A lower bound for homogeneous algebraic branching programs

Sunaina Pati, Chennai Mathematical Institute, Chennai

November 6, 2025

We discuss Professor Mrinal Kumar's *A quadratic lower bound for homogeneous algebraic branching programs* here. We saw in classes and as discussed in Baur and Strassen's paper that any arithmetic circuit which computes the polynomial $P_{n,d} = x_1^d + \cdots + x_n^d$ has at least gates $\Omega(n \log d)$. In this paper, Mrinal works with a special kind of ABP called homogenous ABP (we still need to formally define what ABPs are!) and show that any homogeneous ABP computing $P_{n,d}$ will have atleast $\Omega(nd)$ many vertices. Let us first define what ABPs are.

Definition (Algebraic Branching Program)**.** An *algebraic branching program (ABP)* in variables $\{x_1, x_2, \ldots, x_n\}$ over a field $\mathbb{F}$ is a directed acyclic graph (DAG) with a designated starting vertex s (having in-degree zero) and a designated end vertex t (having out-degree zero). Each edge between any two vertices is labeled by an affine form from $\mathbb{F}[x_1, x_2, \ldots, x_n]$, i.e., a polynomial of the form $\alpha_0 + \alpha_1 x_1 + \cdots + \alpha_n x_n \quad$ where $\alpha_i \in \mathbb{F}$.

Let π be a path in the ABP. The *weight* of the path π, denoted by $w(\pi)$, is the polynomial obtained by taking the product of all the affine (or linear) forms labeling the edges along π.

The polynomial computed at a particular vertex v in the ABP, denoted by $[v]$, is defined as the sum of the weights of all paths from the source vertex s to v, i.e.,

$$[v] = \sum_{\pi : s \rightsquigarrow v} w(\pi).$$

In particular, the polynomial computed by the ABP is $[t]$, where t is the sink (end) vertex.
The *size* of the ABP is the number of vertices in the underlying directed acyclic graph. An ABP is said to be *homogeneous* if the polynomial computed at every node of the ABP is homogeneous.

The main motivation for considering homogeneous ABPs is that they allow for a particularly well-structured or "nice" layering of vertices, a layering that takes advantage of the homogeneity property to separate degrees across layers in a clean way.

At a high level, the proof proceeds as follows. Each such nice layering contains at least $\lfloor n/2 \rfloor$ vertices, and we can show that there must be at least $(d-1)$ layers, each disjoint from the others. So, we get that the total number of vertices in any homogeneous ABP computing $P_{n,d}$ is at least $\lceil n/2 \rceil (d-1) + 2$, where the additional 2 accounts for the source and sink vertices.

Moreover, over an algebraically closed field $\mathbb{F}$, there exists a homogeneous ABP of exactly this size that computes $P_{n,d}$. The construction is as follows: consider $\lceil \frac{n}{2} \rceil$ disjoint paths from the source vertex s to the sink vertex t. The i-th path computes the polynomial $x_{2i}^d + x_{2i+1}^d$. Since $\mathbb{F}$ is algebraically closed, we can factor $x_{2i}^d + x_{2i+1}^d$ into d linear terms and use these as the edge weights of the ABP along the i-th path.

So, we will have to carefully look at two things:

- the manner in which the vertices can be layered or partitioned, and
- why each such partition must contain at least $\lceil n/2 \rceil$ vertices.

The first one follows by a few combinatorial argument and the second one follows from basic variety dimension argument with the help of Algebraic Geometry.

We first deal with how we can partition the vertices.

1 Partitioning the vertices

Intuitively, while computing the polynomial $P_{n,d}$, the higher-degree terms generated within a homogeneous ABP eventually cancel out with other higher-degree terms, and hence do not contribute to the final output. This leads to a natural question:

If there exists a homogeneous ABP that computes $P_{n,d}$, can we construct an ABP that also computes $P_{n,d}$ but never produces "during computation" any term of degree greater than d?

The notion of "also computing" higher-degree terms is captured by the following definition.

Definition (Formal Degree)**.** We say that an ABP has *formal degree at most d* if the number of non-constant edge weights on any path from the source vertex s to the sink vertex t is at most d. More generally, we define the *formal degree* of any vertex v in an ABP to be the maximum number of non-constant edge weights along any path from s to v.

Mrinal shows in his paper that any homogeneous ABP computing a degree-d polynomial can be transformed into an ABP of formal degree d, without increasing the number of vertices.

Lemma 1.1. *Let B be a homogeneous ABP with r vertices that computes a homogeneous polynomial P of degree d. Then there exists an ABP B' computing P such that B' has at most r vertices and has formal degree at most d.*

Then, the partitioning the vertices based on the degree of the polynomial being computed at the vertex. Formally, for any $i \in \{1, 2, 3, \ldots, d-1\}$, let $S_i = \{u_1, u_2, \ldots, u_m\}$ be the set of all vertices in B that compute polynomials of degree equal to i. Then, there exist polynomials $h_1, h_2, \ldots, h_m$ and R such that

$$P_{n,d} = \sum_{j=1}^{m} [u_j] \cdot h_j + R,$$

where each h_j corresponds to the polynomial computed along paths from u_j to the sink vertex t, and R corresponds to the contribution of all paths from s to t that do not pass through any vertex u_j.

Moreover, by the above lemma, if there is a homogeneous ABP computing $P_{n,d}$, then there is a ABP with formal degree atmost d and is computing $P_{n,d}$ with the same number of vertices. Hence, we will prove the following.

Theorem 1.1. *Let B be an algebraic branching program of formal degree at most d over $\mathbb{C}$ which computes the polynomial $P_{n,d}(x)$. Then, the number of vertices in B is at least $\Omega(nd)$.*

We now prove lemma 1.1. The proof idea is deleting edges or modify it when degree is becoming more than d.

Proof. (of lemma 1.1) We start with the ABP B and construct a new ABP B' by modifying or deleting some of the edge weights in B such that the polynomial computed by B' is the same as that computed by B. Moreover, B' will have the additional property that for every vertex v, the degree of the homogeneous polynomial computed at v equals the formal degree of v.

The proof proceeds by induction on the vertices of B, processed in a topological order (which exists because the underlying graph is a DAG), that is, we process a vertex v only after processing every vertex u such that (u, v) is an edge in B.

Base case. The base case is trivial, since there is nothing to do for the starting vertex s.

Induction step. Suppose we are processing a vertex v. Let $u_1, u_2, \ldots, u_m$ be all the vertices such that (u_j, v) is an edge in B. Let the weight of (u_j, v) be $\ell_j + \alpha_j$, where ℓ_j is a homogeneous linear form (possibly identically zero) and α_j is a constant. Let d_{u_j} denote the degree of the polynomial $[u_j]$. Then we have

$$[v] = \sum_{j=1}^{m} [u_j] \cdot (\ell_j + \alpha_j).$$

We separate the summation based on the degrees d_v and d_{u_j}:

$$[v] = \sum_{j:d_v<d_{u_j}} [u_j] \cdot (\ell_j + \alpha_j) + \sum_{j:d_v=d_{u_j}} [u_j] \cdot (\ell_j + \alpha_j) + \sum_{j:d_v>d_{u_j}} [u_j] \cdot (\ell_j + \alpha_j).$$

Since both $[v]$ and $[u_j]$ are homogeneous polynomials, and $[v]$ has degree d_v, the higher degree terms are not contributing anything, so we can simply put 0 in place of them:

$$[v] = \sum_{j:d_v<d_{u_j}} [u_j] \cdot 0 + \sum_{j:d_v=d_{u_j}} [u_j] \cdot (\alpha_j) + \sum_{j:d_v>d_{u_j}} [u_j] \cdot (\ell_j + \alpha_j).$$

Thus, in the new ABP B', we modify the edges as follows:

- For every vertex u_j such that $d_{u_j} > d_v$, we delete the edge (u_j, v).
- For every vertex u_j such that $d_{u_j} = d_v$, and the edge (u_j, v) has weight $\ell_j + \alpha_j$, we relabel the edge with α_j.

And this works. □

Now, let us try to work on the proof of theorem 1.1, the main theorem. The way Mrinal proves it is by first agruing that the polynomial R in $P_{n,d} = \sum_{j=1}^{m} [u_j] \cdot h_j + R$ has degree atmost $d - 1$ and then uses some AG lemmas to show that m is atleast $\lceil n/2 \rceil$.

Infact, he states the following.

Lemma 1.2. *Let B be an algebraic branching program of formal degree at most d with b vertices, which computes an n-variate polynomial P of degree d. For any $i \in \{1, 2, 3, \ldots, d-1\}$, let $S_i = \{u_1, u_2, \ldots, u_m\}$ be the set of all vertices in B that compute polynomials of degree exactly i.*

Then, there exist polynomials $h_1, h_2, \ldots, h_m$ and R, each of degree at most $d-1$, such that

$$P = \sum_{j=1}^{m} [u_j] \cdot h_j + R.$$

We will prove this lemma in some time.

2 Lower bound of partition size

By the above lemma, we have that for any $i \in \{1, 2, 3, \ldots, d-1\}$, we have $P = \sum_{j=1}^{m} [u_j] \cdot h_j + R$. We want to show that $m > n/2$.

Let us try to look at a simplified version where we remove R, that is $R = 0$, and try to approach this and see how we can solve it.

Claim 2.1. *For every set $\{Q_1, Q_2, \ldots, Q_k, R_1, R_2, \ldots, R_k\}$ of homogeneous polynomials of degree at least 1, if*

$$P_{n,d} = \sum_{i=1}^{k} Q_i \cdot R_i,$$

then $k \geq \frac{n}{2}$.

The proof wil use some basic properties of varities which we shall state here:

- For any set $\{Q_1, Q_2, \ldots, Q_t\}$ of polynomials in $\mathbb{F}[x_1, x_2, \ldots, x_n]$, we define the affine variety (or simply, the variety) of $Q_1, Q_2, \ldots, Q_t$ in $\mathbb{F}^n$ as

$$V(Q_1, Q_2, \ldots, Q_t) = \{a \in \mathbb{F}^n \mid \forall i \in [t],\, Q_i(a) = 0\}.$$

- For any set $\{Q_1, Q_2, \ldots, Q_t\}$ of polynomials in $\mathbb{F}[x_1, x_2, \ldots, x_n]$, the ideal generated by $Q_1, Q_2, \ldots, Q_t$ is defined as

$$I(Q_1, Q_2, \ldots, Q_t) = \left\{ \sum_{i=1}^{t} R_i \cdot Q_i \;:\; R_i \in \mathbb{F}[x_1, x_2, \ldots, x_n] \right\}.$$

- For any variety $V \subseteq \mathbb{F}^n$, the ideal associated with V, denoted by $I(V)$, is defined as

$$I(V) = \{R \in \mathbb{F}[x_1, x_2, \ldots, x_n] \mid \forall a \in V,\, R(a) = 0\}.$$

- Let $\mathbb{F}$ be an algebraically closed field, and let S be a set of polynomials in n variables over $\mathbb{F}$ such that $|S| \leq n$. Let $V = V(S)$ be the set of common zeros of the polynomials in S, i.e.,

$$V = \{a \in \mathbb{F}^n \mid \forall f \in S,\, f(a) = 0\}.$$

  If V is non-empty, then the dimension of $V(S)$ is at least $n - |S|$.

- Let $\mathbb{F}$ be an algebraically closed field, and let $V_1, V_2 \subseteq \mathbb{F}^n$ be affine varieties such that $V_1 \subseteq V_2$. Then,

$$\dim(V_1) \leq \dim(V_2).$$

- Let $\mathbb{F}$ be an algebraically closed field, and let $V \subseteq \mathbb{F}^n$ be an affine variety. Then, the dimension of V is zero if and only if V is finite.

Now, we will prove Claim 2.1

Proof. Consider the set of zeros Q_i's and R_i's. That is, look at $V_1 = V(Q_1, ;Q_k, R_1, \dots, R_k)$. This set is non empty and Q_i and R_i's are homogenous of degree 1, so vector 0 is in the variety. And so, the dimension of $V_1 \geq n - 2k$. To show, $k \geq n/2$. We will show that there is a variety V_2 such that $V_1 \subset V_2$ and dimension of V_2 is 0, which will imply dimension of V_1 is atmost 0 which will imply our claim. To show V_2 is of dimension 0, we will show that V_2 is finite.

We claim that any element of V_1 will be a root of $P_{n,d}$ of multiplicity atleast 2. A better way to say the above statement is, any element of V_1 will be element of variety $V_2 = V(\frac{\partial P(n,d)}{\partial x_j}, \forall j \in [n]) = V(dx_1^{d-1}, \dots, dx_n^{d-1})$. To see this, note that,

$$\frac{\partial P(n,d)}{\partial x_j} = \sum_{i=1}^{k} \frac{\partial Q_i}{\partial x_j} \cdot R_i + \sum_{i=1}^{k} Q_i \cdot \frac{\partial R_i}{\partial x_j}.$$

But clearly, $V_2 = V(dx_1^{d-1}, \dots, dx_n^{d-1}) = \{(0, \dots, 0)\}$. So $|V_2|$ is finite. And we are done. □

Note that, we didn't use the homgenity of Q_i's and R_i's at all. What we used is that they do not have a constant term.

Now, we try to prove $m > n/2$ for general R. The idea is almost the same, except now we will be dealing with the variety $V_2 = V(\frac{\partial P(n,d)}{\partial x_j} - \frac{\partial R}{\partial x_j}, \forall j \in [n])$. Turns out this variety is also finite and hence dimension 0 (We showed it in Lemma 2.2). And hence, we will get that dimension of V_2 is 0. Let us formally state what we discussed above.

Lemma 2.1. *Let $\{Q_1, Q_2, \dots, Q_k, R_1, R_2, \dots, R_k\}$ be a set of polynomials in $\mathbb{C}[x]$ such that the set of their common zeros*

$$V = V(Q_1, Q_2, \dots, Q_k, R_1, R_2, \dots, R_k)$$

is non-empty. Let P be any polynomial in $\mathbb{C}[x]$ of degree at most $d-1$, such that

$$P(n,d) = P + \sum_{i=1}^{k} Q_i \cdot R_i.$$

Then, $k \geq \frac{n}{2}$.

Proof. Since $V = V(Q_1, Q_2, \dots, Q_k, R_1, R_2, \dots, R_k)$ is non-empty, the dimension of V is atleast $n - 2k$. Note that,

$$\frac{\partial P(n,d)}{\partial x_j} = \frac{\partial P}{\partial x_j} + \sum_{i=1}^{k} \frac{\partial Q_i}{\partial x_j} \cdot R_i + \sum_{i=1}^{k} Q_i \cdot \frac{\partial R_i}{\partial x_j}.$$

Hence,

$$V(Q_1, Q_2, \dots, Q_k, R_1, R_2, \dots, R_k) \subseteq V\left(dx_j^{d-1} - \frac{\partial P}{\partial x_j} : \forall j \in [n]\right).$$

But note that,

$$V\left(dx_j^{d-1} - \frac{\partial P}{\partial x_j} : \forall j \in [n]\right) = V\left(x_j^{d-1} - \frac{\partial P}{d\partial x_j} : \forall j \in [n]\right).$$

We get that $\frac{\partial P}{d\partial x_j}$ is atmost $d-2$ degree polynomial, so by lemma 2.2, we get that the variety V_2 is finite and hence 0 dimension and so $k \geq n/2$. □

Lemma 2.2. *Let d be a positive natural number. For every choice of polynomials $g_1, g_2, \ldots, g_n \in \mathbb{C}[x]$ of degree at most $d-1$, the cardinality of variety*

$$V(x_1^d - g_1,\, x_2^d - g_2,\, \ldots,\, x_n^d - g_n)$$

is finite.

Note that this helps in proving lemma 2.1 as put $d = d-1$ and $g_i = \frac{\partial P}{d\partial x_i}$. And degree of P is atmost $d-1$, so degree of $\frac{\partial P}{d\partial x_i}$ is atmost $d-2$. Now, we present the proof.

Proof. Let $V = V(x_1^d - g_1,\, x_2^d - g_2,\, \ldots,\, x_n^d - g_n)$. We show that the cardinality of V is at most

$$T = \binom{n + n(d-1)}{n}.$$

We prove this via contradiction. Suppose the cardinality of V is larger than T. Then we focus our attention on an arbitrary subset $S \subseteq V$ of size equal to $T+1$.

Now, consider the linear space of polynomial functions from S to $\mathbb{C}$. Two polynomials functions are equall if they take the same value at every element in S.

Clearly, the dimension of this linear space must be at least $T+1$, since the indicator function of every point in S can be expressed as a sufficiently high-degree polynomial(by lagrange interpolation), and these polynomials are linearly independent. The dimension of these function is bounded by the cardinality of the quotient ring $\frac{\mathbb{C}[x_1,\ldots,x_n]}{I(S)}$

We now argue that the dimension of the linear space of all polynomial functions from V to $\mathbb{C}$ (and therefore from S to $\mathbb{C}$) is upper bounded by T. This completes the proof by contradiction. Again, the dimension of the polynomial functions from V to $\mathbb{C}$ is bounded by the cardinality of the ring $\frac{\mathbb{C}[x_1,\ldots,x_n]}{I(V)}$.

We claim, that any polynomial P can be expressed as polynomial P' of degree less than $n(d-1)$ mod $I(V)$.

Claim 2.2. *Let P be any polynomial in $\mathbb{C}[x]$. Then there exists a polynomial P' of degree at most $n(d-1)$ such that*

$$P = P' \mod I(V).$$

Proof. Clearly, for every $i \in [n]$, we have $x_i^d - g_i \in I(V)$.

If degree of P greater than $n(d-1)$ then done. If not, then for every $x_i \in x$, we can obtain a new polynomial from P by replacing every occurrence of x_i^d by g_i is equivalent to P modulo the ideal $I(V)$.

Hence, we can keep on performing this replacement while still maintaining equivalence modulo the ideal I. This process terminates eventually, since each occurrence of x_i^d is being replaced by a polynomial of strictly smaller degree.

Let P'' be the polynomial obtained when the process terminates. It follows that the individual degree of every variable x_i in P'' is upper bounded by $d-1$, and hence the total degree of P'' is at most $n(d-1)$. This proves the claim. □

Therefore, the space of all polynomial functions from V to $\mathbb{C}$ is spanned by a subset of polynomials in $\mathbb{C}[x]$ of degree at most $n(d-1)$. Hence, the dimension of this linear space is at most the number of monomials of degree at most $n(d-1)$ in n variables, which is equal to T.

□

So, we showed that each partition will be of size $n/2$. And clearly the way we have paritioned, we will have each partitioning disjoint from each other.

Now, we prove Lemma 1.2 and we will be done.

Proof. (Proof of Lemma 1.2) As discussed before h_j is the polynomial given by the sum of weights of all paths from u_j to t.

Now, we claim that the degree of h_j is at most $d - d_{u_j}$. This follows from the fact that if the degree of h_j were larger than $d - d_{u_j}$, then the formal degree of t would be larger than d, which would contradict the hypothesis that B is of formal degree at most d.

Note that deleting S_i from the ABP will not affect the polynomial R. Now, we show that degree of R is atmost $d-1$.

Let B' be the ABP obtained from B by deleting all vertices in the set S_i. Denote by R_t the polynomial computed by B'.

Consider any path $(s, v_1, v_2, \ldots, v_k, t)$ from s to t in B'. Note that all these vertices also appear in the original ABP B. Let v_j be the first vertex along this path whose degree in B is at least $i+1$. By construction, the degree of v_{j-1} in B must then be at most $i-1$, since all vertices of degree exactly i were deleted in forming B'.

It follows that the degree of the monomials in the weight of the path $(s, v_1, v_2, \ldots, v_k, t)$ is at most $(i-1) + \ell + 1$, where ℓ denotes the maximum number of non-constant edge weights on any subpath from v_j to t in B.

We now claim that $\ell \leq d - i - 1$. Indeed, if $\ell \geq d - i$, then since degree of $[v_j]$ is atleast $i+1$ then there exists a path from s to v_j with atleast $i+1$ non constant edges and so consider that path of s to v_j and then this path from v_j to t. So we get a path from s to t with atleast $d+1$ non constant edges. A contradiction. □

So, we are done with the proof of Theorem 1.1.

3 Generalising this proof for general ABPs

Homogeneosity (and formal degree) was used significantly in all the proofs.

However, Mrinal et all, did prove the same quadratic bound for general ABPs which shows that any sufficiently small ABP computing the polynomial $\sum_{i=1}^{n} x_i^n$ can be depth-reduced to an essentially homogeneous ABP of comparable size, which computes the polynomial $\sum_{i=1}^{n} x_i^n + \varepsilon(x)$, for a structured "error polynomial" $\varepsilon(x)$. And then, they show that same lower bound will hold for this polynomial too.