Sums of read-once formulas: How many summands are necessary?☆

Meena Mahajan[*,a], Anuj Tawari[a]

[a] *The Institute of Mathematical Sciences IMSc, HBNI, Chennai, India.*

Abstract

An arithmetic read-once formula (ROF) is a formula (circuit of fan-out 1) over $+, \times$ where each variable labels at most one leaf. Every multilinear polynomial can be expressed as the sum of (possibly exponentially many) ROFs. In this work, we prove, for certain multilinear polynomials, a tight lower bound on the number of summands in such an expression.

Key words: Arithmetic formulas, read-once polynomials, hardness of representation

1. Introduction

Read-once formulas (ROF) are formulas (circuits of fan-out 1) in which each variable appears at most once. A formula computing a polynomial that depends on all its variables must read each variable at least once. Therefore, ROFs compute some of the simplest possible functions that depend on all of their variables. The polynomials computed by such formulas are known as read-once polynomials (ROPs). Since every variable is read at most once, ROPs are multilinear. (A polynomial is said to be multilinear if the individual degree of each variable is at most one.) But not every multilinear polynomial is a ROP. For example, $x_1x_2 + x_2x_3 + x_1x_3$.

We investigate the following question: Given an n-variate multilinear polynomial, can it be expressed as a sum of at most k ROPs? It is easy

☆Preliminary versions of this paper appear in the Proceedings of the 11th Computer Science Symposium in Russia CSR, LNCS Vol. 9691 pp 266–279 and as Technical Reports on ECCC and CoRR [12].

*Corresponding Author

to see that every bivariate multilinear polynomial is a ROP. Any tri-variate multilinear polynomial can be expressed as a sum of 2 ROPs. With a little thought, we can obtain a sum-of-3-ROPs expression for any 4-variate multilinear polynomial. An easy induction on n then shows that any n-variate multilinear polynomial, for $n \geq 4$, can be written as a sum of at most $3 \times 2^{n-4}$ ROPs; see Proposition 5. Also, the sum of two multilinear monomials is a ROP, so any n-variate multilinear polynomial with M monomials can be written as the sum of $\lceil M/2 \rceil$ ROPs. We ask the following question: Does there exist a strict hierarchy among k-sums of ROPs? Formally,

Problem 1. Consider the family of n-variate multilinear polynomials. For $1 < k \leq 3 \times 2^{n-4}$, is $\sum^k \cdot$ROP strictly more powerful than $\sum^{k-1} \cdot$ROP? If so, what explicit polynomials witness the separations?

We answer this affirmatively for $k \leq \lceil n/2 \rceil$. In particular, for $k = \lceil n/2 \rceil$, there exists an explicit n-variate multilinear polynomial which cannot be written as a sum of less than k ROPs but it admits a sum-of-k-ROPs representation.

Note that n-variate ROPs are computed by linear sized formulas. Thus if an n-variate polynomial p is in $\sum^k \cdot$ROP, then p is computed by a formula of size $O(kn)$ where every intermediate node computes a multilinear polynomial. Since superpolynomial lower bounds are already known for the model of multilinear formulas [14], we know that for those polynomials (including the determinant and the permanent), a $\sum^k \cdot$ROP expression must have k at least quasi-polynomial in n. However the best upper bound on k for these polynomials is only exponential in n, leaving a big gap between the lower and upper bound on k. A lesser but still significant gap also exists in the known exponential lower bound for sums of ROPs; in [13] it is shown that a certain polynomial, explicitly described by Raz and Yehudayoff in [15], requires $2^{\Omega(n^{1/3}/\log n)}$ ROP summands, while 2^n summands is anyway sufficient. On the other hand, our lower bound is provably tight.

A counting argument (see Proposition 7) shows that a random multilinear polynomial requires exponentially many ROPs; there are multilinear polynomials requiring $k = \Omega(2^n/n^2)$. Our general upper bound on k is $O(2^n)$, leaving a gap between the lower and upper bound. One challenge is to close this gap.

A natural question to ask is whether stronger lower bounds than the above result can be proven. In particular, to separate $\sum^{k-1} \cdot$ROP from $\sum^k \cdot$ROP, how many variables are needed? Our hierarchy result says that

$2k-1$ variables suffice, but there may be simpler polynomials (with fewer variables) witnessing this separation. We demonstrate another technique which improves upon the previous result for $k=3$, showing that 4 variables suffice. In particular, we show that over the field of reals, there exists an explicit multilinear 4-variate multilinear polynomial which cannot be written as a sum of 2 ROPs. This lower bound is again tight, as there is a sum of 3 ROPs representation for every 4-variate multilinear polynomial.

Our results and techniques

We now formally state our main results.

The first main result establishes the strict hierarchy among k-sums of ROPs.

Theorem 1. *For each $n \geq 1$, the n-variate degree $n-1$ symmetric polynomial S_n^{n-1} cannot be written as a sum of less than $\lceil n/2 \rceil$ ROPs, but it can be written as a sum of $\lceil n/2 \rceil$ ROPs.*

The idea behind the lower bound is that if $g = S_n^{n-1}$ can be expressed as a sum of less than $\lceil n/2 \rceil$ ROFs, then one of the ROFs can be eliminated by taking partial derivative with respect to one variable and substituting another by a field constant. We then use the inductive hypothesis to arrive at a contradiction. This approach necessitates a stronger hypothesis than the statement of the theorem, and we prove this stronger statement in Lemma 18 as part of Theorem 21.

This result separates $\sum^3 \cdot$ROP from $\sum^2 \cdot$ROP via the polynomials S_5^4 and S_6^5. Our second main result shows that $\sum^3 \cdot$ROP is also separated from $\sum^2 \cdot$ROP by a 4-variate multilinear polynomial.

Theorem 2. *There is an explicit 4-variate multilinear polynomial f which cannot be written as the sum of 2 ROPs over $\mathbb{R}$.*

The proof of this theorem mainly relies on a structural lemma (Lemma 25) for sum of 2 read-once formulas. In particular, we show that if f can be written as a sum of 2 ROPs then one of the following must be true:

1. Some 2-variate restriction is a linear polynomial.
2. There exist variables $x_i, x_j \in \mathrm{Var}(f)$ such that the polynomials x_i, x_j, $\partial_{x_i}(f)$, $\partial_{x_j}(f)$, 1 are linearly dependent.
3. We can represent f as $f = l_1 \cdot l_2 + l_3 \cdot l_4$ where (l_1, l_2) and (l_3, l_4) are variable-disjoint linear forms.

Checking the first two conditions is easy. For the third condition we use the commutator of f, introduced in [16], to find one of the l_i's. The knowledge of one of the l_i's suffices to determine all the linear forms. Finally, we construct a 4-variate polynomial which does not satisfy any of the above mentioned conditions. This construction does not work over algebraically closed fields. We do not yet know how to construct an explicit 4-variate multilinear polynomial not expressible as the sum of 2 ROPs over such fields, or even whether such polynomials exist.

Related work

Despite their simplicity, ROFs have received a lot of attention both in the arithmetic as well as in the Boolean world [8, 5, 3, 4, 16, 17]. The most fundamental question that can be asked about polynomials is the polynomial identity testing (PIT) problem: Given an arithmetic circuit $\mathcal{C}$, is the polynomial computed by $\mathcal{C}$ identically zero or not. PIT has a randomized polynomial time algorithm: Evaluate the polynomial at random points. It is not known whether PIT has a deterministic polynomial time algorithm. In 2004, Kabanets and Impagliazzo established a connection between PIT algorithms and proving general circuit lower bounds [10]. Similar results are known for some restricted classes of arithmetic circuits, for instance, constant-depth circuits [6, 1]. However, consider the case of multilinear formulas. Even though strong lower bounds are known for this model, there is no efficient deterministic PIT algorithm. (Notice that multilinear depth 3 circuits are a special case of this model.) For this reason, PIT was studied for the weaker model of sum of read-once formulas.

Shpilka and Volkovich gave a deterministic PIT algorithm for the sum of a small number of ROPs [17]. Interestingly, their proof uses a lower bound for a weaker model, that of 0-justified ROFs (setting some variables to zero does not kill any other variables). In particular, they show that the polynomial $\mathcal{M}_n = x_1x_2\cdots x_n$, consisting of just a single monomial, cannot be represented as a sum of less than $n/3$ weakly justified ROPs. More recently, Kayal showed that if $\mathcal{M}_n$ is represented as a sum of powers of low degree (at most d) polynomials, then the number of summands is at least $\exp(\Omega(n/d))$ [11]. This lower bound, along with the arguments in [17], yields a sub-exponential time PIT algorithm for multilinear polynomials. This can be further extended to arbitrary polynomials written as sum of powers of low degree polynomials, using the ideas in [7]. Our lower bound from Theorem 1 is independent of both these lower bounds (0-justified ROFs from [17], and sums of powers of low-degree polynomials from [11]) and is provably tight. An interesting question is whether it can be used to give a PIT algorithm for sums of k ROPs, when k is linear in n.

Similar to ROPs, one may also study read-restricted formulas. For any number k, RkFs are formulas that read every variable at most k times. For $k \geq 2$, RkFs need not be multilinear, and thus are strictly more powerful than ROPs. However, even when restricted to multilinear polynomials, they are more powerful; in [2], Anderson, Melkebeek and Volkovich show that there is a multilinear n-variate polynomial in R2F requiring $\Omega(n)$ summands when written as a sum of ROPs.

Organization

The paper is organized as follows. In Section 2 we give the basic definitions and notations. In Section 3, we establish Theorem 1. showing that the hierarchy of k-sums of ROPs is proper. In Section 4 we establish Theorem 2, showing an explicit 4-variate multilinear polynomial that is not expressible as the sum of two ROPs. We conclude in Section 5 with some further questions that are still open.

2. Preliminaries

For a positive integer n, we denote $[n] = \{1, 2, \ldots, n\}$. For a polynomial f, by $\mathrm{Var}(f)$ we mean the set of variables occurring in f. For a polynomial $f(x_1, x_2, \ldots, x_n)$, a variable x_i and a field element α, we denote by $f \mid_{x_i=\alpha}$ the polynomial resulting from setting $x_i = \alpha$. Let f be an n-variate polynomial. We say that g is a k-variate restriction of f if g is obtained by setting some variables in f to field constants and $|\mathrm{Var}(g)| \leq k$. A set of polynomials $f_1, f_2, \ldots, f_k$ over the field $\mathbb{F}$ is said to be linearly dependent if there exist constants $\alpha_1, \alpha_2, \ldots, \alpha_k$ such that $\displaystyle\sum_{i \in [k]} \alpha_i f_i = 0$.

The n-variate degree k elementary symmetric polynomial, denoted S_n^k, is defined as follows:

$$S_n^k(x_1, \ldots, x_n) = \sum_{A \subseteq [n], |A| = k} \prod_{i \in A} x_i.$$

A circuit is a directed acyclic graph with variables and field constants labeling the leaves, field operations $+, \times$ labeling internal nodes, and a designated sink node (a node with out-degree zero). Each node naturally computes a polynomial; the polynomial at the designated sink node is the polynomial computed by the circuit. If the underlying undirected graph is a tree, then the circuit is called a formula. A formula is said to be read-k if each variable appears as a leaf label at most k times.

For read-once formulas, it is more convenient to use the following "normal form" from [17].

Definition 3 (Read-once formulas [17]). *A read-once arithmetic formula (ROF) over a field $\mathbb{F}$ in the variables $\{x_1, x_2, \ldots, x_n\}$ is a binary tree as follows. The leaves are labeled by variables and internal nodes by $\{+, \times\}$. In addition, every node is labeled by a pair of field elements $(\alpha, \beta) \in \mathbb{F}^2$. Each input variable labels at most once leaf. The computation is performed in the following way. A leaf labeled by x_i and (α, β) computes $\alpha x_i + \beta$. If a node v is labeled by $\star \in \{+, \times\}$ and (α, β) and its children compute the polynomials f_1 and f_2, then v computes $\alpha(f_1 \star f_2) + \beta$.*

We say that f is a read-once polynomial (ROP) if it can be computed by a ROF, and is in $\sum^k \cdot$ROP if it can be expressed as the sum of at most k ROPs.

Definition 4. *Let $\mathbb{F}$ be a field, and let f be a polynomial in $\mathbb{F}[x_1, \ldots, x_n]$. By* SummandsROP(f) *we denote the minimum $k \in \mathbb{N}$ such that $f \in \sum^k \cdot$ROP.*

Proposition 5. *For every n-variate multilinear polynomial f,* $\mathrm{SummandsROP(f)} \leq \lceil 3 \times 2^{\mathrm{n}-4} \rceil$.

Proof For $n = 1, 2, 3$ this is easy to see.

For $n = 4$, let $f(X)$ be given by the expression $\sum_{S \subseteq [4]} a_S x_S$, where x_S denotes the monomial $\prod_{i \in S} x_i$. We want to express f as $f_1 + f_2 + f_3$, where each f_i is an ROP. If there are no degree 2 terms, we use the following:

$$
\begin{aligned}
f_1 &= a_\emptyset + a_1 x_1 + a_2 x_2 + a_3 x_3 + a_4 x_4 \\
f_2 &= x_1 x_2 (a_{123} x_3 + a_{124} x_4) \\
f_3 &= x_3 x_4 (a_{134} x_1 + a_{234} x_2 + a_{1234} x_1 x_2)
\end{aligned}
$$

Otherwise, assume without loss of generality that $a_{13} \neq 0$. Then define

$$
\begin{aligned}
f_1 &= \left[\sum_{S \subseteq [2]} a_S \prod_{i \in S} x_i\right] + \left[\sum_{\emptyset \neq S \subseteq \{3,4\}} a_S \prod_{i \in S} x_i\right] \\
f_2 &= (a_{13}x_1 + a_{23}x_2 + a_{123}x_1x_2) \cdot \left(\frac{a_{14}}{a_{13}}x_4 + x_3 + \frac{a_{134}}{a_{13}}x_3x_4\right) \\
f_3 &= x_2x_4 \left[\left(a_{24} - \frac{a_{14}a_{23}}{a_{13}}\right) + x_1\left(a_{124} - \frac{a_{14}a_{123}}{a_{13}}\right)\right. \\
&\qquad \left. + x_3\left(a_{234} - \frac{a_{134}a_{23}}{a_{13}}\right) + x_1x_3\left(a_{1234} - \frac{a_{134}a_{123}}{a_{13}}\right)\right]
\end{aligned}
$$

Since any bivariate multilinear polynomial is a ROP, each f_i is indeed an ROP.

For $n > 4$, express f as $x_n g + h$ where $g = \partial_{x_n} f$ and $h = f\mid_{x_n=0}$, and use induction, along with the fact that g does not have variable x_n. $\square$

Proposition 6. *For every n-variate multilinear polynomial f with M monomials,* $\mathrm{SummandsROP(f)} \leq \lceil \frac{\mathrm{M}}{2} \rceil$.

Proof For $S \subseteq [n]$, let x_S denote the multilinear monomial $\prod_{i \in S} x_i$. For any $S, T \subseteq [n]$, the polynomial $ax_S + bx_T$ equals $x_{S\cap T}(ax_{S\setminus T} + bx_{T\setminus S})$ and hence is an ROP. Pairing up monomials in any way gives the $\lceil \frac{M}{2} \rceil$ bound. $\square$

Proposition 7. *Fix any field $\mathbb{F}$. There exists a family of multilinear polynomials $(f_n)_{n>0}$ with each $f_n \in \mathbb{F}[x_1, \ldots, x_n]$ such that* $\mathrm{SummandsROP(f_n)} = \Omega\left(\frac{2^{\mathrm{n}}}{\mathrm{n}^2}\right)$.

Proof Let $\mathcal{M}$ denote the set of multilinear polynomials in $\mathbb{F}[x_1, \ldots, x_n]$ where each coefficient is either zero or one. Then $|\mathcal{M}| = 2^{2^n}$. We will show that unless $s \in \Omega\left(\frac{2^n}{n^2}\right)$, the number of polynomials in $\mathcal{M}$ computable by $\sum^s \cdot$ROF is strictly less than this.

We use the strategy from [9]; a similar strategy was also used in [18]. Using notation from [9], we call a circuit or formula with no field constants a *skeleton*. From any circuit or formula, we can obtain a skeleton by simply replacing each occurrence of a field element by a fresh variable. Our counting proceeds as follows:

Fix any $s \in \mathbb{N}$. Define the following quantities.

N_1: the number of distinct skeletons arising from $\sum^s \cdot$ROF formulas on n variables. Each skeleton computes a polynomial in the variables $X \cup Z$, where $X = \{x_i \mid i \in [n]\}$ and $Z = \{z_i \mid i \in [t]\}$ for some $t \in O(ns)$.

N_2: the number of polynomials from $\mathcal{M}$ computable by a single skeleton on appropriate instantiation of the z variables.

Then $\sum^s \cdot$ROF expressions can compute at most $N_1 \times N_2$ polynomials in $\mathcal{M}$.

First, we estimate N_1. Note that a $\sum^s \cdot$ROF formula has at most $3ns$ gates apart from the top $+$ gates. (We implicitly unfold an ROF gate f labeled $(\circ, \alpha, \beta)$ and with children g, h into a small sub-formula $\alpha \times (g \circ h) + \beta$, and then replace α, β by fresh z variables.) We use a generous over-estimate for N_1, namely, the number of skeletons of circuits of size $3ns$. We have n variables in X and t variables in Z. Each node in the skeleton can be labeled in at most $n+t+2$ ways (a variable or a gate type), and its children can be chosen in at most $(3ns)^2$ ways. Hence the number of skeletons is no more than $\left[(n+t+2)(3ns)^2\right]^{3ns}$. Since $t = O(ns)$, we conclude that $N_1 = 2^{O(ns(\log n + \log s))}$.

Estimating N_2 is trickier because the field may not be finite, and thus a single skeleton can give rise to infinitely many polynomials. However, we are interested only in polynomials from the finite set $\mathcal{M}$. This can be bounded using a dimension argument as used in [9]. In particular, we use the following result proved in [9]:

Lemma 8 (Lemma 3.5 in [9]). *Let $\mathbb{F}$ be a field. Let $F : \mathbb{F}^n \to \mathbb{F}^m$ be a polynomial map of degree $d > 0$, that is, $F = (F_1, \ldots, F_m)$, each F_i is of degree d. Then $|F(\mathbb{F}^n) \cap \{0,1\}^m| \leq (2d)^n$*

We have a given fixed skeleton corresponding to some $\sum^s \cdot$ROF. It computes some polynomial $\psi(X, Z)$, with $|X| = n$, $|Z| = t$, $t = O(ns)$. By the nature of ROF, ψ is multilinear, and hence can be written in the form

$$\psi(X, Z) = \sum_{S \subseteq [n]} \left(c_S(Z) \prod_{i \in S} x_i \right)$$

where each coefficient $c_S(Z)$ is a multilinear polynomial. These 2^n coefficient polynomials form our polynomial map $F : \mathbb{F}^t \to \mathbb{F}^{2^n}$. Since each coefficient polynomial is multilinear, it has total degree at most t. Hence, from Lemma 8, we conclude that at most $(2t)^t$ 0-1 tuples are produced by this map. Thus the given skeleton can compute at most $(2t)^t$ polynomials from $\mathcal{M}$. Since $t = O(ns)$, we obtain $N_2 = 2^{O(ns(\log n + \log s))}$.

Now that we have estimated N_1 and N_2, we can bound SummandsROP. Assume that for all polynomials $f \in \mathcal{M}$, SummandsROP(f) $\leq$ s. Then $\sum^s \cdot$ROF contains all of $\mathcal{M}$. Hence $N_1 \times N_2 \geq |\mathcal{M}|$, implying $s \geq \Omega\left(\frac{2^n}{n^2}\right)$. □

Remark 1. The above proof works over all fields. However, if the field is finite, Lemma 8 is not needed; the following direct combinatorial argument suffices, and in fact shows a slightly better bound $\mathrm{SummandsROP(f_n)} = \Omega\left(\frac{2^n}{n \log n}\right)$.

A single ROF is a binary tree with at most n leaves, and with labels at each node. A leaf is labeled by a single x variable and a pair of field elements, and an internal node is labeled by a gate type ($+$ or $\times$) and a pair of field elements. The number of binary trees with at most n leaves is $2^{O(n)}$. If the field size is q, then the number of labelings per tree is at most $(nq^2)^n(2q^2)^n$. Hence the number of ROFs is no more than $2^{O(n \log n)}$. A $\sum^s \cdot$ROF formula can be obtained by choosing an ROF for each of the s positions; hence there are at most $2^{O(sn \log n)}$ distinct formulas. This is less than $|\mathcal{M}|$ unless $s = \Omega\left(\frac{2^n}{n \log n}\right)$.

The partial derivative of a polynomial is defined naturally over continuous domains. The definition can be extended in more than one way over finite fields. However, for multilinear polynomials, these definitions coincide. We consider only multilinear polynomials in this paper, and the following formulation is most useful for us: The partial derivative of a polynomial $p \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ with respect to a variable x_i, for $i \in [n]$, is given by $\partial_{x_i}(p) \triangleq p \mid_{x_i=1} - p \mid_{x_i=0}$. For multilinear polynomials, the sum, product, and chain rules continue to hold.

Fact 9 (Useful Fact about ROPs [17]). *The partial derivatives of ROPs are also ROPs.*

Proposition 10 (3-variate ROPs). *Let $f \in \mathbb{F}[x_1, x_2, x_3]$ be a 3-variate ROP. Then there exists $i \in [3]$ and $a \in \mathbb{F}$ such that* $\deg(f \mid_{x_i=a}) \leq 1$.

Proof Assume without loss of generality that $f = f_1(x_1) \star f_2(x_2, x_3) + c$ where $\star \in \{+, \times\}$ and $c \in \mathbb{F}$. If $\star = +$, then for all $a \in \mathbb{F}$, $\deg(f \mid_{x_2=a}) \leq 1$. If $\star = \times$, $\deg(f \mid_{f_1=0}) \leq 1$. □

We will also be dealing with a special case of ROFs called multiplicative ROFs defined below:

Definition 11 (Multiplicative Read-once formulas). *A ROF is said to be a multiplicative ROF if it does not contain any addition gates. We say that f is a multiplicative ROP if it can be computed by a multiplicative ROF.*

Fact 12 ([17] (Lemma 3.10)). *A ROP p is a multiplicative ROP if and only if for any two variables $x_i, x_j \in \mathrm{Var}(p)$, $\partial_{x_i}\partial_{x_j}(p) \neq 0$.*

Multiplicative ROPs have the following useful property, observed in [17]. (See Lemma 3.13 in [17]. For completeness, and since we refer to the proof later, we include a proof sketch here.)

Lemma 13 ([17]). *Let g be a multiplicative ROP with $|\mathrm{Var}(g)| \geq 2$. For every $x_i \in \mathrm{Var}(g)$, there exists $x_j \in \mathrm{Var}(g) \setminus \{x_i\}$ and $\gamma \in \mathbb{F}$ such that $\partial_{x_j}(g) \mid_{x_i=\gamma} = 0$.*

Proof Let φ be a multiplicative ROF computing g. Pick any $x_i \in \mathrm{Var}(g)$. As $|\mathrm{Var}(\varphi)| = |\mathrm{Var}(g)| \geq 2$, φ has at least one gate. Let v be the unique neighbour (parent) of the leaf labeled by x_i, and let w be the other child of v. We denote by $P_v(\bar{x})$ and $P_w(\bar{x})$ the ROPs computed by v and w. Since v is a $\times$ gate and we use the normal form from Defintion 3, P_v is of the form $(\alpha x_i + \beta) \times P_w$ for some $\alpha \neq 0$.

Replacing the output from v by a new variable y, we obtain from φ another multiplicative ROF ψ in the variables $\{y\} \cup \mathrm{Var}(g) \setminus \mathrm{Var}(P_v)$. Let ψ compute the polynomial Q; then $g = Q \mid_{y=P_v}$.

Note that the sets $\mathrm{Var}(Q), \{x_i\}, \mathrm{Var}(P_w)$ are non-empty and disjoint, and form a partition of $\{y, x_1, \ldots, x_n\}$.

By the chain rule, for every variable $x_j \in \mathrm{Var}(P_w)$ we have:

$$\partial_{x_j}(g) = \partial_y(Q) \cdot \partial_{x_j}(P_v) = \partial_y(Q) \cdot (\alpha x_i + \beta) \cdot \partial_{x_j}(P_w)$$

It follows that for $\gamma = -\beta/\alpha$, $\partial_{x_j}(g) \mid_{x_i=\gamma} = 0$. □

Along with partial derivatives, another operator that we will find useful is the commutator of a polynomial. The commutator of a polynomial has previously been used for polynomial factorization and in reconstruction algorithms for read-once formulas, see [16].

Definition 14 (Commutator [16]). *Let $P \in \mathbb{F}[x_1, x_2, \ldots, x_n]$ be a multilinear polynomial and let $i, j \in [n]$. The commutator between x_i and x_j, denoted $\triangle_{ij} P$, is defined as follows.*

$$\triangle_{ij} P = \left(P \mid_{x_i=0, x_j=0}\right) \cdot \left(P \mid_{x_i=1, x_j=1}\right) - \left(P \mid_{x_i=0, x_j=1}\right) \cdot \left(P \mid_{x_i=1, x_j=0}\right)$$

The following property of the commutator will be useful to us.

Lemma 15. *Let $f = l_1(x_1, x_2) \cdot l_2(x_3, x_4) + l_3(x_1, x_3) \cdot l_4(x_2, x_4)$ where the l_i's are linear polynomials. Then l_2 divides $\triangle_{12}(f)$.*

Proof First, we show that $\triangle_{12}(l_3 \cdot l_4) = 0$. Assume $l_3 = Cx_1 + m$ and $l_4 = Dx_2 + n$ where $C, D \in \mathbb{F}$ and m, n are linear polynomials in x_3, x_4 respectively. By definition, $\triangle_{12}(l_3 \cdot l_4) = mn(C+m)(D+n) - m(D+n)(C+m)n = 0$.

Now we write $\triangle_{12} f$ explicitly. Let $l_1 = ax_1 + bx_2 + c$. By definition,

$$\begin{aligned}
\triangle_{12} f &= \triangle_{12}(l_1 l_2 + l_3 l_4) \\
&= (cl_2 + mn)((a+b+c)l_2 + (C+m)(D+n)) - \\
&\quad ((b+c)l_2 + m(D+n)) \cdot ((a+c)l_2 + n(C+m)) \\
&= l_2^2(c(a+b+c) - (a+c)(b+c)) \\
&\quad + l_2(c(C+m)(D+n) + mn(a+b+c) - n(b+c)(C+m) - m(a+c)(D+n))
\end{aligned}$$

It follows that l_2 divides $\triangle_{12} f$. □

3. A proper separation in the $\sum^k \cdot$ROP hierarchy

This section is devoted to proving Theorem 1.

We prove the lower bound for S_n^{n-1} by induction. This necessitates a stronger induction hypothesis, so we will actually prove the lower bound for a larger class of polynomials. For any $\alpha, \beta \in \mathbb{F}$, we define the polynomial $\mathcal{M}_n^{\alpha,\beta} = \alpha S_n^n + \beta S_n^{n-1}$.

Proposition 16. *$\mathcal{M}_n^{\alpha,\beta}$ has the following recursive structure:*

$$\begin{aligned}
(\mathcal{M}_n^{\alpha,\beta}) \mid_{x_n=\gamma} &= \mathcal{M}_{n-1}^{\alpha\gamma+\beta, \beta\gamma} \quad . \\
\partial_{x_n}(\mathcal{M}_n^{\alpha,\beta}) &= \mathcal{M}_{n-1}^{\alpha,\beta} \quad .
\end{aligned}$$

Proof

$$\begin{aligned}
\mathcal{M}_n^{\alpha,\beta} &= \alpha S_n^n + \beta S_n^{n-1} = \alpha \left(\prod_{j\in[n]} x_j \right) + \beta \left(\sum_{i\in[n]} \left[\prod_{j\in[n]\setminus\{i\}} x_j \right] \right) \\
&= \alpha x_n \left(\prod_{j\in[n-1]} x_j \right) + \beta \left(\sum_{i\in[n-1]} x_n \left[\prod_{j\in[n-1]\setminus\{i\}} x_j \right] \right) + \beta \left(\prod_{j\in[n-1]} x_j \right) \\
&= \alpha x_n S_{n-1}^{n-1} + \beta x_n S_{n-1}^{n-2} + \beta S_{n-1}^{n-1}.
\end{aligned}$$

$$\begin{aligned}
\text{Hence } (\mathcal{M}_n^{\alpha,\beta})\mid_{x_n=\gamma} &= (\alpha\gamma+\beta) S_{n-1}^{n-1} + \beta\gamma S_{n-1}^{n-2} = \mathcal{M}_{n-1}^{\alpha\gamma+\beta,\beta\gamma} \\
\text{and } \partial_{x_n}(\mathcal{M}_n^{\alpha,\beta}) &= \alpha S_{n-1}^{n-1} + \beta S_{n-1}^{n-2}.
\end{aligned}$$

□

We show below that each $\mathcal{M}_n^{\alpha,\beta}$ is expressible as the sum of $\lceil n/2 \rceil$ ROPs (Lemma 17); however, for any non-zero $\beta \in \mathbb{F}$, $\mathcal{M}_n^{\alpha,\beta}$ cannot be written as the sum of fewer than $\lceil n/2 \rceil$ ROPs (Lemma 18). At $\alpha = 0$, $\beta = 1$, we get S_n^{n-1}, the simplest such polynomials, establishing Theorem 1.

First we establish the upper bound.

Lemma 17. *For any field $\mathbb{F}$ and $\alpha, \beta \in \mathbb{F}$, the polynomial $f = \alpha S_n^n + \beta S_n^{n-1}$ can be written as a sum of at most $\lceil n/2 \rceil$ ROPs.*

Proof For n odd, this follows immediately from Proposition 6.

If n is even, say $n = 2k$, then define the following polynomials:

$$\begin{aligned}
\text{for } i \in [k-1], \quad f_i &= (x_{2i-1} + x_{2i}) \cdot \left(\prod_{\substack{k\in[n] \\ k\neq 2i,2i-1}} x_k \right) \\
f_k &= (\beta x_{2k-1} + \beta x_{2k} + \alpha x_{2k-1} x_{2k}) \cdot \left(\prod_{\substack{m\in[n] \\ k\neq 2k,2k-1}} x_m \right).
\end{aligned}$$

Then we have $f = \beta(f_1 + f_2 + \ldots + f_{k-1}) + f_k$.

Note that each f_i is an ROP; for $i < k$ this is immediate, and for $i = k$, the factor involving x_{2k-1} and x_{2k} is bivariate multilinear and hence an ROP. Thus we have a representation of f as a sum of $k = \lceil n/2 \rceil$ ROPs. $\square$

The following lemma shows that the above upper bound is indeed optimal.

Lemma 18. *Let $\mathbb{F}$ be a field. For every $\alpha \in \mathbb{F}$ and $\beta \in \mathbb{F} \setminus \{0\}$, the polynomial $\mathcal{M}_n^{\alpha,\beta} = \alpha S_n^n + \beta S_n^{n-1}$ cannot be written as a sum of $k < n/2$ ROPs.*

Proof The proof is by induction on n. The cases $n = 1, 2$ are easy to see. We now assume that $k \geq 1$ and $n > 2k$. Assume to the contrary that there are ROPs $f_1, f_2, \ldots, f_k$ over $\mathbb{F}[x_1, x_2, \ldots, x_n]$ such that $f \triangleq \sum_{m \in [k]} f_m = \mathcal{M}_n^{\alpha,\beta}$.

The main steps in the proof are as follows:

1. Show using the inductive hypothesis that for all $m \in [k]$ and $a, b \in [n]$, $\partial_{x_a} \partial_{x_b}(f_m) \neq 0$.
2. Conclude that for all $m \in [k]$, f_m must be a multiplicative ROP. That is, the ROF computing f_m does not contain any addition gate.
3. Use the multiplicative property of f_k to show that f_k can be eliminated by taking partial derivative with respect to one variable and substituting another by a field constant. If this constant is non-zero, we contradict the inductive hypothesis.
4. Otherwise, use the sum of (multiplicative) ROPs representation of $\mathcal{M}_n^{\alpha,\beta}$ to show that the degree of f can be made at most $(n-2)$ by setting one of the variables to zero. This contradicts our choice of f since $\beta \neq 0$.

We now proceed with the proof.

Claim 19. *For all $m \in [k]$ and $a, b \in [n]$, $\partial_{x_a} \partial_{x_b}(f_m) \neq 0$.*

Proof Suppose to the contrary that $\partial_{x_a} \partial_{x_b}(f_m) = 0$. Assume without loss of generality that $a = n$, $b = n-1$, $m = k$, so $\partial_{x_n}\partial_{x_{n-1}}(f_k) = 0$. Then,

$$\begin{aligned}
\mathcal{M}_n^{\alpha,\beta} = f &= \sum_{m=0}^{k} f_m && \text{(by assumption)}\\
\partial_{x_n}\partial_{x_{n-1}}(\mathcal{M}_n^{\alpha,\beta}) &= \sum_{m=0}^{k} \partial_{x_n}\partial_{x_{n-1}}(f_m) && \text{(by additivity of partial derivative)}\\
\mathcal{M}_{n-2}^{\alpha,\beta} &= \sum_{m=0}^{k-1} \partial_{x_n}\partial_{x_{n-1}}(f_m) && \text{(recursive structure of } \mathcal{M}_n \text{ from Proposition 16,}\\
& && \text{and since } \partial_{x_n}\partial_{x_{n-1}}(f_k) = 0)
\end{aligned}$$

Thus $\mathcal{M}_{n-2}^{\alpha,\beta}$ can be written as the sum of $k-1$ polynomials, each of which is a ROP (by Fact 9). By the inductive hypothesis, $2(k-1) \geq (n-2)$. Therefore, $k \geq n/2$ contradicting our assumption. □

From Claim 19 and Fact 12, we can conclude:

Observation 20. *For all $m \in [k]$, f_m is a multiplicative ROP.*

Observation 20 and Lemma 13 together imply that for each $m \in [k]$ and $a \in [n]$, there exist $b \neq a \in [n]$ and $\gamma \in \mathbb{F}$ such that $\partial_{x_b}(f_m) \mid_{x_a=\gamma} = 0$. There are two cases to consider.

First, consider the case when for some m, a and the corresponding b, γ, it turns out that $\gamma \neq 0$. Assume without loss of generality that $m = k$, $a = n-1$, $b = n$, so that $\partial_{x_n}(f_k) \mid_{x_{n-1}=\gamma} = 0$. (For other indices the argument is symmetric.) Then

$$\begin{aligned}
\mathcal{M}_n^{\alpha,\beta} &= \sum_{i\in[k]} f_i && \text{(by assumption)}\\
\partial_{x_n}(\mathcal{M}_n^{\alpha,\beta}) \mid_{x_{n-1}=\gamma} &= \sum_{i\in[k]} \partial_{x_n}(f_i) \mid_{x_{n-1}=\gamma} && \text{(by additivity of partial derivative)}\\
\mathcal{M}_{n-1}^{\alpha,\beta} \mid_{x_{n-1}=\gamma} &= \sum_{i\in[k-1]} \partial_{x_n}(f_i) \mid_{x_{n-1}=\gamma} && \text{(since } \gamma \text{ is chosen as per Lemma 13)}\\
\mathcal{M}_{n-2}^{\alpha\gamma+\beta,\beta\gamma} &= \sum_{i\in[k-1]} \partial_{x_n}(f_i) \mid_{x_{n-1}=\gamma} && \text{(recursive structure of } \mathcal{M}_n \text{ from Proposition 16)}
\end{aligned}$$

Therefore, $\mathcal{M}_{n-2}^{\alpha\gamma+\beta,\beta\gamma}$ can be written as a sum of at most $k-1$ polynomials, each of which is a ROP (Fact 9). By the inductive hypothesis, $2(k-1) \geq n-2$ implying that $k \geq n/2$ contradicting our assumption.

(Note: the term $\mathcal{M}_{n-2}^{\alpha\gamma+\beta,\beta\gamma}$ is what necessitates a stronger induction hypothesis than working with just $\alpha = 0, \beta = 1$.)

It remains to handle the case when for all $m \in [k]$ and $a \in [n]$, the corresponding value of γ to some x_b (as guaranteed by Lemma 13) is 0. Examining the proof of Lemma 13, this implies that each leaf node in any of the ROFs can be made zero only by setting the corresponding variable to zero. That is, the linear forms at all leaves are of the form $a_i x_i$.

Since each φ_m is a multiplicative ROP, setting $x_n = 0$ makes the variables in the polynomial computed at the sibling of the leaf node $a_n x_n$ redundant. Hence setting $x_n = 0$ reduces the degree of each f_m by at least 2. That is, $\deg(f \mid_{x_n=0}) \leq n-2$. But $\mathcal{M}_n^{\alpha,\beta} \mid_{x_n=0}$ equals $\mathcal{M}_{n-1}^{\beta,0} = \beta S_{n-1}^{n-1}$, which has degree $n-1$, contradicting the asusmption that $f = \mathcal{M}_n^{\alpha,\beta}$. $\square$

Combining the results of Lemma 18 and Lemma 17, we obtain the following theorem. At $\alpha = 0, \beta = 1$, it yields Theorem 1.

Theorem 21. *For each $n \geq 1$, any $\alpha \in \mathbb{F}$ and any any $\beta \in \mathbb{F} \setminus \{0\}$, the polynomial $\alpha S_n^n + \beta S_n^{n-1}$ is in $\sum^k \cdot$ROP but not in $\sum^{k-1} \cdot$ROP, where $k = \lceil n/2 \rceil$.*

4. A family of 4-variate multilinear polynomials not in $\sum^2 \cdot$ROP

This section is devoted to proving Theorem 2. We want to find an explicit 4-variate multilinear polynomial that is not expressible as the sum of 2 ROPs.

Note that the proof of Theorem 1 does not help here, since the polynomials separating $\sum^2 \cdot$ROP from $\sum^3 \cdot$ROP have 5 or 6 variables. One obvious approach is to consider other combinations of the symmetric polynomials. This fails too; we can show that all such combinations are in $\sum^2 \cdot$ROP.

Proposition 22. *For every choice of field constants a_i for each $i \in \{0, 1, 2, 3, 4\}$, the polynomial $\sum_{i=0}^4 a_i S_4^i$ can be expressed as the sum of two ROPs.*

Proof Let $g = \sum_i a_i S_4^i$. We obtain the expression for g in different ways in 4 different cases.

Case	Expression
$a_2 = a_3 = 0$	$g = \; a_0 + a_1 S_4^1 + a_4 S_4^4$
$a_2 = 0$; $a_3 \neq 0$	$g = \; \left(a_1 + a_3x_1x_2)(x_3 + x_4 + \frac{a_4}{a_3}x_3x_4)\right) + \left((a_1 + a_3x_3x_4)(x_1 + x_2 - \frac{a_1a_4}{a_3^2})\right) + c$
$a_2 \neq 0$; $a_2a_4 = a_3^2$	$a_2g = \; (a_1 + a_2(x_1 + x_2) + a_3x_1x_2)(a_1 + a_2(x_3 + x_4) + a_3x_3x_4) + (a_2^2 - a_1a_3)(x_1x_2 + x_3x_4)) + c$
$a_2 \neq 0$; $a_2a_4 \neq a_3^2$	$a_2g = \; (a_1 + a_2(x_1 + x_2) + a_3x_1x_2)(a_1 + a_2(x_3 + x_4) + a_3x_3x_4) + \left(x_1x_2 + \frac{a_2^2 - a_1a_3}{a_2a_4 - a_3^2}\right)((a_2a_4 - a_3^2)x_3x_4 + a_2^2 - a_1a_3) + c$

In the above, c is an appropriate field constant, and can be added to any ROP. Notice that the first expression is a sum of two ROPs since it is the sum of a linear polynomial and a single monomial. All the other expressions have two summands, each of which is a product of variable-disjoint bivariate polynomials (ignoring constant terms). Since every bivariate polynomial is a ROP, these representations are also sums of 2 ROPs. $\square$

Instead, we define a polynomial that gives carefully chosen weights to the monomials of S_4^2. Let $f^{\alpha,\beta,\gamma}$ denote the following polynomial:

$$f^{\alpha,\beta,\gamma} = \alpha \cdot (x_1x_2 + x_3x_4) + \beta \cdot (x_1x_3 + x_2x_4) + \gamma \cdot (x_1x_4 + x_2x_3).$$

To keep notation simple, we will omit the superscript when it is clear from the context. In the theorem below, we obtain necessary and sufficient conditions on α, β, γ under which f can be expressed as a sum of two ROPs.

Theorem 23 (Hardness of representation for sum of 2 ROPs). *Let f be the polynomial $f^{\alpha,\beta,\gamma} = \alpha\cdot(x_1x_2+x_3x_4)+\beta\cdot(x_1x_3+x_2x_4)+\gamma\cdot(x_1x_4+x_2x_3)$. The following are equivalent:*

1. *f is not expressible as the sum of two ROPs over $\mathbb{F}$.*
2. *α, β, γ satisfy all the three conditions C1, C2, C3 listed below.*

   C1: $\alpha\beta\gamma \neq 0$.

   C2: $(\alpha^2 - \beta^2)(\beta^2 - \gamma^2)(\gamma^2 - \alpha^2) \neq 0$.

   C3: *None of the equations $X^2 - d_i = 0$, $i \in [3]$, has a root in $\mathbb{F}$, where*

$$\begin{aligned} d_1 &= (+\alpha^2 - \beta^2 - \gamma^2)^2 - (2\beta\gamma)^2 \\ d_2 &= (-\alpha^2 + \beta^2 - \gamma^2)^2 - (2\alpha\gamma)^2 \\ d_3 &= (-\alpha^2 - \beta^2 + \gamma^2)^2 - (2\alpha\beta)^2 \end{aligned}$$

Remark 2. 1. It follows, for instance, that $2(x_1x_2 + x_3x_4) + 4(x_1x_3 + x_2x_4) + 5(x_1x_4 + x_2x_3)$ cannot be written as a sum of 2 ROPs over reals, yielding Theorem 2.

2. If $\mathbb{F}$ is an algebraically closed field, then for every α, β, γ, condition C3 fails, and so every $f^{\alpha,\beta,\gamma}$ can be written as a sum of 2 ROPs. However we do not know if there are other examples, or whether all multilinear 4-variate polynomials are expressible as the sum of two ROPs.
3. Even if $\mathbb{F}$ is not algebraically closed, condition C3 fails if for each $a \in \mathbb{F}$, the equation $X^2 = a$ has a root.

Our strategy for proving Theorem 23 is a generalization of an idea used in [19]. While Volkovich showed that 3-variate ROPs have a nice structural property in terms of their partial derivatives and commutators, we show that the sums of two 4-variate ROPs have at least one nice structural property in terms of their bivariate restrictions, partial derivatives, and commutators. Then we show that provided α, β, γ are chosen carefully, the polynomial $f^{\alpha,\beta,\gamma}$ will not satisfy any of these properties and hence cannot be a sum of two ROPs.

To prove Theorem 23, we first consider the easier direction, $1 \Rightarrow 2$, and prove the contrapositive.

Lemma 24. *If α, β, γ do not satisfy all of C1,C2,C3, then the polynomial f can be written as a sum of 2 ROPs.*

Proof **C1 false:** If any of α, β, γ is zero, then by definition f is the the sum of at most two ROPs.
C2 false: Without loss of generality, assume $\alpha^2 = \beta^2$, so $\alpha = \pm\beta$. Then f is computed by $f = \alpha \cdot (x_1 \pm x_4)(x_2 \pm x_3) + \gamma \cdot (x_1x_4 + x_2x_3)$.
C1 true; C3 false: Without loss of generality, the equation $X^2 - d_1 = 0$ has a root τ. We try to express f as

$$\alpha(x_1 - ax_3)(x_2 - bx_4) + \beta(x_1 - cx_2)(x_3 - dx_4).$$

The coefficients for x_3x_4 and x_2x_4 force $ab = 1$, $cd = 1$, giving the form

$$\alpha(x_1 - ax_3)(x_2 - \frac{1}{a}x_4) + \beta(x_1 - cx_2)(x_3 - \frac{1}{c}x_4).$$

Comparing the coefficients for x_1x_4 and x_2x_3, we obtain the constraints

$$-\frac{\alpha}{a} - \frac{\beta}{c} = \gamma; \qquad -\alpha a - \beta c = \gamma$$

Expressing a as $\frac{-\gamma-\beta c}{\alpha}$, we get a quadratic constraint on c; it must be a root of the equation

$$Z^2 + \frac{-\alpha^2 + \beta^2 + \gamma^2}{\beta\gamma} Z + 1 = 0.$$

Using the fact that $\tau^2 = d_1 = (-\alpha^2 + \beta^2 + \gamma^2)^2 - (2\beta\gamma)^2$, we see that indeed this equation does have roots. The left-hand size splits into linear factors, giving

$$(Z - \delta)(Z - \frac{1}{\delta}) = 0 \quad \text{where} \quad \delta = \frac{\alpha^2 - \beta^2 - \gamma^2 + \tau}{2\beta\gamma}.$$

It is easy to verify that $\delta \neq 0$ and $\delta \neq -\frac{\gamma}{\beta}$ (since $\alpha \neq 0$). Further, define $\mu = \frac{-(\gamma+\beta\delta)}{\alpha}$. Then μ is well-defined (because $\alpha \neq 0$) and is also non-zero. Now setting $c = \delta$ and $a = \mu$, we have satisfied all the constraints and so we can write f as the sum of 2 ROPs as follows:

$$f = \alpha(x_1 - \mu x_3)(x_2 - \frac{1}{\mu}x_4) + \beta(x_1 - \delta x_2)(x_3 - \frac{1}{\delta}x_4).$$

□

Now we consider the harder direction: $2 \Rightarrow 1$. Again, we consider the contrapositive. We first show (Lemma 25) a structural property satisfied by every polynomial in $\sum^2 \cdot$ROP: it must satisfy at least one of the three properties $C1', C2', C3'$ described in the lemma. We then show (Lemma 26) that under the conditions $C1, C2, C3$ from the theorem statement, f does not satisfy any of $C1', C2', C3'$; it follows that f is not expressible as the sum of 2 ROPs.

Lemma 25. *Let g be a 4-variate multilinear polynomial over the field $\mathbb{F}$ which can be expressed as a sum of 2 ROPs. Then at least one of the following conditions is true:*

C1': *There exist $i, j \in [4]$ and $a, b \in \mathbb{F}$ such that $g \mid_{x_i=a, x_j=b}$ is linear.*

C2': *There exist $i, j \in [4]$ such that $x_i, x_j, \partial_{x_i}(g), \partial_{x_j}(g), 1$ are linearly dependent.*

C3': *$g = l_1 \cdot l_2 + l_3 \cdot l_4$ where l_is are linear forms, l_1 and l_2 are variable-disjoint, and l_3 and l_4 are variable-disjoint.*

Proof Let φ be a sum of 2 ROFs computing g. Let v_1 and v_2 be the children of the topmost $+$ gate. The proof is in two steps. First, we reduce to the case when $|\text{Var}(v_1)| = |\text{Var}(v_2)| = 4$. Then we use a case analysis to show that at least one of the aforementioned conditions hold true. In both steps, we will repeatedly use Proposition 10, which showed that any 3-variate ROP can be reduced to a linear polynomial by substituting a single variable with a field constant. We now proceed with the proof.

Suppose $|\text{Var}(v_1)| \leq 3$. Applying Proposition 10 first to v_1 and then to the resulting restriction of v_2, one can see that there exist $i, j \in [4]$ and $a, b \in \mathbb{F}$ such that $g \mid_{x_i=a, x_j=b}$ is a linear polynomial. So condition $C1'$ is satisfied.

Now assume that $|\text{Var}(v_1)| = |\text{Var}(v_2)| = 4$. Depending on the type of gates of v_1 and v_2, we consider 3 cases.

Case 1: Both v_1 and v_2 are $\times$ gates. Then g can be represented as $M_1 \cdot M_2 + M_3 \cdot M_4$ where (M_1, M_2) and (M_3, M_4) are variable-disjoint ROPs.

Suppose that for some i, $|\text{Var}(M_i)| = 1$. Then, $g \mid_{M_i \to 0}$ is a 3-variate restriction of f and is clearly an ROP. Applying Proposition 10 to this restriction, we see that condition $C1'$ holds.

Otherwise each M_i has $|\text{Var}(M_i)| = 2$.

Suppose (M_1, M_2) and (M_3, M_4) define distinct partitions of the variable set. Assume without loss of generality that $g = M_1(x_1, x_2) \cdot M_2(x_3, x_4) + M_3(x_1, x_3) \cdot M_4(x_2, x_4)$. If all M_is are linear forms, it is clear that condition $C3'$ holds. If not, assume that M_1 is of the form $l_1(x_1) \cdot m_1(x_2) + c_1$ where l_1, m_1 are linear forms and $c_1 \in \mathbb{F}$. Now $g \mid_{l_1 \to 0} = c_1 \cdot M_2(x_3, x_4) + M_3'(x_3) \cdot M_4(x_2, x_4)$. Either set x_3 to make M_3' zero, or, if that is not possible because M_3' is a non-zero field constant, then set $x_4 \to b$ where $b \in \mathbb{F}$. In both cases, by setting at most 2 variables, we obtain a linear polynomial, so $C1'$ holds.

Otherwise, (M_1, M_2) and (M_3, M_4) define the same partition of the variable set. Assume without loss of generality that $g = M_1(x_1, x_2) \cdot M_2(x_3, x_4) + M_3(x_1, x_2) \cdot M_4(x_3, x_4)$. If one of the M_is is linear, say without loss of generality that M_1 is a linear form, then $g \mid_{M_4 \to 0}$ is a 2-variate restriction which is also a linear form, so $C1'$ holds. Otherwise, none of the M_is is a linear form. Then each M_i can be represented as $l_i \cdot m_i + c_i$ where l_i, m_i are univariate linear forms and $c_i \in \mathbb{F}$. We consider a 2-variate restriction which sets l_1 and m_4 to 0. (Note that $\text{Var}(l_1) \cap \text{Var}(m_4) = \emptyset$.) Then the resulting polynomial is a linear form, so $C1'$ holds.

Case 2: Both v_1 and v_2 are $+$ gates. Then g can be written as $f = M_1 +$

$M_2 + M_3 + M_4$ where (M_1, M_2) and (M_3, M_4) are variable-disjoint ROPs.

Suppose (M_1, M_2) and (M_3, M_4) define distinct partitions of the variable set.

Suppose further that there exists M_i such that $|\mathrm{Var}(M_i)| = 1$. Without loss Of generality, $\mathrm{Var}(M_1) = \{x_1\}$, $\{x_1, x_2\} \subseteq \mathrm{Var}(M_3)$, and $x_3 \in \mathrm{Var}(M_4)$. Any setting to x_2 and x_4 results in a linear polynomial, so $C1'$ holds.

So assume without loss of generality that $g = M_1(x_1, x_2) + M_2(x_3, x_4) + M_3(x_1, x_3) + M_4(x_2, x_4)$. Then for $a, b \in \mathbb{F}$, $g\mid_{x_1=a,x_4=b}$ is a linear polynomial, so $C1'$ holds.

Otherwise, (M_1, M_2) and (M_3, M_4) define the same partition of the variable set. Again, if say $|\mathrm{Var}(M_1)| = 1$, then setting two variables from M_2 shows that $C1'$ holds. So assume without loss of generality that $g = M_1(x_1, x_2)+M_2(x_3, x_4)+M_3(x_1, x_2)+M_4(x_3, x_4)$. Then for $a, b \in \mathbb{F}$, $g\mid_{x_1=a,x_3=b}$ is a linear polynomial, so again $C1'$ holds.

Case 3: One of v_1, v_2 is a $+$ gate and the other is a $\times$ gate. Then g can be written as $g = M_1 + M_2 + M_3 \cdot M_4$ where (M_1, M_2) and (M_3, M_4) are variable-disjoint ROPs. Suppose that $|\mathrm{Var}(M_3)| = 1$. Then $g\mid_{M_3 \to 0}$ is a 3-variate restriction which is a ROP. Using Proposition 10, we get a 2-variate restriction of g which is also linear, so $C1'$ holds. The same argument works when $|\mathrm{Var}(M_4)| = 1$. So assume that M_3 and M_4 are bivariate polynomials.

Suppose that (M_1, M_2) and (M_3, M_4) define distinct partitions of the variable set. Assume without loss of generality that $g = M_1+M_2+M_3(x_1, x_2)\cdot M_4(x_3, x_4)$, and x_3, x_4 are separated by M_1, M_2. Then $g\mid_{M_3 \to 0}$ is a 2-variate restriction which is also linear, so $C1'$ holds.

Otherwise (M_1, M_2) and (M_3, M_4) define the same partition of the variable set. Assume without loss of generality that $g = M_1(x_1, x_2)+M_2(x_3, x_4)+M_3(x_1, x_2) \cdot M_4(x_3, x_4)$. If M_1 (or M_2) is a linear form, then consider a 2-variate restriction of g which sets M_4 (or M_3) to 0. The resulting polynomial is a linear form. Similarly if M_3 (or M_4) is of the form $l \cdot m + c$ where l, m are univariate linear forms, then we consider a 2-variate restriction which sets l to 0 and some $x_i \in \mathrm{Var}(M_4)$ to a field constant. The resulting polynomial again is a linear form. In all these cases, $C1'$ holds.

The only case that remains is that M_3 and M_4 are linear forms while M_1 and M_2 are not. Assume that $M_1 = (a_1x_1 + b_1)(a_2x_2 + b_2) + c$ and $M_3 = a_3x_1 + b_3x_2 + c_3$. Then $\partial_{x_1}(g) = a_1(a_2x_2 + b_2) + a_3M_4$ and $\partial_{x_2}(g) = (a_1x_1 + b_1)a_2 + b_3M_4$. It follows that $b_3 \cdot \partial_{x_1}(g) - a_3 \cdot \partial_{x_2}(g) + a_1a_2a_3x_1 - a_1a_2b_3x_2 = a_1b_2b_3 - b_1a_2a_3 \in \mathbb{F}$, and hence the polynomials x_1, x_2, $\partial_{x_1}(g)$,

$\partial_{x_2}(g)$ and 1 are linearly dependent. Therefore, condition $C2'$ of the lemma is satisfied. $\square$

Lemma 26. *If α, β, γ satisfy conditions $C1, C2, C3$ from the statement of Theorem 23, then the polynomial $f^{\alpha,\beta,\gamma}$ does not satisfy any of the properties $C1', C2', C3'$ from Lemma 25.*

Proof $\mathbf{C1} \Rightarrow \neg\mathbf{C1'}$: Since $\alpha\beta\gamma \neq 0$, f contains all possible degree 2 monomials. Hence after setting $x_i = a$ and $x_j = b$, the monomial $x_k x_l$ where $k, l \in [4]\backslash\{i, j\}$ still survives.

$\mathbf{C2} \Rightarrow \neg\mathbf{C2'}$: The proof is by contradiction. Assume to the contrary that for some i, j, without loss of generality say for $i = 1$ and $j = 2$, the polynomials $x_1, x_2, \partial_{x_1}(f), \partial_{x_2}(f), 1$ are linearly dependent. Note that $\partial_{x_1}(f) = \alpha x_2 + \beta x_3 + \gamma x_4$ and $\partial_{x_2}(f) = \alpha x_1 + \gamma x_3 + \beta x_4$. This implies that the vectors $(1, 0, 0, 0, 0)$, $(0, 1, 0, 0, 0)$, $(0, \alpha, \beta, \gamma, 0)$, $(\alpha, 0, \gamma, \beta, 0)$ and $(0, 0, 0, 0, 1)$ are linearly dependent. This further implies that the vectors (β, γ) and (γ, β) are linearly dependent. Therefore, $\beta = \pm\gamma$, contradicting C2.

$\mathbf{C1} \wedge \mathbf{C2} \wedge \mathbf{C3} \Rightarrow \neg\mathbf{C3'}$: Suppose, to the contrary, that $C3'$ holds. That is, f can be written as $f = l_1 \cdot l_2 + l_3 \cdot l_4$ where (l_1, l_2) and (l_3, l_4) are variable-disjoint linear forms. By the preceding arguments, we know that f does not satisfy $C1'$ or $C2'$.

First consider the case when (l_1, l_2) and (l_3, l_4) define the same partition of the variable set. Assume without loss of generality that $\mathrm{Var}(l_1) = \mathrm{Var}(l_3)$, $\mathrm{Var}(l_2) = \mathrm{Var}(l_4)$, and $|\mathrm{Var}(l_1)| \leq 2$. Setting the variables in l_1 to any field constants yields a linear form, so f satisfies C1', a contradiction.

Hence it must be the case that (l_1, l_2) and (l_3, l_4) define different partitions of the variable set. Since all degree-2 monomials are present in f, each pair x_i, x_j must be separated by at least one of the two partitions. This implies that both partitions have exactly 2 variables in each part. Assume without loss of generality that $f = l_1(x_1, x_2) \cdot l_2(x_3, x_4) + l_3(x_1, x_3) \cdot l_4(x_2, x_4)$.

At this point, we use properties of the commutator of f; recall Definition 14. By Lemma 15, we know that l_2 divides $\triangle_{12} f$. We compute $\triangle_{12} f$ explicitly for our candidate polynomial:

$$\begin{aligned}\triangle_{12} f &= (\alpha x_3 x_4)(\alpha + (\beta + \gamma)(x_3 + x_4) + \alpha x_3 x_4) \\ &\quad - (\beta x_4 + \gamma x_3 + \alpha x_3 x_4)(\beta x_3 + \gamma x_4 + \alpha x_3 x_4) \\ &= -\beta\gamma(x_3^2 + x_4^2) + (\alpha^2 - \beta^2 - \gamma^2) x_3 x_4\end{aligned}$$

Since l_2 divides $\triangle_{12}f$, $\triangle_{12}f$ is not irreducible but is the product of two linear factors. Since $\triangle_{12}f(0,0) = 0$, at least one of the linear factors of $\triangle_{12}f$ must vanish at $(0,0)$. Let $x_3 - \delta x_4$ be such a factor. Then $\triangle_{12}(f)$ vanishes not only at $(0,0)$, but whenever $x_3 = \delta x_4$. Substituting $x_3 = \delta x_4$ in $\triangle_{12}f$, we get

$$-\delta^2\beta\gamma - \beta\gamma + \delta(\alpha^2 - \beta^2 - \gamma^2) = 0$$

Hence δ is of the form

$$\delta = \frac{-(\alpha^2 - \beta^2 - \gamma^2) \pm \sqrt{(\alpha^2 - \beta^2 - \gamma^2)^2 - 4\beta^2\gamma^2}}{-2\beta\gamma}$$

Hence $2\beta\gamma\delta - (\alpha^2 - \beta^2 - \gamma^2)$ is a root of the equation $X^2 - d_1 = 0$, contradicting the assumption that C3 holds.

Hence it must be the case that $C3'$ does not hold. □

With this, the proof of Theorem 23 is complete.

The conditions imposed on α, β, γ in Theorem 23 are tight and irredundant. Below we give some explicit examples over the field of reals.

1. $f = 2(x_1x_2 + x_3x_4) + 2(x_1x_3 + x_2x_4) + 3(x_1x_4 + x_2x_3)$ satisfies conditions C1 and C3 from the Theorem but not C2; $\alpha = \beta$. A $\sum^2 \cdot$ROP representation for f is $f = 2(x_1 + x_4)(x_2 + x_3) + 3(x_1x_4 + x_2x_3)$.
2. $f = 2(x_1x_2 + x_3x_4) - 2(x_1x_3 + x_2x_4) + 3(x_1x_4 + x_2x_3)$ satisfies conditions C1 and C3 but not C2; $\alpha = -\beta$. A $\sum^2 \cdot$ROP representation for f is $f = 2(x_1 - x_4)(x_2 - x_3) + 3(x_1x_4 + x_2x_3)$.
3. $f = (x_1x_2 + x_3x_4) + 2(x_1x_3 + x_2x_4) + 3(x_1x_4 + x_2x_3)$ satisfies conditions C1 and C2 but not C3. A $\sum^2 \cdot$ROP representation for f is $f = (x_1 + x_3)(x_2 + x_4) + 2(x_1 + x_2)(x_3 + x_4)$.

5. Conclusions

1. We have seen in Proposition 5 that every n-variate multilinear polynomial ($n \geq 4$) can be written as the sum of $3 \times 2^{n-4}$ ROPs. The counting argument from Proposition 7 shows that there exist multilinear polynomials f requiring exponentially many ROPs summands; if $f \in \sum^k \cdot$ROP then $k = \Omega(2^n/n^2)$. Our general upper bound on k is $O(2^n)$, leaving a small gap between the lower and upper bound. What is the true tight bound? Can we find explicit polynomials where exponentially large k is necessary and sufficient in any $\sum^k \cdot$ROP expression?

One such example is the polynomial defined by Raz and Yehudayoff in [15]; as shown in [13], k must be exponential in $\Omega(n^{1/3}/\log n)$. But we do not know whether this value of k is asymptotically tight.

2. We have shown in Theorem 1 that for each k, $\sum^k \cdot$ROP can be separated from $\sum^{k-1} \cdot$ROP by a polynomial on $2k-1$ variables. Can we separate these classes with fewer variables? Note that any separating polynomial must have $\Omega(\log k)$ variables.
3. In particular, can 4-variate multilinear polynomials separate sums of 3 ROPs from sums of 2 ROPs over every field? If not, what is an explicit example?
4. We now understand ROPs and ROFs very well, [19]. However, our understanding of sums of ROPs is not so good. Can we at least characterise $\sum^2 \cdot$ROPs?

References

[1] Agrawal, M., Vinay, V., 2008. Arithmetic circuits: A chasm at depth four. In: 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA. pp. 67–75.
URL https://doi.org/10.1109/FOCS.2008.32

[2] Anderson, M., van Melkebeek, D., Volkovich, I., 2015. Deterministic polynomial identity tests for multilinear bounded-read formulae. Computational Complexity 24 (4), 695–776.

[3] Bshouty, D., Bshouty, N. H., 1998. On interpolating arithmetic read-once formulas with exponentiation. Journal of Computer and System Sciences 56 (1), 112–124.
URL http://dx.doi.org/10.1006/jcss.1997.1550

[4] Bshouty, N. H., Cleve, R., 1998. Interpolating arithmetic read-once formulas in parallel. SIAM Journal on Computing 27 (2), 401–413.
URL http://dx.doi.org/10.1137/S009753979528812X

[5] Bshouty, N. H., Hancock, T. R., Hellerstein, L., 1995. Learning boolean read-once formulas over generalized bases. J. Comput. Syst. Sci. 50 (3), 521–542.
URL http://dx.doi.org/10.1006/jcss.1995.1042

[6] Dvir, Z., Shpilka, A., Yehudayoff, A., 2009. Hardness-randomness tradeoffs for bounded depth arithmetic circuits. SIAM Journal on Computing 39 (4), 1279–1293.
URL https://doi.org/10.1137/080735850

[7] Forbes, M. A., 2015. Deterministic divisibility testing via shifted partial derivatives. In: IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015. pp. 451–465.
URL https://doi.org/10.1109/FOCS.2015.35

[8] Hancock, T. R., Hellerstein, L., 1991. Learning read-once formulas over fields and extended bases. In: Warmuth, M. K., Valiant, L. G. (Eds.), Proceedings of the Fourth Annual Workshop on Computational Learning Theory, COLT 1991, Santa Cruz, California, USA, August 5-7, 1991. Morgan Kaufmann, pp. 326–336.
URL http://dl.acm.org/citation.cfm?id=114867

[9] Hrubes, P., Yehudayoff, A., 2011. Arithmetic complexity in ring extensions. Theory of Computing 7 (1), 119–129.
URL http://dx.doi.org/10.4086/toc.2011.v007a008

[10] Kabanets, V., Impagliazzo, R., 2004. Derandomizing polynomial identity tests means proving circuit lower bounds. Computational Complexity 13 (1-2), 1–46.
URL http://dx.doi.org/10.1007/s00037-004-0182-6

[11] Kayal, N., Koiran, P., Pecatte, T., Saha, C., 2015. Lower bounds for sums of powers of low degree univariates. In: Halldórsson, M. M., Iwama, K., Kobayashi, N., Speckmann, B. (Eds.), Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I. Vol. 9134 of Lecture Notes in Computer Science. Springer, pp. 810–821.
URL http://dx.doi.org/10.1007/978-3-662-47672-7_66

[12] Mahajan, M., Tawari, A., 2016. Sums of read-once formulas: How many summands suffice? In: Proceedings of 11th International Computer Science Symposium in Russia, CSR 2016, LNCS 9691. Springer, pp. 266–279, eCCC technical report 2015-204, CoRR abs/1603.02605.

[13] Ramya, C., Rao, B. V. R., 2016. Sum of products of read-once formulas. In: 36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2016, December 13-15, 2016, Chennai, India. pp. 39:1–39:15.
URL `https://doi.org/10.4230/LIPIcs.FSTTCS.2016.39`

[14] Raz, R., 2009. Multi-linear formulas for permanent and determinant are of super-polynomial size. Journal of the ACM 56 (2).
URL `http://doi.acm.org/10.1145/1502793.1502797`

[15] Raz, R., Yehudayoff, A., 2008. Balancing syntactically multilinear arithmetic circuits. Computational Complexity 17 (4), 515–535.
URL `https://doi.org/10.1007/s00037-008-0254-0`

[16] Shpilka, A., Volkovich, I., 2014. On reconstruction and testing of read-once formulas. Theory of Computing 10, 465–514.
URL `http://dx.doi.org/10.4086/toc.2014.v010a018`

[17] Shpilka, A., Volkovich, I., 2015. Read-once polynomial identity testing. Computational Complexity 24 (3), 477–532, (combines results from papers in RANDOM 2009 and STOC 2008).

[18] Shpilka, A., Yehudayoff, A., 2010. Arithmetic circuits: A survey of recent results and open questions. Foundations and Trends in Theoretical Computer Science 5 (3-4), 207–388.
URL `http://dx.doi.org/10.1561/0400000039`

[19] Volkovich, I., Feb. 2016. Characterizing arithmetic read-once formulae. ACM Transactions on Computation Theory 8 (1), 2:1–2:19.