# On the Geometry of the Multiplication Table Modulo N

Ofer Barasofsky

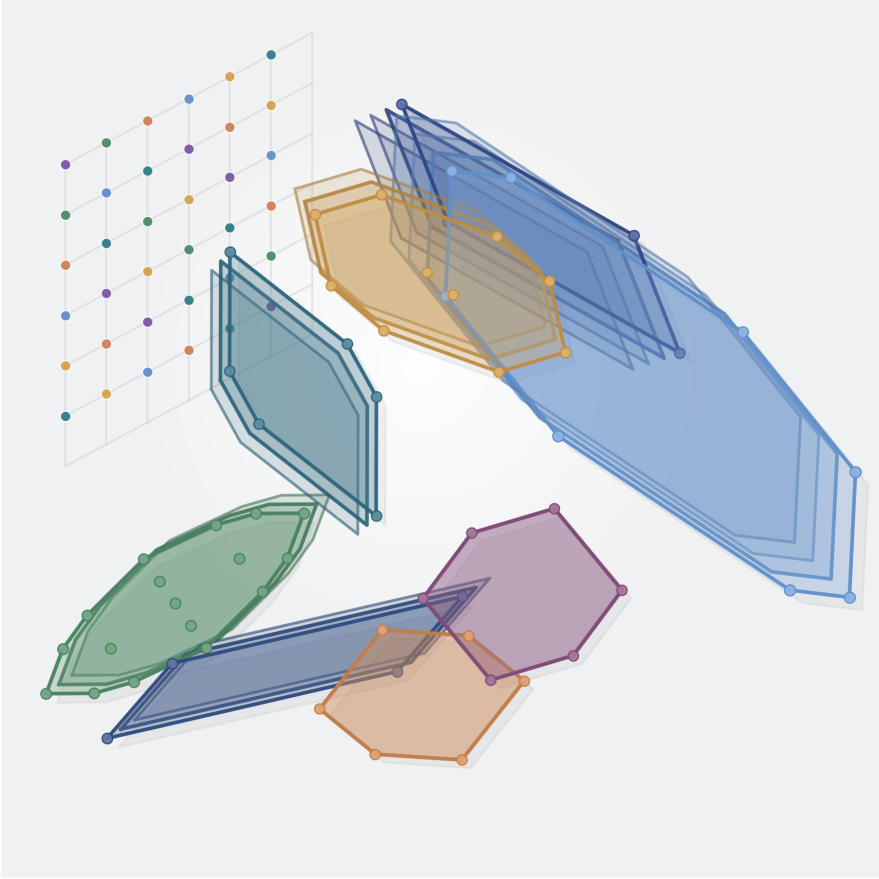

# Contents

# 1 Preface

The multiplication table is one of the first structured objects encountered in arithmetic. It appears simple enough to be taken for granted: a grid of products, a device for learning calculation, something to be mastered and then set aside. Yet when its entries are viewed through the lens of modular arithmetic, the familiar grid begins to reveal a surprising geometric structure.

Take the products on $1, \ldots, N-1$, reduce them modulo $N$, and place equal residues on the lattice $\{1, \ldots, N-1\}^2$. The points corresponding to a fixed residue $a$ are precisely the solutions of the congruence

$$xy \equiv a \pmod{N}.$$

Each residue class therefore traces a discrete curve across the lattice—a finite modular hyperbola contained within the square window of the table. These single-class sets are already part of an established literature on modular hyperbolas. This book returns to them from a different angle: it studies the full residue family at once and asks what their convex hulls and areas reveal about the arithmetic of the modulus.

What emerges from this simple construction is a geometry hidden inside the multiplication table itself. Residue classes arrange themselves into shapes with striking symmetries. The behavior of those shapes changes noticeably between prime and composite moduli. Some residues stretch across the table in long diagonal structures; others cluster tightly into compact forms. When the convex hulls of these point sets are measured and compared, a new numerical landscape begins to appear.

My own path into this subject began visually rather than formally. I started by drawing small examples and grouping together entries in the table that shared the same residue. Very quickly the patterns refused to behave like mere numerical coincidences. The table kept producing envelopes, symmetries, and geometric contrasts that suggested a deeper organizing principle. What had seemed like a simple arithmetic grid began to behave like a geometric object in its own right.

Over more than twenty years the project developed slowly and irregularly. It began with hand-drawn sketches, continued through computer experiments, and eventually grew into

large-scale computations and more systematic questions. Some early conjectures survived only in modified form, and others did not survive at all. I have not tried to erase that history entirely. A failed conjecture can still illuminate which definitions matter, which comparisons are meaningful, and which problems deserve to remain open.

For that reason this book moves between two modes. It aims to be mathematically precise, but it also pauses when necessary to explain notation and ideas that might otherwise interrupt the geometric narrative. The subject sits at the intersection of modular arithmetic, combinatorics, convex geometry, and lattice methods. I do not assume that the reader arrives fluent in all of those languages. When a concept becomes important, it will be introduced carefully rather than taken for granted.

There is also a broader mathematical tradition behind the project. Many important developments have begun with the discovery that a familiar arithmetic object secretly carries a geometric interpretation. In a smaller and more stubbornly discrete setting, the multiplication table appears to belong to that same lineage. The aim here is not to decorate arithmetic with pictures, but to let geometry organize patterns that are already present in the arithmetic itself.

Recent advances in computational tools, including AI-assisted systems, have made it possible to explore these patterns at a scale that would have been difficult to sustain alone. They have helped generate examples, test conjectures, reorganize computations, and check arguments. They do not replace proof, and they certainly do not remove the need for mathematical judgment, but they do make long exploratory investigations more feasible. The mathematical direction and the final responsibility remain mine.

## 1.1 Acknowledgments

I am grateful to Professor Norman Wildberger for the suggestion that the residue-area data should also be studied through polynomials. I also thank Dr. Alex Eizenberg, PhD, Senior Lecturer (retired) at Azrieli College of Engineering Jerusalem, for first introducing the multiplication table modulo $N$ to me; Dr. Alex Reznik of Azrieli College of Engineering Jerusalem, for listening seriously to these ideas in their vague early form, helping guide that early intuition into more productive channels, and accepting me into his class on basic quantum mechanics; Professor Daniel Zajfman and members of the mathematics faculty

at the Weizmann Institute of Science, for reading a very early short version of these ideas despite its rough language and many mistakes, and responding in a gracious and encouraging spirit; Moshe Newman, for reading, commenting, checking some results, and offering encouragement and further ideas; and my uncle, Gideon Barasofsky, for introducing me to Pick's theorem.

This book does not claim to close the subject. Its purpose is more modest and, I hope, more useful: to establish a clear language for the geometry of the multiplication table modulo $N$, to record the first precise results that arise from that perspective, and to place the reader close enough to the patterns that further exploration becomes natural.

## 2 Introduction

For a fixed integer $N \geq 2$, consider the array whose $(x, y)$-entry is the residue of $xy$ modulo $N$, where $1 \leq x, y \leq N-1$. This is the **multiplication table modulo** $N$, which we abbreviate as **MTMN**. The starting data is therefore entirely finite and arithmetic: one studies multiplication on the positive integers $1, \ldots, N-1$ and then asks what survives after passing to residues. Each residue class – that is, each possible remainder $a \in \{0, 1, \ldots, N-1\}$ – determines a set of lattice points

$$A_{N,a} = \{(x, y) \in \{1, \ldots, N-1\}^2 : xy \equiv a \pmod{N}\}.$$

Figure 1 shows a concrete example: the multiplication table modulo $5$, with the convex hull of each residue class overlaid. This is the basic geometric picture that the rest of the book develops.

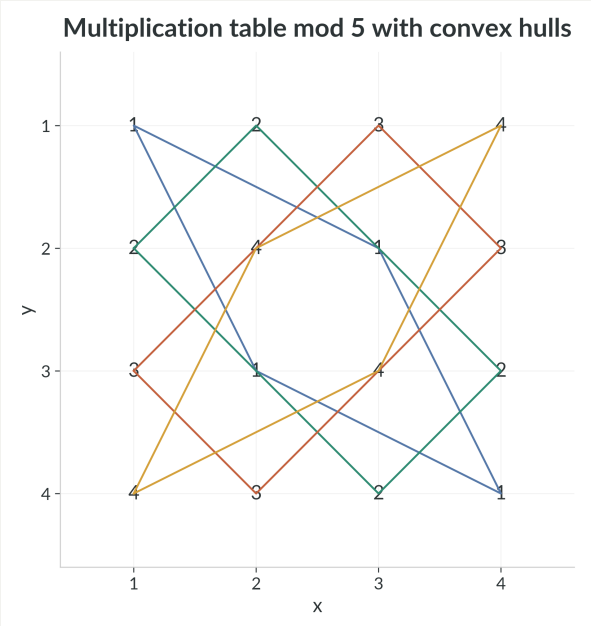


Figure 1: Multiplication table modulo $5$ with convex hulls of all residue classes. Each colored polygon encloses the lattice points sharing a common residue.

The central geometric quantity of this book is the area of the convex hull of $A_{N,a}$.

**Definition 2.1** (Main area functions). For $N \geq 2$ and $a \in \{0, 1, \ldots, N-1\}$, define

$$S(N, a) := \mathsf{Area}(\mathsf{conv}(A_{N,a})).$$

The total area sum is

$$S(N) := \sum_{a=0}^{N-1} S(N, a).$$

Each fixed set $A_{N,a}$ is a finite modular hyperbola inside the positive-integer lattice window. Sets of this form have already been studied under that name in the number-theory literature [1]. The distinctive emphasis of MTMN is different: one studies the whole family of these congruence classes at once, compares their convex hulls residue by residue, and treats the resulting area data as a geometric object in its own right.

There is also a second historical backdrop nearby: the classical Erdős multiplication table problem for ordinary integer products. In that setting one asks how many distinct integers appear in an $N \times N$ multiplication table, and later work such as Koukoulopoulos's generalized multiplication-table paper develops the asymptotic distinct-product side of that story [2]. The guiding contrast is not that MTMN answers the same question in a modular setting, but that it studies a different global invariant. The classical problem concerns distinct ordinary products and shows that this quantity grows more slowly than the naive quadratic scale. MTMN instead studies the summed residue-wise convex-hull area $S(N)$, whose first global theorem already has cubic order of growth. The thematic connection is that both problems ask how much global structure survives inside a multiplication table once multiplication creates many arithmetic coincidences.

That point of view quickly opens into a larger story. A single residue class already has visible Euclidean shape; the whole table can be built border by border from simple modular progressions; and the arithmetic split between coprime and non-coprime indices produces two genuinely different geometric regimes. The next chapters first place this picture beside earlier work on modular hyperbolas, then develop the basic geometry of the residue classes, and then use the zero class to show how divisor geometry and the continuous hy-

perbola $y = N/x$ meet in one exact area decomposition. After that, the later chapters turn those geometric pictures into exact language: nonzero coprime classes become permutation plots, support functions encode hulls direction by direction, outer rings become exact boundary-layer models, and residue-area packages begin to organize the data algebraically.

Between the discrete set $A_{N,a}$ and its convex hull there is also a natural straight-line intermediate object, treated in Chapter 5 under the name **Connecting Residue Points**. That same chapter introduces the layer-by-layer construction of the table, the first complete worked examples ($N = 5$ and $N = 6$), the divisor-controlled zero-class hull, and the hyperbolic correction term $\Delta_N$, so the reader sees concrete geometry before the later exact formulas and asymptotic questions.

Before turning to the wider program, it helps to look at one residue class by itself. In Figure 2 the set $A_{7,2}$ appears as a small finite cloud whose convex hull is already a visible polygon. This is the local picture that the whole subject keeps refining: first identify the arithmetic point set, then understand which of those points survive as geometric vertices.

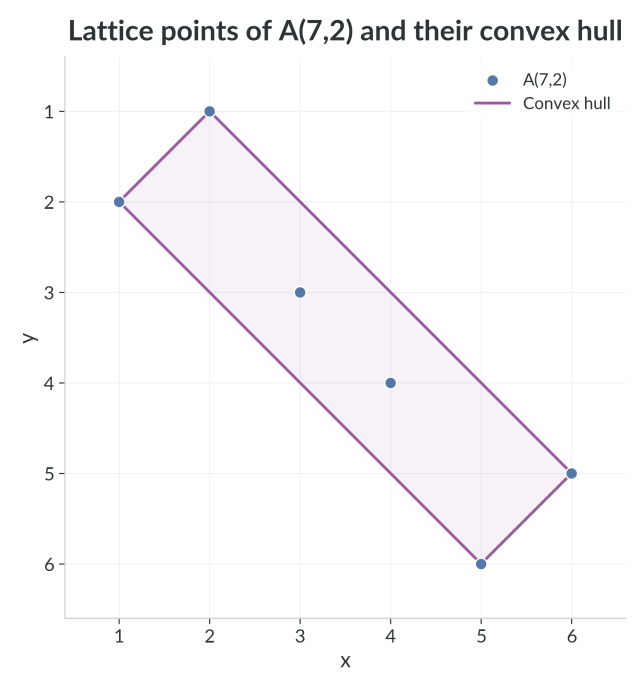


Figure 2: Lattice points of $A_{7,2}$ and their convex hull. The blue dots are the lattice points satisfying $xy \equiv 2 \pmod 7$, and the red boundary traces the convex hull.

Two structural contrasts are worth isolating from the start. First, prime and composite moduli already behave differently: primes give the clean permutation-plot case on the whole nonzero grid, while composites split into a unit-side geometry and a zero-divisor geometry. Second, the border-by-border construction of the table will later turn into exact lower models: the first and second boundary layers are not arbitrary truncations, but the first residue-wise pieces of the same geometric object. The zero residue class will become the first place where the composite side admits a complete geometric description: its hull can be read exactly from the divisors of $N$.

Figure 3 highlights the first of these contrasts.

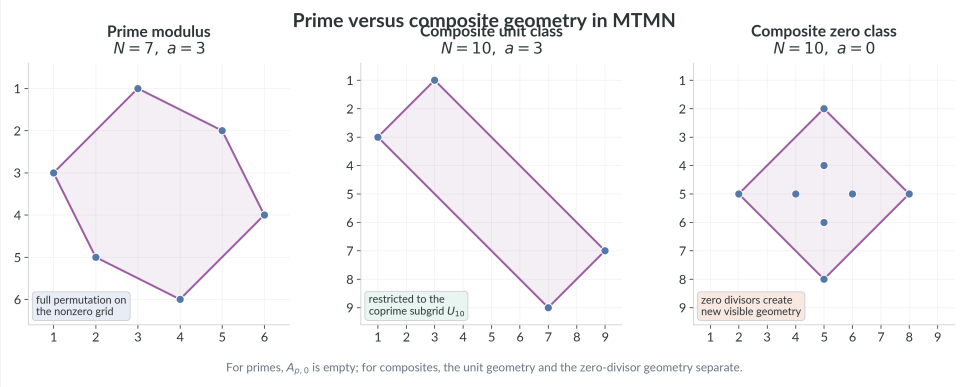


Figure 3: Prime and composite moduli already separate into different geometric regimes. Prime nonzero classes occupy the full nonzero grid, while composite moduli exhibit both a restricted coprime subgrid and a visible zero-divisor class.

These definitions lead immediately to a broader research program.

1. Which lattice points of $A_{N,a}$ become vertices of the convex hull, and how can that extremal condition be read arithmetically?
2. Can one compute $S(N, a)$ and $S(N)$ exactly, or at least by exact boundary-layer models that isolate the first geometric contributions residue by residue?
3. How does the factorization profile of $N$ control the geometry, especially on the composite side where zero-divisor classes produce divisor rectangles, envelopes, and other new shapes?
4. What is the right combinatorial or graph-theoretic language for these residue-class pictures, and which geometric features does it actually preserve?
5. What is the large-$N$ growth of $S(N)$? Does the total area really live on a cubic scale, is the sharper asymptotic $S(N) \sim N^3$ true, and what do the reciprocal and weighted series reveal about that growth?
6. What happens for deeper boundary layers, for unions of several layers, and for the interfaces where new hull vertices first appear?
7. How far does the construction extend beyond the planar multiplication table: to higher-dimensional multiplication cubes, to general dimension $d$, or even to modular level sets of functions more general than multiplication?

8. What should count as "dynamics" in MTMN: varying the modulus $N$, varying the residue $a$, rotating support directions, or growing the table layer by layer?

The present book answers part of this program. In particular:

- it defines the basic geometric language of the subject;
- it gives exact values for small examples;
- it proves that every $S(N,a)$ and $S(N)$ is a nonnegative integer;
- it gives a sharp prime/composite criterion for the degeneracy of the zero class;
- it describes the full zero-class hull exactly in terms of divisor rectangles, a divisor-envelope formula, and a hyperbolic correction identity;
- it records an exact support-function formula for $S(N,a)$;
- it develops an exact first-boundary model $S^{(1)}(N,a)$ and $S^{(1)}(N)$;
- it develops an exact second-boundary model for odd $N$;
- it proves that the total area $S(N)$ is trapped between explicit cubic lower and upper models, so $S(N)$ has cubic order of growth;
- it uses that cubic scale to frame the sharper asymptotic question $S(N) \sim N^3$ and the deficiency $N^3 - S(N)$;
- it packages the residue-area profile into a polynomial whose simplest evaluations isolate the zero class algebraically;
- it proves that the boundary model already implies convergence of

$$\sum_{N=4}^{\infty} \frac{1}{S(N)} \qquad \text{and} \qquad \sum_{N=4}^{\infty} \frac{N}{S(N)}$$

  by comparison;
- it gives a rigorous numerical enclosure for $\sum_{N=4}^{\infty} \frac{1}{S(N)}$ and rules out the historical guess $\sqrt{e} - 1$.

The exposition that follows moves through the subject in that same order: first background and literature context, then the geometry of residue classes and the zero class, then exact support-function formulas, and finally the boundary-layer models suggested by the layer-by-layer construction.

# 3 Background and notation

## 3.1 Conventions on integers and residues

Throughout the book, the modulus $N$ and the coordinates $x, y$ are positive integers. Residues are represented by the set

$$\{0, 1, \ldots, N-1\},$$

so the residue label $a$ may be zero even though the lattice coordinates are positive.

Later, when we compare modular classes with ordinary hyperbolas, we will do so inside the same positive window by tracking the exact products $a + kN$ that occur there.

## 3.2 Modular arithmetic

When we write

$$xy \equiv a \pmod{N},$$

we mean that $N$ divides $xy - a$. Equivalently, $xy$ and $a$ leave the same remainder upon division by $N$. The set of residues modulo $N$ is

$$\mathbb{Z}/N\mathbb{Z} = \{0, 1, \ldots, N-1\}$$

with arithmetic performed modulo $N$. This congruence language was organized systematically by Gauss in *Disquisitiones Arithmeticae* [3].

We will say that a number $u$ is *coprime to* $N$ if $\gcd(u, N) = 1$, meaning that $u$ and $N$ share no common factor greater than $1$. Equivalently, $u$ is *invertible modulo* $N$: there exists some residue $v$ with $uv \equiv 1 \pmod{N}$.

## 3.3 Lattice points and convex hulls

A *lattice point* in the plane is simply a point of $\mathbb{Z}^2$. The finite square

$$\{1, \ldots, N-1\}^2$$

may be viewed as a Euclidean lattice window. For any finite set $E \subset \mathbb{R}^2$, the *convex hull* $\mathrm{conv}(E)$ is the smallest convex set containing $E$. Geometrically, if one places nails at the points of $E$ and stretches a rubber band around them, the rubber band traces the boundary of $\mathrm{conv}(E)$.

### 3.4 Area

For a polygon with ordered vertices $(x_1, y_1), \ldots, (x_m, y_m)$ listed cyclically (with the convention that $(x_{m+1}, y_{m+1}) = (x_1, y_1)$), the *shoelace formula* states

$$\mathrm{Area}(P) = \frac{1}{2}\left|\sum_{j=1}^{m}(x_j y_{j+1} - x_{j+1} y_j)\right|.$$

This will be our main exact tool once the hull vertices are known. For an MTMN example, the residue class

$$A_{5,1} = \{(1,1), (2,3), (3,2), (4,4)\}$$

has convex-hull vertices in cyclic order. The companion picture marks those four points and the cyclic hull order used in the computation:

$$(1,1),\ (3,2),\ (4,4),\ (2,3).$$

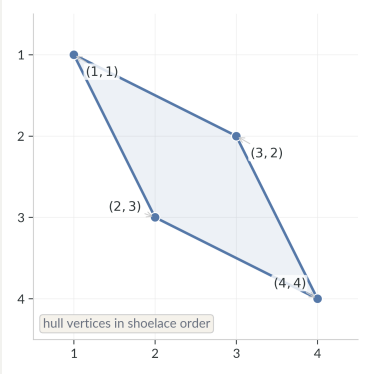


Figure 4: The residue class $A_{5,1}$ and the convex hull used in the shoelace computation.

Therefore

$$\begin{aligned}&\mathrm{Area}(\mathrm{conv}(A_{5,1}))\\ &= \tfrac{1}{2}\left|(1{\cdot}2 + 3{\cdot}4 + 4{\cdot}3 + 2{\cdot}1) - (1{\cdot}3 + 2{\cdot}4 + 4{\cdot}2 + 3{\cdot}1)\right|\\ &= \tfrac{1}{2}|28 - 22| = 3.\end{aligned}$$

So $S(5,1) = 3$.

A second standard fact, useful when the polygon is already explicit as a lattice polygon, is

*Pick's theorem* [4]:

$$\text{Area}(P) = I(P) + \frac{B(P)}{2} - 1,$$

where $I(P)$ is the number of interior lattice points and $B(P)$ is the number of lattice points on the boundary of $P$. Pick's theorem computes the area of any simple lattice polygon – convex or not – purely by counting lattice points, with no need for coordinate decomposition or trigonometry. We will use it mainly as a structural tool: it relates area to lattice counts, but it does not by itself tell us which lattice points form the hull.

Figure 5 shows two examples. On the left, a triangle with vertices $(0, 0)$, $(4, 0)$, $(1, 3)$ has $B = 8$ boundary points and $I = 3$ interior points, giving Area $= 3+8/2-1 = 6$. On the right, an L-shaped non-convex polygon has $B = 12$ and $I = 0$, giving Area $= 0 + 12/2 - 1 = 5$. The second example illustrates that Pick's theorem applies even when the shape is not convex.

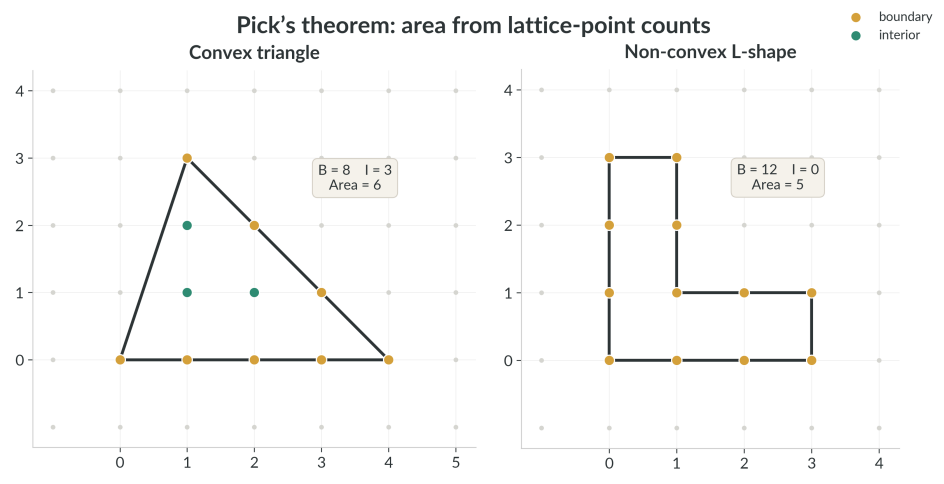


Figure 5: Two lattice polygons counted by Pick's theorem. Left: a convex triangle. Right: a non-convex L-shape. In both cases, area follows from counting boundary (orange) and interior (green) lattice points.

Historically, the formula goes back to Pick's 1899 paper [4].

### 3.5 Indicator notation and exponential filters

For a statement $C$, the indicator function $\mathbf{1}_C$ equals $1$ when $C$ is true and $0$ otherwise. Thus

$$\mathbf{1}_{xy\equiv a \pmod N} = \begin{cases} 1, & xy \equiv a \pmod N, \\ 0, & \text{otherwise.} \end{cases}$$

A standard Fourier identity on the finite cyclic group $\mathbb{Z}/N\mathbb{Z}$ is

$$\mathbf{1}_{xy\equiv a \pmod N} = \frac{1}{N} \sum_{r=0}^{N-1} \exp\left( \frac{2\pi i r(xy-a)}{N} \right).$$

This identity is exact: if $xy - a \equiv 0 \pmod N$, every exponential equals $1$ and the average is $1$; otherwise the $N$ roots of unity cancel to $0$.

For example, when $N = 4$ the four 4th roots of unity are $1,\ i,\ -1,\ -i$. Their sum is

$$1 + i + (-1) + (-i) = 0,$$

because opposite terms cancel in pairs. The same cancellation occurs for any $N$: the $N$-th roots of unity are equally spaced around the unit circle, and unless every term equals $1$ (which happens exactly when $xy \equiv a$), their symmetric arrangement forces the sum to vanish.

# 4 Modular hyperbolas in context

This book studies the lattice sets

$$A_{N,a} = \{(x, y) \in \{1, \dots, N-1\}^2 : xy \equiv a \pmod{N}\}$$

and the convex hulls they determine. For fixed modulus and fixed residue, such a set is already a familiar object in the number-theory literature: it is usually called a *modular hyperbola*.

We will use that phrase when it helps connect MTMN to earlier work. At the same time, we will keep the name **MTMN** for the project-wide point of view of this book, because the present subject is not only one congruence class at a time. It is the geometry of the whole residue family

$$\{A_{N,0}, A_{N,1}, \dots, A_{N,N-1}\}$$

inside one multiplication table.

This chapter is only a guide to mathematical placement. It is not meant to be a full survey.

## 4.1 What earlier work already studies

There is already a substantial literature on modular hyperbolas. A convenient entry point is Igor E. Shparlinski's survey *Modular hyperbolas* [1], which collects many of the main themes and references. At a broad level, that literature studies how the points satisfying

$$xy \equiv a \pmod{m}$$

are distributed, how they cluster inside boxes, how they behave geometrically, and how their behavior changes in higher-dimensional analogues. In that literature the modulus is often written as $m$ (and sometimes later as $n$); throughout this book it is the same modulus that we denote by $N$.

Some of those questions are close to the concerns of this book. Convex-hull questions, for example, are not new here. Konyagin and Shparlinski [5] study the number of vertices of convex hulls of points on modular hyperbolas. Ford, Khan, and Shparlinski [6] study

geometric properties of the unit modular hyperbola

$$xy \equiv 1 \pmod{n}$$

and obtain bounds connected with extremal points and convex-hull behavior. Other work studies concentration of points in small regions [7], distances or coordinate differences [8], and multidimensional modular hyperbolas [9].

Read through the lens of Shparlinski's survey, the standard output of this literature is usually asymptotic rather than exact [1]. One fixes a single congruence class, most often in the coprime regime $\gcd(a, m) = 1$, and then asks for point counts in boxes, discrepancy estimates, concentration bounds, or upper and lower bounds for geometric statistics such as vertex counts, coordinate differences, or distances. The characteristic tools are exponential sums, especially Kloosterman sums, together with discrepancy methods and other analytic estimates. Even when the language is geometric, the aim is typically to understand the large-scale distribution of one modular hyperbola rather than to write exact formulas for a whole residue family.

There is also nearby literature on primitive or visible points, meaning lattice points with coprime coordinates and hence visible from the origin, for congruence-defined curves over finite fields [10]. Those papers do not always use exactly the same MTMN setting, but they show that coprimality already carries its own arithmetic-geometric language in the surrounding literature.

## 4.2 The usual geometric realization

Most of this literature fixes a single congruence class and studies it in the usual positive residue window, typically inside a box such as

$$\{0, 1, \ldots, m-1\}^2$$

or

$$\{1, \ldots, m-1\}^2.$$

That is a natural first realization, and it is exactly the positive lattice window used in the main chapters of this book.

Another common feature is the emphasis on the coprime, or invertible, case

$$\gcd(a, m) = 1.$$

In that regime both coordinates are units modulo $m$, and a large body of analytic and combinatorial machinery becomes available. This case is mathematically important, and it remains part of the background of MTMN as well.

### 4.3 How the present book differs

The distinctive emphasis of this book lies elsewhere. Rather than isolating one congruence class and one geometric statistic, it keeps the full residue family in view and compares classes residue by residue. Its main numerical quantities are the convex-hull areas

$$S(N, a) = \mathrm{Area}(\mathrm{conv}(A_{N,a}))$$

and

$$S(N) = \sum_{a=0}^{N-1} S(N, a).$$

This already shifts the question from the geometry of one modular hyperbola to the geometry of the whole multiplication table.

The book also gives special attention to the zero class and to non-coprime behavior. On that side, factorization and zero divisors create geometry that is invisible if one looks only at the unit case. The divisor-rectangle description of the zero class is the clearest example: there the hull can be read directly from the divisors of $N$, so the arithmetic of compositeness becomes part of the visible Euclidean structure.

The difference is therefore mainly one of emphasis. Much of the earlier literature studies one usually coprime residue class at a time and extracts asymptotic information from exponential-sum estimates. MTMN focuses on exact discrete geometric formulas in the fixed positive lattice window, treats all residues simultaneously, and includes the zero-divisor side from the start. In particular, exact area formulas, exact boundary-layer models,

and the divisor-controlled hull of $A_{N,0}$ are not the usual targets of the analytic modular-hyperbola program even though they begin from the same congruence-defined point sets.

Later chapters do not replace that window by a competing one. Instead they use the same positive embedding to study divisor envelopes, support functions, and boundary layers.

So the purpose of this book is not to reintroduce a known object under a new name. Its purpose is to develop a residue-family, area-centered, and composite-sensitive program around an object that earlier literature has already studied from other angles.

# 5 The basic geometry of MTMN

## 5.1 Definition of the residue classes

For each residue $a$, the set $A_{N,a}$ is a finite subset of the square lattice. The collection

$$\{A_{N,0},\ A_{N,1},\ \ldots,\ A_{N,N-1}\}$$

partitions the entire lattice square $\{1,\ldots,N-1\}^2$ according to the residue of the product $xy$ modulo $N$.

> *Remark* 5.1 (Two different roles often confused in notation). Throughout this document, the residue class is denoted by $a$, not by $i$. The symbol $i$ is reserved for the imaginary unit $i^2 = -1$. This keeps the congruence variable and the complex unit unambiguous.

Figure 6 shows the full multiplication table modulo 7: each cell displays the residue $xy$ mod 7, and the convex hull of every residue class is overlaid.

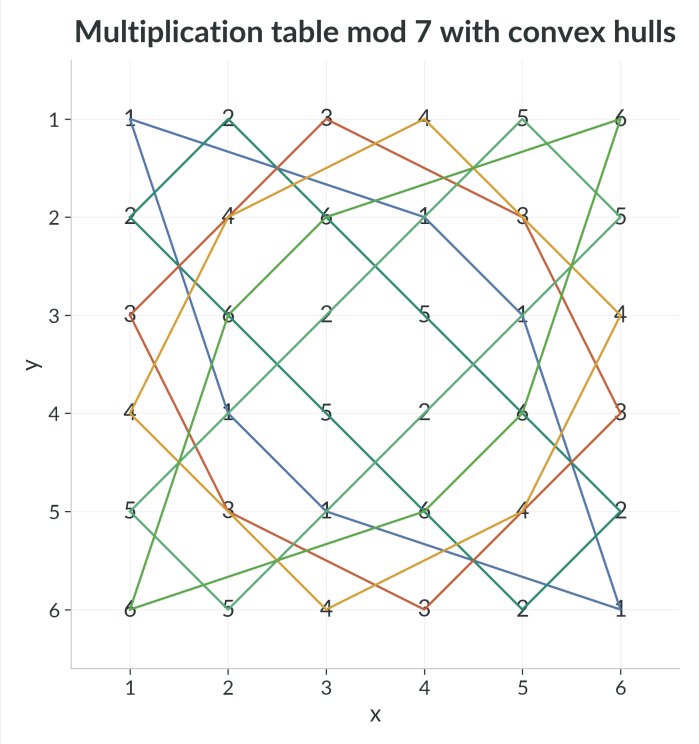


Figure 6: The multiplication table modulo 7 with convex hulls of all residue classes overlaid. Each colored polygon encloses the lattice points sharing a common residue.

## 5.2 Connecting Residue Points

Before one passes to convex hulls, there is already a natural straight-line object attached to a residue class. The informal picture is what one might casually call "connect the dots," but here we will use the more precise phrase **Connecting Residue Points**.

For a fixed residue class $A_{N,a}$, join every pair of points in $A_{N,a}$ by a straight segment and take the union of all those segments:

$$L_{N,a} := \bigcup_{p,q \in A_{N,a}} [p, q].$$

Thus

$$A_{N,a} \subseteq L_{N,a} \subseteq \mathrm{conv}(A_{N,a}).$$

Read from left to right, this is a chain from the smallest object to the largest one in the set-theoretic sense of containment: the discrete residue set $A_{N,a}$, then the union of all connecting segments $L_{N,a}$, and finally the full convex hull. Here the symbol $\subseteq$ allows equality, just as $\leq$ allows equality for numbers. The word "larger" is therefore about inclusion, not about visible richness of structure.

In fact, the middle object can carry more geometric detail than the hull, because the hull fills everything inside its outer boundary. Here "outer envelope" means exactly that outer wrapper. For instance, if three points are non-collinear, meaning that they do not all lie on one straight line, then Connecting Residue Points gives only the three edges of the triangle, while the convex hull also contains the interior.

This book will still focus mainly on convex hulls and their areas, but naming this intermediate layer now makes its purpose explicit: it records the finer segment structure that is discarded when one passes from the residue set to its hull.

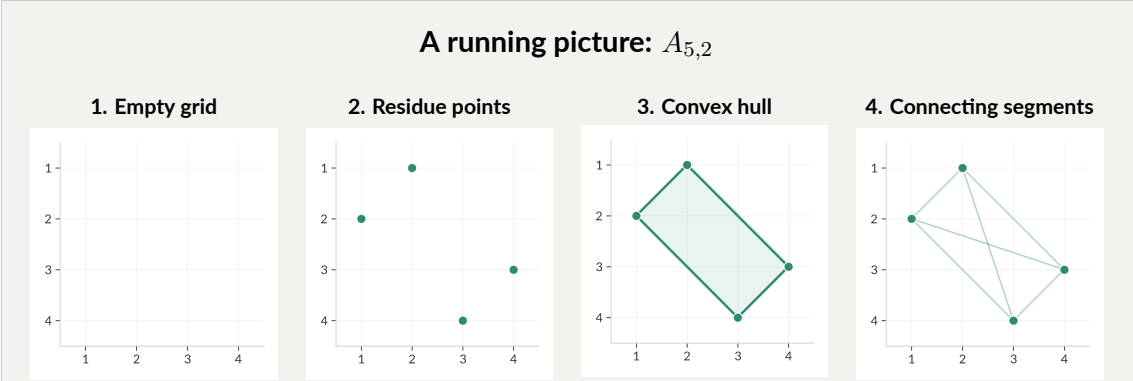


Figure 7: A concrete progression for $A_{5,2}$: the empty lattice window, the residue points, the convex hull, and the full Connecting Residue Points object.

## 5.3 How to draw MTMN layer by layer

One does not need to compute every product $xy$ separately in order to draw the table. Fix a row index $k$. The $k$-th row is

$$k,\ 2k,\ 3k,\ \ldots,\ (N-1)k \pmod{N},$$

so after the first entry every new value is obtained by adding $k$ modulo $N$.

**Proposition 5.1** (A row determines its whole boundary layer). *For $1 \le k \le \lfloor (N-1)/2 \rfloor$, the row $N-k$ is the reverse of the row $k$. More precisely,*

$$(N-k)y \equiv -ky \equiv k(N-y) \pmod{N}.$$

*By symmetry the same relation holds for the columns. Consequently, once the $k$-th row is known, the entire $k$-th boundary layer of the table is determined.*

*Proof.* Fix a column index $y$. The entry in row $N-k$ and column $y$ is

$$(N-k)y \equiv k(N-y) \pmod{N}.$$

The right-hand side is exactly the entry in row $k$ and column $N-y$. Therefore the value appearing in row $N-k$ at position $y$ is the value appearing in row $k$ at the reflected position $N-y$. As $y$ runs through $1, \dots, N-1$, the number $N-y$ runs through $N-1, \dots, 1$, so row $N-k$ is row $k$ read in reverse order.

Since the multiplication table is symmetric across the diagonal ($xy = yx$), the same reasoning applies to columns. The $k$-th boundary layer consists of the top row $k$, the bottom row $N-k$, the left column $k$, and the right column $N-k$. Hence one modular progression and its reverse determine the whole boundary ring.

For example, when $N = 7$ and $k = 2$, the row-$2$ progression is

$$2,\ 4,\ 6,\ 1,\ 3,\ 5,$$

while the row-$5$ progression is

$$5,\ 3,\ 1,\ 6,\ 4,\ 2,$$

which is its reverse. The second boundary layer is therefore already encoded in that one row. ∎

The practical rule is therefore simple. First write the row-$k$ progression across the top side of the $k$-th layer. Then continue around the same ring: the right and left sides use the

reversed order, while the bottom side uses the forward order when read from right to left. In this way a single repeated-addition sequence fills an entire border.

For $N = 7$, the first row gives

$$1,\ 2,\ 3,\ 4,\ 5,\ 6,$$

so the first boundary is immediate. The second row gives

$$2,\ 4,\ 6,\ 1,\ 3,\ 5,$$

which fills the second boundary by the same rule. The third row gives

$$3,\ 6,\ 2,\ 5,\ 1,\ 4,$$

which fills the third boundary and therefore completes the table. Because the row indices $1$, $2$, and $3$ are all coprime to $7$, each of those rows is a rearrangement of the nonzero residues. The prime picture is therefore built by successive permutations.

Figure 8 shows that prime case from start to finish: first boundary, second boundary, third boundary, then the completed table. By the time the full picture is on the page, one can already see the first informal glimpse of permutation geometry: every nonzero row and every nonzero column contains each nonzero residue exactly once.

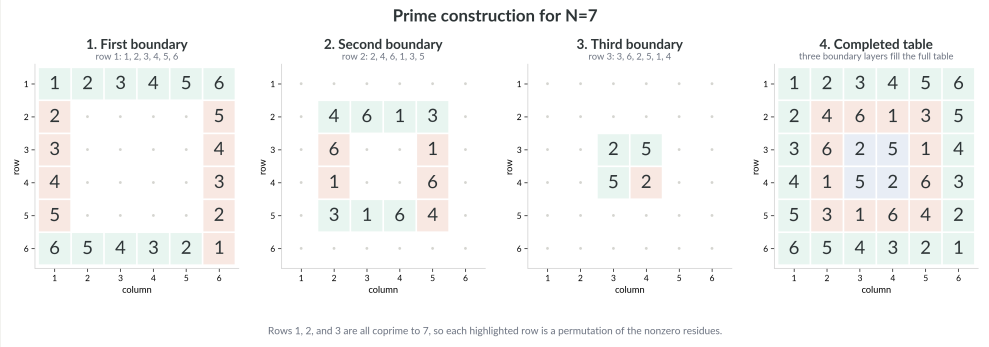


Figure 8: Building MTMN for $N = 7$ one boundary layer at a time. The four panels show the first boundary, the second boundary, the third boundary, and the completed table.

Only after the $N = 7$ story is complete is it helpful to contrast it with a composite modulus. For $N = 6$, the first row still gives a full nonzero progression, but the second row is

$$2,\ 4,\ 0,\ 2,\ 4,$$

so repetition and a zero appear immediately on the second boundary. Here the reason is arithmetic: $2$ and $6$ share the common factor $2$, so $2 \cdot 3 \equiv 0 \pmod 6$. Once a row or column contains $0$, it can no longer list the nonzero residues exactly once.

Figure 9 shows the same construction for $N = 6$: first boundary, second boundary, the central entry, and the completed table. The procedure is unchanged, but the arithmetic is different.

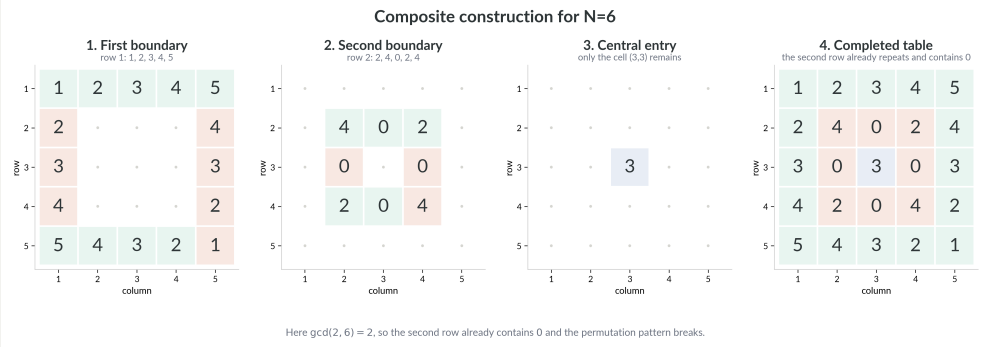


Figure 9: Building MTMN for $N = 6$ from start to finish. The four panels show the first boundary, the second boundary where repetition appears, the central entry, and the completed table.

Figure 10 isolates the directional rule around one boundary ring.

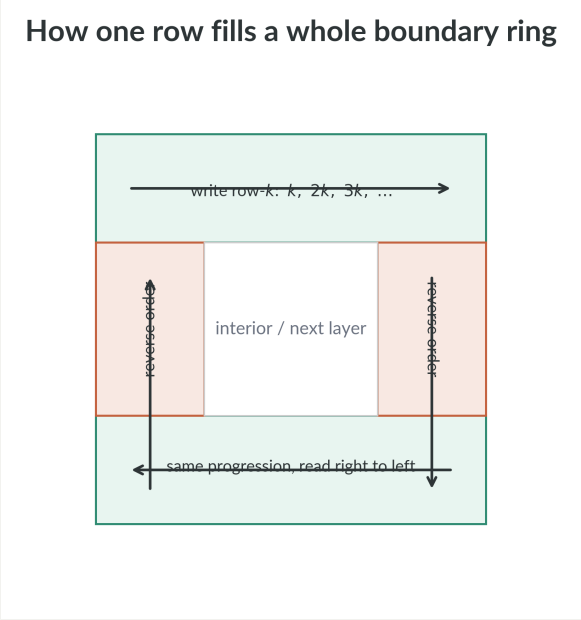


Figure 10: Directional rule for filling a boundary layer. Write the row-$k$ progression across the top, use the reverse order on the two side edges, and read the bottom edge from right to left so that the same forward progression wraps consistently around the ring.

This contrast isolates the key arithmetic rule. A number $k$ is *coprime to* $N$ if $\gcd(k, N) = 1$, meaning that $k$ and $N$ share no common factor greater than $1$. Equivalently, $k$ is *invertible modulo* $N$: there exists some $k^{-1}$ with $kk^{-1} \equiv 1 \pmod{N}$. When $k$ is coprime to $N$, multiplication by $k$ rearranges the nonzero residues. When $k$ is not coprime to $N$, a zero appears and the permutation pattern breaks.

This border-by-border picture is more than a drawing trick. For a fixed residue $a$, one may keep only those cells on the first ring whose value is $a$; that residue-wise subset will become the first boundary model $B^{(1)}_{N,a}$. Keeping the next ring gives the second boundary model $B^{(2)}_{N,a}$. In this sense the boundary program of Chapters 7 and 8 is already visible in the construction itself: the early layers of the table are the first explicit geometric approximations to the full residue class.

## 5.4 Two complete first examples

Before turning to the general propositions, it is worth seeing the first prime and composite cases in full. These two moduli are small enough that one can hold the whole table, all residue classes, and the area data in view at once. They will serve as anchors for the later abstract statements.

### 5.4.1 The Case $N = 5$

Since $5$ is prime, every nonzero index is coprime to $5$, so this is the cleanest first example.

Figure 11 shows the raw multiplication table modulo $5$.

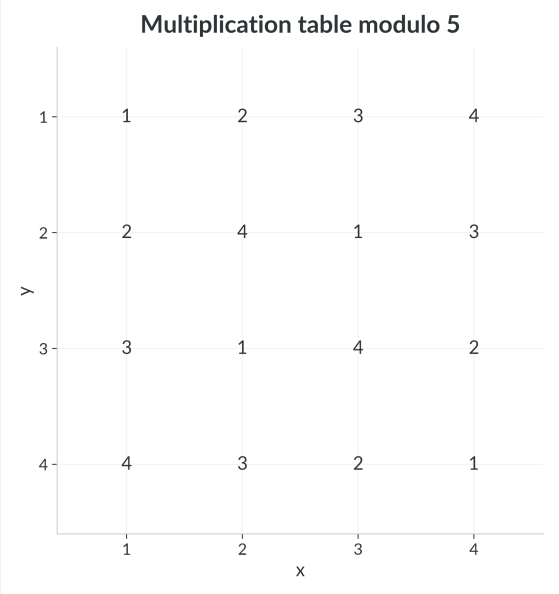


Figure 11: Raw multiplication table modulo $5$, showing only the residue values.

Figure 12 separates the residue classes. Each panel is labeled by $a$ and by the corresponding area $S(5, a)$.


Figure 12: Residue classes $A_{5,a}$ and their convex hulls. Each panel title records $a$ and $S(5, a)$.

Figure 13 overlays the nondegenerate convex hulls on the original table, so the geometry can be compared directly with the residue values.

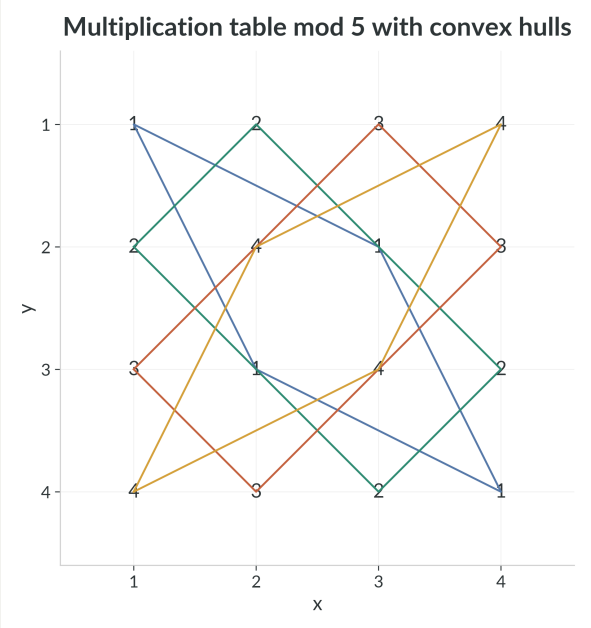


Figure 13: Multiplication table modulo $5$ with convex hulls of all nondegenerate residue classes overlaid.

| $a$ | 0 | 1 | 2 | 3 | 4 | total |
|---|---|---|---|---|---|---|
| $S(5,a)$ | 0 | 3 | 4 | 4 | 3 | $S(5)=14$ |

### 5.4.2 The Case $N = 6$

For $N = 6$, zero divisors already appear, so the composite case changes both the arithmetic and the geometry.

Figure 14 shows the raw multiplication table modulo $6$.

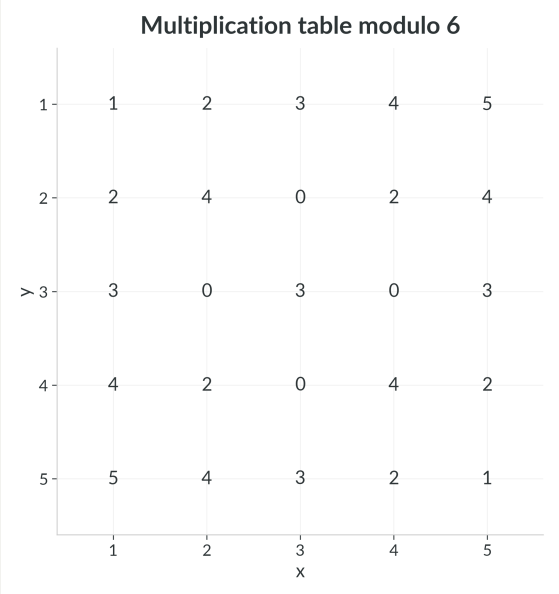


Figure 14: Raw multiplication table modulo $6$, showing only the residue values.

Figure 15 separates the individual residue classes. Here the zero-divisor behavior is already visible in the shapes and in the degenerate cases $S(6,1) = S(6,5) = 0$.

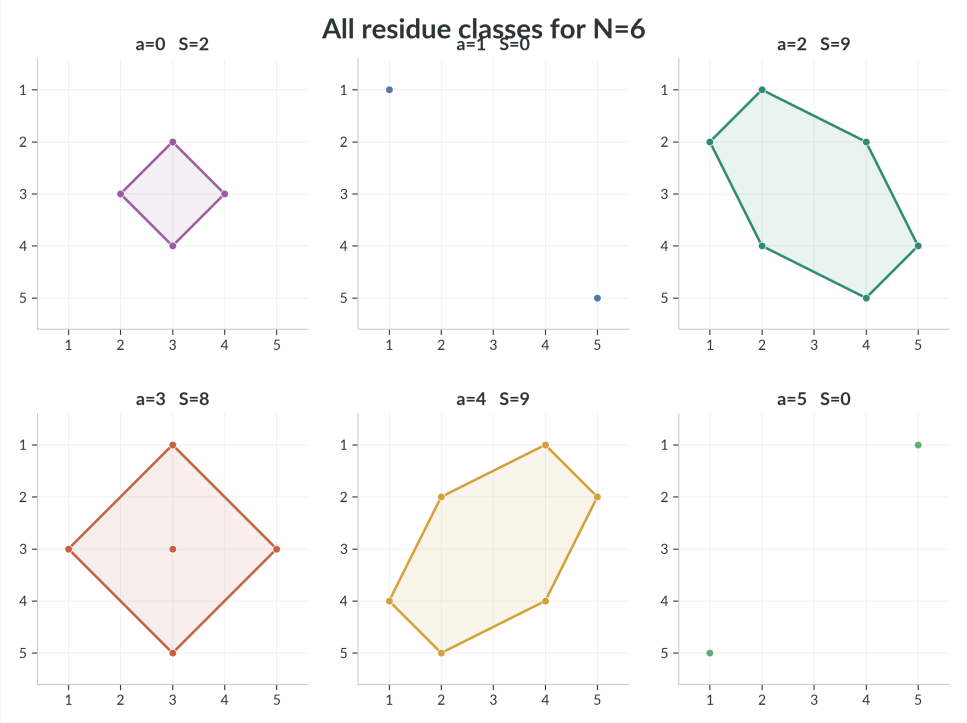


Figure 15: Residue classes $A_{6,a}$ and their convex hulls. Each panel title records $a$ and $S(6, a)$.

Figure 16 overlays all nondegenerate convex hulls on the original numbers.

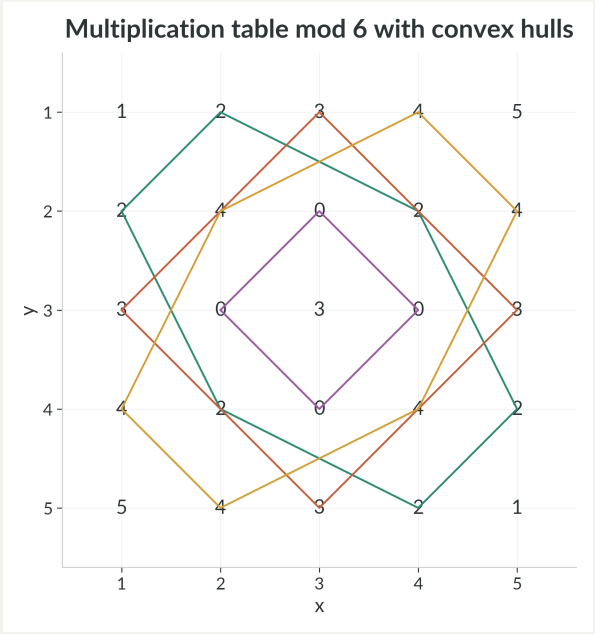


Figure 16: Multiplication table modulo $6$ with convex hulls of all nondegenerate residue classes overlaid.

| $a$ | 0 | 1 | 2 | 3 | 4 | 5 | total |
|---|---|---|---|---|---|---|---|
| $S(6,a)$ | 2 | 0 | 9 | 8 | 9 | 0 | $S(6)=28$ |

## 5.5 A table of initial values

For reference, the exact totals for $4 \le N \le 10$ are:

| $N$ | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|
| $S(N)$ | 2 | 14 | 28 | 70 | 108 | 205 | 334 |

A more detailed table appears in the appendix.

## 5.6 Euclidean symmetries of the residue classes

**Proposition 5.2** (Transpose and central symmetries). *For every $N \geq 2$ and residue $a$, the set $A_{N,a}$ is invariant under the maps*

$$(x, y) \mapsto (y, x), \qquad (x, y) \mapsto (N - x, N - y).$$

*Consequently* conv$(A_{N,a})$ *is symmetric across the diagonal line $x = y$ and under the half-turn (rotation by $180°$) about the center point $(N/2, N/2)$.*

*Proof.* If $(x, y) \in A_{N,a}$, then $xy \equiv a \pmod{N}$. Since multiplication is commutative, $yx \equiv a \pmod{N}$, so $(y, x) \in A_{N,a}$. Also

$$(N - x)(N - y) \equiv xy \pmod{N},$$

so $(N - x, N - y) \in A_{N,a}$ as well.

Both maps are affine maps, meaning linear transformations possibly followed by a translation. More specifically, the transpose is linear, and $(x, y) \mapsto (N - x, N - y)$ is the half-turn about the center $(N/2, N/2)$. We use the elementary fact that affine maps carry convex hulls to convex hulls, so these symmetries of the finite set $A_{N,a}$ become the corresponding symmetries of conv$(A_{N,a})$. ■

**Corollary 5.3** (Complementary-residue reflection). *For every $N \geq 2$ and residue $a$,*

$$S(N, a) = S(N, (-a) \bmod N).$$

*Proof.* The reflection

$$(x, y) \mapsto (x, N - y)$$

preserves Euclidean area and satisfies

$$xy \equiv a \pmod{N} \implies x(N - y) \equiv -xy \equiv -a \pmod{N}.$$

Hence it carries $A_{N,a}$ onto $A_{N,(-a) \bmod N}$, so the two convex hulls have equal area. For a

concrete example, when $N = 7$ and $a = 2$, the points

$$(1, 2),\ (2, 1),\ (5, 6),\ (6, 5) \in A_{7,2}$$

are reflected to

$$(1, 5),\ (2, 6),\ (5, 1),\ (6, 2) \in A_{7,5}.$$

The whole residue class transforms in exactly that way, so the two hulls have equal area. ■

**Corollary 5.4.** *For every $N \geq 2$ and residue $a$, the quantity $S(N, a)$ is a nonnegative integer. Consequently $S(N)$ is a nonnegative integer for every $N \geq 2$.*

*Proof.* If $\mathrm{conv}(A_{N,a})$ is degenerate (a single point or a line segment, so it has no area), then $S(N, a) = 0$. Otherwise, the half-turn from Proposition~5.2 pairs every boundary lattice point with a distinct partner, so the boundary count $B$ is even. Pick's theorem then gives

$$S(N, a) = I + \frac{B}{2} - 1,$$

which is an integer since $I$ is an integer and $B$ is even. Area is nonnegative, so $S(N, a) \in \mathbb{Z}_{\geq 0}$. Summing over $a$ yields the same for $S(N)$. ■

## 5.7 Coprimality and permutation geometry

A permutation is first a one-dimensional object: a rearrangement of a list with no repetitions. If one then plots the input $x$ against the permuted output $\pi(x)$, that same rearrangement becomes a two-dimensional scatter plot with exactly one dot in every row and every column. This is the sense in which we will use the phrase *permutation plot*. The higher-dimensional research program will ask for analogues of this picture in larger multiplication cubes, but the planar case already captures the essential idea.

**Proposition 5.5** (Rows and columns permute exactly at coprime indices)**.** *Let $1 \leq k \leq N - 1$. The $k$-th row and the $k$-th column are permutations of the nonzero residues $1, \ldots, N - 1$ if and only if $\gcd(k, N) = 1$.*

*Proof.* Suppose first that $\gcd(k, N) = 1$. Then there exists a residue $k^{-1}$ such that

$$kk^{-1} \equiv 1 \pmod{N}.$$

If two entries in the $k$-th row coincide, say

$$ky_1 \equiv ky_2 \pmod{N},$$

then multiplying by $k^{-1}$ gives $y_1 \equiv y_2 \pmod{N}$. Since $y_1, y_2 \in \{1, \ldots, N-1\}$, this forces $y_1 = y_2$. Therefore the entries

$$k,\ 2k,\ \ldots,\ (N-1)k \pmod{N}$$

are all distinct.

Next, none of those entries is $0$. Indeed, if there existed $y \in \{1, \ldots, N-1\}$ with $ky \equiv 0 \pmod{N}$, then multiplying by $k^{-1}$ would give $y \equiv 0 \pmod{N}$, impossible. So the row contains $N-1$ distinct nonzero residues, hence exactly the set $\{1, \ldots, N-1\}$.

Conversely, if $\gcd(k, N) = d > 1$, then $N/d$ is an integer with $1 \leq N/d \leq N-1$, and

$$k \cdot \frac{N}{d} \equiv 0 \pmod{N},$$

so the row contains $0$ and therefore cannot be a permutation of the nonzero residues. The same argument applies to the column. ■

**Proposition 5.6** (Coprime residue classes form permutation plots on the coprime subgrid)**.** *Assume* $\gcd(a, N) = 1$*. Then each row indexed by a number coprime to* $N$ *contains exactly one point of* $A_{N,a}$*, each column indexed by a number coprime to* $N$ *contains exactly one point of* $A_{N,a}$*, and no point of* $A_{N,a}$ *lies in a non-coprime row or column. Hence* $A_{N,a}$ *is a permutation plot on the coprime subgrid.*

*Proof.* Suppose $(x, y) \in A_{N,a}$, so $xy \equiv a \pmod{N}$. If $x$ shared a common factor $d > 1$ with $N$, then $d$ would divide both $xy$ and $N$, hence also $a$, contradicting $\gcd(a, N) = 1$. Thus every point of $A_{N,a}$ lies in a row indexed by a number coprime to $N$. The same argument

shows that every point also lies in a column indexed by a number coprime to $N$.

Now fix a row index $x$ with $\gcd(x, N) = 1$. By Proposition~5.5, that row is a permutation of the nonzero residues $1, \ldots, N-1$. In particular, the residue $a$ appears exactly once in the row, so there is exactly one point of $A_{N,a}$ in it. The same argument applies to columns. Therefore, on the coprime subgrid, the residue class has exactly one point in each row and each column, which is exactly the permutation-plot property. ■

The prime case is now the cleanest special case. If $N = p$ is prime, then every nonzero index is coprime to $p$. Therefore every nonzero residue class $A_{p,a}$ is a permutation plot on the whole nonzero grid. This does not solve the convex-hull problem by itself, but it shows that prime nonzero classes already come with a rigid row-and-column structure.

Figure 17 shows the prime example in the two ways that matter here. On the left, one coprime row of the multiplication table is read as a permutation of the nonzero residues. On the right, the residue class $A_{11,3}$ is drawn on the lattice, where the geometric content is that there is exactly one point in each nonzero row and each nonzero column.

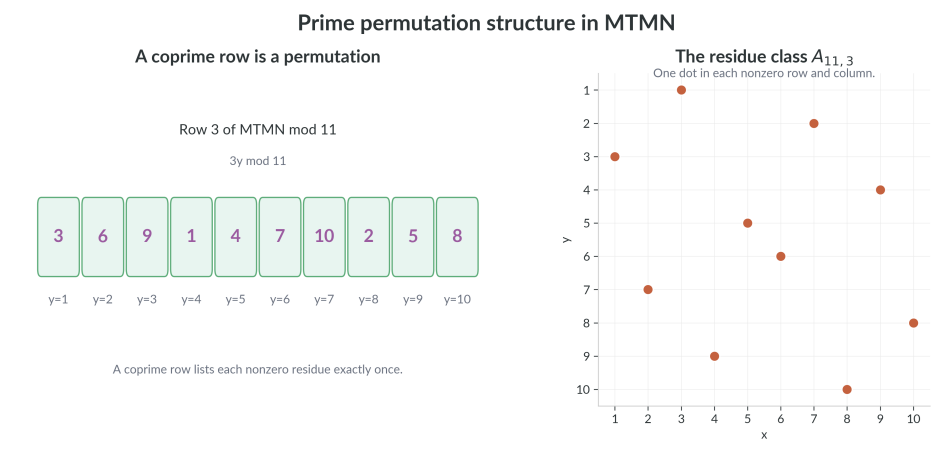


Figure 17: Prime permutation structure in two views. Left: a coprime row of the table is a permutation of the nonzero residues. Right: the residue class $A_{11,3}$ places exactly one point in each nonzero row and each nonzero column.

This rigid picture also has a modern discrepancy-theoretic counterpart. Blomer, Risager,

and Shparlinski study normalized inverse pairs and prove upper and lower bounds for their box, ball, and isotropic discrepancy, explicitly noting visible small-scale cellular structure and deviations from random behavior [11]. Their setting is different from MTMN: they vary the modulus, aggregate only inverse pairs, and work on the unit side. Our setting fixes one modulus $N$, studies one residue class inside a finite lattice window, and then compares the whole residue family including zero divisors. But the geometric moral is the same. Coprime classes are not random point clouds; arithmetic imposes a rigid permutation geometry that remains visible even before one asks for exact hull formulas.

## 5.8 The zero class and zero-divisor geometry

### 5.8.1 The zero class and compositeness

**Proposition 5.7** (The zero class detects compositeness). *If $p$ is prime, then $A_{p,0}$ is empty. If $N$ is composite, then $A_{N,0}$ is nonempty.*

*Proof.* If $p$ is prime and $1 \le x, y \le p-1$, then neither $x$ nor $y$ is divisible by $p$, so $xy \not\equiv 0 \pmod{p}$. Hence $A_{p,0} = \varnothing$.

If $N$ is composite, then there exist integers $d, e$ with $1 < d < N$, $1 < e < N$, and $N = de$. Then

$$de = N \equiv 0 \pmod{N},$$

so $(d, e) \in A_{N,0}$. ∎

The class $A_{N,0}$ is therefore the residue class where the same obstruction becomes most concentrated: zero divisors can multiply to $0 \pmod{N}$. This is one of the clearest differences between prime and composite moduli.

### 5.8.2 The exceptional modulus $N = 4$

The point-set statement does not yet determine the geometry of the hull once the zero class becomes nonempty. The smallest composite modulus already shows the subtlety. For $N = 4$ one has

$$A_{4,0} = \{(2, 2)\},$$

so the zero class is present but still degenerate:

$$S(4,0) = 0.$$

Thus nonemptiness alone does not force positive area.

Figure 18 marks this exceptional case directly on the $N = 4$ residue grid. The lone point $(2,2)$ is present, but a single point still has zero area.

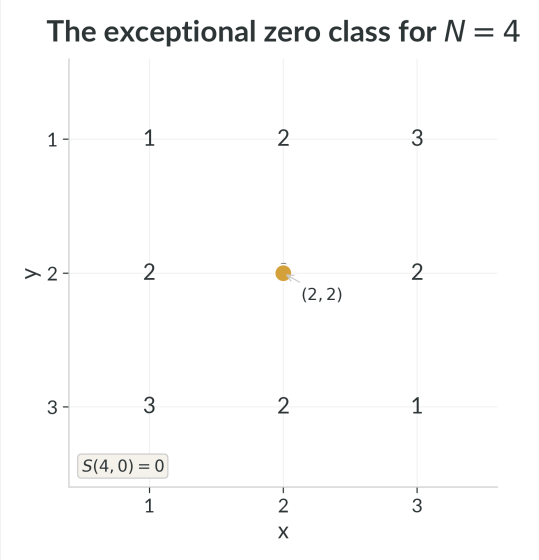


Figure 18: The exceptional zero class for $N = 4$. The residue grid is shown with the unique point $(2,2) \in A_{4,0}$ highlighted.

### 5.8.3 When does the zero class have positive area?

**Theorem 5.8** (Zero-class degeneracy criterion)**.** *For every $N > 4$,*

$$S(N,0) = 0 \iff N \text{ is prime}.$$

*Proof.* If $N$ is prime, there are no zero divisors modulo $N$, so $A_{N,0} = \varnothing$ by Proposition~5.7. Hence $S(N,0) = 0$.

Now assume that $N$ is composite and $N > 4$.

If $N$ is even, the four points

$$(2, N/2), \quad (N/2, 2), \quad (N-2, N/2), \quad (N/2, N-2)$$

all lie in $A_{N,0}$, since each product is divisible by $N$. They are distinct because $N > 4$ implies $2 \neq N/2 \neq N-2$. These four points are the vertices of a rhombus centered at $(N/2, N/2)$, and its diagonals both have positive length. Therefore the rhombus has positive area, so $S(N, 0) > 0$.

If $N$ is odd composite, then there exists a proper divisor $d$ of $N$ with $1 < d < N$. Set $e = N/d$. Then $1 < e < N$ as well, and the four points

$$(d, e), \quad (N-d, e), \quad (d, N-e), \quad (N-d, N-e)$$

all lie in $A_{N,0}$, since each product is congruent to $0 \pmod{N}$. These four points are distinct because $N$ is odd, so neither $d = N - d$ nor $e = N - e$ can occur. They are the vertices of a rectangle of side lengths $N - 2d$ and $N - 2e$. Since $1 < d, e < N$, both lengths are positive. Hence the rectangle has positive area, so $S(N, 0) > 0$.

Thus every composite $N > 4$ satisfies $S(N, 0) > 0$. Therefore, for every $N > 4$,

$$S(N, 0) = 0 \iff N \text{ is prime.}$$

■

As a quick numerical check, the even composite case $N = 6$ gives the four points

$$(2, 3),\ (3, 2),\ (4, 3),\ (3, 4),$$

which form the diamond of area $2$. The odd composite case $N = 9$ gives the four points

$$(3, 3),\ (6, 3),\ (3, 6),\ (6, 6),$$

which form a $3 \times 3$ square of area $9$. In both cases the theorem predicts exactly what one sees in the pictures.

The theorem is intentionally elementary. It does not yet describe the whole hull of the zero class; it only shows that one well-chosen divisor pair already forces positive area whenever $N$ is composite and larger than $4$. In that sense it is the natural first payoff of the zero-divisor side of MTMN.

Figure 19 collects the first small cases. The prime example $N = 5$ has no zero class at all; $N = 4$ is the lone composite exception; $N = 6$ and $N = 9$ are the first even and odd composite moduli with genuinely two-dimensional zero-class geometry.

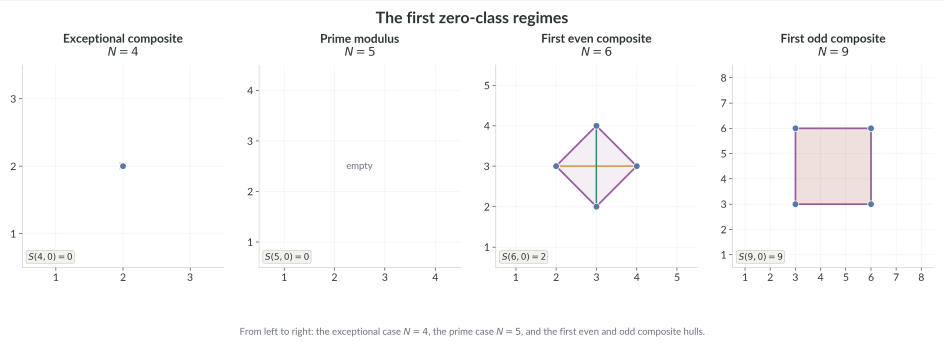


Figure 19: Small zero-class examples. The prime modulus $N = 5$ has $A_{5,0} = \varnothing$; the exceptional modulus $N = 4$ has a single zero-class point; the first even and odd composite cases $N = 6$ and $N = 9$ already have positive-area zero-class hulls.

The chapter has already taught us to read geometry from a small number of strategically chosen lattice points on special rows and columns. The proof of Theorem~5.8 follows exactly that habit. It also suggests a more structural question: once the zero class is non-degenerate, can one describe its entire hull directly from the divisor data of $N$?

### 5.8.4 Why divisors matter

Zero entries do not appear randomly in the multiplication table. They are controlled by complementary divisibility. Suppose $d \mid N$ and $x$ is a multiple of $d$. Then, in order for $xy$ to be divisible by $N$, the missing divisibility must come from $y$, and the complementary factor is $N/d$. In the zero residue class, divisor data in one coordinate therefore forces complementary divisor data in the other coordinate.

This is the source of the rectangle geometry below. A proper divisor does not merely produce one witness point such as $(d, N/d)$; it produces an entire divisor-controlled family of zero-class points.

**Definition 5.1** (Proper divisors and divisor rectangles). A *proper divisor* of $N$ is a divisor $d$ with $1 < d < N$. Many authors also say *nontrivial divisor*; in this book the two phrases mean the same thing.
For each proper divisor $d$ of $N$, define the corresponding *divisor rectangle*

$$R_d := [d,\ N-d] \times \left[\frac{N}{d},\ N - \frac{N}{d}\right].$$

Here the symbol $\times$ means Cartesian product: $R_d$ is the set of all points $(x, y)$ whose $x$-coordinate lies in the first interval and whose $y$-coordinate lies in the second. If one interval collapses to a single value, then the rectangle degenerates to a line segment; if both intervals collapse, it degenerates to a single point.

For $N = 6$, the proper divisors are $2$ and $3$. They produce

$$R_2 = [2, 4] \times \{3\}, \qquad R_3 = \{3\} \times [2, 4].$$

The braces $\{3\}$ mean that the corresponding coordinate is fixed at the single value $3$, so $R_2$ is a horizontal segment and $R_3$ is a vertical segment. Their convex hull is the small diamond that forms the full zero-class hull. It is also useful to notice the direction of these degenerate rectangles: small divisors tend to produce wide, low horizontal objects, while the complementary divisors produce tall, narrow vertical ones. For $N = 6$, $R_2$ runs left to right and $R_3$ runs bottom to top, and transpose symmetry exchanges those two directions.

Figure 20 shows this first composite example in detail. The two degenerate divisor rectangles appear as one horizontal and one vertical segment, and their convex hull is already two-dimensional.

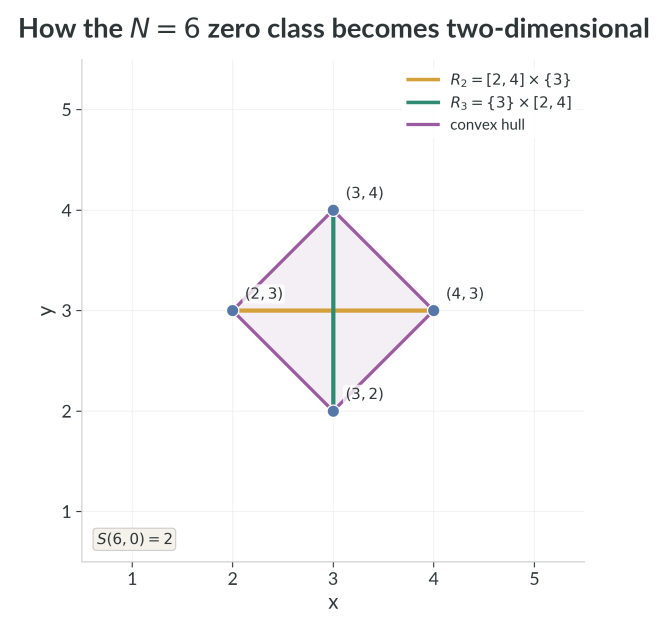


Figure 20: The zero class for $N = 6$. The divisor rectangles $R_2 = [2,4] \times \{3\}$ and $R_3 = \{3\} \times [2,4]$ degenerate to a horizontal and a vertical segment. Their convex hull is the diamond-shaped zero-class hull.

The smallest divisor gives a natural seed shape, but unless the divisor structure is very sparse, intermediate divisors can create additional hull vertices and enlarge the true zero-class hull. The first instructive example is $N = 12$. The extreme divisor pair $(2, 6)$ and $(6, 2)$ suggests a dashed seed rhombus, but the additional divisor points $(3, 4)$ and $(4, 3)$, together with their reflected partners, push the hull farther out.

Figure 21 shows exactly where those extra vertices come from. The point of the picture is not only that the hull is larger than the seed rhombus, but that the extra size is created by intermediate divisors rather than by the smallest divisor alone.

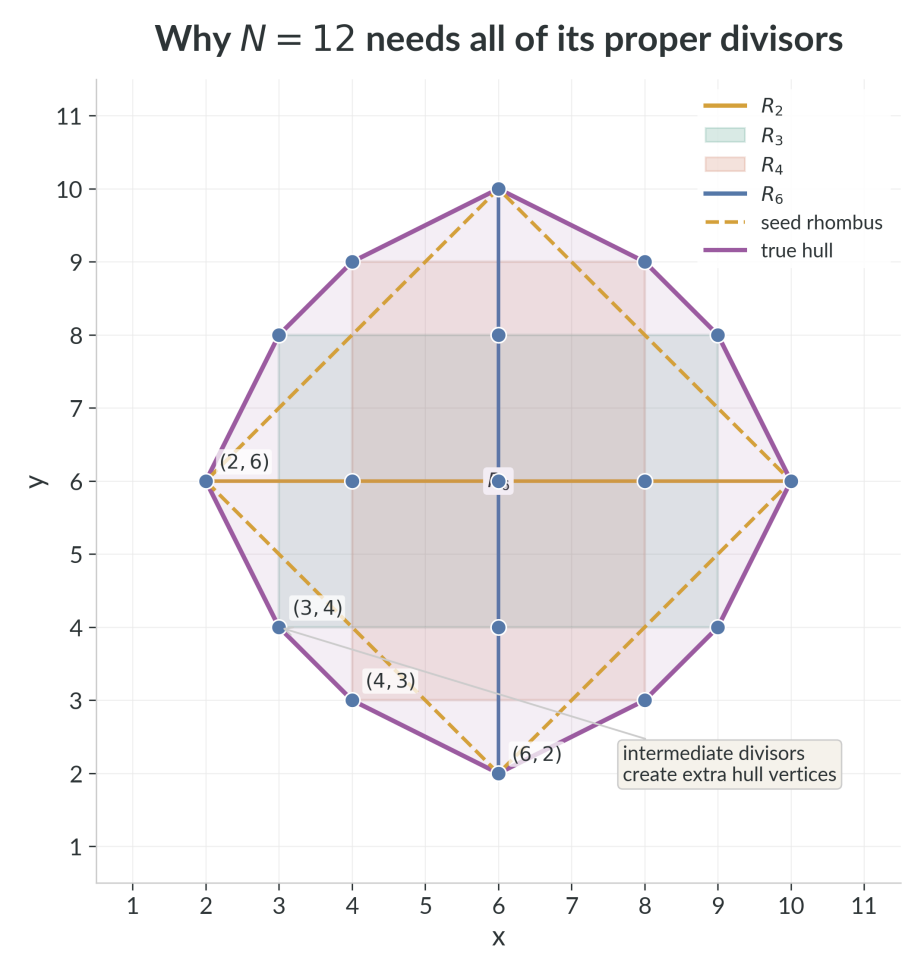


Figure 21: The zero class for $N = 12$ is built from all proper divisors, not only the smallest and largest pair. The dashed seed rhombus comes from the extreme divisor points $(2, 6)$ and $(6, 2)$, while the rectangles $R_2$, $R_3$, $R_4$, and $R_6$ together determine the true hull.

### 5.8.5 Structural anatomy of the zero class

**Theorem 5.9** (Exact divisor-rectangle description of the zero class)**.** *Assume that $N$ is composite. Then*

$$\mathrm{conv}(A_{N,0}) = \mathrm{conv}\left( \bigcup_{\substack{d \mid N \\ 1<d<N}} R_d \right).$$

*In words: every point of the zero class lies in at least one divisor rectangle, and every divisor rectangle already lies inside the zero-class hull. So the two convex hulls are the same.*

*Proof.* We prove the two containments separately.

First we show

$$\mathrm{conv}(A_{N,0}) \subseteq \mathrm{conv}\left( \bigcup_{\substack{d \mid N \\ 1<d<N}} R_d \right).$$

Let $(x, y) \in A_{N,0}$, so $N \mid xy$. Set $d := \gcd(x, N)$. Write $x = dx'$, $N = dN'$, with $\gcd(x', N') = 1$. Since $N \mid xy$, we obtain $dN' \mid dx'y$, so $N' \mid x'y$. Because $\gcd(x', N') = 1$, there exists an inverse of $x'$ modulo $N'$, and therefore $N' \mid y$.

We claim that $d$ is a proper divisor of $N$. If $d = 1$, then $N' = N$, so $N \mid y$, impossible because $1 \leq y \leq N - 1$. Also $d < N$ because $d \mid x$ and $x \leq N - 1$. Hence $1 < d < N$.

Now $d \mid x$ and $x < N$, so

$$d \leq x \leq N - d.$$

Similarly, $N' = N/d$ divides $y$ and $y < N$, so

$$\frac{N}{d} \leq y \leq N - \frac{N}{d}.$$

Therefore $(x, y) \in R_d$. Since $(x, y)$ was arbitrary,

$$A_{N,0} \subseteq \bigcup_{\substack{d \mid N \\ 1<d<N}} R_d,$$

and hence

$$\mathrm{conv}(A_{N,0}) \subseteq \mathrm{conv}\left( \bigcup_{\substack{d \mid N \\ 1<d<N}} R_d \right).$$

This proves the first containment.

For the other containment, fix a proper divisor $d$ of $N$. The four corners

$$\left(d, \frac{N}{d}\right), \quad \left(N-d, \frac{N}{d}\right), \quad \left(d, N-\frac{N}{d}\right), \quad \left(N-d, N-\frac{N}{d}\right)$$

all lie in $A_{N,0}$, because each coordinate product is divisible by $N$. Since $R_d$ is the convex hull of these four corners, possibly degenerate to a segment or a point, it follows that

$$R_d \subseteq \mathrm{conv}(A_{N,0}).$$

In other words, each divisor rectangle is generated by four genuine zero-class points, so the whole rectangle already lies inside the zero-class hull.

Taking the union over all proper divisors and then taking convex hulls yields

$$\mathrm{conv}\left( \bigcup_{\substack{d \mid N \\ 1<d<N}} R_d \right) \subseteq \mathrm{conv}(A_{N,0}).$$

This proves the other containment. Combining the two containments proves the theorem. ■

**Worked check:** $N = 12$, $(8, 9)$. $d = \gcd(8, 12) = 4$, $N' = 12/4 = 3$. Because $3 \mid 9$, the divisibility step succeeds. Also $1 < 4 < 12$, and $4 \le 8 \le 8$, $3 \le 9 \le 9$. So $(8, 9) \in [4, 8] \times [3, 9] = R_4$. This is the theorem's rectangle-membership argument with concrete numbers filled in.

*Remark* 5.2 (Four geometric objects). It is important to distinguish four different sets:

$$A_{N,0}, \qquad \bigcup_d R_d, \qquad \mathrm{conv}(A_{N,0}), \qquad \mathrm{conv}\left(\bigcup_d R_d\right).$$

The first is a finite lattice set. The second is a filled planar set built from divisor rectangles. The theorem says that the last two convex sets are equal. It does **not** say that the discrete zero class itself equals the rectangle union.

### 5.8.6 Reading the lower edge of the zero-class hull

The rectangle theorem uses every point of every divisor rectangle. One can describe the same hull more economically by keeping only the lower corners of those rectangles. The lower polygonal chain, or broken line, through those corners controls the whole hull. Here a broken line means a connected chain of straight segments.

We write $\ell_N(x)$ for the height of that lower broken line at horizontal position $x$. In plain language, $\ell_N(x)$ tells us how low the zero-class hull sits above the $x$-axis at the chosen value of $x$.

**Definition 5.2** (Divisor points and the lower edge of the hull). Assume that $N$ is composite, and let $p$ be the smallest proper divisor of $N$:

$$p := \min\{d > 1 : d \mid N\}.$$

Define the *divisor point set*

$$E_N := \left\{ \left(d, \frac{N}{d}\right), \left(N-d, \frac{N}{d}\right) : 1 < d < N,\ d \mid N \right\}.$$

For $N = 12$ this gives

$$E_{12} = \{(2,6),\ (3,4),\ (4,3),\ (6,2),\ (8,3),\ (9,4),\ (10,6)\}.$$

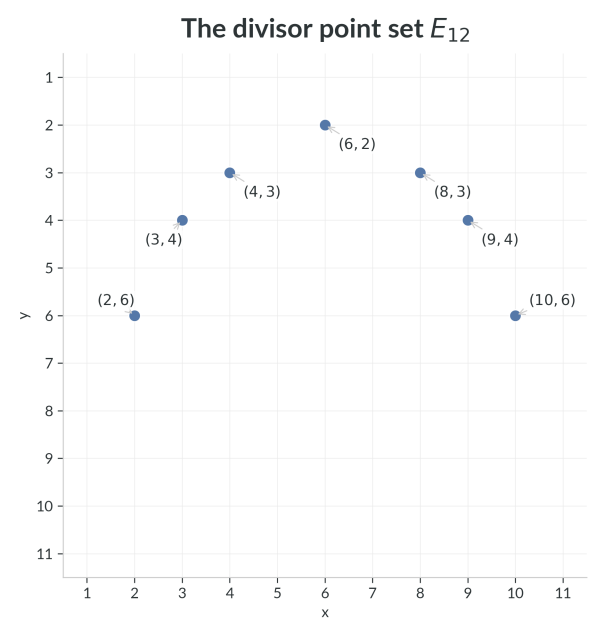


Figure 22: The divisor point set $E_{12}$, with every point labeled on the lattice.

For $p \le x \le N - p$, let $\ell_N(x)$ denote the $y$-coordinate of the lower edge of conv$(E_N)$ at horizontal position $x$. Equivalently, $\ell_N$ is the broken line obtained by tracing that lower edge from left to right. This lower edge is also called the lower convex envelope of the divisor points, meaning simply the lower boundary traced by their convex hull.

Figure 23 is the picture to keep in mind. The blue divisor points determine a lower orange broken line; that lower line is $y = \ell_{12}(x)$. The upper orange line is just its reflected partner, and the shaded region between them is the full zero-class hull.

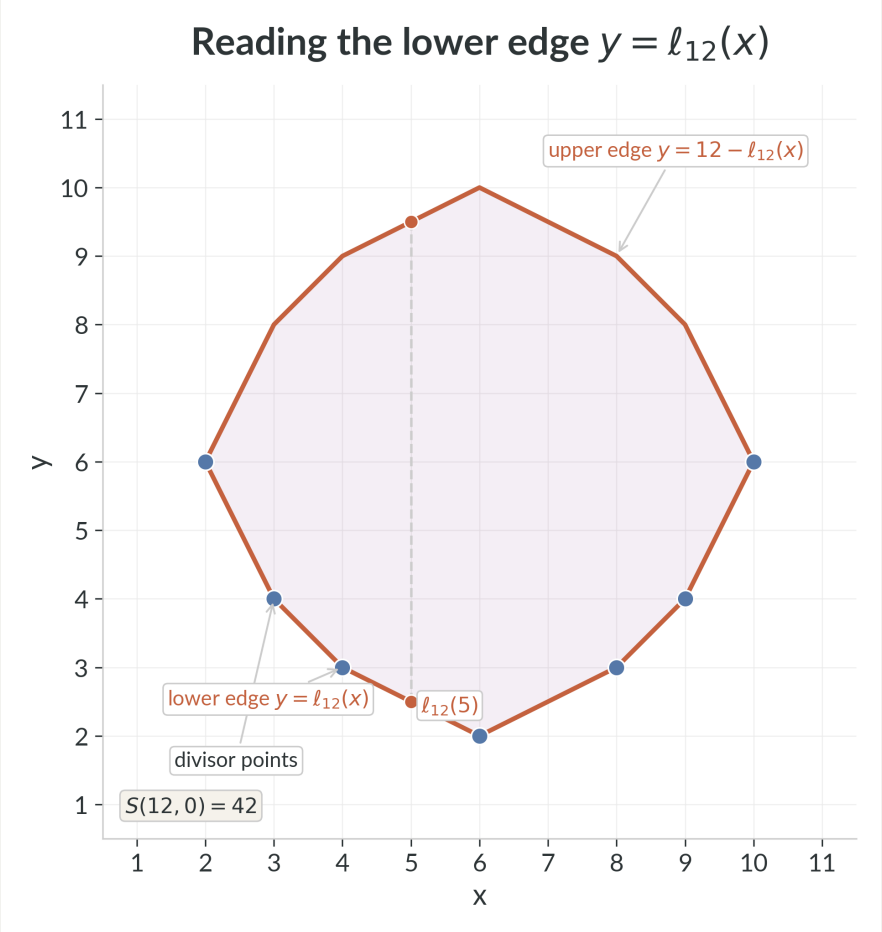


Figure 23: How to read $\ell_N(x)$ in practice. For $N = 12$, the blue divisor points determine a lower broken line. The value $\ell_{12}(x)$ is simply the height of that lower line at the chosen horizontal coordinate $x$. Reflecting that lower line through the center of the square produces the upper edge of the hull, so the shaded region between the two lines is exactly the zero-class hull.

The divisor points already encode the exact lower edge of the zero-class hull. Because $E_N$ is symmetric under $x \mapsto N - x$, the function $\ell_N$ satisfies

$$\ell_N(N - x) = \ell_N(x).$$

By the central symmetry from Proposition~5.2, the upper boundary of the zero-class hull

is therefore the reflected graph

$$y = N - \ell_N(x).$$

Transpose symmetry identifies the left and right portions of the hull across the diagonal line $x = y$; central symmetry reflects the lower hull to the upper hull.

**Proposition 5.10** (The hull as the region between two broken lines)**.** *Assume that $N$ is composite. Then*

$$\mathrm{conv}(A_{N,0}) = \{(x, y) : p \le x \le N - p,\ \ell_N(x) \le y \le N - \ell_N(x)\}.$$

*Proof.* Let

$$\mathcal{R}_N := \bigcup_{\substack{d \mid N \\ 1 < d < N}} R_d.$$

By Theorem~5.9, it is enough to describe $\mathrm{conv}(\mathcal{R}_N)$.

The key idea is simple: for a fixed $x$-coordinate, the lowest point of the hull is obtained by sliding downward inside the divisor rectangles until one reaches their lower edges. Those lower-edge points are exactly the points that build $\mathrm{conv}(E_N)$.

Take any point $(x, y) \in \mathrm{conv}(\mathcal{R}_N)$. Then

$$(x, y) = \sum_{j=1}^{m} \lambda_j (x_j, y_j)$$

for some convex combination with $(x_j, y_j) \in \mathcal{R}_N$. For each $j$, there exists a proper divisor $d_j$ such that $(x_j, y_j) \in R_{d_j}$. The point

$$\left(x_j, \frac{N}{d_j}\right)$$

lies on the lower edge of the same rectangle $R_{d_j}$, so it has the same $x$-coordinate and no larger $y$-coordinate. Moreover, it lies on the horizontal segment joining

$$\left(d_j, \frac{N}{d_j}\right) \quad \text{and} \quad \left(N - d_j, \frac{N}{d_j}\right),$$

so it belongs to conv$(E_N)$.

Therefore the convex combination

$$\left( x, \sum_{j=1}^{m} \lambda_j \frac{N}{d_j} \right)$$

lies in conv$(E_N)$, has the same $x$-coordinate as $(x, y)$, and has $y$-coordinate at most $y$. This shows that, for each fixed $x$, the smallest possible $y$-coordinate in conv$(\mathcal{R}_N)$ is exactly the lower boundary of conv$(E_N)$, namely $\ell_N(x)$.

Because conv$(A_{N,0}) =$ conv$(\mathcal{R}_N)$ and the zero class is centrally symmetric by Proposition~5.2, the upper boundary of conv$(\mathcal{R}_N)$ is the reflection of the lower boundary under $(x, y) \mapsto (N - x, N - y)$, hence it is given by $y = N - \ell_N(x)$. Since every divisor rectangle satisfies $p \le x \le N - p$, the same interval is the full horizontal domain of the hull. The stated description follows. ■

Figure 24 isolates the two Euclidean symmetries used here.

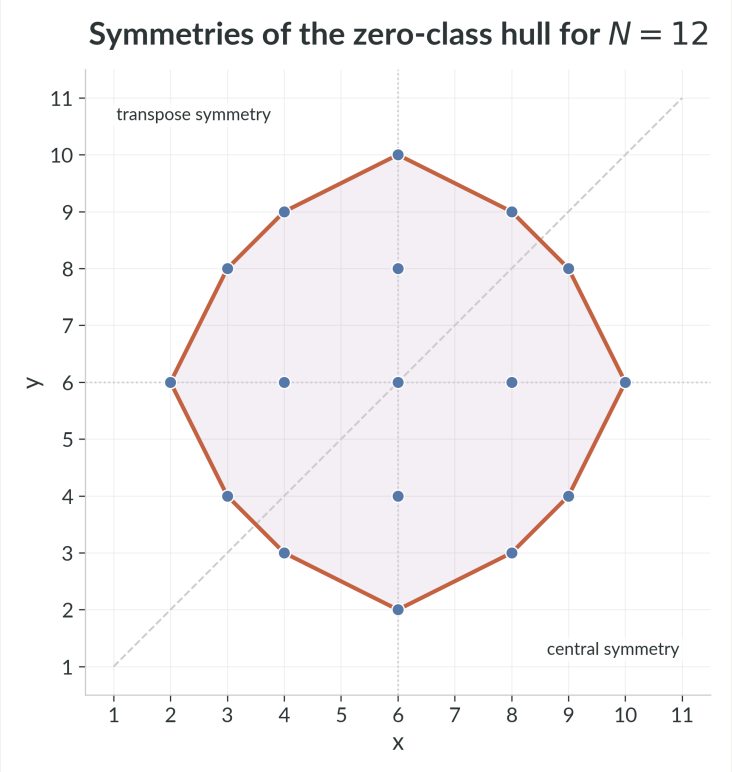


Figure 24: The zero-class hull inherits both transpose symmetry and central symmetry. The diagonal line $x = y$ identifies the left and right portions of the hull, while the half-turn about $(N/2, N/2)$ reflects the lower hull to the upper hull.

The same divisor points may also be viewed as samples on the hyperbola $y = N/x$. Figure 25 uses that picture only as a guide to intuition: the exact hull comes from the finite divisor data, not from the full continuous curve.

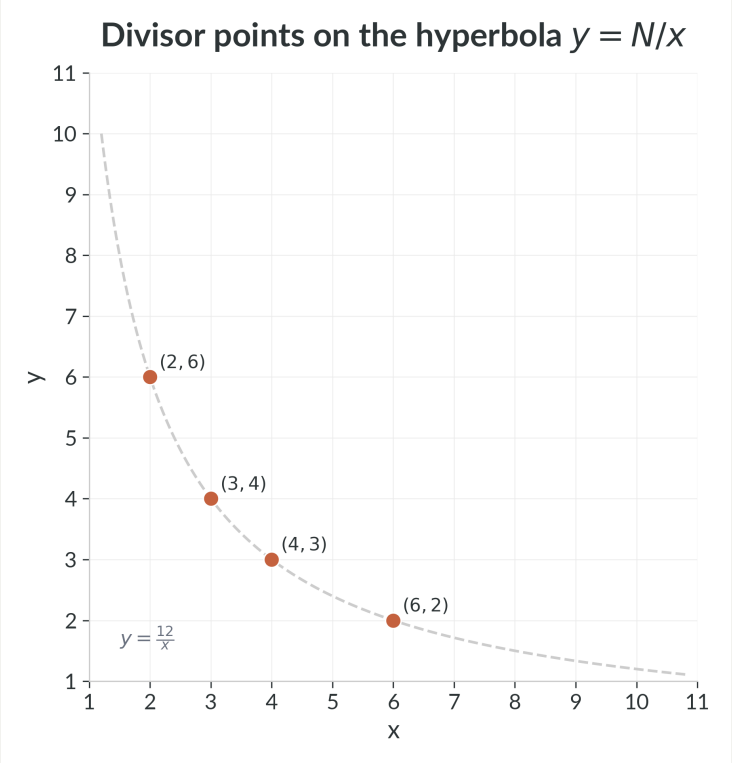


Figure 25: The divisor points $(d, N/d)$ for $N = 12$ lie on the sampled hyperbola $y = N/x$. The continuous hyperbola is only a guide to intuition; the exact hull comes from the finite divisor data, not from the full curve.

**Corollary 5.11** (Integral formula for the zero-class area)**.** *Assume that $N$ is composite. Then*

$$S(N, 0) = \int_p^{N-p} \big(N - 2\ell_N(x)\big)\, dx.$$

*Proof.* By Proposition~5.10, the vertical thickness of the hull at horizontal coordinate $x$ is

$$\big(N - \ell_N(x)\big) - \ell_N(x) = N - 2\ell_N(x).$$

Integrating this thickness across the interval $[p, N - p]$ gives the area. ∎

For example, when $N = 9$, the only proper divisor is $3$, so the lower edge is the constant

line $\ell_9(x) = 3$ on the interval $[3, 6]$. The integral becomes

$$\int_3^6 (9 - 2 \cdot 3)\, dx = \int_3^6 3\, dx = 9,$$

which matches $S(9, 0) = 9$.

**Corollary 5.12** (Trapezoidal formula for the zero-class area)**.** *Assume that $N$ is composite, and let the lower-hull vertices of* conv$(E_N)$ *be*

$$v_1 = (x_1, y_1),\ v_2 = (x_2, y_2),\ \ldots,\ v_r = (x_r, y_r), \qquad x_1 = p,\ x_r = N - p,$$

*listed from left to right. Then*

$$S(N, 0) = \sum_{j=1}^{r-1} (x_{j+1} - x_j)\big(N - y_j - y_{j+1}\big).$$

*Proof.* On each interval $[x_j, x_{j+1}]$, the lower boundary is the line segment joining $(x_j, y_j)$ to $(x_{j+1}, y_{j+1})$, and the upper boundary is its reflected partner. The vertical thickness at the two endpoints is

$$N - 2y_j \qquad \text{and} \qquad N - 2y_{j+1}.$$

Hence the area of the strip over that interval is the area of a trapezoid:

$$\frac{x_{j+1} - x_j}{2}\big((N - 2y_j) + (N - 2y_{j+1})\big) = (x_{j+1} - x_j)\big(N - y_j - y_{j+1}\big).$$

Summing over all consecutive lower-hull segments gives the formula. ∎

Figure 26 shows one such strip. Each linear piece of the lower hull contributes one reflected trapezoid, and the full area is the sum of those elementary pieces.

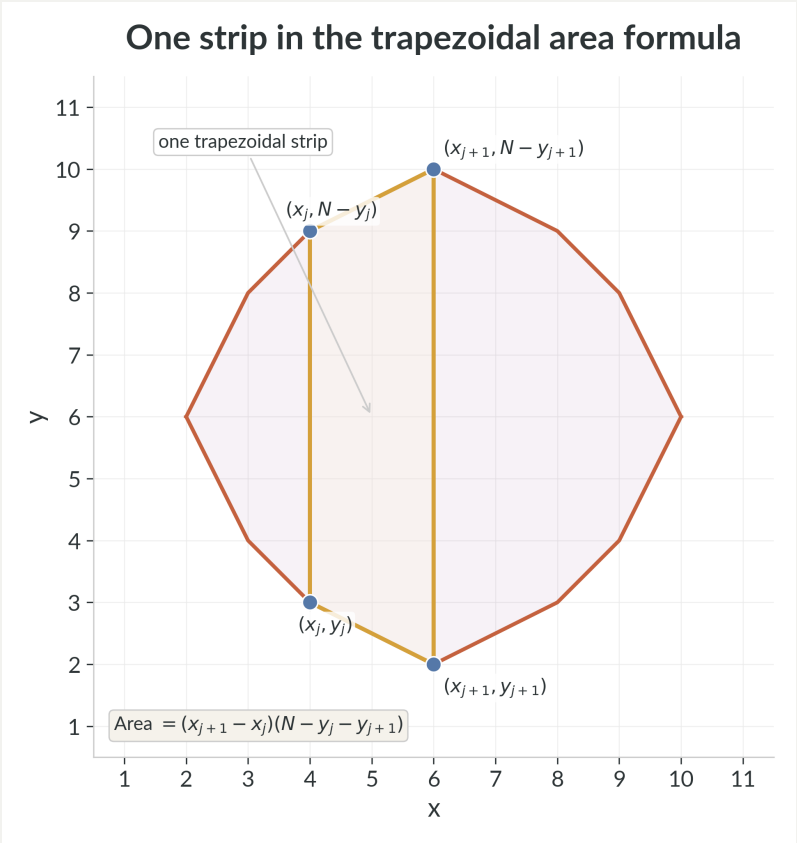


Figure 26: One trapezoidal strip in the zero-class area formula. The lower edge runs from $(x_j, y_j)$ to $(x_{j+1}, y_{j+1})$, the upper edge is its central-symmetry reflection, and the enclosed strip has area $(x_{j+1} - x_j)(N - y_j - y_{j+1})$. Summing these strips over all consecutive lower-hull vertices gives $S(N, 0)$.

For $N = 12$, the lower-hull vertices are

$$(2, 6),\ (3, 4),\ (4, 3),\ (6, 2),\ (8, 3),\ (9, 4),\ (10, 6).$$

Therefore Corollary~5.12 gives

$$\begin{aligned} S(12,0) &= (3-2)(12-6-4)+(4-3)(12-4-3)+(6-4)(12-3-2) \\ &\quad +(8-6)(12-2-3)+(9-8)(12-3-4)+(10-9)(12-4-6) \\ &= 2+5+14+14+5+2 \\ &= 42. \end{aligned}$$

This agrees with the exact hull area computed directly from the lattice picture.

The zero class is now understood in two complementary ways. Theorem~5.8 gives a clean prime/composite classification of when the hull is degenerate. Theorem~5.9 and Corollaries~5.11–5.12 explain the full composite geometry: $S(N,0)$ is the area of a completely described divisor-controlled hull.

There is one further comparison that now becomes irresistible. The decisive support points

$$\left(d, \frac{N}{d}\right)$$

lie on the continuous hyperbola

$$y = \frac{N}{x},$$

but the exact hull does not follow that curve. It follows the polygonal divisor envelope. The next question is therefore not whether the hyperbola matters; it is exactly how much area is created when one replaces the smooth curve by the arithmetic polygon.

## 5.9 The hyperbola gap on the informative half

What is true is that the lower-hull support points lie on the hyperbola $y = N/x$. What is not true is that every point of the lower hull lies on that hyperbola. The hull is the polygonal lower convex envelope of the sampled divisor points, so its edges are straight secants rather than curved arcs. Figure 27 shows this directly for $N = 12$: the vertices lie on the curve, but the orange segments do not.

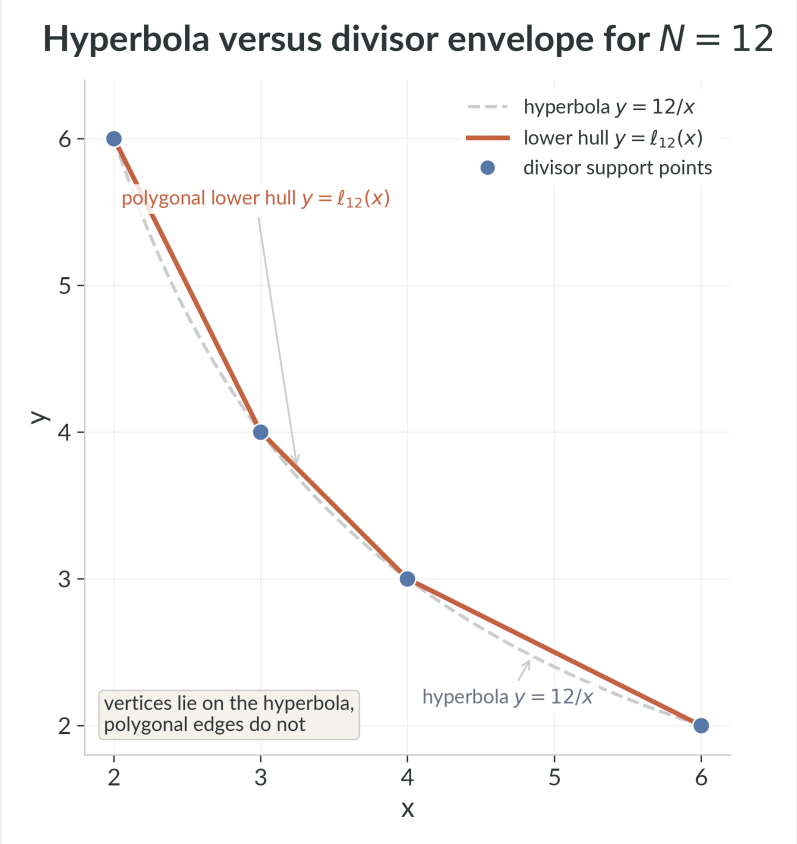


Figure 27: For $N = 12$, the lower hull $y = \ell_{12}(x)$ is the polygonal envelope through divisor support points on the hyperbola $y = 12/x$. The vertices lie on the hyperbola, but the polygonal edges lie above it.

**Proposition 5.13** (The divisor envelope lies above the hyperbola). *Assume that $N$ is composite. Then on the informative left half,*

$$\ell_N(x) \geq \frac{N}{x} \qquad \text{for } p \leq x \leq \frac{N}{2}.$$

*Proof.* The function

$$f(x) = \frac{N}{x}$$

has second derivative

$$f''(x) = \frac{2N}{x^3} > 0$$

for $x > 0$, so its graph is convex. Therefore, whenever one takes two divisor points

$$\left(d_j, \frac{N}{d_j}\right), \qquad \left(d_{j+1}, \frac{N}{d_{j+1}}\right)$$

that appear consecutively on the left lower hull, the secant line joining them lies above the graph of $y = N/x$ on the whole interval $[d_j, d_{j+1}]$.

Those secant lines are exactly the linear pieces of $\ell_N$ until the symmetry point is reached. If $N$ is odd, there is one final interval from the last left-side divisor vertex to $x = N/2$ on which the lower hull is horizontal. Since $x \mapsto N/x$ is decreasing, the hyperbola also lies below that horizontal segment. Hence every piece of the left lower hull lies on or above the hyperbola, so

$$\ell_N(x) \geq \frac{N}{x}$$

throughout $[p, N/2]$. ■

This order of subtraction is therefore the correct one: the hyperbola is the lower continuous model, while the divisor envelope is the higher polygonal arithmetic model.

**Definition 5.3** (The left-half hyperbola gap)**.** Assume that $N$ is composite. Define

$$\Delta_N := \int_p^{N/2} \left(\ell_N(x) - \frac{N}{x}\right) dx.$$

By Proposition~5.13, the integrand is everywhere nonnegative, so

$$\Delta_N \geq 0.$$

Geometrically, $\Delta_N$ is the area between the polygonal lower hull and the continuous hyperbola on the informative half of the picture. The shaded region in Figure 28 is exactly that area for $N = 12$.

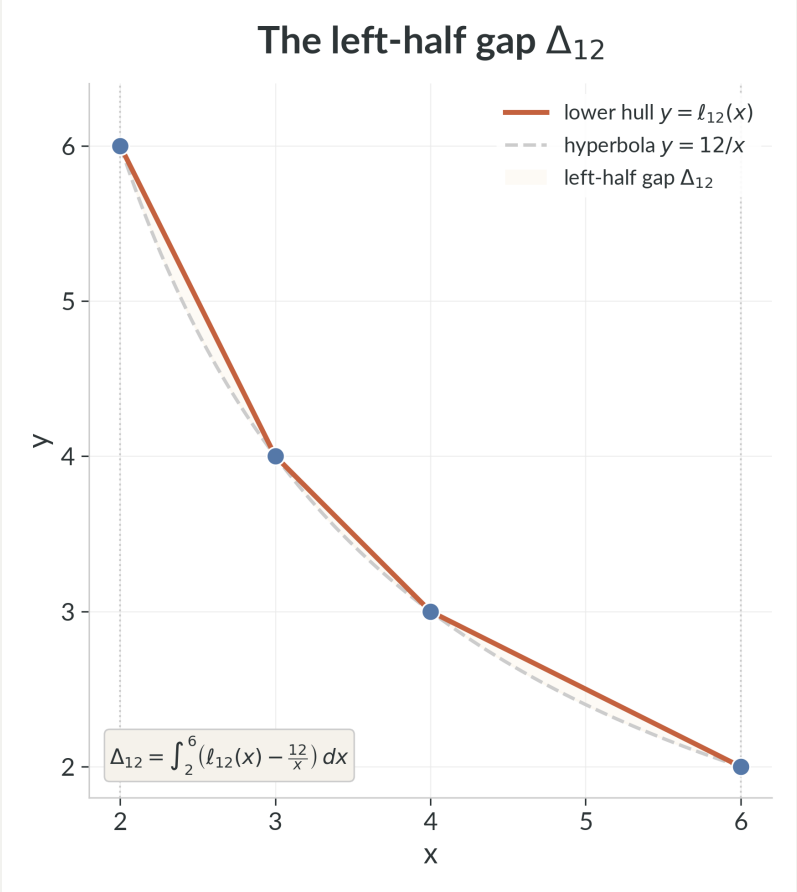


Figure 28: The full left-half gap for $N = 12$. The shaded region is $\Delta_{12} = \int_2^6 (\ell_{12}(x) - 12/x)\,dx$.

## 5.10 Decomposing the gap segment by segment

Because $\ell_N$ is piecewise linear, the total gap breaks into the sum of the smaller regions trapped between each hull segment and the hyperbola.

Let

$$p = x_1 < x_2 < \cdots < x_t = \frac{N}{2}$$

be the breakpoints of the left-half graph of $\ell_N$, and write

$$h_j := \ell_N(x_j).$$

On each interval $[x_j, x_{j+1}]$, the lower hull is the line

$$L_j(x) = h_j + \frac{h_{j+1} - h_j}{x_{j+1} - x_j}(x - x_j).$$

**Proposition 5.14** (Piecewise decomposition of the hyperbola gap)**.** *With the notation above,*

$$\Delta_N = \sum_{j=1}^{t-1} \int_{x_j}^{x_{j+1}} \left( L_j(x) - \frac{N}{x} \right) dx.$$

*Proof.* On each interval $[x_j, x_{j+1}]$, the lower hull agrees with the linear function $L_j$. Splitting the integral for $\Delta_N$ across those consecutive intervals gives the stated sum. ∎

When both endpoints of one interval are genuine divisor vertices, the line has a particularly simple form. Suppose

$$\left( d_j, \frac{N}{d_j} \right), \qquad \left( d_{j+1}, \frac{N}{d_{j+1}} \right)$$

are consecutive left-side divisor vertices joined by one segment of the lower hull. Then on $[d_j, d_{j+1}]$,

$$L_j(x) = \frac{N}{d_j} + \frac{\frac{N}{d_{j+1}} - \frac{N}{d_j}}{d_{j+1} - d_j}(x - d_j).$$

This is the local picture isolated in Figure 29.

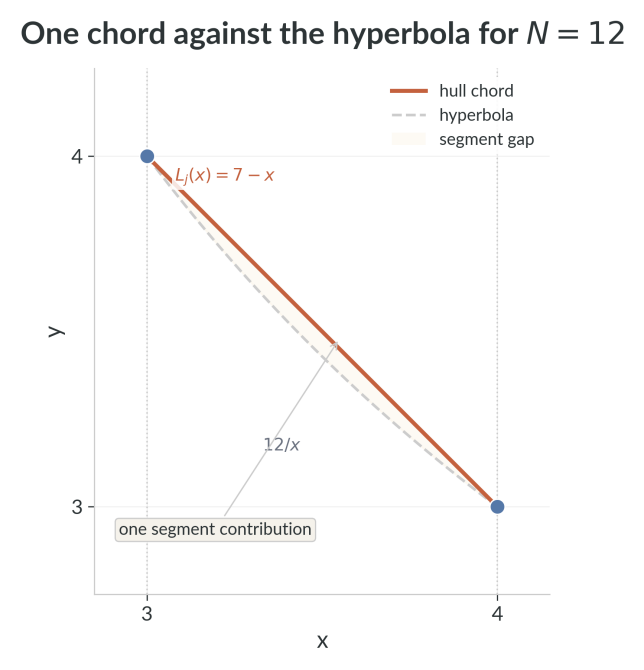


Figure 29: One segment contribution for $N = 12$. The shaded region is the area between one hull chord and the hyperbola on a single interval.

**Proposition 5.15** (A closed formula for one divisor-to-divisor segment). *In the divisor-to-divisor situation above,*

$$\int_{d_j}^{d_{j+1}} L_j(x)\,dx = \frac{d_{j+1} - d_j}{2}\left(\frac{N}{d_j} + \frac{N}{d_{j+1}}\right),$$

$$\int_{d_j}^{d_{j+1}} \frac{N}{x}\,dx = N \ln\left(\frac{d_{j+1}}{d_j}\right),$$

*and therefore*

$$\int_{d_j}^{d_{j+1}} \left(L_j(x) - \frac{N}{x}\right) dx = \frac{d_{j+1} - d_j}{2}\left(\frac{N}{d_j} + \frac{N}{d_{j+1}}\right) - N \ln\left(\frac{d_{j+1}}{d_j}\right).$$

*Proof.* The integral of a line over an interval is the area of the corresponding trapezoid, so

$$\int_{d_j}^{d_{j+1}} L_j(x)\,dx = \frac{d_{j+1}-d_j}{2}\left(\frac{N}{d_j}+\frac{N}{d_{j+1}}\right).$$

Also

$$\int_{d_j}^{d_{j+1}} \frac{N}{x}\,dx = N\big[\ln x\big]_{d_j}^{d_{j+1}} = N\ln\left(\frac{d_{j+1}}{d_j}\right).$$

Subtracting gives the segment formula. ■

*Remark* 5.3 (A ratio form). On a divisor-to-divisor interval it is natural to write

$$r_j := \frac{d_{j+1}}{d_j}.$$

Then

$$\frac{d_{j+1}-d_j}{2}\left(\frac{N}{d_j}+\frac{N}{d_{j+1}}\right) = N\cdot\frac{r_j-1}{2}\left(1+\frac{1}{r_j}\right),$$

so the same segment contribution becomes

$$N\,F(r_j), \qquad F(r) := \frac{r-1}{2}\left(1+\frac{1}{r}\right) - \ln r.$$

Thus each divisor-to-divisor gap contributes an amount determined only by the ratio of successive divisor abscissas, scaled by $N$.

*Remark* 5.4 (If the midpoint is not itself a divisor point). For odd $N$, the last left-half interval runs from the final divisor vertex to $x = N/2$ and is horizontal rather than divisor-to-divisor. Proposition 5.14 still applies exactly, and the same trapezoid-minus-log calculation still evaluates that last piece. The ratio formula above is simply the cleanest form for the intervals whose endpoints are both sampled divisor points.

## 5.11 Exact relation with the zero-class area

The point of $\Delta_N$ is that it measures the exact difference between the true hull area and the area one would obtain by pretending that the lower boundary were the continuous

hyperbola itself. By symmetry, the left half already determines the full area, as Figure 30 makes visible.

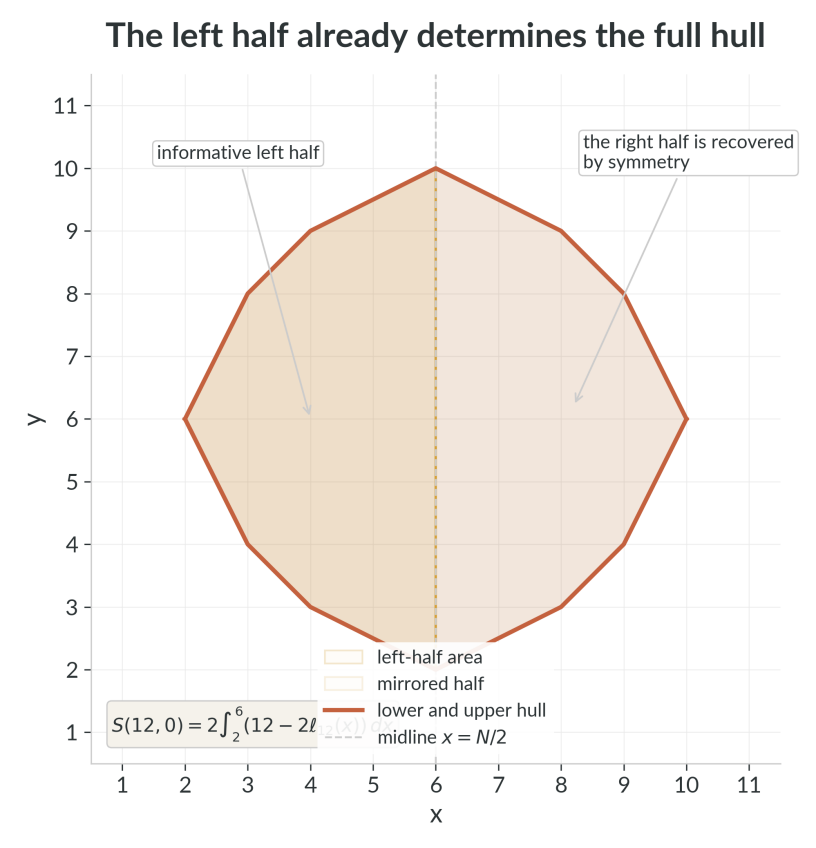


Figure 30: The informative left half and its mirror. By symmetry, the left-half thickness data already determines the whole zero-class hull.

**Theorem 5.16** (Exact hyperbolic decomposition of the zero-class area)**.** *Assume that* $N$ *is composite. Then*

$$S(N,0) = N(N-2p) - 4N \ln\left(\frac{N}{2p}\right) - 4\Delta_N.$$

*Equivalently,*

$$\Delta_N = \frac{1}{4}\left[N(N-2p) - 4N\ln\left(\frac{N}{2p}\right) - S(N,0)\right].$$

*Proof.* Corollary~5.11 gives

$$S(N,0) = \int_p^{N-p} \big(N - 2\ell_N(x)\big)\,dx.$$

Because $\ell_N(N-x) = \ell_N(x)$, the integrand is symmetric about $x = N/2$, so

$$S(N,0) = 2\int_p^{N/2} \big(N - 2\ell_N(x)\big)\,dx.$$

Now write

$$\ell_N(x) = \frac{N}{x} + \left(\ell_N(x) - \frac{N}{x}\right).$$

Then

$$N - 2\ell_N(x) = \left(N - \frac{2N}{x}\right) - 2\left(\ell_N(x) - \frac{N}{x}\right).$$

Substituting into the integral gives

$$S(N,0) = 2\int_p^{N/2} \left(N - \frac{2N}{x}\right) dx - 4\int_p^{N/2} \left(\ell_N(x) - \frac{N}{x}\right) dx.$$

The second integral is exactly $\Delta_N$. The first is

$$2\left[Nx - 2N\ln x\right]_p^{N/2} = N(N-2p) - 4N\ln\left(\frac{N}{2p}\right).$$

Combining the two terms proves the stated identity. ∎

*Remark* 5.5 (Baseline and correction). The term

$$N(N-2p) - 4N\ln\left(\frac{N}{2p}\right)$$

is the area obtained by replacing the exact thickness

$$N - 2\ell_N(x)$$

with the smoother hyperbolic thickness

$$N - \frac{2N}{x}$$

on the informative half and then reflecting by symmetry. The exact area is that hyperbolic baseline minus the correction $4\Delta_N$. This is not an approximation claim. It is an exact geometric decomposition of the zero-class area into a continuous term and an arithmetic polygonal correction.

This comparison sits beside the same classical hyperbolic backdrop as the Dirichlet divisor problem, where one studies lattice-point data attached to the region under $xy = N$, equivalently under the graph $y = N/x$ [12]. The object measured here is different, however. The term $\Delta_N$ is not a summatory divisor error term, but the exact area between the smooth hyperbola and the polygonal divisor envelope that governs the zero-class hull for one fixed modulus.

The pointwise version of that correction is drawn in Figure 31: the exact thickness curve sits above the hyperbolic baseline by twice the local gap $\ell_N(x) - N/x$.

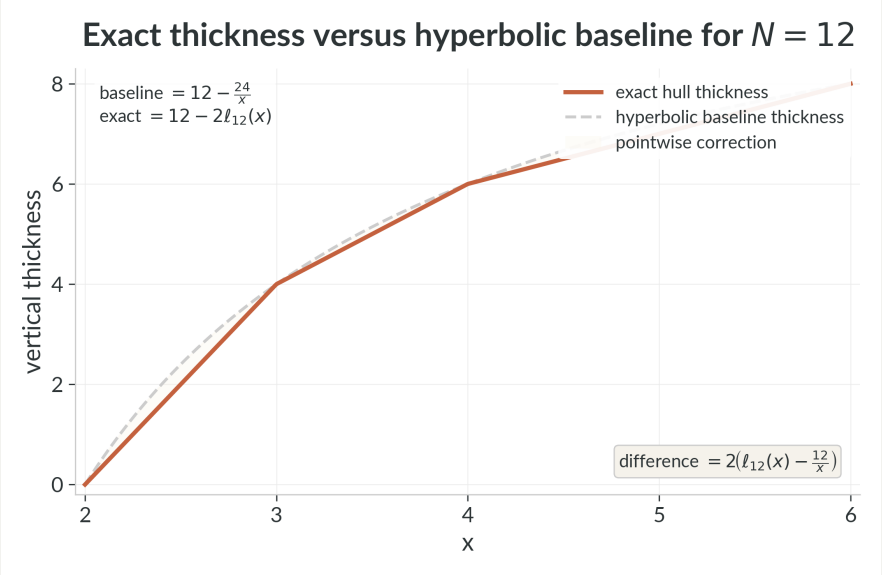


Figure 31: The exact hull thickness $N - 2\ell_{12}(x)$ and the hyperbolic baseline thickness $12 - 24/x$. Their difference is twice the pointwise gap between the divisor envelope and the hyperbola.

## 5.12 Worked example: $N = 12$

For $N = 12$, the informative left-half lower-hull vertices are

$$(2, 6),\ (3, 4),\ (4, 3),\ (6, 2).$$

Hence

$$\ell_{12}(x) = 10 - 2x \qquad \text{for } 2 \le x \le 3,$$

$$\ell_{12}(x) = 7 - x \qquad \text{for } 3 \le x \le 4,$$

and

$$\ell_{12}(x) = 5 - \frac{x}{2} \qquad \text{for } 4 \le x \le 6.$$

Therefore

$$\Delta_{12} = \int_2^3 \left(10 - 2x - \frac{12}{x}\right) dx + \int_3^4 \left(7 - x - \frac{12}{x}\right) dx + \int_4^6 \left(5 - \frac{x}{2} - \frac{12}{x}\right) dx.$$

Evaluating each piece gives

$$\int_2^3 \left(10 - 2x - \frac{12}{x}\right)\,dx = 5 - 12\ln\left(\frac{3}{2}\right),$$

$$\int_3^4 \left(7 - x - \frac{12}{x}\right)\,dx = \frac{7}{2} - 12\ln\left(\frac{4}{3}\right),$$

and

$$\int_4^6 \left(5 - \frac{x}{2} - \frac{12}{x}\right)\,dx = 5 - 12\ln\left(\frac{3}{2}\right).$$

Summing yields

$$\Delta_{12} = \frac{27}{2} - 12\left(2\ln\left(\frac{3}{2}\right) + \ln\left(\frac{4}{3}\right)\right) = \frac{27}{2} - 12\ln 3.$$

The hyperbolic baseline term is

$$H_{12} := 12(12-4) - 48\ln 3 = 96 - 48\ln 3.$$

Theorem~5.16 therefore gives

$$\begin{aligned} S(12,0) &= H_{12} - 4\Delta_{12} \\ &= (96 - 48\ln 3) - 4\left(\frac{27}{2} - 12\ln 3\right) \\ &= 42. \end{aligned}$$

This matches the earlier trapezoidal calculation exactly. So the two pictures agree perfectly: the full hull area is $42$, and the quantity

$$\Delta_{12} = \frac{27}{2} - 12\ln 3$$

is precisely the arithmetic correction that converts the smooth hyperbolic baseline into that exact polygonal area.

The zero class is now understood in three compatible ways: as a divisor-rectangle hull, as a region between the broken lines $y = \ell_N(x)$ and $y = N - \ell_N(x)$, and as a hyperbolic baseline corrected by the nonnegative gap term $\Delta_N$. The next chapter changes the language rather

than the object: it keeps the same finite residue classes, but recasts their areas through support functions and exact integral formulas.

# 6 Exact formulas using convex geometry

The preceding chapters keep the geometry concrete: residue classes are drawn in the positive lattice window, and the zero class is described both by its divisor envelope and by its exact comparison with the continuous hyperbola $y = N/x$. This chapter does not change the underlying object of study. The set $A_{N,a}$ remains the same finite arithmetic set cut out by a congruence. What changes here is the language. We translate the same finite data into support functions, integrals, and exponential expressions in order to obtain exact but implicit formulas.

## 6.1 The support function

Before introducing the notation, it helps to fix the geometric picture. Choose a direction $\theta$. Now look at all lines perpendicular to that direction and slide one of them across the plane until it just touches the convex hull from the outside. The touching point is a point where the quantity

$$x\cos\theta + y\sin\theta$$

is as large as possible. So the support function is simply a record of how far out the set reaches in each direction.

Figure 32 shows this in a concrete case. For the diagonal direction $\theta = \pi/4$, maximizing $x\cos\theta + y\sin\theta$ is the same as maximizing $x + y$, so one sees the supporting line and the winning hull points directly.

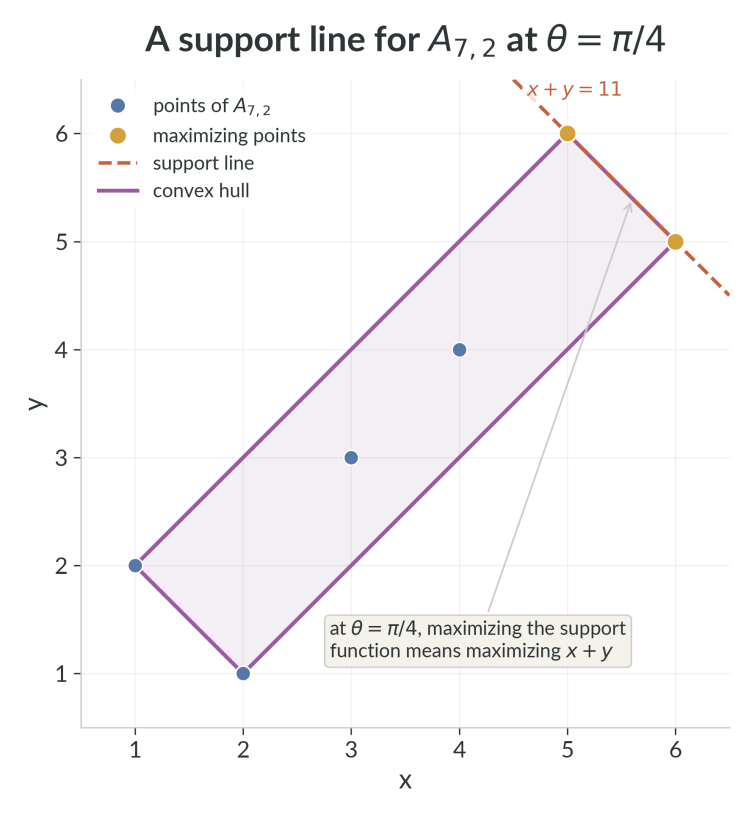


Figure 32: A support line for $A_{7,2}$ at the diagonal direction $\theta = \pi/4$. The dashed line is the outermost line of the form $x + y =$ constant that still touches the hull, so the highlighted points are exactly the maximizers of the support function in that direction.

The diagonal directions $\theta = \pm\pi/4$ are also where the support-function viewpoint meets earlier extremal work on modular inverses. For the unit hyperbola $xy \equiv 1 \pmod{n}$, maximizing $|x - y|$ is exactly the problem of pushing a support line in the $\pm\pi/4$ directions, and Ford, Khan, Shparlinski, and Yankov study the resulting extremal quantity $M(n) = \max\{|a - b| : ab \equiv 1 \pmod{n}\}$ by analytic and divisor-distribution methods [13]. Their focus is asymptotic and one-directional: the unit class, one extremal statistic, and upper/lower bounds. The present chapter keeps that same geometric mechanism but asks for the full support function $h_{N,a}(\theta)$ for every residue class and then integrates it into exact area formulas.

For a compact convex set $K \subset \mathbb{R}^2$, its *support function* is

$$h_K(\theta) = \max_{(x,y)\in K} (x\cos\theta + y\sin\theta).$$

Because the maximum of a linear functional does not change when passing from a set to its convex hull,

$$h_{\mathrm{conv}(E)}(\theta) = \max_{(x,y)\in E} (x\cos\theta + y\sin\theta)$$

for every finite set $E$.

For the modular set $A_{N,a}$ we therefore define

$$h_{N,a}(\theta) := \max_{\substack{1\le x,y\le N-1 \\ xy\equiv a \pmod N}} (x\cos\theta + y\sin\theta).$$

This quantity automatically selects points on the convex hull; interior points never maximize a supporting linear functional.

**Example 6.1** (A support direction for $A_{7,2}$)**.** The residue class

$$A_{7,2} = \{(1,2),(2,1),(3,3),(4,4),(5,6),(6,5)\}.$$

At the diagonal direction $\theta = \pi/4$, shown in @fig:support-direction-n7-a2, one has

$$x\cos\theta + y\sin\theta = \frac{x+y}{\sqrt{2}},$$

so maximizing the support function is the same as maximizing $x+y$. Among the six points above, the largest value is

$$x + y = 11,$$

attained at $(5,6)$ and $(6,5)$. Therefore

$$h_{7,2}(\pi/4) = \frac{11}{\sqrt{2}}.$$

This simple calculation shows what the support function is really tracking: as the direction rotates, different hull points take turns maximizing the same linear functional.

*Remark* 6.1 (A geometric encoding of a finite arithmetic problem). The support function is not a new object replacing $A_{N,a}$; it is an exact Euclidean encoding of the same finite congruence class. The difficulty of the problem remains arithmetic: one must still understand which lattice points become extremal.

## 6.2 Support-function area formula

Once the support function is known in every direction, it encodes the whole convex polygon. A standard fact from convex geometry then turns that directional data into area.

A standard fact from planar convex geometry is the identity

$$\mathrm{Area}(K) = \frac{1}{2}\int_0^{2\pi} \left(h_K(\theta)^2 - h_K'(\theta)^2\right) d\theta,$$

valid for polygons and more generally for sufficiently regular convex bodies.

Applying this to $K = \mathrm{conv}(A_{N,a})$ yields the exact formula

$$S(N,a) = \frac{1}{2}\int_0^{2\pi} \left(h_{N,a}(\theta)^2 - h_{N,a}'(\theta)^2\right) d\theta. \tag{6.1}$$

On any interval of angles where the same vertex $(x_0, y_0)$ remains the maximizer, the support function is simply

$$h_{N,a}(\theta) = x_0 \cos\theta + y_0 \sin\theta,$$

so on that interval

$$h_{N,a}'(\theta) = -x_0 \sin\theta + y_0 \cos\theta.$$

Thus the derivative records how the same supporting point is seen from nearby directions. The real difficulty in formula (6.1) lies in the switching: one must determine exactly when the maximizing lattice point changes as $\theta$ rotates.

*Remark* 6.2 (What formula (6.1) does and does not solve). Formula (6.1) is exact, but it does not yet provide a closed arithmetic formula in $N$ and $a$. The difficult part is hidden in the maximization defining $h_{N,a}(\theta)$: one still needs to know which lattice point wins in each direction $\theta$.

### 6.3 An exponential form

Using the exponential indicator, one may write

$$F_{N,a}(\theta,\lambda) := \sum_{x=1}^{N-1}\sum_{y=1}^{N-1} \mathbf{1}_{xy\equiv a \pmod N}\, e^{\lambda(x\cos\theta + y\sin\theta)}.$$

Then

$$F_{N,a}(\theta,\lambda) = \frac{1}{N}\sum_{r=0}^{N-1} e^{-2\pi i r a/N} \sum_{x=1}^{N-1}\sum_{y=1}^{N-1} \exp\left(\frac{2\pi i r x y}{N} + \lambda(x\cos\theta + y\sin\theta)\right).$$

By the log-sum-exp identity,

$$h_{N,a}(\theta) = \lim_{\lambda\to\infty} \frac{1}{\lambda}\log F_{N,a}(\theta,\lambda).$$

Substituting this into (6.1) yields a fully analytic representation of $S(N,a)$ in terms of exponentials. This is exact, but it is still implicit. It should be read as an exact filter for finite arithmetic data, not as a replacement for the arithmetic or geometric viewpoints developed in the surrounding chapters.

# 7 The first-boundary model

## 7.1 Definition

The *first boundary*, or outer frame, of the table is the union of the first row, last row, first column, and last column. Define

$$B^{(1)}_{N,a} := A_{N,a} \cap \big((\{1, N-1\} \times \{1, \dots, N-1\}) \cup (\{1, \dots, N-1\} \times \{1, N-1\})\big).$$

Set

$$S^{(1)}(N,a) := \mathsf{Area}(\mathsf{conv}(B^{(1)}_{N,a})), \qquad S^{(1)}(N) := \sum_{a=0}^{N-1} S^{(1)}(N,a).$$

By construction,

$$B^{(1)}_{N,a} \subseteq A_{N,a}, \qquad S^{(1)}(N,a) \le S(N,a), \qquad S^{(1)}(N) \le S(N).$$

Thus $S^{(1)}$ is an exact lower-bound model.

From the layer-by-layer drawing viewpoint of Chapter 5, this is the residue-wise version of the first ring of the table. One first draws the full outer frame of MTMN; then $B^{(1)}_{N,a}$ keeps only those outer-frame cells whose entry is the residue $a$. The key point is that the model is not an arbitrary truncation. It is the first geometric approximation already forced on us by the border-by-border construction of the table itself.

> *Remark* 7.1 (Historical origin of the boundary program). Long before the present notation was fixed, the border-only quantity was the first tractable surrogate for the full hull area. In that earlier form it already suggested studying the reciprocal series $\sum 1/S(N)$ through simpler outer-layer models. The current first-boundary formalism keeps that historical intuition but turns it into an exact residue-wise construction.

## 7.2 Exact shape of the first boundary

**Theorem 7.1** (Exact first-boundary formula). *For $N \geq 2$,*

$$S^{(1)}(N, 0) = 0,$$

*and for $1 \leq a \leq N - 1$,*

$$S^{(1)}(N, a) = 2(a - 1)(N - a - 1).$$

*Equivalently, the four boundary points are*

$$(1, a), \quad (a, 1), \quad (N - 1, N - a), \quad (N - a, N - 1),$$

*and their convex hull is a parallelogram, degenerate (collapsing to a point or a line segment, so it has no area) when $a = 1$ or $a = N - 1$.*

*Proof.* On the first row $x = 1$, the congruence $xy \equiv a \pmod{N}$ becomes $y \equiv a \pmod{N}$, so the boundary point is $(1, a)$. On the first column the point is $(a, 1)$. On the last row $x = N - 1 \equiv -1 \pmod{N}$, so $-y \equiv a \pmod{N}$, hence $y \equiv N - a \pmod{N}$, giving $(N - 1, N - a)$. Similarly one gets $(N - a, N - 1)$ on the last column.

The oriented area of the parallelogram spanned by the vectors

$$(a, 1) - (1, a) = (a - 1, 1 - a)$$

and

$$(N - 1, N - a) - (1, a) = (N - 2, N - 2a)$$

is

$$\left| \det \begin{pmatrix} a - 1 & N - 2 \\ 1 - a & N - 2a \end{pmatrix} \right| = 2(a - 1)(N - a - 1),$$

which is therefore the hull area. ∎

The identity $S^{(1)}(N, 0) = 0$ should not be confused with the full zero-class geometry. Chapter 5 shows that for every composite modulus $N > 4$, the full zero-class hull has positive

area. What vanishes here is only the contribution from the outer frame: the zero-divisor geometry of $A_{N,0}$ is an interior phenomenon rather than a first-boundary one.

Figure 33 illustrates the four boundary points and their parallelogram for a specific example.

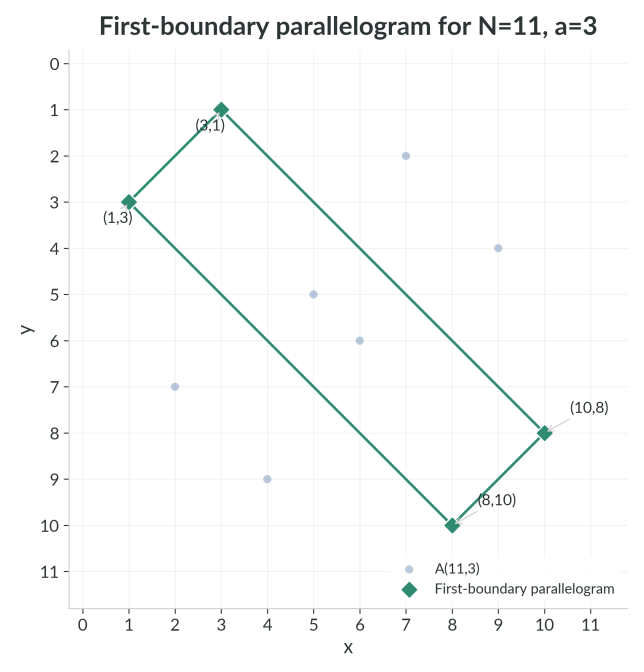


Figure 33: First-boundary parallelogram for $N = 11$, $a = 3$. The four points $(1, 3)$, $(3, 1)$, $(10, 8)$, $(8, 10)$ form a parallelogram. The faded blue dots are the remaining interior lattice points of $A_{11,3}$.

## 7.3 Total first-boundary sum

**Theorem 7.2** (Exact total first-boundary sum). *For every* $N \geq 2$,

$$S^{(1)}(N) = \frac{(N-3)(N-2)(N-1)}{3}.$$

*Proof.* Using the exact first-boundary formula above,

$$S^{(1)}(N) = \sum_{a=1}^{N-1} 2(a-1)(N-a-1).$$

Set $b = a - 1$, so $b = 0, 1, \ldots, N-2$. Then

$$S^{(1)}(N) = 2 \sum_{b=0}^{N-2} b\big((N-2) - b\big).$$

Let $m = N - 2$. Using

$$\sum_{b=0}^{m} b = \frac{m(m+1)}{2}, \qquad \sum_{b=0}^{m} b^2 = \frac{m(m+1)(2m+1)}{6},$$

we obtain

$$S^{(1)}(N) = 2\left(m \cdot \frac{m(m+1)}{2} - \frac{m(m+1)(2m+1)}{6}\right) = \frac{m(m+1)(m-1)}{3},$$

which is exactly

$$\frac{(N-3)(N-2)(N-1)}{3}.$$

■

## 7.4 The cubic scale of the total area

Recall that $S(N, a)$ is the convex-hull area of the residue class $A_{N,a}$, and that

$$S(N) = \sum_{a=0}^{N-1} S(N, a).$$

The exact first-boundary formula already forces a genuine global growth law for this total. The point is not only that one gets matching exponents after a squeeze. The statement below says that the accumulated hull area of all residue classes cannot collapse below cubic order, even though individual classes may be sparse, degenerate, or arithmetically irregular.

**Theorem 7.3** (Cubic-order bounds for the total area)**.** *For every* $N \geq 2$*,*

$$\frac{(N-3)(N-2)(N-1)}{3} \leq S(N) \leq N(N-2)^2.$$

*In particular there exist positive constants* $c_1$ *and* $c_2$ *such that*

$$c_1 N^3 \leq S(N) \leq c_2 N^3$$

*for all sufficiently large* $N$*. Equivalently, in the usual asymptotic notation,*

$$S(N) = \Theta(N^3).$$

*Proof.* The lower bound is immediate from the exact lower-model inequality

$$S^{(1)}(N) \leq S(N)$$

together with the total first-boundary formula just proved. For the upper bound, fix a residue $a$. The whole set $A_{N,a}$ lies in the square

$$[1, N-1] \times [1, N-1],$$

whose area is $(N-2)^2$. Hence conv$(A_{N,a})$ also lies in that square, so

$$S(N, a) \leq (N-2)^2.$$

Summing over the $N$ residue classes gives

$$S(N) \leq N(N-2)^2.$$

■

This makes $N^3$ the right first scale for the total area. There are $N$ residue classes, each class lives in a box of quadratic area scale, and summing such contributions naturally suggests a cubic total. But that is only a scale heuristic, not a proof. A large point set can still have a thin or degenerate hull, arithmetic collisions can force strong nonuniformity, and

zero-divisor structure can distort the geometry from one residue to the next. The theorem matters precisely because it shows that these local irregularities do not destroy cubic growth in the aggregate.

One conceptual clarification is worth making explicit. All of the hulls conv$(A_{N,a})$ lie in the same ambient square, but $S(N)$ is the sum of their areas, not the area of their union. So there is no contradiction between a quadratic Euclidean window and a cubic total: the same square is being counted once for each residue label.

### 7.5 The sharp cubic question

Once the exponent is fixed, the next question is the leading constant. The exact computations now available point consistently to the sharper asymptotic

$$S(N) \sim N^3.$$

Equivalently, one asks whether

$$\frac{S(N)}{N^3} \to 1$$

as $N \to \infty$.

**Conjecture (sharp cubic asymptotic).** As $N \to \infty$,

$$S(N) \sim N^3.$$

Representative exact values are:

| $N$ | $S(N)/N^3$ |
|---|---|
| 1000 | 0.97073535 |
| 2000 | 0.98325402 |
| 3000 | 0.98725631 |
| 5000 | 0.99214422 |

These finite computations do not prove the conjecture, but they support it strongly. A

related diagnostic points in the same direction: because

$$S^{(1)}(N) = \frac{(N-3)(N-2)(N-1)}{3} \sim \frac{N^3}{3},$$

the conjecture $S(N) \sim N^3$ would imply

$$\frac{S(N)}{S^{(1)}(N)} \to 3.$$

The exact data is consistent with that drift as well.

The first-boundary model is especially revealing in this light. One explicit geometric layer already contributes a full cubic amount:

$$S^{(1)}(N) \sim \frac{N^3}{3}.$$

So if the conjecture $S(N) \sim N^3$ is true, then the outer frame is not the whole story but the first explicit third of a much fuller cubic phenomenon.

## 7.6 The deficiency

Once $N^3$ is recognized as the right main scale, the natural second-order object is not $S(N)$ itself but its missing area:

$$D(N) := N^3 - S(N).$$

The sharp cubic conjecture is equivalent to the statement

$$D(N) = o(N^3).$$

that is,

$$\frac{D(N)}{N^3} \to 0.$$

In that form the problem changes character. One no longer asks whether the total area grows cubically; that is already settled. One asks how much of the cubic budget is lost, why it is lost, and how that loss is distributed across residue classes.

Even the crude box cap shows that some deficit is unavoidable. Since

$$S(N) \le N(N-2)^2 = N^3 - 4N^2 + 4N,$$

one always has

$$D(N) \ge 4N^2 - 4N.$$

So the deficiency is not expected to vanish. The real question is whether it is smaller than the cubic main term, and how sharply one can describe it. This is where the subtlety of the problem seems to live: in which points become actual hull vertices, which boundary layers dominate, and how the missing area aggregates across the residue families.

## 7.7 Global series consequences

The same asymptotic picture also organizes the two most natural global series formed from $S(N)$. The reciprocal series is the main one. A weighted companion reflects the same large-$N$ geometry through a slower-decaying lens, but it remains secondary to the area problem itself.

### 7.7.1 The reciprocal series

Once the total area sum $S(N)$ has been introduced, it is natural to ask whether varying the modulus leaves behind any single global quantity. The simplest one is obtained by taking reciprocals and summing over all moduli:

$$C := \sum_{N=4}^{\infty} \frac{1}{S(N)}.$$

This series gives the greatest weight to small moduli, where the total area is smallest, and discounts the larger moduli according to how quickly $S(N)$ grows. So two immediate questions arise: does the series converge, and if it does, can one say anything exact about its value?

The first-boundary model already answers the analogous question completely, because $S^{(1)}(N)$ is an explicit cubic polynomial. That exact model then becomes a comparison tool for the full series.

**Theorem 7.4** (Exact evaluation of the first-boundary reciprocal series)**.** *The series*

$$\sum_{N=4}^{\infty} \frac{1}{S^{(1)}(N)}$$

*converges and has exact value*

$$\sum_{N=4}^{\infty} \frac{1}{S^{(1)}(N)} = \frac{3}{4}.$$

*Proof.* By the total first-boundary formula just proved,

$$\frac{1}{S^{(1)}(N)} = \frac{3}{(N-3)(N-2)(N-1)}.$$

Let $k = N - 3$. Then $k \geq 1$ and

$$\sum_{N=4}^{\infty} \frac{1}{S^{(1)}(N)} = \sum_{k=1}^{\infty} \frac{3}{k(k+1)(k+2)}.$$

Using the telescoping identity

$$\frac{1}{k(k+1)(k+2)} = \frac{1}{2}\left(\frac{1}{k(k+1)} - \frac{1}{(k+1)(k+2)}\right),$$

we get

$$\sum_{k=1}^{M} \frac{3}{k(k+1)(k+2)} = \frac{3}{2}\sum_{k=1}^{M}\left(\frac{1}{k(k+1)} - \frac{1}{(k+1)(k+2)}\right) = \frac{3}{2}\left(\frac{1}{2} - \frac{1}{(M+1)(M+2)}\right).$$

Letting $M \to \infty$ gives $3/4$. ∎

**Corollary 7.5** (Convergence of the full reciprocal series)**.** *The series defining*

$$C = \sum_{N=4}^{\infty} \frac{1}{S(N)}$$

*converges.*

*Proof.* By construction $S(N) \geq S^{(1)}(N)$ for every $N$, so the total first-boundary formula gives

$$0 \leq \frac{1}{S(N)} \leq \frac{1}{S^{(1)}(N)} = \frac{3}{(N-3)(N-2)(N-1)}.$$

The right-hand side is summable by the theorem above, so the full series converges by comparison. ■

**Proposition 7.6** (Exact first-boundary tail formula). *For every integer $M \geq 4$,*

$$\sum_{N=M+1}^{\infty} \frac{1}{S^{(1)}(N)} = \frac{3}{2(M-2)(M-1)}.$$

*Proof.* The proof of the theorem above already gives the exact partial-sum identity

$$\sum_{N=4}^{M} \frac{1}{S^{(1)}(N)} = \frac{3}{4} - \frac{3}{2(M-2)(M-1)}.$$

Subtracting this from the total value $3/4$ yields the stated tail formula. ■

The point of the tail formula is that it immediately turns any finite exact computation of the values $S(N)$ into a rigorous interval for the full constant $C$.

**Corollary 7.7** (Rigorous enclosure from finitely many exact values). *For every integer $M \geq 4$, let*

$$P_M := \sum_{N=4}^{M} \frac{1}{S(N)}.$$

*Then*

$$P_M \leq C \leq P_M + \frac{3}{2(M-2)(M-1)}.$$

*Proof.* Since all terms are nonnegative,

$$0 \leq C - P_M = \sum_{N=M+1}^{\infty} \frac{1}{S(N)}.$$

Also $S(N) \geq S^{(1)}(N)$ for every $N$, so

$$\sum_{N=M+1}^{\infty} \frac{1}{S(N)} \leq \sum_{N=M+1}^{\infty} \frac{1}{S^{(1)}(N)}.$$

Apply the tail formula above. ■

**Proposition 7.8** (A rigorous enclosure for the reciprocal constant)**.** *An exact integer computation gives*

$$\sum_{N=4}^{2000} \frac{1}{S(N)} = 0.648623982413541356.$$

*Consequently*

$$0.648623982413541356 \leq C \leq 0.648624357976698263.$$

*Proof.* The displayed partial sum is the output of an exact integer computation of the finitely many values $S(N)$ for $4 \leq N \leq 2000$. Apply the preceding corollary with $M = 2000$. The tail contribution is bounded by

$$\frac{3}{2(1998)(1999)} = 0.000000375563156954,$$

which gives the stated interval after addition. ■

The width of this interval is below $4 \times 10^{-7}$. So the reciprocal constant is already sharply localized numerically, even though no symbolic formula is known.

*Remark* 7.2 (A historical numerical guess)*.* Earlier stages of the project were guided for a time by the speculation

$$C = \sqrt{e} - 1.$$

The enclosure above rules this out, because

$$\sqrt{e} - 1 = 0.648721270700\ldots > 0.648624357976698263 \geq C.$$

So the old guess cannot be correct.

From the asymptotic viewpoint developed above, the significance of the reciprocal series is clear. If the sharp cubic conjecture is correct, then

$$\frac{1}{S(N)} \sim \frac{1}{N^3},$$

and therefore

$$\sum_{N=M+1}^{\infty} \frac{1}{S(N)} \sim \frac{1}{2M^2}.$$

That would explain conceptually why the reciprocal constant converges so cleanly: the first-boundary tail bound already has the right decay exponent, even if it is not expected to give the sharp leading constant for the full series.

What remains unknown is the symbolic nature of $C$. Computation together with the exact tail control above gives a rigorous enclosure, but it does not explain why the constant takes the value it does. To go further one would need sharper structural control of the total areas themselves: which lattice points become true hull vertices in general, which exact-product layers dominate the boundary globally, and above all how the deficiency $N^3 - S(N)$ is built up across the residue classes. In that sense the reciprocal series is a natural global observable attached to $S(N)$: the present theory already controls it nontrivially, but its exact symbolic value remains open.

### 7.7.2 A weighted companion

The same large-$N$ picture has a natural weighted companion:

$$W := \sum_{N=4}^{\infty} \frac{N}{S(N)}.$$

Because the extra factor $N$ removes one power of the cubic growth, this series is more sensitive to the sharp asymptotic while still remaining summable.

**Theorem 7.9** (Exact evaluation of the weighted first-boundary series)**.** *The series*

$$\sum_{N=4}^{\infty} \frac{N}{S^{(1)}(N)}$$

*converges and has exact value*

$$\sum_{N=4}^{\infty} \frac{N}{S^{(1)}(N)} = \frac{15}{4}.$$

*Proof.* By the total first-boundary formula,

$$\frac{N}{S^{(1)}(N)} = \frac{3N}{(N-3)(N-2)(N-1)}.$$

Let $k = N - 3$. Then $k \geq 1$ and

$$\sum_{N=4}^{\infty} \frac{N}{S^{(1)}(N)} = \sum_{k=1}^{\infty} \frac{3(k+3)}{k(k+1)(k+2)}.$$

Partial fractions give

$$\frac{3(k+3)}{k(k+1)(k+2)} = \frac{9}{2k} - \frac{6}{k+1} + \frac{3}{2(k+2)}.$$

Therefore

$$\sum_{k=1}^{M} \frac{3(k+3)}{k(k+1)(k+2)} = \frac{9}{2}H_M - 6H_{M+1} + 6 + \frac{3}{2}H_{M+2} - \frac{9}{4},$$

where $H_M$ is the $M$-th harmonic number. Since

$$H_{M+1} = H_M + \frac{1}{M+1}, \qquad H_{M+2} = H_M + \frac{1}{M+1} + \frac{1}{M+2},$$

the harmonic terms cancel and one gets

$$\sum_{k=1}^{M} \frac{3(k+3)}{k(k+1)(k+2)} = \frac{15}{4} - \frac{9}{2(M+1)} + \frac{3}{2(M+2)}.$$

Letting $M \to \infty$ gives $15/4$. ∎

**Corollary 7.10** (Convergence of the weighted full series). *Because* $S(N) \geq S^{(1)}(N)$

*for all $N$, one has*

$$0 \leq \frac{N}{S(N)} \leq \frac{N}{S^{(1)}(N)}.$$

*Hence the series*

$$\sum_{N=4}^{\infty} \frac{N}{S(N)}$$

*converges by comparison with the preceding theorem.*

If the conjectural asymptotic $S(N) \sim N^3$ is true, then

$$\frac{N}{S(N)} \sim \frac{1}{N^2},$$

so the weighted tail should satisfy

$$\sum_{N=M+1}^{\infty} \frac{N}{S(N)} \sim \frac{1}{M}.$$

Thus the weighted series is useful as a secondary diagnostic of the same cubic picture. Exact computation together with the first-boundary comparison currently places its full constant in a narrow rigorous interval just above $3.001$, but the point here is not the decimal expansion itself. The main point is that the weighted series, like the reciprocal series, reflects the same unresolved asymptotic geometry of $S(N)$.

# 8 The second-boundary model

## 8.1 Definition

The second boundary uses only rows and columns $2$ and $N-2$:

$$B^{(2)}_{N,a} := A_{N,a} \cap \big((\{2, N-2\} \times \{1, \dots, N-1\}) \cup (\{1, \dots, N-1\} \times \{2, N-2\})\big).$$

Set

$$S^{(2)}(N,a) := \mathsf{Area}(\mathsf{conv}(B^{(2)}_{N,a})), \qquad S^{(2)}(N) := \sum_{a=0}^{N-1} S^{(2)}(N,a).$$

Again one has the lower bound

$$S^{(2)}(N,a) \le S(N,a), \qquad S^{(2)}(N) \le S(N).$$

From the construction viewpoint of Chapter 5, this is the residue-wise version of the second ring of the table. After the outer frame has been drawn, the next boundary layer comes from rows and columns $2$ and $N-2$; $B^{(2)}_{N,a}$ records only those second-layer cells whose entry is $a$.

In plainer language: first ignore the outer frame, then look only at the next ring in from the edge, and finally keep only the cells whose residue is $a$. The notation $B^{(2)}_{N,a}$ is just a compact way to say "the second ring of the residue class $A_{N,a}$."

Figure 34 shows that second ring before any residue selection is made.

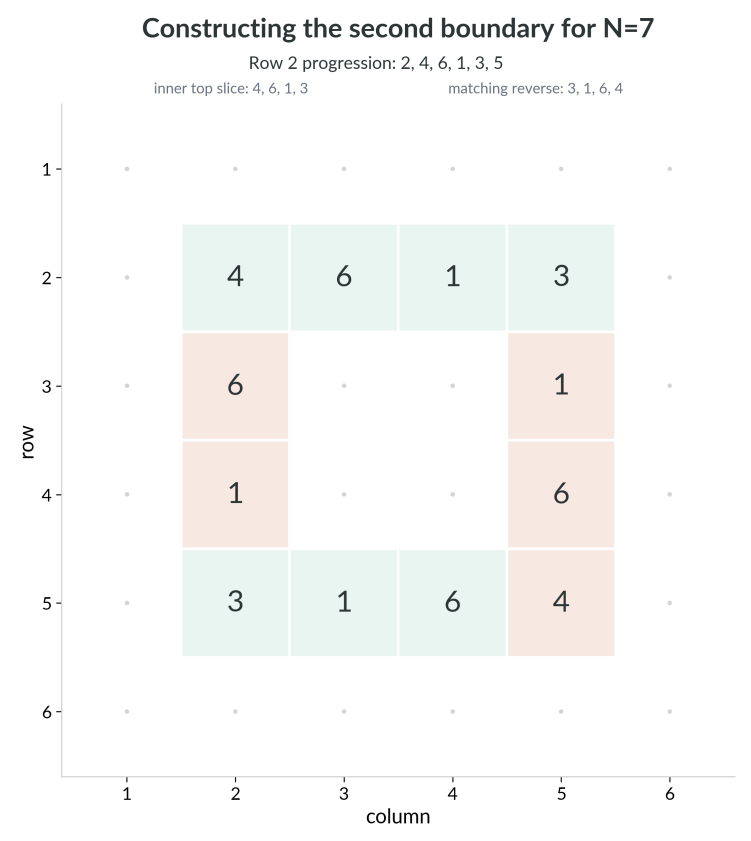


Figure 34: The second boundary is literally the second ring of the MTMN square. For $N = 7$, the row-$2$ progression fills that entire inner ring once the forward and reverse orderings are wrapped around the four sides.

## 8.2 Exact formula for odd $N$

When $N$ is odd, the residue $2$ is invertible modulo $N$. So for each nonzero residue $a$, there is exactly one residue $b$ with

$$2b \equiv a \pmod{N}.$$

The symbol $b$ is simply a convenient name for that unique hit on the second boundary. We now write it as

$$b \equiv 2^{-1}a \pmod{N}, \qquad 1 \le b \le N-1.$$

In other words, the theorem below says: once that single number $b$ is known, the four second-boundary points are immediately determined, and so is the area of their hull.

**Theorem 8.1** (Second-boundary formula for odd $N$)**.** *Assume $N$ is odd. Then*

$$S^{(2)}(N, 0) = 0,$$

*and for $1 \le a \le N-1$,*

$$S^{(2)}(N, a) = 2\,|b-2|\,|N-b-2|, \qquad b \equiv 2^{-1}a \pmod{N}.$$

*The corresponding boundary points are*

$$(2, b), \quad (b, 2), \quad (N-2, N-b), \quad (N-b, N-2).$$

*Proof.* If $x = 2$, the congruence becomes $2y \equiv a \pmod{N}$, so $y \equiv b \pmod{N}$. Since $N$ is odd, $2$ has a unique inverse modulo $N$, hence the point is $(2, b)$. Symmetrically one gets $(b, 2)$. On the rows and columns $N-2 \equiv -2 \pmod{N}$, the congruence becomes $-2y \equiv a \pmod{N}$, hence $y \equiv N-b \pmod{N}$, producing the other two points.

Now apply the linear change of coordinates

$$u = x + y, \qquad v = x - y.$$

Its Jacobian determinant has absolute value $2$, so area in $(x, y)$-space equals one half the area in $(u, v)$-space. The four points become the corners of the rectangle with

$$u \in \{b+2,\ 2N-b-2\}, \qquad v \in \{\pm(b-2)\}.$$

Therefore the rectangle in $(u, v)$-space has side lengths

$$2|N-b-2|, \qquad 2|b-2|.$$

Dividing by $2$ gives the area in $(x, y)$-space:

$$S^{(2)}(N, a) = \frac{1}{2} \cdot 2|N - b - 2| \cdot 2|b - 2| = 2|b - 2|\,|N - b - 2|.$$

■

*Remark* 8.1 (Why the coordinates $u = x + y$, $v = x - y$ are natural). The linear forms

$$u = x + y, \qquad v = x - y$$

already appear in the modular-hyperbola literature as coordinate sums and coordinate differences. Bower, Evans, Luo, and Miller study the cardinalities of the corresponding sumsets and difference sets for reduced modular hyperbolas, and from their geometric viewpoint $\#S_2(a; n)$ and $\#D_2(a; n)$ count how many lines of slope $-1$ and $1$ meet the class [@bowerevansluomiller2012]. The present chapter uses exactly the same linear forms in a more rigid way: for the second boundary they become literal rectangle coordinates, so the side lengths in $(u, v)$-space give an exact area formula. Their paper counts diagonal hits for a full reduced class; here the same coordinates produce an exact Euclidean hull for one boundary layer.

For the zero residue, this theorem gives $S^{(2)}(N, 0) = 0$ for every odd $N$. That does not mean the full zero class is degenerate. Chapter 5 shows that for every odd composite modulus $N > 4$, the full hull of $A_{N,0}$ still has positive area. The second boundary misses that geometry because the divisor-driven zero class lives deeper in the table than the single layer $x, y \in \{2, N - 2\}$.

**Theorem 8.2** (Total second-boundary sum for odd $N$). *If $N$ is odd, then*

$$S^{(2)}(N) = \frac{(N - 3)(N^2 - 9N + 32)}{3}.$$

*Proof.* As $a$ runs over $1, \dots, N-1$, so does $b \equiv 2^{-1}a \pmod{N}$. Hence

$$S^{(2)}(N) = \sum_{b=1}^{N-1} 2\,|b-2|\,|N-b-2|.$$

For odd $N$, the absolute values simplify piecewise:

$$S^{(2)}(N) = 2(N-3) + 2(N-3) + \sum_{b=3}^{N-3} 2(b-2)(N-b-2).$$

Evaluating the quadratic sum gives

$$S^{(2)}(N) = \frac{(N-3)(N^2-9N+32)}{3}.$$

A direct expansion verifies the identity. ∎

*Remark* 8.2 (Even $N$). When $N$ is even, $2$ is not invertible modulo $N$, so the second boundary no longer admits the same one-parameter description. Three new phenomena appear immediately:
1. some residues give no second-boundary points at all; 2. even residues may contribute multiple points on a single side; 3. the clean indexing by $b \equiv 2^{-1}a \pmod{N}$ disappears.
So the even case is not merely an odd case with extra bookkeeping. It is the first place where the boundary-layer program visibly feels the arithmetic failure of invertibility.

*Remark* 8.3 (Series built from the second-boundary totals). For odd $N$, the explicit cubic formula for $S^{(2)}(N)$ makes reciprocal and weighted series

$$\sum_{\substack{N\geq 5\\ N \text{ odd}}} \frac{1}{S^{(2)}(N)}, \qquad \sum_{\substack{N\geq 5\\ N \text{ odd}}} \frac{N}{S^{(2)}(N)}$$

natural next targets in the broader series-and-constants program. Their convergence is immediate from the cubic growth, but unlike the first-boundary case no simple closed form is currently known.

## 8.3 Geometric comparison

Figure 35 packages the whole boundary program for one residue class: select the first two layers inside $A_{N,a}$, identify the first-boundary hull in $(x, y)$, pass to rectangles in $(u, v)$, and then take the hull of the first two layers together.

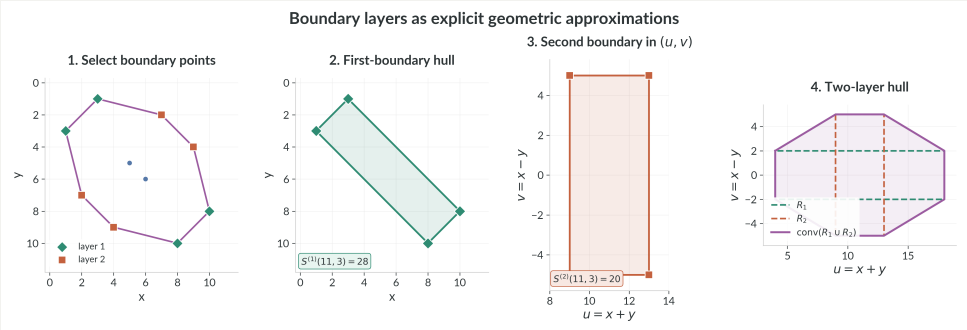


Figure 35: A one-page view of the first two boundary models. Starting from the full residue class, one selects the first two boundary layers, reads the second layer as a rectangle in $(u, v)$-coordinates, and then takes the hull of the two-layer union.

Figure 36 compares the full residue set with the first- and second-boundary models for one sample residue class.

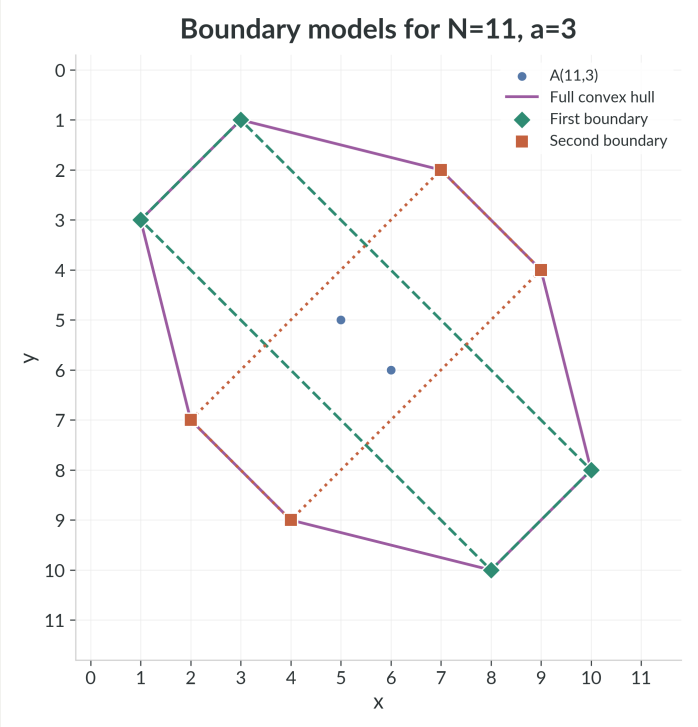


Figure 36: For $N = 11$ and $a = 3$, the boundary models provide explicit inner geometric approximations to the full convex hull. The dashed green parallelogram is the first-boundary hull; the dotted orange parallelogram is the second-boundary hull; the solid red curve is the full convex hull.

As a further illustration, Figure 37 shows a single residue class for $N = 7$, $a = 3$ with both its hull and boundary points marked.

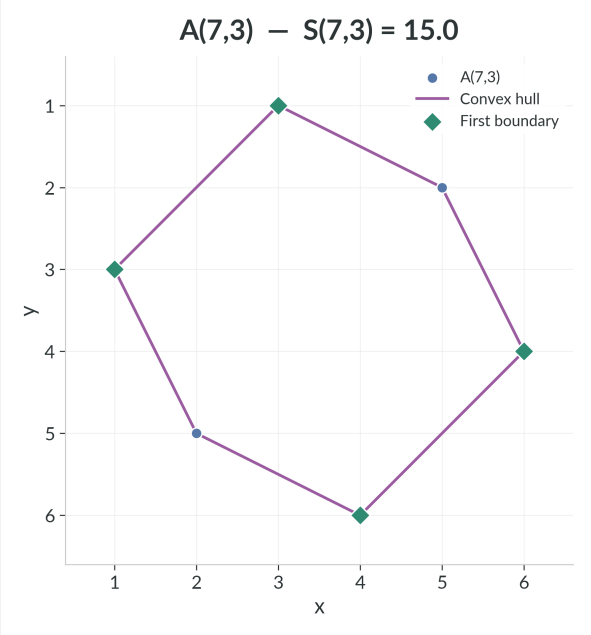


Figure 37: Residue class $A_{7,3}$ with convex hull and first-boundary points.

# 9 Using the first two layers together

## 9.1 Definition

One may also take the union of the first two layers:

$$B_{N,a}^{(\le 2)} := A_{N,a} \cap \big( (\{1,2,N-2,N-1\} \times \{1,\dots,N-1\}) \cup (\{1,\dots,N-1\} \times \{1,2,N-2,N-1\}) \big).$$

Define

$$S^{(\le 2)}(N,a) := \mathrm{Area}(\mathrm{conv}(B_{N,a}^{(\le 2)})), \qquad S^{(\le 2)}(N) := \sum_{a=0}^{N-1} S^{(\le 2)}(N,a).$$

In plainer language, $B_{N,a}^{(\le 2)}$ means: keep only the outer frame and the next ring in, then retain only the cells whose residue is $a$. So this two-layer model is the first place where the residue class is allowed to use more than one boundary layer at once.

Like the one-layer models, this two-layer quantity is a natural source not only of geometric approximations but also of new sequences, totals, and later series questions.

## 9.2 Rectangles after the transformation $u = x + y$, $v = x - y$

The change of coordinates

$$u = x + y, \qquad v = x - y$$

is introduced for one simple reason: in the original $(x, y)$-plane, the boundary hulls are slanted parallelograms, while in the $(u, v)$-plane they become ordinary axis-aligned rectangles. That makes it much easier to see how the first two layers interact.

For odd $N$, the first boundary and the second boundary each become axis-aligned rectangles in $(u, v)$-space. The first boundary has half-widths

$$\alpha_1 = N - a - 1, \qquad \beta_1 = a - 1,$$

and the second boundary has half-widths

$$\alpha_2 = |N - b - 2|, \qquad \beta_2 = |b - 2|, \qquad b \equiv 2^{-1}a \pmod{N}.$$

Thus the combined two-layer hull is the convex hull of two centered rectangles. This reduces the problem to elementary convex geometry.

This is also the point where the coordinate-sum and coordinate-difference literature becomes visibly relevant. Bower, Evans, Luo, and Miller study the arithmetic sizes of the sets of values taken by $x+y$ and $x-y$ on reduced modular hyperbolas [14]. In the present chapter we do not count those values. We fix one residue class and use the same linear forms as actual coordinates, so that the first two MTMN layers become centered rectangles and their interaction becomes an exact convex-hull calculation.

Figure 38 and Figure 39 show the first and second boundary sets for $N = 11$, $a = 3$, in both coordinate systems.

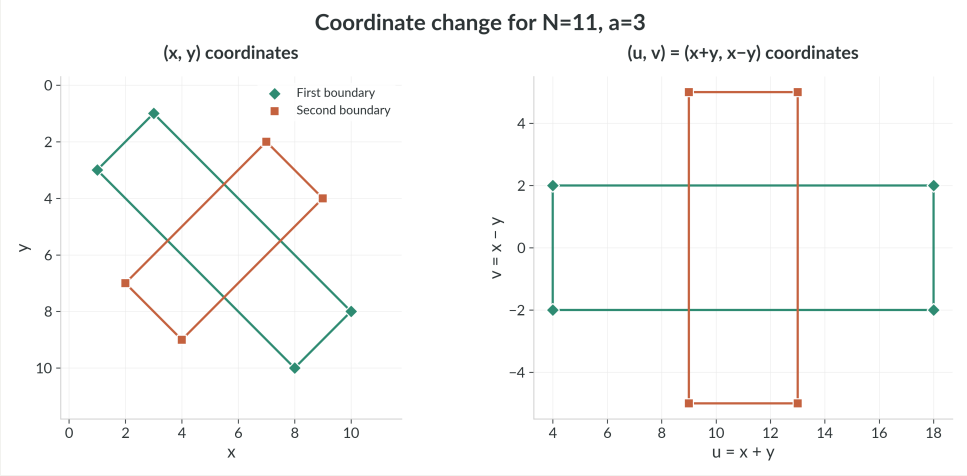


Figure 38: The first and second boundary sets for $N = 11$, $a = 3$, in the original $(x, y)$ coordinates (left) and after the linear change of variables $(u, v) = (x + y, x - y)$ (right). In $(u, v)$-space, each boundary layer becomes an axis-aligned rectangle.

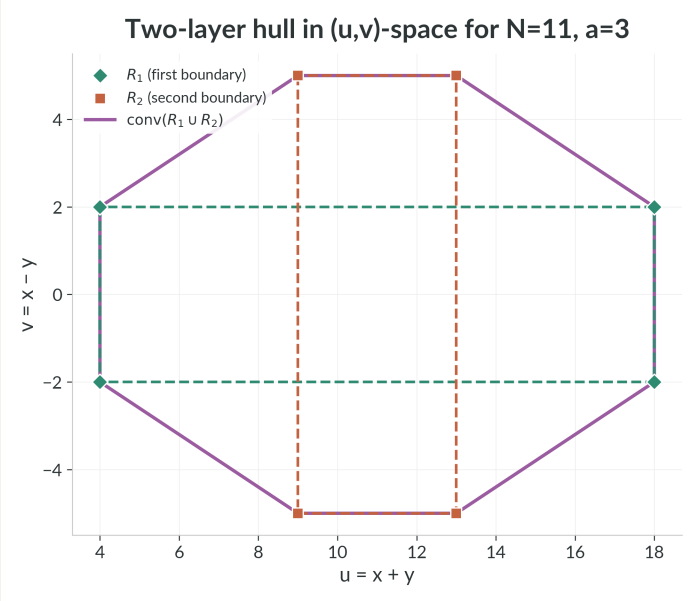


Figure 39: The two rectangles $R_1$ and $R_2$ in $(u, v)$-space for $N = 11$, $a = 3$, together with the convex hull of their union. This hull is the combined two-layer model.

**Proposition 9.1** (Piecewise formula for the two-layer odd-$N$ model). *Assume $N$ is odd. Let $R_1$ and $R_2$ be the centered rectangles in $(u, v)$-space with half-widths $(\alpha_1, \beta_1)$ and $(\alpha_2, \beta_2)$ as above. Then*

$$S^{(\le 2)}(N, a) = \frac{1}{2}\,\mathrm{Area}(\mathrm{conv}(R_1 \cup R_2)).$$

*In words: once one passes to $(u, v)$-coordinates, the two-layer model is nothing more than the convex hull of two ordinary rectangles, followed by the factor $1/2$ that converts area back to the original coordinates.*

*In particular:*

*(a) if one rectangle contains the other, then $S^{(\le 2)}(N, a)$ is just the larger of $S^{(1)}(N, a)$ and $S^{(2)}(N, a)$;*

*(b) if the rectangles cross, say $\alpha_1 \geq \alpha_2$ and $\beta_2 \geq \beta_1$, then*

$$S^{(\leq 2)}(N, a) = 2\alpha_1\beta_1 + (\alpha_1 + \alpha_2)(\beta_2 - \beta_1).$$

*Idea of proof.* In the first quadrant of $(u, v)$-space, the convex hull of two centered axis-aligned rectangles is either a rectangle or a pentagon with vertices

$$(0,0), \quad (\alpha_1, 0), \quad (\alpha_1, \beta_1), \quad (\alpha_2, \beta_2), \quad (0, \beta_2).$$

Its area is elementary to compute. One then uses central symmetry and the Jacobian factor $1/2$ when passing back to $(x, y)$-space. ∎

# 10 Residue-area polynomials

## 10.1 Why package the residue-area profile?

For a fixed modulus $N$, the list

$$\big(S(N,0), S(N,1), \ldots, S(N,N-1)\big)$$

contains more information than the total $S(N)$ alone. The simplest way to keep that whole residue-by-residue profile visible is to package it into a polynomial.

This is not meant to create a new geometry of its own. The point of the package is more modest and more useful: certain evaluations and derivatives combine the coefficients in exact ways that would be harder to see from the raw list. In particular, the zero-class coefficient $S(N,0)$ is no longer just a formal constant term. By Chapter 5, it is the area of a completely described divisor-envelope hull, and that same chapter rewrites it as a hyperbolic baseline minus an exact arithmetic correction. The polynomial package therefore gives several algebraic ways to recover the area of one geometric object that is already understood in detail.

Before writing the formal definition, it helps to read the polynomial as a bookkeeping device. Start with the coefficient list

$$\big(S(N,0), S(N,1), \ldots, S(N,N-1)\big).$$

Then place $S(N,a)$ in front of $x^a$. Nothing geometric is being discarded; the same data is simply being packed into one algebraic object.

Figure 40 summarizes that reading. The polynomial is the middle box, and the later identities come from plugging in special values of $x$ or differentiating.

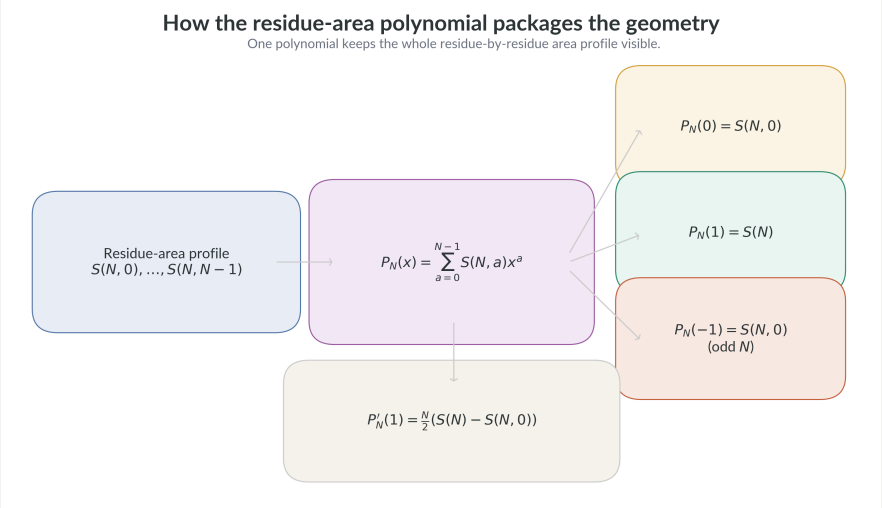


Figure 40: How to read the residue-area polynomial. The coefficient list is packaged into one polynomial, and the most important evaluations and derivatives recover distinguished geometric quantities such as $S(N,0)$ and $S(N)$.

## 10.2 Definition and immediate evaluations

**Definition 10.1** (Residue-area polynomial)**.** For each $N \geq 2$, define

$$P_N(x) := \sum_{a=0}^{N-1} S(N,a)x^a.$$

Equivalently,

$$P_N(x) = S(N,0) + \sum_{a=1}^{N-1} S(N,a)x^a.$$

So the definition should be read very literally: $P_N(x)$ is just the residue-area list written as a polynomial, with the area of the zero class sitting in the constant term.

For the three moduli that will serve as running examples here, one has

$$P_6(x) = 2 + 9x^2 + 8x^3 + 9x^4,$$

$$P_7(x) = 12x + 8x^2 + 15x^3 + 15x^4 + 8x^5 + 12x^6,$$

$$P_9(x) = 9 + 27x + 12x^2 + 32x^3 + 27x^4 + 27x^5 + 32x^6 + 12x^7 + 27x^8.$$

**Proposition 10.1** (Values at $0$ and $1$)**.** *For every $N \geq 2$,*

$$P_N(0) = S(N, 0), \qquad P_N(1) = S(N).$$

*Proof.* Writing

$$P_N(x) = S(N, 0) + \sum_{a=1}^{N-1} S(N, a)x^a$$

makes the constant term explicit, so

$$P_N(0) = S(N, 0).$$

Also

$$P_N(1) = \sum_{a=0}^{N-1} S(N, a) = S(N).$$

■

The first identity should now be read geometrically: $P_N(0)$ is the area of the zero-class divisor hull from Chapter 5, not merely the coefficient indexed by $a = 0$.

## 10.3 Symmetry and reflected coefficients

The basic residue symmetry from Chapter 5 is

$$S(N, a) = S(N, (-a) \bmod N).$$

For $1 \leq a \leq N - 1$, this is the same as

$$S(N, a) = S(N, N - a).$$

**Proposition 10.2** (Reflected coefficient form). *For every* $N \geq 2$,

$$P_N(x) = S(N,0) + \sum_{a=1}^{\lfloor (N-1)/2 \rfloor} S(N,a)\big(x^a + x^{N-a}\big) + \begin{cases} S(N,N/2)x^{N/2}, & N \text{ even}, \\ 0, & N \text{ odd}. \end{cases}$$

*Proof.* Pair the coefficient of $x^a$ with the coefficient of $x^{N-a}$. The symmetry

$$S(N,a) = S(N,N-a)$$

turns the pair into

$$S(N,a)x^a + S(N,N-a)x^{N-a} = S(N,a)\big(x^a + x^{N-a}\big).$$

If $N$ is even, the middle residue $a = N/2$ is unpaired and contributes its own term. ∎

In plain language, the nonzero coefficients come in mirrored pairs: once the constant term $S(N,0)$ is set aside, the coefficient list reads the same from left and right.

*Remark* 10.1 (Not quite a palindromic polynomial). The full polynomial is usually not palindromic, because the constant term $S(N,0)$ does not participate in the nonzero reflection. What *is* reflected is the nonzero coefficient block

$$\big(S(N,1), S(N,2), \ldots, S(N,N-1)\big).$$

For prime moduli, $S(p,0) = 0$, so the constant term disappears and the symmetry becomes especially transparent.

## 10.4 The zero class as a geometric anchor

The derivative at $x = 1$ gives a weighted version of the coefficient sum: larger residue labels count more heavily. The identity below shows exactly how the zero-class area corrects that weighted total.

**Proposition 10.3** (Weighted first moment). *For every $N \geq 2$,*

$$P_N'(1) = \sum_{a=1}^{N-1} a\, S(N,a) = \frac{N}{2}\big(S(N) - S(N,0)\big).$$

*Proof.* Let

$$T := \sum_{a=1}^{N-1} a\, S(N,a).$$

The $a = 0$ term vanishes automatically, so $T = P_N'(1)$. Now substitute $b = N - a$. Using the reflection symmetry on the nonzero residues,

$$T = \sum_{b=1}^{N-1} (N-b) S(N,b) = N \sum_{b=1}^{N-1} S(N,b) - \sum_{b=1}^{N-1} b\, S(N,b).$$

Hence

$$T = N\big(S(N) - S(N,0)\big) - T,$$

so

$$2T = N\big(S(N) - S(N,0)\big).$$

Therefore

$$P_N'(1) = T = \frac{N}{2}\big(S(N) - S(N,0)\big).$$

■

This identity is one of the main reasons to introduce the polynomial viewpoint here. It shows that $S(N,0)$ is not merely one more coefficient and not merely a prime/composite signal. It is the exact correction term between the total area and the first weighted moment of the residue-area profile. Because Chapter 5 gives a full divisor-geometric description of $S(N,0)$, the derivative identity singles out a quantity whose geometry is already completely understood.

For the running examples,

$$P_7'(1) = \frac{7 \cdot 70}{2} = 245,$$

while

$$P_6'(1) = \frac{6(28-2)}{2} = 78.$$

The difference is exactly the zero-class area term.

### 10.5 Odd evaluation at $x = -1$ and a factorization corollary

For odd $N$, evaluating at $x = -1$ turns the polynomial into an alternating sum. The nonzero coefficients then cancel in reflected pairs, leaving only the zero-class term behind.

**Proposition 10.4** (Odd moduli). *For every odd $N \geq 3$,*

$$P_N(-1) = S(N, 0).$$

*Proof.* Pair each nonzero residue $a$ with $N - a$. When $N$ is odd, these two integers have opposite parity, so

$$(-1)^a + (-1)^{N-a} = 0.$$

Using $S(N, a) = S(N, N - a)$, every nonzero pair cancels in the alternating sum. The only surviving term is the constant term $S(N, 0)$. ■

For example,

$$P_9(-1) = 9 - 27 + 12 - 32 + 27 - 27 + 32 - 12 + 27 = 9 = S(9, 0).$$

So in the odd composite case, the evaluation at $x = -1$ recovers the same zero-class area rather than vanishing.

**Corollary 10.5** (Odd factorization criterion). *Assume that $N$ is odd. Then*

$$x \mid P_N(x) \iff S(N, 0) = 0,$$

$$x + 1 \mid P_N(x) \iff S(N, 0) = 0,$$

*and hence*

$$x(x+1) \mid P_N(x) \iff S(N,0) = 0.$$

*Proof.* The factor $x$ divides $P_N(x)$ exactly when the constant term vanishes, so by the values-at-$0$ proposition this is equivalent to $S(N,0) = 0$. For odd $N$, the proposition above gives

$$P_N(-1) = S(N,0),$$

so $x+1$ divides $P_N(x)$ exactly when $S(N,0) = 0$. The combined statement follows immediately. ■

**Corollary 10.6** (Odd prime factorization corollary)**.** *For odd $N > 4$,*

$$x(x+1) \mid P_N(x) \iff N \text{ is prime}.$$

*Proof.* Combine the previous corollary with the sharp zero-class degeneracy criterion from Chapter 5. ■

This is an exact and elegant algebraic consequence of the zero-class geometry, but it should be read in the right scale. It explains the two linear factors forced by the vanishing of the zero-class area. It does **not** yet explain the remaining factorization of $P_N(x)$ over $\mathbb{Q}$.

## 10.6 Brief remark on roots of unity

Because the residue label $a$ lives in the cyclic group $\mathbb{Z}/N\mathbb{Z}$, evaluations of $P_N(x)$ at roots of unity and discrete Fourier methods are natural later directions. We do not make them part of the present chapter. One possible later bridge is to compare the Fourier grouping by residue classes with an exact-product-layer decomposition of the same modular class by values $a + kN$ inside the positive window. The point of the polynomial viewpoint here, however, remains the exact structure already visible from symmetry, evaluation, the first derivative, and the distinguished zero-class coefficient.

# 11 A research program

This chapter is intentionally separated from the proved core. The earlier chapters establish exact results; the present one records the most natural next questions and extensions without presenting them as settled theory.

The first chapters now suggest a sharper research spine than a simple list of open problems. The immediate goal is to strengthen the two-dimensional static theory, then extend the boundary-layer program, and only after that reopen the higher-dimensional and dynamic branches.

All of the branches below are meant in a purely mathematical sense. Physics analogies, smooth ambient pictures, and unsupported shape heuristics may remain part of the historical background, but they are not part of the present theorem-bearing program unless they are translated into precise mathematical statements.

## 11.1 Cubic order and the next asymptotic problem

Chapter 7 now gives a genuine first asymptotic theorem for the total area:

$$\frac{(N-3)(N-2)(N-1)}{3} \leq S(N) \leq N(N-2)^2.$$

So $S(N)$ has cubic order of growth. That settles the exponent, but it does not settle the leading constant. The sharp conjectural refinement is

$$S(N) \sim N^3,$$

and the natural second-order object is the deficiency

$$D(N) := N^3 - S(N).$$

From that viewpoint the next asymptotic frontier is no longer to prove cubic growth again, but to understand how the missing area is created, how it is distributed across residue classes, and which geometric or arithmetic mechanisms are responsible for it.

## 11.2 Vertex characterization

The main unresolved structural question is to characterize those lattice points of $A_{N,a}$ that become vertices of the convex hull. The support-function formula shows that a point $(x_0, y_0)$ lies on the hull exactly when there exists a direction $(\cos\theta, \sin\theta)$ such that

$$x_0 \cos\theta + y_0 \sin\theta \geq x \cos\theta + y \sin\theta$$

for every $(x, y) \in A_{N,a}$. The difficulty is to translate this geometric extremality into usable arithmetic conditions in terms of $N$, $a$, $x$, and $y$.

This also sharpens a literature connection already noted in Chapter 3. Konyagin and Shparlinski [5] study how many vertices convex hulls of modular-hyperbola point sets can have. The MTMN program asks the finer question of which arithmetic points become those vertices and what structural event forces each change of supporting maximizer.

A useful near-term reformulation is to let the support direction $\theta$ rotate and record when the maximizing lattice point changes. This does not yet solve the vertex problem, but it turns the same question into a concrete switching problem: which point wins on each angular interval, and what arithmetic event forces the next switch?

One especially relevant predecessor is the work of Ford, Khan, Shparlinski, and Yankov on the maximal difference between an element and its inverse [13]. In support-function language, their quantity is a diagonal-direction extremum for the unit class $xy \equiv 1 \pmod{n}$: it asks how far the inverse plot reaches in the $x - y$ direction, equivalently where a support line at angle $-\pi/4$ first touches. Their paper gives asymptotic bounds and heuristics for that one-direction problem. MTMN asks a different follow-up question: the switching structure of support maximizers as $\theta$ rotates, across all residues and not only for the unit class.

## 11.3 Euclidean symmetries versus arithmetic structure

The residue classes already carry genuine Euclidean symmetries: transpose symmetry, central symmetry, and complementary-residue reflections. A separate future task is to organize the same sets arithmetically, for example through unit-group actions and multiplication fibers. These are not the same kind of structure. Euclidean isometries preserve the

actual shape in the plane; arithmetic equivalences may preserve only congruence data or fiber structure. Keeping these notions separate should prevent later confusion.

## 11.4 Coprimality and permutation geometry

When $a$ is coprime to $N$, the residue class $A_{N,a}$ is a permutation plot on the coprime subgrid

$$U_N = \{u \in \{1, \dots, N-1\} : \gcd(u, N) = 1\}.$$

There is exactly one point in each row and each column indexed by $U_N$, and no points in the other rows or columns. In the prime case $N = p$, this coprime subgrid is the whole nonzero grid, so each nonzero residue class is the graph of the permutation

$$y \equiv ax^{-1} \pmod{p}.$$

This gives a rigid starting point not only for primes but for the broader coprime part of MTMN. The next questions are whether permutation-plot language can help with vertex characterization, support directions, or asymptotic shape, and which features of these permutations are geometrically relevant.

Recent discrepancy work makes that rigidity quantitative from a different angle. Blomer, Risager, and Shparlinski prove upper and lower bounds for the box, ball, and isotropic discrepancy of modular inverse points, emphasizing their visible cellular structure and their deviations from random point sets [11]. Their setting varies the modulus and aggregates inverse pairs up to $c \leq X$, so it is not an exact fixed-$N$ MTMN theory. But it does give rigorous backing to the visual intuition that coprime residue classes carry structured permutation geometry rather than pseudorandom scatter. A later MTMN theory should explain how that rigidity interacts with support directions and hull vertices inside one fixed multiplication table.

## 11.5 Factorization profile and zero-divisor geometry

The asymptotic problem for $S(N)$ should no longer be phrased only as "prime versus composite." The zero class is now understood exactly: Chapter 5 describes $\mathrm{conv}(A_{N,0})$ in terms of divisor rectangles and the lower envelope of divisor points, then rewrites its area as a hy-

perbolic baseline minus an exact arithmetic correction. So the next questions begin after the mere existence of zero-divisor geometry. One now wants to understand how repeated prime factors, numbers of distinct prime factors, and finer zero-divisor structure influence the *shape* of that hull, the behavior of the nonzero classes, and ultimately the total area $S(N)$.

A practical first step is to build a small atlas grouped by factorization type – prime, prime power, squarefree composite, and moduli with heavier repeated factors – so that geometric changes are compared by arithmetic profile rather than only by the size of $N$.

## 11.6 Product layers and boundary hyperbolas

A related exploratory viewpoint reorganizes the ordinary embedding by exact products. For a fixed class $A_{N,a}$, one may study the occurring-layer set

$$\mathcal{C}_{N,a} := \{c \in \mathbb{Z}_{>0} : \exists (x, y) \in A_{N,a} \text{ with } xy = c\},$$

the multiplicities

$$\mu_{N,a}(c) := \#\{(x, y) \in A_{N,a} : xy = c\},$$

and, more importantly, the subset of those layers that actually meet the hull or the lower hull.

This leads immediately to several structural questions. How large can $\mathcal{C}_{N,a}$ be? Which layers carry most of the points? Which layers meet the hull, and which remain invisible because they stay strictly in the interior? The boundary version is the most important: identify the exact hyperbolas $xy = c$ that contain hull vertices or lower-hull vertices. That is a direct generalization of the divisor-envelope story from the zero class, now posed beyond the solved zero residue.

Here the Chapter 3 literature should stay in view. Ford, Khan, and Shparlinski [6] study geometric and extremal behavior of the unit modular hyperbola, while Konyagin and Shparlinski [5] study convex-hull vertex counts. The MTMN layer program asks a finer boundary question: after one modular class is decomposed by exact products $a+kN$, which of those layers actually survive as visible geometry?

There is also a broader additive-multiplicative backdrop. Hart, Iosevich, and Solymosi prove finite-field sum-product estimates using incidence bounds derived from Kloosterman sums [15]. Their theorem is not about a fixed modular hyperbola in the positive lattice window, and it does not address exact hull layers or zero divisors. But it formalizes a general principle that matters here as well: additive structure and multiplicative structure cannot both remain arbitrarily compressed. Since the MTMN layer program studies multiplicative slices $xy = a + kN$ together with additive coordinates such as $x + y$ and $x - y$, any later theory of dominant hyperbolas should fit within that broader sum-product tension rather than treating these two structures as independent.

## 11.7 Dominant hyperbolas, compositeness, and the unit side

The next step is to understand dominance. A natural first definition uses lower-hull vertex count, and that immediately suggests a research fork between composite residues and unit residues.

On the composite side, the current computations suggest a strong base-layer principle. If $k$ is composite and $\gcd(k, N) = 1$, the layer $xy = k$ often appears to dominate the lower hull because its factor pairs already provide a divisor scaffold before the higher layers $k + N, k + 2N, \ldots$ enter. The first task is to explain when that picture is genuinely universal, when it fails, and what arithmetic feature marks the exceptional cases.

On the unit side, the class $a = 1$ seems to behave differently. The base layer $xy = 1$ appears to stay in the dominant set for odd $N$, but usually only as part of a tie rather than as a unique winner. This suggests that unit classes may spread their lower boundary across several layers more evenly than composite residues do. One then wants to know whether this is specific to $a = 1$, whether other coprime classes behave similarly, and what replaces factorization as the governing geometric mechanism.

These questions also touch the existing harmonic-analysis language. The exponential filter from Chapter 2 and the roots-of-unity remarks later in the book group the data by residue classes modulo $N$. The hyperbola-layer viewpoint groups the same class by exact products $a + kN$. Understanding any later bridge between those two decompositions would be a substantial project in its own right.

Because the project already has an interactive explorer for residue classes, hulls, boundary layers, support directions, and hyperbola overlays, this branch is experimentally approachable. Those tools do not replace proofs, but they make the new questions concrete enough to test on small moduli before one looks for a theorem.

## 11.8 Evaluation of $\sum 1/S(N)$

The first-boundary model proves convergence, and Chapter 7 now gives a rigorous numerical enclosure for

$$\sum_{N=4}^{\infty} \frac{1}{S(N)}$$

using exact computation together with the first-boundary tail identity. Writing

$$C := \sum_{N=4}^{\infty} \frac{1}{S(N)}, \qquad P_M := \sum_{N=4}^{M} \frac{1}{S(N)},$$

one has

$$0 \le C - P_M \le \sum_{N=M+1}^{\infty} \frac{1}{S^{(1)}(N)} = \frac{3}{2(M-2)(M-1)}.$$

In particular, the exact computation through $M = 2000$ yields the rigorous enclosure

$$0.648623982413541356 \le C \le 0.648624357976698263.$$

This already rules out the old guess $C = \sqrt{e} - 1$. What remains open is not numerical localization but symbolic understanding: one still lacks sharp enough structural control of $S(N)$ to identify the exact constant or explain its arithmetic nature. Sharper asymptotics for $S(N)$, a better understanding of which points become hull vertices, or an arithmetic decomposition of the total area are the natural next routes.

The weighted companion series

$$\sum_{N=4}^{\infty} \frac{N}{S(N)}$$

converges for the same cubic-scale reason and provides a slower global diagnostic of the same picture, but the main asymptotic challenge remains the structure of $S(N)$ itself.

## 11.9 Series, constants, and sequence searches

The reciprocal series is only one entry point into a broader numerical program. Because the first-boundary total is an explicit cubic polynomial, one already gets exact constants such as

$$\sum_{N=4}^{\infty} \frac{1}{S^{(1)}(N)} = \frac{3}{4}, \qquad \sum_{N=4}^{\infty} \frac{N}{S^{(1)}(N)} = \frac{15}{4}.$$

This suggests a family of related questions:

- what are the exact values of other boundary-model series?
- which weighted full-MTMN series converge?
- which stable numerical constants deserve to be recorded even before they are explained?
- which derived sequences or partial sums are worth comparison with OEIS or other sequence tables?

Such sequence and constant searches should remain heuristic rather than evidentiary, but they are mathematically legitimate ways of organizing data and spotting hidden structure.

## 11.10 Higher boundary layers and planar interfaces

The first and second layers already produce exact cubic models. The next natural task is to study the $k$-th boundary layer and the union of the first $k$ layers. The change of variables $(u, v) = (x + y, x - y)$ hints that deeper layers may continue to generate structured convex objects, perhaps sums or hulls of rectangles and more general centrally symmetric polygons. A related planar problem is to understand the interfaces between successive layers: where new vertices first appear, how the supporting directions change, and how much area each new layer contributes.

Before that larger $k$-layer program, the even-$N$ second boundary is already an instructive missing local model. Once $2$ ceases to be invertible, some residues disappear from the layer, other residues hit a side more than once, and the clean one-parameter odd-$N$ description breaks down. Resolving that case would clarify how noninvertibility first distorts the boundary program.

## 11.11 Periodic and toric viewpoints

The Euclidean window $\{1, \dots, N-1\}^2$ is not the only natural home for MTMN. Before restricting to that square, the modular equation lives on the periodic grid $(\mathbb{Z}/N\mathbb{Z})^2$. This suggests two later viewpoints. The first is an ambient periodic viewpoint, useful for organizing modular coordinates before they are embedded in the plane. The second is an algebraic-torus viewpoint for prime moduli, where nonzero coordinates lie in $(\mathbb{F}_p^\times)^2$ and the condition $xy = a$ becomes a multiplicative-level constraint. These viewpoints should enrich the subject, but they should not replace the present Euclidean convex-hull geometry.

## 11.12 Higher-dimensional MTMN and slice questions

Once the planar theory is stronger, one can define a $d$-dimensional multiplication table modulo $N$ by imposing

$$x_1 \cdots x_d \equiv a \pmod{N}$$

inside the discrete cube $\{1, \dots, N-1\}^d$. Several structural questions then arise. Coordinate permutations always preserve the residue class. Full coordinate inversion sends $a$ to $(-1)^d a$, so parity matters already at the level of symmetry. For prime moduli, fixing one coordinate suggests slice identifications of the form

$$x_1 \cdots x_{d-1} \equiv ac^{-1} \pmod{p},$$

which hints that many higher-dimensional slices may be arithmetically related even when they are not literally congruent in Euclidean space.

Chapter 3 already recalled Shparlinski's work on multidimensional modular hyperbolas [9]. That literature is the natural backdrop here, even though the present MTMN program asks for higher-dimensional slice and convex-geometric questions in its own language.

Shparlinski's survey is especially useful here because it makes clear what the established higher-dimensional program usually looks like: discrepancy bounds, distribution in boxes, and character-sum estimates for one congruence class at a time [1]. An MTMN extension would emphasize different questions: exact slice identifications, higher-dimensional boundary-layer models, and convex-geometric formulas across the full residue family, in-

cluding composite-sensitive and zero-divisor phenomena that disappear if one only studies invertible classes.

## 11.13 Lattice counting: Pick and Ehrhart

Pick's theorem [4] already explains why central symmetry forces the areas to be integers in dimension two, but a larger lattice-counting program remains open. Once the relevant polygons or polytopes are explicit enough, one can ask whether boundary counts, interior counts, and eventually Ehrhart-style counting functions help organize the area and volume data. This should be treated as a later structural tool, not as a substitute for identifying the actual hulls.

A nearby but distinct model comes from the ordinary multiplication-table problem. Limbach, Scheidweiler, and Triesch encode finite multiplication alphabets by exponent sets, compare their iterated sumsets with Ehrhart polynomials of explicit polytopes, and prove eventual polynomiality for generalized product counts [16]. That is not a modular-hyperbola result and it does not identify residue-class hulls. Its relevance is methodological: it shows that multiplication tables can be reorganized as lattice-point problems in explicit polytopes. For MTMN, Pick and Ehrhart should enter only after the relevant zero-class polygons or higher-dimensional hulls have been identified exactly; the counting theory would then complement the geometric identification problem rather than substitute for it.

## 11.14 Choosing a notion of dynamics

The archives repeatedly point toward "dynamics," but that word should not be used until one specifies what is evolving. At least four mathematically distinct possibilities are visible:

- layer-growth dynamics, where one adds successive boundary layers
- support-direction dynamics, where the maximizer changes as $\theta$ varies
- residue dynamics, where the residue class $a$ varies for fixed $N$
- modulus dynamics, where the entire table changes as $N$ changes

These are all legitimate future directions, but they should not be conflated. The present book develops only the static two-dimensional geometry and uses the research program to indicate where each dynamic branch could begin.

# 12 Appendix

This appendix collects supporting material that remains mathematically useful but does not belong on the main theorem spine of the monograph. Nothing here is discarded; it is gathered behind the core chapters so that the central exact results stay visible while secondary lenses, data tables, and figure-led observations remain available.

## 12.1 The symmetric embedding as a secondary lens

Earlier stages of this project often recentered the nonzero residues around the origin. That construction is mathematically reasonable, and it did help expose some exact-product layers quickly. It is kept here for completeness, for the interactive explorer, and for the concept pages, but it is **not** the main geometry of this book.

**Definition 12.1** (Symmetric embedding)**.** Let

$$I_N := \{t \in \mathbb{Z} : -N/2 < t \le N/2\}, \qquad I_N^\times := I_N \setminus \{0\}.$$

For $N \ge 2$ and $a \in \{0, 1, \dots, N-1\}$, define

$$A_{N,a}^{\mathsf{sym}} := \{(x, y) \in (I_N^\times)^2 : xy \equiv a \pmod N\}.$$

This preserves the modular congruence class. What changes is only the Euclidean placement of the representatives.

### 12.1.1 Why one might try it

The main attraction is visual economy. Exact-product layers that are far apart in the positive window can collapse into a much shorter centered list, so the family of curves $xy = a + kN$ may look cleaner at first glance.

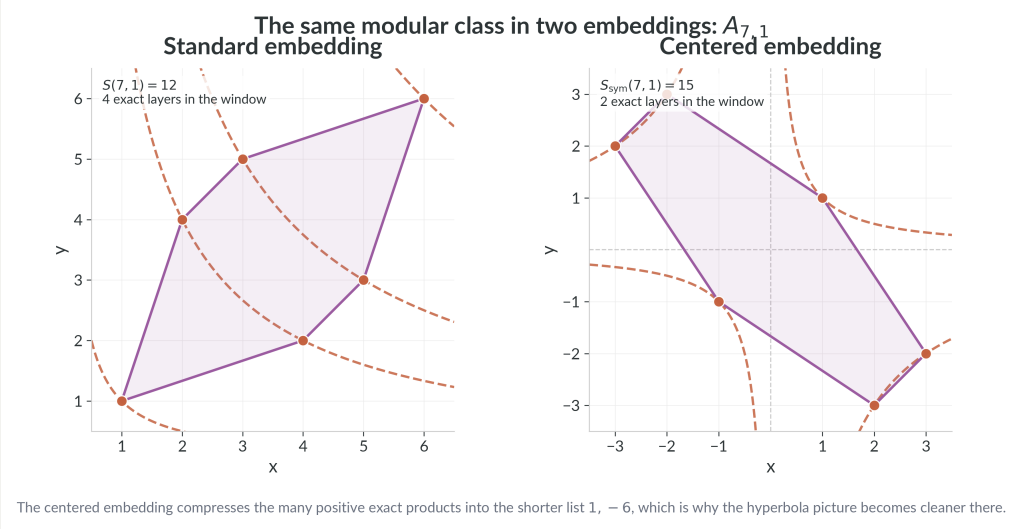


Figure 41: The same residue class in the ordinary and centered windows. The centered picture packs the exact layers more tightly, which explains why this embedding initially looked attractive as a hyperbola-friendly visualization.

The embedding also preserves genuine arithmetic information. The congruence condition is unchanged, the number of points in each residue class is unchanged, and for odd $N$ the centered window has exact half-turn symmetry about the origin. That last point is mathematically appealing: for odd moduli the map

$$(x, y) \longmapsto (-x, -y)$$

really does preserve the centered window and the centered residue class.

So the symmetric embedding is not a mistake or a gimmick. It is a reasonable alternate realization of the same modular data.

### 12.1.2 Why it is not the main geometry

The first problem is parity. For even $N$, the representatives $\pm N/2$ coincide modulo $N$, so one must choose one and discard the other. Our convention keeps $+N/2$. That means exact central symmetry already fails at the level of the coordinate window.

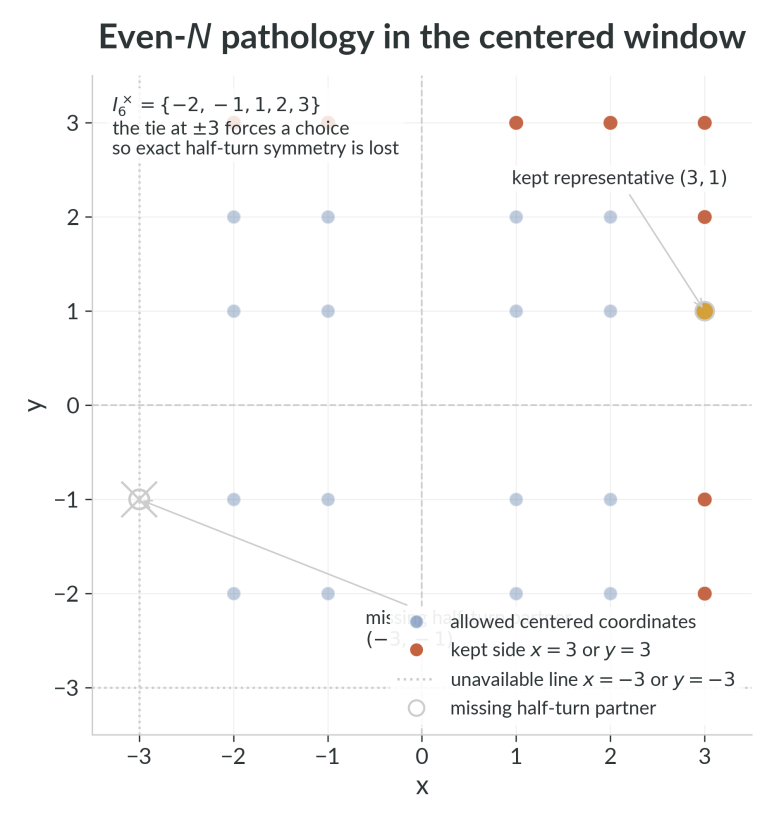


Figure 42: For even $N$, the centered window must choose between the representatives $\pm N/2$. Here the choice of $+3$ and the loss of $-3$ break exact half-turn symmetry before any residue class is even drawn.

The second problem is more geometric. In the main chapters, the zero residue class lives inside the multiplication-table square, and its natural symmetries are transpose symmetry and the half-turn about $(N/2, N/2)$. The divisor-rectangle story, the lower envelope, and the compositeness story are all native to that ordinary window. The centered realization rearranges that geometry rather than clarifying it.

The third problem is that compactness is not the same as boundary relevance. In centered coordinates the exact hyperbolas often look neatly packed, but they can cut through the shape rather than support it naturally. The picture below is typical: several centered layers are visible, yet some of them pass through the interior instead of tracing the actual

boundary geometry.

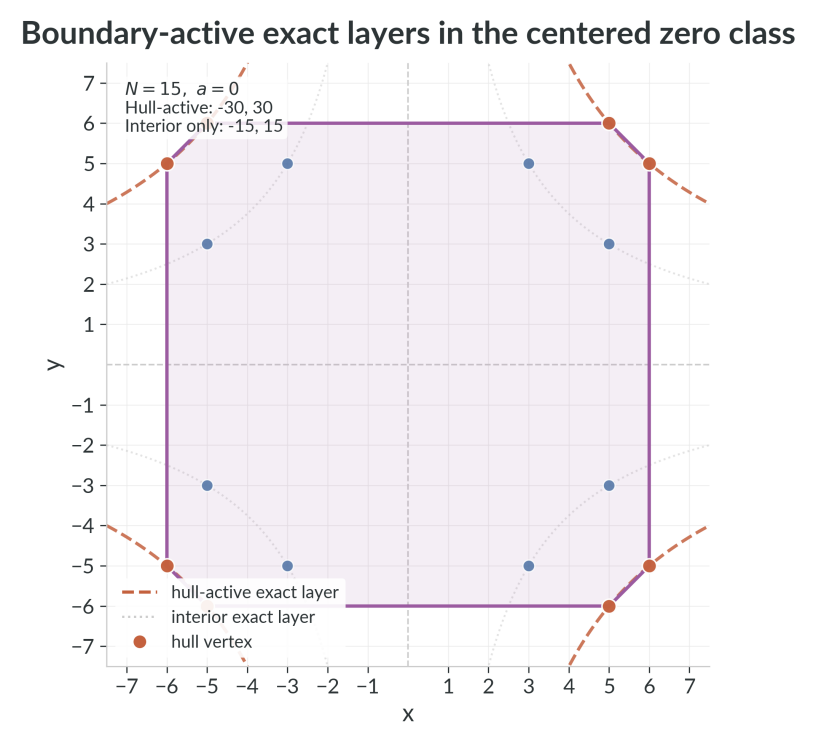


Figure 43: In the centered zero class for $N = 15$, the layers $xy = \pm 30$ carry hull vertices, while the layers $xy = \pm 15$ pass through the interior. The centered curves are visible, but visibility here is weaker than natural boundary support.

For these reasons the symmetric embedding is not adopted as the primary geometry of the book. It remains worth mentioning because it preserves the modular class exactly, it sometimes packs exact layers very efficiently, and it records a real stage in the development of the project. We keep it here in that spirit: not as a major payoff, and not as something to erase, but as an interesting alternate lens that turned out to be secondary.

## 12.2 Residue-by-residue values for $4 \leq N \leq 10$

| $N$ | $(S(N,0), S(N,1), \ldots, S(N,N-1))$ | $S(N)$ |
|---|---|---|
| 4 | $(0,\ 0,\ 2,\ 0)$ | 2 |
| 5 | $(0,\ 3,\ 4,\ 4,\ 3)$ | 14 |
| 6 | $(2,\ 0,\ 9,\ 8,\ 9,\ 0)$ | 28 |
| 7 | $(0,\ 12,\ 8,\ 15,\ 15,\ 8,\ 12)$ | 70 |
| 8 | $(8,\ 0,\ 22,\ 16,\ 24,\ 16,\ 22,\ 0)$ | 108 |
| 9 | $(9,\ 27,\ 12,\ 32,\ 27,\ 27,\ 32,\ 12,\ 27)$ | 205 |
| 10 | $(18,\ 32,\ 41,\ 24,\ 45,\ 32,\ 45,\ 24,\ 41,\ 32)$ | 334 |

## 12.3 What the figures suggest

The exact totals $S(N)$ and the boundary models $S^{(1)}(N)$ and $S^{(2)}(N)$ already show several qualitative patterns in a single plot.

- The true total $S(N)$ grows rapidly and irregularly with $N$.
- The first-boundary model is cubic and simple enough to evaluate exactly.
- The odd-$N$ second-boundary model is also cubic, but with a different coefficient structure.
- Boundary models are not merely lower bounds; they capture a substantial fraction of the total area and therefore deserve to be treated as genuine geometric approximations.

The same figures also motivate the later numerical program. Once the boundary totals are explicit and the full totals can be computed exactly, one can form reciprocal sums, weighted sums, and derived constants. Those directions should remain secondary to the proved core, but they are worth preserving as visual evidence for what the theorem-bearing chapters are pointing toward.

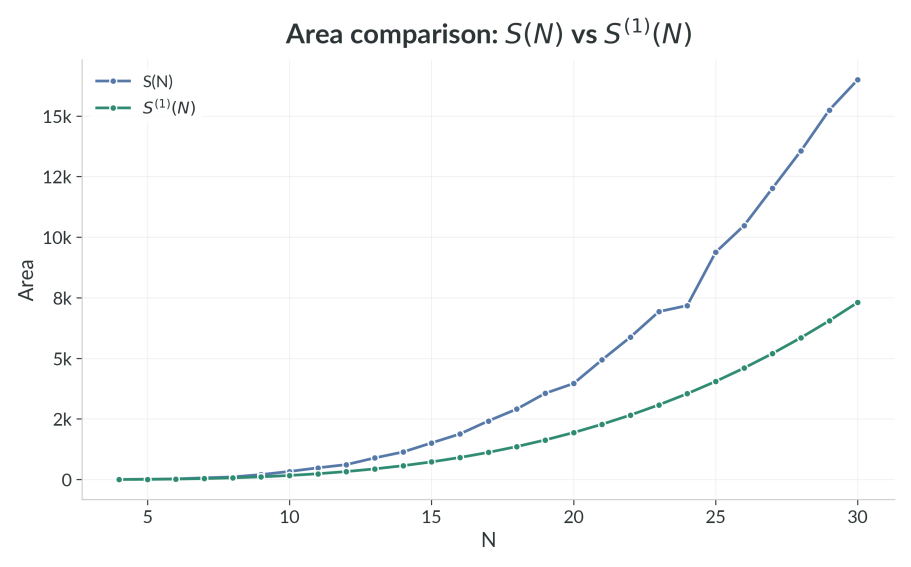


Figure 44: Comparison of the exact area sum $S(N)$ with the explicit lower models $S^{(1)}(N)$ and, for odd $N$, $S^{(2)}(N)$. The visible gap reflects interior contributions not captured by the first boundary alone.

An individual residue class also shows the layered approximation clearly. Even for small $N$, the first boundary often captures the essential outer shape while leaving the finer interior geometry to later chapters.

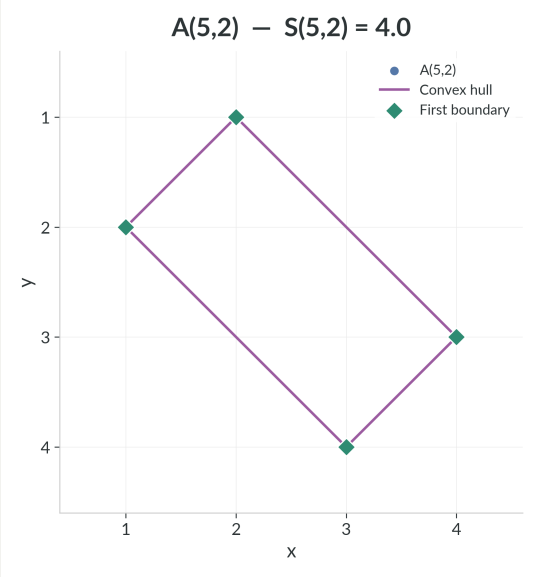


Figure 45: Residue class $A_{5,2}$ with its convex hull and first-boundary points marked.

## 12.4 Summary of exact formulas and identities established in this document

$$S(N,a) = \text{Area}(\text{conv}(A_{N,a})),$$

$$S(N) = \sum_{a=0}^{N-1} S(N,a),$$

$$S(N,0) = 0 \iff \big(N \text{ is prime}\big) \text{ or } N = 4,$$

$$\text{conv}(A_{N,0}) = \text{conv}\left( \bigcup_{\substack{d \mid N \\ 1<d<N}} [d,\, N-d] \times \left[ \frac{N}{d},\, N - \frac{N}{d} \right] \right) \qquad (N \text{ composite}),$$

$$S^{(1)}(N,a) = \begin{cases} 0, & a = 0, \\ 2(a-1)(N-a-1), & 1 \le a \le N-1, \end{cases}$$

$$S^{(1)}(N) = \frac{(N-3)(N-2)(N-1)}{3},$$

$$\sum_{N=4}^{\infty} \frac{1}{S^{(1)}(N)} = \frac{3}{4},$$

$$S^{(2)}(N,a) = \begin{cases} 0, & a = 0, \\ 2\,|b-2|\,|N-b-2|, & N \text{ odd},\ 1 \le a \le N-1,\ b \equiv 2^{-1}a \pmod{N}, \end{cases}$$

$$S^{(2)}(N) = \frac{(N-3)(N^2-9N+32)}{3} \qquad (N \text{ odd}),$$

$$P_N(x) = \sum_{a=0}^{N-1} S(N,a)x^a,$$

$$P_N(0) = S(N,0), \qquad P_N(1) = S(N),$$

$$P_N'(1) = \frac{N}{2}\big(S(N) - S(N,0)\big),$$

$$P_N(-1) = S(N,0) \qquad (N \text{ odd}).$$

## 13 Closing note

The multiplication table modulo $N$ has turned out to be richer than its elementary definition suggests. In these pages, several parts of the subject have moved from experiment to exact description often enough to justify a real theory: the areas $S(N, a)$ and $S(N)$ are integral, the first outer frame already gives exact and nontrivial formulas, the second boundary for odd $N$ has its own explicit model, and the zero residue class now admits both a complete divisor-controlled hull description and an exact hyperbolic baseline-plus-correction identity. The book also places these questions more honestly beside the existing literature on modular hyperbolas.

One pleasure of MTMN is that the pictures do not have to be discarded once the theorems begin. On the contrary, they keep their authority. The small examples are not merely suggestive sketches; they lead toward exact support formulas, boundary layers, divisor envelopes, and clean arithmetic distinctions between unit geometry and zero-divisor geometry. This is often the moment when a project becomes mathematically shareable: the experiments still matter, but they no longer stand alone.

Much remains open. For general residue classes, the full hull problem is still subtle. The first two boundary layers are now explicit, but deeper layers and their interfaces remain largely unexplored. The total area sum already has a genuine first asymptotic theorem: it is trapped between explicit cubic bounds and therefore has cubic order of growth. The reciprocal series and its weighted companion reflect that same large-$N$ picture, and the reciprocal constant already admits a rigorous enclosure. Yet the sharper leading constant, the deficiency $N^3 - S(N)$, and the structural reason behind the missing area remain unclear. One still wants to understand which exact product layers ever become boundary-active, when one layer dominates, and how compositeness and coprimality divide that story. Beyond the plane lie higher-dimensional multiplication tables, slice relations, and new questions about what should count as geometry in those settings.

I hope the reader leaves with the sense that MTMN is not merely an archive of observations, nor yet a finished edifice, but a usable entrance to a subject. It is still possible to compute a new example, draw a better picture, prove a cleaner statement, or discover that some modest-looking residue class has been hiding a theorem all along. If the multiplica-

tion table seemed exhausted in school, it may still have a few geometric surprises left in it.

## References


[1] I. E. Shparlinski, "Modular hyperbolas," *Japanese Journal of Mathematics*, vol. 7, 2012.

[2] D. Koukoulopoulos, "On the number of integers in a generalized multiplication table," *Journal für die reine und angewandte Mathematik*, vol. 689, pp. 33–99, 2014, doi: 10.1515/crelle-2012-0064.

[3] C. F. Gauss, *Disquisitiones arithmeticae*. Leipzig: Gerh. Fleischer, 1801. Available: https://archive.org/details/disquisitionesa00gaus

[4] G. Pick, "Geometrisches zur zahlenlehre," *Sitzungsberichte des deutschen naturwissenschaftlich-medicinischen Vereines fur Bohmen "Lotos" in Prag*, vol. 19, pp. 311–319, 1899, Available: https://hdl.handle.net/10622/02638381-D295-4F97-95C3-57EF49FE3D89

[5] S. V. Konyagin and I. E. Shparlinski, "On the convex hull of the points on modular hyperbolas." 2010. Available: https://arxiv.org/abs/1012.1444

[6] K. Ford, M. R. Khan, and I. E. Shparlinski, "Geometric properties of points on modular hyperbolas," *Proceedings of the American Mathematical Society*, vol. 138, 2010.

[7] J. Cilleruelo and M. Z. Garaev, "Concentration of points on two and three dimensional modular hyperbolas and applications," *Geometric and Functional Analysis*, vol. 21, 2011.

[8] I. E. Shparlinski and A. Winterhof, "On the number of distances between the coordinates of points on modular hyperbolas," *Journal of Number Theory*, vol. 128, 2008.

[9] I. E. Shparlinski, "On the distribution of points on multidimensional modular hyperbolas," *Proceedings of the Japan Academy, Series A, Mathematical Sciences*, vol. 83, 2007.

[10] I. E. Shparlinski and J. F. Voloch, "Visible points on curves over finite fields," *Bulletin of the Polish Academy of Sciences: Mathematics*, vol. 55, 2007.

[11] V. Blomer, M. S. Risager, and I. E. Shparlinski, "Triple sums of kloosterman sums and the discrepancy of modular inverses." 2025. Available: https://arxiv.org/abs/2411.17823

[12] A. Ivić, E. Krätzel, M. Kühleitner, and W. G. Nowak, “Lattice points in large regions and related arithmetic functions: Recent developments in a very classic topic.” 2004. Available: https://arxiv.org/abs/math/0410522

[13] K. Ford, M. R. Khan, I. E. Shparlinski, and C. L. Yankov, “On the maximal difference between an element and its inverse in residue rings,” *Proceedings of the American Mathematical Society*, vol. 133, no. 12, pp. 3463–3468, 2005, doi: 10.1090/S0002-9939-05-07962-1.

[14] A. Bower, R. Evans, V. Luo, and S. J. Miller, “Coordinate sum and difference sets of $d$-dimensional modular hyperbolas.” 2012. Available: https://arxiv.org/abs/1212.2930

[15] D. Hart, A. Iosevich, and J. Solymosi, “Sum-product estimates in finite fields via kloosterman sums.” 2006. Available: https://arxiv.org/abs/math/0609426

[16] A. M. Limbach, R. Scheidweiler, and E. Triesch, “Effective khovanskii, ehrhart polytopes, and the erdős multiplication table problem.” 2025. Available: https://arxiv.org/abs/2503.23578