The Complexity of the Extended GCD Problem

George Havas[1,*] and Jean-Pierre Seifert[2,**]

[1] Centre for Discrete Mathematics and Computing
Department of Computer Science and Electrical Engineering
The University of Queensland, Queensland 4072, Australia
havas@csee.uq.edu.au
[2] Department of Mathematics and Computer Science, University of Frankfurt
60054 Frankfurt on the Main, P.O. Box 11 19 32, Germany
seifert@cs.uni-frankfurt.de

Abstract. We undertake a thorough complexity study of the following fundamental optimization problem, known as the ℓ_p-norm shortest extended GCD multiplier problem:

> given $a_1, \ldots, a_n \in \mathbb{Z}$, find an *$\ell_p$-norm shortest* gcd *multiplier for* $a_1, \ldots, a_n$, i.e., a vector $\mathbf{x} \in \mathbb{Z}^n$ with minimum $(\sum_{i=1}^n |x_i|^p)^{1/p}$ satisfying $\sum_{i=1}^n x_i a_i = \gcd(a_1, \ldots, a_n)$.

First, we prove that the shortest GCD multiplier problem (in its feasibility recognition form) is NP-complete for every ℓ_p-norm with $p \in \mathbb{N}$. This gives an affirmative answer to a conjecture raised by Havas and Majewski. We then strengthen this negative result by ruling out even polynomial-time algorithms which only approximate an ℓ_p-norm shortest gcd multiplier within a factor $n^{1/(p \log^{\gamma} n)}$, for γ an arbitrary small positive constant, under the widely accepted complexity theory assumption $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{\mathrm{poly}(\log n)})$.

For positive results we focus on the ℓ_2-norm GCD multiplier problem. We show that approximating this problem within a factor of $\sqrt{n}$ is very unlikely NP-hard by placing it in $\mathsf{NP} \cap \mathsf{coAM}$ through a simple constant-round interactive proof system. This result is complemented by a polynomial-time algorithm which computes an ℓ_2-norm shortest gcd multiplier up to a factor of $2^{(n-1)/2}$.

This study is motivated by the importance of extended gcd calculations in applications in computational algebra and number theory. Our results rest upon the close connection between the hardness of approximation and the theory of interactive proof systems.

Key Words. Approximation algorithms, computational problems of diophantine equations, extended gcd computations, NP-hardness, probabilistically checkable proofs, interactive proof systems

* Partially supported by the Australian Research Council.
** Partially supported by a UQ High Quality Research Grant.

M. Kutyłowski, L. Pacholski, T. Wierzbicki (Eds.): MFCS'99, LNCS 1672, pp. 103–113, 1999.

1 Introduction

Extended gcd computation is of particular interest in number theory and in computational linear algebra, in both of which it takes a basic role in fundamental algorithms (see [14] for some references). The gcd problem for more than two numbers is interesting in its own right. Furthermore, it has important applications in its extended form where corresponding multipliers are required, for example in computing canonical normal forms of integer matrices ([13, 10, 15]). Of particular value are methods which find good solutions to the extended gcd problem with respect to some measure. Optimization problems for extended GCD calculation with respect to the ℓ_0-metric and the max-norm (ℓ_∞) have been proved NP-complete [19] and approximation for these measures is handled in [21]. Interestingly, the ℓ_0-metric problem is solvable in average polynomial time [12].

It is widely believed that some NP-optimization problems cannot be solved efficiently, i.e., in time polynomial in the input length of the problem. However, in many practical applications, including the extended gcd problem, approximate solutions of such problems do suffice. Thus, there has been much work done in studying the complexity of finding approximate solutions for NP-optimization problems.

It is desirable to have approximation algorithms such that the value of the returned solution is within a small factor of the optimum solution of the problem. For minimization problems, the worst-case ratio of the value of the solution returned by the approximation algorithm to the optimum solution is called the *approximation factor* of the approximation algorithm. It is useful to understand the quality of an approximation algorithm relative to the best that can be expected to be available in polynomial time. Thus, we may want to guarantee that a constructed approximation algorithm is best possible in that no substantially better approximation factor can be achieved, unless certain complexity theoretical assumptions are wrong.

In this context it is natural to study the complexity of the following optimization problem:

SHORTEST GCD MULTIPLIER in ℓ_p-norm (SGCDM_p)
INSTANCE: n numbers $a_1, \ldots, a_n \in \mathbb{Z}$
SOLUTION: A vector $\mathbf{x} \in \mathbb{Z}^n$ such that $\sum_{i=1}^{n} x_i a_i = \gcd(a_1, \ldots, a_n)$
MEASURE: The ℓ_p-norm $\|\mathbf{x}\|_p := (\sum_{1 \le i \le n} |x_i|^p)^{1/p}$ of the vector $\mathbf{x}$

Our strongest result states that, unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{\mathrm{poly}(\log n)})$, SGCDM_p cannot be approximated within a factor $n^{1/(p \log^{\gamma} n)}$, where γ is an arbitrary small positive constant. To prove this inapproximability result we start by constructing a gap-preserving reduction from the MIN TOTAL LABEL COVER problem in ℓ_1-norm to the SGCDM_1 problem. In contrast we show that, for the ℓ_2-norm, approximation within a factor of $\sqrt{n}$ is very unlikely NP-hard using an interactive proof system. We conclude by analyzing the approximation quality of an existing algorithm.

2 Preliminaries

Our proofs closely follow ideas in [1, 21]. First, we briefly introduce some notation (see [3]).

Definition 1. For an input I of a (positive-valued) minimization problem Π whose optimal solution has value $opt_\Pi(I)$, an algorithm A (which produces a solution with value $A(I)$) is said to *approximate* $opt_\Pi(I)$ *within a factor* $f(I)$ iff $opt_\Pi(I) \leq A(I) \leq f(I)\, opt_\Pi(I)$, where $f(I) \geq 1$.

Definition 2. Let Π and Π' be two minimization problems and $g,\ g' \geq 1$. A *gap-preserving reduction from* Π *to* Π' *with parameters* $((c,g),(c',g'))$ is a polynomial-time transformation τ mapping every instance I of Π to an instance $I' = \tau(I)$ of Π' such that for the optima $opt_\Pi(I)$ and $opt_{\Pi'}(I')$ of I and I', respectively, the following hold:

$$opt_\Pi(I) \leq c \Longrightarrow opt_{\Pi'}(I') \leq c' \quad \text{and}$$
$$opt_\Pi(I) > c \cdot g \Longrightarrow opt_{\Pi'}(I') > c' \cdot g',$$

where c, g and c', g' depend on the instance sizes $|I|$ and $|I'|$, respectively.

3 Hardness of Shortest $\mathbb{Z}$-Solution of Linear System

To prove our inapproximability result we construct a gap-preserving reduction from the Min Total Label Cover problem in ℓ_1-norm to the SGCDM_1 problem via the problems Shortest $\mathbb{Z}$-Solution of Linear System in ℓ_1-norm and Shortest Diophantine Equation Solution in ℓ_1-norm.

3.1 The Min Total Label Cover Problem

In the following $G = (V_1, V_2, E)$ denotes a bipartite graph, $\mathcal{B}$ a set of labels for the vertices in $V_1 \cup V_2$, and for every $e \in E$ there exists a partial function $\rho_e : \mathcal{B} \to \mathcal{B}$ describing the admissible pairs of labels. We adapt the notation of [2, 1].

Definition 3. A *labeling* of $G = (V_1, V_2, E)$ is a pair $(\mathcal{P}_1, \mathcal{P}_2)$ of functions $\mathcal{P}_i : V_i \to 2^{\mathcal{B}}$, $i = 1, 2$, which assign to each vertex in $V_1 \cup V_2$ a possibly empty set of labels.

Definition 4. Let $(\mathcal{P}_1, \mathcal{P}_2)$ be a labeling of $G = (V_1, V_2, E)$ and let $e = (v_1, v_2)$, $v_1 \in V_1$, $v_2 \in V_2$, be an edge of G. We call $e = (v_1, v_2)$ *covered* iff $\mathcal{P}_1(v_1) \neq \emptyset$, $\mathcal{P}_2(v_2) \neq \emptyset$ and for all labels $b_2 \in \mathcal{P}_2(v_2)$ there exists a label $b_1 \in \mathcal{P}_1(v_1)$ such that $\rho_e(b_1) = b_2$. A labeling $(\mathcal{P}_1, \mathcal{P}_2)$ of $G = (V_1, V_2, E)$ is called a *total-cover* of G iff every edge of G is covered by the labeling $(\mathcal{P}_1, \mathcal{P}_2)$.

Definition 5. The ℓ_1-*cost* of a labeling $(\mathcal{P}_1, \mathcal{P}_2)$ for a graph $G = (V_1, V_2, E)$ is defined as $cost(\mathcal{P}_1, \mathcal{P}_2) = \sum_{\substack{v_j \in V_i \\ i=1,2}} |\mathcal{P}_i(v_j)|$.

Definition 6. MIN TOTAL LABEL COVER in ℓ_1-norm (MINTLC_1)
INSTANCE: A d-regular bipartite graph $G = (V_1, V_2, E)$, a set of labels $\mathcal{B} = \{1, \ldots, \mathcal{N}\}$, $\mathcal{N} \in \mathbb{N}_+$, and for every edge $e \in E$ a partial function $\rho_e : \mathcal{B} \to \mathcal{B}$ such that $\rho_e^{-1}(1) \neq \emptyset$ for the distinguished label $1 \in \mathcal{B}$
SOLUTION: A total-cover $(\mathcal{P}_1, \mathcal{P}_2)$ of G
MEASURE: The ℓ_1-cost $cost(\mathcal{P}_1, \mathcal{P}_2)$ of the total-cover $(\mathcal{P}_1, \mathcal{P}_2)$

Remark 1. We can always ensure the existence of a total-cover with ℓ_1-cost at most $(|V_1| + 1)\mathcal{N}$; we simply let $\mathcal{P}_1(v_1) = \mathcal{B}$ for all $v_1 \in V_1$ and $\mathcal{P}_2(v_2) = \{1\}$ for all $v_2 \in V_2$.

The MIN TOTAL LABEL COVER in ℓ_1-norm is explicitly due to Khanna, Sudan and Trevisan [17] and a similar form of the following Lemma is implicitly proved in Lund and Yannakakis [18].

Lemma 1.

1. *For every constant $g \geq 1$ there exists a polynomial-time transformation τ from* 3-SAT *to* MIN TOTAL LABEL COVER *such that, for all instances I:*

$$I \in \text{3-SAT} \Longrightarrow \exists \text{ total-cover } (\mathcal{P}_1, \mathcal{P}_2) \text{ of } \tau(I) : cost(\mathcal{P}_1, \mathcal{P}_2) = 1 \cdot (|V_1| + |V_2|)$$
$$I \notin \text{3-SAT} \Longrightarrow \forall \text{ total-cover } (\mathcal{P}_1, \mathcal{P}_2) \text{ of } \tau(I) : cost(\mathcal{P}_1, \mathcal{P}_2) > g \cdot (|V_1| + |V_2|).$$

2. *There exists a quasi-polynomial-time, i.e.,* $\mathsf{DTIME}(n^{\text{poly}(\log n)})$, *transformation τ from* 3-SAT *to* MIN TOTAL LABEL COVER *such that, for all instances I:*

$$I \in \text{3-SAT} \Longrightarrow \exists \text{ total-cover } (\mathcal{P}_1, \mathcal{P}_2) \text{ of } \tau(I) : cost(\mathcal{P}_1, \mathcal{P}_2) = 1 \cdot (|V_1| + |V_2|)$$
$$I \notin \text{3-SAT} \Longrightarrow \forall \text{ total-cover } (\mathcal{P}_1, \mathcal{P}_2) \text{ of } \tau(I) : cost(\mathcal{P}_1, \mathcal{P}_2) > g \cdot (|V_1| + |V_2|),$$

where $g = 2^{\log^{1-\gamma}|\tau(I)|}$ with γ an arbitrary small positive constant.

Proof. Both statements 1. and 2. are simultaneously proved by combining ideas of Arora and Lund [2] and Lund and Yannakakis [18] with the recent result of Raz [20]. The proof is deferred to the full version. □

3.2 SHORTEST $\mathbb{Z}$-SOLUTION OF LINEAR SYSTEM

SHORTEST $\mathbb{Z}$-SOLUTION OF LINEAR SYSTEM in ℓ_1-norm (SZSLS_1)
INSTANCE: A linear system $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ of m equations in n variables where $\boldsymbol{A}$ is a rational $m \times n$ matrix and $\mathbf{b}$ an m-dimensional rational vector
SOLUTION: A nonzero vector $\mathbf{x} \in \mathbb{Z}^n$ satisfying $\boldsymbol{A}\mathbf{x} = \mathbf{b}$
MEASURE: The ℓ_1-norm $\|\mathbf{x}\|_1 := \sum_{1 \leq i \leq n} |x_i|$ of the vector $\mathbf{x}$

Theorem 1. *There exists a polynomial-time transformation τ from* MIN TOTAL LABEL COVER *to* SHORTEST $\mathbb{Z}$-SOLUTION OF LINEAR SYSTEM *such that, for all instances I and for all $g \geq 1$:*

$$opt_{\text{MinTLC}_1}(I) = 1 \cdot (|V_1| + |V_2|) \Longrightarrow opt_{\text{SZSLS}_1}(\tau(I)) = 1 \cdot (|V_1| + |V_2|)$$
$$opt_{\text{MinTLC}_1}(I) > g \cdot (|V_1| + |V_2|) \Longrightarrow opt_{\text{SZSLS}_1}(\tau(I)) > g \cdot (|V_1| + |V_2|).$$

Proof. From a given MIN TOTAL LABEL COVER instance $I = (V_1, V_2, E, \rho, \mathcal{B}, \mathcal{N})$ we construct a linear system of equations $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ with $\boldsymbol{A}$ an $m \times n$ matrix of entries $\{0,1\}$, $\mathbf{b}$ an m-dimensional all-one vector, $m = |E|(\mathcal{N}+1)$ and $n = |V_1|\mathcal{N} + |V_2|\mathcal{N}$.

For every pair (v,b) with $v \in V_1 \cup V_2$ and $b \in \mathcal{B}$ we define a column vector $\mathbf{a}_{v,b} \in \{0,1\}^m$ of $\boldsymbol{A}$ as follows. The $|E|(\mathcal{N}+1)$ coordinates of $\mathbf{a}_{v,b}$ are split into $|E|$ blocks of *e-projections* $\mathbf{u}_e(\mathbf{a}_{v,b})$ — one $(\mathcal{N}+1)$-length block for every edge $e \in E$. In particular, we define for every $(v_2, b_2) \in V_2 \times \mathcal{B}$

$$\mathbf{u}_e(\mathbf{a}_{v_2,b_2}) := \begin{cases} \mathbf{e}_{b_2} & \text{iff } e \text{ is incident to } v_2 \\ \mathbf{0} & \text{otherwise} \end{cases}$$

and for every $(v_1, b_1) \in V_1 \times \mathcal{B}$

$$\mathbf{u}_e(\mathbf{a}_{v_1,b_1}) := \begin{cases} \mathbf{1} - \mathbf{e}_{\rho_e(b_1)} & \text{iff } e \text{ is incident to } v_1 \text{ and } \rho_e(b_1) \neq \emptyset \\ \mathbf{0} & \text{otherwise} \end{cases}$$

where $\mathbf{e}_j$, $j = 1, \ldots, \mathcal{N}$, denotes the j^{th}-unit vector and $\mathbf{0}$, $\mathbf{1}$ the all-zero, all-one vector in $\mathbb{R}^{\mathcal{N}+1}$, respectively.

Next, we define the right-hand side of our linear system — the vector $\mathbf{b}$ — as the vector having 1 in each of its $|E|(\mathcal{N}+1)$ coordinates.

Now let $\mathbf{x} = \sum x_{v,b}\mathbf{a}_{v,b}$ be an integer linear combination of the column vectors $\mathbf{a}_{v,b}$. Then, assigning every vertex v a label b iff $x_{v,b} \neq 0$ defines a labeling $(\mathcal{P}_1^{\mathbf{x}}, \mathcal{P}_2^{\mathbf{x}})$ *induced by the vector* $\mathbf{x}$. From [1, Corollary 12] it follows that any such $\mathbf{x}$ induces a total-cover of (V_1, V_2, E). Thus, any solution $\mathbf{x} \in \mathbb{Z}^{|V_1|\mathcal{N}+|V_2|\mathcal{N}}$ of the linear system $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ induces a total-cover of (V_1, V_2, E). Therefore, we deduce that for any solution $\mathbf{x} \in \mathbb{Z}^n$ of $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ we must have $\|\mathbf{x}\|_1 \geq opt_{\text{MinTLC}_1}(I)$.

On the other hand, assume now that $opt_{\text{MinTLC}_1}(I) = (|V_1| + |V_2|)$ and let $(\mathcal{P}_1, \mathcal{P}_2)$ denote the corresponding labeling. Then, the vector $\mathbf{x}$ given by

$$\begin{aligned} x_{v_j,\mathcal{P}_i(v_j)} &:= 1 \quad \forall v_j \in V_i,\ i = 1,2 \\ x_{v_j,b} &:= 0 \quad \forall v_j \in V_i,\ \forall b \in \mathcal{B} \setminus \mathcal{P}_i(v_j),\ i = 1,2 \end{aligned}$$

is a feasible solution of the linear system $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ satisfying $\|\mathbf{x}\|_1 = (|V_1| + |V_2|)$.

The reduction from the given instance I of MIN TOTAL LABEL COVER to the above constructed linear system $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ is feasible in time polynomial in the dimension of $\boldsymbol{A}$ which in turn is polynomial in $|I|$. Clearly, the above reduction, say τ, is gap-reserving with parameters $((|V_1| + |V_2|, g), (|V_1| + |V_2|, g))$. □

4 Hardness of Approximating SGCDM$_1$

4.1 A new Aggregation Lemma

The following Lemma establishes for the first time a polynomial-time reduction from a system of inhomogeneous linear equations to a single equation with identical ℓ_p-bounded solution set. Previously, it was only known to hold by a result of Kannan [16] for the case of identical ℓ_∞-bounded solution sets.

Lemma 2. *Let $p, q \in \mathbb{N} \cup \{\infty\}$ with $1/p + 1/q = 1$, $\boldsymbol{A}$ be an integral $m \times n$ matrix, $\|\boldsymbol{A}\|_{\infty,q} := \max_{1 \le i \le m} (\sum_{1 \le j \le n} |a_{i,j}|^q)^{1/q}$ and $\mathbf{b}$ be an integral m-dimensional vector. Then*

$$B_\mu \cap \{\mathbf{x} \in \mathbb{Z}^n : \boldsymbol{A}\mathbf{x} = \mathbf{b}\} = B_\mu \cap \left\{\mathbf{x} \in \mathbb{Z}^n : \sum_{i=0}^{m-1} \sum_{j=0}^{n-1} k^i a_{ij} x_j = \sum_{i=0}^{m-1} k^i b_i\right\}$$

where B_μ denotes the n-dimensional ball of ℓ_p-radius μ centered at the origin and $k \ge \|\boldsymbol{A}\|_{\infty,q}\mu + \|\mathbf{b}\|_\infty + 1$ an integer.

Proof. Denote the two sets by S_m and S_1, respectively. Clearly, $S_m \subseteq S_1$. To prove the reverse inclusion, suppose that there exists an element $\mathbf{x} \in S_1$ not satisfying at least one equation of $\boldsymbol{A}\mathbf{x} = \mathbf{b}$. Let $\boldsymbol{A} =: [\mathbf{a}_1, \ldots, \mathbf{a}_m]^\top$ and let $i_{\max}$ denote the largest index for which $\langle \mathbf{a}_i, \mathbf{x} \rangle \neq b_i$. As $\|\mathbf{x}\|_p \le \mu$ we have (using Hölder's inequality) $|\langle \mathbf{a}_i, \mathbf{x} \rangle - b_i| \le \|\boldsymbol{A}\|_{\infty,q}\mu + \|\mathbf{b}\|_\infty \le k - 1$ and since $\mathbf{x} \in S_1$ we must have $\sum_{i=0}^{m-1} k^i(\langle \mathbf{a}_i, \mathbf{x} \rangle - b_i) = 0$. By definition of $i_{\max}$ this yields $\sum_{i=0}^{i_{\max}-1} k^i(\langle \mathbf{a}_i, \mathbf{x} \rangle - b_i) = -k^{i_{\max}}(\langle \mathbf{a}_{i_{\max}}, \mathbf{x} \rangle - b_{i_{\max}})$ with a nonzero right-hand side implying that the left-hand side is also nonzero. Now the left-hand side is both a multiple of $k^{i_{\max}}$ and in absolute value bounded by $k^{i_{\max}} - 1$, a contradiction. □

4.2 Hardness of Approximating Shortest Diophantine Equation Solution or Aggregation Part I

Shortest Diophantine Equation Solution in ℓ_1-norm (SDES_1)
INSTANCE: An equation $x_1 a_1 + \cdots + x_n a_n = b$ with $a_1, \ldots, a_n, b \in \mathbb{Z}$
SOLUTION: A vector $\mathbf{x} \in \mathbb{Z}^n$ such that $\sum_{i=1}^n x_i a_i = b$
MEASURE: The ℓ_1-norm $\|\mathbf{x}\|_1 := \sum_{1 \le i \le n} |x_i|$ of the vector $\mathbf{x}$

Theorem 2. *There exists a polynomial-time transformation τ from (a restricted subset of)* Shortest $\mathbb{Z}$-Solution of Linear System *to* Shortest Diophantine Equation Solution *such that:*

1. *for all instances I and for all $c, g \ge 1$:*

$$opt_{\mathrm{SZSLS}_1}(I) = c \Longrightarrow opt_{\mathrm{SDES}_1}(\tau(I)) = c$$
$$opt_{\mathrm{SZSLS}_1}(I) > c \cdot g \Longrightarrow opt_{\mathrm{SDES}_1}(\tau(I)) > c \cdot g$$

2. *the constructed instance $a'_1 x_1 + \cdots + a'_n x_n = b'$ of SDES_1 satisfies*

$$x a'_{j^*} = b' \text{ for some } j^* \in \{1, \ldots, n\} \text{ and some } x \in \mathbb{Z}.$$

Proof. Consider the linear system $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ constructed in the reduction of the proof of Theorem 1. Recall that for the underlying d-regular graph $G = (V_1, V_2, E)$ with label set $\mathcal{B}$, $|\mathcal{B}| = \mathcal{N}$, the matrix $\boldsymbol{A}$ is a $m \times n$ matrix of entries $\{0, 1\}$ with $m = |E|(\mathcal{N} + 1)$ and $n = |V_1|\mathcal{N} + |V_2|\mathcal{N}$, and that the vector $\mathbf{b}$ has only 1-entries in its coordinates.

Let $\mathbf{a}_{j^*} := \mathbf{a}_{v_2,b_2}$ for some $v_2 \in V_2$, $b_2 \in \mathcal{B}$ be a nonzero column vector of $\boldsymbol{A}$. Since the graph G is d-regular, $\mathbf{a}_{j^*}$ has exactly d 1-entries. Let σ be the permutation shifting all 1-entries of $\mathbf{a}_{j^*}$ to its first d coordinates and note that $\sigma(\mathbf{b}) = \mathbf{b}$.

Aggregating now the (solution equivalent) linear system $\sigma(\boldsymbol{A})\mathbf{x} = \mathbf{b}$ via Lemma 2 yields the inhomogenous Diophantine equation $\sum_{j=0}^{|V_1|\mathcal{N}+|V_2|\mathcal{N}-1} a'_j x_j = b'$ with $a'_j := \sum_{i=0}^{|E|(\mathcal{N}+1)-1} k^i a_{\sigma(i),j}$ and $b' := \sum_{i=0}^{|E|(\mathcal{N}+1)-1} k^i b_i$. In particular, we have $a'_{j^*} = \sum_{i=0}^{d-1} k^i = \frac{k^d-1}{k-1}$ and $b' = \sum_{i=0}^{|E|(\mathcal{N}+1)-1} k^i = \frac{k^{|E|(\mathcal{N}+1)}-1}{k-1}$, and from the regularity of G, i.e., $d|V_2| = |E|$, it follows that (k^d-1) divides $(k^{|E|(\mathcal{N}+1)}-1)$, hence a'_{j^*} divides b'.

Thus, we have constructed a gap-preserving reduction from the linear system $\boldsymbol{A}\mathbf{x} = \mathbf{b}$ to an inhomogenous Diophantine equation instance $a'_1x_1 + \cdots + a'_nx_n = b'$ with parameters $((c,g),(c,g))$ such that there is an index $j^* \in \{1,\ldots,n\}$ satisfying $a'_{j^*} \mid b'$. □

4.3 The Final Reduction or Aggregation Part II

The basic construction The MIN TOTAL LABEL COVER problem in ℓ_1-norm is not approximable within a factor $2^{\log^{1-\gamma} n}$, where γ is an arbitrary small positive constant, unless NP $\subseteq$ DTIME($n^{\text{poly}(\log n)}$). From the gap-preserving reductions we have that the same inapproximability factor holds for the problems SZSLS_1 and SDES_1.

Theorem 3. *There exists a polynomial-time transformation τ from (a restricted subset of)* SDES_1 *to* SGCDM_1 *such that, for all instances I and for all $c, g \geq 1$:*

$$opt_{\text{SDES}_1}(I) = c \Longrightarrow opt_{\text{SGCDM}_1}(\tau(I)) = c+1$$
$$opt_{\text{SDES}_1}(I) > c \cdot g \Longrightarrow opt_{\text{SGCDM}_1}(\tau(I)) > c \cdot g + 1.$$

Proof. We start with the instance $a'_1x_1 + \cdots + a'_nx_n = b'$ of SDES_1 constructed in the above Theorem 2 and consider for an arbitrary integer $c \in \mathbb{Z} \setminus \{0\}$ the linear system

$$\begin{aligned} cx_{n+1} &= c \\ a'_1x_1 + \cdots + a'_nx_n - b'x_{n+1} &= 0 \end{aligned}$$

which forces the variable x_{n+1} to take on the value 1. Now, we fix $g \geq 1$ and apply Lemma 2 with suitable k to this linear system, obtaining the single equation

$$ka'_1x_1 + \cdots + ka'_nx_n + (c - kb')x_{n+1} = c.$$

We observe that the right-hand side c in the last equation was an arbitrarily chosen integer and that b' satisfies (by Theorem 2)

$$xa'_{i^*} = b' \text{ for some } i^* \in \{1, \ldots, n\} \text{ and some } x \in \mathbb{Z} . \qquad (*)$$

This will give us the desired gap-preserving reduction τ. Namely, we choose $c = \gcd(a'_1, \cdots, a'_n)$. By $(*)$ this implies

$$\begin{aligned} &\gcd(ka'_1, \cdots, ka'_n, (c - kb')) \\ &= \gcd(ka'_1, \cdots, ka'_{i^*-1}, ka'_{i^*+1}, \ldots, ka'_n, \gcd(ka'_{i^*}, (c - kxa'_{i^*}))) \\ &= c. \end{aligned}$$

Since the variable x_{n+1} is forced to take on the value 1, it is obvious from the above construction that the reduction τ with $\tau(I) = \{ka'_1, \cdots, ka'_n, (c - kb')\}$ satisfies the claim of the Theorem. □

Treating other norms Close inspection of the whole proof shows that it also implies the following general result.

Theorem 4. *There exists a polynomial-time transformation τ from* SDES_1 *to* SGCDM_p *such that, for all instances I and for all $c, g \geq 1$:*

$$\begin{aligned} opt_{\mathrm{SDES}_1}(I) = c &\Longrightarrow opt_{\mathrm{SGCDM}_p}(\tau(I)) = \sqrt[p]{c+1} \\ opt_{\mathrm{SDES}_1}(I) > c \cdot g &\Longrightarrow opt_{\mathrm{SGCDM}_p}(\tau(I)) > \sqrt[p]{c \cdot g + 1}. \end{aligned}$$

4.4 Implications on the complexity of SGCDM_p

We are now ready to present our main result.

Theorem 5.

1. *Unless* NP ⊆ P, *there exists no polynomial-time algorithm which approximates the* Shortest GCD Multiplier *problem in ℓ_p-norm within a factor of k, where $k \geq 1$ is an arbitrary constant.*
2. *Unless* NP ⊆ DTIME$(n^{\mathrm{poly}(\log n)})$, *there exists no polynomial-time algorithm which approximates the* Shortest GCD Multiplier *problem in ℓ_p-norm within a factor of $n^{1/(p \log^{\gamma} n)}$, for γ an arbitrary small positive constant.*

Proof. Combining the gap-preserving reductions of Lemma 1 (first statement), Theorem 1, Theorem 2 and Theorem 4 gives for all instances I:

$$\begin{aligned} I \in \text{3-Sat} &\Longrightarrow opt_{\mathrm{SGCDM}_p}(\tau(I)) = \sqrt[p]{|V_1| + |V_2| + 1} \\ I \notin \text{3-Sat} &\Longrightarrow opt_{\mathrm{SGCDM}_p}(\tau(I)) > \sqrt[p]{g(|V_1| + |V_2|) + 1}, \end{aligned}$$

where $g \geq 1$ can be any positive constant. Therefore, given a polynomial-time algorithm approximating the SGCDM_p problem within a factor of $k = \sqrt[p]{g}$ would enable us to decide 3-Sat in polynomial-time, thus proving the first claim.

The second claim follows similarly, via the second statement of Lemma 1 instead of its first statement. □

Corollary 1. *For every ℓ_p-norm, the Extended GCD problem is* NP*-complete (in its feasibility recognition form).*

5 Limits on the Hardness of Approximating SGCDM_2

Theorem 6. *The $\sqrt{n/O(\log n)}$-*SHORTEST GCD MULTIPLIER *in ℓ_2-norm is not* NP*-hard unless the Polynomial-Time Hierarchy collapses to its second level.*

Proof. We only give a brief outline of the proof, which uses results from Babai [4], Boppana et al. [6], Goldreich and Goldwasser [7], Goldwasser et al. [8] and Goldwasser and Sipser [9].

Along the thread of the above Theorem and Boppana et al. [6] we have to show a constant-round interactive proof system for $\mathrm{co}g(n)$-SGCDM_2 with $g = \sqrt{n/O(\log n)}$. Notice that the input to $\mathrm{co}g(n)$-SGCDM_2 is a tuple $(\mathbf{a}, l)$ with $\mathbf{a} \in \mathbb{Z}^n$ and $l \in \mathbb{R}$ and we have to show that yes-instances (i.e., $opt_{\mathrm{SGCDM}_2}(\mathbf{a}) > g \cdot l$) are always accepted, whereas no-instances (i.e., $opt_{\mathrm{SGCDM}_2}(\mathbf{a}) \leq l$) are accepted with probability bounded away from 1. The following simple interactive proof system establishes all that we need for an input tuple $(\mathbf{a}, l)$:

$V(\mathbf{a}, l)$		$P(\mathbf{a}, l)$
compute basis $B \stackrel{\text{def}}{=} [\mathbf{b}_1, \ldots, \mathbf{b}_{n-1}]$ of lattice $\mathbb{Z}^n \cap \mathbf{a}\mathbb{R}^\perp$		
compute $\mathbf{v} \in \mathbb{Z}^n$ with $\langle \mathbf{v}, \mathbf{a} \rangle = \gcd(\mathbf{a})$		
let $M := \max\{\|\mathbf{b}_1\|, \ldots, \|\mathbf{b}_{n-1}\|, \|\mathbf{v}\|\}$		
pick $\boldsymbol{r} \in_R \{\boldsymbol{p} \in L(B) : \|\boldsymbol{p}\|_2 \leq 2^n \cdot M\}$		
pick $\boldsymbol{\eta} \in_R \{\boldsymbol{p} \in \mathbb{R}^n : \|\boldsymbol{p}\|_2 \leq \frac{l \cdot g}{2}\}$		
pick $\sigma \in_R \{0, 1\}$		
$\mathbf{x} := \boldsymbol{r} + \sigma\mathbf{v} + \boldsymbol{\eta}$		
	$\xrightarrow{\mathbf{v}, \mathbf{x}}$	
	P computes $\tau := \begin{cases} 0 \text{ iff } \mathrm{dist}(\mathbf{x}, L(B)) < \mathrm{dist}(\mathbf{x}, L(B) + \mathbf{v}) \\ 1 \text{ otherwise} \end{cases}$	
	$\xleftarrow{\tau}$	
accept iff $\tau = \sigma$		

Following Goldreich and Goldwasser [7], the claimed properties of the above protocol are witnessed below by two Lemmas, whose easy proofs are omitted.

Lemma 3. *If $opt_{\mathrm{SGCDM}_2}(\mathbf{a}) > g \cdot l$ then there exists a prover P such that V accepts with probability 1 when interacting with P.*

Lemma 4. *Let $c > 0$ and $g(n) \geq \sqrt{n/(c \log n)}$, if $opt_{\mathrm{SGCDM}_2}(\mathbf{a}) \leq l$ then, for any $\widetilde{P}$ interacting with V, the verifier V accepts with probability at most $1 - 1/n^{2c}$.*

Finally, transforming the above protocol with the methods of Babai [4] and Goldwasser and Sipser [9] into an AM proof system yields for the gap $g = \sqrt{n/O(\log n)}$ the inclusion $g(n)$-$\mathrm{SGCDM}_2 \in \mathsf{NP} \cap \mathsf{coAM}$. The rest of the proof is straightforward. □

6 A LLL-based approximation algorithm

Here we give an overview of a polynomial time approximation algorithm for the ℓ_2-norm. This and related algorithms are studied in [14]:

compute basis $B \stackrel{\text{def}}{=} [\mathbf{b}_1, \ldots, \mathbf{b}_{n-1}]$ of lattice $\mathbb{Z}^n \cap \mathbf{a}\mathbb{R}^\perp$
compute multiplier $\mathbf{v} \in \mathbb{Z}^n$ with $\langle \mathbf{v}, \mathbf{a} \rangle = \gcd(\mathbf{a})$
apply the LLL-reduction to the basis $[\mathbf{b}_1, \ldots, \mathbf{b}_{n-1}]$
compute $\mathbf{x}' \in L(B)$ with $\|\mathbf{x}' - \mathbf{v}\| \leq 2^{(n-1)/2} \min_{\mathbf{u} \in L(B)} \|\mathbf{u} - \mathbf{v}\|$
  by applying Babai's nearest plane algorithm to the reduced basis
  and the target vector $\mathbf{v}$
output short multiplier $\mathbf{x} := \mathbf{v} - \mathbf{x}'$

Theorem 7. *On inputs $a_1, \ldots, a_n \in \mathbb{Z}$ the above algorithm computes in polynomial time a vector $\mathbf{x}$ with $\langle \mathbf{x}, \mathbf{a} \rangle = \gcd(\mathbf{a})$ satisfying*

$$\|\mathbf{x}\| \leq 2^{(n-1)/2} \cdot opt_{\mathrm{SGCDM}_2}(a_1, \ldots, a_n).$$

Proof. The lattice $\mathbb{Z}^n \cap \mathbf{a}\mathbb{R}^\perp$ constitutes an $n-1$ dimensional lattice in $\mathbb{Z}^n$. As shown by Babai [5], the nearest plane heuristic run on the LLL-reduced basis of $[\mathbf{b}_1, \ldots, \mathbf{b}_{n-1}]$ and the target vector $\mathbf{v}$ delivers a vector $\mathbf{x}' \in L(B)$ satisfying $\|\mathbf{x}' - \mathbf{v}\| \leq 2^{(n-1)/2} \min_{\mathbf{u} \in L(B)} \|\mathbf{u} - \mathbf{v}\|$. The Theorem follows. □

7 Conclusions

We have shown that the shortest GCD multiplier problem (in its feasibility recognition form) is NP-complete for every ℓ_p-norm with $p \in \mathbb{N}$. Furthermore, approximating an ℓ_p-norm shortest GCD multiplier within a factor $n^{1/(p \log^\gamma n)}$, for γ an arbitrary small positive constant is quasi-NP-hard.

On the positive side, we specialize to the ℓ_2-norm. Approximating the solution to the ℓ_2-norm problem within a factor of $\sqrt{n}/\sqrt{O(\log n)}$ is very unlikely NP-hard. We see that there is a transition in theoretical complexity at around $\sqrt{n}$: within a factor of $(\sqrt{n})^{1/(\log^\gamma n)}$ is quasi-NP-hard; within a factor of $\sqrt{n}/\sqrt{O(\log n)}$ is unlikely to be NP-hard.

The best known polynomial time algorithm for the ℓ_2-norm problem achieves a factor of $2^{(n-1)/2}$. It remains to be seen whether this substantial gap between around $\sqrt{n}$ and $2^{(n-1)/2}$ can be reduced or whether a LLL-based algorithm can be better analyzed in some circumstances.

References

1. S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, April 1997.

2. S. Arora and C. Lund. Hardness of approximation. In D. Hochbaum, editor, *Approximation Algorithms for NP-hard Problems*, chapter 11. PWS Publishing, 1996.
3. G. Ausiello, P. Crescenzi, and M. Protasi. Approximate solutions of NP optimization problems. *Theoretical Computer Science*, 150:1–55, 1995.
4. L. Babai. Trading group theory for randomness. In *Proc. 17th Ann. ACM Symp. on Theory of Computing*, pages 421–429, 1985.
5. L. Babai. On Lovasz' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–13, 1986.
6. R. Boppana, J. Håstad, and S. Zachos. Does coNP have short interactive proofs? *Information Processing Letters*, 25(2):127–132, 1987.
7. O. Goldreich and S. Goldwasser. On the limits of non-approximability of lattice problems. In *Proc. 30th ACM Symp. Theory of Computing*, pages 1–9, 1998.
8. S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, February 1989.
9. S. Goldwasser and M. Sipser. Private coins versus public coins in interactive proof systems. *Advances in Computing Research 5: Randomness and Computation*, 1989.
10. G. Havas and B. S. Majewski. Hermite normal form computation for integer matrices. *Congressus Numerantium*, 105:184–193, 1994.
11. G. Havas and B. S. Majewski. Extended gcd calculation. *Congressus Numerantium*, 111:104–114, 1995.
12. G. Havas and B. S. Majewski. A hard problem which is almost always easy. In *Algorithms and Computation*, Lecture Notes in Computer Science 1004, 216–223, 1995.
13. G. Havas and B.S. Majewski. Integer matrix diagonalization. *Journal of Symbolic Computation*, 24:399–408, 1997.
14. G. Havas, B. S. Majewski and K. R. Matthews. Extended gcd and Hermite normal form algorithms via lattice basis reduction. *Experimental Mathematics*, 7:125–136, 1998.
15. C. S. Iliopoulos. Worst case complexity bounds on algorithms for computing the canonical structure of finite abelian groups and the Hermite and Smith normal forms of an integer matrix. *SIAM Journal on Computing*, 18:658–669, 1989.
16. R. Kannan. Polynomial-time aggregation of integer programming problems. *J. ACM*, 30:133–145, 1983.
17. S. Khanna, M. Sudan and L. Trevisan. Constraint Satisfaction: The Approximability of Minimization Problems. In *Proceedings of the 12th IEEE Conference on Computational Complexity*, pages 282-296, 1997.
18. C. Lund and M. Yannakakis. On the hardness of minimization problems. *J. ACM*, 41:960–981, 1994.
19. B. S. Majewski and G. Havas. The complexity of greatest common divisor computations. In *Algorithmic Number Theory*, pages 184–193. Springer, 1994. LNCS 877.
20. R. Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998.
21. C. Rössner and J.-P. Seifert. The complexity of approximate optima for greatest common divisor computations. In *Algorithmic Number Theory*, pages 184–193. Springer, 1996. LNCS 1122.