Electronic Colloquium on Computational Complexity, Report No. 211 (2026)

Computing Modular Factorials Below the Square-Root Barrier

Yann Tal
yanntal@mail.tau.ac.il

September 2026

Abstract

Given a prime p, an integer $0 \le n < p$, and a divisor $q \mid (1 + p + p^2)$, we compute $n! \bmod p$ in expected bit complexity

$$\widetilde{O}\left(q^c + \frac{\sqrt{p}}{q^{1/4}}\right)$$

for some absolute constant $c \ge 1$. More generally, the construction applies when $q \mid \Phi_r(p)$, where Φ_r is the r-th cyclotomic polynomial and r is any fixed odd prime power. Combining these constructions, for every fixed $0 < \epsilon < 1/2$, we obtain an expected bit complexity of $p^{1/2-\delta_\epsilon+o(1)}$, with $\delta_\epsilon > 0$, for at least a $1-\epsilon$ fraction of primes up to X, for all sufficiently large X. Under the same divisor conditions, the method extends to moduli p^k with an additional factor polynomial in k.

The method recovers ratios of factorials modulo p from Jacobi sums, character sums over finite fields. We use Lenstra and Silverberg's algorithm to reconstruct these sums up to a root of unity from their ideal factorizations and products with their complex conjugates. We determine this root using van Wamelen's criterion. Lattice rounding selects nearly equal parts while keeping the remaining factorials small, so each recursive step uses only one smaller factorial and short interval products. The only randomized steps are Las Vegas constructions of the finite-field representations and multiplicative characters.

Related constructions use suitable divisors of $p-1$ for moduli p and p^2, and of $p+1$ for modulus p. A separate deterministic algorithm uses modular halving and simultaneous evaluation to improve the Bostan–Gaudry–Schost bound modulo p and p^2 by a factor of $\sqrt{\log p / \log\log p}$, without any divisor assumption.

ISSN 1433-8092

Contents

1 Introduction

1.1 Earlier Algorithms and Main Results

Let p be prime and let $0 \le n < p$. We study the computation of

$$n! = 1 \cdot 2 \cdots n \pmod{p},$$

with $0! = 1$. The inputs are written in binary, so their length is $O(\log p)$, even though the product may contain nearly p factors. For $n \ge 1$, direct multiplication uses $n - 1$ operations in $\mathbb{F}_p$, the field of residues modulo p. The question is whether arithmetic relations among factorials allow us to avoid most of those multiplications. [1] Throughout, running times include all work for a single input and count bit operations unless field operations are specified. We write $\widetilde{O}$ to suppress logarithmic factors in the running time; polynomial dependence on the exponent of the modulus is kept separate.

The starting point is the fast product method of Bostan, Gaudry, and Schost [1]. Its square-root dependence is already visible in the simpler baby-step/giant-step construction, which computes $n! \bmod p$ using

$$\widetilde{O}(\sqrt{n})$$

field operations. To illustrate the idea, suppose for simplicity that n is a perfect square, and define

$$f(x) = (x+1)(x+2)\cdots(x+\sqrt{n}).$$

Then

$$f(0) = 1 \cdot 2 \cdots \sqrt{n},$$
$$f(\sqrt{n}) = (\sqrt{n}+1)\cdots 2\sqrt{n},$$

and similarly for the remaining consecutive blocks. Therefore,

$$\prod_{i=0}^{\sqrt{n}-1} f(i\sqrt{n}) \equiv n! \pmod{p}.$$

The polynomial f has degree $\sqrt{n}$ and is evaluated at $\sqrt{n}$ points. A product tree constructs f, and fast multipoint evaluation computes all these values together using

$$\widetilde{O}(\sqrt{n})$$

field operations. This procedure also uses $\widetilde{O}(\sqrt{n})$ field elements of workspace. One may reduce the space usage by working with shorter blocks, but this increases the number of evaluations and hence the running time. Wilson's theorem also allows one to apply the same method to the complementary product:

$$n! \equiv -\left(\prod_{j=n+1}^{p-1} j\right)^{-1} \pmod{p}. \tag{1}$$

Thus the elementary bound is

$$\widetilde{O}\left(1 + \sqrt{\min\{n, p-1-n\}}\right),$$

which remains $\widetilde{O}(\sqrt{p})$ in the worst case.

[1] The direct variant that lists the primes up to n and multiplies their powers in $n!$ saves at most logarithmic factors.

Our algorithms retain this fast interval-product routine, but arrange that it is applied only to much shorter intervals. The missing factorial values are recovered from algebraic relations involving multiplicative character sums. A divisor $q \mid (p-1)$ determines both the spacing of these values and the cost of the algebraic computation: increasing q shortens the remaining products, but makes the character sums more expensive to recover. The first theorem states this trade-off; the mechanism is explained in Section 1.2.

Theorem 1.1. *There exists a deterministic algorithm that, given a prime p and natural numbers q, n, where $q \mid (p-1)$ and $0 \le n < p$, computes*

$$n! \pmod p$$

in time

$$\widetilde{O}\left(q^c + \sqrt{\frac{p}{q}}\right),$$

for some absolute constant $c \ge 1$.

The two terms in the running time are balanced when

$$q = \Theta\Big(p^{1/(2c+1)}\Big).$$

Whenever $p-1$ has a divisor of this size, the worst-case running time becomes

$$\widetilde{O}\left(p^{\frac{c}{2c+1}}\right),$$

and therefore achieves the exponent

$$\frac{c}{2c+1} = \frac{1}{2} - \frac{1}{4c+2} < \frac{1}{2},$$

breaking the square-root barrier.

The exponent c comes from the algebraic reconstruction, whose main ingredient is the polynomial-time polarized-ideal recovery algorithm of Lenstra and Silverberg [2].

Computing factorials modulo p^2 also has arithmetic applications. For a prime p, the Wilson quotient is

$$W_p = \frac{(p-1)! + 1}{p}.$$

Consequently, $(p-1)! \bmod p^2$ determines $W_p \bmod p$, and p is a Wilson prime precisely when

$$(p-1)! \equiv -1 \pmod{p^2}.$$

Costa, Gerbicz, and Harvey [3] used this relation in a large-scale search for Wilson primes. Their setting treats many primes simultaneously. Here we retain a worst-case, single-input problem.

Theorem 1.2. *Let c be the absolute constant from Theorem 1.1. There exists a deterministic algorithm such that, given a prime p and natural numbers q, n, where $q \mid (p-1)$ and $0 \le n < p^2$, the algorithm computes*

$$n! \pmod{p^2}$$

in time

$$\widetilde{O}\left(q^c + \sqrt{\frac{p}{q}}\right).$$

The preceding results rely on a useful divisor of $p-1$. Such a divisor need not exist in the desired range. We therefore seek similar relations using divisors of other explicitly computable integers associated with p. These constructions use divisors of integers obtained by evaluating cyclotomic polynomials at p, and extend to moduli p^k with arbitrary $k \geq 1$.

Let $\Phi_r(X)$ denote the r-th cyclotomic polynomial, whose roots are the primitive r-th roots of unity. Let $\varphi(\cdot)$ denote Euler's totient function: $\varphi(r)$ counts the integers in $\{1, \ldots, r\}$ coprime to r. Thus r is the cyclotomic index, whereas the degree of $\Phi_r(X)$ is $\varphi(r)$. For an odd prime power $r = \ell^e$,

$$\Phi_r(X) = \sum_{j=0}^{\ell-1} X^{j\ell^{e-1}}, \qquad \varphi(r) = \ell^{e-1}(\ell - 1).$$

For example, $\Phi_3(p) = 1 + p + p^2$ and $\Phi_9(p) = 1 + p^3 + p^6$. The same index r is the extension degree in $\mathbb{F}_{p^r}/\mathbb{F}_p$. The integer k specifies the target modulus p^k. Only r is fixed in the following theorem.

Theorem 1.3 (Odd prime-power cyclotomic indices)**.** *Fix an odd prime power r. There is a constant $c_r \geq 1$, depending only on r, with the following property. Given a prime p, a divisor $q \mid \Phi_r(p)$, an integer $k \geq 1$, an integer $0 \leq n < p^k$, a representation of $\mathbb{F}_{p^r}$, and a compatible order-q character, a deterministic algorithm computes $n!$ mod p^k in*

$$\widetilde{O}_r\left(k^{O_r(1)}\left[q^{c_r} + \left(\frac{\Phi_r(p)}{q}\right)^{1/(2\varphi(r))}\right]\right)$$

bit operations. The exponent in $k^{O_r(1)}$ and all implied constants depend only on r.

The order-q character χ_q and reduction maps in the theorem are chosen together: modulo the selected prime above p, the character value at $x \neq 0$ reduces to $x^{-(p^r-1)/q}$. For each divisor $m \mid q$, we use the character power $\chi_m = \chi_q^{q/m}$ and restrict the same reduction map to the subfield generated by the m-th roots of unity. These choices are described in Section 2.1.

For fixed k, the factor $k^{O_r(1)}$ is constant. Its polynomial dependence also permits $k = p^{o(1)}$ without changing a fixed exponent saving in p. The proof accounts for this dependence in Section 7.6.

For $r = \ell^e$, the factor ℓ, when it divides q, is removed before reconstruction. It occurs in $\Phi_r(p)$ at most once, so this changes the stated bounds only by constants depending on r. If the remaining divisor is one, the algorithm uses direct interval products. These reductions are justified in Lemma 2.2. For fixed r, the interval term in Theorem 1.3 is

$$p^{1/2}q^{-1/(2\varphi(r))}\left(1 + O_r\left(\frac{1}{p}\right)\right).$$

For fixed k, the indices $r = 3, 5, 9$, for example, give the respective bounds

$$\widetilde{O}\left(q^{c_3} + p^{1/2}q^{-1/4}\right), \qquad \widetilde{O}\left(q^{c_5} + p^{1/2}q^{-1/8}\right),$$
$$\widetilde{O}\left(q^{c_9} + p^{1/2}q^{-1/12}\right).$$

More generally, for fixed k and $q = p^{\delta+o(1)}$, the exponent for index r is

$$\max\left\{c_r\delta,\ \frac{1}{2} - \frac{\delta}{2\varphi(r)}\right\}.$$

The two terms balance when

$$q = p^{\frac{\varphi(r)}{2c_r\varphi(r)+1}+o(1)},$$

giving the exponent

$$\frac{c_r\varphi(r)}{2c_r\varphi(r)+1} = \frac{1}{2} - \frac{1}{4c_r\varphi(r)+2}.$$

A representation of $\mathbb{F}_{p^r}$ and the order-q character can also be constructed by a Las Vegas procedure. Given the factorization of q, sample an element of $\mathbb{F}_{p^r}^{\times}$, raise it to $(p^r-1)/q$, and verify that the result has order q. The expected number of trials is $q/\varphi(q)$. The resulting element fixes the reduction map, which in turn determines the character as above. Factoring q deterministically within $q^{O(1)}$ work is sufficient for the stated bound. Thus these constructions give the same expected asymptotic running time.

Remark 1.4 (Cyclotomic fallback). For every fixed $0 < \epsilon < 1/2$, Proposition 1.7 allows us to choose $O(\log(1/\epsilon))$ fixed odd prime-power indices so that, for all sufficiently large X, their algorithms give a positive exponent saving for at least a $1-\epsilon$ proportion of the primes $p \le X$. The indices and the saving are fixed once ϵ is chosen. The largest index is $O(\log(1/\epsilon)\log\log(1/\epsilon))$ as ϵ tends to zero. Section 1.3 also discusses wider divisor windows through a random-congruence and conditional-independence model.

Section 4.3 extends the prime-modulus construction to divisors of $p+1$. A uniform refinement modulo p and p^2 appears in Section 4.4.

1.2 The Technique

Suppose first that $q \mid (p-1)$, and set

$$K = \frac{p-1}{q}.$$

The first construction recovers factorials at multiples of K from $K!$ and an algebraic quantity associated with p, q, and the chosen multiple. Once a nearby value $(aK)!$ is known, the factors between aK and n give $n!$. After Wilson complementation, the nearest such boundary is at distance at most $K/2$, so this last step costs $\widetilde{O}(\sqrt{K})$. For example, when $p = 101$, $q = 5$, and $n = 47$, we recover 40! from 20! and then multiply the seven factors from 41 through 47.

The relation making this possible is a specialization of Young's congruence between Jacobi sums and multinomial coefficients [4, Theorem 2.2]:

$$(aK)! \equiv (-1)^{a-1}\rho_{\mathfrak{p}}(J_a(\chi))\,(K!)^a \pmod{p}, \qquad 1 \le a < q. \tag{2}$$

Here χ is a multiplicative character of order q: it assigns q-th roots of unity to nonzero field elements and respects multiplication. The Jacobi sum $J_a(\chi)$ combines products of these values over tuples whose entries add to one. The map $\rho_{\mathfrak{p}}$ reduces this algebraic integer modulo a chosen prime above p. The choices satisfy $\rho_{\mathfrak{p}}(\chi(x)) = x^{-K}$ for $x \neq 0$, the relation used in the congruence. The definitions appear in Section 2; the point needed now is that one shorter factorial and one character-sum residue determine $(aK)!$.

Since $q \mid (p-1)$, the field $\mathbb{F}_p$ contains all q-th roots of unity, so $\Phi_q(X)$ factors into distinct linear factors over $\mathbb{F}_p$. Equivalently, p splits completely in $\mathbb{Q}(\zeta_q)$, where ζ_q is a primitive q-th root of unity; we therefore call this the *split case*. In particular, the reduction of a Jacobi sum at a prime above p takes values in $\mathbb{F}_p$.

Computing $K!$ and the final correction costs $\widetilde{O}(\sqrt{p/q})$. We reconstruct the Jacobi sum in $\widetilde{O}(q^c)$ time. These are the two contributions to Theorem 1.1. The remaining question is how the algebraic reconstruction avoids the long sum.

High-level idea. For an ordinary nonzero integer, its prime factorization determines it up to a sign. For an algebraic integer, the corresponding ideal factorization leaves a larger ambiguity: multiplication by any unit, that is, an algebraic integer whose inverse is also integral. Stickelberger's theorem determines the ideal generated by a Jacobi sum. We also know the product of the sum with its complex conjugate. Together, these constraints leave only a root of unity undetermined. We recover a generator with these properties, identify the remaining root of unity, and retain products in factored form until they are reduced modulo the chosen prime.

The algorithmic steps. Reconstructing a two-character Jacobi sum has three steps.

1. Construct its principal ideal from the prime-ideal factorization given by Stickelberger's theorem.
2. Use the deterministic algorithm of Lenstra and Silverberg to find a generator with the prescribed product with its complex conjugate [2]. This pair of constraints is called a *polarization*; it reduces the unit ambiguity to a root of unity.
3. Determine that root using van Wamelen's criterion [5, Theorem 9 and Section 7], then reduce the resulting algebraic value. The criterion uses sums at proper divisors of the character order, which are computed first.

Only van Wamelen's root-selection criterion is used; its shortest-vector search is replaced by polarized-ideal recovery.

1.2.1 Jacobi-sum reconstruction

This subsection states the algebraic subroutine used by the factorial algorithms. Readers may continue to Section 1.2.2 on a first reading and return here after Section 2, where the fields, reduction maps, and product representations are defined.

For positive u, v, the two-character Jacobi sum is

$$\mathcal{J}_q(u,v) = J_p(\chi^u, \chi^v) = \sum_{x \in \mathbb{F}_p} \chi^u(x)\chi^v(1-x),$$

where each character is extended by zero at the origin. These sums are the building blocks for the a-fold sum in (2):

$$J_{u+v}(\chi) = J_u(\chi)J_v(\chi)\mathcal{J}_q(u,v), \qquad u+v<q. \tag{3}$$

An addition chain builds a from 1 by repeatedly adding two previously obtained integers. Binary doubling and addition use $O(\log a)$ steps, so (3) expresses $J_a(\chi)$ through $O(\log a)$ two-character sums. Their factors and integer exponents are stored separately; this is the product representation in the theorem below. For local evaluation, each factor is reduced modulo p or p^2 before the products are formed.

Theorem 1.5 (Jacobi-sum reconstruction)**.** *There exists a deterministic algorithm that, given a prime p, a divisor $q \mid (p-1)$, an integer $1 \le a < q$, and the character χ and prime ideal $\mathfrak{p}$ appearing in* (2)*, outputs a product representation of*

$$J_a(\chi) \in \mathcal{O}_{\mathbb{Q}(\zeta_q)}$$

of total description length $q^{O(1)}\,\mathrm{polylog}(p)$, in

$$\widetilde{O}(q^c)$$

bit operations for some absolute constant $c \geq 1$, using space polynomial in q and $\log p$. Consequently, for each $k \in \{1, 2\}$, it also computes

$$\rho_{\mathfrak{p}^k}(J_a(\chi)) \in \mathbb{Z}/p^k\mathbb{Z}$$

within the same asymptotic bound.

Enlarging c if necessary, the same constant is used in Theorems 1.1 and 1.2. The bound follows from polynomial-time polarized recovery and root selection in representations of size polynomial in q and $\log p$.

For a fixed odd prime power r, the same construction recovers a family of sums over $\mathbb{F}_{p^r}$, using the character powers $\chi_m = \chi_q^{q/m}$ for $m \mid q$, in total time $\widetilde{O}(q^{c_r})$. A character power χ_q^a has order $q/\gcd(a, q)$, so its value is obtained from the member with that order. We call the order the *conductor*. This is why the algorithm allows composite q without restricting a to be coprime to q. The family is constructed once and reused throughout the factorial recursion.

1.2.2 Divisors of $\Phi_r(p)$

For a divisor of $p - 1$, the algorithm recovers a factorial at a nearby block boundary. A divisor of $1 + p + p^2$ gives a different kind of relation: a binomial coefficient whose upper argument is near n and whose two lower factorial arguments are near $n/2$. The identity

$$z! = \binom{z}{x} x!\,(z - x)!$$

then reduces the problem to smaller factorials. The useful choice is to make x and $z - x$ close, so that one recursive value determines both factorials through a short interval product. We first explain this choice in two dimensions, then pass to higher cyclotomic indices.

The case $r = 3$. If the given divisor of $1 + p + p^2$ is divisible by 3, first divide it by 3, as justified by Lemma 2.2. Let $q > 1$ denote the resulting divisor, so that $\gcd(q, 3) = 1$, and set

$$D_3 = \frac{1 + p + p^2}{q}.$$

Here D_3 governs a congruence between two base-p digits, rather than the spacing of factorial block boundaries. We choose an integer $1 \leq a < q$ for which

$$aD_3 = pA + B, \qquad 1 \leq B \leq A < p.$$

The reconstructed character sum over $\mathbb{F}_{p^3}$, after a known sign correction, determines

$$\binom{A-1}{B-1} = \frac{(A-1)!}{(B-1)!\,(A-B)!} \pmod{p}.$$

This is the specialization of Proposition 5.1 to $r = 3$. The numerator argument is $A - 1$, while the two denominator arguments are $B - 1$ and $A - B$; their sum is $A - 1$. We therefore seek A near $n + 1$ and B near $n/2 + 1$.

A nearby lattice point. The pairs we may choose form the lattice

$$\left\{(A,B) \in \mathbb{Z}^2 : pA + B \equiv 0 \pmod{D_3}\right\}.$$

It has index, and hence determinant, D_3 in $\mathbb{Z}^2$, since $(A,B) \mapsto pA + B \bmod D_3$ is onto. The cyclotomic congruence $p^2+p+1 \equiv 0 \pmod{D_3}$ gives additional geometric structure. Let $\omega = e^{2\pi i/3}$, and identify (A,B) with $B + A\omega$ in the Eisenstein integers $\mathbb{Z}[\omega]$. Multiplication by ω, a rotation by 120 degrees, sends

$$(A,B) \longmapsto (B-A,-A).$$

It preserves the lattice, because

$$p(B-A) - A = p(pA+B) - (p^2+p+1)A \equiv 0 \pmod{D_3}.$$

Thus the lattice has a rotational symmetry in its Eisenstein embedding. A nonzero vector of length $O(\sqrt{D_3})$, found by two-dimensional lattice reduction, and its rotated vector give two independent directions of the same length. Expressing a target in these directions and rounding the two coefficients reaches a lattice point within $O(\sqrt{D_3})$ in each coordinate. The change between the Eisenstein embedding and the integer coordinates has bounded distortion, so the same error bound holds for A and B. Lemma 5.2 proves this rounding statement and its higher-dimensional form.

Apply the rounding procedure near $(n+1, n/2+1)$. After Wilson complementation reduces the target to $n \le (p-1)/2$, and with the stopping rule for small arguments used in Section 5, the resulting digits satisfy

$$\begin{aligned} A - 1 &= n + O(\sqrt{D_3}), \\ B - 1 &= n/2 + O(\sqrt{D_3}), \quad A - B = n/2 + O(\sqrt{D_3}). \end{aligned}$$

The congruence makes $a = (pA+B)/D_3$ an integer. The digit bounds give $0 < pA + B < p^2 < qD_3$, so $1 \le a < q$ and the reconstruction applies.

One recursive call for both factorials. The recursive call computes only $(B-1)! \bmod p$. The other factorial is obtained from the short ratio

$$\frac{(A-B)!}{(B-1)!} = \begin{cases} \displaystyle\prod_{j=B}^{A-B} j, & A - B \ge B - 1, \\ \displaystyle\left(\prod_{j=A-B+1}^{B-1} j\right)^{-1}, & A - B < B - 1, \end{cases} \pmod{p}.$$

Its length is $|A - 2B + 1| = O(\sqrt{D_3})$. Squaring the returned $(B-1)!$, multiplying by this ratio, and then by the known binomial coefficient gives $(A-1)!$. A final interval correction of length $|n - (A-1)| = O(\sqrt{D_3})$ gives $n!$.

For example, take $p = 541$, $q = 13963$, and $n = 180$. Then $D_3 = 21$, and the lattice pair $(A,B) = (182, 91)$ satisfies $pA + B = 4693D_3$. The corresponding relation is

$$180! \equiv \binom{181}{90} (90!)^2 \frac{91}{181} \pmod{541}.$$

Once the binomial coefficient has been reconstructed, only 90! is computed recursively; the other corrections are the factor 91 and the inverse of 181.

Each correction costs $\widetilde{O}(D_3^{1/4})$, by the interval-product routine. The recursive argument is $n/2 + O(\sqrt{D_3})$, and the routine stops at arguments of size $O(\sqrt{D_3})$. It therefore has logarithmic depth, with only one child at each stage. The same family of reconstructed sums serves every stage, giving total time

$$\widetilde{O}\Big(q^{c_3} + D_3^{1/4}\Big) = \widetilde{O}\Big(q^{c_3} + p^{1/2}q^{-1/4}\Big).$$

From two dimensions to $r-1$ dimensions. Now let r be an odd prime, let $q > 1$ divide $\Phi_r(p)$ with $\gcd(q, r) = 1$, and set $D_r = \Phi_r(p)/q$. The two digits B, A are replaced by $r-1$ increasing digits:

$$aD_r = x_0 + x_1 p + \cdots + x_{r-2}p^{r-2}, \qquad 1 \le x_0 \le x_1 \le \cdots \le x_{r-2} < p.$$

The reconstructed value now determines the multinomial coefficient

$$\binom{x_{r-2} - 1}{x_0 - 1,\ x_1 - x_0,\ \ldots,\ x_{r-2} - x_{r-3}}.$$

Its parts are successive digit differences, and its total is $x_{r-2} - 1$. To make all parts nearly equal, we round the target vector

$$\left(1 + \frac{n}{r-1},\ 1 + \frac{2n}{r-1},\ \ldots,\ 1 + n\right)$$

to the congruence lattice

$$\left\{(x_0, \ldots, x_{r-2}) \in \mathbb{Z}^{r-1} : \sum_{i=0}^{r-2} x_i p^i \equiv 0 \pmod{D_r}\right\}.$$

Its determinant is D_r. In the power basis of $\mathbb{Z}[\zeta_r]$, multiplication by ζ_r preserves the congruence, just as the Eisenstein rotation did for $r = 3$. The translates of one short nonzero vector give $r-1$ independent short directions. The rounding error is consequently $O_r(D_r^{1/(r-1)})$ in each digit.

The total is therefore $n + O_r(D_r^{1/(r-1)})$, and all $r-1$ parts are $n/(r-1) + O_r(D_r^{1/(r-1)})$. One recursive factorial again suffices: short products give the other part factorials, the multinomial coefficient gives the total factorial, and a final short product reaches $n!$. The dimension $r-1$, which equals $\varphi(r)$ for prime r, determines the interval cost $\widetilde{O}(D_r^{1/(2(r-1))})$.

Odd prime-power indices. For $r = \ell^e$, set $h = \ell^{e-1}$ and $d = \varphi(r) = (\ell-1)h$, with $D_r = \Phi_r(p)/q$. The lattice now has dimension d, rather than $r-1$, and rounding error $O_r(D_r^{1/d})$. The identity

$$\frac{p^r - 1}{q} = (p^h - 1)D_r$$

arranges the base-p digits into h chains, with consecutive positions in a chain separated by h. Each chain gives a multinomial coefficient with $\ell - 1$ parts, and the reconstructed value is their product. We place the total of one chain near n and its parts near $n/(\ell-1)$; the other chains have arguments $O_r(D_r^{1/d})$. Their factorials are computed by short products. Within the large chain, one recursive value still determines every part factorial through short corrections.

For example, at $r = 9$ the six digits give

$$\binom{x_3 - 1}{x_0 - 1}\binom{x_4}{x_1}\binom{x_5}{x_2}.$$

We choose $x_3 - 1 = n + O(D_9^{1/6})$ and $x_0 - 1 = n/2 + O(D_9^{1/6})$, while $x_1, x_2, x_4, x_5 = O(D_9^{1/6})$. The second and third binomial coefficients are obtained directly from short factorials and divided out. The first is then used in the same one-child recursion as for $r = 3$. The interval cost is

$$\widetilde{O}(D_r^{1/(2d)}) = \widetilde{O}\Big(p^{1/2} q^{-1/(2\varphi(r))}\Big),$$

which, together with reconstruction in $\widetilde{O}(q^{c_r})$ time, explains the bound modulo p.

Higher powers of p. To work modulo p^k, the recursive call retains the shifted factorial polynomial

$$P_N(Z) = \prod_{j=1}^{N} (j + pZ) \pmod{p^k}, \qquad 0 \le N < p.$$

Its constant term is $N!$, and its other values give the shifted products needed by the lifted relation. The coefficient of Z^j is divisible by p^j, so the polynomial has degree less than k modulo p^k. In the recursive branch, values at k nearby lattice points with interpolation arguments distinct modulo p determine it. One recursive call returns the whole smaller polynomial.

The lifted relation also contains a quotient of p-adic gamma values. Its contribution is computed from the same identity at small factorial arguments. For fixed r, the number of relevant coefficients and all interpolation costs are polynomial in k. This gives the factor $k^{O_r(1)}$ in Theorem 1.3. If $n \ge p$, counting the factors of p either proves that the answer is zero or reduces the problem to shifted products with indices below p.

The split construction modulo p^2 is simpler. Its extra term is a harmonic-sum difference that a roots-of-unity filter evaluates using $q - 1$ modular exponentiations, so the original split bound is unchanged.

1.3 How restrictive is the assumption on q?

For each cyclotomic index, we consider the cost of finding a divisor in the chosen range and the proportion of primes for which such a divisor exists. We treat the split case $q \mid (p-1)$ first, followed by the higher cyclotomic constructions.

The split case. Theorems 1.1 and 1.2 take a suitable divisor $q \mid (p-1)$ as part of the input. Given the factorization of $p - 1$, a near-optimal divisor can be selected efficiently.

Proposition 1.6. *Given the complete prime factorization of an integer $M \ge 2$, an integer target $1 \le T \le M$, and a rational $0 < \delta \le 1$, one can deterministically find a divisor $q \mid M$ such that*

$$\frac{1}{1+\delta} \max\{d : d \mid M,\ d \le T\} \le q \le T$$

in time polynomial in the binary input length and $1/\delta$.

Proof. Write

$$M = \ell_1 \ell_2 \cdots \ell_s,$$

where prime factors are repeated according to multiplicity. Since every $\ell_i \ge 2$, one has $s \le \log_2 M$. Set

$$\gamma = \frac{\delta}{2s}$$

and begin with $L_0 = \{1\}$. After processing ℓ_i, merge

$$L_{i-1} \quad \text{and} \quad \{\ell_i x : x \in L_{i-1},\ \ell_i x \le T\}.$$

Sort the result, retain its smallest member, and thereafter retain an element only when it exceeds the previously retained element by a factor greater than $1+\gamma$. This is the usual trimmed-list approximation for Subset Sum, applied multiplicatively [6].

Inductively, every divisor $x \le T$ of $\ell_1 \cdots \ell_i$ has a retained representative y satisfying

$$y \le x \le (1+\gamma)^i y.$$

Consequently, the largest member of L_s is within a factor

$$(1+\gamma)^s \le \exp(\delta/2) \le 1+\delta$$

of the largest divisor below T. Consecutive retained elements differ by a factor greater than $1+\gamma$, so every list has

$$O\Big(1 + \frac{\log T}{\gamma}\Big) = O\left(\frac{(\log M)^2}{\delta}\right)$$

members. This proves the claimed running time. □

Armed with Proposition 1.6, we can turn a factorization into a useful divisor. Set

$$T = p^{1/(2c+1)}.$$

Choose a fixed window satisfying

$$0 < \alpha < \frac{1}{2c+1} < \beta < \frac{1}{2c}.$$

Applying Proposition 1.6 to T, and also to $(p-1)/T$ before taking complementary divisors, with $\delta = 1$, finds a divisor at most T within a factor of two of the largest such divisor, and a divisor at least T within a factor of two of the smallest such divisor. The two calls use the integer parts of T and $(p-1)/T$. We choose the candidate minimizing

$$q^c + \sqrt{p/q}.$$

Its cost is within a constant factor of the minimum over all divisors of $p-1$: below T, the decreasing term dominates, and above T, the increasing term dominates. The factor-two approximations change these terms by bounded factors. A divisor in the chosen window therefore ensures the same exponent saving for the selected candidate.

Starting from p alone, the preliminary arithmetic can also be carried out deterministically within the stated bound. Harvey and Hittmeir factor an integer N in $N^{1/5+o(1)}$ time [7]. Once q and its factorization are known, an element of order q in $\mathbb{F}_p^\times$ can be constructed deterministically using least power nonresidues; Burgess-type bounds give a search of $p^{1/4+o(1)}$ candidates [8]. As shown in the proof of Theorem 1.1, the inequality $q^c + \sqrt{p/q} \ge p^{1/3}$ holds for every $q \ge 1$. These setup costs are therefore covered by the stated running time.

Alternatively, the preliminary arithmetic has a rigorous Las Vegas implementation with expected $p^{o(1)}$ running time. Lenstra and Pomerance factor $p-1$ in expected time

$$\exp\Big((1+o(1))\sqrt{\log p \log\log p}\Big) = p^{o(1)}$$

[9]. After choosing q by Proposition 1.6, sample $x \in \mathbb{F}_p^\times$ and set

$$\eta = x^{(p-1)/q}.$$

This is uniform in the subgroup of order q. Given the factorization of q, it has order q precisely when

$$\eta^{q/\ell} \neq 1$$

for every prime $\ell \mid q$. The probability of this event is $\varphi(q)/q$, so the expected number of trials is

$$\frac{q}{\varphi(q)} = O(\log \log q).$$

Every claimed factor and every order condition is verified before the output is accepted; randomness therefore affects only the running time.

Consequently, whenever $p-1$ has a divisor in the chosen window, selecting a divisor with the same exponent saving, obtaining its factorization, and constructing an element of its order can all be carried out from p alone in expected $p^{o(1)}$ time.

Results of Koukoulopoulos [10] give a positive proportion of primes for which $p-1$ has a divisor in fixed intermediate-size windows in the applicable ranges.

Higher cyclotomic constructions. Fix an odd prime power r, and let M_r be the largest divisor of $\Phi_r(p)$ coprime to r. The conductors used here are the divisors $q > 1$ of M_r. For every such q, and every $m > 1$ dividing q, one has

$$\mathrm{ord}_m(p) = r,$$

where $\mathrm{ord}_m(p)$ is the multiplicative order of p modulo m; see Lemma 2.3. For fixed k, the two terms in the higher cyclotomic bound balance when

$$q = p^{\frac{\varphi(r)}{2c_r\varphi(r)+1}+o(1)}.$$

For fixed k, a saving of p^ε over the square-root bound is obtained whenever

$$p^{2\varphi(r)\varepsilon} \leq q \leq p^{(1/2-\varepsilon)/c_r}.$$

Fix a window with

$$0 < \alpha_r < \frac{\varphi(r)}{2c_r\varphi(r)+1} < \beta_r < \frac{1}{2c_r}.$$

Every divisor of M_r in $[p^{\alpha_r}, p^{\beta_r}]$ gives, for fixed k, an exponent saving of at least

$$\min\left\{\frac{\alpha_r}{2\varphi(r)}, \frac{1}{2} - c_r\beta_r\right\} > 0.$$

For every $c_r \geq 1$, one may choose

$$\alpha_r = \frac{\varphi(r)}{4c_r\varphi(r)+1}, \qquad \beta_r = \frac{2\varphi(r)}{4c_r\varphi(r)+1}. \tag{4}$$

These endpoints satisfy $\beta_r/\alpha_r = 2$ and give a saving of $1/(8c_r\varphi(r)+2)$. Their fixed ratio accommodates the dependence of the reconstruction cost on r.

Given the factorization of $\Phi_r(p)$, remove its factors dividing r. Use the two calls to Proposition 1.6 from the split case, with $M = M_r$ and $T = p^{\varphi(r)/(2c_r\varphi(r)+1)}$, and choose the candidate minimizing

$$q^{c_r} + p^{1/2}q^{-1/(2\varphi(r))}.$$

The same monotonicity argument gives a constant-factor approximation to the minimum cost. This selection uses prime factors with their multiplicities, including divisors formed from several factors and from prime powers.

For fixed r, the integer $\Phi_r(p)$ has $O_r(\log p)$ bits. The Las Vegas factorization algorithm of Lenstra and Pomerance therefore finds its factorization in expected $p^{o(1)}$ time [9]. A representation of $\mathbb{F}_{p^r}$ and the order-q character can likewise be constructed by Las Vegas procedures, with every result verified before it is used. The deterministic theorem takes the representation and character as inputs; the randomized setup affects only the running time.

The use of all divisors is helpful when M_r has many small prime factors. Suppose every prime factor of M_r is at most p^η, where $0 < \eta < 1/(2c_r)$ is fixed. Set

$$\beta = \frac{\varphi(r) + \eta}{2c_r\varphi(r) + 1}.$$

Multiply the prime factors, with multiplicity, until the next multiplication would first exceed p^β. The preceding product q satisfies

$$p^{\beta-\eta} < q \le p^\beta.$$

For sufficiently large p, the crossing occurs because $\log M_r = \varphi(r)\log p + O_r(1)$; the removed part divides the fixed prime underlying r. Substitution in the running-time bound gives

$$p^{\frac{c_r(\varphi(r)+\eta)}{2c_r\varphi(r)+1}+o(1)},$$

whose exponent is below $1/2$. In particular, when every prime factor is $p^{o(1)}$, one can select $q = p^{\varphi(r)/(2c_r\varphi(r)+1)+o(1)}$.

An unconditional density bound. Smaller windows keep the ratio of their endpoint exponents equal to two while bringing products of candidate primes into the range of the Bombieri–Vinogradov theorem. A finite-moment argument then gives the following density bound.

Proposition 1.7. *Fix $t \ge 1$, distinct odd prime-power indices $r_1, \dots, r_t$, and their reconstruction exponents $c_{r_i} \ge 1$. As $X \to \infty$, the proportion of primes $p \in [X, 2X]$ for which none of the $\Phi_{r_i}(p)$ has a prime divisor $q \nmid r_i$ satisfying*

$$p^{1/(24tc_{r_i})} \le q \le p^{1/(12tc_{r_i})} \tag{5}$$

is at most $2^{1-t} + o(1)$. All indices and exponents are fixed before X tends to infinity. Each successful divisor gives, for fixed k, a saving in the exponent of p of at least

$$\frac{1}{48t\max_{1\le i\le t}\{c_{r_i}\varphi(r_i)\}}. \tag{6}$$

Proof. Use candidate primes λ in the frozen windows

$$(2X)^{1/(24tc_{r_i})} \le \lambda \le X^{1/(12tc_{r_i})}.$$

They lie inside (5) for every $p \in [X, 2X]$. Exclude the finitely many primes dividing $r_1 \cdots r_t$. For each candidate λ, let Ω_λ be the union of the nonzero residue classes of order r_i for the indices whose windows contain λ. Distinct orders give disjoint classes, so

$$|\Omega_\lambda| = \sum_{\substack{i:\lambda \text{ in the } i\text{-th window} \\ \lambda \equiv 1 \pmod{r_i}}} \varphi(r_i).$$

Let $Z_X(p)$ count candidate primes for which $p \bmod \lambda \in \Omega_\lambda$. The failure event in the proposition is contained in $\{Z_X = 0\}$. Mertens' theorem in arithmetic progressions gives

$$\sum_\lambda \frac{|\Omega_\lambda|}{\lambda - 1} = t \log 2 + o(1). \tag{7}$$

Here λ ranges over primes, and the constant in the error may depend on the fixed indices and exponents.

For each $j \le 4t$, a product of j distinct candidate primes is at most $X^{j/(12t)} \le X^{1/3}$. The Chinese remainder theorem turns their joint condition into $\prod_{\nu=1}^{j} |\Omega_{\lambda_\nu}|$ reduced residue classes modulo that product. This number is bounded in terms of the fixed indices. Each unordered set of candidate primes determines a different squarefree product. The Bombieri–Vinogradov theorem, in its form with a maximum over reduced residue classes, therefore bounds the total prime-counting error over these products by $o(X/\log X)$ [11, Section 1]. The same bound applies to the interval $[X, 2X]$ by subtraction and partial summation.

It follows that the binomial moments, averaged over primes in $[X, 2X]$, satisfy

$$\mathbb{E}\binom{Z_X}{j} = \frac{(t \log 2)^j}{j!} + o(1) \qquad (0 \le j \le 4t). \tag{8}$$

Indeed, the main term is the j-th elementary symmetric sum of $|\Omega_\lambda|/(\lambda - 1)$. Their sum is given by (7), and their maximum tends to zero, since the lower endpoints are positive fixed powers of X. Terms with repeated primes therefore contribute $o(1)$.

Even Bonferroni truncation and (8) give

$$\Pr(Z_X = 0) \le \sum_{j=0}^{4t} \frac{(-t \log 2)^j}{j!} + o(1) \le 2^{-t} + \frac{(t \log 2)^{4t+1}}{(4t+1)!} + o(1) \le 2^{1-t} + o(1).$$

For the last inequality, use $(4t+1)! \ge ((4t+1)/e)^{4t+1}$ and $(e \log 2/4)^{4t+1} < 2^{-t}$.

Finally, a successful q in the i-th window satisfies

$$q^{c_{r_i}} \le p^{1/(12t)}, \qquad p^{1/2} q^{-1/(2\varphi(r_i))} \le p^{1/2 - 1/(48 t c_{r_i} \varphi(r_i))}.$$

The latter is the larger exponent. Taking the worst case over the finite list proves (6). □

For $0 < \epsilon < 1$, take $t = \lceil \log_2(4/\epsilon) \rceil$. The proportion of unsuccessful primes in $[X, 2X]$ is then at most $\epsilon/2 + o(1)$. Partitioning the primes up to X into dyadic intervals gives the following consequence: for all sufficiently large X, at least a $1 - \epsilon$ proportion of the primes $p \le X$ have a successful divisor in one of these windows. Choosing the first t odd prime-power indices gives largest index $O(t \log(t+1))$. For each fixed ϵ, these indices and the positive saving in (6) are independent of p and X.

Scan the integers up to $p^{1/(12t)}$, retaining divisors in the appropriate windows that are coprime to the corresponding index. For the fixed list this takes $p^{1/(12t)+o(1)}$ time and includes every prime

factor used in the proof. Each test evaluates the cyclotomic polynomial modulo the candidate. The search cost is absorbed by the factorial bound on the successful windows. The arithmetic is deterministic given the field and character, whose Las Vegas construction has the expected cost stated in Theorem 1.3.

Heuristic predictions for larger divisors. For every fixed $0 < \epsilon < 1/2$, the preceding density bound uses smaller divisors to give a positive exponent saving for at least a $1 - \epsilon$ fraction of primes up to X, for all sufficiently large X. We now consider the larger divisor ranges in (4), where a probabilistic model predicts stronger savings.

As p varies over primes, let E_r denote the event that M_r has a divisor in $[p^{\alpha_r}, p^{\beta_r}]$. These events describe the possibility of moving to the next index when no useful divisor has been found at an earlier one.

For a candidate prime $\lambda \nmid rp$,

$$\lambda \mid \Phi_r(p) \quad \Longleftrightarrow \quad \operatorname{ord}_\lambda(p) = r.$$

This is possible only if $\lambda \equiv 1 \pmod{r}$. There are $\varphi(r)$ elements of order r among the $\lambda - 1$ nonzero residue classes. The random-congruence model therefore assigns probability $\varphi(r)/(\lambda - 1)$ to this divisibility condition. For any fixed finite collection of candidate primes, the Chinese remainder theorem and the prime number theorem in arithmetic progressions justify equidistribution among the corresponding reduced residue classes [12].

For fixed r, Mertens' theorem for primes in arithmetic progressions [13] gives

$$\sum_{\substack{p^{\alpha_r} \le \lambda \le p^{\beta_r} \\ \lambda \equiv 1 \pmod{r}}} \frac{\varphi(r)}{\lambda - 1} = \log \frac{\beta_r}{\alpha_r} + o(1).$$

A Poisson model for these divisibility conditions predicts a useful individual factor with probability $1 - \alpha_r/\beta_r + o(1)$. Products of factors can only enlarge the divisor event E_r. Thus the same model gives a lower estimate for its probability, bounded away from zero when the ratios β_r/α_r are uniformly bounded away from one. For (4), the prime-factor prediction is $1/2 + o(1)$, for every fixed index. The absolute endpoints may move downward as c_r grows without changing this reciprocal-prime mass. Here r is fixed as p tends to infinity.

For successive indices, we make the further heuristic assumption

$$\Pr\Big(E_{r_j} \mid E_{r_1}^c, \dots, E_{r_{j-1}}^c\Big) \ge \delta_0$$

for some fixed $\delta_0 > 0$. Under this model, the probability that the first t indices all fail is at most

$$(1 - \delta_0)^t \le e^{-\delta_0 t}.$$

Hence $O(\log(1/\epsilon))$ attempts give success probability at least $1 - \epsilon$. One may try the odd prime-power indices in increasing order,

$$r = 3, 5, 7, 9, 11, 13, 17, 19, 23, 25, \dots.$$

They include the odd primes, so the largest of the first t indices is $O(t \log t)$. It is therefore enough to consider

$$r = O\left(\log \frac{1}{\epsilon} \, \log\log \frac{1}{\epsilon}\right).$$

For every fixed ϵ, only finitely many fixed-degree algorithms are used, so Theorem 1.3 applies. With (4), their common guaranteed saving is the minimum of $1/(8c_r\varphi(r) + 2)$ over this finite list, which is fixed before p tends to infinity.

1.4 A remark about Factorials and Integer Factoring

The relation between factorial evaluation and integer factoring begins with Strassen's observation [14] that the baby-step/giant-step product method can also factor integers. If $N = p_1 p_2$ with $p_1 < p_2$, then

$$p_1 \le \lfloor \sqrt{N} \rfloor < p_2$$

and therefore

$$\gcd\Big(\lfloor \sqrt{N} \rfloor ! \bmod N, N\Big) = p_1.$$

Computing the factorial residue by baby-step/giant-step takes $\widetilde{O}(N^{1/4})$ time.

The connection extends beyond this reduction. Lipton [15] observed that a polynomial-time construction of polylogarithmic-length, division-free straight-line programs for $n!$ would give a polynomial-time factoring algorithm. One may compute $y! \bmod N$ for several values of y and use binary search on $\gcd(y!, N)$ to recover a factor. Conversely, one may ask whether a fast factoring algorithm implies a fast algorithm for $n! \bmod p$. No such implication is known.

The straight-line complexity of factorials has also been studied in the Blum–Shub–Smale model [16]. Shub and Smale [17] considered the relaxed task of computing a nonzero multiple $m_n n!$ by a short straight-line program and related its hardness to $P_{\mathbb{C}} \neq NP_{\mathbb{C}}$. Using ideas from Lenstra's elliptic-curve factoring algorithm [18], Cheng [19] gave, under a conjecture on smooth integers in short intervals, a randomized construction of such a program of length

$$\exp\Big(O\Big(\sqrt{\log n \log\log n}\Big)\Big).$$

The multiplier m_n is unspecified and may itself be divisible by p.

Factoring nevertheless helps with composite moduli. Given

$$N = \prod_i p_i^{s_i},$$

the Chinese remainder theorem reduces $n! \bmod N$ to computations modulo the prime powers $p_i^{s_i}$. The higher cyclotomic algorithms apply componentwise whenever the corresponding divisor conditions hold, with polynomial dependence on each s_i for a fixed index r. If $s_i = p_i^{o(1)}$, a fixed saving in the exponent of p_i is preserved by Theorem 1.3. The residues are then recombined modulo N.

1.5 Organization of the Paper

Section 2 develops the cyclotomic notation, local reduction maps, Gauss and Jacobi sums, polarized ideals, product representations, p-adic gamma values, and fast interval products. It also states the two algorithmic ingredients used in the reconstruction.

Section 3 gives the reconstruction routines for the split Jacobi sums and for the higher cyclotomic family $\Theta_{r,m}$, $m \mid q$. Section 4 proves the split prime-modulus algorithm, with subsections on divisors of $p+1$ and a uniform refinement modulo p and p^2. Section 5 develops the multinomial identity, cyclotomic lattice, and recursive algorithm for the higher cyclotomic constructions. The preprocessing to locate and construct the relevant cyclotomic objects, the explicit window choices, and the smaller-window density argument are discussed in Section 1.3.

Section 6 lifts the split construction to modulus p^2. Section 7 lifts the higher cyclotomic recursion to modulus p^k by shifted-factorial interpolation and a p-adic gamma factor, and accounts for the polynomial dependence on k in Section 7.6. Finally, Section 8 summarizes the results and the

remaining quantitative and distributional questions. The appendices contain the finite congruence calculations, Stickelberger factorizations, higher cyclotomic multinomial identity, and lifted gamma identities. Sections G and H contain the additional reconstruction and interval-routing proofs.

2 Preliminaries

The overview separates the computation into algebraic reconstruction and short products. We now specify the objects that connect those two stages: first the fields and reduction maps, then the character-sum identities, and finally their computational representations. For the higher cyclotomic constructions $\widetilde{O}_r$ hides factors polynomial in $\log p$, $\log(q+1)$, and $\log(k+1)$, with constants depending on the fixed index r. Powers of k, including those arising from the bit length $k \log p$ of the modulus, are counted separately.

2.1 Cyclotomic Fields and Local Reduction

A character sum is first recovered in a number field, but the factorial algorithm needs its value in a finite ring. The choice of a root of unity in that ring specifies the reduction map. Fixing this choice is necessary: different choices can give different residues of the same algebraic number.

For an integer $m \geq 1$, let $\Phi_m(X)$ denote the m-th cyclotomic polynomial, let ζ_m be a primitive m-th root of unity, and set

$$L_m = \mathbb{Q}(\zeta_m), \qquad \mathcal{O}_m = \mathbb{Z}[\zeta_m].$$

The degree $[L_m : \mathbb{Q}]$ is $\varphi(m)$, where $\varphi(\cdot)$ is Euler's totient function. For a prime power $r = \ell^e$,

$$\Phi_r(X) = \sum_{j=0}^{\ell-1} X^{j\ell^{e-1}}.$$

Complex conjugation in a cyclotomic field is denoted by an overline. For a finite free R-algebra B, we use $\mathrm{Tr}_{B/R}(x)$ for the trace of the R-linear map $y \mapsto xy$, following Lenstra and Silverberg's trace convention [2, Section 3].

Let p be a prime with $p \nmid m$, and set

$$f_m = \mathrm{ord}_m(p),$$

with $f_1 = 1$ at the trivial conductor. Choose an element $\eta_m \in \mathbb{F}_{p^{f_m}}^{\times}$ of order m. Evaluation at η_m defines a surjective homomorphism

$$\rho_{\mathfrak{p}_m} : \mathcal{O}_m \longrightarrow \mathbb{F}_{p^{f_m}}, \qquad g(\zeta_m) \longmapsto g(\eta_m),$$

and we write

$$\mathfrak{p}_m = \ker \rho_{\mathfrak{p}_m}.$$

Thus $\mathfrak{p}_m$ is a prime of residue degree f_m above p. For $t \in (\mathbb{Z}/m\mathbb{Z})^{\times}$, let

$$\sigma_t(\zeta_m) = \zeta_m^t, \qquad \mathfrak{p}_{m,t} = \sigma_t^{-1}(\mathfrak{p}_m).$$

If $\mathcal{H}_m = \langle p \rangle \subseteq (\mathbb{Z}/m\mathbb{Z})^{\times}$, then $\mathfrak{p}_{m,t}$ depends only on the coset $t\mathcal{H}_m$, and these cosets index the primes of L_m above p.

For $k \geq 1$, reduction modulo $\mathfrak{p}_m^k$ is denoted by

$$\rho_{\mathfrak{p}_m^k} : \mathcal{O}_m \longrightarrow \mathcal{O}_m/\mathfrak{p}_m^k.$$

The prime p is unramified in L_m. If $g_m(X)$ is the irreducible factor of $\Phi_m(X)$ mod p determined by η_m, and $\widetilde{g}_{m,k}(X)$ is its Hensel lift modulo p^k, then

$$\mathcal{O}_m/\mathfrak{p}_m^k \cong (\mathbb{Z}/p^k\mathbb{Z})[X]/(\widetilde{g}_{m,k}(X)).$$

In the split case $m \mid (p-1)$, one has $f_m = 1$. The chosen root $\eta_m \in \mathbb{F}_p$ has a unique lift $\eta_{m,k}$ mod p^k as an m-th root of unity, and the reduction map becomes

$$\rho_{\mathfrak{p}_m^k} : \mathcal{O}_m \longrightarrow \mathbb{Z}/p^k\mathbb{Z}, \qquad \zeta_m \longmapsto \eta_{m,k}.$$

The reduction maps also apply to elements integral at $\mathfrak{p}_m$, through their images in the local ring.

When one conductor q is fixed, the choices at its divisors are made compatibly. Namely, for $m \mid q$, set

$$\zeta_m = \zeta_q^{q/m}, \qquad \eta_m = \eta_q^{q/m}.$$

Then η_m has order m, and the corresponding reduction maps agree with the natural embeddings among the fields L_m.

In the split construction, where $q \mid (p-1)$, we use the abbreviations

$$K = \frac{p-1}{q}, \qquad \eta = \eta_q, \qquad \mathfrak{p} = \mathfrak{p}_q.$$

We also write $\mathfrak{p}_t = \mathfrak{p}_{q,t}$, and, when working modulo p^2, $\widetilde{\eta} = \eta_{q,2}$.

2.2 Characters, Gauss Sums, and Jacobi Sums

Multiplicative characters encode products of field elements by products of roots of unity. Gauss and Jacobi sums combine this multiplicative information with additive relations. The reconstruction algorithm in Section 3 will compute the sums from their algebraic properties.

Let $Q = p^f$, let $m \mid (Q-1)$, and choose the compatible prime $\mathfrak{p}_m$ and root η_m above. Reduction identifies the m-th roots of unity in L_m with $\langle \eta_m \rangle \subseteq \mathbb{F}_Q^\times$. We define the character $\chi_m : \mathbb{F}_Q^\times \to \langle \zeta_m \rangle$ by

$$\rho_{\mathfrak{p}_m}(\chi_m(x)) = x^{-(Q-1)/m},$$

and extend it to $\mathbb{F}_Q$ by $\chi_m(0) = 0$. Following van Wamelen [5, Section 1], we extend every character power, including the trivial character, by zero at the origin. If a conductor q is fixed and $m \mid q$, then

$$\chi_m = \chi_q^{q/m}.$$

The trace map from $\mathbb{F}_Q$ to $\mathbb{F}_p$ is given by

$$\mathrm{Tr}_{\mathbb{F}_Q/\mathbb{F}_p}(x) = \sum_{i=0}^{f-1} x^{p^i}.$$

The standard additive character of $\mathbb{F}_Q$ is

$$\psi_Q(x) = \exp\!\left(\frac{2\pi i}{p}\,\mathrm{Tr}_{\mathbb{F}_Q/\mathbb{F}_p}(x)\right).$$

For a nontrivial multiplicative character θ, define the Gauss sum

$$G_Q(\theta) = \sum_{x \in \mathbb{F}_Q} \theta(x)\psi_Q(x).$$

For multiplicative characters $\theta_1, \dots, \theta_s$, define the Jacobi sum

$$J_Q(\theta_1, \dots, \theta_s) = \sum_{\substack{x_1,\dots,x_s \in \mathbb{F}_Q \\ x_1+\cdots+x_s=1}} \prod_{j=1}^{s} \theta_j(x_j).$$

Young defines both sums with a leading minus sign [4, Equations (2.4) and (2.7)].

For integers u, v, we abbreviate

$$\mathcal{J}_{Q,m}(u, v) = J_Q(\chi_m^u, \chi_m^v).$$

Only the residue classes of u and v modulo m matter. When $Q = p$, we omit Q from the notation.

If none of $u, v, u+v$ is divisible by m, then

$$\mathcal{J}_{Q,m}(u, v)\overline{\mathcal{J}_{Q,m}(u, v)} = Q. \tag{9}$$

The degenerate cases are

$$\mathcal{J}_{Q,m}(u, v) = \begin{cases} Q-2, & m \mid u \text{ and } m \mid v, \\ -1, & \text{one of } u, v, \text{ but not both, is divisible by } m, \\ -\chi_m(-1)^u, & m \nmid u \text{ and } u+v \equiv 0 \pmod{m}. \end{cases} \tag{10}$$

Proposition 2.1 (Gauss–Jacobi identities)**.** *Suppose that $\theta_1, \dots, \theta_s$ and $\theta_1 \cdots \theta_s$ are nontrivial. Then*

$$J_Q(\theta_1, \dots, \theta_s) = \frac{G_Q(\theta_1) \cdots G_Q(\theta_s)}{G_Q(\theta_1 \cdots \theta_s)}.$$

Moreover,

$$G_Q(\theta^p) = G_Q(\theta).$$

In the split setting $Q = p$, let $\chi = \chi_q$ and define

$$J_a(\chi) = J_p(\underbrace{\chi, \dots, \chi}_{a \text{ times}}), \qquad 1 \le a < q.$$

Then

$$J_a(\chi) = \frac{G_p(\chi)^a}{G_p(\chi^a)}, \qquad J_a(\chi)\overline{J_a(\chi)} = p^{a-1},$$

and, for positive u, v with $u + v < q$,

$$J_{u+v}(\chi) = J_u(\chi) J_v(\chi) \mathcal{J}_{p,q}(u, v). \tag{11}$$

These identities are standard; see, for example, [20, Chapters 1 and 2]. A proof of the formulas used in the split construction is included in Section B.

Before using higher cyclotomic characters, we remove the bounded factor of the conductor that could lower their residue degree. This keeps all nontrivial divisors of the remaining conductor in the same finite field.

Lemma 2.2. *Let $r = \ell^e$ be an odd prime power. The factor of $\Phi_r(p)$ supported on ℓ is either one or ℓ. Removing that factor from the given conductor changes the higher cyclotomic running-time bound only by a constant depending on r.*

Proof. Set $h = \ell^{e-1}$. Modulo ℓ, $\Phi_r(p) \equiv (p^h - 1)^{\ell-1} \pmod{\ell}$, so divisibility by ℓ implies $p^h \equiv 1 \pmod{\ell}$. In that case the binomial expansion of $((1+t)^\ell - 1)/t$, with $t = p^h - 1$, gives $\Phi_r(p) \equiv \ell \pmod{\ell^2}$. Thus the removed factor is at most ℓ. The term q^{c_r} decreases, and the interval term increases by at most $\ell^{1/(2\varphi(r))}$. The compatible character at the remaining conductor is the corresponding power of the given character. If that conductor is one, then $q \le \ell$, and direct interval products satisfy the stated bound for fixed k, with polynomial overhead in k. For $n \ge p$, separating multiples of p and using Section 7.5 gives the same conclusion. The small-prime branch used there also has polynomial cost in k. □

For the higher cyclotomic constructions, let $r = \ell^e$ be an odd prime power, and set

$$h = \ell^{e-1}, \qquad d = \varphi(r) = r - h, \qquad Q = p^r, \qquad D_r = \frac{\Phi_r(p)}{q},$$

where q now denotes the conductor after Lemma 2.2. We may assume $q > 1$ and $\gcd(q, r) = 1$, since the other cases use direct computation.

Lemma 2.3. *For every $m > 1$ dividing q,*

$$\gcd(m, p^h - 1) = 1, \qquad \operatorname{ord}_m(p) = r.$$

Proof. The identity

$$\Phi_r(p) = 1 + p^h + \cdots + p^{(\ell-1)h}$$

gives $\Phi_r(p) \equiv \ell \pmod{p^h - 1}$. Since m is coprime to ℓ, it follows that $\gcd(m, p^h - 1) = 1$. Also, $p^r \equiv 1 \pmod{m}$, so $\operatorname{ord}_m(p)$ divides $r = \ell^e$. Every proper divisor of r divides h, and would therefore imply $m \mid p^h - 1$. The order must be r. □

The compatible characters $\chi_m = \chi_q^{q/m}$, $m \mid q$, can therefore all be viewed on $\mathbb{F}_{p^r}^\times$. For $m > 1$, define

$$\Theta_{r,m} = \mathcal{J}_{p^r,m}(1, p^h - 1),$$

and let $\Theta_{r,1} = p^r - 2$, the value at the trivial conductor. For $u \in (\mathbb{Z}/m\mathbb{Z})^\times$, the Gauss–Jacobi identity and Frobenius invariance give

$$\sigma_u(\Theta_{r,m}) = G_{p^r}(\chi_m^{(p^h-1)u}). \tag{12}$$

Moreover,

$$\Theta_{r,m}\overline{\Theta_{r,m}} = p^r \qquad (m > 1).$$

The family $(\Theta_{r,m})_{m|q}$ determines the Gauss sums used by the recursion, including those for character powers of smaller conductor. The relation between the cyclotomic quotient and the Gauss exponent is

$$\frac{p^r - 1}{q} = (p^h - 1)D_r.$$

2.3 The p-adic Gamma Function

The next formulas explain how local values of Gauss sums determine factorial ratios. The p-adic gamma function interpolates products with the multiples of p omitted. It is therefore well suited to separating a character sum's known valuation from its remaining unit, which is the part used by the factorial algorithm.

We follow the notation of Long and Ramakrishna [21, Section 2] for Morita's p-adic gamma function and its reflection formula. For a positive integer n, the function is defined by

$$\Gamma_p(n) = (-1)^n \prod_{\substack{1 \le j < n \\ p \nmid j}} j,$$

and extends continuously to a function $\Gamma_p : \mathbb{Z}_p \to \mathbb{Z}_p^\times$, where $\mathbb{Z}_p$ is the ring of p-adic integers and $\mathbb{Z}_p^\times$ its unit group. It satisfies

$$\frac{\Gamma_p(x+1)}{\Gamma_p(x)} = \begin{cases} -x, & x \in \mathbb{Z}_p^\times, \\ -1, & x \in p\mathbb{Z}_p, \end{cases}$$

and, for odd p, the reflection identity

$$\Gamma_p(x)\Gamma_p(1-x) = (-1)^{a_0(x)},$$

where $a_0(x) \in \{1, \dots, p\}$ is congruent to x mod p.

To state the Gross–Koblitz formula [22], let $Q = p^f$, and let ω_Q be the Teichmüller character of $\mathbb{F}_Q^\times$: its value at $x \neq 0$ is the $(Q-1)$-st root of unity reducing to x. Extend the chosen local embedding to include $\zeta_p = e^{2\pi i/p}$, and choose π with

$$\pi^{p-1} = -p, \qquad \zeta_p \equiv 1 + \pi \pmod{\pi^2}.$$

This choice agrees with the additive character ψ_Q. For $0 < a < Q-1$, write the base-p expansion

$$a = a_0 + a_1 p + \cdots + a_{f-1} p^{f-1}, \qquad s_p(a) = a_0 + \cdots + a_{f-1}.$$

The Gross–Koblitz formula is

$$G_Q(\omega_Q^{-a}) = -\pi^{s_p(a)} \prod_{i=0}^{f-1} \Gamma_p\left(\left\langle \frac{ap^i}{Q-1} \right\rangle\right), \tag{13}$$

where $\langle x \rangle \in [0,1)$ denotes the fractional part of x. Its first-order consequence is Stickelberger's congruence

$$G_Q(\omega_Q^{-a}) \equiv -\frac{\pi^{s_p(a)}}{a_0! \cdots a_{f-1}!} \pmod{\pi^{s_p(a)+p-1}}. \tag{14}$$

The higher-power algorithms use Taylor expansions of quotients of p-adic gamma values. The function is locally analytic on $\mathbb{Z}_p$, so modulo p^k only finitely many Taylor coefficients are relevant. The number needed and the work to determine them are polynomial in k for each fixed r. We use the standard derivative estimates collected in [21, Section 2]; the finite-precision statements are proved when the lifted recursion is introduced.

2.4 Stickelberger Factorizations

We rewrite van Wamelen's form of Stickelberger's factorization [5, Theorem 3] with our character convention and indexing of the primes above p. The exponents below count carries and can be computed before the sum itself is known.

Let $m \mid (Q-1)$, assume $Q = p^{f_m}$ with $f_m = \mathrm{ord}_m(p)$, and let $\mathcal{H}_m = \langle p \rangle \subseteq (\mathbb{Z}/m\mathbb{Z})^\times$. Choose representatives for the cosets in $(\mathbb{Z}/m\mathbb{Z})^\times/\mathcal{H}_m$, and replace u and v by their representatives in $\{0, \dots, m-1\}$. For $s \in (\mathbb{Z}/m\mathbb{Z})^\times$, represented by an integer in $\{1, \dots, m-1\}$, set

$$c_s(u,v) = \left\lfloor \frac{(u+v)s}{m} \right\rfloor - \left\lfloor \frac{us}{m} \right\rfloor - \left\lfloor \frac{vs}{m} \right\rfloor.$$

For nondegenerate u, v, define the orbit exponent

$$e_t(u,v) = \sum_{j=0}^{f_m-1} c_{[tp^j]_m}(u,v),$$

where $[x]_m \in \{1,\dots,m-1\}$ is the representative of a nonzero class modulo m. With the inverse Galois-action convention used above, Stickelberger's factorization is

$$(\mathcal{J}_{Q,m}(u,v)) = \prod_{t\in(\mathbb{Z}/m\mathbb{Z})^\times/\mathcal{H}_m} \mathfrak{p}_{m,t}^{e_t(u,v)}. \tag{15}$$

In the split case $m \mid (p-1)$, one has $f_m = 1$, and (15) becomes

$$(\mathcal{J}_m(u,v)) = \prod_{t\in(\mathbb{Z}/m\mathbb{Z})^\times} \mathfrak{p}_{m,t}^{c_t(u,v)}. \tag{16}$$

For the multifold sum, the binary factors telescope.

Proposition 2.4 (Split multifold factorization)**.** *For every $1 \le a < q$,*

$$(J_a(\chi)) = \prod_{t\in(\mathbb{Z}/q\mathbb{Z})^\times} \mathfrak{p}_{q,t}^{\lfloor at/q\rfloor}.$$

A derivation is included in Section C.

2.5 Polarized Ideals and Product Representations

The ideal factorization alone leaves a unit ambiguity. We now include the conjugation product to reduce that ambiguity to finitely many roots of unity, and specify how the resulting algebraic values are stored.

A CM-field is a totally imaginary quadratic extension of a totally real number field. Its nontrivial automorphism over the maximal real subfield is complex conjugation. The cyclotomic fields $\mathbb{Q}(\zeta_m)$, for $m > 2$, are CM-fields. An order in a CM-field that is stable under complex conjugation is a CM-order; this is the case used below.

Following the notation of Lenstra and Silverberg [2, Theorem 1.3], let A be a CM-order. A *polarized ideal* is a pair (I, w), where I is an invertible fractional A-ideal and w is totally positive (positive under every real embedding) in the fixed real subalgebra, such that

$$I\overline{I} = wA.$$

It is principal as a polarized ideal if there exists α satisfying

$$I = \alpha A, \qquad \alpha\overline{\alpha} = w.$$

If α and β generate the same polarized ideal, then α/β is a unit whose image under every complex embedding has absolute value one. Kronecker's theorem therefore implies that α/β is a root of unity.

Following the formal-product approach of Kirchner [23, Definition 8], we use a *product representation* of the form

$$\alpha = \prod_{j=1}^{s} \alpha_j^{z_j},$$

where the factors α_j are given by rational coordinates in a fixed basis and the exponents z_j are integers. Its size is the total bit length of the factors and exponents. Such a representation may describe an algebraic number of large height without expanding the product.

In the split construction, an addition chain represents $J_a(\chi)$ as a product of $O(\log a)$ two-character Jacobi sums. In the higher cyclotomic construction, one reconstructs the sums $\Theta_{r,m}$, $m \mid q$, from their ideals and conjugation products. Equation (12), conductor reduction, and Galois conjugation then identify the Gauss sums whose local values enter the factorial identities.

2.6 Computational Representations and Local Evaluation

The algorithms need residues of algebraic products. The following representation lets us apply a Galois automorphism and then evaluate each factor in the local ring. Keeping the precision as an explicit parameter will also be useful when k grows.

A number field is represented by a monic irreducible defining polynomial over $\mathbb{Z}$, together with an integral basis of its ring of integers. Ideals are represented by integer bases, usually in Hermite normal form. When a prime-ideal factorization is already known, it is stored as a list of prime ideals and exponents.

For a cyclotomic field, we use

$$L_m \cong \mathbb{Q}[X]/(\Phi_m(X)), \qquad \mathcal{O}_m = \mathbb{Z}[\zeta_m],$$

with power basis

$$1, \zeta_m, \ldots, \zeta_m^{\varphi(m)-1}.$$

The cyclotomic polynomial, multiplication and conjugation maps, and the power basis can all be constructed in $m^{O(1)}$ bit operations. The logarithm of the absolute discriminant is $O(\varphi(m) \log m)$ [24, Chapter 2].

If the residue degree is one, the prime $\mathfrak{p}_{m,t}$ is represented by

$$(p, \zeta_m - \eta_m^t).$$

In general, if $g_{m,t}(X)$ is the irreducible factor of $\Phi_m(X) \bmod p$ corresponding to $\mathfrak{p}_{m,t}$, then

$$\mathfrak{p}_{m,t} = (p, g_{m,t}(\zeta_m)).$$

All ideals and algebraic numbers appearing in the Stickelberger factorizations above have descriptions of length polynomial in m and $\log Q$.

For nonzero $x \in L_m$, let $v_{\mathfrak{p}_m}(x)$ be the exponent of $\mathfrak{p}_m$ in the principal fractional ideal (x). We also use the p-adic valuation v_p, normalized by $v_p(p) = 1$.

Lemma 2.5 (Local evaluation of product representations)**.** *Let $k \geq 1$, and let*

$$\alpha = \prod_{j=1}^{s} \alpha_j^{z_j}$$

be a product representation in L_m. If every factor has valuation zero at $\mathfrak{p}_m$, then

$$\rho_{\mathfrak{p}_m^k}(\alpha) \in \mathcal{O}_m/\mathfrak{p}_m^k$$

can be computed deterministically in time polynomial in m, k, $\log p$, and the representation length.

Proof. Clear the rational-coordinate denominators of a factor, writing it as A/b, with $A \in \mathcal{O}_m$ and a positive integer b. If $t = v_p(b)$, the valuation assumption gives $v_{\mathfrak{p}_m}(A) = t$. Compute the representative of A of degree less than f_m in the Hensel-lifted quotient modulo p^{k+t}. Its coefficients are divisible by p^t; divide them by p^t and multiply by the inverse of b/p^t modulo p^k. This computes the image of the factor in the local quotient. The integer t is bounded by the bit length of b, so the additional precision is polynomial in the input length.

Each resulting residue is a unit. Negative exponents are therefore handled by inversion, and the product is evaluated by binary exponentiation without expansion in L_m. □

If an integral element α has known valuation $v_{\mathfrak{p}_m}(\alpha) = e$, then its normalized local unit

$$p^{-e}\alpha \pmod{\mathfrak{p}_m^k}$$

is obtained from $\alpha \bmod \mathfrak{p}_m^{e+k}$. Since p is a uniformizer in the localization at $\mathfrak{p}_m$, multiplication by p^e identifies $\mathcal{O}_m/\mathfrak{p}_m^k$ with $\mathfrak{p}_m^e/\mathfrak{p}_m^{e+k}$. This is the form of local reduction used for the normalized Gauss sums in the higher cyclotomic constructions.

For every binary sum used in a split addition chain, the exponent of the distinguished prime $\mathfrak{p}_{q,1}$ in (16) is zero because $u + v < q$. Consequently, Lemma 2.5 applies factor by factor.

2.7 Fast Interval Products and Shifted Factorials

After the character-sum residue has been obtained, only short consecutive products remain. We use the same product routine over finite fields and over the coefficient algebras needed at higher precision. Its cost is stated in algebra operations so that the cost of representing those operations can be counted separately.

Lemma 2.6 (Bostan–Gaudry–Schost)**.** *Let $k \geq 1$, and let A belong to $\mathbb{F}_p$, $\mathbb{Z}/p^k\mathbb{Z}$, or a finite free algebra of specified finite rank over one of these rings. For $0 \leq M < p$, the interval product*

$$\prod_{j=1}^{M}(A + j)$$

can be computed using $\widetilde{O}(1 + \sqrt{M})$ operations in the chosen algebra.

Proof. The case $M = 0$ is immediate. Otherwise, set $B = \lceil \sqrt{M} \rceil$ and form

$$F(X) = \prod_{j=1}^{B}(X + j).$$

Evaluate F at the starting points of the complete blocks by a product tree and a remainder tree. The remaining incomplete block is multiplied directly. The polynomial divisors in the remainder tree are monic, so the same argument works over the stated rings. See Bostan, Gaudry, and Schost [1]. □

Define the shifted factorial polynomial

$$P_M(Z) = \prod_{j=1}^{M}(j + pZ).$$

For every integer $k \geq 1$, the coefficient of Z^j is divisible by p^j, so reduction modulo p^k gives

$$\deg(P_M(Z) \bmod p^k) < k.$$

All shifted interval products in the higher-power algorithms are computed in

$$(\mathbb{Z}/p^k\mathbb{Z})[Z]/(Z^k).$$

This algebra has rank k. With dense coefficient representations, ordinary quadratic polynomial multiplication uses a polynomial number of bit operations in k and $\log p$. Thus the interval bound remains $\widetilde{O}(1+\sqrt{M})$ in algebra operations, with polynomial precision overhead in bit operations.

2.8 Algorithmic Ingredients

We now have the representations needed by the two reconstruction algorithms. The first ingredient recovers an algebraic generator from its ideal and conjugation product; the second removes the root of unity that this recovery leaves undetermined.

Theorem 2.7 (Lenstra–Silverberg)**.** *Let A be a CM-order, let I be an invertible fractional A-ideal, and let w be totally positive with*

$$I\overline{I} = wA.$$

There is a deterministic polynomial-time algorithm that decides whether there exists α satisfying

$$I = \alpha A, \qquad \alpha\overline{\alpha} = w,$$

and returns such an α when one exists.

This is the polarized-principal-ideal theorem of Lenstra and Silverberg [2], building on the Gentry–Szydlo method [25]. Kirchner gives a complementary treatment for cyclotomic and, more generally, CM-orders [23]. Since the cyclotomic orders, ideals, and polarizations occurring here have description length polynomial in m and $\log Q$, its running time is polynomial in these two parameters.

The preceding theorem gives a generator with the right polarization, but its modular residue may still differ from the desired one. The next criterion resolves this remaining ambiguity using the Jacobi sums at the proper divisors of the conductor.

Lemma 2.8 (Root-of-unity orientation)**.** *Let $Q = p^f$, let $m > 2$ divide $Q - 1$, and suppose the values of $\mathcal{J}_{Q,m'}(u, v)$ are known for every proper divisor $m' \mid m$. If*

$$\beta = \xi\,\mathcal{J}_{Q,m}(u, v)$$

is known for some root of unity $\xi \in L_m$, then $\mathcal{J}_{Q,m}(u, v)$ can be recovered deterministically in time polynomial in m, f, and $\log p$.

Proof. The degenerate cases are evaluated by (10). Assume henceforth that none of $u, v, u + v$ is divisible by m. The field L_m contains at most $\operatorname{lcm}(m, 2) \le 2m$ roots of unity. Following van Wamelen's polynomial construction [5, Section 5], combine each candidate $\omega\beta$ at conductor m with the known Jacobi sums at the proper divisors of m. The Chinese remainder isomorphism

$$\mathbb{Q}[X]/(X^m - 1) \cong \prod_{m' \mid m} \mathbb{Q}[X]/(\Phi_{m'}(X))$$

produces a unique polynomial

$$A_\omega(X) = \sum_{j=0}^{m-1} a_j X^j$$

of degree less than m with these prescribed components. Van Wamelen proves that a unique candidate has integral coefficients and satisfies

$$\sum_{j=0}^{m-1} j a_j \equiv \begin{cases} 0 \pmod m, & m \text{ odd}, \\ \dfrac{Q-1}{2}\bigl(\gcd(m,u)+\gcd(m,v)\bigr) \pmod m, & m \text{ even}. \end{cases}$$

That candidate is $\mathcal{J}_{Q,m}(u,v)$ [5, Theorem 9 and Section 7].

There are at most $2m$ candidates and $\tau(m) = m^{o(1)}$ cyclotomic components [26], where $\tau(m)$ counts the divisors of m. Polynomial Chinese remaindering, integrality testing, and verification of the congruence therefore take time polynomial in m, f, and $\log p$. □

3 Jacobi-Sum Reconstruction

The two factorial algorithms use different Jacobi sums, but their reconstruction follows the same pattern. Stickelberger's theorem determines the principal ideal generated by each sum, its product with its complex conjugate determines a polarization, and Theorem 2.7 recovers a generator with that polarization. The remaining root-of-unity ambiguity is then resolved by Lemma 2.8.

Armed with these two ingredients, we can reconstruct the values needed by the factorial algorithms. The procedures below follow the order in which the work is performed. This also separates the family constructed once from the local evaluations repeated during the recursion.

3.1 The split reconstruction

Let $q \mid (p-1)$, let $1 \le a < q$, and retain the compatible character χ and prime $\mathfrak{p}$ from Section 2. An addition chain reduces the multifold sum to a logarithmic number of binary sums. Processing their conductors in increasing order ensures that all inputs to the orientation criterion are available at every step.

Input. The prime p, the divisor q, the index a, and the compatible pair $(\chi, \mathfrak{p})$.

Output. A product representation of $J_a(\chi)$, together with its reductions modulo $\mathfrak{p}$ and $\mathfrak{p}^2$.

Procedure.

1. Choose an addition chain
$$1 = a_0 < a_1 < \cdots < a_s = a,$$
where every a_i, $i \ge 1$, is the sum of two earlier terms. A binary addition chain has length $O(\log a)$.

2. For every relation
$$a_i = a_{r_i} + a_{t_i}$$
in the addition chain, reconstruct the binary Jacobi sum $\mathcal{J}_q(a_{r_i}, a_{t_i})$. To do so, process the divisors $m \mid q$ in increasing order. At conductor m, let $u, v \in \{0, \dots, m-1\}$ be the reductions of a_{r_i}, a_{t_i} modulo m.

   If $\mathcal{J}_m(u,v)$ is degenerate, evaluate it using (10). Otherwise form the Stickelberger ideal
$$I_m(u,v) = \prod_{w \in (\mathbb{Z}/m\mathbb{Z})^\times} \mathfrak{p}_{m,w}^{c_w(u,v)}.$$

Apply Theorem 2.7 to the polarized ideal

$$(I_m(u,v), p).$$

The returned generator differs from $\mathcal{J}_m(u,v)$ by a root of unity. Use Lemma 2.8, together with the Jacobi sums already computed at the proper divisors of m, to determine $\mathcal{J}_m(u,v)$. Continue until $m = q$.

3. Starting from $J_1(\chi) = 1$, follow the addition chain. Whenever $a_i = a_{r_i} + a_{t_i}$, form

$$J_{a_i}(\chi) = J_{a_{r_i}}(\chi) J_{a_{t_i}}(\chi) \mathcal{J}_q(a_{r_i}, a_{t_i}).$$

Retain these products in factored form, combining repeated factors by adding their integer exponents. There are $O(\log a)$ binary factors, and every multiplicity is at most a, so its exponent has $O(\log a)$ bits. At the final step this gives a product representation of $J_a(\chi)$.

4. Apply Lemma 2.5 to the final product representation with $k = 1$ and $k = 2$.

Theorem 3.1 (Split Jacobi reconstruction)**.** *The preceding procedure deterministically constructs a product representation of $J_a(\chi)$, together with its reductions modulo $\mathfrak{p}$ and $\mathfrak{p}^2$.*

Proof. We verify the numbered steps in turn. Fix one relation in the addition chain and one conductor $m \mid q$. In the nondegenerate case, (16) gives

$$I_m(u,v) = (\mathcal{J}_m(u,v)),$$

while (9) gives

$$\mathcal{J}_m(u,v)\overline{\mathcal{J}_m(u,v)} = p.$$

If β is returned by polarized recovery, then

$$\frac{\beta}{\overline{\mathcal{J}_m(u,v)}}$$

is a unit whose image under every complex embedding has absolute value one. Kronecker's theorem therefore shows that this quotient is a root of unity.

The divisors of q are processed in increasing order. The degenerate values are explicit, and at every nondegenerate conductor all proper-conductor values used in Lemma 2.8 have already been computed. An induction over the divisors of q therefore proves that Step 2 returns the correct top-conductor binary Jacobi sum for every relation in the addition chain.

A second induction, this time along the addition chain, proves that the factored expression constructed in Step 3 at index a_i is $J_{a_i}(\chi)$. Finally,

$$a_{r_i} + a_{t_i} = a_i < q$$

for every top-conductor relation, and hence

$$c_1(a_{r_i}, a_{t_i}) = 0.$$

Every top-conductor factor therefore has valuation zero at the distinguished prime $\mathfrak{p}$, so Lemma 2.5 applies factor by factor in Step 4. □

3.2 Reconstruction for odd prime-power indices

Let $r = \ell^e$, $h = \ell^{e-1}$, and q satisfy the higher cyclotomic conditions of Section 2. Retain the compatible characters $\chi_m = \chi_q^{q/m}$ and primes $\mathfrak{p}_m$, for $m \mid q$. Here one compatible family replaces the addition chain. A later recursion index may lower the character order, but its conductor remains in this family.

Input. The prime p, the divisor q, the fixed index r, a representation of $\mathbb{F}_{p^r}$, and the compatible character and prime above p.

Output. The family $(\Theta_{r,m})_{m|q}$, together with the normalized local values of the Gauss sums requested by the factorial recursion.

Procedure.

1. Compute the divisors of q and process them in increasing order. Start with $\Theta_{r,1} = p^r - 2$. For every $m > 1$, use the orbit exponents from (15) to form

$$I_m = \prod_{t \in (\mathbb{Z}/m\mathbb{Z})^\times/\langle p \rangle} \mathfrak{p}_{m,t}^{e_t(1,p^h-1)}.$$

   Apply Theorem 2.7 to (I_m, p^r). Its returned generator differs from $\Theta_{r,m}$ by a root of unity. Apply Lemma 2.8, using the values already obtained at the proper divisors of m, to determine $\Theta_{r,m}$.

2. When the recursion requests the value indexed by $1 \le a < q$, set

$$m = \frac{q}{\gcd(a,q)}, \qquad a' = \frac{a}{\gcd(a,q)}.$$

   Then $m > 1$, $\gcd(a', m) = 1$, and $\chi_q^a = \chi_m^{a'}$. By Lemma 2.3, compute the unit

$$u \equiv a'(p^h - 1)^{-1} \pmod{m}.$$

   Equation (12) gives

$$G_{p^r}(\chi_q^a) = \sigma_u(\Theta_{r,m}). \tag{17}$$

   Apply this automorphism, reduce at the local precision used by the factorial algorithm, and divide by the known power of p to obtain the normalized unit.

The first step is performed once. The second gives the local values requested during the factorial recursion: a bounded number per stage modulo p, and $O_r(k)$ per stage in the interpolation routine modulo p^k. Lowering the conductor preserves the Gauss exponent and its digits:

$$a\frac{p^r - 1}{q} = a'\frac{p^r - 1}{m}. \tag{18}$$

Theorem 3.2 (Jacobi reconstruction for odd prime-power indices)**.** *The preceding procedure deterministically reconstructs $(\Theta_{r,m})_{m|q}$. For every $1 \le a < q$, the normalized local unit associated with $G_{p^r}(\chi_q^a)$ can then be computed in time polynomial in q, k, and $\log p$ when the unit is requested modulo p^k.*

Proof. By Lemma 2.3, the characters χ_m, $\chi_m^{p^h-1}$, and $\chi_m^{p^h}$ are nontrivial for every $m > 1$. The Stickelberger and norm identities give

$$I_m = (\Theta_{r,m}), \qquad \Theta_{r,m}\overline{\Theta_{r,m}} = p^r.$$

Polarized recovery therefore produces a generator whose quotient by $\Theta_{r,m}$ is a root of unity. Since the proper divisors have already been processed, Lemma 2.8 determines the root. Induction over the divisors proves the first step.

For the second step, a' and $p^h - 1$ are units modulo m, so u is a Galois index. Compatibility and (12) prove (17). The local maps at conductors m and q agree on $\mathcal{O}_m$. Applying the automorphism and removing the known local valuation are polynomial-time operations in the representations of Section 2. To evaluate a unit modulo p^k, the precision before this removal is at most p^{r+k}. Lemma 2.5 therefore gives polynomial dependence on k as well as on q and $\log p$. □

3.3 Complexity of the reconstruction

At conductor m, the field arithmetic and ideals have descriptions of length $m^{O(1)}\,\mathrm{polylog}(p)$. In the split case the Stickelberger exponents belong to $\{0, 1\}$; for a fixed odd prime-power index they belong to $\{0, 1, \dots, r\}$. Constructing the ideals and converting them to integer bases therefore have polynomial bit complexity.

Choose a common exponent $C > 1$ covering reconstruction, orientation, and the local operations at one conductor. Summing over the divisors gives

$$\sum_{m|q} m^C = q^C \sum_{j|q} j^{-C} \le q^C \sum_{j\ge 1} j^{-C} = O(q^C).$$

The split addition chain has $O(\log q)$ relations, so its total cost is $\widetilde{O}(q^c)$ for a suitable absolute c. This proves Theorem 1.5.

For a fixed odd prime-power index, the whole family is constructed in $\widetilde{O}_r(q^{c_r})$ time for a suitable $c_r \ge 1$, independently of the later precision. Its algebraic representations have $q^{O(1)}\,\mathrm{polylog}(p)$ bits. Local arithmetic at precision p^k has polynomial dependence on k. One such evaluation therefore costs $\widetilde{O}_r(q^{c_r}k^{O_r(1)})$, after increasing c_r if necessary. The remaining precision costs are accounted for in Section 7.6.

4 The Split Prime-Modulus Algorithm

Retain $K = (p-1)/q$. By Theorem 3.1, the algebraic stage is complete. We are now ready to use its result in a factorial algorithm. The routine of Section 3.1 computes the residue appearing in (2). The remaining part of the algorithm is elementary: choose the block boundary nearest to the desired factorial, evaluate the factorial at that boundary, and multiply or divide by the intervening factors.

4.1 The lower-half routine

We first describe the routine for an index in the lower half of the interval. At this stage the compatible character and prime ideal are assumed to have been fixed.

Input. An integer

$$0 \le m \le \frac{p-1}{2}.$$

Output. The residue $m! \bmod p$.

Procedure.

1. Choose an integer b for which bK is a nearest multiple of K to m, breaking a tie toward the smaller multiple.
2. If $b = 0$, compute $m!$ directly with Lemma 2.6 and stop.
3. If $b > 0$, run the split reconstruction routine of Section 3.1 with index b, and use (2) to compute
$$(bK)! \equiv (-1)^{b-1} \rho_{\mathfrak{p}}(J_b(\chi))\,(K!)^b \pmod{p}.$$
The value $K!$ is computed once with Lemma 2.6.
4. Move from $(bK)!$ to $m!$ by multiplying the factors between bK and m, or by dividing by them when $m < bK$.

Proposition 4.1. *The preceding routine computes $m! \bmod p$ deterministically in*
$$\widetilde{O}\Big(q^c + \sqrt{p/q}\Big)$$
bit operations.

Proof. Since $m \le qK/2$, the nearest block boundary chosen in Step 1 satisfies
$$0 \le b < q, \qquad |m - bK| \le \frac{K}{2}.$$
If $b = 0$, then $m \le K/2$, and the direct computation in Step 2 costs $\widetilde{O}(\sqrt{K})$ operations.

Suppose that $1 \le b < q$. By Theorem 1.5, Step 3 computes $\rho_{\mathfrak{p}}(J_b(\chi))$ in $\widetilde{O}(q^c)$ bit operations. The interval-product algorithm computes $K!$ in $\widetilde{O}(\sqrt{K})$ operations, and exponentiation by b contributes only $O(\log q)$ further field operations. Equation (2) therefore gives $(bK)! \bmod p$.

If $bK < m$, Step 4 uses
$$m! = (bK)! \prod_{j=bK+1}^{m} j.$$
If $m < bK$, it uses
$$m! = (bK)! \left(\prod_{j=m+1}^{bK} j \right)^{-1} \pmod{p}.$$
All factors are nonzero modulo p, because $m, bK < p$. In either case the interval has length at most $K/2$, so the correction costs $\widetilde{O}(\sqrt{K})$ operations. Since
$$\sqrt{K} = \sqrt{\frac{p-1}{q}} \le \sqrt{\frac{p}{q}},$$
the claimed bound follows. □

4.2 Wilson complementation and proof of the theorem

The lower-half routine contains all the nontrivial factorial work. Wilson complementation extends it to the other half, and the final part of the proof accounts for constructing the character and its reduction map from the stated input.

Proof of Theorem 1.1. The cases $p = 2$, $n = 0$, and $n = p - 1$ are immediate. If $q \le 2$, the standard interval-product algorithm already gives the claimed bound, so assume $q \ge 3$.

If

$$n \le \frac{p-1}{2},$$

apply Proposition 4.1 with $m = n$.

Now suppose that $n > (p-1)/2$, and set

$$m = p - 1 - n.$$

Then $0 \le m < (p-1)/2$, and

$$\prod_{j=n+1}^{p-1} j = \prod_{j=1}^{m} (p-j) \equiv (-1)^m m! \pmod{p}.$$

Wilson's theorem therefore gives

$$n! \equiv (-1)^{m+1} (m!)^{-1} \pmod{p}.$$

Thus one call to the lower-half routine also handles the upper half. The final inversion and sign take only polylog(p) bit operations.

It remains to account for the construction of the order-q element and the associated character. As explained in Section 1.3, the deterministic setup costs

$$p^{1/4+o(1)} + q^{O(1)}.$$

Enlarge c, if necessary, so that the polynomial dependence on q is absorbed by q^c. For every $1 \le q \le p-1$, either

$$q \ge p^{1/(2c+1)},$$

in which case $q^c \ge p^{c/(2c+1)}$, or

$$q < p^{1/(2c+1)},$$

in which case

$$\sqrt{p/q} > p^{c/(2c+1)}.$$

Since $c \ge 1$, one has $c/(2c+1) \ge 1/3$, so the stated running time also absorbs $p^{1/4+o(1)}$.

Combining the lower-half routine, Wilson complementation, and the setup gives

$$\widetilde{O}\left(q^c + \sqrt{p/q}\right)$$

bit operations. □

4.3 Divisors of $p+1$

The reconstruction method also applies to divisors of $p+1$. A character on $\mathbb{F}_{p^2}$ whose restriction to $\mathbb{F}_p$ is quadratic gives a grid of binomial coefficients for the factorial recursion.

Proposition 4.2. *There is an absolute constant $C \geq 1$ with the following property. Let p be an odd prime, let $q \geq 4$ divide $p+1$, and suppose that $(p+1)/q$ is odd. Given a representation of $\mathbb{F}_{p^2}$ and a compatible character of order $2q$, a deterministic algorithm computes $n! \bmod p$, for every $0 \leq n < p$, in*

$$\widetilde{O}\Big(q^C + \sqrt{p/q}\Big)$$

bit operations. The representation and character can be constructed by Las Vegas procedures within the same expected bound.

Proof. Set $K = (p+1)/q$, $h = (p-1)/2$, and let χ be the compatible character of order $2q$. Its restriction to $\mathbb{F}_p^\times$ is the quadratic character. For odd a, $1 \leq a < q$, define

$$E_a = \sum_{\substack{x \in \mathbb{F}_{p^2} \\ \mathrm{Tr}_{\mathbb{F}_{p^2}/\mathbb{F}_p}(x) = 1}} \chi(x)^a, \qquad k_a = \frac{aK-1}{2}.$$

Appendix G reconstructs the compatible family at the divisors of $2q$ in total $\widetilde{O}(q^C)$ time and gives

$$\rho_{\mathfrak{p}}(E_a) = \binom{h}{k_a}, \qquad \binom{2k_a}{k_a} = (-4)^{k_a}\rho_{\mathfrak{p}}(E_a) \quad \text{in } \mathbb{F}_p. \tag{19}$$

Wilson complementation reduces the target to $0 \leq n \leq h$. If $n \leq 2K$, use Lemma 2.6. Otherwise choose an odd $a \in [1, q)$ nearest to $(n+1)/K$, and let $z = aK - 1$. The grid spacing is $2K$, so $|z - n| \leq K$. Its endpoints cover the target range: the smallest grid value is $K-1$, and the largest is $p - K \geq h$. One recursive call for $(z/2)!$, followed by

$$z! = (-4)^{z/2}\rho_{\mathfrak{p}}(E_a)((z/2)!)^2 \pmod p,$$

and an interval correction of length at most K, gives $n!$. The child satisfies $z/2 \leq (n+K)/2 < 3n/4$, so there are $O(\log p)$ stages. All direct products have length $O(K)$. The reconstruction cost and Lemma 2.6 therefore give the stated bound.

For the initial choices, construct a quadratic extension, factor $2q$ within $q^{O(1)}$ work, and sample $x \in \mathbb{F}_{p^2}^\times$ uniformly. Test the order of $x^{(p^2-1)/(2q)}$ at each prime divisor of $2q$. The expected number of trials is $2q/\varphi(2q)$; all tests precede acceptance. This gives the Las Vegas assertion. □

The condition that $(p+1)/q$ is odd says that q contains the full power of 2 dividing $p+1$. For $q = p^{\theta+o(1)}$, the running-time exponent is $\max\{C\theta, (1-\theta)/2\}$, with balanced value $C/(2C+1)$.

4.4 A uniform logarithmic refinement

Modular halving and simultaneous evaluation improve the uniform bound for factorials modulo p and p^2, without a divisor assumption.

Proposition 4.3. *Given a prime $p \geq 5$, deterministic algorithms compute $n! \bmod p$, for $0 \leq n < p$, and $n! \bmod p^2$, for $0 \leq n < p^2$, in*

$$O\Big(\sqrt{p}\,(\log p)^{3/2}\sqrt{\log\log p}\Big)$$

bit operations. The same bound computes the Wilson quotient modulo p.

For either modulus, this improves the general $O(\sqrt{p}\,(\log p)^2)$ Bostan–Gaudry–Schost bound by a factor of order $\sqrt{\log p/\log\log p}$.

The case $n = p-1$ modulo p^2 gives a logarithmic improvement for the problem in Harvey's 2026 Arizona Winter School notes [27, Problem 6.1.8 and Appendix A.1(a)], where he proposes factorial sieving. Here we use modular halving and simultaneous evaluation of shifted factorials.

For $0 \leq a \leq b < p$ and $z \in \mathbb{F}_p$, separating even and odd indices gives

$$\begin{aligned}\prod_{a<j\leq b}(j+pz) \equiv 2^{b-a}&\left(\prod_{\lfloor a/2\rfloor<j\leq\lfloor b/2\rfloor}(j+pz/2)\right)\\ &\times\left(\prod_{\lfloor (a+p)/2\rfloor<j\leq\lfloor (b+p)/2\rfloor}(j+p(z-1)/2)\right)\pmod{p^2}.\end{aligned}\tag{20}$$

The shifts $z/2$ and $(z-1)/2$ are taken in $\mathbb{F}_p$. We stop a branch when its numerical interval enters a prefix of length $O(p\log\log p/\log p)$, and apply only $O(\log p)$ rounds. The remaining branches have much smaller total length. Endpoints and shifts for both collections are constructed in $p^{1/4}\,\mathrm{polylog}(p)$ time.

At each stopped endpoint m, the polynomial $P_m(Z) \bmod p^2$ has degree at most one. The multiple-output theorem of Bostan, Gaudry, and Schost [1, Theorem 15] computes these coefficient pairs together, after which each branch uses its own shift. Batched interval products handle the remaining branches. With the $O(b\log b)$-time integer multiplication algorithm of Harvey and van der Hoeven [28], this gives the stated bit bound. Modulo p, the shifts disappear and the same computation uses only the prefix factorials $m! \bmod p$. Appendix H contains the proof and the reduction for indices at least p.

5 Higher Cyclotomic Algorithms Modulo a Prime

Fix an odd prime power $r = \ell^e$, and retain

$$h = \ell^{e-1}, \qquad d = \varphi(r) = (\ell-1)h, \qquad D_r = \frac{\Phi_r(p)}{q},$$

where $q > 1$, $q \mid \Phi_r(p)$, and $\gcd(q, r) = 1$. The routine of Section 3.2 constructs $(\Theta_{r,m})_{m|q}$ once. Conductor reduction and Galois conjugation then provide all the local values used below. Three steps remain: convert a local value into a multinomial relation, choose its arguments so that the parts are useful, and arrange the factorial computation with only one recursive child.

5.1 The multinomial identity

The identity below specifies what a reconstructed local value tells us about factorials. The digit conditions ensure that multiplication by $p^h - 1$ produces controlled borrows in base p. Within each chain, successive digit differences become the parts of a multinomial coefficient.

Suppose that $1 \le a < q$ and

$$aD_r = \sum_{i=0}^{d-1} x_i p^i, \qquad 1 \le x_u \le x_{u+h} \le \cdots \le x_{u+(\ell-2)h} < p \quad (0 \le u < h). \tag{21}$$

Define

$$\begin{aligned} y_0 &= x_0 - 1, & y_u &= x_u & &(1 \le u < h),\\ y_i &= x_i - x_{i-h} & & & &(h \le i < d),\\ Y_0 &= x_{d-h} - 1, & Y_u &= x_{d-h+u} & &(1 \le u < h). \end{aligned} \tag{22}$$

Then

$$Y_u = \sum_{j=0}^{\ell-2} y_{u+jh}, \qquad Y = \sum_{u=0}^{h-1} Y_u = \sum_{i=0}^{d-1} y_i.$$

All parts and all Y_u lie in $\{0, \dots, p-1\}$.

The digit calculation in Section D shows that $G_{p^r}(\chi_q^a)$ has valuation d at the chosen prime above p. Let U_a be the reduction of its normalized local unit:

$$G_{p^r}(\chi_q^a) \equiv -p^d U_a \pmod{\mathfrak{p}_q^{d+1}}.$$

The reconstruction routine computes this value by (17), also when a is not coprime to q.

Proposition 5.1. *With the notation above,*

$$\prod_{u=0}^{h-1} \binom{Y_u}{y_u, y_{u+h}, \dots, y_{u+(\ell-2)h}} = (-1)^Y U_a^{-1} \quad \text{in } \mathbb{F}_p.$$

Proof. The base-p digits of

$$a\frac{p^r - 1}{q} = (p^h - 1)aD_r$$

are

$$p - 1 - y_0, \dots, p - 1 - y_{d-1}, Y_0, \dots, Y_{h-1}.$$

Their sum is $d(p-1)$. Since $d = \varphi(r)$ is even, the first-order Gross–Koblitz congruence gives

$$U_a^{-1} = \left(\prod_{u=0}^{h-1} Y_u!\right)\left(\prod_{i=0}^{d-1} (p-1-y_i)!\right) \quad \text{in } \mathbb{F}_p.$$

Wilson's identity $(p-1-y)!\, y! = (-1)^{y+1}$, together with $\sum_i y_i = Y$, transforms the right side into

$$(-1)^Y \frac{\prod_u Y_u!}{\prod_i y_i!}.$$

This proves the proposition. □

5.2 The cyclotomic lattice

We use the multinomial identity at arguments near a balanced split. The cyclotomic relation gives the needed control: rotations of one short vector produce a full set of independent short directions. The following lemma turns that structure into a deterministic rounding algorithm.

Lemma 5.2. *Fix an integer $r \geq 3$, and set $d = \varphi(r)$. There is an integer $C_r \geq 1$ and a deterministic algorithm with the following property. Given p, a divisor $D \mid \Phi_r(p)$, and a rational target vector $t \in \mathbb{Q}^d$, it returns $x \in \mathbb{Z}^d$ satisfying*

$$\sum_{i=0}^{d-1} x_i p^i \equiv 0 \pmod{D}$$

and

$$\|x - t\|_\infty \leq C_r D^{1/d}.$$

For fixed r, the running time is polynomial in $\log p$ and the bit length of t. All target vectors used below have bit length polynomial in $\log p$.

Proof. Let Λ be the lattice defined by the congruence in the lemma. The coefficient of the first coordinate is one, so the defining map $\mathbb{Z}^d \to \mathbb{Z}/D\mathbb{Z}$ is surjective and

$$\det \Lambda = D.$$

Construct an integer basis of Λ in Hermite normal form and apply LLL reduction [29]. Since d is fixed, the first reduced vector $v \neq 0$ satisfies

$$\|v\|_2 = O_r(D^{1/d}).$$

Identify coefficient vectors with elements of $\mathbb{Z}[\zeta_r]$ in the power basis. Multiplication by ζ_r preserves Λ: evaluation at p multiplies the defining congruence by p, while reduction by $\Phi_r(\zeta_r) = 0$ changes the value at p by a multiple of $\Phi_r(p)$, and hence by a multiple of D. Let T be the corresponding integer matrix. The vectors

$$v, Tv, \ldots, T^{d-1}v$$

are linearly independent. Indeed, a rational relation among them would, after division by the nonzero element represented by v, give a polynomial of degree less than d vanishing at ζ_r.

The powers of T have operator norm bounded in terms of r, so these vectors have norm $O_r(D^{1/d})$. They generate a full-rank sublattice of Λ. Express t in this basis and round each coordinate to the nearest integer. The resulting lattice point differs from t by at most half the sum of the basis-vector sup norms, and hence by at most $C_r D^{1/d}$ in the sup norm. All operations take place in fixed dimension and have bit complexity polynomial in $\log p$ and the bit length of the rational target.

We will also use the following consequence of this construction. The absolute determinant of the short basis is at most a constant, depending only on r, times D, by the determinant bound in terms of column lengths. Its index in Λ is therefore bounded only in terms of r. Since $D \mid \Phi_r(p)$ implies $p \nmid D$, this basis is nonsingular modulo p whenever p exceeds that fixed index bound. □

5.3 The recursive factorial routine

Armed with the multinomial identity and the lattice rounding lemma, we can choose one chain of large, nearly equal parts and keep the other chains small. Only the first chain contributes a recursive call. The other factorials are removed by direct short products.

Set

$$E = \left\lceil C_r D_r^{1/d} \right\rceil .$$

The following routine returns $N! \bmod p$ for $0 \le N \le (p-1)/2$. The family $(\Theta_{r,m})_{m|q}$ is constructed on the first call that needs it and retained thereafter.

Procedure.

1. If $E \ge p/(8r)$ or $N \le 8rE$, compute $N!$ directly with Lemma 2.6 and stop.

2. Apply Lemma 5.2 to the target t defined by

$$t_{jh} = 1 + \frac{(j+1)N}{\ell - 1}, \quad 0 \le j \le \ell - 2,$$
$$t_{u+jh} = 4(j+1)E, \qquad 1 \le u < h, \quad 0 \le j \le \ell - 2.$$

   Let x be the resulting lattice point. Define its parts and totals by (22), and set

$$a = \frac{\sum_{i=0}^{d-1} x_i p^i}{D_r}.$$

3. Use Section 3.2 to compute U_a. Make one recursive call to obtain $y_0!$. Obtain every other $y_{jh}!$ from $y_0!$ by an interval correction. Compute the factorials $y_i!$ for $h \nmid i$, and $Y_u!$ for $u > 0$, directly with Lemma 2.6.

4. Use Proposition 5.1 to compute

$$Y_0! = (-1)^Y U_a^{-1} \frac{\displaystyle\prod_{i=0}^{d-1} y_i!}{\displaystyle\prod_{u=1}^{h-1} Y_u!} \pmod p,$$

   and obtain $N!$ by one final interval correction. Empty products are interpreted as one.

For an original input $n > (p-1)/2$, apply the routine to $p-1-n$ and use Wilson complementation as in Section 4.

5.4 Correctness and complexity

With the local values from Theorem 3.2, we now verify the two points needed for the saving: lattice rounding must preserve the digit inequalities, and every factorial other than the single recursive child must be obtained from a short product.

Theorem 5.3. *For the fixed odd prime power r and the divisor q specified above, given a representation of $\mathbb{F}_{p^r}$ and a compatible order-q character, the preceding deterministic algorithm computes $n! \bmod p$ in*

$$\widetilde{O}\left(q^{c_r} + D_r^{1/(2\varphi(r))}\right)$$

bit operations for a constant $c_r \ge 1$.

Proof. Wilson complementation reduces the analysis to $N \le (p-1)/2$. Under either stopping condition, $1 + \sqrt{N} = O_r(\sqrt{E})$, since $E \ge 1$. Thus the direct branch satisfies the stated bound.

Otherwise $E < p/(8r)$ and $N > 8rE$. In the chain indexed by $u = 0$, target gaps are $N/(\ell-1) > 8E$. In every other chain they are $4E$. Rounding changes each coordinate by at most E, so all chains are strictly increasing and start above zero. The last coordinate of the first chain is at most $1 + N + E < p$. Every coordinate in another chain is at most $(4(\ell-1)+1)E < p$. Thus (21) holds.

The same estimates give

$$\begin{aligned} y_{jh} &= \frac{N}{\ell-1} + O(E), \quad Y_0 = N + O(E), \\ y_i &= O_r(E) \quad (h \nmid i), \quad Y_u = O_r(E) \quad (u > 0). \end{aligned}$$

The lattice congruence makes a an integer, and

$$0 < \sum_{i=0}^{d-1} x_i p^i < p^d < \Phi_r(p) = qD_r$$

shows that $1 \le a < q$. The local value U_a is therefore available by conductor reduction, and the multinomial identity applies. All factorial arguments are below p, so every division in the routine is by a nonzero residue.

All direct products and interval corrections have length $O_r(E)$. There are only a constant number at each stage, giving $\widetilde{O}(E^{1/2})$ work per recursive call. The single child satisfies

$$y_0 \le \frac{N}{\ell-1} + E < \frac{3N}{4},$$

so the recursion has logarithmic depth. The reconstructed family and all local evaluations contribute $\widetilde{O}(q^{c_r})$, and $E^{1/2} = O_r(D_r^{1/(2\varphi(r))})$. This proves the bound. □

6 The Split Algorithm Modulo p^2

Return to the split setting $q \mid (p-1)$, with $K = (p-1)/q$, and retain the character and local reduction maps from Section 2.6. The same product representation can be reduced modulo $\mathfrak{p}^2$. The lifted factorial identity adds a harmonic term, which we evaluate using a roots-of-unity filter. Define

$$J_a^{(2)} := \rho_{\mathfrak{p}^2}(J_a(\chi)) \in \mathbb{Z}/p^2\mathbb{Z}.$$

For $0 \le s < p$, let

$$H_s = \sum_{j=1}^{s} \frac{1}{j} \pmod{p}, \qquad H_0 = 0.$$

All inverses are taken in $\mathbb{F}_p$.

6.1 The Central Congruence Modulo p^2

Proposition 6.1. *For every $1 \le a < q$,*

$$(aK)! \equiv (-1)^{a-1} J_a^{(2)} (K!)^a \left(1 + \frac{ap}{q}(H_{aK} - H_K)\right) \pmod{p^2}.$$

The proof is given in Section F. It uses the same finite specialization of Young's Jacobi–multinomial congruence as before, followed by an elementary block expansion modulo p^2.

6.2 Evaluating the Harmonic Term

A roots-of-unity filter evaluates the harmonic difference using $q-1$ modular exponentiations, whose cost is absorbed by the algebraic reconstruction.

For $1 \le t < q$, define

$$Q_t = \frac{(1-\widetilde{\eta}^t)^{p-1}-1}{p} \pmod p.$$

The exponentiation is performed modulo p^2. Since $1-\eta^t \neq 0$ in $\mathbb{F}_p$, Fermat's theorem makes the quotient well defined.

Lemma 6.2. *For every $1 \le a < q$,*

$$H_{aK} - H_K \equiv \sum_{t=1}^{q-1} (\eta^{at}-\eta^t) Q_t \pmod p.$$

Proof. The congruence

$$\binom{p}{m} \equiv p\frac{(-1)^{m-1}}{m} \pmod{p^2}, \qquad 1 \le m < p,$$

and $\widetilde{\eta}^p = \widetilde{\eta}$ give

$$(1-\widetilde{\eta}^t)^p \equiv 1-\widetilde{\eta}^t - p\sum_{m=1}^{p-1}\frac{\eta^{tm}}{m} \pmod{p^2}.$$

Dividing by $1-\widetilde{\eta}^t$ yields

$$\sum_{m=1}^{p-1}\frac{\eta^{tm}}{m} \equiv -(1-\eta^t)Q_t \pmod p. \tag{23}$$

Now

$$H_{aK} - H_K = \sum_{b=1}^{a-1}\sum_{u=1}^{K}\frac{1}{bK+u}.$$

Since $q(bK+u) \equiv qu-b \pmod p$,

$$\sum_{u=1}^{K}\frac{1}{bK+u} = q\sum_{\substack{1\le m<p\\ m\equiv -b \pmod q}}\frac{1}{m}.$$

Insert the roots-of-unity filter

$$\mathbf{1}_{m\equiv -b \pmod q} = \frac{1}{q}\sum_{t=0}^{q-1}\eta^{t(m+b)}.$$

The term $t=0$ vanishes because $H_{p-1} = 0 \pmod p$. For $t \neq 0$,

$$\sum_{b=1}^{a-1}\eta^{bt} = \frac{\eta^t-\eta^{at}}{1-\eta^t}.$$

Combining this identity with (23) gives the result. □

6.3 Proof of the Theorem Modulo p^2

We are now ready to prove Theorem 1.2. The reconstructed residue gives the block-boundary value through the lifted congruence, and Lemma 6.2 makes its correction inexpensive.

Proof of Theorem 1.2. The case $p = 2$ is immediate. For $q = 1$, use Lemma 2.6 when $n < p$, and the reduction for larger indices at the end of this proof when $n \geq p$. Assume henceforth that p is odd and $q \geq 2$. The compatible character and local representation can be constructed within the bound established in the proof of Theorem 1.1.

First let $0 \leq n < p$. Choose the nearer block boundary bK among those with $0 \leq b < q$. Except in the final block $(q-1)K < n < p$, the correction interval has length at most $K/2$; in that final block we take $b = q - 1$, so its length is at most K. This includes the endpoint $n = p - 1$.

If $b = 0$, Lemma 2.6 computes $n! \bmod p^2$ directly. For $1 \leq b < q$, the same lemma computes $K!$ and the correction interval in $\widetilde{O}(\sqrt{p/q})$ time. Theorem 1.5 computes

$$J_b^{(2)} = \rho_{\mathfrak{p}^2}(J_b(\chi))$$

in $\widetilde{O}(q^c)$ time. Lemma 6.2 evaluates the harmonic term from $q - 1$ modular exponentiations, a cost $\widetilde{O}(q)$ absorbed by $\widetilde{O}(q^c)$. Proposition 6.1 then gives $(bK)! \bmod p^2$, and the short correction interval gives $n! \bmod p^2$.

Now suppose $p \leq n < 2p$, and write $n = p + s$, $0 \leq s < p$. Then

$$\begin{aligned}(p+s)! &= p(p-1)! \prod_{j=1}^{s} (p+j) \\ &\equiv -p\, s! \pmod{p^2}.\end{aligned}$$

It remains to compute $s! \bmod p$, and Theorem 1.1 does so within the same bound. Finally, if $n \geq 2p$, both p and $2p$ divide $n!$, so $n! \equiv 0 \pmod{p^2}$.

The total deterministic running time is

$$\widetilde{O}\left(q^c + \sqrt{p/q}\right),$$

as claimed. □

Remark 6.3 (Wilson quotients). Let $F \in \{0, \ldots, p^2 - 1\}$ represent $(p-1)! \bmod p^2$. Wilson's theorem gives $p \mid F + 1$, and hence

$$W_p \equiv \frac{F+1}{p} \pmod{p}.$$

Thus Theorem 1.2 computes the Wilson quotient modulo p. The uniform bound of Proposition 4.3 applies to the same computation for every prime.

7 Higher Cyclotomic Algorithms Modulo Prime Powers

We now lift the recursion of Section 5 by returning a shifted factorial polynomial from each call. Its values give the shifted products in the lifted identity. Recall from Section 2.7 that

$$P_N(Z) = \prod_{j=1}^{N} (j + pZ).$$

Its coefficient of Z^j is divisible by p^j, so its reduction modulo p^k has degree less than k.

Fix an odd prime power $r = \ell^e$, let $q > 1$ divide $\Phi_r(p)$ with $\gcd(q, r) = 1$, and let $k \geq 2$ be part of the input. Retain

$$h = \ell^{e-1}, \qquad d = \varphi(r), \qquad D_r = \frac{\Phi_r(p)}{q}, \qquad E = \left\lceil C_r D_r^{1/d} \right\rceil.$$

The family $(\Theta_{r,m})_{m|q}$ is reconstructed when it is first needed and then reused. Each Gauss sum is obtained by the conductor reduction in (17).

The totals outside the distinguished chain remain small, so their shifted factorials can be evaluated directly. Several nearby lattice points satisfying the digit inequalities then give enough values to recover the remaining polynomial by interpolation. To keep track of precision, choose a sufficiently large integer K_r, depending only on r, and set

$$C_{r,k} = K_r(k + 1). \tag{24}$$

In the following proofs, K_r increases by factors depending only on r. This linear threshold bounds both the auxiliary factorial arguments and the error from nearby lattice points.

7.1 The lifted multinomial identity

The same reconstructed Gauss sum is available at any requested local precision. The purpose of the next identity is to separate its normalized unit into short shifted products and a gamma factor. The latter will be determined from small instances of the identity.

Retain the digits, parts, and totals satisfying (21) and (22), including $Y = \sum_u Y_u$. In the completion at the chosen prime above p, define

$$\mathcal{U}_a = -p^{-d} G_{p^r}(\chi_q^a).$$

The Gross–Koblitz formula identifies this normalized unit with a product of Morita gamma values, so $\mathcal{U}_a \in \mathbb{Z}_p^\times$. To compute it modulo p^k, reduce the corresponding conjugate of $\Theta_{r,m}$ modulo $\mathfrak{p}_m^{d+k}$, and remove its known valuation d, as in Lemma 2.5.

The Gross–Koblitz arguments can be ordered cyclically as $\gamma_0, \ldots, \gamma_{r-1}$, with $\gamma_0 = a/q$ and $\gamma_r = \gamma_0$, such that

$$\begin{aligned} p\gamma_{i+1} &= \gamma_i + p - 1 - y_i \quad (0 \leq i < d), \\ p\gamma_{d+u+1} &= \gamma_{d+u} + Y_u \qquad\quad (0 \leq u < h). \end{aligned}$$

These arguments have denominators dividing q. Their numerators are computed by cyclic digit rotation before reduction modulo p^k, so these divisions by p preserve the working precision. Set

$$z_i = \gamma_{i+1} - 1 \quad (0 \leq i < d), \qquad w_u = -\gamma_{d+u+1} \quad (0 \leq u < h).$$

Summing the cyclic relations gives

$$\sum_{i=0}^{d-1} z_i = \sum_{u=0}^{h-1} w_u. \tag{25}$$

Define the gamma factor

$$\mathcal{G}_p(\mathbf{z}, \mathbf{w}) = \frac{\prod_{i=0}^{d-1} \Gamma_p(pz_i)}{\prod_{u=0}^{h-1} \Gamma_p(pw_u)}.$$

Proposition 7.1 (Lifted multinomial identity)**.** *For odd p, with the notation above,*

$$\mathcal{U}_a = (-1)^Y \mathcal{G}_p(\mathbf{z},\mathbf{w}) \frac{\prod_{i=0}^{d-1} P_{y_i}(z_i)}{\prod_{u=0}^{h-1} P_{Y_u}(w_u)} \quad \text{in } \mathbb{Z}_p^\times. \tag{26}$$

The proof is given in Section E. It uses the Gross–Koblitz formula, Morita's functional equation, and reflection. Modulo p, the gamma factor is one and the identity reduces to Proposition 5.1.

7.2 The gamma factor

The relation (25) removes the linear term from the local gamma expansion. For the Robert–Zuber expansion, we use the coefficients λ_m as in Long and Ramakrishna [21, Theorem 7]. Let $\log_p$ denote the p-adic logarithm, defined by $\log_p(1+T) = \sum_{j\ge1}(-1)^{j+1}T^j/j$ for $T \in p\mathbb{Z}_p$. For $p \ge 5$, the expansion gives, on $p\mathbb{Z}_p$,

$$\log_p \Gamma_p(X) = \lambda_0 X - \sum_{m\ge1} \frac{\lambda_m}{2m(2m+1)} X^{2m+1},$$

where $\lambda_1 \in \mathbb{Z}_p$ and $p\lambda_m \in \mathbb{Z}_p$ for $m \ge 2$. Consequently,

$$\log_p \mathcal{G}_p(\mathbf{z},\mathbf{w}) = -\sum_{m\ge1} \frac{\lambda_m p^{2m+1}}{2m(2m+1)} \left(\sum_{i=0}^{d-1} z_i^{2m+1} - \sum_{u=0}^{h-1} w_u^{2m+1} \right). \tag{27}$$

Only $O(k)$ coefficients can contribute modulo p^k in the range used by the algorithm. In particular,

$$\mathcal{G}_p(\mathbf{z},\mathbf{w}) \equiv 1 \pmod{p^3} \qquad (p \ge 5).$$

Thus the gamma factor is one modulo p^2 and p^3; the first possible nontrivial term occurs modulo p^4.

The remaining coefficients are determined from small instances of (26). The following size condition is tested before that construction is attempted.

Lemma 7.2 (Computing the gamma factor)**.** *For fixed r, choose K_r in (24) sufficiently large. If $E < p/C_{r,k}$, one can determine the coefficients needed to evaluate $\mathcal{G}_p(\mathbf{z},\mathbf{w}) \bmod p^k$ from $O(k)$ instances of (26). Every part y_i and every total Y_u in these instances is at most $C_{r,k}E$ and strictly less than p. After reconstruction of $(\Theta_{r,m})_{m|q}$, the additional bit cost is at most*

$$\widetilde{O}_r\Big(k^{O_r(1)}[q^{c_r} + E^{1/2}]\Big)$$

where c_r and the polynomial exponent depend only on r.

The proof is given in Section E. It uses a grid of $(k+1)^d$ candidates to find independent equations, but evaluates the reconstructed algebraic values at only $O(k)$ selected points. For fixed r, both counts are polynomial in k. Choosing K_r large enough ensures $p > \max\{5, 2k+1\}$ whenever the small-E condition holds. The other cases use the direct branch.

7.3 Nearby lattice points for interpolation

To recover a polynomial of degree less than k, we need k values at arguments with pairwise unit differences. We obtain them by short translations of one rounded lattice point. Only Y_0 is placed near the desired index; the other chains remain near zero. Near the upper endpoint, a downward shift keeps every rounded digit below p.

The cyclic relations imply the following reduction for the interpolation argument:

$$w_0 \equiv \begin{cases} -x_0 \pmod p, & h = 1, \\ x_{d-h+1} \pmod p, & h > 1. \end{cases} \tag{28}$$

For $h = 1$, use $w_0 = -\gamma_0 = -a/q$ and $aD_r \equiv x_0$, $qD_r \equiv 1 \pmod p$. For $h > 1$, the relation for γ_{d+1} gives $-\gamma_{d+1} \equiv Y_1 = x_{d-h+1} \pmod p$.

Lemma 7.3. *For fixed r, increase K_r in (24) if necessary, and suppose*

$$0 \le N < p, \qquad N > C_{r,k}E, \qquad E < p/C_{r,k}.$$

One can deterministically find k lattice points $x^{(v)}$, $0 \le v < k$, satisfying (21), whose parts and totals satisfy

$$y_{jh}^{(v)} = \frac{N}{\ell - 1} + O_r((k+1)E), \quad Y_0^{(v)} = N + O_r((k+1)E),$$
$$y_i^{(v)} = O_r((k+1)E) \quad (h \nmid i), \quad Y_u^{(v)} = O_r((k+1)E) \quad (u > 0).$$

The arguments $w_0^{(v)}$ are pairwise distinct modulo p. For fixed r, the construction takes time polynomial in k and $\log p$.

Proof. Use the short full-rank sublattice from Lemma 5.2. Its index in the congruence lattice is bounded in terms of r. Since $p \nmid D_r$, its basis is nonsingular modulo p once p exceeds that bound. The small-E hypothesis ensures this after increasing K_r. Choose a short basis vector whose coordinate specified in (28) is nonzero modulo p.

Round a target to the lattice and add $0, 1, \dots, k-1$ times this vector. If bE, with b depending only on r, bounds the vector's supremum norm, set $B = 1 + (k-1)b$. The total rounding and translation error is then at most BE in each coordinate, with $B = O_r(k+1)$. Use the target t

$$t_{jh} = 1 + \frac{(j+1)(N - 2BE)}{\ell - 1}, \quad 0 \le j \le \ell - 2,$$
$$t_{u+jh} = (j+1)(2B+2)E, \qquad 1 \le u < h, \quad 0 \le j \le \ell - 2.$$

For sufficiently large K_r, the first coordinate in each chain is positive and every successive gap remains positive after an error of $2BE$. The last coordinate in the first chain is at most

$$1 + N - 2BE + BE \le N < p.$$

Every other coordinate is bounded by $((\ell - 1)(2B+2) + B)E$, which is below p after increasing K_r. Thus all points satisfy (21). Their parts and totals have the stated sizes.

The chosen coordinate changes by a nonzero residue with each translation. Taking $K_r \ge 1$ ensures $p > k$, so (28) gives pairwise unit differences among the $w_0^{(v)}$. All operations take place in fixed dimension with rational targets of bit length polynomial in $\log p + \log(k+1)$. □

7.4 The recursive interpolation algorithm

We are now ready to combine the lifted identity, gamma factor, and interpolation points. The routine returns the whole polynomial $P_N(Z) \bmod p^k$, so one recursive call gives every shifted value needed at its parent. Choose K_r large enough for the preceding lemmas. The Jacobi family and gamma coefficients are constructed only if the direct branch fails, and are then retained.

Input. An integer $0 \le N < p$, with p, q, r, k, the finite-field representation, and the compatible character fixed as in Theorem 1.3.

Output. The polynomial $P_N(Z) \bmod p^k$, whose constant term is $N! \bmod p^k$.

Procedure.

1. If $E \ge p/C_{r,k}$ or $N \le C_{r,k}E$, compute P_N directly by Lemma 2.6 in
$$(\mathbb{Z}/p^k\mathbb{Z})[Z]/(Z^k)$$
and return it.

2. On the first call reaching this step, reconstruct $(\Theta_{r,m})_{m|q}$ using Section 3.2, and determine the gamma coefficients by Lemma 7.2. Use Lemma 7.3 to construct the k points. For each point, compute its index a_v, parts $y_i^{(v)}$, totals $Y_u^{(v)}$, arguments $z_i^{(v)}, w_u^{(v)}$, and local unit $\mathcal{U}_{a_v} \bmod p^k$.

3. Set $M = \lfloor N/(\ell-1) \rfloor$, and make one recursive call to compute $P_M(Z) \bmod p^k$. For each point, evaluate this polynomial at $z_{jh}^{(v)}$, and obtain $P_{y_{jh}^{(v)}}(z_{jh}^{(v)})$ by a short shifted interval correction. Compute the remaining $P_{y_i^{(v)}}(z_i^{(v)})$, for $h \nmid i$, and $P_{Y_u^{(v)}}(w_u^{(v)})$, for $u > 0$, directly by Lemma 2.6.

4. Rearrange (26) to obtain
$$P_{Y_0^{(v)}}(w_0^{(v)}) \equiv (-1)^{\sum_u Y_u^{(v)}} \mathcal{G}_p(\mathbf{z}^{(v)}, \mathbf{w}^{(v)}) \mathcal{U}_{a_v}^{-1} \cdot \frac{\prod_{i=0}^{d-1} P_{y_i^{(v)}}(z_i^{(v)})}{\prod_{u=1}^{h-1} P_{Y_u^{(v)}}(w_u^{(v)})} \pmod{p^k}.$$
A final shifted interval correction from $Y_0^{(v)}$ to N gives $P_N(w_0^{(v)})$.

5. Interpolate the polynomial of degree less than k from these k values, and return it. The interpolation matrix is invertible by Lemma 7.3.

Proposition 7.4. *For fixed r and every $k \ge 1$, the preceding routine computes $P_N(Z) \bmod p^k$ for $0 \le N < p$ in*
$$\widetilde{O}_r\Big(k^{O_r(1)}[q^{c_r} + D_r^{1/(2\varphi(r))}]\Big)$$
bit operations, where c_r and the polynomial exponent are chosen sufficiently large in terms of r alone.

Proof. For $k = 1$, the polynomial is constant and Theorem 5.3 applies. Suppose $k \geq 2$. We prove correctness by induction on N. The direct branch follows from Lemma 2.6. Otherwise both size conditions in the preceding lemmas hold. The selected lattice points satisfy (21), and

$$0 < \sum_{i=0}^{d-1} x_i^{(v)} p^i < p^d < \Phi_r(p) = qD_r$$

ensures $1 \leq a_v < q$. The reconstructed family gives the local value for every such a_v.

Since $\ell \geq 3$, the recursive child satisfies $M \leq N/2 < N$. Induction gives P_M. Every shifted correction factor has the form $j + pz$, with $1 \leq j < p$, and is a unit modulo p^k. Thus all divisions are valid. The lifted identity gives the values of $P_{Y_0^{(v)}}$, the interval corrections give the values of P_N, and interpolation recovers the polynomial because its degree is less than k and the arguments have pairwise unit differences.

Under either stopping condition,

$$1 + \sqrt{N} = O_r\Big((k+1)^{1/2} E^{1/2}\Big).$$

The direct branch therefore has the stated form of cost after counting arithmetic in the rank-k coefficient algebra. At a recursive stage, there are $O_r(k)$ interval products or corrections, each of length $O_r((k+1)E)$. Polynomial evaluation and interpolation use polynomially many operations in k, and the single-child recursion has logarithmic depth. The gamma factor is computed once by Lemma 7.2. Reconstruction and all local reductions have polynomial dependence on k, as established in Section 3. Collecting these factors into $k^{O_r(1)}$, and using $E^{1/2} = O_r(D_r^{1/(2\varphi(r))})$, proves the bound. The uniform bit-cost details appear in Section 7.6. □

7.5 Indices at least p

For $n \geq p$, we first count the powers of p in $n!$. This determines whether the answer is zero modulo p^k and limits the remaining work.

Let $0 \leq n < p^k$. Compute

$$v_p(n!) = \sum_{j \geq 1} \left\lfloor \frac{n}{p^j} \right\rfloor$$

by Legendre's formula. If $v_p(n!) \geq k$, return zero. Otherwise, $\lfloor n/p \rfloor < k$, so

$$n < kp. \tag{29}$$

If $p \leq K_r(k+1)$, direct multiplication therefore involves only $O_r((k+1)^2)$ factors. Its bit cost is polynomial in k and $\log p$, including the formation of the modulus.

For the remaining primes, $p > k$. Write

$$n = ap + m, \qquad 0 \leq a < k, \qquad 0 \leq m < p.$$

Then $n < kp < p^2$, so $v_p(n!) = a$. Separating the multiples of p gives the integer identity

$$n! = p^a a! \left(\prod_{b=0}^{a-1} P_{p-1}(b) \right) P_m(a). \tag{30}$$

The factor multiplying p^a is a unit and is needed only modulo p^{k-a}. Compute the two polynomials P_{p-1} and P_m at this precision and evaluate them at the indicated integers. There are at most k evaluations and k factors in the direct computation of $a!$. At precision one the polynomials are constant, and Theorem 5.3 gives their values.

7.6 Uniform dependence on the modulus exponent

The preceding construction uses higher precision in several places. We now account for all of them with r fixed and k variable. This is the last ingredient needed for Theorem 1.3.

Arithmetic and representations. An element modulo p^k has $O(k \log p)$ bits. A shifted factorial polynomial has at most k coefficients, so its dense representation has $O(k^2 \log p)$ bits. Ring operations, unit inversion, polynomial evaluation, and interpolation consequently have polynomial bit cost in k and $\log p$. Using the interval-product routine over this coefficient algebra multiplies its $\widetilde{O}(1+\sqrt{M})$ operation count by a polynomial precision factor.

The Jacobi family is reconstructed independently of k. Its local evaluation uses precision at most p^{d+k}, and Lemma 2.5 bounds this work polynomially in q, k, and $\log p$. By increasing c_r once, these costs are bounded by q^{c_r} times a fixed polynomial in k and $\log p$.

Choosing the gamma relations. The proof of Lemma 7.2 uses $s = \lfloor (k-1)/2 \rfloor$ unknown coefficients and a grid of $(k+1)^d$ points. Because $d = \varphi(r)$ is fixed, forming the evaluation matrix modulo p and selecting s independent rows has polynomial cost in k and $\log p$. The reconstructed Gauss sums and the short products are evaluated only at those selected rows. There are $O_r(k)$ such products, and their lengths are $O_r((k+1)E)$.

The digit inequalities and the conditions for an invertible evaluation system are all covered by $C_{r,k} = K_r(k+1)$. To see the dependence, the auxiliary grid uses $H = 2 + dkb$, where b depends only on r, and the interpolation translations have error at most $(1+(k-1)b)E$. Their coordinate bounds are therefore linear in $k+1$. The remaining conditions, including the fixed sublattice index bound and $p > 2k+1$, are covered by the same choice of K_r.

Recursion and finite precision. There are k interpolation arguments per stage and only one recursive child, with index at most half that of its parent. Thus the recursion has $O(\log p)$ stages, each with polynomial precision overhead. The gamma logarithm and exponential use $O(k)$ terms. On the branch where they are needed, $p > 2k+1$, so the retained denominators are units. Their omitted terms have valuation at least k, by the bounds in Section E.

If a size condition fails, direct interval products cost at most $k^{O(1)}\widetilde{O}_r(E^{1/2})$. The small-prime branch has polynomial cost in k by (29). For indices at least p, (30) uses only two recursive polynomial computations and $O(k)$ evaluations. Reading n, computing its valuation, and writing the residue take polynomial time in $k \log p$, since $n < p^k$.

Every precision-dependent factor is a polynomial in k whose degree depends only on r. Combining these bounds gives

$$\widetilde{O}_r\Big(k^{O_r(1)}[q^{c_r} + E^{1/2}]\Big) = \widetilde{O}_r\Big(k^{O_r(1)}[q^{c_r} + D_r^{1/(2\varphi(r))}]\Big).$$

The polynomial exponent and the implied constants depend only on r.

Proof of Theorem 1.3. Apply Lemma 2.2. The discarded factor changes only the constant depending on r. If the remaining conductor is one, direct interval products, with the small-prime branch and block decomposition above, satisfy the stated bound.

For a conductor greater than one, $k = 1$ is Theorem 5.3. For $k \geq 2$ and $n < p$, take the constant term returned by Proposition 7.4. For general $n < p^k$, use Section 7.5. The preceding cost analysis is uniform in k and proves the stated bound. Finally, for fixed r,

$$D_r^{1/(2\varphi(r))} = p^{1/2} q^{-1/(2\varphi(r))} \left(1 + O_r\Big(\frac{1}{p}\Big)\right),$$

which gives the equivalent form stated in the introduction. □

8 Conclusion and Open Problems

We have obtained two complementary constructions for modular factorials. The split construction uses a divisor $q \mid (p-1)$ and computes modulo p and p^2 in

$$\widetilde{O}\left(q^c + \sqrt{p/q}\right)$$

bit operations. The higher construction uses a fixed odd prime-power index r and a divisor $q \mid \Phi_r(p)$, and, for every $k \geq 1$, computes modulo p^k in

$$\widetilde{O}_r\left(k^{O_r(1)}\left[q^{c_r} + \left(\frac{\Phi_r(p)}{q}\right)^{1/(2\varphi(r))}\right]\right).$$

The split algorithm is deterministic. The higher algorithm is deterministic once the finite-field representation and character are fixed; their Las Vegas construction gives the same expected bound. The two constructions use the same polarized reconstruction of Jacobi sums, but their factorial decoders are different: nearby block boundaries in the split case, and a product of multinomial relations with one recursive chain in the higher cyclotomic constructions. At each divisor $m \mid q$, the sums use $\chi_m = \chi_q^{q/m}$, so all orders use the same choices of roots of unity and reduction maps.

Section 4.3 gives a further prime-modulus construction for $q \geq 4$ dividing $p+1$, with $(p+1)/q$ odd. The underlying interval-product routine also gives the uniform refinement modulo p and p^2 in Proposition 4.3.

Several questions remain.

Explicit reconstruction exponents. The first question is to determine useful numerical values for c and c_r. A sharper analysis may exploit the explicit carry patterns in the Stickelberger exponents. More substantially, the desired output is only a local residue, whereas the present reconstruction works with a global algebraic number. It is therefore natural to ask whether that residue can be computed without constructing the full ideal or its generator.

Distribution of useful cyclotomic divisors. For every fixed $0 < \epsilon < 1/2$, Proposition 1.7 gives the saving in (6) for at least a $1-\epsilon$ proportion of the primes up to X, for all sufficiently large X, using $O(\log(1/\epsilon))$ fixed indices. Density bounds for the wider windows in (4) would give stronger savings. It would also be useful to account for conductors formed from several smaller prime factors.

A uniform polynomial improvement. Proposition 4.3 gives a logarithmic improvement for every prime. Can one deterministically compute $n! \bmod p$ in $p^{1/2-\delta+o(1)}$ bit operations for some absolute $\delta > 0$, for every prime p and $0 \leq n < p$? Further cyclotomic indices or other families of factorial relations may lead to such an algorithm.

The complexity-theoretic status of modular factorials. When p and n are given in binary, sequential multiplication computes $n! \bmod p$ using polynomial space. The decision problem asking whether a proposed residue equals $n! \bmod p$ therefore belongs to PSPACE.[2] Does it belong to the polynomial hierarchy?

[2] In fact, it belongs to $\mathsf{P}^{\#\mathsf{P}}$.

There is also an algorithmic similarity with the discrete logarithm problem. The standard deterministic baby-step/giant-step algorithm over $\mathbb{F}_p^\times$ uses $\widetilde{O}(\sqrt{p})$ bit operations and bits of workspace in the worst case, while its answers can be verified efficiently. It is natural to ask whether the methods developed here can also be used to improve deterministic discrete-logarithm algorithms.

Acknowledgments

The author thanks Amnon Ta-Shma for valuable advice, and Amir Shpilka for a supportive working environment and for drawing attention to modular factorials.

The use of divisors of higher cyclotomic polynomials evaluated at p was proposed by the author and developed with the assistance of OpenAI's GPT-5.6. In particular, these discussions helped discover the lattice-point construction that turned the idea into an algorithm, beginning with the case $q \mid (1 + p + p^2)$. GPT-6 was subsequently used to refine the exposition and simplify some of the proofs.

References

[1] Alin Bostan, Pierrick Gaudry, and Éric Schost. Linear recurrences with polynomial coefficients and application to integer factorization and Cartier–Manin operator. *SIAM Journal on Computing*, 36(6):1777–1806, 2007. doi: 10.1137/S0097539704443793.

[2] Hendrik Willem Lenstra, Jr. and Alice Silverberg. Testing isomorphism of lattices over CM-orders. *SIAM Journal on Computing*, 48(4):1300–1334, 2019. doi: 10.1137/17M115390X.

[3] Edgar Costa, Robert Gerbicz, and David Harvey. A search for Wilson primes. *Mathematics of Computation*, 83(290):3071–3091, 2014. doi: 10.1090/S0025-5718-2014-02800-7.

[4] Paul Thomas Young. On Jacobi sums, multinomial coefficients, and p-adic hypergeometric functions. *Journal of Number Theory*, 52(1):125–144, 1995. doi: 10.1006/jnth.1995.1060.

[5] Paul van Wamelen. Jacobi sums over finite fields. *Acta Arithmetica*, 102(1):1–20, 2002. doi: 10.4064/aa102-1-1. URL `https://www.impan.pl/shop/en/publication/transaction/download/product/83146`.

[6] Oscar H. Ibarra and Chul E. Kim. Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of the ACM*, 22(4):463–468, 1975. doi: 10.1145/321906.321909.

[7] David Harvey and Markus Hittmeir. A log-log speedup for exponent one-fifth deterministic integer factorisation. *Mathematics of Computation*, 91(335):1367–1379, 2022. doi: 10.1090/mcom/3708.

[8] Enrique Treviño. The least k-th power non-residue. *Journal of Number Theory*, 149:201–224, 2015. doi: 10.1016/j.jnt.2014.10.019.

[9] Hendrik Willem Lenstra, Jr. and Carl Pomerance. A rigorous time bound for factoring integers. *Journal of the American Mathematical Society*, 5(3):483–516, 1992. doi: 10.1090/S0894-0347-1992-1137100-0.

[10] Dimitris Koukoulopoulos. Divisors of shifted primes. *International Mathematics Research Notices*, 2010(24):4585–4627, 2010. doi: 10.1093/imrn/rnq045.

[11] R. C. Vaughan. The Bombieri–Vinogradov theorem. Lecture notes, Pennsylvania State University, n.d. URL `https://personal.science.psu.edu/rcv4/Bombieri.pdf`. Section 1.

[12] Harold Davenport. *Multiplicative Number Theory*, volume 74 of *Graduate Texts in Mathematics*. Springer, New York, 3 edition, 2000. doi: 10.1007/978-1-4757-5927-3. Revised by Hugh L. Montgomery.

[13] Kenneth S. Williams. Mertens' theorem for arithmetic progressions. *Journal of Number Theory*, 6(5):353–359, 1974. doi: 10.1016/0022-314X(74)90032-8.

[14] Volker Strassen. Einige resultate über berechnungskomplexität. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 78(1):1–8, 1976. URL `https://eudml.org/doc/146659`.

[15] Richard J. Lipton. *Factoring and Factorials*, pages 153–156. Springer, Boston, MA, 2010. doi: 10.1007/978-1-4419-7155-5_33.

[16] Lenore Blum, Michael Shub, and Steve Smale. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bulletin of the American Mathematical Society*, 21(1):1–46, 1989. doi: 10.1090/S0273-0979-1989-15750-9.

[17] Michael Shub and Steve Smale. On the intractability of Hilbert's Nullstellensatz and an algebraic version of "NP $\neq$ P?". *Duke Mathematical Journal*, 81(1):47–54, 1995. doi: 10.1215/S0012-7094-95-08105-8.

[18] Hendrik Willem Lenstra, Jr. Factoring integers with elliptic curves. *Annals of Mathematics*, 126(3):649–673, 1987. doi: 10.2307/1971363.

[19] Qi Cheng. On the ultimate complexity of factorials. *Theoretical Computer Science*, 326(1–3): 419–429, 2004. doi: 10.1016/j.tcs.2004.06.020.

[20] Bruce C. Berndt, Ronald J. Evans, and Kenneth S. Williams. *Gauss and Jacobi Sums*. Canadian Mathematical Society Series of Monographs and Advanced Texts. John Wiley & Sons, New York, 1998. ISBN 978-0-471-12807-6.

[21] Ling Long and Ravi Ramakrishna. Some supercongruences occurring in truncated hypergeometric series. *Advances in Mathematics*, 290:773–808, 2016. doi: 10.1016/j.aim.2015.11.043.

[22] Benedict H. Gross and Neal Koblitz. Gauss sums and the p-adic Γ-function. *Annals of Mathematics*, 109(3):569–581, 1979. doi: 10.2307/1971226.

[23] Paul Kirchner. Algorithms on ideal over Complex Multiplication order, 2016. URL `https://arxiv.org/abs/1602.09037`.

[24] Lawrence C. Washington. *Introduction to Cyclotomic Fields*, volume 83 of *Graduate Texts in Mathematics*. Springer, New York, 2 edition, 1997. doi: 10.1007/978-1-4612-1934-7.

[25] Craig Gentry and Michael Szydlo. Cryptanalysis of the revised NTRU signature scheme. In *Advances in Cryptology—EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 299–320. Springer, 2002. doi: 10.1007/3-540-46035-7_20.

[26] Severin Wigert. Sur l'ordre de grandeur du nombre des diviseurs d'un entier. *Arkiv för Matematik, Astronomi och Fysik*, 3(18):1–9, 1907.

[27] David Harvey. Counting points on hyperelliptic curves over finite fields. Arizona Winter School lecture notes, 2026. URL `https://swc-math.github.io/aws/2026/2026HarveyNotes.pdf`.

[28] David Harvey and Joris van der Hoeven. Integer multiplication in time $O(n \log n)$. *Annals of Mathematics*, 193(2):563–617, 2021. doi: 10.4007/annals.2021.193.2.4.

[29] Arjen K. Lenstra, Hendrik Willem Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982. doi: 10.1007/BF01457454.

[30] S. A. Katre. Gauss–Jacobi sums and Stickelberger's theorem. In S. D. Adhikari, S. A. Katre, and Dinesh S. Thakur, editors, *Cyclotomic Fields and Related Topics*, pages 75–92. Bhaskaracharya Pratishthana, Pune, 2000.

A Proof of the Central Congruence

We prove (2) by reducing the Jacobi sum to a sum of powers in $\mathbb{F}_p$. The summation will leave one multinomial coefficient, which gives the desired factorial ratio. Retain $K = (p-1)/q$ and $1 \le a < q$, so $q \ge 2$. By the character convention in Section 2,

$$\rho_{\mathfrak{p}}(\chi(x)) = x^{-K} \qquad (x \in \mathbb{F}_p^{\times}).$$

Set

$$s = p - 1 - K.$$

Then $x^{-K} = x^s$ for $x \neq 0$, while both sides are interpreted as zero at $x = 0$. It follows that

$$\rho_{\mathfrak{p}}(J_a(\chi)) = \sum_{\substack{x_1,\dots,x_a \in \mathbb{F}_p \\ x_1+\cdots+x_a=1}} \prod_{i=1}^{a} x_i^s.$$

To remove the condition on the sum of the variables, use the identity

$$\mathbf{1}_{y=0} = 1 - y^{p-1} \qquad (y \in \mathbb{F}_p).$$

The right-hand side is one at zero and zero elsewhere. Therefore,

$$\rho_{\mathfrak{p}}(J_a(\chi)) = \sum_{x_1,\dots,x_a \in \mathbb{F}_p} \prod_{i=1}^{a} x_i^s \left(1 - (x_1 + \cdots + x_a - 1)^{p-1}\right).$$

The contribution from the first term vanishes, since

$$\sum_{x \in \mathbb{F}_p} x^s = 0$$

for $0 < s < p - 1$.

In the multinomial expansion of the remaining term, let k_0 denote the exponent of -1, and let k_i denote the exponent of x_i. A term survives the summation over x_i only if

$$p - 1 \mid s + k_i.$$

Since

$$0 < s + k_i < 2(p-1),$$

this forces

$$s + k_i = p - 1, \qquad k_i = K$$

for every i. Consequently,

$$k_0 = p - 1 - aK,$$

so only the term with these exponents remains. Using

$$\sum_{x \in \mathbb{F}_p} x^{p-1} = -1,$$

we obtain

$$\rho_{\mathfrak{p}}(J_a(\chi)) = -(-1)^{p-1-aK}(-1)^a \frac{(p-1)!}{(p-1-aK)!(K!)^a}.$$

Finally,

$$\frac{(p-1)!}{(p-1-aK)!} = \prod_{j=1}^{aK}(p-j) \equiv (-1)^{aK}(aK)! \pmod{p}.$$

Combining the signs gives

$$\rho_{\mathfrak{p}}(J_a(\chi)) = (-1)^{a-1}\frac{(aK)!}{(K!)^a},$$

which is equivalent to (2).

B Gauss–Jacobi Identities

We prove the identities from Proposition 2.1 using the character and sign conventions of Section 2. The first step groups the terms by their sum; the second uses the map $x \mapsto x^p$, which permutes the finite field and preserves its trace.

Proof of Proposition 2.1. Expand the product of the Gauss sums and group the summands according to

$$t = x_1 + \cdots + x_s.$$

For $t \neq 0$, the change of variables $x_i = ty_i$ shows that the inner sum is

$$(\theta_1 \cdots \theta_s)(t)J_Q(\theta_1, \ldots, \theta_s).$$

The contribution from $t = 0$ vanishes: scaling all variables by an element on which the product character is nontrivial preserves the equation $x_1 + \cdots + x_s = 0$ and multiplies the sum by a nontrivial scalar. Summing over t gives

$$G_Q(\theta_1) \cdots G_Q(\theta_s) = J_Q(\theta_1, \ldots, \theta_s)G_Q(\theta_1 \cdots \theta_s),$$

which proves the Gauss–Jacobi identity.

To prove Frobenius invariance, write

$$G_Q(\theta^p) = \sum_{x \in \mathbb{F}_Q} \theta(x)^p \psi_Q(x).$$

Substitute $x = y^{p^{f-1}}$, where $Q = p^f$. The field trace is invariant under Frobenius, and

$$\theta(y^{p^{f-1}})^p = \theta(y)^{p^f} = \theta(y).$$

Hence $G_Q(\theta^p) = G_Q(\theta)$.

The split formulas follow by taking all $\theta_i = \chi$. The magnitude identity for nontrivial Gauss sums gives

$$J_a(\chi)\overline{J_a(\chi)} = p^{a-1}.$$

For the addition formula, the three Gauss–Jacobi expressions give

$$\begin{aligned} J_u(\chi)J_v(\chi)\mathcal{J}_q(u,v) &= \frac{G_p(\chi)^u}{G_p(\chi^u)}\frac{G_p(\chi)^v}{G_p(\chi^v)}\frac{G_p(\chi^u)G_p(\chi^v)}{G_p(\chi^{u+v})} \\ &= \frac{G_p(\chi)^{u+v}}{G_p(\chi^{u+v})} = J_{u+v}(\chi). \end{aligned}$$

For the higher cyclotomic construction, recall from (12) that $\Theta_{r,m} = \mathcal{J}_{p^r,m}(1, p^h - 1)$. For $m > 1$ dividing q, the same identities give

$$\begin{aligned}\sigma_u(\Theta_{r,m}) &= J_{p^r}(\chi_m^u, \chi_m^{(p^h-1)u}) \\ &= \frac{G_{p^r}(\chi_m^u) G_{p^r}(\chi_m^{(p^h-1)u})}{G_{p^r}(\chi_m^{p^h u})} = G_{p^r}(\chi_m^{(p^h-1)u}).\end{aligned}$$

The last equality follows from Frobenius invariance. The three characters are nontrivial by Lemma 2.3, so $\Theta_{r,m}\overline{\Theta_{r,m}} = p^r$. □

C The Stickelberger Factorization

We derive the multifold factorization from the two-character formula (16). Retain the notation $\mathcal{J}_q(j, v) = J_p(\chi^j, \chi^v)$ from the preliminaries. The following product makes the exponents telescope.

Lemma C.1. *For every $1 \le a < q$,*

$$J_a(\chi) = \prod_{j=1}^{a-1} \mathcal{J}_q(j, 1),$$

where the product is empty when $a = 1$.

Proof. For $a = 1$, both sides equal 1. Suppose $2 \le a < q$. Since χ^j and χ^{j+1} are nontrivial for $1 \le j \le a - 1$, the two-character Gauss–Jacobi relation gives

$$\mathcal{J}_q(j, 1) = \frac{G_p(\chi^j) G_p(\chi)}{G_p(\chi^{j+1})}.$$

Multiplying for $j = 1, \ldots, a - 1$, the intermediate Gauss sums cancel:

$$\prod_{j=1}^{a-1} \mathcal{J}_q(j, 1) = \frac{G_p(\chi)^a}{G_p(\chi^a)} = J_a(\chi).$$

□

Proposition C.2. *For every $1 \le a < q$,*

$$(J_a(\chi)) = \prod_{t \in (\mathbb{Z}/q\mathbb{Z})^\times} \mathfrak{p}_t^{\lfloor at/q \rfloor}.$$

Proof. With the choice $\mathfrak{p}_t = \sigma_t^{-1}(\mathfrak{p})$ from Section 2.1, the two-character factorization gives

$$(\mathcal{J}_q(j, v)) = \prod_{t \in (\mathbb{Z}/q\mathbb{Z})^\times} \mathfrak{p}_t^{c_t(j,v)},$$

where

$$c_t(j, v) = \left\lfloor \frac{(j + v)t}{q} \right\rfloor - \left\lfloor \frac{jt}{q} \right\rfloor - \left\lfloor \frac{vt}{q} \right\rfloor$$

[30, Proposition 4(ii), pp. 83–84]. Katre uses the negative of our two-character Jacobi sum, which does not change the generated principal ideal.

By Lemma C.1,

$$(J_a(\chi)) = \prod_{j=1}^{a-1} (\mathcal{J}_q(j,1)).$$

Hence the exponent of $\mathfrak{p}_t$ is

$$\begin{aligned}\sum_{j=1}^{a-1} c_t(j,1) &= \sum_{j=1}^{a-1} \left(\left\lfloor \frac{(j+1)t}{q} \right\rfloor - \left\lfloor \frac{jt}{q} \right\rfloor - \left\lfloor \frac{t}{q} \right\rfloor \right) \\ &= \left\lfloor \frac{at}{q} \right\rfloor,\end{aligned}$$

The intermediate floor terms cancel, and $\lfloor t/q \rfloor = 0$ because $1 \le t < q$. □

Corollary C.3. *For an odd prime-power index, for every $m > 1$ dividing q,*

$$(\Theta_{r,m}) = \prod_{t \in (\mathbb{Z}/m\mathbb{Z})^\times / \langle p \rangle} \mathfrak{p}_{m,t}^{e_t(1,p^h-1)}.$$

Proof. Apply (15) to $\mathcal{J}_{p^r,m}(1, p^h - 1)$. Each prime above p corresponds to one orbit under multiplication by p modulo m, so its exponent is the sum of the carry counts along that orbit. □

D The Higher Cyclotomic Multinomial Identity

The factorial arguments in Proposition 5.1 come from the base-p expansion of the integer $a(p^r-1)/q$. We first find this expansion and then apply Stickelberger's congruence.

Retain $r = \ell^e$, $h = \ell^{e-1}$, and $d = (\ell - 1)h$. The numbers x_i in (21) are ordinary integer base-p digits of aD_r, each between zero and $p-1$. The numbers y_i and Y_u in (22) are the multinomial parts and their totals. Their relation to the digits of the Gauss exponent is given by the next lemma.

Lemma D.1. *The base-p digits of*

$$a\frac{p^r-1}{q} = (p^h - 1)aD_r$$

are, in increasing order of the powers of p,

$$p - 1 - y_0, \ldots, p - 1 - y_{d-1}, Y_0, \ldots, Y_{h-1}.$$

Their sum is $d(p-1)$.

Proof. Multiplication by p^h shifts the base-p expansion by h positions. Subtracting aD_r starts a borrow at the units position, since $x_0 > 0$. The inequalities $x_i \ge x_{i-h}$ for $i \ge h$ keep that borrow through the first d positions; the positive entry x_{d-h} ends it at position d. Thus

$$\begin{aligned}(p^h - 1)aD_r &= (p - x_0) + \sum_{i=1}^{h-1} (p - 1 - x_i)p^i \\ &\quad + \sum_{i=h}^{d-1} (p - 1 + x_{i-h} - x_i)p^i \\ &\quad + (x_{d-h} - 1)p^d + \sum_{u=1}^{h-1} x_{d-h+u} p^{d+u} \\ &= \sum_{i=0}^{d-1} (p - 1 - y_i)p^i + \sum_{u=0}^{h-1} Y_u p^{d+u}.\end{aligned}$$

The chain inequalities ensure that each digit is between zero and $p-1$. The identities $Y_u = \sum_{j=0}^{\ell-2} y_{u+jh}$ give the asserted digit sum. □

Proof of Proposition 5.1. The compatible character is $\chi_q = \omega_{p^r}^{-(p^r-1)/q}$. By Lemma D.1 and (14),

$$G_{p^r}(\chi_q^a) \equiv -\frac{\pi^{d(p-1)}}{\prod_{u=0}^{h-1} Y_u! \prod_{i=0}^{d-1}(p-1-y_i)!} \pmod{\pi^{(d+1)(p-1)}}.$$

Since $d = \varphi(r)$ is even and $\pi^{p-1} = -p$, the numerator is p^d. After division by $-p^d$,

$$U_a^{-1} = \prod_{u=0}^{h-1} Y_u! \prod_{i=0}^{d-1}(p-1-y_i)! \quad \text{in } \mathbb{F}_p.$$

Wilson's congruence gives

$$\prod_{i=0}^{d-1}(p-1-y_i)! = \frac{(-1)^{Y+d}}{\prod_{i=0}^{d-1} y_i!} = \frac{(-1)^Y}{\prod_{i=0}^{d-1} y_i!}.$$

Substitution gives the factorial ratio in Proposition 5.1. When a is not coprime to q, (18) shows that the lower-conductor computation uses the same digits and the same normalization. □

E The Lifted Higher Cyclotomic Identity

We first express the normalized Gauss sum through the shifted factorials used in Section 7. We then determine the remaining gamma factor from instances with small factorial arguments. These are the two statements in Proposition 7.1 and Lemma 7.2.

E.1 The shifted-factorial identity

Retain the parts y_i, totals Y_u, and shifted factorials $P_N(Z) = \prod_{j=1}^{N}(j+pZ)$ from Section 7. Let

$$\mu_i = p-1-y_i \quad (0 \le i < d), \qquad \mu_{d+u} = Y_u \quad (0 \le u < h).$$

These are the integer base-p digits of $a(p^r-1)/q$ found in Lemma D.1. Rotating this finite digit list orders the rational arguments in the Gross–Koblitz formula as follows:

$$\gamma_0 = \frac{a}{q}, \qquad p\gamma_{i+1} = \gamma_i + \mu_i \quad (0 \le i < r), \qquad \gamma_r = \gamma_0.$$

Each γ_i has denominator dividing q, and hence belongs to $\mathbb{Z}_p$. The γ_i are gamma-function arguments, whereas the μ_i are the integer digits. Recall the shifts $z_i = \gamma_{i+1} - 1$ and $w_u = -\gamma_{d+u+1}$. For $i < d$, the cyclic relation becomes

$$\gamma_i = y_i + 1 + p(\gamma_{i+1} - 1) = y_i + 1 + pz_i.$$

Applying Morita's functional equation from pz_i to $pz_i + y_i + 1$ gives

$$\Gamma_p(\gamma_i) = (-1)^{y_i+1}\Gamma_p(pz_i)P_{y_i}(z_i). \tag{31}$$

For each $0 \le u < h$, the remaining relations give

$$1 - \gamma_{d+u} = Y_u + 1 - p\gamma_{d+u+1} = Y_u + 1 + pw_u.$$

Functional equation and reflection therefore yield

$$\Gamma_p(\gamma_{d+u}) = \frac{1}{\Gamma_p(pw_u)P_{Y_u}(w_u)}. \tag{32}$$

Reflection and the functional equation each contribute the sign $(-1)^{Y_u+1}$, so their signs cancel, also when $Y_u = 0$.

The Gross–Koblitz formula and the evenness of d give

$$G_{p^r}(\chi_q^a) = -p^d \prod_{i=0}^{r-1} \Gamma_p(\gamma_i), \qquad \mathcal{U}_a = \prod_{i=0}^{r-1} \Gamma_p(\gamma_i).$$

Multiplying (31) and (32), and using $\sum_{i<d}(y_i + 1) = Y + d$, proves (26).

Finally, summing all cyclic relations gives

$$(p-1)\sum_{i=0}^{r-1} \gamma_i = \sum_{i=0}^{r-1} \mu_i = d(p-1).$$

Consequently,

$$\sum_{i=0}^{d-1} z_i - \sum_{u=0}^{h-1} w_u = \sum_{i=0}^{r-1} \gamma_i - d = 0,$$

which proves (25).

E.2 Computing the gamma factor

We prove Lemma 7.2 by finding the gamma coefficients for the fixed prime p once, then reusing them at every stage of the factorial recursion. The values needed to determine them come from small instances of (26). Throughout, $E < p/C_{r,k}$, with $C_{r,k} = K_r(k+1)$, as on the corresponding branch of the algorithm. Choosing K_r sufficiently large ensures $p > \max\{5, 2k+1\}$, since $E \ge 1$.

For $k \le 3$, the gamma factor is one modulo p^k. Suppose $k \ge 4$. By (27),

$$\log_p \mathcal{G}_p(\mathbf{z}, \mathbf{w}) = \sum_{m \ge 1} b_m S_m(\mathbf{z}, \mathbf{w}),$$

where

$$b_m = -\frac{\lambda_m p^{2m+1}}{2m(2m+1)}, \qquad S_m(\mathbf{z}, \mathbf{w}) = \sum_{i=0}^{d-1} z_i^{2m+1} - \sum_{u=0}^{h-1} w_u^{2m+1}.$$

The unknowns will be $b_m \bmod p^k$; the polynomials S_m are known from the chosen arguments. This gives a linear system directly over $\mathbb{Z}/p^k\mathbb{Z}$.

Terms needed modulo p^k. Set $s = \lfloor (k-1)/2 \rfloor$. For $2 \le m \le k$, the assumption on p makes $2m(2m+1)$ a unit, and $p\lambda_m \in \mathbb{Z}_p$ gives $v_p(b_m) \ge 2m$. For $m > k$,

$$v_p(b_m) \ge 2m - v_p(2m(2m+1)) \ge m > k,$$

using $2m(2m+1) \le 5^m$ for $m \ge 2$ and $p \ge 5$. Also, $v_p(b_1) \ge 3$. Hence

$$\log_p \mathcal{G}_p(\mathbf{z}, \mathbf{w}) \equiv \sum_{m=1}^{s} b_m S_m(\mathbf{z}, \mathbf{w}) \pmod{p^k}. \tag{33}$$

All omitted terms vanish modulo p^k, and every polynomial retained in the sum has degree at most k.

A grid of small lattice points. We need enough instances to determine these coefficients while keeping their factorial arguments small. Let Λ be the congruence lattice from Lemma 5.2, with its short vectors $v, Tv, \dots, T^{d-1}v$. Their supremum norms are at most bE for a fixed integer $b \ge 1$, and their index in Λ is bounded in terms of r. After increasing K_r, the basis is nonsingular modulo p.

Set $H = 2 + dkb$, and round the target

$$((2H+2)E,\, 2(2H+2)E, \dots, d(2H+2)E)$$

to a point $x^{(0)} \in \Lambda$. Use the grid

$$x^{(0)} + \sum_{i=0}^{d-1} t_i T^i v, \qquad t_i \in \{0, 1, \dots, k\}.$$

Every grid point differs from that target by at most HE. Its coordinates are positive and increasing, so in particular they are increasing along each of the h chains. Its last coordinate is at most $(d(2H+2)+H)E$. Since this coefficient is $O_r(k+1)$, a choice of K_r depending only on r bounds every digit, part, and total Y_u by $C_{r,k}E < p$.

The lattice congruence gives an integer $a = (\sum_i x_i p^i)/D_r$, and

$$0 < \sum_i x_i p^i < p^d < \Phi_r(p) = qD_r$$

shows that $1 \le a < q$. The lifted identity is therefore available at every grid point.

Choosing independent equations. We now show that the grid contains enough independent equations for the unknown coefficients. Modulo p, the cyclic relations give

$$(z_0, \dots, z_{d-1}) \equiv (y_1, \dots, y_{d-1}, -x_{d-h}),$$
$$(w_0, \dots, w_{h-1}) \equiv (Y_1, \dots, Y_{h-1}, -x_0).$$

Use the coordinate change given by the first entry and the successive differences in each chain:

$$\xi_u = x_u \quad (0 \le u < h), \qquad \xi_i = x_i - x_{i-h} \quad (h \le i < d).$$

Adding successive differences recovers the x_i, so this change is invertible. Substitution in S_m gives

$$S_m \equiv \sum_{u=0}^{h-1} \left\{ \sum_{j=0}^{\ell-2} \xi_{u+jh}^{2m+1} - \left(\sum_{j=0}^{\ell-2} \xi_{u+jh} \right)^{2m+1} \right\} \pmod{p}.$$

Every monomial in the m-th polynomial has total degree $2m+1$. That polynomial is nonzero: the coefficient of $\xi_0^{2m}\xi_h$ is $-(2m+1)$, which does not vanish because $\ell \ge 3$ and $p > 2k+1$. The nonzero

polynomials have distinct total degrees, so they are linearly independent. The invertible changes from the grid parameters to x, and then to ξ, preserve this independence.

The resulting polynomials have degree at most k in each grid variable. Interpolating successively in each variable shows that a polynomial with these bounds that vanishes on $\{0,1,\dots,k\}^d$ is zero. Therefore the matrix obtained by evaluating $S_1,\dots,S_s$ on the grid has rank s modulo p. Gaussian elimination selects s rows with nonzero determinant modulo p, and that determinant is also invertible modulo p^k. The full matrix has $(k+1)^d$ rows and $s = O(k)$ columns. Forming it modulo p and selecting the rows therefore has polynomial bit cost in k and $\log p$ for fixed r.

Computing the coefficients. For every selected point, compute $\mathcal{U}_a$ mod p^k using the appropriate member of the Jacobi family. Evaluate its short shifted products and rearrange (26):

$$\mathcal{G}_p(\mathbf{z},\mathbf{w}) \equiv (-1)^Y \mathcal{U}_a \frac{\displaystyle\prod_{u=0}^{h-1} P_{Y_u}(w_u)}{\displaystyle\prod_{i=0}^{d-1} P_{y_i}(z_i)} \pmod{p^k}.$$

Every denominator is a unit, and each product has length $O_r((k+1)E)$. There are $s = O(k)$ selected points, each evaluated using $O_r(1)$ interval products and one local algebraic evaluation. These algebraic evaluations are performed only for the selected rows; choosing the rows uses polynomial arithmetic over $\mathbb{F}_p$.

The gamma factor belongs to $1 + p^3\mathbb{Z}_p$, as shown in (27). Compute its p-adic logarithm by the series for $\log(1+X)$, retaining the first $k-1$ terms modulo p^k. The inverse operation uses the series for $\exp(X)$, again through degree $k-1$. The retained denominators are units because $p > 2k+1$. For omitted logarithm terms, the valuation is at least $3j - v_p(j)$; for exponential terms it is at least $3j - v_p(j!)$. Both are at least k for $j \geq k$, using $p \geq 5$. Thus these evaluations have polynomial bit cost in k and $\log p$. The selected equations (33) form a linear system over $\mathbb{Z}/p^k\mathbb{Z}$ with invertible determinant. Solving it determines $b_1,\dots,b_s$ mod p^k. Later values of the gamma factor follow by evaluating the finite sum and applying the p-adic exponential.

The $O_r(k)$ short products cost $k^{O(1)}\widetilde{O}_r(E^{1/2})$ bit operations, including arithmetic modulo p^k. Local reductions cost $k^{O_r(1)}\widetilde{O}_r(q^{c_r})$ in total. The grid and linear-system computations are polynomial in k and $\log p$ for fixed r. Combining these polynomial costs proves Lemma 7.2.

F Proof of the Central Congruence Modulo p^2

This appendix proves Proposition 6.1. The proof first specializes Young's congruence between Jacobi sums and multinomial coefficients and then expands the resulting quotient modulo p^2.

F.1 Young's congruence

Young's theorem relates generalized Jacobi sums to quotients of multinomial coefficients. We first state the part of the theorem that is needed here.

Retain the Teichmüller character ω_p from Section 2. Its value at $x \neq 0$ is the $(p-1)$-st root of unity reducing to x modulo p. Following Young's sign convention [4, Equation (2.7)], let

$$J_{\mathrm{Y}}(\psi_1,\dots,\psi_s) := -\sum_{\substack{x_1,\dots,x_s\in\mathbb{F}_p \\ x_1+\cdots+x_s=1}} \prod_{j=1}^{s} \psi_j(x_j).$$

Thus J_{Y} is the negative of the Jacobi sum convention used in this paper.

Adapting Young's parametrization [4, Theorem 2.2], let

$$\alpha_1, \ldots, \alpha_s \in [0,1) \cap \mathbb{Q}$$

have denominators not divisible by p, and suppose that

$$A_j = (p-1)\alpha_j$$

is an integer for every j. Let

$$\alpha = \alpha_1 + \cdots + \alpha_s,$$

and assume that $0 < \alpha < 1$. The resulting congruence modulo p^2 is

$$\frac{\binom{(p^2-1)\alpha}{(p^2-1)\alpha_1, \ldots, (p^2-1)\alpha_s}}{\binom{(p-1)\alpha}{(p-1)\alpha_1, \ldots, (p-1)\alpha_s}} \equiv (-1)^s J_{\mathrm{Y}}\left(\omega_p^{-A_1}, \ldots, \omega_p^{-A_s}\right) \pmod{p^2}, \tag{34}$$

provided that the Jacobi sum on the right is a unit at the chosen prime above p. We check this condition for $J_a(\chi)$ below.

We now specialize (34) to the Jacobi sum occurring in this paper.

Lemma F.1. *For every $2 \le a < q$,*

$$\frac{\binom{a(p+1)K}{\underbrace{(p+1)K, \ldots, (p+1)K}_{a \text{ times}}}}{\binom{aK}{\underbrace{K, \ldots, K}_{a \text{ times}}}} \equiv (-1)^{a-1} J_a^{(2)} \pmod{p^2}.$$

The multinomial coefficient in the denominator is invertible modulo p^2.

Proof. In (34), take

$$s = a, \qquad \alpha_1 = \cdots = \alpha_a = \frac{1}{q}.$$

Since $q \mid (p-1)$,

$$A_1 = \cdots = A_a = \frac{p-1}{q} = K.$$

Moreover,

$$\alpha = \frac{a}{q} < 1.$$

For the lower multinomial coefficient,

$$(p-1)\alpha_j = K, \qquad (p-1)\alpha = aK.$$

For the upper multinomial coefficient,

$$(p^2-1)\alpha_j = \frac{p^2-1}{q} = (p+1)K,$$

and

$$(p^2-1)\alpha = a(p+1)K.$$

Thus the quotient in (34) is the quotient in the lemma.

It remains to identify the character and verify the unit condition. For every $x \in \mathbb{F}_p^\times$, the local image of $\chi(x)$ is the unique q-th root of unity modulo p^2 reducing to x^{-K}. The Teichmüller value $\omega_p(x)^{-K}$ has the same properties, so $\chi = \omega_p^{-K}$ under the chosen local embedding. By the Stickelberger factorization,

$$v_{\mathfrak{p}}(J_a(\chi)) = \left\lfloor \frac{a}{q} \right\rfloor = 0.$$

Hence the Jacobi sum is not divisible by the distinguished prime, so (34) applies modulo p^2.

Finally, Young's Jacobi sum has the opposite sign from ours:

$$J_{\mathrm{Y}}\left(\omega_p^{-K}, \dots, \omega_p^{-K}\right) = -J_a(\chi).$$

Since $s = a$, the right-hand side of (34) becomes

$$(-1)^a(-J_a(\chi)) = (-1)^{a-1}J_a(\chi).$$

Reducing through $\rho_{\mathfrak{p}^2}$ gives the claimed congruence.

Since $aK < p$, neither $(aK)!$ nor $K!$ is divisible by p. Therefore,

$$\binom{aK}{K, \dots, K} = \frac{(aK)!}{(K!)^a}$$

is invertible modulo p^2. □

F.2 Expansion of the block products

After separating the multiples of p, we are left with complete blocks of $p-1$ factors and one shorter block. Recall $H_m = \sum_{j=1}^{m} j^{-1} \bmod p$ from Section 6. The next lemma shows how these harmonic sums enter the product modulo p^2.

Lemma F.2. *Let p be odd. For every $0 \le m < p$,*

$$\prod_{\substack{1 \le j \le (p+1)m \\ p \nmid j}} j \equiv ((p-1)!)^m \, m! \, (1 + mpH_m) \pmod{p^2}.$$

Proof. The product consists of m complete blocks followed by one final block:

$$\prod_{\substack{1 \le j \le (p+1)m \\ p \nmid j}} j = \prod_{h=0}^{m-1} \prod_{j=1}^{p-1} (hp + j) \prod_{j=1}^{m} (mp + j).$$

For a complete block,

$$\begin{aligned} \prod_{j=1}^{p-1} (hp + j) &= (p-1)! \prod_{j=1}^{p-1} \left(1 + \frac{hp}{j}\right) \\ &\equiv (p-1)! \, (1 + hpH_{p-1}) \pmod{p^2}. \end{aligned}$$

Pairing j with $p-j$ gives

$$H_{p-1} = 0 \pmod{p},$$

so every complete block is congruent to $(p-1)!$ modulo p^2.

For the final block,

$$\begin{aligned}\prod_{j=1}^{m}(mp+j) &= m!\prod_{j=1}^{m}\left(1+\frac{mp}{j}\right)\\ &\equiv m!\,(1+mpH_m) \pmod{p^2}.\end{aligned}$$

Multiplying the blocks proves the lemma. □

F.3 The factorial congruence

Proof of Proposition 6.1. The case $a=1$ follows from $J_1(\chi)=1$, so assume

$$2 \le a < q.$$

Separating the multiples of p in the two multinomial coefficients gives

$$\frac{\underbrace{\binom{a(p+1)K}{(p+1)K,\dots,(p+1)K}}_{a\text{ times}}}{\underbrace{\binom{aK}{K,\dots,K}}_{a\text{ times}}} = \frac{\displaystyle\prod_{\substack{1\le j\le a(p+1)K\\ p\nmid j}} j}{\left(\displaystyle\prod_{\substack{1\le j\le (p+1)K\\ p\nmid j}} j\right)^{a}}.$$

Applying Lemma F.2 with $m=aK$ and $m=K$ gives

$$\begin{aligned}\frac{\binom{a(p+1)K}{(p+1)K,\dots,(p+1)K}}{\binom{aK}{K,\dots,K}} &\equiv \frac{(aK)!}{(K!)^a}\,\frac{1+aKpH_{aK}}{(1+KpH_K)^a}\\ &\equiv \frac{(aK)!}{(K!)^a}\,(1+aKp(H_{aK}-H_K)) \pmod{p^2}.\end{aligned}$$

Since $K=(p-1)/q$, we have $Kp \equiv -p/q \pmod{p^2}$. Combining the preceding expansion with Lemma F.1 therefore gives

$$(-1)^{a-1}J_a^{(2)} \equiv \frac{(aK)!}{(K!)^a}\left(1-\frac{ap}{q}(H_{aK}-H_K)\right) \pmod{p^2}.$$

Since

$$\frac{ap}{q}(H_{aK}-H_K)$$

is divisible by p, its square is zero modulo p^2, and hence

$$\left(1-\frac{ap}{q}(H_{aK}-H_K)\right)^{-1} \equiv 1+\frac{ap}{q}(H_{aK}-H_K) \pmod{p^2}.$$

Multiplying by this inverse and by $(K!)^a$ proves

$$(aK)! \equiv (-1)^{a-1}J_a^{(2)}(K!)^a\left(1+\frac{ap}{q}(H_{aK}-H_K)\right) \pmod{p^2}.$$

□

G Reconstruction of the Trace-One Sums

We prove the reconstruction used in Proposition 4.2 and then show that its reduction gives the stated binomial coefficient. The reconstruction follows Section 3: determine the principal ideal and the product with the complex conjugate, then select the remaining root of unity. Here the sums run over the elements $x \in \mathbb{F}_{p^2}$ with trace $\mathrm{Tr}(x) = x + x^p = 1$.

Let $q \geq 4$ divide $p+1$, with $K = (p+1)/q$ odd, and set $M = 2q$. Choose compatible characters $\chi_m = \chi_M^{M/m}$ on $\mathbb{F}_{p^2}$, and compatible primes $\mathfrak{p}_m$, for $m \mid M$, as in Section 2.6. Write

$$\mathcal{E}_m = \sum_{\mathrm{Tr}(x)=1} \chi_m(x), \qquad \mathcal{E}_1 = p.$$

Values at proper divisors of the character order are computed first and used when selecting the root of unity at a larger order.

G.1 Principal ideals and conjugation products

If $m \mid p+1$ and $m > 1$, the restriction of χ_m to $\mathbb{F}_p^\times$ is trivial. Choose $z \neq 0$ of trace zero. The nonzero elements of trace zero are cz, with $c \in \mathbb{F}_p^\times$. For each $t \neq 0$, multiplication by t maps the elements of trace one to those of trace t, while leaving the character values unchanged. Their sum is therefore $\mathcal{E}_m$. Summing the character over the whole field gives

$$0 = (p-1)\chi_m(z) + (p-1)\mathcal{E}_m, \qquad \mathcal{E}_m = -\chi_m(z).$$

Since $z^p = -z$, compatibility also gives $\chi_m(z) = (-1)^{(p+1)/m}$. Hence

$$\mathcal{E}_m = (-1)^{(p+1)/m+1},$$

so an integer parity test determines these proper-conductor values.

Now suppose $m \nmid p+1$. Since $m \mid 2q$, the integer m is divisible by 4 and

$$p \equiv m/2 - 1 \pmod{m}.$$

The character χ_m, and each of its conjugates, restricts to the quadratic character λ of $\mathbb{F}_p^\times$, whose values are 1 on squares and -1 on nonsquares. Grouping a Gauss sum according to the trace gives

$$G_{p^2}(\chi_m) = \mathcal{E}_m G_p(\lambda). \tag{35}$$

Multiplication by a nonsquare in $\mathbb{F}_p^\times$ preserves the elements of trace zero and changes the sign of their character sum, so that sum is zero. The elements of trace $t \neq 0$ have character sum $\lambda(t)\mathcal{E}_m$. The usual absolute values of nontrivial Gauss sums then imply

$$\mathcal{E}_m \overline{\mathcal{E}_m} = p. \tag{36}$$

The same argument holds after every embedding of $\mathbb{Q}(\zeta_m)$.

For $u \in (\mathbb{Z}/m\mathbb{Z})^\times$, choose its representative in $\{1, \ldots, m-1\}$. The sum of the base-p digits of $u(p^2-1)/m$ is $(p-1)/2$ when $u < m/2$, and $3(p-1)/2$ when $u > m/2$. To see the two digits directly, write $2u(p+1)/m = 2v+1$ and $h = (p-1)/2$. Then

$$\frac{u(p^2-1)}{m} = \begin{cases} (h-v) + vp, & v \leq h, \\ (p+h-v) + (v-1)p, & v > h. \end{cases}$$

The coefficients of 1 and p are the units digit and the next base-p digit, respectively. By (14) and (35), the valuations of $\mathcal{E}_m$ at the primes over p are consequently zero or one. In the inverse-action convention of Section 2.1, this gives

$$(\mathcal{E}_m) = \prod_{\substack{u \in (\mathbb{Z}/m\mathbb{Z})^\times / \langle p \rangle \\ [u]_m > m/2}} \mathfrak{p}_{m,u}. \tag{37}$$

The condition on the representative is constant on each Frobenius orbit: for odd u, multiplication by p sends u to $m/2 - u$ modulo m, preserving the lower and upper halves separately. The order of p modulo m is at most two; order one, which can occur at $m = 4$, causes no change to the argument. There are no prime ideal factors away from p, by (36).

The ideal in (37) and the polarization p are explicitly computable. Applying Theorem 2.7 gives a generator differing from $\mathcal{E}_m$ only by an m-th root of unity. The next step chooses that root classically.

G.2 Selecting the root of unity

As in Lemma 2.8, values at all proper divisors allow us to test the possible roots of unity by polynomial Chinese remaindering. For the present sums, the additional test is a congruence for the derivative of that polynomial at one.

Lemma G.1. *Suppose $m \mid M$ and $m \nmid p + 1$, and that $\mathcal{E}_d$ is known for every proper divisor $d \mid m$. From any $\beta = \xi\mathcal{E}_m$, with ξ a root of unity in $\mathbb{Q}(\zeta_m)$, one can recover $\mathcal{E}_m$ deterministically in time polynomial in m and $\log p$.*

Proof. For each of the m roots of unity ω, use polynomial Chinese remaindering to form the unique polynomial $A_\omega(X)$ of degree less than m whose remainder modulo Φ_m represents $\omega\beta$, and whose remainder modulo each Φ_d, for $d \mid m$ with $d < m$, represents the already known value $\mathcal{E}_d$. Select the candidate satisfying

$$A_\omega(X) \in \mathbb{Z}[X], \qquad A'_\omega(1) \equiv 0 \pmod{m}. \tag{38}$$

We adapt van Wamelen's integrality argument [5, Lemmas 3–4 and Theorems 8–9] to the character values on the set of elements of trace one.

For the correct root, the coefficient of X^j counts the elements of trace one on which the character equals ζ_m^j:

$$A(X) = \sum_{j=0}^{m-1} \#\{x : \mathrm{Tr}(x) = 1,\ \chi_m(x) = \zeta_m^j\} X^j.$$

There are p elements of trace one, so $A(1) = p$. They are the roots of $X^p + X - 1$, whose product is one. The product of their character values is therefore $\zeta_m^{A'(1)} = 1$, which gives $A'(1) \equiv 0 \pmod{m}$. This explains both tests in (38).

Suppose another candidate B has integer coefficients and satisfies $B(\zeta_m) = \eta\mathcal{E}_m$, with $\eta \neq 1$. Its remainders modulo Φ_d agree with those of A for every proper divisor $d \mid m$, so

$$\frac{X^m - 1}{\Phi_m(X)} \text{ divides } B(X) - A(X) \quad \text{in } \mathbb{Z}[X].$$

Evaluate at ζ_m and take absolute norms, that is, the absolute values of the products over all conjugates. For each prime $\ell \mid m$, the valuation of the norm of the left-hand polynomial is

$\varphi(m)/(\ell-1)$. The norm of $\mathcal{E}_m$ is $p^{\varphi(m)/2}$, coprime to m. If η has order $s>1$, the absolute norm of $1-\eta$ is

$$\begin{cases} \ell^{\varphi(m)/\varphi(s)}, & s \text{ is a power of } \ell, \\ 1, & s \text{ is not a prime power.} \end{cases}$$

These norm formulas are the ones proved in the cited lemmas. If m has two distinct prime divisors, these norm divisibilities cannot hold. If m has only one prime divisor, then m is a power of two; the inequality $\varphi(m)/\varphi(s) \geq \varphi(m)$ forces $s=2$. Hence the only remaining possibility is $\eta=-1$.

In this last case, every proper divisor of m divides $m/2$, so Chinese remaindering gives

$$B(X) \equiv X^{m/2}A(X) \pmod{X^m-1}.$$

Taking derivatives at one modulo m,

$$B'(1)-A'(1) \equiv \frac{m}{2}A(1) = \frac{m}{2}p \not\equiv 0 \pmod{m}.$$

Thus the second condition in (38) excludes it. The correct candidate is unique. There are m candidates and at most m cyclotomic components, all of degree at most m, so the construction and the tests have the stated polynomial cost. □

Process all divisors of M in increasing order. For a divisor of $p+1$, use the explicit value above; for the other divisors, use (37), polarized recovery, and Lemma G.1. This constructs the whole compatible family. All its field degrees, ideal exponents, and coefficient lengths are polynomially bounded in q and $\log p$. Summing over conductors, as in Section 3.3, gives $\widetilde{O}(q^C)$ time for an absolute C.

For a requested odd $a<q$, let $m = M/\gcd(a,M)$. The corresponding trace-one sum is the conjugate $\sigma_{a/\gcd(a,M)}(\mathcal{E}_m)$. These values are integral algebraic numbers with polynomial-size representations; their reductions at the chosen prime use only polynomial-time local arithmetic.

G.3 The binomial congruence

The same base-p expansion now identifies the reduction of the reconstructed sum. For the character of order $M=2q$ and an odd $a<q$, retain $h=(p-1)/2$ and $k_a=(aK-1)/2$. The integer exponent has the base-p expansion

$$a\frac{p^2-1}{2q} = aKh = k_a p + (h-k_a).$$

Using (14) with one compatible choice of π,

$$G_{p^2}(\chi_M^a) \equiv -\frac{\pi^h}{k_a!(h-k_a)!}, \qquad G_p(\lambda) \equiv -\frac{\pi^h}{h!} \pmod{\pi^{h+p-1}}.$$

Divide both congruences by $-\pi^h$, and then take their quotient. Equation (35) gives

$$\rho_{\mathfrak{p}}(E_a) = \frac{h!}{k_a!(h-k_a)!} = \binom{h}{k_a}.$$

This lies in $\mathbb{F}_p$ even though the reduction was computed in $\mathbb{F}_{p^2}$. Finally, $\binom{h}{k} \equiv (-1)^k 4^{-k}\binom{2k}{k} \pmod{p}$ proves (19) and completes the proof of the classical reconstruction used in Proposition 4.2.

H Proof of the Uniform Logarithmic Refinement

We prove Proposition 4.3 modulo p^2 by moving most factors into a short prefix and evaluating the resulting shifted products together. The intervals still outside the prefix are evaluated in a second batch. Reduction of the final residue modulo p gives the bound for $n! \bmod p$. We first establish the two product-evaluation bounds, then analyze the routing and choose its parameters.

Throughout this appendix, p is an odd prime and $L = \lceil \log_2 p \rceil$. Integer multiplication uses $O(b \log b)$ bit operations on b-bit inputs [28]. In particular, multiplication of polynomials of degree less than $d \le p$ over $\mathbb{Z}/p^2\mathbb{Z}$ costs $O(dL^2)$ bit operations by Kronecker substitution: the coefficients are packed into integers so that integer multiplication gives their polynomial product [1, Lemma 3].

For $0 \le a \le b < p$ and $z \in \mathbb{F}_p$, write

$$P(a, b; z) = \prod_{a<j\le b} (j + pz) \pmod{p^2}.$$

The value is independent of the integer representative of z. The shifted factorial polynomial from Section 2.7 has the form

$$P_m(Z) \equiv A_m + B_m Z \pmod{p^2}, \qquad 0 \le m < p, \tag{39}$$

since the coefficient of Z^j is divisible by p^j. In particular, $A_m \equiv m! \pmod{p^2}$ and $p \mid B_m$. Thus all shifts can be evaluated from the same two coefficients at each endpoint.

Simultaneous prefix evaluation.

Lemma H.1. *Suppose $H < p$, $2\sqrt{H} + 1 < p$, and at most $H^{3/8}$ positive indices at most H are specified in increasing order, including H. The pairs (A_m, B_m) at all these indices can be computed deterministically in $O(\sqrt{H}\, L^2)$ bit operations.*

Proof. The coefficient pairs satisfy

$$\begin{pmatrix} A_m \\ B_m \end{pmatrix} = \begin{pmatrix} m & 0 \\ p & m \end{pmatrix} \begin{pmatrix} A_{m-1} \\ B_{m-1} \end{pmatrix} \pmod{p^2}, \qquad \begin{pmatrix} A_0 \\ B_0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}. \tag{40}$$

Apply [1, Theorem 15] to this degree-one matrix recurrence over $\mathbb{Z}/p^2\mathbb{Z}$, with its fixed parameter $\epsilon = 1/10$. The case $H = 1$ is immediate; for $H > 1$, the number of requested indices is less than $H^{2/5}$. The integers assumed invertible in that theorem are at most $2\sqrt{H} + 1 < p$, so they are units modulo p^2. The auxiliary product and its inverse can be formed within the same bound, using [1, Lemmas 6–7]. The matrix dimension is two, and each ring element has $O(L)$ bits. The bit cost in Theorem 15 is therefore dominated by polynomial multiplication at degree $O(\sqrt{H})$, giving $O(\sqrt{H}\, L^2)$. The pair at index zero is already known. □

Once these pairs are available, a stopped interval is evaluated by

$$P(a, b; z) \equiv \frac{A_b + B_b z}{A_a + B_a z} \pmod{p^2}. \tag{41}$$

The denominator reduces to $a! \neq 0$ modulo p. All branch shifts use the same collection of coefficient pairs.

For computation modulo p, one may instead perform all arithmetic directly in $\mathbb{F}_p$. Since $p \mid B_m$, the recurrence (40) reduces to $A_m = mA_{m-1} \pmod{p}$, with $A_0 = 1$, and (41) becomes $b!/a! \pmod{p}$. The same selected-term theorem computes these prefix factorials in $O(\sqrt{H}\, L^2)$ bit operations. All shifts can be omitted; the routing and the bounds on the interval lengths are unchanged.

Batched interval products. The second routine evaluates many short products together, allowing a different shift for each interval.

Lemma H.2. *Given s disjoint integer intervals $(a_i, b_i]$ in $\{1, \dots, p-1\}$, of total length U, and shifts $z_i \in \mathbb{F}_p$, all the products $P(a_i, b_i; z_i)$ can be computed deterministically in*

$$O\Big((s + 1 + \sqrt{U})L^6\Big)$$

bit operations.

Proof. Split each interval into at most L consecutive pieces with power-of-two lengths, and group pieces of the same length. Consider a group of v pieces, each of length M. For a power of two $b \le M$, construct $\prod_{j=1}^{b}(X + j)$ over $\mathbb{Z}/p^2\mathbb{Z}$, and evaluate it at all starting points of the vM/b complete blocks. A block with lower endpoint a and shift z uses the starting point $a + pz$. Product and remainder trees, followed by multiplication of the values belonging to each piece, cost

$$O\Big((b + vM/b + 1)L^3\Big)$$

bit operations. When $v \le M$, choose b within a factor of two of $\sqrt{vM}$; otherwise take $b = M$. This gives $O((v + 1 + \sqrt{vM})L^3)$.

There are at most L groups. The total piece count is at most sL, and the sum of their lengths is U. Cauchy–Schwarz bounds the sum of the square-root terms by $\sqrt{LU}$. Sorting the descriptions and recovering the original interval products fit within the stated bound. All polynomial remainder operations divide by monic polynomials, so they apply over $\mathbb{Z}/p^2\mathbb{Z}$ as in Lemma 2.6. Every shifted factor reduces to an integer in $\{1, \dots, p-1\}$, so each resulting product is a unit modulo p^2. □

Routing the shifted products. Let $B = 2^h$, with $h \ge 1$, and set $H = \lfloor p/B \rfloor$. The routing acts on the integer indices $1, \dots, p-1$ through the permutation

$$T(x) = \begin{cases} x/2, & x \text{ even}, \\ (x + p)/2, & x \text{ odd}. \end{cases}$$

Each branch consists of an interval $(a, b]$ and its shift z. For $0 \le n < p$, start with $P(0, n; 0)$. At each stage, retain the part whose indices lie in $[1, H]$ for the final prefix evaluation, and apply (20) to the rest. The even branch changes its shift from z to $z/2$, and the odd branch to $(z-1)/2$, in $\mathbb{F}_p$. Indeed, for $x = 2y$ or $x = 2y - p$, respectively,

$$x + pz \equiv 2(y + pz/2) \quad \text{or} \quad x + pz \equiv 2(y + p(z-1)/2) \pmod{p^2}.$$

Each replacement contributes the number of its factors to the known exponent of 2. Continue for t halving rounds or until no active branch remains, retaining any prefix portion created by the final round.

At a fixed depth the intervals of indices are disjoint, since T is a permutation. Intervals retained at different depths may overlap; keep their shifts separately. The intervals left after depth t are disjoint and lie outside the prefix. The total number of retained and remaining branches is $O(2^t)$. We obtain

$$n! \equiv 2^E \prod_i P(a_i, b_i; z_i) \prod_v P(c_v, d_v; w_v) \pmod{p^2}, \tag{42}$$

where $b_i \le H$ and $c_v \ge H$. The exponent E, endpoints, and shifts are constructed using $O(2^t L^3)$ bit operations, including sorting the stopped endpoints. Each shift occupies $O(L)$ bits, and $E \le t(p-1)$.

Length of the remaining intervals. Let U be the total length of the intervals $(c_v, d_v]$. If y lies in one of these intervals, its earlier indices in the routing are obtained by repeated doubling modulo p. None entered the prefix, so

$$[2^j y]_p > H \qquad (0 \le j \le t),$$

where the representative belongs to $\{1, \dots, p-1\}$. Consider the ordinary binary expansion of the rational number y/p. Doubling modulo one shifts this expansion by one position. Its next h bits are zero precisely when the corresponding residue is in $[1, H]$. Since p is odd, the expansion does not terminate, and the condition above excludes h consecutive zeros among its first $t + h$ bits.

We count these binary strings using independent bits, each equally likely to be zero or one. The expected waiting time for h consecutive zeros is $2^{h+1} - 2 < 2B$. For instance, this follows by solving the recurrence for the expected remaining time after j trailing zeros, with $0 \le j \le h$. Markov's inequality shows that a block of $4B$ digits avoids the run with probability less than $1/2$. Applying this to disjoint blocks, at most

$$2^{t+h}\, 2^{-\lfloor (t+h)/(4B) \rfloor}$$

binary strings of length $t + h$ avoid the run. Fixing these initial bits specifies an interval of length $2^{-(t+h)}$, which contains at most $p2^{-(t+h)} + 1$ points of the form y/p. Consequently,

$$U \le 2p \exp\!\left(-\frac{(\log 2)t}{4B}\right) + 2^t B. \tag{43}$$

The count depends only on the routed indices and therefore applies with the branch shifts retained.

Parameters and running time. Take $t = \lfloor L/4 \rfloor$, and let B be the largest power of two not exceeding

$$\frac{t}{128\lceil \log_2(t+2) \rceil}.$$

If this bound is less than two, or the size conditions of Lemma H.1 fail, use the ordinary interval algorithm modulo p^2. For sufficiently large p, these parameters give

$$B = \Theta(L/\log L), \qquad H = \Theta(p \log L/L), \qquad 2^t = O(p^{1/4}).$$

Equation (43) then gives

$$U = O\!\left(p/L^{32} + p^{1/4}L\right). \tag{44}$$

Collect the distinct positive endpoints of the stopped intervals, and add H. There are $O(p^{1/4})$ such endpoints, all at most H. For sufficiently large p, their number is less than $H^{3/8}$, and $2\sqrt{H} + 1 < p$. By Lemma H.1, all the coefficient pairs (A_m, B_m) are obtained in one computation costing $O(\sqrt{H}\, L^2)$. Equation (41) then evaluates each stopped product at its branch shift, with total additional work $p^{1/4} \operatorname{polylog}(p)$.

There are $O(p^{1/4})$ intervals remaining outside the prefix. Apply Lemma H.2 to them with their respective shifts, and multiply the factors in (42). The total bit cost is at most

$$O\!\left(\sqrt{H}\, L^2 + (p^{1/4} + 1 + \sqrt{U})L^6\right) = O\!\left(\sqrt{p}\, L^{3/2}\sqrt{\log L}\right).$$

The second term is of smaller order than the first. This proves the bound for $0 \le n < p$, including $n = p - 1$.

Indices at least p and Wilson quotients. For $p \le n < 2p$, write $n = p + s$, with $0 \le s < p$. Separating the factor p and using Wilson's theorem gives

$$(p+s)! = p(p-1)! \prod_{j=1}^{s} (p+j) \equiv -p\, s! \pmod{p^2}.$$

Compute $s!$ by the preceding routine and use its residue modulo p. For $2p \le n < p^2$, the factors p and $2p$ make $n!$ zero modulo p^2. Finally, if $F \in \{0, \ldots, p^2 - 1\}$ represents the computed value of $(p-1)! \bmod p^2$, then

$$W_p \equiv \frac{F+1}{p} \pmod{p}.$$

The division is an integer division because $p \mid F + 1$. These operations take polynomial time in $\log p$, completing the proof of Proposition 4.3.

ECCC ISSN 1433-8092
https://eccc.weizmann.ac.il