Small-Bias Spaces for Group Products

Raghu Meka* and David Zuckerman*
Department of Computer Science, University of Texas at Austin
{raghu, diz}@cs.utexas.edu

Abstract

Small-bias, or ϵ-biased, spaces have found many applications in complexity theory, coding theory, and derandomization. We generalize the notion of small-bias spaces to the setting of group products. Besides being natural, our extension captures some of the difficulties in constructing pseudorandom generators for constant-width branching programs - a longstanding open problem. We provide an efficient deterministic construction of small-bias spaces for solvable groups. Our construction exploits the fact that solvable groups have nontrivial normal subgroups that are abelian and builds on the construction of Azar et al. [AMN98] for abelian groups. For arbitrary finite groups, we give an efficient deterministic construction achieving constant bias. We also construct pseudorandom generators fooling linear functions mod p for primes p.

1 Introduction

In this work we generalize the notion of small-bias spaces to the setting of group products. Small-bias, or ϵ-biased, spaces over $\mathbb{Z}_2$ have been very useful in constructions of various pseudorandom objects. In particular, they are used in the construction of almost k-wise independent spaces ([NN93]), which in turn have many applications such as universal sets ([LY94], [BEG$^+$94]). An application of interest to us is that ϵ-biased spaces fool branching programs of width two. Can we generalize this observation to fool constant-width branching programs? Our extension of small-bias spaces to finite groups besides being interesting on its own, could be useful for constructing pseudorandom generators for small width branching programs. We address the problem of explicitly constructing such small-bias spaces over finite groups, and give an efficient deterministic construction for solvable groups and a partial solution to the problem for arbitrary finite groups.

Constructing pseudorandom generators for constant-width branching programs is a fundamental problem with many applications in circuit lower bounds and derandomization. The problem is largely open even for strongly restricted classes such as width three *read-once permutation branching programs* (ROPBPs) - branching programs where no variable is read more than once and the edges between any two layers with the same label define a permutation.

Our notion of small-bias spaces is motivated by the following group-theoretic formulation of the problem of constructing pseudorandom generators for constant-width ROPBPs. Consider the edges between layers i and $i+1$ of a width-w ROPBP. By relabeling the nodes if necessary, we may assume that the permutation corresponding to the edges labeled 0 is the identity permutation. Then the permutation corresponding to the label 1 is some permutation $g_i \in S_w$, where S_w denotes

*Supported in part by NSF Grant CCF-0634811 and THECB ARP Grant 003658-0113-2007.

the permutation group on w elements. Under this description, if the variable read by the branching program at layer i is x_i, the computation performed by the ROPBP can be written as follows: on input $(b_1, \ldots, b_n) \in \{0,1\}^n$, output $g_1^{b_1} g_2^{b_2} \ldots g_n^{b_n} \in \mathbb{S}_w$.

Thus, pseudorandom generators for width w ROPBPs are equivalent to functions $P : \{0,1\}^r \to \{0,1\}^n$ that fool *products of group elements* in the sense that, for all $g_1, \ldots, g_n \in \mathbb{S}_w$, the distributions of $g_1^{b_1} g_2^{b_2} \ldots g_n^{b_n}$ with $b \in_u \{0,1\}^n$ and $g_1^{P(y)_1} g_2^{P(y)_2} \ldots g_n^{P(y)_n}$ with $y \in_u \{0,1\}^r$ are ϵ-close in variation distance (for a multi-set S, $x \in_u S$ denotes a uniformly sampled element of S.).

In this work we consider a dual of the above problem. A PRG for ROPBPs outputs Boolean exponents that fool products of arbitrary group elements raised to these exponents. Our notion of ϵ-biased space outputs group elements that fool products of these elements raised to arbitrary exponents. For convenience we use the max norm instead of variation distance.

Definition 1.1 *A multi-set $S \subseteq G^n$ is an ϵ-biased space for products over G* [1] *if for all $b_1, \ldots, b_n \in \{0,1\}$ not all zero, and every $h \in G$*

$$\left| \mathsf{Pr}_{\boldsymbol{g} \in_u S} [\, g_1^{b_1} g_2^{b_2} \ldots g_n^{b_n} = h \,] - \frac{1}{|G|} \right| \leq \epsilon.$$

Remark. The definition can naturally be extended to non-binary exponents $b_1, \ldots, b_n \in [|G|]$ and arbitrary permutations $\pi : [n] \to [n]$, where we look at products of the form $g_{\pi(1)}^{b_1} g_{\pi(2)}^{b_2} \cdots g_{\pi(n)}^{b_n}$. In this abstract we only consider the definition above for simplicity. Our results extend straightforwardly to arbitrary permutations π as well as for non-binary powers $b_i \in [|G|]$, provided $\gcd(b_1, \ldots, b_n, |G|) = 1$.

When the group G is the additive group $\mathbb{Z}_2$, the definition above coincides with the usual notion of small-bias spaces over $\mathbb{Z}_2$ of Naor and Naor [NN93]. Besides being a natural generalization of ϵ-biased spaces over $\mathbb{Z}_2$, the definition above captures some of the difficulties involved in constructing PRGs for constant-width ROPBPs, as PRGs for constant-width ROPBP imply ϵ-biased spaces for finite groups.

Theorem 1.1 *Given a PRG $G : \{0,1\}^r \to \{0,1\}^n$ for width w ROPBP with error at most ϵ and running time $t(n, \epsilon)$, for every group $H \subseteq \mathbb{S}_w$ there exists an algorithm with running time $O(t(n,\epsilon)2^r)$ that outputs a 2ϵ-biased set over H of size poly$(2^r, n, 1/\epsilon)$.*

We defer the proof to Section 6.

Remark. Azar et al. [AMN98] characterize ϵ-biased spaces for abelian groups in terms of the characters of the group. One could generalize this definition to non-abelian groups using the irreducible characters or irreducible representations of the group. However, there does not seem to be any connection between such objects and pseudorandom generators for constant-width read-once branching programs, our original motivation. As far as we know, a notion of small-bias spaces for finite groups in terms of irreducible representations is incomparable to our notion of small-bias

[1] By convention, ϵ-biased spaces with no explicit mention of a group will correspond to ϵ-biased spaces over $\mathbb{Z}_2$. For brevity, we will refer to *small-bias spaces for products over G* simply as *small-bias spaces over G*.

spaces and to pseudorandom generators for constant-width ROPBPs.

By the probabilistic method it can be shown that for any group G, ϵ-biased sets of size $O(|G|n/\epsilon^2)$ exist. The constructions of Naor and Naor [NN93], Alon et al. [AGHP92], Azar et al. [AMN98] give explicit polynomial size ϵ-biased spaces for abelian groups. However, the problem seems to be considerably harder for non-abelian groups and the techniques of [NN93, AGHP92, AMN98] fail when the group is non-abelian. We prove the following results for general finite groups.

Theorem 1.2 *Let G be a finite group. There exists a deterministic algorithm running in time $poly(n)$ that takes as input n and outputs a set S of size $poly(n)$ that is α-biased over G, where $\alpha < 1/|G|$ is a fixed constant depending on $|G|$.*

Theorem 1.3 *Let G be a finite solvable group. There exists a deterministic algorithm running in time $poly(n, 1/\epsilon)$ that takes as input n, ϵ and outputs a set S of size $poly(n, 1/\epsilon)$ that is ϵ-biased over G.*

Our constructions are based on the ϵ-biased spaces for abelian groups of Azar et al. The construction for solvable groups is recursive and uses the fact that every solvable group has a nontrivial normal subgroup that is abelian.

It is instructive to examine the dual objects - PRGs and ϵ-biases spaces - for simpler families of constant-width ROPBPs such as the class of linear functions modulo a prime p. For this case, our notion of ϵ-biased spaces with group $G = \mathbb{Z}_p$ corresponds to the usual notion of ϵ-biased spaces over $\mathbb{Z}_p$ except that our notion assumes the linear functions have $\{0, 1\}$ coefficients. As far as we know there were no previous efficient constructions of PRGs for the ROPBPs corresponding to the family of linear functions modulo a prime p.

Definition 1.2 *A function $G : \{0,1\}^r \to \{0,1\}^n$ is said to be an ϵ-pseudorandom bit generator (ϵ-PBG) for sums mod p, if for every $v \in \mathbb{F}_p^n$, $v \neq 0$ and all $a \in \mathbb{F}_p$*

$$|\mathsf{Pr}_{z \in_u \{0,1\}^r} [\, \langle v, G(z) \rangle = a \,] - \mathsf{Pr}_{x \in_u \{0,1\}^n} [\, \langle v, x \rangle = a \,]| \leq \epsilon, \tag{1}$$

where the inner product is taken over $\mathbb{F}_p$.

Note that the existence of an efficient ϵ-PBG $G : \{0,1\}^r \to \{0,1\}^n$ with $r = O(\log n + \log(1/\epsilon))$ does not follow from the known constructions of ϵ-biased spaces for $\mathbb{Z}_p$. (The main difference is that ϵ-biased spaces mod p, by definition, take values in $\mathbb{Z}_p$.) We present a construction of ϵ-PBG with $r = O(\log n + \log(1/\epsilon))$ based on pseudorandom generators for low-degree polynomials obtained in [Vio08, Lov08, BV07]. Recently, independent of our work, Lovett et al. [LRTV09] constructed ϵ-PBG with better dependence on the field size p which also works for composite moduli. However, for our intended application, the field size is always a constant and the construction below is optimal up to constant factors.

Theorem 1.4 *For all $\epsilon > 0$ and primes p, there exists an efficient ϵ-PBG for $\mathbb{F}_p$, $G : \{0,1\}^r \to \{0,1\}^n$, with seed length $r = O(\log n + \log(1/\epsilon))$.*

Observe that a pseudorandom generator for width 3 read-once branching programs gives both an ϵ-biased set over $\mathbb{F}_2$ and an ϵ-PBG for $\mathbb{F}_3$. Motivated in part by our construction of ϵ-biased space for the permutation group on three elements - $\mathbb{S}_3$, a solvable group - we conjecture that a weak converse of the above statement holds.

Conjecture 1.5 *Let $G_1 : \{0,1\}^r \to S$ generate uniform samples from a ϵ-biased set $S \subseteq \{0,1\}^n$. Let $G_2 : \{0,1\}^r \to \{0,1\}^n$ be a ϵ-PBG for $\mathbb{F}_3$. Then, the sum $G_1 \oplus G_2 : \{0,1\}^r \times \{0,1\}^r \to \{0,1\}^n$ defined by $(G_1 \oplus G_2)(z_1, z_2) = G_1(z_1) \oplus G_2(z_2)$ is pseudorandom with respect to width 3 read-once permutation branching programs.*

We also provide an example showing that the sum of two constant-bias spaces over $\mathbb{Z}_2$ does not fool linear functions mod 3; in particular, the sum of two constant-bias spaces does not fool width 3 ROPBPs. Reingold and Vadhan [RV06] had asked whether the sum of two $n^{-O(1)}$-biased spaces fools logspace. Although we do not resolve the question, we remark that previously there was no known example ruling out the possibility that the sum of two constant-biased spaces gives a *hitting set* for logspace computations.

Theorem 1.6 *There exists an absolute constant α, $0 < \alpha < 1/2$, such that for all $n > 0$, there exists an α-biased space $S \subseteq \{0,1\}^n$, such that the dimension of the span of $S \oplus S = \{x \oplus y : x, y \in S\}$ viewed as a subset of the vector space $\mathbb{F}_3^n$ is $o(n)$. In particular, there exists a linear function $f \bmod 3$ such that f is constant on $S \oplus S$.*

2 Previous Work and Preliminaries

We first present the notions of pseudorandom generators for small width branching programs and small-bias spaces over abelian groups.

Definition 2.1 *A branching program (BP) of width w and length t ((w,t)-BP) is a rooted, layered directed acyclic graph with $t+1$ layers and at most w nodes at each layer. The nodes (internal nodes) at a layer $j \leq t$ in the graph are labeled with a variable x_i and have two outgoing edges each, corresponding to the two possible values of x_i. The nodes at the last layer (leaf nodes) are labeled either 0 or 1. An instance $x \in \{0,1\}^n$ defines a unique directed path through the branching program starting at the root and following the outgoing arc from internal nodes labeled by the value of the variable at that node. The output of the branching program is the label of the leaf node reached by this path.*

A branching program is read-once (ROBP) if no variable occurs more than once on any path from the root to a leaf. A branching program is a permutation branching program (PBP) if any two edges pointing to the same node have different labels.

Definition 2.2 *A pseudorandom generator (PRG) for width w BPs with error ϵ is a function $G : \{0,1\}^r \to \{0,1\}^n$ such that for every (w,t)-BP A with $t = poly(n)$, $A(U_n)$ is ϵ-close to $A(G(U_r))$, where U_k denotes the uniform distribution on $\{0,1\}^k$. Pseudorandom generators for ROBP and read-once permutation branching programs (ROPBP) are defined similarly.*

Constructing pseudorandom generators for constant-width branching programs with seed length $r = O(\log n + \log(1/\epsilon))$ is a fundamental open problem. It is known that ϵ-biased spaces over $\mathbb{F}_2$ fool width two branching programs ([SZ], [BDVY08]). Generalizing this observation, Bogdanov et al. [BDVY08] show that PRGs for degree k polynomials over GF(2) fool width two branching programs that are allowed to read up to k bits at each internal node.

However, for width more than two little is known. In fact, by Barrington's theorem ([Bar86]) constructing pseudorandom generators for width five branching programs would imply lower bounds

for the circuit class NC^1 - a longstanding open problem in complexity theory. In view of the above, focus on the problem has been restricted to the class of read-once branching programs. Most known PRGs for ROBPs are based on their relation to space-bounded computations; nonuniform logspace computations in particular are equivalent to polynomial-width ROBPs. Even for width three ROBPs, the best generators are the much more powerful generators for logspace machines of Nisan [Nis92] and Impagliazzo et al. [INW94] that achieve a seed-length of $O(\log^2 n)$.

Constructing pseudorandom generators for logspace-computations with logarithmic seed-length is an outstanding open problem with progress being relatively slow. The main nontrivial results are those of [AKS87], [NZ96], [RR99], [Rei08]. In particular, Nisan and Zuckerman [NZ96] give a generator with seed-length $O(\log n)$ for logspace machines that use polylogarithmic randomness and Reingold [Rei08] gives a logspace algorithm for undirected st-connectivity.

The notion of ϵ-biased spaces over $\mathbb{Z}_2$ was introduced by Naor and Naor [NN93] who also gave efficient constructions of such spaces of size $poly(n, 1/\epsilon)$. Subsequently, Alon et al. [AGHP92] and Azar et al. [AMN98] obtained efficient constructions that work for arbitrary abelian groups.

In our construction of ϵ-biased spaces over solvable groups we make use of the fact that for abelian groups we can construct polynomial size sets that are *strongly* ϵ-biased in the following sense.

Definition 2.3 *Let N be an abelian group. A set $S \subseteq N^n$ is strongly ϵ-biased in N if, for all non-empty $I = \{i_1, \ldots, i_k\} \subseteq [n]$, automorphisms $\Phi_{i_1}, \ldots, \Phi_{i_k} : N \to N$, and $h \in N$,*

$$\left| \mathsf{Pr}_{\boldsymbol{g} \in_u S} [\, \Phi_{i_1}(g_{i_1}) \Phi_{i_2}(g_{i_2}) \cdots \Phi_{i_k}(g_{i_k}) = h \,] - \frac{1}{|N|} \right| \leq \epsilon.$$

To get intuition for the above definition, consider the case when N is the cyclic group $\{1, \omega, \omega^2, \ldots, \omega^{p-1}\}$ with ω a p'th root of unity for p prime. Then, the automorphisms of N are functions of the form $\Phi_a : N \to N$ defined by $\Phi_a(\omega^x) = \omega^{ax \bmod p}$, for $a \not\equiv 0 \bmod p$. Thus strongly ϵ-biased spaces for N in this case correspond to pseudorandom sets for linear functions mod p. The explicit constructions of ϵ-biased spaces of Azar et al. are in fact strongly ϵ-biased.

Theorem 2.1 ([AMN98]) *For every $d > 0$, there exists a deterministic algorithm running in time $poly(n, 1/\epsilon)$ that takes as input n, ϵ and outputs a set S of size $poly(n, 1/\epsilon)$ that is strongly ϵ-biased in $\mathbb{Z}_d$.*

Proof: Follows from the fact that small-bias spaces of Azar et al. fool the irreducible characters of $\mathbb{Z}_d$.□

As a corollary we obtain strongly ϵ-biased sets for all abelian groups.

Corollary 2.2 *For an abelian group N, there exists a deterministic algorithm running in time $poly(n, 1/\epsilon)$ that takes as input n, ϵ and outputs a set S of size $poly(n, 1/\epsilon)$ that is strongly ϵ-biased in N.*

Proof: Follows from Theorem 2.1 and the fact that abelian groups are isomorphic to direct products of cyclic groups.□

3 Constant-Bias Spaces for Arbitrary Groups

We now give a construction that achieves constant bias and works for arbitrary finite groups. We use the efficient constructions of small-bias spaces for $\mathbb{Z}_{|G|}$ given by Azar et al.

Proof of Theorem 1.2: Let $S \subseteq [|G|]^n$ be a $1/(2|G|)$-biased space in $\mathbb{Z}_{|G|}$ of size $poly(n)$ as given by setting $\epsilon = 1/(2|G|)$ in Theorem 2.1. Consider the set

$$T = \{(g^{x_1}, g^{x_2}, \ldots, g^{x_n}) : g \in G, (x_1, \ldots, x_n) \in S\}. \tag{2}$$

We claim that the set T is α-biased for a constant $\alpha < 1/|G|$. We will use the following lemma.

Lemma 3.1 *For any l with $\gcd(l, |G|) = 1$, the random variable $X = g^l$, where g is uniform in G, is uniformly distributed in G.*

Proof of lemma: The lemma follows from the fact that the map $\phi : G \to G$, $\phi(x) = x^l$ is bijective. For, if $g_1^l = g_2^l$, then for a, b such that $al + b|G| = 1$, $g_1 = g_1^{al+b|G|} = g_2^{al+b|G|} = g_2$.□

Fix a sequence $b_1, \ldots, b_n \in \{0, 1\}$ and let $I = \{i_1, \ldots, i_k\} = \{i : b_i \neq 0\} \subseteq [n]$. Let $Y(g, x) = g^{x_{i_1} + \ldots + x_{i_k}}$. Note that for a fixed $x = (x_1, \ldots, x_n)$, if $\gcd(x_{i_1} + \ldots + x_{i_k}, |G|) = 1$, then by Lemma 3.1 $Y(U_G, x)$ is uniformly distributed in G, where U_G is the uniform distribution over G. Further, since $x \in_u S$ and S is $1/2|G|$-biased, $Pr[\gcd(x_{i_1} + \ldots + x_{i_k}, |G|) = 1] \geq \phi(|G|)(1/|G| - 1/2|G|)$, where ϕ is the Euler function. Thus, for a fixed $h \in G$,

$$Pr_{g \in_u G, x \in_u S}[Y(g, x) = h] \geq \frac{\beta}{|G|},$$

where $\beta = \phi(|G|)/2|G|$. Therefore T is α-biased in G, where $\alpha = (1 - \beta)/|G| < 1/|G|$ and $|T| = |G||S| = poly(n)$.□

For abelian groups G, given a set that achieves constant bias, we can combine several independent copies of the constant bias space to obtain a space with smaller bias. The construction of ϵ-biased spaces in [NN93], at a high level, takes this approach. However, for non-abelian groups it is not clear how to perform such *amplification*. In particular, we ask the following question:

Question 3.2 *Let T be α-biased over G. Define*

$$T^k = \{(g_{11}g_{21}\cdots g_{k1},\ g_{12}g_{22}\cdots g_{k2},\ \ldots,\ g_{1n}g_{2n}\cdots g_{kn}) : g_i = (g_{i1}, g_{i2}, \ldots, g_{in}) \in T, 1 \leq i \leq k\}.$$

Then, is T^k ϵ-biased over G for $\epsilon = \alpha^{\Omega(k)}$?

For abelian groups the answer to the above question is yes, but the answer is not clear for non-abelian groups. An answer to the question even for the specific constant-bias space of equation (2) would be very interesting.

4 Small-Bias Spaces for Solvable Groups

We now address the case of solvable groups and prove Theorem 1.3. Our construction of ϵ-biased spaces is recursive, by using the fact that every solvable group G has a nontrivial abelian subgroup, say N, that is normal in the group G. We use the known constructions of ϵ-biased spaces for the abelian group N and combine them with an ϵ-biased space for the factor group G/N which can be obtained recursively, since G/N is also solvable. We first present some preliminaries from group theory.

Let N be a nontrivial normal subgroup of G and let $H = G/N$ be the factor group of N in G. Without loss of generality assume that the factor group H is given by elements of G which are in distinct cosets of N in G. Note that the representatives of H may not form a subgroup in G. In case of ambiguity in group operations we will denote multiplication in H by $\circ$.

Lemma 4.1 *Every $g \in G$ can be written uniquely as $g = nh$, where $n \in N$ and $h \in H$.*

The following lemma gives us a way to separate a *mixed* product $n_1h_1n_2h_2 \dots n_kh_k$ with $n_i \in N$, $h_i \in H$ into products of elements in N and H respectively.

Lemma 4.2 *Let $g_1 = n_1h_1, g_2 = n_2h_2, \dots, g_k = n_kh_k$, with $n_i \in N$ and $h_i \in H$. Let $\boldsymbol{h} = (h_1, \dots, h_k)$. Then,*

$$g_1g_2 \dots g_k = (n_1h_1)(n_2h_2) \dots (n_kh_k) = n_1\, \Phi_{1,\boldsymbol{h}}(n_2)\, \Phi_{2,\boldsymbol{h}}(n_3)\, \dots\, \Phi_{k-1,\boldsymbol{h}}(n_k)\, a_{\boldsymbol{h}}\, (h_1 \circ h_2 \circ \dots \circ h_k),$$

where $\Phi_{i,\boldsymbol{h}} : N \to N$ is an automorphism for $1 \le i \le k-1$, and $a_{\boldsymbol{h}} \in N$ depends only on $h_1, \dots, h_k$.

Proof: For $1 \le i \le k-1$, define $\Phi_{i,\boldsymbol{h}} : N \to N$ by $\Phi_{i,\boldsymbol{h}}(n) = (h_1 \dots h_i)\, n\, (h_1 \dots h_i)^{-1}$. Since N is a normal subgroup of G, $\Phi_{i,\boldsymbol{h}}$ are automorphisms on N. Observe that,

$$(n_1h_1)(n_2h_2) \dots (n_kh_k) = n_1\, \Phi_{1,\boldsymbol{h}}(n_2)\, \Phi_{2,\boldsymbol{h}}(n_3)\, \dots\, \Phi_{k-1,\boldsymbol{h}}(n_k)\, (h_1h_2 \dots h_k).$$

Further, $h_1h_2 \dots h_k$ and $h_1 \circ h_2 \circ \dots \circ h_k$ (as elements of G) lie in the same coset of N. Thus, there exists $a_{\boldsymbol{h}} \in N$ depending only on $h_1, h_2, \dots, h_k$ such that $h_1h_2 \dots h_k = a_{\boldsymbol{h}}\, (h_1 \circ h_2 \circ \dots \circ h_k)$. The lemma now follows. $\square$

Definition 4.1 *A group G is said to be solvable if there exist subgroups $G = N_0 \supset N_1 \supset N_2 \supset \dots \supset N_r = (e)$ such that each N_i is normal in N_{i-1} and N_{i-1}/N_i is abelian.*

The following properties of solvable groups can be found, for instance, in [Her75].

Lemma 4.3 ([Her75]) *Let G be a solvable group. Then,*

- *For a normal subgroup $N \subseteq G$, the factor group G/N is solvable.*
- *G contains a nontrivial abelian subgroup N which is normal in G.*

Proof of Theorem 1.3: Let G be a solvable group. Let N be a nontrivial abelian subgroup of G that is also normal and let $H = G/N$. Such an N is guaranteed to exist by Lemma 4.3. As before, we assume that the factor group H is given by elements of G which are in distinct cosets of N in G, with group operation of H denoted by $\circ$.

Lemma 4.4 *Let $S \subseteq N^n$ be strongly ϵ-biased in N and let $T \subseteq H^n$ be ϵ-biased in H. Let*

$$S \times T = \{(n_1h_1, \ldots, n_nh_n) : (n_1, \ldots n_n) \in S, (h_1, h_2, \ldots, h_n) \in T\} \subseteq G^n. \tag{3}$$

Then, the set $S \times T$ is ϵ-biased in G.

Given the above lemma, Theorem 1.3 follows from Corollary 2.2 and induction, as H is a solvable group with $|H| < |G|$. We now calculate the exact dependence of the size of the small-bias space on $n, \epsilon, |G|$.

For an abelian group N, and sufficiently large n, Azar et al. give a strongly ϵ-biased set of size $(cn/\epsilon)^{2\log|N|}$, where c is an absolute constant independent of $n, \epsilon, |N|$. Combining the above with Lemma 4.4, for sufficiently large n, we get an ϵ-biased space for G of size $(cn/\epsilon)^{2\log|G|}$. In general, using a similar estimate of Azar et al. we get a bound of $(n/\epsilon)^{O(\log|G|)}$. $\square$

Proof of Lemma 4.4: Fix $b_1, \ldots, b_n \in \{0,1\}$ and let $I = \{i : b_i \neq 0\}$. Without loss of generality, let $I = \{1, \ldots, k\}$. For $\boldsymbol{n} = (n_1, \ldots, n_n) \in N^n$ and $\boldsymbol{h} = (h_1, \ldots, h_n) \in H^n$, let $X(\boldsymbol{n}, \boldsymbol{h}) = (n_1h_1)(n_2h_2)\ldots(n_kh_k)$. Using the notation of Lemma 4.2, let $X(\boldsymbol{n}, \boldsymbol{h}) = Y_{\boldsymbol{h}}(\boldsymbol{n})Z(\boldsymbol{h})$, where

$$\begin{aligned} Y_{\boldsymbol{h}}(\boldsymbol{n}) &= n_1\Phi_{1,\boldsymbol{h}}(n_2)\Phi_{2,\boldsymbol{h}}(n_3)\ldots\Phi_{k-1,\boldsymbol{h}}(n_k)a_{\boldsymbol{h}} \in N, \\ Z(\boldsymbol{h}) &= (h_1 \circ h_2 \circ \cdots \circ h_k) \in H. \end{aligned}$$

Let $g_0 = n_0h_0 \in G$ with $n_0 \in N, h_0 \in H$. Then, for a fixed $\boldsymbol{h}$, since S is strongly ϵ-biased in N,

$$|\mathsf{Pr}_{\boldsymbol{n} \in_u S}[Y_{\boldsymbol{h}}(\boldsymbol{n}) = n_0] - \frac{1}{|N|}| \leq \epsilon. \tag{4}$$

Further, since T is ϵ-biased in H,

$$|\mathsf{Pr}_{\boldsymbol{h} \in_u T}[Z(\boldsymbol{h}) = h_0] - \frac{1}{|H|}| \leq \epsilon. \tag{5}$$

Therefore,

$$\begin{aligned} \mathsf{Pr}_{\boldsymbol{n} \in_u S, \boldsymbol{h} \in_u T}[X(\boldsymbol{n}, \boldsymbol{h}) = g_0] &= \mathsf{Pr}_{\boldsymbol{n} \in_u S, \boldsymbol{h} \in_u T}[Y_{\boldsymbol{h}}(\boldsymbol{n}) = n_0 \wedge Z(\boldsymbol{h}) = h_0] \\ &= \sum_{\boldsymbol{h}: Z(\boldsymbol{h}) = h_0} \frac{1}{|T|} \mathsf{Pr}_{\boldsymbol{n} \in_u S}[Y_{\boldsymbol{h}}(\boldsymbol{n}) = n_0] \\ &\geq \sum_{\boldsymbol{h}: Z(\boldsymbol{h}) = h_0} \frac{1}{|T|} \left(\frac{1}{|N|} - \epsilon\right) \quad \text{from equation (4)} \\ &\geq \left(\frac{1}{|N|} - \epsilon\right) \mathsf{Pr}_{h \in_u T}[Z(\boldsymbol{h}) = h_0] \\ &\geq \left(\frac{1}{|N|} - \epsilon\right)\left(\frac{1}{|H|} - \epsilon\right) \quad \text{from equation (5)} \\ &\geq \frac{1}{|G|} - \beta, \end{aligned}$$

where $\beta = \epsilon/|N| + \epsilon/|H| - \epsilon^2/|G|$. As the above argument is applicable for all non-empty $I \subseteq [n]$, and $\beta < \epsilon$ for $|N|, |H| \geq 2$, we get that $S \times T$ is ϵ-biased in G.$\square$

4.1 Width 3 Branching Programs

We now study the particular case of width 3 ROPBPs and present some motivation for our Conjecture 1.5. Let $\mathbb{S}_3$ be the symmetric group on three elements $\{1,2,3\}$. Let $a \in \mathbb{S}_3$ be the transposition (12), $b \in \mathbb{S}_3$ be the cyclic-shift (123), and e be the identity permutation. Then, the group $N = \{1, b, b^2\} \cong \mathbb{Z}_3$ is an abelian subgroup of $\mathbb{S}_3$ that is also normal. The factor group S_3/N is isomorphic to the group $\{1, a\} \cong \mathbb{Z}_2$. Thus, for the special case of S_3, the construction presented in the previous section becomes,

$$S \times T = \{(a^{x_1}b^{y_1}, a^{x_2}b^{y_2}, \ldots, a^{x_n}b^{y_n}) : (x_1, \ldots, x_n) \in S, (y_1, \ldots, y_n) \in T\},$$

where $S \subseteq \{0,1\}^n$ is ϵ-biased over $\mathbb{Z}_2$ and $T \subseteq \{0,1,2\}^n$ is ϵ-biased over $\mathbb{Z}_3$. This provides some motivation for Conjecture 1.5. Also, note that any pseudorandom generator for width three permutation branching programs must be pseudorandom with respect to linear functions mod 2 and mod 3 - a property satisfied by the generator of (1.5).

5 Pseudorandom Bit Generators for Modular Sums

As the existence of ϵ-PBG as required in Conjecture 1.5 does not follow directly from known constructions of ϵ-biased spaces, we provide an efficient construction of ϵ-PBG of size $poly(n, 1/\epsilon)$ below. Our construction is based on the pseudorandom-generators for low-degree polynomials obtained by Viola [Vio08].

Proof of Theorem 1.4: Suppose that p is an odd prime; for $p = 2$, ϵ-PBG can be obtained from ϵ-biased spaces straightforwardly. For the rest of this section, the arithmetic is over $\mathbb{F}_p$ and let $\oplus$ denote addition mod 2. To motivate our construction, let $v \in \mathbb{F}_p^n, v \neq 0$. We consider two cases depending on the support size of v. Let C be a large enough constant depending on p such that for all $v \in \mathbb{F}_p$ with $|support(v)| \geq C\log(1/\epsilon)$

$$|\mathsf{Pr}_{x \in_u \{0,1\}^n}[\,\langle v, x\rangle = a\,] - \frac{1}{p}| \leq \frac{\epsilon}{3}, \tag{6}$$
$$|\mathsf{Pr}_{y \in_u \mathbb{F}_p^n}[\,\sum_i v_i y_i^{p-1} = a\,] - \frac{1}{p}| \leq \frac{\epsilon}{3}.$$

Case 1: $|support(v)| \leq C\log(1/\epsilon)$. This case can be handled by a generator $H_1 : \{0,1\}^r \to \{0,1\}^n$ that generates a ϵ-almost $C\log(1/\epsilon)$-wise independent distribution. Such generators with $r = O(\log n + \log(1/\epsilon))$ are given in [NN93].

Case 2: $|support(v)| > C\log(1/\epsilon)$. Let $H : \{0,1\}^r \to \mathbb{F}_p^n$ be such that for every degree at most $p-1$ polynomial $P : \mathbb{F}_p^n \to \mathbb{F}_p$, and $a \in \mathbb{F}_p$,

$$|\mathsf{Pr}_{z \in_u \{0,1\}^r}[\,P(H(z)) = a\,] - \mathsf{Pr}_{y \in_u \mathbb{F}_p^n}[\,P(y) = a\,]| \leq \frac{\epsilon}{3}. \tag{7}$$

Pseudorandom generators for low-degree polynomials as above with $r = O(\log n + \log(1/\epsilon))$ were given by Viola building on the works of Bogdanov and Viola [BV07], Lovett [Lov08]. Define $H_2 : \{0,1\}^r :\to \{0,1\}^n$ by $H_2(z) = (y_1^{p-1}, \ldots, y_n^{p-1})$, where $(y_1, \ldots, y_n) = H(z)$. We will show

that H_2 satisfies the conditions of ϵ-PBG. For $z \in \{0,1\}^r$, we have $\langle v, H_2(z)\rangle = \sum_i v_i (H(z)_i)^{p-1} = P_v(H(z))$, where P_v is the degree $p-1$ polynomial $P_v(y) = \sum_i v_i y_i^{p-1}$. Since H fools degree $p-1$ polynomials over $\mathbb{F}_p^n$,

$$|\mathsf{Pr}_{z\in_u\{0,1\}^r}[\,\langle v, H_2(z)\rangle = a\,] - \mathsf{Pr}_{y\in_u\mathbb{F}_p^n}[\,P_v(y) = a\,]| = |\mathsf{Pr}_{z\in_u\{0,1\}^r}[\,P_v(H(z)) = a\,] - \mathsf{Pr}_{y\in_u\mathbb{F}_p^n}[\,P_v(y) = a\,]| \leq \frac{\epsilon}{3}. \quad (8)$$

From equations (6), (8), for $v \in \mathbb{F}_p^n$ with $|support(v)| > C\log(1/\epsilon)$,

$$|\mathsf{Pr}_{z\in_u\{0,1\}^r}[\,\langle v, H_2(z)\rangle = a\,] - \frac{1}{p}| \leq \frac{2\epsilon}{3}. \quad (9)$$

We now combine the generators H_1, H_2 from the above cases to obtain $(H_1 \oplus H_2) : \{0,1\}^{2r} \to \{0,1\}^n$ by defining $(H_1 \oplus H_2)(z_1, z_2) = H_1(z_1) \oplus H_2(z_2)$. Observe that for $b, c \in \{0,1\}$, $b \oplus c = b + c + (p-2)bc \mod p$. Let $v \in \mathbb{F}_p^n, v \neq 0$. Then,

$$\begin{aligned}
\langle v, (H_1 \oplus H_2)(z_1, z_2)\rangle &= \sum_i v_i (H_1(z_1)_i \oplus H_2(z_2)_i) \\
&= \sum_i v_i \left(H_1(z_1)_i + H_2(z_2)_i + (p-2)H_1(z_1)_i H_2(z_2)_i\right) \\
&= \sum_i v_i H_1(z_1)_i + \sum_i v_i (1 + (p-2)H_1(z_1)_i) H_2(z_2)_i \\
&= \langle v, H_1(z_1)\rangle + \langle v(z_1), H_2(z_2)\rangle,
\end{aligned}$$

where $v(z_1)$ is the vector defined by $v(z_1)_i = v_i(1 + (p-2)H_1(z_1)_i)$. Note that $|support(v)| = |support(v(z_1))|$. Fix $a \in \mathbb{F}_p$ and suppose $|support(v)| > C\log(1/\epsilon)$. Then, for a fixed z_1, by equation (9)

$$|\mathsf{Pr}_{z_2\in_u\{0,1\}^r}[\,\langle v(z_1), H_2(z_2)\rangle = a\,] - \frac{1}{p}| \leq \frac{2\epsilon}{3}.$$

Since z_1, z_2 are chosen independently, from the above equation and equation (6)

$$|\mathsf{Pr}_{z_1,z_2\in_u\{0,1\}^r}[\,\langle v, (H_1 \oplus H_2)(z_1, z_2)\rangle = a\,] - \mathsf{Pr}_{x\in_u\{0,1\}^n}[\,\langle v, x\rangle = a\,]| \leq \epsilon.$$

Proceeding similarly for the case $|support(v)| \leq C\log(1/\epsilon)$, it follows that the above inequality holds for all $v \in \mathbb{F}_p^n, v \neq 0$ and $a \in \mathbb{F}_p$. Hence $H_1 \oplus H_2$ is an ϵ-PBG for sums mod p.□

6 Relation to Branching Programs

Here we prove that PRGs for constant-width ROPBP imply small-bias spaces for finite groups.

Proof of Theorem (1.1)**:** Assume we are given a PRG fooling constant-width ROPBPs. We want to construct an ϵ-biased space for a group H. Since finite groups are isomorphic to subgroups of the permutation groups, we can assume H to be a subgroup of the symmetric group $\mathbb{S}_w$ on w elements. Let $G : \{0,1\}^r \to \{0,1\}^n$ fool width w ROPBPs of length n with error at most ϵ. Consider the following procedure for generating a sequence in H^n:

1. Generate ϵ-almost k-wise independent sequences $(g_1, \ldots, g_n)$, $(h_1, \ldots, h_n) \in H^n$ for $k = O(\log(1/\epsilon))$ to be chosen later. For $k = O(\log n)$, Naor and Naor [NN93], Alon et al. [AGHP92] give efficient constructions of almost k-wise independent sequences using $O(\log n + \log(1/\epsilon))$ bits of randomness.
2. Choose $y \in_u \{0,1\}^r$ and output the sequence $(g_1 h_1^{G(y)_1}, g_2 h_2^{G(y)_2}, \ldots, g_n h_n^{G(y)_n})$.

Note that the procedure uses $O(\log n + \log(1/\epsilon)) + r$ bits of randomness. We will show that the multi-set of sequences generated by the above procedure is a $O(\epsilon)$-biased space over H. We need the following lemmas.

Let $g_1, \ldots, g_l \in H$, for l to be chosen later. Call a sequence of group elements $(h_1, \ldots, h_l)$ *complete* if $\{g_1 h_1^{x_1} g_2 h_2^{x_2} \cdots g_l h_l^{x_l} : x_i \in \{0,1\}\} = H$.

Lemma 6.1 *There exists a constant c such that for $l = c|H|^2$, a sequence $(h_1, \ldots, h_l) \in_u H^l$ is complete with probability at least $1/2$.*

Proof For $1 \leq i \leq l$, let $S_i = \{g_1 h_1^{x_1} \cdots g_i h_i^{x_i} : x_j \in \{0,1\}\}$ and let random variable $X_i = |S_i|$. Note that given $h_1, \ldots, h_i$ and a $g \in H$, $g \notin S_i$, $\mathsf{Pr}_{h_{i+1} \in_u H}[g \in S_{i+1} \mid h_1, \ldots, h_i] = X_i/|H| \geq 1/|H|$. Thus, if $X_i < |H|$, then $\mathsf{Pr}[X_{i+1} \geq X_i + 1 \mid X_i] \geq 1/|H|$. The lemma now follows.□

Lemma 6.2 *For any group H and $0 < \epsilon < 1/2$, there exists $k = O(\log(1/\epsilon))$, such that for all $t > k$ and $g_1, \ldots, g_t \in H$ the following holds. For $(x_1, \ldots, x_t) \in_u \{0,1\}^t$ and $h_1, \ldots, h_t \in H$ chosen from an ϵ-almost k-wise independent distribution, the distribution of $g_1 h_1^{x_1} g_2 h_2^{x_2} \ldots g_t h_t^{x_t}$ is 4ϵ-close in variation distance to the uniform distribution on H.*

Proof Let $l = c|H|^2$ be as in Lemma 6.1. Let $k = 4ml\log(1/\epsilon)$, for m to be chosen later, and partition $(h_1, \ldots, h_k)$ into $4m\log(1/\epsilon)$ blocks of length l each. Then, by Lemma 6.1 and Chernoff bounds, for $(h_1, \ldots, h_k) \in H^k$ chosen from an ϵ-almost k-wise independent distribution, with probability at least $1 - (\exp(-\Omega(m\log(1/\epsilon)) + \epsilon)$, $m\log(1/\epsilon)$ of the $4m\log(1/\epsilon)$ blocks will be complete for $g_1, \ldots, g_k$.

Note that for any complete sequence $(h_{i_1}, \ldots, h_{i_l})$ the distribution of $g_{i_1} h_{i_1}^{x_1} g_{i_2} h_{i_2}^{x_2} \ldots g_{i_l} h_{i_l}^{x_l}$ for $x \in_u \{0,1\}^l$ is at least $\alpha = (1 - 1/2^l)$-close in variation distance to the uniform distribution on H. Thus, with probability at least $1-(\exp(-\Omega(m\log(1/\epsilon)))+\epsilon)$, the distribution of $g_1 h_1^{x_1} g_2 h_2^{x_2} \ldots g_t h_k^{x_k}$ with $x_i \in_u \{0,1\}$ is $(1-1/2^l)^{m\log(1/\epsilon)}$-close in variation distance to the uniform distribution on H. The lemma now follows by taking $m = O(2^l)$.□

Let $k = O(\log(1/\epsilon))$ be such that the above lemma holds for H. Let $I = \{i_1, \ldots, i_t\} = \{i : b_i \neq 0\} \subseteq [n]$. We consider two cases.

(a) $|I| = t \leq k$: Since, $(g_1, \ldots, g_n)$ is chosen independently of $(h_1, \ldots, h_n)$ and is ϵ-almost k-wise independent, the distribution of $g_{i_1} h_{i_1}^{G(y)_{i_1}} g_{i_2} h_{i_2}^{G(y)_{i_2}} \ldots g_{i_t} h_{i_t}^{G(y)_{i_t}}$ is ϵ-close to the uniform distribution on H.

(b) $|I| > k$: By relabeling the nodes according to the g_i, we can construct a width w ROPBP of length at most n such that on input $x_1, \ldots, x_t, \ldots, x_n$ the output is $g_{i_1} h_{i_1}^{x_{i_1}} g_{i_2} h_{i_2}^{x_{i_2}} \ldots g_{i_t} h_{i_t}^{x_{i_t}} \in \mathbb{S}_w$. Since G fools ROPBPs of width w and length at most n, we have for every $\pi \in \mathbb{S}_w$,

$$\left| \mathsf{Pr}_{y \in_u \{0,1\}^r}[g_{i_1} h_{i_1}^{G(y)_{i_1}} g_{i_2} h_{i_2}^{G(y)_{i_2}} \ldots g_{i_t} h_{i_t}^{G(y)_{i_t}} = \pi] - \mathsf{Pr}_{x \in_u \{0,1\}^n}[g_{i_1} h_{i_1}^{x_{i_1}} g_{i_2} h_{i_2}^{x_{i_2}} \ldots g_{i_t} h_{i_t}^{x_{i_t}} = \pi] \right| \leq \epsilon.$$

Now, by lemma 6.2 when $x \in_u \{0,1\}^n$, the distribution of $g_{i_1} h_{i_1}^{x_1} g_{i_2} h_{i_2}^{x_2} \dots g_{i_t} h_{i_t}^{x_t}$ is 4ϵ-close to the uniform distribution on H. Therefore, for every $\pi \in \mathbb{S}_w$,

$$\left| \mathsf{Pr}_{y \in_u \{0,1\}^r} [g_{i_1} h_{i_1}^{G(y)_{i_1}} g_{i_2} h_{i_2}^{G(y)_{i_2}} \dots g_{i_t} h_{i_t}^{G(y)_{i_t}} = \pi] - \frac{1}{|H|} \right| \leq 5\epsilon.$$

It follows that the generator defined above is $O(\epsilon)$-biased over H.$\square$

7 Sum of Constant-Bias Spaces Does Not Fool Width 3

We now show that the sum of two constant-bias spaces over $\mathbb{Z}_2$ does not fool width 3 branching programs and prove Theorem 1.6. We do this by constructing a constant-bias space S over $\mathbb{Z}_2$ such that $S \oplus S$ is contained in a subspace of dimension $o(n)$ in $\mathbb{F}_3^n$. To avoid confusion in the following let $+$ denote addition in $\mathbb{F}_3$ and $\oplus$ denote addition in $\mathbb{F}_2$. For a set $T \subseteq \mathbb{F}_3^n$, let $d_3(T)$ denote the dimension of span of T in $\mathbb{F}_3^n$ and let $T \odot T = \{x \odot y = (x_1 y_1, \dots, x_n y_n) : x = (x_1, \dots, x_n), y = (y_1, \dots, y_n) \in T\}$. We'll use the following lemmas.

Lemma 7.1 *For any $T \subseteq \mathbb{F}_3^n$, $d_3(T \odot T) \leq d_3(T)^2$.*

Proof: If $u_1, \dots, u_k \in \mathbb{F}_3^n$ span T, then the $\binom{k}{2} + k$ vectors $u_i \odot u_j$ span $T \odot T$. $\square$

Lemma 7.2 *Let $T \subseteq \{0,1\}^n$. Then $d_3(T \oplus T) \leq 2d_3(T) + d_3(T)^2$.*

Proof: Observe that for $x, y \in \{0,1\}$, $x \oplus y = x + y + xy$. Therefore the dimension of span of $T_1 \oplus T_2$ is at most the dimension of span of $T + T + T \odot T$. The lemma now follows from Lemma 7.1. $\square$

Proof of Theorem 1.6: Let $n = \binom{d}{5}$. We will denote vectors $x \in \mathbb{F}_3^n$, by $(x_I)_{I \in \mathcal{C}}$, where $\mathcal{C} = \binom{[d]}{5}$ is the collection of subsets of $[d]$ of size 5. Let $p : \mathbb{F}_3^5 \to \mathbb{F}_3$ be the degree two multi-variate polynomial defined by

$$p(y_1, y_2, y_3, y_4, y_5) = (y_1 + y_2 + y_3 + y_4 + y_5)^2.$$

Let $q : \mathbb{F}_2^5 \to \mathbb{F}_2$ be the degree five multi-variate polynomial defined by

$$q(y_1, y_2, y_3, y_4, y_5) = \bigoplus_i y_i \oplus \bigoplus_{i \neq j} y_i y_j \oplus \bigoplus_{i,j,k,l \text{ distinct}} y_i y_j y_k y_l \oplus y_1 y_2 y_3 y_4 y_5.$$

Our construction is based on the observation - which can be verified by direct computation - that evaluated over the set $\{0,1\}^5$ the polynomials p and q are identical. That is, for all $(y_1, \dots, y_5) \in \{0,1\}^5$, $p(y_1, \dots, y_5) = q(y_1, \dots, y_5)$. Now, let

$$S = \{(p(y_{i_1}, y_{i_2}, y_{i_3}, y_{i_4}, y_{i_5}))_{\{i_1,i_2,i_3,i_4,i_5\} \in \mathcal{C}} : (y_1, \dots, y_d) \in \{0,1\}^d\}.$$

Let $T = \{(y_{i_1} + y_{i_2} + y_{i_3} + y_{i_4} + y_{i_5})_{\{i_1,i_2,i_3,i_4,i_5\} \in \mathcal{C}} : (y_1, \dots, y_d) \in \{0,1\}^d\}$. Now, $d_3(T) \leq d$ and $S \subseteq T \odot T$. Therefore, by Lemma 7.1 $d_3(S) \leq d^2$. However, the dimension of the span of S viewed as a subset of $\mathbb{F}_2^n$ is $n = \binom{d}{5}$. In fact, for any non-zero $\boldsymbol{\alpha} \in \mathbb{F}_2^n$, $\{\langle \boldsymbol{\alpha}, x \rangle : x \in S\} = \{q_{\boldsymbol{\alpha}}(y) : y \in \mathbb{F}_2^d\}$, where $q_{\boldsymbol{\alpha}} : \mathbb{F}_2^d \to \mathbb{F}_2$ is a non-constant polynomial of degree 5. For,

$$q_{\boldsymbol{\alpha}}(y) = \sum_{I \in \mathcal{C}} \boldsymbol{\alpha}_I \prod_{i \in I} y_i + R_{\boldsymbol{\alpha}}(y),$$

where $R_{\boldsymbol{\alpha}}$ is a degree at most four polynomial. Since $\boldsymbol{\alpha} \neq 0$, $q_{\boldsymbol{\alpha}}$ has degree five. Using the fact that the minimum distance of the Reed Muller code of degree 5 over $\mathbb{F}_2$ is 1/32, we get that for $a \in \{0,1\}$,

$$\mathsf{Pr}_{x \in_u S}[\langle \boldsymbol{\alpha}, x \rangle = a] = \mathsf{Pr}_{y \in_u \mathbb{F}_2^d}[\, q_{\boldsymbol{\alpha}}(y) = a \,] \leq \frac{31}{32}.$$

Thus, S is ϵ-biased over $\mathbb{F}_2$ for $\epsilon = 15/32$. Further, from Lemma 7.2 $d_3(S \oplus S) \leq 2d^2 + d^4 = o(n)$. □

Acknowledgements. We thank Russell Impagliazzo for useful discussions and the anonymous referees for helpful comments.

References

[AGHP92] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta, *Simple construction of almost k-wise independent random variables*, Random Struct. Algorithms **3** (1992), no. 3, 289–304.

[AKS87] Miklós Ajtai, János Komlós, and Endre Szemerédi, *Deterministic simulation in logspace*, STOC, 1987, pp. 132–140.

[AMN98] Yossi Azar, Rajeev Motwani, and Joseph Naor, *Approximating probability distributions using small sample spaces*, Combinatorica **18** (1998), no. 2, 151–171.

[Bar86] David A. Mix Barrington, *Bounded-width polynomial-size branching programs recognize exactly those languages in* NC^1, STOC, 1986, pp. 1–5.

[BDVY08] Andrej Bogdanov, Zeev Dvir, Elad Verbin, and Amir Yehudayoff, *Pseudorandomness for width 2 branching programs*, Manuscript, 2008.

[BEG+94] Manuel Blum, William S. Evans, Peter Gemmell, Sampath Kannan, and Moni Naor, *Checking the correctness of memories*, Algorithmica **12** (1994), no. 2/3, 225–244.

[BV07] Andrej Bogdanov and Emanuele Viola, *Pseudorandom bits for polynomials*, FOCS, 2007, pp. 41–51.

[Her75] I Herstein, *Topics in algebra*, Wiley, 1975.

[INW94] Russell Impagliazzo, Noam Nisan, and Avi Wigderson, *Pseudorandomness for network algorithms*, STOC, 1994, pp. 356–364.

[Lov08] Shachar Lovett, *Unconditional pseudorandom generators for low degree polynomials*, STOC, 2008, pp. 557–562.

[LRTV09] Shachar Lovett, Omer Reingold, Luca Trevisan, and Salil Vadhan, *Pseudorandom bit generators that fool modular sums*, RANDOM, 2009.

[LY94] Carsten Lund and Mihalis Yannakakis, *On the hardness of approximating minimization problems*, J. ACM **41** (1994), no. 5, 960–981.

[Nis92] Noam Nisan, *Pseudorandom generators for space-bounded computation*, Combinatorica **12** (1992), no. 4, 449–461.

[NN93] Joseph Naor and Moni Naor, *Small-bias probability spaces: Efficient constructions and applications*, SIAM J. Comput. **22** (1993), no. 4, 838–856.

[NZ96] Noam Nisan and David Zuckerman, *Randomness is linear in space*, J. Comput. Syst. Sci. **52** (1996), no. 1, 43–52.

[Rei08] Omer Reingold, *Undirected connectivity in log-space*, J. ACM **55** (2008), no. 4.

[RR99] Ran Raz and Omer Reingold, *On recycling the randomness of states in space bounded computation*, STOC, 1999, pp. 159–168.

[RV06] Omer Reingold and Salil Vadhan, personal communication, 2006.

[SZ] Michael Saks and David Zuckerman, unpublished manuscript.

[Vio08] Emanuele Viola, *The sum of d small-bias generators fools polynomials of degree d*, IEEE Conference on Computational Complexity, 2008, pp. 124–127.