\mathrm{rad}(T_1T_2T_3)) > d$. Thus,

$$C \equiv 0,$$

already as element of $\mathbb{F}[x_1, \ldots, x_n]$.

Since Theorem 2.2.1 does not hold for more than 3 summands, or over arbitrary $\mathbb{F}$, the previous approach has to be tweaked for these cases (cf. [KS07]). □

[7] Note that $\mathbb{F}[x_1, \ldots, x_n]/(l)$ is also a unique factorization domain.

A simplified version of this algorithm finally boils down to Algorithm 2 [Sax09]. Cf. [KS07, Section 3.2] for a better explanation of the algorithm, on how to select each $\mathfrak{I}$, and when a specific $\mathfrak{I}$ is considered useful.

Algorithm 2: Non–blackbox PIT

Input: Circuit $C = T_1 + \ldots + T_k$
Output: Output TRUE iff C is zero, FALSE otherwise
begin
  repeat

$$\mathfrak{I} \leftarrow \{(f_1, \ldots, f_l) \,|\, l \in [k-1], \forall i \in [l] :$$
$$f_i \text{ is a maximal factor of some } T_j \text{ s.t.}$$
$$f_i \text{ is not a zero-divisor modulo } (0, f_1, \ldots, f_{i-1}) \text{ and}$$
$$f_i \text{ is power of a linear polynomial modulo } \operatorname{rad}(0, f_1, \ldots, f_{i-1})\}$$

    forall $I \in \mathfrak{I}$ **do**
      if *C tests "bad" with respect to I*[a] **then**
        output FALSE;
      end
    end
  until *no useful $\mathfrak{I}$ left* ;
  output TRUE;
end

[a] E.g., but not only, $C \neq 0 \mod I$.

2.3 $k^3 logd$ rank bound over arbitrary fields

In 2008, Saxena and Seshadhri presented the rank bound for simple, minimal depth–3 circuits $\Sigma\Pi\Sigma(k, d, n)$ over an arbitrary field $\mathbb{F}$ of $(k^3 \log d)$ [SS09]. This bound is actually almost optimal, since the proof suggests the following identity C of rank $(\log_2 d + 2)$ over $\mathbb{F}_2$:

$$\begin{aligned} C(x_1, \ldots, x_r) := & \prod_{b_1,\ldots,b_{r-1} \in \mathbb{F}_2, b_1+\cdots+b_{r-1}=1} (b_1 x_1 + \ldots + b_{r-1} x_{r-1}) \\ & + \prod_{b_1,\ldots,b_{r-1} \in \mathbb{F}_2, b_1+\cdots+b_{r-1}=0} (x_r + b_1 x_1 + \ldots + b_{r-1} x_{r-1}) \\ & + \prod_{b_1,\ldots,b_{r-1} \in \mathbb{F}_2, b_1+\cdots+b_{r-1}=1} (x_r + b_1 x_1 + \ldots + b_{r-1} x_{r-1}) \\ & = 0. \end{aligned}$$

This is a $\Sigma\Pi\Sigma(3, d, r)$ identity with $k = 3$. We will also only outline the proof of

their main theorem for the case that $k = 3$. The theorem is the following:[8]

Theorem 2.3.1 (Theorem 2, [SS09]). *Let C be a simple, minimal $\Sigma\Pi\Sigma(k, d, n)$ circuit computing the zero polynomial. Then*

$$\operatorname{rank}(C) \leq O(k^3 \log d).$$

For the proof, we will need the following definitions and lemmas:

Definition 2.3.2. A *list of linear forms* is a multiset (of linear forms) with some ordering. The specific ordering is not important, as long as it allows to distinguish between repeated elements in the list.

Definition 2.3.3. Let U and V be two lists of linear forms and π a bijection between them. U and V are called *similar*, if, for all $u \in U$, there is a $c \in \mathbb{F}^\times$ such that $u = c\pi(u)$.

Definition 2.3.4 (Matching. Definition 12, [SS09]). Let U, V be two lists of linear forms, and I be a form. We call a bijection π an *I–matching between U and V*, if for all $l \in U$, there is a $c \in \mathbb{F}^\times$ and $v \in \operatorname{span}(I)$ such that $\pi(l) = (cl + v)$.

For I, we will be picking different linear forms l of T_1 and count the l-matchings between T_2 and T_3:

Lemma 2.3.5 (Lemma 14, [SS09]). *Let U, V be two lists of linear forms, each of size d, i.e. $|U| = |V| = d$. Let $I_1, \ldots, I_r$ be linearly independent linear forms such that for all $i \in [r]$, there is an I_i-matching π_i between U and V.*

If $r > (\log_2 d + 2)$, then U and V are similar lists.

Proof. We will prove this by contradiction, using the following combinatorial process: View the setup as a bipartite graph $G = (U, V, E)$, with the edges E referencing the matchings π_i. I.e., for every π_i and $u \in U$, there is an edge in E labeled $\pi_i(u)$ that connects u with its *I_i-neighbor* $\pi(u)$.

We will now describe an iterative process over the linear forms I_i in order to build a *basis* B that will span all $(U \cup V)$. In each round, the set J_i will be composed of those linear forms that we have considered until then, i.e.

$$J_i = \bigcup_{j \in [i]} I_j.$$

Initialize B with one element $u_0 \in U$. In the i^{th} round, we simply add I_i to B. Notice that we thereby cover forms of $(U \cup V)$ with $\operatorname{span}(B)$, e.g., in the first round, we covered u_0 and its I_1-neighbor $\pi_1(u)$, since $\pi_1(u)$ is a linear combination of u and I_1. Thus, at any time, if a form l is covered, its neighbor $\pi_i(i)$ is also covered. So there are always equally many forms of U and V covered.

[8] For a full proof, cf. [SS09]

We now pick the smallest subset $S \subseteq [r]$, such that $(\{u_0\} \cup \bigcup_{j \in S} I_j)$ is a set of linear *dependent* forms. Without loss of generality, and with $i_0 = |S|$, let $S = [i_0]$. Then, by construction, $u_0, I_1, \ldots, I_{i_0 - 1}$ are independent, and thus $u_0 \in \operatorname{span}(J_{i_0})$. Then, since all $I_1, \ldots, I_r$ are independent, $\operatorname{span}(J_{i_0})$ must be independent of $I_{i_0+1}, \ldots, I_r$.

Let us examine the i^{th} step[9] more closely. Take a u' that has been covered in the last round (i.e. in the $(i-1)^{\text{th}}$ round). Notice that its I_i-neighbor will not have been covered until this (the i^{th}) round, due to the following: Say $u' \in U$, then $v' = \pi_i(u') = (cu' + v) \in V$, with $c \in \mathbb{F}^\times$, and $v \in \operatorname{span}(I_i), v \neq 0$. If v' had already been covered, then $v \in \operatorname{span}(B)$. But since I_i is independent of B, that would mean that $v \in (\operatorname{span}(I_i) \cap \operatorname{span}(B)) = \{0\}$, which is a contradiction.

Hence, this i^{th} round doubles the amount of covered forms.

Since $|U| = |V| = d$, this doubling can happen at most $\log_2 d$ times, and thereby

$$r - 2 \leq \log_2 d.$$

This is a contradiction, therefore U and V must be similar. □

Proof of Theorem 2.3.1. For circuits of fanin 3, the theorem follows directly from Lemma 2.3.5: Take the linear forms of T_1 and T_2 to be the linear forms of U and V respectively. The matching forms I_i will be those of T_3. By the lemma, there can be at most $(\log_2 d + 2)$ linear independent forms in T_3. Repeat this by taking the matching forms to be of T_2 and T_1 to get a rank bound

$$\operatorname{rank}(C) = O(\log_2 d).$$

For higher fanins, cf. [SS09]. □

2.4 k^k rank bound over $\mathbb{R}$

The above (almost optimal) rank bound of $(k^3 \log d)$ is valid for any field $\mathbb{F}$. For fields of characteristic 0, better results may be achieved, as shown by Kayal and Saraf in 2009 [KS09]. They proved that over $\mathbb{R}$ (or rather any extension of $\mathbb{Q}$ that is embedded in $\mathbb{R}$), identities have rank bound k^k. Notice that this bound is independent of the degree d, and thus a lot closer to the bound $O(k)$ conjectured by Dvir and Shpilka than the one for arbitrary fields.

The proofs are based heavily on incidence geometry, especially the *Sylvester–Gallai Theorem*, and its higher dimensional versions. We will focus on these geometric theorems in Chapter 3, and on the proof of Kayal and Saraf's rank bound in Chapter 4.

[9] $i \notin \{1, r\}$

Chapter 3. Sylvester–Gallai Theorems

In this chapter, we will present those results from incidence geometry that proved to be very useful in handling identities. They are all based on the *Sylvester–Gallai Theorem*.

In 1893, Sylvester asked whether it is possible to find a set of non–collinear points in $\mathbb{R}^2$ such that, for each line passing through two points of the set, the line contains a third point of the set. In 1933, Gallai showed that such a set does not exist: The *Sylvester–Gallai Theorem* asserts that in every set of points spanning a 2–dimensional linear space over $\mathbb{R}$, there are two points such that the line passing through these points does not meet any third point of the initial set.

Later, in 1965, Hansen extended this to higher dimension [Han65] by proving that any set of points spanning a space isomorphic to $\mathbb{R}^n$ contains a subset of points spanning a hyperplane of $\mathbb{R}^n$ in such a way that all but one of the points in the subset only span an $(n-2)$–dimensional subspace, and only by adding the last point does the dimension increases to a hyperplane.

Serre proposed a similar problem in 1966: whether such a configuration as given by Sylvester must be coplanar over the complex projective space. This was proven by Kelly in 1986 [Kel86].

In the meantime, several proofs for both problems were given. A very elementary proof of the Sylvester–Gallai Theorem is presented in Borwein and Moser's survey paper [BM90]. Unfortunately, this proof does not extend easily to the complex case. Therefore, we will be following Elkies, Pretorius and Swanepoel's proof [EPS06], which sets up the case over $\mathbb{R}$ in such a way that one can prove the complex case with the same arguments, which facilitates the proof of the latter dramatically.

Almost as a corollary to Hansen's Theorem, Bonnice and Edelstein [BE67] proved that for a finite set of points spanning (a space isomorphic to) $\mathbb{R}^{2n}$, there is a subset of exactly $(n+1)$ points that span an n–dimensional subspace.

This fact will be heavily relied on by Kayal and Saraf's *Hyperplane Decomposition Lemma* [KS09], which states that, given a finite set of points spanning $\mathbb{R}^n$, there is a "core" subspace of lower dimension, say $n_{\text{core}} < n$, such that this core subspace can be extended at least $\frac{n-n_{\text{core}}}{2}$ times into disjoint subspaces, each of dimension exactly $(n_{\text{core}}+1)$.

Remark 3.0.1. When discussing geometric properties, we work in the projective space, i.e. we restrict ourselves to the following:

- For points in $\mathbb{R}^n$, *we do not allow a point to be in the origin*

$$O = (0, \ldots, 0)^T.$$

- For any two *distinct* points in $a, b \in \mathbb{R}^n$, *we do not allow the two points to be on the same line passing through the origin. I.e., there is no* $c \in \mathbb{R}$*, such that* $a = cb$.

We can construct this situation in the following way: For each point $P_i \in \mathbb{R}^n$, we consider the line

$$L(P_i) = \{\lambda P_i \mid \lambda \in \mathbb{R}\}.$$

Then, we select a hyperplane $V \subset \mathbb{R}^n$ as follows:

1. $V \cap O = \emptyset$,
2. $V \cap L(P_i) \neq \emptyset$ for all points P_i.

We now can map a point P_i into V by the linear projection $\pi : \mathbb{R}^n \to V$, with

$$\pi(P_i) = L(P_i) \cap V.$$

Remark 3.0.2. In the rest of this chapter, we will often loosely[1] refer to d–dimensional linear spaces as $\mathbb{R}^d$.

Definition 3.0.3. A *Sylvester–Gallai configuration* (SG configuration) is a finite set of non–collinear points, such that, for each line passing through two points, there is a third point of the set on this line.

Definition 3.0.4. For $n, d \in \mathbb{N}$, $n < d$, and a d–dimensional linear space V, define the following:

- An *n–flat*, is an affine n–dimensional subspace of V.
- A *hyperplane* is a $(d-1)$–flat in V.
- Let S be a finite set of points that span V. A hyperplane H of V is called *ordinary* (with respect to S), if H is spanned by $(S \cap H)$ and there exists a $P \in H$, such that $((S \setminus \{P\}) \cap H)$ spans a $(d-2)$–dimensional flat of V. P is called the *leader*, $((S \setminus \{P\}) \cap H)$ the *follower*.
- A n–flat F of V is called an *elementary* n–flat, if $|F \cap S| = (n+1)$, and F is spanned by exactly these $(n+1)$ points..

From the definition, the following is immediately obvious: Let F be an elementary flat of V, spanned by a set of points S. Then F is an ordinary flat, with every point in S being a leader.

Definition 3.0.5. We can associate a d–dimensional linear space V with its *dual* V^*, a *duality map* taking k–flats in V to $(d-k-1)$–flats in V^*.

[1] By isomorphism.

For the proof of the Sylvester–Gallai Theorem, the following will be of importance:

Remark 3.0.6. Notice that the dual of a point is a hyperplane and vice–versa.

For $\mathbb{R}^2$ and its dual, with a duality map π, say we have points $P_1, P_2, \ldots, P_n$ that are all on the line L. Then notice the following:

$$\bigcap_{i\in[n]} \pi(P_i) = \pi(L). \tag{3.1}$$

In words: The dual of the points P_i (which will be lines) all intersect in one point: the dual of the line L (which will be a point).[2]

3.1 The Sylvester–Gallai Theorem

The *Sylvester–Gallai Theorem* asserts that one cannot arrange a set of points in $\mathbb{R}^2$ in such a way that no line in $\mathbb{R}^2$ meets exactly two points of this set. There are various different proofs known for this. We have chosen this specific proof since it translates rather easily into a similar statement over $\mathbb{C}$. A very simple proof of the Sylvester–Gallai Theorem, using only basic geometric properties in $\mathbb{R}^2$, can be found in Borwein and Moser's survey paper [BM90].

Theorem 3.1.1 (Sylvester–Gallai)**.** *There is no Sylvester–Gallai configuration in* $\mathbb{R}^2$*.*

In order to prove this, we will need to prove the following theorem first.

Theorem 3.1.2. *Say we have the following triangles:*

In a plane, three non–collinear points P_1, P_2, and P_3 are given. The lines joining all pairs of points form a triangle, and for all $i \in [3]$, we label the line opposite to P_i with L_i. Now, for all all distinct $i, j, k \in [3]$, let L_{jk} be a (third) line passing through P_i such that $L_{jk} \neq L_j$ and $L_{jk} \neq L_k$.[3] (We pick L_{jk} and L_{kj} as the same line.)

Then the following holds:

- *Either, for all distinct $i, j, k \in [3]$,, L_i is parallel to L_{jk}.*
- *Or, the area of $\Delta(L_1, L_2, L_3)$ is larger than the area of at least one of the following:*

$$\begin{array}{ccc} \Delta(L_{12}, L_2, L_3) & \Delta(L_1, L_{12}, L_3) & \Delta(L_1, L_2, L_{13}) \\ \Delta(L_{13}, L_2, L_3) & \Delta(L_1, L_{23}, L_3) & \Delta(L_1, L_2, L_{23}) \\ \Delta(L_1, L_{12}, L_{13}) & \Delta(L_{12}, L_2, L_{23}) & \Delta(L_{13}, L_{23}, L_3) \end{array}$$

[2] This can be naturally extended to higher dimensional linear spaces and higher dimensional flats. E.g., in $\mathbb{R}^d$, if we have a set of points all contained in a hyperplane, the dual of the points (now hyperplanes) will all intersect in a point (the dual of the hyperplane).

[3] See Figure 3.1.

This is a reformulation of a proposition of Rainwater, which was proven by Diananda and Bager [RDB61]. Following the lead of Elkies, Pretorius, and Swanepoel [EPS06], we will give a slightly more complicated version of the proof, since this will carry over very well to the complex case. Note that in order to prove both cases of Theorem 3.1.2, it is enough to show that the area of $L_1L_2L_3$ cannot be less than that of all the other triangles, and that, if equality holds, the appropriate lines must be parallel.

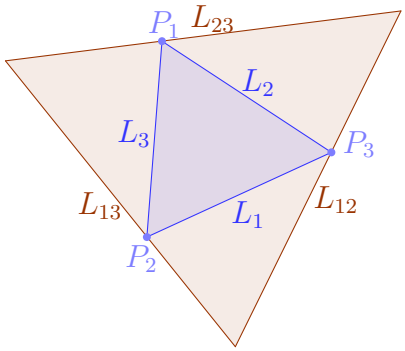

Figure 3.1: Example of triangles in Theorem 3.1.2

Proof. Without loss of generality, assume that we are working in $\mathbb{P}^2$, on the affine plane identified with $\sum_{i=1}^{3} x_i = 1$ in $\mathbb{R}^3$,[4] such that the lines L_i have the equation $x_i = 0$. (Note that the multiplicity of determinants ensures that a minimal volume triangle remains so.) Then, the points P_i are exactly $(\delta_{i,1}, \delta_{i,2}, \delta_{i,3})^T$.[5]

For all distinct $i, j, k \in [3]$, let L_{jk} and L_{kj} be such that they are the same line[6], and notice that L_{jk} passes through P_i. Thus, the equation for L_{jk} turns out to be

$$\alpha x_j + \beta x_k = 0. \tag{3.2}$$

This can be transformed[7] into $\frac{\alpha}{\beta} x_j + x_k = 0$. Set $\alpha_{jk} := \frac{\alpha}{\beta}$ and, respectively, $\alpha_{kj} := \frac{\beta}{\alpha}$. Thus, for any $j \neq k$, we have

$$\alpha_{jk}\alpha_{kj} = 1. \tag{3.3}$$

For any distinct i, j, k, note that $\alpha_{jk} = 1$ iff L_i is parallel to L_{jk}.[8]

For $A, B, C \in \mathbb{R}^3$, the area of a triangle $\Delta(A, B, C)$ is $\frac{\sqrt{3}}{2} |\det[ABC]|$. We can therefore easily compute the area of the nine triangles given above:

[4] I.e., while the coordinates given are in $\mathbb{R}^3$, everything is *actually* happening in (a linear space isomorphic to) $\mathbb{R}^2$. Therefore, all coordinates and equations are understood to be $\cap\{x_1+x_2+x_3 = 1\}$.

[5] $\delta_{i,j}$ is 1 iff $i = j$, 0 else.

[6] This is allowed by the (loose) restrictions to L_{jk} in the theorem.

[7] Note that $\alpha, \beta \neq 0$, since otherwise the line would be parallel to L_j, L_k, respectively.

[8] By equation (3.2)

The triangle $\Delta(L_1, L_{12}, L_3)$ has the following vertices:

$$\begin{aligned} P_1' &= \left((1-\alpha_{12})^{-1}, -\alpha_{12}(1-\alpha_{12})^{-1}, 0\right)^T, \\ P_2' &= P_2, \\ P_3' &= P_3. \end{aligned}$$

Thus, $|\det[P_1', P_2', P_3']| = \left|(1-\alpha_{12})^{-1}\right|$.

Let us assume that the area of $L_1L_2L_3$ is smaller or equal to the area of $L_1L_{12}L_3$. Then, via $|(1-\alpha_{12})^{-1}| \geq 1$, the following holds:

$$|1-\alpha_{12}| \leq 1.$$

We can repeat the same argument with, e.g., the triangle $L_1L_{12}L_{13}$ in order to obtain

$$|1-\alpha_{12}-\alpha_{13}| \leq 1.$$

Thus, we obtain for each of the nine triangles an inequality:[9]

$$|1-\alpha_{ij}| \leq 1, \tag{3.4}$$
$$|1-\alpha_{ij}-\alpha_{ik}| \leq 1. \tag{3.5}$$

Because of equation (3.4), we know that $\alpha_{ij} \geq 0$. By (3.3) and the AGM inequality[10], it follows that $\alpha_{ij} + \alpha_{ji} \geq 2$, with equality iff $\alpha_{ij} = \alpha_{ji} = 1$. Thus, the following holds:

$$\sum_{i,j\in[3],i\neq j} \alpha_{ij} \geq 6.$$

Using (3.5), we know that $\alpha_{ij} + \alpha_{ji} \leq 2$, and thus

$$\sum_{i,j\in[3],i\neq j} \alpha_{ij} \leq 6.$$

Therefore, if the area of $L_1L_2L_3$ is not strictly greater than the area of one of the given triangles, all α_{ij} must be equal to 1, and thereby L_i must be parallel to L_{jk} for all distinct $i, j, k \in [3]$. □

Now, we have all tools necessary to prove the Sylvester–Gallai Theorem:

Proof of Theorem 3.1.1. Say, for contradiction, that we have a SG configuration in $\mathbb{R}^2$. Via a duality map π, we obtain a set of lines L_i with intersection points P_j. Note that by (3.1) at least three lines pass through each P_j.

[9] (for all distinct $i, j, k \in [3]$)

[10] The AGM inequality states that for $x_1, \ldots, x_n \in \mathbb{R}^{>0}$, the geometric mean is less or equal to the arithmetic mean, i.e.

$$\sqrt[n]{x_1 \cdots x_n} \leq \frac{x_1 + \ldots + x_n}{n},$$

with equality iff $x_1 = \ldots = x_n$. For a proof, cf. e.g. [AZ01, Chapter 16].

Remember that, originally, we neither had points in the origin, nor two points that were multiples of each other. It therefore is easy to see that, when we pass over into the dual, we do not have any two lines that are parallel to each other.

Since all lines are non–parallel, they form triangles. Choose a triangle with minimum area. Without loss of generality, let the lines forming this triangle be called L_1, L_2, and L_3 and the vertices be called P_1, P_2, and P_3. Since there is a third line passing through each vertex we can consider us to be in the situation of Theorem 3.1.2.[11] Now, since our lines are all non–parallel, we get a contradiction to our choice of triangle, since we find a triangle of even smaller area. □

Corollary 3.1.3. *For all $n \geq 2$, there is no SG configuration in $\mathbb{R}^n$.*

Proof. Suppose a SG configuration S exists. Pick 3 non–collinear co–planar points, say A, B, and C, in S. Since S is a SG configuration, so is $(S \cap \text{span}(A, B, C))$. But that would mean that we have a SG configuration in a 2–flat of $\mathbb{R}^n$ (which is isomorphic to $\mathbb{R}^2$). This contradicts Theorem 3.1.1. □

3.2 Sylvester–Gallai over $\mathbb{C}$

We now can apply the same arguments as before over $\mathbb{C}$. Unfortunately, we get a slightly weaker result: Instead of a SG configuration being *collinear*, as over $\mathbb{R}$, over $\mathbb{C}$ it will only be *coplanar*. An example for a SG configuration in a plane are the nine points of inflexion of the following cubic curve [Cox48]:

$$x^3 + y^3 + z^3 = xyz$$

Definition 3.2.1. For $z \in \mathbb{C}$, let $\text{Re}(z)$ be the real part of z, and $\text{Im}(z)$ the imaginary part. I.e., for $z = a + bi$, $a, b \in \mathbb{R}$, let

$$\text{Re}(z) = a,$$
$$\text{Im}(z) = b.$$

Theorem 3.2.2. *Every Sylvester–Gallai configuration in $\mathbb{C}^n$ is coplanar.*

By the same arguments as over $\mathbb{R}$, it is enough to prove this fact for $n = 3$.

Proof. Suppose we have a SG configuration S in $\mathbb{C}^3$. Again, we take its dual and call this S^*.[12]

Without loss of generality, we work on the affine plane in $\mathbb{C}^4$ given by

$$\sum_{i \in [4]} x_i = 1.$$

[11] Strictly speaking, we would need to map $\mathbb{R}^2$ to $\{x_1 + x_2 + x_3 = 1\} \cap \mathbb{R}^3$.

[12] Note that due to our setting, there are no parallel lines.

Analogous to a triangle being a three-tuple of lines in $\mathbb{R}^2$, we call a four-tuple of distinct planes Π_1, Π_2, Π_3, and Π_4 a *tetrahedron*. Its vertices are $P_j = \bigcap_{k \neq j} \Pi_k$, for $j \in [4]$. We can then define the "volume" to be $|\det[P_1, P_2, P_3, P_4]|$.[13]

We select a tetrahedron $\Pi_1\Pi_2\Pi_3\Pi_4$ of minimal volume and adjust the base coordinates such that its vertices P_j are the standard unit vectors

$$e_j = (\delta_{j,1}, \delta_{j,2}, \delta_{j,3}, \delta_{j,4})^\tau.$$

Then Π_j has the equation $x_j = 0$.

For $j, k \in [4], j < k$, choose Π_{jk} to be a plane of S^* such that $(\Pi_j \cap \Pi_k) \subset \Pi_{jk}$. Fix $\Pi_{kj} := \Pi_{jk}$. Then Π_{jk} is specified by

$$\alpha x_j + \beta x_k = 0$$

for $\alpha, \beta \neq 0$. Thus, for each $\alpha_{jk} = \frac{\alpha}{\beta}$, we get the equations

$$\alpha_{jk}\alpha_{kj} = 1 \tag{3.6}$$

for all $j, k \in [4], j \neq k$.

Compare the volume of $\Pi_1\Pi_2\Pi_3\Pi_4$ with that of, e.g., $\Pi_1\Pi_{12}\Pi_3\Pi_4$ with vertices P'_1, P'_2, P'_3, and P'_4 being:

$$\begin{aligned} P'_1 &= \left((1-\alpha_{12})^{-1}, -\alpha_{12}(1-\alpha_{12})^{-1}, 0, 0\right)^\tau, \\ P'_2 &= P_2, \\ P'_3 &= P_3, \\ P'_4 &= P_4. \end{aligned}$$

Then we obtain

$$|1 - \alpha_{12}| \leq 1.$$

Considering $\Pi_1\Pi_{12}\Pi_{13}\Pi_4$, we obtain

$$|1 - \alpha_{12} - \alpha_{13}| \leq 1.$$

Considering $\Pi_1\Pi_{12}\Pi_{13}\Pi_{14}$, we obtain

$$|1 - \alpha_{12} - \alpha_{13} - \alpha_{14}| \leq 1.$$

Thus we get 28 equations in 12 variables, for all distinct $j, k, k_1, k_2, k_3 \in [4]$:

$$|1 - \alpha_{jk}| \leq 1, \tag{3.7}$$

$$|1 - \alpha_{jk_1} - \alpha_{jk_2}| \leq 1, \tag{3.8}$$

$$|1 - \alpha_{jk_1} - \alpha_{jk_2} - \alpha_{jk_3}| \leq 1. \tag{3.9}$$

As over $\mathbb{R}$, in the following we will present two lemmas that allow us to fix

$$\alpha_{13} = \alpha_{24} = 1.$$

This then contradicts our choice of Π_{13} and Π_{24} since they must be parallel, finishing the proof of the theorem. □

[13]Notice that the constant factor of $\frac{\sqrt{3}}{2}$ for the area of triangles was of no importance in the proofs over reals. The same will happen here.

Let ρ be the cube root of unity, $\rho = e^{\frac{2\pi i}{3}}$. Note that $|\rho| = 1$.

Lemma 3.2.3. *For all distinct $j, k \in [4]$, $\alpha_{jk} \in \mathbb{C}$ satisfy (3.6)–(3.9) iff*

$$\begin{aligned}\alpha_{13} = \alpha_{31} = \alpha_{24} = \alpha_{42} &= 1,\\ \alpha_{12} = \alpha_{23} = \alpha_{34} = \alpha_{31} &= -\rho,\\ \alpha_{14} = \alpha_{43} = \alpha_{32} = \alpha_{21} &= -\rho^2.\end{aligned}$$

Proof. From (3.6) we get

$$|\alpha_{jk}| + |\alpha_{kj}| \geq 2,$$

with equality iff $|\alpha_{jk}| = |\alpha_{kj}| = 1$.

Also, the following holds:

$$\sum_{j,k\in[4],j\neq k} |\alpha_{jk}| \geq 12.$$

As over $\mathbb{R}$, we show that the upper bound of this sum is also 12 in the next lemma. This finishes the proof of this lemma. □

Lemma 3.2.4. *Let $\beta_1, \beta_2, \beta_3 \in \mathbb{C}$ be such that for all $S \subseteq [3]$,*

$$\left|1 - \sum_{n\in S} \beta_n\right| \leq 1.$$

Then $\sum_{n\in[3]} |\beta_n| \leq 3$, with equality iff

$$\{\beta_1, \beta_2, \beta_3\} = \{1, -\rho, -\rho^2\}.$$

Proof. Note that since $|1 - \beta_n| \leq 1$, it must be that $\mathrm{Re}\,(\beta_n) \geq 0$ for all $n \in [3]$, with equality iff $\beta_n = 0$.

Without loss of generality, order the β_n such that β_2 lies in between β_1 and β_3. Then the following is a set of vertices of a hexagon lying within $\{z \in \mathbb{C} : |1 - z| \leq 1\}$:

$$\{0, \beta_1, \beta_1 + \beta_2, \beta_1 + \beta_2 + \beta_3, \beta_2 + \beta_3, \beta_3\}.$$

Its circumference is equal to $\left(2\sum_{n\in[3]} |\beta_n|\right)$. But since it is contained in a circle of radius 1, it can be at most 6, with equality iff it is a regular hexagon with its vertices touching the border of the circle.[14]

The lemma follows from this. □

[14] See Figure 3.2.

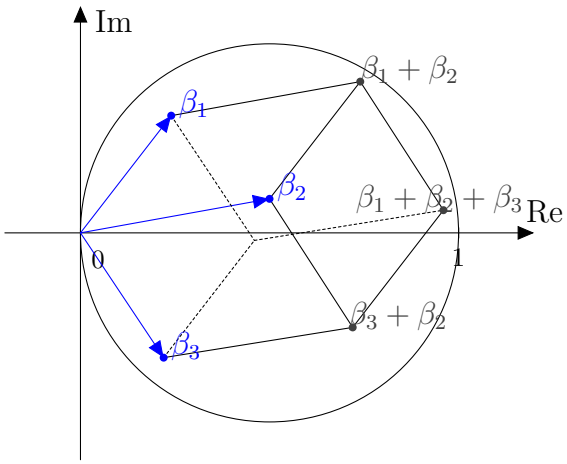

Figure 3.2: Proof of Lemma 3.2.4

3.3 Extending the Sylvester–Gallai Theorem to higher dimensional subspaces

The Sylvester–Gallai Theorem shows that for all sets of non–collinear points (in $\mathbb{R}^n$), there is a line (i.e., a 1–flat) containing exactly 2 of the points. Can a similar statement be given for higher dimensional subspaces? We will show that for a non-elementary configuration of points in $\mathbb{R}^d$, there is a subset spanning an ordinary hyperplane of $\mathbb{R}^d$.

Let us begin with a few definitions. Note that in this chapter, we will often index the geometric objects by their dimension. E.g., a hyperplane of $\mathbb{R}^n$ could be called H_{n-1}.

We will concentrate on *finite sets of points spanning* $\mathbb{R}^d$ and their properties:

Definition 3.3.1. We call $V \subseteq \mathbb{R}^d$ an *affine space* if it is a translation of a linear space, i.e. there exists a $V' \subseteq \mathbb{R}^d$ and a vector $u \in \mathbb{R}^d$, such that

$$V = u + V' = \{u + v \mid v \in V'\}.$$

We set $dim(V) = dim(V')$.

An affine space is constructed as follows: Let $S \subset \mathbb{R}^d$ be a set of points. Let n be the maximum number of linear independent points found in S. Without loss of generality, call these n points $s_0, \ldots, s_{n-1}$. Then we call the affine space spanned by these points affine-span(S), which is the following:

$$\text{affine-span}(S) = s_0 + \text{span}\left(\{(s_1 - s_0), \ldots, (s_{n-1} - s_0)\}\right).$$

It follows from the construction that the dimension of the affine span of n linearly independent points is exactly $(n-1)$:

$$\dim(\text{affine-span}(S)) = n - 1.$$

On the other hand, an affine space S' of dimension n' must be affinely spanned by at least $(n'+1)$ points:

$$|S'| \geq n' + 1.$$

Definition 3.3.2. Let S be a finite set of points spanning $\mathbb{R}^d$. For $i \in [d]$, denote the set of i–flats in $\mathbb{R}^d$ spanned by subsets of points in S by Γ_i, i.e.

$$\Gamma_i := \{\text{affine-span}(S') \mid S' \subseteq S, \dim(\text{affine-span}(S')) = i\}.$$

Note that $|S'| \geq (i+1)$.

We will call

$$\Gamma := \bigcup_{i \in [d]} \Gamma_i$$

the *configuration* Γ (with respect to S).

Remark 3.3.3. Note that $\Gamma_0 = S$ and $\Gamma_d = \mathbb{R}^d$.

Also, since S spans $\mathbb{R}^d$, no Γ_i is empty.

Definition 3.3.4. For all $i \in [d]$, $F_i \in \Gamma_i$ will be split into polyhedral domains by the subspaces $F_{i-1} \in \Gamma_{i-1}$. The closures of these domains will be called the *i–dimensional cells* of Γ.[15]

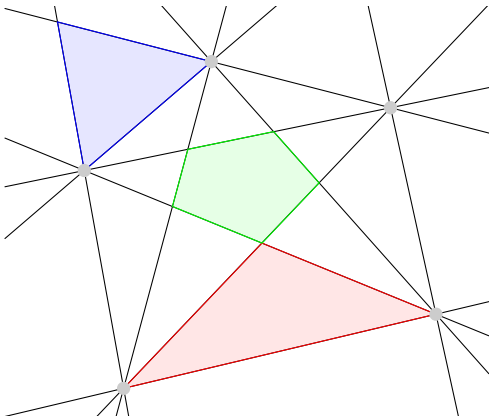

Figure 3.3: Example of 2–dimensional cells

Remark 3.3.5. Note that every cell is convex.[16]

Definition 3.3.6. For a set of points $S \subset \mathbb{R}^d$, the *convex hull H of S* is the intersection of all convex sets that contain S, i.e.

$$H = \bigcap_{C \subset \mathbb{R}^d \text{ convex},\ S \subseteq C} C.$$

[15] See Figure 3.3.

[16] I.e., for two points a and b of a cell, the connecting line $\{(ta + (1-t)b) \mid t \in [0,1]\}$ is fully contained in the cell.

Definition 3.3.7. A *d–dimensional simplex* is the convex hull of a set of exactly $(d+1)$ vertices that span a d–dimensional space.

Remark 3.3.8. A simplex is the natural extension of a triangle into higher dimensions.

Definition 3.3.9. Let A_{d-1}, B_{d-1} be two non–parallel hyperplanes of $\mathbb{R}^d$ intersecting in L_{d-2}. Then L_{d-2} splits both hyperplanes in two disjoint parts, i.e.

$$A_{d-1} \setminus L_{d-2} = A_1 \,\overline{\cup}\, A_2,$$
$$B_{d-1} \setminus L_{d-2} = B_1 \,\overline{\cup}\, B_2.$$

These parts form exactly two distinct *wedges* $W_{1,d}, W_{2,d}$ of dimension d:

$$W_{1,d} = \text{convex} - \text{hull}(A_1, B_1) \cup \text{convex} - \text{hull}(A_2, B_2)$$
$$W_{2,d} = \text{convex} - \text{hull}(A_1, B_2) \cup \text{convex} - \text{hull}(A_2, B_1)$$

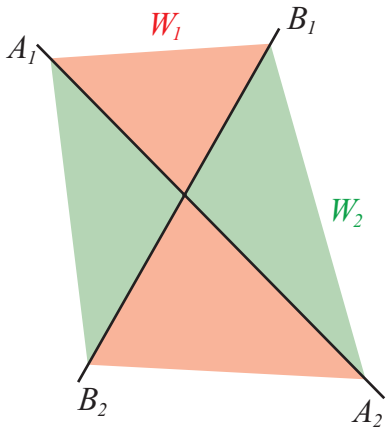

Figure 3.4: Example of the two wedges W_1 and W_2 between two lines in $\mathbb{R}^2$

Lemma 3.3.10. *Let S be a finite set of points spanning $\mathbb{R}^d$. Let Σ_d be a closed d–dimensional simplex, such that its vertices are points from S. Let P_0 be a point in S such that it is not in the simplex, i.e.*

$$\Sigma_d \cap P_0 = \emptyset.$$

Then there is a $(d-2)$–dimensional face F_{d-2} of Σ_d and a hyperplane

$$H_{d-1} = \text{affine-span}(P_0, F_{d-2})$$

of $\mathbb{R}^d$, such that the following holds:

$$\Sigma_d \cap H_{d-1} = F_{d-2}.$$

Proof. Consider the set of $(d-1)$–dimensional faces of Σ_d. Each of these faces is contained in a hyperplane of $\mathbb{R}^d$. Let this set of hyperplanes be called $\mathcal{H}$. Note that none of these hyperplanes are parallel to another, since Σ_d is a simplex. Thus, any two hyperplanes in $\mathcal{H}$ form wedges in $\mathbb{R}^d$. Since Σ_d is the intersection of all possible wedges, and since $P_0 \notin \Sigma_d$, there must be a wedge W such that P_0 is not in this wedge.

Notice that the hyperplane spanned by P_0 and the intersection of the bounding hyperplanes of W satisfy the requirements of the lemma. □

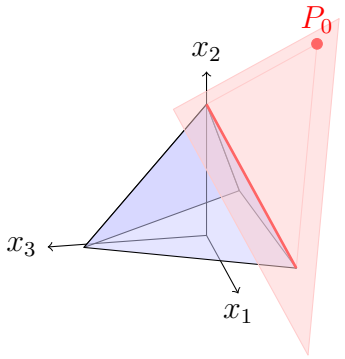

Figure 3.5: Example of Lemma 3.3.10 in $\mathbb{R}^3$

Definition 3.3.11. For two points $A, B \in \mathbb{R}^n$, the *line* $\overline{AB}$ is the set

$$\overline{AB} := \{A + t(B - A) \mid t \in \mathbb{R}\} = \text{affine-span}\{A, B\}.$$

The points in the set with $0 \leq t \leq 1$ form the *interior section* of the line, the others the *exterior section*.

Lemma 3.3.12. *Let S be a finite set of points spanning $\mathbb{R}^d$ with configuration Γ. Let H_{d-1} be a hyperplane spanned by a subset of S. Let A_0 be a point of S such that it is not contained in H_{d-1}, i.e.*

$$H_{d-1} \cap \{A_0\} = \emptyset.$$

Let $\sigma_{d-1} \subset H_{d-1}$ be a $(d-1)$–dimensional cell of Γ.

Let P_0 be a point in H_{d-1} such that the line $\overline{A_0P_0}$ does not intersect σ_{d-1}, and Q_0 be a point contained in σ_{d-1}.

Then both sections of the line $\overline{P_0Q_0}$ intersect at least one of the hyperplanes spanned by A_0 and a $(d-2)$–dimensional face of σ_{d-1}.

Proof. Consider the set

$$\left\{\overline{A_0X_0} \mid X_0 \in \sigma_{d-1}\right\}.$$

This is a $(d-1)$–dimensional convex cone. Since P_0 is outside and Q_0 is inside the cone, both sections of the line $\overline{P_0Q_0}$ obviously must intersect with the boundary of the cone.

Since every point on the boundary of this cone is contained in a hyperplane spanned by A_0 and a $(d-2)$–dimensional face of σ_{d-1}, the lemma follows. □

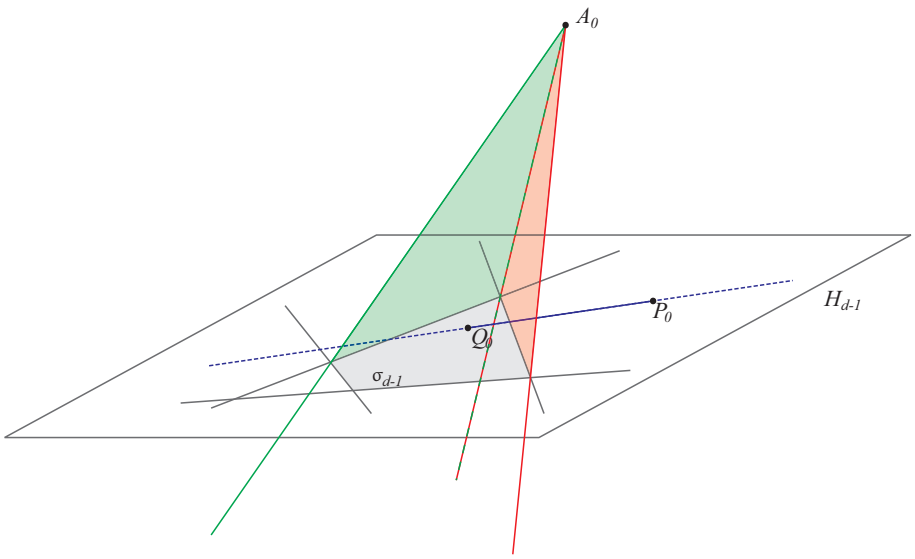

Figure 3.6: Example of Lemma 3.3.12 in $\mathbb{R}^3$

Theorem 3.3.13 (Hansen [Han65])**.** *Let $d \geq 3$, and let Γ be a non–elementary configuration in $\mathbb{R}^d$. Then there exists an ordinary hyperplane of $\mathbb{R}^d$ spanned by a subset of Γ_0.*

Proof. We prove this via induction. For $d = 2$ it is the basic Sylvester–Gallai Theorem 3.1.1. Assume that the theorem holds for $(d-1)$.

Let δ_d be a d–dimensional cell of Γ. Let σ_d be a d–dimensional simplex containing δ_d, such that it contains no point of Γ_0 other than its vertices:

$$\sigma_d \cap \left(\Gamma_0 \setminus \left\{V_i^0 \mid V_i^0 \text{ is a vertex of } \sigma_d\right\}\right) = \emptyset.$$

Such a σ_d exists, because of the following: Take any simplex[17] covering δ_d. If it contains a point of Γ_0 other than those of δ_d, then note that the simplex is cut into smaller simplexes by the hyperplanes spanned by this point and the vertices of the original simplex. One of these new simplexes contains δ_d.[18] If this new simplex still contains unwanted points of Γ_0, repeat the procedure.[19]

Note that if Γ were elementary, $\mathbb{R}^d$ would be spanned by $(d+1)$ points, which would be exactly the vertices of δ_d and σ_d (which would be equal). But since it is non–elementary, there must be a point of Γ_0 outside of σ_d. Thus, by Lemma 3.3.10, there is a hyperplane $B_{d-1} \in \Gamma_{d-1}$, spanned by this point and a $(d-2)$–dimensional face of δ_d, say S_{d-2}, such that

$$B_{d-1} \cap \delta_d \subseteq B_{d-1} \cap \sigma_d = S_{d-2}.$$

[17] whose edges coincide with the hyperplanes

[18] δ_d must be in exactly one of these simplexes by the definition of a *cell*.

[19] See Figure 3.7.

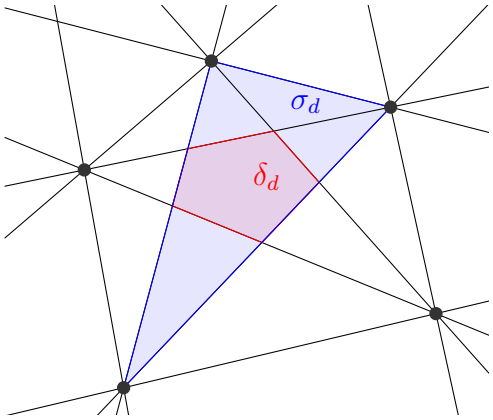

Figure 3.7: Example of cell δ_d and a covering simplex σ_d (for $d = 2$)

If B_{d-1} is elementary, then this is the hyperplane mentioned in the theorem and the proof is done. Thus, assume that it is not elementary. Later in the proof, we will need a non–elementary hyperplane not intersecting the interior of δ_d. B_{d-1} is such a hyperplane. Figure 3.8 shows such a possible B_{d-1}.

Select a point $P_0 \in \Gamma_0$ not in B_{d-1}, and an interior point of δ_d (not necessarily in Γ_0), such that the line L_1 between these two points[20] does not intersect any element of Γ_{d-2} apart from in the point P_0 itself, i.e.

$$L_1 \cap \left(\bigcup \Gamma_{d-2} \right) = \{P_0\} .$$

An example of this (and the following) setting is shown in Figure 3.9.

Traveling on L_1, starting from P_0 (either towards δ_d, or, if there is no hyperplane between P_0 and δ_d, away from δ_d), mark the first non–elementary hyperplane encountered as Q_{d-1}, the point of intersection respectively as Q_0. In other words, the interior section of $\overline{P_0 Q_0}$ neither intersects with δ_d, nor with any non–elementary hyperplane in Γ_{d-1}.

Now, the point Q_0 is located in a $(d-1)$–dimensional cell (within Q_{d-1}). Call this cell δ_{d-1}. Further, let γ_d be the convex cone formed by P_0 and δ_{d-1}:

$$\gamma_d = \left\{ \overline{P_0 X_0} \mid X_0 \in \delta_{d-1} \right\} .$$

Notice that

$$\delta_d \subset \gamma_d .$$

Since we assumed the theorem to be true for $(d-1)$, there must be an ordinary $(d-2)$–dimensional hyperplane of Q_{d-1}. I.e., there is a point $C_0 \in Q_{d-1} \cap \Gamma_0$ and a $(d-3)$–dimensional subspace $S_{d-3} \in \Gamma_{d-3}$ such that

$$\text{affine-span}(C_0, S_{d-3}) \in \Gamma_{d-2},$$

[20] i.e. the inner section

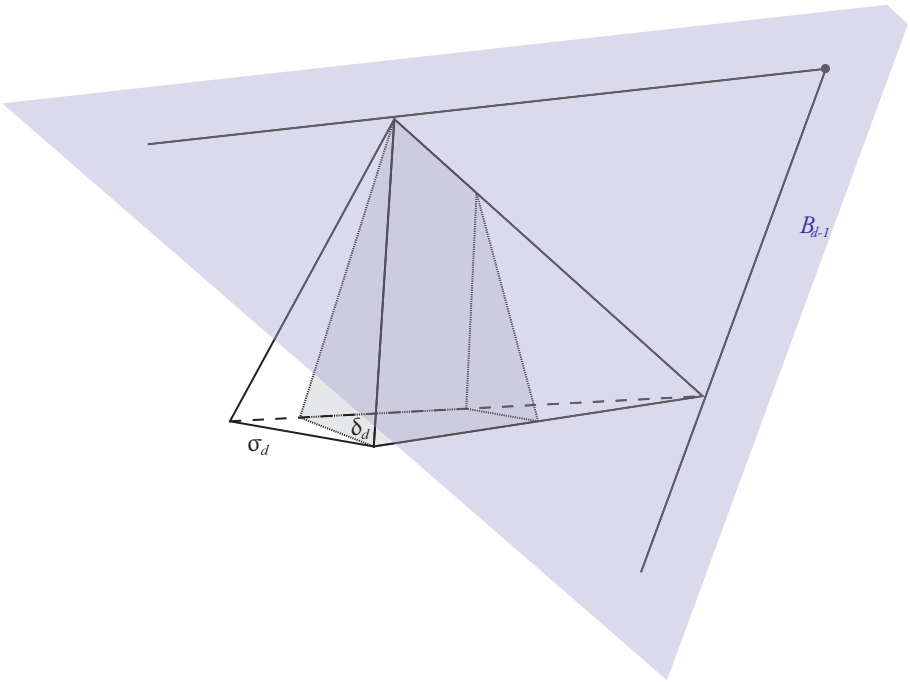

Figure 3.8: A non–elementary hyperplane B_{d-1} (for $d = 3$)

with

$$(\text{affine-span}(C_0, S_{d-3}) \cap (\Gamma_0 \setminus \{C_0\})) \subset S_{d-3}$$

and

$$(\text{affine-span}(C_0, S_{d-3}) \cap \delta_{d-1}) \subset S_{d-3}.$$

If we set

$$S_{d-2} = \text{affine-span}(P_0, S_{d-3}) \in \Gamma_{d-2},$$

we can examine the hyperplane affine-span$(C_0, S_{d-3}) \in \Gamma_{d-3}$.

Let

$$S_{d-1} = \text{affine-span}(C_0, S_{d-2}) \in \Gamma_{d-1}.$$

Remember that γ_d is the cone formed by P_0 and the cell δ_{d-1}. Then, because of

$$(S_{d-1} \cap \gamma_d) \subset S_{d-2}, \tag{3.10}$$

it follows that

$$(S_{d-1} \cap \delta_d) \subset S_{d-2}. \tag{3.11}$$

Now, if all elements of $(S_{d-1} \setminus \{C_0\}) \cap \Gamma_0$ lie within S_{d-2}, then S_{d-1} is an ordinary $(d-1)$–dimensional hyperplane, and the proof is finished.

Therefore, assume that S_{d-1} is not ordinary, i.e. there exists a point

$$A_0 \in ((S_{d-1} \setminus \{C_0\}) \cap \Gamma_0),$$

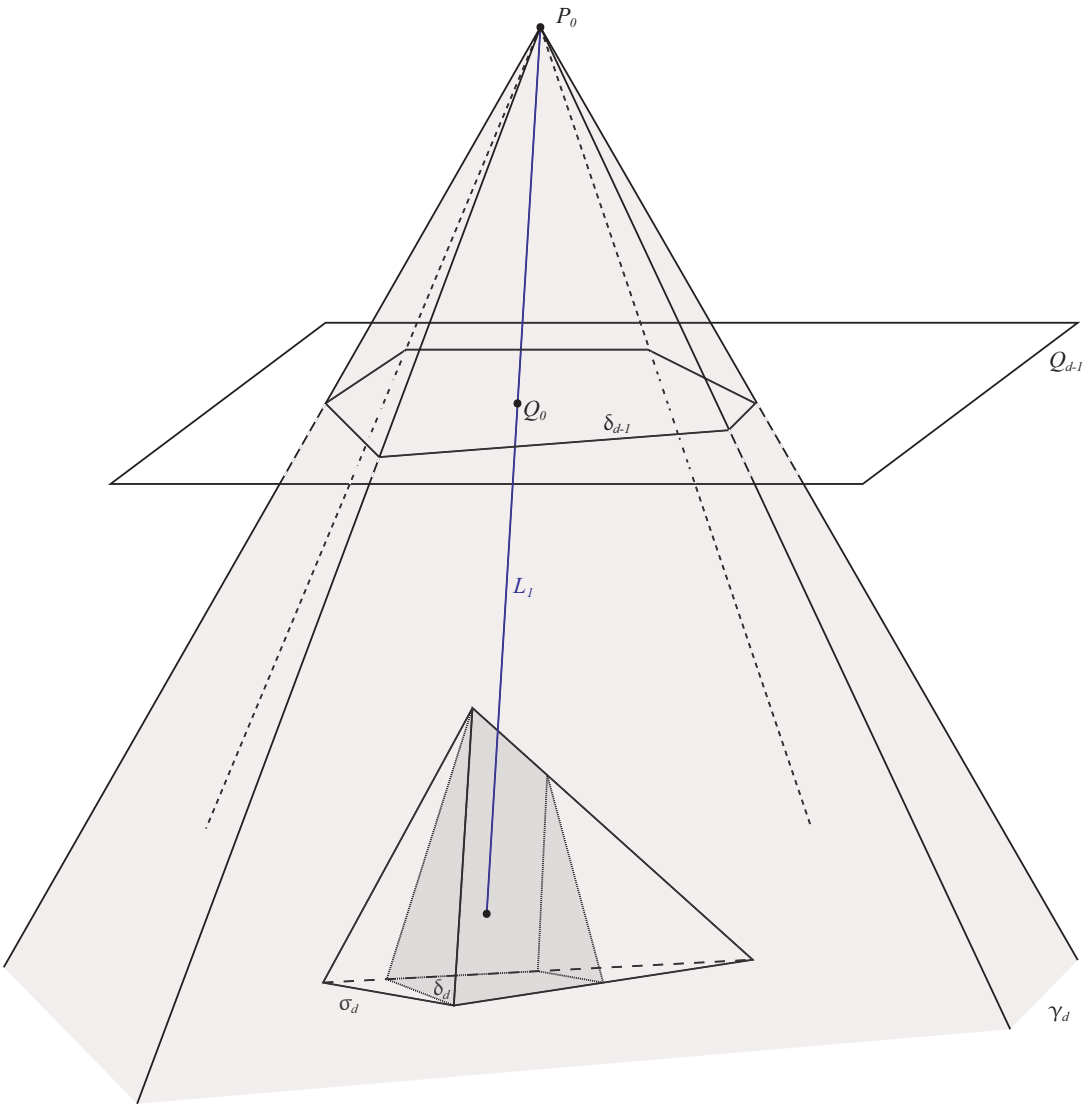

Figure 3.9: Example setting for Hansen's proof (for $d = 3$)

such that

$$A_0 \notin S_{d-2}.$$

From

- $S_{d-1} \cap Q_{d-1} = \text{affine-span}(C_0, S_{d-3})$,
- $\text{affine-span}(C_0, S_{d-3})$ is ordinary, and
- $S_{d-3} \subset S_{d-2}$,

it follows that

$$A_0 \notin Q_{d-1}.$$

Also, because of (3.10) and the fact that $A_0 \notin S_{d-2}$, it follows that

$$A_0 \notin \gamma_d,$$

fulfilling the requirements of Lemma 3.3.12, since the line $\overline{P_0 A_0}$ does not intersect δ_{d-1}.

Let $T_{d-2} \in \Gamma_{d-2}$ be the $(d-2)$–dimensional face of δ_{d-1} found by Lemma 3.3.12, i.e.

$$\text{affine-span}(A_0, T_{d-2}) \in \Gamma_{d-1}.$$

Notice that $\text{affine-span}(A_0, T_{d-2})$ intersects with L_1 precisely in that segment of L_1 which does not meet δ_d. It follows by the choice of Q_0 that $\text{affine-span}(A_0, T_{d-2})$ is an elementary hyperplane, thus finishing the proof.

Thus, it remains to be shown that $\text{affine-span}(A_0, T_{d-2}) \cap \delta_d \subset T_{d-2}$. But, since $Q_{d-1} \cap \text{affine-span}(P_0, T_{d-2}) = T_{d-2}$ and both $Q_{d-1} \in \Gamma_{d-1}$ and $\text{affine-span}(P_0, T_{d-2}) \in \Gamma_{d-1}$, they do not intersect the interior of δ_d. Thus, δ_d is contained in one of the wedges created by Q_{d-1} and $\text{affine-span}(P_0, T_{d-2})$. And, since $\text{affine-span}(A_0, T_{d-2})$ is contained in the other wedge, it cannot have points in common with δ_d, other than exactly on T_{d-2}. □

The main theorem in this chapter is the following of Bonnice and Edelstein [BE67]:

Theorem 3.3.14. *Let S be a finite set of points spanning $\mathbb{R}^{2k}$. Then there is a subset $S' \subset S$ spanning an elementary k–flat of $\mathbb{R}^{2k}$.*

Proof. We prove this by induction.

Case $k = 1$. This is equivalent to the Sylvester–Gallai Theorem, as each finite set of points spanning $\mathbb{R}^2$ has an elementary 1–flat, i.e. a line, containing exactly 2 points of the set.

Assume the statement to hold for $(k-1)$.

Case k. By Theorem 3.3.13, there is an ordinary $(2k-1)$–dimensional hyperplane spanned by a point $A_0 \in S$ and a $(2k-2)$–dimensional flat S_{2k-2}. Notice that S_{2k-2} has, by induction, an elementary $(k-1)$–dimensional flat, say S_{k-1}. As before,

$$\text{affine-span}(A_0, S_{k-1})$$

is an elementary k–flat of $\mathbb{R}^{2k}$ with regards to S. □

3.4 Hyperplane Decomposition

As stated in the introduction to this chapter, the Hyperplane Decomposition Lemma finds a "core" subspace with helpful properties when applied to identity testing.

Remark 3.4.1. Note that for Kayal and Saraf, hyperplanes are subspaces of *any* lower dimension. In our case, since a hyperplane is always of dimension one less than the original space, it would make more sense to call it the *Subspace Decomposition Lemma*.

Lemma 3.4.2 (Hyperplane Decomposition Lemma [KS09])**.** *Let S be a finite set of points spanning $\mathbb{R}^m$. Choose a subset $S_{\text{core}} \subset S$, and let $H_{\text{core}} = \text{affine-span}(S_{\text{core}})$ with $m_{\text{core}} = \dim(H_{\text{core}})$.*

Then there exist subspaces $H_1, \ldots, H_r \subset \mathbb{R}^m$, with $r \geq \frac{m - m_{\text{core}}}{2}$, such that, for $H = \text{affine-span}\{H_i\}_{i \in [r]}$, the following holds:

1. *Each H_i contains H_{core} and is of dimension $(m_{\text{core}} + 1)$:*

$$\forall\ i \in [r] : H_{\text{core}} \subset H_i, \dim(H_i) = m_{\text{core}} + 1.$$

2. *All $(H_i \setminus H_{\text{core}})$ are disjoint. I.e., $\dim(H) = m_{\text{core}} + r$, and for $R \subseteq [r]$, $P_i \in (H_i \setminus H_{\text{core}}) \forall\ i \in R$, the dimension of the space spanned by the P_i is exactly $(|R| - 1)$:*

$$\dim\left(\text{affine-span}\{P_i\}_{i \in R}\right) = |R| - 1.$$

3. *The "additional" dimension of an H_i is spanned by a point of S:*

$$\forall\ i \in [r] : (H_i \setminus H_{\text{core}}) \cap S \neq \emptyset.$$

4. *For all $P \in (H \cap S)$, there is an $i \in [r]$ such that $P \in H_i$, i.e. every point of S that lies within H also lies within an H_i.*

Proof. Let

$$\forall\ P \in (S \setminus H_{\text{core}}) : H_P = \text{affine-span}(P, H_{\text{core}}),$$
$$\mathcal{H} = \{H_P\}_{P \in (S \setminus H_{\text{core}})},$$

and Q be a point in H_{core}. Let H' be an $(m - m_{\text{core}})$–flat of $\mathbb{R}^m$ such that

$$H' \cap H_{\text{core}} = Q,$$
$$\forall\ H_P : H_P \cap H' = L_P,$$

with L_P being lines (i.e., $\dim(L_P) = 1$) such that

$$Q \subseteq L_P.$$

The lines L_P are chosen such that they are contained in H'. We call the pencil of these lines

$$\mathcal{L} = \{L_P\}_{P \in (S \setminus H_{\text{core}})}.$$

Note that no two lines of $\mathcal{L}$ are parallel.

Let H'' be a hyperplane of H' (i.e., $\dim(H'') = (m - m_{\text{core}} - 1)$ and $H'' \subset H'$) and R_P points, such that

- $\forall\ L_P : R_P = (L_P \cap H'')$,
- $H'' = \text{affine-span}(R_P)_{R_P}$,
- $H'' \cap H_{\text{core}} = \emptyset$.

Let

$$S_{\mathcal{L}} = \{R_P\}_{P \in (S \setminus H_{\text{core}})}$$

be the points spanning H'' outside of H_{core}.

For some $r \geq \frac{m - m_{\text{core}}}{2}$, let $Q_1, \ldots, Q_r \in S_{\mathcal{L}}$ be the points spanning an $((r-1)$–dimensional) elementary flat H_{elem} as stated in Hansen's Theorem (Theorem 3.3.13). Let $L_1, \ldots, L_r \in \mathcal{L}$ be the lines corresponding to $Q_1, \ldots, Q_r$.

Since $H'' \cap H_{\text{core}} = \emptyset$ and $H_{\text{elem}} \subset H''$, it must be that $H_{\text{elem}} \cap H_{\text{core}} = \emptyset$.

Then, the following holds:

$$\text{affine-span}\,\{L_i\}_{i \in [r]} = \text{affine-span}(Q, H_{\text{elem}}).$$

For all $i \in [r]$, let

$$H_i = \text{affine-span}(H_{\text{core}}, L_i).$$

Then, each H_i is in $\mathcal{H}$.

Thus, with

$$H = \text{affine-span}\,\{H_i\}_{i \in [r]},$$

it remains to be shown that the properties given hold:

1. That H_{core} is contained in H_i directly follows from the definition of H_i. And since $\dim(L_i) = 1$, and L_i intersects H_{core} in exactly one point, L_i expands the dimension of H_i with regards to H_{core} by exactly 1.

2. Since $\dim(H_{\text{elem}}) = (r-1)$ and $(H_{\text{elem}} \cap H_{\text{core}}) = \emptyset$, it follows that

$$\dim(\text{affine-span}\,(H_{\text{core}}, H_{\text{elem}})) = (m_{\text{core}} + r).$$

Note that

$$\begin{aligned} H &= \text{affine-span}\,\{H_i\}_{i \in [r]} \\ &= \text{affine-span}\left(H_{\text{core}} \cup \{H_i\}_{i \in [r]}\right) \\ &= \text{affine-span}\left(H_{\text{core}} \cup \{Q_i\}_{i \in [r]}\right) \\ &= \text{affine-span}\,(H_{\text{core}}, H_{\text{elem}}). \end{aligned}$$

Thus,

$$\dim(H) = m_{\text{core}} + r.$$

Let R and P_i be as in the property. Then from $H_i = \text{affine-span}\,(P_i, H_{\text{core}})$ follows affine-span $\{H_i\}_{i\in[R]} = \text{affine-span}\,\{H_i\}_{i\in[R]}$. By property 1 and the first part of this property, it must be that $\dim(\text{affine-span}\,\{H_i\}_{i\in[R]}) = m_{\text{core}} + |R|$.

Since $\dim(H_{\text{core}}) = m_{\text{core}}$, the dimension of affine-span $\{P_i\}_{i\in[R]}$ must then be $(|R| - 1)$.

3. Since $H_i \in \mathcal{H}$, and every hyperplane contained in $\mathcal{H}$ is spanned by H_{core} and a point of $(S \setminus H_{\text{core}})$, this property holds.

4. Let P be a point of $(S \setminus H)$ such that P does not lie within any H_i. Then $H_P \in \mathcal{H}$ would be such that

$$H_P \subset H,$$
$$H_P \neq H_i$$

for all $i \in [r]$. Let L_P again be the line intersecting H_P with H'.

Note that L_P cannot be contained in any H_i, since otherwise H_P would be contained in this H_i. So $L_P \in \mathcal{L}$ such that $L_P \subset (H \setminus H_{\text{core}})$.

Thus, $L_P \subset \text{affine-span}\,\{L_i\}_{i\in[r]}$. Call $Q_P = (H'' \cap L_P)$.

Then $Q_P \in \text{affine-span}\,\{Q_i\}_{i\in[r]}$, and thus $Q_1, \ldots, Q_r$ do not span an *elementary* hyperplane, which is a contradiction. □

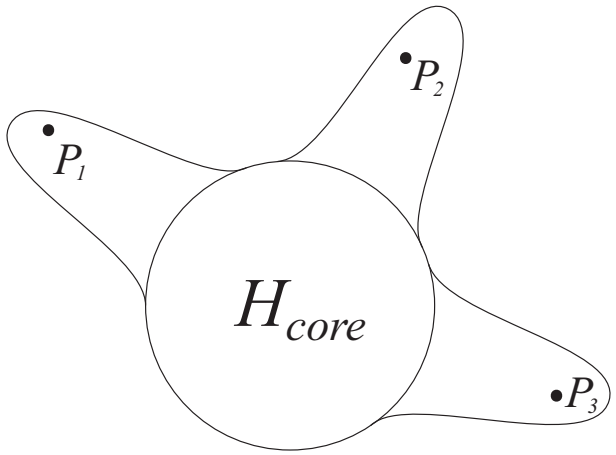

Figure 3.10: Example of Lemma 3.4.2

Chapter 4. Depth–3 Identities

In Chapter 3, we studied geometric properties of sets of points in n-dimensional spaces. In this chapter, we now will show a way of translating circuits into such sets, such that we then can work with these geometric properties in order to prove the final rank bound.

The main idea is simply to translate each linear form into a point in $\mathbb{P}^n$ by using the coefficients within the linear form as coefficients. With the help of this, we will then present a way to reduce the amount of multiplication gates within a circuit of high pairwise–rank in such a way that the resulting circuit also has a high pairwise–rank.

In the next part, we will be studying circuits with top fanin 3. We will introduce the Desmic Conjecture, and thereby show how we can use geometric settings to find identities. Moreover, we will propose new ideas to work with such translations.

The last part of this chapter will prove the rank bound of $k^{O(k)}$.

4.1 From Circuits to Geometry

As in the previous chapters, let C be a depth–3 circuit of the form

$$\begin{aligned} C &= \sum_{i\in[k]} T_i \\ &= \sum_{i\in[k]} \prod_{j\in[d]} l_{ij} \\ &= \sum_{i\in[k]} \prod_{j\in[d]} \sum_{\kappa\in[n]} a_{ij\kappa} x_\kappa. \end{aligned}$$

Then, for each of the linear forms l_{ij}, we can define a point P_{ij} in $\mathbb{P}^n$:

$$P_{ij} = (a_{ij1}, \ldots, a_{ijn})^\tau.$$

Let the multiset of all these points be called S:

$$S = \{P_{ij} \mid i \in [k], j \in [d]\}.$$

Let us label each T_i with a different color i. Then we can also label each point P_{ij} with the color i. Notice that each point can have several colors, since the same linear form may appear in various gates.

So for every point P_{ij}, we have a multiset $K(P_{ij})$, containing the colors (with multiplicity) of this point. Let $C_{P_{ij}}$ be the set of colors of the point P_{ij},[1] and $\text{mult}(P_{ij}, i)$ be the multiplicity of the color i in the multiset $K(P_{ij})$.

On the other hand, let S_i denote the multi-subset of S containing only those points that are colored i. The symmetric difference of two such multisets is

$$S_i \Delta S_j = \{P \in S \mid \text{mult}(P, i) \neq \text{mult}(P, j)\}.$$

With this setting, for all $i, j \in [k], i \neq j$, notice the following translations between circuits and their geometric counterparts:

$$\text{rank}(C) \geq B \longleftrightarrow \dim(\text{span}(S)) \geq B - 1,$$
$$\text{pairwise-rank}(C) \geq A \longleftrightarrow \dim(\text{span}(S_i \Delta S_j)) \geq A - 1.$$

We will now introduce a rather technical definition, which will be used in the proof of the Fanin Reduction Lemma (Lemma 4.2.4).

Definition 4.1.1. Given the setting above and

- a flat F_{core} spanned by points in S,
- a set $\mathcal{F} = \{F \mid S' \subseteq S, F = \text{span}(S'), F_{\text{core}} \subseteq F \text{ as sets of points}\}$ of flats being spanned by points of S, and each of these flats fully contains the flat F_{core},
- a pair of colors $i \neq j$, and
- an $A \in \mathbb{N}$,

we say that a flat $F \in \mathcal{F}$ *splits* the colors i, j, if there is a point

$$P \in ((F \setminus F_{\text{core}}) \cap (S_i \Delta S_j)).^2$$

We say that the pair of colors i, j is *over-split* (with regards to F_{core}, $\mathcal{F}$, and A) if $\mathcal{F}$ contains more than A splitting flats, else we call the pair *under-split*.

4.2 Fanin Reduction

The Fanin Reduction Lemma offers a way to "remove" multiplication gates from a simple circuit with high pairwise–rank, and obtaining a new simple circuit that still has high pairwise–rank.

Definition 4.2.1. Let $C \in \Sigma\Pi\Sigma(k)$ be of the form $C = \sum_{i \in k} A_i$. Then the *pairwise–rank of* C is

$$\text{pairwise-rank}(C) := \min_{1 \leq i < j \leq k} \{\text{rank}\left(\text{sim}(A_i + A_j)\right\}.$$

If $k = 1$, we set $\text{pairwise-rank}(C) = 1$.

[1] I.e., $C_{P_{ij}}$ is $K(P_{ij})$ without multiplicity.

[2] View F and F_{core} as sets of points, i.e. $(F \setminus F_{\text{core}}) = \{P' \mid P' \in F, \nexists\, F' \in F_{\text{core}} : P' \in F'\}$.

Definition 4.2.2. For a circuit C containing a linear form l, let $C|_{l=0}$ be a circuit obtained as follows: Let $\pi : \mathbb{R}^n \to \mathbb{R}^n$ be of $\text{rank}(\pi) = (n-1)$, and $\ker(\pi) = \text{span}(l)$. Apply π on C to obtain a circuit $C|_{l=0} = \pi(C)$.

Remark 4.2.3. With this definition, notice the following:

- For two different linear transformations, e.g. π and π', $\pi(C)$ and $\pi'(C)$ are equivalent[3]. Therefore, we will simply refer to *the* circuit $C|_{l=0}$.
- Since the linear form l occurring in the circuit C is transformed to zero, the whole gate containing l vanishes. Thus, $C|_{l=0}$ contains strictly fewer gates than C.
- For a circuit C containing a linear form l, if $C \equiv 0$, then also $C|_{l=0} \equiv 0$.

Lemma 4.2.4 (Fanin Reduction Lemma [KS09])**.** *Let $k, A \in \mathbb{N}$, $B = 3(A+1)k^2$. Let C be a simple $\Sigma\Pi\Sigma(k)$ circuit with* pairwise-rank$(C) \geq A$ *and* rank$(C) \geq B$.
Then there is a linear form l within C for which pairwise-rank$(C|_{l=0}) \geq A$.

Proof. Following the notation of Section 4.1, since pairwise-rank$(C) \geq A$, for all $i, j \in [k], i \neq j$,

$$\dim(\text{span}(S_i \Delta S_j)) \geq A - 1$$

holds. Therefore, for every $i \neq j$, pick a set of A points $S_{ij} \subseteq (S_i \Delta S_j)$, such that these span an elementary $(A-1)$-flat:

$$\dim(\text{span}(S_{ij})) = A - 1.$$

Denote the union of these subsets by S_{core}:

$$S_{\text{core}} = \bigcup_{i,j \in [k], i \neq j} S_{ij}.$$

We can now apply the Hyperplane Decomposition Lemma 3.4.2 with $F_{\text{core}} = \text{span}(S_{\text{core}})$. We then receive flats $F_1, \ldots, F_r$ for an $r \geq \frac{m - m_{\text{core}}}{2} = \frac{B - k^2 A}{2}$ with $m_{\text{core}} = \dim(F_{\text{core}})$, each of dimension $(m_{\text{core}} + 1)$.

We now select a flat $F' \in \{F_i\}_{i \in [r]}$, such that for all points of S within $(F' \setminus F_{\text{core}})$ the multiplicity of a pair of colors i and j is equal. This must exist, since each of the (at most) k^2 pairs of under-split colors has at most A splitting flats, but there are $r \geq (A+1)k^2$ flats to choose from. Thus, for all undersplit $\{i, j\}$, and for all $P \in (S \cap (F' \setminus F_{\text{core}}))$, we have $\text{mult}(P, i) = \text{mult}(P, j)$.

Let l be the linear form corresponding to such a point with equal multiplicity, i.e. $\pi(l) = P \in (S \cap (F' \setminus F_{\text{core}}))$. This will be the form used to reduce the rank as required by the lemma. Thus, we still need to show that pairwise-rank$(C|_{l=0}) \geq A$. I.e., for all $A_i|_{l=0}$ and $A_j|_{l=0}$ that remained non–zero, the it has to be shown that the following holds:

$$\text{rank}(\text{sim}(A_i|_{l=0} + A_j|_{l=0})) \geq A.$$

[3]I.e., there is an invertible linear τ, such that $\pi(C) = \tau(\pi'(C))$.

Notice that if C_P contained all colors, l would be contained in all gates of C, contradicting the simplicity of C. Also, if C_P were to contain $(k-1)$ colors, then $C|_{l=0}$ would only consist of a single multiplication gate, and the proof would be finished. We therefore only consider $|C_P| \leq (k-2)$.

Thus, we can select two colors i and j that are not in C_P. The proof then finishes by distinguishing two cases:

Case 1: The pair i and j is over-split. I.e., there exist at least A flats other than F' splitting i and j. Without loss of generality, let these be $F_1, \ldots, F_A$. For $t \in [A]$, let $Q_t \in F_t$ such that $\text{mult}(Q_t, i) \neq \text{mult}(Q_t, j)$. Let $S_l = \{Q_t\}_{t\in[A]}$.

We can show that

$$\dim(\text{span}(S_l^\pi)) = A - 1$$

by the following: By Lemma 3.4.2, it follows that $\dim(\text{span}(\{P\} \cup S_l)) = A$. Since 0 is not in $\text{span}(\{P\} \cup S_l))$, we have $\dim(\text{span}(\{0, P\} \cup S_l)) = A + 1$. Then, from $\dim(\ker(\pi)) = 1$, $\pi(0) = 0$, and $\pi(P) = 0$, it directly follows that $\dim(\text{span}(S_l^\pi)) = A - 1$.

In order to show that $\dim(\text{span}((S_i \Delta S_j)^\pi))$, it is now enough to show that $S_l^\pi \subseteq (S_i \Delta S_j)^\pi)$. For that, take a $Q_t^\pi \in S_l^\pi$, and the respective $Q_t \in S_l$, such that $\pi(Q_t) = Q_t^\pi$.[4] Since $\text{mult}(Q_t, i) \neq \text{mult}(Q_t, j)$, it follows that $\text{mult}(Q_t^\pi, i) \neq \text{mult}(Q_t^\pi, j)$, and thereby $Q_t^\pi \in (S_i \Delta S_j)^\pi$. Since this holds for all $t \in [A]$, it must be that $S_l^\pi \subseteq (S_i \Delta S_j)^\pi)$.

This concludes the proof for the first case.

Case 2: The pair i and j is under-split. I.e., for all points $Q \in (S \cap (F' \setminus F_{\text{core}})) : \text{mult}(Q, i) = \text{mult}(Q, j)$.

Let

$$\{Q_1, \ldots, Q_s\} = (S_i \Delta S_j)^{\text{core}} = F_{\text{core}} \cap (S_i \Delta S_j).$$

Note that (by choice of F_{core}) $\dim(\text{span}((S_i \Delta S_j)^{\text{core}})) \geq (A - 1)$. As in Case 1, we will show that this $(S_i \Delta S_j)^{\text{core}}$ contains a subset of points which span a $(A-1)$-dimensional space, which will conclude the proof: Let

$$(S_i \Delta S_j)^{\text{core},\pi} = \{\pi(Q) \mid Q \in (S_i \Delta S_j)^{\text{core}}\}.$$

As before, by $\text{span}((S_i \Delta S_j)^{\text{core},\pi}) \subseteq F_{\text{core}}$ and $P \notin F_{\text{core}}$, we have

$$\dim(\text{span}(\{P\} \cup (S_i \Delta S_j)^{\text{core}})) = A,$$

and thereby

$$\dim(\text{span}((S_i \Delta S_j)^{\text{core},\pi})) = A - 1.$$

It now only remains to be shown that $(S_i \Delta S_j)^{\text{core},\pi} \subseteq (S_i \Delta S_j)^{\text{core}}$: For $t \in [s]$, let $\pi(Q_t) = Q_t^\pi \in (S_i \Delta S_j)^{\text{core},\pi}$. Notice that the line combining P with Q_t intersects

[4] This Q_t must be unique: Say it is not. Then let Q_t' be another such point that maps to Q_t^π. Then P, Q_t, and Q_t' are collinear. But this is a contradiction to Lemma 3.4.2.

F_{core} in exactly one point.[5] Thus, any point Q'_t (other than Q_t) for which $\pi(Q'_t) = Q^\pi_t$ must lie within $(F' \setminus F_{\text{core}})$, and then $\text{mult}(Q'_t, i) = \text{mult}(Q'_t, j)$. By this and the definition of $(S_i \Delta S_j)^{\text{core}}$, Q_t is then the only preimage of Q^π_t for which $\text{mult}(Q_t, i) \neq \text{mult}(Q_t, j)$ holds. But thereby we get $\text{mult}(Q^\pi_t, i) \neq \text{mult}(Q^\pi_t, j)$, and thus $Q^\pi_t \in (S_i \Delta S_j)^{\text{core},\pi}$.

Therefore,

$$(S_i \Delta S_j)^{\text{core},\pi} \subseteq (S_i \Delta S_j)^{\text{core}},$$

completing the proof. □

4.3 Fanin 3 Circuits

Before studying arbitrary depth–3 circuits, we will restrict ourselves to depth–3 circuits of fanin 3 in order to better understand the correlation between circuits and their geometric representation. We will, e.g., only have to deal with 3 colors, since we restrict ourselves to 3 products: $C = T_1 + T_2 + T_3$.

4.3.1 An Identity of $\text{rank}_{\mathbb{R}} = 4$

The Desmic Conjecture [Bor83] presents a configuration of three sets of points spanning $\mathbb{R}^3$, such that any line intersecting two of these sets also intersects the third. Using the 12–point Desmic configuration, we are able to use this as a guide to finding an identity of *fanin 3* and *depth 4*. It can be shown that this is the *only* such configuration of less than 27 points, and therefore this identity is also the *unique* identity with these properties.

Figure 4.1 shows the 12–point Desmic configuration, with the 3 different sets of 4 points each: red, green, and blue points. Notice the labels on the various vertices: The axes are labeled as such. Then, we pick one vertex, and label it y. Then, the rest of the vertices are simply labeled by their geometric configuration in space (with regards to the vertex y).

This can be directly converted into an identity C of rank 4 (note that the labels are grouped into 3 multiplication gates by their respective color):

$$\begin{aligned} C(y, x_1, x_2, x_3) &= y(y + x_1 + x_2)(y + x_1 + x_3)(y + x_2 + x_3) \\ &\quad -(y + x_1)(y + x_2)(y + x_3)(y + x_1 + x_2 + x_3) \\ &\quad +x_1 x_2 x_3(2y + x_1 + x_2 + x_3) \\ &= 0. \end{aligned}$$

[5] Since the line lies entirely within F', and since $\dim(F') = (\dim(F_{\text{core}}) + 1)$.

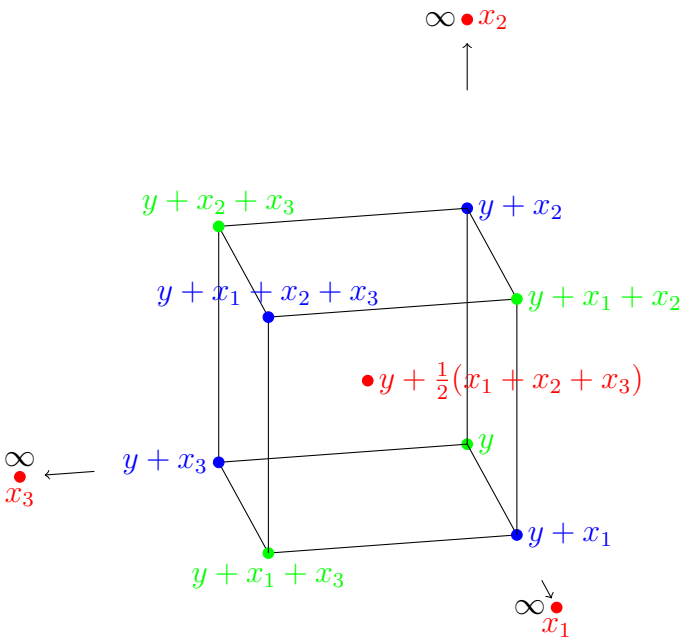

Figure 4.1: 12–point Desmic conjecture

4.3.2 No Identity of $\text{rank}_{\mathbb{C}} = 5$?

In the complex projective plane, and with ω being the 3^{rd} root of unity, i.e. $\omega = e^{2\pi i/3}$, it is obvious that the nine points of inflection

$$\begin{array}{ccc} (1:-1:0) & (0:1:-1) & (-1:0:1) \\ (1:-\omega:0) & (0:1:-\omega) & (-\omega:0:1) \\ (1:-\omega^2:0) & (0:1:-\omega^2) & (-\omega^2:0:1) \end{array}$$

of the cubic curve defined by

$$x^3 + y^3 + z^3 = xyz.$$

form a SG configuration [Cox48].

We might expect that this helps us to find an identity over $\mathbb{C}$ of rank $= 5$, as the Desmic Conjecture does over $\mathbb{R}$. But after working with this configuration, this seems unlikely, and thus we give the following:

Conjecture 4.3.1. *There is no identity of* $\text{rank}_{\mathbb{C}} = 5$ *and fanin 3.*

4.3.3 Geometric Questions over Fanin 3: Matchings

We will now further study the relationships between geometry and circuits of depth 3 and top fanin 3.

The main tool we will be using is a *matching* between two sets of linear forms. We are mainly interested in the amount of matchings we can find for two such sets, and suggest bounds for several types of matchings.

Definition 4.3.2. We call a non–zero linear form l a *matching* between two (multi)sets U and V of linear forms, if it induces a mapping of U onto V via

$$m_l : U \xrightarrow[u \mapsto (c_u u + d_u l)]{} V,$$

with $c_u \in \mathbb{F}^\times$, $d_u \in \mathbb{F}$. We call c_u the *scaling factor* of u to m_l.

For convenience, we will also use l for the map m_l.

Definition 4.3.3. Two linear forms l and m are *similar* ($l \sim m$), if there exists a $c \in \mathbb{F}^\times$ such that $l = cm$. Two matchings are *similar* if their underlying linear forms are similar.

Two matchings are *coprime* if their underlying linear forms are coprime.

Lemma 4.3.4. *Let U, V be two sets of linear forms, each of* rank $= 2$*. Then there can be infinitely many coprime matchings between U and V.*

Proof. Without loss of generality, let $U = \{x_1, x_2\}$ and $V = \{(x_1 + x_2), (x_1 - x_2)\}$. Then $l = (x_1 + cx_2)$ is a matching for all $c \neq 0$, because

$$x_1 + x_2 = \frac{1}{c} l + \left(1 - \frac{1}{c}\right) x_1,$$
$$x_1 - x_2 = l + (-1 - c)\, x_2.$$

□

Lemma 4.3.5. *If $u_1, u_2, u_3 \in U$ are linearly independent, then for all $v_1, v_2, v_3 \in V$ there exists at most one linear form l, up to coprimality, that maps*

$$u_1 \overset{l}{\mapsto} v_1,$$
$$u_2 \mapsto v_2,$$
$$u_3 \mapsto v_3.$$

Proof. Let $S := \text{span}(u_1, v_1) \cap \text{span}(u_2, v_2) \cap \text{span}(u_3, v_3)$. Then $l \in S$, and since $\text{rank}(S) = 2$ iff $\text{span}(u_1, v_1) = \text{span}(u_2, v_2) = \text{span}(u_3, v_3) = S$, it follows that $\text{rank}(S) < 2$, since u_1, u_2, u_3 are linearly independent.

Thus, there can be only one l up to coprimality. □

Example 4.3.6. Let us look at matchings of $U = \{x_1, x_2, x_3\}$ onto $V = \{\alpha x_1 + x_2 + x_3, x_1 + \alpha x_2 + x_3, x_1 + x_2 + \alpha x_3\}$, over $\mathbb{F}[x_1, x_2, x_2]$. We can find 4 coprime matchings:

$$l_1 = (x_1 + x_2 + x_3),$$
$$l_2 = (x_1 + \alpha x_2 + \alpha x_3),$$
$$l_3 = (\alpha x_1 + x_2 + \alpha x_3),$$
$$l_4 = (\alpha x_1 + \alpha x_2 + x_3).$$

Let us further study matchings m between $U = \{x_1, x_2, x_3, x_4\}$ and V. We can find different types of matchings:

Type 1 $m \notin \text{span}(U)$: Then, without loss of generality, we can view this as $m = x_5$, so

$$U \overset{x_5}{\mapsto} V = \{x_1 + \alpha_1 x_5, \ldots, x_4 + \alpha_4 x_5\}.$$

There cannot be any other matching l, because

$$l \in \bigcap_{i \in [4]} \text{span}\left(x_i, l(x_i)\right),$$

but $\text{span}(x_i, l(x_i)) \neq \text{span}(x_j, l(x_j))$ for $i \neq j$.

Type 2 $m = \sum_{i \in [4]} x_i$: Then, V will be of the form

$$\begin{aligned} V = \{&\alpha_1 x_1 + x_2 + x_3 + x_4, \\ &x_1 + \alpha_2 x_2 + x_3 + x_4, \\ &x_1 + x_2 + \alpha_3 x_3 + x_4, \\ &x_1 + x_2 + x_3 + \alpha_4 x_4\} \end{aligned}$$

and it follows that $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4$.

Type 3 $m = \sum_{i \in [3]} x_i$: Then, without loss of generality, V will be of the form

$$\begin{aligned} V = \{&\alpha_1 x_1 + x_2 + x_3, \\ &x_1 + \alpha_2 x_2 + x_3, \\ &x_1 + x_2 + \alpha_3 x_3, \\ &x_1 + x_2 + x_3 + \alpha x_4\} \end{aligned}$$

and it again follows that $\alpha_1 = \alpha_2 = \alpha_3$.

Notice that for every type, there is, up to coprimality, at most one matching possible.

Definition 4.3.7. For a set X and an element x (not necessarily in X), we say that $x \in_s X$ if there exists a $c \in \mathbb{R}, c \neq 0$, such that $cx \in X$.

Let us further study the method of viewing the linear forms as vectors. We use the coefficients of the linear forms as entries in the vectors. We can then apply methods known from linear algebra. Let l be a linear form. By abuse of notation, we also call its associated vector l. We can then make the following statement:

Definition 4.3.8. Let l, l' be linear forms viewed as vectors in $\mathbb{R}^n$. Following the natural definitions in linear algebra, define

- the inner product $(l, l') = \sum_{i \in [n]} l_i l'_i \in \mathbb{R}$,

- $|l|$ as $|l|^2 = (l, l)$,
- and

$$l_{\|l'} := \left(l, \frac{l'}{|l'|}\right) \frac{l'}{|l'|}$$
$$l_{\perp l'} := l - l_{\|l'}.$$

Remark 4.3.9. Using this notation, let us view a matching m between $U = \{u_1, \ldots, u_d\}$ and $V = \{v_1, \ldots, v_d\}$. Then we can see, by looking at the image of u_i, that

$$m_{\perp u_i} \in_s \{v_{1\perp u_i}, \ldots, v_{d\perp u_i}\}.$$

Theorem 4.3.10. *If* $\operatorname{rank} U \geq 3$*, then there are at most* (d^2+d-2) *matchings between* U *and a* V*, where* $d = |U| = |V|$*.*

Proof. Let m be a matching from $U = \{u_1, \ldots, u_d\}$ to $V = \{v_1, \ldots, v_d\}$. Assume without loss of generality that u_1, u_2, u_3 are linearly independent.

If $m \notin \operatorname{span}(u_1, u_2)$, say, $m \in u + \operatorname{span}(u_1, u_2)$ where $u \notin \operatorname{span}(u_1, u_2)$, then the images respectively of u_1 and u_2 suggest that

$$m_{\perp u_1} \sim (u_2 + w_1), \tag{4.1}$$
$$m_{\perp u_2} \sim (u_1 + w_2) \tag{4.2}$$

for some $w_1, w_2 \notin \operatorname{span}(u_1, u_2)$. And m is already fixed (up to similarity) by these two properties, since for $m = au_1 + bu_2 + cw$, with $w_1 = dw$ and $w_2 = d'w$, it follows

$$\frac{b}{c} = \alpha,$$
$$\frac{a}{c} = \beta,$$

and thus

$$m = c(\beta u_1 + \alpha u_2 + u).$$

By (4.1) and (4.2), these are d^2 possibilities for m.

Now, if $m \in \operatorname{span}(u_1, u_2)$, then (4.1) and (4.2) will be

$$m_{\perp u_1} \in_s \{\alpha u_2, \ldots\},$$
$$m_{\perp u_2} \in_s \{\beta u_1, \ldots\}.$$

Then, without loss of generality, select the subsets $U' = \{u_3 + \operatorname{span}_i(u_1, u_2)\}_{i=1}^{k}$ of U and $V' = \{u_3 + \operatorname{span}_i(u_1, u_2)\}_{i=k+1}^{2k}$ of V. Then, the matchings between U' and V' are already all the matchings between U and V that are in $\operatorname{span}(u_1, u_2)$. Notice that $m \mid (\prod(V') - \prod(U'))$, where $\prod(X)$ for a set X is the product of all its elements. Since U' and V' are sets of at most $(d-2)$ linear forms, the degree of the polynomial $(\prod(V') - \prod(U'))$ is at most $(d-2)$. Taking the matchings in $\operatorname{span}(u_1, u_2)$ into consideration, we get at most $(d-2)$ matchings in this case.

Thus, there are at most $(d^2 + d - 2)$ matchings between U and V. □

We can further use this method to view the problem geometrically. Let us first only look at $\text{rank}(U) = 3$, that means that our vectors are all in $\mathbb{R}^3$. Let these vectors be the generators of lines passing through the origin, i.e. for a linear form l we have, using its corresponding vector, a line $\{al : a \in \mathbb{R}^3\}$. The vectors of matchings between two linear forms l_1 and l_1' will, by the definition of matchings, have to be on the plane $\text{span}(l_1, l_1')$. Thus, for sets of linear forms $U = \{l_1, l_2, \ldots\}$ and $V = \{l_1', l_2', \ldots\}$ and a matching m, the planes $\text{span}(l_i, m(l_i))$ must intersect in one line, generated by the vector m, c.f. figure (4.2).

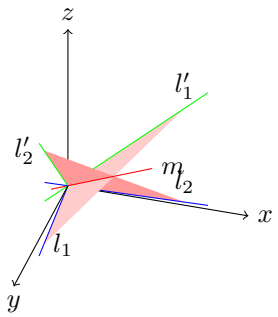

Figure 4.2: Matching m in $\mathbb{R}^3$

Note that all objects intersect the origin, and we are not interested in depth information. Thus, let us draw the unit ball, i.e. the ball with radius 1, around the origin. Now, every line, or rather linear form, intersects this ball exactly twice, and we can find a plane cutting the ball in half with exactly one intersection of every linear form on each side. We continue to work with the projection on one of these halves. Thus, our lines turn into points, our planes turn into lines, and we lose one dimension, so we can work in $\mathbb{P}^2$. After flattening out the projection, figure (4.2) could look something like figure (4.3).

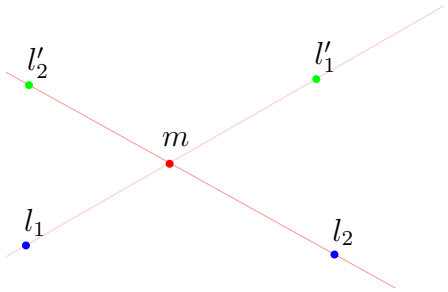

Figure 4.3: Matching m in $\mathbb{P}^2$

Definition 4.3.11. The *convex hull* of a set of points in $\mathbb{R}^2$ is the smallest polygon, such that all points are inside or on the polygon, and the polygon does not have an inner angle greater than π. In other words, any line intersecting the polygon must intersect exactly twice, or be tangent to an edge or corner.

Remark 4.3.12. For a set of n points in $\mathbb{R}^2$, the convex hull can have at most n corners.

Up to now, we differentiated between points coming from U and V by coloring them appropriately. Let us ignore the colors now, and study intersections of lines between arbitrarily picked points. Obviously, there will be at least as many matchings found as if one were to consider the coloring.

First, let us study the case that $|U| = |V| = 3$ and where the convex hull is a hexagon. We can always find a matching in the center of the hexagon, i.e. the intersection of the lines between opposite corners. The most matchings we can find in this case is 4: one inside and three outside of the convex hull. This situation is displayed in figure (4.4).

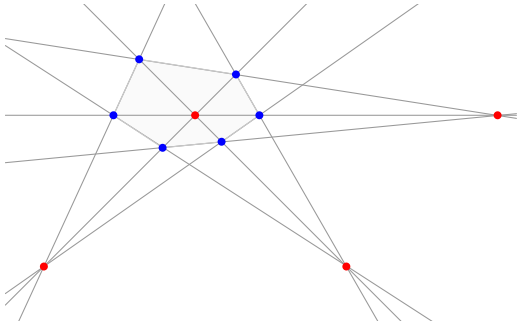

Figure 4.4: Matchings for $|U| = |V| = 3$

Lemma 4.3.13. *There are U and V, with* $\operatorname{rank}(U \bigcup V) = 3$ *and $d = |U| = |V|$, such that there are at least d matchings between U and V.*

Proof. Let us put a regular d–gon in the (x, y)–plane with its midpoint on the origin. Mark the disjoint sets of parallel lines with P_i, and the corners of the d–gon counter-clockwise with v_i. Notice that each line in any set P_k contains a v_i and v_j such that $i \neq j \mod 2$.

We can then rotate the d–gon around the x and y axes, such that each P_i creates a matching. Since there are d sets of parallel lines, we get d matchings.

Consider the equation

$$(y + v_1)(y + v_3) \cdots (y + v_{2d-1}) - (y + v_2)(y + v_4) \cdots (y + v_{2d}).$$

The matchings show that $(v_1 - v_2), (v_2 - v_3), \ldots, (v_{2d-1} - v_{2d})$ all divide this. Thus, we have an identity $T_1 - T_2 + T_3 = 0$, with

$$\begin{aligned} T_1 &:= (y + v_1)(y + v_3) \cdots (y + v_{2d-1}) \\ T_2 &:= (y + v_2)(y + v_4) \cdots (y + v_{2d}) \\ T_3 &:= (v_1 - v_2)(v_2 - v_3) \cdots (v_{2d-1} - v_{2d}), \end{aligned}$$

of degree d and rank 3.

If d is odd, then the following also holds:

$$y \mid (T_1 - T_2)$$

Thus, there exist $d+1$ coprime matchings. □

Lemma 4.3.14. *Let U and V be sets of linear forms, such that $d := |U| = |V| > 4$ and $r := \text{rank}(U) = \text{rank}(V) \geq 4$. Then, with n being the number of matchings between U and V, it can be that*

$$n \geq \frac{|U|}{\text{rank}(U)}.$$

Proof. Let $U = U_1 \overline{\cup} \ldots \overline{\cup} U_{r-2}$ and $V = V_1 \overline{\cup} \ldots \overline{\cup} V_{r-2}$. Let $v_1, \ldots, v_{2d} \in \text{span}(x_1, x_2)$, such that they form a regular $2d$–gon, and let

$$U_i := \{y_i + v_1, y_i + v_3, \ldots\}$$
$$V_i := \{y_i + v_2, y_i + v_4, \ldots\}$$

Then $\text{rank}(U) = \text{rank}(V) = r$ and $|U| = |V| = (r-2)d$.

From the rank = 3 case, we know that there can be d or more coprime matchings between U_i and V_i. □

We now formulate some conjectures over $\mathbb{R}$:

Conjecture 4.3.15. *Let U and V be sets of linear forms, such that $d := |U| = |V| > 4$ and $\text{rank}(U) = \text{rank}(V) \geq 4$. Then there are strictly less than d coprime matchings between U and V.*

Conjecture 4.3.16. *Let $U \subset \mathbb{R}^2$ with $d := |U|$ finite. Then there are not more than $\frac{d}{2}$ distinct matchings outside of the convex hull of U.*

We conjecture even a stronger version of Conjecture 4.3.15:

Conjecture 4.3.17. *For any two sets U and V of linear forms, with $|U| = |V| \geq 4$, there cannot be more than $(log_2|U| + 2)$ non–similar matchings.*

Example 4.3.18 (Matchings on the hypercube)**.** Let U and V be sets of linear forms with $\text{rank}(U) = \text{rank}(V) \geq 4$. Let $\mathbb{F}$ be such that $\text{char}(\mathbb{F}) \neq 2$.

Obviously, if $|U| \neq |V|$, there cannot be any matchings.

Let us view $\mathbb{F}[y, x_1, \ldots, x_n]$. We build the sets

$$U = \left\{ y + \sum_{i \in I} x_i \right\}_{I, |I| even}$$

and

$$V = \left\{ y + \sum_{i \in I} x_i \right\}_{I, |I| odd}.$$

Say, our linear form inducing a matching is $(a_0 y + \sum_{i=1}^{n} a_i x_i)$ with coefficients $a_j \in \mathbb{F}, j = 0, \ldots, n$. So, for all $u \in U$, we have to find a $v \in V$ with corresponding c, d, such that

$$\begin{aligned} v &= c\left(y + \sum_{I} x_i\right) + d\left(a_o y + \sum_{i=1}^{n} a_i x_i\right) \\ v' &= c'\left(y + \sum_{I'} x_i\right) + d'\left(a_o y + \sum_{i=1}^{n} a_i x_i\right) \\ &\vdots \end{aligned}$$

Thus, we can build a corresponding system of equations, e.g. for our $U = \{y, y + x_1 + x_2, y + x_1 + x_3, \ldots\}$:

$$\begin{aligned} c + da_0 &= 1, [da_1, da_2, da_3, \ldots] \\ c' + d'a_0 &= 1, [c' + d'a_1, c' + d'a_2, d'a_3, \ldots] \\ c'' + d''a_0 &= 1, [c'' + d''a_1, d''a_2, c'' + d''a_3, \ldots] \\ &\vdots \end{aligned}$$

with each right part having an odd number of 1's, the rest all 0.

Now, we study whether such an $l = a_0 y + \sum_{i=1}^{n} a_i x_i$ exists.

On the left column we have $c + da_0 = c' + d'a_0 = \ldots = 1$. If $a_0 = 0$, then

$$c = c' = \ldots = 1. \tag{4.3}$$

Else,

$$c^{(i)} = (1 - d^{(i)} a_0) \neq 1 \tag{4.4}$$

as $d^{(i)} \in \mathbb{F}^{\times}$.

If a_i and a_j, $i \neq j$, are non–zero, choose a row where the columns i and j do not have a c–summand. (This is always possible due to our choice of U.) Now, notice that the coefficients of x_i and x_j for elements in V can only be 0 or 1, with 0 being impossible in this case, since $d \in \mathbb{F}^{\times}$. Therefore, a_i and a_j must be equal, and this row's respective d must be

$$d = \frac{1}{a_i}. \tag{4.5}$$

Now choose a row where exactly one of the columns i or j has a c–summand (without loss of generality the i'th column). As above, since $da_j = 1$, $a_i = a_j$, and $c + da_i \in \{0, 1\}$, it follows that

$$c = -1. \tag{4.6}$$

(Remember that $c \in \mathbb{F}^{\times}$!) If $a_0 = 0$, this contradicts (4.3), thus it follows that $a_0 \neq 0$ and $d = d' = \ldots = \frac{2}{a_0}$ from (4.4).

It also follows that if $a_0 = 0$, we cannot have more than one a_i non–zero (from (4.6)), whence we are done.

If $a_0 \neq 0$ (and thus $c = -1$ and $d = \frac{2}{a_0}$), for all $i \in [n]$, we get (by again choosing a row without a c–summand in the i'th column) $a_i = \frac{1}{d}$. This can be normalized to $a_0 = 2$, $d = 1$, $a_i = 1$.

Note that for even n, we have a line with c–summands in all columns. If we take $a_0 = 2$, $d = 1$, $a_i = 1$, $c = -1$, we get

$$\begin{aligned} & (-1)\left(y + \sum_{i=1}^{n} x_i\right) + \left(2y + \sum_{i=1}^{n} x_i\right) \\ = \; & (-1+2)y + (-1+1)x_1 + (-1+1)x_2 + \ldots \\ = \; & y \notin V. \end{aligned}$$

Therefore, for even n, a_0 can only be 0.

It is easy to verify that for even n, the forms x_i ($i \in [n]$) yield matchings, while for odd n, $(2y + x_1 + \ldots + x_n)$ is another matching.

Remark 4.3.19. Note that the case $|U| = |V| = 4$ is the hypercube as described in 4.3.1.

4.4 Higher Fanins

The following is the main result of Kayal and Saraf [KS09].

Theorem 4.4.1 (Rank Bound for $\Sigma\Pi\Sigma(k)$ Circuits)**.** *Let $c : \mathbb{N} \to \mathbb{N}, c(k) = 3^k((k+1)!)^2$. Let C be a simple, minimal $\Sigma\Pi\Sigma(k)$ circuit s.t. $C \equiv 0$.*

Then

$$\operatorname{rank}(C) \leq c(k).$$

Proof. Let C be a simple and minimal circuit,

$$C = \sum_{i \in [k]} A_i = \sum_{i \in [k]} \sum_{j \in [d]} l_{ij}.$$

We prove the theorem via induction: It is obviously true for $k = 1$. Assume the theorem to hold for $k - 1$. Then, for k, we can distinguish between two cases:

- pairwise-rank$(C) < (c(k) - c(k-1))$:

  Thus, there are $i, j \in [k], i \neq j$, for which

  $$\operatorname{rank}(\operatorname{sim}(A_i + A_j)) < (c(k) - c(k-1)).$$

  Without loss of generality, let $\operatorname{sim}(A_i + A_j)$ be a polynomial in the variables $X_1, \ldots, X_t$, with $t = \operatorname{rank}(\operatorname{sim}(A_i + A_j))$.

  For all $\iota \in [t]$, let $\alpha_\iota \in_{\mathrm{R}} [0, 1]$, and replace the variable X_ι with $\alpha_\iota Z$. Then, $\operatorname{sim}(A_i + A_j)$ is a polynomial in one variable, Z.

Let C' be the circuit C with the same adjustments, i.e.

$$C(X_1, \ldots, X_n) = C'(\alpha_1 Z, \ldots, \alpha_t Z, X_{t+1}, \ldots, X_n).$$

Then,

- $C' \equiv 0$ (since $C \equiv 0$),
- $\text{rank}(C') = k - 1$ (since A_i and A_j are merged into one gate),
- C' is minimal (since C is minimal),
- $\text{rank}(\text{sim}(C')) > c(k-1)$.

Thus, it follows that $\text{sim}(C')$ is a simple, minimal $\Sigma\Pi\Sigma(k-1)$ circuit computing 0 with $\text{rank}(C') > c(k-1)$. This contradicts our assumption.[6] Therefore, this case cannot happen.

Let us therefore study the second case:

- $\text{pairwise-rank}(C) \geq (c(k) - c(k-1))$:

  Thus, for all $i, j \in [k]$, the following holds:

$$\begin{aligned} \text{rank}(\text{sim}(A_i + A_j)) &\geq c(k) - c(k-1) \\ &> \frac{c(k)}{2} \\ &> c(k-1) + 1. \end{aligned}$$

  We can now apply the Fanin Reduction Lemma (4.2.4) with $A = (c(k-1)+1)$ and $B = c(k)$, in order to find a linear form l in C, such that for $C' = C|_{l=0}$, we have

$$\text{pairwise-rank}(C') \geq c(k-1) + 1.$$

  I.e., for any subset $S \subseteq [k]$, the following holds:

$$\text{rank}(\sum_{i \in S} A_i) \geq c(k-1) + 1.$$

  Let $A'_i = A_i|_{l=0}$. Since $C' = C|_{l=0} = \sum_{i \in [k]} A_i|_{l=0} = \sum_{i \in S} A'_i = 0$, pick the smallest S possible such that the equation holds.[7]

  With this S, $\text{sim}(\sum_{i \in S} A'_i)$ is a minimal, simple $\Sigma\Pi\Sigma(k-1)$ circuit[8] computing the zero polynomial with

$$\text{rank}(\text{sim}(\sum_{i \in S} A'_i)) \geq c(k-1) + 1,$$

  contradicting the assumption.

[6] That the theorem is valid for $(k-1)$.

[7] Such an S will have at most $(k-1)$ elements, since one of the k-many A_i is set to zero by $\cdot|_{l=0}$ and thus does not contribute to the sum.

[8] Actually, it is a $\Sigma\Pi\Sigma(k')$ circuit for $k' = |S| \leq (k-1)$.

Therefore, the rank of C cannot be larger than $c(k)$:

$$\text{rank}(C) \leq c(k). \qquad \square$$

Chapter 5. Future work

By finding new and improving the previous algebraic and combinatorial results, Saxena and Seshadhri have proposed a rank bound of $3k^2$ [SS10].[1] This result will improve the previously best known rank bound from an *exponential* $k^{O(k)}$ into a *quadratic* $O(k^2)$ over reals. Dvir and Shpilka conjectured polynomial rank bound in k, and actually showed identities of rank $O(k)$. Thus, Saxena and Seshadhri's rank bound is almost an optimal bound.

5.1 Quadratic rank bound

We will give a short overview of the basic theorems (without proof) and definitions, following the 3-step order as proposed in [SS10]. Note that these work over arbitrary fields $\mathbb{F}$, except for the bounds given in the third step, which only work over reals since they are strictly connected to the Sylvester-Gallai Theorems.

Matching the Gates in an Identity. For a base field $\mathbb{F}$, let R be the ring of polynomials in n variables over $\mathbb{F}$, i.e. $R := \mathbb{F}[x_1, \ldots, x_n]$. Let $L(R)$ be the set of all linear forms in R, i.e. $L(R) := \left\{ \sum_{i \in [n]} a_i x_i \mid a_1, \ldots, a_n \in \mathbb{F} \right\}$.

Theorem 5.1.1 (Theorem 5 [SS10]: Matching-nucleus)**.** *Let $C = \sum_{i \in [k]} T_i$ be a minimal $\Sigma\Pi\Sigma(k, d)$-circuit computing the zero polynomial. Then there exists a linear subspace $K \subseteq L(R)$, such that*

- $\text{rank}(K) < k^2$, *and*
- $\forall i \in [k]$, *there is a K-matching π_i between T_1 and T_i.*

This linear subspace K is called mat-nucleus of C.

This splits C into a part with terms that are within mat-nucleus, having *low* rank k^2, and the term not within mat-nucleus, which are *similar*[2].

Certificate for Linear Independence of Gates. In addition to the above setting, let $L_K(T_i) := (L(T_i) \cap K)$. For a list S of linear forms, let $M(S) := \prod_{l \in S} l$ (or, since S may be empty, $M(\emptyset) := 1$).

[1] As before, k denotes the fanin of the simple, minimal depth-3 circuit C, representing the polynomial computing zero.

[2] I.e., for each term l in T_1 and not in mat-nucleus, there is a term in mat-nucleus matching l to a term in each of the remaining $T_2, \ldots, T_k$.

Theorem 5.1.2 (Theorem 7 [SS10]: Nucleus)**.** *Let $C = \sum_{i\in[k]} T_i$ be a minimal $\Sigma\Pi\Sigma(k,d)$-circuit computing the zero polynomial. Let $\{T_i \mid i \in \mathcal{I}\}$ be a maximal set of linearly independent terms, with $1 \leq k' := \mathcal{I} < k$. Then there exists a linear subspace $K \subseteq L(R)$, such that*

- $\operatorname{rank}(K) < 2k^2$.
- *$\forall i \in [k]$, there is a K-matching π_i between T_1 and T_i.*
- *$\forall i \in \mathcal{I}$, let $K_i := M(L_k(T_i))$. The terms $\{K_i \mid i \in \mathcal{I}\}$ are linearly independent.*

This linear subspace K is called nucleus of C.

In addition to Theorem 5.1.1, it now holds that if C is strongly minimal[3], then the nucleus identity is also strongly minimal.

Invoking Sylvester-Gallai Theorems. Saxena and Seshadhri came up with the notion of SG_k-*closed* set, which will allow to take a slightly different approach on the Sylvester-Gallai Theorem:

Definition 5.1.3 (Theorem 9, 11 [SS10]: SG_k-closed)**.** Let $S \subset \mathbb{P}^{n-1}$. Then, S is called SG_k-*closed*, if for every set of k linearly independent points $V \subseteq S$, $\operatorname{span}(V)$ contains at least $(k+1)$ points of S.

Define the following operator: $\mathrm{SG}_k(\mathbb{F}, m)$ is the largest possible rank of an SG_k-closed set of at most m points in $\mathbb{P}^{n-1}$.

Theorem 5.1.4 (Theorem 10 [SS10]: Sylvester-Gallai for higher dimension)**.** *Let $k \in \mathbb{N}$ and $S \subset \mathbb{R}^n$ be a finite set of points. If S is SG_k-closed, then*

$$\operatorname{rank}(S) \leq 2(k-1).$$

Note that this is actually only a rephrased version of Theorem 3.3.13, and it is therefore currently known to be only true over $\mathbb{R}$. From this, it follows that $\mathrm{SG}_k(\mathbb{R}, m) \leq 2(k-1)$.

One further definition will lead to the main theorem, the rank bound.

Definition 5.1.5 (Definition 14 [SS10]: Independent fanin)**.** Let $C = \sum_{i\in[k]T_i}$ be a $\Sigma\Pi\Sigma(k,d)$ circuit. The *ind-fanin of C* is the size of the maximal $\mathcal{I} \subseteq [k]$ such that $\{T_i\}_{i\in\mathcal{I}}$ are linearly independent polynomials.

Theorem 5.1.6 (Theorem 15 [SS10]: Final bound)**.** *Let $|\mathbb{F}| > d$. Let C be a simple, minimal, ind-fanin k', $\Sigma\Pi\Sigma(k,d)$ circuit. Then the rank of C is at most*

$$2k^2 + (k-k')\mathrm{SG}_{k'}(\mathbb{F}, d).$$

[3] I.e., $T_1, \ldots, T_{k-1}$ are linearly independent.

5.2 Conclusion

The aforementioned result by Saxena and Seshadhri [SS10] is extremely close to Dvir and Shpilka's $O(k)$ rank conjecture, so one might pose the question whether further intense study towards an even better bound would still be fulfilling.

Actually improving the rank bound to $O(k)$ might still be interesting, since it may yield further methods for depth–3 PIT that will become useful when studying circuits of depth 4.[4] Since we currently do not have too many tools for higher depths, it might therefore still be worthwhile first tackling depth-3 circuits.

[4] Agrawal and Vinay have shown that an efficient blackbox algorithm for depth–4 circuits also solves PIT for "low" degree circuits of *all* depths greater than 4 [AV08].

List of Figures

List of Algorithms

Index

Bibliography

[Agr05] M. Agrawal. Proving lower bounds via pseudo-random generators. *FSTTCS 2005: Foundations of Software Technology and Theoretical Computer Science*, 3821:92–105, 2005.

[AV08] M. Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *IEEE 49th Annual IEEE Symposium on Foundations of Computer Science, 2008. FOCS'08*, pages 67–75, 2008.

[AZ01] M. Aigner and G. M. Ziegler. *Proofs from THE BOOK*. Springer, second edition, 2001.

[BE67] W. E. Bonnice and M. Edelstein. Flats associated with finite sets in $\mathbb{P}^d$. *Nieuw archief voor wiskunde*, (3), XV:11–14, 1967.

[BM90] P. Borwein and W. O. J. Moser. A survey of Sylvester's problem and its generalizations. *Aequationes Mathematicae*, 40(1):111–135, 1990.

[Bor83] P. Borwein. The desmic conjecture. *Journal of Combinatorial Theory, Series A*, 35(1):1–9, 1983.

[Cox48] H. S. M. Coxeter. A problem of collinear points. *American Mathematical Monthly*, 55:26–28, 1948.

[dB09] M. de Bondt. Another generalization of Mason's ABC–theorem. *Arxiv preprint arXiv:0707.0434*, 2009.

[DS05] Z. Dvir and A. Shpilka. Locally decodable codes with 2 queries and polynomial identity testing for depth 3 circuits. In *Proceedings of the thirty–seventh annual ACM symposium on Theory of computing*, pages 592–601. ACM, 2005.

[EPS06] N. Elkies, L. M. Pretorius, and K. J. Swanepoel. Sylvester–Gallai theorems for complex numbers and quaternions. *Discrete and computational geometry*, 35(3):361–373, 2006.

[Han65] S. Hansen. A generalization of a theorem of Sylvester on the lines determined by a finite point set. *Math. Scand*, 16:175–180, 1965.

[Kel86] L. M. Kelly. A resolution of the Sylvester–Gallai problem of J.–P. Serre. *Discrete and Computational Geometry*, 1(1):101–104, 1986.

[KI04] V. Kabanets and R. Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1):1–46, 2004.

[KS01] A. R. Klivans and D. A. Spielman. Randomness efficient identity testing of multivariate polynomials. In *Proceedings of the thirty–third annual ACM symposium on Theory of computing*, pages 216–223. ACM, 2001.

[KS07] N. Kayal and N. Saxena. Polynomial identity testing for depth 3 circuits. *Computational Complexity*, 16(2):115–138, 2007.

[KS09] N. Kayal and S. Saraf. Blackbox Polynomial Identity Testing for Depth 3 Circuits. *Electronic Colloquium on Computational Complexity*, 2009.

[Mas84] R. C. Mason. Diophantine equations over function fields. *London Mathematical Society Lecture Note Series*, 96, 1984.

[RDB61] J. Rainwater, P. H. Diananda, and A. Bager. 4908, Partition of a Triangle. *American Mathematical Monthly*, 68(4):386–387, 1961.

[Sax09] N. Saxena. Progress on polynomial identity testing. *Bulletin of the EATCS no*, 99:49–79, 2009.

[Sch80] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM (JACM)*, 27(4):701–717, 1980.

[SS09] N. Saxena and C. Seshadhri. An almost optimal rank bound for depth–3 identities. In *Proceedings of the 24th Annual Conference on Computational Complexity (CCC)*, pages 137–148, 2009.

[SS10] N. Saxena and C. Seshadhri. From Sylvester–Gallai configurations to rank bounds: Improved black–box identity test for depth–3 circuits. *Arxiv preprint arXiv:1002.0145*, 2010.

[Sto81] W. W. Stothers. Polynomial identities and Hauptmoduln. *The Quarterly Journal of Mathematics*, 32(2):349–370, 1981.

[Zip79] R. Zippel. Probabilistic algorithms for sparse polynomials. *Symbolic and algebraic computation*, pages 216–226, 1979.