

Non-Abelian Analogs of Lattice Rounding

Evgeni Begelfor
Department of Computer Science
The Hebrew University of Jerusalem
`begelfor@gmail.com`

Stephen D. Miller*
Department of Mathematics
Rutgers University
`millier@math.rutgers.edu`

Ramarathnam Venkatesan
Microsoft Research
Redmond, WA and Bangalore, India
`venkie@microsoft.com`

January 12, 2015

Abstract

Lattice rounding in Euclidean space can be viewed as finding the nearest point in the orbit of an action by a discrete group, relative to the norm inherited from the ambient space. Using this point of view, we initiate the study of non-abelian analogs of lattice rounding involving matrix groups. In one direction, we give an algorithm for solving a normed word problem when the inputs are random products over a basis set, and give theoretical justification for its success. In another direction, we prove a general inapproximability result which

*Supported by NSF grant DMS-1201362.

essentially rules out *strong approximation algorithms* (i.e., whose approximation factors depend only on dimension) analogous to LLL in the general case.

Keywords: lattice rounding, matrix groups, norm concentration, Lyapunov exponents, word problems, inapproximability.

1 Introduction

Given a basis $\{a_i\}_{i=1}^n$ of a lattice $L \subset \mathbb{R}^n$ and a vector $y \in \mathbb{R}^n$, the Lattice Rounding Problem (LRP) in Euclidean space asks to find $\arg \min_{z \in L} \|z - y\|_2$, that is, a vector $z \in L$ nearest to y . This problem is very closely related to the lattice basis reduction problem of finding a good basis for L , which informally is to find another basis $\{b_i\}_{i=1}^n$ for L whose elements are as orthogonal as possible. The motivation is that given such a good basis $\{b_i\}_{i=1}^n$, LRP may be easy. To wit, if $L = \mathbb{Z}^n$ a good basis is trivial to find, and LRP can be solved by coordinate-wise rounding. For general L and bases $\{a_i\}_{i=1}^n$, one has NP-hardness results for exact and approximate versions of LRP [1, 4], and their study is an active area of research.

The presumed hardness of these problems also has led to constructions of cryptosystems. This typically involves three main ingredients:

- (a) **Good Basis.** Generation of a basis $\{b_i\}_{i=1}^n$ for L that is good in the sense that LRP is easy relative to it on inputs randomly chosen from some distribution ν .
- (b) **Bad Basis.** Generation of a suitable matrix $M \in SL_n(\mathbb{Z})$ such that LRP with respect to ν is hard relative to the basis $\{a_i\}_{i=1}^n$, where $a_i = Mb_i$.
- (c) **Public Key System.** One keeps the good basis as the private key and the bad basis as a public key, and designs an encryption or signature scheme such that an attack on it would entail solving LRP relative to a bad basis.

This paper presents a non-abelian generalization of lattice rounding, and some steps in the direction of ingredients (a) and (b). Our generalization starts with the viewpoint of \mathbb{R}^n as an additive abelian group and L as a

discrete subgroup: LRP is equivalent to finding the nearest point to z (in the ambient metric) to the orbit of the origin under the action of L . This viewpoint can be extended to a larger class of groups, and spaces upon which they act. For example, one could consider a Lie group such as the $n \times n$ invertible matrices $G = GL_n(\mathbb{R})$, and a discrete subgroup Γ ; this direction quickly leads to rich mathematical theory connected with dynamics and automorphic forms. In this case one could choose ambient metrics on G related to a variety of matrix norms.

Another direction is to consider the action of G on some space X endowed with its own metric. For example, $G = GL_n(\mathbb{R})$ acts on the vector space $X = \mathbb{R}^n$ or even the projective space $\mathbb{R}P^{n-1}$ by the usual multiplication of vectors by matrices. Let Γ as before denote a subgroup of G . A non-abelian analog of lattice rounding asks to find the closest point in the Γ -orbit of a fixed vector in \mathbb{R}^n , where the closeness is measured using some natural metric on vectors (but not on matrices, although we do make a restriction on word length for practical reasons).

Alternatively, if Γ and X are themselves endowed with a discrete structure (e.g., Γ consists of integral matrices and X consists of integral vectors), we can instead study the problem of recognizing elements of a Γ -orbit. To address items (a) and (b) above, is natural to ask if one can develop analogous positive algorithms for rounding with good bases and, conversely, negative results for general subgroups Γ in $GL_n(\mathbb{R})$. One naive approach would be to modify a generating set $\{g_1, \dots, g_r\}$ by successively replacing a generator g_i by $g_i g_j^c$, where $j \neq i$ and $c \in \mathbb{Z}$. In the abelian case such repeated modifications generate any change of lattice basis. However, in the non-abelian case there are some geometric constraints (such as coarse quasi-isometry) which may at times dull the effects of such a change. We do not investigate this direction here.

In Section 3 we consider the Word Problem on Vectors (3.3), for which we propose the Norm Reduction Algorithm (3.4). The analysis of the latter leads to well-studied mathematical and algorithmic topics. For example, multiplying random elements of Γ times a fixed vector can be viewed as a *generalized Markov chain* (using more than one matrix); the growing vector norms of these products is itself a generalization of the law of large numbers (the case of $n = 1$). Additionally, the conditions for the success of our Norm Reduction Algorithm depend on an analog of the spectral or norm gap in Markov chains: it requires instead a gap between *Lyapunov exponents* (see (4.8)).

Some remarks on our generalization

The generalization of LRP from lattices L in \mathbb{R}^n to finitely-generated subgroups $\Gamma = \langle S \rangle$ in $GL_n(\mathbb{R})$ is neither unique nor straightforward. Here we seek to make a distinction between our norms and the word-length metric, since the latter already appears in the existing literature in combinatorial group theory and the study of special groups (e.g., braid groups [8]) from algorithmic and cryptographic points of view. We informally outline a few issues that guide our formulation.

Full (or at least large) dimensionality: We would like our discrete subgroups to not be contained inside some subgroup of much smaller dimension of the ambient group. In \mathbb{R}^n one typically assumes the lattice L has full rank, or least has relatively large rank. Its natural matrix analogue is to require the *Zariski closure* of $\Gamma = \langle S \rangle$ be the full group (or at least correspond to a subgroup having a significant fraction of the dimension of the full group). By definition, this means that the full group is the only group containing S which can be defined as the common zeroes of a set of polynomial equations. This ensures $\Gamma = \langle S \rangle$ is non-abelian in as general way as possible.

For example, if S has only diagonal matrices it cannot generate any non-abelian group, and its Zariski closure is at most an n -dimensional subgroup of the n^2 -dimensional group $GL_n(\mathbb{R})$. In fact, by considering commuting diagonal matrices one can embed subset-sum type problems and get NP-hardness results. Note that matrices composed of 2×2 blocks along the diagonal can generate non-abelian groups that essentially describe simultaneous problems in dimension 2; nevertheless, the Post Correspondence Problem can be embedded as a word problem over 4×4 matrices with 2×2 blocks, proving the undecidability of the latter [16]. However, certain problems can actually become easier in the non-abelian setting: for example, finding the order of a random element in S_n is much easier than in $(\mathbb{Z}/p\mathbb{Z})^*$.

Metrics: The distinction between the word length metric and ambient matrix norm is discussed in some detail in Section 2 below. The former depends on the generating set S . In general these can be very different notions of distance, which makes our study difficult – yet is key to potential cryptographic applications. We use the Furstenberg-Kesten theory [6, 7, 13] of random matrix products to correlate the two (in a probabilistic sense) in certain situations, which is analogous to the “good basis” situation described in (a) above.

Finite co-volume and compactness If L has full rank, then $L \backslash \mathbb{R}^n$

is a compact, finite-volume quotient. However, neither property necessarily extends to the quotients $\Gamma \backslash G$ in many important examples of Γ and G . Thus we do not impose this requirement. Some further comments are given just below in the beginning of the following section.

Outline of this paper

Section 2 contains some background about different metrics on Lie groups and their discrete subgroups. Section 3 introduces the statements of the word problems that motivate our results, as well as the Norm Reduction Algorithm (3.4), which is rigorously analyzed in Theorem 4.1. The Closest Group Element Problem is also given in section 3, along with the statement of its inapproximability result Theorem 3.1. The analysis of the Norm Reduction Algorithm is performed in Section 4 using results in dynamical systems. Some experimental results on the algorithm are also presented in Section 4.5. The proof of Theorem 3.1 is given in Section 5; it demonstrates a polynomial time reduction from the Traveling Salesman Problem.

We would like to thank Anthony Bloch, Hillel Furstenberg, Nathan Keller, Peter Sarnak, Adi Shamir, Boaz Tsaban, and Akshay Venkatesh for their helpful comments.

2 Background

Just as a lattice $L = \langle a_1, \dots, a_n \rangle$ is additively generated by its basis $\{a_i\}$, the subgroups $\Gamma = \langle g_1, \dots, g_k \rangle$ we consider will be finitely generated. A crucial difference, however, is that the quotient of \mathbb{R}^n by L is a compact n -dimensional torus with finite volume under the usual Lebesgue measure on \mathbb{R}^n (for example, the quotient $\mathbb{Z}^n \backslash \mathbb{R}^n$). However, this fails to be true for nice examples such as $GL_n(\mathbb{Z}) \backslash GL_n(\mathbb{R})$ or even $SL_n(\mathbb{Z}) \backslash SL_n(\mathbb{R})$, both of which are noncompact under the natural group invariant metric inherited from G (the latter quotient, however, does have finite volume). The theory and construction of both compact and noncompact discrete subgroups of Lie groups involves numerous beautiful subtleties (see [15, 24]); we do not restrict ourselves to these objects in this paper.

There are two natural notions of size in Γ , and by extension to the Γ -orbit of any basepoint $x \in X$:

1. *Word length metric:* If $S = \{g_1, \dots, g_k\}$ is a generating set of Γ as above, any element $w \in \Gamma$ can be expressed as a finite word in the alphabet $S \cup S^{-1}$. There may be many possibilities for such a word, taking into account relations amongst the g_i (including the trivial relation $g_i g_i^{-1} = 1$). The minimal such length among all such expressions is the *word length of w with respect to S* .

The ability to efficiently compute the word length of w enables one to efficiently write it as a minimal length word, simply by successively checking which of the expressions $g_i^{\pm 1} w$ reduces the word length by one. Finding the word length depends of course on the generating set S , which is analogous to the basis of a lattice. In analogy with ingredients (a), (b), and (c) above for Euclidean lattices, we want the word length to be difficult for typical generating sets S of Γ , yet at the same time easy for some “good bases” S ; moreover, we would like to be able to transform each “good base” into a seemingly bad one.

2. *Inherited metric:* Fundamental to lattice reduction and rounding is the notion of metric on the ambient space. Natural metrics on G and X therefore can be used to give generalizations of lattice rounding. Combining this with word length results in problems such as the following: given $\ell \in \mathbb{N}$, $\Gamma \subset GL_n(\mathbb{R})$, and vectors y and $z \in \mathbb{R}^n$, find $\gamma \in \Gamma$ such that $\|gy - z\|_2$ is minimized over all $\gamma \in \Gamma$ with word length at most ℓ . Thus the length parameter ℓ is used to complement (rather than to duplicate) the ambient metric.

Though we do not present any cryptographic systems here, generalizations of attacks on existing cryptosystems motivate studying rounding problems in more general settings than lattices in \mathbb{R}^n alone. With some performance enhancing additions, the lattice reduction algorithm LLL [12] has long become a valuable tool in cryptanalysis [11], and typically is more effective than the provable guarantees attached to it indicate alone. Starting with the original attack of Shamir [21], some very effective attacks have been discovered. The attacks are often based on the Shortest Vector Problem in lattices: given a basis for L , find a nonzero vector in L with minimal norm. In polynomial time, the LLL algorithm finds a vector within a factor of $2^{n/4}$ of being the shortest, a *strong bound* – i.e., one which depends only on the dimension of the lattice, and not on the sizes of the entries in the lattice basis themselves. Babai’s rounding algorithm [2] – which is based on LLL

– also has this feature for solving lattice rounding problems in Euclidean space. The fact that this bound depends only on the dimension is crucial for attacks.

In contrast, we prove in Theorem 3.1 that the analogous question of rounding products of matrices cannot have a polynomial time strong approximation algorithm¹ – unless P=NP. This is done by creating a polynomial time reduction to the Traveling Salesman Problem, which has a similar inapproximability result. Thus a strong approximation algorithm like LLL for rounding in matrix groups is unlikely to exist.

3 Some non-abelian problems and an algorithm

We study problems that arise out of group actions on normed spaces, where we are concerned with the action of group elements that have short expressions relative to a given basis or generating set. We now proceed to formally define these problems and state some known results.

We shall work with $GL_d(\mathbb{R})$, the group of all invertible $d \times d$ real matrices, and often with subsets that have integer entries. Given $g_1, \dots, g_k \in GL_d(\mathbb{R})$, we consider the possible products of these matrices up to a certain length bound, and whether or not they can be recognized as such. The word problem is the algorithmic task of representing a given matrix in this semigroup as a product of the generators:

<p>Word Problem <u>INPUT:</u> Matrices g_1, \dots, g_k and $x \in GL_d(\mathbb{R})$. <u>OUTPUT:</u> An integer $\ell > 0$ and indices $1 \leq s_1, \dots, s_\ell \leq k$ such that $g_{s_1}g_{s_2} \cdots g_{s_\ell} = x$, if such a solution exists.</p>	(3.1)
--	-------

This word problem is known to be unsolvable when $d \geq 4$ [17]; however, there is an algorithm for specifically constructed generators when $d = 2$ [9] (the case of $d = 3$ is open). It becomes NP-hard for $d \geq 4$ if we bound the

¹where the approximating factor is a polynomial time computable function of the dimension.

word length ℓ , as we do for all our problems in the rest of the paper:

<p>Bounded Word Problem</p> <p><u>INPUT</u>: An integer $L > 0$, and matrices g_1, \dots, g_k and $x \in GL_d(\mathbb{R})$.</p> <p><u>OUTPUT</u>: Indices $1 \leq s_1, \dots, s_L \leq k$ such that $g_{s_1}g_{s_2} \cdots g_{s_L} = x$, if such a solution exist.</p>	(3.2)
---	-------

This problem can be modified to allow for words of length $\leq L$.

We now define another related problem, in which the matrices act on vectors:

<p>Word Problem on Vectors.</p> <p><u>INPUT</u>: An integer $L > 0$, matrices $g_1, \dots, g_k \in GL_d(\mathbb{R})$ with integer entries, and nonzero vectors $v, w \in \mathbb{Z}^d$.</p> <p><u>OUTPUT</u>: An integer $\ell \leq L$ and indices $1 \leq s_1, \dots, s_\ell \leq k$ such that $g_{s_1}g_{s_2} \cdots g_{s_\ell}v = w$, if such a solution exists.</p>	(3.3)
---	-------

Typically we are interested in instances where $\ell = L$ and the indices s_j are chosen independently and uniformly at random from the above interval. Using the ambient norm on Euclidean space, we present the following algorithm for this problem:

<p>Norm Reduction Algorithm:</p> <p>Let $j = 0$, and t be a fixed parameter.</p> <p>repeat</p> <p style="padding-left: 20px;">$j = j + 1$</p> <p style="padding-left: 20px;">$s_j = \arg \min_i \ g_i^{-1}w\$</p> <p style="padding-left: 20px;">$w = g_{s_j}^{-1}w$</p> <p>until $w = v$ or $j = L - t$.</p> <p>Solve for s_{L-t+1}, \dots, s_L by exhaustive search.</p>	(3.4)
--	-------

We include the option of exhaustive search for the final t steps in case the algorithm performs worse on smaller words than on larger ones. Another possibility is to use a memory-length look-ahead algorithm such as in [19, §7]. The Norm Reduction Algorithm is rigorously analyzed in the next section, where it is related to a maximal likelihood algorithm. Its success depends

on some mild yet complicated conditions on generators g_1, \dots, g_k that come from dynamics. Theorem 4.1 in the next section gives a rigorous upper bound on the error probability of this algorithm. We give a successful numerical example in Table 1 in Section 4.5, along with how Theorem 4.1's constants pertain to it.

One can also define a related rounding problem, whose analysis and algorithms are quite similar. Instead, we will focus on the following *matrix rounding* question: finding a short word in a semigroup closest to a given one (with an length constraint imposed for practical reasons).

<p>Closest Group Element Problem (CGEP) <u>INPUT</u>: A positive integer $L > 0$, and matrices g_1, \dots, g_k and $z \in GL_d(\mathbb{R})$. <u>OUTPUT</u>: The closest word of length $\leq L$ in the g_i to z.</p>	(3.5)
--	-------

Though the problem can be stated for various notions of distance, we will use the sum-of-squares matrix distance

$$\left\| (a_{ij}) - (b_{ij}) \right\|^2 = \sum_{i,j} |a_{ij} - b_{ij}|^2 \quad (3.6)$$

in studying this problem.

Our main result about the CGEP problem is the following negative result, which comes close to ruling out the existence of an algorithm such as LLL that approximates the closest element up to a constant factor depending only on the dimension. In the following we denote by $CGEP(g_1, \dots, g_k, z, L)$ the solution to the CGEP problem as above.

Theorem 3.1. *Let $f : \mathbb{Z}_{>0} \rightarrow [1, \infty)$ be a polynomial time computable function. If there exists a polynomial time algorithm A which, given the input of a CGEP problem as in (3.5), always outputs a word w' of length $\leq L$ in the g_i such that*

$$\|w' - z\| \leq f(d) \|CGEP(g_1, \dots, g_k, z, L) - z\|, \quad (3.7)$$

then $P = NP$.

It is an interesting open problem whether or not the approximation factor can instead depend on the sizes of the entries.

4 Maximum Likelihood Algorithms

In this section we give and analyze a simple algorithm to solve the Word Problem on Vectors (3.3): try to reduce the norm at each step, or put differently, attempt to use the norm as a proxy for word length. This involves studying some background from dynamics related to random products of matrices, first studied by Furstenberg and Kesten [6, 7]. Our results are sensitive to certain conditions related to the generators, which we describe before stating our result. These are discussed thoroughly in the book [13], which serves as a general reference for background material on the topic of this section. In addition, several of the techniques and arguments in this section are taken from [13].

Let $S = \{g_1, \dots, g_k\}$ denote a finite subset of $G = GL_d(\mathbb{R})$, and $T = \langle S \rangle$ the semigroup it generates. Throughout this section we will use $\|g\|$ to denote the operator norm of a matrix g . We make the following standing assumptions on the set S throughout this section:

- A1.** T is *contracting* in the sense of [3, Definition III.1.3]. This means that T has a sequence of matrices M_1, M_2, \dots such that $M_n/\|M_n\|$ converges to a rank 1 matrix. It is readily seen (using Jordan canonical form) that this condition holds automatically if S (or even T) contains a matrix with an eigenvalue strictly larger than its others in modulus.
- A2.** T is *strongly irreducible*: there is no finite union of proper vector subspaces of \mathbb{R}^d which is stabilized by each element of T . Equivalently, the same statement holds with T replaced by the *group* generated by S ([3, p. 48]).
- A3.** The operator norms $\|g_j^{-1}g_i\|$, $j \neq i$, are all at least some constant $N > 1$.

We prove the following result about the probability of success of the Norm Reduction Algorithm (3.4). This gives a strong indication (along with numerical testing) that norm reduction is a suitable algorithm for solving the Word Problem on Vectors (3.3). It is also often possible to show that the group generated by S is free by deriving a quantitative version of the well-known Ping-Pong Lemma. We do not address these issues in this version of the paper.

Theorem 4.1. *Let $S = \{g_1, \dots, g_k\}$ be a fixed subset of $GL_d(\mathbb{R})$ and v a fixed nonzero vector in \mathbb{Z}^d . Assume properties **A1-3**. Then there exists positive quantities α , B , and C such that if h is a random product² of length L elements of S , the Norm Reduction Algorithm (3.4) recovers v from hv (i.e., solves the Word Problem on Vectors (3.3)) with probability at least*

$$1 - C(L - t)(|S| - 1)N^{-\alpha},$$

where N is as defined in assumption **A3** and the parameter t in the algorithm is taken to be at least $B \log N$.

Roughly speaking, the algorithm succeeds for long enough words when the operator norms $\|g_j^{-1}g_i\|$ are themselves sufficiently large. Though the constant N is readily computable from the generating set S , the numerical values of C and α are unfortunately more subtle. We are unable to rigorously prove that C is reasonably small, or that α is somewhat large. (It is not clear that these statistics of $\langle S \rangle$ are even computable in general; see [5, 20, 22].) In particular, one cannot directly take $N \rightarrow \infty$ to get the above error estimate to decay to zero, without possibly simultaneously affecting α . However, in concrete examples of generating sets it is possible to make heuristic estimates of the values of N and α from the proof. We give such an example in Section 4.5, in which numerical estimates for these constants give a small error probability in Theorem 4.1. Our experiments on this example are vastly better: the algorithm was successful in nearly all trials we tested for $L \leq 1000$ (see Table 1).

4.1 Motivation for the algorithm and its analysis

Recall the Word Problem on Vectors (3.3), in which the matrices in S are assumed to be integral. One is given $L \in \mathbb{N}$ and vectors v and $w = hv \in \mathbb{Z}^d$, where h is an unknown word of length at most L in S ; the problem is to find some word h' of length at most L in S such that $w = h'v$. Were we to have a concrete description of ν as a product $f\lambda$, where f is an easily computable function, we could attempt to solve for h using the following maximum likelihood algorithm:

²I.e., $h = g_{i_1} \cdots g_{i_L}$ where i_1, \dots, i_L are each chosen independently and uniformly from $\{1, \dots, k\}$.

Idealized Algorithm:Let $j = 0$ **repeat** $j = j + 1$ $s_j = \arg \max_i \frac{f(g_i^{-1}w)}{\|g_i^{-1}w\|^d |\det g_i|}$ $w = g_{s_j}^{-1}w$ **until** $w = v$ or $j = L$.

Recall the notation $\arg \max_i$ denotes a value of i which maximizes the expression it precedes. The particular expression here represents the change in local density under the map $w \mapsto g_i^{-1}w$. The numerator accounts for the difference between ν and λ , while the denominator represents the change in the uniform measure λ . If successful, the algorithm produces h' as $g_{s_1}g_{s_2}\cdots$, possibly reconstructing h . However, it is impractical to assume that f is easily computable. Because of this limitation, we instead use the simpler, more practical Norm Reduction Algorithm (3.4). It is tantamount to pretending f equals 1 and that the matrices have determinant 1, meaning that we seek to minimize $\|g_j^{-1}w\|$ at each stage.

In effect, the Norm Reduction Algorithm (3.4) uses the norm as a height function, and proceeds by descent to shorten the word length of h each time. Of course, a direct way to measure the word length would be preferable. The relationship between word length and matrix norm has been studied by several authors, e.g., [10, 14].

To study the distribution of elements of T and their orbits in \mathbb{R}^d , we need to define some measures. We let $\mu = \mu_S$ denote the Dirac measure of S on G , meaning that it gives mass $\frac{1}{|S|}$ to each element. Given two measures μ_1, μ_2 on G , their convolution is defined as the unique measure $\mu_1 * \mu_2$ satisfying

$$\int_G f(x) d\mu_1 * \mu_2(x) = \int_G \int_G f(xy) d\mu_1(x) d\mu_2(y) \quad \text{for all } f \in C(G), \quad (4.1)$$

the continuous functions on G . To simplify notation we sometimes write $\mu_1 * \mu_2$ simply as $\mu_1\mu_2$; for example the n -fold convolution of μ with itself will be denoted as μ^n (it is the measure giving mass $|S|^{-n}$ to each product of n elements taken from S , allowing repetitions). We can also define the convolution of μ with any measure ρ on \mathbb{RP}^{d-1} : $\mu * \rho$ is the unique measure

satisfying

$$\int_{\mathbb{RP}^{d-1}} f(x) d\mu * \rho(x) = \int_G \int_{\mathbb{RP}^{d-1}} f(Mx) d\rho(x) d\mu(M) \quad \text{for all } f \in C(\mathbb{RP}^{d-1}). \quad (4.2)$$

To be concrete, we identify measures on \mathbb{RP}^{d-1} with measures on the unit sphere in \mathbb{R}^d that are invariant under the antipodal map. Typically the uniform measure λ on \mathbb{RP}^{d-1} is not stabilized by convolution with μ , unless the matrices in S are orthogonal. However, there exist measures ν on \mathbb{RP}^{d-1} which are μ -invariant:

$$\mu * \nu = \nu \quad (4.3)$$

(see [6, Lemma 1.2]). Under certain conditions more can be said about ν , such as its regularity properties. This measure is not always uniquely determined by S , but assumptions **A1** and **A2** however guarantee the uniqueness of the μ -invariant measure in our setting (see [3, Theorem III.4.3.(iii)]).

The main step in the proof of Theorem 4.1 involves estimating measures of the subsets of vectors in \mathbb{RP}^{d-1} which get contracted by the operators $g_j^{-1}g_i$. Indeed, let p_j equal the probability that the algorithm obtains the wrong value for g_{s_j} at the j -th step. One has that $p_j = \frac{1}{k} \sum_{i \leq k} p_{ji}$, where p_{ji} is the probability of error in the j -th step, conditioned on the correct answer equaling g_i . In terms of the measure δ_v , the Dirac measure of $v \in \mathbb{RP}^{d-1}$, this probability can be computed as

$$p_{ji} = \mu^{j-1} \delta_v(B_i), \quad (4.4)$$

where $\mu^{j-1} \delta_v$ denotes $\mu^{j-1} * \delta_v$ and

$$\begin{aligned} B_i &= \{x \in \mathbb{RP}^{d-1} \mid \exists r \neq i \text{ such that } \|g_r^{-1}g_i x\| < \|x\|\} \\ &= \cup_{r \neq i} B_{r,i}, \end{aligned} \quad (4.5)$$

with

$$B_{r,i} = \{x \in \mathbb{RP}^{d-1} \mid \|g_r^{-1}g_i x\| < \|x\|\}. \quad (4.6)$$

Thus the error probability in Theorem 4.1 is

$$\text{Prob}_{\text{Error}} \leq \sum_{j=t+1}^L p_j = \frac{1}{k} \sum_{\substack{t < j \leq L \\ 1 \leq i \leq k}} \mu^{j-1} \delta_v(B_i) \leq \frac{1}{k} \sum_{\substack{t < j \leq L \\ 1 \leq r \neq i \leq k}} \mu^{j-1} \delta_v(B_{r,i}). \quad (4.7)$$

The proof therefore amounts to estimates on $\mu^{j-1} \delta_v(B_{r,i})$, which are given in the following subsections.

4.2 Lyapunov Exponents

In the remainder of this section, we shall need some technical results and concepts from the literature on random products of matrices. For the reader's convenience we have chosen to cite background results in the book [3] wherever possible, while at the same time attempting to correctly attribute the original source of the results. The top two Lyapunov exponents γ_1, γ_2 of S are defined through the following limits (see [3, p. 6]):

$$\begin{aligned}\gamma_1 &= \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}\{\log \|h\| \mid h \in S^n\} &= \frac{1}{n} \int_G \log \|M\| d\mu^n \\ \gamma_1 + \gamma_2 &= \lim_{n \rightarrow \infty} \frac{1}{n} \mathbb{E}\{\log \|\wedge^2 h\| \mid h \in S^n\} &= \frac{1}{n} \int_G \log \|\wedge^2 M\| d\mu^n,\end{aligned}\tag{4.8}$$

where $\wedge^2 g$ is the operator on $\wedge^2 \mathbb{R}^d$ given by $x \wedge y \mapsto gx \wedge gy$ and $\|\cdot\|$ denotes the operator norm (the general Lyapunov exponents are likewise defined inductively through higher exterior powers). Not only do these limits exist, but in fact a theorem of Furstenberg and Kesten [7] asserts that the individual terms in the above sets are close to those limits with probability one as $n \rightarrow \infty$. Under assumptions **A1** and **A2** one has separation between these top two Lyapunov exponents:

$$\gamma_1 > \gamma_2 \tag{4.9}$$

([3, Theorem III.6.1]). We remark that computing or even approximating the Lyapunov exponents is in general difficult [22].

We shall use the following variant of (4.8), which involves the action of a random product on \mathbb{RP}^{d-1} .

Proposition 4.2. (*Furstenberg* [6]; see [3, Corollary III.3.4.(iii)].) *Under assumption **A2** one has that*

$$\frac{1}{n} \mathbb{E}\{\log \|hx\| \mid h \in S^n\} = \frac{1}{n} \int_G \log \left(\frac{\|Mx\|}{\|x\|} \right) d\mu^n \longrightarrow \gamma_1 \tag{4.10}$$

uniformly for $x \in \mathbb{RP}^{d-1}$.

Consequently,

$$\limsup_{n \rightarrow \infty} \sup_{x \neq 0} \frac{1}{n} \int_G \log \left(\frac{\|Mx\|}{\|x\|} \right) d\mu^n = \gamma_1. \tag{4.11}$$

Following [3, p. 55] we use the natural angular distance

$$\delta(x, y) = \frac{\|x \wedge y\|}{\|x\| \|y\|} = \sqrt{1 - \frac{\langle x, y \rangle^2}{\|x\|^2 \|y\|^2}}, \quad (4.12)$$

which is a metric on \mathbb{RP}^{d-1} . It satisfies the following estimate:

Proposition 4.3. (See [3, Proposition III.6.4(ii)].) For any $x, y \in \mathbb{RP}^{d-1}$,

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \int_G \log \left(\frac{\delta(Mx, My)}{\delta(x, y)} \right) d\mu^n(M) \leq \gamma_2 - \gamma_1 < 0. \quad (4.13)$$

Proof. By (4.12)

$$\begin{aligned} \frac{\delta(Mx, My)}{\delta(x, y)} &= \frac{\|M(x \wedge y)\|}{\|x \wedge y\|} \frac{\|x\|}{\|Mx\|} \frac{\|y\|}{\|My\|} \\ &\leq \|\wedge^2 M\| \frac{\|x\|}{\|Mx\|} \frac{\|y\|}{\|My\|}, \\ \frac{1}{n} \log \frac{\delta(Mx, My)}{\delta(x, y)} &\leq \frac{1}{n} \log \|\wedge^2 M\| - \frac{1}{n} \log \frac{\|Mx\|}{\|x\|} - \frac{1}{n} \log \frac{\|My\|}{\|y\|}. \end{aligned} \quad (4.14)$$

The proposition follows by integrating this inequality over M , and appealing to (4.8) and (4.10). \square

4.3 Cocycle integrals

We have just seen that the integrand

$$s(M, (x, y)) = \log \frac{\delta(Mx, My)}{\delta(x, y)} \quad (4.15)$$

in (4.13) tends to be negative on S^n . Our next goal is to show that the integral of an exponential of it is accordingly smaller than 1. Writing z as shorthand for (x, y) , define

$$\mathcal{S}(n) = \sup_z \int_G e^{\alpha s(M, z)} d\mu^n(M), \quad (4.16)$$

which exists for any $\alpha > 0$ since S is finite. It is proven in [3, p. 104] that

$$\mathcal{S}(n + m) \leq \mathcal{S}(n) \mathcal{S}(m), \quad (4.17)$$

using the cocycle identity

$$s(g_1 g_2, v) = s(g_1, g_2 v) + s(g_2, v) \quad (4.18)$$

and a simple change of variables. According to [3, Lemma III.5.4], any matrix $M \in G$ satisfies the inequality

$$\left| \log \|\wedge^2 M\| \right| \leq 2 \ell(M), \quad (4.19)$$

where

$$\ell(M) = \max\{\log \|M\|, \log \|M^{-1}\|, 0\}. \quad (4.20)$$

It follows from (4.14) that

$$s(M, z) \leq \log \|\wedge^2 M\| + 2 \log \|M^{-1}\| \leq 4 \ell(M). \quad (4.21)$$

If ℓ_{max} denotes $\max\{\ell(g) | g \in S\}$, then

$$s(M, z) \leq 4 n \ell_{max} \quad (4.22)$$

on $S^n =$ the support of μ^n , independently of z .

Proposition 4.4. (See [13, Theorem 1] and [3, Proposition V.2.3].) *For $\alpha > 0$ sufficiently small, there exists $n_0 > 0$ and $\rho < 1$ such that*

$$\int_G \left(\frac{\delta(Mx, My)}{\delta(x, y)} \right)^\alpha d\mu^n(M) \leq \rho^n \quad (4.23)$$

for all $x \neq y \in \mathbb{RP}^{d-1}$, and $n \geq n_0$.

Proof. The inequality

$$e^x \leq 1 + x + \frac{x^2}{2} e^{|x|}$$

and (4.22) imply that

$$e^{\alpha s(M, z)} \leq 1 + \alpha s(M, z) + 8 \alpha^2 n^2 \ell_{max}^2 e^{4n\alpha \ell_{max}} \quad (4.24)$$

for $M \in S^n$. Thus the lefthand side of (4.23), which is the integral of $e^{\alpha s(M, z)} d\mu^n(M)$ over G , is bounded by

$$1 + \alpha \int_G s(M, z) d\mu^n(M) + 8 \alpha^2 n^2 \ell_{max}^2 e^{4n\alpha \ell_{max}}. \quad (4.25)$$

Proposition 4.3 asserts that for any $\varepsilon > 0$ there exists n' sufficiently large so that

$$\sup_z \int_G s(M, z) d\mu^n(M) \leq n(\gamma_2 - \gamma_1 + \varepsilon) \quad (4.26)$$

for all $n \geq n'$, and so

$$\mathcal{S}(n) \leq 1 + n\alpha(\gamma_2 - \gamma_1 + \varepsilon) + 8\alpha^2 n^2 \ell_{max}^2 e^{4n\alpha\ell_{max}} \quad (4.27)$$

for such n . In particular, if ε and α are sufficiently small, the righthand side of (4.26) is negative and $\mathcal{S}(n') < 1$. Repeated applications of the subadditivity property (4.17) show that $\mathcal{S}(kn' + m) \leq \mathcal{S}(n')^k \mathcal{S}(m)$ for $1 \leq m \leq n'$, which implies the proposition. \square

4.4 Estimate on $\mu^{j-1}\delta_v(B_{r,i})$

This subsection contains the mathematical core of the argument, a Hölder estimate relating the measures $\mu^{j-1}\delta_v$ and ν . For any $\varepsilon > 0$ and closed subset $U \subset \mathbb{RP}^{d-1}$, define a function $f = f_{\varepsilon,U}$ on \mathbb{RP}^{d-1} by

$$f(x) = \max \left\{ 1 - \frac{\delta(x, U)}{\varepsilon}, 0 \right\}. \quad (4.28)$$

Proposition 4.5. *For $0 < \alpha < 1$ the function f satisfies the bound*

$$\frac{|f(x) - f(y)|}{\delta(x, y)^\alpha} \leq \varepsilon^{-\alpha} \quad (4.29)$$

uniformly in $x, y \in \mathbb{RP}^{d-1}$.

Note: the expression on the lefthand side of (4.5) appears in [13, p. 106], where it is use to create a Banach space norm.

Proof. The result is immediate if either x and y are both in U , or both distance at least ε from U ; likewise it is immediate if one of them lies in U and the other lies distance at least ε from U . We may therefore assume, without loss of generality, that $0 < \delta(x, U) < \varepsilon$.

If $y \in U$, the quotient equals $\varepsilon^{-1}\delta(x, U)^{1-\alpha} < \varepsilon^{-\alpha}$. If $0 < \delta(y, U) < \varepsilon$, $|f(x) - f(y)|/\delta(x, y)^\alpha = \varepsilon^{-1}|\delta(x, U) - \delta(y, U)|/\delta(x, y)^\alpha \leq \varepsilon^{-1}|\delta(x, U) - \delta(y, U)|^{1-\alpha} \leq \varepsilon^{-\alpha}$, using the inequality

$$|\delta(x, U) - \delta(y, U)| \leq \delta(x, y). \quad (4.30)$$

For the remaining case $\delta(y, U) \geq \varepsilon$ we again use (4.30) to deduce $|f(x) - f(y)|/\delta(x, y)^\alpha = \varepsilon^{-1}|\varepsilon - \delta(x, U)|/\delta(x, y)^\alpha \leq \varepsilon^{-1}|\varepsilon - \delta(x, U)|^{1-\alpha} \leq \varepsilon^{-\alpha}$. \square

Proposition 4.6. (See [3, p. 107]) Consider the function f defined in terms of the set U and constant $\varepsilon > 0$ in (4.28). For α sufficiently small, there exists $n_0 > 0$ and $\rho < 1$ such that

$$\int_G f(Mv) d\mu^n(M) - \int_{\mathbb{RP}^{d-1}} f(y) d\nu(y) \leq \varepsilon^{-\alpha} \rho^n \quad (4.31)$$

for all $n \geq n_0$.

Proof. In fact, the present argument shows this inequality holds when the lefthand side of (4.31) is replaced by its absolute value, though we shall not need this. After ν by $\mu^n * \nu = \nu$ in the second integral, the lefthand side equals

$$\begin{aligned} & \int_G f(Mv) d\mu^n(M) - \int_{\mathbb{RP}^{d-1}} \int_G f(My) d\mu^n(M) d\nu(y) \\ &= \int_{\mathbb{RP}^{d-1}} \int_G f(Mv) d\mu^n(M) d\nu(y) - \int_{\mathbb{RP}^{d-1}} \int_G f(My) d\mu^n(M) d\nu(y) \\ &= \int_{\mathbb{RP}^{d-1}} \int_G (f(Mv) - f(My)) d\mu^n(M) d\nu(y) \\ &\leq \int_{\mathbb{RP}^{d-1}} \int_G \left(\sup_{x,y} \frac{|f(x) - f(y)|}{\delta(x, y)^\alpha} \right) \delta(Mv, My)^\alpha d\mu^n(M) d\nu(y) \\ &\leq \int_{\mathbb{RP}^{d-1}} \int_G \left(\sup_{x,y} \frac{|f(x) - f(y)|}{\delta(x, y)^\alpha} \right) \left(\frac{\delta(Mv, My)}{\delta(v, y)} \right)^\alpha d\mu^n(M) d\nu(y), \end{aligned} \quad (4.32)$$

the last inequality holding because $\delta(\cdot, \cdot) \leq 1$. The result now follows from Propositions 4.4 and 4.5. \square

We will eventually apply this to sets containing the $B_{r,i}$ from (4.6), which are all of the form

$$\{x \in \mathbb{RP}^{d-1} \mid \|Ax\| < \|x\|\} \quad (4.33)$$

for some $A \in GL(d, \mathbb{R})$ of norm greater than 1. Given such a matrix A , let $w = w_A \in \mathbb{R}^d$ be a unit vector such that $\|Aw\| = \|A\|$.

Proposition 4.7. $\|Ax\| < \|x\| \implies \|\langle x, w \rangle Aw\| \leq \|x\|$.

Proof. Let z be a vector perpendicular to w . For all $t \in \mathbb{R}$ we have that

$$\|Aw\|^2 \geq \frac{\|A(w + tz)\|^2}{\|w + tz\|^2} = \frac{\|Aw\|^2 + 2t\langle Aw, Az \rangle + t^2\|Az\|^2}{\|w\|^2 + t^2\|z\|^2}, \quad (4.34)$$

and so this last expression must have a local maximum at $t = 0$. In particular, its t -derivative at $t = 0$ must vanish, i.e., $\langle Aw, Az \rangle = 0$. Therefore if a vector $x \in \mathbb{R}^d$ is decomposed as $x = \langle x, w \rangle w + z$ for some $z \perp w$, then $Ax = A\langle x, w \rangle w + Az$ is again an orthogonal decomposition. It follows that $\|Ax\| \geq \|A\langle x, w \rangle w\|$, proving the proposition. \square

We now return to bounding $\mu^{j-1}\delta_v(B_i)$ in order to get an error estimate in (4.7). The sets $B_{r,i}$ are of the form (4.33), with $A = g_r^{-1}g_i$. We now fix r and i . By Proposition 4.7,

$$B_{r,i} \subset U = \left\{ x \in \mathbb{RP}^{d-1} \mid \frac{|\langle x, w \rangle|}{\|x\|} \leq \|A\|^{-1} \right\}. \quad (4.35)$$

Proposition 4.6 now shows that

$$\mu^{j-1}\delta_v(B_{r,i}) \leq \mu^{j-1}\delta_v(U) \leq \int_{\mathbb{RP}^{d-1}} f(y) d\nu(y) + \varepsilon^{-\alpha} \rho^{j-1}, \quad (4.36)$$

where $\varepsilon > 0$ is arbitrary and $f = f_{\varepsilon,U}$ is the function (4.28). The last integral is bounded by $\nu(U')$, where

$$\begin{aligned} U' &= \{x \in \mathbb{RP}^{d-1} \mid \delta(x, U) < \varepsilon\} \\ &= \{x \in \mathbb{RP}^{d-1} \mid \exists y \text{ with } \delta(x, y) < \varepsilon \text{ and } \frac{|\langle y, w \rangle|}{\|y\|} \leq \|A\|^{-1}\}. \end{aligned} \quad (4.37)$$

Here w , as above, represents a unit vector such that $\|Aw\| = \|A\|$. Using (4.12), this last condition on $|\langle y, w \rangle|$ can be restated as $\delta(y, w) \geq \sqrt{1 - \|A\|^{-2}}$. U' is in turn contained in the set

$$\begin{aligned} U'' &= \{x \in \mathbb{RP}^{d-1} \mid \delta(x, w) \geq \sqrt{1 - \|A\|^{-2}} - \varepsilon\} \\ &= \{x \in \mathbb{RP}^{d-1} \mid \frac{\langle x, w \rangle}{\|x\|} \leq \sqrt{1 - (\sqrt{1 - \|A\|^{-2}} - \varepsilon)^2}\} \end{aligned} \quad (4.38)$$

by the triangle inequality.

We now quote a result of Guivarc'h and Raugi (see [3, Theorem VI.2.1]) which immediately implies a bound on the ν -measure of U'' through the Chebyshev inequality. The comments in the proof of this Theorem on [3, p. 156] indicate that the exponent α has the same source as the one in Proposition 4.4 above, and thus may be taken to have the same value.

Theorem 4.8. (*Guivarc'h and Raugi*) Under assumptions **A1** and **A2** there exists constants $\alpha > 0$ and $K > 0$ such that

$$\int_{\mathbb{RP}^{d-1}} \left| \frac{\langle x, y \rangle}{\|x\|} \right|^{-\alpha} d\nu(x) \leq K \quad (4.39)$$

uniformly in y .

Applying the Chebyshev inequality to this with $y = w$, one gets

$$\int_{\mathbb{RP}^{d-1}} f(y) d\nu(y) \leq \nu(U'') \leq K \left(1 - (\sqrt{1 - \|A\|^{-2}} - \varepsilon)^2 \right)^{\alpha/2}. \quad (4.40)$$

Therefore using (4.36) and assumption **A3**, we bounded the $\text{Prob}_{\text{Error}}$ probability from (4.7) by

$$\text{Prob}_{\text{Error}} \leq \frac{1}{k} \sum_{\substack{t < j \leq L \\ 1 \leq r \neq i \leq k}} \left[K \left(1 - (\sqrt{1 - N^{-2}} - \varepsilon)^2 \right)^{\alpha/2} + \varepsilon^{-\alpha} \rho^{j-1} \right]. \quad (4.41)$$

The expression inside the large parentheses is

$$\frac{1}{N^2} - \varepsilon^2 + 2\varepsilon \sqrt{1 - \frac{1}{N^2}} < \frac{1}{N^2} + 2\varepsilon.$$

We now specify ε to be $\frac{3}{2N^2}$, so that the error is bounded by

$$\text{Prob}_{\text{Error}} \leq \frac{1}{k} \sum_{\substack{t < j \leq L \\ 1 \leq r \neq i \leq k}} (K 2^\alpha N^{-\alpha} + \varepsilon^{-\alpha} \rho^{j-1}). \quad (4.42)$$

Take $t = \lceil 1 + \frac{\log(3^\alpha K N^{-3\alpha})}{\log \rho} \rceil$, so that

$$K 2^\alpha N^{-\alpha} > \varepsilon^{-\alpha} \rho^{j-1} \quad \text{for } j \geq t \quad (4.43)$$

and

$$\text{Prob}_{\text{Error}} \leq \frac{1}{k} \cdot (L - t) k(k - 1) \cdot 2^{\alpha+1} K N^{-\alpha}. \quad (4.44)$$

This completes the proof of Theorem 4.1.

4.5 Numerical Examples

Example 1: where Norm Reduction works well

We now present an example of the algorithm in practice, for dimension $d = 3$ and the generating set $S = \{g_1, g_2, g_3\}$, where

$$g_1 = \begin{pmatrix} -9 & -59 & 30 \\ 11 & 66 & -32 \\ 3 & 21 & -11 \end{pmatrix}, \quad g_2 = \begin{pmatrix} 444 & -31 & -363 \\ -110 & 7 & 90 \\ -1271 & 90 & 1039 \end{pmatrix}, \quad \text{and } g_3 = \begin{pmatrix} 9 & 31 & 33 \\ -91 & -303 & -310 \\ -35 & -116 & -118 \end{pmatrix}. \quad (4.45)$$

These matrices were chosen randomly among those with integral entries in a

L	Number of Attempts	Number of Successes
2	10,000	10,000
10	10,000	9,998
50	10,000	9,978
100	10,000	9,963
200	10,000	9,936
1,000	1,000	1,000

Table 1: Numerical results with generating set S from (4.45).

bounded range. In all our tests we ran the algorithm with the parameter $t = 0$, i.e., not allowing for brute force search for the final steps. The parameter N in this example is ≈ 12157.1 . We ran several numerical trials of the Norm Reduction Algorithm (3.4) on the Word Problem on Vectors (3.3) with the vector $v = (1, 0, 0)$, almost all of which were successful (see Table 1).

The error term (4.44) is bounded by the one given in Theorem 4.1 if C is taken to be $4K$. In this typical example, the invariant measure ν and its approximations $\mu^n * \delta_v$ are supported near the eigenvectors for the g_i corresponding to their maximal eigenvalue. Recall that the constant K comes from the measure of the set U'' , which in (4.38) is related to points in \mathbb{RP}^2 which have δ -distance very close to 1 from the direction of maximal stretching of the six matrices $g_r^{-1}g_i$. We computed that these 18 pairs of δ -distances range between .33 and .98, far from 1 on the scale of $1/N$. Since C can be large only if these distances are much closer to 1, we concluded that C is small – under some heuristics, we computed its value to be below 7.

To estimate the value of α , we recall its origin in Proposition 4.4 comes from bounds on the quantities $\mathcal{S}(n)$ (4.16). We numerically estimated that $\mathcal{S}(2) < .83$ for $\alpha = 0.4$. This was done by approximating that maximum

using a mesh. While that is no guarantee of an accurate estimate for the maximum, it is worth noting that the values to be maximized were typically much smaller. Also, using $\mathcal{S}(n)$ for larger values of n would result in a better estimate for α . With this value of $\alpha = .4$, the probability in Theorem 4.1 is less than 1 only for small values of $L - t$. However, that estimate is certainly an overestimate for other reasons: for one thing, the proof estimates the error probability at each step, and multiplies this individual estimate by the number of steps to obtain the final estimate. The actual error probability is likely to be far smaller. The combination of this potential to improve the estimates, along with the excellent performance of the Norm Reduction Algorithm (3.4) in practice, demonstrates its usefulness in attacking the Word Problem on Vectors (3.3).

Example 2: where Norm Reduction does not work well

The algorithm does not perform well when one of the generators is orthogonal. In this example we take $S' = \{g'_1, g_2, g_3\}$, where $g'_1 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix}$ and g_2, g_3 are as defined in (4.45). With this one change (but otherwise the same conditions as in Example 1) the outcomes were much worse, and are summarized in Table 2.

L	Number of Attempts	Number of Successes
2	10,000	10,000
10	10,000	4,404
50	10,000	86
100	10,000	2
200	10,000	0
1000	1000	0

Table 2: Numerical results with generating set S' .

5 Rounding and the Traveling Salesman Problem

In this section we show how algorithms to solve the Closest Group Element Problem (3.5) can be easily converted to solve the Traveling Salesman Problem

lem (TSP), and in particular prove Theorem 3.1.

Definition 5.1. *Traveling Salesman Problem (on graphs).* Given a complete graph on n vertices whose edges have positive integer weights, find a Hamiltonian cycle which has minimal total weight (i.e., sum of its edge weights).

The above formulation is more general than the metric TSP problem, in that the edge weights do not need to obey the triangle inequality. The TSP problem is NP-hard, as is the simpler problem of finding a Hamiltonian cycle whose total weight is within a constant factor of the minimum [23, Theorem 3.6].

We shall now describe how to convert any instance of TSP into a Closest Group Element Problem (3.5). First we set some notation for the TSP problem. Let $w_e = w_{ij} = w_{ji}$ be the weight of the directed edge $e = (i, j)$ connecting the i -th and j -th vertices. Let m be an *a priori* lower bound for the total weight of the shortest Hamiltonian cycle (for example, m can be n times the lowest edge weight), and M be an upper bound (for example, the weight of any Hamiltonian cycle). Let m_0 denote the minimal total weight, which is unknown (and hence which we do not use in setting parameters). Since the weights are positive integers, one may of course assume that $m_0, m \geq 1$. The edge weight unit can be rescaled without affecting the solution to the TSP problem: accordingly we shall replace the above parameters by mT, MT , and m_0T , where $T > 0$ is a parameter that will be chosen later. After this rescaling, one has that

$$\text{any cycle weight less than } m_0T + T \text{ is minimal.} \quad (5.1)$$

In particular, there is no loss of generality in assuming that $M \geq m_0 + 1$. Given an edge $e = (i, j)$, let v_e denote the row vector of length n which has all zeroes except for K 's in positions i and j , where K is a parameter that will be chosen later. Let E_{ij} denote the $n \times n$ matrix which has all 0 entries except a 1 in the (i, j) -th position. Let $\beta > \alpha \geq 0$ be parameters (to be specified later), and $M_e = M_{ij} = \alpha I + \beta E_{ij}$. We set $d = 2n + 3$ and define $d \times d$ matrices for each directed edge by

$$g_e = g_{ij} = \begin{pmatrix} M_e & & & & \\ & 1 & & & \\ & & v_e & & \\ & & & I_{n \times n} & \\ & & & & 1 & w_e \\ & & & & & & 1 \end{pmatrix} \quad (5.2)$$

(the blocks in this matrix are of sizes $n, 1, n, 1,$ and 1 , respectively; we have as well used the convention that blank entries are zero). Note that $E_{ij} \neq E_{ji}$,

and consequently $M_{ij} \neq M_{ji}$ and $g_{ij} \neq g_{ji}$. The Zariski closure of the group (or semigroup) generated by $\{g_{i,j} | i \neq j\}$ contains $GL_n(\mathbb{R})$, embedded into the $n \times n$ block in the upper left corner, and satisfies the large dimensionality constraint of Section 1.

If h_1, \dots, h_ℓ are all square matrices of the same size, let $\prod_{i \leq \ell} h_i$ denote the product $h_1 \cdots h_\ell$. If e_1, \dots, e_ℓ are edges, then

$$g_{e_1} g_{e_2} \cdots g_{e_\ell} = \begin{pmatrix} \prod M_{e_r} & & & \\ & 1 & \sum v_{e_r} & \\ & & I_{n \times n} & \\ & & & 1 & \sum w_{e_r} \\ & & & & & 1 \end{pmatrix}. \quad (5.3)$$

We shall now see how features of this matrix are related to the total weights of Hamiltonian cycles. First of all, $\sum_{r \leq \ell} v_{e_r}$ equals $[2K \dots 2K]$ (i.e., a vector of all $2K$'s) if and only if the edges e_1, \dots, e_ℓ touch each vertex exactly twice. The entry $\sum_{r \leq \ell} w_{e_r}$ is of course the total weight of the path, if indeed e_1, \dots, e_r trace out a path. The product

$$\prod_{r \leq \ell} M_{e_r} = \prod_{r \leq \ell} (\alpha I + \beta E_{i_r j_r}) = \sum_{(\varepsilon_1, \dots, \varepsilon_\ell) \in \{0,1\}^\ell} \alpha^{\ell - (\varepsilon_1 + \dots + \varepsilon_\ell)} \beta^{\varepsilon_1 + \dots + \varepsilon_\ell} \prod_{r \leq \ell} E_{i_r j_r}^{\varepsilon_r} \quad (5.4)$$

helps detect such a path. The last product is zero unless the edges e_r for which $\varepsilon_r = 1$ trace out a connected path; if they do, the product equals E_{ij} , where i is the first value of i_r for which $\varepsilon_r = 1$ and j is the last value of j_r for which $\varepsilon_r = 1$. Note that if $\alpha = 0$, the only nonzero term is the one for $\varepsilon_1 = \varepsilon_2 = \dots = \varepsilon_\ell = 1$: then the product $\prod_{r \leq \ell} M_{e_r} = \beta^\ell E_{i_1 j_\ell}$ if the edges e_1, e_2, \dots, e_ℓ trace out a connected path, but is zero otherwise. Thus in the extreme case $\alpha = 0$, tracing out a connected path is equivalent to the nonvanishing of this product. Unfortunately, however, the matrices are only invertible if $\alpha > 0$. We will mainly be concerned with the case of $\alpha > 0$ because of its relevance to the Closest Group Element Problem (3.5), but include some comments about the $\alpha = 0$ case as well. In fact, the extra parameters α and β are needed simply to adapt features of the simpler $\alpha = 0$ case to noninvertible matrices.

Proposition 5.2. *The (i, j) -th entry of $\prod_{r \leq \ell} M_{e_r}$ satisfies the bound*

$$\left(\prod_{r \leq \ell} M_{e_r} \right)_{ij} \leq \alpha \beta^{\ell-1} 2^\ell \quad (5.5)$$

if the edges e_1, e_2, \dots, e_ℓ do not trace out a path, and

$$\left(\prod_{r \leq \ell} M_{e_r} \right)_{ij} - \beta^\ell \delta_{i=i_1} \delta_{j=j_\ell} \leq \alpha \beta^{\ell-1} 2^\ell \quad (5.6)$$

if they do.

Proof. In either case, the expressions to be bounded are the matrix entries of the sum on the righthand side of (5.4), except for the term corresponding to $\varepsilon_1 = \varepsilon_2 = \dots = \varepsilon_\ell = 1$ (which only comes up in (5.6) anyhow). The matrix entries of a product of $E_{i_r j_r}$ are all ≤ 1 , so the sum is bounded by $(\alpha + \beta)^\ell - \beta^\ell \leq \alpha \beta^{\ell-1} 2^\ell$. \square

Let $\varepsilon > 0$ be a parameter (which will be specified later). The Closest Group Element Problem (3.5) derived from this TSP instance is the following, assuming $\alpha > 0$ (if $\alpha = 0$, it is the verbatim rounding problem for *semigroups*):

Find the closest product of length $\leq n$ of the g_e 's to the matrix

$$z = \begin{pmatrix} \beta^n E_{11} + \varepsilon I_n & & \\ & 1 & 2K \cdots 2K \\ & & I_{n \times n} \\ & & & 1 & 0 \\ & & & & 1 \end{pmatrix} \text{ in the matrix norm (3.6).} \quad (5.7)$$

The block structure of the matrices allows us to compute the distance of a product in terms of the features described after (5.3):

$$\left\| \prod_{r \leq \ell} g_{e_r} - z \right\|^2 = \left\| \prod_{r \leq \ell} M_{e_r} - \beta^n E_{11} - \varepsilon I_n \right\|^2 + \left\| \sum_{r \leq \ell} v_{e_r} - [2K \cdots 2K] \right\|^2 + \left(\sum_{r \leq \ell} w_{e_r} \right)^2, \quad (5.8)$$

where we again stress that $\|\cdot\|$ refers to the norm (3.6) for the rest of this section.

Proposition 5.3. (Note that $\ell = n$ in parts (A) and (C).)

(A) If the edges e_1, e_2, \dots, e_n trace out a Hamiltonian cycle starting and ending at the first vertex, then

$$\left\| \prod_{r \leq \ell} M_{e_r} - \beta^n E_{11} - \varepsilon I_n \right\|^2 \leq (n \alpha \beta^{n-1} 2^n)^2 + (n \varepsilon)^2 \quad (5.9)$$

and consequently

$$\left\| \prod_{r \leq n} g_{e_r} - z \right\|^2 \leq (n\alpha\beta^{n-1}2^n)^2 + (n\varepsilon)^2 + (\text{weight of cycle})^2. \quad (5.10)$$

(B) If the edges e_1, e_2, \dots, e_ℓ do not touch each vertex exactly twice, then

$$\left\| \prod_{r \leq \ell} g_{e_r} - z \right\|^2 \geq K^2. \quad (5.11)$$

(C) If the edges e_1, e_2, \dots, e_n do not trace out a path beginning and ending at vertex 1, then

$$\left\| \prod_{r \leq n} M_{e_r} - \beta^n E_{11} - \varepsilon I_n \right\|^2 \geq \left| \left(\prod_{r \leq n} M_{e_r} \right)_{11} - \beta^n - \varepsilon \right|^2 \geq (\varepsilon + \beta^n - \alpha\beta^{n-1}2^n)^2. \quad (5.12)$$

Proof. The inequality (5.9) in part (A) is an immediate consequence of (5.6) and the triangle inequality. It then implies (5.10) because the middle term on the righthand side of (5.8) vanishes when the path enters and exists each vertex exactly once.

On the other hand, failure to touch each vertex exactly twice means one of the vector entries for the middle term in (5.8) will be at least K , showing that the righthand side of (5.11) is at least K^2 (in fact by parity considerations it will be at least $2K^2$). This demonstrates part (B). Part (C) is likewise a consequence of Proposition 5.2. \square

Proposition 5.4. *Suppose*

1. $(n\alpha\beta^{n-1}2^n)^2 + (n\varepsilon)^2 < mT^2$
2. $K \geq MT$
3. $\varepsilon + \beta^n - \alpha\beta^{n-1}2^n \geq MT$.

Then any word of length $\leq n$ in the g_e closest to z has the form $g_{e_1}g_{e_2} \cdots g_{e_n}$, where the edges e_1, e_2, \dots, e_n trace out a Hamiltonian cycle of shortest total weight that begins and ends at the first vertex.

Proof. We shall use all three parts of the previous Proposition. Part (A) and property 1 imply that if e_1, e_2, \dots, e_n is the shortest Hamiltonian cycle and $h_1 = g_{e_1}g_{e_2} \cdots g_{e_n}$, then (5.8) implies

$$\begin{aligned} \|h_1 - z\|^2 &\leq (n\alpha\beta^{n-1}2^n)^2 + (n\varepsilon)^2 + (m_0T)^2 \\ &< mT^2 + m_0^2T^2 \leq T^2(m_0 + 1)^2, \end{aligned} \quad (5.13)$$

because $m \leq m_0$.

Part (B) and property 2 imply that a path which does not touch each vertex exactly twice has

$$M^2T^2 \leq \left\| \prod_{r \leq \ell} g_{e_r} - z \right\|^2. \quad (5.14)$$

Since we have assumed $M \geq m_0 + 1$, the word $\prod_{r \leq \ell} g_{e_r}$ cannot be closest to z . In particular, the closest word to z must be a product of length exactly n (otherwise the edges it is formed from do not touch each vertex exactly twice). Part (C) and property 3 likewise show that the edges of the closest word trace out a path beginning and ending at 1.

Thus the closest word comes from a Hamiltonian cycle. We now must show that it comes from the Hamiltonian cycle of lowest total weight. Indeed, suppose that $h = \prod_{r \leq n} g_{e_r}$ comes from a Hamiltonian cycle and $\|h - z\| < \|h_1 - z\|$. By (5.8) and (5.13) we must have

$$(\text{total weight of } h\text{'s path})^2 \leq \|h - z\|^2 < \|h_1 - z\|^2 \leq T^2(m_0 + 1)^2, \quad (5.15)$$

and property (5.1) shows that this path is minimal – a contradiction. \square

Proposition 5.5. *Suppose edges e_1, e_2, \dots, e_n trace out a Hamiltonian cycle starting and ending at the first vertex, and whose total weight is $\leq m_0TA$ for some $A \geq 1$ (that is, within a factor A of being minimal). Suppose furthermore that*

$$(n\alpha\beta^{n-1}2^n)^2 + (n\varepsilon)^2 \leq (mAT)^2, \quad (5.16)$$

which is a consequence of the first assumption of Proposition 5.4 since $m, A \geq 1$. If $h = g_{e_1}g_{e_2} \cdots g_{e_n}$ (respectively, h') is the word formed from this cycle (respectively, a minimal cycle), then

$$\|h - z\| \leq \sqrt{2}A \|h' - z\|. \quad (5.17)$$

Proof. By (5.10) one has

$$\begin{aligned} \|h - z\|^2 &\leq (n\alpha\beta^{n-1}2^n)^2 + (n\varepsilon)^2 + (\text{weight of path})^2 \leq \\ &m^2 A^2 T^2 + m_0^2 A^2 T^2 \leq 2(m_0 A T)^2. \end{aligned} \quad (5.18)$$

The result follows because $\|h' - z\| \geq m_0 T$ by (5.8). \square

The conditions of the previous Propositions can be achieved with matrix entries that are polynomially-sized in the input of the TSP instance. For example, the following parameter choices are easily checked to satisfy them.

Proposition 5.6. *Properties 1, 2, and 3 of Proposition 5.4 as well as (5.16) hold under the following parameter choices.*

(i) $\alpha = 1$, $\beta = \max(2^{n+4}, 4M^2, \frac{n^2 2^{2n+1}}{m})$, $\varepsilon = \frac{1}{2n}$, $T = \beta^{n-1/2}$, and $K = MT$. In this case the matrices g_e all have determinant 1.

(ii) $T = 1$, $K = M$, $\beta = (2M)^{1/n}$, $\alpha = \frac{\sqrt{m/2}}{n^{2n}\beta^{n-1}}$, and $\varepsilon = \frac{\sqrt{m/2}}{n}$. In this case the matrices all have determinant α^n (and are hence invertible).

(iii) $\alpha = 0$, $\beta = M^{1/n}$, $\varepsilon = 0$, $K = M$, and $T = 1$. In this case the matrices are not invertible.

Since the entries in these matrices g_{ij} and z are polynomially sized, Theorem 3.1 then follows immediately from Proposition 5.5 and the corresponding inapproximability of the Traveling Salesman Problem on graphs [23, Theorem 3.6].

References

- [1] Sanjeev Arora, László Babai, Jacques Stern, and Z. Sweedyk, *The hardness of approximate optima in lattices, codes, and systems of linear equations*, J. Comput. System Sci. **54** (1997), no. 2, 317–331. 34th Annual Symposium on Foundations of Computer Science (Palo Alto, CA, 1993).
- [2] L. Babai, *On Lovász' lattice reduction and the nearest lattice point problem*, Combinatorica **6** (1986), no. 1, 1–13.
- [3] Philippe Bougerol and Jean Lacroix, *Products of random matrices with applications to Schrödinger operators*, Progress in Probability and Statistics, vol. 8, Birkhäuser Boston Inc., Boston, MA, 1985.
- [4] I. Dinur, G. Kindler, and S. Safra, *Approximating CVP to within almost-polynomial factors is NP-hard*, Symposium on Foundations of Computer Science, IEEE, New York, 1998, pp. 99–111.

- [5] Elena Fuchs, *The ubiquity of thin groups*, Thin groups and superstrong approximation, Math. Sci. Res. Inst. Publ., vol. 61, Cambridge Univ. Press, Cambridge, 2014, pp. 73–92. MR3220885
- [6] Harry Furstenberg, *Noncommuting random products*, Trans. Amer. Math. Soc. **108** (1963), 377–428.
- [7] H. Furstenberg and H. Kesten, *Products of random matrices*, Ann. Math. Statist **31** (1960), 457–469.
- [8] D. Garber, S. Kaplan, M. Teicher, B. Tsaban, and U. Vishne, *Probabilistic solutions of equations in the braid group*, Advances in Applied Mathematics **35** (2005), 323–334. <http://arxiv.org/abs/math/0404076>.
- [9] Yuri Gurevich and Paul Schupp, *Membership problem for the modular group*, SIAM J. Comput. **37** (2007), no. 2, 425–459 (electronic).
- [10] M. Gromov, *Asymptotic invariants of infinite groups*, Geometric group theory, Vol. 2 (Sussex, 1991), London Math. Soc. Lecture Note Ser., vol. 182, Cambridge Univ. Press, Cambridge, 1993, pp. 1–295.
- [11] Antoine Joux and Jacques Stern, *Lattice reduction: a toolbox for the cryptanalyst*, J. Cryptology **11** (1998), no. 3, 161–185.
- [12] A. K. Lenstra, H. W. Lenstra Jr., and L. Lovász, *Factoring polynomials with rational coefficients*, Math. Ann. **261** (1982), no. 4, 515–534.
- [13] Émile Le Page, *Théorèmes limites pour les produits de matrices aléatoires*, Probability measures on groups (Oberwolfach, 1981), Lecture Notes in Math., vol. 928, Springer, Berlin, 1982, pp. 258–303 (French).
- [14] Alexander Lubotzky, Shahar Mozes, and M. S. Raghunathan, *The word and Riemannian metrics on lattices of semisimple groups*, Inst. Hautes Études Sci. Publ. Math. **91** (2000), 5–53 (2001).
- [15] G. A. Margulis, *Discrete subgroups of semisimple Lie groups*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 17, Springer-Verlag, Berlin, 1991.
- [16] A. A. Markov, *The theory of Algorithms*, Israel Program for Scientific Translation, Jerusalem, 1961.
- [17] K. A. Mihailova, *The occurrence problem for direct products of groups*, Mat. Sb. (N.S.) **70 (112)** (1966), 241–251 (Russian).
- [18] Walter Rudin, *Functional analysis*, 2nd ed., International Series in Pure and Applied Mathematics, McGraw-Hill Inc., New York, 1991.
- [19] D. Ruinskiy, A. Shamir, and B. Tsaban, *Length-based cryptanalysis: The case of Thompson’s Group*, Journal of Mathematical Cryptology **1** (2007), 359–372.
- [20] Peter Sarnak, *Notes on thin matrix groups*, Thin groups and superstrong approximation, Math. Sci. Res. Inst. Publ., vol. 61, Cambridge Univ. Press, Cambridge, 2014, pp. 343–362. MR3220897

- [21] Adi Shamir, *A polynomial time algorithm for breaking the basic Merkle-Hellman cryptosystem*, Crypto, 1982, pp. 279–288.
- [22] J. Tsitsiklis and V. Blondel, *The Lyapunov exponent and joint spectral radius of pairs of matrices are hard - when not impossible - to compute and to approximate*, Mathematics of Control, Signals, and Systems **10** (1997), 31–40. Correction in **10**, p. 381.
- [23] Vijay V. Vazirani, *Approximation algorithms*, Springer-Verlag, Berlin, 2001. MR1851303 (2002h:68001)
- [24] Robert J. Zimmer, *Ergodic theory and semisimple groups*, Monographs in Mathematics, vol. 81, Birkhäuser Verlag, Basel, 1984.