Symposium on Theoretical Aspects of Computer Science 2010 (Nancy, France), pp. 59-70
www.stacs-conf.org

THE REMOTE POINT PROBLEM, SMALL BIAS SPACES, AND EXPANDING GENERATOR SETS

V. ARVIND AND SRIKANTH SRINIVASAN

The Institute of Mathematical Sciences,
C.I.T. Campus, Chennai 600 113, India.

E-mail address, V. Arvind: arvind@imsc.res.in

E-mail address, Srikanth Srinivasan: srikanth@imsc.res.in

Abstract. Using ε-bias spaces over $\mathbb{F}_2$, we show that the Remote Point Problem (RPP), introduced by Alon et al [APY09], has an NC^2 algorithm (achieving the same parameters as [APY09]). We study a generalization of the Remote Point Problem to groups: we replace $\mathbb{F}_2^n$ by $\mathcal{G}^n$ for an arbitrary fixed group $\mathcal{G}$. When $\mathcal{G}$ is Abelian we give an NC^2 algorithm for RPP, again using ε-bias spaces. For nonabelian $\mathcal{G}$, we give a deterministic polynomial-time algorithm for RPP. We also show the connection to construction of expanding generator sets for the group $\mathcal{G}^n$. All our algorithms for the RPP achieve essentially the same parameters as [APY09].

1. Introduction

Valiant, in his celebrated work [V77] on circuit lower bounds for computing linear transformations $A : \mathbb{F}^n \longrightarrow \mathbb{F}^m$ for a field $\mathbb{F}$, initiated the study of rigid matrices. If explicit rigid matrices of certain parameters can be constructed it would result in superlinear lower bounds for logarithmic depth linear circuits over $\mathbb{F}$. This problem and the construction of such rigid matrices has remained elusive for over three decades.

Alon, Panigrahy and Yekhanin [APY09] recently proposed a problem that appears to be of intermediate difficulty. Given a subspace L of $\mathbb{F}_2^n$ by its basis and a number $r \in [n]$ as input, the problem is to compute in deterministic polynomial time a point $v \in \mathbb{F}_2^n$ such that $\Delta(u, v) \geq r$ for all $u \in L$, where $\Delta(u, v)$ is the Hamming distance. They call this the *Remote Point Problem*. The point v is said to be r-far from the subspace L.

1998 ACM Subject Classification: Algorithms and Complexity Theory.
Key words and phrases: Small Bias Spaces, Expander Graphs, Cayley Graphs, Remote Point Problem.

27th Symposium on Theoretical Aspects of Computer Science, Nancy, 2010
Editors: Jean-Yves Marion, Thomas Schwentick
Leibniz International Proceedings in Informatics (LIPIcs), Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany
Digital Object Identifier: 10.4230/LIPIcs.STACS.2010.2444

Alon et al [APY09] give a nice polynomial time-bounded (in n) algorithm for computing a $v \in \mathbb{F}_2^n$ that is $c \log n$-far from a given subspace L of dimension $n/2$ and c is a fixed constant. For L such that $\dim(L) = k < n/2$ they give a polynomial-time algorithm for computing a point $v \in \mathbb{F}_2^n$ that is $\frac{cn \log k}{k}$-far from L.

Results of this paper. In [AS09a] we recently investigated the problem of proving circuit lower bounds in the presence of help functions. Specifically, one of the problems we consider is proving lower bounds for constant-depth Boolean circuits which can take a given set of (arbitrary) help functions $\{h_1, h_2, \cdots, h_m\}$ at the input level, where $h_i : \{0,1\}^n \longrightarrow \{0,1\}$ for each i. Proving explicit lower bounds for this model would allow us to separate EXP from the polynomial-time many-one closure of nonuniform AC^0. We show that it suffices to find a polynomial-time solution to the Remote Point Problem for parameters $k = 2^{(\log \log n)^c}$ and $r = \frac{n}{2^{(\log \log n)^d}}$ for all constants c and d. Unfortunately, the parameters of the Alon et al algorithm are inadequate for our application.

However, motivated by this connection, in the present paper we carry out a more detailed study of the Remote Point Problem as an algorithmic question. We briefly summarize our results.

1. The first question we address is whether we can give a deterministic parallel (i.e. NC) algorithm for the problem — Alon et al's algorithm is inherently sequential as it is based on the method of conditional probabilities and pessimistic estimators.

It turns out an element of an ε-bias space for suitably chosen ε is a solution to the Remote Point Problem which gives us an NC algorithm quite easily.

2. Since the RPP for $\mathbb{F}_2^n$ can be solved using small bias spaces, it naturally leads us to address the problem in a more general group-theoretic setting.

In the generalization we study we will replace $\mathbb{F}_2$ with an arbitrary fixed finite group $\mathcal{G}$ such that $|\mathcal{G}| \geq 2$. Hence we will have the n-fold product group $\mathcal{G}^n$ instead of the vector space $\mathbb{F}_2^n$.

Given elements $x = (x_1, x_2, \ldots, x_n), y = (y_1, y_2, \ldots, y_n)$ of $\mathcal{G}^n$, let $\Delta(x, y) = |\{i \mid x_i \neq y_i\}|$. I.e. $\Delta(x, y)$ is the *Hamming distance* between x and y. Furthermore, for $S \subseteq \mathcal{G}^n$, let $\Delta(x, S)$ denote $\min_{y \in S} \Delta(x, y)$.

We now define the *Remote Point Problem (RPP) over a finite group* $\mathcal{G}$. The input is a subgroup $\mathcal{H}$ of $\mathcal{G}^n$, where $\mathcal{H}$ is given by a generating set, and a number $r \in [n]$. The problem is to compute in deterministic polynomial (in n) time an element $x \in \mathcal{G}^n$ such that $\Delta(x, H) > r$. The results we show in this general setting are the following.

(a) The Remote Point Problem over any *Abelian group* $\mathcal{G}$ has an NC^2 algorithm for $r = O(\frac{n \log k}{k})$ and $k \leq n/2$, where $k = \log_{|\mathcal{G}|} |\mathcal{H}|$.
(b) Over an arbitrary group $\mathcal{G}$ the Remote point problem has a polynomial-time algorithm for $r = O(\frac{n \log k}{k})$ and $k \leq n/2$, where $k = \log_{|\mathcal{G}|} |\mathcal{H}|$.

The parallel algorithm stated in part(a) above is based on ε-bias space constructions for finite Abelian groups described in Azar et al [AMN98]. The sequential algorithm stated in part(b) above is a group-theoretic generalization of the Alon et al algorithm for $\mathbb{F}_2^n$ [APY09].

Due to lack of space, some proofs have been omitted. They may be found in the full version which has been published as an ECCC report [AS09b].

2. Preliminaries

Fix a finite group $\mathcal{G}$ such that $|\mathcal{G}| \geq 2$. Given any $x \in \mathcal{G}^n$, let $wt(x)$ denote the number of coordinates i such that $x_i \neq 1$, where 1 is the identity of the group $\mathcal{G}$. By $B(r)$, we will refer to the set of $x \in \mathcal{G}^n$ such that $wt(x) \leq r$. Given a subset S of $\mathcal{G}^n$, $B(S, r)$ will denote the set $S \cdot B(r) = \{sx \mid s \in S, x \in B(r)\}$. Clearly, for any $S \subseteq \mathcal{G}^n$ and any $x \in \mathcal{G}^n$, $x \in B(S, r)$ if and only if $\Delta(x, S) \leq r$. We say that x is *r-close* to S if $x \in B(S, r)$ and *r-far* from S if $x \notin B(S, r)$.

The *Remote Point Problem (RPP) over* $\mathcal{G}$ is defined to be the following algorithmic problem:

> INPUT: A subgroup $\mathcal{H}$ of $\mathcal{G}^n$ (given by its generators) and an $r \in \mathbb{N}$.
> OUTPUT: An $x \in \mathcal{G}^n$ such that $x \notin B(\mathcal{H}, r)$.

Clearly, there are inputs to the above problem where no solution can be found. But the input instances of the kind that we will study will clearly have a solution (in fact, a random point of $\mathcal{G}^n$ will be a solution with high probability).

Given a subgroup $\mathcal{H}$ of $\mathcal{G}^n$, denote by $\delta(\mathcal{H})$ the quantity $\log_{|\mathcal{G}|} |\mathcal{H}|$. We will call $\delta(\mathcal{H})$ the *dimension of $\mathcal{H}$ in $\mathcal{G}^n$*.

We say that the RPP over $\mathcal{G}$ has a $(k(n), r(n))$-algorithm if there is an efficient algorithm that solves the Remote Point Problem when given as input a subgroup $\mathcal{H}$ of $\mathcal{G}^n$ of dimension at most $k(n)$ and an r that is bounded by $r(n)$. (Here, 'efficient' can correspond to polynomial time or some smaller complexity class.)

A simple counting argument shows that there is a valid solution to the RPP over $\mathcal{G}$ on inputs $(\mathcal{H}, r)$ where $\delta(\mathcal{H}) + r \leq n(1 - \frac{H(r/n)}{\log |G|} - \varepsilon)$, for any fixed $\varepsilon > 0$ (where $H(\cdot)$ denotes the binary entropy function). However, the best known deterministic solution to the RPP – from [APY09] – is a polynomial time $(k, \frac{cn \log k}{k})$-algorithm which works over $\mathbb{F}_2^n$ (i.e, the group $\mathcal{G}$ involved is the additive group of the field $\mathbb{F}_2$).

2.1. Some Group-Theoretic Algorithms

We introduce basic definitions and review some group-theoretic algorithms. Let $\mathrm{Sym}(\Omega)$ denote the group of all permutations on a finite set Ω of size m. In this section we use G, H etc. to denote *permutation groups on* Ω, which are simply subgroups of $\mathrm{Sym}(\Omega)$.

Let G be a subgroup of $\mathrm{Sym}(\Omega)$. For a subset $\Delta \subseteq \Omega$ denote by $G_{\{\Delta\}}$ the *point-wise stabilizer* of Δ. I.e $G_{\{\Delta\}}$ is the subgroup consisting of exactly those elements of G that fix each element of Δ.

Theorem 2.1 (Schreier-Sims)**.** [Lu93]

(1) *If a subgroup G of $\mathrm{Sym}(\Omega)$ is given by a generating set as input along with the subset Δ there is a polynomial-time (sequential) algorithm for computing a generator set for $G_{\{\Delta\}}$.*
(2) *If a subgroup G of $\mathrm{Sym}(\Omega)$ is given by a generating set as input, then there is a polynomial time algorithm for computing $|G|$.*
(3) *Given as input a permutation $\sigma \in \mathrm{Sym}(\Omega)$ and a generator set for a subgroup G of $\mathrm{Sym}(\Omega)$, we can test in deterministic polynomial time if σ is an element of G.*

We are also interested in a special case of this problem which we now define. A subset $\Gamma \subseteq \Omega$ is an *orbit* of G if $\Gamma = \{\sigma(i) \mid \sigma \in G\}$ for some $i \in \Omega$. Any subgroup G of $\mathrm{Sym}(\Omega)$ partitions Ω into orbits (called G-orbits).

For a constant $b > 0$, a subgroup G of $\mathrm{Sym}(\Omega)$ is defined to be a *b-bounded permutation group* if every G-orbit is of size at most b.

In [MC87], McKenzie and Cook studied the parallel complexity of *Abelian* permutation group problems. Specifically, they gave an NC^3 algorithm for testing membership in an Abelian permutation group given by a generator set and for computing the order of an Abelian permutation group. When restricted to b-bounded Abelian permutation groups, the algorithms of [MC87] for these problems are actually NC^2 algorithms. We formally state their result and derive a consequence.

Theorem 2.2 ([MC87])**.** *There is an NC^2 algorithm for membership testing in a b-bounded Abelian permutation group G given by a generator set.*

We now consider problems over $\mathcal{G}^n$, for a fixed finite group $\mathcal{G}$. We know from basic group theory that every group $\mathcal{G}$ is a permutation group acting on itself. I.e. every $\mathcal{G}$ can be seen as a subgroup of $\mathrm{Sym}(\mathcal{G})$, where $\mathcal{G}$ acts on itself by left (or right) multiplication. Therefore, $\mathcal{G}^n$ can be easily seen as a permutation group on the set $\Omega = \mathcal{G} \times [n]$ and hence, $\mathcal{G}^n$ can be considered a subgroup of $\mathrm{Sym}(\Omega)$. Furthermore, notice that each subset $\mathcal{G} \times \{i\}$ is an orbit of this group $\mathcal{G}^n$. Hence, $\mathcal{G}^n$ is a b-bounded permutation group contained in $\mathrm{Sym}(\Omega)$, where $b = |\mathcal{G}|$. Finally, if $\mathcal{G}$ is an Abelian group, then so is this subgroup of $\mathrm{Sym}(\Omega)$. We have the following lemma as an easy consequence of Theorem 2.2.

Lemma 2.3. *Let $\mathcal{G}$ be Abelian. There is an NC^2 algorithm that takes as input a generator set for some subgroup $\mathcal{H}$ of $\mathcal{G}^n$ and an $x \in \mathcal{G}^n$, and accepts iff $x \in \mathcal{H}$.*

Given any $y = (y_1, y_2, \ldots, y_i) \in \mathcal{G}^i$ with $1 \leq i \leq n$ and any $S \subseteq \mathcal{G}^n$, let S_y denote the set $\{x \in S \mid x_j = y_j \text{ for } 1 \leq j \leq i\}$.

Lemma 2.4. *Let $\mathcal{G}$ be any fixed finite group. There is a polynomial time algorithm that takes as input a subgroup $\mathcal{H}$ of $\mathcal{G}^n$, where $\mathcal{H}$ is given by generators, and a $y \in \mathcal{G}^i$ with $1 \leq i \leq n$, and computes $|\mathcal{H}_y|$.*

Proof. Let $\mathcal{K} = \{(x_1, x_2, \ldots, x_n) \in \mathcal{H} \mid x_1 = x_2 = \cdots = x_i = 1\}$, where 1 denotes the identity element of $\mathcal{G}$. Clearly, $\mathcal{K}$ is a subgroup of $\mathcal{H}$. The set $\mathcal{H}_y$, if nonempty, is simply a coset of $\mathcal{K}$ and thus, we have $|\mathcal{H}_y| = |\mathcal{K}|$. To check if $\mathcal{H}_y$ is nonempty, we consider the map $\pi_i : \mathcal{G}^n \to \mathcal{G}^i$ that projects its input onto its first i coordinates; note that $\mathcal{H}_y$ is nonempty iff the subgroup $\pi_i(\mathcal{H})$ contains y, which can be checked in polynomial time by point (3) of Theorem 2.1 (here, we are identifying $\mathcal{G}^n$ with a subgroup of $\mathrm{Sym}(\mathcal{G} \times [n])$ as above). If $y \notin \pi_i(\mathcal{H})$, the algorithm outputs 0. Otherwise, we have $|\mathcal{H}_y| = |\mathcal{K}|$ and it suffices to compute $|\mathcal{K}|$. But $\mathcal{K}$ is simply the point-wise stabilizer of the set $\mathcal{G} \times [i]$ in $\mathcal{H}$, and hence $|\mathcal{K}|$ can be computed in polynomial time by points (1) and (2) of Theorem 2.1. ∎

3. Expanding Cayley Graphs and the Remote Point Problem

Fix a group $\mathcal{G}$ such that $|\mathcal{G}| \geq 2$, and consider an instance of the RPP over $\mathcal{G}$. The main idea that we develop in this section is that if we have a (symmetric) expanding generator set S for the group $\mathcal{G}^n$ with appropriate expansion parameters then for a subgroup $\mathcal{H}$ of $\mathcal{G}^n$ such that $\delta(\mathcal{H}) \leq k$ some element of S will be r-far from H, for suitable k and r.

We review some definitions related to expander graphs (e.g. see the survey of Hoory, Linial, and Wigderson [HLW06]). An undirected multigraph $G = (V, E)$ is an (n, d, α)-graph for $n, d \in \mathbb{N}$ and $\alpha > 0$ if $|V| = n$, the degree of each vertex is d, and the second largest value $\lambda(G)$ from among the absolute values of eigenvalues of $A(G)$ – the adjacency matrix of the graph G – is bounded by αd.

A *random walk* of length $t \in \mathbb{N}$ on an (n, d, α)-graph $G = (V, E)$ is the output of the following random process: a vertex $v_0 \in V$ of picked uniformly at random, and for $0 \leq i < t$, if v_i has been picked, then v_{i+1} is obtained by selecting a neighbour v_{i+1} uniformly at random (i.e a random edge out of v_i is picked, and v_{i+1} is chosen to be the other endpoint of the edge); the output of the process is $(v_0, v_1, \ldots, v_t)$. We now state an important result regarding random walks on expanders (see [HLW06, Theorem 3.6] for details).

Lemma 3.1. *Let $G = (V, E)$ be an (n, d, α)-graph and $B \subseteq V$ with $|B| \leq \beta n$. Then, the probability that a random walk $(v_0, v_1, \ldots, v_t)$ is entirely contained inside B (i.e, $v_i \in B$ for each i) is bounded by $(\beta + \alpha)^t$.*

Let $\mathcal{H}$ be a group and S a *symmetric* multiset of elements from $\mathcal{H}$. I.e. there is a bijection of multisets $\varphi : S \to S$ such that $\varphi(s) = s^{-1}$ for each $s \in S$. We define the Cayley graph $C(\mathcal{H}, S)$ to be the (multi)graph G with vertex set $\mathcal{H}$ and edges of the form (x, xs) for each $x \in \mathcal{H}$ and each $s \in S$; since S is symmetric, we consider $C(\mathcal{H}, S)$ to be an undirected graph by identifying the edges (x, xs) and $(xs, (xs)\varphi(s))$, for each x and s.

We now show a lemma that will help relate generators of expanding Cayley graphs on $\mathcal{G}^n$ and the RPP over $\mathcal{G}$. In what follows, let S be a symmetric multiset of elements from $\mathcal{G}^n$; let G denote the Cayley graph $C(\mathcal{G}^n, S)$; and let N, D denote $|\mathcal{G}|^n$ and $|S|$ (counted with repetitions) respectively.

Lemma 3.2. *Assume S as above is such that G is an (N, D, α)-graph, where $\alpha \leq \frac{1}{n^d}$, for some fixed $d > 0$. Then, given any subgroup $\mathcal{H}$ of $\mathcal{G}^n$ such that $\delta(\mathcal{H}) \leq 2n/3$, we have $\frac{|S \cap \mathcal{H}|}{|S|} \leq \frac{1}{n^{d/2}}$ for large enough n (where the elements of $S \cap \mathcal{H}$ are counted with repetitions).*

Proof. Let $S' = S \cap \mathcal{H}$ and let $\eta = |S'|/|S|$. We want an upper bound on η. Consider a random walk $(x_0, x_1, \ldots, x_t)$ of length t on the graph G (the exact value of t will be fixed later). Let $\mathcal{B}$ denote the following event: there is a $y \in \mathcal{G}^n$ such that all the vertices $x_0, x_1, \ldots, x_t$ are all contained in the coset $y\mathcal{H}$ of $\mathcal{H}$. Let p denote the probability that $\mathcal{B}$ occurs.

We will first lower bound p. At each step of the random walk, a random $s_i \in S$ is chosen and x_{i+1} is set to $x_i s_i$. If these s_i all happen to belong to S', then the cosets $x_i\mathcal{H}$ and $x_{i+1}\mathcal{H}$ are the same for all i and hence, the event $\mathcal{B}$ does occur. Hence, $p \geq \eta^t$.

We now upper bound p. Fix any coset $y\mathcal{H}$ of the subgroup $\mathcal{H}$. Since the dimension of $\mathcal{H}$ in $\mathcal{G}^n$ is bounded by $2n/3$, we have $|y\mathcal{H}| = |\mathcal{H}| \leq |\mathcal{G}|^{2n/3} \leq 2^{-n/3}|\mathcal{G}^n|$. That is, the coset $y\mathcal{H}$ is a very small subset of $\mathcal{G}^n$. Applying Lemma 3.1, we see that the probability that the random walk $(x_0, x_1, \ldots, x_t)$ is completely contained inside this coset is bounded by $(2^{-n/3} + n^{-d})^t \leq \frac{2^t}{n^{dt}}$, for large enough n. As the total number of cosets of $\mathcal{H}$ is bounded by $|\mathcal{G}|^n$, an application of the union bound tells us that p is upper bounded by $|\mathcal{G}|^n \frac{2^t}{n^{dt}} \leq \frac{|\mathcal{G}|^{n+t}}{n^{dt}}$. Setting $t = \frac{2n}{d \log_{|G|} n - 2}$ we see that p is at most $\frac{1}{n^{dt/2}}$.

Putting the upper and lower bounds together, we see that $\eta^t \leq \frac{1}{n^{dt/2}}$ and hence, $\eta \leq \frac{1}{n^{d/2}}$. This completes the proof. ∎

We follow the structure of the algorithm for the RPP over $\mathbb{F}_2$ in [APY09]. We first describe their $(n/2, c\log n)$-algorithm for the RPP, followed by our own algorithm. We then describe how they extend this algorithm to a $(k, \frac{cn \log k}{k})$-algorithm for any $k \leq n/2$; the same procedure works for our algorithm also.

The $(n/2, c\log n)$-algorithm proceeds as follows. On an input instance consisting of a subgroup V (which is a subspace of $\mathbb{F}_2^n$) of dimension at most $n/2$ and an $r \leq c\log n$,

(1) The algorithm first computes a collection of $m = n^{O(c)}$ subspaces $V_1, V_2, \ldots, V_m$, each of dimension at most $2n/3$ such that $B(V, c\log n) \subseteq \bigcup_{i=1}^m V_i$.
(2) The algorithm then finds an $x \in \mathbb{F}_2^n$ such that $x \notin \bigcup_i V_i$. (This is done using a method similar to the method of pessimistic estimators introduced by Raghavan [Rag88].)

Our algorithm will proceed exactly as the above algorithm in the first step. The second step of our algorithm will be different (assuming that the group $\mathcal{G}$ is Abelian). We first state Step 1 of the algorithm of [APY09] in greater generality:

Lemma 3.3. *Let $\mathcal{G}$ be any fixed finite group with $|\mathcal{G}| \geq 2$. For any constant $c > 0$ and large enough n, the following holds. Given any subgroup $\mathcal{H}$ of $\mathcal{G}^n$ such that $\delta(\mathcal{H}) \leq \frac{n}{2}$, there is a collection of $m \leq n^{10c}$ subgroups $\mathcal{H}_1, \mathcal{H}_2, \ldots, \mathcal{H}_m$ such that $B(\mathcal{H}, c\log n) \subseteq \bigcup_{i=1}^m \mathcal{H}_i$, and*

$\delta(\mathcal{H}_i) \leq 2n/3$ *for each* i*. Moreover, there is a logspace algorithm that, when given as input* $\mathcal{H}$ *as a set of generators, produces generators for the subgroups* $\mathcal{H}_1, \mathcal{H}_2, \ldots, \mathcal{H}_m$*.*

Proof. The proof follows exactly as in [APY09]. We reproduce it here for completeness and to analyze the complexity of the procedure.

Let 1 denote the identity element of $\mathcal{G}$. For each $S \subseteq [n]$, let $\mathcal{G}^n(S)$ denote the subgroup of $\mathcal{G}^n$ consisting of those x such that $x_i = 1$ for each $i \notin S$. Note that $\delta(\mathcal{G}^n(S)) = |S|$. Also note that for each $S \subseteq [n]$, the group $\mathcal{G}^n(S)$ is a normal subgroup; in particular, this implies that the set $\mathcal{K} \cdot \mathcal{G}^n(S)$ is a subgroup of $\mathcal{G}^n$ whenever $\mathcal{K}$ is a subgroup of $\mathcal{G}^n$.

Partition the set $[n]$ into $\ell \leq 10c \log n$ sets of size at most $\lceil \frac{n}{10c \log n} \rceil$ each – we will call these sets $S_1, S_2, \ldots, S_\ell$. For each $A \subseteq [\ell]$ of size $\lceil c \log n \rceil$, let $\mathcal{K}_A$ denote the subgroup $\mathcal{G}^n(\bigcup_{i \in A} S_i)$. Note that the number of such subgroups is at most $2^\ell \leq n^{10c}$. Also, for each A as above, $\delta(\mathcal{K}_A) = |\bigcup_{i \in A} S_i| \leq \left(\frac{n}{10c \log n} + 1\right)(c \log n + 1) < \frac{n}{9}$, for large enough n.

Consider any $x \in B(c \log n)$ (i.e, an element x of $\mathcal{G}^n$ s.t $wt(x) \leq c \log n$). We know that $x \in \mathcal{G}^n(S)$ for some S of size at most $c \log n$. Hence, it can be seen that $x \in \mathcal{G}^n(\bigcup_{i \in A} S_i)$ for some A of size $\lceil c \log n \rceil$; this shows that $B(c \log n) \subseteq \bigcup_A \mathcal{K}_A$. Therefore, we see that $B(\mathcal{H}, c \log n) = \mathcal{H} B(c \log n) \subseteq \bigcup_A \mathcal{H}\mathcal{K}_A$.

For each $A \subseteq [\ell]$ of size $\lceil c \log n \rceil$, let $\mathcal{H}_A$ denote the subgroup $\mathcal{H}\mathcal{K}_A$ (note that this is indeed a subgroup, since $\mathcal{K}_A$ is a normal subgroup). Moreover, the cardinality of this subgroup is bounded by $|\mathcal{H}| \cdot |\mathcal{K}_A| \leq |\mathcal{G}|^{n/2} |\mathcal{G}|^{n/9} < |\mathcal{G}|^{2n/3}$; hence, $\delta(\mathcal{H}_A) \leq 2n/3$. Thus, the collection of subgroups $\{\mathcal{H}_A\}_A$ satisfies all the properties mentioned in the statement of the lemma. That a set of generators for this subgroup can be computed in deterministic logspace – for some suitable choice of $S_1, S_2, \ldots, S_\ell$ – is a routine check from the definition of the subgroups $\{\mathcal{K}_A\}_A$. This completes the proof of the lemma. ■

Using Lemma 3.3, we are able to efficiently "cover" $B(\mathcal{H}, c \log n)$ for any small subgroup $\mathcal{H}$ of $\mathcal{G}^n$ by a union of small subgroups. Therefore, to find a point that is $c \log n$-far from $\mathcal{H}$, it suffices to find a point $x \in \mathcal{G}^n$ not contained in any of the covering subgroups. To do this, we note that if S is a multiset containing elements from $\mathcal{G}^n$ such that $C(\mathcal{G}^n, S)$ is a Cayley graph with good expansion, then S must contain such an element. This is formally stated below.

Lemma 3.4. *For any constant* $c > 0$ *and large enough* $n \in \mathbb{N}$*, the following holds. Let* S *be any multiset of elements of* $\mathcal{G}^n$ *such that* $\lambda(C(\mathcal{G}^n, S)) < \frac{1}{n^{20c}}$*. Then, for* $m \leq n^{10c}$ *and any collection* $\mathcal{H}_1, \mathcal{H}_2, \ldots, \mathcal{H}_m$ *of subgroups such that* $\delta(\mathcal{H}_i) \leq 2n/3$ *for each* i*, there is some* $s \in S$ *such that* $s \notin \bigcup_i \mathcal{H}_i$*.*

Proof. The proof follows easily from Lemma 3.2. Given any $i \in [m]$, we know, from Lemma 3.2, that $|S \cap \mathcal{H}_i| < \frac{|S|}{n^{10c}}$ (where the elements of the multisets are counted with repetitions). Hence, $|S \cap \bigcup_i \mathcal{H}_i| \leq \sum_i |S \cap \mathcal{H}_i| < \frac{m|S|}{n^{10c}} \leq |S|$. Therefore, there must be some $s \in S$ such that $s \notin \bigcup_i \mathcal{H}_i$. ■

Therefore, to find a point x that is $c\log n$-far from the subspace $\mathcal{H}$, it suffices to construct an S such that $C(\mathcal{G}^n, S)$ is a sufficiently good expander, find the covering subgroups $\mathcal{H}_i$ ($i \in [m[$), and then to find an $s \in S$ that does not lie in any of the $\mathcal{H}_i$. We follow the above approach to give an efficient parallel algorithm for the RPP in the case that $\mathcal{G}$ is an Abelian group. For arbitrary groups, we show that the method of [APY09] yields a polynomial time algorithm.

4. Remote Point Problem for Abelian Groups

Fix an Abelian group $\mathcal{G}$. Recall that a *character* χ of $\mathcal{G}^n$ is a homomorphism from $\mathcal{G}^n$ to $\mathbb{C}_1^*$, the multiplicative subgroup of the complex numbers of absolute value 1. For $\varepsilon > 0$, a distribution μ over $\mathcal{G}^n$ is said to be ε-biased if, given any non-trivial character χ of $\mathcal{G}^n$, $|\mathbf{E}_{x\sim\mu}[\chi(x)]| \leq \varepsilon$.

A multiset S consisting of elements from $\mathcal{G}^n$ is said to be an *ε-biased space in $\mathcal{G}^n$* if the uniform distribution over S is an ε-biased distribution.

It can be checked that a multiset consisting of $(\frac{n}{\varepsilon})^{O(1)}$ independent, uniformly random elements from $\mathcal{G}^n$ form an ε-biased space with high probability. Explicit ε-biased spaces were constructed for the group $\mathbb{F}_2^n$ by Naor and Naor in [NN93]; further constructions were given by Alon et al. in [AGHP92]. Explicit constructions of ε-biased spaces in $\mathbb{Z}_d^n$ were given by Azar et al. in [AMN98]. We observe that this last construction yields a construction for all Abelian groups $\mathcal{G}^n$, when $\mathcal{G}$ is of constant size. We first state the result of [AMN98] in a form that we will find suitable.

Theorem 4.1. *For any fixed d, there is an* NC^2 *algorithm that does the following. On input n and $\varepsilon > 0$ (both in unary), the algorithm produces a symmetric multiset $S \subseteq \mathbb{Z}_d^n$ of size $O((\frac{n}{\varepsilon})^2)$ such that S is an ε-biased space in $\mathbb{Z}_d^n$.*

Proof. It is easy to see that the ε-biased space construction in [AMN98] can be implemented in deterministic logspace (and hence in NC^2). If the space S obtained is not symmetric, we can consider the multiset that is the disjoint union of S and S^{-1}, which is also easily seen to be ε-biased. ∎

Remark 4.2. We note that the definition of small bias spaces in [AMN98] differs somewhat from our own definition above. But it is easy to see that an ε-bias space in $\mathbb{Z}_d^n$ in the sense of [AMN98] is a $(d\varepsilon)$-bias space according to our definition above.

Remark 4.3. In a recent paper, Meka and Zuckerman [MZ09] observe, as we do below, that the construction of [AMN98] gives small bias spaces for any arbitrary Abelian group $\mathcal{G}$. Nevertheless, we present our own proof of this fact, since the small bias spaces that follow from our proof are of *smaller* size. Specifically, our proof shows how to explicitly construct sample spaces of size $O\left(\frac{n^2}{\varepsilon^2}\right)$, whereas the relevant result in [MZ09] only produces small bias spaces of size $O\left((\frac{n}{\varepsilon})^b\right)$, where b is some constant that depends on $\mathcal{G}$ (and can be as large as $\Omega(\log|\mathcal{G}|)$).

Lemma 4.4. *For any fixed group $\mathcal{G}$, there is an* NC^2 *algorithm which, on input n and $\varepsilon > 0$ in unary, produces a symmetric multiset $S \subseteq \mathcal{G}^n$ of size $O((\frac{n}{\varepsilon})^2)$ such that S is an ε-biased space in $\mathcal{G}^n$.*

Proof. By the Fundamental Theorem of finite Abelian groups, $\mathcal{G} \cong \mathbb{Z}_{d_1} \oplus \mathbb{Z}_{d_2} \oplus \cdots \oplus \mathbb{Z}_{d_k}$, for positive integers $d_1, d_2, \ldots, d_k$ such that $d_1 \mid d_2 \mid \cdots \mid d_k$. Let $\mathcal{G}_0$ denote $\mathbb{Z}_{d_k}^k$. Note that for any $s, t \in \mathbb{N}$, $\mathbb{Z}_s \cong \mathbb{Z}_{st}/\mathbb{Z}_t$. Hence, we see that that $\mathcal{G} \cong \mathcal{G}_0/\mathcal{H}$, where $\mathcal{H}$ is the subgroup $\mathbb{Z}_{e_1} \oplus \mathbb{Z}_{e_2} \oplus \cdots \oplus \mathbb{Z}_{e_k}$, and $e_i = d_k/d_i$ for each $i \in [k]$. Therefore, $\mathcal{G}^n \cong \mathcal{G}_0^n/\mathcal{H}^n$. Let $\pi : \mathcal{G}_0^n \to \mathcal{G}^n$ be the natural onto homomorphism with kernel $\mathcal{H}^n$. Note that π is just the projection map and can easily be computed in NC^2.

Since $\mathcal{G}_0^n \cong \mathbb{Z}_{d_k}^{nk}$, by Theorem 4.1, there is an NC^2 algorithm that constructs a symmetric multiset $S_0 \subseteq \mathcal{G}_0^n$ of size $O(\left(\frac{kn}{\varepsilon}\right)^2)$ such that S_0 is an ε-biased space in $\mathcal{G}_0^n$. We claim that the multiset $S = \pi(S_0)$ is a symmetric ε-biased space in $\mathcal{G}^n$. To see this, consider any non-trivial character χ of $\mathcal{G}^n$; note that $\chi_0 = \chi \circ \pi$ is a non-trivial character of $\mathcal{G}_0^n$. We have

$$\left| \mathop{\mathbf{E}}_{x \sim S}[\chi(x)] \right| = \left| \mathop{\mathbf{E}}_{x_0 \sim S_0}[\chi(\pi(x_0))] \right| = \left| \mathop{\mathbf{E}}_{x_0 \sim S_0}[\chi_0(x)] \right| \leq \varepsilon$$

where the first equality follows from the definition of S, and the last inequality follows from the fact that S_0 is an ε-biased space in $\mathcal{G}_0^n$. Since χ was an arbitrary non-trivial character of $\mathcal{G}^n$, we have proved that S is indeed an ε-biased space in $\mathcal{G}^n$. It is easy to see that S is symmetric. Finally, note that S can be computed in NC^2. This completes the proof. ∎

Finally, we mention a well-known connection between small bias spaces in $\mathcal{G}^n$ and Cayley graphs over $\mathcal{G}^n$ (e.g. see Alon and Roichman [AR94]).

Lemma 4.5. *Given any symmetric multiset $S \subseteq \mathcal{G}^n$, the Cayley graph $C(\mathcal{G}^n, S)$ is an $(|\mathcal{G}|^n, |S|, \alpha)$-graph iff S is an α-biased space.*

Lemmas 4.5 and 4.4 have the following easy consequence:

Lemma 4.6. *For any Abelian group $\mathcal{G}$, there is an* NC^2 *algorithm which, on unary inputs n and $\alpha > 0$, produces a symmetric multiset $S \subseteq \mathcal{G}^n$ of size $O((\frac{n}{\alpha})^2)$ such that $C(\mathcal{G}^n, S)$ is a $(|\mathcal{G}|^n, |S|, \alpha)$-graph.*

Putting the above statement together with the results of Section 3, we have the following.

Theorem 4.7. *For any constant $c > 0$, the RPP over $\mathcal{G}$ has an* NC^2 *$(n/2, c \log n)$-algorithm.*

Proof. Let $\mathcal{H}$ denote the input subgroup. By Lemma 3.3, there is a logspace (and hence NC^2) algorithm that computes a collection of $m = n^{O(c)}$ many subgroups $\mathcal{H}_1, \mathcal{H}_2, \ldots, \mathcal{H}_m$ such that $B(\mathcal{H}, c \log n) \subseteq \bigcup_{i=1}^m \mathcal{H}_i$ and $\delta(\mathcal{H}_i) \leq 2n/3$ for each $i \in [m]$. Now, fix any multiset $S \subseteq \mathcal{G}^n$ such that the Cayley graph $C(\mathcal{G}^n, S)$ is a $(|\mathcal{G}|^n, |S|, \alpha)$-graph, where $\alpha = \frac{1}{2n^{20c}}$; by Lemma 4.6, such an S can be constructed in NC^2. It follows from Lemma 3.4 that there is some $s \in S$ such that $s \notin \bigcup_{i=1}^m \mathcal{H}_i$. Finally, by Lemma 2.3, there is an NC^2 algorithm to test if each $s \in S$ belongs to $\mathcal{H}_i$, for any $i \in [m]$. Hence, we can find out (in parallel) exactly which $s \in S$ do not belong to any of the $\mathcal{H}_i$ and output one of them. The output element s is surely $c \log n$-far from $\mathcal{H}$. ∎

Let $\mathcal{G}$ be Abelian. We observe that a method of [APY09], coupled with Theorem 4.7, yields an efficient $(k, \frac{cn \log k}{k})$-algorithm for any constant $c > 0$, and $k \leq n/2$.

Theorem 4.8. *Let $c > 0$ be any constant. If $\mathcal{G}$ is an Abelian group, then the RPP over $\mathcal{G}$ has an* NC^2 *$(k, \frac{cn \log k}{k})$-algorithm for any $k \leq n/2$.*

Proof. Given as input a subgroup $\mathcal{H}$ such that $\delta(\mathcal{H}) = k \leq n/2$, the algorithm partitions $[n]$ as $[n] = \bigcup_{i=1}^{m} T_i$, where $2k \leq |T_i| < 4k$ for each i; note that $m \geq n/4k$. Let $\mathcal{H}_i$ denote the subgroup obtained when $\mathcal{H}$ is projected onto the coordinates in T_i. Since $\delta(\mathcal{H}_i) \leq k \leq |T_i|/2$, we can, by Theorem 4.7, efficiently find a point $x_i \in \mathcal{G}^{|T_i|}$ that is at least $4c \log k$-far from $\mathcal{H}_i$. Putting these x_i together in the natural way, we obtain an $x \in \mathcal{G}^n$ that is $\frac{cn \log k}{k}$-far from the subgroup $\mathcal{H}$.

Since $\mathcal{G}$ is Abelian, using the algorithm of Theorem 4.7, the x_i can all be computed in parallel in NC^2. Hence, the entire procedure can be performed in NC^2. ■

5. RPP over General Groups

Let $\mathcal{G}$ denote some fixed finite group. We can generalize the polynomial-time algorithm of [APY09], described for $\mathbb{F}_2$, to compute a point $x \in \mathcal{G}^n$ that is $c \log n$-far from a given input subgroup $\mathcal{H}$ such that $\delta(\mathcal{H}) \leq n/2$. We only state this result below and refer the interested reader to the full version [AS09b] for details.

Theorem 5.1. *For any constant $c > 0$, the RPP over $\mathcal{G}$ has a polynomial time $(n/2, c \log n)$-algorithm.*

Analogous to Theorem 4.8, we have the following solution to RPP for general groups.

Theorem 5.2. *Let $c > 0$ be any constant. For any $\mathcal{G}$, the RPP over $\mathcal{G}$ has a polynomial time $(k, \frac{cn \log k}{k})$-algorithm for any $k \leq n/2$.*

Proof. The construction is exactly the same as in the proof of Theorem 4.8. The only difference is that we will apply the algorithm of Theorem 5.1. In this case, the x_i can all be found in deterministic polynomial time. Hence, the entire procedure gives us a polynomial-time algorithm. ■

6. Limitations of expanding sets

In the previous sections, we have shown how generators for expanding Cayley graphs on $\mathcal{G}^n$, where $\mathcal{G}$ is a fixed finite group, can help solve the RPP over $\mathcal{G}$. In particular, we have the following easy consequence of Lemmas 3.3 and 3.4.

Corollary 6.1. *For any constant $c > 0$, large enough n, and any symmetric multiset $S \subseteq \mathcal{G}^n$ such that $\lambda(C(\mathcal{G}^n, S)) < \frac{1}{n^{20c}}$, the following holds. If $\mathcal{H}$ is any subgroup of $\mathcal{G}^n$ such that $\delta(\mathcal{H}) \leq n/2$, there is some $s \in S$ such that $s \notin B(\mathcal{H}, c \log n)$.*

It makes sense to ask if the parameters in Corollary 6.1 are far from optimal. Is it true that any polynomial-sized symmetric multiset $S \subseteq \mathcal{G}^n$ with good enough expansion properties is $\omega(\log n)$-far from every subgroup of dimension at most $n/2$? We can show that this is not true. Formally, we can prove:

Theorem 6.2. *For any constant $c > 0$ and large enough n, there is a symmetric multiset $S \subseteq \mathbb{F}_2^n$ such that $\lambda(C(\mathbb{F}_2^n, S)) \leq \frac{1}{n^c}$ but there is a subspace L of dimension $n/2$ such that $S \subseteq B(L, 20c \log n)$.*

It is well known that for any family of d-regular multigraphs G $\lambda(G) = \Omega(1/\sqrt{d})$ (see e.g. [HLW06, Theorem 5.3]). As a consequence of this lower bound it follows for any fixed group $\mathcal{G}$ and any multiset $S \subseteq \mathcal{G}^n$ that $\lambda(C(\mathcal{G}, S)) = \Omega(1/\sqrt{|S|})$. Hence, the above theorem tells us that just the expansion properties of $C(\mathbb{F}_2^n, S)$ for any poly(n)-sized S are not sufficient to guarantee $\omega(\log n)$-distance from every subspace of dimension $n/2$. The proof of the above statement can be found in the full version [AS09b].

7. Discussion

For the remote point problem over an Abelian group $\mathcal{G}$, we have shown how expanding generating sets for Cayley graphs of $\mathcal{G}^n$ can be used to obtain deterministic NC^2 algorithms. A natural question is whether we can obtain a similar algorithm for non-Abelian $\mathcal{G}$. Note that Lemma 3.4 holds in the non-Abelian setting too. Hence, in order to obtain an NC^2-algorithm for the RPP over arbitrary non-Abelian $\mathcal{G}$ along the lines of our algorithm for Abelian groups, we need to be able to check (in NC^2) for membership in $\mathcal{G}^n$, and we need to be able to construct small multisets S of $\mathcal{G}^n$ such that $C(\mathcal{G}^n, S)$ has sufficiently good expansion properties. Luks' work [Lu86] yields an NC^4 test for membership in $\mathcal{G}^n$ for arbitrary $\mathcal{G}$. Building on that, there is also an NC^2 membership test for $\mathcal{G}^n$ [AKV05]. However, we are unable to compute a (good enough) expanding generator set for the group $\mathcal{G}^n$ in deterministic NC or even in deterministic polynomial time.

Acknowledgements

We are grateful to Noga Alon and Sergey Yekhanin for interesting comments. In particular, Alon pointed out to us that Lemma 3.2 has an alternative proof using the expander mixing lemma. We thank the anonymous referees for their comments and suggestions.

References

[AGHP92] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple construction of almost k-wise independent random variables. *Random Struct. Algorithms*, 3(3):289–304, 1992.

[AKV05] V. Arvind, Piyush P. Kurur, T. C. Vijayaraghavan. Bounded Color Multiplicity Graph Isomorphism is in the #L Hierarchy. In *IEEE Conference on Computational Complexity 2005*: 13-27.

[APY09] Noga Alon, Rina Panigrahy, and Sergey Yekhanin. Deterministic approximation algorithms for the nearest codeword problem. In *APPROX-RANDOM*, pages 339–351, 2009.

[AR94] Noga Alon, Yuval Roichman. Random Cayley Graphs and Expanders. *Random Structures and Algorithms,* 5(2): 271-285 (1994).

[AS09a] V. Arvind and Srikanth Srinivasan. Circuit Complexity, Help Functions and the Remote point problem. manuscript.

[AS09b] V. Arvind and Srikanth Srinivasan. The Remote Point Problem, Small Bias Spaces, and Expanding Generator Sets ECCC Report TR09-105. Can be found at `http://eccc.hpi-web.de/report/2009/105/`

[AMN98] Yossi Azar, Rajeev Motwani, and Joseph Naor. Approximating probability distributions using small sample spaces. *Combinatorica*, 18(2):151–171, 1998.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S)*, 43:439–561, 2006.

[Lu86] Eugene M. Luks. Parallel algorithms for permutation groups and graph isomorphism. In *FOCS*, pages 292–302, 1986.

[Lu93] Eugene M. Luks. Permutation groups and polynomial time computation. *Groups and Computation I,* DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol 11, 139-174, 1993.

[MC87] Pierre McKenzie and Stephen Cook. The parallel complexity of Abelian permutation group problems. *SIAM Journal on Computing*, 16(5):880-909, 1987.

[MZ09] Raghu Meka and David Zuckerman. Small-Bias Spaces for Group Products. *APPROX-RANDOM 2009*: 658-672.

[NN93] Joseph Naor and Moni Naor. Small-bias probability spaces: Efficient constructions and applications. *SIAM J. Comput.*, 22(4):838–856, 1993.

[Rag88] Prabhakar Raghavan. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *Journal of Computer and System Sciences*, 37(2):130 – 143, 1988.

[Rei08] Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4), 2008.

[V77] Leslie G. Valiant. Graph-Theoretic Arguments in Low-Level Complexity. *Proceedings Mathematical Foundations of Computer Science,* LNCS vol. 53: 162-176, Springer 1977.