1 Fractured Lattices, Integer Programming, and Diophantine Approximation*

dedicated to Prof. Dr. Rainer Burkard (Graz)
on the occasion of his 60th birthday

Günter Rote

Institut für Informatik, Freie Universität Berlin, Takustraße 9, D-14195 Berlin, Germany, *rote@inf.fu-berlin.de*

Abstract. For a d-dimensional lattice Λ and a d-dimensional vector a we consider the sequence of point sets

$$A_n := \{\, ia + x \mid x \in \Lambda,\ 0 \le i < n \,\}$$

for increasing values of n. For the values of n when a new shortest nonzero vector appears in A_{n+1}, the set A_n has a structure of a *perturbed lattice*, i. e., each point is in some small neighborhood of a (unique) lattice. We use this structure for a recursive approach to finding best approximations in fixed dimensions, with applications to integer programming.

2000 Mathematics Subject Classification: 11K60
Key words: best simultaneous Diophantine approximation, geometry of numbers.

1.1 Introduction

Integer Programming is of course closely connected to integer lattices and the geometry of numbers. Especially the algorithms in small dimensions which have theoretical performance guarantees are based on shortest vectors and good bases for integer lattices or at least approximations of them. The most notable instance is the polynomial-time algorithm of H. W. Lenstra [9] for integer programming in fixed dimensions.

For integer programming in two dimensions, there are a number of different algorithms, the nicest (from a geometric viewpoint) being perhaps the algorithm of Kanamuru, Nishizeki and Asano [5]. All these algorithms, in one way or another, eventually boil down to some sort of continued fraction expansion or greatest common divisor computation. This is no surprise, since the greatest common divisor of two numbers a and b can be formulated as an integer programming problem

$$\min\{\, ax + by \mid ax + by \ge 1,\ x, y \in \mathbb{Z} \,\}$$

On the other hand, it has been shown that integer programming in two variables is no more difficult than greatest common divisor computation, at least as far as the "number-theoretic" complexity is concerned, the dependence on the size of the input numbers. (This does not account for the "combinatorial" (or "geometric") complexity due to the fact that the problem may have a large number of constraints and the feasible region may be a complicated polytope, in higher dimensions.)

Eisenbrand and Rote [3] introduced a novel parametric approach to integer programming in two variables. They reduced it to the following key problem.

> *The Parametric Shortest Vector Problem.*
> Let a two-dimensional lattice Λ and a parameter ε be given. Find the smallest factor l such that scaling the y-coordinate by l produces a lattice whose shortest vector (in the max-norm) has length ε.

Of course, the problem has no solution if the x-axis contains a vector shorter than ε. Otherwise, it is easy to see that l can be calculated by solving the following problem, see Figure 1.1.

* submitted for publication to *Monatshefte für Mathematik*, December 2002

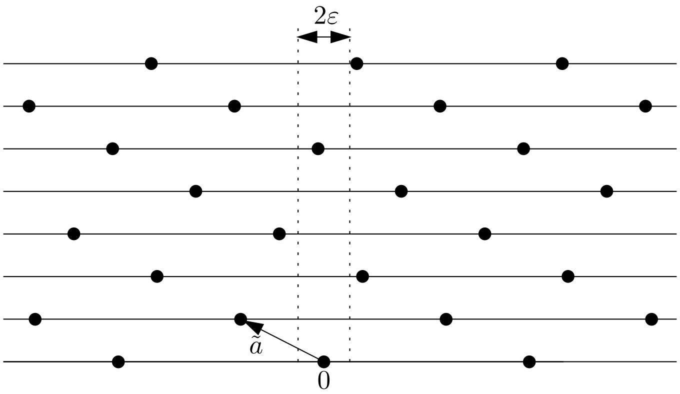

Fig. 1.1. The Best Approximation Problem for a two-dimensional lattice Λ

> Let a two-dimensional lattice Λ be given. Find the lowest lattice point of the upper half-plane inside a given vertical strip of width 2ε that is centered at the origin.

In [3], the problem is solved by direct application of an algorithm of Schönhage [12] for best rational approximation.

This paper arose from an effort to extend this approach to higher dimensions. We will discuss the higher-dimensional version of the problem and derive some structural properties. We will also give an algorithm, although its running time is not competitive with the best other algorithms [4].

We formulate the higher-dimensional problem a follows.

> The *Best Approximation Problem.*
> Let a $(d+1)$-dimensional lattice $\tilde{\Lambda} \subset \mathbb{R}^{d+1}$ and an error tolerance ε be given. Find the lattice point with smallest positive x_{d+1}-coordinate that lies within a distance ε of the x_{d+1}-axis.

The distance from the x_{d+1}-axis can be measured by some arbitrary norm, but we will stick mostly to the usual Euclidean norm

This problem is just a reformulation of the classical Diophantine approximation problem: We want to approximate a set of d numbers $\alpha_1, \dots, \alpha_d$ by rational numbers $x_1/Q, \dots, x_d/Q$ with a common denominator Q. If we multiply the approximating numbers by Q, we see that this problem amounts to finding an integer vector $(x_1, \dots, x_d, Q) \in \mathbb{Z}^d$ that lies close to the line generated by the vector $(\alpha_1, \dots, \alpha_d, 1)$. We just have to transform this line into the x_{d+1}-axis leaving the other axes fixed, and measure the "distance" from the axis appropriately, namely in the max-norm, to get our formulation of the Best Approximation Problem.

We can split the lattice into layers as follows, see Figure 1.1. We assume that the sublattice $\Lambda := \{\, x \in \tilde{\Lambda} \mid x_{d+1} = 0 \,\}$ is d-dimensional. This is no loss of generality if $\tilde{\Lambda}$ consist of rational points. Let $\tilde{a}$ be a vector which, together with Λ, generates $\tilde{\Lambda}$. Then $\tilde{\Lambda}$ consists of horizontal *level* $\Lambda + i\tilde{a}$, for $i \in \mathbb{Z}$. We write $\tilde{a} = (a, a_{d+1})$ with $a \in \mathbb{R}^d$, and assume $a_{d+1} > 0$. Then the Best Approximation Problem amounts to the following problem:

> Find the lowest level $\Lambda + i\tilde{a}$ $(i \geq 1)$ which contains a point within ε of the x_{d+1}-axis.

Since the x_{d+1}-coordinate is irrelevant for measuring the distance from the x_{d+1}-axis, we obtain the following equivalent formulation in $\mathbb{R}^d$.

> Find the smallest $i \geq 1$ such that $\Lambda + ia$ contains a point within ε of the origin.

We denote the union of the first $n-1$ layers by

$$A_n := \{\, ia + x \mid x \in \Lambda,\ 0 \le i < n \,\}$$

Hence we arrive at the following problem formulation, which is the one we will use in the paper.

> Find the smallest $n \geq 1$ such that A_n contains a nonzero point of length less than ε.

(This formulation is not completely equivalent, because Λ might already contain a point of length at most ε, or the desired short point at level i might be the zero vector. These cases have to be checked separately.)

1.1.1 Related Work

The Shortest Vector Problem. We solve the Best Approximation Problem in a sequence of rounds, repeatedly increasing n and looking for a vector in A_n that is shorter than the previously found shortest vector. Thus, during the course of the algorithm, we will also find the shortest non-zero vector of $\tilde{\Lambda}$: Either this vector lies in Λ, or it lies in one of the other levels $\Lambda + ai$. In that level, it must clearly be the closest point to the x_{d+1}-axis, and there can be no point which is as close to the x_{d+1}-axis in any of the levels below the i-th level. The process can be stopped as soon as the distance from the hyperplane of the current level to the origin is larger than the shortest vector found so far.

The shortest vector problem, and the related problem of finding a good lattice basis, are important for other areas like algebraic number theory or the analysis of pseudo-random number generators. The best general algorithm in arbitrary dimension d is the algorithm of Lovász [8], which computes in polynomial time a vector which is at most $2^{\frac{d-1}{2}}$ times longer than the shortest vector. This algorithm starts with some lattice basis, and repeatedly performs lattice reductions in the two-dimensional lattice spanned by pairs of basis vectors.

Schnorr [11] has extended this algorithm so that it looks at k-tuples of vectors and reduces the corresponding k-dimensional sublattices. The algorithm runs in a polynomial number of iterations for fixed k and achieves an approximation factor of $(6k^2)^{\frac{d}{2k}}$. Thus, any progress in lattice reduction algorithms for fixed dimension k has an impact on algorithms for arbitrary dimensions.

In fixed dimension, the problems of best approximation, of computing a short lattice vector, of computing a "good" lattice basis or even of enumerating all reduced lattice basis (in a very weak sense) are all computationally equivalent, see [4, Section 6]. There running time is within a constant factor of each other.

Continued Fractions. Multi-dimensional continued fraction are also related to best approximations, see [1] for an overview. However, in this area, it is customary to define some procedure (like the classical Jacobi-Perron algorithm), and then try to analyze its behavior. In contrast to this, the algorithm that we describe is found by specifying its behavior clearly in geometric terms.

Quasiperiodic tilings. From the above description, the set A_n arises as a section of the $(d+1)$-dimensional lattice $\tilde{\Lambda}$ between two parallel hyperplanes, projected to a d-dimensional space. Such projections of slices of higher-dimensional lattices are also used to generate quasiperiodic point sets, tilings, and quasicrystals by the cut-and-project method, see [13,10].

1.1.2 Overview.

In Section 1.2 we will discuss how the set A_n evolves as n increases. In particular, we will be interested in the points when a new shortest nonzero vector appears, because this is a candidate for the solution of the Best Approximation Problem. In particular, we will define the structure of a *fractured lattice.* In Section 1.3 we will use the properties that we have found for defining an algorithm that solves the Best Approximation Problem recursively by reducing it to lower-dimensional problems.

1.2 The Structure of A_n

In the following, the concepts will be illustrated with examples in the plane, but unless otherwise stated, they hold in any dimension. We will use the language that is appropriate for the planar case; for example, we will speak of disks when balls would be the proper term for higher dimensions.)

Figure 1.2 shows the set A_{27} for the two-dimensional lattice Λ generated by the vectors $(1, 0)$ and $(0, 1)$ and for the vector a indicated in the figure. Since the point set is periodic with the period of Λ, we only show the range $[-0.5, 0.5]^2$ around the origin. The last point added is point 26, and it is the nearest neighbor of the origin among the points added so far.

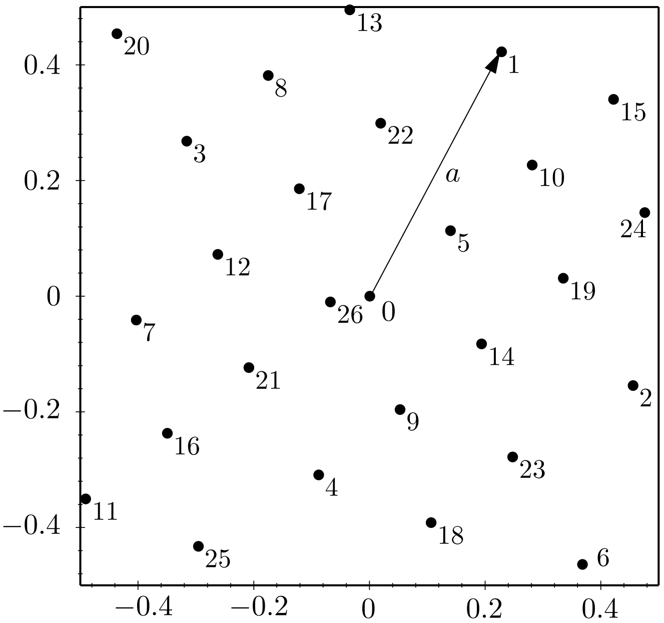

Fig. 1.2. The point set A_{27}. The point labeled i is $ia \bmod \Lambda$

We will regard A_n as a set of n vectors $0, a, 2a, \ldots, (n-1)a$ modulo Λ. The vector $x \bmod \Lambda$ will denote the congruence class $x + \Lambda$ and we can perform vector additions and multiplications by integer scalars with these congruence classes as with usual vectors. We denote congruence by $x \equiv y \pmod{\Lambda}$, and we will usually omit $(\mathrm{mod} \Lambda)$ because the lattice Λ is clear from the context.

It is visually apparent that the point set in Figure 1.2 has a somewhat lattice-like regular structure. This structure is shown in Figure 1.3. We call this a *fractured lattice.* A fractured lattice is determined by a *fracture vector* f, and instead of d basis vectors, it is characterized by d pairs of “basis vectors” $\{u_1, u_1+f\}, \ldots, \{u_d, u_1+f\}$.

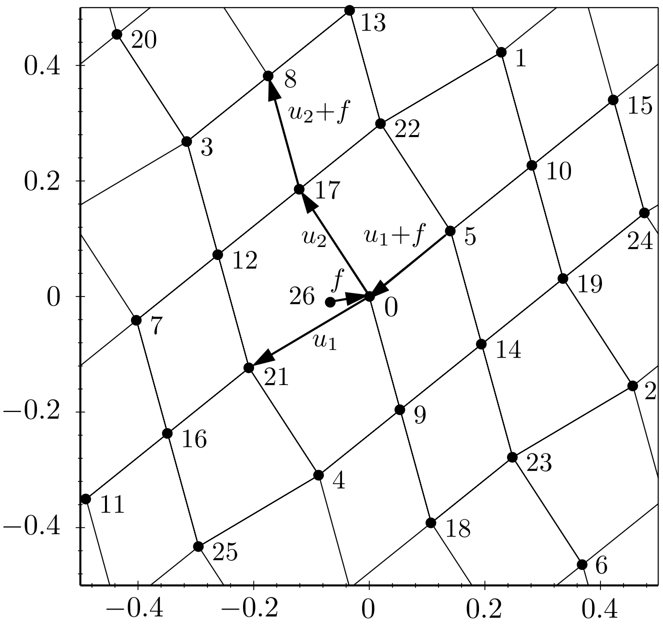

Fig. 1.3. The point set A_{26} and the fractured lattice

Lemma 1. *Assume that the vectors* $0, a, 2a, \dots, (n-1)a$ *are distinct modulo* Λ. *Then the set* A_n *has the structure of a fractured lattice, in the following sense*: *For every* x *and for every* $i = 1, \dots, d$, *the following holds.*

- $x + u_i \in A_n$ *or* $x + (u_i + f) \in A_n$, *but not both.*
- $x - u_i \in A_n$ *or* $x - (u_i + f) \in A_n$, *but not both.*

We call these $2d$ *points the neighbors of* x. *Moreover,* A_n *is connected by this neighborhood relation.*

For $f = 0$, this definition coincides with the usual concept of a lattice.

Proof. Let $z_1, \dots, z_d$ be positive integers less than n with $\gcd(z_1, \dots, z_d, n) = 1$. Then the vectors $u_i \equiv z_i a$ and $f \equiv -na$ will form the fractured lattice. To see this, let $x \equiv ja \in A_n$, for some $0 \le j < n$. If $j + z_i < n$, then $(j + z_i)a \equiv x + u_i \in A_n$. Otherwise, $0 \le j + z_i - n \le n$, and $(j + z_i - n)a = x + (u_i + f) \in A_n$. Both possibilities cannot hold simultaneously. Otherwise we would have two vectors in A_n whose difference is f: $ja + f \equiv j'a$, with $0 \le j, j' < n$. If $j < j'$ it follows that $(n-1)a \equiv (j' - j - 1)a$, contrary to the assumption that all points in A_n are distinct. If $j > j'$, then $0 \equiv (j' - j + n)a \in A_n$, again a contradiction. The case where we want to subtract u_i or $u_i + f$ from x works similarly.

To show that any two vectors $ja, j'a \in A_n$ are connected, note that the equation

$$j + k_1 z_1 + k_2 z_2 + \cdots k_d z_d \equiv j' \pmod{n}$$

has an integer solution $(k_1, \dots, k_d)$ by assumption. Thus we can transform ja into $j'a$ by repeatedly adding or subtracting u_i of $u_i + f$ (as appropriate), the given number of times $|k_i|$. The result is a vector $j''a \in A_n$ with $0 \le j'' < n$ and $j'' \equiv j'$ (mod n), and hence it must be $j'a$. □

This proof is quite trivial. Note that we did not exclude the possibility $z_1 = z_2 = \cdots = z_d$, and hence the basis vectors u_i need not even be distinct. However, the lemma might not give what is suggested by the picture in Figure 1.3. The fracture vector $f \equiv -na \bmod \Lambda$ might be very long compared to the “basis vectors” u_i, and hence the structure is not really visible. However, when na is a best approximation, then we may choose some "good" basis vectors and the structure becomes apparent.

When na is a best approximation, then the set A_n is also a *perturbed lattice*, in the sense that the points can be moved slightly so that they form a real lattice. Figure 1.4 shows this lattice for the point set of Figure 1.3.

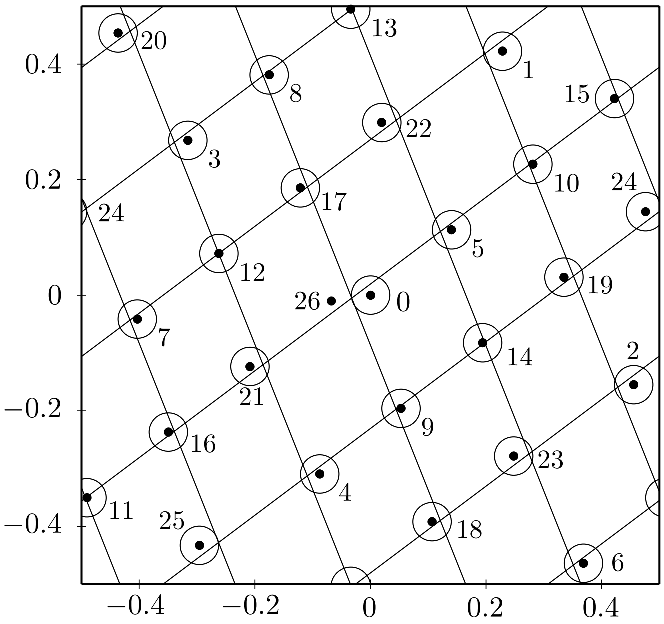

Fig. 1.4. The point set A_{26} is approximated by the lattice $\hat{\Lambda}$ shifted by f', which is indicated by the grid of lines. The little disks indicate the largest perturbation error ε

This lattice has been introduced by Larcher [7], and the following theorem is essentially given in [2] (apart from the statement about the bijection).

First we state an easy observation.

Lemma 2. *The length of the shortest vector of A_n is a lower bound for the distance between two points of A_n*

Proof. The distance between two points ja and $j'a$ is $\|ja - j'a\| = |j - j'|a$, which is the length of some vector of A_n. □

Let $n_0 = 1 < n_1 < n_2 < \cdots < n_{t-1}$ be the sequence of indices n when A_{n+1} contains a non-zero vector which is shorter than the previous shortest vector (a best approximating vector). We define the last index n_t as the first level $\Lambda + ia$ which contains the zero vector; the set A_n will remain stable for $n \geq n_t$. This last index must exist when the data are rational. If a or Λ are not rational, then the sequence $n_0, n_1, \ldots$ may be infinite.

Theorem 1. *For $n = n_1, n_2, \ldots$, let $f \equiv na$ be the shortest non-zero vector in A_{n+1}. (For $n = n_t$, let $f = 0$.) Then there is a d-dimensional lattice $\hat{\Lambda}$ and a bijection π between A_n and $\hat{\Lambda}$ such that*

$$\|\hat{x} + f' - x\| \leq \varepsilon, \tag{1.1}$$

for all $x \in A_n$ and $\hat{x} = \pi(x)$, with

$$f' = f \cdot \frac{n-1}{2n} \text{ and } \varepsilon = \|f\|.$$

Moreover, the given value ε is the smallest value for which such a lattice exists, and the lattice $\hat{\Lambda}$ is the unique lattice that fulfills (1.1) *for any $\varepsilon < \|f\|$.*

Proof. $\hat{\Lambda}$ is generated by Λ and the vector $\hat{a} := a - f/n$. Since $n\hat{a} = na - f \equiv 0$, it is indeed a discrete lattice, with at most n congruence classes modulo Λ. The mapping $\pi\colon A_n \to \hat{\Lambda}$ is given by $\pi(ja) := j\hat{a}$, and we get $x - \hat{x} \equiv ja - j\hat{a} = jf/n$, for $j = 0, 1, \ldots, n-1$. To balance this error between the two extreme cases $j = 0$ and $j = n-1$, we shift $\hat{x}$ by $f' = \frac{1}{2} \cdot (n-1)f/n$:

$$\hat{x} + f - x \equiv (\tfrac{n-1}{2} - j)f/n,$$

for some $0 \le j \le n-1$, and the length of this vector is at most ε. (This is true for any norm.)

The mapping π must be injective, since two points ja and $j'a$ that are mapped to the same point are within distance ε of a common point, and hence $\|ja - j'a\| \le 2\varepsilon < \|f\|$. By Lemma 2, this contradicts the assumption that f is the shortest vector. □

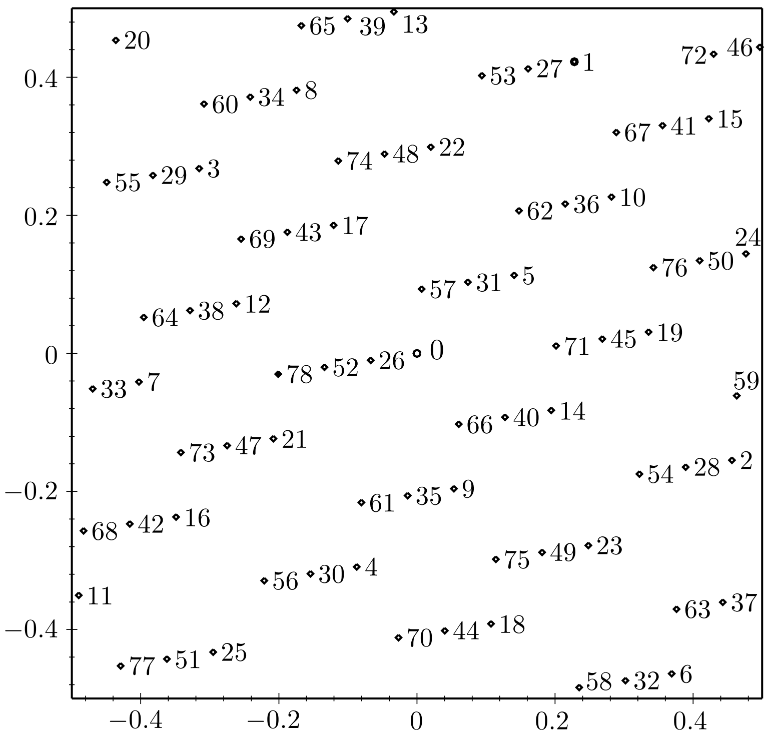

Fig. 1.5. The point set A_{79}. Each point $x = 0, a, 2a, \ldots, 25a$ has received two neighbors $x + f$ and $x + 2f$, and point 0 has just received its third neighbor

1.3 Inductive construction of good fractured lattice bases

We will now see how a good "basis" for a fractured lattice or for $\hat{\Lambda}$ can help us to understand the further evolution of the set A_n. Consider the situation in Figure 1.2. Point 0 has just received a new nearest neighbor $f \equiv na$. The next point $(n+1)a = 28a$ will be a neighbor of a, and then the points $2a$, $3a$, and so on will get neighbors, by the relation $(n+j)a \equiv ja + f$, see Figure 1.5. Once all points $0, \ldots, (n-1)a$

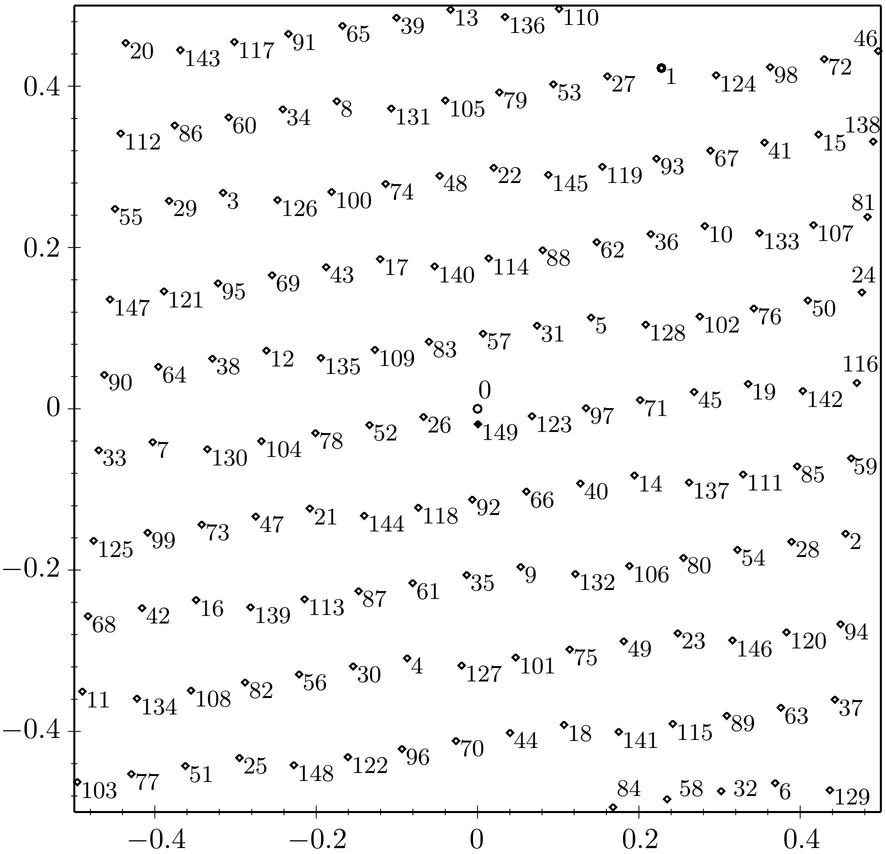

Fig. 1.6. The point set A_{150}

have received their neighbors, point 0 will receive a second-degree neighbor $52a$, which is a neighbor of $26a$, and similarly for the other points. Then there will be third-degree neighbors, and so on. Although this structure has some regularity, it does not approximate a lattice. Only when a new nearest neighbor appears (this is $149a$ in Figure 1.6), another approximate lattice structure emerges.

Now let us imagine that we stand at the origin and wait until a point arrives which has a smaller length than $f = na$, the shortest vector so far. A typical situation is shown in Figure 1.7. All points of A_n start "shooting" in the direction of f, by successively generating neighbors at distance $f, 2f, 3f, \ldots$. We are interested in the first point which hits the small disc of radius $\|f\|$ around the origin. The set A_n can be divided into $(d-1)$-dimensional layers according to the fractured lattice structure. If the layers are sufficiently well "separated", as in Figure 1.7, we can ask for the first *layer* whose shooting ray in direction f hits the small disk. A schematic picture is shown in Figure 1.8. The shooting direction is vertical downward, we have arranged each layer on a horizontal line and placed the layers on successive levels. The origin with the target disc lies on the horizontal *ground line.*

Now, the points in this drawing *form a lattice*: The difference vectors between neighbors are of the form $\pm u_i$ or $\pm(u_i + f)$, by the fractured lattice structure. The irregularity introduced by the fracture vector f is *projected away* in this process, since we are only interested in rays parallel to f through the points. This observation is the main geometric insight on which our algorithm is based.

The situation in Figure 1.8 is now almost the same as the picture of the Best Approximation Problem in two dimensions shown in Figure 1.1. In other words, *the*

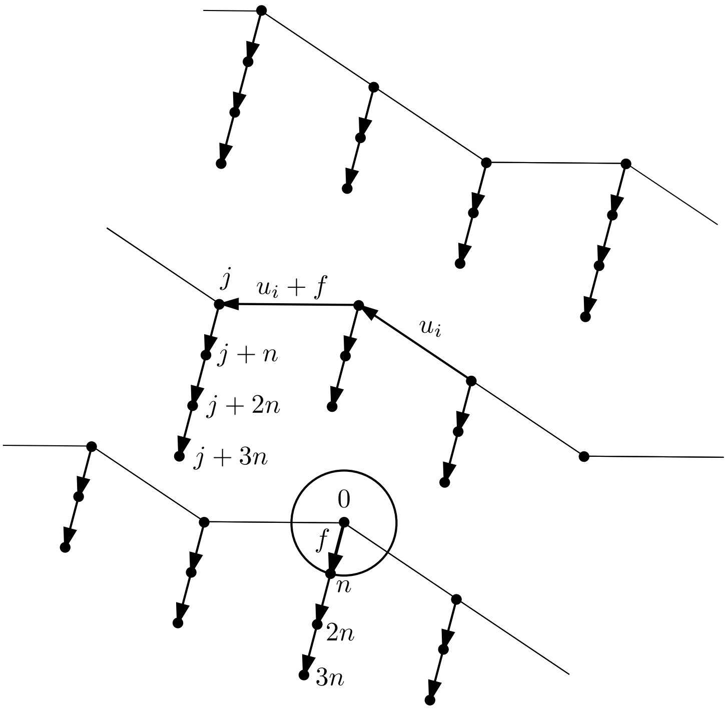

Fig. 1.7. The points of A_n are arranged in successive layers and shoot in direction f

problem that we face is of the same kind as the original problem, but one dimension lower.

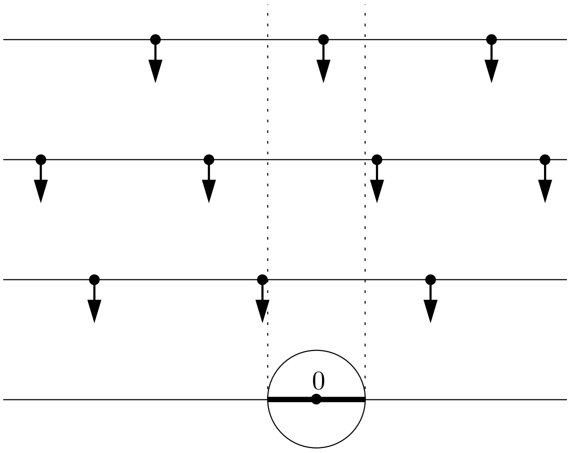

Fig. 1.8. A schematic picture of the shooting process

There are some differences, however. If we ask which *ray* in direction f is the first one that hits the disk, we can replace the disk by its intersection with the ground line (the thick segment in Figure 1.8). In higher dimensions, this target is not a segment but a $(d-1)$-dimensional ball. Thus we would get *exactly* the Best Approximation Problem in one dimension lower. However, the neighbors are generated for the points of A_n in discrete steps and don't fill a ray densely. It may happen that some ray in direction f passes through the disc but none of the points fall in the disk, see Figure 1.9a. We call this a *phase error* because it depends on the locations of the equidistantly placed points on the ray.

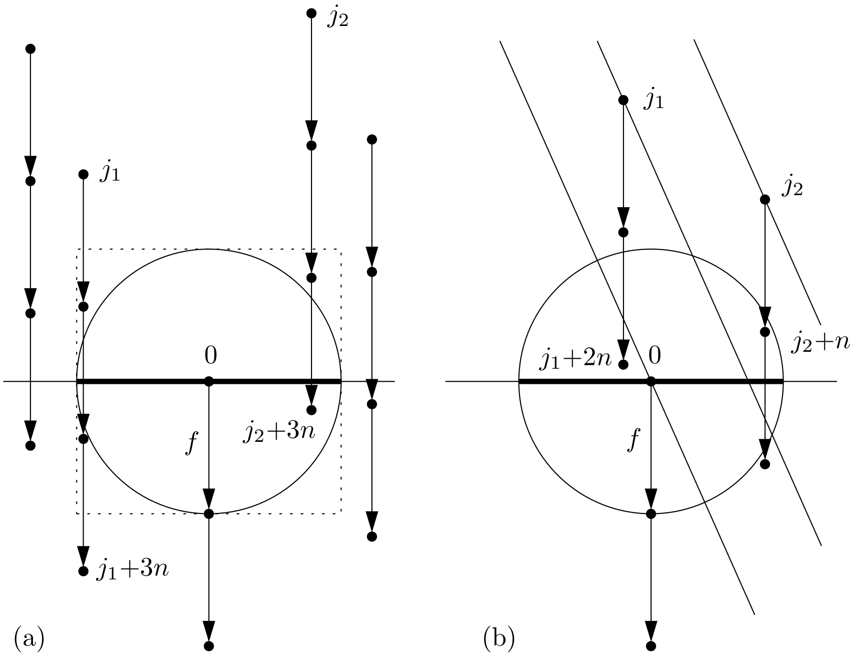

Fig. 1.9. (**a**) A phase error. The first point whose ray in direction f hits the disk is j_1, but the points $j_1+n, j_1+2n, \ldots$ miss the disk. The first point that really falls in the disk comes later from j_2. (**b**) A sequence error. The point j_1 lies on a lower layer than j_2 but the sequence of points starting at j_2 hits the disk first

There is another source of error, a *sequence* error. It may happen that the levels are not so clearly separated, and a point on a later layer hits the disk first, see Figure 1.9b. This phenomenon may occur when the layers form a small angle with the shooting direction. It may also happen that sequences of points emanating from two points on the same layer of A_n hit the disk in the wrong order.

To deal with these problems, we sometimes have to continue the procedure even after finding the "first" ray that hits the horizontal segment, in order to build a safety margin. The *Repeated Best Approximation Problem* is like the Best Approximation Problem defined in Section 1.1, but after the lowest lattice point within a distance ε of the x_{d+1}-axis is found, we may also ask for the second-lowest, the third-lowest, and so on (up to some constant bound that depends only on the dimension).

> The *Repeated Best Approximation Problem*
> Enumerate the points $ia+x$, $i>0$, $x \in \Lambda$, with $\|ia+x\| < \varepsilon$ in order of increasing i.

In our application of this problem, we will have $\varepsilon \le \min\{\, \|x\| \mid x \in \Lambda, x \neq 0 \,\}$. Then the number of points with the same value of i is bounded by a constant that depends only on the dimension, by Lemma 2.

Lemma 3. *In each dimension, there is some constant bound C_d^1 on the number of times that a phase error can occur.*

Proof. This is a volume packing argument. For each ray that hits the vertical projection of the disk, there must be a point generated that falls into the bounding box vertical cylinder circumscribed about the disk. (This appears as a circumscribed square in Figure 1.9a). By Lemma 2, these points must have a minimum distance of $\|f\|$ as long as no point lies within the disk. The box can only contain a constant

```
Input: A lattice basis $u_1, \dots, u_d$, a vector $a$, and a threshold $\varepsilon$.
Output: A sequence of points $x = z_1 u_1 + \cdots + z_d u_d + ia$ for the Repeated Best
Approximation Problem.
Initialization: Set $f := a$;
Main loop:
    Select one vector of the current basis. (For convenience, let us assume it is $u_d$.)
    Let $u'_1, \dots, u'_d$ denote the projection of $u_1, \dots, u_d$ onto the hyperplane perpendicular to $f$.
    Solve the $(d-1)$-dimensional Repeated Best Approximation Problem for
    $u'_1, \dots, u'_{d-1}$, $a' := u'_d$ and $\varepsilon' := \|f\|$.
    This will enumerate the successive points
        $y = z_1 u'_1 + \cdots + z_d u'_d$
    with $\|y\| < \|f\|$ in order of increasing $z_d > 0$.
    For each point $y$, find all $i \in \mathbb{Z}$ such that
        $x = z_1 u_1 + \cdots + z_d u_d + if$
    satisfies $\|x\| < \|f\|$.
    If enough points have been generated to be sure that the point $x$ with minimum
    $i$ has been found, then let $f$ be this vector.
    If $\|f\| > \varepsilon$, update the basis $u_1, \dots, u_d$ and repeat the loop.
Output the generated points $x$ with $\|x\| < \varepsilon$ in order of increasing $i$.
```

Fig. 1.10. The recursive algorithm for the Repeated Best Approximation Problem

number of points with minimum separation $\|f\|$. (This argument can be adapted to other metrics. In the $(d-1)$-dimensional problem one works with a projection of a d-dimensional unit ball, or with a suitable enclosing body that is easier to handle.) □

The recursive algorithm is schematically shown in Figure 1.10.

To take care of the sequence error, we have to relate the order of the points

$$x = z_1 u_1 + \cdots z_d u_d + if$$

with $\|x\| < \|f\|$ according to z_d, in which the points are generated, to the desired order according to i. We use the following lemma.

Lemma 4. *Let $x = z_1 u_1 + \cdots + z_d u_d + if$ and $x' = z'_1 u_1 + \cdots + z'_d u_d + i'f$ be two points with $\|x\| \le \|f\|$ and $\|x'\| \le \|f\|$, and let $z'_d \ge z_d$. Then $i' \ge i - C$ for $C = 3\|u^*_d\| \cdot \|f\|$, where $u^*_1, \dots, u^*_d$ is a basis of the dual lattice.*

Proof. Let $y = z_1 u_1 + \cdots + z_d u_d$ and $y' = z'_1 u_1 + \cdots + z'_d u_d$ be the corresponding points of the lattice. Then we have $z_d = \langle y, u^*_d \rangle$ and $z'_d = \langle y', u^*_d \rangle$, and $if - i'f = (x - y) - (x' - y') = (x - x') - (y - y')$. By taking the inner product with u^*_d, we obtain $i - i' \le z_d - z'_d + \|x - x'\| \cdot \|u^*_d\|$. We have $\|x - x'\| \le 2\|f\|$, and we have to add $\|f\|$ because of the error bound between the point set A_n and the lattice (Theorem 1). □

Since f is shorter than the shortest vector of the lattice, the product $\|u^*_d\| \cdot \|f\|$ is bounded by a constant which depends only on the dimension.

It follows that, in fixed dimension, the algorithm has to wait only a constant number of rounds after the first point $x = z_1 u_1 + \cdots + z_d u_d + if$ with $\|x\| \le \|f\|$ (the point x with smallest z_d) has been found before it can start to output x. Thus the algorithm has to accommodate a buffer for a limited number of points. In addition, for some points y which are generated, there may be no integer value i such that $\|z_1u_1+\cdots+z_du_d+if\| < \|f\|$. The occurrence of these phase errors is also bounded, by Lemma 3. Thus we can conclude:

Lemma 5. *In one iteration of the main loop, there is a constant number of calls to the $(d-1)$-dimensional Repeated Best Approximation Problem.* □

The number of iterations is equal to the number of best approximations, which is known to be logarithmic in the size of the input numbers [6].

Theorem 2. *For fixed d, the algorithm reduces the d-dimensional Repeated Best Approximation Problem with input numbers of s bits to $O(s)$ instances of the $(d-1)$-dimensional Repeated Best Approximation Problem.*

Therefore, the d-dimensional Repeated Best Approximation Problem is solved in $O(s^d)$ steps. □

1.4 Conclusion

The algorithm that we have given is not competitive with the fastest theoretical algorithm in fixed dimension [4], which reduced the problem to only $O(\log^d s)$ greatest common divisor computations. However, the simple and in some sense, canonical structure may be the key to an algorithm in the style of Schönhage's algorithm for the two-dimensional problem [12], which solves the problem only approximately for the leading half of the bits of the numbers, and then extends this solution to the full numbers (the "half-GCD" approach). A canonical procedure like the continued fraction expansion (the Euclidean algorithm) seems crucial for the success of this approach.

We have not analyzed the constants in the algorithm. They depend on relations between a lattice and its dual lattice and on packing arguments, for which precise estimates are hard to obtain, even in three dimensions.

References

1. A. J. Brentjes, *Multi-dimensional continued fraction algorithms*, Mathematical Centre Tracts, vol. 145, Mathematisch Centrum, Amsterdam, 1981.
2. Nicolas Chevallier, *Meilleures approximations diophantiennes d'un élément du tore $\mathbb{T}^2$ et géométrie de la suite des multiples de cet élément*, Acta Arith. **78** (1996), 19–35.
3. Friedrich Eisenbrand and Günter Rote, *Fast 2-variable integer programming*, IPCO 2001—Proceedings of the 8th Conference on Integer Programming and Combinatorial Optimization, Utrecht (K. Aardal and B. Gerards, eds.), Lecture Notes in Computer Science, vol. 2081, Springer-Verlag, 2001, pp. 78–89.
4. ———, *Fast reduction of ternary quadratic forms*, Cryptography and Lattices — International Conference, CaLC 2001 (Joseph H. Silverman, ed.), Lecture Notes in Computer Science, vol. 2146, Springer-Verlag, 2001, pp. 32–44.
5. Naoyoshi Kanamuru, Takao Nishizeki, and Tetsuo Asano, *Efficient enumeration of grid points in a convex polygon and its application to integer programming*, Internat. J. Comput. Geom. Appl. **4** (1994), 69–85.
6. Jeff C. Lagarias, *Best simultaneous diophantine approximations. I: Growth rates of best approximation denominators*, Trans. Am. Math. Soc. **272** (1982), 545–554.
7. G. Larcher, *On the distribution of s-dimensional Kronecker sequences*, Acta Arithmetica **51** (1988), 335–347.
8. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász, *Factoring polynomials with rational coefficients*, Math. Ann. **261** (1982), 514–534.
9. H. W. Lenstra, *Integer programming with a fixed number of variables*, Math. Oper. Res. **8** (1983), 538–548.

10. Yves Meyer, *Quasicrystals, diophantine approximation and algebraic numbers*, Beyond Quasicrystals. Papers of the winter school, Les Houches, France, March 7–18, 1994. (Françoise Axel and D. Gratias, eds.), Springer, Berlin, 1995, pp. 3–16.
11. C. P. Schnorr, *A hierarchy of polynomial time lattice basis reduction algorithms*, Theoret. Comput. Sci. **53** (1987), 201–224.
12. Arnold Schönhage, *Schnelle Berechnung von Kettenbruchentwicklungen*, Acta Informatica **1** (1971), 139–144.
13. Marjorie Senechal, *Quasicrystals and geometry*, Cambridge University Press, 1996.