Exact Coset Sampling for Quantum Lattice Algorithms

Yifan Zhang
Princeton University
yifzhang@princeton.edu

September 30, 2025

Abstract

We give a simple and provably correct replacement for the contested "domain-extension" in Step 9 of a recent windowed-QFT lattice algorithm with complex-Gaussian windows [Chen, 2024]. As acknowledged by the author, the reported issue is due to a periodicity/support mismatch when applying domain extension to only the first coordinate in the presence of offsets. Our drop-in subroutine replaces domain extension by a pair-shift difference that cancels all unknown offsets exactly and synthesizes a uniform cyclic subgroup (a zero-offset coset) of order P inside $(\mathbb{Z}_{M_2})^n$. A subsequent QFT enforces the intended modular linear relation by plain character orthogonality. The sole structural assumption is a residue-accessibility condition enabling coherent auxiliary cleanup; no amplitude periodicity is used. The unitary is reversible, uses $\mathrm{poly}(\log M_2)$ gates, and preserves upstream asymptotics.

Project Page: https://github.com/yifanzhang-pro/quantum-lattice

1 Introduction

Fourier Sampling-based quantum algorithms for lattice problems typically engineer a structured superposition whose Fourier transform reveals modular linear relations. A recent proposal of a windowed quantum Fourier transform (QFT) with complex-Gaussian windows by Chen [2024] follows this paradigm and, after modulus splitting and CRT recombination, arrives at a joint state whose n coordinate registers (suppressing auxiliary workspace) are of the explicit affine form

$$|\phi_8.f\rangle = \sum_{j\in\mathbb{Z}} \alpha(j) \, \big| \, 2D^2 j \, b_1^* \, \big| \, 2D^2 j \, \boldsymbol{b}^*_{[2..n]} + \boldsymbol{v}^*_{[2..n]} \mod M_2 \big\rangle, \tag{1.1}$$

where $M_2 := D^2 P$ with $P = \prod_{\eta=1}^{\kappa} p_\eta$ the product of distinct odd primes, $\gcd(D, P) = 1$, $\alpha(j) = \exp\big(\frac{2\pi i}{M_2}(aj^2 + bj + c)\big)$ is a known quadratic envelope from the windowed-QFT stage,[1] $\boldsymbol{b}^* = (b_1^*, \dots, b_n^*) \in \mathbb{Z}^n$ (with $b_1^* = p_2 \cdots p_\kappa$ in the concrete pipeline of Chen [2024]), and the offset vector $\boldsymbol{v}^* \in \mathbb{Z}^n$ has unknown entries (often $v_1^* = 0$ by upstream normalization). The algorithmic goal is to sample a vector $\boldsymbol{u} \in \mathbb{Z}_{M_2}^n$ satisfying the modular linear relation

$$\langle \boldsymbol{b}^*, \boldsymbol{u} \rangle \equiv 0 \pmod{P}, \tag{1.2}$$

[1] The sum over j is effectively finite due to the upstream window; we omit a global normalization constant, which plays no role in our arguments.

from which the hidden information is recovered by standard linear algebra over the CRT factors.

The published Step 9 of Chen [2024] seeks to implement Eq. (1.2) by a "domain extension" applied only to the first coordinate, justified by a periodicity-of-amplitude heuristic. However, the domain-extension lemma invoked there presupposes global P-periodicity of the amplitude, while the presence of offsets $\boldsymbol{v}^*$ breaks this premise: extending one coordinate alone changes the support and misaligns it with the intended $\mathbb{Z}_P$-fiber. As acknowledged by the author, the resulting state does not enforce Eq. (1.2) once offsets are present.

In this work, we give a simple, reversible subroutine that substitutes Step 9 and restores correctness without appealing to amplitude periodicity. The core idea is a pair-shift difference that cancels offsets exactly and synthesizes a uniform cyclic coset of order P inside $(\mathbb{Z}_{M_2})^n$; a plain QFT then enforces Eq. (1.2) by character orthogonality. Formally, we prepare a uniform label $T \in \mathbb{Z}_P$, realize the difference register $\mathbf{Z} \equiv -2D^2 T\, \boldsymbol{b}^* \pmod{M_2}$, and (coherently) erase T. This produces an exactly uniform superposition over a cyclic subgroup of size P contained in the $\mathbb{Z}_P$-component of $(\mathbb{Z}_{M_2})^n$. Applying $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$ to $\mathbf{Z}$ yields outcomes exactly supported on Eq. (1.2) and uniform over that set; the quadratic phase $\alpha(j)$ and the offsets $\boldsymbol{v}^*$ play no role in the support.

We require only a mild residue-accessibility condition: for each prime $p_\eta \mid P$, some coordinate of $\boldsymbol{b}^*$ is nonzero modulo p_η. Equivalently, the map $T \mapsto T\,\boldsymbol{b}^* \pmod{P}$ is injective. This assumption is used solely to erase T coherently; no amplitude periodicity is assumed anywhere. The unitary is realized with classical reversible modular arithmetic (no QFT-based adders) in $\mathrm{poly}(\log M_2)$ gates and preserves the upstream phase envelope $\alpha(j)$. It is drop-in compatible with the CRT and windowed-QFT bookkeeping of Chen [2024].

Conceptually, the subroutine embeds $\mathbb{Z}_P$ into $(\mathbb{Z}_{M_2})^n$ via $T \mapsto -2D^2 T\,\boldsymbol{b}^*$ and averages uniformly over that orbit. Offsets cancel because we only manipulate basis registers and then take a difference between a shifted and an unshifted copy; the resulting uniform coset lives entirely in the $\mathbb{Z}_P$-component of $(\mathbb{Z}_{M_2})^n$ (since $M_2 = D^2 P$ and $2D^2$ is a unit modulo P). By standard Pontryagin duality for finite abelian groups, the QFT of a uniform coset has support on the annihilator, which here is precisely the hyperplane Eq. (1.2). Section 3 gives the concrete circuit and a proof of exact correctness.

Our analysis explains why one-coordinate domain extension cannot be justified under offsets: Lemma 2.17 of Chen [2024] requires global P-periodicity, which is violated post-Step 8 once $\boldsymbol{v}^* \neq \mathbf{0}$. The proposed replacement avoids any periodicity argument, works entirely at the level of subgroup cosets, and recovers the intended constraint by an elementary orthogonality calculation. By synthesizing and Fourier-sampling a uniform subgroup coset rather than extending an index, we operate at the group-structure level and sidestep support misalignment entirely in the presence of offsets.

Organization. Section 2 introduces notation and states the residue-accessibility condition. Section 3 gives the Step $9^\dagger$ circuit, the cleanup, and a proof of exact correctness. It explains how we keep phases fixed. Section 4 records gate counts, complexities, and variants. Appendix A contains explanations about the mechanics behind offset cancellation, the cyclic coset, and the orthogonality check. Appendix B proves state factorization, Appendix C lists a gate-level skeleton, and Appendix D defines the scope of determinism.

2 Preliminaries

Notation. For $q \in \mathbb{N}$, $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ with representatives in $(-\frac{q}{2}, \frac{q}{2}]$. Vectors are bold; inner product is $\langle \cdot, \cdot \rangle$. All modular arithmetic on registers is modulo $M_2 = D^2 P$ unless noted. We write $\boldsymbol{x}_{[2..n]} := (x_2, \dots, x_n)$ for coordinate slices. Throughout, for each prime $p_\eta \mid P$ we let $i(\eta)$ denote the lexicographically first index $i \in \{1, \dots, n\}$ with $\Delta_i \not\equiv 0 \pmod{p_\eta}$ (equivalently, $b_i^* \not\equiv 0 \pmod{p_\eta}$ since $2D^2$ is a unit). This choice is fixed once and for all and is implementable by a reversible priority encoder (see Step $9^\dagger$.4).

Quantum tools. We use standard primitives: $\mathrm{QFT}_{\mathbb{Z}_q}$ in $\mathrm{poly}(\log q)$ gates and reversible modular addition/multiplication. We distinguish two routines:

(i) Coordinate evaluator U_{coords}, the reversible arithmetic block that writes the coordinate registers appearing in Eq. (1.1) on basis input j:

$$U_{\mathrm{coords}} : \ |j\rangle |\mathbf{0}\rangle \longmapsto |j\rangle |\mathbf{X}(j)\rangle .$$

We call U_{coords} only on basis inputs (here $j = 0, 1$) to harvest data.

(ii) Arithmetic evaluator U_{prep}, a separate phase-free reversible circuit that never invokes U_{coords} again and that, with read-only access to harvested basis data (V, Δ), computes

$$|j\rangle |\mathbf{0}\rangle \ \mapsto \ |j\rangle \,|\, V + j \cdot \Delta \mod M_2 \,\rangle .$$

Concretely, we first call U_{coords} on $j = 0, 1$ to obtain $V := \mathbf{X}(0)$ and $W := \mathbf{X}(1)$, set $\Delta := W - V \pmod{M_2}$, and thereafter realize U_{prep} by double-and-add plus modular additions (Toffoli/Peres-style classical reversible circuits; no QFT-based adders). Because U_{prep} is a permutation of computational basis states, applying it on superpositions introduces no data-dependent phases. Reversibility/garbage is handled by standard uncomputation. In the optional constant-adder path of Step $9^\dagger$.4 one may use $(2D^2 b^*_{i(\eta)})^{-1} \bmod p_\eta$ if a classical description of $\mathbf{b}^* \bmod P$ is available; the default path uses only $\Delta \equiv 2D^2\, \mathbf{b}^*$.

Lemma 2.1 (Existence of a basis-callable coordinate evaluator)**.** Any unitary implementation that produces Eq. (1.1) necessarily contains a reversible arithmetic block that maps $|j\rangle |\mathbf{0}\rangle \mapsto |j\rangle |\mathbf{X}(j)\rangle$ (possibly with workspace later uncomputed). We denote such a block by U_{coords} and call it only on basis inputs.

Assumption 2.2 (Basis-callable coordinate evaluator; run-local determinism)**.** Within a single circuit execution, the coordinate evaluator U_{coords} uses fixed classical constants so that the basis outputs $\mathbf{X}(0)$ and $\mathbf{X}(1)$ are reproducible. We harvest (V, Δ) inside the same run prior to any superposition-time step: $V := \mathbf{X}(0)$ and $\Delta := \mathbf{X}(1) - \mathbf{X}(0)$. The arithmetic evaluator U_{prep} used during superpositions performs only classical reversible (Toffoli/Peres) arithmetic and never calls U_{coords} on a superposed input. Harvested registers (V, Δ) are treated as read-only basis data.

Security/indistinguishability note. If an external oracle were to return $\mathbf{X}(j+T)$ from $\mathbf{X}(j)$ for arbitrary T with the same offset, then—as in LWE with reused noise—subtracting two outputs would reveal the offset-free difference and compromise indistinguishability. Our construction never assumes such an oracle. All calls to U_{coords} are intra-run basis calls that reuse the very arithmetic that prepared Eq. (1.1); across runs, upstream randomness need not preserve the same offset.

Implementation note. (i) Harvest (V, Δ) within the same run before any superposition-time step, and keep them as read-only basis data. The coordinate evaluator U_{coords} is never applied to a superposed input. (ii) The evaluator U_{prep} is implemented with classical reversible (Toffoli/Peres) adders/multipliers only; we do not use QFT-based adders, ensuring no data-dependent phase is introduced on superpositions.

Lemma 2.3 (Phase discipline)**.** If all superposition-time arithmetic in Steps $9^{\dagger}.1$–$9^{\dagger}.4$ is realized by classical reversible circuits (no QFT-based adders) and U_{coords} is never applied on a superposed input, then no additional data-dependent phase is imprinted beyond the fixed quadratic envelope $\alpha(j)$ produced upstream.

Proof. Classical reversible adders/multipliers implement permutations of the computational basis; thus they preserve amplitudes and phases. Avoiding U_{coords} on superpositions prevents reintroduction of state-preparation phases. □

Remark. QFT-based adders would, in general, introduce data-dependent phases through controlled rotations; these are precisely the kind of envelope phases one must avoid in the windowed-QFT regime that produced $\alpha(j)$ upstream. In our construction, U_{coords} is never applied to a superposed input.

Within a single run, one could measure the harvested basis registers $V = \mathbf{X}(0)$ and $\Delta = \mathbf{X}(1) - \mathbf{X}(0)$ and hence recover $\mathbf{v}^*$ and $2D^2\mathbf{b}^*$ classically. Our default path simply does not require such measurement; we retain (V, Δ) as basis data to maintain phase discipline. If an implementation is happy to expose $\mathbf{b}^*$ classically, the constant-adder variant (Remark 3.4) applies verbatim and further simplifies cleanup. No indistinguishability claim is made or needed here.

Arithmetic evaluator and finite difference Δ. Let U_{prep} be the reversible arithmetic evaluator of $\mathbf{X}(\cdot)$ as above, and define

$$\Delta \;:=\; \mathbf{X}(1) - \mathbf{X}(0) \pmod{M_2},$$

harvested once via basis calls $j = 0, 1$. Because $\mathbf{X}(j)$ depends only on $j \bmod P$, this same Δ equals $\mathbf{X}(J{+}1) - \mathbf{X}(J)$ for any classical J, but we do not recompute it; Δ is treated as read-only basis data. In all cases, $\Delta \equiv 2D^2\,\mathbf{b}^* \pmod{M_2}$. We will use Δ to compute T from $\mathbf{Z}$ without any classical knowledge of $\mathbf{b}^*$.

Where $\mathbf{X}(j)$ comes from in Chen [2024]. In Chen's nine-step pipeline, after modulus splitting P and CRT recombination, the state denoted $|\varphi_7\rangle$ (and the discussion immediately before Step 8 there) contains a coordinate block of the explicit affine form

$$\big(2D^2 j\, b_1^* \;\big|\; 2D^2 j\, \boldsymbol{b}^*_{[2..n]} + \boldsymbol{v}^*_{[2..n]}\big) \pmod{M_2},$$

up to an orthogonal $\frac{M_2}{2}$–coset index k and a global quadratic phase in j (the "Karst-wave" envelope). If we retain just this coordinate block (suppressing k), rename the surviving (effectively finite) loop variable as j, and ignore global phases, we obtain exactly Eq. (1.1). In the notation used throughout our paper,

$$\mathbf{X}(j) \;:=\; V + j\Delta \;\equiv\; 2D^2 j\, \boldsymbol{b}^* + \boldsymbol{v}^* \pmod{M_2},$$

with $V = \boldsymbol{v}^*$ and $\Delta = 2D^2\boldsymbol{b}^*$ harvested once via basis calls $j = 0, 1$ to the preparation/evaluator block U_{coords} (Prop. 2.4). The optional label $J \equiv j \pmod{P}$ that we carry in Section 3 is precisely the CRT-reduced index present after Chen's Step 8. No periodicity-of-amplitude assumption is used here, only the affine computational-basis content of the coordinate registers.

Explicit construction of U_{prep} without classical $\boldsymbol{b}^*, \boldsymbol{v}^*$. We now give a stand-alone construction of the reversible arithmetic evaluator $U_{\text{prep}} : |j\rangle |\mathbf{0}\rangle \mapsto |j\rangle |\mathbf{X}(j)\rangle$ that does not require any classical knowledge of $\boldsymbol{b}^*$ or $\boldsymbol{v}^*$.

Proposition 2.4 (Harvest-on-basis & arithmetic re-evaluation)**.** Let U_{coords} be the coordinate evaluator from Lemma 2.1. Invoke it once each on the basis inputs $j = 0$ and $j = 1$ (with all ancillas restored to $|0\rangle$) to obtain two program registers in the computational basis:

$$V := \mathbf{X}(0) = \boldsymbol{v}^* \pmod{M_2}, \qquad \Delta := \mathbf{X}(1) - \mathbf{X}(0) \equiv 2D^2\, \boldsymbol{b}^* \pmod{M_2}.$$

This harvest occurs within the same run, before any superposition-time step, and uses no mid-circuit measurement. Now define a separate reversible arithmetic evaluator U_{prep} that acts on $|j\rangle |\mathbf{0}\rangle$ (with read-only access to V, Δ) by computing

$$|j\rangle |\mathbf{0}\rangle \;\mapsto\; |j\rangle\, |\, V + j \cdot \Delta \;\bmod M_2\, \rangle\, .$$

This evaluator performs no phase kickback (Toffoli/Peres-style arithmetic; no QFT adders) and never invokes U_{coords} again; hence any quadratic phases created during the windowed–QFT stage remain unaffected. The multiplication $j \cdot \Delta$ is implemented reversibly by a standard double-and-add routine that treats Δ as data (not as a hard-coded constant) without mutating it: if $j = \sum_\ell j_\ell 2^\ell$ in binary, perform for each bit ℓ the controlled update "if j_ℓ=1 then add R_ℓ", where $R_0 := \Delta$ and $R_\ell := 2R_{\ell-1} \pmod{M_2}$ is maintained in a scratch register; Δ itself remains unchanged and the R_ℓ ladder is uncomputed at the end. Finally add $V \pmod{M_2}$.

Lemma 2.5 (Efficiency and independence from classical secrets)**.** Construction 2.4 realizes a unitary U_{prep} with gate complexity $O(n \log P \cdot \text{poly}(\log M_2))$. It uses only reversible modular additions/doublings and treats (V, Δ) as basis registers obtained from U_{coords}; no classical description of $\boldsymbol{b}^*$ or $\boldsymbol{v}^*$ is required. The reversible double-and-add uses one scratch register R to hold R_ℓ and uncomputes it at the end; Δ is never modified. Computing per-prime modular inverses during cleanup via a reversible extended Euclidean algorithm costs $O((\log p_\eta)^2)$ gates per p_η (or $\widetilde{O}(\log p_\eta)$ with half-GCD). Re-evaluating $\mathbf{X}(\cdot)$ at $J{+}T$ therefore consists of invoking the arithmetic evaluator on the input label $J{+}T$, without imprinting any additional phases.

Proof. The schoolbook double-and-add uses $O(\log P)$ additions per coordinate, each in poly$(\log M_2)$ gates; n coordinates contribute the stated factor. All operations are on computational-basis registers (V, Δ) and do not assume knowledge of their numeric values. As U_{coords} is the known reversible subroutine already used to produce Eq. (1.1), preparing (V, Δ) once is efficient; after preparation, U_{prep} can be called repeatedly at different inputs (e.g., $J{+}T$ in Step $9^\dagger$.2). *Note.* Multiplication by the data vector Δ via double-and-add performs $O(\log P)$ controlled additions per coordinate, never mutates Δ, and uncomputes the scratch ladder R_ℓ exactly. □

Remark 2.6. If a classical description of $\boldsymbol{b}^* \bmod P$ happens to be available, one may replace the data-multiplication by a constant adder using $2D^2T\, \boldsymbol{b}^*$ as in Remark 3.4; this is optional and not used in our default path.

Lemma 2.7 (Affine register form)**.** For all j in the implicit finite window (from the windowed-QFT stage), the coordinate registers immediately before Step 9 have the exact affine form

$$\mathbf{X}(j) \equiv 2D^2 j\, \mathbf{b}^* + \mathbf{v}^* \pmod{M_2},$$

and the window affects only the amplitudes $\alpha(j)$, not the computational-basis contents. In particular, $\mathbf{X}(j{+}1) - \mathbf{X}(j) \equiv \Delta \pmod{M_2}$ for all j, hence $\mathbf{X}(j) \equiv V + j\Delta \pmod{M_2}$.

Default J-free realization. If one prefers to avoid carrying J, the construction can be simplified as follows: after harvesting Δ as basis data, skip the re-evaluation of $\mathbf{X}(j+T)$ and directly allocate $\mathbf{Z}$ and set

$$\mathbf{Z} \leftarrow -T \cdot \Delta \pmod{M_2}$$

by a double-and-add with read-only access to Δ. The subsequent cleanup (computing T' from $\mathbf{Z}$ and uncomputing it) proceeds unchanged. This variant removes the need for $\mathbf{Y}$ and J entirely.

Injectivity condition. We will use the following natural assumption, which enables coherent coset synthesis by allowing us to uncompute the shift parameter T from the difference register. Without it, T cannot be erased from the rest of the state, and Fourier sampling on $\mathbf{Z}$ alone becomes uniform over $\mathbb{Z}_{M_2}^n$ (i.e., it does not enforce Eq. (1.2) with constant success probability).

Definition 2.8 (Residue accessibility)**.** For each prime $p_\eta \mid P$, there exists a coordinate $i(\eta) \in \{1, \dots, n\}$ such that the entry $b^*_{i(\eta)}$ is not a multiple of p_η, i.e., $b^*_{i(\eta)} \not\equiv 0 \pmod{p_\eta}$.

This condition holds with overwhelming probability for the lattice instances considered in [Chen, 2024]; any given instance can be checked efficiently, and coordinates can be permuted if necessary. Importantly, this assumption is needed only for the cleanup that erases T coherently. If the cleanup is skipped, then regardless of whether Definition 2.8 holds, applying QFT to $\mathbf{Z}$ alone yields the uniform distribution on $\mathbb{Z}_{M_2}^n$ (the T-branches remain orthogonal and do not interfere). When Definition 2.8 holds, T is a function of $\mathbf{Z} \bmod P$, enabling coherent erasure and the interference that enforces Eq. (1.2). It implies that the map $T \mapsto T\,\boldsymbol{b}^* \pmod{P}$ from $\mathbb{Z}_P$ to $(\mathbb{Z}_P)^n$ is injective. To see this, if $T\,\boldsymbol{b}^* \equiv \mathbf{0} \pmod{P}$, then for each η, the condition $b^*_{i(\eta)} \not\equiv 0 \pmod{p_\eta}$ (equivalently, $\Delta_{i(\eta)} \not\equiv 0 \pmod{p_\eta}$ since $\Delta \equiv 2D^2\mathbf{b}^*$ and $2D^2$ is a unit mod p_η) forces $T \equiv 0 \pmod{p_\eta}$. By the Chinese Remainder Theorem, this implies $T \equiv 0 \pmod{P}$. Conversely, if Definition 2.8 fails for some p_η, then $b^*_i \equiv 0 \pmod{p_\eta}$ for all i, so every T multiple of p_η lies in the kernel of $T \mapsto T\,\boldsymbol{b}^* \bmod P$; hence injectivity fails. Thus, Definition 2.8 is equivalent to the injectivity of this map and to the recoverability of T from $\mathbf{Z} \bmod P$.

Remark 2.9 (Random-instance bound)**.** Because $b^*_1 = p_2 \cdots p_\kappa$, we have $b^*_1 \not\equiv 0 \pmod{p_1}$ and $b^*_1 \equiv 0 \pmod{p_\eta}$ for all $\eta \geq 2$. If, for each prime p_η, the remaining coordinates $(b^*_2, \dots, b^*_n) \bmod p_\eta$ are close to uniform over $(\mathbb{Z}_{p_\eta})^{n-1}$ (as in typical reductions), then for $\eta = 1$ the accessibility condition holds deterministically, while for each $\eta \geq 2$ we have

$$\Pr[\, b^*_i \equiv 0 \text{ for all } i \bmod p_\eta \,] \;=\; \Pr[\, b^*_2 \equiv \cdots \equiv b^*_n \equiv 0 \bmod p_\eta \,] \;=\; p_\eta^{-(n-1)}.$$

A union bound therefore yields

$$\Pr[\text{residue accessibility fails for some } p_\eta] \;\leq\; \sum_{\eta=2}^{\kappa} p_\eta^{-(n-1)},$$

which is negligible once $n \geq 2$ and the p_η are moderately large (for $n = 2$, the sum still decays with the prime sizes).

Proposition 2.10 (Cleanup necessity and consequence)**.** Let $|\Phi_3\rangle$ be the joint state immediately after forming $\mathbf{Z}$ (Eq. (3.1)) but before auxiliary cleanup. If T remains entangled with $\mathbf{Z}$, then Fourier sampling on $\mathbf{Z}$ alone is uniform over $(\mathbb{Z}_{M_2})^n$, irrespective of $\boldsymbol{v}^*$ and the phases $\alpha(j)$. Under Definition 2.8, T is a function of $\mathbf{Z} \bmod P$ and can be erased coherently; the resulting pure state factors as in Eq. (3.2), enabling interference that enforces Eq. (1.2).

Proof. Tracing out $(\mathbf{X}, \mathbf{Y}, T)$ before cleanup leaves the classical mixture $\rho_{\mathbf{Z}} = \frac{1}{P}\sum_{t\in\mathbb{Z}_P} |-2D^2 t\boldsymbol{b}^*\rangle\langle -2D^2 t\boldsymbol{b}^*|$. Since $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}|z\rangle$ has a uniform measurement distribution for every basis state $|z\rangle$, any convex mixture of basis states yields a uniform measurement on $(\mathbb{Z}_{M_2})^n$. Thus, before cleanup, Fourier sampling enforces no constraint. When Definition 2.8 holds, t is a (CRT-)function of $\mathbf{Z}$ mod P; we reversibly compute t from $\mathbf{Z}$, zero the original T, restore $\mathbf{Y}$ to $\mathbf{X}(j)$ using the evaluator U_{prep} (the b^*-free path), uncopy, and uncompute the auxiliary arithmetic. The post-cleanup state factors as in Eq. (3.2), so subsequent Fourier sampling interferes across the T-branches and enforces Eq. (1.2). Full details are in Appendix B. □

3 The new Step $9^\dagger$: pair-shift difference and exact coset synthesis

3.1 Idea in one line

Make a second copy of the coordinate registers, coherently shift it by a uniform $T \in \mathbb{Z}_P$ along the $\boldsymbol{b}^*$ direction, and subtract. The subtraction cancels the unknown offsets $\boldsymbol{v}^*$ and leaves a clean difference register $-2D^2T\,\boldsymbol{b}^* \pmod{M_2}$. Because T is uniform, this is an exactly uniform superposition over a cyclic subgroup of order P (the $\mathbb{Z}_P$-fiber in the CRT decomposition $\mathbb{Z}_{M_2} \cong \mathbb{Z}_{D^2} \times \mathbb{Z}_P$) indexed by T. A QFT on this coset yields Eq. (1.2) exactly, by plain character orthogonality. The pseudo code is shown in Algorithm 1.

3.2 Method

We present two realizations of Step $9^\dagger$ (we adopt the J-free route as the canonical default; the re-evaluation route is optional). Neither realization assumes an oracle that, from $\mathbf{X}(j)$ alone, produces $\mathbf{X}(j{+}T)$:

Default J-free route: Steps $9^\dagger.2'$ and $9^\dagger.4$ only; no $\mathbf{Y}$ register is ever allocated and Step $9^\dagger.1$ is not used. This route forms $\mathbf{Z}$ directly from Δ; no evaluation of $\mathbf{X}(j{+}T)$ occurs and no "$As{+}e \mapsto As'{+}e$" oracle is assumed.

Re-evaluation route: Steps $9^\dagger.1$–$9^\dagger.4$; this route allocates $\mathbf{Y}$ and uses a label $J \equiv j \pmod P$ (carried from preparation).

We begin with the input state $|\phi_8.f\rangle$ from Eq. (1.1). We prepare a register for $T \in \mathbb{Z}_P$ in the uniform superposition $\frac{1}{\sqrt{P}}\sum_{t\in\mathbb{Z}_P}|t\rangle$, e.g., preferably by independent $\mathrm{QFT}_{\mathbb{Z}_{p_\eta}}$ with CRT wiring (an exact realization); a monolithic $\mathrm{QFT}_{\mathbb{Z}_P}$ is also possible. (Only in the re-evaluation route do we also append $\mathbf{Y} \in (\mathbb{Z}_{M_2})^n$; the default J-free route does not allocate $\mathbf{Y}$.)

Step $9^\dagger$.1 (copy). Use CNOT or modular addition gates to coherently copy the coordinate registers into $\mathbf{Y}$. This basis-state copying does not violate the no-cloning theorem.

$$\sum_j \alpha(j)\,|\mathbf{X}(j)\rangle\,|\mathbf{0}\rangle \;\longmapsto\; \sum_j \alpha(j)\,|\mathbf{X}(j)\rangle\,|\mathbf{X}(j)\rangle\,,$$

where for brevity we write $\mathbf{X}(j) := \left(2D^2 j\, b_1^* \;\middle|\; 2D^2 j\, \boldsymbol{b}^*_{[2..n]} + \boldsymbol{v}^*_{[2..n]}\right)$ modulo M_2.

Remark 3.1 (Copying basis states does not violate no-cloning)**.** Let U_{add} act coordinatewise by $U_{\mathrm{add}}|x\rangle|y\rangle = |x\rangle|x+y\rangle \pmod{M_2}$. This is a permutation of the computational basis and hence unitary. In particular, $U_{\mathrm{add}}|x\rangle|0\rangle = |x\rangle|x\rangle$, so computational-basis states are copied exactly. For a superposition $|\psi\rangle = \sum_j \alpha(j)\,|\mathbf{X}(j)\rangle$, linearity gives

$$U_{\text{add}}\Big(\sum_j \alpha(j)\,|\mathbf{X}(j)\rangle\,|\mathbf{0}\rangle\Big) = \sum_j \alpha(j)\,|\mathbf{X}(j)\rangle\,|\mathbf{X}(j)\rangle\,,$$

which is entangled and *not* $|\psi\rangle \otimes |\psi\rangle$ unless $|\psi\rangle$ is a single basis vector. Thus Step $9^{\dagger}.1$ does not implement a universal cloner; it coherently copies classical (commuting) information, in agreement with the no-cloning [Wootters and Zurek, 1982, Dieks, 1982] and no-broadcasting theorems [Barnum et al., 1996]; see also [Nielsen and Chuang, 2010].

Step $9^{\dagger}$.2. The re-evaluation (copy-shift-difference) variant is presented in Subsection 3.3. The default path is the J-free shift in Step $9^{\dagger}.2'$ below.
Step $9^{\dagger}$.2' (J-free shift). *Alternative to (and simpler than) Step $9^{\dagger}$.2.* Skip $\mathbf{Y}$ and J altogether and directly set

$$\mathbf{Z} \leftarrow -T \cdot \Delta \ (\text{mod}\, M_2),$$

using the double-and-add with Δ as read-only data. Equivalently,

$$\mathbf{Z} \ = \ -2D^2 T\,\boldsymbol{b}^* \ \ (\text{mod } M_2). \tag{3.1}$$

Proceed to Step $9^{\dagger}$.4 for cleanup. This variant removes the need for $\mathbf{Y}$ and J entirely.

Step $9^{\dagger}$.4 (mandatory auxiliary cleanup). The residue accessibility assumption (Definition 2.8) ensures that T can be computed as a function of $\mathbf{Z}$ mod P (uniquely by CRT). Default (b^*-free) path: recall the harvested finite difference $\Delta := \mathbf{X}(1) - \mathbf{X}(0) \equiv 2D^2\,\mathbf{b}^*\ (\text{mod}\, M_2)$, obtained once from literal basis inputs $j = 0, 1$. Do not invoke U_{coords} again. For each prime $p_\eta \mid P$, reduce $(\Delta, \mathbf{Z})$ modulo p_η and fix once and for all the lexicographically smallest index $i(\eta) \in \{1, \ldots, n\}$ with $\Delta_{i(\eta)} \not\equiv 0 \pmod{p_\eta}$ (equivalently, $b^*_{i(\eta)} \not\equiv 0 \pmod{p_\eta}$ since $2D^2$ is a unit). We implement this choice by a reversible priority encoder over the predicates $[\Delta_i \not\equiv 0 \ (\text{mod} p_\eta)]$, write $i(\eta)$ into an ancilla, and uncompute all scan flags afterward; thus the selection is deterministic, reversible, and measurement-free. Then compute into a fresh auxiliary register T', the residues

$$T' \equiv -\Delta_{i(\eta)}^{-1}\, Z_{i(\eta)} \pmod{p_\eta},$$

using a modular inversion subroutine controlled on the predicate $[\Delta_{i(\eta)} \not\equiv 0]$; this avoids undefined inversions. The inverses $\Delta_{i(\eta)}^{-1}$ mod p_η are computed on the fly (e.g., reversible extended Euclidean algorithm) and require no classical knowledge of $\mathbf{b}^*$. Finally recombine the residues via a reversible CRT—either a naive Garner mixed-radix scheme (quadratic in κ) or a remainder/product-tree CRT (near-linear $O(\kappa \log \kappa)$)—with precomputed constants depending only on P. Keep the intermediate digits so they can be uncomputed in reverse; this recovers $T' \in \mathbb{Z}_P$. Here is the detailed cleanup steps in the J-free (default) branch:

(i) Compute T' from $(\mathbf{Z}, \Delta)$ via per-prime inversions and reversible CRT.

(ii) Set $T \leftarrow T - T'$ so that $T = 0$.

(iii) Erase T' by applying the inverse of its computation from $\mathbf{Z}$.

These steps leave $\mathbf{Z}$ unchanged and require no classical access to $\boldsymbol{b}^*$. The cleanup for the re-evaluation variant is given below in Subsection 3.3.

Reversibility note: CRT recombination can be implemented (i) by a reversible Garner mixed-radix scheme in $O(\kappa^2)$ modular operations, or (ii) by a reversible remainder/product-tree CRT in $O(\kappa \log \kappa)$ modular operations; both use constants depending only on (p_η) and are reversible when the ancilla trail is retained, so the subsequent uncomputation is exact. After these actions, the global state factorizes with a coherent superposition on $\mathbf{Z}$.[2]

Lemma 3.2 (Recovering T from $\mathbf{Z}$)**.** Under Definition 2.8 and Eq. (3.1), let $\Delta := \mathbf{X}(J+1) - \mathbf{X}(J) \equiv 2D^2\, \mathbf{b}^* \pmod{M_2}$. For each p_η, after reducing modulo p_η, fix the lexicographically smallest $i(\eta)$ with $\Delta_{i(\eta)} \not\equiv 0 \pmod{p_\eta}$ and let $c_\eta := \Delta_{i(\eta)}^{-1} \bmod p_\eta$. Then $T \equiv -c_\eta\, Z_{i(\eta)} \pmod{p_\eta}$ for all η, and the unique $T \in \mathbb{Z}_P$ is obtained by CRT recombination.

Proof. Immediate from $Z_{i(\eta)} \equiv -2D^2 T\, b^*_{i(\eta)} \pmod{p_\eta}$ and $\Delta_{i(\eta)} \equiv 2D^2 b^*_{i(\eta)} \pmod{p_\eta}$, which give $Z_{i(\eta)} \equiv -T\, \Delta_{i(\eta)} \pmod{p_\eta}$ and hence $T \equiv -\Delta_{i(\eta)}^{-1} Z_{i(\eta)} \pmod{p_\eta}$. ☐

After Step $9^\dagger$.4 we have the *factorized* state

$$\left(\sum_j \alpha(j)\ |\mathrm{junk}(j)\rangle \right) \otimes \frac{1}{\sqrt{P}} \sum_{T \in \mathbb{Z}_P} \Big| -2D^2 T\, \boldsymbol{b}^* \mod M_2 \Big\rangle_{\mathbf{Z}}, \tag{3.2}$$

where "junk(j)" denotes registers independent of $\mathbf{Z}$ that we will never touch again.

Why one-coordinate domain extension fails. Consider the map $j \mapsto \mathbf{X}(j)$ in Eq. (1.1) with offsets. Any one-coordinate domain-extension rule that prolongs only the first coordinate while holding the others modulo P is valid only when the entire state amplitude is P-periodic in the extended index. Offsets break this premise: the last $n-1$ coordinates shift by $2D^2 j\, \boldsymbol{b}^*_{[2..n]} + \boldsymbol{v}^*_{[2..n]}$, whose P-periodicity depends on the unknown $\boldsymbol{v}^*$ and cannot be assumed. As in the paper's own DCP caution, replacing j by a longer register while keeping (j mod P) in the other coordinates changes the instance (cf. $|j\rangle\, |(j \bmod 2)x - y\rangle \neq |j\rangle\, |jx - y\rangle$).

Fourier sampling. Apply $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$ to the entire $\mathbf{Z}$-register block and measure $\boldsymbol{u} \in \mathbb{Z}_{M_2}^n$. The outcome distribution is analyzed next.

Algorithm 1 Step $9^\dagger$ — *Default J-free route* (no copy step)

Require: Registers $\mathbf{X} \in (\mathbb{Z}_{M_2})^n$ as in Eq. (1.1); harvested $\Delta = \mathbf{X}(1) - \mathbf{X}(0)$.
1: Prepare $T \in \mathbb{Z}_P$ in $\frac{1}{\sqrt{P}} \sum_{t \in \mathbb{Z}_P} |t\rangle$.
2: **Set** $\mathbf{Z} \leftarrow -T \cdot \Delta \pmod{M_2}$ (double-and-add; read-only Δ)
3: **Auxiliary cleanup:** compute $T' \leftarrow f(\mathbf{Z}, \Delta)$ by per-prime inversions and reversible CRT; set $T \leftarrow T - T'$ (so $T = 0$); uncompute T' from $\mathbf{Z}$ by inverting its construction.
4: Apply $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$ to $\mathbf{Z}$; measure $\boldsymbol{u} \in \mathbb{Z}_{M_2}^n$.
5: Output $\boldsymbol{u}$; by Theorem 3.9 (given Definition 2.8), it satisfies $\langle \boldsymbol{b}^*, \boldsymbol{u} \rangle \equiv 0 \pmod{P}$.

The re-evaluation route (which uses Step $9^\dagger$.1) is given next and in Subsection 3.3.

[2] This cleanup is necessary for correctness; see Prop. 2.10.

Algorithm 2 Step $9^\dagger$ — *Re-evaluation route* (uses Step $9^\dagger$.1 copy)

Require: Registers $\mathbf{X} \in (\mathbb{Z}_{M_2})^n$ (Eq. (1.1)); label $J \equiv j \pmod P$; harvested (V, Δ) for U_{prep}.

1: Prepare $T \in \mathbb{Z}_P$ in $\frac{1}{\sqrt{P}} \sum_{t \in \mathbb{Z}_P} |t\rangle$.

2: **($9^\dagger$.1 Copy)** Copy $\mathbf{X}$ to $\mathbf{Y}$ via modular adds.

3: **($9^\dagger$.2 Shift)** Evaluate U_{prep} at $J+T$ into $\mathbf{Y}$ to get $\mathbf{X}(j{+}T)$.

4: **($9^\dagger$.3 Difference)** Set $\mathbf{Z} \leftarrow \mathbf{X} - \mathbf{Y} \pmod{M_2}$.

5: **($9^\dagger$.4 Cleanup)** Compute $T' \leftarrow f(\mathbf{Z}, \Delta)$; update $\mathbf{Y} \leftarrow \mathbf{Y} + \big(\mathbf{X}(J{+}T{-}T') - \mathbf{X}(J{+}T)\big)$; set $T \leftarrow T - T'$; uncopy $\mathbf{Y}$; uncompute T' from $\mathbf{Z}$.

6: Apply $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$ to $\mathbf{Z}$; measure $\boldsymbol{u}$.

7: **return** $\boldsymbol{u}$.

3.3 Re-evaluation variant for Step $9^\dagger$

Optional index label (retained from the windowed–QFT stage). For one realization of our pair–shift difference and cleanup without any classical knowledge of the full vector $\boldsymbol{b}^*$, it can be convenient to retain a small label register $J \in \mathbb{Z}_P$ with $J \equiv j \bmod P$ from the state-preparation routine that produces Eq. (1.1). This is operationally free: we simply refrain from uncomputing the j-label modulo P while preparing the coordinate registers. Crucially, $\mathbf{X}(j) = (2D^2 j\, \boldsymbol{b}^* + \boldsymbol{v}^*) \bmod M_2$ depends only on $j \bmod P$ because $2D^2 P \equiv 0 \pmod{M_2}$; hence a label in $\mathbb{Z}_P$ suffices to re-evaluate the preparation. In this re-evaluation route one uses J to re-evaluate the same reversible preparation map at $j + T$, and in cleanup (below) we use J to realize a b^*-free erasure of T.

Step $9^\dagger$.1 (copy). Use CNOT or modular addition gates to coherently copy the coordinate registers into $\mathbf{Y}$:

$$\sum_j \alpha(j)\, |\mathbf{X}(j)\rangle\, |\mathbf{0}\rangle \;\mapsto\; \sum_j \alpha(j)\, |\mathbf{X}(j)\rangle\, |\mathbf{X}(j)\rangle\,.$$

Step $9^\dagger$.2 (pair-evaluation shift). Using the arithmetic evaluator U_{prep} of Prop. 2.4, compute into $\mathbf{Y}$ the value corresponding to $j + T$ without reproducing any phases:

$$(\mathbf{X}(j), \mathbf{Y} = \mathbf{X}(j), J, T) \;\longmapsto\; (\mathbf{X}(j), \mathbf{Y} = \mathbf{X}(j+T), J, T),$$

where $J \equiv j \pmod P$ and $j+T$ is treated as an integer (all arithmetic inside the preparation circuit is modulo M_2). Equivalently,

$$\mathbf{Y} = \big(2D^2(j+T)b_1^* \;\big|\; 2D^2(j+T)\boldsymbol{b}^*_{[2..n]} + \boldsymbol{v}^*_{[2..n]}\big).$$

Remark 3.3 (No classical knowledge of $\boldsymbol{b}^*$ is required)**.** This step uses the arithmetic evaluator that computes $V + j\Delta$ with read-only data (V, Δ) and therefore never forms $2D^2T\, \boldsymbol{b}^*$ as an explicit classical constant and never modifies the pre-existing quadratic phase profile $\alpha(\cdot)$.

Remark 3.4 (Constant-adder realization when $\boldsymbol{b}^*$ is known)**.** If a classical description of $\boldsymbol{b}^*$ modulo P is available, one may instead implement this step by adding the constant vector $2D^2T\, \boldsymbol{b}^*$ coordinatewise (mod M_2). Only $\boldsymbol{b}^*$ mod P is needed, since $2D^2$ annihilates the $\mathbb{Z}_{D^2}$ component.

Step $9^\dagger$.3 (difference; offset cancellation). Compute the coordinatewise difference $\mathbf{Z} := \mathbf{X} - \mathbf{Y}$ (mod M_2) into a fresh n-register block:

$$\mathbf{Z} \leftarrow \mathbf{X} - \mathbf{Y} \pmod{M_2},$$

so that $\mathbf{Z} \equiv -2D^2T\,\boldsymbol{b}^* \pmod{M_2}$ and the unknown offsets $\boldsymbol{v}^*$ cancel exactly.

Step $9^\dagger$.4 (cleanup; re-evaluation variant). With residue accessibility (Definition 2.8), compute T' from $(\mathbf{Z}, \Delta)$ by per-prime inversions and CRT, then:

(i) Without modifying $\mathbf{Z}$, coherently update $\mathbf{Y}$ from $\mathbf{X}(j+T)$ to $\mathbf{X}(j+T-T')$ by re-evaluating U_{prep} on input $J+T-T'$ and subtracting the previously computed value $\mathbf{X}(j+T)$:

$$\mathbf{Y} \leftarrow \mathbf{Y} + \big(\mathbf{X}(J+T-T') - \mathbf{X}(J+T)\big) \pmod{M_2}.$$

(ii) Set $T \leftarrow T - T'$, so $T = 0$ and hence $\mathbf{Y} = \mathbf{X}(j)$.

(iii) Uncopy by applying the inverse of the copy to map $(\mathbf{X}, \mathbf{Y}) \mapsto (\mathbf{X}, \mathbf{0})$.

(iv) Erase T' by applying the inverse of its computation from $\mathbf{Z}$.

Remark 3.5 (Implementation note (index label availability))**.** If this re-evaluation route is used, the implementation must expose (and not uncompute) a computational-basis register $J \equiv j \pmod P$ during the superposition-time steps.

Remark 3.6 (Alternative when $\boldsymbol{b}^*$ is known modulo P)**.** One may undo the shift on $\mathbf{Y}$ using the constant adder $\mathbf{Y} \leftarrow \mathbf{Y} - 2D^2T'\,\boldsymbol{b}^*$. Here, invertibility of $2D^2$ modulo each p_η follows from oddness and $\gcd(D, P) = 1$.

Variant: pair-evaluation without classical $\boldsymbol{b}^*$. Let U_{prep} denote the arithmetic evaluator that sends $|j\rangle\,|\mathbf{0}\rangle \mapsto |j\rangle\,|\mathbf{X}(j)\rangle$ using (V, Δ) (suppressing ancillary work registers). Retain a label $J \equiv j$ (mod P). Then implement Step $9^\dagger$.2 as follows:

1. Compute $J + T$ in place (mod P).
2. Run U_{prep} on input $J + T$ into Y to obtain $\mathbf{X}(j+T)$.
3. (Optionally) restore J by subtracting T.

The subsequent difference $Z \leftarrow X - Y$ yields $Z \equiv -2D^2T\,\boldsymbol{b}^* \pmod{M_2}$, with the offsets cancelling identically. This realization needs no classical access to $\boldsymbol{b}^*$ (nor to $\boldsymbol{v}^*$).

Implementation note. In practice, set $\Delta = \mathbf{X}(1) - \mathbf{X}(0)$ (harvested once) and reduce $(\Delta, \mathbf{Z})$ modulo each p_η in parallel. For each prime, choose the lexicographically smallest coordinate $i(\eta)$ with $\Delta_i \not\equiv 0 \pmod{p_\eta}$ (deterministic and reversible), compute $\Delta_{i(\eta)}^{-1} \pmod{p_\eta}$ via a reversible extended Euclidean algorithm, and form $T_\eta \equiv -\Delta_{i(\eta)}^{-1} Z_{i(\eta)} \pmod{p_\eta}$. Recombine the residues by a reversible CRT (e.g., Garner mixed-radix). As D and all p_η are odd with $\gcd(D, P) = 1$, the factors 2 and D^2 are units modulo every p_η, and residue accessibility guarantees the existence of at least one invertible coordinate per prime. Keep T' as a dedicated scratch register that is not modified by any other step until it is uncomputed by inverting its computation from **Z**. For preparing $\frac{1}{\sqrt{P}}\sum_{t\in\mathbb{Z}_P}|t\rangle$, the per-prime preparation $\bigotimes_\eta \frac{1}{\sqrt{p_\eta}}\sum_{t_\eta\in\mathbb{Z}_{p_\eta}}|t_\eta\rangle$ followed by CRT wiring is exact and avoids approximation issues associated with a monolithic $\mathrm{QFT}_{\mathbb{Z}_P}$; this mirrors the modulus-splitting/CRT bookkeeping already used in Chen [2024]. The unit factor -2 in the generator is immaterial (any fixed unit modulo P yields the same annihilator); we keep it to match Eq. (1.1).

3.4 Exact correctness

Lemma 3.7 (Cyclic embedding)**.** Under Definition 2.8, the map $\phi : \mathbb{Z}_P \to (\mathbb{Z}_{M_2})^n$ given by $\phi(T) = -2D^2T\,\boldsymbol{b}^* \pmod{M_2}$ is an injective group homomorphism. Hence, its image is a cyclic subgroup of order P, and the state in Eq. (3.2) is uniform over a subgroup-coset of size P.

Proof. Homomorphism is immediate. For injectivity, reduce modulo P: if $\phi(T) \equiv \mathbf{0}$, then $2D^2T\,\boldsymbol{b}^* \equiv \mathbf{0} \pmod P$. Since $2D^2$ is a unit modulo P and by Definition 2.8 some coordinate of $\boldsymbol{b}^*$ is a unit modulo each p_η, we must have $T \equiv 0 \pmod{p_\eta}$ for all η. The Chinese Remainder Theorem gives $T \equiv 0 \pmod P$. Moreover, under the CRT decomposition $\mathbb{Z}_{M_2} \cong \mathbb{Z}_{D^2} \times \mathbb{Z}_P$, the image of ϕ lies entirely in the $\mathbb{Z}_P$-component (the $\mathbb{Z}_{D^2}$ projection is 0), and residue accessibility guarantees that, for each p_η, some coordinate has order p_η. Hence the subgroup has order exactly $\prod_\eta p_\eta = P$. □

Lemma 3.8 (Exact orthogonality from a CRT-coset)**.** Consider the uniform superposition over the CRT-coset generated by $\boldsymbol{b}^*$:

$$|\Psi\rangle = \frac{1}{\sqrt{P}}\sum_{T\in\mathbb{Z}_P}\left|-2D^2T\,\boldsymbol{b}^* \mod M_2\right\rangle.$$

After $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$, the amplitude of $\boldsymbol{u} \in \mathbb{Z}_{M_2}^n$ is

$$A(\boldsymbol{u}) \;=\; \frac{1}{\sqrt{M_2^n}}\cdot\frac{1}{\sqrt{P}}\sum_{T=0}^{P-1}\exp\!\Big(\tfrac{2\pi i}{M_2}\left\langle -2D^2T\,\boldsymbol{b}^*,\,\boldsymbol{u}\right\rangle\Big) \;=\; \frac{1}{\sqrt{M_2^n}}\cdot\frac{1}{\sqrt{P}}\sum_{T=0}^{P-1}\Big(\exp\tfrac{2\pi i}{P}\cdot(-2)\,\langle\boldsymbol{b}^*,\boldsymbol{u}\rangle\Big)^T.$$

Only the $\mathbb{Z}_P$-component of $\boldsymbol{u}$ influences the sum over T (the $\mathbb{Z}_{D^2}$ projection cancels since $M_2 = D^2P$). Because P is odd, 2 is invertible modulo P. Hence $A(\boldsymbol{u}) = 0$ unless $\langle\boldsymbol{b}^*,\boldsymbol{u}\rangle \equiv 0 \pmod P$, in which case $|A(\boldsymbol{u})| = \sqrt{P}/M_2^{n/2}$ (up to a global phase). Consequently, the measurement outcomes are exactly supported on Eq. (1.2) and are uniform over that set; indeed,

$$\#\{\boldsymbol{u}\in(\mathbb{Z}_{M_2})^n :\ \langle\boldsymbol{b}^*,\boldsymbol{u}\rangle \equiv 0 \pmod P\} \;=\; \frac{M_2^n}{P}.$$

Since each feasible $\boldsymbol{u}$ occurs with probability P/M_2^n and there are M_2^n/P of them, the total probability sums to 1.

Proof. Let $r := \exp\!\big(\frac{2\pi i}{M_2}\cdot(-2D^2)\,\langle\boldsymbol{b}^*,\boldsymbol{u}\rangle\big) = \exp\!\big(-\frac{2\pi i}{P}\cdot 2\,\langle\boldsymbol{b}^*,\boldsymbol{u}\rangle\big)$. Because P is odd, 2 is a unit modulo P, and only the $\mathbb{Z}_P$-component of the phase contributes to the sum over T (the $\mathbb{Z}_{D^2}$-component cancels since $M_2 = D^2P$). Note also that $r^P = \exp\!\big(-\frac{2\pi i}{M_2}\,2D^2P\,\langle\boldsymbol{b}^*,\boldsymbol{u}\rangle\big) = 1$ for all $\boldsymbol{u}$, so the geometric sum over $T\in\mathbb{Z}_P$ always collapses to either 0 or P. Since $M_2 = D^2P$, we have $\frac{-2D^2}{M_2} \equiv -\frac{2}{P} \pmod 1$, i.e., only the P-component of the phase matters in the sum over T; this is exactly why the base of the geometric progression is $e^{\frac{2\pi i}{P}(-2)\langle\boldsymbol{b}^*,\boldsymbol{u}\rangle}$. Because P is odd, 2 is invertible mod P. Thus $r = 1$ iff $\langle\boldsymbol{b}^*,\boldsymbol{u}\rangle \equiv 0 \pmod P$. The sum $\sum_{T=0}^{P-1} r^T$ is P if $r = 1$ and 0 otherwise; multiplying by the prefactor $M_2^{-n/2}P^{-1/2}$ gives the stated amplitude magnitude. □

At each prime p_η, Definition 2.8 guarantees that the linear form $\boldsymbol{u} \mapsto \langle \boldsymbol{b}^*, \boldsymbol{u} \rangle$ has rank 1 over $\mathbb{Z}_{p_\eta}$, so the solution set on $(\mathbb{Z}_{p_\eta})^n$ has size p_η^{n-1}. By CRT this gives P^{n-1} solutions on the $\mathbb{Z}_P$-part, while the $\mathbb{Z}_{D^2}$-parts are unconstrained and contribute $(D^2)^n$, yielding a total of $(D^2)^n P^{n-1} = M_2^n / P$.

Group-theoretic perspective. For a finite abelian group G and a subgroup $H \le G$, the QFT on the uniform superposition over any coset of H produces uniform support on the annihilator $H^\perp \subseteq \widehat{G}$. Taking $G = (\mathbb{Z}_{M_2})^n$, $H = \langle -2D^2 \boldsymbol{b}^* \rangle$, and identifying $\widehat{G} \cong G$ via the standard pairing, we recover Lemma 3.8 with $H^\perp = \{\boldsymbol{u} : \langle \boldsymbol{b}^*, \boldsymbol{u} \rangle \equiv 0 \pmod{P}\}$. The overall sign is immaterial since -1 is a unit modulo P.

Theorem 3.9 (Step $9^\dagger$ is correct)**.** Assume Assumption 2.2 and Definition 2.8. Starting from Eq. (1.1), after executing either (i) the default J-free route (Steps $9^\dagger.2'$ and $9^\dagger.4$), or (ii) the re-evaluation route (Steps $9^\dagger.1$–$9^\dagger.4$), the state factors as in Eq. (3.2). In all cases, U_{coords} is never applied on superpositions. Applying $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$ to the $\mathbf{Z}$-register and measuring yields $\boldsymbol{u} \in \mathbb{Z}_{M_2}^n$ uniformly distributed over the solutions of Eq. (1.2). The offsets $\boldsymbol{v}^*$ and the quadratic phases $\alpha(j)$ do not affect the support or uniformity of the measured $\boldsymbol{u}$.

Proof. Eq. (3.1) shows $\mathbf{Z}$ depends only on T, not on j or $\boldsymbol{v}^*$. Under Definition 2.8, Step $9^\dagger.4$ erases T and yields the factorization Eq. (3.2); the part carrying $\alpha(j)$ is in registers disjoint from $\mathbf{Z}$. By Lemma 3.8, Fourier sampling of $\mathbf{Z}$ yields Eq. (1.2) uniformly. Neither $\boldsymbol{v}^*$ nor $\alpha(j)$ enters that calculation. □

Remark 3.10 (Approximate QFTs)**.** In practice, $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$ will be implemented approximately. Let a single-register QFT be U and an implementation be $\widetilde{U}$ with $\|U - \widetilde{U}\|_{\text{op}} \le \varepsilon_1$. A telescoping argument gives

$$\left\| U^{\otimes n} - \widetilde{U}^{\otimes n} \right\|_{\text{op}} \;\le\; n\, \varepsilon_1.$$

Consequently, for any input state, the output state's ℓ_2 error is at most $n\varepsilon_1$, and for any measurement, the induced total-variation distance between the ideal and realized outcome distributions is at most $n\varepsilon_1$. If one prefers a single parameter, write $\varepsilon_n := \|U^{\otimes n} - \widetilde{U}^{\otimes n}\|_{\text{op}} \le n\varepsilon_1$, and the leakage mass is $\le \varepsilon_n$. The support (solutions to Eq. (1.2)) remains the ideal annihilator; approximation affects only leakage probability, not the constraint itself.

Remarks. (i) No amplitude periodicity is used anywhere. (ii) The offsets $\boldsymbol{v}^*$ are canceled exactly by construction; no knowledge of their residues is required. (iii) The residue accessibility condition (Definition 2.8) is operationally necessary. It enables the erasure of T from the rest of the state, which ensures that a coherent uniform coset forms on the $\mathbf{Z}$ register. Without it, the Fourier sampling step would fail, as discussed in Section 4. (iv) Edge case $n = 1$: with $b_1^* = p_2 \cdots p_\kappa$, the condition in Definition 2.8 cannot hold (it vanishes modulo every p_η for $\eta \ge 2$), consistent with upstream requirements that $n \ge 2$. (v) The optional J-free realization (Step $9^\dagger.2'$) produces the same $\mathbf{Z}$ and avoids carrying index labels or re-evaluation ancillas. (vi) The factor 2 in the generator $-2D^2\, T\, \boldsymbol{b}^*$ is inessential: any fixed unit modulo P yields the same annihilator condition. We keep the factor 2 to align with the upstream normalization in Eq. (1.1).

Connection back to Chen [2024]. Under the CRT viewpoint, Step $9^\dagger$ replaces the domain-extension-on-one-coordinate maneuver with a coset synthesis that is agnostic to offsets. Conceptually, we embed $\mathbb{Z}_P$ into $(\mathbb{Z}_{M_2})^n$ via $T \mapsto -2D^2T\boldsymbol{b}^*$, average uniformly over the orbit, and then read off the annihilator by QFT. This directly yields the intended linear relation modulo P without invoking amplitude periodicity across heterogeneous coordinates.

4 Complexity and variants

Complexity. Copying registers and reversible modular adders and multipliers over $\mathbb{Z}_{M_2}$ use $O(\mathrm{poly}(\log M_2))$ gates. The shift $\mathbf{Z} \leftarrow \mathbf{Z} - 2D^2T\,\boldsymbol{b}^*$ costs $O\big(n\,\mathrm{poly}(\log M_2)\big)$. Computing $\mathbf{Z} = \mathbf{X} - \mathbf{Y}$ is linear in n. Uncomputing T needs κ modular reductions and inverses in $\mathbb{Z}_{p_\eta}$ and one CRT recombination. A reversible extended Euclid for one inverse costs $O\big((\log p_\eta)^2\big)$ gates, or $\widetilde{O}(\log p_\eta)$ with half-GCD. CRT recombination works with either a Garner mixed-radix scheme in $O(\kappa^2)$ modular steps, or a remainder and product tree in $O(\kappa \log \kappa)$ steps. Word sizes stay in $\mathrm{poly}(\log P)$, and we keep all intermediate digits for clean uncomputation. The transform $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$ costs $O\big(n\,\mathrm{poly}(\log M_2)\big)$.

The subroutine matches the time and success bounds of Chen [2024]. No amplitude amplification is needed. The support on the target coset is exact and uniform.

The method does not need a periodic amplitude function or any phase flattening. All dependence on j and on $\boldsymbol{v}^*$ stays in registers that are disjoint from $\mathbf{Z}$. These terms do not affect the Fourier sample.

If residue accessibility fails. If Definition 2.8 fails for some prime p_η, the map $T \mapsto T\,\boldsymbol{b}^* \pmod P$ has a nontrivial kernel. Then T is not a function of $\mathbf{Z}$ mod P. Coherent erasure of T is not possible. Fourier sampling on $\mathbf{Z}$ alone becomes uniform over $\mathbb{Z}_{M_2}^n$ and does not force Eq. (1.2). Two paths remain:

1. Enforce the condition modulo $P' = \prod_{\eta \in \mathcal{I}} p_\eta$, where accessibility holds. Handle the missing primes by adding one or more auxiliary directions or by a short unimodular re-basis so that each missing prime is accessible in at least one coordinate. Then rerun the coset step for those primes. The measured $\boldsymbol{u}$ then obeys $\langle \boldsymbol{b}^*, \boldsymbol{u} \rangle \equiv 0 \pmod{P'}$ exactly and is free modulo the other primes. Downstream linear algebra can consume this partial set and repeat after fixing the rest.

2. Use a postselection fallback. First unshift $\mathbf{Y}$ by the known T, that is, apply $\mathbf{Y} \leftarrow \mathbf{Y} - 2D^2T\,\boldsymbol{b}^*$. Then apply QFT^{-1} to T and keep the zero frequency. The outcome is a coherent uniform coset on $\mathbf{Z}$ without computing T from $\mathbf{Z}$. The zero frequency appears with probability $1/P$. Amplitude amplification raises this rate to $\Theta(1)$ at a cost of $\Theta(\sqrt{P})$ queries.

We adopt Definition 2.8. It gives deterministic cleanup with no postselection cost.

Alternative modulus choices. Under Definition 2.8 we can compute the coset label $J = T$ from $\mathbf{Z}$ mod P. Applying $\mathrm{QFT}_{\mathbb{Z}_P}$ to J produces a flat spectrum over $\mathbb{Z}_P$, but this step alone does not force Eq. (1.2). A safe route is to map J back into $\mathbf{Z}$ by $-2D^2J\,\boldsymbol{b}^* \pmod{M_2}$ and then apply $\mathrm{QFT}_{\mathbb{Z}_{M_2}}^{\otimes n}$, identical to the main path. We keep the J-free variant for clarity.

5 Conclusion

We presented a reversible Step $9^\dagger$ that (i) cancels unknown offsets exactly, (ii) synthesizes a coherent, uniform CRT-coset state without amplitude periodicity, and (iii) yields the intended modular linear relation via an exact character-orthogonality argument. The subroutine is simple to implement, asymptotically light, and robust. We expect the pair-shift difference pattern to be broadly useful in windowed-QFT pipelines whenever unknown offsets obstruct clean CRT lifting.

Acknowledgment

We are grateful to all who provided constructive discussions and helpful feedback.

References

Howard Barnum, Carlton M Caves, Christopher A Fuchs, Richard Jozsa, and Benjamin Schumacher. Noncommuting mixed states cannot be broadcast. *Physical Review Letters*, 76(15):2818, 1996.

Yilei Chen. Quantum algorithms for lattice problems. *Cryptology ePrint Archive*, 2024.

DGBJ Dieks. Communication by epr devices. *Physics Letters A*, 92(6):271–272, 1982.

Michael A Nielsen and Isaac L Chuang. *Quantum computation and quantum information*. Cambridge university press, 2010.

William K Wootters and Wojciech H Zurek. A single quantum cannot be cloned. *Nature*, 299(5886): 802–803, 1982.

Appendices

A Mechanics inside Step $9^{\dagger}$

Offset cancellation. Write

$$\mathbf{X}(j) = \left(2D^2 j\, b_1^* \;\middle|\; 2D^2 j\, \boldsymbol{b}^*_{[2..n]} + \boldsymbol{v}^*_{[2..n]}\right), \qquad \mathbf{X}(j+T) = \left(2D^2 (j+T)\, b_1^* \;\middle|\; 2D^2 (j+T)\, \boldsymbol{b}^*_{[2..n]} + \boldsymbol{v}^*_{[2..n]}\right).$$

Then

$$\mathbf{X}(j) - \mathbf{X}(j+T) \equiv -2D^2 T\, \boldsymbol{b}^* \pmod{M_2},$$

so the offset $\boldsymbol{v}^*$ vanishes identically.

Uniform CRT coset on Z. After Step $9^{\dagger}.4$ we have erased T from the rest. A uniform superposition over $T \in \mathbb{Z}_P$ maps by

$$T \longmapsto -2D^2 T\, \boldsymbol{b}^* \pmod{M_2}$$

to a coherent uniform coset on $\mathbf{Z}$ of length P. No amplitude reweighting appears. The image is cyclic of order P by Lemma 3.7.

Orthogonality check. For any $\boldsymbol{u}$ the phase base is

$$r = \exp\Big(-\tfrac{2\pi i}{M_2}\, 2D^2\, \langle \boldsymbol{b}^*, \boldsymbol{u} \rangle \Big).$$

We have

$$r^P = \exp\Big(-\tfrac{2\pi i}{M_2}\, 2D^2 P\, \langle \boldsymbol{b}^*, \boldsymbol{u} \rangle \Big) = 1,$$

with $M_2 = D^2 P$. So the P-term geometric sum collapses exactly. Equivalently,

$$\frac{-2D^2}{M_2} \equiv -\frac{2}{P} \pmod{1},$$

which makes the reduction to phases modulo P explicit.

B Proof of State Factorization

For completeness, we show that the state after cleanup (Step $9^{\dagger}.4$) factors as claimed, and we contrast it with the pre-cleanup mixed state on $\mathbf{Z}$ (this also makes Prop. 2.10 fully formal). Let the joint state after Step $9^{\dagger}.2$ be

$$|\Phi_2\rangle \;=\; \frac{1}{\sqrt{P}} \sum_{t \in \mathbb{Z}_P} \sum_{j} \alpha(j)\; |\mathbf{X}(j)\rangle_{\mathbf{X}}\, |\mathbf{X}(j) + 2D^2 t\, \boldsymbol{b}^*\rangle_{\mathbf{Y}}\, |t\rangle_T\,.$$

Computing $\mathbf{Z} \leftarrow \mathbf{X} - \mathbf{Y}$ gives

$$|\Phi_3\rangle \;=\; \frac{1}{\sqrt{P}} \sum_{t} \sum_{j} \alpha(j)\; |-2D^2 t\, \boldsymbol{b}^*\rangle_{\mathbf{Z}}\, |\mathbf{X}(j)\rangle_{\mathbf{X}}\, |\mathbf{X}(j) + 2D^2 t\, \boldsymbol{b}^*\rangle_{\mathbf{Y}}\, |t\rangle_T\,.$$

Tracing out $(\mathbf{X}, \mathbf{Y}, T)$ at this point leaves the mixed state

$$\rho_{\mathbf{Z}} \;=\; \frac{1}{P} \sum_{t \in \mathbb{Z}_P} |-2D^2 t\, \boldsymbol{b}^*\rangle\langle -2D^2 t\, \boldsymbol{b}^*|\,,$$

since the different t-branches are orthogonal in the T-register. Under Definition 2.8, Step $9^\dagger.4$ computes t from $\mathbf{Z}$ mod P and uncomputes the original T-register (and $\mathbf{X}, \mathbf{Y}$), yielding the factorized pure state

$$\Big(\sum_j \alpha(j)\,|\text{junk}(j)\rangle\Big) \;\otimes\; \frac{1}{\sqrt{P}} \sum_{t \in \mathbb{Z}_P} |-2D^2 t\, \boldsymbol{b}^*\rangle_{\mathbf{Z}}\,,$$

which is exactly Eq. (3.2). □

C Gate skeleton for the shift and difference

Route map. Items (1), (2), and (4) below are used only in the *re-evaluation route*; the *J-free route* uses item (3) directly to form $\mathbf{Z} \leftarrow -T \cdot \Delta$ and skips copy/difference. Cleanup (item (5)) applies to both routes (with the re-evaluation sub-steps when $\mathbf{Y}$ is present).

Each coordinate uses the same pattern (we suppress the index):

1. Copy: CNOTs (or modular adds) from X into Y.
2. Shift (optional re-evaluation route): add $2D^2 b^* \cdot T$ into Y via a controlled modular adder with precomputed $2D^2 b^* \pmod{M_2}$.
3. Shift (default J-free): set $Z \leftarrow -T \cdot \Delta \pmod{M_2}$ using double-and-add with Δ as read-only data (no classical access to $\boldsymbol{b}^*$).
4. Difference: set $Z \leftarrow X - Y$ using a modular subtractor; this can overwrite X if desired.
5. Cleanup: use the harvested $\Delta \leftarrow X(1) - X(0)$; compute $T' \leftarrow f(Z, \Delta)$ into an auxiliary by, for each p_η, choosing a coordinate with $\Delta_i \not\equiv 0 \pmod{p_\eta}$, inverting Δ_i modulo p_η, and CRT-recombining; if using the optional route, update $Y \leftarrow Y + \big(X(J + T - T') - X(J + T)\big)$ via the reversible evaluator U_{prep}; set $T \leftarrow T - T'$; if using the optional route, apply the inverse of the copy to clear Y; uncompute T' from Z. (All steps preserve Z.)

Phase discipline. All arithmetic inside U_{prep} uses classical reversible (Toffoli/Peres) adders/multipliers; no QFT-based adders are used. This ensures that applying U_{prep} on superpositions introduces no data-dependent phases.
Determinism across invocations. Basis calls to U_{coords} (such as $0, 1$ or $J, J{+}1$) use fixed classical constants within a single run so that $\mathbf{X}(\cdot)$ is reproducible as computational-basis data.

Variant: pair-evaluation without classical $\boldsymbol{b}^*$. Let U_{prep} denote the arithmetic evaluator that sends $|j\rangle\,|\mathbf{0}\rangle \mapsto |j\rangle\,|\mathbf{X}(j)\rangle$ using the harvested (V, Δ) (suppressing ancillary work registers). Retain a label $J \equiv j \pmod{P}$. Then implement Step $9^\dagger.2$ as follows:

1. Compute $J + T$ in place (mod P).
2. Run U_{prep} on input $J + T$ into Y to obtain $\mathbf{X}(j + T)$.

3. (Optionally) restore J by subtracting T.

The subsequent difference $Z \leftarrow X - Y$ yields $Z \equiv -2D^2 T\, \boldsymbol{b}^* \pmod{M_2}$, with the offsets cancelling identically. This realization needs no classical access to $\boldsymbol{b}^*$ (nor to $\boldsymbol{v}^*$).

Implementation note. In practice, set $\Delta = \mathbf{X}(1) - \mathbf{X}(0)$ (harvested once) and reduce $(\Delta, \mathbf{Z})$ modulo each p_η in parallel. For each prime, choose the lexicographically smallest coordinate $i(\eta)$ with $\Delta_i \not\equiv 0 \pmod{p_\eta}$ (deterministic and reversible), compute $\Delta_{i(\eta)}^{-1} \pmod{p_\eta}$ via a reversible extended Euclidean algorithm, and form $T_\eta \equiv -\Delta_{i(\eta)}^{-1} Z_{i(\eta)} \pmod{p_\eta}$. Recombine the residues by a reversible CRT (e.g., Garner mixed-radix), keeping the mixed-radix digits and running-product moduli so they can be uncomputed exactly in reverse. Since $\gcd(D, P) = 1$ and each p_η is odd, the factors 2 and D^2 are units modulo every p_η, and residue accessibility guarantees the existence of at least one invertible coordinate per prime. Keep T' as a dedicated scratch register that is not modified by any other step until it is uncomputed by inverting its computation from $\mathbf{Z}$. For preparing $\frac{1}{\sqrt{P}} \sum_{t \in \mathbb{Z}_P} |t\rangle$, the per-prime preparation $\bigotimes_\eta \frac{1}{\sqrt{p_\eta}} \sum_{t_\eta \in \mathbb{Z}_{p_\eta}} |t_\eta\rangle$ followed by CRT wiring is exact and avoids approximation issues associated with a monolithic $\mathrm{QFT}_{\mathbb{Z}_P}$; this mirrors the modulus-splitting/CRT bookkeeping already used in Chen [2024]. The unit factor -2 in the generator is immaterial (any fixed unit modulo P yields the same annihilator); we keep it to match Eq. (1.1).

D Run-local determinism

A run is one coherent execution from the start of state preparation up to (and including) Step $9^\dagger$. Within a run, the coordinate evaluator U_{coords} uses a fixed set of classical constants (including any classical values obtained by earlier measurements in the same run, such as y', z', h^* in Chen [2024]). Hence, the basis outputs $\mathbf{X}(0)$ and $\mathbf{X}(1)$ are reproducible within that run. We harvest

$$V := \mathbf{X}(0), \qquad \Delta := \mathbf{X}(1) - \mathbf{X}(0) \equiv 2D^2\, \boldsymbol{b}^* \pmod{M_2},$$

once on literal inputs $j = 0, 1$ and then treat (V, Δ) as read-only basis data.

All superposition-time arithmetic (copy/shift/difference/cleanup) is implemented by classical reversible circuits (no QFT-based adders), so it is a permutation of computational-basis states and introduces no data-dependent phase (Lemma 2.3). We never call U_{coords} on a superposed input.

Approximate QFTs may be used for standard transforms; their approximation error is tracked separately (Remark after Theorem 3.9) and is unrelated to determinism of (V, Δ).

Across different runs, the upstream randomness, offsets, and even the arithmetic constants used by U_{coords} may change. Our proofs do not assume that (V, Δ) are identical across runs, nor do they assume any global seeding, device-level determinism, or that the overall global phase is fixed. The only place determinism is needed is to ensure that the single-run harvest (V, Δ) is well-defined and then reused verbatim by U_{prep} in that same run.

Under this scope, the cleanup step can always compute T' from $(\mathbf{Z}, \Delta)$ when Definition 2.8 holds, guaranteeing the factorization in Eq. (3.2). If desired, one may even measure (V, Δ) early and cache them as classical strings; this does not affect correctness or phases because we never feed U_{coords} with a superposition.