Separating the Power of Monotone Span Programs over Different Fields

Amos Beimel[*] Enav Weinreb[†]

Abstract

Monotone span programs are a linear-algebraic model of computation. They are equivalent to linear secret sharing schemes and have various applications in cryptography and complexity. A fundamental question is how the choice of the field in which the algebraic operations are performed effects the power of the span program. In this paper we prove that the power of monotone span programs over finite fields of different characteristics is incomparable; we show a super-polynomial separation between any two fields with different characteristics, answering an open problem of Pudlák and Sgall 1998. Using this result we prove a super-polynomial lower bound for monotone span programs for a function in $\mathrm{uniform}-\mathcal{NC}^2$ *(and therefore in* $\mathcal{P}$*), answering an open problem of Babai, Wigderson, and Gál 1999. (All previous lower bounds for monotone span programs were for functions not known to be in* $\mathcal{P}$*.) Finally, we show that quasi-linear schemes, a generalization of linear secret sharing schemes introduced in Beimel and Ishai 2001, are stronger than linear secret sharing schemes. In particular, this proves, without any assumptions, that non-linear secret sharing schemes are more efficient than linear secret sharing schemes.*

1 Introduction

The relation between computational complexity and linear algebra is an important research direction with two main avenues. On one hand, algebraic techniques were used to prove lower bounds in combinatorics [BF92, GR01, Juk01] and complexity, e.g., [Smo87, MS82, Raz90]. On the other hand, algebraic computational models, which capture the essence of linear algebra, were defined. Such models include, for example, arithmetic circuit, Boolean circuits with MOD_p gates, and the Blum-Shub-Smale model of computation [BSS89].

In this paper we discuss the algebraic computational model of span programs, introduced by Karchmer and Wigderson [KW93]. Intuitively, span programs capture the power of basic linear algebraic operations – the rank and dependency of a set of vectors. More specifically, a monotone span program is presented as a matrix over some field, with rows labelled by variables. The span program accepts an input if the rows whose variables are satisfied by the input span a fixed nonzero vector. The size of a span program is its number of rows. A detailed definition is given in Section 2.

This paper deals with the role of the field in algebraic models of computation. Part of the specification of algebraic models of computation, in particular span programs, is the field in which the arithmetic operations are performed. A fundamental question is how the choice of the field, and especially its characteristic, effects the power of the model. As different fields may differ substantially in their structure, especially when the characteristics of the fields are different, it would be natural to expect computational models defined over different fields to differ significantly in their power. The major result separating the power of algebraic models of computation over different fields was the seminal paper by Smolensky for bounded depth circuits with MOD_p gates [Smo87]. Lower bounds related to the characteristic of the field are also known for polynomial calculus proofs [BI99]. However, the power of the field in algebraic models of computation is yet to be fully understood.

Our Results. The main contribution of this paper is showing that the power of monotone span programs over finite fields of different characteristic is incomparable. Prior to this work, the best separation known for monotone span programs, was a logarithmic separation for the threshold function [KW93].[1] In this paper we show a super-polynomial separation between any two fields with different characteristics, answering an open problem of [PS98]. That is, for every fixed prime number p we describe a function which has a small monotone span program over the field with p elements, but requires a monotone span program of size $n^{\Omega(\sqrt{\log n})}$ over any field whose characteristic is not p (including fields with characteristic 0).

Our second contribution concerns the functions for which lower bounds for monotone span programs have been proved. The best known lower bound for monotone span programs, proved by Gál [Gál98], is $n^{\Omega(\log n)}$ (improving previous results of [BGP97, BGW99]). However, all the

[*]Dept. of Computer Science, Ben-Gurion University, Beer-Sheva 84105, Israel. beimel@cs.bgu.ac.il.

[†]Dept. of Computer Science, Ben-Gurion University, Beer-Sheva 84105, Israel. weinrebe@cs.bgu.ac.il. Partially supported by a Kreitman Foundation Fellowship.

[1]It was known that span programs over finite fields with the same characteristic basically have the same power.

known super-polynomial lower bounds were for functions in $\mathcal{NP}$, not known to be in $\mathcal{P}$. We show a lower bound of $n^{\Omega(\sqrt{\log n})}$ for a function in $\mathrm{uniform}-\mathcal{NC}^2$ (and therefore in $\mathcal{P}$), thus answering an open problem of [BGW99].[2]

Our third contribution concerns secret sharing schemes, which are an important tool in cryptography, introduced by Blakley [Bla79], Shamir [Sha79], and Ito, Saito, and Nishizeki [ISN87]. A *secret sharing scheme* enables a dealer to share a secret among a set of parties, such that only some pre-defined authorized subsets will be able to reconstruct the secret from their shares. The authorized sets correspond to a monotone Boolean function $f : \{0,1\}^n \to \{0,1\}$, where n is the number of parties and the authorized subsets are the subsets with their characteristic vectors in $f^{-1}(1)$. The efficiency of a secret sharing scheme is the overall size of the shares given to the parties. Monotone span programs are equivalent to a subclass of secret sharing schemes called *"linear secret sharing schemes*." Monotone span programs were also used in other cryptographic applications, e.g., [NPR99, CDM00]. Beimel and Ishai [BI01] showed functions that, under plausible assumptions, have no efficient linear secret sharing scheme but yet have an efficient non-linear secret sharing scheme. Furthermore, they introduced the class of quasi-linear secret sharing schemes. In this paper we show that quasi-linear schemes are stronger than linear schemes. In particular, this proves, without any assumptions, that non-linear schemes are more efficient than linear schemes.

Highlights of the Techniques. Proving a separation between the power of two models of computation requires a function with both a lower bound for one model, and an upper bound for the other. To get the lower bound for monotone span program over a certain field, we use the method of [Gál98], which is based on [Raz90]. In the center of Gál's method is a matrix whose rank over this field is much larger than its combinatorial cover number. To get the upper bound for the same function for monotone span programs over another field, we require the cover to have an additional property which is related to the characteristic of the field. As an example, for GF(2) we require that each entry of the matrix is covered by an odd number of rectangles. Our use of combinatorial covers and their properties is borrowed from communication complexity (see [KN97] for background on communication complexity). In particular, we use ideas similar to [DKMW03], where they considered the model of counting communication complexity.

The main technical contribution of this paper is in constructing such a matrix and in proving that it satisfies the desired properties. In particular, the matrix we construct checks whether two linear subspaces over GF(p) have nontrivial intersection. Not surprisingly, the matrix reflects linear algebraic computations over GF(p), which are difficult to simulate over fields with characteristics different than p.

[2] We note that every function which has a polynomial monotone $\mathcal{NC}^1$ circuit has a polynomial monotone span program, and every function which has a polynomial span program over a small field has a polynomial $\mathcal{NC}^2$ circuit.

Organization. In Section 2 we supply some preliminaries. In Section 3 we give a general method for proving a separation between the power of monotone span programs over fields with different characteristics. Next, in Section 4 we apply this general method to achieve a separation of $n^{\Omega(\sqrt{\log n})}$ for an explicit function. Finally, in Section 5, we use this separation to exhibit a monotone function in $\mathrm{uniform}-\mathcal{NC}^2$ that has no polynomial monotone span program, and to prove that there exist secret sharing schemes stronger than the linear secret sharing schemes.

2 Preliminaries

We start with the definition of our main computational model – span programs.

Definition 2.1 (Span Program [KW93]) *A* span program *over a field F is a triplet $\widehat{M} = \langle M, \rho, \vec{v}\rangle$, where M is a matrix over F, $\vec{v}$ is a non-zero row vector called the target vector (it has the same number of coordinates as the number of columns in M), and ρ is a labelling of the rows of M by literals from $\{x_1, \ldots, x_n, \overline{x}_1, \ldots, \overline{x}_n\}$ (every row is labelled by one literal, and the same literal can label many rows).*

A span program accepts or rejects an input by the following criterion. For every input $u \in \{0,1\}^n$ define the submatrix M_u of M consisting of those rows whose labels are satisfied by the assignment u. The span program $\widehat{M}$ accepts *u if and only if $\vec{v} \in \mathrm{span}(M_u)$, i.e., some linear combination of the rows of M_u gives the vector $\vec{v}$. A span program* computes *a Boolean function f if it accepts exactly those inputs u where $f(u) = 1$. The* size *of $\widehat{M}$ is the number of rows in M.*[3]

A span program is called monotone *if the labels of the rows are only positive literals $\{x_1, \ldots, x_n\}$. Monotone span programs compute only monotone functions, and every monotone Boolean function can be computed by a monotone span program. The size of the smallest monotone span program over F that computes f is denoted by* $\mathrm{mSP}_F(f)$.

Combinatorial Rectangles and Covers. Combinatorial rectangles and covers are a useful tool in communication complexity, and are used in this work in a similar way. Let X and Y be arbitrary finite sets. A combinatorial rectangle is a set $X_0 \times Y_0$, where $X_0 \subseteq X$ and $Y_0 \subseteq Y$. A *cover* of $X \times Y$ is a set $\mathcal{R}$ of rectangles such that every pair $\langle x, y\rangle \in X \times Y$ belongs to at least one rectangle in $\mathcal{R}$.

Let M be a Boolean $|X| \times |Y|$ matrix such that the rows of M are indexed by the elements of X, and the columns of M are indexed by the elements of Y. We say that a rectangle $R_0 = X_0 \times Y_0$, where $X_0 \subseteq X$ and $Y_0 \subseteq Y$, is a *monochromatic* rectangle if there exists a $b \in \{0,1\}$ such that for every $x \in X_0$ and $y \in Y_0$, it holds that $M[x, y] = b$.

[3] The choice of the fixed non-zero vector $\vec{v}$ does not effect the size of the span program. It is always possible to replace $\vec{v}$ by another vector $\vec{v}'$ via a change of basis without changing the function computed and the size of the span program. Most often $\vec{v}$ is chosen to be the $\vec{1}$ vector (with all entries equal 1).

If $b = 1$ we call R_0 a 1-*rectangle*, and if $b = 0$ we call R_0 a 0-*rectangle*. We say that a cover $\mathcal{R}$ is a *monochromatic* cover of M if every rectangle $R \in \mathcal{R}$ is a monochromatic rectangle. If $\mathcal{R}$ is a set of 1-rectangles that cover all the 1-entries of M, then $\mathcal{R}$ is called a 1-*cover* of M. If $\mathcal{R}$ is a set of 0-rectangles that cover all the 0-entries of M, we call $\mathcal{R}$ a 0-*cover* of M.

Linear Subspaces. We use basic linear algebra to find a function that is easy for span programs over one field and hard for span programs over another field. For a prime number p, we denote by GF(p) the unique finite field with p elements.

Let k be a positive integer, and let p be a prime. Denote by $V_k^{2k}(p)$ the set of all k-dimensional subspaces of GF$(p)^{2k}$, and denote by $v_k^{2k}(p)$ the number of such subspaces, that is, $v_k^{2k}(p) = \left|V_k^{2k}(p)\right|$. To prove our result, we count the number of subspaces satisfying a certain property. Towards this aim, we will use the following easy algebraic claim. We say that two linear spaces U and W are different if there exists a vector $\vec{v}$ such that $\vec{v} \in U$ and $\vec{v} \notin W$ or vice versa.

Claim 2.2 *Let k be positive integer, F be a field, and M be a matrix with k rows such that $\mathrm{rank}_F(M) = k$. Let T_1, T_2 be matrices with k rows each, where $T_1 \neq T_2$. Define M_1 (respectively, M_2) to be the matrix resulting from concatenating the matrix T_1 (respectively, T_2) to M, that is $M_i = (M|T_i)$ for $i \in \{1, 2\}$. Then, the linear spaces spanned by the rows of M_1 and M_2 are different.*

Proof: Since $T_1 \neq T_2$, there exists an index $j \in \{1, \ldots, k\}$, such that the rows $T_1[j]$ and $T_2[j]$ are different. Let $\vec{r} = M_1[j]$, that is, $\vec{r}$ is the jth row of M_1. We show that $\vec{r}$ is not spanned by the rows of M_2. Assume there exist a combination of the rows of M_2 that spans $\vec{r}$. That is, $\vec{r} = \sum_{i=1}^{k} \alpha_i M_2[i]$ for some $\alpha_1, \ldots, \alpha_k \in F$. Let m be the number of columns in M, and consider the restriction of the above sum to the first m coordinates. It holds that $M[j] = \sum_{i=1}^{k} \alpha_i M[i]$. Since M has k rows and $\mathrm{rank}_F(M) = k$, we get that $\alpha_j = 1$ and $\alpha_i = 0$ for every $i \neq j$. Thus, $\vec{r} = M_2[j]$, that is, $M_1[j] = M_2[j]$, contradicting the fact that $T_1[j] \neq T_2[j]$. □

One application of Claim 2.2 is the following corollary, which gives a lower bound on $v_k^{2k}(p)$.

Corollary 2.3 *Let k be a positive integer, and let p be a prime. Then $v_k^{2k}(p) \geq p^{k^2}$.*

Proof: Let I_k be the $k \times k$ unit matrix, T be an arbitrary $k \times k$ matrix over GF(p), and M_1 be the $k \times 2k$ matrix that is a concatenation of I_k and T. There are p^{k^2} different choices of T, and therefore p^{k^2} different ways to construct M_1. By Claim 2.2, each such M_1 represents a different element of $V_k^{2k}(p)$, and thus $v_k^{2k}(p) \geq p^{k^2}$. □

It is easy to see that $v_k^{2k}(p) < p^{2k^2}$, since this is the number of ways to choose *any* k vectors from GF$(p)^{2k}$. Thus, we have $p^{k^2} \leq v_k^{2k}(p) < p^{2k^2}$.

We will denote by $\vec{e}_j$ the jth unit vector, that is, the vector that is 1 in the jth coordinate, and 0 in all the others. We say that a non-zero vector has a *leading* 1, if the first non zero coordinate in the vector is 1. Let p be a prime, ℓ be a positive integer, and U be a subspace of dimension ℓ over GF(p). Then, the number of vectors with a leading 1 in U is $\frac{p^\ell - 1}{p-1}$. We will denote by $\mathrm{char}(F)$ the characteristic of the field F.

3 The General Method for Separation

We want to construct a function that is hard for monotone span programs over fields with characteristic different than p, and easy for monotone span programs over GF(p), where p is a prime. We use the method of [Gál98] to get the lower bound for monotone span programs over fields with characteristic different than p. In the center of this method is a matrix with a large gap between its rank and the size of its monochromatic cover. To get a small upper bound for monotone span programs over GF(p), we shell require the cover to have an additional property which we call 1-mod-p, that is, for every entry of the matrix, the number of rectangles covering it is equivalent to 1 modulo p. Generally speaking, the number of variables in f, the function we prove the separation for, is equal to the number of rectangles in a cover. A detailed description is given below.

3.1 The Lower Bound

Let M be a matrix, and let $\mathcal{R}$ be a monochromatic cover of M. Recall that $\mathcal{R}$ is a set of rectangles. Denote $n = |\mathcal{R}|$, that is, $\mathcal{R} = \{R_1, \ldots, R_n\}$, where $R_i = X_i \times Y_i$. A vector in $\{0,1\}^n$ can be viewed as a characteristic vector of a subset of $\mathcal{R}$. Throughout the paper, we identify each such vector with its corresponding subset. We define two subsets of $\{0,1\}^n$. The first set is $\mathrm{Acc} = \{\langle z_1, \ldots, z_n\rangle :$ There is a row x in M s.t. $z_i = 1$ iff $x \in X_i$ for every $i \in \{1, \ldots, n\}\}$. Thus, a subset of $\mathcal{R}$ is in Acc if and only if it contains exactly all the rectangles covering one of the rows of M. The second set is $\mathrm{Rej} = \{\langle z_1, \ldots, z_n\rangle :$ There is a column y in M s.t. $z_i = 1$ iff $y \notin Y_i$ for every $i \in \{1, \ldots, n\}\}$. Hence, Rej contains subsets of rectangles containing *all but* the rectangles covering one column of M. An example for Acc and Rej is described in Figure 1.

The lower bound is achieved using the following theorem which is implicit in [Gál98]:

Theorem 3.1 ([Gál98]) *Let M be a Boolean matrix, $\mathcal{R}$ be a monochromatic cover of M, and* Acc *and* Rej *as defined above. If $f : \{0,1\}^n \to \{0,1\}$ is a monotone function such that $f(x) = 1$ for every $x \in \mathrm{Acc}$, and $f(y) = 0$ for every $y \in \mathrm{Rej}$ then $\mathrm{mSP}_F(f) \geq \mathrm{rank}_F(M)$, for every field F.*

That is, we get the lower bound for every function f accepting Acc, and rejecting Rej. Note that there are no requirements concerning inputs $z \notin (\mathrm{Acc} \cup \mathrm{Rej})$, except for monotonicity.

3.2 The Upper Bound

To prove a gap between the power of monotone span programs over the different fields, we need the cover $\mathcal{R}$ to be a

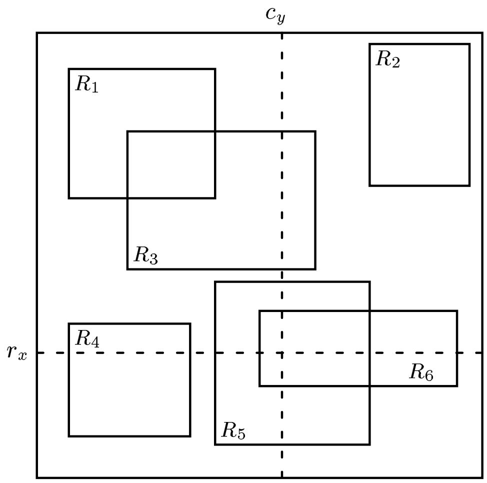

Figure 1. An illustration of elements in the sets Acc and Rej. The set x in Acc corresponding to r_x is $x = \{R_4, R_5, R_6\}$, the rectangles that cover r_x. The set y in Rej corresponding to c_y is $y = \{R_1, R_2, R_4\}$, the rectangles that *do not* cover c_y. Note that the rectangles in the figure *do not* form a cover.

monochromatic $1\text{-}\mathrm{mod}\text{-}p$ *cover*, according to the following definition:

Definition 3.2 *Let M be a Boolean matrix. A set $\mathcal{R}$ of combinatorial rectangles is called a* monochromatic $1\text{-}\mathrm{mod}\text{-}p$ cover *of M, if $\mathcal{R}$ is a monochromatic cover of M, and, for each entry of M, the number of rectangles covering it is equivalent to* 1 *modulo p.*

Given a small monochromatic $1\text{-}\mathrm{mod}\text{-}p$ cover of M, we construct a monotone span program over GF(p) that accepts Acc and rejects Rej. The gap will hold for the function computed by this span program.

Consider the following monotone span program $\widehat{P}$ over GF(p). The program $\widehat{P}$ associates a row with each rectangle of $\mathcal{R}$, and a column with each column of the matrix M. The row associated with the rectangle $R_i = X_i \times Y_i$ is 1 in the column labelled by y if $y \in Y_i$, that is, if the rectangle R_i covers the column y in M. Otherwise, this entry in $\widehat{P}$ is 0. Note that $size(\widehat{P}) = n$, that is, there is exactly one row for each variable.

Lemma 3.3 *The program $\widehat{P}$ accepts every $x \in \mathrm{Acc}$ and rejects every $y \in \mathrm{Rej}$.*

Proof: We first prove that $\widehat{P}$ accepts every $x \in \mathrm{Acc}$. Specifically, we will show that since $\mathcal{R}$ is a $1\text{-}\mathrm{mod}\text{-}p$ cover, the sum of the rows labelled by the rectangles of x is the vector $\vec{1}$, and thus x is accepted by $\widehat{P}$. That is, we show that for every column of $\widehat{P}$, the rows labelled by x sum to 1 in this column. Towards this goal, fix a column c_y. Since $x \in \mathrm{Acc}$, there exists a row r_x in M, such that x is the characteristic vector of the set of rectangles covering r_x. According to the definition of $\widehat{P}$, for every rectangle $R \in x$, the entry $\langle R, c_y\rangle$ of $\widehat{P}$ is 1 if and only if R covers c_y. On the other hand, $R \in x$ if and only if R covers the row r_x. Thus, the sum over the rows of $\widehat{P}$ associated with x in the column c_y is exactly the number of rectangles covering both c_y and r_x, that is, the number of rectangles covering the entry $\langle r_x, c_y\rangle$ in M. Since $\mathcal{R}$ is a $1\text{-}\mathrm{mod}\text{-}p$ cover, this number is 1 modulo p. Thus, the sum of the rows labelled by x is the vector $\vec{1}$, and x is accepted by $\widehat{P}$.

Let $y \in \mathrm{Rej}$. We show that there is no linear combination of the rows labelled by the rectangles of y that give the vector $\vec{1}$. Since $y \in \mathrm{Rej}$, there is a column c_y of M that is not covered by any of the rectangles in the subset of $\mathcal{R}$ associated with y. Hence, all the rows of $\widehat{P}$ corresponding to rectangles from y are 0 in the column associated with c_y. Therefore, every combination of the rows labelled by y is 0 in this column. Thus, the vector $\vec{1}$ is not a linear combination of these rows, and y is rejected by $\widehat{P}$. $\square$

Combining Theorem 3.1 and Lemma 3.3, we get the main theorem of this section:

Theorem 3.4 (Separation Theorem) *Let M be a Boolean matrix, and let $\mathcal{R}$ be a monochromatic* $1\text{-}\mathrm{mod}\text{-}p$ *cover of M of size n. Then there exists a function f, with n variables, such that* $\mathrm{mSP}_{\mathrm{GF}(p)}(f) = n$ *and* $\mathrm{mSP}_F(f) \geq \mathrm{rank}_F(M)$ *for every field F.*

Proof: Denote by f_P the function computed by $\widehat{P}$. By Lemma 3.3, f_P accepts Acc and rejects Rej, and thus by Theorem 3.1 $\mathrm{mSP}_F(f_P) \geq \mathrm{rank}_F(M)$. On the other hand, $size(\widehat{P}) = n$ and thus $\mathrm{mSP}_{\mathrm{GF}(p)}(f) = n$. $\square$

4 The Linear Subspaces Zero Intersection Function

In this section we show an explicit matrix, with a high rank over fields with characteristic different than p, and a small monochromatic $1\text{-}\mathrm{mod}\text{-}p$ cover. Thus, by Theorem 3.4 we get a function f with a super-polynomial gap between $\mathrm{mSP}_{\mathrm{GF}(p)}(f)$ and $\mathrm{mSP}_F(f)$ where F is a field such that $\mathrm{char}(F) \neq p$. We define the desired matrix in two steps: in the first step we define the matrix M_{ZI}, and prove it has full rank over fields with $\mathrm{char} \neq p$. In the second step we use M_{ZI} to define another matrix, M_{LZI}, which has both a high rank over fields with $\mathrm{char} \neq p$, and a small monochromatic $1\text{-}\mathrm{mod}\text{-}p$ cover.

Let k be a positive integer and p be a prime.[4] The *Zero Intersection* (ZI) function determines whether the intersection of two k-dimensional linear subspaces of $\mathrm{GF}(p)^{2k}$ is the subspace $\left\{\vec{0}\right\}$. More formally, define $\mathrm{ZI}_k^p : V_k^{2k}(p) \times V_k^{2k}(p) \to \{0,1\}$ as follows: $\mathrm{ZI}_k^p(U, W) = 1$, where U and W are subspaces in $V_k^{2k}(p)$, if and only if $\dim(U \cap W) = 0$.

[4]Through this section the reader should think of k as small. That is, we construct a function with n variables and $k \approx \sqrt{\log n}$.

Recall that the intersection of any two linear subspaces is a linear subspace.

We represent ZI_k^p by a $v_k^{2k}(p) \times v_k^{2k}(p)$ matrix denoted $M_{\mathrm{ZI}_k^p}$. Each row and each column of $M_{\mathrm{ZI}_k^p}$ is labelled by a subspace $U \in V_k^{2k}(p)$, and each entry $M_{\mathrm{ZI}_k^p}[U,W]$ is equal to $\mathrm{ZI}_k^p(U,W)$. Denote by r_U the row in $M_{\mathrm{ZI}_k^p}$ associated with the subspace $U \in V_k^{2k}(p)$. We will use ZI instead of ZI_k^p, and M_{ZI} instead of $M_{\mathrm{ZI}_k^p}$, when k and p are clear from the context.

4.1 Analyzing the Rank of M_{ZI}

The next theorem shows that M_{ZI} has full rank over any field with $\mathrm{char} \neq p$.

Theorem 4.1 *Let k be a positive integer, p be a prime, and F be a field such that $\mathrm{char}(F) \neq p$. Then, $M_{\mathrm{ZI}_k^p}$ has full rank over F.*

Proof: To prove that the matrix has full rank, it is sufficient to show that any unit vector is spanned by the rows of the matrix. Recall that the columns of the matrix are labelled by subspaces from $V_k^{2k}(p)$. For every $U \in V_k^{2k}(p)$ we consider the unit vector $\vec{e}_U \in \mathrm{GF}(p)^{v_k^{2k}(p)}$ and show that it is spanned by the rows of M_{ZI}. Specifically, we show a combination of the rows of the matrix spanning $\vec{e}_U$ having a special structure: The coefficient of $\vec{r}_Z$, the row labelled by $Z \in V_k^{2k}(p)$, depends only on the dimension of the subspace $U \cap Z$. More precisely, we show there are constants $\alpha_0, \ldots, \alpha_k \in F$, such that

$$\vec{e}_U = \sum_{d=0}^{k} \alpha_d \sum_{\substack{W \in V_k^{2k}(p) \\ \dim(Z \cap U) = d}} \vec{r}_Z. \tag{1}$$

Fix $W \in V_k^{2k}(p)$, and consider c_W, the column of M_{ZI} associated with W. We have to show that with the appropriate constants $\alpha_0, \ldots, \alpha_k \in F$, the above expression is 0 in this column if $W \neq U$, and is 1 if $W = U$. Computing the sum in the column c_W, we add α_d for every subspace Z such that $\mathrm{ZI}(Z,W) = 1$ (i.e., $\dim(Z \cap W) = 0$) and $\dim(Z \cap U) = d$. This motivates the following definition:

Definition 4.2 *Let $U, W \in V_k^{2k}(p)$ be subspaces, and let ℓ be an integer such that $\dim(U \cap W) = \ell$. Define $H_k^p(\ell, d)$ to be the* number *of subspaces $Z \in V_k^{2k}(p)$ such that $\dim(U \cap Z) = d$ and $\dim(W \cap Z) = 0$.*

¿From symmetry arguments, the number $H_k^p(\ell, d)$ is independent of the choice of U and W. We will write $H_k(\ell, d)$ instead of $H_k^p(\ell, d)$, when p is clear from the context.

To summarize, we need to show there are constants $\alpha_0, \ldots, \alpha_k \in F$ such that:

1. For each $0 \leq \ell \leq k-1$, it holds that $\sum_{d=0}^{k} \alpha_d \cdot H_k(\ell, d) = 0$. That is, the sum over any column labelled with $W \neq U$ equals 0, where for a subspace $W \in V_k^{2k}(p)$ such that $\dim(U \cap W) = \ell$, the relevant equation is the ℓ-th equation.
2. $\sum_{d=0}^{k} \alpha_d \cdot H_k(k, d) = 1$. That is, the sum over the column associated with U is 1.

Putting things differently, we view the numbers $H_k(\ell, d)$ for $\ell, d \in \{0, \ldots, k\}$ as a $(k+1) \times (k+1)$ matrix over F.[5] According to the above conditions we have to prove there are $\alpha_0, \ldots, \alpha_k \in F$ such that $H_k \langle \alpha_0, \alpha_1, \ldots, \alpha_k \rangle^{\perp} = \langle 0, 0, \ldots, 1 \rangle^{\perp}$. We show that H_k is invertible over F, and thus we can find $\alpha_0, \ldots, \alpha_k$ using H_k^{-1} as follows. $\langle \alpha_0, \alpha_1, \ldots, \alpha_k \rangle^{\perp} = H_k^{-1} \langle 0, 0, \ldots, 1 \rangle^{\perp}$.

In the next two claims, we show that H_k is upper-left triangular, where the numbers on the secondary diagonal are non-zero in F, thus H_k has full rank over F.

Claim 4.3 *Let k be a positive integer, ℓ and d be non-negative integers, p be a prime, and H_k be as above. If $\ell + d > k$ then $H_k^p(\ell, d) = 0$.*

Proof: Let $U, W \in V_k^{2k}(p)$, where $\dim(U \cap W) = \ell$. We have to show that since $\ell + d > k$ there is no subspace $Z \in V_k^{2k}(p)$, such that $\dim(Z \cap U) = d$ and $\dim(Z \cap W) = 0$. Assume toward contradiction that there exists such Z. Let $B_{U \cap W} = \langle \vec{w}_1, \ldots, \vec{w}_\ell \rangle$ be a basis of the subspace $U \cap W$. Let $B_{U \cap Z} = \langle \vec{z}_1, \ldots, \vec{z}_d \rangle$ be a basis for $U \cap Z$. Consider the set of vectors $X = B_{U \cap W} \cup B_{U \cap Z}$. First note that $X \subseteq U$, that is, all the vectors in X are in the subspace U. Since $\dim(U) = k$ and $|X| = \ell + d > k$, the set X must be linearly dependent. Thus, there must be a nontrivial combination of the vectors of X, giving the vector $\vec{0}$, that is, $\sum_{i=1}^{\ell} \lambda_i \vec{w}_i + \sum_{i=1}^{d} \delta_i \vec{z}_i = \vec{0}$. Since both $B_{U \cap W}$ and $B_{U \cap Z}$ are linearly independent, the non-zero vector $\vec{v} = \sum_{i=1}^{\ell} \lambda_i \vec{w}_i$ is spanned by both $B_{U \cap W}$ and $B_{U \cap Z}$. Since $U \cap W \subseteq W$ and $U \cap Z \subseteq Z$, we get that $\vec{v} \in W \cap Z$ and thus, $\dim(W \cap Z) > 0$, contradicting the assumption that $\dim(W \cap Z) = 0$. ■(Claim 4.3)

We shell need the following notation for the next claim: Let $B = \langle \vec{v}_1, \ldots, \vec{v}_{2k} \rangle$ be a basis of $\mathrm{GF}(p)^{2k}$. Let $Z \in V_k^{2k}(p)$ and $B_Z = \langle \vec{z}_1, \ldots, \vec{z}_k \rangle$ be a basis for Z, such that for every $i \in \{1, \ldots, k\}$ we have $z_i = \sum_{j=1}^{2k} \beta_{i,j} \vec{v}_j$. Then we call the $k \times 2k$ matrix $(\beta_{i,j})$ the *representation matrix* of B_Z *according to* B.

Claim 4.4 *Let k be a positive integer, ℓ and d be non-negative integers, and p be a prime. If $\ell + d = k$ then $H_k^p(\ell, d) = p^{\ell(k+d)}$.*

Proof: Let $U, W \in V_k^{2k}(p)$ be any subspaces such that $\dim(U \cap W) = \ell$. We must show that the number of subspaces Z such that $\dim(Z \cap U) = d$ and $\dim(Z \cap W) = 0$ is $p^{\ell(k+d)}$. We will first define the term *canonic representation* of a subspace in $V_k^{2k}(p)$. Next, we will show that each subspace Z such that $\dim(Z \cap U) = d$ and $\dim(Z \cap W) = 0$ has a canonic representation. Then we will show that every different canonic representation is associated with a different

[5] Since $H_k(\ell, d)$ may be a number not in F, we will replace it by $H_k(\ell, d) \bmod c$, where c is the characteristic of F. If the characteristic of F is 0, $H_k(\ell, d)$ will always be in F.

subspace Z such that $\dim(Z \cap U) = d$ and $\dim(Z \cap W) = 0$. Thus, the number of such subspaces is equal to the number of different canonic representations. To complete the proof, we will show that the number of such canonic representations is $p^{\ell(k+d)}$. The canonic representation is defined according to a specific basis of $\mathrm{GF}(p)^{2k}$. Consider a basis $B_{U,W}$ of $\mathrm{GF}(p)^{2k}$ defined as follows:

$$B_{U,W} = \langle \vec{v}_1, \ldots, \vec{v}_\ell, \vec{u}_1, \ldots, \vec{u}_d, \vec{w}_1, \ldots, \vec{w}_d, \vec{x}_1, \ldots, \vec{x}_\ell \rangle$$

where:

- $\langle \vec{v}_1, \ldots, \vec{v}_\ell \rangle$ is a basis of $U \cap W$. Recall that $\dim(U \cap W) = \ell$.
- $\langle \vec{u}_1, \ldots, \vec{u}_d \rangle$ is an expansion of $\langle \vec{v}_1, \ldots, \vec{v}_\ell \rangle$ to a basis of U. Recall that $\dim(U) = k$ and $d + \ell = k$.
- $\langle \vec{w}_1, \ldots, \vec{w}_d \rangle$ is an expansion of $\langle \vec{v}_1, \ldots, \vec{v}_\ell \rangle$ to a basis of W. Recall that $\dim(W) = k$ as well.
- $\langle \vec{x}_1, \ldots, \vec{x}_\ell \rangle$ is an expansion of $\langle \vec{v}_1, \ldots, \vec{v}_\ell, \vec{u}_1, \ldots, \vec{u}_{k-\ell}, \vec{w}_1, \ldots, \vec{w}_{k-\ell} \rangle$ to a basis of $\mathrm{GF}(2)^{2k}$. Here there are ℓ vectors since $2k - (\ell + d + d) = \ell$.

We say that a subspace $Z \in V_k^{2k}(p)$ has a *canonic representation* according to $B_{U,W}$ if it has a basis whose representation matrix according to $B_{U,W}$ is as described in Figure 2. The matrix in Figure 2 is a $k \times 2k$ matrix. Each entry in zones (b), (g), and (h) must be 0. The entries in zones (d) and (f) must form the unit matrices I_ℓ and I_d respectively. Each entry in zones (a), (c), and (e) can take any value from $\mathrm{GF}(p)$. First we show that every subspace $Z \in V_k^{2k}(p)$ such

(a) ?	(b) 0	(c) ?	(d) I_ℓ
(e) ?	(f) I_d	(g) 0	(h) 0
$v_1, \ldots, v_\ell$	$u_1, \ldots, u_d$	$w_1, \ldots, w_d$	$x_1, \ldots, x_\ell$

Figure 2. A canonic representation of a subspace $Z \in V_k^{2k}(p)$ with $\dim(U \cap Z) = d$ and $\dim(W \cap Z) = 0$.

that $\dim(Z \cap U) = d$ and $\dim(Z \cap W) = 0$ has a canonic representation according to $B_{U,W}$. Let $Y = Z \cap U$. Note that $\dim(Y) = d$. Let $B_Y = \langle \vec{y}_1, \ldots, \vec{y}_d \rangle$ be a basis of Y, and let $B_Z = \langle \vec{y}_1, \ldots, \vec{y}_d, \vec{z}_1, \ldots, \vec{z}_\ell \rangle$ be an expansion of B_Y to a basis of Z. Consider M_Z, the representation matrix of B_Z according to $B_{U,W}$. Since $Y \subseteq U$, all the entries in the zones (g) and (h) are 0 as required. We claim that we can perform elementary operations on the lower part of M_Z so that we get the matrix I_d in zone (f). Otherwise, we would get a row $\vec{r}$ that is $\vec{0}$ in zone (f), but this would leave all the non-zero entries of $\vec{r}$ in zone (e). Since zone (e) represents the basis vectors from $U \cap W$, this would mean $\dim(Z \cap W) > 0$, contradicting the properties of Z. It is left to set zone (d) to I_ℓ and all the entries in zone (b) to 0. Setting all the entries in zone (b) to 0 can be done by elementary operations on the upper part of M_Z using the rows from the lower part, which now form the unit matrix I_d in zone (f). (This would change the entries in zone (a), but we have no constraints on this zone.) We claim that we can set zone (d) to be I_ℓ by elementary operations on the upper part of M_Z. Otherwise we would get a row $\vec{r}$ that is all zero in zone (d). Thus $\vec{r}$ has non-zero entries only in zones (a) and (c), but then it again implies that $\vec{r}$ represents a vector from W, contradicting the fact that $\dim(Z \cap W) = 0$.

Next we prove that *every* subspace $Z \in V_k^{2k}(p)$ which can be represented in the above canonic form, satisfies $\dim(Z \cap W) = 0$ and $\dim(Z \cap U) = d$. Let M_Z be a canonic representation of Z according to $B_{U,W}$. Since M_Z has I_ℓ and I_d as sub-matrices, we have $\mathrm{rank}_{\mathrm{GF}(p)} M_Z = k$ and thus $Z \in V_k^{2k}(p)$. Now suppose $\dim(Z \cap W) > 0$. Then we can span a vector $w \in W$ by the rows of M_Z. This vector has to be zero in the coordinates labelled by $\vec{u}_1, \ldots, \vec{u}_d$, and by $\vec{x}_1, \ldots, \vec{x}_\ell$, but this cannot be done by a non-trivial combination of the rows of M_Z. Thus, $\dim(W \cap Z) = 0$. The lower part of M_Z is non-zero only in coordinates labelled by vectors from U, and since it has I_d as a sub-matrix, we get that $\dim(Z \cap U) \geq d$. Now suppose that $\dim(Z \cap U) = d' > d$. Then we have $\dim(Z \cap U) = d'$, $\dim(Z \cap W) = 0$, and $\dim(U \cap W) = \ell$, where $\ell + d' > \ell + d = k$, which is impossible by Claim 4.3. Therefore, $\dim(U \cap Z) = d$.

To complete the proof, we show that any two subspace who have different canonic representations over $B_{U,W}$ are different. To see that, note that the matrix $S = \begin{pmatrix} 0 & I_\ell \\ I_d & 0 \end{pmatrix}$ is a sub-matrix of any canonic representation. The matrix S is clearly of rank k, and thus, by Claim 2.2 any two subspaces with different canonic representation are different.

Therefore, when constructing a subspace Z, with $\dim(Z \cap U) = d$ and $\dim(Z \cap W) = 0$, the freedom in only in the entries marked with '?' in Figure 2. Since there are p possibilities for every such entry, and the number of such entries is $(k \cdot \ell) + (\ell \cdot d) = \ell(k + d)$, we conclude that $H_k(\ell, d) = p^{\ell(k+d)}$. ■(claim 4.4)

Since the characteristic of F is different than p, every power of p is non zero over F. Therefore, as argued above, we proved that H_k has full rank over F, and the theorem follows. □(Theorem 4.1)

In Corollary 2.3 we proved that $v_k^{2k}(p) \geq p^{k^2}$. Since M_{ZI_k} is a $v_k^{2k}(p) \times v_k^{2k}(p)$ matrix, $\mathrm{rank}_F(M_{\mathrm{ZI}_k}) \geq p^{k^2}$.

4.2 A Small 1-mod-p Cover for the Zeros of M_{ZI}

To apply Theorem 3.4 on an explicit matrix, we need this matrix to have a small monochromatic 1-mod-p cover. We next show that there is a small 1-mod-p cover for the 0's of M_{ZI}. We do not know if there exists a small 1-mod-p cover for the 1's of M_{ZI}. Thus, we are not able to use M_{ZI} directly, and we use it in Section 4.3 to build the matrix

$M_{\rm LZI}$, which has a small 1-mod-p cover for both the 1's and the 0's.

To give some intuition on the cover of $M_{\rm LZI}$ we show a 1-mod-p cover for the 0's of $M_{\rm ZI}$ of size less than p^{2k}. This should be compared to the number of rows in $M_{\rm ZI}$ which is $p^{\Theta(k^2)}$. Define the cover $\mathcal{R}$ as follows: Let $\vec{v} \in GF(p)^{2k}$ be a vector with a leading 1, that is, the first non-zero coordinate of $\vec{v}$ is 1. We add the rectangle $R_{\vec{v}} = X_{\vec{v}} \times Y_{\vec{v}}$ to the cover $\mathcal{R}$, where:

$$X_{\vec{v}} = \left\{U \in V_k^{2k}(p) : \vec{v} \in U\right\}, \text{ and}$$

$$Y_{\vec{v}} = \left\{W \in V_k^{2k}(p) : \vec{v} \in W\right\}.$$

That is, $R_{\vec{v}}$ contains the rows and the columns of $M_{\rm ZI}$ labelled by subspaces that contain the vector $\vec{v}$. The rectangle $R_{\vec{v}}$ is a 0-rectangle, since for each $U \in X_{\vec{v}}$ and $W \in Y_{\vec{v}}$ it holds that $\vec{v} \in U \cap W$, hence $\dim(U \cap W) \neq 0$, and thus $\mathrm{ZI}(U, W) = 0$. We claim $\mathcal{R}$ is a 1-mod-p cover of the 0's of $M_{\rm ZI}$. Let $\langle U, W\rangle$ be an entry of $M_{\rm ZI}$, such that $\mathrm{ZI}(U, W) = 0$. Then $\dim(U \cap W) > 0$. Therefore, the entry $\langle U, W\rangle$ is covered by any rectangle $R_{\vec{v}}$ such that $\vec{v} \in U \cap W$. Since $U \cap W$ is a linear subspace of $GF(p)^{2k}$, it has $\frac{p^\ell - 1}{p-1}$ vectors with a leading 1, where $\ell = \dim(U \cap W) \geq 1$. Since $\frac{p^\ell-1}{p-1} \equiv \frac{-1}{-1} \equiv 1 \pmod p$, the number of rectangles covering the entry $\langle U, W\rangle$ is equivalent to 1 modulo p. Since there are $\frac{p^{2k}-1}{p-1}$ different vectors with a leading 1 in $\mathrm{GF}(p)^{2k}$, the size of the 0-cover is $\frac{p^{2k}-1}{p-1}$.

4.3 The List Version of the Zero Intersection Function

To get a matrix with a high rank over fields with characteristic different than p, and a small monochromatic 1-mod-p cover, we define the function LZI, the list version of the Zero Intersection function. The idea of using the list version of functions has been used in communication complexity [MS82] (see, e.g., [KN97]). Define $\mathrm{LZI}_k^p : (V_k^{2k}(p))^k \times (V_k^{2k}(p))^k \to \{0,1\}$ as follows:

$$\mathrm{LZI}_k^p(\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle) = 1 \iff$$

$$\exists i \in \{1 \ldots k\} \text{ such that } \mathrm{ZI}_k^p(A_i, B_i) = 1.$$

That is, LZI_k^p gets k instances of ZI_k^p, and outputs the value 1 iff ZI_k^p outputs 1 on at least one of the given instances. The matrix $M_{\rm LZI}$, representing LZI, is defined in a similar way to $M_{\rm ZI}$. The next two lemmas show that $M_{\rm LZI}$ has a small 1-mod-p cover.

Lemma 4.5 *There is a monochromatic* 1-mod-p *cover of the* 0*'s of* $M_{\rm LZI}$ *of size smaller than* p^{2k^2}.

Proof: We build the 0-cover $\mathcal{R}_0$ of the 0's of $M_{\rm LZI}$ in a similar way to the 0-cover for $M_{\rm ZI}$ built in Section 4.2. Let $\langle \vec{v}_1, \ldots, \vec{v}_k\rangle \in (\mathrm{GF}(p)^{2k})^k$ be a tuple of k vectors from $\mathrm{GF}(p)^{2k}$, each with a leading 1. The rectangle in $\mathcal{R}_0$ corresponding to $\langle \vec{v}_1, \ldots, \vec{v}_k\rangle$ is $R = X \times Y$ where:

$$X = \{\langle A_1, \ldots, A_k\rangle \in (V_k^{2k}(p))^k : \vec{v}_i \in A_i \text{ for each } i \in \{1, \ldots, k\}\}, \text{ and}$$

$$Y = \{\langle B_1, \ldots, B_k\rangle \in (V_k^{2k}(p))^k : \vec{v}_i \in B_i \text{ for each } i \in \{1, \ldots, k\}\}.$$

First we show R is a 0-rectangle. If $\langle A_1, \ldots, A_k\rangle \in X$ and $\langle B_1, \ldots, B_k\rangle \in Y$, then $\vec{v}_i \in A_i \cap B_i$ for every $i \in \{1, \ldots, k\}$, and thus $\mathrm{ZI}(A_i, B_i) = 0$ for every $i \in \{1, \ldots, k\}$. Therefore, $\mathrm{LZI}(\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle) = 0$.

Next we show that for every 0-entry of $M_{\rm LZI}$, the number of rectangles covering it is equivalent to 1 modulo p. Let $\langle\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle\rangle \in (V_k^{2k}(p))^k \times (V_k^{2k}(p))^k$ such that $\mathrm{LZI}(\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle) = 0$. The entry $\langle\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle\rangle$ is covered by any rectangle associated with a tuple of k non-zero vectors $\langle \vec{v}_1, \ldots, \vec{v}_k\rangle$, such that $\vec{v}_i \in A_i \cap B_i$, for every $i \in \{1, \ldots, k\}$, and has a leading 1. Since $A_i \cap B_i$ is a linear subspace, the number of vectors with a leading 1 in $A_i \cap B_i$ is $\frac{p^{\ell_i}-1}{p-1}$ where $\ell_i = \dim(A_i \cap B_i) \geq 1$. Thus, the number of rectangles covering $\langle\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle\rangle$ is a product of numbers that are equivalent to 1 modulo p, and therefore is equivalent to 1 modulo p itself.

The number of 0-rectangles in $\mathcal{R}_0$ is the number of tuples of k vectors with a leading 1 from $\mathrm{GF}(p)^{2k}$, that is, $(\frac{p^{2k}-1}{p-1})^k < p^{2k^2}$. (This is much smaller than the number of rows in $M_{\rm LZI}$, which is $p^{\Theta(k^3)}$.) ■

Now we show the cover $\mathcal{R}_1$ for the 1's of $M_{\rm LZI}$. The natural way to do it would be to associate a rectangle $R = X \times Y$ with each pair $\langle i, U\rangle$, such that $i \in \{1, \ldots, k\}$, and $U \in V_k^{2k}(p)$, where:

$$X = \left\{\langle A_1, \ldots, A_k\rangle \in (V_k^{2k}(p))^k : A_i = U\right\}, \text{ and}$$

$$Y = \left\{\langle B_1, \ldots, B_k\rangle \in (V_k^{2k}(p))^k : \dim(U \cap B_i) = 0\right\}.$$

That is, any input pair having $ZI(A_i, B_i) = 1$ in the ith instance, will be covered by the rectangle associated with i and A_i. Clearly, R is a 1-rectangle. We show $\mathcal{R}_1$ is a 1-cover of $M_{\rm LZI}$. Let $\langle\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle\rangle \in (V_k^{2k}(p))^k \times (V_k^{2k}(p))^k$ such that $\mathrm{LZI}(\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle) = 1$. Then there exist an index $i \in \{1, \ldots, k\}$ such that $\dim(A_i \cap B_i) = 0$. Thus, the entry $\langle\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle\rangle$ is covered by the rectangle associated with $\langle i, A_i\rangle$.

The problem with this choice of $\mathcal{R}_1$ is that it is not a 1-mod-p cover. For example, if $\langle A_1, \ldots, A_k\rangle$ and $\langle B_1, \ldots, B_k\rangle$ have exactly p instances $\langle A_i, B_i\rangle$ such that $\mathrm{ZI}(A_i, B_i) = 1$, then the number of rectangles covering the entry $\langle\langle A_1, \ldots, A_k\rangle, \langle B_1, \ldots, B_k\rangle\rangle$ will be equivalent to 0 modulo p. To solve this problem, we require i to be the index of the *first* instance of ZI, such that $\mathrm{ZI}(A_i, B_i) = 1$.

Lemma 4.6 *There is a monochromatic* 1-mod-p *cover for the* 1*'s of* $M_{\rm LZI}$ *of size smaller than* p^{4k^2}.

Proof: Associate a rectangle $R = X \times Y$ with any pair $\langle\langle \vec{v}_1,\ldots,\vec{v}_{i-1}\rangle, U\rangle$, where $\langle \vec{v}_1,\ldots,\vec{v}_{i-1}\rangle$ is a tuple of $i-1$ vectors with a leading 1 from GF$(p)^{2k}$ where $1 \le i \le k$, and $U \in V_k^{2k}(p)$ is a subspace. The sets X and Y are defined as follows:

$$X = \{\langle A_1,\ldots,A_k\rangle \in (V_k^{2k}(p))^k : \vec{v}_j \in A_j \text{ for each } j \in \{1,\ldots,i-1\} \text{ and } A_i = U\}, \text{ and}$$

$$Y = \{\langle B_1,\ldots,B_k\rangle \in (V_k^{2k}(p))^k : \vec{v}_j \in B_j \text{ for each } j \in \{1,\ldots,i-1\} \text{ and } \dim(B_i \cap U) = 0\}.$$

To see that R is a 1-rectangle take $\langle A_1,\ldots,A_k\rangle \in X$ and $\langle B_1,\ldots,B_k\rangle \in Y$. Then, $\dim(A_i \cap B_i) = \dim(U \cap B_i) = 0$, and thus $\mathrm{ZI}(A_i,B_i) = 1$. Therefore, $\mathrm{LZI}(\langle A_1,\ldots,A_k\rangle,\langle B_1,\ldots,B_k\rangle) = 1$.

We next show that for every 1-entry of M_{LZI}, the number of rectangles covering it is equivalent to 1 modulo p. Let $\langle\langle A_1,\ldots,A_k\rangle,\langle B_1,\ldots,B_k\rangle\rangle \in (V_k^{2k}(p))^k \times (V_k^{2k}(p))^k$ such that $\mathrm{LZI}(\langle A_1,\ldots,A_k\rangle,\langle B_1,\ldots,B_k\rangle) = 1$. Let i be the smallest index such that $\dim(A_i \cap B_i) = 0$. Then the entry $\langle\langle A_1,\ldots,A_k\rangle,\langle B_1,\ldots,B_k\rangle\rangle$ is covered by any rectangle associated with a pair $\langle\langle \vec{v}_1,\ldots,\vec{v}_{i-1}\rangle, A_i\rangle$, such that $\vec{v}_j \in A_j \cap B_j$ for every $j \in \{1,\ldots,i-1\}$. Since the number of vectors with a leading 1 in $A_j \cap B_j$ for every $j \in \{1,\ldots,i\}$ is equivalent to 1 modulo p, the number of such rectangles is equivalent to 1 modulo p as well.

The size of $\mathcal{R}_1$ is smaller than the number of ways to choose k vectors with a leading 1 from GF$(p)^{2k}$, and a subspace from $V_k^{2k}(p)$, and thus is smaller than $p^{2k^2} \cdot v_k^{2k}(p) < p^{4k^2}$. ■

By taking the union of the 0-cover from Lemma 4.5 and the 1-cover from Lemma 4.6 we get the following corollary.

Corollary 4.7 *M_{LZI} has a monochromatic $1\text{-mod-}p$ cover of size smaller than p^{5k^2}.*

We proved in Theorem 4.1 that $\mathrm{rank}_F(M_{\mathrm{ZI}_k}) \ge p^{k^2}$. We analyze the rank of M_{LZI_k} over F, given $\mathrm{rank}_F(M_{\mathrm{ZI}_k})$. The following lemma is implied by the properties of the Kronecker product.

Lemma 4.8 *Let k be a positive integer and let p be a prime. Then $\mathrm{rank}_F(M_{\mathrm{LZI}^p_k}) = p^{\Omega(k^3)}$.*

We are ready to prove our main result:

Theorem 4.9 (Main Result) *Let p be a fixed prime. Then there exist a family of functions $\{f_n\}_{n\in\mathcal{N}}$, such that $\mathrm{mSP}_{\mathrm{GF}(p)}(f_n) = n$ and for every field F with characteristic different than p, it holds that $\mathrm{mSP}_F(f_n) = n^{\Omega(\sqrt{\log n})}$.*

Proof: For a positive number k, denote by n_k the size of the monochromatic $1\text{-mod-}p$ cover for M_{LZI} given by Corollary 4.7. We first show f_n for each n of the form $n = n_k$ for some positive k. According to Corollary 4.7, M_{LZI_k} has a monochromatic $1\text{-mod-}p$ cover of size n, which is smaller than p^{5k^2}. According to Lemma 4.8, we have that $\mathrm{rank}_F(M_{\mathrm{LZI}_k}) = p^{\Omega(k^3)}$. In terms of n, we have

$$n^{\sqrt{\log_p n}} \le (p^{5k^2})^{\sqrt{\log_p(p^{5k^2})}} = (p^{5k^2})^{\sqrt{5k^2}}.$$

By Theorem 3.4 we get that there is a function f_n in n variables, such that $\mathrm{mSP}_{\mathrm{GF}(p)}(f) = n$ and $\mathrm{mSP}_F(f) \ge p^{\Omega(k^3)} = n^{\Omega(\sqrt{\log n})}$. The last equality holds since p is a constant. By padding arguments, the result holds for every value of n. □

5 A Super-polynomial Lower Bound for a Function in $\mathrm{uniform} - \mathcal{NC}^2$

In this section we show a monotone function that is computable by uniform-$\mathcal{NC}^2$ circuits, and does not have a polynomial monotone span program over any field.[6] For comparison, all the previous super-polynomial lower bounds are for function not known to be in $\mathcal{P}$.

Denote by $f^2 = \{f^2_n\}_{n\in\mathcal{N}}$ the family of functions given by Theorem 4.9 for $p = 2$. Denote by $f^3 = \{f^3_n\}_{n\in\mathcal{N}}$ the family of functions given by Theorem 4.9 for $p = 3$. Define the family of functions $f = \{f_{2n}\}_{n\in\mathcal{N}}$ to be $f_{2n}(x_1,\ldots,x_n,y_1,\ldots,y_n) = f^2_n(x_1,\ldots,x_n) \wedge f^3_n(y_1,\ldots,y_n)$.

We show a $\mathrm{uniform} - \mathcal{NC}^2$ family of circuits for f. Let $\widehat{P_2}$ be the monotone span program over GF(2) that computes f^2. Since $size(\widehat{P_2}) = n$, and since linear algebra over fixed finite fields is in log-space uniform-$\mathcal{NC}^2$ [Ber84, Mul87, BDHM92, KW93], there exists an $\mathcal{NC}^2$ circuit C_2 that computes f^2. Similarly, there exists an $\mathcal{NC}^2$ circuit C_3 that computes f^3. Thus, the $\mathcal{NC}^2$ circuit $C = C_2 \wedge C_3$ computes f.

The problem with the circuit C, as described, is that it is not uniform. The reason is that the number of columns in the monotone span programs derived from Theorem 4.9 is super-polynomial. It is known that there exists an equivalent monotone span program in which the number of columns does not exceed the number of rows. However, the mere *existence* of a monotone span program with a small number of columns does not yield a uniform-$\mathcal{NC}^2$ circuit. To get uniform circuits we have to show an *explicit* monotone span program with a small number of columns that can be generated in space $O(\log n)$. We do this in Section 5.1.

We next show that f has no monotone span program over any field. Assume there is a polynomial monotone span program $\widehat{Q}$ that computes f over some field F. Let c be the characteristic of F. If $c \ne 2$ then the restriction of $\widehat{Q}$ to inputs of the form $x_1,\ldots,x_n \cdot 1^n$, gives a new monotone span program $\widehat{Q}_2$ of polynomial size over F that computes f^2 (as any restriction of a function with a small monotone span program has a small monotone span program [KW93]), contradicting the fact that f^2 has no polynomial monotone span program over fields with characteristic different that 2. If $c = 2$ then $c \ne 3$ and we get the contradiction for f^3 in a similar way. Thus,

[6] In this paper uniform means log-space uniform.

Theorem 5.1 *There exist a family of monotone functions* $\{f_n\}_{n\in\mathcal{N}}$ *that is computable by a* $\text{uniform} - \mathcal{NC}^2$ *family of circuits, having* $\text{mSP}_F(f_n) = n^{\Omega(\sqrt{\log n})}$ *for every field* F.

5.1 Reducing the Number of Columns

In Theorem 4.9 we introduced a function f_P such that $\text{mSP}_{\text{GF}(p)}(f_n) = n$ and $\text{mSP}_F(f_n) = n^{\Omega(\sqrt{\log n})}$. In this section we want to construct a family of uniform-$\mathcal{NC}^2$ circuits for f_P.

It is known that any function that has a polynomial monotone span program has a family of $\mathcal{NC}^2$ circuits. Since any monotone span program with m rows that computes a function f has an equivalent monotone span program with no more than m columns, we can deduce the existence of a family of $\mathcal{NC}^2$ circuits that computes f. However, we want a uniform family of circuits. Since any transformation from a monotone span program with an arbitrary number of columns to an equivalent program with a smaller number of columns has to go over all the columns of the big original program, we cannot use the generic span program for f_P, as presented in Section 3.4. In this section we show a monotone span program with a linear number of both rows and columns, that accepts Acc and rejects Rej. We show that the span program can be generated in space $O(\log n)$, and by this we ensure the uniformity of the $\mathcal{NC}^2$ circuits.

Let $\mathcal{R}_{\text{LZI}}$ be the monochromatic 1-mod-p cover of M_{LZI} described in Corollary 4.7, and consider the following monotone span program $\widehat{S}$: The program $\widehat{S}$ has a column for each k-tuple $\langle \vec{v}_1, \ldots, \vec{v}_k\rangle \in (\text{GF}(p)^{2k})^k$ where each $\vec{v}_i$ is a vector with a leading 1 from $\text{GF}(p)^{2k}$. Thus the number of columns in $\widehat{S}$ is smaller than the number of rectangles in $\mathcal{R}_{\text{LZI}}$, and hence is linear in the number of variables.

Recall that in $\mathcal{R}_{\text{LZI}}$ there are two types of rectangles:

- 0-rectangles. A 0-rectangle for every k-tuple of vectors $\langle \vec{v}_1, \ldots, \vec{v}_k\rangle \in (\text{GF}(p)^{2k})^k$, each with a leading 1.
- 1-rectangles. We associated a 1-rectangle $R = X \times Y$, with any pair $\langle\langle \vec{v}_1, \ldots, \vec{v}_{i-1}\rangle, U\rangle$, where $\langle \vec{v}_1, \ldots, \vec{v}_{i-1}\rangle$ is a tuple of $i-1$ vectors with a leading 1 from $\text{GF}(p)^{2k}$, where $1 \le i \le k$, and $U \in V_k^{2k}(p)$ is a subspace.

Every rectangle is assigned a row in $\widehat{S}$. Let R be a rectangle in $\mathcal{R}_{\text{LZI}}$, and let c be a column in $\widehat{S}$ labelled with the tuple $\langle \vec{v}_1, \ldots, \vec{v}_k\rangle$. Then the value of the entry $\widehat{S}[R, c]$ is defined as follows:

For a 0-rectangle R, let $\langle \vec{u}_1, \ldots, \vec{u}_k\rangle$ be the k tuple of vectors associated with R. We set $\widehat{S}[R, c] = 1$ if $\vec{u}_i = \vec{v}_i$ for every $i \in \{1, \ldots, k\}$. Otherwise, $\widehat{S}[R, c] = 0$.

For a 1-rectangle R, let $\langle\langle \vec{u}_1, \ldots, \vec{u}_{i-1}\rangle, U_i\rangle$ be the $(i-1)$-tuple of vectors and the subspace associated with R. In this case set $\widehat{S}[R, c] = 1$ if $\vec{u}_j = \vec{v}_j$ for every $j \in \{1, \ldots, i-1\}$ and $v_i \notin U_i$. Otherwise $\widehat{S}[R, c] = 0$.

By putting the rows corresponding to 0-rectangles in the upper part of $\widehat{S}$, the upper block of $\widehat{S}$ is in fact the unit matrix I. To compute an entry in the lower part of $\widehat{S}$, we only have to check if a vector in $\text{GF}(p)^{2k}$ belongs to a subspace, where $k = O(\sqrt{\log n})$. This can be easily done in space $O(\log n)$. Thus, $\widehat{S}$ can be generated in log-space. The proof of the next lemma is omitted due to lack of space.

Lemma 5.2 *The program* $\widehat{S}$ *accepts every* $x \in \text{Acc}$ *and rejects every* $y \in \text{Rej}$.

5.2 Span Programs and Secret Sharing Schemes

Secret sharing schemes, introduced by Blakley [Bla79], Shamir [Sha79], and Ito, Saito, and Nishizeki [ISN87], are a cryptographic tool allowing a dealer to share a secret between a set of parties such that only some pre-defined authorized subset of parties can reconstruct the secret from their shares. The reader is referred to [Sim92] and [Sti92] for a more formal and detailed discussion on secret sharing schemes.

The authorized sets in a secret sharing scheme are described by a monotone Boolean function $f : \{0,1\}^n \to \{0,1\}$, where n is the number of parties and the authorized subsets are the subsets with their characteristic vectors in $f^{-1}(1)$. Most of the known secret sharing schemes are linear schemes, that is, schemes in which the shares are a linear combination of the secret and some random field elements. Linear schemes are equivalent to monotone span programs where the total size of the shares is related to the size of the corresponding monotone span program. Beimel and Ishai [BI01] showed functions that, under plausible assumptions, have no efficient linear secret sharing scheme but yet have an efficient non-linear secret sharing scheme. However, prior to this work, no secret sharing schemes were proved more powerful than linear schemes, without any assumptions.

A quasi-linear secret sharing scheme [BI01] is obtained by composing a finite number of linear secret sharing schemes, possibly over different fields. Beimel and Ishai [BI01] have shown that under the assumption that the power of monotone span programs over different fields is incomparable, quasi-linear schemes are super-polynomially stronger than linear schemes. Their proof is very similar to the proof of Theorem 5.1. That is, the functions described in Theorem 5.1 have, by definition, a small quasi-linear secret sharing scheme but cannot have a small linear scheme.

Theorem 5.3 *There is an explicit family of functions* $\{f_n\}_{n\in\mathcal{N}}$ *such that the complexity of every linear secret sharing scheme for the family is* $n^{\Omega(\sqrt{\log n})}$*, and yet the family has a polynomial quasi-linear secret sharing scheme.*

Acknowledgments. We thank Yinnon Haviv for his very valuable help and Anna Gál for many helpful discussions.

References

[BDHM92] G. Buntrock, C. Damm, U. Hertrampf, and C. Meinel. Structure and importance of the logspace-mod class. *Math. Systems Theory*, 25:223–237, 1992.

[Ber84] S. J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Inform. Process. Lett.*, 18:147–150, 1984.

[BF92] L. Babai and P. Frankl. *Linear Algebra Methods in Combinatorics*. University of Chicago, 1992. Preliminary Version 2.

[BGP97] A. Beimel, A. Gál, and M. Paterson. Lower bounds for monotone span programs. *Computational Complexity*, 6(1):29–45, 1997.

[BGW99] L. Babai, A. Gál, and A. Wigderson. Superpolynomial lower bounds for monotone span programs. *Combinatorica*, 19(3):301–319, 1999.

[BI99] E. Ben-Sasson and R. Impagliazzo. Random CNF's are Hard for the Polynomial Calculus. In *40th FOCS*, pages 415–421, 1999.

[BI01] A. Beimel and Y. Ishai. On the power of nonlinear secret-sharing. In *Conf. on Computational Complexity*, pages 188 – 202, 2001.

[Bla79] G. R. Blakley. Safeguarding cryptographic keys. In *Proc. of the 1979 AFIPS National Computer Conference*, volume 48 of *AFIPS Conference proceedings*, pages 313–317. AFIPS Press, 1979.

[BSS89] L. Blum, M. Shub, and S. Smale. On a theory of computation and complexity over the real numbers; NP completeness, recursive functions and universal machines. *Bulletin of the AMS (new series)*, 21(1):1–46, 1989.

[CDM00] R. Cramer, I. Damgård, and U. Maurer. General secure multi-party computation from any linear secret-sharing scheme. In B. Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *LNCS*, pages 316–334. Springer, 2000.

[DKMW03] C. Damm, M. Krause, C. Meinel, and S. Waack. On relations between counting communication complexity classes. *J. of Computer and System Sciences*, 2003. To appear. Prelimenary version: STACS '92.

[Gál98] A. Gál. A characterization of span program size and improved lower bounds for monotone span programs. In *30th STOC*, pages 429–437, 1998.

[GR01] C. Godsil and G. Royle. *Algebraic Graph Theory*, volume 207 of *Graduate Texts in Mathematcs*. Springer, 2001.

[ISN87] M. Ito, A. Saito, and T. Nishizeki. Secret sharing schemes realizing general access structure. In *Proc. of the IEEE Global Telecommunication Conf., Globecom 87*, pages 99–102, 1987. Journal version: Multiple Assignment Scheme for Sharing Secret. *J. of Cryptology*, 6(1):15-20, 1993.

[Juk01] S. Jukna. *Extremal Combinatorics with Applications in Computer Science*. Texts in Theoretical Comp. Sci. Springer-Verlag, 2001.

[KN97] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge Univ. Press, 1997.

[KW93] M. Karchmer and A. Wigderson. On span programs. In *Proc. of the 8th Structure in Complexity Theory*, pages 102–111, 1993.

[MS82] K. Mehlhorn and E. M. Schmidt. Las vegas is better than determinism in vlsi and distributed computing. In *Proc. of the 14th STOC*, pages 330–337, 1982.

[Mul87] K. Mulmuley. A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. *Combinatorica*, 7:101–104, 1987.

[NPR99] M. Naor, B. Pinkas, and O. Reingold. Distributed pseudo-random functions and KDCs. *LNCS*, 1592:327–337, 1999.

[PS98] P. Pudlák and J. Sgall. Algebraic models of computation and interpolation for algebraic proof systems. In *Proof Complexity and Feasible Arithmetic*, volume 39 of *DIMACS Series in Discrete Mathematics and Theor. Comp. Sci.*, pages 279–296. AMS, 1998.

[Raz90] A. A. Razborov. Applications of matrix methods to the theory of lower bounds in computational complexity. *Combinatorica*, 10(1):81–93, 1990.

[Sha79] A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

[Sim92] G. J. Simmons. An introduction to shared secret and/or shared control and their application. In G. J. Simmons, editor, *Contemporary Cryptology, The Science of Information Integrity*, pages 441–497. IEEE Press, 1992.

[Smo87] R. Smolensky. Algebraic methods in the theory of lower bounds for boolean circuit complexity. *In Proceedings of the 19th STOC*, pages 77–82, 1987.

[Sti92] D. R. Stinson. An explication of secret sharing schemes. *Designs, Codes and Cryptography*, 2:357–390, 1992.