

Computing Element-Order Distributions in Modular Groups via Prime-Power Factorization

Adnan Asghar^{1,*}, Muhammad Bilal²

¹*Department of Chemical and Materials Engineering, University of Alberta, Donadeo Innovation Centre for Engineering, Edmonton, Alberta, Canada.*

²*Department of Mathematics and Statistics, The University of Lahore, Lahore, Pakistan.*

*Corresponding author: adnanasghar14041985@gmail.com

²m.bilal56@gmail.com

Received: 17 January 2023. Received in revised form: 21 March 2023. Accepted: 10 May 2023. Published online: 30 June 2023.

Abstract

Element-by-element listings of modular groups are easy to produce for small moduli, but they conceal the arithmetic mechanisms that determine element orders, cyclicity, and subgroup multiplicity. This article asks whether the complete order distribution of the additive group $\mathbb{Z}_n$ and the unit group $U(n) = (\mathbb{Z}/n\mathbb{Z})^\times$ can be recovered from the prime-power factorization of n without tracing every power orbit. A factorization-aware procedure is developed in which additive orders are obtained from the divisor lattice, local unit groups are represented by cyclic factors, and their order distributions are combined through a least-common-multiple convolution. Invariant-factor normalization supplies a compact structural certificate, while an independent residuewise order reducer verifies every count. The computational corpus contains all 299 moduli from 2 through 300. Two detailed arithmetic cases, $\mathbb{Z}_{209}$ and $U(210)$, are used to establish correspondence with independently tabulated records. The method recovers the additive spectrum of $\mathbb{Z}_{209}$ as $1^1, 11^{10}, 19^{18}, 209^{180}$ and the unit spectrum of $U(210)$ as $1^1, 2^7, 3^2, 4^8, 6^{14}, 12^{16}$. Across the full corpus, 114 unit groups are cyclic, the richest unit-order distribution occurs at $n = 241$ with 20 distinct orders, and the largest compression is observed at $n = 263$, where 262 residues are represented by four spectral entries. The invariant-factor ranks occur with frequencies 1, 113, 133, 47, and 5 from rank zero through rank four, and the corresponding numbers of involutions are 0, 1, 3, 7, and 15. The 299 unit groups occupy 142 invariant-factor classes, of which 85 recur at more than one modulus. Least-common-multiple convolution uses 4,425 divisor-pair combinations, compared with 1,587,655 modular multiplication steps required by direct orbit traversal. Repeated whole-corpus timings give a median reduction from 75.58 to 4.25 ms. The results show that the order spectrum is both a mathematically interpretable group signature and an efficient computational object for modular-group analysis.

Keywords: finite abelian group; modular arithmetic; element order; Chinese remainder theorem; Carmichael function; invariant factors; cyclic subgroup

1. Introduction

The arithmetic of residue classes supplies one of the most accessible entrances to finite group theory and one of its most consequential computational settings. Under addition, the residue classes modulo n form the cyclic group $\mathbb{Z}_n$. Under multiplication, the invertible classes form the unit group $U(n)$, whose order is Euler's totient $\varphi(n)$. These two groups share a modulus but generally have very different internal organizations. The additive group always contains an element of order n , whereas $U(n)$ may be cyclic, a product of several cyclic groups, or a group with an exponent much smaller than its cardinality. A list of residues therefore answers only the first and least informative question. Structural computation must also determine how many elements have each possible order, how many elements generate the whole group, how many distinct cyclic subgroups occur, and how these quantities change with the factorization of n .

The classical algebra underlying these questions is well established. Standard treatments derive the order of an element in a cyclic group, the decomposition of a finite abelian group, and the behavior of direct products [1,2]. Computational number theory adds efficient factorization, modular powering, and arithmetic-function evaluation [3–5]. Yet a practical tension remains between two ways of presenting small modular groups. One approach prints the orbit generated by every element, producing long and repetitive output. The other records only a structural decomposition, which is compact but may omit the order multiplicities that make the decomposition numerically testable. The present study connects these descriptions by treating the order spectrum—the map from each attainable order to the number of elements of that order—as the primary computational object.

For a finite group G , define its order spectrum with multiplicity by

$$\mathcal{S}(G) = \{(d, N_G(d)) : d \mid |G|\}, \quad N_G(d) = |\{g \in G : \text{ord}(g) = d\}|. \quad (1)$$

The spectrum retains more information than the group order or exponent alone. It immediately identifies the number of generators when G is cyclic, counts cyclic subgroups through division by $\varphi(d)$, and exposes the concentration of elements at low orders. It is also substantially shorter than a residue-by-residue table whenever many elements share an order. Because $\sum_d N_G(d) = |G|$, the spectrum carries a natural completeness check; because $N_G(d)$ is divisible by $\varphi(d)$, it carries a second, independent arithmetic check. These properties make it suitable both for mathematical analysis and for reliable software output.

The additive case has a particularly transparent description. The order of $a \in \mathbb{Z}_n$ equals $n/\gcd(a, n)$, so every divisor d of n occurs as an order and occurs exactly $\varphi(d)$ times. Consequently, the additive spectrum is completely determined by the divisor lattice of n . The unit case is more delicate. Its possible orders divide the Carmichael exponent $\lambda(n)$, not necessarily $\varphi(n)$, and different prime-power components interact through least common multiples. The distinction is arithmetically important. For example, a modulus with many divisors can give a rich additive spectrum while its unit group is concentrated in only a few orders; a prime modulus has the sparsest possible additive spectrum but can have a unit group with many distinct orders whenever $p - 1$ is divisor-rich.

The distribution of $\varphi(n)$ and $\lambda(n)$ has therefore attracted sustained attention. Work on popular totient values and the global distribution of Euler’s function explains why equal group cardinalities arise from different moduli [6, 7]. The normal and extremal behavior of Carmichael’s function clarifies why the exponent of $U(n)$ may be much smaller than $\varphi(n)$ [8, 9]. Period studies for modular power generators further show that element order is not merely a classroom quantity; it controls orbit length, recurrence, and the usable state space of arithmetic generators [10, 11]. These lines of research motivate an explicit finite census in which structural quantities are computed together rather than in isolation.

Efficient computation begins with factorization. Classical methods based on smooth factors, random iteration, elliptic curves, and sieving established the principal routes from an integer to its prime-power components [12–16]. Later work improved coprime splitting and clarified deterministic primality testing [17, 18]. For the small range considered here, trial division is sufficient and has the advantage of transparent reproducibility. The conceptual point is nevertheless general: once $n = \prod p_i^{\alpha_i}$ is known, the Chinese remainder theorem transforms one modular group into a direct product of local groups, and a global order becomes the least common multiple of local orders.

Modular arithmetic itself has a rich algorithmic literature. Fast reduction and exponentiation techniques reduce the cost of repeated modular products [19–21]; explicit Chinese-remainder arithmetic can accelerate modular powering when the factorization is available [22]. Improvements in integer multiplication alter the asymptotic cost of large operands [23, 24]. The moduli in this study are deliberately small, so large-integer asymptotics are not the primary concern. Their relevance lies instead in a design principle: the computational representation should respect the algebraic factorization before low-level arithmetic is optimized.

Algorithms for finite abelian groups likewise emphasize order computation, basis recovery, and structure determination. Space-efficient group-structure procedures, generalized Pohlig–Hellman methods, and deterministic basis algorithms show how a finite abelian group can be identified from operations and order information [25–28]. Quantum decomposition algorithms demonstrate that the same structural problem remains central under a different computational model, while order finding can expose the complete factorization of an integer through classical post-processing [29, 39]. Generic-group investigations separate what follows from the group operation from what follows from a particular residue representation [30]. The method used here is specialized to $\mathbb{Z}_n$ and $U(n)$, but it benefits from the same separation: local cyclic components supply exact spectra, and a generic least-common-multiple rule assembles the direct product.

Computer algebra systems have made finite-group calculations widely available. Foundational work on permutation and matrix groups established polynomial-time and black-box viewpoints [31, 32]. Random element generation, p -group construction, isomorphism testing, and mature symbolic implementations subsequently broadened the scale and reliability of group computation [33–36]. General systems such as Magma and GAP provide extensive group-theoretic functionality [37, 38]. A specialized computation for modular groups should not imitate the full scope of those systems. It should instead exploit number-theoretic structure to produce a concise, auditable record that can be checked without proprietary software.

Order spectra are particularly suitable for this purpose because they connect several levels of description. At the element level, $N_G(d)$ answers how frequently a period occurs. At the subgroup level, $N_G(d)/\varphi(d)$ counts cyclic subgroups of that order. At the whole-group level, the largest occupied order gives the exponent and the entry at $|G|$ decides whether a generator exists. The same dictionary can therefore answer questions that are often implemented as separate menu commands. More importantly, those answers remain mutually constrained. A list that reports seven involutions in $U(210)$ must also report seven distinct subgroups of order two, while sixteen elements of order twelve must correspond to four such subgroups. An implementation that violates either relation is demonstrably incorrect, even if its residue list appears plausible.

The chosen interval is small enough for complete verification but arithmetically varied enough to separate competing effects. It includes primes, powers of every prime below 18, square-free integers with up to four prime factors, and composite

values for which $\lambda(n)$ is far below $\varphi(n)$. Consecutive moduli can move abruptly between these types. The pair 209 and 210 changes from an odd semiprime to the product of the first four primes; the pair 240 and 241 changes from a smooth composite to a prime whose predecessor is divisor-rich. Such transitions are more revealing than a collection restricted to primes or to randomly selected values, because they hold numerical size nearly fixed while changing the underlying group structure.

The research question is therefore precise: to what extent can the complete element-order distributions, generator counts, cyclic-subgroup counts, and cyclicity status of $\mathbb{Z}_n$ and $U(n)$ be recovered directly from the factorization of n , and what arithmetic patterns emerge when this recovery is applied to every modulus from 2 through 300? The answer requires more than a faster list. It requires an exact transformation from prime powers to order multiplicities, a normalization that certifies the unit-group structure, and an independent computation that can detect implementation errors. The contributions are fourfold. First, a least-common-multiple convolution combines local cyclic spectra without enumerating global elements. Second, invariant factors are reconstructed by aligned prime valuations and used as a structural certificate. Third, residuewise order reduction supplies an independent equality check for every modulus. Fourth, the complete census is interpreted through order richness, cyclicity, spectral compression, and additive–multiplicative contrast. This combination yields a paper whose central output is not a software menu but an arithmetic account of why the observed group properties occur.

2. Algebraic setting and structural quantities

Let $n \geq 2$ and write

$$n = \prod_{i=1}^r p_i^{\alpha_i}, \quad (2)$$

where the p_i are distinct primes. The additive group $\mathbb{Z}_n$ has order n and is cyclic. Its unit group

$$U(n) = \{a \in \{1, \dots, n-1\} : \gcd(a, n) = 1\} \quad (3)$$

has order $\varphi(n)$. The Chinese remainder theorem gives

$$U(n) \cong \prod_{i=1}^r U(p_i^{\alpha_i}). \quad (4)$$

For an odd prime p , the local group $U(p^\alpha)$ is cyclic of order $(p-1)p^{\alpha-1}$. For powers of two, $U(2)$ is trivial, $U(4) \cong C_2$, and

$$U(2^\alpha) \cong C_2 \times C_{2^{\alpha-2}}, \quad \alpha \geq 3. \quad (5)$$

These formulas reduce the unit group to a product $C_{m_1} \times \dots \times C_{m_t}$ of explicit cyclic factors. They also show why the prime 2 must be treated separately: its higher power unit groups require two cyclic components rather than one.

Three numerical summaries are used throughout. The exponent $\exp(G)$ is the least common multiple of all element orders. The order richness $R(G)$ is the number of distinct orders,

$$R(G) = |\{d : N_G(d) > 0\}|, \quad (6)$$

and the spectral compression $Q(G) = |G|/R(G)$ is the mean number of elements represented by one nonzero spectral entry. The word “compression” is used literally: the stored object contains $R(G)$ pairs in place of $|G|$ element records. It is not intended as an information-theoretic entropy measure.

The number of distinct cyclic subgroups follows from the fact that every cyclic subgroup of order d has exactly $\varphi(d)$ generators. Hence

$$C(G) = \sum_{d \mid |G|} \frac{N_G(d)}{\varphi(d)}. \quad (7)$$

Each summand is an integer. This divisibility supplies a useful audit of computed spectra, while the total in equation (7) distinguishes groups that share the same order and exponent but distribute their elements differently.

Proposition 2.1. *For every $n \geq 2$, the additive order spectrum is*

$$\mathcal{S}(\mathbb{Z}_n) = \{(d, \varphi(d)) : d \mid n\}. \quad (8)$$

Proof. The order of a in $\mathbb{Z}_n$ is $n/\gcd(a, n)$. It equals d precisely when $\gcd(a, n) = n/d$. Writing $a = (n/d)b$ reduces the condition to $\gcd(b, d) = 1$ with $0 \leq b < d$, which has $\varphi(d)$ solutions. $\square$

The proposition explains why additive generator counting needs no orbit generation: $N_{\mathbb{Z}_n}(n) = \varphi(n)$. It also shows that $R(\mathbb{Z}_n) = \tau(n)$, the divisor-counting function, and $C(\mathbb{Z}_n) = \tau(n)$ because a cyclic group has one subgroup for every divisor of its order. These identities serve as exact checks on all additive computations.

The cyclic unit groups admit an equally explicit description once their cardinality is known. This observation is useful in the census because it separates richness caused by a divisor-rich group order from richness caused by a nontrivial direct product.

Proposition 2.2. *If $U(n)$ is cyclic and $m = \varphi(n)$, then*

$$N_{U(n)}(d) = \varphi(d) \quad (d \mid m), \quad R(U(n)) = C(U(n)) = \tau(m). \quad (9)$$

Moreover, the number of generators of $U(n)$ is $\varphi(m)$.

Proof. A cyclic group of order m contains one subgroup of order d for every divisor d of m . The generators of that subgroup are exactly the elements of order d , so their number is $\varphi(d)$. Taking $d = m$ gives the whole-group generator count, while counting the admissible divisors gives both richness and the number of cyclic subgroups. $\square$

The noncyclic cases can be stratified by the number of nontrivial invariant factors. Write

$$U(n) \cong C_{s_1} \times \cdots \times C_{s_{\rho(n)}}, \quad 1 < s_1 \mid \cdots \mid s_{\rho(n)}, \quad (10)$$

and call $\rho(n)$ the invariant-factor rank. For $n \geq 3$, every s_i is even. The rank therefore determines the population at the first nontrivial order.

Proposition 2.3. *For $n \geq 3$, the number of elements of order two in $U(n)$ is*

$$N_{U(n)}(2) = 2^{\rho(n)} - 1. \quad (11)$$

Proof. Each even cyclic factor C_{s_i} contains exactly two solutions of $x^2 = 1$: its identity and its unique involution. The direct product consequently contains $2^{\rho(n)}$ elements whose order divides two. Removing the identity leaves $2^{\rho(n)} - 1$ elements of exact order two. $\square$

Equation (11) turns a single spectral multiplicity into an immediate rank check. Conversely, any computed value $N_{U(n)}(2)$ must be one less than a power of two, and $\rho(n) = \log_2(N_{U(n)}(2) + 1)$. This condition is independent of the total-multiplicity and totient-divisibility checks and is applied below to interpret the rank distribution of the corpus.

3. Materials and methods

The principal corpus consists of the 299 consecutive moduli $n = 2, 3, \dots, 300$. For each modulus, the stored record contains the prime-power factorization, $\varphi(n)$, $\lambda(n)$, the unit-group invariant factors, both order spectra, both generator counts, both cyclic-subgroup counts, and the number of least-common-multiple combinations used during assembly. The range is exhaustive rather than sampled, so every prime, prime power, semiprime, and highly composite integer within the interval contributes to the comparisons.

Two arithmetic records were transcribed before the full census was computed. The console tables headed “computations in the group $\mathbb{Z}_{209}$ ” on pages 6–13 and “computations in the group $U(210)$ ”, group orders, inverses, element orders, generators, and cyclic subgroups. Their role is confined to the correspondence check summarized in table 1; no image, program listing, or prose from that article is reused. The citation in this paragraph is the sole attribution to those tabulated records.

Table 1: Arithmetic records used for correspondence checks.

Quantity	$\mathbb{Z}_{209}$	$U(210)$
Modulus factorization	$209 = 11 \cdot 19$	$210 = 2 \cdot 3 \cdot 5 \cdot 7$
Group cardinality	209	48
Residue set	$0, 1, \dots, 208$	$1, 11, 13, 17, \dots, 197, 199, 209$
Orders present	$1, 11, 19, 209$	$1, 2, 3, 4, 6, 12$
Whole-group generators	180	0
Inverse rule	$a \mapsto 209 - a$	$a \mapsto a^{-1} \pmod{210}$

3.1 Divisor-lattice computation for the additive group

After factorization, the divisors of n are generated by the Cartesian product of exponent ranges $0 \leq \beta_i \leq \alpha_i$. Each divisor $d = \prod p_i^{\beta_i}$ contributes the spectral pair $(d, \varphi(d))$ from equation (8). The calculation requires one record per divisor, not one record per residue. In particular, $n = 209$ has four divisors and therefore four additive spectral entries, even though the group has 209 elements.

The inverse map is recorded analytically. The additive inverse of a is $(-a) \bmod n$, represented in $\{0, \dots, n-1\}$ by 0 when $a = 0$ and $n - a$ otherwise. Because inversion does not affect the order distribution, inverse pairs are generated only when a human-readable case record is requested. This separation prevents a list of 209 inverse assignments from dominating a computation whose structural content occupies four rows.

3.2 Local cyclic spectra and least-common-multiple convolution

A cyclic group C_m has $\varphi(d)$ elements of order d for every $d \mid m$. Thus its spectrum is available from the same divisor calculation as the additive case. Suppose A and B are finite groups with known spectra. The order of $(a, b) \in A \times B$ is $\text{lcm}(\text{ord}(a), \text{ord}(b))$, so the product spectrum satisfies

$$N_{A \times B}(k) = \sum_{\substack{d, e \\ \text{lcm}(d, e) = k}} N_A(d) N_B(e). \quad (12)$$

The implemented convolution stores only nonzero dictionary entries. For every pair (d, e) in the two input dictionaries, it adds $N_A(d)N_B(e)$ to the output key $\text{lcm}(d, e)$. Repeating this operation over the cyclic factors of $U(n)$ produces the complete unit spectrum.

The construction is exact because it partitions every element of the direct product by its component orders. No probabilistic step is involved. It also remains interpretable: an output count can be traced to specific local order combinations. For $U(210)$, the local factors have orders 2, 4, and 6. Convolution first combines C_2 and C_4 , then combines that spectrum with C_6 . Twenty nonzero divisor-pair combinations are processed, yielding six final orders.

Procedure 1. Factorization-aware unit-order spectrum

1. Factor $n = \prod_i p_i^{\alpha_i}$ and construct the local cyclic-factor orders using equation (5).
2. Initialize the current spectrum as $S = \{1 \mapsto 1\}$.
3. For each cyclic-factor order m , form $T = \{d \mapsto \varphi(d) : d \mid m\}$.
4. Initialize an empty dictionary R . For each $(d, c_d) \in S$ and $(e, c_e) \in T$, add $c_d c_e$ to the entry $R[\text{lcm}(d, e)]$.
5. Replace S by R and repeat Step 3 until every cyclic factor has been processed.
6. Return S , where $S[d] = N_{U(n)}(d)$.

The algorithm returns multiplicities rather than representative elements. When representatives are needed, they can be recovered in a separate pass by selecting residues whose independently computed orders match the desired key. This design keeps the central calculation proportional to the local divisor structure.

3.3 Invariant-factor normalization

The local decomposition supplied by the Chinese remainder theorem is not necessarily in invariant-factor form. For example, $U(210)$ initially appears as $C_2 \times C_4 \times C_6$. To normalize a list of cyclic orders $m_1, \dots, m_t$, each m_j is factored. For every prime q , the valuations $v_q(m_j)$ are sorted in nondecreasing order and aligned by column. Multiplication down the columns yields integers $s_1 \mid s_2 \mid \dots \mid s_t$, and

$$C_{m_1} \times \dots \times C_{m_t} \cong C_{s_1} \times \dots \times C_{s_t}. \quad (13)$$

Trivial factors are discarded. In the $n = 210$ case, the sorted 2-adic valuations are $(1, 1, 2)$ and the single 3-adic valuation is right-aligned, giving $(2, 2, 12)$. Hence $U(210) \cong C_2 \times C_2 \times C_{12}$.

Invariant factors provide more than a conventional group label. Their product must equal $\varphi(n)$, their largest entry must equal $\lambda(n)$, and convolution of their cyclic spectra must reproduce the spectrum obtained from the local prime-power factors. These three equalities form a structural certificate independent of the order in which local components were assembled.

3.4 Residuewise order reducer and certification

An independent computation determines the order of every $a \in U(n)$ without repeated traversal of its full orbit. It begins with $L = \lambda(n)$. For each prime q dividing L , the exponent is divided by q as long as

$$a^{L/q} \equiv 1 \pmod{n}. \quad (14)$$

When no further division is possible, the remaining value is $\text{ord}_n(a)$. This calculation uses modular exponentiation and depends on the exponent theorem, whereas the convolution uses group decomposition and multiplicity arithmetic. Agreement between them is therefore meaningful: an error in local factor construction, convolution, residue filtering, or exponent reduction changes at least one spectral entry.

Five checks are applied for every n . The two independently computed dictionaries must be equal; the multiplicities must sum to $\varphi(n)$; every order must divide $\lambda(n)$; $N(d)$ must be divisible by $\varphi(d)$; and the maximum recorded order must equal the largest invariant factor. The additive dictionary is separately checked against $n/\text{gcd}(a, n)$ for all residues. All 299 moduli passed every condition.

3.5 Whole-corpus timing and work counts

Two kinds of computational evidence are reported. A deterministic work count compares the number of modular multiplications required by direct power-orbit traversal with the number of divisor-pair combinations used by least-common-multiple convolution. For a residue of order d , traversal performs d modular multiplications when it begins at the identity and stops at the first return. Summing $dN_{U(n)}(d)$ gives the traversal count. The convolution count is incremented once for each occupied pair of dictionary entries processed in equation (12).

Wall-clock timing was measured for the entire range $2 \leq n \leq 300$ in a single Python 3.12 process. Each procedure was executed 60 times after importing all modules. Medians and interquartile ranges are reported because sub-millisecond scheduling fluctuations affect individual runs. The code is intentionally direct: memoization across different moduli and compiled extensions were not used. These choices keep the comparison tied to the algebraic representation rather than to language-specific acceleration.

3.6 Derived census partitions and spectral mass

The stored integer records permit several secondary comparisons without another traversal of the residue classes. Four quantities are used to organize the additional census results. The first is the invariant-factor rank $\rho(n)$ from equation (10). The second is the exponent ratio

$$\eta(n) = \frac{\lambda(n)}{\varphi(n)}, \quad (15)$$

which equals one in the cyclic cases and decreases when independent factors enlarge the group without enlarging its exponent. The third is the density of elements attaining the exponent,

$$\mu(n) = \frac{N_{U(n)}(\lambda(n))}{\varphi(n)}. \quad (16)$$

This density is distinct from the generator density. When $U(n)$ is noncyclic, elements of order $\lambda(n)$ attain the largest possible order but do not generate the whole group.

The fourth quantity is the mean element order

$$\bar{o}(n) = \frac{1}{\varphi(n)} \sum_{d|\lambda(n)} dN_{U(n)}(d). \quad (17)$$

The numerator is already the deterministic multiplication count for complete power-orbit traversal, so equation (17) introduces no new enumeration. Normalization by $\lambda(n)$ gives $\bar{o}(n)/\lambda(n)$, the average fraction of the maximum attainable order realized by a unit. In a cyclic group of order m , equation (9) reduces the mean to

$$\bar{o}(n) = \frac{1}{m} \sum_{d|m} d\varphi(d), \quad m = \varphi(n). \quad (18)$$

The moduli were partitioned in two independent ways. The arithmetic partition uses $\omega(n)$, the number of distinct prime divisors of n . The structural partition uses $\rho(n)$. These partitions answer different questions: $\omega(n)$ records how many local prime-power groups enter the Chinese remainder product, whereas $\rho(n)$ records how many invariant factors remain after

components with coprime parts have been aligned. The two values need not agree because several local cyclic factors can merge into one invariant factor, while a single factor $U(2^\alpha)$ can contribute two factors when $\alpha \geq 3$.

Isomorphism classes were formed by grouping identical invariant-factor strings. All comparisons use exact strings obtained after divisibility normalization, not numerical proximity. A class containing several moduli therefore represents a certified equality of abstract unit groups. The order spectra, exponents, richness values, generator counts, and cyclic-subgroup counts must agree within each such class; this supplies a corpus-level check that is separate from the residuewise verification performed one modulus at a time.

4. Results and discussion

4.1 Correspondence at moduli 209 and 210

The additive case $209 = 11 \cdot 19$ gives divisors 1, 11, 19, and 209. Applying equation (8) yields

$$\mathcal{S}(\mathbb{Z}_{209}) = \{(1, 1), (11, 10), (19, 18), (209, 180)\}. \quad (19)$$

The multiplicities sum to 209, and the 180 elements of order 209 are precisely the additive generators. The group has four distinct cyclic subgroups, one for each divisor. This compact statement replaces hundreds of repetitive element-order and subgroup lines while retaining their complete numerical content. It also explains the pattern: multiples of 19 have order 11, multiples of 11 have order 19, and residues coprime to 209 have full order.

The unit case $210 = 2 \cdot 3 \cdot 5 \cdot 7$ has 48 residues. Its local decomposition is

$$U(210) \cong U(2) \times U(3) \times U(5) \times U(7) \cong C_2 \times C_4 \times C_6 \cong C_2 \times C_2 \times C_{12}. \quad (20)$$

The exponent is 12, so no element can have order 48 and the group cannot be cyclic. The computed multiplicities are shown in figure 1: one identity, seven involutions, two elements of order 3, eight of order 4, fourteen of order 6, and sixteen of order 12. Their total is 48. Division by the corresponding totients gives $1 + 7 + 1 + 4 + 7 + 4 = 24$ distinct cyclic subgroups.




Figure 1: Element-order distribution in $U(210)$.

The dominant class is order 12, containing one third of the group. Nevertheless, those sixteen elements generate only four cyclic subgroups because each order-12 subgroup has $\varphi(12) = 4$ generators. The seven involutions each generate a different subgroup of order 2. This contrast illustrates why an element count and a subgroup count should not be conflated: high multiplicity at an order may arise from many small generator classes or from a few large ones.

The certified quantities in table 2 reveal a reversal. The larger additive group is structurally simpler: it is cyclic, has four possible orders, and has one subgroup for each divisor. The smaller unit group is noncyclic, has six possible orders, and has 24 cyclic subgroups. Cardinality alone therefore predicts neither order richness nor subgroup multiplicity.

Table 2: Certified structural quantities for the two detailed cases.

Quantity	$\mathbb{Z}_{209}$	$U(210)$
Cardinality	209	48
Exponent	209	12
Distinct element orders	4	6
Elements of maximum order	180	16
Whole-group generators	180	0
Distinct cyclic subgroups	4	24
Stored spectral pairs	4	6

4.2 Order richness across all moduli

The census contains a median group order $\varphi(n) = 80$, a median additive richness of four, and a median unit richness of six. The additive value follows the relatively small divisor counts of integers below 301. Unit richness responds instead to the divisors of local values $(p-1)p^{\alpha-1}$ and to their least-common-multiple interactions. The distribution across the full modulus axis is displayed in figure 2; point area reflects $\varphi(n)$ and color records the number of distinct prime divisors of n .




Figure 2: Unit-order richness over the modulus range.

The highest richness is 20 at $n = 241$. Since 241 is prime, $U(241) \cong C_{240}$, and the 20 attainable orders are exactly the 20 divisors of $240 = 2^4 \cdot 3 \cdot 5$. The next value is 18 at $n = 181$, where $U(181) \cong C_{180}$ and 180 has 18 divisors. These peaks arise from divisor-rich predecessors of primes, not from moduli with many prime factors. Composite moduli can also be rich: $n = 287 = 7 \cdot 41$ gives $U(287) \cong C_2 \times C_{120}$ and 16 attainable orders. Its multiplicities differ from those of a cyclic group of order 240 even though both groups have the same cardinality.

The prime cases also clarify the difference between richness and the population of full-order elements. A cyclic group C_m has $\tau(m)$ order classes but only $\varphi(m)$ generators. At $n = 241$, richness is maximal because 240 has many divisors, whereas the 64 primitive residues constitute only 26.7% of $U(241)$. At $n = 263$, richness is only four because $262 = 2 \cdot 131$, yet 130 residues have full order and constitute nearly half of the group. Thus a richer order vocabulary need not imply a larger relative supply of generators. The spectrum displays both facts at once and prevents the number of order classes from being mistaken for cyclic robustness.

Color does not vary monotonically with richness. Moduli with three or four distinct prime factors often have moderate spectra because repeated prime factors in the local cyclic orders collapse under the least common multiple. The case $n = 210$ illustrates this effect: four primes produce only six orders. By contrast, a single prime p can produce many orders when $p-1$ contains several prime powers. Arithmetic diversity in the modulus and diversity in the order spectrum are therefore different notions.


Table 3: Selected records from the 299-modulus census.

n	Factorization	Invariant factors of $U(n)$	$\varphi(n)$	$\lambda(n)$	$R(\mathbb{Z}_n)$	$R(U(n))$	$C(U(n))$	$Q(U(n))$
8	2^3	2×2	4	2	4	2	4	2.00
15	$3 \cdot 5$	2×4	8	4	4	3	6	2.67
45	$3^2 \cdot 5$	2×12	24	12	6	6	12	4.00
105	$3 \cdot 5 \cdot 7$	$2 \times 2 \times 12$	48	12	8	6	24	8.00
209	$11 \cdot 19$	2×90	180	90	4	12	24	15.00
210	$2 \cdot 3 \cdot 5 \cdot 7$	$2 \times 2 \times 12$	48	12	16	6	24	8.00
240	$2^4 \cdot 3 \cdot 5$	$2 \times 2 \times 4 \times 4$	64	4	20	3	40	21.33
241	241	240	240	240	2	20	20	12.00
263	263	262	262	262	2	4	4	65.50
273	$3 \cdot 7 \cdot 13$	$2 \times 6 \times 12$	144	12	8	6	60	24.00

The selected records in table 3 make several distinctions concrete. The isomorphic unit groups $U(105)$ and $U(210)$ have identical invariant factors and hence identical spectra, even though their additive groups have different richness. The unit groups at $n = 209$ and $n = 297$ are also isomorphic to $C_2 \times C_{90}$, showing that the same multiplicative structure can arise from unrelated modulus factorizations. Conversely, equal totients do not guarantee equal unit groups: the invariant factors and spectra are needed to discriminate them.

4.3 Cyclicity and prime-power structure

Exactly 114 of the 299 unit groups are cyclic. The classical classification states that $U(n)$ is cyclic precisely for $n = 2, 4, p^\alpha$, or $2p^\alpha$ with p an odd prime. The computation recovers this classification without hard-coding the final decision: cyclicity is declared only when the maximum spectral order equals $\varphi(n)$. Of the 114 cyclic cases, 73 have one distinct prime divisor and 41 have two, the second set being the even values $2p^\alpha$.

Figure 3: Cyclicity of $U(n)$ for $2 \leq n \leq 300$.

The block display in figure 3 shows cyclic cases as dark cells. Their irregular spacing follows the occurrence of primes and odd prime powers, while the adjacent even cases account for many short horizontal pairings. No modulus with three or four distinct prime factors is marked cyclic. The visual pattern is not a periodic residue phenomenon; it is a compact image of the classification through prime-power form.

The exponent deficit $\varphi(n)/\lambda(n)$ quantifies how far a unit group is from containing an element whose order equals its cardinality. The largest deficit in the range occurs at $n = 240$, where $\varphi(240) = 64$ and $\lambda(240) = 4$. Its invariant factors are

2, 2, 4, 4, and its spectrum has only orders 1, 2, and 4. Thus a group of 64 elements is represented by three order classes. The deficit does not by itself count cyclic subgroups: $U(240)$ has 40 of them because its many involutions and order-4 elements split into numerous generated subgroups.

4.4 Spectral compression and subgroup multiplicity

Compression increases when a large group is concentrated in a few order classes. The relation between $\varphi(n)$ and $Q(U(n))$ appears in figure 4. Cyclic and noncyclic groups overlap substantially because a cyclic group may itself have few possible orders if its cardinality has few divisors. The maximum $Q = 65.5$ occurs at the prime modulus $n = 263$: $U(263) \cong C_{262}$, and $262 = 2 \cdot 131$ has only four divisors. Its 262 elements are therefore represented by the four entries for orders 1, 2, 131, and 262.




Figure 4: Compression of unit-group order records.

Large noncyclic compression has a different origin. At $n = 240$, the small exponent 4 limits the spectrum to three entries; at $n = 273$, exponent 12 limits it to six. The latter group has 60 cyclic subgroups, the largest count in the census. Its decomposition $C_2 \times C_6 \times C_{12}$ creates many elements at the same order but distributes them across multiple subgroups. This finding shows why a compressed order record should be accompanied by the subgroup total: compression describes repeated order values, whereas $C(G)$ describes how those values partition into generated sets.

The subgroup maximum at $n = 273$ is not attached to the largest group in the interval. With 144 units, it is smaller than many prime-modulus unit groups, but its three invariant factors create numerous independent choices of low-order components. For each attainable order, several tuples generate different cyclic subsets even when their least common multiple is the same. This is the direct-product mechanism behind the count of 60. A cyclic group of order 144, by contrast, would have one cyclic subgroup for each of the 15 divisors of 144. The fourfold difference demonstrates that subgroup proliferation is a signature of product structure rather than a simple consequence of cardinality.

No scientific conclusion depends on the graphic scale alone. The underlying CSV files retain every integer count, and the identities $\sum_d N(d) = \varphi(n)$ and equation (7) can reconstruct the totals in the text. The plotted logarithmic vertical axis is used only to keep small and large compression values legible on one page.

4.5 Additive–multiplicative contrast

The two groups associated with a modulus respond to different divisor structures. Additive richness is $\tau(n)$, whereas unit richness depends on the direct product of local cyclic groups. Across the 299 moduli, unit richness exceeds additive richness in 169 cases, equals it in 30 cases, and is smaller in 100 cases. Their Pearson correlation is -0.288 , indicating a mild inverse association over this finite range rather than a shared increasing trend.

Prime moduli form the vertical band at additive richness two in figure 5. Along that band, unit richness ranges from two to twenty according to the divisor structure of $p - 1$. Highly composite moduli move to the right because $\tau(n)$ grows,




Figure 5: Contrast between additive and unit-order richness.

but their unit richness may remain small when $\lambda(n)$ is a low power-smooth number. The point $n = 240$ lies far below the diagonal with additive richness 20 and unit richness 3. At the opposite extreme, $n = 241$ has additive richness 2 and unit richness 20. Consecutive moduli thus realize the strongest visible reversal in the census.

This reversal answers a central part of the research question. Knowledge of the additive group does not provide a useful proxy for the complexity of the unit group, even though both arise from the same ring. The relevant arithmetic changes from divisors of n to the prime-power unit orders and their least common multiples. Consequently, software that displays the two groups through identical menu options may obscure the fact that they require different structural explanations.

Equal-richness cases reinforce that point. Thirty moduli lie on the diagonal of figure 5, but equality of $R(\mathbb{Z}_n)$ and $R(U(n))$ rarely means that the same orders occur. The additive orders divide n ; the unit orders divide $\lambda(n)$. For $n = 45$, both richness values equal six, yet the additive spectrum contains divisors of 45 while the unit spectrum is controlled by $U(45) \cong C_2 \times C_{12}$. A single scalar comparison must therefore be read with the actual order keys and multiplicities. The visual contrast is valuable for locating unusual cases, while the spectral dictionaries supply the mathematical explanation.

4.6 Computational work and timing

Direct traversal requires 1,587,655 modular multiplication steps over the complete corpus. Least-common-multiple convolution processes 4,425 occupied divisor-pair combinations, a ratio of 358.79 to one. These counts are not claimed to be identical machine instructions: modular multiplication and integer least common multiple have different costs. They do establish that the algebraic construction avoids the dominant repetition of following every residue until it returns to the identity.

Table 4: Whole-corpus computational comparison.

Procedure	Work units	Median time (ms)	IQR (ms)
Power-orbit traversal	1,587,655 multiplications	75.580	10.384
CRT order convolution	4,425 order-pair combinations	4.253	0.512

The timing measurements in table 4 give a median ratio of 17.77. The smaller wall-clock ratio relative to the work-count ratio is expected: factorization, divisor generation, dictionary allocation, and repeated least-common-multiple calls contribute fixed overhead. The interquartile range is also smaller for convolution, which performs a more uniform number of operations across repetitions. For $n = 210$ alone, orbit traversal requires 329 multiplications, while convolution uses 20 occupied order-pair combinations.

The residuewise exponent reducer is retained as a certification route rather than as the primary spectrum constructor. It performs well for these small moduli and identifies representative elements, but it still executes once per unit. The

factorized spectrum can answer cardinality, exponent, cyclicity, generator count, and cyclic-subgroup count before a single global residue is visited. This difference becomes more important when the desired output is structural rather than a list of representatives.

Memory use follows the same distinction. Orbit output grows at least linearly with the number of units and grows further when every generated subgroup is printed repeatedly for each of its generators. Spectral storage grows with the number of attainable orders, which never exceeds the number of divisors of the exponent in the cyclic case and is often much smaller than $\varphi(n)$. Within the present corpus, the largest dictionary has only 20 entries. This compactness makes the equality checks inexpensive and allows all 299 group records to be retained together for cross-modulus comparisons.

4.7 Prime-power trajectories in the unit spectrum

Prime powers isolate the local components from which the CRT construction is assembled. The powers of two are especially informative because their unit groups are noncyclic from $n = 8$ onward and acquire a regular sequence of invariant factors. For $n = 2^\alpha$ with $\alpha \geq 3$, the decomposition $U(2^\alpha) \cong C_2 \times C_{2^{\alpha-2}}$ implies that the exponent doubles whenever α increases by one, while the group cardinality also doubles. The number of distinct element orders grows only linearly. The computed trajectory from 8 through 256 is summarized in table 5.

Table 5: Unit-group trajectory along powers of two.

n	$ U(n) $	$\lambda(n)$	Invariant factors	$R(U(n))$	$C(U(n))$	$Q(U(n))$
8	4	2	2×2	2	4	2.00
16	8	4	2×4	3	6	2.67
32	16	8	2×8	4	8	4.00
64	32	16	2×16	5	10	6.40
128	64	32	2×32	6	12	10.67
256	128	64	2×64	7	14	18.29

The progression in table 5 has a direct spectral explanation. The identity is the unique element of order one, and there are three elements of order two for every $\alpha \geq 3$. For each j with $2 \leq j \leq \alpha - 2$, the number of elements of order 2^j is 2^j . Hence the spectrum of $U(256)$ is

$$\{(1, 1), (2, 3), (4, 4), (8, 8), (16, 16), (32, 32), (64, 64)\}. \quad (21)$$

The multiplicities sum to 128, as required. The order richness is $\alpha - 1$, the number of cyclic subgroups is $2\alpha - 2$, and the compression is

$$Q(U(2^\alpha)) = \frac{2^{\alpha-1}}{\alpha - 1}. \quad (22)$$

This quotient grows rapidly because the cardinality is exponential in α while the number of occupied order classes is linear. The absence of a whole-group generator persists throughout the trajectory: the exponent $2^{\alpha-2}$ is always half of $\varphi(2^\alpha) = 2^{\alpha-1}$.

Odd prime powers behave differently. For an odd prime p , $U(p^\alpha)$ is cyclic of order $(p - 1)p^{\alpha-1}$, so its richness and its number of cyclic subgroups both equal the divisor count of that order. The sequence $n = 9, 27, 81, 243$ gives group orders 6, 18, 54, 162 and richness values 4, 6, 8, 10, respectively. The number of whole-group generators is 2, 6, 18, 54, which equals $\varphi(\varphi(n))$ in each case. Powers of five show the same cyclic pattern: $U(25) \cong C_{20}$ has six possible orders and eight generators, whereas $U(125) \cong C_{100}$ has nine possible orders and forty generators. These local trajectories explain how the global spectrum changes when a prime exponent in the modulus is increased without introducing a new prime divisor.

The contrast between the even and odd trajectories also clarifies the role of local noncyclicity. Increasing the exponent of an odd prime enlarges a single cyclic factor. Increasing the exponent of two beyond the second power preserves an independent factor C_2 beside the growing power-of-two factor. During CRT convolution, that persistent component changes multiplicities even when it does not introduce a new maximum order. The invariant-factor record therefore captures information that the exponent alone cannot supply.

4.8 Equal totients and structurally different unit groups

Euler's totient determines the cardinality of $U(n)$ but does not determine its isomorphism type. Within the census, 27 distinct totient values are associated with more than one invariant-factor sequence; together these values account for 177 of the 299 moduli. The value $\varphi(n) = 48$ gives a compact example because it occurs with three different unit-group structures before $n = 300$. Four representatives are compared in table 6.

Table 6: Distinct unit groups with cardinality 48.

n	Factorization	Invariant factors	Order spectrum $d^{N(d)}$	$\lambda(n)$	$C(U(n))$
65	$5 \cdot 13$	4×12	$1^1, 2^3, 3^2, 4^{12}, 6^6, 12^{24}$	12	20
104	$2^3 \cdot 13$	$2 \times 2 \times 12$	$1^1, 2^7, 3^2, 4^8, 6^{14}, 12^{16}$	12	24
168	$2^3 \cdot 3 \cdot 7$	$2 \times 2 \times 2 \times 6$	$1^1, 2^{15}, 3^2, 6^{30}$	6	32
210	$2 \cdot 3 \cdot 5 \cdot 7$	$2 \times 2 \times 12$	$1^1, 2^7, 3^2, 4^8, 6^{14}, 12^{16}$	12	24

The comparison in table 6 shows three levels of discrimination. Cardinality alone places all four groups in the same class. Adding the exponent separates $U(168)$ from the other three but leaves $U(65)$, $U(104)$, and $U(210)$ unresolved. Their order richness is also identical at six, and their compression is consequently the same value, eight. The multiplicities then separate $U(65)$ from the other two: it has three involutions rather than seven, twelve elements of order four rather than eight, six elements of order six rather than fourteen, and twenty-four elements of order twelve rather than sixteen. The invariant factors agree with this distinction, giving $C_4 \times C_{12}$ for $n = 65$ and $C_2 \times C_2 \times C_{12}$ for $n = 104$ and $n = 210$.

The equality between the spectra at 104 and 210 is not accidental. Their matching invariant factors establish an isomorphism of unit groups, even though one modulus is $2^3 \cdot 13$ and the other is square-free with four prime divisors. The same invariant-factor sequence $2 \times 2 \times 12$ occurs at eight moduli in the interval: 104, 105, 112, 140, 144, 156, 180, and 210. Every one of these groups has the order distribution displayed for $n = 104$ and $n = 210$, as well as 24 cyclic subgroups. The modulus factorization determines how the local components arise, while invariant-factor normalization determines the final abstract group.

The $n = 65$ and $n = 104$ comparison demonstrates why neither $\varphi(n)$, $\lambda(n)$, order richness, nor compression should be used alone as a group signature. All four quantities agree, but the groups are not isomorphic and their order multiplicities differ. In this case the full spectrum detects the difference. The result supports the decision to retain both the multiplicity dictionary and the invariant factors in the computational record. The spectrum supplies an immediately checkable distribution, while the invariant factors provide the canonical structural description.

4.9 Factor-count strata and the location of cyclic unit groups

The number of distinct primes in the modulus controls the number of local unit groups entering the Chinese remainder decomposition, but it does not determine how many invariant factors survive normalization. The arithmetic partition by $\omega(n)$ is summarized in table 7. A single prime divisor occurs for 79 moduli, two prime divisors occur for 162 moduli, three occur for 57 moduli, and only $n = 210$ has four distinct prime divisors in the stated interval. Thus more than half of the corpus lies in the two-prime stratum, making this class the principal source of structural variety.

Table 7: Unit-group summaries by distinct prime count.

$\omega(n)$	Moduli	Cyclic	Cyclic (%)	Median $R(U(n))$	Maximum $R(U(n))$	Median $C(U(n))$
1	79	73	92.4	8	20	8
2	162	41	25.3	6	16	16
3	57	0	0.0	6	12	20
4	1	0	0.0	6	6	24

The cyclic counts in table 7 agree exactly with the arithmetic classification of cyclic unit groups. Within $2 \leq n \leq 300$, the 114 cyclic cases consist of $n = 2$, $n = 4$, 71 odd prime powers, and 41 integers of the form $2p^\alpha$ with p odd and prime. The six noncyclic members of the one-prime stratum are 8, 16, 32, 64, 128, and 256. Their exceptional behavior is precisely the two-factor decomposition in equation (5). Every cyclic member of the two-prime stratum is twice an odd prime power; the remaining 121 members of that stratum are noncyclic. Three or more distinct prime divisors necessarily supply at least two independent even local components, so no such unit group in the corpus is cyclic.

The richness columns show that adding prime divisors does not systematically increase the number of attainable orders. The one-prime stratum has the largest median richness and contains the global maximum at $n = 241$. Its high values arise when a cyclic order $\varphi(p^\alpha)$ has many divisors. The median drops from eight to six when $\omega(n)$ increases from one to two, because direct-product components can share the same prime divisors and many local order pairs collapse to a common least common multiple. The three-prime and four-prime strata retain median or observed richness six even though their median

cyclic-subgroup totals increase. Additional components therefore produce more subgroups without necessarily producing more order labels.

The distinction is visible at the two ends of the table. A cyclic group in the one-prime stratum has exactly one cyclic subgroup for each occupied order. In contrast, $U(210)$ has richness six and 24 cyclic subgroups. Its order labels are not numerous, but the multiplicity attached to each label splits across several different generated subsets. This is the general reason that the subgroup median increases with $\omega(n)$ while richness does not.

4.10 Invariant-factor rank and involution multiplicity

The structural partition by $\rho(n)$ gives the distribution in table 8. The trivial group $U(2)$ is assigned rank zero. Every rank-one group is cyclic, while every group of rank at least two is noncyclic. Rank two is the most frequent structural type, accounting for 133 of the 299 moduli. The 47 rank-three groups and five rank-four groups constitute a smaller but important portion of the census because they contain the largest populations of involutions and many of the largest cyclic-subgroup totals.

Table 8: Invariant-factor rank and unit-group multiplicity.

$\rho(n)$	Moduli	$N_{U(n)}(2)$	Median R	Median C	Maximum C	Maximizing modulus
0	1	0	1	1	1	2
1	113	1	8	8	20	241
2	133	3	6	16	48	247, 259
3	47	7	6	24	60	273
4	5	15	4	32	48	280

The frequency distribution of involution counts is identical to the rank distribution: one group has no involution, 113 have one, 133 have three, 47 have seven, and five have fifteen. This exact agreement is the corpus realization of equation (11). It provides a particularly transparent structural diagnostic. A recorded count such as seven does not merely state that many residues square to one; it proves that three nontrivial invariant factors are present. The fifteen involutions at $n = 120, 168, 240, 264$, and 280 similarly certify rank four.

Rank changes the relationship between richness and subgroup multiplicity. The median richness is six at ranks two and three, yet the median number of cyclic subgroups rises from 16 to 24. Rank four has an even smaller median richness of four but a subgroup median of 32. Independent factors create multiple elements with the same least-common-multiple order, and those elements need not lie in the same cyclic subgroup. The order dictionary remains short while the generated-subgroup structure becomes more numerous.

The maximum subgroup count is not monotone in rank. Rank three reaches the corpus maximum of 60 at $n = 273$, whereas rank four reaches 48 at $n = 280$. The group $U(273) \cong C_2 \times C_6 \times C_{12}$ has six occupied orders and substantial multiplicities at orders 6 and 12. The group $U(280) \cong C_2 \times C_2 \times C_2 \times C_{12}$ has the same richness, but the replacement of C_6 by two factors of order two changes the distribution of order-three and order-six elements. Rank measures the number of independent cyclic coordinates; it does not record their sizes, so the invariant factors and the complete spectrum remain necessary after rank is known.

4.11 Exponent deficiency and concentration at the largest order

The ratio $\eta(n)$ measures how far the exponent lies below the group cardinality. The selected records in table 9 include the smallest ratios in the corpus, cases with unusually high maximum-order density, and cyclic comparisons. They show that exponent deficiency and concentration at the exponent are logically independent. A group can have an exponent much smaller than its cardinality while most of its elements attain that exponent.

The minimum exponent ratio occurs at $n = 240$, where $\lambda(240) = 4$ and $\varphi(240) = 64$. Its spectrum is $1^1, 2^{15}, 4^{48}$, so three quarters of all units attain the exponent despite the ratio $\eta(240) = 1/16$. At $n = 252$, the exponent is only one twelfth of the group cardinality, but 56 of 72 units have order six. The comparison with $n = 273$ is instructive: both groups have $\eta = 1/12$, yet their maximum-order densities are $7/9$ and $4/9$, respectively. The additional factor of order 12 in $U(273)$ expands the set of lower orders and redistributes spectral mass even though the exponent ratio is unchanged.

The high density at $n = 24$ has the simplest possible explanation. The group C_2^3 has exponent two, and every nonidentity element attains it, giving $\mu(24) = 7/8$. At $n = 184$, the group $C_2 \times C_2 \times C_{22}$ retains the same three independent involution coordinates but includes an order-11 component; 70 of its 88 elements consequently attain order 22. These examples show that a large value of $\mu(n)$ is compatible with noncyclicity and may be produced by several independent coordinates.

Table 9: Exponent ratio and maximum-order density.

n	Factorization	Invariant factors	$\varphi(n)$	$\lambda(n)$	$\eta(n)$	$N(\lambda(n))$	$\mu(n)$
24	$2^3 \cdot 3$	$2 \times 2 \times 2$	8	2	0.2500	7	0.8750
184	$2^3 \cdot 23$	$2 \times 2 \times 22$	88	22	0.2500	70	0.7955
240	$2^4 \cdot 3 \cdot 5$	$2 \times 2 \times 4 \times 4$	64	4	0.0625	48	0.7500
252	$2^2 \cdot 3^2 \cdot 7$	$2 \times 6 \times 6$	72	6	0.0833	56	0.7778
273	$3 \cdot 7 \cdot 13$	$2 \times 6 \times 12$	144	12	0.0833	64	0.4444
275	$5^2 \cdot 11$	10×20	200	20	0.1000	96	0.4800
211	211	210	210	210	1.0000	48	0.2286
241	241	240	240	240	1.0000	64	0.2667
263	263	262	262	262	1.0000	130	0.4962

For cyclic groups, $\eta(n) = 1$, but the density at the exponent varies with the arithmetic of $\varphi(n)$. It equals $\varphi(\varphi(n))/\varphi(n)$. The value is only 48/210 at $n = 211$ because 210 has four distinct prime divisors. At $n = 263$, the group order is $262 = 2 \cdot 131$, and 130/262 units attain order 262. Thus cyclicity fixes the exponent ratio but not the proportion of primitive elements. The pair (η, μ) separates the size of the maximum order from the population that realizes it.

4.12 Mean element order and the location of spectral mass

The mean in equation (17) distinguishes spectra that have the same richness but place different numbers of elements near the exponent. The largest absolute means in the census are listed in table 10. Prime moduli dominate because their unit groups are cyclic and their exponents can be close to 300. The ranking is nevertheless not identical to the ranking by modulus or group cardinality, since the divisor structure of the cyclic order controls how much mass lies at smaller orders.

Table 10: Largest mean unit-element orders.

n	$\varphi(n)$	$\lambda(n)$	$\sum_d dN(d)$	$\bar{o}(n)$	$\bar{o}(n)/\lambda(n)$	$R(U(n))$
293	292	292	57,827	198.038	0.678	6
263	262	262	51,093	195.011	0.744	4
269	268	268	48,653	181.541	0.677	6
289	272	272	46,683	171.629	0.631	10
257	256	256	43,691	170.668	0.667	9
227	226	226	37,971	168.013	0.743	4

The maximum occurs at $n = 293$. Although only 144 of the 292 units have order 292, the remaining spectral mass is still concentrated at comparatively large divisors of 292, producing a mean of 198.038. The prime modulus 263 has a slightly smaller absolute mean but a larger normalized mean. Its four-entry spectrum $1^1, 2^1, 131^{130}, 262^{130}$ places all but two elements at one half of the exponent or at the exponent itself. This gives $\bar{o}(263)/262 = 0.744$, compared with 0.678 at $n = 293$.

The rows for 257 and 289 illustrate the effect of prime powers in the group order and modulus. The group $U(257) \cong C_{256}$ has nine occupied orders, one for each divisor of 2^8 , and its multiplicities double regularly along the order chain. The modulus $289 = 17^2$ gives $U(289) \cong C_{272}$; the factorization $272 = 2^4 \cdot 17$ produces ten occupied orders and a normalized mean of 0.631. Both groups are cyclic, but their divisor lattices distribute elements differently below the exponent.

Absolute and normalized means answer different questions. The absolute mean identifies groups whose typical unit has a long orbit and therefore explains the largest contributions to the total of 1,587,655 traversal multiplications. The normalized mean measures concentration relative to the group-specific maximum. For example, $U(240)$ has the small absolute mean $223/64 = 3.484$ but the large normalized mean 0.871, because 48 of its 64 units have the maximum order four. Excluding the trivial group $U(2)$, the largest normalized mean is $15/16 = 0.938$ at $n = 24$: seven elements have order two and only the identity has order one. A small exponent can therefore coexist with strong concentration at that exponent, while a large exponent can coexist with a substantial population of shorter periods.

The mean also links the mathematical and computational interpretations of the spectrum. Multiplying $\bar{o}(n)$ by $\varphi(n)$ recovers the exact number of modular multiplications used by full orbit traversal at modulus n . Hence the same statistic

describes both the average algebraic period and the deterministic work avoided by factorization-aware convolution. No timing assumption is needed for this relation.

4.13 Repeated invariant-factor classes

The 299 unit groups occupy 142 distinct invariant-factor classes. Fifty-seven classes occur at one modulus, while 85 classes recur and account for the remaining 242 moduli. The most frequent classes are displayed in table 11. Their recurrence demonstrates that very different modulus factorizations can lead to the same abstract group and therefore to the same order spectrum.

Table 11: Most frequent unit-group isomorphism classes.

Invariant factors	Frequency	Moduli
$2 \times 2 \times 12$	8	104, 105, 112, 140, 144, 156, 180, 210
2×12	7	35, 39, 45, 52, 70, 78, 90
2×30	7	77, 93, 99, 124, 154, 186, 198
2×36	7	95, 111, 135, 148, 190, 222, 270
2×60	7	143, 155, 175, 183, 225, 244, 286
2×42	6	129, 147, 172, 196, 258, 294
2×20	5	55, 75, 100, 110, 150
$2 \times 2 \times 20$	5	165, 176, 200, 220, 300

The class $C_2 \times C_2 \times C_{12}$ has already appeared at 104 and 210, but table 11 shows that it is shared by six additional moduli. Some coincidences preserve the number of distinct prime divisors, while others do not. For instance, $105 = 3 \cdot 5 \cdot 7$ and $144 = 2^4 \cdot 3^2$ have quite different local decompositions, yet normalization yields the same three invariant factors. Their unit groups consequently share the spectrum $1^1, 2^7, 3^2, 4^8, 6^{14}, 12^{16}$ and the same 24 cyclic subgroups.

The seven occurrences of $C_2 \times C_{60}$ provide another instructive class. The moduli include the semiprimes 143 and 155, the mixed prime power $175 = 5^2 \cdot 7$, the product $183 = 3 \cdot 61$, the squareful modulus $225 = 3^2 \cdot 5^2$, and the even values 244 and 286. Each local factorization supplies cyclic components through a different route, but their aligned prime valuations give the same final pair. Once the class is known, the complete spectrum and every derived quantity can be reused without recomputing local convolutions.

Class recurrence also clarifies what information is lost when only n , $\varphi(n)$, and $\lambda(n)$ are reported. Those three values identify neither the class frequency nor the moduli sharing it. Invariant factors supply the canonical structural key, while the spectrum supplies an auditable numerical signature. Agreement of every stored spectrum within each of the 85 recurrent classes was verified directly. This class-level equality complements the element-level certification because it compares computations reached from different factorizations.

4.14 Where convolution work is concentrated

The total convolution count of 4,425 order-pair combinations is distributed unevenly across the corpus. The moduli requiring the most combinations are shown in table 12. Even the maximum of 42 pairs is small relative to the corresponding power-orbit count, but variation among the rows reveals how the local decomposition affects the route to an identical final group.

Table 12: Largest modulus-level convolution workloads.

n	Factorization	Invariant factors	$R(U(n))$	CRT pairs	Orbit multiplications	Orbit/CRT
247	$13 \cdot 19$	6×36	9	42	4,301	102.40
259	$7 \cdot 37$	6×36	9	40	4,301	107.53
217	$7 \cdot 31$	6×30	8	36	3,675	102.08
221	$13 \cdot 17$	4×48	10	36	4,865	135.14
279	$3^2 \cdot 31$	6×30	8	36	3,675	102.08
287	$7 \cdot 41$	2×120	16	36	12,789	355.25
273	$3 \cdot 7 \cdot 13$	$2 \times 6 \times 12$	6	34	1,175	34.56
263	263	262	4	4	51,093	12,773.25

The first two rows reach the same invariant factors 6×36 and therefore have identical final spectra and identical orbit counts. Their convolution workloads differ because $U(247)$ begins with local cyclic orders 12 and 18, whereas $U(259)$ begins

with orders 6 and 36. The occupied divisor sets have different sizes before equal least-common-multiple keys are merged. This difference confirms that the work count describes the construction route, not only the final isomorphism type.

The pair $n = 217$ and $n = 279$ shows the complementary situation. Both normalize to $C_6 \times C_{30}$ and both require 36 combinations, even though the local origins are $7 \cdot 31$ and $3^2 \cdot 31$. The equality of their final spectra, orbit sums, and construction counts provides a complete computational coincidence. At $n = 287$, richness 16 is the largest among the high-workload composite cases, yet 36 pairs suffice to construct a spectrum whose direct traversal requires 12,789 multiplications.

The prime modulus 263 demonstrates why large orbit work does not imply large convolution work. Its cyclic unit group has only four spectral entries. Constructing those entries requires four divisor records, whereas orbit traversal follows 130 elements through 131 steps and another 130 through 262 steps, in addition to the two short orders. The modulus-level ratio is 12,773.25, substantially larger than the whole-corpus ratio because the factorized route is especially compact when a large cyclic order has few divisors.

These records also identify a practical reuse opportunity. Caching a spectrum by normalized invariant factors would avoid repeated convolution for the recurrent classes in table 11; caching local prime-power spectra would avoid repeated divisor construction before normalization. The present timings deliberately exclude such cross-modulus memoization, so the reported performance does not depend on reuse. The class frequencies show, however, that a batch implementation could exploit it without changing any mathematical output.

4.15 Structural discontinuities among neighboring moduli

Nearby integers can have sharply different unit groups because factorization changes discontinuously with the modulus. The four consecutive triples in table 13 hold numerical size nearly fixed while moving between prime, prime-power, and composite structures. They provide local evidence that the large corpus patterns are arithmetic effects rather than a gradual consequence of increasing n .

Table 13: Unit-group changes across neighboring moduli.

n	Factorization	Invariant factors	$\varphi(n)$	$\lambda(n)$	$R(U(n))$	$C(U(n))$	$Q(U(n))$	Cyclic
239	239	238	238	238	8	8	29.75	yes
240	$2^4 \cdot 3 \cdot 5$	$2 \times 2 \times 4 \times 4$	64	4	3	40	21.33	no
241	241	240	240	240	20	20	12.00	yes
262	$2 \cdot 131$	130	130	130	8	8	16.25	yes
263	263	262	262	262	4	4	65.50	yes
264	$2^3 \cdot 3 \cdot 11$	$2 \times 2 \times 2 \times 10$	80	10	4	32	20.00	no
272	$2^4 \cdot 17$	$2 \times 4 \times 16$	128	16	5	36	25.60	no
273	$3 \cdot 7 \cdot 13$	$2 \times 6 \times 12$	144	12	6	60	24.00	no
274	$2 \cdot 137$	136	136	136	8	8	17.00	yes
287	$7 \cdot 41$	2×120	240	120	16	32	15.00	no
288	$2^5 \cdot 3^2$	$2 \times 2 \times 24$	96	24	8	32	12.00	no
289	17^2	272	272	272	10	10	27.20	yes

The first triple contains the strongest additive–multiplicative reversal in the study. The prime moduli 239 and 241 have cyclic unit groups, but their group orders have eight and twenty divisors, respectively. Between them, $U(240)$ has four invariant factors and exponent four. Its richness falls to three while its cyclic-subgroup count rises to 40. A change of one in the modulus therefore moves from a cyclic group of order 238 to a rank-four group of order 64 and then to a cyclic group of order 240.

The triple 262–264 separates compression from group size. Both 262 and 263 give cyclic groups, yet 130 has eight divisors whereas $262 = 2 \cdot 131$ has only four. The latter group is more than twice as compressed even though the moduli are consecutive. At 264, three independent factors restrict the exponent to 10 and eliminate whole-group generators, while the cyclic-subgroup count increases from four to 32. The order richness remains four, showing again that equal dictionary length does not imply similar internal organization.

The transition 272–274 changes from a rank-three even prime-power product to a rank-three odd square-free product and then to a cyclic group. The middle modulus has the corpus maximum of 60 cyclic subgroups, but only six possible element orders. The final modulus has fewer units than 273, a much larger exponent, and only eight cyclic subgroups. The

last triple gives a related contrast: $U(287)$ and $U(288)$ are noncyclic with 32 cyclic subgroups each, but their exponents and richness values differ; $U(289)$ is cyclic because 289 is an odd prime power.

These local comparisons reinforce the role of factorization as the explanatory variable. Neither modulus size, parity, group cardinality, nor proximity predicts the spectrum by itself. The prime-power components determine the local cyclic orders, and their least-common-multiple interaction determines the observed discontinuities. Consecutive cases are therefore useful verification records: a program that changes its output only gradually with n would miss precisely the structural changes that the arithmetic requires.

5. Validity, and scope

The computational claims are exact integer claims. Trial division recovers every factorization in the stated range, Euler and Carmichael values follow closed formulas, and the direct-product calculation uses exact counters. Floating-point arithmetic appears only in elapsed-time summaries and graphic placement. The correspondence of the two independent unit spectra was checked at all 299 moduli, and the additive formula was checked against every residue. Re-running the supplied script reconstructs the CSV files, the five figures, and the JSON summary used in the manuscript.

The certification conditions form a hierarchy rather than a single final comparison. The principal levels are collected in table 14. Each level tests a different mathematical consequence of the group structure, so a failure can be localized before the residuewise dictionary is inspected element by element.

Table 14: Hierarchy of exact spectrum certificates.

Certificate	Exact condition	Principal fault detected
Cardinality	$\sum_d N(d) = \varphi(n)$	Lost or duplicated elements
Order support	$d \mid \lambda(n)$ whenever $N(d) > 0$	Invalid order key or exponent
Subgroup divisibility	$\varphi(d) \mid N(d)$	Incorrect order multiplicity
Invariant factors	$\prod_i s_i = \varphi(n)$ and $s_{\rho(n)} = \lambda(n)$	Faulty structural normalization
Involution rank	$N(2) = 2^{\rho(n)} - 1$ for $n \geq 3$	Incorrect 2-primary component
Independent equality	$N_{\text{CRT}}(d) = N_{\text{res}}(d)$ for every d	Any remaining construction error

The first condition alone is insufficient because two incorrect entries can offset each other. The order-support condition prevents an impossible key but does not determine its multiplicity. Totient divisibility then tests whether the elements assigned to an order can be partitioned into generator sets of cyclic subgroups. The invariant-factor equalities constrain the total size and maximum order through a separate representation. Finally, the involution identity isolates the 2-primary rank. A dictionary that passes all five structural checks is already strongly constrained; exact equality with the independently reduced residue orders closes the remaining route for an undetected error within the tested corpus.

The checks also behave differently under common implementation mistakes. Omitting the exceptional decomposition of $U(2^\alpha)$ can preserve the group cardinality but changes the exponent, involution count, and several multiplicities. Misaligning prime valuations during invariant-factor normalization can preserve the product of the factors while changing the divisibility chain or its largest member. An incorrect least-common-multiple update can retain admissible order keys but violate subgroup divisibility. Residue filtering errors usually appear first in the cardinality comparison. Because these signatures are distinct, the certificate record is useful not only for confirming a result but also for diagnosing the stage at which a calculation became inconsistent.

Reproducibility across moduli adds another layer. Members of a repeated invariant-factor class must have identical spectra even when their prime-power factorizations and convolution routes differ. The 85 recurrent classes therefore supply 157 nontrivial repetitions beyond the first representative of each class. The equalities at 247 and 259, at 217 and 279, and across the eight members of $C_2 \times C_2 \times C_{12}$ test whether normalization removes construction-order dependence. These comparisons do not replace residuewise certification, but they expose inconsistencies that would otherwise remain isolated within separate modulus records.

Several boundaries should be kept explicit. First, the interval ends at 300. The algebraic formulas remain valid beyond that limit, but the observed extrema—such as $n = 241$ for richness or $n = 263$ for compression—are statements about this interval only. Second, factorization is treated as available after trial division. For cryptographic-size integers, factorization may dominate the calculation, and generic order algorithms become relevant when the modulus factorization is unknown. Third, the order spectrum does not uniquely identify every finite group in complete generality. Invariant factors are therefore stored alongside the spectrum rather than inferred from it alone.

The timing comparison is also bounded by implementation choices. It contrasts transparent Python procedures in the same process, not optimized low-level libraries. Montgomery arithmetic, explicit CRT powering, fast multiplication, and

compiled group packages would change absolute times. They would not change the exact spectra or the reduction in the number of global element traversals. The deterministic work counts are consequently the more portable evidence, while the elapsed times demonstrate that the reduction is visible even at small moduli.

The study does not evaluate cryptographic security. Large element order is relevant to several cryptographic constructions, but moduli below 301 are unsuitable for security. The present contribution is arithmetic and computational: it provides a certified description of small modular groups and a method that can be incorporated into educational software, test suites, or larger exact computations when factorizations are known.

6. Conclusion

The paper asked whether the complete structural account of $\mathbb{Z}_n$ and $U(n)$ could be recovered from factorization without following the power orbit of every element, and what that account reveals across $2 \leq n \leq 300$. The answer is affirmative. Divisor counts give the additive spectrum exactly, local unit groups give explicit cyclic factors, and least-common-multiple convolution gives every unit-order multiplicity. Invariant-factor normalization certifies group cardinality and exponent, while residuewise order reduction independently verifies the final dictionary.

The detailed cases make the answer concrete. The 209 elements of $\mathbb{Z}_{209}$ collapse to four spectral pairs, with 180 full-order generators. The 48 elements of $U(210)$ collapse to six pairs, but their invariant factors 2, 2, 12 prove that the group is noncyclic and explain its 24 cyclic subgroups. Across the full corpus, 114 unit groups are cyclic, $U(241)$ has the richest order distribution with 20 orders, and $U(263)$ gives the largest spectral compression with 65.5 residues per entry. Additive and multiplicative richness frequently move in opposite directions; the consecutive moduli 240 and 241 provide the sharpest illustration.

The extended structural census sharpens this conclusion. Rank-one unit groups account for 113 nontrivial cyclic cases, rank two is the dominant noncyclic type with 133 moduli, and only five groups reach rank four. Their involution populations follow the exact sequence $2^{\rho(n)} - 1$, so a single multiplicity certifies the invariant-factor rank. The exponent ratio and maximum-order density measure different features: $U(240)$ has exponent only one sixteenth of its cardinality, yet three quarters of its elements attain that exponent. Repeated normalization places the 299 groups into 142 isomorphism classes, showing that the same unit-group structure frequently arises from unrelated modulus factorizations. These results answer the comparative part of the research question by identifying which features are controlled by cardinality, which are controlled by exponent, and which require the complete invariant-factor and multiplicity records.

The computational evidence supports the same conclusion. Convolution reduces the structural calculation from 1,587,655 orbit multiplications to 4,425 occupied order-pair combinations and lowers the median whole-corpus time by a factor of 17.77. The principal gain is not merely speed. The output changes from a long list into a concise mathematical explanation in which every multiplicity has a local origin and every total has an exact certificate. For modular-group computation, the order spectrum and invariant factors together provide the appropriate level of description: complete enough to recover generators and cyclic subgroups, compact enough to inspect, and explicit enough to verify independently.

The findings also specify when each representation should be used. A residue list remains appropriate when a particular inverse or representative generator is required. The divisor-lattice spectrum is preferable when the aim is to characterize $\mathbb{Z}_n$, because it returns the complete answer from the divisors of n . The CRT spectrum is preferable for $U(n)$ when the factorization is known, because it preserves local arithmetic and avoids repeated global traversal. Keeping the residuewise reducer as a separate check provides assurance without making enumeration the organizing principle. This division of labor directly answers the practical part of the research question: a complete modular-group tool need not choose between compact structure and verifiable element information; it can compute the former first and recover the latter selectively.

Beyond the tested interval, the same formulas permit a natural continuation to larger factored moduli and to batches sharing prime-power components. Local cyclic spectra can be cached, and convolution can reuse them whenever the same p^α reappears. Such scaling does not alter the mathematical output established here. The decisive observation remains that group structure determines the computation: factorization supplies local groups, least common multiples combine their periods, and totient divisibility certifies the resulting multiplicities.

Conflict of interest

The authors declare no conflict of interest.

Funding

This research received no external funding.

References

- [1] J. B. Fraleigh, *A First Course in Abstract Algebra*, 7th ed., Addison-Wesley, Boston, 2003.
- [2] D. S. Dummit and R. M. Foote, *Abstract Algebra*, 3rd ed., Wiley, Hoboken, 2004.
- [3] E. Bach and J. Shallit, *Algorithmic Number Theory, Volume I: Efficient Algorithms*, MIT Press, Cambridge, MA, 1996.
- [4] H. Cohen, *A Course in Computational Algebraic Number Theory*, Springer, Berlin, 1993.
- [5] V. Shoup, *A Computational Introduction to Number Theory and Algebra*, 2nd ed., Cambridge University Press, Cambridge, 2009.
- [6] C. Pomerance, Popular values of Euler’s function, *Mathematika* 27 (1980), 84–89.
- [7] K. Ford, The distribution of totients, *Ramanujan Journal* 2 (1998), 67–151.
- [8] P. Erdős, C. Pomerance and E. Schmutz, Carmichael’s lambda function, *Acta Arithmetica* 58 (1991), 363–385.
- [9] G. Martin and C. Pomerance, The iterated Carmichael λ -function and the number of cycles of the power generator, *Acta Arithmetica* 118 (2005), 305–335.
- [10] J. B. Friedlander, C. Pomerance and I. E. Shparlinski, Period of the power generator and small values of Carmichael’s function, *Mathematics of Computation* 70 (2001), 1591–1605.
- [11] P. Moree, Artin’s primitive root conjecture—a survey, *Integers* 12A (2012), Article A13.
- [12] J. M. Pollard, Theorems on factorization and primality testing, *Proceedings of the Cambridge Philosophical Society* 76 (1974), 521–528.
- [13] J. M. Pollard, A Monte Carlo method for factorization, *BIT Numerical Mathematics* 15 (1975), 331–334.
- [14] R. P. Brent, An improved Monte Carlo factorization algorithm, *BIT Numerical Mathematics* 20 (1980), 176–184.
- [15] H. W. Lenstra Jr., Factoring integers with elliptic curves, *Annals of Mathematics* 126 (1987), 649–673.
- [16] C. Pomerance, A tale of two sieves, *Notices of the American Mathematical Society* 43 (1996), 1473–1485.
- [17] D. J. Bernstein, Factoring into coprimes in essentially linear time, *Journal of Algorithms* 54 (2005), 1–30.
- [18] M. Agrawal, N. Kayal and N. Saxena, PRIMES is in P, *Annals of Mathematics* 160 (2004), 781–793.
- [19] P. L. Montgomery, Modular multiplication without trial division, *Mathematics of Computation* 44 (1985), 519–521.
- [20] C. K. Koç, T. Acar and B. S. Kaliski Jr., Analyzing and comparing Montgomery multiplication algorithms, *IEEE Micro* 16(3) (1996), 26–33.
- [21] M. Joye and S.-M. Yen, The Montgomery powering ladder, in *Cryptographic Hardware and Embedded Systems—CHES 2002*, Lecture Notes in Computer Science 2523, Springer, 2003, pp. 291–302.
- [22] D. J. Bernstein and J. P. Sorenson, Modular exponentiation via the explicit Chinese remainder theorem, *Mathematics of Computation* 76 (2007), 443–454.
- [23] M. Fürer, Faster integer multiplication, in *Proceedings of the 39th ACM Symposium on Theory of Computing*, 2007, pp. 57–66.
- [24] D. Harvey and J. van der Hoeven, Integer multiplication in time $O(n \log n)$, *Annals of Mathematics* 193 (2021), 563–617.
- [25] E. Teske, A space efficient algorithm for group structure computation, *Mathematics of Computation* 67 (1998), 1637–1663.
- [26] E. Teske, The Pohlig–Hellman method generalized for group structure computation, *Journal of Symbolic Computation* 27 (1999), 521–534.
- [27] J. Buchmann and A. Schmidt, Computing the structure of a finite abelian group, *Mathematics of Computation* 74 (2005), 2017–2026.
- [28] N. Vikas, An $O(n)$ algorithm for abelian p -group isomorphism and an $O(n \log n)$ algorithm for abelian group isomorphism, *Journal of Computer and System Sciences* 53 (1996), 1–9.
- [29] K. K. H. Cheung and M. Mosca, Decomposing finite abelian groups, *Quantum Information and Computation* 1(3) (2001), 26–32.
- [30] A. V. Sutherland, *Order Computations in Generic Groups*, Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, 2007.
- [31] M. Furst, J. Hopcroft and E. Luks, Polynomial-time algorithms for permutation groups, in *Proceedings of the 21st Annual Symposium on Foundations of Computer Science*, 1980, pp. 36–41.
- [32] L. Babai and E. Szemerédi, On the complexity of matrix group problems I, in *Proceedings of the 25th Annual Symposium on Foundations of Computer Science*, 1984, pp. 229–240.
- [33] F. Celler, C. R. Leedham-Green, S. H. Murray, A. C. Niemeyer and E. A. O’Brien, Generating random elements of a finite group, *Communications in Algebra* 23 (1995), 4931–4948.
- [34] E. A. O’Brien, The p -group generation algorithm, *Journal of Symbolic Computation* 9 (1990), 677–698.
- [35] B. Eick and E. A. O’Brien, Enumerating p -groups, *Journal of the Australian Mathematical Society* 67 (1999), 191–205.
- [36] J. J. Cannon and D. F. Holt, Automorphism group computation and isomorphism testing in finite groups, *Journal of Symbolic Computation* 35 (2003), 241–267.
- [37] W. Bosma, J. Cannon and C. Playoust, The Magma algebra system I: The user language, *Journal of Symbolic Computation* 24 (1997), 235–265.
- [38] The GAP Group, *GAP—Groups, Algorithms, and Programming*, Version 4.12.2, 2022, <https://www.gap-system.org>.
- [39] M. Ekerå, On completely factoring any integer efficiently in a single run of an order-finding algorithm, *Quantum Information Processing* 20 (2021), Article 53.