Polynomial Identity Testing

Random evaluation, the Schwartz–Zippel lemma, and applications to perfect matching

How can randomness reveal an algebraic identity?

The basic question

Polynomial Identity Testing (PIT)

Given an arithmetic formula representing

$$p(x_1, \ldots, x_n) \in F[x_1, \ldots, x_n],$$

decide whether p is the zero polynomial.

For example, consider the compact formula

$$\Phi(x_1, \ldots, x_{100}) = \left(\sum_{i=1}^{100} x_i\right)^{100} - \sum_{i=1}^{100} x_i^{100}.$$

- Is every coefficient of Φ zero?
- Does the answer depend on the field?
- Can we decide this without expanding the formula?

Expansion gives a deterministic—but enormous—test

Every polynomial of total degree at most d has the form

$$p(x_1, \ldots, x_n) = \sum_{\substack{\alpha \in \mathbb{Z}^n_{\geq 0} \\ |\alpha| \leq d}} c_\alpha x^\alpha, \qquad x^\alpha = \prod_{i=1}^{n} x_i^{\alpha_i}, \qquad |\alpha| = \sum_{i=1}^{n} \alpha_i.$$

It is the zero polynomial exactly when

$$c_\alpha = 0 \quad \text{for every } \alpha.$$

The obstacle

The number of possible monomials is

$$\binom{n+d}{d}.$$

Characteristic matters—and so does the exponent

The expression with exponent 100 is *not* zero over $\mathbb{F}_5$. Indeed, setting $x_1 = x_2 = 1$ and all other variables to zero gives

$$2^{100} - 1 - 1 \equiv 1 - 2 \equiv -1 \not\equiv 0 \pmod 5.$$

A correct characteristic-5 example is

$$\Psi(x_1, \ldots, x_{100}) = \left(\sum_{i=1}^{100} x_i\right)^5 - \sum_{i=1}^{100} x_i^5.$$

Over $\mathbb{Q}$

The coefficient of $x_1^4 x_2$ is

$$\binom{5}{1} = 5 \neq 0.$$

Over $\mathbb{F}_5$

The Frobenius identity gives

$$(a+b)^5 = a^5 + b^5.$$

The computational goal

We work with arithmetic formulas built from

$$\text{variables}, \quad \text{field constants}, \quad +, \quad -, \quad \times, \quad \text{and possibly powers.}$$

Desired algorithm

Decide whether the polynomial represented by the formula is identically zero in time polynomial in the formula size, the number of variables, and a degree bound.

- Expanding all coefficients is deterministic but can take exponential time.
- No deterministic polynomial-time algorithm is known for unrestricted arithmetic formulas or circuits.
- A remarkably simple randomized polynomial-time test is available.

Part I

Roots of polynomials

A zero polynomial is not merely a zero function

Zero polynomial

$$p(x) = a_0 + a_1 x + \cdots + a_d x^d$$

is the zero polynomial when

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

This is stronger than requiring

$$p(a) = 0 \qquad \text{for every } a \in F.$$

Finite-field example

Over $\mathbb{F}_q$, the polynomial

$$p(x) = x^q - x$$

A nonzero univariate polynomial has few roots

Root bound

Let F be a field, and let

$$p(x) = a_0 + a_1 x + \cdots + a_d x^d \in F[x]$$

have degree at most d. If p vanishes at $d+1$ distinct field elements, then p is the zero polynomial.

Equivalently, every nonzero polynomial of degree at most d has at most d distinct roots.

The field assumption is important: it lets us cancel every nonzero difference $a_i - a_j$.

Proof of the root bound: repeatedly factor out roots

Suppose that

$$p(a_1) = p(a_2) = \cdots = p(a_{d+1}) = 0$$

for distinct $a_1, \ldots, a_{d+1} \in F$.
By the factor theorem,

$$p(x) = (x - a_1)q_1(x).$$

For every $j > 1$,

$$0 = p(a_j) = (a_j - a_1)q_1(a_j).$$

Because $a_j - a_1 \neq 0$ and F is a field,

$$q_1(a_j) = 0.$$

Thus $a_2, \ldots, a_{d+1}$ are roots of q_1, so the same argument may be repeated.

Proof of the root bound: the degree contradiction

Repeating the factorization yields

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

for some $q_{d+1}(x) \in F[x]$.
If p were nonzero, then q_{d+1} would also be nonzero, and

$$\deg p = (d + 1) + \deg q_{d+1} \geq d + 1.$$

This contradicts $\deg p \leq d$. Therefore

$$\boxed{p \equiv 0.}$$

Immediate probabilistic consequence

If X is uniform on a finite set $S \subseteq F$, then for every nonzero p of degree at most d,

Part II

The Schwartz–Zippel lemma

Random evaluation rarely hits a root

DeMillo–Lipton–Schwartz–Zippel lemma

Let

$$p(x_1, \ldots, x_n) \in F[x_1, \ldots, x_n]$$

be a nonzero polynomial of total degree at most d. Let $S \subseteq F$ be finite, and choose

$$X_1, \ldots, X_n \quad \text{independently and uniformly from } S.$$

Then

$$\boxed{\Pr[p(X_1, \ldots, X_n) = 0] \leq \frac{d}{|S|}.}$$

The number of variables does not appear in the bound; only the *total degree* and the size of the sampling set matter.

The Schwartz–Zippel bound is tight

Assume $d \leq |S|$, and choose distinct $a_1, \ldots, a_d \in S$. Consider the univariate polynomial

$$p(x_1, \ldots, x_n) = \prod_{j=1}^{d} (x_1 - a_j).$$

This polynomial is nonzero and has total degree d. Moreover,

$$p(X_1, \ldots, X_n) = 0 \quad \Longleftrightarrow \quad X_1 \in \{a_1, \ldots, a_d\}.$$

Hence

$$\Pr[p(X_1, \ldots, X_n) = 0] = \frac{d}{|S|}.$$

Conclusion

Without additional assumptions about p, the factor $d/|S|$ cannot be improved.

Proof by induction: isolate the final variable

We prove the lemma by induction on the number n of variables.

- For $n = 1$, the claim is exactly the univariate root bound.
- Assume the result holds for polynomials in $n - 1$ variables.

Write p as a polynomial in x_n:

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

where $p_k \not\equiv 0$ and k is maximal.
Since p has total degree at most d,

$$\deg p_i \leq d - i, \qquad \text{and in particular} \qquad \deg p_k \leq d - k.$$

The first way to fail: the leading coefficient vanishes

Let

$$\boldsymbol{X}' = (X_1, \ldots, X_{n-1})$$

and define the bad event

$$B = \{p_k(\boldsymbol{X}') = 0\}.$$

The polynomial p_k is nonzero and has degree at most $d - k$. By the induction hypothesis,

$$\Pr[B] \leq \frac{d-k}{|S|}.$$

Interpretation

Event B is the event that, after fixing the first $n-1$ variables, the apparent degree in x_n drops below k.

Outside the bad event, only k choices can be roots

Fix any

$$\boldsymbol{a} = (a_1, \ldots, a_{n-1}) \in S^{n-1}$$

such that $p_k(\boldsymbol{a}) \neq 0$.
Then

$$q_{\boldsymbol{a}}(t) := p(a_1, \ldots, a_{n-1}, t)$$

is a nonzero univariate polynomial of degree exactly k.
Therefore $q_{\boldsymbol{a}}$ has at most k roots in S, so

$$\Pr[p(\boldsymbol{a}, X_n) = 0 \mid \boldsymbol{X}' = \boldsymbol{a}] \leq \frac{k}{|S|}.$$

This bound holds for every fixed $\boldsymbol{a}$ outside B.

Combining the two possible causes of failure

Split according to whether B occurs:

$$\begin{aligned}\Pr[p(\boldsymbol{X}', X_n) = 0] &= \Pr[p = 0 \ \wedge \ B] + \Pr[p = 0 \ \wedge \ B^c] \\ &\leq \Pr[B] + \Pr[p = 0 \ \wedge \ B^c].\end{aligned}$$

Averaging the previous conditional bound over all $\boldsymbol{a} \notin B$ gives

$$\Pr[p = 0 \ \wedge \ B^c] \leq \frac{k}{|S|}.$$

Consequently,

$$\Pr[p(\boldsymbol{X}', X_n) = 0] \leq \frac{d-k}{|S|} + \frac{k}{|S|} = \boxed{\frac{d}{|S|}}.$$

This completes the induction.

Part III

A randomized PIT algorithm

Polynomial identity testing by one random evaluation

Suppose a formula computes p and we know $\deg p \leq d$.

One trial

1. Choose a finite $S \subseteq F$ with $|S| \geq 2d$.
2. Sample $X_1, \ldots, X_n$ independently and uniformly from S.
3. Evaluate the original formula at $(X_1, \ldots, X_n)$.
4. If the value is nonzero, output NONZERO; otherwise output ZERO.

No coefficient expansion is performed: the formula is evaluated gate by gate.

A syntactic degree bound is enough

A degree upper bound can be computed recursively from the formula:

Gate	Syntactic degree
field constant	0
x_i	1
$f + g$ or $f - g$	$\max\{\deg f, \deg g\}$
$f \cdot g$	$\deg f + \deg g$
f^m	$m \deg f$

Cancellation may make the true degree smaller, but an upper bound is all Schwartz–Zippel requires.

Degree 100

For a degree-100 formula, sampling from any set S of size 200 gives error probability at most

$$\frac{100}{200} = \frac{1}{2}.$$

The error is one-sided

If $p \equiv 0$

Every evaluation is zero.

$$\Pr[\text{output NONZERO}] = 0.$$

The algorithm never falsely claims that a zero polynomial is nonzero.

If $p \not\equiv 0$

A random evaluation may accidentally hit a root.

$$\Pr[\text{output ZERO}] \leq \frac{d}{|S|} \leq \frac{1}{2}.$$

A single nonzero evaluation is a definitive certificate that the polynomial is nonzero.

Independent repetition amplifies success probability

Repeat the entire test independently K times.

- Output NONZERO as soon as any evaluation is nonzero.
- Output ZERO only if all K evaluations are zero.

If $p \not\equiv 0$, independence gives

$$\Pr[\text{all } K \text{ trials miss}] \leq \left(\frac{d}{|S|}\right)^K.$$

For $|S| \geq 2d$,

$$\Pr[\text{error}] \leq 2^{-K}.$$

In particular,

$$2^{-100} < 8 \times 10^{-31}.$$

The algorithm still has one-sided error after amplification.

What if the field is too small?

The sampling set must satisfy $S \subseteq F$, so

$$|S| \leq |F|.$$

If $|F| < 2d$, the base field does not contain a sufficiently large sampling set.

Standard remedy

Evaluate the same polynomial over an extension field

$$E/F \qquad \text{with} \qquad |E| \geq 2d.$$

The coefficients of p embed naturally into E, and a polynomial that is nonzero over F remains nonzero over E.

For the applications below, we instead choose a prime field large enough from the outset.

Part IV

Bipartite perfect matching

Perfect matchings select one edge at every vertex

Let

$$G = (L \cup R, E), \qquad L = \{\ell_1, \ldots, \ell_n\}, \qquad R = \{r_1, \ldots, r_n\}.$$

A perfect matching contains n pairwise vertex-disjoint edges and covers every vertex.

Encode the graph by a matrix of variables

Give every possible edge (ℓ_i, r_j) its own variable x_{ij}. Define the Edmonds matrix

$$T_{ij} = \begin{cases} x_{ij}, & (\ell_i, r_j) \in E, \\ 0, & (\ell_i, r_j) \notin E. \end{cases}$$

For the preceding graph, the zero pattern is

$$T = \begin{pmatrix} x_{11} & x_{12} & x_{13} & 0 \\ x_{21} & 0 & x_{23} & 0 \\ 0 & x_{32} & 0 & x_{34} \\ x_{41} & x_{42} & x_{43} & 0 \end{pmatrix}.$$

The determinant is a polynomial of total degree n:

$$\det T = \sum_{\pi \in S_n} \operatorname{sgn}(\pi) \prod_{i=1}^{n} T_{i,\pi(i)}.$$

A determinant term selects one edge per row and column

Fix a permutation $\pi \in S_n$. Its determinant term is

$$\operatorname{sgn}(\pi) T_{1,\pi(1)} T_{2,\pi(2)} \cdots T_{n,\pi(n)}.$$

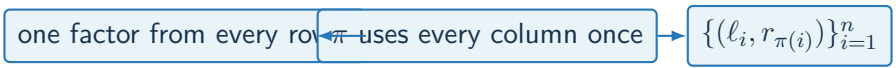

If every selected entry is nonzero,
every vertex has exactly one selected edge.

Thus a nonzero determinant term is exactly a perfect matching, together with the sign of its permutation.

Why the determinant is nonzero exactly when a matching exists

If a nonzero term exists

Suppose

$$\prod_{i=1}^{n} T_{i,\pi(i)} \neq 0.$$

Then $(\ell_i, r_{\pi(i)}) \in E$ for every i. Since π is bijective, these edges cover every vertex exactly once. They form a perfect matching.

If a perfect matching exists

A perfect matching has the form

$$M = \{(\ell_i, r_{\pi(i)}) : i = 1, \ldots, n\}$$

The randomized bipartite matching algorithm

Choose a prime $p > 2n$ and work over $\mathbb{F}_p$.

One trial

1. For every edge $(\ell_i, r_j) \in E$, choose $X_{ij} \in \mathbb{F}_p$ independently and uniformly.
2. Replace x_{ij} by X_{ij}, obtaining a numerical matrix T'.
3. Compute $\det T'$, for example by Gaussian elimination.

- If $\det T' \neq 0$, the graph definitely has a perfect matching.
- If no perfect matching exists, $\det T \equiv 0$, so every trial returns zero.
- If a matching exists, Schwartz–Zippel gives

$$\Pr[\det T' = 0] \leq \frac{n}{p} < \frac{1}{2}.$$

Each trial uses $O(n^3)$ field operations; repetition reduces the false-negative probability exponentially.

Part V

Perfect matching in general graphs

A symmetric adjacency matrix does not work

For a non-bipartite graph, one might try a symmetric matrix with $A_{ij} = A_{ji} = x_{ij}$.
Consider a triangle:

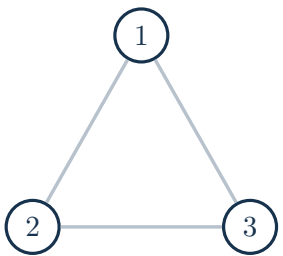

$$A = \begin{pmatrix} 0 & x_{12} & x_{13} \\ x_{12} & 0 & x_{23} \\ x_{13} & x_{23} & 0 \end{pmatrix},$$

$$\det A = 2x_{12}x_{13}x_{23} \not\equiv 0$$

over fields of characteristic not 2.

The triangle has no perfect matching. Its two directed orientations contribute the same monomial with the same sign.
We need the two orientations of every odd cycle to cancel.

The Tutte matrix introduces antisymmetric signs

Let $G = (V, E)$, where $V = \{1, \dots, n\}$. For every edge $\{i, j\} \in E$ with $i < j$, introduce one variable x_{ij}.
Define the Tutte matrix T by

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

Thus

$$T_{ji} = -T_{ij}, \qquad T_{ii} = 0, \qquad T^{\mathsf{T}} = -T.$$

We work over a field of characteristic different from 2; later we choose a prime $p > 2n$.

The theorem we want

Tutte-matrix criterion

For a graph G,

$$\det T \not\equiv 0 \quad \Longleftrightarrow \quad G \text{ has a perfect matching.}$$

Once this is proved, the algorithm is immediate:

1. Randomly evaluate the variables x_{ij}.
2. Compute the determinant of the evaluated matrix.
3. Apply Schwartz–Zippel and independent repetition.

The main issue is cancellation: several permutations may now produce the same monomial.

A permutation decomposes into disjoint directed cycles

For example,

$$\pi = (1\,4\,3)(2\,5)(6\,8\,9\,7).$$

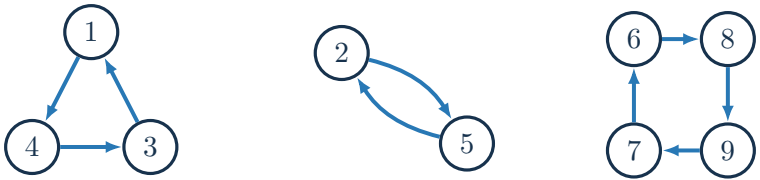

Every vertex appears in exactly one cycle, because every vertex has one image and one preimage under a permutation.

Nonzero determinant terms are spanning cycle covers

The term indexed by π is

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

This term is nonzero exactly when

$$\{i, \pi(i)\} \in E \qquad \text{for every } i.$$

Therefore a nonzero term corresponds to a spanning collection of vertex-disjoint directed cycles in G.

- A cycle of length at least 3 becomes an ordinary undirected cycle after directions are forgotten.
- A 2-cycle $i \to j \to i$ uses one undirected edge $\{i, j\}$ in both directions.
- Fixed points cannot occur because $T_{ii} = 0$.

A perfect matching gives a permutation of 2-cycles

Suppose

$$M = \{\{u_1, v_1\}, \ldots, \{u_m, v_m\}\}, \qquad n = 2m,$$

is a perfect matching.
Define

$$\pi_M = (u_1\, v_1)(u_2\, v_2)\cdots(u_m\, v_m).$$

$$\pi_M = (1\,6)(2\,4)(3\,5)$$

Every matching edge is traversed once in each direction, so its variable appears squared.

The matching monomial has positive coefficient

Write each matching edge as $\{u_r, v_r\}$ with $u_r < v_r$. Then

$$T_{u_r,v_r} T_{v_r,u_r} = x_{u_r v_r}(-x_{u_r v_r}) = -x_{u_r v_r}^2.$$

Since π_M consists of m transpositions,

$$\operatorname{sgn}(\pi_M) = (-1)^m.$$

Hence

$$\begin{aligned}
\operatorname{Term}(\pi_M) &= \operatorname{sgn}(\pi_M) \prod_{r=1}^{m} T_{u_r,v_r} T_{v_r,u_r} \\
&= (-1)^m (-1)^m \prod_{r=1}^{m} x_{u_r v_r}^2 \\
&= \prod_{\{i,j\} \in M} x_{ij}^2.
\end{aligned}$$

No other permutation produces that matching monomial

Consider

$$m_M = \prod_{\{i,j\} \in M} x_{ij}^2.$$

If a determinant term produces m_M, then for every $\{i, j\} \in M$,

- the variable x_{ij} must occur twice;
- the only entries containing it are T_{ij} and T_{ji};
- therefore the permutation must satisfy

$$\pi(i) = j, \qquad \pi(j) = i.$$

This forces $\pi = \pi_M$.

Thus m_M appears exactly once, with coefficient $+1$, and cannot be canceled.

Therefore

$$G \text{ has a perfect matching} \quad \Longrightarrow \quad \det T \not\equiv 0.$$

Without a matching, every cycle cover has an odd cycle

Assume G has no perfect matching, and take any $\pi \in \mathcal{P}$.
Suppose every cycle in the decomposition of π had even length. For a cycle

$$C = (v_1\, v_2\, \ldots\, v_{2r}),$$

select the alternating edges

$$\{v_1, v_2\}, \{v_3, v_4\}, \ldots, \{v_{2r-1}, v_{2r}\}.$$

These edges form a perfect matching of the vertices of C. For a 2-cycle, this simply selects its single underlying edge.
Since the cycles are vertex-disjoint and cover every vertex, taking these alternating matchings over all cycles would produce a perfect matching of G, a contradiction.
Therefore every $\pi \in \mathcal{P}$ contains at least one odd cycle, necessarily of length at least 3.

Every even cycle contains an alternating matching

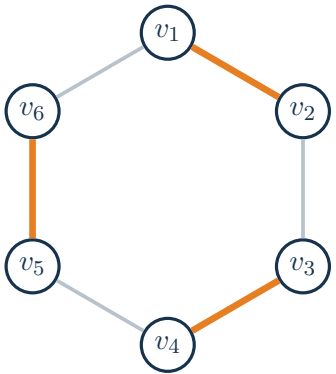

selected edges:
$\{v_1, v_2\}$,
$\{v_3, v_4\}$,
$\{v_5, v_6\}$

If all components of a spanning cycle cover are even, their alternating matchings combine into a perfect matching of the entire graph.

Define a canonical odd-cycle reversal

For each $\pi \in \mathcal{P}$:

1. Find the smallest-labelled vertex that lies on an odd cycle.
2. Let C be the unique cycle containing this vertex.
3. Reverse the direction of C, leaving every other cycle unchanged.

Denote the resulting permutation by

$$\theta(\pi).$$

If

$$C = (v_1\, v_2\, \ldots\, v_\ell), \qquad \ell \text{ odd},$$

then θ replaces it by

$$C^{-1} = (v_1\, v_\ell\, v_{\ell-1}\, \ldots\, v_2).$$

Since G is undirected, every reversed edge is still an edge of G. Thus $\theta(\pi) \in \mathcal{P}$.

Reversing an odd cycle preserves the monomial

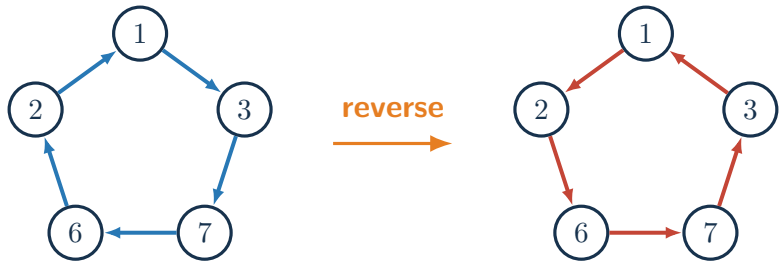

$C = (1\,3\,7\,6\,2)$ $\qquad$ $C^{-1} = (1\,2\,6\,7\,3)$

The same undirected edge variables occur in both terms; only the orientations—and therefore the matrix-entry signs—change.

Reversal does not change the permutation sign

The sign of an r-cycle is

$$(-1)^{r-1}.$$

A cycle and its reversal have the same length. More generally, if $\mathcal{C}(\pi)$ is the collection of cycles of π, then

$$\operatorname{sgn}(\pi) = \prod_{D \in \mathcal{C}(\pi)} (-1)^{|D|-1}.$$

The permutations π and $\theta(\pi)$ have exactly the same cycle lengths. Therefore

$$\boxed{\operatorname{sgn}(\theta(\pi)) = \operatorname{sgn}(\pi).}$$

The sign change will instead come from the skew-symmetry $T_{ji} = -T_{ij}$.

Reversing an odd cycle flips the matrix-entry sign

Let C be the selected odd cycle, with $|C| = \ell$. Define its contribution

$$P_C(\pi) = \prod_{v \in C} T_{v,\pi(v)}.$$

On the reversed cycle, $\theta(\pi)$ follows every edge backward. After reindexing the product,

$$\begin{aligned} P_C(\theta(\pi)) &= \prod_{v \in C} T_{v,\theta(\pi)(v)} \\ &= \prod_{v \in C} T_{\pi(v),v} \\ &= \prod_{v \in C} \left(-T_{v,\pi(v)}\right) \\ &= (-1)^{\ell} P_C(\pi) \\ &= -P_C(\pi) \end{aligned}$$

The entire determinant term changes sign

Contributions from all cycles other than C remain unchanged. Combining this with the permutation-sign calculation,

$$\begin{aligned}\operatorname{Term}(\theta(\pi)) &= \operatorname{sgn}(\theta(\pi)) \prod_{i=1}^{n} T_{i,\theta(\pi)(i)} \\ &= \operatorname{sgn}(\pi) \left(- \prod_{i=1}^{n} T_{i,\pi(i)} \right) \\ &= - \operatorname{Term}(\pi).\end{aligned}$$

Thus π and $\theta(\pi)$ contribute the same monomial with opposite coefficients.

The reversal map is a sign-reversing involution

Reversing the same cycle twice returns its original orientation:

$$\theta(\theta(\pi)) = \pi.$$

The rule selects the same cycle on the second application because reversal does not change

- the vertex set of any cycle,
- the length or parity of any cycle, or
- the smallest vertex lying on an odd cycle.

Since the selected odd cycle has length at least 3,

$$\theta(\pi) \neq \pi.$$

Therefore $\mathcal{P}$ is partitioned into disjoint pairs

$$\{\pi, \theta(\pi)\},$$

and each pair satisfies

$$\mathrm{Term}(\pi) + \mathrm{Term}(\theta(\pi)) = 0.$$

If there is no perfect matching, every term cancels

Under the assumption that G has no perfect matching, every $\pi \in \mathcal{P}$ contains an odd cycle, so θ is defined on all of $\mathcal{P}$.
Hence

$$\begin{aligned}\det T &= \sum_{\pi \in \mathcal{P}} \operatorname{Term}(\pi) \\ &= \sum_{\{\pi, \theta(\pi)\}} \left(\operatorname{Term}(\pi) + \operatorname{Term}(\theta(\pi))\right) \\ &= 0.\end{aligned}$$

Therefore

$$G \text{ has no perfect matching} \quad \Longrightarrow \quad \det T \equiv 0.$$

Together with the unique matching-monomial argument, this proves the Tutte-matrix criterion.

The randomized algorithm for general graphs

Let G have n vertices. Choose a prime $p > 2n$.

One trial

1. For every edge variable x_{ij}, independently sample $X_{ij} \in \mathbb{F}_p$.
2. Form the evaluated skew-symmetric matrix T'.
3. Compute $\det T'$ by Gaussian elimination.

Since $\det T$ has total degree at most n:

- If G has no perfect matching, then $\det T'$ is always zero.
- If G has a perfect matching, then

$$\Pr[\det T' = 0] \leq \frac{n}{p} < \frac{1}{2}.$$

- After K independent trials, the false-negative probability is less than 2^{-K}.

Each trial takes $O(n^3)$ field operations.

The full chain of ideas

1. A nonzero univariate degree-d polynomial has at most d roots.
2. Induction extends this to multivariate polynomials:

$$\Pr[p(\boldsymbol{X}) = 0] \leq \frac{d}{|S|}.$$

3. Random evaluation gives a fast one-sided-error identity test.
4. Independent repetition reduces error exponentially.
5. Determinants turn combinatorial structures into polynomials:

$$\begin{aligned} \text{Edmonds matrix} &\longleftrightarrow \text{bipartite perfect matchings,} \\ \text{Tutte matrix} &\longleftrightarrow \text{perfect matchings in general graphs.} \end{aligned}$$

6. In the non-bipartite case, odd-cycle reversal pairs and cancels every cycle cover that does not arise from a perfect matching.

Randomness turns algebraic identity into efficient computation.

Evaluate once, bound the chance of failure, and amplify.