Mathematical Approaches to Lower Bounds: Complexity of Proofs and Computation
2022-Jul-04

Multiplicities in GCT: What is in #P and what is not?

Christian Ikenmeyer

UNIVERSITY OF LIVERPOOL

Valiant's conjecture (often phrased as determinant vs permanent)

A **homogeneous algebraic branching program**:

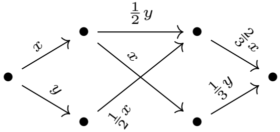

Computes $h = \sum_{s\text{-}t\text{-path } P} \prod_{e \in P} \ell(e)$

Let h be a homogeneous degree d polynomial. The **homogeneous ABP width complexity** $\mathsf{c}(h)$ is defined as the smallest n such that there exist a **homogeneous** width n ABP computing h.

Valiant's conjecture (often phrased as determinant vs permanent)

A **homogeneous algebraic branching program**:

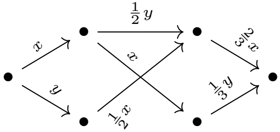

Computes $h = \sum_{s\text{-}t\text{-path } P} \prod_{e \in P} \ell(e)$

Let h be a homogeneous degree d polynomial. The **homogeneous ABP width complexity** $\mathsf{c}(h)$ is defined as the smallest n such that there exist a **homogeneous** width n ABP computing h.

Clearly, $\mathsf{c}(h)$ is always finite: Compute each monomial independently.

Valiant's conjecture (often phrased as determinant vs permanent)

A **homogeneous algebraic branching program**:

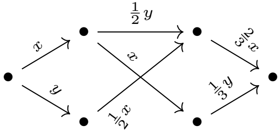

Computes $h = \sum_{s\text{-}t\text{-path } P} \prod_{e \in P} \ell(e)$

Let h be a homogeneous degree d polynomial. The **homogeneous ABP width complexity** $\mathsf{c}(h)$ is defined as the smallest n such that there exist a **homogeneous** width n ABP computing h.

Clearly, $\mathsf{c}(h)$ is always finite: Compute each monomial independently.

Remark: The noncommutative version of c is completely understood (Nisan 1991).

Valiant's conjecture (often phrased as determinant vs permanent)

A **homogeneous algebraic branching program**:

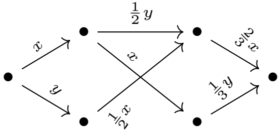

Computes $h = \sum_{s\text{-}t\text{-path } P} \prod_{e \in P} \ell(e)$

Let h be a homogeneous degree d polynomial. The **homogeneous ABP width complexity** $\mathsf{c}(h)$ is defined as the smallest n such that there exist a **homogeneous** width n ABP computing h.

Clearly, $\mathsf{c}(h)$ is always finite: Compute each monomial independently.

Remark: The noncommutative version of c is completely understood (Nisan 1991).

Let $\mathrm{per}_m := \sum_{\pi \in \mathfrak{S}_m} x_{1,\pi(1)} x_{2,\pi(2)} \cdots x_{m,\pi(m)}$.

Valiant's conjecture

The sequence $\mathsf{c}(\mathrm{per}_m)$ grows superpolynomially. In other words, VBP $\neq$ VNP.

Theorem (Grenet 2012): $\mathsf{c}(\mathrm{per}_m) \leq \binom{m}{\lfloor m/2 \rfloor}$.

A simpler machine model: Waring rank

Every homogeneous degree d polynomial h can be written as a sum of d-th powers of homogenoeous linear polynomials ℓ_i:

$$h = \sum_{i=1}^{r} (\ell_i)^d.$$

The smallest r possible is called the Waring rank $\mathrm{WR}(h)$ of h.

A simpler machine model: Waring rank

Every homogeneous degree d polynomial h can be written as a sum of d-th powers of homogenoeous linear polynomials ℓ_i:

$$h = \sum_{i=1}^{r} (\ell_i)^d.$$

The smallest r possible is called the Waring rank $\mathrm{WR}(h)$ of h.

Waring rank measures the size of ABPs of special format:

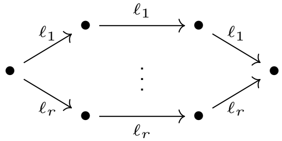

A simpler machine model: Waring rank

Every homogeneous degree d polynomial h can be written as a sum of d-th powers of homogenoeous linear polynomials ℓ_i:

$$h = \sum_{i=1}^{r} (\ell_i)^d.$$

The smallest r possible is called the Waring rank $\mathrm{WR}(h)$ of h.

Waring rank measures the size of ABPs of special format:

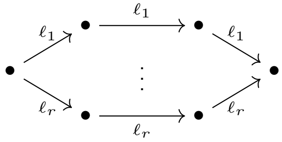

Example:
$6x^2y = (x+y)^3 + (y-x)^3 - 2y^3$, hence $\mathrm{WR}(x^2y) \leq 3$. In fact, $\mathrm{WR}(x^2y) = 3$.

A simpler machine model: Waring rank

Every homogeneous degree d polynomial h can be written as a sum of d-th powers of homogenoeous linear polynomials ℓ_i:

$$h = \sum_{i=1}^{r} (\ell_i)^d.$$

The smallest r possible is called the Waring rank $\mathrm{WR}(h)$ of h.

Waring rank measures the size of ABPs of special format:

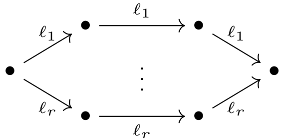

Example:
$6x^2y = (x+y)^3 + (y-x)^3 - 2y^3$, hence $\mathrm{WR}(x^2y) \leq 3$. In fact, $\mathrm{WR}(x^2y) = 3$.

Waring rank for cubic polynomials: matrix multiplication exponent [Chiantini, Hauenstein, I, Landsberg, Ottaviani, 2017]

$$\omega = \liminf_{n\to\infty} \left\{ \log_n \mathrm{WR}\left(\sum_{i,j,k=1}^{n} x_{i,j}x_{j,k}x_{k,i} \right) \right\}$$

Example:

$h = x^2y = \frac{1}{6}((x+y)^3 + (y-x)^3 - 2y^3)$, $\mathrm{WR}(h) = 3$.

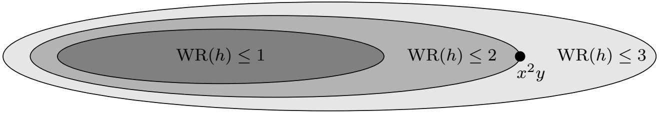

Example:

$h = x^2y = \frac{1}{6}((x+y)^3 + (y-x)^3 - 2y^3)$, $\mathrm{WR}(h) = 3$.

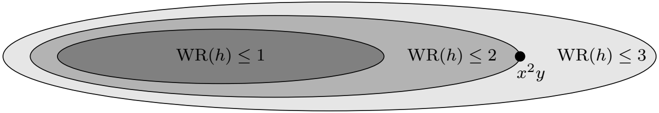

$$\frac{1}{3\varepsilon}\Big((x+\varepsilon y)^3 - x^3\Big) = x^2y + \varepsilon xy^2 + \frac{\varepsilon^2}{3}y^3$$

$$\Big\downarrow \varepsilon \to 0$$

$$x^2y$$

Example:
$h = x^2y = \frac{1}{6}((x+y)^3 + (y-x)^3 - 2y^3)$, $\mathrm{WR}(h) = 3$.

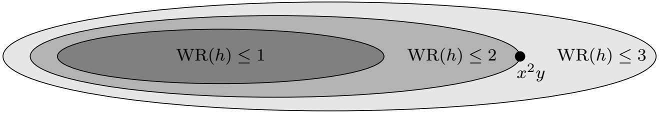

$$\frac{1}{3\varepsilon}\left((x+\varepsilon y)^3 - x^3\right) = x^2y + \varepsilon xy^2 + \frac{\varepsilon^2}{3}y^3 \xrightarrow{\varepsilon \to 0} x^2y$$

This makes determining $\mathrm{WR}(x^2y)$ subtle!
If a continuous function Δ vanishes on all h with $\mathrm{WR}(h) \leq 2$, then f also vanishes on x^2y.

Definition (border Waring rank)

The **border Waring rank** $\underline{\mathrm{WR}}(h)$ is defined as the smallest n such that h can be approximated arbitrarily closely by polynomials of Waring rank $\leq n$. (e.g., $\underline{\mathrm{WR}}(x^2y) = 2$)

Example:
$h = x^2y = \frac{1}{6}((x+y)^3 + (y-x)^3 - 2y^3)$, $\mathrm{WR}(h) = 3$.

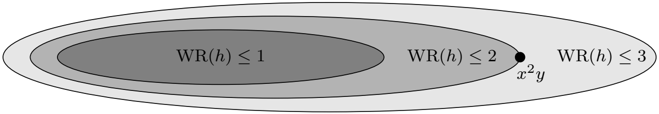

$$\frac{1}{3\varepsilon}\Big((x+\varepsilon y)^3 - x^3\Big) = x^2y + \varepsilon xy^2 + \frac{\varepsilon^2}{3}y^3 \xrightarrow{\varepsilon \to 0} x^2y$$

This makes determining $\mathrm{WR}(x^2y)$ subtle!
If a continuous function Δ vanishes on all h with $\mathrm{WR}(h) \le 2$, then f also vanishes on x^2y.

Definition (border Waring rank)

The **border Waring rank** $\underline{\mathrm{WR}}(h)$ is defined as the smallest n such that h can be approximated arbitrarily closely by polynomials of Waring rank $\le n$. (e.g., $\underline{\mathrm{WR}}(x^2y) = 2$)

Main advantage: $X_n := \{h \mid \underline{\mathrm{WR}}(h) \le n\} = \overline{\{h \mid \mathrm{WR}(h) \le n\}}$ is **closed**, so it is guaranteed that non-membership $p \notin X_n$ can be proved by finding a continuous function Δ that vanishes on X_n, but does not vanish on p.

Example:
$h = x^2y = \frac{1}{6}((x+y)^3 + (y-x)^3 - 2y^3)$, $\mathrm{WR}(h) = 3$.

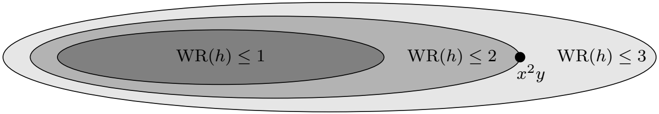

$$\frac{1}{3\varepsilon}\left((x+\varepsilon y)^3 - x^3\right) = x^2y + \varepsilon xy^2 + \frac{\varepsilon^2}{3}y^3 \quad \Big\downarrow \varepsilon \to 0 \quad x^2y$$

This makes determining $\mathrm{WR}(x^2y)$ subtle!
If a continuous function Δ vanishes on all h with $\mathrm{WR}(h) \le 2$, then f also vanishes on x^2y.

Definition (border Waring rank)

The **border Waring rank** $\underline{\mathrm{WR}}(h)$ is defined as the smallest n such that h can be approximated arbitrarily closely by polynomials of Waring rank $\le n$. (e.g., $\underline{\mathrm{WR}}(x^2y) = 2$)

Main advantage: $X_n := \{h \mid \underline{\mathrm{WR}}(h) \le n\} = \overline{\{h \mid \mathrm{WR}(h) \le n\}}$ is **closed**, so it is guaranteed that non-membership $p \notin X_n$ can be proved by finding a continuous function Δ that vanishes on X_n, but does not vanish on p.

Example: Consider the 3-dim vector space $\mathbb{C}[x,y]_2$. Let
$X_1 := \{h \mid \underline{\mathrm{WR}}(h) \le 1\} = \{h \mid \exists \alpha, \beta \in \mathbb{C} : h = (\alpha x + \beta y)^2\}$

Example:
$h = x^2y = \frac{1}{6}((x+y)^3 + (y-x)^3 - 2y^3)$, $\mathrm{WR}(h) = 3$.

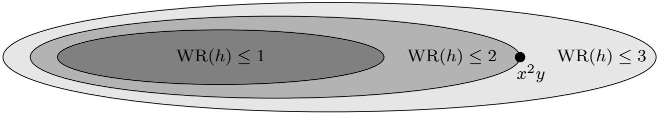

$$\frac{1}{3\varepsilon}\left((x+\varepsilon y)^3 - x^3\right) = x^2y + \varepsilon xy^2 + \frac{\varepsilon^2}{3}y^3 \;\xrightarrow{\varepsilon \to 0}\; x^2y$$

This makes determining $\mathrm{WR}(x^2y)$ subtle!
If a continuous function Δ vanishes on all h with $\mathrm{WR}(h) \le 2$, then f also vanishes on x^2y.

Definition (border Waring rank)

The **border Waring rank** $\underline{\mathrm{WR}}(h)$ is defined as the smallest n such that h can be approximated arbitrarily closely by polynomials of Waring rank $\le n$. (e.g., $\underline{\mathrm{WR}}(x^2y) = 2$)

Main advantage: $X_n := \{h \mid \underline{\mathrm{WR}}(h) \le n\} = \overline{\{h \mid \mathrm{WR}(h) \le n\}}$ is **closed**, so it is guaranteed that non-membership $p \notin X_n$ can be proved by finding a continuous function Δ that vanishes on X_n, but does not vanish on p.

Example: Consider the 3-dim vector space $\mathbb{C}[x,y]_2$. Let
$X_1 := \{h \mid \underline{\mathrm{WR}}(h) \le 1\} = \{h \mid \exists \alpha, \beta \in \mathbb{C} : h = (\alpha x + \beta y)^2\} = \{ax^2 + bxy + cy^2 \mid b^2 - 4ac = 0\}$

Example:
$h = x^2y = \frac{1}{6}((x+y)^3 + (y-x)^3 - 2y^3)$, $\mathrm{WR}(h) = 3$.

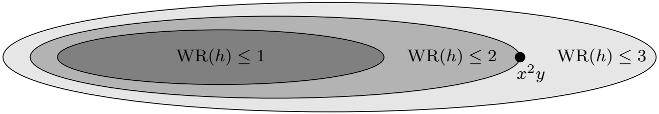

$$\frac{1}{3\varepsilon}\Big((x+\varepsilon y)^3 - x^3\Big) = x^2y + \varepsilon xy^2 + \frac{\varepsilon^2}{3}y^3 \quad \xrightarrow{\varepsilon \to 0} \quad x^2y$$

This makes determining $\mathrm{WR}(x^2y)$ subtle!
If a continuous function Δ vanishes on all h with $\mathrm{WR}(h) \le 2$, then f also vanishes on x^2y.

Definition (border Waring rank)

The **border Waring rank** $\underline{\mathrm{WR}}(h)$ is defined as the smallest n such that h can be approximated arbitrarily closely by polynomials of Waring rank $\le n$. (e.g., $\underline{\mathrm{WR}}(x^2y) = 2$)

Main advantage: $X_n := \{h \mid \underline{\mathrm{WR}}(h) \le n\} = \overline{\{h \mid \mathrm{WR}(h) \le n\}}$ is **closed**, so it is guaranteed that non-membership $p \notin X_n$ can be proved by finding a continuous function Δ that vanishes on X_n, but does not vanish on p.

Example: Consider the 3-dim vector space $\mathbb{C}[x,y]_2$. Let
$X_1 := \{h \mid \underline{\mathrm{WR}}(h) \le 1\} = \{h \mid \exists \alpha, \beta \in \mathbb{C} : h = (\alpha x + \beta y)^2\} = \{ax^2 + bxy + cy^2 \mid b^2 - 4ac = 0\}$
A lower bound: $\underline{\mathrm{WR}}(xy) \ge 2$, because $1^2 - 4 \cdot 0 \cdot 0 = 1 \ne 0$.

- Such $\Delta = b^2 - 4ac$ is sometimes called a **polynomial obstruction** or a **separating polynomial**.
- X_k has more structure and we can obtain more information about polynomial obstructions.

Definition ($\Sigma\wedge\Sigma$)

Let $\Sigma\wedge\Sigma$ be the set of seq. of polynomials whose degree, number of variables, and Waring rank is polynomially bounded.

Definition ($\Sigma\wedge\Sigma$)

Let $\Sigma\wedge\Sigma$ be the set of seq. of polynomials whose degree, number of variables, and Waring rank is polynomially bounded.

Open questions: • $\overline{\Sigma\wedge\Sigma} \stackrel{?}{=} \Sigma\wedge\Sigma$.

Definition ($\Sigma\wedge\Sigma$)

Let $\Sigma\wedge\Sigma$ be the set of seq. of polynomials whose degree, number of variables, and Waring rank is polynomially bounded.

Open questions: • $\overline{\Sigma\wedge\Sigma} \stackrel{?}{=} \Sigma\wedge\Sigma$. • $\overline{\mathsf{VF}} \stackrel{?}{=} \mathsf{VF}$. • $\overline{\mathsf{VBP}} \stackrel{?}{=} \mathsf{VBP}$. • $\overline{\mathsf{VP}} \stackrel{?}{=} \mathsf{VP}$. • $\overline{\mathsf{VNP}} \stackrel{?}{=} \mathsf{VNP}$.

Definition ($\Sigma\wedge\Sigma$)

Let $\Sigma\wedge\Sigma$ be the set of seq. of polynomials whose degree, number of variables, and Waring rank is polynomially bounded.

Open questions: • $\overline{\Sigma\wedge\Sigma} \stackrel{?}{=} \Sigma\wedge\Sigma$. • $\overline{\mathsf{VF}} \stackrel{?}{=} \mathsf{VF}$. • $\overline{\mathsf{VBP}} \stackrel{?}{=} \mathsf{VBP}$. • $\overline{\mathsf{VP}} \stackrel{?}{=} \mathsf{VP}$. • $\overline{\mathsf{VNP}} \stackrel{?}{=} \mathsf{VNP}$.

Comparing a class with its closure:

- [Nisan 1991] showed that $\mathsf{VBP}_{\text{non-comm}} = \overline{\mathsf{VBP}_{\text{non-comm}}}$.

Definition ($\Sigma\wedge\Sigma$)

Let $\Sigma\wedge\Sigma$ be the set of seq. of polynomials whose degree, number of variables, and Waring rank is polynomially bounded.

Open questions: • $\overline{\Sigma\wedge\Sigma} \stackrel{?}{=} \Sigma\wedge\Sigma$. • $\overline{\mathsf{VF}} \stackrel{?}{=} \mathsf{VF}$. • $\overline{\mathsf{VBP}} \stackrel{?}{=} \mathsf{VBP}$. • $\overline{\mathsf{VP}} \stackrel{?}{=} \mathsf{VP}$. • $\overline{\mathsf{VNP}} \stackrel{?}{=} \mathsf{VNP}$.

Comparing a class with its closure:

- [Nisan 1991] showed that $\mathsf{VBP}_{\text{non-comm}} = \overline{\mathsf{VBP}_{\text{non-comm}}}$.
- [Bringmann I Zuiddam 2018] prove $\mathsf{VBP}_2 \subsetneq \overline{\mathsf{VBP}_2}$ (algebraic branching programs of width 2)

Definition ($\Sigma\wedge\Sigma$)

Let $\Sigma\wedge\Sigma$ be the set of seq. of polynomials whose degree, number of variables, and Waring rank is polynomially bounded.

Open questions: • $\overline{\Sigma\wedge\Sigma} \stackrel{?}{=} \Sigma\wedge\Sigma$. • $\overline{\mathsf{VF}} \stackrel{?}{=} \mathsf{VF}$. • $\overline{\mathsf{VBP}} \stackrel{?}{=} \mathsf{VBP}$. • $\overline{\mathsf{VP}} \stackrel{?}{=} \mathsf{VP}$. • $\overline{\mathsf{VNP}} \stackrel{?}{=} \mathsf{VNP}$.

Comparing a class with its closure:

- [Nisan 1991] showed that $\mathsf{VBP}_{\text{non-comm}} = \overline{\mathsf{VBP}_{\text{non-comm}}}$.
- [Bringmann I Zuiddam 2018] prove $\mathsf{VBP}_2 \subsetneq \overline{\mathsf{VBP}_2}$ (algebraic branching programs of width 2)

Placing a closure of a class into a larger non-closed class ("de-bordering"):

- [Bläser Dörfler I 2020] prove $\mathsf{c}(p) \leq \underline{\mathrm{WR}}(p)$. Hence $\overline{\Sigma\wedge\Sigma} \subseteq \mathsf{VBP}$ (which was discovered earlier by Forbes).

Definition ($\Sigma\wedge\Sigma$)

Let $\Sigma\wedge\Sigma$ be the set of seq. of polynomials whose degree, number of variables, and Waring rank is polynomially bounded.

Open questions: • $\overline{\Sigma\wedge\Sigma} \stackrel{?}{=} \Sigma\wedge\Sigma$. • $\overline{\mathsf{VF}} \stackrel{?}{=} \mathsf{VF}$. • $\overline{\mathsf{VBP}} \stackrel{?}{=} \mathsf{VBP}$. • $\overline{\mathsf{VP}} \stackrel{?}{=} \mathsf{VP}$. • $\overline{\mathsf{VNP}} \stackrel{?}{=} \mathsf{VNP}$.

Comparing a class with its closure:

- [Nisan 1991] showed that $\mathsf{VBP}_{\text{non-comm}} = \overline{\mathsf{VBP}_{\text{non-comm}}}$.
- [Bringmann I Zuiddam 2018] prove $\mathsf{VBP}_2 \subsetneq \overline{\mathsf{VBP}_2}$ (algebraic branching programs of width 2)

Placing a closure of a class into a larger non-closed class ("de-bordering"):

- [Bläser Dörfler I 2020] prove $\mathsf{c}(p) \leq \underline{\mathrm{WR}}(p)$. Hence $\overline{\Sigma\wedge\Sigma} \subseteq \mathsf{VBP}$ (which was discovered earlier by Forbes).
- [Dutta Dwivedi Saxena 2021] prove $\overline{\Sigma^k\Pi\Sigma} \subseteq \mathsf{VBP}$ for any constant k, using the logarithmic derivative.

Definition ($\Sigma\wedge\Sigma$)

Let $\Sigma\wedge\Sigma$ be the set of seq. of polynomials whose degree, number of variables, and Waring rank is polynomially bounded.

Open questions: • $\overline{\Sigma\wedge\Sigma} \stackrel{?}{=} \Sigma\wedge\Sigma$. • $\overline{\mathsf{VF}} \stackrel{?}{=} \mathsf{VF}$. • $\overline{\mathsf{VBP}} \stackrel{?}{=} \mathsf{VBP}$. • $\overline{\mathsf{VP}} \stackrel{?}{=} \mathsf{VP}$. • $\overline{\mathsf{VNP}} \stackrel{?}{=} \mathsf{VNP}$.

Comparing a class with its closure:

- [Nisan 1991] showed that $\mathsf{VBP}_{\text{non-comm}} = \overline{\mathsf{VBP}_{\text{non-comm}}}$.
- [Bringmann I Zuiddam 2018] prove $\mathsf{VBP}_2 \subsetneq \overline{\mathsf{VBP}_2}$ (algebraic branching programs of width 2)

Placing a closure of a class into a larger non-closed class ("de-bordering"):

- [Bläser Dörfler I 2020] prove $\mathsf{c}(p) \leq \underline{\mathrm{WR}}(p)$. Hence $\overline{\Sigma\wedge\Sigma} \subseteq \mathsf{VBP}$ (which was discovered earlier by Forbes).
- [Dutta Dwivedi Saxena 2021] prove $\overline{\Sigma^k\Pi\Sigma} \subseteq \mathsf{VBP}$ for any constant k, using the logarithmic derivative.
- Open question: Is $\overline{\mathsf{VF}} \subseteq \mathsf{VNP}$ or not?

Mulmuley-Sohoni strengthening of Valiant's conjecture:

Valiant's conjecture

The sequence $\mathsf{c}(\mathrm{per}_m)$ grows superpolynomially. In other words, $\mathsf{VBP} \neq \mathsf{VNP}$.

Mulmuley-Sohoni conjecture

The sequence $\underline{\mathsf{c}}(\mathrm{per}_m)$ grows superpolynomially. In other words, $\mathsf{VNP} \not\subseteq \overline{\mathsf{VBP}}$ or equivalently $\overline{\mathsf{VBP}} \neq \overline{\mathsf{VNP}}$.

Fundamental open question: Are the two conjectures equivalent?

Let $X_n := \{h \mid \underline{\mathsf{c}}(h) \le n\} = \overline{\{h \mid \mathsf{c}(h) \le n\}}$.

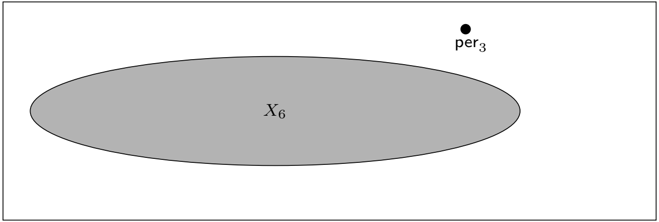

would imply $\underline{\mathsf{c}}(\mathsf{per}_3) > 6$

[Bläser, I, Lysikov, Pandey, Schreyer, SODA21]: In general, orbit **closure** containment is NP-hard (minrank of tensors).

X_n is a projective variety

Let $X_n := \overline{\{h \mid \mathsf{c}(h) \leq n\}}$.

A subset $\mathcal{X} \subseteq \mathbb{C}^N$ is a **projective variety** if there exist finitely many homogeneous polynomials $\Delta_1, \dots, \Delta_k$ such that

$$h \in \mathcal{X} \quad \text{iff} \quad \Delta_1(h) = \Delta_2(h) = \cdots = \Delta_k(h) = 0.$$

X_n is a projective variety (by Chevalley's theorem).

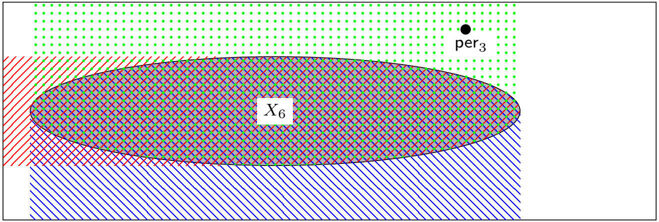

would imply $\underline{\mathsf{c}}(\mathrm{per}_3) > 6$

X_n is a projective variety

Let $X_n := \overline{\{h \mid \mathsf{c}(h) \leq n\}}$.

A subset $\mathcal{X} \subseteq \mathbb{C}^N$ is a **projective variety** if there exist finitely many homogeneous polynomials $\Delta_1, \ldots, \Delta_k$ such that

$$h \in \mathcal{X} \quad \text{iff} \quad \Delta_1(h) = \Delta_2(h) = \cdots = \Delta_k(h) = 0.$$

X_n is a projective variety (by Chevalley's theorem).

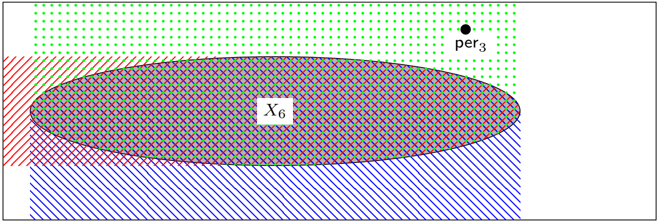

would imply $\underline{\mathsf{c}}(\text{per}_3) > 6$

Fundamental Conclusion: All border complexity lower bounds can be proved via **polynomials**

$\text{per}_m \notin X_n$ iff there exists a homogeneous polynomial Δ with

- $\Delta(h) = 0$ for all $h \in X_n$ and
- $\Delta(\text{per}_m) \neq 0$.

Meta-complexity (algebraic natural proofs): What can be said about the complexity of the Δ_i?

Complexity lower bounds via highest weight polynomials

Definition (highest weight polynomial)

A function Δ is called a **highest weight polynomial** of weight $\lambda = (\lambda_1, \ldots, \lambda_N)$, if

- Δ is invariant under the action of upper triangular matrices with 1s on the diagonal
- and Δ gets rescaled by $\alpha_1^{\lambda_1} \cdots \alpha_N^{\lambda_N}$ under the action of diagonal matrices $\mathrm{diag}(\alpha_1, \ldots, \alpha_N)$.

Recall: Want Δ vanishing on X_n and $\Delta(\mathrm{per}_m) \neq 0$.

Theorem (representation theory)

If $\mathrm{per}_m \notin X_n$, then there exists a highest weight polynomial Δ such that $A\Delta$ vanishes on X_n and $A\Delta(\mathrm{per}_m) \neq 0$ for a generic matrix A.

Complexity lower bounds via highest weight polynomials

Definition (highest weight polynomial)

A function Δ is called a **highest weight polynomial** of weight $\lambda = (\lambda_1, \ldots, \lambda_N)$, if

- Δ is invariant under the action of upper triangular matrices with 1s on the diagonal
- and Δ gets rescaled by $\alpha_1^{\lambda_1} \cdots \alpha_N^{\lambda_N}$ under the action of diagonal matrices $\mathrm{diag}(\alpha_1, \ldots, \alpha_N)$.

Recall: Want Δ vanishing on X_n and $\Delta(\mathrm{per}_m) \neq 0$.

Theorem (representation theory)

If $\mathrm{per}_m \notin X_n$, then there exists a highest weight polynomial Δ such that $A\Delta$ vanishes on X_n and $A\Delta(\mathrm{per}_m) \neq 0$ for a generic matrix A.

Fundamental Conclusion: All border complexity lower bounds can be proved via highest weight polynomials

If complexity lower bounds exist, then there exist highest weight polynomials proving them.

[Bläser, Dörfler, I, CCC21]: Given a highest weight polynomial Δ, succinctly encoded as a semistandard Young tableau, then it is #P-hard to evaluate Δ at a fixed point of Waring rank 3.

Mulmuley and Sohoni's heuristic attempt: Occurrence Obstructions

Proposition (a coarse technique for finding complexity lower bounds: occurrence obstructions)

If there exists λ such that for a generic matrix A we have

- for **all (!)** highest weight polynomials Δ of weight λ: $A\Delta$ vanishes on X_n
- there exists a highest weight polynomial Δ of weight λ such that $A\Delta(\text{per}_m) \neq 0$

then $\text{per}_m \notin X_n$.

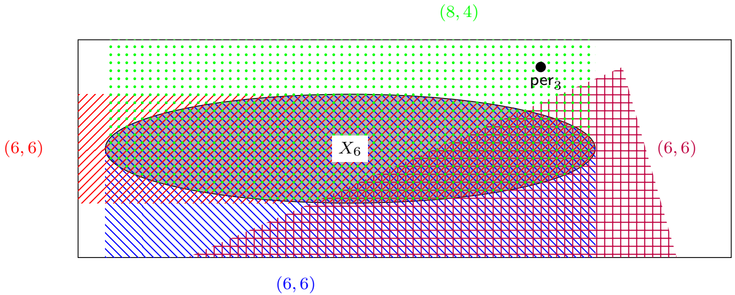

[I, Panova FOCS16] and [Bürgisser, I, Panova FOCS16]: In a non-homogeneous setting, no occurrence obstructions exist. Nothing is known about the homogeneous setting.

More general heuristic attempt: Multiplicity obstructions

Let $Y_m := \overline{\{\mathsf{per}_m(A\vec{x})\}}$

$\mathsf{per}_m \in X_n \quad \Leftrightarrow \quad Y_m \subseteq X_n$

More general heuristic attempt: Multiplicity obstructions

Let $Y_m := \overline{\{\mathrm{per}_m(A\vec{x})\}}$
$\mathrm{per}_m \in X_n \quad \Leftrightarrow \quad Y_m \subseteq X_n$

Def.: The **multiplicity** $\mathrm{mult}_\lambda(\mathbb{C}[X_n])$ is defined as the dimension of the space of highest weight polynomials of weight λ restricted to X_n.

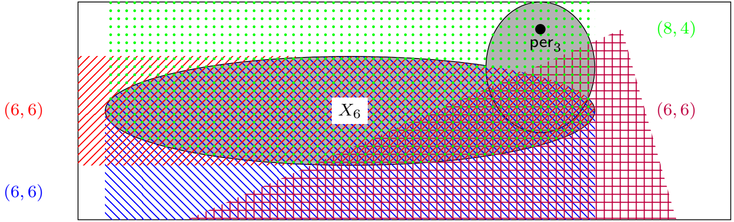

More general heuristic attempt: Multiplicity obstructions

Let $Y_m := \overline{\{\mathrm{per}_m(A\vec{x})\}}$

$\mathrm{per}_m \in X_n \quad \Leftrightarrow \quad Y_m \subseteq X_n$

Def.: The **multiplicity** $\mathrm{mult}_\lambda(\mathbb{C}[X_n])$ is defined as the dimension of the space of highest weight polynomials of weight λ restricted to X_n.

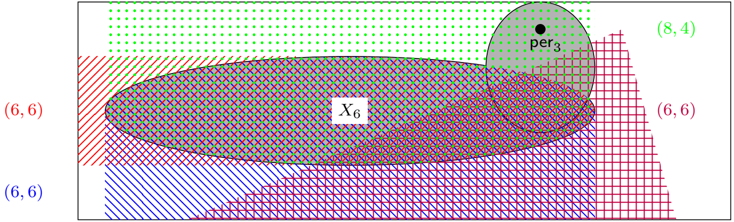

If $\mathrm{mult}_\lambda(\mathbb{C}[Y_m]) > \mathrm{mult}_\lambda(\mathbb{C}[X_n])$, then $\mathrm{per}_m \notin X_n$.

Not much is known about these multiplicities!

- [Dörfler, I, Panova ICALP19] There are finite homogeneous settings where multiplicity obstructions are stronger than occurrence obstructions.
- [I, Kandasamy STOC20] Multiplicity obstructions can be created in homogeneous toy settings without contructing the highest weight polynomials.

The original GCT approach (Mulmuley-Sohoni):

1. The multiplicities are easier to study than the polynomials.
2. Oftentimes the multiplicities in representation theory have well-understood combinatorial descriptions (Littlewood-Richardson coefficient).
3. Positivity of the Littlewood-Richardson coefficient can be decided in polynomial time (using a combinatorial algorithm).
4. So maybe this works as well for $\mathrm{mult}_\lambda(\mathbb{C}[Y_m])$ and $\mathrm{mult}_\lambda(\mathbb{C}[X_n])$.
5. Then one could analyze the algorithms and construct an input λ with $\mathrm{mult}_\lambda(\mathbb{C}[Y_m]) > 0 = \mathrm{mult}_\lambda(\mathbb{C}[X_n])$.

The original GCT approach (Mulmuley-Sohoni):

1. The multiplicities are easier to study than the polynomials.
2. Oftentimes the multiplicities in representation theory have well-understood combinatorial descriptions (Littlewood-Richardson coefficient).
3. Positivity of the Littlewood-Richardson coefficient can be decided in polynomial time (using a combinatorial algorithm).
4. So maybe this works as well for $\mathrm{mult}_\lambda(\mathbb{C}[Y_m])$ and $\mathrm{mult}_\lambda(\mathbb{C}[X_n])$.
5. Then one could analyze the algorithms and construct an input λ with $\mathrm{mult}_\lambda(\mathbb{C}[Y_m]) > 0 = \mathrm{mult}_\lambda(\mathbb{C}[X_n])$.

We know by now that

- About 3.: the positivity of coefficients is often NP-hard [I Mulmuley Walter 2017], [Fischer I 2020].
- About 2.: Connections to classical questions in algebraic combinatorics!

Closely related multiplicities in Stanley's "Positivity Problems and Conjectures in Algebraic Combinatorics" (2000):

Problem 9

Find a combinatorial interpretation of the plethysm coefficient $a_\lambda(d, n)$.

Problem 10

Find a combinatorial interpretation of the Kronecker coefficient $k(\lambda, \mu, \nu)$.

Problem 11

Find a combinatorial interpretation of the Schubert coefficient.

In other words, is there a **combinatorial proof** that these numbers are in $\mathbb{N}$.

Closely related multiplicities in Stanley's "Positivity Problems and Conjectures in Algebraic Combinatorics" (2000):

Problem 9

Find a combinatorial interpretation of the plethysm coefficient $a_\lambda(d, n)$.

Problem 10

Find a combinatorial interpretation of the Kronecker coefficient $k(\lambda, \mu, \nu)$.

Problem 11

Find a combinatorial interpretation of the Schubert coefficient.

In other words, is there a **combinatorial proof** that these numbers are in $\mathbb{N}$.

Recall that $f : \{0,1\}^* \to \mathbb{N}$ is in #P if there exists a polytime Turing machine M with

$$\forall w \in \{0,1\}^* : \ \#\mathsf{acc}_M(w) = f(w)$$

Closely related multiplicities in Stanley's "Positivity Problems and Conjectures in Algebraic Combinatorics" (2000):

Problem 9

Find a combinatorial interpretation of the plethysm coefficient $a_\lambda(d, n)$.

Problem 10

Find a combinatorial interpretation of the Kronecker coefficient $k(\lambda, \mu, \nu)$.

Problem 11

Find a combinatorial interpretation of the Schubert coefficient.

In other words, is there a **combinatorial proof** that these numbers are in $\mathbb{N}$.

Recall that $f : \{0,1\}^* \to \mathbb{N}$ is in #P if there exists a polytime Turing machine M with

$$\forall w \in \{0,1\}^* : \ \#\mathsf{acc}_M(w) = f(w)$$

Problems 9, 10, 11 are in GapP $=$ #P$-$#P, and all these are nonnegative. Are they in #P?

The problem #Sperner

Input: A Boolean circuit C with $3\lceil \log n \rceil$ inputs and 2-bit output.

The problem #Sperner

Input: A Boolean circuit C with $3\lceil \log n \rceil$ inputs and 2-bit output.
C describes a coloring of a side length n triangular grid.

- Positions are (x, y, z) with $x + y + z = n$
- color$(x, y, z) = \begin{cases} \text{red} & \text{if } C(x,y,z) = (0,0) \\ \text{blue} & \text{if } C(x,y,z) = (0,1) \\ \text{green} & \text{if } C(x,y,z) \in \{(1,0),(1,1)\} \end{cases}$

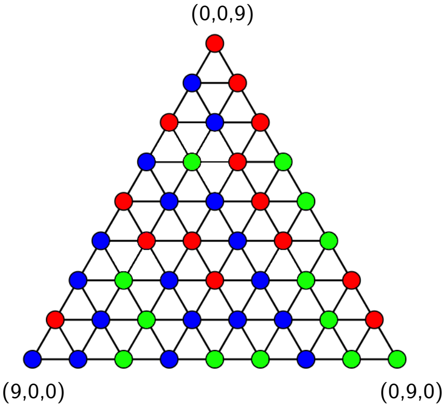

The problem #Sperner

Input: A Boolean circuit C with $3\lceil \log n \rceil$ inputs and 2-bit output.
C describes a coloring of a side length n triangular grid.

- Positions are (x, y, z) with $x + y + z = n$
- $\text{color}(x, y, z) = \begin{cases} \text{red} & \text{if } C(x, y, z) = (0, 0) \\ \text{blue} & \text{if } C(x, y, z) = (0, 1) \\ \text{green} & \text{if } C(x, y, z) \in \{(1, 0), (1, 1)\} \end{cases}$

We enforce constraints at the corners and sides:

- Corners:
  $\text{color}(n, 0, 0) = \text{blue}$,
  $\text{color}(0, n, 0) = \text{green}$,
  $\text{color}(0, 0, n) = \text{red}$
- Sides:
  $\text{color}(i, j, 0) \in \{\text{blue,green}\}$,
  $\text{color}(i, 0, j) \in \{\text{blue,red}\}$,
  $\text{color}(0, i, j) \in \{\text{green,red}\}$

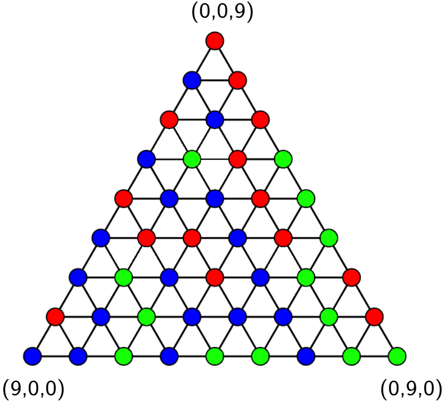

#Sperner

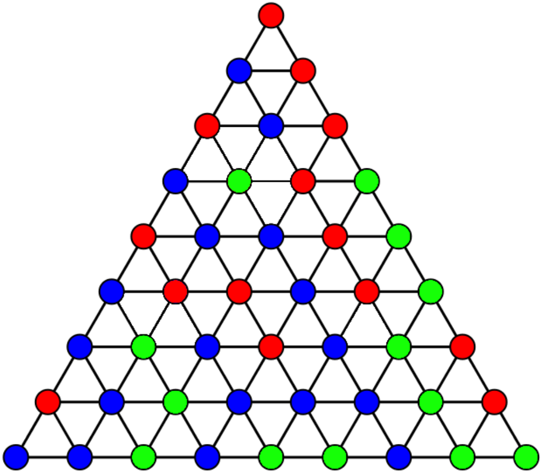

#Sperner(C) = the number of rainbow triangles in the coloring given by C.

#Sperner

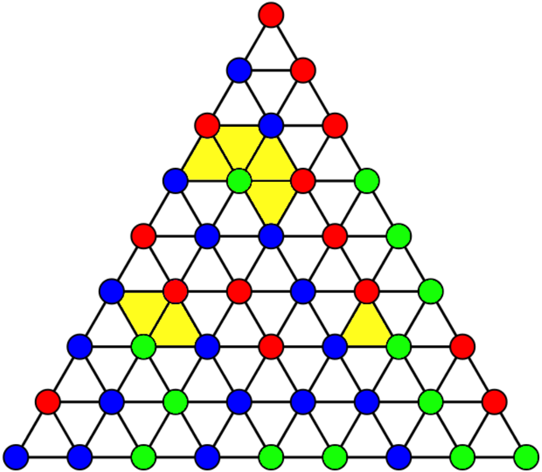

#Sperner(C) = the number of rainbow triangles in the coloring given by C.

$\#$Sperner$(C) =$ the number of rainbow triangles in the coloring given by C.

Clearly $\#$Sperner $\in \#$P.

$\#\text{Sperner}(C) =$ the number of rainbow triangles in the coloring given by C.

Clearly $\#\text{Sperner} \in \#\text{P}$.

Sperner's Lemma

$\forall C \in \{0,1\}^*: \quad \#\text{Sperner}(C) \geq 1.$

Maybe $\#\text{Sperner}-1$ is another candidate for being in $\text{GapP} \setminus \#\text{P}$?

$\#\text{Sperner} - 1 \in \#\text{P}?$

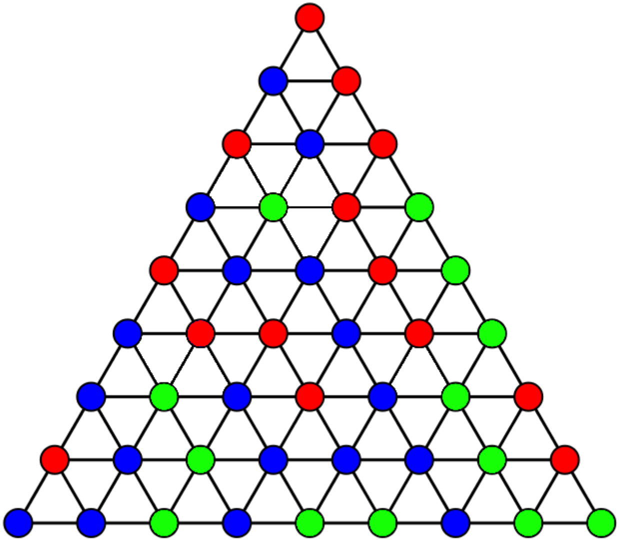

#Sperner $- 1 \in$ #P?

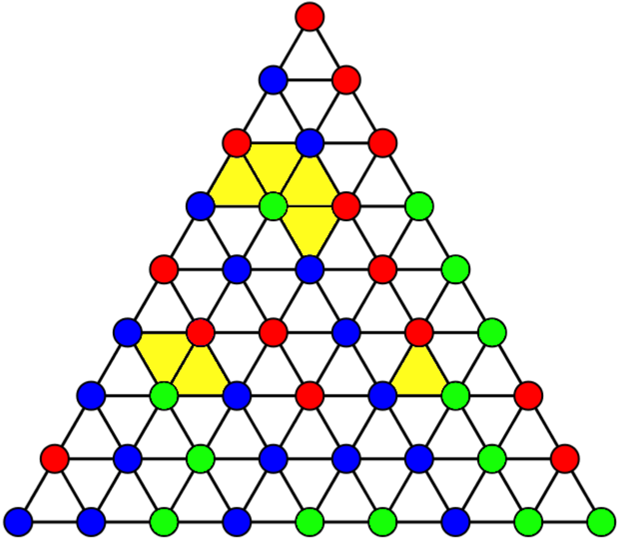

#Sperner − 1 ∈ #P?

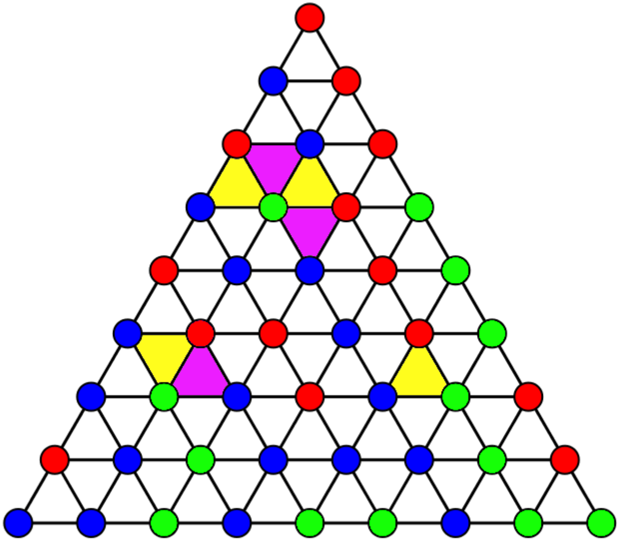

#Sperner $- 1 \in$ #P?

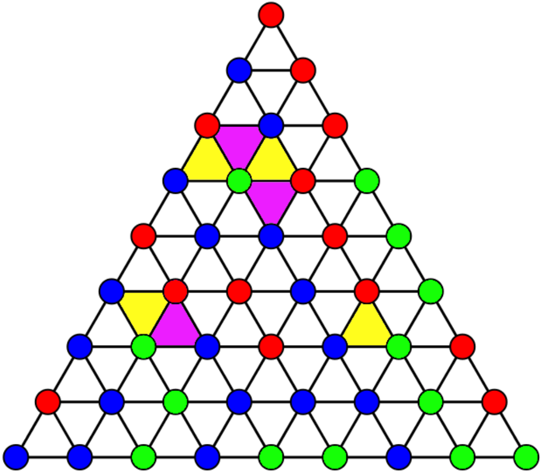

#Sperner $- 1 \in$ #P, because we can ignore the yellow triangles and count the purple triangles twice!

When does such a trick work? An easier version:

Given two #P functions f, g with the property that $f + 1 = g$.

When does such a trick work? An easier version:

Given two #P functions f, g with the property that $f + 1 = g$. Then $f + g - 1 = 2g \in \#\mathsf{P}$.

When does such a trick work? An easier version:

Given two #P functions f, g with the property that $f + 1 = g$. Then $f + g - 1 = 2g \in \#\mathrm{P}$.

- $f + g - 1$ and $2g$ are the same element in the quotient ring $\mathbb{Q}[f, g]/(f + 1 - g)$.
- doubling is a closure property of #P.

When does such a trick work? An easier version:

Given two $\#$P functions f, g with the property that $f + 1 = g$. Then $f + g - 1 = 2g \in \#$P.

- $f + g - 1$ and $2g$ are the same element in the quotient ring $\mathbb{Q}[f, g]/(f + 1 - g)$.
- doubling is a closure property of $\#$P.

The functional closure properties of $\#$P

When does such a trick work? An easier version:

Given two #P functions f, g with the property that $f+1=g$. Then $f+g-1=2g \in \#\mathrm{P}$.

- $f+g-1$ and $2g$ are the same element in the quotient ring $\mathbb{Q}[f,g]/(f+1-g)$.
- doubling is a closure property of #P.

The functional closure properties of #P

Every polynomial $\varphi \in \mathbb{Q}[f_1,\ldots,f_k]$ has a unique expression over the **binomial basis**:

$$\varphi(f_1,\ldots,f_k) = \sum_{\vec{a}\in\mathbb{N}^k} c_{\vec{a}} \binom{f_1}{a_1}\cdots\binom{f_k}{a_k}.$$

If all $c_{\vec{a}} \in \mathbb{N}$, then we say that φ is **binomial-good**.

For example, $f^2 g = 2\binom{f}{2}\binom{g}{1} + \binom{f}{1}\binom{g}{1}$.

When does such a trick work? An easier version:

Given two #P functions f, g with the property that $f+1=g$. Then $f+g-1=2g \in$ #P.

- $f+g-1$ and $2g$ are the same element in the quotient ring $\mathbb{Q}[f,g]/(f+1-g)$.
- doubling is a closure property of #P.

The functional closure properties of #P

Every polynomial $\varphi \in \mathbb{Q}[f_1, \ldots, f_k]$ has a unique expression over the **binomial basis**:

$$\varphi(f_1, \ldots, f_k) = \sum_{\vec{a} \in \mathbb{N}^k} c_{\vec{a}} \binom{f_1}{a_1} \cdots \binom{f_k}{a_k}.$$

If all $c_{\vec{a}} \in \mathbb{N}$, then we say that φ is **binomial-good**.

For example, $f^2 g = 2\binom{f}{2}\binom{g}{1} + \binom{f}{1}\binom{g}{1}$.

Theorem (Hertrampf, Vollmer, Wagner, 1995)

The binomial-good polynomials are exactly the relativizing functional closure properties of #P.

When does such a trick work? An easier version:

Given two #P functions f, g with the property that $f + 1 = g$. Then $f + g - 1 = 2g \in$ #P.

- $f + g - 1$ and $2g$ are the same element in the quotient ring $\mathbb{Q}[f, g]/(f + 1 - g)$.
- doubling is a closure property of #P.

The functional closure properties of #P

Every polynomial $\varphi \in \mathbb{Q}[f_1, \ldots, f_k]$ has a unique expression over the **binomial basis**:

$$\varphi(f_1, \ldots, f_k) = \sum_{\vec{a} \in \mathbb{N}^k} c_{\vec{a}} \binom{f_1}{a_1} \cdots \binom{f_k}{a_k}.$$

If all $c_{\vec{a}} \in \mathbb{N}$, then we say that φ is **binomial-good**.

For example, $f^2 g = 2\binom{f}{2}\binom{g}{1} + \binom{f}{1}\binom{g}{1}$.

Theorem (Hertrampf, Vollmer, Wagner, 1995)

The binomial-good polynomials are exactly the relativizing functional closure properties of #P.

Given an ideal I, a coset $\varphi + I$ is called **binomial-good** if $\varphi + I$ contains a binomial-good representative.

- [I, Pak 2022]: For a large class of ideals, the Sperner trick works if and only if $\varphi + I$ is binomial-good.
- If the ideal is generated by linear polynomials, then checking if $\varphi + I$ is binomial-good reduces to integer programming.

This gives an algorithmic way of finding out when exactly "Sperner-type tricks" work. With some extra work (i.e., simulating #P instances in TFNP search problems) we get a TFNP-like inclusion diagram [I, Pak 2022]:

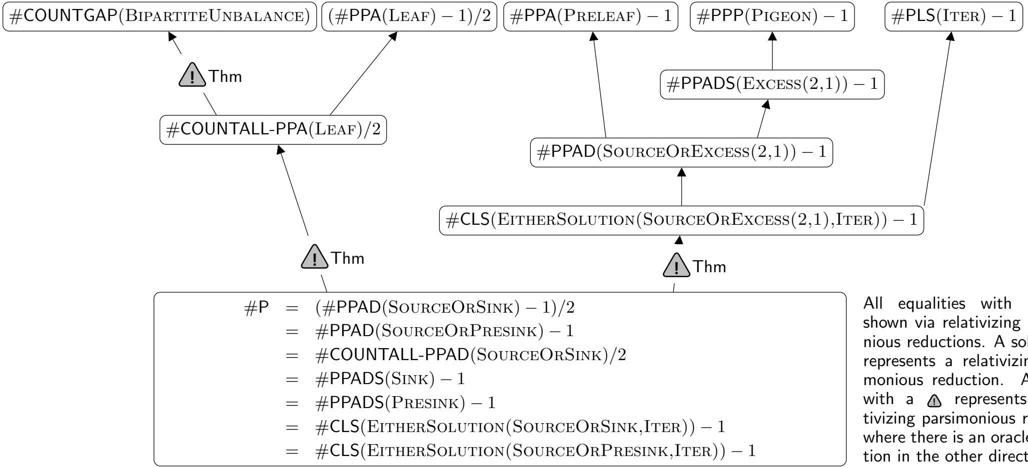

All equalities with #P are shown via relativizing parsimonious reductions. A solid arrow represents a relativizing parsimonious reduction. An arrow with a ⚠ represents a relativizing parsimonious reduction where there is an oracle separation in the other direction.

This gives an algorithmic way of finding out when exactly "Sperner-type tricks" work. With some extra work (i.e., simulating #P instances in TFNP search problems) we get a TFNP-like inclusion diagram [I, Pak 2022]:

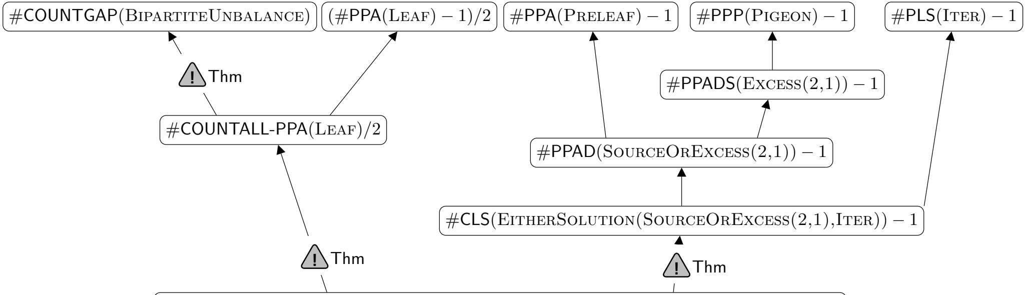

$$
\begin{aligned}
\#\mathsf{P} &= (\#\mathsf{PPAD}(\textsc{SourceOrSink}) - 1)/2 \\
&= \#\mathsf{PPAD}(\textsc{SourceOrPresink}) - 1 \\
&= \#\mathsf{COUNTALL\text{-}PPAD}(\textsc{SourceOrSink})/2 \\
&= \#\mathsf{PPADS}(\textsc{Sink}) - 1 \\
&= \#\mathsf{PPADS}(\textsc{Presink}) - 1 \\
&= \#\mathsf{CLS}(\textsc{EitherSolution}(\textsc{SourceOrSink},\textsc{Iter})) - 1 \\
&= \#\mathsf{CLS}(\textsc{EitherSolution}(\textsc{SourceOrPresink},\textsc{Iter})) - 1
\end{aligned}
$$

All equalities with #P are shown via relativizing parsimonious reductions. A solid arrow represents a relativizing parsimonious reduction. An arrow with a ⚠ represents a relativizing parsimonious reduction where there is an oracle separation in the other direction.

Thank you for your attention!