On Routing in Circulant Graphs

Jin-Yi Cai[1,⋆], George Havas[2,⋆⋆], Bernard Mans[3], Ajay Nerurkar[1,⋆⋆⋆], Jean-Pierre Seifert[4], and Igor Shparlinski[3,†]

[1] Department of Computer Science and Engineering,
State University of NY at Buffalo, Buffalo, NY 14260-2000, USA.
{cai,apn}@cs.buffalo.edu.

[2] Centre for Discrete Mathematics and Computing,
Department of Computer Science and Electrical Engineering,
The University of Queensland, Qld 4072, Australia.
havas@csee.uq.edu.au.

[3] Department of Computing,
Macquarie University, Sydney, NSW 2109, Australia.
{bmans,igor}@mpce.mq.edu.au

[4] Department of Mathematics and Computer Science,
University of Frankfurt, 60054 Frankfurt on the Main, Germany.
seifert@cs.uni-frankfurt.de

Abstract. We investigate various problems related to circulant graphs – finding the shortest path between two vertices, finding the shortest loop, and computing the diameter. These problems are related to shortest vector problems in a special class of lattices. We give matching upper and lower bounds on the length of the shortest loop. We claim NP-hardness results, and establish a worst-case/average-case connection for the shortest loop problem. A pseudo-polynomial time algorithm for these problems is also given. Our main tools are results and methods from the geometry of numbers.

Keywords: Circulant graphs, Shortest paths, Loops, Diameter, Lattices

1 Introduction

Circulant graphs have a vast number of applications in telecommunication networking, VLSI design and distributed computation [5,21,23]. We study various routing problems in circulant graphs such as finding the shortest path between two vertices, finding the shortest loop and the diameter. We establish relations between these routing problems and the problem of finding the shortest vector in the L_1-norm in a lattice.

We recall that an n-vertex *circulant graph* G is a graph whose adjacency matrix $A = (a_{ij})_{i,j=0}^{n-1}$ is a circulant. That is, the ith row of A is the cyclic shift

⋆ Supported in part by NSF grant CCR-9634665 and a J. S. Guggenheim Fellowship.
⋆⋆ Supported in part by ARC grants A49702337 and A49801415.
⋆⋆⋆ Supported in part by NSF grant CCR-9634665.
† Supported in part by ARC grant A69700294.

T. Asano et al. (Eds.): COCOON'99, LNCS 1627, pp. 360–369, 1999.

of the zeroth row by i, so that $a_{ij} = a_{0,j-i}$ for $i, j = 0, \ldots, n-1$ (adjacency matrix subscripts are taken modulo n). We consider undirected graphs, that is, $a_{ij} = a_{ji}$ for $i, j = 0, \ldots, n-1$ and there are no self-loops ($a_{ii} = 0$ for $i = 0, \ldots, n-1$).

Therefore with every circulant graph one can associate a set S of positive integers which shows which pairs of nodes are connected. Two nodes u and v of the graph are connected if and only if $\pm(u-v) \pmod n \in S$. Given S we denote the corresponding n-vertex graph by $\langle S \rangle_n$. We denote the group of residues modulo n by $\mathbb{Z}_n$. Without loss of generality, we always assume that $S \subseteq \mathbb{Z}_n$.

Let $S = \{s_1, \ldots, s_m\} \subseteq \mathbb{Z}_n$ with $1 \le s_i \le n/2$ for $i = 1, \ldots, m$. Given two vertices u and v of the graph $\langle S \rangle_n$, a path from u to v can be described by an integer vector $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbb{Z}^m$ such that there are x_i steps of the form $w \to w + s_i$ if $x_i \ge 0$, or there are $|x_i|$ steps of the form $w \to w - s_i$ if $x_i < 0$, for $i = 1, \ldots, m$. It is easy to see that any permutation of a sequence of steps in a path from u to v produces another path from u to v.

Starting from u we arrive at v if and only if $\sum_{i=1}^m x_i s_i \equiv v - u \pmod n$. It is natural to aim to minimize the L_1-norm $|\mathbf{x}|_{L_1}$ which is the length of the shortest path between u and v. For any given n and $S = \{s_1, \ldots, s_m\}$, this minimal number depends only on $u - v$. Accordingly for $w \in \mathbb{Z}_n$ we denote by $\Omega(w) = \Omega_{n,S}(w)$ the set of solutions $\mathbf{x} = (x_1, \ldots, x_m) \in \mathbb{Z}^m$ of the congruence $\sum_{i=1}^m x_i s_i \equiv w \pmod n$. Define $D_{n,S}(w) = \min\{|\mathbf{x}|_{L_1} : \mathbf{x} \in \Omega(w), \mathbf{x} \ne 0\}$. The condition $\mathbf{x} \ne 0$ is relevant only when $w = 0$ and in this case $L(n, S) = D_{n,S}(0)$ is the length of the shortest loop. We note that $D(n, S) = \max_{1 \le w \le n-1} D_{n,S}(w)$ is the diameter of $\langle S \rangle_n$.

Let $a_1, \ldots, a_n$ be linearly independent vectors in $\mathbb{R}^n$. The set of all integral linear combinations of the a_i forms a (n-dimensional) lattice L, and the a_i are called a basis of L. A lattice can also be abstractly defined as a discrete additive subgroup of $\mathbb{R}^n$. The *determinant* of a lattice L is denoted by $\det(L)$. If $a_1, \ldots, a_n$ is a basis of L, then $\det(L) = |\det(a_1, \ldots, a_n)|$. It is invariant under a change of basis. We remark that $\Omega(0)$ is a lattice and $\Omega(w)$ is a shifted (affine) lattice, and so we can apply existing tools for studying short vectors in lattices and their shifts.

We formulate the following problems.

Shortest-Path: Given a set $S = \{s_1, \ldots, s_m\} \subseteq \mathbb{Z}_n$ and a residue $w \in \mathbb{Z}_n$, $w \ne 0$, find $D_{n,S}(w)$ and a vector $\mathbf{x} \in \Omega(w)$ for which $D_{n,S}(w) = |\mathbf{x}|_{L_1}$.

Shortest-Loop: Given a set $S = \{s_1, \ldots, s_m\} \subseteq \mathbb{Z}_n$, find $L(n, S)$ and a vector $\mathbf{x} \in \Omega(0)$ for which $L(n, S) = |\mathbf{x}|_{L_1}$.

Diameter: Given a set $S = \{s_1, \ldots, s_m\} \subseteq \mathbb{Z}_n$, find $D(n, S)$.

For general graphs these are well known problems. Efficient polynomial time algorithms have been developed for various shortest path problems. However, for the class of circulant graphs, there is an important distinction to be made, and that concerns the natural input size to a problem. For a general graph it is common to consider that the input size is of order n^2, which is the number of elements in the adjacency matrix. However, any circulant graph can be described by only m integers $1 \le s_1, \ldots, s_m \le n/2$. In this representation the input size is of order $m \log n$. Thus polynomial time algorithms for general graphs may exhibit exponential complexity in the special case of circulant graphs.

Despite quite active interest in the above problems, motivated by various applications of circulant graphs, very few algorithmic results are known except for some special partial cases such as $m = 2$, see [10,11,19,30]. These papers use quite elementary number theoretic considerations. Thus one of the purposes of this paper is to introduce new techniques in this area, namely the relatively modern theory of geometric lattices [1,2,3,4,6,7,8,13,15,16,17,18,20,22,24,27,28,29] as well as more classical tools from the geometry of numbers [9] and combinatorial number theory [12,14].

First, we give estimates on the shortest loop length in terms of associated lattices. We provide upper and lower bounds which are close, within a factor of approximately two.

Then, we briefly discuss the case when the input n is given in unary. We provide an algorithm specifically tailored for the circulant graphs which solves the *Shortest-Loop* problem in polynomial time for this input measure. In contrast, we show that the *Shortest-Path* problem is NP-hard in the context of the more concise representation.

Finally, we claim that three lattice problems, which are believed to be hard, are reducible to an average-case instance of the *Shortest Vector Problem* **(SVP)** for the homogeneous lattice family defined by circulant graphs, under a certain distribution on such lattices.

2 Estimates of Shortest Vectors

In this section we give estimates for the shortest loop length $L(n, S)$ in terms of the lattice defined by the congruence $\sum_{i=1}^{m} a_i x_i \equiv 0 \pmod{n}$. For simplicity we focus on the case where the modulus n is a prime p. This can be generalized to a composite modulus via the Chinese Remainder Theorem. In fact taking the lattice view, it is natural to consider the more general setting where a set of s congruences is given, $A\mathbf{x} \equiv 0 \pmod{p}$ with $\mathbf{x} \in \mathbb{Z}^m$, where $A = (a_{ij}) \in \mathbb{Z}_p^{s \times m}$, which defines a lattice L of dimension m. The case with the circulant graph is $s = 1$. This lattice is of dimension m, independent of s, since $p \cdot e_i \in L$, for all i, where e_i is the ith canonical basis vector (for $\mathbb{Z}^m$). Note that in general $p \cdot e_i$ do not form a basis for L.

We now compute the determinant of L. The map $\mathbf{x} \mapsto A\mathbf{x} \pmod{p}$ defines a homomorphism from $\mathbb{Z}^m$ to $\mathbb{Z}_p^s$, the kernel of which is our lattice L. Thus the cardinality $|\mathbb{Z}^m/L| \le p^s$ (equality holds if and only if A has rank m over $\mathbb{Z}_p$). Hence $\det(L) \le p^s$, with equality if and only if A has rank m over $\mathbb{Z}_p$.

Let $B_m(r) = \{\mathbf{x} \in \mathbb{R}^m \; : \; |\mathbf{x}|_{L_1} \le r\}$ be the L_1-norm ball of radius r in m-dimensional space. Then the volume $\mathrm{vol}(B_m(r)) = r^m \cdot \frac{2^m}{m!}$.

By Minkowski's Theorem, see Theorem 2 of Chapter 3 of [9], if $\mathrm{vol}(B_m(r)) \ge 2^m \det(L)$, then there is a non-zero vector $x \in L \cap B_m(r)$. Because $2^m p^s \ge 2^m \det(L)$, a sufficient condition on r is $r \ge \sqrt[m]{m!} p^{s/m} \approx \frac{m}{e} p^{s/m}$. (The approximation is based on $\sqrt[m]{m!} = \frac{m}{e}(1 + o(1))$ as $m \to \infty$.) Thus we have

Theorem 1. *For any system* $A\mathbf{x} \equiv 0 \pmod{p}$, *where* $A \in \mathbb{Z}_p^{s\times m}$, *there is a solution* $\mathbf{x} \in \mathbb{Z}^m$, $\mathbf{x} \neq 0$, *and* $|\mathbf{x}|_{L_1} \leq \sqrt[m]{m!}p^{s/m}$. *In particular, for any circulant graph* $\langle S\rangle_p$, $L(p,S) \leq \sqrt[m]{m!}p^{1/m}$, *where* $m = |S|$.

Let $N_m(r) = |\{\mathbf{x} \in \mathbb{Z}^m : x_i \geq 0, \text{ and } x_1 + \cdots + x_m \leq r\}|$. It is elementary to show that $N_m(r) = \binom{r+m}{m}$. Thus $\mathbb{Z}^m \cap B_m(r) = \{\mathbf{x} \in \mathbb{Z}^m : |\mathbf{x}|_{L_1} \leq r\}$ has cardinality at most $2^m\binom{r+m}{m}$.

Now to each non-zero $\mathbf{x} \in \mathbb{Z}^m \cap B_m(r)$, there are exactly $p^{m-1} - 1$ non-zero coefficient sequences $(a_1, \ldots, a_m) \in \mathbb{Z}_p^m$ for which $\sum_{i=1}^m a_i x_i \equiv 0 \pmod{p}$. Altogether this can account for at most $(2^m\binom{r+m}{m} - 1)(p^{m-1} - 1)$ non-zero sequences $(a_1, \ldots, a_m) \in \mathbb{Z}_p^m$. Thus if $p^m - 1 > (2^m\binom{r+m}{m} - 1)(p^{m-1} - 1)$, then some $(a_1, \ldots, a_m) \neq 0$ has no solution with norm at most r. The following bound is certainly a sufficient condition for this: $p \geq 2^m\binom{r+m}{m}$. Since $\binom{r+m}{m} \leq ((r+m)e/m)^m$, after some simple calculation, we have the sufficient condition $r < m\left(\frac{1}{2e}p^{1/m} - 1\right)$.

Theorem 2. *There are circulant graphs* $\langle S\rangle_p$ *with* $L(p,S) \geq m\left(\frac{1}{2e}p^{1/m} - 1\right)$.

We note that this lower bound almost matches the upper bound in Theorem 1, within a factor of about 2. Also the analysis here can be given in terms of a matrix $A \in \mathbb{Z}_p^{s\times m}$, in which case, we replace the quantity $p^{1/m}$ by $p^{s/m}$.

Unfortunately the above considerations do not seem to apply to the non-homogeneous case. Nevertheless, some number theoretic results of [12,14] allow us to deal with this case as well. In particular, we obtain an upper bound on $D(n,S)$, which, in some cases, is quite tight.

Theorem 3. *For any circulant graph* $\langle S\rangle_p$ *with* $m \geq 0.5\left\lfloor(4p-7)^{1/2}\right\rfloor + 1$, $D(p,S) \leq 0.5\left\lfloor(4p-7)^{1/2} + 1\right\rfloor \sim p^{1/2}$, *where* $|S| = m$.

Proof. For any set $T \subseteq \mathbb{Z}_p^*$ of cardinality $|T| = \left\lfloor(4p-7)^{1/2}\right\rfloor + 1$, it is shown in [12] that any element $w \in \mathbb{Z}_p$ can be represented in the form $w \equiv \sum_{t\in W} t \pmod{p}$ for some subset $W \subseteq T$ of cardinality at most $|T|/2$. It is obvious that the cardinality of the set $R = S \cup -S$ is at least $2m - 1$. Thus selecting an arbitrary subset $T \subseteq R$ of cardinality $|T| = \left\lfloor(4p-7)^{1/2}\right\rfloor + 1 \leq 2m - 1$ we obtain the required result. □

We remark that if $m \sim p^{1/2}$ and $S = \{1, \ldots, m\}$ then, obviously, $D(p,S) \geq p/m \sim p^{1/2}$, thus Theorem 3 is tight for such circulant graphs.

3 Polynomial Time Algorithms for Small n

When n is given as a unary input and the time complexity is measured in terms of n, there are well known polynomial time algorithms to find both the length of the shortest path between any two vertices of a graph with n vertices, as well as a shortest path itself. This applies to circulant graphs as a special case. However, even in this case, due to the symmetry present in any circulant graph, considerable savings can be realized in time complexity, compared to using a general graph algorithm.

We concentrate on the *Shortest-Loop* problem. The definition of a *loop* here is specifically tailored for the circulant graph, and is not merely a specialization of a loop in general graphs. This is because there is a "commutativity" of the underlying group $\mathbb{Z}_n$, and so we exclude certain "trivial loops". For example, suppose there are two distinct step sizes s_1 and $s_2 \in S$. The closed path $\langle 0, s_1, s_1+s_2, s_2, 0\rangle$, using steps $s_1, s_2, -s_1, -s_2$ respectively, would be considered a loop in a general graph, but is excluded here.

Assume we are given $S = \{s_1, \ldots, s_m\} \subseteq \mathbb{Z}_n$. The idea of the following algorithm for finding loops is to try to compute all shortest paths connecting 0 and some $is_k \pmod n$, using only steps from $S' = \{s_1, \ldots, s_{k-1}\}$. We do this for all i and then for all k, and finally take the minimum.

Fix k, $1 \le k \le m$. Let $g_k = \gcd(s_k, n)$. Let $s'_k = s_k/g_k, n_k = n/g_k$. Then $\gcd(s'_k, n_k) = 1$, and we can solve integral linear equations $s'_k x - n_k y = \pm 1$. The general solution has the form $x = \pm[x_0 + tn_k]$ and $y = \pm[y_0 + ts'_k]$, for some x_0 and y_0, and where all $t \in \mathbb{Z}$ are admissible. Choose t so that the absolute value $|x| = |x_0 + tn_k|$ is minimum. This will set $|x| \le \lfloor n_k/2 \rfloor$ and the corresponding t is essentially unique. We denote this minimizing x by x_k.

Next, for each j, $1 \le j \le \lfloor n_k/2 \rfloor$, we compute the shortest path from 0 to $js_k \bmod(n)$, in the graph $\langle S' \rangle_n$ defined by $S' = \{s_1, \ldots, s_{k-1}\}$. The minimum path length in this graph is $D_{n,S'}(js_k)$.

$$\text{Let } f(i,k) = \begin{cases} j + D_{n,S'}(js_k) & \text{if } i \equiv \pm js_k \pmod n, 1 \le j \le \lfloor n_k/2 \rfloor, \\ \infty & \text{otherwise, that is, } g_k \nmid i, \end{cases} \tag{1}$$

where $1 \le i \le n-1$, and let $f(0,k) = n_k$.

Note that, for $1 \le i \le n-1$, $g_k = \gcd(s_k, n) \mid i$ if and only if there is some j, where $1 \le j \le \lfloor n_k/2 \rfloor$, such that $i \equiv \pm js_k \pmod n$. Moreover, in this case, such a $\pm j$ is unique, (except in the case where n_k is even and $i = \frac{n_k}{2} s_k$, in which case there are exactly two values $+\frac{n_k}{2}$ and $-\frac{n_k}{2}$.)

To see this, one direction is trivial: if $i \equiv \pm js_k \pmod n$ for some j, then $g_k = \gcd(s_k, n) \mid i$. Now suppose $g_k \mid i$. Then i belongs to the subgroup generated by g_k in $\mathbb{Z}_n$, which is also generated by s_k. This subgroup has order n_k. Hence there is some j within the specified range $1 \le j \le \lfloor n_k/2 \rfloor$, such that

$$i \equiv \pm js_k \pmod n. \tag{2}$$

$j \ne 0$ since $i \not\equiv 0 \pmod n$. In the general case when $g_k \mid i$, except when n_k is even and $i = \frac{n_k}{2} s_k$, the uniqueness of such a $\pm j$ is quite obvious. For otherwise we have $((\pm j) - (\pm j'))s_k \equiv 0 \pmod n$ which implies that $n_k \mid ((\pm j) - (\pm j'))$. And the only possibility for this to hold is $j = j' = n_k/2$ but they took opposite signs in (2), and $i = \frac{n_k}{2} s_k$. Indeed in this case the $\pm j$ in (2) is not unique, but nonetheless the value of $f(i,k)$ is uniquely defined.

This shows that the function $f(i,k)$ is well defined, and given any algorithm for finding shortest paths, this gives us an algorithm to compute $f(i,k)$. In particular this gives a polynomial time algorithm for $f(i,k)$.

Now we claim that the minimum loop size

$$L(n, S) = \min_{0 \le i \le n-1, 1 \le k \le m} f(i, k). \tag{3}$$

Clearly the minimum value on the right in (3) produces the length of some loop. Since all $s_k \ne 0$, using s_k alone, one can form a loop with exactly $f(0, k) = n_k$ steps (and no less). Thus the minimum is finite. In the case $1 \le i \le n - 1$, the value $f(i, k)$ gives the length of some loop, if such a loop exists, that consists of step sizes $s_1, \ldots, s_k$ only, with a non-zero number of step size s_k.

Let $\mathcal{L}$ be a loop achieving $L(n, S)$. We show that our minimum value on the right of (3) is at least as small. This establishes (3).

If $\mathcal{L}$ involves only one step size s_k, then $f(0, k) = n_k$ gives the shortest loop in this case. Let $\mathcal{L}$ involve more than one step size, and let k be the largest ℓ such that $\mathcal{L}$ involves a non-zero number x_ℓ of steps of size s_ℓ. Since $\mathcal{L}$ is minimal, $s_k x_k \not\equiv 0 \pmod{n}$, for otherwise, by deleting the steps involving s_k we still have a non-trivial loop.

Let i be the least modulus of $-s_k x_k \pmod{n}$, where $1 \le i \le n - 1$, then $\sum_{\ell=1}^{k-1} s_\ell x_\ell \equiv i \pmod{n}$, and x_ℓ is the number of steps in $\mathcal{L}$ involving s_ℓ. Since the subgroup of $\mathbb{Z}_n$ generated by s_k has order n_k, there is a j, $1 \le j \le \lfloor n_k/2 \rfloor$, such that $-s_k x_k \equiv \pm j s_k \pmod{n}$. This implies that $-x_k \equiv \pm j \pmod{n_k}$. Since $|\mathbf{x}|_{L_1}$ is minimal, $x_k = \pm j$.

After taking away the steps involving s_k, what remains is a path from 0 to $i \equiv -s_k x_k \pmod{n}$ using only steps $s_1, \ldots, s_{k-1}$, the minimum length of which has been found in the equation (1) defining $f(i, k)$.

This completes the description of our algorithm to compute $L(n, S)$. Now, we estimate its time complexity. Let $\langle S_l \rangle_n$ be the graph defined by $S_l = \{s_1, \ldots, s_l\}$, where $1 \le l \le m$. To compute $f(i, k)$, where $i \equiv \pm j s_k \pmod{n}$, we need to know the value of $D_{n,S_{k-1}}(j s_k)$. We first do a breadth-first search on $\langle S_l \rangle_n$, for $l = 1, \ldots, m$, to compute the length of the shortest path from 0 to every other vertex in each of these graphs. This can be done in a total time of $O(nm^2)$. The rest of the computation takes time $O(nm \log^{O(1)} n)$. Therefore we have the following result

Theorem 4.
The Shortest-Loop Problem is solvable in time $O(nm^2 + nm \log^{O(1)} n)$.

A more careful design of the algorithm can reduce the time to $O(nm)$.

4 NP-hardness for the *Shortest-Path* Problem

In this section we claim that the *Shortest-Path* problem is NP-hard. We can also show that this problem is NP-hard to approximate within a factor $n^{1/\log\log n}$ using the result of [13] and its recent improvement due to Ran Raz [26]. This will be discussed in the full paper, which will also contain omitted proofs. We first state a decision problem version of the *Shortest-Path* problem. This problem is NP-complete, and the NP-hardness of the optimization problem follows easily.

Shortest-Path Decision Problem(SPDP)
Instance: Given in binary an integer n, a set $S = \{s_1, \ldots, s_m\} \subseteq \mathbb{Z}_n$, a residue $w \in \mathbb{Z}_n$, $w \neq 0$, and a bound b.
Question: Is $D_{n,S}(w) \leq b$?
Size: $\text{length}(n) + \sum_{i=1}^{m} \text{length}(s_i) + \text{length}(w) + \text{length}(b)$.

Theorem 5. *The Shortest-Path Decision Problem (SPDP) is NP-complete.*

The proof is by reducing the well-known Exact-3-Cover NP-complete decision problem to the Shortest-Path Decision Problem.

5 The Hardness for Homogeneous Systems

In the last section we claimed that the optimization problem of finding $D_{n,S}(w)$ for a general right hand side w is NP-hard. This corresponds to the SVP in a special class of *affine* lattices. The homogeneous case where $w = 0$ is the SVP in a special class of *homogeneous* lattices, namely those definable by a single congruence of the form $\sum_{i=1}^{m} s_i x_i \equiv 0 \pmod{n}$. In this section, we focus on a worst-case/average-case connection.

There has been a considerable amount of work on the SVP for homogeneous lattices. Van Emde Boas [29] showed that it is NP-hard for the L_∞-norm (see also [20]). The NP-hardness for every other L_p-norm, especially for the L_2-norm, was open until a recent breakthrough by Ajtai [2], who showed that the problem for the L_2-norm is NP-hard under randomized reductions. Moreover Ajtai [2] also showed that it is NP-hard to approximate the shortest non-zero vector in an m-dimensional lattice within a factor of $(1 + 2^{-m^k})$, for a sufficiently large constant k. This was improved by Cai and Nerurkar [8] to a factor of $\left(1 + \frac{1}{m^\epsilon}\right)$, for any $\epsilon > 0$. They [8] also noted explicitly that the proof works for arbitrary L_p-norms, $1 \leq p < \infty$. Micciancio further improved this factor to $2^{1/p} - \epsilon$ for the L_p-norm, $p \geq 1$ [24]. Approximations to the closest vector in a lattice were considered in [4,13]. However, in all these reductions, the lattices constructed do not fall into the category of lattices defined in terms of circulant graphs, that is, definable by a single congruence. Nevertheless, we claim that the SVP for such lattices is, in fact, NP-hard under randomized reductions. We do so by reducing the general SVP to it. Such a reduction was given previously by Paz and Schnorr [25] but our result uses more elementary ideas. We have the following theorem:

Theorem 6. *The SVP for the class of (homogeneous) lattices defined by circulant graphs is NP-hard under randomized reductions. Moreover, it remains NP-hard to find an approximate shortest vector in this class of lattices within a factor of $2 - \epsilon$, for any constant $\epsilon > 0$.*

Ajtai, in a separate work [1], established a reduction from the worst-case complexity of some problems believed to be hard, to the average-case complexity of the approximate SVP for a certain class of random lattices. This worst-case/average-case connection is the only such provable reduction known for a problem in NP believed to be hard and has generated a lot of interest [3,6,7,15,16,17,18]. We state three problems which are believed to be hard for any constants c_1, c_2 and c_3.

(P1) Find $\lambda_1(L)$, the length of a shortest non-zero vector in an m-dimensional lattice L, up to a polynomial factor m^{c_1}.

(P2) Find the shortest non-zero vector in an m-dimensional lattice, where the shortest vector is unique up to a polynomial factor m^{c_2}.

(P3) Find a basis in an m-dimensional lattice whose maximum length is the smallest possible, up to a polynomial factor m^{c_3}.

Let $\mathbb{Z}_q^{r\times l}$ denote the set of $r \times l$ matrices over $\mathbb{Z}_q$. For every r, l, q, $\Omega_{r,l,q}$ denotes the uniform distribution on $\mathbb{Z}_q^{r\times l}$. For every $X \in \mathbb{Z}_q^{r\times l}$, the set $\Lambda(X) = \{y \in \mathbb{Z}^l \mid Xy \equiv 0 \pmod{q} \}$ defines a lattice of dimension l. $\Lambda_{r,l,q}$ denotes the probability space of lattices consisting of $\Lambda(X)$ by choosing X according to $\Omega_{r,l,q}$. Let $q = r^c$ be an arbitrary but fixed polynomial of r. By Minkowski's Theorem, see Theorem 2 of Chapter 3 of [9], it can be proved that, $\forall c\ \exists c'$ such that $\forall \Lambda(X) \in \Lambda_{r,c'r,r^c}$, $\lambda_1(\Lambda(X)) \le r$. (In fact, this bound can be improved to $\Theta(r^{1/2+\epsilon})$.) Let $l = c'r$ where c' depends on c as indicated above. Let $\Lambda = \Lambda_{r,l,q}$. Ajtai [1] showed:

Theorem 7. *Let γ be any constant. If there is a probabilistic polynomial time algorithm $\mathcal{A}$ such that, with non-trivial probability $1/r^{O(1)}$, $\mathcal{A}$ finds a short non-zero vector v, $|v|_{L_2} \le r^\gamma$, for a uniformly chosen lattice in the class Λ indexed by r, then there is a probabilistic polynomial time algorithm $\mathcal{B}$ that solves the three problems* **(P1), (P2)** *and* **(P3)** *in the worst case, with high probability, for some constants c_1, c_2 and c_3.*

We claim here that these problems are also reducible to an average-case instance of the SVP for the lattice family defined by circulant graphs, under a certain distribution on such lattices. This is not quite an NP-hardness proof, since these problems **(P1), (P2)** and **(P3)** are not known to be NP-hard. In fact there is evidence that they are not. Goldreich and Goldwasser showed that approximating the shortest lattice vector within a factor of $O(\sqrt{m/\log m})$ is not NP-hard assuming the polynomial time hierarchy does not collapse [15]. Cai showed that finding an $m^{1/4}$-unique shortest lattice vector is not NP-hard unless the polynomial time hierarchy collapses [6]. On the other hand, these problems have resisted all attempts at finding a (probabilistic) polynomial time algorithm to date. The best polynomial time algorithm is essentially the L^3 basis reduction algorithm [22] which achieves an approximation factor of $2^{m/2}$ for the SVP. Schnorr improves this factor to $(1+\epsilon)^m$, still exponential in m and with the running time depending badly on ϵ in the exponent [27]. Thus it is reasonable to assume that these problems **(P1), (P2)** and **(P3)** are computationally hard.

Our proof establishes that for a certain distribution of circulant graphs, or equivalently for the special class of lattices, the average-case complexity of approximating the shortest vector in a lattice within any polynomial factor is at least as hard as these three problems in the worst-case. We use a construction that reduces the problem of finding a short vector in a lattice $\Lambda(X)$ to the problem of finding a short solution vector to an appropriate congruence.

Theorem 8. *Let α be any number. Given a vector $\mathbf{v} \in \Lambda(X)$, $0 < |\mathbf{v}|_{L_1} \leq \alpha$, a solution vector $\mathbf{x}$ to the congruence $\sum_{i=1}^{l+r} s_i x_i \equiv 0 \pmod n$, with $0 < |\mathbf{x}|_{L_1} \leq (r+1)\alpha$, can be computed in polynomial time, and given a solution $\mathbf{x}$ to the congruence, $0 < |\mathbf{x}|_{L_1} \leq \beta$, a vector $\mathbf{v} \in \Lambda(X)$ with $0 < |\mathbf{v}|_{L_1} \leq \beta$, can be computed in linear time.*

We now apply this reduction to show that, to solve a congruence of the form $\sum_{i=1}^{r} s_i x_i \pmod n$ is as hard on the average, under a suitable distribution on the s_i and n, as problems **(P1)**, **(P2)** and **(P3)** are in the worst case. We get the following worst-case/average-case connection.

Theorem 9. *Let $\gamma > 2.5$ be any constant in Theorem 7 and let $l = \Theta(r)$ as above. If there is a probabilistic polynomial time algorithm $\mathcal{C}$ such that, with non-trivial probability $1/r^{O(1)}$, $\mathcal{C}$ finds a solution vector x of L_1-norm length at most r^γ for the congruence $\sum_{i=1}^{l+r} s_i x_i \equiv 0 \pmod n$, where the s_i and n are chosen according to distribution D_{r,r^γ}, then there is a probabilistic polynomial time algorithm $\mathcal{D}$ that solves the three problems* **(P1)**, **(P2)** *and* **(P3)** *in the worst case, with high probability, for some constants c_1, c_2 and c_3.*

References

1. M. Ajtai. Generating hard instances of lattice problems. In *Proc. 28th Annual ACM Symposium on the Theory of Computing* (1996) 99–108.
2. M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions. In *Proc. 30th Annual ACM Symposium on the Theory of Computing* (1998) 10–19.
3. M. Ajtai and C. Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proc. 29th ACM Symposium on Theory of Computing* (1997) 284–293.
4. S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. In *Proc. 34th IEEE Symposium on Foundations of Computer Science (FOCS)*, 724-733, 1993.
5. J.-C. Bermond, F. Comellas and D. F. Hsu. Distributed loop computer networks: A survey. *Journal of Parallel and Distributed Computing* **24** (1995) 2–10.
6. J-Y. Cai. A relation of primal-dual lattices and the complexity of shortest lattice vector problem. *Theoretical Computer Science* **207** (1998) 105–116.
7. J-Y. Cai and A. Nerurkar. An improved worst-case to average-case connection for lattice problems. In *Proc. 38th IEEE Symposium on Foundations of Computer Science* (1997) 468–477.
8. J-Y. Cai and A. Nerurkar. Approximating the SVP to within a factor $\left(1 + \frac{1}{\dim^\epsilon}\right)$ is NP-hard under randomized reductions. In *Proc. 13th Annual IEEE Conference on Computational Complexity* (1998) 46–55.
9. J. W. S. Cassels. *An introduction to the geometry of numbers*. Springer-Verlag, 1959.
10. N. Chalamaiah and B. Ramamurty. Finding shortest paths in distributed loop networks. *Inform. Proc. Letters* **67** (1998) 157–161.
11. Y. Cheng and F. K. Hwang. Diameters of weighted loop networks. *J. of Algorithms* **9** (1988) 401–410.

12. J. A. Dias da Silva and Y. O. Hamidoune. Cyclic spaces for Grassmann derivatives and additive theory *Bull. Lond. Math. Soc.* **26** (1994) 140–146.
13. I. Dinur, G. Kindler and S. Safra. Approximating-CVP to within almost-polynomial factors is NP-hard. 1998.
14. P. Erdös and H. Heilbronn On the addition of residue classes mod p *Acta Arithm.* **9** (1964) 149–159.
15. O. Goldreich and S. Goldwasser. On the limits of non-approximability of lattice problems. In *Proc. 30th Annual ACM Symposium on the Theory of Computing* (1998) 1–9.
16. O. Goldreich, S. Goldwasser and S. Halevi. Collision-free hashing from lattice problems. 1996. Available as TR96-042 from *Electronic Colloquium on Computational Complexity* at `http://www.eccc.uni-trier.de/eccc/`.
17. O. Goldreich, S. Goldwasser and S. Halevi. Eliminating decryption errors in the Ajtai-Dwork cryptosystem. *Advances in Cryptology - CRYPTO '97* (editor B. Kaliski Jr.), Lecture Notes in Computer Science 1294 (Springer Verlag, 1997) 105–111.
18. O. Goldreich, S. Goldwasser and S. Halevi. Public-key cryptosystems from lattice reduction problems. Advances in Cryptology - CRYPTO '97 (editor B. Kaliski Jr.), Lecture Notes in Computer Science 1294 (Springer Verlag, 1997) 112–131.
19. D. J. Guan. Finding shortest paths in distributed loop networks. *Inform. Proc. Letters* **65** (1998) 255–260.
20. J. C. Lagarias. The computational complexity of simultaneous diophantine approximation problems. In *Proc. 23rd IEEE Symposium on Foundations of Computer Science* (1982) 32–39.
21. F. T. Leighton. *Introduction to parallel algorithms and architectures: Arrays, trees, hypercubes*. M. Kaufmann, 1992.
22. A. K. Lenstra, H. W. Lenstra and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen* **261** (1982) 515–534.
23. B. Mans. Optimal Distributed algorithms in unlabeled tori and chordal rings. *Journal on Parallel and Distributed Computing* **46** (1997) 80–90.
24. D. Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. In *Proc. 39th IEEE Symposium on Foundations of Computer Science* (1998) 92–98.
25. A. Paz and C. P. Schnorr. Approximating integer lattices by lattices with cyclic factor groups. Automata, Languages and Programming, 14th International Colloquium, *Lecture Notes in Computer Science* 267 (Springer-Verlag, 1987) 386-393.
26. R. Raz. Personal communication.
27. C. P. Schnorr. A hierarchy of polynomial time basis reduction algorithms. *Theoretical Computer Science* **53** (1987) 201–224.
28. J-P. Seifert and J. Blömer. On the complexity of computing short linearly independent vectors and short bases in a lattice. To appear in the proceedings of STOC 1999.
29. P. van Emde Boas. Another NP-complete partition problem and the complexity of computing short vectors in lattices. *Technical Report 81-04*, Mathematics Department, University of Amsterdam, 1981.
30. J. Žervonik and T. Pisanski. Computing the diameter in multi-loop networks. *J. of Algorithms* **14** (1993) 226–243.