SOME PROBLEMS INVOLVING RAZBOROV-SMOLENSKY POLYNOMIALS

David Mix Barrington

Computer and Information Science Department
University of Massachusetts

COINS Technical Report 90-59

Some Problems Involving Razborov-Smolensky Polynomials

David Mix Barrington
July 18, 1990

1. Abstract

Several recent results in circuit complexity theory have used a representation of boolean functions by polynomials over finite fields. Our current inability to extend these results to superficially similar situations may be related to properties of these polynomials which do not extend to polynomials over general finite rings. Here we pose a number of conjectures on the behavior of such polynomials over rings, and present some partial results toward proving them.

2. Introduction

2.1. Polynomials and Circuit Complexity

The representation of boolean functions as polynomials over the boolean algebra $Z_2 = \{0, 1\}$ dates back to early work in switching theory [Sh38]. A formal language L can be identified with the family of functions $f_i : Z_2^i \rightarrow Z_2$, where $f_i(x_1, \ldots, x_i) = 1$ iff $x_1 \ldots x_1 \in L$, and each of these functions can be written as a polynomial in the variables x_1. We can consider algebraic formulas or circuits with inputs $\{x_1, \ldots, x_n\}$ and the usual complexity measures of size and depth, and get a complexity theory with an algebraic character. Since the binary AND and XOR functions can be simulated in constant size and depth by the binary AND and OR functions and vice versa, this is essentially the usual complexity theory of boolean circuits and formulas.

The algebraic view of circuit complexity has led to some new insights. Skyum and Valiant [SV81] introduced another complexity measure on functions, the degree of the polynomials. Since the variables are idempotent, the degree of a term is the number of distinct variables occurring in it, and of course the degree of a polynomial is the maximum degree of any of its terms. As one example, they defined the class pdC of languages whose polynomial families have both degree and circuit size bounded by polynomials in n, a class now known as $LOGCFL$. Their work emphasized non-uniform complexity classes and very simple reductions between languages.

Razborov [Ra87] used this algebraic view to develop a new lower bound technology for boolean circuits. He noted that the functions in the class $AC^0[2]$ (those functions computable by constant-depth, poly-size, unbounded fan-in circuits of AND, OR, and XOR gates) share a certain algebraic property. Each such function can be approximated by a polynomial of relatively low degree (the function and the polynomial agree expect on a small fraction of the inputs). He showed that the MAJORITY function cannot be so approximated, and thus confirmed the general belief that MAJORITY cannot be computed with such circuits.

Smolensky [Sm87] extended Razborov's method in two ways. First (along with Barrington [Ba86]) he developed an algebraic setting in which MOD-p, for a particular prime p, was a primitive operation instead of MOD-2. Given a field F of characteristic p, he defined the algebra of polynomials over F in variables $\{x_1, \ldots, x_n\}$ satisfying the identities $x_i^2 = x_i$. We will call this the Razborov-Smolensky ring of order n over F, $RS_n(F)$. A function from Z_2^n to F is represented by a unique polynomial in $RS_n(F)$, and (following Razborov) functions computable by constant-depth, poly-size, unbounded fan-in circuits of AND, OR, and MOD-p gates (the class $AC^0[p]$) can be approximated by low-degree polynomials in $RS_n(F)$ whenever F is of characteristic p. Razborov's method can then be extended [Ba86] to show that MAJORITY is not in $AC^0[p]$.

Smolensky's second and more important contribution was to show that the iterated multiplication function of F cannot be so approximated. This not only provides a simpler proof that MAJORITY is not in $AC^0[p]$, but shows that none of the functions MOD-q (with q prime to p) are in that class. (In retrospect, Smolensky's methods provide also provide a reasonably simple way to an earlier seminal result. This is the theorem of Furst-Saxe-Sipser [FSS84] and Ajtai [Aj83] that the functions MOD-q cannot be computed with only AND and OR gates (are not in the class AC^0)).

It is intriguing that these methods provide no way to place bounds on the computing power of such circuits with AND, OR, and MOD-q gates for composite q. Indeed no non-trivial bounds on this class $AC^0[q]$ are known (it might be equal to NP) although no surprising functions are known to be in it. The union of $AC^0[q]$ for all integers q is called ACC^0 [MT89] (also earlier called ACC [BT88]). It is conjectured [Ba89] that MAJORITY is not in ACC^0 and thus that ACC^0 is a strict subset of NC^1 (circuits of depth $O(\log n)$ and fan-in two). Yao [personal communication] has recently shown that functions in ACC^0 can be computed by threshold circuits of depth 3 and quasipolynomial ($2^{(\log n)^{O(1)}}$) size, showing that either ACC^0 is small or such threshold circuits are very powerful.

2.2. The Programs-Over-Monoid Model

The same methods were subsequently used to obtain further lower bounds in circuit complexity, which are most easily stated using the language of programs over finite monoids [BT88]. A monoid is a set with an associative binary operation and an identity. In classical algebraic automata theory, the computation of a finite-state machine is viewed as an iterated multiplication in a particular finite monoid (the transformations of states of the machine, under the operation of composition). Every finite-state computable language has a particular minimal monoid (called the syntactic monoid) which must be contained within any finite-state machine which recognizes the language. A complexity theory of the regular languages has been developed, where one language is viewed as harder than another if its syntactic monoid contains that of the first. ("Contains" has a well-defined meaning — a monoid A contains B (or B *divides* A) if there is a monoid homomorphism from a submonoid of A onto B.)

Furthermore, there is a structure theory of finite monoids due originally to Krohn and Rhodes [KRT68]. Just as every integer is a product of primes, and every finite group (by the Jordan-Hölder Theorem) can be built up from simple groups in a certain way, every finite monoid can be constructed from particular building blocks using particular operations. The building blocks are the simple groups (Z_p for prime p and various non-abelian groups) and one other monoid which is not a group. The operations are division and the *wreath product* — for details see [BT88] and [BST90].

To get natural subsets of the set of finite monoids (and hence natural subclasses of the regular languages), we can restrict which building blocks may be used in the wreath product operation. In this way we can get the classes of solvable groups (only groups of the form Z_p), all groups (only simple groups), aperiodic monoids (only the non-group component), solvable monoids (the non-group component and Z_p's), or all monoids (all components). For each such set of monoids we can look at the class of regular languages whose syntactic monoids lie within that set.

What does this all this have to do with circuit complexity? It turns out that the structure of these classes of monoids and regular languages bears a close relationship to that of various circuit complexity classes [BT88]. Given a monoid M and a set $\{x_1, \ldots, x_n\}$ of boolean inputs, define an *instruction* as the name of an input and a map from Z_2 to M (i.e., two elements of M). The *yield* of an instruction (i, a, b), say, is a if $x_i = 0$ and b if $x_i = 1$. A *program* over M is a sequence of instructions, and its yield is the ordered product of the elements of M yielded by each instruction. A program thus defines a map from Z_2^n into M, or a boolean function on n variables if we divides the elements of M into "accepting" and "rejecting".

Barrington [Ba89] showed that poly-length programs over the non-abelian simple

group A_5 (which correspond to poly-length width-5 branching programs obeying certain restrictions) can be constructed to recognize any language in the circuit complexity class NC^1. NC^1 is easily seen to be powerful enough to simulate polynomial-length programs over any finite monoid, so it is equal to the class of languages recognized by such programs. The circuit complexity classes AC^0, $AC^0[p]$, and ACC^0 discussed above (each of which is a subclass of NC^1), are exactly the classes of languages recognized by poly-length programs over aperiodic monoids, p-monoids, and solvable monoids respectively [BT88]. (A p-monoid is one built up from the non-group component and Z_p in the Krohn-Rhodes framework, where p is prime.)

The connection between circuits and monoids is somewhat understood — the operations of modular and threshold counting occur in the same places in both settings, and the placing of one operation above another in a circuit corresponds to the wreath product operations. The classes of regular languages can also be thought of as very uniform versions of the corresponding circuit complexity classes, especially when both are considered as expressibility classes in the sense of Immerman [BIS90]. Other complexity classes defined by programs have also been shown to be of interest. Szegedy [Sz90] has shown that the languages recognized by programs over abelian monoids are exactly the languages of constant symmetric communication complexity. Bedard, Lemieux, and McKenzie [BLM90] have defined an extension of the model to programs over finite *groupoids* (binary algebras, not necessarily associative) and shown that poly-length programs over groupoids recognize exactly the class $LOGCFL$.

2.3. Polynomials and Programs Over Groups

Barrington, Straubing and Thérien [BST90] have begun the study of the computational power of programs over finite groups. It is known that poly-length programs over any non-solvable group compute exactly NC^1, and that poly-length programs over any solvable group compute only languages in ACC^0. Since we believe that $ACC^0 \neq NC^1$, we believe that programs over solvable monoids are relatively limited in power. The power of solvable monoids can be studied in two stages: the power of solvable groups and the possible interactions between the group and non-group components of a monoid.

In particular, it is conjectured [BST90] that no poly-length program over a solvable group can compute the AND of all n variables. Given our polynomial language, this is fairly easily seen to be true for p-groups. A program over a p-group, of whatever length, can be simulated by a polynomial of constant degree over Z_p. If, for some f of constant degree, $f = a$ iff $x_1x_2 \ldots x_n = 1$, then the polynomial $1 - (f - a)^{p-1}$ is also of constant degree and must be exactly $x_1x_2 \ldots x_n$, which is impossible unless n

is a constant.

Even moving to nilpotent groups (direct products of p_i-groups, in general involving more than one p_i), things are not so easy. In the case of programs over Z_6, for example, we can convert such a program into one polynomial over Z_2 and one over Z_3, each of constant degree. We know about the sets defined by each of these two polynomials, but how do we rule out the possibility that two of these sets intersect in the singleton $(1, 1, \ldots, 1)$? A Ramsey argument [BST90] shows that this is impossible for "sufficiently large", but still constant, n.

Over groups which are solvable but not nilpotent, it is possible for a program of exponential length to compute the AND of all the variables. For some of these groups, methods have been devised to show that exponential length is necessary, but in general this remains open. The known subcases are groups which divide the wreath product of a p-group and an abelian group [BST90], and dihedral groups [Straubing, personal communication]. A plausible but unproven conjecture about the ring $RS_n(Z_p)$ would give extend the exponential lower bound to groups dividing the wreath product of a p-group and a nilpotent group.

These constructions, like Smolensky's, use the ring $RS_n(F)$ for a field F of characteristic p having elements of the desired multiplicative order. Rather than degree, the complexity measure used on polynomials is the number of terms needed to form them by addition, in two different bases of $RS_n(F)$.

The Razborov-Smolensky ring of polynomials can be defined over an arbitrary ring as well as over a field. Many properties, such as the unique representation of any function from Z_2^n to R by a polynomial, still hold, but some do not. One which does not is the one we used above to convert from one way to represent a set by a polynomial to another:

Definition: A polynomial f in $RS_n(R)$ *weakly represents* a set $A \subseteq Z_2^n$ if for some $a \in R$, $f(\mathbf{x}) = a$ iff $\mathbf{x} \in A$. A polynomial f *strongly represents* a set A if $f(\mathbf{x}) = 1$ for $\mathbf{x} \in A$ and $f(\mathbf{x}) = 0$ otherwise.

Fact: If f weakly represents A with value a over a finite field F, then $1-(f-a)^{|F|-1}$ strongly represents A and has degree at most a constant multiple of that of f.□

We know that this exact property can fail for polynomials over a ring which is not a field. Consider $R = Z_6$, and let $f = x_1 + x_2 + \ldots + x_n$. This linear polynomial weakly represents the set $A = \{\mathbf{x} : x_1 + x_2 + \ldots x_n = 0 \bmod 6\}$. But the unique polynomial strongly representing A over Z_6 can be calculated (using methods in [BST90] or [Sm87]) — it is congruent mod 3 to the MOD-2 function and congruent mod 2 to the MOD-3 function. It has degree n, the maximal possible degree.

The inability to do this conversion occurs as a roadblock constantly in the algebraic

approach to circuit lower bounds. In this paper we would like to raise the possibility that some relationship between weak and strong representation might hold for other reasons, and that this might allow these algebraic methods to be extended.

3. The Small Image-Set Conjecture

Let us first consider attempting to represent a small set with a polynomial of as small a degree as possible. For strong representation degree n is needed because the unique polynomial representing that function has degree n (and an exponential number of terms). With degree d we can strongly represent a set of size 2^{-d}, but not one any smaller (we'll prove this later). We can do somewhat better if we are willing to settle for weak representation. For example, in Z_m we can have $m-1$ terms of degree $\lceil n/(m-1) \rceil$, each giving one if its variables are all one and zero otherwise. The sum of these terms gives $m-1$ iff all n varibles are one. Our first conjecture says that this construction is the best possible.

Conjecture 1: (Small Image-Set Conjecture, Weak Form) Any nonempty set weakly represented by a polynomial of degree d over R has size at least $2^{-d(|R|-1)}$.

But actually something stronger seems to be true. If we consider the number of values taken on by the polynomial, we have a continuum between two (strong representation) and $|R|$ (weak representation). Using examples similar to those above and conjecturing that they are best possible, we are led to:

Conjecture 2: (Small Image-Set Conjecture) Any nonempty set weakly represented by a polynomial of degree d which takes on r values has size at least $2^{-d(r-1)}$.

Observations: The SIS Conjecture is trivially true if $r = 1$, $d = 0$, or if $n \leq d(r-1)$.□

To prove the conjecture in some more interesting special cases, let us consider the following construction. Let f be a polynomial of degree at most d on n variables, taking on at most r values. Write f as $g + hx_n$ so that g and h are independent of x_n, and let $g' = g + h$. Note that h is of degree at most $d-1$. The value taken on by f are exactly that of g (if $x_n = 0$) or of g' (if $x_n = 1$), so each of these functions is r-valued. If both take on the same r values, and the SIS Conjecture holds for smaller values of n, then it continues to hold for this n. This is because a set weakly represented by f in this case is the union of two subsets of Z_2^{n-1}, each weakly represented by a polynomial which is degree d and r-valued. But it is entirely possible for g, say, to take on only $r-1$ values, so that some value is taken on only with $x_n = 1$. In this case the size of the subset of Z_2^n on which f takes that value would be half of that

of the subset of Z_2^{n-1} where g' takes it on. In two special cases, however, we can see carry out an inductive proof of the Conjecture.

Fact 1: The SIS Conjecture holds in the case $r = 2$.

Proof: Let f, g, h, g' be as above. The argument above suffices if g and g' each take on both the values taken on by f. But because $r = 2$, if either g or g' fails to do so, it must be a constant. But then the other differs from a constant by h and is thus of degree at most $d - 1$. If we also induct on d, we can assume that a nonempty subset of Z_2^{n-1} weakly represented by g or g' has size at least $2^{-(d-1)}$, so that the same set viewed as a subset of 2^n has size at least 2^{-d}.□

Fact 2: The SIS Conjecture holds in the case $d = 1$.

Proof: Again we use the terminology above. If $d = 1$, then h must be a constant. This means that g and g' take on the same number of values. If this number is r, then each takes on all the values of f and the argument above suffices. Otherwise each is at most $(r-1)$-valued, and weakly represents sets of size only at least $2^{-(r-2)}$ in Z_2^{n-1}, which are of size $2^{-(r-1)}$ in Z_2^n.□

Fact 3: The SIS Conjecture holds when R is a field.

Proof: Here we can show that a set weakly represented by an r-valued f of degree d is also weakly represented by a two-valued polynomial of degree $d(r-1)$. (Over a field, this is the same as being strongly represented by a polynomial of that degree, because we can divide by any constant.) The desired result will then follow from Fact 1 above.

Suppose $f = a$ on the set in question, and on other points takes on some value in $\{b_1, \ldots, b_{r-1}\}$. Consider the product over all i of $f - b_i$. It equals zero off the set in question, and equals a particular nonzero constant on it — the product over all i of $a - b_i$. □

Fact 4: Conjecture 1 (the weaker form of the SIS Conjecture) holds when $R = Z_{p^k}$ for some prime p.

Proof: A trick of Chandra, Stockmeyer, and Vishkin [CSV84], relating the MOD-p^k operation to the MOD-p operation, can be adapted to this setting. The sum modulo p^k of a set of boolean terms is zero iff the modulo p sum of the terms is zero and the modulo p^{k-1} sum of all the products of p-tuples of the terms is also zero. Using this fact, we can take a degree d polynomial f over Z^{p^k} and create a degree $d(p^k - 1)$ polynomial f' over Z_p such that $f = 0$ iff $f' = 0$. □

This means that the full SIS conjecture holds for the case $r = p^k$, but it is not known to hold for general r. In the case $r = 3$, if the range of f is $\{0, a, b\}$, the degree $2d$ polynomial $(f - a)(f - b)$ is 2-valued and tests $f = 0$, unless $ab = 0$. The latter

is possible only if p divides both a and b. But in that case f can be divided by p (a polynomial takes on all its values in a subgroup iff every term of it does), without changing its degree. In fact, for $r \leq 3$ and $r = p^k$, an r valued polynomial of degree d over Z_{p^k} can be simulated by a degree $d(r-1)$ polynomial over Z_p. It is not know whether this holds for all r (it is interesting that so far whether it does appears to be independent of k). Note also that the cases shown here suffice to prove the SIS conjecture for $R = Z_4$.

This makes $R = Z_6$ the smallest ring for which the conjecture is unknown. We don't know whether a quadratic over Z_6 can weakly represent a singleton set with $n = 11$ (it can for $n = 10$). We don't even know whether a three-valued quadratic can weakly represent a singleton with $n = 5$ ($x_1x_2 + x_3x_4$ suffices for $n = 4$).

The quadratic case has an alternate representation as a graph problem. Label both the vertices and edges of an undirected graph with elements of R, and consider the sums of the vertex and edge labels in a clique. What is the largest graph for which every nonempty clique has a nonzero sum? Heath and Pemmaraju [personal communication] have looked at such graphs (for cyclic R) and defined "tight" graphs to be those where (1) every nonempty clique sum is nonzero and (2) the number of distinct clique sums is minimal over all graphs of that size satisfying (1). They conjecture that tight graphs with r distinct nonzero clique sums have $2r$ vertices. Also, they conjecture that the r values of a tight graph must form an arithmetic progression. These conjectures are independent of the modulus, suggesting that they may have a combinatorial rather than an algebraic proof.

4. The Intersection Conjecture

One way to examine the sets of input values which can be defined by specific types of polynomials is to see how they intersect. For example, suppose that we have k equations $f_i = 0$, each of degree d. The SIS conjecture would give us limits on how small the locus of each equations could be, but would place no limits on the size of the intersection. Over a field, we could convert each of the equations into a two-valued one and then multiply them together, giving a single equation $F = 0$ which is true iff $f_i = 0$ for all i. To what extent is this a general phenomenon for rings?

We know that if R is not a field, then two sets which are each strongly represented by a polynomial of low degree can intersect in a set whose strong representation has high degree. Let $R = Z_6$ and consider two sets of inputs: those summing to zero mod 2 and those summing to zero mod 3. The strong representations of these sets over Z_6 are linear and quadratic, respectively, but that of their intersection is degree n. In

this case the intersection has a weak representation of low degree. Does this always happen?

Conjecture 3: (Intersection Conjecture, Strong Form) If $\{f_i\}$ is a set of k polynomials of degree d over a finite ring R, then there is a polynomial F of degree kd such that $F = 0$ iff $f_i = 0$ for all i.

Even if this is false, the following weaker statement might still be true:

Conjecture 4: (Intersection Conjecture, Weak Form) If $\{f_i\}$ are as above, and the set S of inputs for which $f_i = 0$ for all i is nonempty, then there is a polynomial F of degree kd such that the set for which $F = 0$ is nonempty and no bigger than S.

Conjecture 3 is easily seen to be true for fields. Over rings of the form Z_{p^k}, any set weakly representable in degree k is strongly representable in degree $d(p^k - 1)$ over Z_p, so the intersection is strongly representable in degree $kd(p^k - 1)$ over Z_p. It must then be of size at least $2^{-kd(p^k-1)}$if it is nonempty, and sets this small are weakly representable in degree kd over Z_{p^k}. So Conjecture 4, at least, is true for these rings as well.

Curiously enough, an *ad hoc* argument gives something very close to the stronger result for a large class of rings:

Proposition: Conjecture 3 holds if k is a power of two and R is a direct product of finite fields (this includes, for example, the case $R = Z_m$ where m is square-free).

Proof: Let R be the direct product of some finite fields $F_1, \ldots, F_r$. For each F_i, choose a quadratic polynomial $a_i z^2 + b_i z + c_i$ which is irreducible over F_i. Let $AP^2 + BPQ + CQ^2$ be the unique polynomial over R whose F_i component, for each i, is $a_i P^2 + b_i PQ + c_i Q^2$. We claim that if P and Q are elements of R, then $AP^2 + BPQ + CQ^2 = 0$ iff $P = 0$ and $Q = 0$. (Clearly it is zero if both P and Q are. If even one of P and Q, say P, is zero, it reduces to CQ^2 which is nonzero for nonzero Q. If both are nonzero, then $P = zQ$ for some $z \in R$, and the polynomial is equal to $(Az^2 + Bz + C)Q^2$. Q^2 must be nonzero in some component because Q is, and the term in parentheses is nonzero for every z in every component.)

Applying this identity recursively, we can convert k equations of degree d (where k is a power of two) to one equation of degree kd, such that the single equation is zero iff all the original ones are. This does not get us an equation of degree $3d$, for example, out of three equations of degree d, though this might be possible. It's also not clear that similar identities should hold in other rings — there do not appear to be any in Z_4.□

The intuition here is that a polynomial of a certain degree has only so much power to isolate a small fraction of the input space. Our conjectures say in effect that the

polynomial can do no better at isolating a small fraction of a subset of that space defined by another polynomial.

5. Making Change in an Abelian Group

A family of k equations over a finite ring R can equally well be thought of as a single equation over the direct product R^k. If the original equations are all of degree d, then so is the composite equation. This opens the possibility of applying the SIS Conjecture to the study of systems of equations. For example, we know that for $d = 1$ and for any R, a family of k linear equations over R gives a single equation over R^k. A set weakly represented by this equation must be of size at most $2^{-(|R|^k-1)}$, because the SIS Conjecture holds with $d = 1$. But if our Independence Conjectures are correct, we can do better. The set of inputs satisfying each component has size at least $2^{-(|R|-1)}$, so the intersection of these sets would be size $2^{-k(|R|-1)}$. Perhaps a stronger form of the SIS Conjecture holds for many-fold direct products? The following problem is an attempt to investigate this. We restrict ourselves to linear polynomials, which means that we can ignore the multiplicative structure of R and consider it as an abelian group.

Definitions: The *boolean span* of a multiset of elements of R is the set of all sums of submultisets of it. The *proper span* of a multiset of nonzero elements is the set of all sums of nonempty submultisets.

Problem: What conditions on a multiset from R force its boolean span to be all of R?

We call this problem "change-making" because it is a variation of a familiar problem: How many coins of a specific type are needed to guarantee having change available for any amount? Rather than have the exact number of pence (in Britain) available for our purchase, we might be satisfied to have the correct number modulo 100, so we would be able to transfer an integral number of pounds. Thus we might want the boolean span of our coins in Z_{100} to be all of Z_{100}.

Lemma: Any multiset of $p - 1$ nonzero elements from Z_p has boolean span Z_p.

Proof: By induction on k, we show that any set of k coins spans a set of size at least $k + 1$ in Z_p, if $k < p$. The $k = 0$ case is trivial. Given k coins spanning a set S and a new coin a, we can make all of S and also the set $S + a = \{s + a : s \in S\}$. $S + a$ has the same size as S and must be different from S (S could be invariant under translation by a nonzero a only if it were a subgroup). So the new boolean span is strictly larger than S, and has size at least $k + 2$ if S has size at least $k + 1$.□

Applying this same argument to other cyclic rings, we get a more complicated result, because these rings might have nontrivial subgroups:

Lemma: Any multiset of $m-1$ elements of Z_m has a nonempty submultiset whose boolean span is a subgroup of Z_m.

Proof: A zero element spans the trivial subgroup. The process from the Lemma above must terminate with S a proper subgroup or continue until the span is all of Z_m.□

For rings which are not cyclic, something like this result appears to hold. The key parameter of the ring is not the order but, roughly, the sum of the orders of the ring's components.

Definition: Let R be a finite abelian group. Let e be the maximum order of any element of R. The reader may verify that R can be written as a direct product $Z_e \times R'$. Define a function on abelian groups by $f(Z_1) = 0$ and $f(R) = (e-1)f(R')$. We will call this function the *capacity* of the group.

Fact: For any R, there is a multiset of $f(R)$ elements such that (1) its boolean span is R, (2) its proper span is $R \setminus \{0\}$.

Proof: To get (1) and (2), write R as $Z_e \times R'$ and let the multiset be $e-1$ copies of $(1, 0_{R'})$ together with $(0, x)$ for each x in the set for R'. □

This means that a linear polynomial over R can weakly represent a set of size $2^{-f(R)}$. We believe that no polynomial can do better.

Conjecture 5: No multiset of more than $f(R)$ elements from R satisfies (1) and (2) above. That is, no linear polynomial over R can weakly represent the function strongly represented by $x_1 x_2 \ldots x_{f(R)}$.

A natural way to try to prove this would be to use induction on subgroups. If we could show that every multiset of $f(R)$ elements has a subset spanning a subgroup H, that subset would be of size $f(H)$ by induction, and the remaining elements could be reinterpreted as elements of R/H. We are guaranteed that $f(R) \leq f(H) + f(R/H)$, so this would suffice. Unfortunately, consider the multiset from $Z_2 \times Z_4$, listed as $\{(0,1),(0,1),(1,1),(1,1)\}$. No subset spans a subgroup (the sum $(1,0)$ cannot be formed). Similar examples exist in many other rings. This revised version, however, is true in so far as we have been able to check:

Conjecture 6: (The Change Conjecture) Any multiset of $f(R)+1$ elements from R has a submultiset whose boolean span is a subgroup. Any multiset of $f(R)$ elements either has such a submultiset or has the following property: any element which is the sum of a submultiset is also the sum of another distinct submultiset.

From this rather convoluted fact we can prove the following fact about weak representation, by the methods we used for the SIS Conjecture in the linear case:

Corollary: (assuming Conjecture 6) Any nonempty set weakly represented by a linear polynomial over R has size at least $2^{-f(R)}$.

Proof: (sketch) Consider adding the linear terms in one by one. The boolean span of the multiset of coefficients seen so far increases as we go. When a variable is added and it does not change, the minimal size of a represented set also does not change. When it does change, this size can decrease by half. We are guaranteed at most $f(R)+1$ changes. If there are exactly $f(R)+1$, we know that on occasion when a variable was added, the last uniquely-formed sum disappeared. The minimal set size could not have decreased on that occasion. The size could thus have been halved only $f(R)$ times.□

These last conjectures imply the weaker form of the Independence Conjecture with $d = 1$, because k equations over R are equivalent to one equation over R^k, and $f(R^k) = kf(R) \leq f(|R| - 1)$.

6. Consequences

Our goal in this study is primarily to extend our understanding of the computational power of Razborov-Smolensky polynomials over general finite rings. We believe that resolving the conjectures here will require techniques which should be of general utility in the study of lower bounds for the related computational models of programs over groups and boolean circuits. But there are some direct consequences which would ensue if certain of our conjectures were to be proven.

A proof of even the weak form of the SIS Conjecture would show that programs over nilpotent groups cannot calculate functions whose polynomials have more than constant degree. In particular they cannot compute the AND function on more than a constant number of variables. Both of these facts are currently known, but the proof of them [BST90] involves a Ramsey argument. Because of this, the "constant" bound on the degree for a given nilpotent group is very large. The new proof would reduce this degree bound to the product of the order of the group and its "nilpotency class" (see, e.g., [BST90]).

It is possible that similar techniques would allow the extension of the exponential lower bound for program length to groups which are the wreath products of nilpotent groups with abelian groups. This is because any such group divides a direct product of groups of the kind we can currently handle. But since other parmeters than degree

are used in the existing lower bound proofs, the new results would have to deal with them.

Our conjectures, even to the extent that they are already proven, have minor consequences in circuit complexity. Consider a circuit of MOD-6 gates computing the AND function (a MOD-6 gate, which has unbounded fan-in, outputs one iff the MOD-6 sum of its inputs is nonzero, and otherwise outputs zero). It is conjectured that such a circuit would require an exponential number of gates (this is implied by the conjecture that programs over solvable groups require exponential length to compute AND). Our results give us a small but nontrivial lower bound on the number of MOD-6 gates on the level of the circuit closest to the inputs, as follows. We would be interested to know if this result can be duplicated by other methods.

If k is the number of such gates, then the row computes a linear function from Z_2^n to Z_6^k (before converting the output of this function to an element of Z_2^k. By the $d = 1$ case of the SIS Conjecture, which we have proven, a nonempty set weakly represented by this linear function has size at least $2^{-(6^k-1)}$. For the circuit to compute the AND function, the first row must behave differently on the input $(1, 1, \ldots, 1)$, so that the polynomial for the first row must weakly represent a singleton set and thus k must be $\Omega(\log n)$. Furthermore, we know that if k is a power of two, a linear polynomial over Z_6^k is equivalent to a polynomial of degree k over Z_6. Thus, the (unproven) SIS conjecture for $R = Z_6$ would tell us that $k = \Omega(n)$. The Change Conjecture for $R = Z_6^k$ would also give this linear lower bound, by the argument presented in the last section.

7. Acknowledgements

Much of this paper represents joint work with Denis Thérien, who suggested a number of these problems. I would like to thank him, the other participants in the 1989 and 1990 McGill Invitational Workshops, my colleagues in the COINS Theory Group, Lenny Heath, and Sriram Pemmaraju for various helpful discussions.

8. References

[Aj83] M. Ajtai, "Σ_1^1 formulae on finite structures", *Annals of Pure and Applied Logic* **24** (1983), 1-48.

[Ba86] D. A. Barrington, "A note on a theorem of Razborov", COINS Technical Report 87-93 (July 1986), University of Massachusetts.

[Ba89] D. A. Barrington, "Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1", *J. Comp. Syst. Sci.* **38:1** (Feb. 1989), 150-164.

[BIS90] D. A. M. Barrington, N. Immerman, and H. Straubing, "On uniformity within NC^1," *J. Comp. Syst. Sci.*, in press. Preliminary version *Structure in Complexity Theory: Third Annual Conference* (Washington: IEEE Computer Society Press, 1988), 47-59.

[BST90] D. A. M. Barrington, H. Straubing, and D. Thérien, "Non-uniform automata over groups", *Information and Computation*, to appear. Preliminary version by Barrington and Thérien, *Proc. 14th ICALP*, (Berlin: Springer Verlag, 1987), 163-173.

[BT88] D. A. M. Barrington and D. Thérien, "Finite monoids and the fine structure of NC^1", *J. ACM* **35:4** (Oct. 1988), 941-952.

[BLM90] F. Bédard, F. Lemieux, and P. McKenzie, "Extensions to Barrington's M-program model", *Structure in Complexity Theory: Fifth Annual Conference*, to appear.

[CSV84] A. K. Chandra, L. J. Stockmeyer, and U. Vishkin, "Constant depth reducibility," *SIAM J. Comp.* **13:2** (1984), 423-439.

[FSS84] M. Furst, J. B. Saxe, and M. Sipser, "Parity, circuits, and the polynomial-time hierarchy", *Math. Syst. Theory* **17** (1984), 13-27.

[KRT68] K. B. Krohn, J. Rhodes, and B. Tilson, in M. A. Arbib, ed., *The Algebraic Theory of Machines, Languages, and Semigroups* (New York: Academic Press, 1968).

[MT89] P. McKenzie and D. Theérien, "Automata theory meets circuit complexity," *Proc. 16th ICALP, Springer Lecture Notes in Computer Science* **372** (1989), 589-602.

[Ra87] A. A. Razborov, "Lower bounds for the size of circuits of bounded depth with basis $\{\&, \oplus\}$", *Mathematicheskie Zametki* **41:4** (April 1987), 598-607 (in Russian). English translation Math. Notes Acad. Sci. USSR **41:4** (Sept. 1987), 333-338.

[Sh38] C. E. Shannon, "A symbolic analysis of relay and switching circuits", *Trans. AIEE* **57** (1938), 713-723.

[SV81] S. Skyum and L. G. Valiant, "A complexity theory based on Boolean algebra", *Proc. 22nd IEEE FOCS* (1981), 244-253.

[Sm87] R. Smolensky, "Algebraic methods in the theory of lower bounds for Boolean circuit complexity", *19th ACM STOC Symp.* (1987), 77-82.

[Sz90] M. Szegedy, "Functions with bounded symmetric communication complexity and circuits with mod m gates", *Proc. 22nd ACM STOC* (1990), 278-286.