Computer Sciene 2: Advanced Algorithms Fall 2026

Lecture 1 — August, 31, 2026

Prof. Jarosław Błasiok *Scribe: Ivana Crescenzi*

1 Overview

This lecture covers Polynomial Identity Testing (PIT): a case where randomness gives an efficient algorithm where no efficient deterministic one is known. We build the tool needed for PIT, namely the the Schwartz Zippel lemma starting from the univariate root bound, turn it into a randomized algorithm with one-sided error and apply it to a combinatorial problem: deciding whether a graph has a perfect matching, first for bipartite graphs and then for general graphs.

2 Math things

Here I will list some useful notations and definitions that are going to be useful in the lesson:

2.1 Zero polynomials and monomial notation

Definition 1 (Zero polynomial)**.** *A polynomial* $p(x) = a_0 + a_1 + \cdots + a_d x^d$ *is a* zero polynomial *if*

$$a_0 = a_1 = \cdots = a_d = 0.$$

Let's not forget this is a condition on coefficients, NOT on the function values!

A polynomial in m variables is a sum of monomials:

$$p(x_1, \ldots, x_m) = \sum_{\alpha} c_\alpha x^\alpha, \qquad x^\alpha := \prod_i x_i^{\alpha_i},$$

where $\alpha = (\alpha_1, \ldots, \alpha_m)$ collects the exponents. E.g. for $\alpha = (2, 0, 1)$, $x^\alpha = x_1^2 x_3$.

Example. $p(x_1, x_2) = 4x_1^2 x_2 + 7x_1 - 3$ has 3 monomials: $\alpha = (2, 1)$, $c_{(2,1)} = 4$; $\alpha = (1, 0)$, $c_{(1,0)} = 7$; $\alpha = (0, 0)$, $c_{(0,0)} = -3$. Saying p is a zero polynomial just means $c_\alpha = 0$ for every α.

2.2 Characteristic of a field

$\mathbb{F}_5$ is the finite field with 5 elements $\{0, 1, 2, 3, 4\}$, addition and multiplication mod 5:

$$3 + 4 = 7 \equiv 2 \pmod 5, \qquad 4 \times 4 = 16 \equiv 1 \pmod 5, \qquad 5 \equiv 0 \pmod 5.$$

More gnerally, for a prime p, $\mathbb{F}_p = \{0, 1, \ldots, p-1\}$ with arithmetic mod p. A field has *characteristic* p if adding 1 to itself p times gives 0; this behaves very differently from $\mathbb{Q}$ or $\mathbb{R}$, which have characteristic 0.

2.3 Binomial coefficients

$$\binom{n}{k} = \frac{n!}{k!(n-k)!}$$

counts the number of ways of choosing k elements out of n, ignoring order. They appear when expanding a binomial:

$$(x+y)^n = \sum_{k=0}^{n} \binom{n}{k} x^{n-k} y^k.$$

Fact 2. *If p is prime, then $\binom{p}{k}$ is divisible by p for every $0 < k < p$.*

2.4 Permutations and the determinant

A permutation of $\{1, \ldots, m\}$ is the following $\sigma : \{1, \ldots, m\} \to \{1, \ldots, m\}$; there are $m!$ of them. For an $m \times m$ matrix M,

$$\det(M) = \sum_{\sigma \in S_m} \operatorname{sgn}(\sigma) \prod_{i=1}^{m} M_{i,\sigma(i)},$$

where the sum ranges over all $m!$ permutations and $\operatorname{sgn}(\sigma) = \pm 1$ depending on how many inversions do we do trough the permutations. For $m = 2$: $\det \begin{pmatrix} a & b \\ c & d \end{pmatrix} = ad - bc$, matching the two permutations of $\{1, 2\}$.

3 Polynomial Identity Testing

We now put the tools from the previous section to work: we state the PIT problem precisely, build up the root bound Theorem and its multivariate generalization that is the Schwartz–Zippel lemma and then we provide the randomized algorithm itself.

3.1 The PIT Problem

We are given an input formula describing a polynomial $p(x_1, \ldots, x_n)$ in n variables and the algorithmic task is to check if p is a zero polynomial (Section 2.1).

Example (trivial).

$$a^2 + b^2 + 2ab - (a+b)^2 = a^2 + b^2 + 2ab - a^2 - 2ab - b^2 = 0.$$

This is clearly a zero polynomial, but even here it took a small algebraic manipulation to see it.

Example (non-trivial).

$$\Big(x_1 + \cdots + x_{100}\Big)^{100} - \Big(x_1^5 + x_2^5 + \cdots + x_{100}^5\Big)^{20}.$$

This formula is in 100 variables, given in compact form. Expanded into a sum of monomials, this would be enormous.

Claim 3. *In* $\mathbb{F}_5$, $(x+y)^5 = x^5 + y^5$.

Proof. Expanding, $(x+y)^5 = x^5 + \binom{5}{1}x^4y + \binom{5}{2}x^3y^2 + \binom{5}{3}x^2y^3 + \binom{5}{4}xy^4 + y^5$. By the Fact in Section 2.3, all four middle coefficients are divisible by 5, hence $\equiv 0$ in $\mathbb{F}_5$, leaving $x^5 + y^5$. □

More generally, for any prime p, in characteristic p: $(x+y)^p = x^p + y^p$. Repeating the argument pairwise,

$$(x_1 + \cdots + x_m)^5 = x_1^5 + \cdots + x_m^5 \quad \text{in } \mathbb{F}_5.$$

If we go back to the example: over $\mathbb{F}_5$,

$$\Big(x_1 + \cdots + x_{100}\Big)^5 = x_1^5 + \cdots + x_{100}^5,$$

and raising both sides to the 20th power,

$$\Big(x_1 + \cdots + x_{100}\Big)^{100} = \Big(x_1^5 + \cdots + x_{100}^5\Big)^{20} \quad \text{in } \mathbb{F}_5.$$

So the difference is 0 as a polynomial over $\mathbb{F}_5$, but bear in mind that over $\mathbb{Q}$ it is not a zero polynomial since there the coefficients don't reduce to zero.

Remark: Over $\mathbb{R}$, a nonzero polynomial has finitely many roots ($x^2 - 4 = 0$ has only $x = \pm 2$). Over finite fields this can fail: there are nonzero polynomials that vanish at *every* field element.

Example. $p(x) = x^5 - x$ over $\mathbb{F}_5$: coefficients $1, 0, 0, 0, -1$, not all zero, yet p vanishes on every element of $\mathbb{F}_5$. Therefore for prime p and integer x, $x^p \equiv x \pmod{p}$.

This is why later statements need an explicit degree bound d together with $|S|$ large enough relative to d since picking S too small relative to the degree risks exactly this trap.

To check by full expansion that p is a zero polynomial we'd have to enumerate every possible monomial. A degree-d polynomial in m variables has up to $\binom{m+d}{d}$ monomials; for $m = d = 100$, $\binom{200}{100} \approx 9 \times 10^{58}$. This is completely inefficient therefore we need a way to test this without expanding the formula.

The whole point is therefore to plug in specific numbers and compute the resulting number. E.g. $(a+b)^2 - a^2 - 2ab - b^2$ at $a = 3, b = 5$ gives $64 - 64 = 0$.

- If p is truly zero, every choice of numbers gives 0.
- If p is nonzero, most choices give something nonzero.

The problem has a deterministic answer, but deterministic (full expansion) is far too slow, that is why we use randomness. Many other problems reduce to PIT and get solved the same way: **plug in random numbers, check the result**.

3.2 The Root Bound

Lemma 4 (Root bound)**.** *Let $\mathbb{F}$ be a field and $p(x) = a_0 + a_1x + \cdots + a_dx^d \in \mathbb{F}[x]$ of degree at most d. If p has $d+1$ distinct roots, then p is the zero polynomial.*

What we are saying here is that a nonzero polynomial of degree at most d has at most d roots. This generalizes the fact that a parabola crosses the x-axis at most twice, to any degree and any field. Also what we are saying is that if we have a polynomial of degree d with $d+1$ zeros then it is a zero polynomial.

Example. $p(x) = x^2 - 5x + 6 = (x-2)(x-3)$, degree 2, exactly 2 roots.

Proof. If a_1 is a root, then $x - a_1$ divides $p(x)$ with no remainder: $p(x) = (x-a_1)q_1(x)$, $\deg q_1 = d-1$.

Suppose p has $d+1$ distinct roots $a_1, \ldots, a_{d+1}$.

1. Factor out a_1: $p(x) = (x - a_1)q_1(x)$.
2. $0 = p(a_2) = (a_2 - a_1)q_1(a_2)$; since $a_2 \neq a_1$ and we're in a field, $q_1(a_2) = 0$. Factor: $q_1(x) = (x - a_2)q_2(x)$.
3. Repeat for $a_3, \ldots, a_{d+1}$, giving

$$p(x) = (x - a_1)(x - a_2)\cdots(x - a_{d+1}) \cdot q'(x)$$

   for some q'.

The product $(x - a_1)\cdots(x - a_{d+1})$ already has degree $d+1$, so unless $q' \equiv 0$ the total degree exceeds d, contradicting $\deg p \leq d$. Hence $q' \equiv 0$ and $p \equiv 0$. □

We assume that the polynomial p is a non zero polynomial of degree at most d therefore by Root Bound Theorem it ha at most d roots. If we pick a set $S \subseteq \mathbb{F}$ the "bad" outcomes (landing exactly on a root and wrongly seeing 0) are at most d of a total of $|S|$ choices.

Corollary 5. *If p is a nonzero polynomial of degree at most d and $S \subseteq \mathbb{F}$ is finite, then out of $|S|$ possible values, at most d can be roots, so*

$$\mathbb{P}[p(\textit{random } x) = 0] \leq \frac{d}{|S|}.$$

3.3 The Schwartz–Zippel Lemma

This is just the multivariate generalization of te Root Bound theorem: for a nonzero polynomial in many variables with degree at most d if we pick each variable indipendently and uniformly from a set S the chance that the whole thing evaluates to 0 is at most $\frac{d}{|S|}$

Lemma 6 (Schwartz–Zippel [1, 2])**.** *Let $p(x_1, \ldots, x_m) \in \mathbb{F}[x_1, \ldots, x_m]$ have total degree at most d. If $x_1, \ldots, x_m$ are drawn independently and uniformly from a finite subset $S \subseteq \mathbb{F}$, then*

$$\mathbb{P}\left[p(x_1, \ldots, x_m) = 0\right] \leq \frac{d}{|S|}.$$

For $p(x, y) = x^2 + y^2 - 1$ ($d = 2$), the zero set over $\mathbb{R}$ is the unit circle. On a grid of candidate points, most miss the circle, so hitting it by chance is unlikely which is consistent with $d/|S|$ being small when $|S| \gg d$.

The bound $d/|S|$ holds for any finite S, regardless of p and cannot is a tight bound. Take $a_1, \ldots, a_d \in S$ distinct and

$$p(x_1, \ldots, x_m) = (x_1 - a_1)(x_1 - a_2) \cdots (x_1 - a_d),$$

degree exactly d. Then $p(X_1, \ldots, X_m) = 0$ iff $X_1 \in \{a_1, \ldots, a_d\}$, probability exactly $d/|S|$.

Take

$$P = \Big(x_1 + \cdots + x_{100}\Big)^{100} - \Big(x_1^5 + \cdots + x_{100}^5\Big)^{20},$$

degree at most $d = 100$, and pick $S = \{0, \ldots, 199\}$, $|S| = 200 = 2d$.

- If $P \equiv 0$, plugging in any numbers gives exactly 0: the algorithm never makes a mistake here.
- If $P \not\equiv 0$, Schwartz–Zippel gives $\mathbb{P}[P = 0] \leq \frac{d}{|S|} \leq \frac{1}{2}$.

So a single evaluation of a nonzero P still has up to 50% chance of hitting a root by accident (a false negative). We can only conclude $P \equiv 0$ if *every* trial gives 0: running K independent trials, the chance of being unlucky K times in a row is $(1/2)^K$, decaying exponentially.

Proof. Induction on the number m of variables.

Base case $m = 1$: this is exactly the root bound.

Inductive step. Group the terms of p by powers of x_m:

$$p(x_1, \ldots, x_m) = \sum_{i=0}^{k} p_i(x_1, \ldots, x_{m-1})\, x_m^i,$$

with k maximal such that $p_k \not\equiv 0$. Since p has total degree $\leq d$, each p_i satisfies $\deg(p_i) \leq d - i$ (the term $p_i x_m^i$ contributes degree $\deg(p_i) + i$ to p).

Illustration. $p(x_1, x_2, x_3) = x_1^2 x_2 + x_1^2 x_3 + x_1 x_2 x_3 + x_3$, degree 3. Grouping by powers of x_1:

$$p = x_1^2(x_2 + x_3) + x_1(x_2 x_3) + x_3,$$

so $k = 2$, $p_2 = x_2 + x_3$ (degree $1 \leq d - k = 1$), $p_1 = x_2 x_3$, $p_0 = x_3$.

Now fix $a_1, \ldots, a_{m-1} \in S$ for the first $m - 1$ variables.

1. p_k has $m - 1$ variables and degree $\leq d - k$, so by induction $\mathbb{P}[p_k(a_1, \ldots, a_{m-1}) = 0] \leq (d - k)/|S|$.

2. If instead $p_k(a_1, \ldots, a_{m-1}) \neq 0$, then $q(x_m) := p(a_1, \ldots, a_{m-1}, x_m)$ is univariate of degree exactly k (its top coefficient is $p_k(a_1, \ldots, a_{m-1}) \neq 0$), so by the root bound $\mathbb{P}[q(x_m) = 0] \leq k/|S|$.

Splitting on whether $p_k(a_1, \ldots, a_{m-1}) = 0$:

$$\mathbb{P}[p = 0] \leq \frac{d-k}{|S|} + \frac{k}{|S|} = \frac{d}{|S|}. \quad \square$$

3.4 A Randomized Algorithm for PIT

Schwartz–Zippel needs a degree bound to set $|S| \geq 2d$. This is computed recursively from the formula's structure, no expansion needed:

Operation	**Resulting degree**
constant	0
variable x_i	1
$f + g$ or $f - g$	$\max\{\deg f, \deg g\}$
$f \cdot g$	$\deg f + \deg g$
f^m	$m \cdot \deg f$

1. Compute an upper bound d on the resulting degree.
2. Choose a finite $S \subseteq \mathbb{F}$ with $|S| \geq 2d$.
3. Sample $(a_1, \ldots, a_m) \in S^m$ uniformly at random.
4. Evaluate the formula gate by gate at this point.
5. Output NonZero if the result is $\neq 0$; otherwise output probably Zero.

Observation 7. *This algorithm makes errors only in one direction. If $p \equiv 0$, the algorithm always outputs 0: it never falsely reports* NonZero*. If $p \not\equiv 0$, Schwartz–Zippel gives $\mathbb{P}[p(\text{random point}) = 0] \leq d/|S| \leq 1/2$, so a single run can still return 0 by accident.*

We conclude $p \equiv 0$ only if every trial returns 0. Over K independent trials, the chance of missing a genuine root K times in a row is $(1/2)^K$: the error decays exponentially in K.

4 Applications

We now apply the randomized PIT algorithm to a problem that seems unrelated to it: deciding whether a graph has a perfect matching, first for bipartite graphs and then for general graphs

4.1 Bipartite Perfect Matching

Let $G = (L \cup R, E)$, $L = \{l_1, \ldots, l_m\}$, $R = \{r_1, \ldots, r_m\}$, edges only between L and R.

Definition 8 (Perfect matching)**.** *A perfect matching contains m pairwise vertex-disjoint edges covering every vertex.*

We want a systematic, fast algorithm, so we reduce this task as a PIT.

Following Edmonds [3], give every possible edge (l_i, r_j) its own variable x_{ij} and set

$$M_{ij} = \begin{cases} x_{ij}, & (l_i, r_j) \in E, \\ 0, & \text{otherwise.} \end{cases}$$

Example. Edges $l_1{-}r_1,\ l_1{-}r_2,\ l_2{-}r_1,\ l_3{-}r_2,\ l_3{-}r_3$:

$$M = \begin{pmatrix} x_{11} & x_{12} & 0 \\ x_{21} & 0 & 0 \\ 0 & x_{32} & x_{33} \end{pmatrix}.$$

Using the determinant formula from Section 2.4, each term $\prod_i M_{i,\sigma(i)}$ picks one entry per row and column, using edges $(l_i, r_{\sigma(i)})$; it's nonzero exactly when all these edges exist.

Worked example. Checking all 6 permutations of $\{1, 2, 3\}$ on the matrix above:

σ	product	
$(1,2,3)$	$M_{11}M_{22}M_{33}$	$M_{22} = 0$
$(1,3,2)$	$M_{11}M_{23}M_{32}$	$M_{23} = 0$
$(2,1,3)$	$M_{12}M_{21}M_{33}$	survives, sign $-$
$(2,3,1)$	$M_{12}M_{23}M_{31}$	$M_{23} = M_{31} = 0$
$(3,1,2)$	$M_{13}M_{21}M_{32}$	$M_{13} = 0$
$(3,2,1)$	$M_{13}M_{22}M_{31}$	$M_{13} = 0$

Only $\sigma = (2,1,3)$ survives: $\det(M) = -x_{12}x_{21}x_{33} \neq 0$, matching $\{l_1{-}r_2,\ l_2{-}r_1,\ l_3{-}r_3\}$.

Theorem 9. $\det(M)$*, as a formal polynomial, is zero iff G has no perfect matching.*

σ survives (contributes a nonzero term) iff every edge it proposes exists in G, so surviving permutations correspond exactly to perfect matchings. $\det(M) = 0$ as a polynomial iff no permutation survives iff G has no matching. If a matching exists, it contributes a monomial in variables no other surviving permutation touches, so it cannot cancel: $\det(M) \not\equiv 0$.

Symbolically computing $\det(M)$ would sum up to $m!$ permutations — exponential again. Same fix as before:

1. Choose a prime $p > 2m$, work over $\mathbb{F}_p$.

2. Sample $X_{ij} \in \mathbb{F}_p$ for each edge, getting numerical M'.
3. Compute $\det(M')$ via Gaussian elimination, $O(m^3)$.
4. $\det(M') \neq 0 \Rightarrow$ matching exists, certainly. If $\det(M') = 0$: since $\det(M)$ has total degree m, when a matching exists,
$$\mathbb{P}[\det(M') = 0] \leq m/p < 1/2,$$
and repeated trials push this down exponentially.

4.2 Perfect Matching in General Graphs: The Tutte Matrix

Non-bipartite graphs have a much more involved deterministic algorithm (Edmonds' blossom algorithm). We push the same PIT reduction further instead.

Try $A_{ij} = A_{ji} = x_{ij}$ for $\{i, j\} \in E$, $A_{ii} = 0$.

Counter-example: the triangle. On $\{1, 2, 3\}$ with all three edges,

$$A = \begin{pmatrix} 0 & x_{12} & x_{13} \\ x_{12} & 0 & x_{23} \\ x_{13} & x_{23} & 0 \end{pmatrix}, \qquad \det(A) = 2x_{12}x_{13}x_{23} \not\equiv 0$$

in characteristic $\neq 2$. But 3 vertices can't have a perfect matching therefore we have a false positive.

The determinant sum has two surviving terms, corresponding to the two orientations of the triangle. With a symmetric matrix both give the same monomial with the same sign, so they add instead of cancelling. We need the two orientations of every odd cycle to cancel; that requires antisymmetry.

Definition 10 (Tutte matrix [4])**.** *For $G = (V, E)$, $V = \{1, \ldots, n\}$, one variable x_{ij} per edge $\{i, j\}$ with $i < j$:*

$$T_{ij} = \begin{cases} x_{ij}, & i < j,\ \{i, j\} \in E, \\ -x_{ji}, & i > j,\ \{i, j\} \in E, \\ 0, & \textit{otherwise.} \end{cases}$$

$T_{ji} = -T_{ij}$, $T_{ii} = 0$: T is skew-symmetric. (Characteristic $\neq 2$; later $p > 2n$.)

Example. Edges $\{1, 2\}, \{2, 3\}, \{2, 4\}, \{3, 4\}$, matching $\{1, 2\}, \{3, 4\}$:

$$T = \begin{pmatrix} 0 & x_{12} & 0 & 0 \\ -x_{12} & 0 & x_{23} & x_{24} \\ 0 & -x_{23} & 0 & x_{34} \\ 0 & -x_{24} & -x_{34} & 0 \end{pmatrix}, \qquad \det(T) = (x_{12}x_{34})^2 \neq 0.$$

Theorem 11 (Tutte-matrix criterion)**.** $\det(T) \not\equiv 0 \iff$ *G has a perfect matching.*

Unlike the bipartite case, different permutations can now produce the same monomial, so cancellation of non-matching terms needs an actual argument.

Every $\pi \in S_n$ splits into disjoint directed cycles. The term

$$\text{Term}(\pi) := \text{sgn}(\pi) \prod_{i=1}^{n} T_{i,\pi(i)}$$

is nonzero iff $\{i, \pi(i)\} \in E$ for all i ($T_{ii} = 0$ rules out fixed points), so nonzero terms correspond to spanning collections of vertex-disjoint cycles.

A perfect matching turns into a permutation by swapping each matched pair. Its term in the determinant uses both directions of every matched edge: the two minus signs (one from the skew-symmetry, one from each swap being an odd permutation) cancel out, so the term survives as a positive product of squared edge-variables. No other permutation can produce that exact monomial, so it can never be cancelled: a matching guarantees $\det(T) \not\equiv 0$.

If G has no perfect matching, every surviving permutation must contain an odd cycle (otherwise its cycles would themselves give a perfect matching). Pair up permutations by reversing the direction of just one odd cycle each time: this keeps the sign the same, but flips every edge inside that cycle and since an odd cycle has an odd number of edges, the term's sign flips overall. So each permutation and its reversed partner cancel exactly. This pairs up everything with nothing left over, so the whole determinant is 0.

Same procedure as before: plug in random numbers from $\mathbb{F}_p$ (with $p > 2n$) for every edge variable, and compute the determinant numerically via Gaussian elimination. If there's no matching, this is always 0. If there is one, Schwartz–Zippel says we get 0 by accident with probability under $1/2$, so repeating the trial a few times makes the algorithm reliable.

5 Summary

1. A nonzero degree-d polynomial has at most d roots; Schwartz–Zippel extends this to many variables.
2. Plugging in random numbers gives a fast identity test that only ever makes errors in one way and repeating it pushes that error down fast.
3. Determinants can encode graph matchings as polynomials, turning matching into a PIT problem: Edmonds matrix for bipartite graphs, Tutte matrix for general graphs.
4. For general graphs, everything that isn't a real matching cancels out in pairs, leaving only genuine matchings behind.

Bibliography.

References

[1] Jacob T. Schwartz. Fast Probabilistic Algorithms for Verification of Polynomial Identities. *J. ACM*, 27(4):701–717, 1980.

[2] Richard Zippel. Probabilistic Algorithms for Sparse Polynomials. In *Symbolic and Algebraic Computation (EUROSAM '79)*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226. Springer, 1979.

[3] Jack Edmonds. Paths, Trees, and Flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965.

[4] W. T. Tutte. The Factorization of Linear Graphs. *Journal of the London Mathematical Society*, 22:107–111, 1947.