My Journey into the Mordell-Weil Theorem: Understanding Rational Points on Elliptic Curves

Yunus Annanepesov

July 10, 2025

Abstract

This paper documents my exploration of the Mordell-Weil Theorem, a truly fundamental result in number theory and geometry. What's so cool about it? It essentially tells us that even though an elliptic curve might have infinitely many rational points, we can actually generate all of them from just a finite starting set. To get there, I'll first walk through some key ideas from group theory – don't worry, I'll keep it as clear as possible! Then, we'll dive into elliptic curves themselves, understanding their equations and how their rational points can actually form a group. The main part of this paper is dedicated to breaking down the proof of the Mordell-Weil Theorem. We'll see how "height functions" act as a way to measure the "complexity" of points, and how the powerful "infinite descent" method, a technique going back to Fermat, helps us prove the theorem. I'll explain each crucial lemma about how heights behave and why a certain quotient group is finite, with step-by-step explanations. My goal is to synthesize and explain these ideas from classic texts in a way that feels clear, engaging, and helps build a solid understanding of this amazing theorem where geometry and number theory beautifully intertwine.

Contents

1 Getting Started: A Look at Group Theory

Before we jump into elliptic curves, I think it's really helpful to get comfortable with some basic ideas from group theory. It gives us the perfect framework for understanding how rational points on an elliptic curve actually behave. So, let's quickly go over some fundamental concepts.

Definition 1.1 (What is a Group?)**.** Imagine you have a set of things, and a way to combine any two of them to get another thing still in that set. If this combination method (which we call an "operation") follows a few simple rules, we call the set and its operation a **group** $(G, \cdot)$. Here are the rules:

1. **Closure:** If you combine any two elements from the set, the result is always still within that set. This means the operation never takes you outside the group. For example, if you add two integers, you always get another integer.

2. **Associativity:** When you combine three or more elements, the way you group them doesn't change the final outcome. For example, $(a \cdot b) \cdot c = a \cdot (b \cdot c)$. This is pretty natural for operations like addition and multiplication.

3. **Identity Element:** There's a special "neutral" element in the set, often called e. When you combine any element with e, the element stays exactly the same. Think of adding zero to a number, or multiplying by one. This element acts like a placeholder that doesn't change anything.

4. **Inverse Element:** For every element in the set, there's another element (its "inverse") that, when combined with the original element, gives you the identity element. For example, the inverse of 5 under addition is -5, because $5 + (-5) = 0$ (which is the identity). The inverse basically "undoes" the operation.

Example 1.2 (Some Groups We Already Know)**.**

- The set of integers $\mathbb{Z}$ with the operation of addition $(+)$ is a classic example of a group.
  - Closure: Adding any two integers gives an integer.
  - Associativity: $(a + b) + c = a + (b + c)$ is always true.
  - Identity: 0 is the identity element, as $a + 0 = a$.
  - Inverse: The inverse of any integer a is $-a$, since $a + (-a) = 0$.
- The set of non-zero rational numbers $\mathbb{Q}^*$ with the operation of multiplication $(\cdot)$ also forms a group. Here, the identity is 1, and the inverse of x is $1/x$.

Definition 1.3 (Abelian Group: When Order Doesn't Matter)**.** A group $(G, \cdot)$ is called an **abelian group** (named after the mathematician Niels Henrik Abel) if, on top of the four group rules, it also has this property:

Definition 1.4 (Subgroup: A Smaller Group Inside a Bigger One)**.** Sometimes, a smaller collection of elements within a larger group can also form a group using the same operation. We call this a **subgroup**. To check if a non-empty subset H of a group G is a subgroup, you just need to verify one simple condition: if you pick any two elements a and b from H, then $a \cdot b^{-1}$ must also be in H. This single check is pretty clever because it automatically makes sure all the other group rules are met within H.

Example 1.5. The set of even integers $2\mathbb{Z} = \{\ldots, -4, -2, 0, 2, 4, \ldots\}$ is a subgroup of the integers $\mathbb{Z}$ under addition. If you take any two even numbers, say $a = 2k$ and $b = 2j$, then $a - b = 2k - 2j = 2(k - j)$, which is also an even number.

Definition 1.6 (Homomorphism: Maps That Keep Structure)**.** A **homomorphism** is a special kind of function (or map) that goes from one group to another. If you have a function f from group G to group H, it's a homomorphism if it "preserves" the group operation. What does that mean? It means if you combine two elements in G and then apply f, you get the same result as if you applied f to each element first and then combined them in H. Mathematically:

$$f(a \cdot_G b) = f(a) \cdot_H f(b)$$

Homomorphisms are super useful because they help us understand how different group structures are connected. They reveal hidden similarities between seemingly different mathematical objects.

Definition 1.7 (Kernel and Image: What a Map Tells Us)**.** Every homomorphism $f : G \to H$ has two important sets associated with it:

Definition 1.8 (Coset and Quotient Group: Grouping by Similarity)**.** Imagine you have a group G and a subgroup H inside it. You can divide G into "chunks" called **cosets**. For any element $a \in G$, the coset containing a is the set $a + H = \{a + h \mid h \in H\}$. Think of it like taking every element in H and "shifting" it by adding a. All elements in a coset are considered "similar" in some way relative to H.

The set of all these different cosets can actually form its own group! This is called the **quotient group** (or factor group), and we write it as G/H. The operation in G/H is pretty natural: you add cosets by adding their representatives: $(a + H) + (b + H) = (a + b) + H$. This works nicely because our group G is abelian.

The **index** of H in G, written as $[G : H]$, is simply the number of different cosets of H in G. If this number is finite, we say the index is finite. For example, if you look at the integers $\mathbb{Z}$ and the subgroup of even integers $2\mathbb{Z}$, the cosets are $0 + 2\mathbb{Z}$ (all even numbers) and $1 + 2\mathbb{Z}$ (all odd numbers). There are two distinct cosets, so $[\mathbb{Z} : 2\mathbb{Z}] = 2$, which is finite.

The Mordell-Weil Theorem, in a nutshell, tells us that the group of rational points on an elliptic curve, $E(\mathbb{Q})$, is a *finitely generated abelian group*. This is a huge deal! It means we can find a small, finite set of "building block" points, and by repeatedly adding and subtracting them, we can get to every single rational point on the curve. It's kind of like how every integer can be generated just from $\{1\}$ by adding or subtracting it. More formally, a big theorem about abelian groups tells us that $E(\mathbb{Q})$ looks like a combination of a finite subgroup (called the **torsion subgroup**, which contains points that eventually loop back to the identity) and a finite number of copies of $\mathbb{Z}$. The number of copies of $\mathbb{Z}$ is called the **rank** of the elliptic curve, and it tells us how "big" the infinite part of the group is.

2 Elliptic Curves: Where Geometry and Algebra Meet

Okay, now that we've got our group theory basics down, let's talk about the main stars of this paper: elliptic curves! These curves aren't just abstract math; they're fascinating places where visual geometry and precise algebra come together in really cool ways.

2.1 The Curve's Equation: The Elegant Weierstrass Form

At its heart, an elliptic curve E is a specific kind of algebraic curve. Imagine a smooth, continuous line that doesn't have any sharp corners or places where it crosses itself. In fancy math terms, it's a "smooth projective curve of genus one" with a special rational point on it, usually called O. The "genus one" thing is a topological property that helps us tell elliptic curves apart from simpler shapes like straight lines or parabolas (which have genus zero).

When we're working with rational numbers ($\mathbb{Q}$), an elliptic curve can be beautifully described by an equation called the **Weierstrass equation**. The most general version looks a bit intimidating:

$$E : y^2 + a_1xy + a_3y = x^3 + a_2x^2 + a_4x + a_6$$

where all those a_i coefficients are rational numbers. But here's a neat trick: for fields like $\mathbb{Q}$ (which don't have a "characteristic" of 2 or 3, meaning $2 \neq 0$ and $3 \neq 0$), we can always simplify this equation using some clever substitutions. This brings us to the much nicer and more common **short Weierstrass form**:

$$E : y^2 = x^3 + Ax + B$$

Here, A and B are just rational numbers. This simplified form is what we'll mostly be using to study elliptic curves over the rationals.

Now, there's one really important condition for this curve to be a "non-singular elliptic curve" – it has to be smooth everywhere. We guarantee this smoothness if its **discriminant**, Δ, is not zero. For our short Weierstrass form, the discriminant is a specific formula involving A and B:

$$\Delta = -16(4A^3 + 27B^2)$$

If $\Delta = 0$, it means the curve has a "singular point" – a spot where it pinches off (like a cusp) or crosses itself (like a node). These singular curves behave very differently, and the Mordell-Weil Theorem only applies to the smooth ones.

Finally, what about that special rational point O? We usually think of it as the "point at infinity." While it might sound a bit mysterious, in projective geometry (which is a different way of looking at geometry where points at infinity are perfectly normal), this point $(0 : 1 : 0)$ is always a smooth point on the curve and plays a super important role in how we "add" points.

2.2 Adding Points: The Geometric Group Law

This is where elliptic curves get truly amazing! One of their most surprising properties is that all the rational points on them, $E(\mathbb{Q})$, actually form an *abelian group* using a special geometric way of "adding" points. This isn't just a convenient mathematical trick; it's a deep algebraic structure that's built right into the curve itself. The point at infinity, O, acts as the "zero" or identity element for this group.

Let me try to explain how this "addition" works for two points P_1 and P_2 on our elliptic curve $E : y^2 = x^3 + Ax + B$:

Doubling a Point ($P_1 = P_2$):

The Identity Point (O):

The fact that this geometric way of adding points actually follows all the group rules (closure, associativity, identity, inverse, and commutativity) is a fundamental idea in elliptic curve theory. Commutativity is pretty easy to see visually because of the curve's symmetry, but associativity, $(P_1 + P_2) + P_3 = P_1 + (P_2 + P_3)$, is a deeper result that usually needs more advanced algebraic geometry to prove.

What's really important is that these geometric operations can be turned into exact algebraic formulas. For points $P_1 = (x_1, y_1)$ and $P_2 = (x_2, y_2)$ on $y^2 = x^3 + Ax + B$, the coordinates of their sum $P_3 = (x_3, y_3)$ are given by: First, calculate the slope λ:

2.3 Points That Loop: The Nagell-Lutz Theorem

Among all the rational points on an elliptic curve, some are special because they have "finite order." This means if you keep adding such a point to itself, you'll eventually get back to the identity point O. For instance, a point P has order 2 if $P + P = O$, which geometrically means P is its own reflection across the x-axis, so its y-coordinate must be 0.

The Nagell-Lutz Theorem is a really elegant and useful result that helps us find these finite-order rational points on an elliptic curve, especially when the coefficients are integers.

Theorem 2.1 (Nagell-Lutz Theorem: Pinpointing Finite-Order Points)**.** *Let $E : y^2 = x^3 + Ax + B$ be a non-singular elliptic curve, where A and B are integers. If $P = (x, y)$ is a rational point on E that has finite order (meaning $nP = O$ for some positive integer n), then two pretty amazing things must be true:*

Part 1: Why x and y have to be integers. *Imagine a rational point $P = (x, y)$. We can write its coordinates in a special "lowest terms" way that's good for elliptic curves: $x = m/n^2$ and $y = k/n^3$, where m, k, n are integers, $n > 0$, and they don't share any common prime factors. This form makes analyzing denominators easier.*

If you plug these into the Weierstrass equation $y^2 = x^3 + Ax + B$ and clear the denominators, you get something like:

$$k^2 = m(m^2 + Amn^2 + Bn^4)$$

Now, let's think: what if n isn't 1? That means there's some prime number p that divides n. Because we wrote x and y in lowest terms, p can't divide m or k. If you look at the equation above modulo p, the terms with n in them disappear. So you're left with $k^2 \equiv m^3 \pmod{p}$. The deeper argument involves looking at how many times a prime p divides the denominators of $P, 2P, 3P, \ldots$. If P has finite order, these denominators can't keep getting "more divisible" by p forever. But if $n \neq 1$, a careful look at the formulas for doubling a point shows that the denominator of $2P$ would actually be divisible by a higher power of p than P's denominator. This would create an "infinite descent" on the p-adic valuation of the denominator, which isn't possible because valuations are just non-negative integers. The only way to avoid this endless descent is if the denominator n is 1, meaning x and y must be integers.

Part 2: Why y^2 divides Δ. *Once we know x and y are integers, the second part of the theorem comes into play. If $y = 0$, then $x^3 + Ax + B = 0$, and these points $(x, 0)$ are exactly the points of order 2. If $y \neq 0$, the proof involves looking at the coordinates modulo primes that divide y. It turns out that if P is a torsion point, its y-coordinate* *has to satisfy the condition that y^2 divides the discriminant Δ. This comes from studying how the group law works in local fields and connecting it to the discriminant, which tells us about the curve's singularities. If y^2 didn't divide Δ, it would imply that the point couldn't have finite order, because its coordinates would keep growing indefinitely with repeated additions.*

3 Measuring Complexity: The Idea of Height

To really get a handle on the Mordell-Weil Theorem, which is all about whether rational points can be "finitely generated," we need a way to quantify how "complicated" or "big" a rational point is. This is where the brilliant idea of a ***height function*** *comes in. Think of it as an arithmetic ruler that gives a positive number to each rational point, telling us about its complexity.*

3.1 The Height of a Rational Number: Our First Ruler

Let's start with the simplest things: rational numbers.

Definition 3.1 (Height of a Rational Number)**.** For any rational number x, we always write it in its simplest form, or "lowest terms," as a fraction m/n. This means m and n are integers, n is positive ($n > 0$), and they don't share any common factors ($\gcd(|m|, |n|) = 1$). The **height** of x, written $H(x)$, is simply the larger of the absolute values of its numerator and denominator:

$$H(x) = \max\{|m|, |n|\}$$

For the special case of $x = 0$, which we can write as $0/1$, we define $H(0) = 1$.

Let's try a few examples to get a feel for this:

Now for a really important property of this height function – its **finiteness property**. This might seem simple, but it's the absolute foundation for our "descent" argument later.

Proposition 3.2 (Height's Finiteness: Only So Many "Simple" Fractions)**.** *For any positive real number M, there are only a finite number of rational numbers x whose height $H(x)$ is less than or equal to M. In other words, the set $\{x \in \mathbb{Q} : H(x) \leq M\}$ is finite.*

3.2 The Height of a Rational Point: Extending Our Ruler to Curves

Now we take our height concept and apply it to the rational points on an elliptic curve.

Definition 3.3 (Height of a Rational Point)**.** Let $E : y^2 = x^3 + Ax + B$ be an elliptic curve, and let $P = (x, y)$ be a rational point on E. The **height** of P, written $H(P)$, is simply the height of its x-coordinate:

$$H(P) = H(x)$$

For the point at infinity O, we define $H(O) = 1$.

For practical math, especially when we're dealing with inequalities and how things grow, it's often easier to use an additive version of the height, called the **logarithmic height** (or sometimes just "small h height"):

$$h(P) = \log H(P)$$

Since $H(P) \geq 1$, the logarithmic height $h(P)$ is always a non-negative real number ($\log 1 = 0$). The finiteness property we just discussed extends directly to rational points: for any M, the set $\{P \in E(\mathbb{Q}) : H(P) \leq M\}$ (or equivalently $\{P \in E(\mathbb{Q}) : h(P) \leq \log M\}$) is finite. This is because if $H(P) \leq M$, then $H(x) \leq M$. By Proposition 3.2, there are only finitely many possible x-coordinates. For each such x, the equation $y^2 = x^3 + Ax + B$ can give us at most two corresponding y-coordinates (which are rational if $x^3 + Ax + B$ is the square of a rational number). So, the total number of rational points with height less than or equal to M is indeed finite. This means that "small" points (in terms of height) are always limited in number.

4 How Heights Behave: The Dance of Growth and Shrinkage

The real power of height functions becomes clear when we see how they change when we apply the group operations on an elliptic curve. The next two lemmas are absolutely central to Mordell's proof; they're the mathematical engine that drives the "descent" argument. They basically tell us how the arithmetic complexity of a point changes when we add it to another point, or, most importantly, when we double it.

Lemma 4.1 (Adding a Fixed Point: Height Grows Predictably)**.** *Let $E : y^2 = x^3 + Ax + B$ be an elliptic curve with integer coefficients A, B. If we pick any fixed rational point $P_0 = (x_0, y_0)$ on E, then there's a constant κ_0 (which depends only on P_0 and the curve's coefficients A, B) such that for every other rational point $P = (x, y) \in E(\mathbb{Q})$:*

$$h(P + P_0) \leq 2h(P) + \kappa_0$$

Remember the formula for the x-coordinate of the sum of two distinct points P and P_0:

$$\xi = \lambda^2 - x - x_0$$

where λ is the slope of the line connecting P and P_0. If $x \neq x_0$, then $\lambda = \frac{y-y_0}{x-x_0}$. We can play around with this expression using the curve equations $y^2 = x^3 + Ax + B$ and $y_0^2 = x_0^3 + Ax_0 + B$. If we subtract the second from the first, we get $y^2 - y_0^2 = (x^3 - x_0^3) + A(x - x_0)$. Factoring both sides, we get $(y - y_0)(y + y_0) = (x - x_0)(x^2 + xx_0 + x_0^2 + A)$. So, the slope can also be written as $\lambda = \frac{x^2+xx_0+x_0^2+A}{y+y_0}$.

Now, substitute this back into the formula for ξ:

$$\xi = \left(\frac{x^2 + xx_0 + x_0^2 + A}{y + y_0}\right)^2 - x - x_0$$

This looks complicated because it still has y. But here's the key idea: ξ can actually be written as a rational function of x (meaning a ratio of two polynomials in x). If we clear denominators and use $y^2 = x^3 + Ax + B$ to get rid of y^2 terms, we find that ξ is *indeed a rational function of x where the highest power of x in both the numerator and denominator is 2. Let's call this rational function $f(x) = \frac{N(x)}{D(x)}$, where $N(x)$ and $D(x)$ are polynomials in x with integer coefficients (after some scaling, since A, B, x_0, y_0 are rational). The maximum degree of $N(x)$ and $D(x)$ is 2.*

There's a powerful general result in number theory (often called the fundamental inequality for heights of rational functions, or specifically, Lemma 3.6 in Silverman and Tate [1]) that says if $f(x) = \frac{\phi(x)}{\psi(x)}$ is a rational function where $\phi(x)$ and $\psi(x)$ are coprime polynomials (meaning they don't share any common roots) with integer coefficients, and $d = \max(\deg \phi, \deg \psi)$, then for any rational number x, the height of $f(x)$ is roughly $H(x)^d$. In terms of logarithmic heights, $h(f(x)) \approx d \cdot h(x)$.

In our case, $d = 2$. So, $h(\xi)$ is approximately $2h(x)$. The constant κ_0 in the lemma accounts for the exact relationship, including any "error" terms that might come from common factors between the numerator and denominator of ξ when x is plugged in in lowest terms, and the specific values of A, B, x_0, y_0. This constant κ_0 is indeed fixed once P_0 and the curve E are fixed.

Special cases like $P = P_0$ (doubling), $P = -P_0$ (sum is O), or $P_0 = O$ (sum is P) are either covered by the general formula or trivially satisfy the inequality. For example, if $P_0 = O$, then $P + O = P$, so $h(P + O) = h(P)$. The inequality $h(P) \leq 2h(P) + \kappa_0$ holds for any positive κ_0 (since $h(P) \geq 0$).

This lemma tells us that adding a fixed point P_0 to any point P on the curve won't make its height spiral out of control. The height of the sum grows at most quadratically with the height of P. This controlled growth is important, but the next lemma is even more crucial for our descent argument.

Lemma 4.2 (Doubling a Point: Height Explodes!). *Let $E : y^2 = x^3 + Ax + B$ be an elliptic curve with integer coefficients A, B. There's a constant κ, which depends only on the coefficients A, B of E, such that for all rational points $P = (x, y) \in E(\mathbb{Q})$:*

$$h(2P) \geq 4h(P) - \kappa$$

The x-coordinate of $2P$ is then given by the duplication formula:

$$\xi = \lambda^2 - 2x = \left(\frac{3x^2 + A}{2y}\right)^2 - 2x$$

To get ξ only in terms of x, we substitute $y^2 = x^3 + Ax + B$:

$$\xi = \frac{(3x^2 + A)^2}{4y^2} - 2x = \frac{(3x^2 + A)^2}{4(x^3 + Ax + B)} - 2x$$

Now, let's combine these terms over a common denominator:

$$\xi = \frac{(3x^2 + A)^2 - 8x(x^3 + Ax + B)}{4(x^3 + Ax + B)}$$

Let's expand the top part (the numerator):

$$(3x^2 + A)^2 = 9x^4 + 6Ax^2 + A^2$$

So, the numerator becomes:

$$9x^4 + 6Ax^2 + A^2 - 8x^4 - 8Ax^2 - 8Bx = x^4 - 2Ax^2 - 8Bx + A^2$$

Thus, ξ can be written as a rational function of x:

$$\xi = \frac{x^4 - 2Ax^2 - 8Bx + A^2}{4(x^3 + Ax + B)}$$

Let $F(x) = x^4 - 2Ax^2 - 8Bx + A^2$ (the numerator) and $G(x) = 4(x^3 + Ax + B)$ (the denominator). Both $F(x)$ and $G(x)$ are polynomials with integer coefficients (since A, B are integers).

The highest degree of these polynomials is $d = \max(\deg F, \deg G) = \max(4, 3) = 4$. This degree of 4 is super important.

Now, we need to make sure that $F(x)$ and $G(x)$ are "coprime" as polynomials. This means they don't share any common roots in the complex numbers. If they did share a common root x_c, that would mean x_c is a root of $G(x)$, so $x_c^3 + Ax_c + B = 0$. This would imply that $(x_c, 0)$ is a point on the curve. If x_c is also a root of $F(x)$, it would mean that the tangent at $(x_c, 0)$ (which is a vertical line) intersects the curve with multiplicity greater than 2, which would mean the curve has a singular point there. But we started by assuming E is a non-singular elliptic curve, which guarantees that $x^3 + Ax + B$ has distinct roots and no common roots with its derivative. This ensures that $F(x)$ and $G(x)$ are indeed coprime polynomials.

Since $F(x)$ and $G(x)$ are coprime polynomials with integer coefficients, and their maximum degree is $d = 4$, we can use that fundamental inequality for heights of rational functions (Lemma 3.6 in Silverman and Tate [1]). This theorem states that there are constants κ_1, κ_2 (which depend only on A and B) such that for any rational number x:

$$d \cdot h(x) - \kappa_1 \leq h\left(\frac{F(x)}{G(x)}\right) \leq d \cdot h(x) + \kappa_2$$

Plugging in $d = 4$, and remembering that $h(2P) = h(\xi) = h(F(x)/G(x))$ and $h(P) = h(x)$, we get:

$$4h(P) - \kappa_1 \leq h(2P) \leq 4h(P) + \kappa_2$$

The lemma specifically asks for a lower bound, so we can just pick $\kappa = \kappa_1$. And that completes the proof! The constant κ just accounts for any small "loss" in height that might happen if there are common factors between the numerator and denominator of ξ when x is plugged in in lowest terms.

This lemma is the real workhorse for our descent argument. It shows that when we keep doubling a rational point on an elliptic curve, its logarithmic height grows pretty fast – roughly four times bigger at each step. This quick, predictable growth is exactly what lets us use the "infinite descent" method: if we have a point that's "too big," we can always find a "smaller" one by repeatedly "halving" it (in the group sense), and this process has to stop eventually because heights can't go on getting infinitely small.

5 The Weak Mordell-Weil Theorem: A Key Finite Step

The last, and maybe most algebraically complex, piece we need for proving the Mordell-Weil Theorem is showing that a particular quotient group is finite. This result is known as the ***Weak Mordell-Weil Theorem****.*

Lemma 5.1 (Weak Mordell-Weil Theorem: Points Grouped into Finite Families)**.** *The index $[E(\mathbb{Q}) : 2E(\mathbb{Q})]$ is finite. To put it simply, if you take all the rational points on an* *elliptic curve E, and then you look at the subgroup formed by doubling all those points ($2E(\mathbb{Q})$), there are only a finite number of distinct "types" or "families" of points. Two points are in the same family if their difference is a point that can be obtained by doubling another point.*

5.1 Curves with a Point of Order Two: The ϕ-Descent Trick

Let's consider an elliptic curve E that can be written in this specific form: $y^2 = x^3 + Ax^2 + Bx$, where A and B are integers. If you quickly check, you'll see a special point: if we plug in $(x, y) = (0, 0)$ into this equation, we get $0^2 = 0^3 + A(0)^2 + B(0)$, which is $0 = 0$. So, $T = (0, 0)$ is definitely a rational point on the curve. What's the "order" of T? To find out, we look at the tangent line at $(0, 0)$. If we differentiate the curve equation implicitly with respect to x: $2y\frac{dy}{dx} = 3x^2 + 2Ax + B$. At the point $(0, 0)$, this equation becomes $0 = B$. If $B \neq 0$, this means the derivative $\frac{dy}{dx}$ is undefined (the tangent is a vertical line). A vertical tangent at $(0, 0)$ means that the line $x = 0$ intersects the curve at $(0, 0)$ twice, and its third intersection point is O (the point at infinity). According to our group law, $T + T = 2T$ is the reflection of this third point across the x-axis. Since the third point is O, and O is its own reflection, we conclude that $2T = O$. So, if $B \neq 0$, $T = (0, 0)$ is a rational point of order two. (If $B = 0$, the curve equation becomes $y^2 = x^3 + Ax^2 = x^2(x + A)$, which has a singularity at $(0, 0)$. Since the Mordell-Weil Theorem applies only to non-singular elliptic curves, we can safely assume $B \neq 0$.)

For these curves $E : y^2 = x^3 + Ax^2 + Bx$ (with $B \neq 0$), we introduce a clever trick: an auxiliary elliptic curve E' defined by a similar equation:

$$E' : y'^2 = x'^3 + A'x'^2 + B'x'$$

where $A' = -2A$ and $B' = A^2 - 4B$. Just like E, for E' to be non-singular, its discriminant $\Delta' = -16B'^2(A'^2 - 4B')$ must be non-zero.

Now, here's the cool part: there are two important homomorphisms that link the groups of rational points $E(\mathbb{Q})$ and $E'(\mathbb{Q})$:

The truly amazing thing about these maps is what happens when you combine them. If you apply ϕ and then ψ, or vice-versa, you end up with the same result as just multiplying the point by 2:

$$\psi \circ \phi(P) = 2P \quad \textit{for all } P \in E(\mathbb{Q})$$

$$\phi \circ \psi(P') = 2P' \quad \textit{for all } P' \in E'(\mathbb{Q})$$

This means that "doubling a point" on $E(\mathbb{Q})$ can be "broken down" into two simpler steps: mapping it to $E'(\mathbb{Q})$ and then mapping it back to $E(\mathbb{Q})$. This decomposition is the core idea behind the "descent via isogeny" method.

5.2 The α Homomorphism: Finding Square-Free Clues

To prove the Weak Mordell-Weil Theorem (Lemma 5.1), we introduce another clever homomorphism, called α. This map takes points on E and tells us something useful about their x-coordinates, specifically what they are modulo squares.

Definition 5.2 (The α Homomorphism)**.** Let $E : y^2 = x^3 + Ax^2 + Bx$ be our elliptic curve. We define a homomorphism $\alpha : E(\mathbb{Q}) \to \mathbb{Q}^*/(\mathbb{Q}^*)^2$ like this:

The fact that α is actually a homomorphism (meaning $\alpha(P_1 + P_2) = \alpha(P_1) \cdot \alpha(P_2)$ in $\mathbb{Q}^*/(\mathbb{Q}^*)^2$) is a non-trivial algebraic property that needs careful checking using the elliptic curve addition formulas.

The cool thing about α is how it connects with ψ: the kernel of α is exactly the image of ψ, so $\ker(\alpha) = \psi(E'(\mathbb{Q}))$. Because of a theorem called the First Isomorphism Theorem for groups, this means that $E(\mathbb{Q})/\psi(E'(\mathbb{Q})) \cong \operatorname{Im}(\alpha)$. So, to show that the index $[E(\mathbb{Q}) : \psi(E'(\mathbb{Q}))]$ is finite, we just need to prove that the image of α, $\operatorname{Im}(\alpha)$, is a finite group.

Let's see why $\operatorname{Im}(\alpha)$ is finite. Take any point $P = (x, y) \in E(\mathbb{Q})$ with $x \neq 0$. From our earlier discussion about the Nagell-Lutz theorem, we know we can write x as m/n^2 and y as k/n^3, where m, k, n are integers, $n > 0$, and they don't share any common prime factors. If you substitute these into the curve equation $y^2 = x^3 + Ax^2 + Bx$ and multiply by n^6 to clear all denominators, you get:

$$k^2 = m(m^2 + Amn^2 + Bn^4)$$

Now, let $d = \gcd(m, m^2 + Amn^2 + Bn^4)$. Since d divides m, it must also divide the whole second factor. So, d must divide their difference:

$$d \mid (m^2 + Amn^2 + Bn^4) - m(m + An^2) = Bn^4$$

Since $x = m/n^2$ is in lowest terms, $\gcd(m, n) = 1$, which means $\gcd(m, n^4) = 1$. Since d divides m and Bn^4, and m has no common factors with n^4, it means d must divide B. This is a really important observation! Since $k^2 = m \cdot (\text{some integer})$, and the only common factors between m and that "some integer" are divisors of B, it implies that m must be of the form b_1u^2 for some integer u, where b_1 is a square-free integer whose prime factors only come from the prime factors of B. Specifically, we can write $m = b_1u^2$ and $m^2 + Amn^2 + Bn^4 = b_2v^2$, where $b_1b_2 = B$ (ignoring squares). Then, the x-coordinate $x = m/n^2 = (b_1u^2)/n^2 = b_1(u/n)^2$. Therefore, $x \pmod{(\mathbb{Q}^*)^2}$ must be equivalent to $b_1 \pmod{(\mathbb{Q}^*)^2}$.

The possible values for b_1 are the square-free divisors of B. Since B is a fixed integer from the curve's equation, it only has a finite number of prime factors. This means there are only a finite number of square-free divisors of B. This directly shows that the image of α, $\operatorname{Im}(\alpha)$, is a finite group. For example, if $B = 12 = 2^2 \cdot 3$, the square-free divisors are $\pm 1, \pm 2, \pm 3, \pm 6$. These are the only possible values for $x \pmod{(\mathbb{Q}^*)^2}$.

We can define a similar homomorphism $\alpha' : E'(\mathbb{Q}) \to \mathbb{Q}^*/(\mathbb{Q}^*)^2$ for the auxiliary curve E'. And similarly, we can show that $\operatorname{Im}(\alpha')$ is also a finite group. The kernel of α' is $\phi(E(\mathbb{Q}))$. So, the index $[E'(\mathbb{Q}) : \phi(E(\mathbb{Q}))]$ is also finite.

5.3 Putting It All Together: Why $E(\mathbb{Q})/2E(\mathbb{Q})$ Is Finite

The very last step in proving the Weak Mordell-Weil Theorem (Lemma 5.1) from our ϕ-descent involves a powerful general group theory result about something called "isogenies." An isogeny between elliptic curves is basically a homomorphism that has a finite kernel. Our maps ϕ and ψ are indeed isogenies.

Lemma 5.3 (Descent Lemma: The Finiteness Connection)**.** *Let A and B be abelian groups. Suppose we have homomorphisms $\phi : A \to B$ and $\psi : B \to A$ such that when you compose them, you get multiplication by 2: $\psi \circ \phi(P) = 2P$ for all $P \in A$, and $\phi \circ \psi(P') = 2P'$ for all $P' \in B$. If:*

Consider this chain of subgroups: $2E(\mathbb{Q}) \subseteq \psi(E'(\mathbb{Q})) \subseteq E(\mathbb{Q})$*. We can express the total index as a product of these smaller indices:*

$$[E(\mathbb{Q}) : 2E(\mathbb{Q})] = [E(\mathbb{Q}) : \psi(E'(\mathbb{Q}))] \cdot [\psi(E'(\mathbb{Q})) : 2E(\mathbb{Q})]$$

We already know that $[E(\mathbb{Q}) : \psi(E'(\mathbb{Q}))]$ *is finite (from point 3 above). So, the remaining task is to show that* $[\psi(E'(\mathbb{Q})) : 2E(\mathbb{Q})]$ *is finite. Remember that* $2E(\mathbb{Q}) = \psi(\phi(E(\mathbb{Q})))$*. So we're interested in the index* $[\psi(E'(\mathbb{Q})) : \psi(\phi(E(\mathbb{Q})))]$*.*

Using properties of indices more directly, we know:

The key insight from the theory of isogenies is that the finiteness of the kernels AND the finiteness of the images of the α *maps (which are related to something called the cokernels of* ϕ *and* ψ*) are exactly what we need to put a bound on the size of* $E(\mathbb{Q})/2E(\mathbb{Q})$*. Specifically, a theorem shows that:*

$$[E(\mathbb{Q}) : 2E(\mathbb{Q})] = \frac{[E(\mathbb{Q}) : \psi(E'(\mathbb{Q}))] \cdot [E'(\mathbb{Q}) : \phi(E(\mathbb{Q}))]}{|\ker(\phi)| \cdot |\ker(\psi)|}$$

Since we've already shown that $[E(\mathbb{Q}) : \psi(E'(\mathbb{Q}))]$ *and* $[E'(\mathbb{Q}) : \phi(E(\mathbb{Q}))]$ *are finite, and* $|\ker(\phi)| = 2$ *and* $|\ker(\psi)| = 2$ *are finite, their product is also finite. Therefore,* $[E(\mathbb{Q}) : 2E(\mathbb{Q})]$ *must be finite. This is the crucial step that connects our number-theoretic analysis of the* α *maps to the group-theoretic structure we need for the descent.*

6 The Grand Finale: Proving the Mordell-Weil Theorem

Now, with all the necessary preliminary results carefully laid out, we're ready to bring them all together and see the full power of the Mordell-Weil Theorem's proof. This proof is a truly magnificent demonstration of the method of infinite descent, brilliantly powered by the properties of height functions we've explored.

Theorem 6.1 (The Mordell-Weil Theorem: You Can Generate All Rational Points from a Finite Set!)**.** *Let* E *be a non-singular elliptic curve defined over the rational numbers* $\mathbb{Q}$*. Then the group of rational points* $E(\mathbb{Q})$ *is a finitely generated abelian group.*

The Heart of the Proof: The Descent Strategy

The proof strategy is a beautiful application of Fermat's famous method of infinite descent. We want to show that all points in Γ *can be constructed from a finite set.*

Step 1: Picking Our "Starting Points" (Coset Representatives). *From the Weak Mordell-Weil Theorem (Lemma 5.1), we know that the quotient group* $\Gamma/2\Gamma$ *is finite. Let's say it has* n *elements, so* $n = [\Gamma : 2\Gamma]$*. This means we can choose a finite collection of points from* Γ*, let's call them* $Q_1, Q_2, \ldots, Q_n$*. Every single point in* Γ *will belong to exactly one of the "families" defined by these* Q_j *points (specifically, one of the cosets* $Q_j + 2\Gamma$*). Think of these* Q_j *as a finite set of "representatives" for all the different "types" of points in* Γ*.*

Step 2: Breaking Down Any Point P***.*** *Now, pick any point* $P \in \Gamma$*. Since* P *must belong to one of these* n *families (cosets), there has to be some representative* Q_{i_1} *(where* i_1 *is an index from* 1 *to* n*) such that* P *is in the same family as* Q_{i_1}*. This means their difference,* $P - Q_{i_1}$*, must be an element of the subgroup* 2Γ *(the points that are doubles of other points). So, we can write:*

$$P - Q_{i_1} = 2P_1$$

for some point $P_1 \in \Gamma$. This P_1 is essentially "half" of the difference between P and its representative Q_{i_1}.

Step 3: Repeating the Descent. *We can do the exact same thing for P_1. Since $P_1 \in \Gamma$, it also belongs to some family. So, there's some $Q_{i_2} \in \{Q_1, \dots, Q_n\}$ such that $P_1 - Q_{i_2} \in 2\Gamma$. This lets us write:*

$$P_1 - Q_{i_2} = 2P_2$$

for some point $P_2 \in \Gamma$.

We can keep repeating this process. For each point P_j in our sequence, we find a representative $Q_{i_{j+1}}$ such that $P_j - Q_{i_{j+1}} = 2P_{j+1}$ for some $P_{j+1} \in \Gamma$. This creates a sequence of points:

$$P_0 = P, P_1, P_2, \dots, P_m, \dots$$

and a corresponding sequence of representatives $Q_{i_1}, Q_{i_2}, \dots, Q_{i_m}, \dots$.

Step 4: Writing P with Our Generators. *Now, let's substitute these relationships back into the first one to express our original point P in terms of the Q_j representatives and the very last point in our sequence: From $P_0 = Q_{i_1} + 2P_1$, we have $P = Q_{i_1} + 2P_1$. Substitute $P_1 = Q_{i_2} + 2P_2$: $P = Q_{i_1} + 2(Q_{i_2} + 2P_2) = Q_{i_1} + 2Q_{i_2} + 4P_2$. If we continue this for m steps, we get a really powerful expression for P:*

$$P = Q_{i_1} + 2Q_{i_2} + 4Q_{i_3} + \dots + 2^{m-1}Q_{i_m} + 2^m P_m$$

This equation shows that any point $P \in \Gamma$ can be written as a sum involving powers of 2 multiplied by elements from our finite set of coset representatives Q_j, plus a final term $2^m P_m$. The key idea now is to show that this sequence of P_m points can't just keep getting "smaller" forever.

Step 5: The Height Descent in Action.

Let's look at how the height of P_j relates to the height of P_{j-1}. We have the relationship $2P_j = P_{j-1} - Q_{i_j}$. From Lemma 4.2, we know that the height of a doubled point grows significantly: $h(2P_j) \geq 4h(P_j) - \kappa$. So, we can write:

$$4h(P_j) - \kappa \leq h(2P_j) = h(P_{j-1} - Q_{i_j})$$

Now, let's use Lemma 4.1 for $h(P_{j-1} - Q_{i_j})$. Since Q_{i_j} is one of our finite set of representatives, there's a constant $\kappa_{Q_{i_j}}$ such that $h(P_{j-1} - Q_{i_j}) \leq 2h(P_{j-1}) + \kappa_{Q_{i_j}}$. To simplify, let κ' be the largest of all such constants for all Q_j and their inverses (since $P - Q_j$ is the same as $P + (-Q_j)$). Then:

$$h(P_{j-1} - Q_{i_j}) \leq 2h(P_{j-1}) + \kappa'$$

Combining these two inequalities, we get:

$$4h(P_j) - \kappa \leq 2h(P_{j-1}) + \kappa'$$

Now, let's rearrange this to get $h(P_j)$ by itself:

$$4h(P_j) \leq 2h(P_{j-1}) + \kappa + \kappa'$$

$$h(P_j) \leq \frac{1}{2}h(P_{j-1}) + \frac{\kappa + \kappa'}{4}$$

Let $C = \frac{\kappa+\kappa'}{4}$. This constant C is fixed and depends only on the elliptic curve E and our finite set of coset representatives $\{Q_1, \ldots, Q_n\}$. Our crucial inequality now becomes:

$$h(P_j) \leq \frac{1}{2}h(P_{j-1}) + C$$

This inequality is the heart of the descent argument. It tells us that if the height of P_{j-1} is big enough (specifically, if $h(P_{j-1}) > 2C$), then the height of P_j will be strictly smaller than $h(P_{j-1})$. Here's why:

$$h(P_j) \leq \frac{1}{2}h(P_{j-1}) + C < \frac{1}{2}h(P_{j-1}) + \frac{1}{2}h(P_{j-1}) = h(P_{j-1})$$

This means that as we repeatedly apply our descent procedure ($P \to P_1 \to P_2 \ldots$), the heights of the points in the sequence $P, P_1, P_2, \ldots$ will rapidly decrease, as long as their height is above that $2C$ limit. Since logarithmic heights are non-negative, this sequence of strictly decreasing heights can't go on forever. It has to stop eventually. That means there must be some integer m where the height of P_m falls below or equals $2C$:

$$h(P_m) \leq 2C$$

Step 6: The Final Conclusion: Finite Generation!

We've shown that any point $P \in \Gamma$ can be written in this form:

$$P = \sum_{j=1}^{m} 2^{j-1}Q_{i_j} + 2^m P_m$$

where:

Therefore, every single point $P \in \Gamma$ can be expressed as a combination (using integer coefficients, which are powers of 2) of elements taken from the finite set $\{Q_1, \ldots, Q_n\} \cup S_{bounded}$. This means that the entire group $\Gamma = E(\mathbb{Q})$ is indeed generated by a finite set of points. And that, my friends, concludes the elegant and profound proof of the Mordell-Weil Theorem!

References

[1] Silverman, J.H. and Tate, J.T., 2015. Rational Points on Elliptic Curves. Springer.

[2] Cassels, J.W.S., 1991. Lectures on Elliptic Curves. Cambridge University Press.

[3] Washington, L.C., 2008. Elliptic Curves: Number Theory and Cryptography. CRC Press.

[4] Knapp, A.W., 1992. Elliptic Curves. Princeton University Press.