Electronic Colloquium on Computational Complexity, Revision 2 of Report No. 171 (2015)

Monotone projection lower bounds from extended formulation lower bounds

Joshua A. Grochow*

December 23, 2016

Abstract

In this short note, we show that the Hamiltonian Cycle polynomial is not a monotone sub-exponential-size projection of the permanent; this both rules out a natural attempt at a monotone lower bound on the Boolean permanent, and shows that the permanent is *not* complete for non-negative polynomials in $\mathsf{VNP}_{\mathbb{R}}$ under monotone p-projections. We also show that the cut polynomials, $\sum_{A \subseteq [n]} \prod_{i \in A, j \notin A} x_{ij}^{q-1}$, and the perfect matching polynomial (or "unsigned Pfaffian") $\frac{1}{2^n n!} \sum_{\pi \in S_{2n}} \prod_{i=1}^{n} x_{\pi(2i-1),\pi(2i)}$ are not monotone p-projections of the permanent. The latter can be interpreted as saying that there is no monotone projection reduction from counting perfect matchings in general graphs to counting perfect matchings in bipartite graphs, putting at least one theorem behind the well-established intuition. As the permanent is universal for monotone formulas, these results also imply exponential lower bounds on the monotone formula size and monotone circuit size of these polynomials. To prove these results we introduce a new connection between monotone projections of polynomials and extended formulations of linear programs that may have further applications.

1 Introduction

The permanent

$$\mathrm{perm}_n(X) = \sum_{\pi \in S_n} x_{1,\pi(1)} x_{2,\pi(2)} \cdots x_{n,\pi(n)}$$

(where S_n denotes the symmetric group of all permutations of $\{1, \ldots, n\}$) has long fascinated combinatorialists [13, 25, 14], more recently physicists [26, 2], and since Valiant's seminal paper [22], has also been a key object of study in computational complexity. Despite its beauty, the permanent has some computational quirks: in particular, although the permanent of integer matrices is $\#\mathsf{P}$-complete and the permanent is VNP-complete in characteristic zero, the permanent *mod 2* is the same as the determinant, and hence can easily be computed. In fact, computing the permanent mod 2^k is easy for any k [22], though the proof is more involved. Modulo any other number n, the permanent of integer matrices is $\mathsf{Mod}_n\mathsf{P}$-complete.

In contrast, the seemingly similar Hamiltonian Cycle polynomial,

$$HC_n(X) = \sum_{n\text{-cycles } \sigma} x_{1,\sigma(1)} x_{2,\sigma(2)} \cdots x_{n,\sigma(n)},$$

*Department of Computer Science, University of Colorado at Boulder, Boulder, CO, USA and Santa Fe Institute, Santa Fe, NM, USA, `jgrochow@santafe.edu`

ISSN 1433-8092

where the sum is only over n-cycles rather than over all permutations, does not have these quirks: The Hamiltonian Cycle polynomial is VNP-complete over any ring R [21] and $\mathsf{Mod}_n\mathsf{P}$-complete for all n (that is, counting Hamiltonian cycles is complete for these Boolean counting classes).

Jukna [11] observed that, over the Boolean semi-ring, if the Hamiltonian Cycle polynomial were a monotone p-projection of the permanent, there would be a $2^{n^{\Omega(1)}}$ lower bound on monotone circuits computing the permanent, a lower bound that still remains open (the current record is still Razborov's $n^{\Omega(\log n)}$ [15]). Even over the real numbers, such a monotone p-projection would give an alternative proof of a $2^{n^{\Omega(1)}}$ lower bound on the permanent (Jerrum and Snir [9] already showed the permanent requires monotone circuits of size $2^{\Omega(n)}$ over $\mathbb{R}$ and over the tropical $(\min, +)$ semi-ring). Here we show that no such monotone reduction exists—over $\mathbb{R}$, nor over the tropical semi-ring, nor over the Boolean semi-ring—by connecting monotone p-projections to extended formulations of linear programs.

We use the same technique to show that the perfect matching polynomial or "unsigned Pfaffian"

$$\frac{1}{2^n n!}\sum_{\pi\in S_{2n}}\prod_{i=1}^{n} x_{\pi(2i-1),\pi(2i)} = \sum_{\substack{\pi\in S_{2n}\\ \pi(1)<\pi(3)<\cdots<\pi(2n-1)\\ \pi(2k-1)<\pi(2k)\ \forall k}}\prod_{i=1}^{n} x_{\pi(2i-1),\pi(2i)}$$

is not a monotone p-projection of the permanent. As the perfect matching polynomial counts perfect matchings in a general graph, and the permanent counts perfect matchings in a bipartite graph, we have:

Corollary 1.1. *Any efficient projection reduction from counting perfect matchings in general graphs to counting perfect matchings in bipartite graphs must be non-monotone (and, therefore, seemingly quite unintuitive!).*

Remark 1.2 (On the Boolean semi-ring)**.** Our results also hold for *formal polynomials* over the Boolean semi-ring $\mathbb{B} = (\{0,1\}, \vee, \wedge)$. Over the Boolean semi-ring, the permanent is the indicator function of the existence of a perfect matching in a bipartite graph, and the unsigned Pfaffian is the indicator function of a perfect matching in a general graph. However, over $\mathbb{B}$, each function is represented by more than one formal polynomial, and we do not yet know how to extend our results to the setting of *functions* over $\mathbb{B}$. See Section 5 for details and specific questions.

We also use the same technique to show that the cut polynomials $\mathrm{Cut}^q = \sum_{A\subseteq[n]}\prod_{i\in A, j\notin A} x_{ij}^{q-1}$ are not monotone p-projections of the permanent. Perhaps the main complexity-theoretic interest in the cut polynomials is that Cut^q *over the finite field* $\mathbb{F}_q$ was (until recently [12]) the only known example of a natural polynomial that is neither in $\mathsf{VP}_{\mathbb{F}_q}$ nor $\mathsf{VNP}_{\mathbb{F}_q}$-complete under a standard complexity-theoretic assumption (that PH doesn't collapse) [3]; there it was also shown that if $\mathsf{VP}_{\mathbb{F}_q} \neq \mathsf{VNP}_{\mathbb{F}_q}$ then such polynomials of intermediate complexity must exist. In that paper, it was asked whether the cut polynomials, considered as polynomials *over the rationals*, were $\mathsf{VNP}_{\mathbb{Q}}$-complete.[1] Although our results don't touch on this question, these previous results motivate the study of these polynomials over $\mathbb{Q}$.

[1] Cut^2 was subsequently shown to be $\mathsf{VNP}_{\mathbb{Q}}$-complete under circuit reductions [5]; its completeness under projections remains open.

Because the permanent is universal for monotone formulas, our lower bounds also imply exponential lower bounds on the monotone algebraic formula size—and, by balancing algebraic circuits, monotone algebraic circuit size—of these polynomials; see Section 4.2.

Finally, we note that our results shed a little more light on the complicatedness of the known VNP-completeness proofs for the permanent [22, 1]. Namely, prior to our result, the fact that the permanent is not hard modulo 2 already implied that any completeness result must use 2 in a "bad" way: for example, dividing by 2 somewhere. This is indeed true of both Valiant's original proof [22] and Aaronson's independent quantum linear-optics proof [1]. One might hope for a classical analogue of Aaronson's quantum proof, using the characterization of BPP in terms of stochastic matrices as a replacement for the characterization of BQP using unitary matrices. However, our result says that any completeness proof for the permanent must use non-monotone reductions, so such a classical analogue is not possible:

Corollary 1.3. *Aaronson's quantum linear optics proof [1] that the permanent is* #P*-hard cannot be replaced by one using classical randomized algorithms in place of quantum algorithms.*

Our results also imply the necessity of the use of negative numbers in Valiant's 4×4 gadget [22, p. 195]. In light of these results, Valiant's 4×4 gadget may perhaps seem less mysterious than the fact that such a gadget exists that is *only* 4×4!

To prove these results, we show that a monotone projection between non-negative polynomials essentially implies that the Newton polytope of one polynomial is an extension of the Newton polytope of the other (Lemma 3.1), and then apply known lower bounds on the extension complexity of certain polytopes. We hope that the connection between Newton polytopes, monotone projections, and extended formulations finds further use.

2 Preliminaries

A polynomial $f(x_1, \ldots, x_n)$ is a *(simple) projection* of a polynomial $g(y_1, \ldots, y_m)$ if f can be constructed from g by replacing each y_i with a constant or with some x_j. The polynomial f is an *affine projection* of g if f can be constructed from g by replacing each y_i with an affine linear function $\pi_i(\vec{x})$. When we say "projection" we mean simple projection. Given two families of polynomials (f_n), (g_n), if there is a function $p(n)$ such that f_n is a projection of $g_{p(n)}$ for all sufficiently large n, then we say that (f_n) is a projection of (g_n) with *blow-up* $p(n)$. If (f_n) is a projection of (g_n) with polynomial blow-up, we say that (f_n) is a *p-projection* of (g_n).

Over any subring of $\mathbb{R}$—or more generally any totally ordered semi-ring (see below)—a *monotone projection* is a projection in which all constants appearing in the projection are non-negative. Monotone p-projection is defined analogously.

To each monomial $x_1^{e_1} \cdots x_n^{e_n}$ we associate its exponent vector $(e_1, \ldots, e_n)$, as a point in $\mathbb{N}^n \subseteq \mathbb{R}^n$. We then have:

Definition 2.1. The *Newton polytope* of a polynomial $f(x_1, \ldots, x_n)$, denoted $\mathrm{Newt}(f)$, is the convex hull in $\mathbb{R}^n$ of the exponent vectors of all monomials appearing in f with non-zero coefficient.

A polytope is *integral* if all its vertices have integer coordinates; note that Newton polytopes are always integral. A *face* of a polytope P is the intersection of P with a linear space L such that none of the interior points of P lie in L.

For a polytope P, let $c(P)$ denote the "complexity" of P, as measured by the minimal number of linear inequalities needed to define P (equivalently, the number of faces of P of dimension $\dim P - 1$). A polytope $Q \subseteq \mathbb{R}^m$ is an *extension* of $P \subseteq \mathbb{R}^n$ if there is an affine linear map $\ell\colon \mathbb{R}^m \to \mathbb{R}^n$ such that $\ell(Q) = P$. The *extension complexity* of P, denoted $xc(P)$, is the minimum complexity of any extension of P (of any dimension): $xc(P) = \min\{c(Q) | Q \text{ is an extension of } P\}$.

The m-th *cycle cover polytope* (also known as the bipartite perfect matching polytope) is the convex hull in $\mathbb{R}^{m^2}$ of the $\{0,1\}$-indicator functions of the directed cycle covers of the complete directed graph with self-loops on m vertices. The cycle cover polytope is the Newton polytope of the permanent, as each monomial in the permanent corresponds to such a cycle cover.

A totally ordered semi-ring (we only consider commutative ones here) is a totally ordered set together with two operations, denoted $(R, \leq, \times, +, 1, 0)$ such that $(R, \times, 1)$ and $(R, +, 0)$ are both commutative monoids, $\times$ distributes over $+$, $0 \times a = 0$ for all a, $a + c \leq b + c$ whenever $a \leq b$, and $ac \leq bc$ whenever $a \leq b$ and $0 \leq c$. An element c of a totally ordered semi-ring is non-negative if $0 \leq c$, and is positive if furthermore $c \neq 0$. We will restrict our attention to *non-zero* totally ordered semi-rings; equivalently, we assume $1 \neq 0$. (There is only one polynomial over the zero semi-ring—the zero polynomial—so that case was of no interest anyways.)

Note that (non-zero) totally ordered semi-rings always have "characteristic zero:" $1 + \cdots + 1 \neq 0$ (any number of times). In a non-zero totally ordered semi-ring, either $0 < 1$ or $1 < 0$; we handle the first case, the second being analogous. If $0 < 1$, then by adding 1 to both sides k times we get that $\sum_{i \in [k]} 1 < \sum_{i \in [k+1]} 1$. If the ring had nonzero characteristic, then for some $k > 1$, the right-hand side here would become zero, and we would get $0 < 1 < 1 + 1 < \cdots < 0$, a contradiction.

The following totally ordered semi-rings are of particular interest:

- The real numbers with its usual ordering and algebraic operations $(\mathbb{R}, \leq, \times, +)$
- The so-called "tropical semi-ring" $(\mathbb{R}, \leq, +, \min)$, which is the real numbers with its usual ordering, where the product is taken to be real addition and the addition operation is taken to be the minimum.
- The Boolean and-or semi-ring $\mathbb{B} = (\{0,1\}, \leq, \wedge, \vee)$, where $0 \leq 1$.

To get a feel for the latter two semi-rings, note that polynomials over the tropical semi-ring generally compute some optimization problem, and over $\mathbb{B}$ generally compute a decision problem. For example, the Hamiltonian Cycle polynomial over the tropical semi-ring computes the Traveling Salesperson Problem, and over $\mathbb{B}$ is the indicator function of the existence of a Hamiltonian cycle. Note that over $\mathbb{R}$, if two formal polynomials compute the same function then they must be identical, but this is not true over the tropical nor Boolean semi-rings.

3 Main Lemma

Lemma 3.1. *Let R be a totally ordered semi-ring, and let $f(x_1, \ldots, x_n)$ and $g(y_1, \ldots, y_m)$ be polynomials over R with non-negative coefficients. If f is a monotone projection of g, then some face of* $\mathrm{Newt}(g)$ *is an extension of* $\mathrm{Newt}(f)$. *In particular, $xc(\mathrm{Newt}(f)) \leq c(\mathrm{Newt}(g))$.*

Proof. Under simple projections, a monomial in the y's maps to some scalar multiple of a monomial in the x's (possibly the empty monomial, resulting in a constant term, or possibly the zero multiple, resulting in zero). Let π be a monotone projection map, defined on the variables y_i, and extended naturally to monomials and terms in the y's. (Recall that a *term* of a polynomial is a monomial together with its coefficient.) Since each term t of g is a monomial multiplied by a positive coefficient, and since π is non-negative, $\pi(t)$ is either zero or a single monomial in the x's with nonzero coefficient. The former situation can happen only if t contains some variable y_i such that $\pi(y_i) = 0$. Let $\ker(\pi)$ denote the set $\{y_i | \pi(y_i) = 0\}$. Thus, for every term t of g that is disjoint from $\ker(\pi)$, $\pi(t)$ actually appears in f—possibly with a different coefficient, but still non-zero—since no two terms can cancel under projection by π.

Let $e_1, \dots, e_m$ be the coordinates on $\mathbb{R}^m$, the ambient space of $\mathrm{Newt}(g)$. Let K denote the subspace of $\mathbb{R}^m$ defined by the equations $e_i = 0$ for each i such that $y_i \in \ker(\pi)$. Let P be the intersection of $\mathrm{Newt}(g)$ with K, considered as a polytope in K; since all vertices of $\mathrm{Newt}(g)$ are non-negative, intersecting $\mathrm{Newt}(g)$ with a coordinate hyperplane, $e_i = 0$, results in a face of $\mathrm{Newt}(g)$, and thus P is a face of $\mathrm{Newt}(g)$. Note that P is exactly the convex hull of the exponent vectors of monomials in g that are disjoint from $\ker(\pi)$. In particular, since π is multiplicative on monomials, it induces a *linear* map ℓ_π from K to $\mathbb{R}^n$ (the ambient space of $\mathrm{Newt}(f)$). By the previous paragraph, the exponent vectors of f are exactly ℓ_π applied to the exponent vectors of monomials in g that are disjoint from $\ker(\pi)$. By the linearity of ℓ_π and the convexity of P and $\mathrm{Newt}(f)$, we have that $\mathrm{Newt}(f) = \ell_\pi(P)$, so P is an extension of $\mathrm{Newt}(f)$. Since P is defined by intersecting $\mathrm{Newt}(g)$ with additional linear equations, the lemma follows. $\square$

Several partial converses to our Main Lemma also hold. Perhaps the most natural and interesting of these is:

Observation 3.2. *Let R be a totally ordered semi-ring. Given any sequence of integral polytopes ($P_n \subseteq \mathbb{R}^n$) such that the poly($n$)-th cycle cover polytope is an extension of P_n along an affine linear map $\ell_n\colon \mathbb{R}^{poly(n)} \to \mathbb{R}^n$ with integer coefficients of polynomial bit-length, there is a sequence of polynomials $(f_n) \in \mathsf{VNP}_R$ such that* $\mathrm{Newt}(f_n) = P_n$ *and f is a monotone p-projection of the permanent.*

Proof. Let C_m denote the m-th cycle cover polytope, let $m(n)$ be a polynomial such that $C_{m(n)}$ is an extended formulation of P_n, and let $b(n)$ be a polynomial upper bound on the bit-length of the coefficients of ℓ_n. Let V_m denote the vertex set of the cycle cover polytope, i.e. the incidence vectors of cycle covers. Define f_n as $\sum_{\vec{e} \in V_m} \vec{y}^{\ell_n(\vec{e})}$, where $\vec{y} = (y_1, \dots, y_n)$, and the vector notation $\vec{y}^{\vec{e}'}$ is defined as $y_1^{e'_1} y_2^{e'_2} \cdots y_n^{e'_n}$. Note that ℓ_n has only integer coefficients by assumption, and each $\vec{e} \in V_m$ is integral, so the vector $\ell_n(\vec{e})$ is also integral, and the above expression is well-defined. By construction, every exponent vector of f_n is in $\ell_n(C_{m(n)}) = P_n$. Conversely, every vertex of P_n is an exponent vector of f_n, for any non-zero totally ordered semi-ring has characteristic zero (see Section 2). (Without noting this, it would be possible that k distinct vertices in V_m would get mapped to the same point under ℓ_n, for k a multiple of the characteristic, and then the corresponding monomials $\vec{y}^{\ell_n(\vec{e})}$ would add up to 0 in f_n.) Thus $\mathrm{Newt}(f_n) = P_n$. Furthermore, f_n is a monotone *nonlinear* projection of the permanent using the map $x_{ij} \mapsto \vec{y}^{\ell((0,0,\dots,1,\dots,0))}$, where the 1 is in the (i, j) position. Using the universality of the permanent and repeated squaring, this can easily be turned into a monotone *simple* projection of the permanent of size $\mathrm{poly}(m(n), b(n))$. $\square$

This can be generalized from the cycle cover polytopes and the permanent to arbitrary integral polytopes and the natural associated polynomial (the sum over all monomials whose exponent vectors are vertices of the polytope), but at the price of using "monomial projections"—in which each variable is replaced by a monomial—rather than simple projections. There ought to be a version of this observation over sufficiently large fields and allowing rational coefficients in ℓ, using Strassen's division trick [20], but the only such versions the author could come up with had so many hypotheses as to seem uninteresting.

4 Applications

4.1 Projection Lower Bounds

Remark 4.1. The following theorems hold over any totally ordered semi-ring, including the Boolean and-or semi-ring, the non-negative real numbers under multiplication and addition, and the tropical semi-ring of real numbers under addition and min. To see that this introduces no additional difficulty, note that over any totally ordered semi-ring R, the Newton polytope of a polynomial over R is still a polytope in a vector space over the real numbers, so standard results on polytopes and the cited results on extension complexity still apply.

Theorem 4.2. *Over any totally ordered semi-ring, the Hamiltonian Cycle polynomial is not a monotone affine p-projection of the permanent; in fact, any monotone affine projection from the permanent to the Hamiltonian Cycle polynomial has blow-up at least* $2^{\Omega(n)}$.

Proof. First, recall that if an n-variable polynomial is an affine projection of the $m \times m$ permanent, then it is a simple projection of the $(n+1)m \times (n+1)m$ permanent. For completeness we recall the brief proof: Let $\pi_{ij}(\vec{x})$ be the affine linear function corresponding to the variable y_{ij} of the $m \times m$ permanent, and write $\pi_{ij} = a_0 + a_1 x_1 + \cdots + a_n x_n$. Let G be the complete directed graph with loops on m vertices and edge weights y_{ij}. Replace the edge (i, j) by $n+1$ parallel edges with weights $a_0, a_1 x_1, \cdots, a_n x_n$. Add a new vertex on each of these parallel edges, splitting each parallel edge into two. For the edge weighted a_0, the two edges have weights $1, a_0$, and for the remaining edges the new edges get weights a_i, x_i. It is a simple and instructive exercise to see that this has the desired effect. Note also that if the original affine projection π was monotone, then so is the constructed simple projection.

Now we show the result for simple projections. If the Hamiltonian Cycle polynomial were a monotone projection of the permanent, then by the Main Lemma, some face of Newt(perm) would be an extension of Newt(HC).

The Newton polytope of the permanent is the cycle cover polytope (see Section 2). The cycle cover polytope can easily be described by the m^2 inequalities saying that all variables $x_{i,j}$ are non-negative, together with the equalities saying that each vertex has in-degree and out-degree exactly 1, namely $\sum_i x_{i,j} = 1$ for all j and $\sum_j x_{i,j} = 1$ for all i (it is easy to see that these are necessary; for sufficiency, see, e. g., [19, Theorem 18.1]). Since equalities do not count towards the complexity of a polytope, we have $c(\text{Newt}(\text{perm}_m)) \leq m^2$.

But the Newton polytope of the n-th Hamiltonian Cycle polynomial is exactly the TSP polytope, which by [17, Corollary 2] requires extension complexity $2^{\Omega(n)}$. $\square$

Theorem 4.3. *Over any totally ordered semi-ring, the perfect matching polynomial (or "unsigned Pfaffian") is not a monotone affine p-projection of the permanent; in fact, any* *monotone affine projection from the permanent to the perfect matching polynomial has blow-up at least* $2^{\Omega(n)}$.

Proof. The proof is the same as for the Hamiltonian Cycle polynomial, using [17, Theorem 1], which gives a lower bound of $2^{\Omega(n)}$ on the extension complexity of the perfect matching polytope, which is the Newton polytope of the perfect matching polynomial. □

Theorem 4.4. *Over any totally ordered semi-ring, for any* q*, the* q*-th cut polynomial is not a monotone affine p-projection of the permanent; in fact, any monotone affine projection from the permanent to the* q*-th cut polynomial has blow-up at least* $2^{\Omega(n)}$.

Proof. Use [6, Theorem 7], which says that $xc(\mathrm{Newt}(\mathrm{Cut}^2)) \geq 2^{\Omega(n)}$, as $\mathrm{Newt}(\mathrm{Cut}^2)$ is the cut polytope. The one additional observation we need is that $\mathrm{Newt}(\mathrm{Cut}^q)$ is just the $(q-1)$-scaled version of $\mathrm{Newt}(\mathrm{Cut}^2)$, and this rescaling does not affect the extension complexity. □

4.2 Monotone Formula and Circuit Lower Bounds

As pointed out by an anonymous reviewer, the universality of the permanent also holds in the monotone setting, so lower bounds on monotone projections from the permanent imply the same lower bounds on monotone formula size, and therefore quasi-polynomially related lower bounds on monotone circuit size. We assume circuits only have gates of bounded fan-in; with unbounded fan-in, rather than losing a factor of a half in the exponent of the exponent, we lose a factor of a third.

Proposition 4.5. *Any polynomial computable by a monotone formula of size* s *is a monotone projection of* perm_{s+1}.

Proof. The proof of the universality of the permanent given in [4, Proposition 2.16] works *mutatis mutandis* in the monotone setting. □

As a consequence of this, Theorems 4.2–4.4 are nearly tight, since every monotone polynomial in n variables of $\mathrm{poly}(n)$ degree can be written as a monotone formula of size $2^{O(n \log n)}$ (write it as a sum of monomials).

Corollary 4.6. *Over any totally ordered semi-ring, any monotone formula computing the Hamiltonian Cycle polynomial, the perfect matching polynomial, or the* q*-th cut polynomial has size at least* $2^{\Omega(n)}$*. Consequently, any monotone circuit computing these polynomials has size at least* $2^{\Omega(\sqrt{n})}$.

For the cut polynomials, we believe this result to be new. For the other polynomials, this provides a new proof of (slightly weaker versions of) previously known lower bounds. Namely, Jerrum and Snir gave a lower bound of $(n-1)((n-2)2^{n-3}+1) = 2^{n+\Omega(\log n)}$ on the monotone circuit size of HC [9, Section 4.4], and a lower bound of $n(2^{n-1}-1)$ on the monotone circuit size of the permanent [9, Section 4.3]. As the permanent is a monotone projection of the perfect matching polynomial—namely, restrict the perfect matching polynomial to a bipartite graph, e. g., by setting $x_{ij} = 0$ whenever i and j have the same parity—the same lower bound holds for the perfect matching polynomial.

Proof. The first part follows by combining Proposition 4.5 with Theorems 4.2–4.4. The second part follows from the fact that monotone circuits of size s can be balanced to have size $\mathrm{poly}(s)$ and depth $O(\log^2 s)$ (the proof in [24] works *mutatis mutandis* in the monotone setting), which can then be converted to monotone formulas of size $s^{O(\log s)} = 2^{O(\log^2 s)}$ by the usual conversion from bounded fan-in circuits to formulas. If there is a monotone circuit of size s computing any of these polynomials, there is thus a monotone formula of size $2^{O(\log^2 s)}$, which must be at least $2^{\Omega(n)}$, so $s \geq 2^{\Omega(n^{1/2})}$. □

5 Open Questions

Despite the common feeling that Razborov's super-polynomial lower bound [15] on monotone Boolean circuits for CLIQUE "finished off" monotone Boolean circuit lower bounds, several natural and interesting question remain. For example, does Directed s-t Connectivity require monotone Boolean circuits of size $\Omega(n^3)$? (A matching upper bound is given by the Bellman–Ford algorithm.) Is there a monotone Boolean reduction from general perfect matching to bipartite perfect matching? A positive answer to the following question would rule out such monotone (projection) reductions.

Open Question 5.1. *Extend Theorem 4.3 from formal polynomials over the Boolean semi-ring to Boolean functions.*

However, there are even easier questions, intermediate between the Boolean function case and the algebraic case considered in this paper; Jukna [10] discusses the notion of one polynomial "counting" another, which means that they agree on all $\{0,1\}$ inputs.

Open Question 5.2. *Prove that no monotone polynomial-size projection of the permanent agrees with the perfect matching polynomial on all* $\{0,1\}$ *inputs ("counts the perfect matching polynomial"). Similarly, prove that no monotone polynomial-size projection of the permanent counts the Hamiltonian cycle polynomial.*

S. Jukna points out (personal communication) that projections of the s-t connectivity polynomial correspond, even in the Boolean setting, to switching-and-rectifier networks, so the known lower bounds on monotone switching-and-rectifier networks (see, e. g., the survey [16]) imply that the Hamiltonian path polynomial and the permanent are not monotone p-projections of the s-t connectivity polynomial, even over the Boolean semi-ring. This helps explain why the only known monotone lower bound on the s-t connectivity polynomial that we are aware of [10] goes by a somewhat roundabout proof: Razborov's lower bound on CLIQUE [15], followed by Valiant's reduction from the clique polynomial to the Hamiltonian path polynomial [21], followed by a standard reduction from Hamiltonian path to counting s-t paths. In the course of discussing this, we were led to the following question; although the motivation for the question has since disappeared, it still seems like an interesting question about polytopes, whose answer may require new methods.

Open Question 5.3 (S. Jukna, personal communication)**.** *Is the m-th s-t path polytope an extension of the n-th TSP polytope (or n-th cycle cover polytope) with $m \leq poly(n)$?*

Since the separation problem for the s-t path polytope is NP-hard (see, e. g., [19, §13.1])—and the cycle cover polytope has low (extension) complexity—answering this question negatively seems to require more subtle understanding of these polytopes than "simply" an extended formulation lower bound.

Another example of a natural polytope question with a similar flavor comes from the cut polynomials. In combination with Bürgisser's results and questions on the cut polynomials [3] (discussed in Section 1), we are led to the following question.

Open Question 5.4. *Is the m-th cut polytope an extension of the n-th TSP polytope, for $m \leq poly(n)$?*

A negative answer would show that Cut^q is not complete for non-negative polynomials in $\mathsf{VNP}_{\mathbb{Q}}$ under monotone p-projections, though as with the example of the permanent, this is not necessarily an obstacle to being VNP-complete under general p-projections. Yet even the monotone completeness of the cut polynomials remains open. In fact, even more basic questions remain open:

Open Question 5.5. *Is every non-negative polynomial in* VNP *a monotone projection of the Hamiltonian Cycle polynomial? Is there any polynomial that is "positive* VNP*-complete" in this sense?*

To relate this to the current proofs of VNP-completeness of HC_n, we need to draw a distinction. Let $\mathsf{VP}_{\mathbb{R}}^{\geq 0}$ denote the polynomial families in $\mathsf{VP}_{\mathbb{R}}$ all of whose coefficients are non-negative, and let $\mathsf{mVP}_{\mathbb{R}}$ ("monotone VP") denote the class of families of polynomials with polynomially many variables, of polynomial degree, and computable by polynomial-size *monotone* circuits over $\mathbb{R}$. Similarly, define $\mathsf{VNP}_{\mathbb{R}}^{\geq 0}$ to be the non-negative polynomials in $\mathsf{VNP}_{\mathbb{R}}$, and $\mathsf{mVNP}_{\mathbb{R}}$ to be the function families of the form $f_n = \sum_{\vec{e} \in \{0,1\}}^{\mathrm{poly}(n)} g_m(\vec{e}, \vec{x})$, where $m \leq \mathrm{poly}(n)$ and $(g_m) \in \mathsf{mVP}_{\mathbb{R}}$.

Valiant's original completeness proof for the Hamiltonian Cycle polynomial [21] is "mostly" monotone: It uses polynomial-size formulas for the coefficients of the monomials (coming from the definition of VNP), but otherwise is entirely monotone. In other words, the proof shows that HC is mVNP-hard under monotone projections. However, we note that it's not clear whether HC is even *in* mVNP! Question 5.5 asks whether HC, or indeed any polynomial, is $\mathsf{VNP}^{\geq 0}$-complete under monotone projections; the question of whether there exist polynomials that are mVNP-complete under monotone projections also seems potentially interesting.

Finally, we ask about stronger notions of monotone reduction, which seem to require a different kind of proof technique. Recall that a *c-reduction* from f to g is a family of algebraic circuits for f with oracle gates for g.

Open Question 5.6. *Do the analogues of Theorems 4.2–4.4 hold for monotone bounded-depth c-reductions in place of affine p-projections? What about weakly-skew or even general monotone c-reductions?*

6 Subsequent Developments

Since the appearance of the preliminary version of this paper [7], our Main Lemma 3.1 has been used to prove that several other polynomials of combinatorial and complexity-theoretic interest are not sub-exponential-size projections of the permanent.

1. The n-th *satisfiability polynomial* over $\mathbb{F}_q$ is a polynomial in $n+8\binom{n}{3}$ variables denoted $X_1, \ldots, X_n$ and $\{Y_c : c \in C_n\}$ where C_n denote the set of clauses on 3 literals in n variables. It is defined as:

$$\mathrm{Sat}_n^q(X,Y) = \sum_{a \in \{0,1\}^n} \left(\prod_{i \in [n]} X_i^{q-1} \right) \left(\prod_{c \in C_n : c(a)=1} Y_c^{q-1} \right)$$

2. A *clow* in an n-vertex graph is a closed walk of length exactly n, in which the minimum-numbered vertex appears exactly once. The n-th *clow polynomial* over $\mathbb{F}_q$ is a polynomial in $\binom{n}{2} + n$ variables X_e for each edge e in the complete undirected graph K_n on n vertices and Y_v for each $v \in [n]$. It is defined as:

$$\mathrm{Clow}_n^q(X,Y) = \sum_{w:\text{clow of length } n} \left(\prod_{e:\text{edges in } w} X_e^{q-1} \right) \left(\prod_{v:\text{distinct vertices in } w} Y_v^{q-1} \right),$$

or more precisely,

$$\mathrm{Clow}_n^q(X,Y) = \sum_{w=[v_0,\ldots,v_{n-1}]} \left(\prod_{i \in [n]} X_{(v_{i-1}, v_{i \bmod n})}^{q-1} \right) \left(\prod_{v \in \{v_0,\ldots,v_{n-1}\}} Y_v^{q-1} \right),$$

where the sum is over clows w and v_0 denotes the minimum-numbered vertex in w.

3. The *clique polynomial* is a polynomial in $\binom{n}{2}$ variables X_e:

$$\mathrm{Clique}_n(X) = \sum_{T \subseteq \binom{[n]}{2} : T \text{ is a clique in } K_n} \prod_{e \in T} X_e.$$

Theorem (Mahajan and Saurabh [12, Theorems 2 and 6])**.** *Over any totally ordered semi-ring, any monotone affine projection from the permanent to* Sat_n^q *or to the clique polynomial requires blow-up at least* $2^{\Omega(\sqrt{n})}$. *Any monotone affine projection from the permanent to* Clow_n^q *requires blow-up at least* $2^{\Omega(n)}$.

As in Section 4.2, we get the following corollary. Again, we note that the lower bound on the clique polynomial over the Boolean semi-ring only works for the formal clique polynomial (in contrast to Razborov's result [15], which works for any monotone Boolean circuit computing the CLIQUE *function*).

Corollary 6.1. *Over any totally ordered semi-ring, any monotone formula computing* Clow_n^q *has size at least* $2^{\Omega(n)}$ *and any monotone circuit computing* Clow_n^q *has size at least* $2^{\Omega(\sqrt{n})}$. *Any monotone formula computing* Sat_n^q *or* Clique_n *has size at least* $2^{\Omega(\sqrt{n})}$ *and any monotone circuit computing these polynomials has size at least* $2^{\Omega(n^{1/4})}$.

For the clow and satisfiability polynomials, we believe this result to be new. For the clique polynomials, this provides a new proof of (a slightly weaker version of) the exponential monotone circuit lower bound due to Schnorr [18].[2]

Acknowledgment. We would like to thank Stasys Jukna for the question that motivated this paper [11], and `cstheory.stackexchange.com` for providing a forum for the question. We also thank Stasys for comments on a draft, pointing out the paper by Schnorr [18], and interesting discussions leading to Questions 5.2 and 5.3. We thank Leslie Valiant

[2]Schnorr showed a $\binom{n}{k} - 1$ lower bound on the monotone circuit size of the the k-th clique polynomial Clique_n^k, the sum over all cliques of size k, rather than all cliques. For $k = n/2$, this lower bound is asymptotically equal to $2^n/\sqrt{\pi n/2}$. The k-th clique polynomial Clique_n^k is the degree k homogeneous component of the clique polynomial Clique_n; by homogenization (implicit in Strassen's work, explicit in Valiant [23, Lemma 2]), any monotone circuit of size s for Clique_n can be converted into a monotone circuit of size $s(n/2+1)^2$ computing $\mathrm{Clique}_n^{n/2}$. Thus Schnorr's result implies a lower bound of $\Omega(2^n/n^{5/2})$ on the monotone circuit complexity of Clique_n.

for an interesting conversation that led to Question 5.5. We thank Ketan Mulmuley and Youming Qiao for collaborating on [8], which is why the author had Newton polytopes on the mind. We thank an anonymous reviewer for pointing out the monotone universality of the permanent and therefore the implications for monotone formula and circuit size (Section 4.2 and Corollary 6.1). The author was supported during this work by an Omidyar Fellowship from the Santa Fe Institute.

References

[1] Scott Aaronson. A linear-optical proof that the permanent is #P-hard. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 467(2136):3393–3405, 2011. doi:10.1098/rspa.2011.0232.

[2] Scott Aaronson and Alex Arkhipov. The computational complexity of linear optics. *Theory Comput.*, 9:143–252, 2013. Initially appeared in STOC '11. doi:10.4086/toc.2013.v009a004.

[3] Peter Bürgisser. On the structure of Valiant's complexity classes. *Discrete Math. Theor. Comput. Sci.*, 3(3):73–94, 1999. Initially appeared in STACS '98.

[4] Peter Bürgisser. *Completeness and reduction in algebraic complexity theory*, volume 7 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2000. doi:10.1007/978-3-662-04179-6.

[5] Nicolas de Rugy-Altherre. A dichotomy theorem for homomorphism polynomials. In *37th International Symposium on Mathematical Foundations of Computer Science*, MFCS '12, pages 308–322, 2012. doi:10.1007/978-3-642-32589-2_29.

[6] Samuel Fiorini, Serge Massar, Sebastian Pokutta, Hans Raj Tiwary, and Ronald de Wolf. Exponential lower bounds for polytopes in combinatorial optimization. *J. ACM*, 62(2):Art. 17, 23, 2015. Initially appeared in STOC '12. doi:10.1145/2716307.

[7] Joshua A. Grochow. Monotone projection lower bounds from extended formulation lower bounds. ECCC Tech. Report TR15-171 and arXiv:1510.08417 [cs.CC], 2015.

[8] Joshua A. Grochow, Ketan D. Mulmuley, and Youming Qiao. Boundaries of VP and VNP. In *43rd International Colloquium on Automata, Languages, and Programming*, ICALP '16, 2016. Preprint of the full version available as arXiv:1605.02815 [cs.CC].

[9] Mark Jerrum and Marc Snir. Some exact complexity results for straight-line computations over semirings. *J. ACM*, 29(3):874–897, 1982. doi:10.1145/322326.322341.

[10] Stasys Jukna. Lower bounds for monotone counting circuits. ECCC Tech. Report TR14-169, 2014.

[11] Stasys Jukna. "Why is Hamilton Cycle so different from permanent?". http://cstheory.stackexchange.com/q/27496/129, 2014.

[12] Meena Mahajan and Nitin Saurabh. Some complete and intermediate polynomials in algebraic complexity theory. In *Proceedings of 11th Computer Science Symposium*

in Russia, volume 9691 of *CSR '16, Lecture Notes in Computer Science*, pages 251–265, 2016. Full version available as ECCC Tech. Report TR16-038. doi:10.1007/978-3-319-34171-2_18.

[13] Henryk Minc. *Permanents*. Cambridge University Press, 1984.

[14] Thomas Muir and William H. Metzler. *A treatise on the theory of determinants*. Dover, New York, 1960.

[15] Alexander A. Razborov. Lower bounds of monotone complexity of the logical permanent function. *Mat. Zametki*, 37(6):887–900, 942, 1985.

[16] Alexander A. Razborov. Lower bounds for deterministic and nondeterministic branching programs. In L. Budach, editor, *Fundamentals of Computation Theory*, volume 529 of *Lecture Notes in Computer Science*, pages 47–60. Springer Berlin Heidelberg, 1991. URL: http://people.cs.uchicago.edu/~razborov/files/fct.pdf, doi:10.1007/3-540-54458-5_49.

[17] Thomas Rothvoß. The matching polytope has exponential extension complexity. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 263–272, New York, NY, USA, 2014. ACM. Preprint available as arXiv:1311.2369 [cs.CC]. doi:10.1145/2591796.2591834.

[18] Claus-Peter Schnorr. A lower bound on the number of additions in monotone computations. *Theoret. Comput. Sci.*, 2(3):305–315, 1976. doi:10.1016/0304-3975(76)90083-9.

[19] Alexander Schrijver. *Combinatorial optimization. Polyhedra and efficiency. Vol. A*, volume 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Paths, flows, matchings, Chapters 1–38.

[20] Volker Strassen. Vermeidung von Divisionen. *J. Reine Angew. Math.*, 264:184–202, 1973.

[21] Leslie G. Valiant. Completeness classes in algebra. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing*, STOC '79, pages 249–261, New York, NY, USA, 1979. ACM. doi:10.1145/800135.804419.

[22] Leslie G. Valiant. The complexity of computing the permanent. *Theoret. Comput. Sci.*, 8(2):189–201, 1979. doi:10.1016/0304-3975(79)90044-6.

[23] Leslie G. Valiant. Negation can be exponentially powerful. *Theoret. Comput. Sci.*, 12(3):303–314, 1980. doi:10.1016/0304-3975(80)90060-2.

[24] Leslie G. Valiant, Sven Skyum, Stuart Berkowitz, and Charles Rackoff. Fast parallel computation of polynomials using few processors. *SIAM J. Comput.*, 12(4):641–644, 1983. doi:10.1137/0212043.

[25] Jacobus H. van Lint and Richard M. Wilson. *A course in combinatorics*. Cambridge University Press, 2001.

[26] Tzu-Chieh Wei and Simone Severini. Matrix permanent and quantum entanglement of permutation invariant states. *J. Math. Phys.*, 51(9), 2010. doi:10.1063/1.3464263.

ECCC ISSN 1433-8092
http://eccc.hpi-web.de