1.1 "Short" solutions to modular equations

1.1.1 A basic example

Let's start by considering the equation

$$26x + y \equiv 0 \bmod 43 \ .$$

We note a few solutions to the equation:

$$(x, y) = (0, 0),\ (-1, 26),\ (0, 43),\ \text{or } (1, 17) \ .$$

It is natural to ask whether we can find "shorter" solutions under some definition of length. For concreteness, we can try to minimize the ℓ_∞ norm, $\max\{|x|, |y|\}$. The solution $(x, y) = (0, 0)$ clearly has minimal length, but it is not very interesting. So, a more refined question asks for non-zero short solutions.

It will be useful to consider the set of all solutions,[1]

$$\mathcal{L} := \{(x, y) \in \mathbb{Z}^2 \ : \ 26x + y \equiv 0 \bmod 43\} \ .$$

Note that, if $(x_1, y_1), (x_2, y_2) \in \mathcal{L}$, then $(-x_1, -y_1), (x_1 + x_2, y_1 + y_2) \in \mathcal{L}$ as well. In other words, $\mathcal{L}$ is a group under coordinate-wise addition. Furthermore, (x, y) is a solution if and only if $(x, y) = a(1, -26) + b(0, 43)$ for some integers a and b.[2] We call

$$\begin{pmatrix} 1 & 0 \\ -26 & 43 \end{pmatrix}$$

a *basis* for the *lattice* $\mathcal{L}$. Note that the basis is not unique. Indeed,

$$\begin{pmatrix} 1 & 1 \\ -26 & 17 \end{pmatrix}$$

is also a basis, which we can verify by writing $(0, 43) = (1, 17) - (1, -26)$.

We can use these observations to find short solutions to the original equation relatively easily by using a Euclidean-like algorithm. In particular, since $(1, -26) \in \mathcal{L}$ and $(1, 17) \in \mathcal{L}$, we see that $(2, -9) = (1, -26) + (1, 17) \in \mathcal{L}$. We then see that $(3, 8) = (2, -9) + (1, 17) \in \mathcal{L}$,

[1] In lectures, we are careful to write column vectors. In these introductory notes, we will sin and fail to distinguish between row and column vectors, in order to avoid writing transposes everywhere or writing ugly bulky in-line math like $\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$.

[2]

$$\begin{aligned} (x, y) \in \mathcal{L} &\iff 26x + y = 43k \ , \ k \in \mathbb{Z} \\ &\iff (x, y) = (x, 43k - 26x) \\ &\iff (x, y) = k(0, 43) + x(1, -26) \ . \end{aligned}$$

and finally that $(5,-1) = (3,8) + (2,-9) \in \mathcal{L}$. It is relatively straightforward to show that there is no non-zero vector in $\mathcal{L}$ shorter than $(5,-1)$, so we have found an optimal solution.[3]

Should we be surprised by this? I.e., is $(5,-1)$ a remarkably short solution to a two-dimensional linear equation mod 43? Perhaps I chose an equation with a planted "short" solution. Or perhaps this is typical. In class, we considered this specific two-dimensional question directly. In the notes, we'll simply move directly to the n-dimensional generalization.

1.1.2 n dimensions

Let's consider the natural n-dimensional generalization of the question from the previous section. For some integer modulus $q \geq 2$ and $a_1, \ldots, a_n \in \mathbb{Z}_q$, we consider solutions to the modular equation

$$\sum_{i=1}^{n} a_i z_i \equiv 0 \bmod q \ .$$

Equivalently, we consider solutions $\boldsymbol{z} \in \mathbb{Z}^n$ to $\langle \boldsymbol{a}, \boldsymbol{z} \rangle \equiv 0 \bmod q$, where $\boldsymbol{a} := (a_1, \ldots, a_n)$. As before, we define the set of all solutions,

$$\mathcal{L} := \{ \boldsymbol{z} \in \mathbb{Z}^n \ : \ \langle \boldsymbol{a}, \boldsymbol{z} \rangle \equiv 0 \bmod q \} \ ,$$

and we observe that $\mathcal{L}$ is a group under addition. Furthermore, there is a basis $(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ such that a vector is in $\mathcal{L}$ if and only if it is an integer linear combination of the basis vectors. We leave it to the reader to find the $\boldsymbol{b}_i$. (You might find it easier to first consider the case when q is prime.)

As before, we are interested in non-zero solutions $\boldsymbol{z} \in \mathcal{L}$ such that $\|\boldsymbol{z}\|_\infty := \max |z_i|$ is as small as possible. How short will the shortest solution be? Note that most vectors in $\boldsymbol{z} \in \mathbb{Z}_q^n$ have $\|\boldsymbol{z}\|_\infty \approx q/2 \cdot (1 - 1/n)$. So, we might expect the shortest solution in $\mathcal{L}$ to have ℓ_∞ norm roughly $\Theta(q)$. But, the following proposition shows that much shorter solutions always exist.

Proposition 1.1. *For any integer modulus $q \geq 2$ and any $\boldsymbol{a} \in \mathbb{Z}^n$, there exists a non-zero vector $\boldsymbol{z} \in \mathbb{Z}^n$ with $\|\boldsymbol{z}\| \leq q^{1/n}$ and $\langle \boldsymbol{a}, \boldsymbol{z} \rangle \equiv 0 \bmod q$.*

Proof. Consider the set of all vectors $\boldsymbol{z} \in \mathbb{Z}^n$ with ℓ_∞ norm at most $r := q^{1/n}/2$. Of course, these vectors typically will not be solutions, but we can still consider their dot products with $\boldsymbol{a}$ mod q. In particular, we consider the set of all such dot products,

$$C := \{ \langle \boldsymbol{a}, \boldsymbol{z} \rangle \bmod q \ : \ \|\boldsymbol{z}\|_\infty \leq r \} \ .$$

Note that $|C| \leq q$. But, there are $(2\lfloor r \rfloor + 1)^n > (2r)^n = q$ vectors in $\mathbb{Z}^n$ with ℓ_∞ norm at most r. So, there must be at least one pair $\boldsymbol{z}_1 \neq \boldsymbol{z}_2$ with $\|\boldsymbol{z}_1\|_\infty, \|\boldsymbol{z}_2\|_\infty \leq r$ but $\langle \boldsymbol{a}, \boldsymbol{z}_1 \rangle \equiv \langle \boldsymbol{a}, \boldsymbol{z}_2 \rangle \bmod q$. Then, $\langle \boldsymbol{a}, \boldsymbol{z}_1 - \boldsymbol{z}_2 \rangle \equiv 0 \bmod q$. So, $\boldsymbol{z}_1 - \boldsymbol{z}_2$ is a non-zero solution with $\|\boldsymbol{z}_1 - \boldsymbol{z}_2\|_\infty \leq 2r$ by triangle inequality. □

[3] This procedure is typically attributed to Gauss. Unfortunately, it does not generalize well to higher dimensions. (Why not?)

In fact, it is easy to see that Proposition 1.1 is quite tight when $\boldsymbol{a}$ is chosen uniformly at random from $\mathbb{Z}_q^n$ for some prime modulus q. In that case, $\langle \boldsymbol{a}, \boldsymbol{z} \rangle$ is uniformly random mod q for any $\boldsymbol{z} \not\equiv \boldsymbol{0} \bmod q$. In particular, the probability that a given vector that is non-zero mod q lands in the lattice is exactly $1/q$, and one can simply apply union bound to see that this is unlikely to happen for any $\boldsymbol{z}$ of length less than, say, $q^{1/n}/4$.

1.1.3 A cryptographic interlude: Ajtai's hash function

Lattices of solutions to modular linear equations are often called "Ajtai lattices," after Miklós Ajtai, who showed beautiful cryptographic and complexity-theoretic applications of such lattices [Ajt96, Ajt98]. Here, we introduce the first such application: Ajtai's family of hash functions. In particular, for $\boldsymbol{a} \in \mathbb{Z}_q^n$, let $h_{\boldsymbol{a}} : \{0,1\}^n \to \mathbb{Z}_q$ be defined as $h_{\boldsymbol{a}}(\boldsymbol{z}) := \langle \boldsymbol{a}, \boldsymbol{z} \rangle \bmod q$. Note that we restrict the input of $h_{\boldsymbol{a}}$ to vectors $\boldsymbol{z}$ with positive coordinates and $\|\boldsymbol{z}\|_\infty \leq 1$. And, note that $h_{\boldsymbol{a}}$ is a compressing function if $2^n > q$.

More importantly, notice how incredibly simple this hash function is! Other cryptographic hash functions are either relatively complicated circuits designed heuristically (e.g., SHA) or use exponentiation over some group. This hash function only uses modular addition and multiplication! (The major downside of this function is the size of the description of $\boldsymbol{a}$, which is $n \log q$.)

The following easy claim shows that Ajtai's hash function is collision resistant as long as it is hard to find short solutions to $\langle \boldsymbol{a}, \boldsymbol{z} \rangle \equiv 0 \bmod q$.

Claim 1.2. *Finding distinct vectors $\boldsymbol{z}_1, \boldsymbol{z}_2 \in \{0,1\}^n$ such that $h_{\boldsymbol{a}}(\boldsymbol{z}_1) = h_{\boldsymbol{a}}(\boldsymbol{z}_2)$ is at least as hard as finding a vector $\boldsymbol{z} \in \mathbb{Z}^n$ with $\|\boldsymbol{z}\|_\infty = 1$ and $\langle \boldsymbol{a}, \boldsymbol{z} \rangle \equiv 0 \bmod q$.*

Proof. Take $\boldsymbol{z} := \boldsymbol{z}_1 - \boldsymbol{z}_2$. We have $\langle \boldsymbol{a}, \boldsymbol{z}_1 \rangle \equiv \langle \boldsymbol{a}, \boldsymbol{z}_2 \rangle \bmod q$, so $\langle \boldsymbol{a}, \boldsymbol{z} \rangle \equiv 0 \bmod q$. Furthermore, $\boldsymbol{z}$ is non-zero with $\boldsymbol{z} \in \{-1, 0, 1\}$ so that $\|\boldsymbol{z}\|_\infty = 1$. □

Later, we will see that finding a short solution $\boldsymbol{z}$ for a randomly selected $\boldsymbol{a}$ is as hard as a certain *worst-case* computational lattice problem. So, Ajtai's hash function is collision-resistant as long as some computational lattice problem is hard in the worst case!

1.2 General lattices and Minkowski's Theorem

Now that we've spent a while beating around the bush, it's time to define a lattice in general. A lattice is the set of integer linear combinations of linearly independent basis vectors $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n \in \mathbb{R}^n$,

$$\mathcal{L} := \{ z_1 \boldsymbol{b}_1 + \cdots + z_n \boldsymbol{b}_n \ : \ z_i \in \mathbb{Z} \} \ .$$

Here's an example in two dimensions, called the hexagonal lattice:

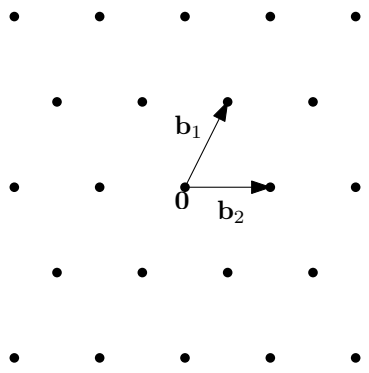

Note in particular that an Ajtai lattice is a lattice, but not every lattice is an Ajtai lattice. (E.g., a lattice generated by non-integral basis vectors is certainly not an Ajtai lattice.) As before, the basis of a lattice is not unique, and finding a "good basis" is an important technique in lattice algorithms.

1.2.1 The fundamental parallelepiped

The basis B of a lattice defines a *fundamental parallelepiped*,[4]

$$\mathcal{P} := \left\{ \sum c_i \boldsymbol{b}_i \ : \ 0 \le c_i < 1 \right\} .$$

The fundamental parallelepiped of the basis for the hexagonal lattice shown above looks like this:

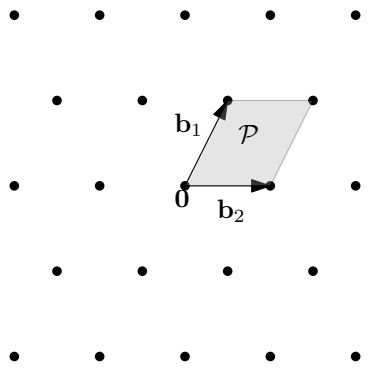

It is immediate from the definition that the fundamental parallelepiped forms a *lattice tiling*. I.e.,

$$\bigcup_{\boldsymbol{x} \in \mathcal{L}} (\mathcal{P} + \boldsymbol{x}) = \mathbb{R}^n ,$$

where the union is disjoint, as in this picture:

[4]For brevity's sake, we will typically not worry about measure-zero changes to sets. For example, we make no distinction between $\mathcal{P}$ and its closure or its interior. (Some authors prefer to make the fundamental parallelepiped symmetric by taking $|c_i| \le 1/2$.)

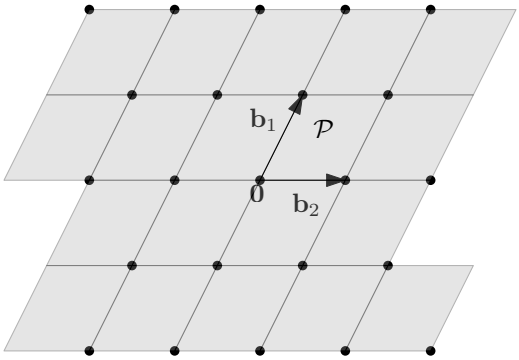

1.2.2 The determinant of a lattice

Again, we would like an upper bound on the length of the shortest non-zero vector in the lattice. (I.e., we would like to argue that a lattice always has a relatively short non-zero vector.) Of course, we can always scale a lattice by an arbitrarily large amount, so we will need to normalize somehow. For this, we will define the lattice determinant. Let $B := (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n)$ be the matrix of some basis for $\mathcal{L}$. Then, we define the determinant of $\mathcal{L}$ as

$$\det(\mathcal{L}) := |\det(B)| \,.$$

At this point, it is not even obvious from this definition that $\det(\mathcal{L})$ is well-defined! Perhaps different bases yield different determinants. The next claim gives a geometric interpretation of the determinant—it shows that $\det(\mathcal{L})^{-1}$ is the number of lattice points per unit volume. Since this characterization only depends on the lattice (and not on the basis), it in particular implies that the determinant is well-defined.

Claim 1.3.

$$\det(\mathcal{L})^{-1} = \lim_{r\to\infty} \frac{|\mathcal{L} \cap [-r/2, r/2]^n|}{r^n} \,.$$

Proof. Note that $\det(\mathcal{L}) := |\det(B)|$ is the volume of the fundamental parallelepiped $\mathcal{P}$ defined by the basis B. And, recall that $\mathcal{P}$ forms a lattice tiling. For large r, the number of parallelepipeds in this tiling that intersect $[-r/2, r/2]^n$ will approach $\mathrm{Vol}([-r/2, r/2]^n)/\mathrm{Vol}(\mathcal{P}) + o(r^n) = r^n/\det(\mathcal{L}) + o(r^n)$. (This is essentially the definition of volume.) The result follows by noting that there is one parallelepiped per lattice point, so the number of parallelepipeds that intersect $[-r/2, r/2]^n$ is $|\mathcal{L} \cap [-r/2, r/2]| + o(r^n)$. □

To check your understanding, see if you can use this claim to show that the determinant of an Ajtai lattice

$$\{\boldsymbol{z} \in \mathbb{Z}^n \;:\; \langle \boldsymbol{a}, \boldsymbol{z} \rangle \equiv 0 \bmod q\}$$

is q as long as $\boldsymbol{a} \not\equiv \mathbf{0} \bmod q$.

1.2.3 Minkowski's theorem

Now that we have an appropriate normalization, we are ready to present (one of) the most important result(s) in the study of lattices: Minkowski's first fundamental theorem.

Minkowski's theorem guarantees that any lattice with small determinant has a short vector. (In particular, it is a very strong generalization of Proposition 1.1.) This result shows up everywhere in the study of lattices—from lattice-based cryptography to lattice algorithms to coding theory to algebraic number theory and the geometry of numbers.

Minkowski's theorem actually works for *any* norm. We will typically be interested in the ℓ_2 or ℓ_∞ norms, but we will still present the more general theorem. Recall that a general norm is defined by its unit ball $K \subset \mathbb{R}^n$, which is a symmetric convex body. I.e., K is bounded with non-empty interior, $K = -K$, and if $\boldsymbol{x}, \boldsymbol{y} \in K$, then so is $\alpha\boldsymbol{x} + (1-\alpha)\boldsymbol{y}$ for any $\alpha \in [0,1]$.

Theorem 1.4 (Minkowski's theorem [Min10])**.** *For any lattice $\mathcal{L} \subset \mathbb{R}^n$, if $K \subset \mathbb{R}^n$ is a symmetric convex body with $\mathrm{Vol}(K/2) \geq \det(\mathcal{L})$, then K contains a non-zero lattice vector.*

Note that Minkowski's theorem is a strict strengthening of Proposition 1.1. In particular, Proposition 1.1 follows by taking K to be the cube $[-q^{1/n}, q^{1/n}]^n$ and recalling that the Ajtai lattices from Proposition 1.1 have determinant q (or determinant one if $\boldsymbol{a} = \boldsymbol{0}$).

Proof of Minkowski's theorem. Suppose that $K/2 + \boldsymbol{x}$ and $K/2 + \boldsymbol{x}'$ have non-empty intersection for some $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{L}$ with $\boldsymbol{x} \neq \boldsymbol{x}'$. Let $\boldsymbol{y} \in (K/2 + \boldsymbol{x}) \cap (K/2 + \boldsymbol{x}')$. I.e., $\boldsymbol{y} - \boldsymbol{x} \in K/2$ and $\boldsymbol{y} - \boldsymbol{x}' \in K/2$. By symmetry and convexity, it follows that $(\boldsymbol{y} - \boldsymbol{x}')/2 - (\boldsymbol{y} - \boldsymbol{x})/2 = \boldsymbol{x}/2 - \boldsymbol{x}'/2 \in K/2$. Therefore, $\boldsymbol{x} - \boldsymbol{x}'$ is a non-zero vector in $K \cap \mathcal{L}$.

So, we may assume that $K/2 + \boldsymbol{x}$ and $K/2 + \boldsymbol{x}'$ have non-empty intersection for all distinct $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{L}$. Then, for any r,

$$S := \bigcup_{\boldsymbol{x} \in \mathcal{L} \cap [-r/2, r/2]^n} (K/2 + \boldsymbol{x})$$

is a disjoint union. So,

$$\begin{aligned} \mathrm{Vol}(S) &= \sum_{\boldsymbol{x} \in \mathcal{L} \cap [-r/2, r/2]^n} \mathrm{Vol}(K/2 + \boldsymbol{x}) \\ &= |\mathcal{L} \cap [-r/2, r/2]^n| \cdot \mathrm{Vol}(K/2) \\ &\geq |\mathcal{L} \cap [-r/2, r/2]^n| \cdot \det(\mathcal{L}) \ . \end{aligned}$$

On the other hand, by the definition of S, $S \subset [-r/2, r/2] + K/2$. So, $\mathrm{Vol}(S) \leq r^n + o(r^n)$. This contradicts Claim 1.3. □

1.2.4 λ_1 and the Euclidean norm

Recall that the Euclidean norm, or the ℓ_2 norm, is given by

$$\|\boldsymbol{x}\| := \left(\sum x_i^2 \right)^{1/2} .$$

Previously, we worked with the ℓ_∞ norm out of convenience, but in this universe, we have a strong preference for the Euclidean norm. For any lattice $\mathcal{L}$, we define

$$\lambda_1(\mathcal{L}) := \inf_{\boldsymbol{x} \in \mathcal{L} \setminus \{\boldsymbol{0}\}} \|\boldsymbol{x}\|$$

to be the minimal (Euclidean) length of a non-zero vector in the lattice. (In the next lecture, we will see that the infimum can be replaced by a minimum.)

To apply Minkowski's theorem to this setting, we will need to know the volume of the unit ball,

$$B_2^n := \{\boldsymbol{x} \in \mathbb{R}^n \ : \ \|\boldsymbol{x}\| \leq 1\} \ .$$

If you have not seen this before, then you might be surprised to learn that it has extremely small volume!

$$\mathrm{Vol}(B_2^n) = \frac{1}{\sqrt{\pi n}} \cdot \left(\frac{2\pi e}{n}\right)^{n/2} \cdot (1 + o(1)) \ .$$

Indeed, a ball of radius $\sqrt{n}/2$ has volume greater than one, and this is tight up to a constant in the radius. From this, we derive the following immediate corollary of Minkowski's theorem.

Corollary 1.5. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$, $\lambda_1(\mathcal{L}) \leq \sqrt{n}\det(\mathcal{L})^{1/n}$.*

In the next lecture, we will use Minkowski's theorem to stumble upon the LLL algorithm, which is unquestionably the most important algorithm in this field.

References

[Min10] Hermann Minkowski. *Geometrie der Zahlen.* B.G. Teubner, 1910.

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems. In *STOC*, pages 99–108. ACM, 1996.

[Ajt98] Miklós Ajtai. The shortest vector problem in ℓ_2 is NP-hard for randomized reductions. In *STOC*, pages 10–19. ACM, 1998.