The Multivariate Schwartz-Zippel Lemma

M. Levent Doğan

Joint work with A. A. Ergür, J. D. Mundo and E. Tsigaridas

Technische Universität Berlin

EuroCG 2020

Würzburg - 18.03.2020

Table of Contents

Introduction and the Main Theorem

Applications

The Algorithm

There is a wide literature on counting number of zeroes of a polynomial on a finite grid thanks to its applications to Polynomial Identity Testing, Incidence Geometry and Extremal Combinatorics.

Theorem (The Schwartz-Zippel-DeMillo-Lipton Lemma)

Let $\mathbb{F}$ *be a field, let* $S \subseteq \mathbb{F}$ *be a finite set and let* $0 \neq p \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ *be a polynomial of degree* d*. Suppose* $|S| > d$ *and let* $S^n := S \times S \times \cdots \times S$*. Then we have*

$$|Z(p) \cap S^n| \leq d|S|^{n-1}$$

where $Z(p) = \{v \in \mathbb{F}^n \mid p(v) = 0\}$ *denotes the zero locus of* p*.*

There is a wide literature on counting number of zeroes of a polynomial on a finite grid thanks to its applications to Polynomial Identity Testing, Incidence Geometry and Extremal Combinatorics.

Theorem (The Schwartz-Zippel-DeMillo-Lipton Lemma)

Let $\mathbb{F}$ be a field, let $S \subseteq \mathbb{F}$ be a finite set and let $0 \neq p \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial of degree d. Suppose $|S| > d$ and let $S^n := S \times S \times \cdots \times S$. Then we have

$$|Z(p) \cap S^n| \leq d|S|^{n-1}$$

where $Z(p) = \{v \in \mathbb{F}^n \mid p(v) = 0\}$ denotes the zero locus of p.

A theorem on the same direction is given by Alon:

Theorem (Alon's Combinatorial Nullstellensatz)

Let $p \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a polynomial of degree $d = \sum_{i=1}^n d_i$ for some positive integers d_i and assume that the coefficient of the monomial $\prod_{i=1}^n x_i^{d_i}$ in p is non-zero. Let $S_i \subseteq \mathbb{F}$ be finite sets with $|S_i| > d_i$ and let $S := S_1 \times S_2 \times \cdots \times S_n$. Then, there exists $v \in S$ such that

$$p(v) \neq 0.$$

In this talk, we want to obtain similar results for *multi-grids*.

Notation

We call a sequence $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ of positive integers a partition of n into m parts if $n = \lambda_1 + \lambda_2 + \cdots + \lambda_m$. In this case, we write $\lambda \underset{m}{\vdash} n$. Given a partition $\lambda \underset{m}{\vdash} n$, we introduce the notation $\overline{x_1} = (x_1, x_2, \ldots, x_{\lambda_1}), \overline{x_2} = (x_{\lambda_1+1}, x_{\lambda_1+2}, \ldots, x_{\lambda_1+\lambda_2})$ and so on.
Given finite sets $S_1 \subseteq \mathbb{F}^{\lambda_1}, S_2 \subseteq \mathbb{F}^{\lambda_2}, \ldots, S_m \subseteq \mathbb{F}^{\lambda_m}$, we call the product

$$S := S_1 \times S_2 \times \cdots \times S_m$$

the multi-grid defined by $S_1, S_2, \ldots, S_m$.

In this talk, we want to obtain similar results for *multi-grids*.

Notation

We call a sequence $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_m)$ of positive integers a partition of n into m parts if $n = \lambda_1 + \lambda_2 + \cdots + \lambda_m$. In this case, we write $\lambda \underset{m}{\vdash} n$. Given a partition $\lambda \underset{m}{\vdash} n$, we introduce the notation $\overline{x_1} = (x_1, x_2, \ldots, x_{\lambda_1}), \overline{x_2} = (x_{\lambda_1+1}, x_{\lambda_1+2}, \ldots, x_{\lambda_1+\lambda_2})$ and so on.

Given finite sets $S_1 \subseteq \mathbb{F}^{\lambda_1}, S_2 \subseteq \mathbb{F}^{\lambda_2}, \ldots, S_m \subseteq \mathbb{F}^{\lambda_m}$, we call the product

$$S := S_1 \times S_2 \times \cdots \times S_m$$

the multi-grid defined by $S_1, S_2, \ldots, S_m$.

Given a multivariate polynomial $p \in \mathbb{C}[\overline{x_1}, \overline{x_2}, \ldots, \overline{x_m}]$, we want to bound number of zeros of p can have on a multi-grid S. It turns out that this task is impossible without imposing some conditions for p.

Example

Let $g_1 \in \mathbb{C}[x_1, x_2] \setminus \mathbb{C}$ and $g_2 \in \mathbb{C}[x_3, x_4] \setminus \mathbb{C}$. For $h_1, h_2 \in \mathbb{C}[x_1, x_2, x_3, x_4]$, set

$$p = g_1 h_1 + g_2 h_2.$$

Observe that $Z(g_1)$ and $Z(g_2)$ are planar curves in $\mathbb{C}^2$ and $Z(p)$ contains $Z(g_1) \times Z(g_2)$. In particular, p can vanish on arbitrarily large Cartesian products!

Example

Let $g_1 \in \mathbb{C}[x_1, x_2] \setminus \mathbb{C}$ and $g_2 \in \mathbb{C}[x_3, x_4] \setminus \mathbb{C}$. For $h_1, h_2 \in \mathbb{C}[x_1, x_2, x_3, x_4]$, set

$$p = g_1 h_1 + g_2 h_2.$$

Observe that $Z(g_1)$ and $Z(g_2)$ are planar curves in $\mathbb{C}^2$ and $Z(p)$ contains $Z(g_1) \times Z(g_2)$. In particular, p can vanish on arbitrarily large Cartesian products!

Definition

Let $\lambda \underset{m}{\vdash} n$. An affine variety $\mathcal{V} \subseteq \mathbb{C}^n$ is called λ-reducible if there exist <u>positive dimensional</u> varieties $\mathcal{V}_i \subseteq \mathbb{C}^{\lambda_i}$ such that

$$\mathcal{V}_1 \times \mathcal{V}_2 \times \cdots \times \mathcal{V}_m \subseteq \mathcal{V}.$$

Otherwise, we say $\mathcal{V}$ is λ-irreducible. A polynomial $p \in \mathbb{C}[x_1, x_2, \ldots, x_n]$ is said to be λ-reducible (resp. λ-irreducible) if the hypersurface $Z(p)$ defined by p is λ-reducible (resp. λ-irreducible).

The Main Theorem

Theorem (D., Ergür, Mundo, Tsigaridas)

Let $\lambda \underset{m}{\vdash} n$ be a partition of n into m parts and let $p \in \mathbb{C}[x_1, x_2, \ldots, x_n]$ be a λ-irreducible polynomial of degree $d \geq 2$. Let $S_i \subseteq \mathbb{C}^{\lambda_i}$ and let $S := S_1 \times S_2 \times \cdots \times S_m$ be the multi-grid defined by S_i. Then, for all $\varepsilon > 0$, we have

$$|Z(p) \cap S| = O_{n,\varepsilon}(d^5 \prod_{i=1}^{m} |S_i|^{1-\frac{1}{\lambda_i+1}+\varepsilon} + d^{2n^4} \sum_{i=1}^{m} \prod_{j \neq i} |S_j|)$$

where $O_{n,\varepsilon}$ notation only hides constants depending on n and ε.

The Main Theorem

Theorem (D., Ergür, Mundo, Tsigaridas)

Let $\lambda \vdash_m n$ be a partition of n into m parts and let $p \in \mathbb{C}[x_1, x_2, \ldots, x_n]$ be a λ-irreducible polynomial of degree $d \geq 2$. Let $S_i \subseteq \mathbb{C}^{\lambda_i}$ and let $S := S_1 \times S_2 \times \cdots \times S_m$ be the multi-grid defined by S_i. Then, for all $\varepsilon > 0$, we have

$$|Z(p) \cap S| = O_{n,\varepsilon}\left(d^5 \prod_{i=1}^{m} |S_i|^{1-\frac{1}{\lambda_i+1}+\varepsilon} + d^{2n^4} \sum_{i=1}^{m} \prod_{j \neq i} |S_j|\right)$$

where $O_{n,\varepsilon}$ notation only hides constants depending on n and ε.

Observation

As long as we check λ-irreducibility over $\mathbb{C}$, the bound works over any subfield of $\mathbb{C}$.

Table of Contents

Introduction and the Main Theorem

Applications

The Algorithm

Point-Line Incidences

Theorem (Szemerédi-Trotter)

Let P be a set of points and L be a set of lines in the real plane, $\mathbb{R}^2$. Let

$$\mathcal{I}(P, L) = \{(p, l) \in P \times L \mid p \in l\}$$

be the set of incidences between P and L. Then

$$|\mathcal{I}(P, L)| = O(|P|^{2/3}|L|^{2/3} + |P| + |L|).$$

Point-Line Incidences

Theorem (Szemerédi-Trotter)

Let P be a set of points and L be a set of lines in the real plane, $\mathbb{R}^2$. Let

$$\mathcal{I}(P, L) = \{(p, l) \in P \times L \mid p \in l\}$$

be the set of incidences between P and L. Then

$$|\mathcal{I}(P, L)| = O(|P|^{2/3}|L|^{2/3} + |P| + |L|).$$

The theorem holds if we replace $\mathbb{R}^2$ with $\mathbb{C}^2$. To our knowledge, the complex version is first proven by Tóth. As our first application, we use the main theorem to recover the above bound, except for ε in the exponent:

Theorem (Cheap Szemerédi-Trotter Theorem)

Let P be a set of points and L be a set of lines in $\mathbb{C}^2$ (or $\mathbb{R}^2$). Then, for any $\varepsilon > 0$, there are at most

$$O(|P|^{2/3+\varepsilon}|L|^{2/3+\varepsilon} + |P| + |L|)$$

incidences between P and L.

Proof.

Let $p = x_1 + x_2x_3 + x_4 \in \mathbb{C}[x_1, x_2, x_3, x_4]$. It is straightforward to show that p is $(2,2)$-irreducible: For $u = (u_1, u_2), v = (v_1, v_2) \in \mathbb{C}^2$, the equations

$$p(u_1, u_2, x_3, x_4) = 0,$$
$$p(v_1, v_2, x_3, x_4) = 0$$

are (affine) linear in x_3, x_4, thus has at most one solution. We deduce that $Z(p)$ cannot contain a 2×2-multi-grid, which implies that p is $(2,2)$-irreducible.

Observe that given a point $z = (z_1, z_2) \in \mathbb{C}^2$ and a line $l : x + ay + b = 0$ with non-zero slope, we have $z \in l$ if and only if $p(z_1, z_2, a, b) = 0$. Thus, using the main theorem, the number of incidences between points in P and lines in L with a non-zero slope is bounded by

$$O(|P|^{2/3+\varepsilon}|L|^{2/3+\varepsilon} + |P| + |L|).$$

Note that there are at most $|P|$ incidences between points in P and lines in L with a zero slope, so the above bound works in general. $\square$

Unit Distance Problem

Erdős's Unit Distance Problem

Given a finite set P of points in $\mathbb{R}^2$, what is the maximum number of pairs $(u, v) \in P \times P$ with $\|u - v\|_2 = 1$?

Erdős conjectured that the number of pairs of points in P with Euclidean distance 1 apart is bounded by $O(|P|^{1+\varepsilon})$ for all $\varepsilon > 0$.

Unit Distance Problem

Erdős's Unit Distance Problem

Given a finite set P of points in $\mathbb{R}^2$, what is the maximum number of pairs $(u, v) \in P \times P$ with $\|u - v\|_2 = 1$?

Erdős conjectured that the number of pairs of points in P with Euclidean distance 1 apart is bounded by $O(|P|^{1+\varepsilon})$ for all $\varepsilon > 0$.

Theorem (Spencer, Szemerédi, Trotter)

Let P be a finite set of points in $\mathbb{R}^2$. Then, the number of pairs in P with Euclidean distance 1 apart is bounded by $O(|P|^{4/3})$.

Unit Distance Problem

Erdős's Unit Distance Problem

Given a finite set P of points in $\mathbb{R}^2$, what is the maximum number of pairs $(u, v) \in P \times P$ with $\|u - v\|_2 = 1$?

Erdős conjectured that the number of pairs of points in P with Euclidean distance 1 apart is bounded by $O(|P|^{1+\varepsilon})$ for all $\varepsilon > 0$.

Theorem (Spencer, Szemerédi, Trotter)

Let P be a finite set of points in $\mathbb{R}^2$. Then, the number of pairs in P with Euclidean distance 1 apart is bounded by $O(|P|^{4/3})$.

Tao and Solymosi studied the complex version of the problem and came up with a similar bound except for the ε in the exponent.

Theorem (Tao, Solymosi)

Let P be a finite set of points in $\mathbb{C}^2$. Then, for all $\varepsilon > 0$, the cardinality of the set

$$\{((u_1, u_2), (v_1, v_2)) \in P \times P \mid (u_1 - v_1)^2 + (u_2 - v_2)^2 = 1\}$$

is bounded by $O(|P|^{4/3+\varepsilon})$.

We reproduce the same bound using the main theorem:

Proof.

Let $p = (x_1 - y_1)^2 + (x_2 - y_2)^2 - 1 \in \mathbb{C}[x_1, x_2, y_1, y_2]$. We first observe that $Z(p)$ contains no 3×3-multi-grid. For any triple $u, v, w \in \mathbb{C}^2$, the system

$$\begin{aligned} p(u_1, u_2, y_1, y_2) &= 0, \\ p(v_1, v_2, y_1, y_2) &= 0, \\ p(w_1, w_2, y_1, y_2) &= 0 \end{aligned}$$

has at most one solution: If u, v, w are on an affine (complex) line, then a direct computation shows that there is no solution. If not, then taking pairwise differences of the equations we get

$$\begin{bmatrix} y_1 & y_2 \end{bmatrix} \cdot \begin{bmatrix} v_1 - u_1 & w_1 - u_1 & w_1 - v_1 \\ v_2 - u_2 & w_2 - u_2 & w_2 - v_2 \end{bmatrix} = 0.$$

Since u, v, w are affinely independent, we deduce that $(y_1, y_2) = (0, 0)$. Thus, p is $(2, 2)$-irreducible and applying the main theorem to $\varepsilon/2$ yields the result. □

Table of Contents

Introduction and the Main Theorem

Applications

The Algorithm

We have a symbolic algorithm providing a solution to the following problem:

Problem

Set $\lambda = (k, k, \ldots, k) \underset{m}{\vdash} n$. Given a polynomial $p \in \mathbb{Q}[\overline{x_1}, \overline{x_2}, \ldots, \overline{x_m}]$ of degree d, are there polynomials $g_i \in \mathbb{Q}[\overline{x_i}] \setminus \mathbb{Q}$ and polynomials $h_i \in \mathbb{Q}[\overline{x_1}, \overline{x_2}, \ldots, \overline{x_m}]$ such that

$$p = g_1 h_1 + g_2 h_2 + \cdots + g_m h_m?$$

Equivalently, given a hypersurface $\mathcal{V} \subseteq \mathbb{C}^n$, do there exist hypersurfaces $\mathcal{V}_i \subseteq \mathbb{C}^k, i = 1, \ldots, m$ such that

$$\mathcal{V}_1 \times \mathcal{V}_2 \times \cdots \times \mathcal{V}_m \subseteq \mathcal{V}?$$

We have a symbolic algorithm providing a solution to the following problem:

Problem

Set $\lambda = (k, k, \ldots, k) \underset{m}{\vdash} n$. Given a polynomial $p \in \mathbb{Q}[\overline{x_1}, \overline{x_2}, \ldots, \overline{x_m}]$ of degree d, are there polynomials $g_i \in \mathbb{Q}[\overline{x_i}] \setminus \mathbb{Q}$ and polynomials $h_i \in \mathbb{Q}[\overline{x_1}, \overline{x_2}, \ldots, \overline{x_m}]$ such that

$$p = g_1 h_1 + g_2 h_2 + \cdots + g_m h_m?$$

Equivalently, given a hypersurface $\mathcal{V} \subseteq \mathbb{C}^n$, do there exist hypersurfaces $\mathcal{V}_i \subseteq \mathbb{C}^k, i = 1, \ldots, m$ such that

$$\mathcal{V}_1 \times \mathcal{V}_2 \times \cdots \times \mathcal{V}_m \subseteq \mathcal{V}?$$

The algorithm detects whether a polynomial $p \in \mathbb{C}[\overline{x}_1, \ldots, \overline{x}_m]$ is λ-irreducible in the special case $\lambda = (k, k, \ldots, k) \underset{m}{\vdash} n$. We leave detecting λ-irreducibility in the general case as an open problem. Suggestions and ideas are welcomed!

Thank you for your attention!