Tensor Encodings for SIMD HSS

Jaehyung Kim

Stanford University
jaehk@cs.stanford.edu

Abstract. We study SIMD packing for the lattice-based homomorphic secret sharing scheme of Boyle–Kohl–Scholl (BKS) over dimension-N cyclotomic rings. A trace construction with alternating tensor encodings matches the $\Theta(\sqrt{N})$ packing of SIMD-HSS by Kim et al. (ePrint 2026/485). Its addition-closed mode uses $O(\log N)$ authenticated automorphisms, two BKS multiplications, and two constant multiplications; an alternating fast path roughly halves these costs. A second construction uses a three-term-progression-free slot set A: homomorphic traces isolate the product coefficients at $2a$ for $a \in A$, and a halving automorphism returns them to a. For fixed k and conductor primes, with balanced prime-power factors, this packs $N^{1-o(1)}$ slots with one BKS multiplication, $O(kN^{1/(2k)})$ authenticated automorphisms, and $O(kN^{1/k})$ constant multiplications. Both constructions support standard (non-entropic) secrets. Two-party executions at 162 and 495 slots confirm the algebra and exact call counts.

Keywords: Homomorphic secret sharing · SIMD · RLWE

1 Introduction

Homomorphic secret sharing (HSS) was introduced by Boyle, Gilboa, and Ishai, who also formulated its public-key form [4]; subsequent work systematized its general multi-input and multi-server syntax [5]. In the two-evaluator public-key setting, the evaluators transform public encodings of private inputs into additive shares of a function value, without interaction during evaluation. The lattice-based construction of Boyle, Kohl, and Scholl (BKS) [6] is particularly attractive for restricted-multiplication straight-line (RMS) programs: its multiplication is essentially distributed RLWE decryption, rather than an FHE multiplication. Its efficiency nevertheless depends strongly on how many independent computations can share one ring operation. Throughout, N denotes the dimension of the ambient cyclotomic ring.

CRT decompositions have long enabled SIMD evaluation in FHE [17], but they do not directly supply the bounded integral semantics needed by lattice HSS [6]. Kim et al. [11] give a SIMD-HSS construction with depth-independent parameters, packing roughly the square root of the ring dimension by arranging bounded integers in coefficient intervals and moving them with authenticated-share linear maps. Their type conversion, however, relies on entropic circular

RLWE with a predetermined set of secret coefficients fixed to zero. This raises two natural questions: can one obtain substantially denser packing while retaining native HSS evaluation, and can one do so without this structured-secret restriction?

We answer both questions affirmatively, with an explicit extraction tradeoff. Neither construction prescribes which secret coefficients vanish; both admit ordinary bounded secret distributions, including independent uniform ternary, under circular RLWE instantiated for the chosen distribution and automorphism family. For any fixed number k of balanced tensor factors, the progression-free construction occupies $N^{1-o(1)}$ coefficients and uses one BKS multiplication and $O(kN^{1/(2k)})$ authenticated automorphisms per SIMD multiplication, while requiring $O(kN^{1/k})$ constant multiplications for coefficient projection.

Lyubashevsky, Peikert, and Regev (LPR) develop powerful tensor decompositions and dual bases for arbitrary cyclotomic rings [15]. Liu and Wang use complementary tensor subrings, dual bases, and relative traces to retain diagonal products in batch bootstrapping [13]. We instantiate this algebraic framework in HSS by maintaining the tensor encodings as authenticated BKS shares. The central technical requirement is to reconcile trace-based coefficient extraction with integral coefficient encodings, particularly outside power-of-two cyclotomics.

We use LPR's unconjugated coefficient power basis and bilinear trace dual, together with integral conductor scaling, so that all extraction maps can be evaluated within the BKS authenticated-share interface. This viewpoint leads to the two constructions summarized next.

1.1 Contributions

Near-linear packing. Our second construction stores values at a Cartesian product of guarded, three-term-progression-free exponent sets [3]. Multiplication creates a tensor convolution; factorwise integral coordinate projectors retain its diagonal; and one global automorphism returns the diagonal to the original slot positions. For fixed k and fixed conductor primes, with balanced prime-power factors, this yields $N^{1-o(1)}$ slots using one base-HSS multiplication and $O(kN^{1/k})$ constant multiplications. Using a depth-two baby-step–giant-step circuit [10], the authenticated-automorphism count is $O(kN^{1/(2k)})$. This establishes the asymptotic feasibility of near-linear HSS packing and improves the per-slot base-HSS and authenticated-automorphism call counts relative to paired Construction 1, with a growing general-ring projector. Section 5 reports the exact resource profile.

Complementary traces and an alternating fast path. Construction 1 instantiates Liu and Wang's complementary-ring, dual-basis, and relative-trace pattern [13] inside BKS HSS. We retain their conductor-clearing mechanism but use LPR's alternative unconjugated power-basis and trace-dual convention [15]. The default representation maintains both encodings as authenticated shares, so arbitrary RMS additions remain immediate; with balanced factors, it packs $\Theta(\sqrt{N})$ values using two base-HSS multiplications, $O(\log N)$ authenticated automorphisms,

and two scalar normalizations. On type-consistent regions, a typed fast path retains one encoding and alternates its type, reducing each multiplication to one base-HSS call, one directional trace, and one scalar normalization. This supplies a compact scalar extraction layer together with an addition-closed default representation.

Removing the entropic-secret restriction. Both constructions avoid the ring conversion responsible for the fixed-zero requirement in prior SIMD HSS [11]. Their only secret-dependent linear transformations are authenticated automorphisms of the ambient ring; all other extraction maps use public constants, additions, and decoding. Consequently, correctness requires a bounded secret, while privacy relies on circular RLWE for the automorphisms used by authenticated key switching; neither argument imposes a coefficient-support pattern. In particular, independent uniform ternary secrets are supported when this assumption holds for that distribution (Proposition 1).

Correctness, security, and implementation. We prove the new tensor identities and extraction circuits, and compose them with the correctness and security interfaces of BKS and SIMD-HSS. An end-to-end implementation exercises the complete two-party HSS evaluation. Executions at 162 and 495 slots reproduce the predicted SIMD outputs and match the derived operation counts; the Construction 1 evaluation also checks the directional fast path.

1.2 Technical Overview

Common interface. Let $\mathcal{R}_{q,M} = \mathbb{Z}_q[X]/(\Phi_M(X))$ and $\boldsymbol{s} = (1, s)^{\mathsf{T}}$. An authenticated share $\langle \boldsymbol{s}z \rangle$ consists of two party-local vectors whose difference, party 1 minus party 0 modulo q, is $\boldsymbol{s}z$. The BKS multiplication interface [6] maps a ciphertext ct_x of a ring element x and $\langle \boldsymbol{s}y \rangle$ to $\langle \boldsymbol{s}xy \rangle$; we call this one *base-HSS multiplication*. The authenticated-automorphism interface of SIMD-HSS [11] maps $\langle \boldsymbol{s}z \rangle$ to $\langle \boldsymbol{s}\sigma(z) \rangle$ by key switching and a local ring automorphism. Multiplication by a public ring element and addition act linearly on the shares, while rerandomization preserves the authenticated value. Our task is therefore algebraic: encode vectors so that one or two ring products can be reduced to componentwise products $\boldsymbol{x} \odot \boldsymbol{y}$ by an ambient-ring linear circuit.

All extraction maps below stay in $\mathcal{R}_{q,M}$. Whenever a trace-dual constant is not integral, we multiply it by its cyclotomic conductor before evaluation and remove the resulting scale by an invertible public scalar modulo q. Beyond the inherited BKS multiplication, the only secret-dependent maps are authenticated automorphisms of the ambient ring; neither construction converts to a different quotient ring or truncates the secret. Accordingly, under circular RLWE for the function family comprising zero and the required automorphism maps, no secret coordinate is required to vanish (Proposition 1); this includes independent uniform ternary secrets when the assumption holds for that distribution and automorphism set.

Construction 1. Following the alternating complementary-ring pattern of Liu and Wang [13], let $M_1 = 2^\alpha$ and $M_2 = 3^\beta$ for $\alpha \geq 2$ and $\beta \geq 1$, and put $N_i = \varphi(M_i)$, $N = N_1 N_2$, and $T = \min(N_1, N_2)$. For q coprime to 6, identify $\mathcal{R}_{q,M_1M_2}$ with $\mathcal{R}_{q,M_1} \otimes \mathcal{R}_{q,M_2}$. Let (v_i) and (w_i) be the two power bases. Using the unconjugated trace-dual convention recorded by LPR [15], choose $v_i^\vee, w_i^\vee$ so that $\mathrm{Tr}_{M_1}(v_i^\vee v_j) = \mathrm{Tr}_{M_2}(w_i^\vee w_j) = \delta_{i,j}$, and put $v_i^\sharp = M_1 v_i^\vee$ and $w_i^\sharp = M_2 w_i^\vee$; these scaled duals are integral. Define

$$\mathcal{E}_A(\boldsymbol{x}) = \sum_{i<T} x_i(v_i \otimes 1), \qquad \mathcal{E}_B(\boldsymbol{x}) = \sum_{i<T} x_i(1 \otimes w_i),$$
$$\mathcal{C}_{A\to B}(\boldsymbol{x}) = \sum_{i<T} x_i(v_i^\sharp \otimes w_i), \quad \mathcal{C}_{B\to A}(\boldsymbol{x}) = \sum_{i<T} x_i(v_i \otimes w_i^\sharp).$$

Every memory vector is maintained in both pure encodings, and every input is encrypted in both connector encodings. Trace duality gives

$$(\mathrm{Tr}_{M_1} \otimes \mathrm{id})\big(\mathcal{C}_{A\to B}(\boldsymbol{x})\mathcal{E}_A(\boldsymbol{y})\big) = M_1 \mathcal{E}_B(\boldsymbol{x} \odot \boldsymbol{y}),$$
$$(\mathrm{id} \otimes \mathrm{Tr}_{M_2})\big(\mathcal{C}_{B\to A}(\boldsymbol{x})\mathcal{E}_B(\boldsymbol{y})\big) = M_2 \mathcal{E}_A(\boldsymbol{x} \odot \boldsymbol{y}).$$

Indeed, each trace replaces $v_i^\sharp v_j$ or $w_i^\sharp w_j$ by a scaled Kronecker delta, eliminating every off-diagonal product. The default construction evaluates both connector–memory products, using two base-HSS multiplications and two constant multiplications by inverse-conductor scalars to restore the paired invariant. The binary trace uses $\alpha - 1$ authenticated automorphisms and the ternary trace uses $2\beta - 1$, so maintaining both encodings costs $\alpha + 2\beta - 2 = O(\log N)$ calls. A type-consistent region may instead retain only $\mathcal{E}_A$ or $\mathcal{E}_B$: one multiplication follows the applicable identity above, changes the type, and costs one base-HSS multiplication, respectively $\alpha - 1$ or $2\beta - 1$ automorphisms, and one scalar normalization. Same-type additions preserve this fast-path invariant. At an incompatible join, the evaluator can either have retained the paired representation or restore the missing type using one public unit connector, the corresponding directional trace, and one scalar normalization. The exact packing is $T = \min(N_1, N_2)$ and becomes $\Theta(\sqrt{N})$ when the factors are balanced.

Construction 2. Fix $k \geq 2$. For pairwise-coprime odd prime powers $m_i = p_i^{e_i}$, put $M = \prod_{i=1}^k m_i$, $n_i = \varphi(m_i)$, and $N = \prod_i n_i$, and assume $\gcd(q, M) = 1$. Use $\mathcal{R}_{q,M} \cong \bigotimes_i \mathcal{R}_{q,m_i}$ and let X_i be the power-basis generator in factor i. Adapting the 3-AP-free diagonal-support idea of Block, Maji, and Nguyen [3], choose for each factor a set $A_i \subseteq \{0, \ldots, L_i - 1\}$ with $L_i = \lfloor (n_i + 1)/2 \rfloor$ such that $b + c = 2a$ for $a, b, c \in A_i$ implies $a = b = c$. For $A = \prod_i A_i$ and $T = |A|$, encode

$$\mathcal{E}(\boldsymbol{x}) = \sum_{\boldsymbol{a} \in A} x_{\boldsymbol{a}} \prod_{i=1}^k X_i^{a_i}.$$

The guard bands ensure $a_i + b_i < n_i$, so multiplication is ordinary tensor convolution without cyclotomic wraparound. A coefficient at $2\boldsymbol{a}$ can arise only from

indices $\boldsymbol{b}, \boldsymbol{c}$ satisfying $b_i + c_i = 2a_i$ in every factor; 3-AP-freeness forces $\boldsymbol{b} = \boldsymbol{c} = \boldsymbol{a}$, so the retained coefficient is $x_{\boldsymbol{a}} y_{\boldsymbol{a}}$.

It remains to remove coefficients away from the doubled support. For the power basis $b_{i,t} = X_i^t$ and its trace dual $b_{i,t}^\vee$, define the factor projector below, where Tr_{m_i} acts on factor i and fixes every other tensor factor:

$$\Pi_i(z) = \sum_{t \in 2A_i} b_{i,t} \, \mathrm{Tr}_{m_i}(b_{i,t}^\vee z), \qquad 2A_i = \{2a : a \in A_i\}.$$

This is a genuine coefficient selector, unlike a bare trace in an odd-prime cyclotomic. If p is odd, ζ_p is a primitive p-th root, $f = \sum_{j=0}^{p-2} a_j \zeta_p^j$, and $0 \le t, j < p-1$, then

$$\mathrm{Tr}(\zeta_p^{-t} f) = p a_t - \sum_{j=0}^{p-2} a_j, \qquad \beta_t = \frac{\zeta_p^{-t} - \zeta_p}{p}, \qquad \mathrm{Tr}(\zeta_p^j \beta_t) = \delta_{j,t}.$$

Thus the trace-dual multiplier cancels the common term. More generally, $m_i b_{i,t}^\vee$ is integral for $m_i = p_i^{e_i}$. Choose a generator g_i of the cyclic group $G_i = (\mathbb{Z}/m_i\mathbb{Z})^*$ and define $\sigma_{i,a}(X_i) = X_i^{g_i^a}$, fixing the other factors. Expanding the trace expresses the scaled projector as

$$m_i \Pi_i(z) = \sum_{a=0}^{n_i - 1} \widehat{c}_{i,a} \, \sigma_{i,a}(z) \qquad (\widehat{c}_{i,a} \in \mathcal{R}_{m_i}).$$

Writing $a = r + j\ell_i$ and twisting each inner-sum constant by the inverse giant step $\sigma_{i,-j\ell_i}$, the baby-step–giant-step rearrangement of Halevi and Shoup [10] evaluates this sum with $\ell_i + \lceil n_i/\ell_i \rceil - 2$ nonidentity authenticated automorphisms and one constant multiplication for each nonzero $\widehat{c}_{i,a}$. After all factor projectors, one multiplication by $M^{-1} \bmod q$ removes their combined scale. Finally, the CRT automorphism $\tau(X_i) = X_i^{2^{-1} \bmod m_i}$ simultaneously maps every retained $X_i^{2a_i}$ back to $X_i^{a_i}$, after which each party locally selects the coefficients indexed by A.

Choosing each A_i using Behrend's progression-free construction [2] gives $|A_i| = n_i^{1-o(1)}$. For fixed k and fixed distinct primes with balanced $n_i = \Theta(N^{1/k})$, we therefore obtain $T = N^{1-o(1)}$ slots. Choosing $\ell_i = \lceil \sqrt{n_i} \rceil$ gives $1 + \sum_i (\ell_i + \lceil n_i/\ell_i \rceil - 2) = O(kN^{1/(2k)})$ authenticated automorphisms, including the final halving, while at most $1 + \sum_i n_i = O(kN^{1/k})$ constant multiplications include the combined normalization.

Table 1 summarizes the resulting resource bounds.

Against paired Construction 1, Construction 2 improves the asymptotic per-slot amortization of base-HSS calls and authenticated automorphisms, while requiring the growing number of constant multiplications shown in the last column. In paired and conversion-free typed evaluation, Construction 1 uses only a constant number of these operations, all for scalar normalization. Thus the table separates the three resources governing the packing–extraction tradeoff: occupied slots, authenticated calls, and constant-multiplication work. Beyond these

Table 1. Logical operation counts for one restricted SIMD multiplication in the common BKS/SIMD-HSS interface; additions and rerandomizations are omitted. Here N is the ambient ring dimension. Construction 1 assumes balanced factors; its conversion-free typed row applies to a type-consistent gate, with $A_{\mathrm{dir}} = \alpha - 1$ for $A \to B$ and $A_{\mathrm{dir}} = 2\beta - 1$ for $B \to A$. Construction 2 assumes fixed k and conductor primes with balanced prime-power factors.

Scheme	Slots	Base-HSS calls	Authenticated automorphisms	Constant multiplications
SIMD-HSS [11]	$\Theta(\sqrt{N})$	1	$O(\log N)$	$O(1)$
Construction 1 (paired)	$\Theta(\sqrt{N})$	2	$O(\log N)$	2
Construction 1 (typed)	$\Theta(\sqrt{N})$	1	A_{dir}	1
Construction 2	$N^{1-o(1)}$	1	$O(kN^{1/(2k)})$	$O(kN^{1/k})$

operation counts, the interval conversion of SIMD-HSS [11] requires an entropic secret with prescribed zero coefficients, whereas neither of our constructions does.

1.3 Related Work

Homomorphic secret sharing. Boyle, Gilboa, and Ishai introduced HSS and its public-key form [4], and Boyle et al. developed the general multi-input and multi-server framework [5]. Boyle, Kohl, and Scholl (BKS) construct public-key lattice HSS for RMS programs from encryption with nearly linear decryption [6]; their authenticated memory representation, KDM input-encryption procedure, and distributed-rounding multiplication form our base layer. We combine this layer with the authenticated-automorphism interface of SIMD-HSS [11], focusing on SIMD representations and authenticated linear maps.

Packed FHE and SIMD HSS. In FHE, the split-plaintext-algebra approach of Smart and Vercauteren [17] supports independent SIMD operations in one ciphertext. BKS [6] obtain CRT-slot SIMD when the plaintext polynomial splits, but bounded slot values need not have bounded coefficient representatives, causing their parameters to deteriorate with multiplicative degree. SIMD-HSS [11] instead uses interval encodings and authenticated automorphisms over a power-of-two negacyclic ring; we adopt its key-switch-then-automorphize interface. Its Type-C-to-A ring conversion requires the transformed secret to have sufficiently low degree, enforced by fixing a subset of secret coefficients to zero. Our extraction circuits remain inside the ambient ring, thereby avoiding this entropic-secret requirement. Construction 1 retains the same square-root packing order, whereas Construction 2 changes the coefficient layout to obtain $N^{1-o(1)}$ slots at the stated constant-multiplication cost.

Cyclotomic tensor and trace methods. LPR [15] develop the powerful tensor basis and record the unconjugated trace dual used here as an alternative basis

convention. Halevi and Shoup [9] use powerful-basis tensor decompositions for coefficient-to-slot transforms; Gentry et al. [8] realize subring-linear maps by multiplication with a public ring element followed by relative trace; Alperin-Sheriff and Peikert [1] exploit tensorial decompositions and structured cyclotomic bases, and homomorphically evaluate relative traces, in their bootstrapping procedure. Closest to Construction 1, Liu and Wang [13] use complementary tensor subrings, alternating encodings, dual-basis cancellation, and trace for batch FHE bootstrapping; their sequel [14] reaches $\widetilde{O}(1)$ amortized FHE multiplications per input LWE ciphertext. We adapt the first framework to authenticated BKS shares: Liu and Wang supply the alternating FHE tensor pattern, while our contribution is its bounded-integral realization as a closed HSS invariant.

Linear transforms and collision-free supports. Halevi and Shoup [10] use a baby-step–giant-step rearrangement to reduce the automorphisms required by homomorphic linear transforms; Construction 2 applies the same rearrangement to the group-algebra form of our coefficient projector. Lipmaa [12] uses progression-free sets to prevent unwanted additive collisions in pairing-based NIZKs. Block, Maji, and Nguyen [3] use a 3-free set so that the coefficient at exponent $2s_i$ in an extension-field product receives only the desired diagonal contribution. Our use follows that collision argument, adding half-length guard bands to rule out cyclotomic wraparound, Cartesian tensor supports, and integral projectors that act on authenticated HSS shares.

The remainder of the paper is organized as follows. Section 2 fixes the cyclotomic and inherited HSS interfaces. Sections 3 and 4 give the two constructions and their proofs, Section 5 reports the implementation and evaluation, and Section 6 summarizes the guarantees and open problems.

2 Preliminaries

We first fix the cyclotomic tensor algebra used by both encodings, then state the BKS multiplication and SIMD-HSS authenticated-automorphism interfaces on which our protocols rely, and finally recall the bounded-integer SIMD and progression-free-set conventions.

Notation. All logarithms are base two. For $T \geq 1$, write $[T] = \{0, \ldots, T-1\}$ and use bold symbols for vectors, indexed from zero. We write $\boldsymbol{x} \odot \boldsymbol{y}$ for componentwise multiplication and $\|z\|_\infty$ for the coefficient infinity norm in the basis under discussion. The security parameter is λ. For a cyclotomic ring $\mathcal{R}_M = \mathbb{Z}[X]/(\Phi_M(X))$, we write $N = \varphi(M)$ for its dimension, equivalently the cardinality of an integral basis. Following BKS, p is the post-round HSS modulus and $q = p\Delta$ is the ciphertext modulus; the primes p_i used later are cyclotomic conductors and are unrelated to p. Ring elements use the power basis unless stated otherwise. When $M = \prod_i m_i$ has pairwise-coprime factors, X_i denotes the root in the i-th tensor factor. An automorphism or trace carrying an index i acts on that factor and fixes the others. Tensor identities are proved integrally and then reduced modulo p or q.

2.1 Cyclotomic tensor rings

For $m \geq 2$, let

$$\mathcal{R}_m = \mathbb{Z}[X]/(\Phi_m(X)), \qquad \mathcal{R}_{q,m} = \mathcal{R}_m/q\mathcal{R}_m, \qquad n_m = \varphi(m).$$

We place the coefficient modulus first in every reduced-ring subscript, and tensor products of reduced rings are taken over $\mathbb{Z}_q$. We identify a residue class with its power-basis representative of degree less than n_m.

Lemma 1 (Cyclotomic tensor decomposition). *Let $m_1, \ldots, m_k$ be pairwise coprime and put $M = \prod_i m_i$. Then*

$$\mathcal{R}_M \cong \bigotimes_{i=1}^{k} \mathcal{R}_{m_i}.$$

Consequently, after reduction modulo any q, the monomials $\prod_i X_i^{a_i}$ with $0 \leq a_i < \varphi(m_i)$ form a $\mathbb{Z}_q$-basis of $\mathcal{R}_{q,M}$.

Proof. Map X_i to X^{M/m_i}, a primitive m_i-th root in $\mathcal{R}_M$. $\gcd_i(M/m_i) = 1$, so Bezout's identity expresses X as a product of integer powers of these images; negative powers are permitted because roots of unity are units. Hence the induced map is surjective. Both sides are free $\mathbb{Z}$-modules with $\prod_i \varphi(m_i) = \varphi(M)$ basis elements, and hence the map is an isomorphism. □

This tensor monomial representation is the *powerful basis* of Lyubashevsky, Peikert, and Regev [15]. Let ζ_m denote the image of X in $\mathcal{R}_m$, let $G_m = (\mathbb{Z}/m\mathbb{Z})^*$ act by $\sigma_a(\zeta_m) = \zeta_m^a$, and write $\mathrm{Tr}_m(z) = \sum_{a \in G_m} \sigma_a(z)$ for the field trace on $\mathbb{Q}(\zeta_m)$. The trace dual $(b_j^\vee)_j$ of a basis $(b_j)_j$ is characterized over $\mathbb{Q}(\zeta_m)$ by

$$\mathrm{Tr}_m(b_i b_j^\vee) = \delta_{i,j}. \tag{1}$$

On a tensor ring, relative traces and automorphisms act factorwise. We use the unconjugated power-basis dual recorded by LPR [15] and clear its denominator by the conductor, following the same principle used in batch bootstrapping [13].

Lemma 2 (Integral scaled dual). *Let $m = p^e$, let $b_j = X^j$ be the power basis of $\mathcal{R}_m$, and let $b_j^\vee$ be its trace dual in $\mathbb{Q}(\zeta_m)$. Then*

$$b_j^\sharp := m b_j^\vee \in \mathcal{R}_m, \qquad \mathrm{Tr}_m(b_i b_j^\sharp) = m\delta_{i,j}.$$

If $\gcd(q, m) = 1$, the normalization is therefore removed after tracing by multiplication with $m^{-1} \bmod q$. Moreover, $\|b_j^\sharp\|_\infty = 2$ when $p = 2$ and $e \geq 2$, while $\|b_j^\sharp\|_\infty \leq 2p$ when p is odd.

Proof. Put $h = p^{e-1}$. From $\Phi_m(Z) = \Phi_p(Z^h)$,

$$\Phi_m'(\zeta_m) = mu/(\zeta_m^h - 1)$$

for a root of unity u. The codifferent of $\mathcal{R}_m$ is the trace-dual module

$$\mathcal{R}_m^\vee = \{y \in \mathbb{Q}(\zeta_m) : \mathrm{Tr}_m(y\mathcal{R}_m) \subseteq \mathbb{Z}\}.$$

Because the power basis is an integral basis, each $b_j^\vee$ belongs to $\mathcal{R}_m^\vee$. By the standard formula for the trace dual of a power order [7, Theorem 3.7],

$$\mathcal{R}_m^\vee = \Phi'_m(\zeta_m)^{-1}\mathcal{R}_m.$$

Thus $b_j^\vee = r_j/\Phi'_m(\zeta_m)$ for some $r_j \in \mathcal{R}_m$. The preceding identity gives

$$\frac{m}{\Phi'_m(\zeta_m)} = u^{-1}(\zeta_m^h - 1) \in \mathcal{R}_m,$$

and hence $mb_j^\vee \in \mathcal{R}_m$. The trace identity follows by scaling (1). Applying the power-basis dual formula of Conrad [7, Theorem 3.7], if $p = 2$ and $n = m/2$, then $\Phi_m(X) = X^n + 1$ and $b_j^\vee = \zeta_m^{-j}/n$. Hence $b_0^\sharp = 2$ and $b_j^\sharp = -2X^{n-j}$ for $1 \le j < n$. For odd p, the same formula expresses $b_j^\vee$ as $c_j(\zeta_m)/\Phi'_m(\zeta_m)$, where $c_j(\zeta_m)$ is the coefficient of X^j in $\Phi_m(X)/(X - \zeta_m)$. Since $\Phi_m(X) = \sum_{r=0}^{p-1} X^{rh}$, each $c_j(\zeta_m)$ contains at most $p - 1$ root monomials, while differentiation gives $m/\Phi'_m(\zeta_m) = \zeta_m(\zeta_m^h - 1)$. Thus $b_j^\sharp$ is a sum of at most $2(p - 1)$ signed root monomials. Each such monomial reduces to the power basis with coefficients in $\{-1, 0, 1\}$, so $\|b_j^\sharp\|_\infty \le 2(p - 1) \le 2p$. □

For a free $\mathbb{Z}$-algebra $\mathcal{R}$ with integral basis $(b_i)_{i<N}$, write

$$b_i b_j = \sum_\ell c_{i,j,\ell} b_\ell.$$

Its coefficient multiplication constant is

$$\mu_\mathcal{R} = \max_\ell \sum_{i,j} |c_{i,j,\ell}|. \tag{2}$$

Lemma 3 (Power-basis growth). *Let $\mathcal{R}_M \cong \bigotimes_i \mathcal{R}_{p_i^{e_i}}$ have dimension N. Multiplication by a factor-basis monomial and a factor automorphism each have coefficient infinity-norm operator bound at most two; both bounds are one on a power-of-two factor. A tensor product of such maps acting on j odd-prime factors therefore has bound at most 2^j. Moreover, $\mu_{\mathcal{R}_M} \le N^2$.*

Proof. For $m = p^e$ and $h = p^{e-1}$, after reducing exponents modulo m, the only nontrivial power-basis reduction is $X^{(p-1)h+t} = -\sum_{r=0}^{p-2} X^{rh+t}$. Within each residue class modulo h, a monomial shift or an exponent permutation places at most one input term in this omitted block, proving the factor-two bounds; for $p = 2$ an automorphism is a signed permutation. Tensoring multiplies operator bounds. Every basis-product structure coefficient is in $\{-1, 0, 1\}$, so the sum in the definition of $\mu_{\mathcal{R}_M}$ has at most N^2 terms. □

2.2 The BKS HSS interface

Syntax. We use the public-key HSS and restricted-multiplication straight-line (RMS) syntax of Boyle, Kohl, and Scholl (BKS) [6]. An RMS program has input wires x_i and memory wires y_j, and supports the instruction types

$$\mathsf{Load}(i,j), \quad \mathsf{MAdd}(i,j,k), \quad \mathsf{IAdd}(i,j,k), \quad \mathsf{Mult}(i,j,k), \quad \mathsf{Out}(i,j).$$

Each instruction instance carries a unique identifier id, suppressed from the display. The instructions load an input, add two memory wires, add two input wires, multiply an input wire by a memory wire, and release a memory wire, respectively. A ring-valued execution is $B_{\mathcal{R}}$-bounded if every input and intermediate memory value has coefficient norm at most $B_{\mathcal{R}}$, where $\log B_{\mathcal{R}} \le \mathrm{poly}(\lambda)$.

A two-party public-key HSS has algorithms

$$\begin{aligned}(\mathsf{pk}, \mathsf{ek}_0, \mathsf{ek}_1) &\leftarrow \mathsf{Gen}(1^\lambda),\\ \mathsf{ct}_x &\leftarrow \mathsf{Enc}(\mathsf{pk}, x),\\ z_b &\leftarrow \mathsf{Eval}(b, \mathsf{ek}_b, (\mathsf{ct}_{x_i})_i, P, r).\end{aligned}$$

Correctness requires $z_0 + z_1 = P(\boldsymbol{x}) \bmod r$ for every admissible polynomial-size program, except with negligible probability; one-evaluator privacy is the standard BKS indistinguishability requirement. We use these definitions and the BKS theorem without restating their experiments.

Internally, following SIMD-HSS [11], put $\boldsymbol{s} = (1, s)^{\mathsf{T}}$ and represent an authenticated value by local vectors satisfying

$$\boldsymbol{t}_1 - \boldsymbol{t}_0 = \boldsymbol{s}z \pmod q.$$

We write this pair as $\langle \boldsymbol{s}z \rangle$ and its party-b component as $\langle \boldsymbol{s}z \rangle_b$. Because the first coordinate of $\boldsymbol{s}$ is one, the first coordinates of the two local vectors form the ordered-difference sharing of z; we call them the *message components*. At output, party 0 negates the centered lift of its decoded message component, restoring the additive BKS convention before both parties reduce modulo the requested output modulus.

Base multiplication. We use the following logical BKS interface:

$$\langle \boldsymbol{s}xy \rangle \leftarrow \mathsf{Mult}(\mathsf{ct}_x, \langle \boldsymbol{s}y \rangle). \tag{3}$$

In BKS, ct_x is generated by the KDM oracle $\mathsf{PKE.OKDM}(\mathsf{pk}, x)$, and party b evaluates Equation (3) from ek_b and $\langle \boldsymbol{s}y \rangle_b$ using distributed decryption. The BKS analysis [6] shows that, for bounded inputs and parameters satisfying its rounding and lifting conditions, this call is correct except with negligible probability and that composing it with additions gives HSS for admissible RMS programs; Lemma 3 of their work proves that the oracle outputs are indistinguishable from zero encryptions for polynomially many queries. We write ReRand for the common PRF rerandomization used by BKS and SIMD-HSS. Each internal circuit

node extends its instruction identifier id with its branch and stage, share component, and purpose; the resulting labels are unique, common masks preserve the ordered difference, and the random-function hybrid gives the uniform marginals required by the cited BKS lemmas.

The original BKS rounding statement is written for the integer and power-of-two cyclotomic rings. The following observation supplies the needed correctness extension for our tensor rings.

Lemma 4 (BKS portability). *Let $\mathcal{R}$ be a commutative free $\mathbb{Z}$-algebra with a fixed integral basis of size N and multiplication constant* (2)*. Then*

$$\|xy\|_\infty \le \mu_{\mathcal{R}} \|x\|_\infty \|y\|_\infty.$$

Moreover, the BKS rounding and lifting lemmas apply coefficientwise to masked shares over $\mathcal{R}$. Consequently, if N and the number of calls are polynomially bounded, while $\mu_{\mathcal{R}}$ and all circuit, secret, and error bounds have polynomial bit length and the stated tails are negligible, the original BKS correctness conditions are met by polynomial-bit moduli and give negligible total error.

Proof. The displayed norm inequality follows by expanding in the basis. This bounds the perturbation entering each distributed decryption: for error e and represented multiplicand y, the only ring-dependent term is ey, whose norm is at most $\mu_{\mathcal{R}}\|e\|_\infty\|y\|_\infty$. The remaining BKS rounding and centered-lifting arguments [6] are coefficientwise and otherwise independent of the multiplication law; applying them to the N coordinates and taking a polynomial union bound proves the claim. □

We call such a parameter choice *BKS-compatible*; this includes the original BKS hypotheses for every call and terminal reconstruction, as well as the coprimality conditions used by our cyclotomic normalizations.

Authenticated automorphisms. SIMD-HSS [11] shows how to evaluate a ring automorphism σ on authenticated shares by switching from s to $\sigma^{-1}(s)$ and then applying σ locally:

$$\mathsf{Aut}_\sigma(\langle \boldsymbol{s} z\rangle) = \sigma\Big(\mathsf{KeySwitch}_{s\to\sigma^{-1}(s)}(\langle \boldsymbol{s} z\rangle)\Big) = \langle \boldsymbol{s}\sigma(z)\rangle. \tag{4}$$

As in SIMD-HSS, $\mathsf{KeySwitch}$ ends with a common PRF shift, so every Aut_σ call returns a rerandomized share; applying σ to that shift preserves its uniformity in the random-function hybrid. The BKS switching argument gives correctness; SIMD-HSS gives security under circular RLWE. For an automorphism set Σ, we use the function family $\mathcal{F}_\Sigma = \{0\}\cup\{x \mapsto \Delta\sigma^{-1}(x) : \sigma \in \Sigma\}$. The circular-RLWE definition of SIMD-HSS already covers polynomially many samples associated with any sequence of functions from $\mathcal{F}_\Sigma$ under one secret.

Proposition 1 (Inherited HSS guarantees). *Let C be an algebraically correct RMS circuit using BKS multiplication, authenticated automorphisms from Σ, and public linear operations. With BKS-compatible parameters, C is correct. For the BKS RLWE instantiation, independently masked auxiliary setup shares and circular RLWE for $\mathcal{F}_\Sigma$ imply one-evaluator privacy.*

Proof. Correctness follows by induction from BKS multiplication and authenticated key switching. For privacy, circular RLWE replaces the switching keys by uniform samples and supplies ordinary RLWE through the zero function; the BKS KDM-oracle hybrid [6] then replaces the inputs by zero encryptions. The independent setup masks complete the simulation. □

We report logical calls and local ring operations separately. With BKS ciphertext dimension two, one base multiplication performs four ciphertext–share ring multiplications and two distributed roundings per evaluator. One authenticated automorphism performs two such multiplications and one rounding. We call multiplication of an authenticated share by a public ring constant a *constant multiplication*. Because an authenticated share has two components, one constant multiplication performs two componentwise constant multiplications per evaluator.

2.3 SIMD encodings and progression-free sets

We follow the bounded-integer SIMD convention of SIMD-HSS [11]. An integral T-SIMD encoding is an injective $\mathbb{Z}$-linear map $\mathcal{E} : \mathbb{Z}^T \to \mathcal{R}_M$ together with a multiplication procedure that maps encodings of $\boldsymbol{x}$ and $\boldsymbol{y}$ to an encoding of $\boldsymbol{x} \odot \boldsymbol{y}$. Cryptographic evaluation uses its reductions modulo p and q; admissibility ensures that these represent the intended bounded integers.

We call a SIMD program–input pair B-admissible if every logical slot value on every input and memory wire has magnitude at most B, with $\log B \le \mathrm{poly}(\lambda)$. The construction-specific parameter arguments translate this logical bound into the ring-coordinate bounds needed by the BKS interface.

A set $A \subseteq \mathbb{Z}$ is *3-AP-free* if

$$a + c = 2b, \quad a, b, c \in A \qquad \Longrightarrow \qquad a = b = c.$$

Behrend's construction [2] implies that, for every sufficiently large L, there is a 3-AP-free $A \subseteq \{0, \dots, L-1\}$ with

$$|A| \ge L \exp(-O(\sqrt{\log L})) = L^{1-o(1)}. \tag{5}$$

3 Construction 1: Alternating Tensor Encodings

Our first construction realizes within BKS HSS the scalar-factor specialization of Liu and Wang's batch-bootstrapping framework [13]: it alternates between complementary tensor subrings, clears trace-dual denominators, and uses relative traces to remove cross terms. By default, each memory vector is kept in two authenticated subring encodings, and two symmetric connector-and-trace paths update these copies in parallel. This paired invariant makes arbitrary memory additions immediate and uses only two scalar normalizations per multiplication. On type-consistent regions, one may retain a single encoding and evaluate only the direction needed at each multiplication. The resulting construction packs

$\Theta(\sqrt{N})$ slots for balanced factors with logarithmic authenticated-automorphism overhead and a constant-size scalar extraction layer. Because every path remains in the ambient ring, neither mode imposes a prescribed-zero pattern on the RLWE secret.

3.1 Algebraic encoding

Let $M_1 = 2^\alpha$ and $M_2 = 3^\beta$ for $\alpha \geq 2$ and $\beta \geq 1$, and put

$$N_1 = \varphi(M_1), \qquad N_2 = \varphi(M_2), \qquad N = N_1 N_2, \qquad T = \min(N_1, N_2).$$

Let q be coprime to 6 and identify $\mathcal{R}_{q,M_1M_2}$ with $\mathcal{R}_{q,M_1} \otimes \mathcal{R}_{q,M_2}$ by Lemma 1. Let $(v_i)_{i<N_1}$ and $(w_i)_{i<N_2}$ be the respective power bases, with trace-dual bases $(v_i^\vee)_{i<N_1}$ and $(w_i^\vee)_{i<N_2}$. Define the integral scaled duals

$$v_i^\sharp = M_1 v_i^\vee \in \mathcal{R}_{M_1}, \qquad w_i^\sharp = M_2 w_i^\vee \in \mathcal{R}_{M_2}$$

as in Lemma 2. These conductor-cleared dual connectors instantiate Liu and Wang's construction [13] with LPR's alternative unconjugated power-basis and trace-dual convention [15].

For $\boldsymbol{x} \in \mathbb{Z}^T$, define two memory encodings and two input encodings in $\mathcal{R}_{M_1M_2}$:

$$\mathcal{E}_A(\boldsymbol{x}) = \sum_{i<T} x_i (v_i \otimes 1), \qquad \mathcal{E}_B(\boldsymbol{x}) = \sum_{i<T} x_i (1 \otimes w_i), \tag{6}$$

$$\mathcal{C}_{A\to B}(\boldsymbol{x}) = \sum_{i<T} x_i (v_i^\sharp \otimes w_i), \qquad \mathcal{C}_{B\to A}(\boldsymbol{x}) = \sum_{i<T} x_i (v_i \otimes w_i^\sharp). \tag{7}$$

Memory wires are maintained in both encodings in (6), while each input is encrypted in both connector encodings in (7). This duplication closes the representation under arbitrary RMS programs, including additions of wires at different multiplicative depths. It is a protocol choice rather than an algebraic necessity: Section 3.3 formalizes a single-encoding fast path for type-consistent regions.

Lemma 5 (Alternating multiplication). *For all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Z}^T$, the following identities hold over $\mathcal{R}_{M_1M_2}$ and hence after reduction modulo p or q:*

$$\mathrm{Tr}_{\mathcal{R}_{M_1M_2}/\mathcal{R}_{M_2}}\big(\mathcal{C}_{A\to B}(\boldsymbol{x})\mathcal{E}_A(\boldsymbol{y})\big) = M_1 \mathcal{E}_B(\boldsymbol{x} \odot \boldsymbol{y}),$$
$$\mathrm{Tr}_{\mathcal{R}_{M_1M_2}/\mathcal{R}_{M_1}}\big(\mathcal{C}_{B\to A}(\boldsymbol{x})\mathcal{E}_B(\boldsymbol{y})\big) = M_2 \mathcal{E}_A(\boldsymbol{x} \odot \boldsymbol{y}).$$

Proof. Expanding the first expression and applying trace duality gives

$$\sum_{i,j<T} x_i y_j \,\mathrm{Tr}_{M_1}(v_i^\sharp v_j) w_i = M_1 \sum_{i<T} x_i y_i w_i.$$

The second identity is identical with the two tensor factors exchanged. □

The decoders select the corresponding power-basis coordinates; equivalently, they apply the modular trace-dual functionals. Equation (1) therefore gives $\mathcal{D}_A(\mathcal{E}_A(\boldsymbol{x})) = \boldsymbol{x}$ and $\mathcal{D}_B(\mathcal{E}_B(\boldsymbol{x})) = \boldsymbol{x}$.

3.2 Homomorphic construction and correctness

The algebraic lemma supplies two complementary relative-trace paths: paired evaluation uses both, whereas typed evaluation uses the direction selected by its input type. We evaluate each relative trace as a tower of subgroup sums, following the homomorphic trace method of Alperin-Sheriff and Peikert [1], with authenticated automorphisms in place of ciphertext automorphisms. Each automorphism produces a missing coset and supplies the fresh mask needed by the ensuing local sum.

To realize factor automorphisms in the full ring, for $a \in G_{M_1}$ let $\widetilde{a} \in G_{M_1 M_2}$ be the CRT lift satisfying $\widetilde{a} \equiv a \pmod{M_1}$ and $\widetilde{a} \equiv 1 \pmod{M_2}$; for $a \in G_{M_2}$ use the analogous lift that is one modulo M_1. We suppress the tildes below. Thus every displayed automorphism acts on its indicated tensor factor and fixes the other one.

The group G_{2^α} is isomorphic to $C_2 \times C_{2^{\alpha-2}}$. Let ρ_1 and γ_1 generate these factors, with the latter factor trivial when $\alpha = 2$. On an authenticated share u, compute

$$u \leftarrow u + \mathsf{Aut}_{\rho_1}(u), \qquad u \leftarrow u + \mathsf{Aut}_{\gamma_1^{2^j}}(u) \quad (0 \le j < \alpha - 2). \tag{8}$$

The result is the trace over the first factor and uses $\alpha - 1$ authenticated automorphisms.

Similarly, $G_{3^\beta} \cong C_2 \times C_{3^{\beta-1}}$. For generators ρ_2 and γ_2, with the latter factor trivial when $\beta = 1$, compute

$$\begin{aligned} &u \leftarrow u + \mathsf{Aut}_{\rho_2}(u), \\ &u \leftarrow u + \mathsf{Aut}_{\gamma_2^{3^j}}(u) + \mathsf{Aut}_{\gamma_2^{2\cdot 3^j}}(u), \quad 0 \le j < \beta - 1. \end{aligned} \tag{9}$$

This gives the trace over the second factor using $2\beta - 1$ authenticated automorphisms. When $\alpha = 2$ or $\beta = 1$, respectively, the corresponding second update is absent; throughout, an empty product is the identity group-algebra operator.

Lemma 6 (Trace evaluation). *Equations* (8) *and* (9) *evaluate the respective relative traces on authenticated shares.*

Proof. The first update sums over the C_2 factor. For the remaining cyclic factor, the successive updates multiply the group-algebra operator by

$$\prod_{j=0}^{\alpha-3} (1 + \gamma_1^{2^j}) = \sum_{r=0}^{2^{\alpha-2}-1} \gamma_1^r$$

or, respectively, by

$$\prod_{j=0}^{\beta-2} (1 + \gamma_2^{3^j} + \gamma_2^{2\cdot 3^j}) = \sum_{r=0}^{3^{\beta-1}-1} \gamma_2^r.$$

These are the binary and ternary digit expansions of every exponent exactly once. Multiplying by the initial C_2 sum therefore yields the sum over the whole Galois

group, i.e., the relative trace. Equation (4) lifts the identity to authenticated shares. Each update contains an authenticated-automorphism output with a fresh common PRF shift, so the updated share is already rerandomized.

Write HTr_{M_1} and HTr_{M_2} for the authenticated trace procedures in (8) and (9), respectively. Each right-hand side in an in-place update is evaluated from the same pre-update value of u. The normalizing scalars M_1^{-1} and M_2^{-1} are invertible modulo q and therefore preserve the fresh-share invariant.

The construction maintains the paired invariant

$$\big(\langle \boldsymbol{s}\mathcal{E}_A(\boldsymbol{y})\rangle, \langle \boldsymbol{s}\mathcal{E}_B(\boldsymbol{y})\rangle\big)$$

for every memory wire $\boldsymbol{y}$. The two connector-and-trace paths produce the opposite encodings in parallel, avoiding any separate conversion between pure A and B encodings. We instantiate this invariant using the BKS multiplication [6] and the authenticated automorphisms of SIMD-HSS [11].

Gen. Generate the public and evaluation keys of the base HSS over $\mathcal{R}_{q,M_1M_2}$ and publish the automorphism keys used in (8) and (9). Let $\mathbf{1} = (1, \ldots, 1)$. Independently of the base-HSS evaluation sharing, and independently for each $z \in \{\mathcal{E}_A(\mathbf{1}), \mathcal{E}_B(\mathbf{1})\}$, sample $\boldsymbol{t}_0^{(z)} \leftarrow \mathcal{R}_{q,M_1M_2}^2$ uniformly, set $\boldsymbol{t}_1^{(z)} = \boldsymbol{t}_0^{(z)} + \boldsymbol{s}z$, and give $\boldsymbol{t}_b^{(z)}$ to party b.

Enc. For every SIMD input $\boldsymbol{x}$, output two base-HSS ciphertexts

$$\mathsf{ct}_{A\to B,\boldsymbol{x}} \leftarrow \mathsf{Enc}(\mathcal{C}_{A\to B}(\boldsymbol{x})), \qquad \mathsf{ct}_{B\to A,\boldsymbol{x}} \leftarrow \mathsf{Enc}(\mathcal{C}_{B\to A}(\boldsymbol{x})).$$

Eval. Associate every memory wire carrying $\boldsymbol{y}$ with the pair

$$\big(\langle \boldsymbol{s}\mathcal{E}_A(\boldsymbol{y})\rangle, \langle \boldsymbol{s}\mathcal{E}_B(\boldsymbol{y})\rangle\big).$$

Memory addition applies ReRand independently to each componentwise sum. Addition of two input wires is componentwise addition of their two connector ciphertext pairs. To load $\boldsymbol{x}$, or to multiply a memory value $\boldsymbol{y}$ by the input $\boldsymbol{x}$, invoke the base multiplication twice and then the appropriate relative trace:

$$\begin{aligned}
\langle \boldsymbol{s}\mathcal{E}_B(\boldsymbol{x}\odot\boldsymbol{y})\rangle &\leftarrow M_1^{-1}\cdot \mathsf{HTr}_{M_1}\left(\mathsf{Mult}(\mathsf{ct}_{A\to B,\boldsymbol{x}}, \langle \boldsymbol{s}\mathcal{E}_A(\boldsymbol{y})\rangle)\right),\\
\langle \boldsymbol{s}\mathcal{E}_A(\boldsymbol{x}\odot\boldsymbol{y})\rangle &\leftarrow M_2^{-1}\cdot \mathsf{HTr}_{M_2}\left(\mathsf{Mult}(\mathsf{ct}_{B\to A,\boldsymbol{x}}, \langle \boldsymbol{s}\mathcal{E}_B(\boldsymbol{y})\rangle)\right).
\end{aligned}$$

For a load instruction, take $\boldsymbol{y} = \mathbf{1}$. For an output instruction with modulus r_{out}, both parties apply the same fixed public decoder, say $\mathcal{D}_A$, to their own message component and take its coefficientwise centered representative modulo q. Party 0 negates the result, and both parties reduce locally modulo r_{out}.

Typed alternating fast path. Write $\overline{A} = B$ and $\overline{B} = A$. Call a region of an RMS program *type-consistent* if each memory wire w is assigned a type $\chi(w) \in \{A, B\}$, every memory addition has two inputs and an output of the same type, and every restricted multiplication satisfies $\chi(w_{\mathrm{out}}) = \overline{\chi(w_{\mathrm{mem}})}$. A load may receive either output type by using the authenticated unit of the opposite type, and an output wire of type τ is decoded with $\mathcal{D}_\tau$. Input addition imposes no further constraint because Enc supplies both connector ciphertexts.

Proposition 2 (Typed alternating evaluation). *On a type-consistent region, Construction 1 may retain only $\langle \boldsymbol{s}\mathcal{E}_{\chi(w)}(\boldsymbol{y})\rangle$ for each memory wire w. A load or multiplication from type A to type B uses one base-HSS multiplication, $\alpha - 1$ authenticated automorphisms, and one scalar normalization; the corresponding operation from B to A uses one base-HSS multiplication, $2\beta - 1$ authenticated automorphisms, and one scalar normalization. This evaluation is correct and satisfies one-evaluator privacy under the same parameter and security hypotheses as the paired construction.*

Proof. For an A-typed source, evaluate $M_1^{-1}\mathsf{HTr}_{M_1}(\mathsf{Mult}(\mathsf{ct}_{A\to B,\boldsymbol{x}}, \langle \boldsymbol{s}\mathcal{E}_A(\boldsymbol{y})\rangle))$; Lemma 5 gives an authenticated B-encoding of the componentwise product, while (8) uses $\alpha - 1$ automorphisms. The second assignment in Eval gives the symmetric B-to-A statement and uses $2\beta - 1$ automorphisms by (9). The same identities with $\boldsymbol{y} = \mathbf{1}$ prove the load rule, same-type additions preserve the invariant by linearity, and $\mathcal{D}_A, \mathcal{D}_B$ decode the respective encodings. Induction proves correctness on the region. The fast path uses a subset of the connector ciphertexts, authenticated automorphisms, and public maps already covered by Proposition 1, so the paired construction's correctness and privacy hypotheses suffice. □

The fast path can be selected locally rather than for an entire program. If a typed wire must meet an incompatible type, the missing encoding can be materialized without a base-HSS multiplication. For $u_A = \langle \boldsymbol{s}\mathcal{E}_A(\boldsymbol{y})\rangle$ and $u_B = \langle \boldsymbol{s}\mathcal{E}_B(\boldsymbol{y})\rangle$, define

$$\mathsf{Conv}_{A\to B}(u_A) = M_1^{-1} \cdot \mathsf{HTr}_{M_1}\left(\mathsf{ReRand}\big(\mathcal{C}_{A\to B}(\mathbf{1}) \cdot u_A\big)\right),$$
$$\mathsf{Conv}_{B\to A}(u_B) = M_2^{-1} \cdot \mathsf{HTr}_{M_2}\left(\mathsf{ReRand}\big(\mathcal{C}_{B\to A}(\mathbf{1}) \cdot u_B\big)\right).$$

Multiplication by the all-ones connector is a constant multiplication. Because this connector need not be invertible modulo q, its product is rerandomized before the first authenticated automorphism. Lemma 5 with $\boldsymbol{x} = \mathbf{1}$ shows that the conversions return the opposite encoding. Retaining the original share promotes a typed wire back to the paired invariant; alternatively, one conversion reconciles a mismatched addition. The two conversions use, respectively, $\alpha - 1$ and $2\beta - 1$ authenticated automorphisms; each uses no base-HSS multiplication and two constant multiplications, one by the all-ones connector and one for scalar normalization. They use the same authenticated automorphisms and public maps as the paired circuit.

Parameter admissibility. Suppose the logical input and memory slots have magnitudes $X, Y \le B$. Lemma 2 bounds the two connector norms by $2X$ and $6X$. Writing $\mu = \mu_{\mathcal{R}_{M_1 M_2}}$, the corresponding products have norms at most $2\mu XY$ and $6\mu XY$, and the two trace circuits have norms at most $2N_1\mu XY$ and $12N_2\mu XY$. Normalization restores an encoding of the logical value, while every factor-automorphic secret image has norm at most $2B_s$ whenever $\|s\|_\infty \le B_s$; this estimate uses no degree or support property of s. These are bounds on represented ring states; the party-local masks remain uniform modulo q. Moreover, both $\mathcal{D}_A$ and $\mathcal{D}_B$ are coefficient selections with operator norm one, so a B-admissible terminal value remains bounded by B. Typed evaluation is a subcircuit of paired evaluation, and the conversion formulas instantiate the same connector-and-trace identities with $\boldsymbol{x} = \mathbf{1}$, so the stated polynomial bounds also cover the optional fast path. Since $\mu \le N^2$, Lemma 4 gives BKS-compatible polynomial-bit parameters for families with $N \le \mathrm{poly}(\lambda)$, polynomially many calls, polynomial-bit secret and error bounds, and negligible error tails.

Theorem 1 (Correctness of Construction 1). *For every B-admissible RMS program–input pair $(P, \boldsymbol{x})$ and every output modulus $r_{\mathsf{out}} \ge 2$, Construction 1 with BKS-compatible parameters produces additive shares of the T coordinatewise evaluations of P modulo r_{out}, except with negligible probability.*

Proof. Induct on the program instructions. The setup shares encode $\mathbf{1}$ in both types. For a load or multiplication, the BKS interface gives the two connector–memory products; Lemmas 6 and 5, followed by the inverse-conductor normalizations, give the two encodings of $\boldsymbol{x} \odot \boldsymbol{y}$. Additions preserve the invariant by linearity, and the fixed decoder $\mathcal{D}_A$ recovers the output. Proposition 1 supplies the probabilistic BKS correctness step. □

Theorem 2 (Security of Construction 1). *Fix any secret distribution D_{sk} supported on $\{s \in \mathcal{R}_{M_1 M_2} : \|s\|_\infty \le B_s\}$. Assume circular RLWE for D_{sk} and the function family $\mathcal{F}_{\Sigma_1}$ of Section 2, where Σ_1 is the set of trace automorphisms used by Construction 1. Then Construction 1 satisfies one-evaluator HSS privacy; in particular, no coordinate of s is required to vanish identically under D_{sk}.*

Proof. Either evaluator's pair of setup shares for the public unit encodings is uniform and independent of the remaining setup view. The authenticated automorphisms inside the trace circuits are covered by circular RLWE for $\mathcal{F}_{\Sigma_1}$; the connector, normalization, addition, and decoding maps are public linear operations. The circuit stays in $\mathcal{R}_{M_1 M_2}$, and neither its algebraic identities nor its norm analysis requires a secret coordinate to vanish. Proposition 1 therefore applies with the trace-automorphism set. □

3.3 Complexity

We distinguish logical HSS instructions from the local ring operations that implement them.

Paired operation counts. Construction 1 packs $T = \min(N_1, N_2)$ slots. In the default paired mode, each SIMD load or multiplication uses two base-HSS multiplications and

$$A_{\text{pair}} = (\alpha - 1) + (2\beta - 1) = \alpha + 2\beta - 2 \tag{10}$$

authenticated automorphisms, together with two constant multiplications by the scalars M_1^{-1} and M_2^{-1}. If N_1 and N_2 are balanced, then $T = \Theta(\sqrt{N})$ and the automorphism overhead is $O(\log N)$. Thus the logical circuit makes two base-HSS calls and A_{pair} authenticated-automorphism calls. More concretely, each party performs

$$4 \cdot 2 + 2A_{\text{pair}} = 2\alpha + 4\beta + 4$$

ciphertext–share ring multiplications and $2 \cdot 2 + A_{\text{pair}} = \alpha + 2\beta + 2$ distributed roundings. The two normalizations act on both authenticated components and hence require four componentwise constant multiplications per party. These counts exclude local additions and PRF expansions. They follow from Lemma 6 and the local BKS costs in Section 2. Since $N_1 = 2^{\alpha-1}$ and $N_2 = 2 \cdot 3^{\beta-1}$, balanced factors satisfy $\alpha, \beta = O(\log N)$.

Fast-path operation counts. Let $g_{A\to B}$ and $g_{B\to A}$ be the numbers of typed loads or multiplications in the two directions. Proposition 2 gives $g_{A\to B} + g_{B\to A}$ base-HSS calls, $(\alpha-1)g_{A\to B} + (2\beta-1)g_{B\to A}$ authenticated-automorphism calls, and one constant multiplication for scalar normalization per gate. Per evaluator, a single A-to-B gate performs $2\alpha + 2$ ciphertext–share ring multiplications, $\alpha + 1$ distributed roundings, and two componentwise constant multiplications; a B-to-A gate performs $4\beta+2$, $2\beta+1$, and two, respectively. Hence a complete two-gate cycle uses two base-HSS calls and $\alpha+2\beta-2$ automorphisms—exactly the online counts of one paired gate. Returning an A-typed wire to the paired invariant costs no base-HSS call, $\alpha - 1$ automorphisms, and two constant multiplications; the reverse promotion uses $2\beta - 1$ automorphisms and two constant multiplications. Typed evaluation is therefore most effective on type-consistent regions long enough to amortize these boundary conversions; same-type additions inside such a region require only local addition and rerandomization, which are excluded from the counts. Generic encryption still emits both connector ciphertexts, and setup retains both trace-key families; the optimization reduces online invocations and authenticated memory, not the public input or key set.

The counts above do not depend on a ring representation. An indicative arithmetic interpretation is given after Theorem 5; Section 5 reports the concrete operation classes separately.

Construction 1 therefore offers two operating modes: paired evaluation provides closure under arbitrary RMS addition, while typed evaluation reduces online work on compatible regions.

4 Construction 2: Progression-Free Tensor Encoding

The second construction replaces the paired representation by a single coefficient encoding. Its algebraic pipeline has three stages: guard bands make multiplica-

tion an ordinary tensor convolution, progression-free supports isolate diagonal products at doubled slot positions, and factorwise trace-dual projectors remove all remaining coefficients. The diagonal-collision principle follows the 3-free-set packing of Block, Maji, and Nguyen [3] for extension-field multiplication; adapting it to HSS requires guard bands, Cartesian support, conductor-cleared authenticated projectors, and one simultaneous halving automorphism. For balanced factors, this yields near-linear packing with a factor-projector circuit using $O(kN^{1/k})$ constant multiplications. All projectors and the halving map act within the ambient ring, so the construction does not require a prescribed-zero secret pattern.

4.1 Encoding and coefficient projection

Fix a constant $k \geq 2$. Let $m_i = p_i^{e_i}$ for distinct odd primes p_i, and put

$$M = \prod_{i=1}^{k} m_i, \qquad n_i = \varphi(m_i), \qquad N = \prod_{i=1}^{k} n_i.$$

Let p and q be odd, let q be coprime to M, and use the tensor representation

$$\mathcal{R}_{q,M} \cong \bigotimes_{i=1}^{k} \mathcal{R}_{q,m_i}.$$

For each i, set $L_i = \lfloor (n_i+1)/2 \rfloor$ and choose a 3-AP-free set $A_i \subseteq \{0, \ldots, L_i - 1\}$. Write $A = \prod_i A_i$ and $T = |A| = \prod_i |A_i|$. We index a SIMD vector by $\boldsymbol{a} = (a_1, \ldots, a_k) \in A$ and define

$$\mathcal{E}(\boldsymbol{x}) = \sum_{\boldsymbol{a} \in A} x_{\boldsymbol{a}} \prod_{i=1}^{k} X_i^{a_i}. \tag{11}$$

The guard band implies $a_i + b_i < n_i$ for all $a_i, b_i \in A_i$; hence multiplying two encoded monomials invokes no cyclotomic reduction.

Using LPR's trace-dual basis [15], we express factorwise coordinate functionals and realize their selected sum by a bounded circuit on authenticated shares. For the i-th factor, let $b_{i,t} = X_i^t$ for $0 \leq t < n_i$, let $b_{i,t}^{\vee}$ be its trace-dual element over $\mathbb{Q}(\zeta_{m_i})$, and reduce this element modulo q when defining

$$\Pi_i(z) = \sum_{t \in 2A_i} b_{i,t} \operatorname{Tr}_{m_i}(b_{i,t}^{\vee} z), \qquad 2A_i = \{2a : a \in A_i\}. \tag{12}$$

Lemma 7 (Coefficient projector). *For $z = \sum_{t=0}^{n_i-1} b_{i,t} z_t$, where z_t lies in the tensor product of the other factors,*

$$\Pi_i(z) = \sum_{t \in 2A_i} b_{i,t} z_t.$$

Moreover, for $G_i = (\mathbb{Z}/m_i\mathbb{Z})^*$,

$$\Pi_i(z) = \sum_{\sigma \in G_i} c_{i,\sigma}\sigma(z), \qquad c_{i,\sigma} = \sum_{t \in 2A_i} b_{i,t}\sigma(b_{i,t}^{\vee}). \tag{13}$$

Proof. Trace duality gives $\mathrm{Tr}_{m_i}(b_{i,t}^{\vee} z) = z_t$ after reduction modulo q; this reduction is defined because q is coprime to m_i. Expanding the trace as a sum over G_i and interchanging the sums gives (13). Projectors on distinct tensor factors commute, because each is the identity on every other factor. □

Odd-prime coefficient extraction. The projector cannot be replaced by a bare trace in an odd-prime cyclotomic. In particular, the coefficient-clearing trace available in the power-of-two setting of SIMD-HSS [11] does not itself provide an odd-prime coefficient selector. Indeed, if ζ_p is a primitive p-th root and $f = \sum_{j=0}^{p-2} a_j \zeta_p^j$, then for $0 \le t < p-1$,

$$\mathrm{Tr}(\zeta_p^{-t} f) = pa_t - \sum_{j=0}^{p-2} a_j, \qquad \beta_t = \frac{\zeta_p^{-t} - \zeta_p}{p}, \qquad \mathrm{Tr}(\zeta_p^j \beta_t) = \delta_{j,t}.$$

The first identity follows from $\mathrm{Tr}(\zeta_p^d) = p-1$ when $p \mid d$ and -1 otherwise. The difference

$$\frac{\mathrm{Tr}(\zeta_p^{j-t}) - \mathrm{Tr}(\zeta_p^{j+1})}{p}$$

equals one precisely when $j = t$, proving the second identity. Thus the bare trace mixes every coefficient through the common term $-\sum_j a_j$, whereas the dual multiplier β_t cancels that term and extracts a_t. Division by p is valid because q is coprime to the conductor.

After the factor projectors, let $d \in G_M$ be the unique CRT solution to $d \equiv 2^{-1} \pmod{m_i}$ for every i, and define the tensor automorphism $\tau = \sigma_d$ by

$$\tau : \quad X_i \longmapsto X_i^{2^{-1} \bmod m_i} \qquad (1 \le i \le k). \tag{14}$$

Because the factors are coprime, τ is a single automorphism of $\mathcal{R}_{q,M}$, rather than one automorphism per factor.

Lemma 8 (Diagonal coefficients). *For all* $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Z}^T$,

$$\tau \circ (\Pi_k \circ \cdots \circ \Pi_1)\big(\mathcal{E}(\boldsymbol{x})\mathcal{E}(\boldsymbol{y})\big) = \mathcal{E}(\boldsymbol{x} \odot \boldsymbol{y})$$

holds in $\mathcal{R}_{q,M}$.

Proof. The guard bands make the product an ordinary multivariate convolution. A term indexed by $\boldsymbol{b}, \boldsymbol{c} \in A$ survives all projectors only if $b_i + c_i = 2a_i$ for some $a_i \in A_i$ in every factor. Since each A_i is 3-AP-free, this forces $\boldsymbol{b} = \boldsymbol{c} = \boldsymbol{a}$. Finally, τ maps every $X_i^{2a_i}$ to $X_i^{a_i}$ simultaneously. □

For example, take $m_1 = 5$, $m_2 = 7$, and $A_1 = A_2 = \{0, 1\}$. The four slots occupy $1, X_1, X_2, X_1X_2$. Their product has no wraparound, and the projectors retain only exponents in $\{0, 2\}$ on each axis. Every retained exponent $(2a_1, 2a_2)$ comes from the square of the slot (a_1, a_2); the other terms have an odd coordinate and are removed. The single automorphism τ then maps $(2a_1, 2a_2)$ back to (a_1, a_2), exhibiting the three stages of the construction.

4.2 Authenticated BSGS coefficient projection

Equation (13) expresses coefficient selection as a linear combination of Galois conjugates. Conductor clearing supplies bounded integral group-algebra coefficients, and the baby-step–giant-step rearrangement of Halevi and Shoup [10] reuses a small set of automorphic images. This gives an authenticated implementation of each conductor-scaled factor projector; a single normalization after all factors removes their combined scale modulo q.

Let $b^{\sharp}_{i,t} = m_i b^{\vee}_{i,t} \in \mathcal{R}_{m_i}$ as in Lemma 2. For $\sigma \in G_i$, put

$$\widehat{c}_{i,\sigma} = \sum_{t \in 2A_i} b_{i,t}\sigma(b^{\sharp}_{i,t}). \tag{15}$$

Then every $\widehat{c}_{i,\sigma}$ is integral and

$$m_i \Pi_i(z) = \sum_{\sigma \in G_i} \widehat{c}_{i,\sigma}\sigma(z). \tag{16}$$

Lemma 9 (Size of the integral projector constants). *For $m_i = p_i^{e_i}$ and the power basis above,*

$$\|b^{\sharp}_{i,t}\|_\infty \le 2p_i, \qquad \|\widehat{c}_{i,\sigma}\|_\infty \le 8p_i|A_i|, \qquad \|\sigma'(\widehat{c}_{i,\sigma})\|_\infty \le 16p_i|A_i|$$

for every $\sigma' \in G_i$.

Proof. The first inequality is Lemma 2. By Lemma 3, the factor automorphism and the subsequent multiplication by $b_{i,t}$ each expand the norm by at most two. Thus every summand in (15) has norm at most $8p_i$; summing the $|2A_i| = |A_i|$ terms proves the second inequality. Applying the same factor-two automorphism bound to $\widehat{c}_{i,\sigma}$ proves the third; BSGS uses it with $\sigma' = \sigma_{i,-j\ell_i}$. □

Because m_i is an odd prime power, G_i is cyclic. Fix a generator g_i, write $\sigma_{i,a}$ for the factor automorphism $X_i \mapsto X_i^{g_i^a}$, with all subscripts a interpreted modulo n_i, choose a baby-step size $1 \le \ell_i \le n_i$, and set $h_i = \lceil n_i/\ell_i \rceil$. Write $\widehat{c}_{i,a} = \widehat{c}_{i,\sigma_{i,a}}$ for $0 \le a < n_i$. On an authenticated share u, first compute

$$U_r = \mathsf{Aut}_{\sigma_{i,r}}(u), \qquad 0 \le r < \ell_i,$$

with $U_0 = u$ evaluated locally. For $0 \le j < h_i$, form

$$V_j = \sum_{\substack{0 \le r < \ell_i,\ r+j\ell_i < n_i \\ \widehat{c}_{i,r+j\ell_i} \neq 0}} \sigma_{i,-j\ell_i}(\widehat{c}_{i,r+j\ell_i})\, U_r, \tag{17}$$

where an automorphism applied to a public constant is computed offline. Because the projector constants need not be invertible modulo q, rerandomize V_j before a giant-step automorphism whenever $j > 0$. Finally, output

$$\widehat{\mathsf{Project}}_i(u) = \mathsf{ReRand}\left(\sum_{j=0}^{h_i-1} \mathsf{Aut}_{\sigma_{i,j\ell_i}}(V_j)\right). \tag{18}$$

The baby step U_0 and the $j = 0$ giant step are local identity operations.

Lemma 10 (Scaled BSGS projector). *Equation* (18) *returns an authenticated share of $m_i\Pi_i(z)$. Its authenticated-automorphism count is*

$$(\ell_i - 1) + (h_i - 1) = \ell_i + \lceil n_i/\ell_i \rceil - 2$$

and it also uses s_i constant multiplications and $O(n_i)$ additions, where

$$s_i = \big|\{a \in \{0, \dots, n_i - 1\} : \widehat{c}_{i,a} \neq 0\}\big|.$$

The s_i products form the BSGS inner sums.

Proof. Applying $\sigma_{i,j\ell_i}$ to the summand indexed by r in (17) gives

$$\widehat{c}_{i,r+j\ell_i}\, \mathsf{Aut}_{\sigma_{i,r+j\ell_i}}(u).$$

Summing over the nonzero terms indexed by r, j enumerates the support of the group-algebra expression exactly once; omitted terms are zero. Thus Equation (16) gives an authenticated share of $m_i\Pi_i(z)$. There are $\ell_i - 1$ nonidentity baby steps and $h_i - 1$ nonidentity giant steps. There is one constant multiplication for each nonzero group-algebra coefficient. Integrality follows from Lemma 2 and is preserved by automorphisms. □

Combining the factor projectors with the single halving automorphism, write

$$A_{\mathrm{pf}} = 1 + \sum_{i=1}^{k} \left(\ell_i + \left\lceil \frac{n_i}{\ell_i} \right\rceil - 2\right) \tag{19}$$

for the total number of projector and global-halving automorphisms in one restricted multiplication.

Within the consecutive rectangular schedules of Equations (17)–(18), we choose ℓ_i to minimize $\ell_i + \lceil n_i/\ell_i \rceil - 2$.

Gen. Generate the base-HSS parameters over $\mathcal{R}_{q,M}$ and publish keys for the baby and giant automorphisms in (17)– (18), together with the global halving automorphism in (14). The public data also contain the guard sets, BSGS schedules, nonzero twisted projector constants, and the combined inverse-conductor normalization. The evaluation shares include an authenticated unit $\langle \boldsymbol{s}\cdot 1\rangle$: sample its party-0 vector uniformly and obtain party 1's vector by adding $\boldsymbol{s}$.

Enc. Encrypt each SIMD input as $\mathsf{ct}_{\boldsymbol{x}} \leftarrow \mathsf{Enc}(\mathcal{E}(\boldsymbol{x}))$.

Eval. Associate a memory wire $\boldsymbol{y}$ with $\langle \boldsymbol{s}\mathcal{E}(\boldsymbol{y})\rangle$. Addition is defined as ReRand of the componentwise sum; input ciphertexts are added componentwise. Because the input and memory encodings are both $\mathcal{E}$, a load computes $\mathsf{Mult}(\mathsf{ct}_{\boldsymbol{x}}, \langle \boldsymbol{s}\rangle)$ using the authenticated unit share in the evaluation key. Restricted multiplication is

$$\begin{aligned} u &\leftarrow \mathsf{Mult}(\mathsf{ct}_{\boldsymbol{x}}, \langle \boldsymbol{s}\mathcal{E}(\boldsymbol{y})\rangle), \\ \widehat{v} &\leftarrow \widehat{\mathsf{Project}}_k(\cdots \widehat{\mathsf{Project}}_1(u)\cdots), \\ v &\leftarrow \left(\prod_{i=1}^{k} m_i\right)^{-1} \cdot \widehat{v}, \\ \langle \boldsymbol{s}\mathcal{E}(\boldsymbol{x}\odot\boldsymbol{y})\rangle &\leftarrow \mathsf{Aut}_\tau(v). \end{aligned}$$

The inverse of $\prod_i m_i$ is taken modulo q and exists by the conductor coprimality assumption; being invertible, this normalization preserves the rerandomized output of the final projector. At output, each party selects the coefficients indexed by A, takes their centered representatives modulo q, and applies the sign convention of Section 2.

Parameter admissibility. Suppose the input and memory encodings have coefficient bounds $X, Y \le B$. Set $K_i = 16p_i|A_i|$. Lemma 9 bounds the constants in every BSGS inner sum by K_i, and each such sum contains at most n_i terms. At the logical level the unnormalized output is exactly $m_i\Pi_i(z)$, of norm at most $m_i\|z\|_\infty$; after all factors, the single inverse-conductor product restores the nonexpansive selector $\Pi_k \circ \cdots \circ \Pi_1$. Lemma 3 bounds the intermediate factor automorphisms and the global halving map, and the corresponding secret images by $2^k B_s$ whenever $\|s\|_\infty \le B_s$; this estimate uses no degree or support property of s. Thus all represented states have norm polynomial in N, $\mu_{\mathcal{R}}$, X, Y, the K_i, and $\prod_i m_i$. Since $\mu_{\mathcal{R}} \le N^2$, Lemma 4 supplies BKS-compatible polynomial-bit parameters for fixed k when N and the call count are polynomial in λ, the logical, secret, conductor, and error bounds have polynomial bit length, and the error tails are negligible.

Theorem 3 (Correctness of Construction 2). *For every B-admissible RMS program–input pair $(P, \boldsymbol{x})$ and every output modulus $r_{\mathsf{out}} \ge 2$, Construction 2 with BKS-compatible parameters produces additive shares of the T coordinatewise evaluations modulo r_{out}, except with negligible probability.*

Proof. Induct on the instructions of P. The BKS interface gives the encoded ring product. Lemma 10 and the combined normalization realize the factor projectors; Lemma 8 then returns the encoding of the componentwise product. Rerandomized memory addition and componentwise input addition preserve the encoding, and coefficient selection decodes it. Proposition 1 supplies the probabilistic BKS correctness step. □

Theorem 4 (Security of Construction 2). *Fix any secret distribution D_{sk} supported on $\{s \in \mathcal{R}_M : \|s\|_\infty \le B_s\}$. Assume circular RLWE for D_{sk} and the function family $\mathcal{F}_{\Sigma_2}$ of Section 2, where Σ_2 contains the baby, giant, and global-halving automorphisms used by Construction 2. Then Construction 2 satisfies one-evaluator HSS privacy; in particular, no coordinate of s is required to vanish identically under D_{sk}.*

Proof. The authenticated unit uses an independent uniform mask, so either evaluator's share is uniform and independent of the secret and remaining setup view. The encoding, projector constants, guard sets, and BSGS schedule are public, and all processing after the BKS multiplications and authenticated automorphisms is local. The projector and halving circuits remain in $\mathcal{R}_M$, and neither their algebraic identities nor their norm analysis requires a secret coordinate to vanish. Proposition 1 applies with the stated automorphism set. □

4.3 Asymptotic parameters

Balancing the factor dimensions makes the effect of k explicit. For every fixed k, Construction 2 packs $N^{1-o(1)}$ values and uses $O(kN^{1/(2k)})$ authenticated automorphisms and $O(kN^{1/k})$ constant multiplications per restricted multiplication.

Theorem 5 (Asymptotic packing and extraction). *Fix k and distinct odd primes $p_1, \dots, p_k$. Let the prime-power exponents vary with λ so that the degrees n_i are balanced, put $N(\lambda) := \prod_i n_i$, and assume $N(\lambda) \le \mathrm{poly}(\lambda)$. Thus $n_i = \Theta(N^{1/k})$. Construction 2 admits guard-banded sets with $T = N^{1-o(1)}$ SIMD slots. With $\ell_i = \lceil \sqrt{n_i} \rceil$, each restricted multiplication uses one base-HSS multiplication,*

$$A_{\mathrm{pf}} = 1 + \sum_{i=1}^{k} \left(\ell_i + \left\lceil \frac{n_i}{\ell_i} \right\rceil - 2 \right) = O(kN^{1/(2k)})$$

authenticated automorphisms, where the leading one is the final global halving. Thus each party performs

$$4 + 2A_{\mathrm{pf}} = O(kN^{1/(2k)})$$

ciphertext–share ring multiplications and $2+A_{\mathrm{pf}}$ distributed roundings. If s_i is the support size of the i-th projector as in Lemma 10, the logical circuit additionally uses exactly

$$C = 1 + \sum_{i=1}^{k} s_i \le 1 + \sum_{i=1}^{k} n_i = O(kN^{1/k})$$

constant multiplications. Equivalently, each party performs $2C$ componentwise constant multiplications. The precomputed projector data comprise $\sum_i s_i$ factor-ring elements and one normalization scalar.

Proof. Equation (5) with $L_i = \Theta(n_i)$ gives

$$T = \prod_i |A_i| \geq \prod_i L_i \exp\left(-O\left(\sum_i \sqrt{\log L_i}\right)\right) = N^{1-o(1)}.$$

Lemma 10 gives the projector counts; Equation (14) adds one automorphism and the initial base-HSS multiplication gives the stated BKS local-operation counts. The nonzero projector terms and the single combined inverse-conductor normalization give the stated constant-multiplication count. Balancing gives the claimed exponents. For fixed conductor primes, the power-basis reduction rule and Lemma 9 give polynomial-bit extraction intermediates. Lemma 4 therefore supplies BKS-compatible parameters of polynomial bit length. □

For factor i, the public projector contains s_i nonzero factor-ring constants, each represented by n_i coefficients modulo q. With balanced factors, their compact coefficient representation therefore contains

$$\sum_i s_i n_i \leq \sum_i n_i^2 = O(kN^{2/k})$$

coefficients modulo q. If q is split into $w(N)$ RNS limbs and each factor transform is padded to m_i coordinates, the table occupies $w(N)\sum_i s_i m_i = O(w(N)kN^{2/k})$ machine-word residues for fixed conductor primes. This excludes the full-ring switching keys from the inherited authenticated-automorphism interface, whose representation-dependent storage scales with A_{pf}. Both setup costs remain relevant even when the constants are transformed offline.

A native load uses four ciphertext–share ring multiplications and two distributed roundings per party; it invokes no projector.

A paired Construction 1 multiplication uses two base-HSS calls, A_{pair} authenticated automorphisms, and two scalar normalizations; Proposition 2 reduces the first and last counts to one on a type-consistent region and invokes only one directional trace. Construction 2 uses one base-HSS call, A_{pf} automorphisms, and C constant multiplications, of which $C-1$ multiply by factor-projector constants embedded in the tensor ring. Thus Construction 2 achieves higher occupancy and better asymptotic amortization of authenticated calls, while Construction 1 provides a compact extraction circuit and clean addition closure.

Under the illustrative cost model in which a ciphertext–share ring multiplication or distributed rounding costs $O(N \log N)$ and a constant multiplication costs $O(N)$, Theorem 5 gives Construction 2 total cost

$$O\Big(kN^{1+1/k} + kN^{1+1/(2k)} \log N\Big)$$

per evaluator, or $N^{1/k+o(1)}$ per occupied slot for fixed k. Under the same model, paired Construction 1 costs $O(N \log^2 N)$ in total and $O(\sqrt{N} \log^2 N)$ per occupied slot. Since these estimates depend on the underlying ring representation and transform strategy, Section 5 reports the constituent operation counts separately.

5 Implementation and Evaluation

We evaluate both constructions at their target ring dimensions with two objectives: to validate the complete two-party HSS computation and to determine its exact algebraic cost. The implementation instantiates the BKS multiplication of Boyle, Kohl, and Scholl [6] and the authenticated-automorphism interface of SIMD-HSS [11], following the syntax of Section 2. Inputs are encrypted by the client, evaluated independently by the two parties, and reconstructed only from the released additive output shares.

The ring backend embeds the powerful basis into a padded product grid and evaluates ring products with separable mixed-radix transforms. Primitive-root coordinates represent the cyclotomic quotient, and the remaining coordinates are fixed to zero. Construction 1 represents the degree-41,472 ring $\mathcal{R}_{512\cdot 243}$ on a 124,416-coordinate grid; Construction 2 represents its degree-45,360 ring on an $81 \cdot 25 \cdot 49 = 99{,}225$-coordinate grid and applies the three projectors factorwise. This realizes the powerful-basis tensor arithmetic of LPR [15]. The implementation uses Lattigo's cryptographic randomness interface [16] together with a dedicated mixed-radix modular-arithmetic backend.

Both instances use ten 54-bit RNS primes, split evenly between p and Δ, and hence a 540-bit ciphertext modulus $q = p\Delta$. Each secret coefficient is sampled independently and uniformly from $\{-1, 0, 1\}$, without prescribed-zero positions, and errors follow an integer-CDF approximation to a discrete Gaussian of width 3.2 conditioned on $[-64, 64]$. Our evaluation reports exact operation counts and end-to-end correctness at these dimensions.

Table 2. Exact operation counts for one restricted SIMD multiplication. The Construction 1 paired row uses the addition-closed representation; the typed rows report its directional fast paths.

		Logical calls		Per-evaluator ring operations		
Mode	Slots	Base mult.	Aut.	CT–share mult.	Rounds	Comp. const. mult.
C1 paired	162	2	17	42	21	4
C1 typed ($A \to B$)	162	1	8	20	10	2
C1 typed ($B \to A$)	162	1	9	22	11	2
C2	495	1	32	68	34	232

One base multiplication accounts for four ciphertext–share products and two roundings per evaluator, while an authenticated automorphism accounts for two products and one rounding. Each logical constant multiplication updates both authenticated components, so the final column counts two componentwise constant multiplications for each such operation. Compared with paired Construction 1, Construction 2 reduces the ciphertext–share count per occupied slot from $42/162 \approx 0.259$ to $68/495 \approx 0.137$, at the cost of increasing the final-column count from $4/162 \approx 0.025$ to $232/495 \approx 0.469$. The typed Construction 1 directions use 0.123 and 0.136 ciphertext–share products per slot and 0.049 and 0.056

automorphisms per slot, respectively, versus 0.137 and 0.065 for Construction 2. Thus Construction 2 offers substantially greater occupancy, whereas typed Construction 1 can require fewer authenticated operations and far fewer componentwise constant multiplications.

For each construction, a full-parameter execution samples $\boldsymbol{x}, \boldsymbol{y}$ independently and uniformly from $\{-3, \dots, 3\}^T$, loads $\boldsymbol{y}$, evaluates the restricted product with an encoding of $\boldsymbol{x}$, and releases additive shares of $(x_t y_t)_{t \in [T]}$ modulo 257. The validation checks the authentication relation after every authenticated automorphism, the decoded output and its support, the operation counters, and the BKS rounding and lifting conditions at every call. All checks pass, and a union bound over the executed calls places the correctness-failure probability below 2^{-180} for each reported execution.

5.1 Construction 1

For conductors $512 = 2^9$ and $243 = 3^5$, the factor degrees are 256 and 162, so every mode supports $T = 162$ slots. The binary and ternary subgroup-sum traces use 8 and 9 authenticated automorphisms, respectively. A paired load or multiplication evaluates both connector products, both traces, and the normalizations by 512^{-1} and 243^{-1}, yielding the first Construction 1 row of Table 2.

We also evaluate an alternating $A \to B \to A \to B$ execution containing a nonzero same-type memory addition. This execution covers both directional multiplications and both output decoders. One $A \to B \to A$ cycle uses two base-HSS multiplications, 17 automorphisms, 42 ciphertext–share multiplications, 21 roundings, and four componentwise constant multiplications per evaluator, matching the cost of one paired instruction. The setup retains both trace-key families, but typed evaluation invokes only the family required by the current direction.

5.2 Construction 2

The factor degrees are 54, 20, and 42, with guarded intervals of lengths 27, 10, and 21. The selected supports are

$$\begin{aligned} A_1 &= \{0, 2, 3, 7, 8, 10, 15, 19, 21, 24, 25\}, \\ A_2 &= \{0, 5, 6, 8, 9\}, \\ A_3 &= \{1, 3, 6, 7, 12, 14, 15, 19, 20\}. \end{aligned}$$

Proposition 3 (Optimality within the reported search family). *The displayed supports are 3-AP-free and give* 495 *slots. This is the maximum in the separable half-interval guard-band family for conductors* 81, 25, *and* 49. *Within the reported full consecutive-rectangle BSGS circuits of Equations* (17)–(18)*, coefficient projection uses the minimum* $13 + 7 + 11 = 31$ *authenticated automorphisms, attained by the schedules in Table 3. The simultaneous halving raises the total to* 32.

Proof. An exhaustive branch-and-bound search considers every subset of each guarded interval; a branch is pruned only when its cardinality cannot exceed the current optimum. Writing $r_3(L)$ for the largest 3-AP-free subset of an interval of length L, the resulting optima are

$$r_3(27) = 11, \qquad r_3(10) = 5, \qquad r_3(21) = 9,$$

and the displayed sets attain these values. Since the support is Cartesian and the three constraints are factorwise, the maximum number of slots is $11 \cdot 5 \cdot 9 = 495$.

For a degree-n_i factor, a full consecutive rectangle of width ℓ uses $\ell - 1$ baby and $\lceil n_i/\ell \rceil - 1$ giant automorphisms. Minimizing $\ell + \lceil n_i/\ell \rceil - 2$ over $1 \le \ell \le n_i$ gives 13, 7, and 11 for $n_i = 54, 20, 42$, attained by the displayed 6×9, 4×5, and 6×7 rectangles. Equation (14) contributes the remaining single automorphism. □

The exhaustive search returns 12, 24, and 18 maximum-cardinality guarded sets for the three factors. Minimizing projector support over these sets gives 54, 19, and 42 nonzero coefficients. Among the resulting candidates, we enumerate the cyclic generators and minimum-call rectangular shapes and select those minimizing the induced coefficient bound. Table 3 reports the selected schedules, where g_i defines $X_i \mapsto X_i^{g_i}$.

Table 3. Projector schedules for the 495-slot Construction 2 instance.

m_i	n_i	g_i	BSGS shape	Support	Aut.
81	54	56	6×9	54	13
25	20	17	4×5	19	7
49	42	47	6×7	42	11

Here support denotes the number of nonzero projector coefficients. The 115 coefficients and one combined inverse-conductor normalization give 116 constant multiplications, or 232 componentwise constant multiplications per evaluator. The implementation verifies each factor projector on the complete power basis and reproduces the automorphism and support counts in Table 3.

6 Conclusion

Our progression-free construction combines guarded supports with integral factorwise projectors to pack $N^{1-o(1)}$ values, for fixed k and fixed conductor primes with balanced prime-power factors. It uses one base-HSS multiplication, improves asymptotic per-slot authenticated-call amortization relative to paired Construction 1, and extracts coefficients with $O(kN^{1/k})$ constant multiplications. The complementary first construction adapts Liu and Wang's alternating tensor framework to authenticated BKS shares and packs $\Theta(\sqrt{N})$ values. Its

paired mode closes arbitrary RMS additions with two base-HSS multiplications, logarithmically many authenticated automorphisms, and two scalar normalizations; its typed fast path uses one base-HSS multiplication and one directional trace on compatible regions. Under the stated circular-RLWE hypotheses, both constructions support independent uniform ternary secrets without requiring selected coefficients to vanish.

The technical bridge is an integral authenticated-share realization of trace-dual coefficient functionals. Composing these maps with the BKS multiplication and SIMD-HSS automorphism interfaces yields correctness and one-evaluator privacy under their corresponding assumptions. The end-to-end implementation validates the algebraic outputs and exact operation counts at the selected 162- and 495-slot instances, including the alternating dependency of the typed Construction 1 path. Together, the constructions expand the HSS packing design space with near-linear occupancy, clean addition closure, compact scalar extraction, and typed online savings.

Acknowledgements

The author was supported by the Stanford Graduate Fellowship. The technical ideas and high-level construction designs in this work were developed by the author. OpenAI's Codex assisted with their formalization and exposition, editorial revision, literature review, and the development and validation of the experimental implementation used in this work. All AI-assisted material was reviewed and verified by the author, who assumes full responsibility for the claims, citations, implementation, and reported results.

References

1. Alperin-Sheriff, J., Peikert, C.: Practical bootstrapping in quasilinear time. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 1–20. Springer, Berlin, Heidelberg (Aug 2013). https://doi.org/10.1007/978-3-642-40041-4_1
2. Behrend, F.A.: On sets of integers which contain no three terms in arithmetical progression. Proceedings of the National Academy of Sciences **32**(12), 331–332 (Dec 1946). https://doi.org/10.1073/pnas.32.12.331
3. Block, A.R., Maji, H.K., Nguyen, H.H.: Secure computation based on leaky correlations: High resilience setting. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 3–32. Springer, Cham (Aug 2017). https://doi.org/10.1007/978-3-319-63715-0_1
4. Boyle, E., Gilboa, N., Ishai, Y.: Breaking the circuit size barrier for secure computation under DDH. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part I. LNCS, vol. 9814, pp. 509–539. Springer, Berlin, Heidelberg (Aug 2016). https://doi.org/10.1007/978-3-662-53018-4_19
5. Boyle, E., Gilboa, N., Ishai, Y., Lin, H., Tessaro, S.: Foundations of homomorphic secret sharing. In: Karlin, A.R. (ed.) ITCS 2018. vol. 94, pp. 21:1–21:21. LIPIcs (Jan 2018). https://doi.org/10.4230/LIPIcs.ITCS.2018.21

6. Boyle, E., Kohl, L., Scholl, P.: Homomorphic secret sharing from lattices without FHE. In: Ishai, Y., Rijmen, V. (eds.) EUROCRYPT 2019, Part II. LNCS, vol. 11477, pp. 3–33. Springer, Cham (May 2019). https://doi.org/10.1007/978-3-030-17656-3_1
7. Conrad, K.: The different ideal. Expository note (2009), https://kconrad.math.uconn.edu/blurbs/gradnumthy/different.pdf
8. Gentry, C., Halevi, S., Peikert, C., Smart, N.P.: Ring switching in BGV-style homomorphic encryption. In: Visconti, I., De Prisco, R. (eds.) SCN 12. LNCS, vol. 7485, pp. 19–37. Springer, Berlin, Heidelberg (Sep 2012). https://doi.org/10.1007/978-3-642-32928-9_2
9. Halevi, S., Shoup, V.: Bootstrapping for HElib. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015, Part I. LNCS, vol. 9056, pp. 641–670. Springer, Berlin, Heidelberg (Apr 2015). https://doi.org/10.1007/978-3-662-46800-5_25
10. Halevi, S., Shoup, V.: Faster homomorphic linear transformations in HElib. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 93–120. Springer, Cham (Aug 2018). https://doi.org/10.1007/978-3-319-96884-1_4
11. Kim, J., Li, H., Lin, H., Liu, Z.: SIMD HSS and aHMAC from interval encoding with application to one-bit-per-gate garbling. Cryptology ePrint Archive, Paper 2026/485 (2026), https://eprint.iacr.org/2026/485
12. Lipmaa, H.: Progression-free sets and sublinear pairing-based non-interactive zero-knowledge arguments. In: Cramer, R. (ed.) TCC 2012. LNCS, vol. 7194, pp. 169–189. Springer, Berlin, Heidelberg (Mar 2012). https://doi.org/10.1007/978-3-642-28914-9_10
13. Liu, F.H., Wang, H.: Batch bootstrapping I: A new framework for SIMD bootstrapping in polynomial modulus. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part III. LNCS, vol. 14006, pp. 321–352. Springer, Cham (Apr 2023). https://doi.org/10.1007/978-3-031-30620-4_11
14. Liu, F.H., Wang, H.: Batch bootstrapping II: Bootstrapping in polynomial modulus only requires $\tilde{O}(1)$ FHE multiplications in amortization. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part III. LNCS, vol. 14006, pp. 353–384. Springer, Cham (Apr 2023). https://doi.org/10.1007/978-3-031-30620-4_12
15. Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-LWE cryptography. In: Johansson, T., Nguyen, P.Q. (eds.) EUROCRYPT 2013. LNCS, vol. 7881, pp. 35–54. Springer, Berlin, Heidelberg (May 2013). https://doi.org/10.1007/978-3-642-38348-9_3
16. Mouchet, C., Troncoso-Pastoriza, J.R., Bossuat, J.P., Hubaux, J.P.: Multiparty homomorphic encryption from ring-learning-with-errors. PoPETs **2021**(4), 291–311 (Oct 2021). https://doi.org/10.2478/popets-2021-0071
17. Smart, N.P., Vercauteren, F.: Fully homomorphic SIMD operations. DCC **71**(1), 57–81 (2014). https://doi.org/10.1007/s10623-012-9720-4