

Lower Bounds for Symmetric Arithmetic Circuits

Anuj Dawar


Department of Computer Science and Technology, University of Cambridge

Joint work with Gregory Wilsenach (ICALP 2020 and ITCS 2022)

ICMS, 7 July 2022

Arithmetic Circuits

An *Arithmetic Circuit* over a field K computes (or represents) a *polynomial* in $K[X]$.



Matrix Inputs

We are often interested in inputs which are entries of *a matrix*.

$$X = \{x_{ij} \mid 1 \leq i \leq m; 1 \leq j \leq n\}$$

Especially, when the input is a *square matrix*, so $m = n$.

$$\text{tr}(X) = \sum_i x_{ii}$$

$$\det(X) = \sum_{\sigma \in \text{Sym}_n} \text{sgn}(\sigma) \prod_{i \in [n]} x_{i\sigma(i)}$$

$$\text{per}(X) = \sum_{\sigma \in \text{Sym}_n} \prod_{i \in [n]} x_{i\sigma(i)}$$

Lower Bounds for Arithmetic Circuits

We have lower bounds for *restricted* classes of circuits computing the permanent.

No monotone family of circuits of sub-exponential size for the permanent.
(Jerrum, Snir 1982)

No sub-exponential size family of depth 3 circuits for the permanent over any finite field.
(Grigoriev, Karpinski 1998)

Both methods also yield similar lower bounds for the *determinant*

We consider upper and lower bounds for *symmetric* circuits computing the determinant and the permanent.

Symmetric Arithmetic Circuits

Suppose C is a circuit computing a polynomial $p \in K[X]$.

Sym_X —the group of *permutations* of X .

Let Γ be a group acting on X (or simply $\Gamma \leq \text{Sym}_X$).


p is Γ -*symmetric* if for all $\pi \in \Gamma$, $p^\pi = p$.

C is Γ -*symmetric* if the action of Γ on the inputs X extends to an *automorphism* of C .

Elementary Symmetric Polynomials

The *elementary symmetric polynomials* on a set of variables X are Sym_X -symmetric.

Example: $\prod_{1 \leq i \leq n} x_i$.



Fan-in matters!

Square Symmetric Action

When the input is a *square matrix* $X = \{x_{ij} \mid 1 \leq i, j \leq n\}$, the full symmetric group is $\text{Sym}_X = \text{Sym}_{[n] \times [n]}$.

The matrix polynomials $\text{tr}(X)$, $\det(X)$ and $\text{per}(X)$ are all invariant under the action of $\text{Sym}_{[n]}$ given by

$$x_{ij}^\pi = x_{\pi(i)\pi(j)}.$$

i.e., *simultaneous row and column permutations*.

We say that these polynomials are *square symmetric*.

Matrix Symmetric Action

The *permanent*

$$\text{per}(X) = \sum_{\sigma \in \text{Sym}_n} \prod_{i \in [n]} x_{i\sigma(i)}$$

is further invariant under *independent row and column permutations*.

That is, under the action of $\text{Sym}_{[n]} \times \text{Sym}_{[n]}$ given by

$$x_{ij}^{(\sigma, \pi)} = x_{\sigma(i)\pi(j)}.$$

We say that $\text{per}(X)$ is *matrix symmetric*.

$\text{tr}(X)$ and $\text{det}(X)$ are not matrix symmetric.

Determinant

The invariance group of

$$\det(X) = \sum_{\sigma \in \text{Sym}_n} \text{sgn}(\sigma) \prod_{i \in [n]} x_{i\sigma(i)}$$

includes

$$D = \{(\sigma, \pi) \in \text{Sym}_{[n]} \times \text{Sym}_{[n]} \mid \text{sgn}(\sigma) = \text{sgn}(\pi)\} \times \mathbb{Z}_2.$$

In particular, it is $\text{Alt}_{[n]} \times \text{Alt}_{[n]}$ -symmetric.

The defining expression yields a circuit with these symmetries, but of $\Omega(n!)$ size.

Circuits for the Determinant

Many different algorithms yield small circuits for the determinant, but they are not often *symmetric*.

e.g. pivot choice is a *symmetry-breaking* operation.

Le Verrier's method shows how to compute $\det(X)$ (for fields of *characteristic 0*) from

$$\operatorname{tr}(X), \operatorname{tr}(X^2), \dots, \operatorname{tr}(X^n).$$

Since each $\operatorname{tr}(X^i)$ can be computed by a small *square-symmetric* circuit, this gives a *polynomial-size, square-symmetric* (i.e. $\operatorname{Sym}_{[n]}$ -symmetric) circuit for the determinant.

Permanent

The defining expression for the permanent yields *matrix-symmetric* circuits of size $\Omega(n!)$.

The smallest known circuits for the permanent are given by *Ryser's formula*:

$$\text{per}(X) = (-1)^n \sum_{S \subseteq [n]} (-1)^{|S|} \prod_{i=1}^n \sum_{j \in S} x_{ij}.$$

This gives a *matrix-symmetric* circuit of size $O(n^2 2^n)$.

Results

Γ	$\{\text{id}\}$	$\text{Sym}_{[n]}$	$\text{Alt}_{[n]} \times \text{Alt}_{[n]}$	$\text{Sym}_{[n]} \times \text{Sym}_{[n]}$
Det	$O(n^4)$	$O(n^4)$ <i>(char 0)</i>	$2^{\Omega(n)}$ <i>(char 0)</i>	N/A
Perm	$O(n^2 2^n)$ VP = VNP?	$2^{\Omega(n)}$ <i>(char 0)</i>	$2^{\Omega(n)}$ <i>(char $\neq 2$)</i>	$2^{\Omega(n)}$ <i>(char $\neq 2$)</i>

Actually, all lower bounds are not just on the *size* of the circuit, but on *orbit size*.

Proof Ingredients – Support Theorem

Any group $\Delta \leq \text{Alt}_A$ with *small index* ($[\text{Alt}_A : \Delta]$) has *small support* *i.e.* a *small* set $S \subset A$ such that any $\pi \in \text{Alt}_A$ which *fixes* S *pointwise* is in Δ .

So, if C is a *small* Γ -symmetric circuit (where Γ is any of $\text{Sym}_A, \text{Alt}_A \times \text{Alt}_B, \text{Sym}_A \times \text{Sym}_B$) then we can associate with each gate g of C , a *small support*

i.e. a *small* set $S \subset A \cup B$ such that any automorphism of C which *fixes* S *pointwise* fixes g .

Aim to show lower bounds on *support size*

- *super-constant* support size implies *super-polynomial* orbit size.
- *linear* support size implies *exponential* orbit size.

Proof Ingredients – Indistinguishable Pairs

Aim to construct, for a polynomial p , a pair of *matrices* M, M' such that

- $p(M) \neq p(M')$
- M and M' cannot be distinguished by circuits with small support.

The matrices we construct are $\{0, 1\}$ -matrices, so can be seen as the *biadjacency* matrices of a *bipartite graph*

$$(A, B, E \subseteq A \times B).$$

Proof Ingredients – Bijection Games

A *two-player game* played on a pair of graphs G and H with k pairs of pebbles (a_i, b_i) .

We fix a group $\Gamma \leq \text{Sym}_{V(H)}$ and an initial bijection $h : V(G) \rightarrow V(H)$. At any point, the pebbles a_i are on elements of $V(G)$ and b_i on elements of $V(H)$.

- *Spoiler* chooses a pair of pebbles a_i and b_i ;
- *Duplicator* chooses a *permutation* $\pi \in \Gamma$ such that for pebbles a_j and $b_j (j \neq i)$, $\pi \circ h(a_j) = b_j$;
- *Spoiler* chooses $a \in V(G)$ and places a_i on a and b_i on $\pi \circ h(a)$.

Spoiler wins if the partial map $a_i \mapsto b_i$ is not a partial isomorphism.

Duplicator wins if it has a strategy to play forever.

If *Duplicator* has a winning strategy, then G and H cannot be distinguished by a Γ -symmetric circuit with support size $\leq k/2$.

Permanent Lower Bound

We construct bipartite graphs $G = (A, B, E)$ and $H = (A, B, E')$ with

- $|A| = |B| = O(k)$
- G and H have *different* numbers of perfect matchings (indeed, they differ by 2^l for some $l > 0$.)
- *Duplicator* wins the k -pebble, $\text{Sym}_A \times \text{Sym}_B$ bijection game on G and H starting with the identity.


Permanent

$$\text{per}(X) = \sum_{\sigma \in \text{Sym}_n} \prod_{i \in [n]} x_{i\sigma(i)}$$

If G is a bipartite graph with *biadjacency matrix* N , then $\text{per}(N)$ is the *number of perfect matchings* in G .

If the *adjacency matrix* of G is M , then

$$\text{per}(M) = \text{per}(N)^2$$



Determinant Lower Bound

We construct a bipartite graph $G = (A, B, E)$ with

- $|A| = |B| = O(k)$
- the bi-adjacency matrix has *non-zero* determinant
- *Duplicator* wins the k -pebble, $\text{Alt}_A \times \text{Alt}_B$ bijection game on two copies of G starting with any bijection swapping two elements of B .

Results

	$\{\text{id}\}$	$\text{Sym}_{[n]}$	$\text{Alt}_{[n]} \times \text{Alt}_{[n]}$	$\text{Sym}_{[n]} \times \text{Sym}_{[n]}$
Det	$O(n^\omega)$	$O(n^3)$ <i>(char 0)</i>	$2^{\Omega(n)}$ <i>(char 0)</i>	N/A
Perm	$O(n^2 2^n)$ VP = VNP?	$2^{\Omega(n)}$ <i>(char 0)</i>	$2^{\Omega(n)}$ <i>(char $\neq 2$)</i>	$2^{\Omega(n)}$ <i>(char $\neq 2$)</i>