

Article

Adaptive Compact Knapsack Cryptography for Post-Quantum Security

Aso Ahmed Majeed ¹

College of Nursing, University of Kirkuk, Kirkuk, Iraq

E-mail: asoalsalihi@uokirkuk.edu.iq

Article information:

Manuscript received: 4 June 2026; **Accepted:** 10 July 2026; **Published:** 15 August 2026

Abstract: Digital signatures based on the compact knapsack problem are attractive post-quantum candidates: they rely solely on integer arithmetic, require no algebraic ring structure, and rest on a hardness assumption distinct from the Module-LWE/Module-SIS problems that underlie the recently standardized ML-DSA (FIPS 204). The three-move identification scheme of Rizo and Draziotis (2023) and its Fiat-Shamir signature establish completeness, special soundness, special honest-verifier zero-knowledge, and EUF-CMA security in the random-oracle model. Its concrete parameters, however, were selected by hand, yielding a comparatively large signature (≈ 31.9 KB); the authors explicitly leave smarter parameter selection as future work. We close this gap. First, we provide an independent, fully reproducible re-implementation of the scheme's lattice attacks (the closest-vector and divide-and-conquer attacks) and of its signature-size accounting, validated against the original tables; our re-implementation reproduces the published success rates to within experimental noise. Second, we cast parameter selection as a constrained optimization problem and design an adaptive bit-allocation optimizer that minimizes signature size subject to soundness, completeness, commitment-unpredictability, key-entropy, and an empirical lattice-attack hardness gate. The optimizer exposes two honest sources of slack in the base parameters—an over-provisioned nonce space and an under-tightened secret space. The resulting configuration $S = S1/4, 1/3$, $S' = S'1/8, 1/8$ shrinks the signature by 19.6% (to 25.6 KB) at identical security: equal soundness (2^{-80}), equal commitment-unpredictability (2^{-24}), equal completeness error (2.15×10^{-6}), empirically equal closest-vector-attack hardness (50.0%), and equal-or-better divide-and-conquer hardness. All code, seeds, and numerical results are released for reproducibility.

Keywords: Post-quantum cryptography; Compact knapsack problem; Fiat-Shamir transform; Lattice attacks; Closest vector problem; Digital signatures; Parameter optimization; Sigma protocols

1. Introduction

The prospect of large-scale quantum computers has made the migration of public-key cryptography to quantum-resistant primitives a global priority. In August 2024 the U.S. National Institute of Standards and Technology (NIST) issued its first post-quantum standards—FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA)[1][2]. Two of the three are lattice-based, and both signature workhorses—ML-DSA and the NTRU-based FN-DSA—rest on structured-lattice assumptions. A healthy cryptographic ecosystem, however, benefits from a *diversity* of hardness assumptions: should a weakness emerge in one family, deployed systems must be able to fall back on primitives whose security rests on an unrelated problem[3][4].

The *compact knapsack problem* offers exactly such an alternative. Given an integer matrix $A \in \mathbb{Z}^{m \times n}$ and a target b , one must find an integer vector x in a prescribed structured set $S \subset \mathbb{Z}^n$ with $Ax^T = b^T$. The problem requires no ring structure and only

integer additions and dot products; it is a multidimensional, matrix-defined generalization of the classical subset-sum/knapsack problem[5][6]. Although early one-dimensional knapsack cryptosystems were broken by lattice reduction, the modern *compact* (matrix, large-coefficient) formulation lifts the instances out of the low-density regime that made those attacks effective[7].

Building on this problem, Rizos and Draziotis[8] constructed a sound three-move identification (ID) scheme and, via the Fiat–Shamir (FS) transform[9][10], a digital signature that is existentially unforgeable under chosen-message attack (EUF-CMA) in the random-oracle model (ROM). The scheme is complete, special sound, and special honest-verifier zero-knowledge (S-HVZK), and is believed post-quantum because no efficient quantum algorithm is known for compact knapsack. Its principal practical drawback is signature size: with the published parameters ($t = 80, m = 4, n = 48, R = 192, S = S_{1/4,1/2}, S' = S'_{1/8,1/4}$) a signature occupies roughly 31.9 KB. The authors note that choosing the parameters more cleverly may reduce signature length, and explicitly leave this to future work.

Contributions. This paper takes up that future-work question and answers it concretely and reproducibly.

1. **Faithful, reproducible re-implementation.** We re-implement the scheme's two lattice attacks—the closest-vector (CVP) attack and the divide-and-conquer (D&C) attack—together with the exact signature-size accounting, in open Python using the fpylll reduction engine. We validate the implementation against the original Tables 1 and 2: the homogeneous space S_1 is recovered at 100% and the structured spaces at exactly 50% (Table 1), and the six-block space of Table 2 is reproduced to within 1–2 percentage points (Section 5).
2. **Formalization of parameter selection.** We cast the choice of the block bit-profile as a constrained optimization problem (Problem 1) whose constraints encode soundness, completeness (the authors' Corollaries 4.2–4.3), commitment unpredictability, key entropy, and *empirical* attack hardness.
3. **An adaptive optimizer with a hardness oracle.** Algorithm 1 searches well-separated bit-profiles and gates each candidate through the reproduced lattice attacks, used as an empirical hardness oracle. This is the adaptive core of the scheme: bits are allocated to blocks only where the attack data show they purchase security.
4. **A concrete, verified improvement.** The optimizer identifies the two honest sources of slack and returns $S = S_{1/4,1/3}, S' = S'_{1/8,1/8}$. This shrinks the signature by 19.6% at identical security across every metric we measure, with completeness independently confirmed by a 2×10^5 -sample Monte-Carlo experiment.

We emphasize that the proposed scheme introduces *no new hardness assumption*: it operates in the same security model as and reuses the same *class* of structured secret space; rather than asserting that hardness is inherited, we *re-measure* it empirically and require the optimized secret space to meet the reference hardness level, removing only provably unnecessary slack elsewhere.

Organization. Section 2 reviews lattice-based Fiat–Shamir signatures, knapsack cryptography, lattice attacks, and parameter optimization. Section 3 recalls the compact-knapsack signature. Section 4 formalizes the optimization problem and presents the adaptive optimizer. Section 5 reports the reproduction study, the hardness frontier, the size reduction, and the security-parity analysis. Section 6 concludes.

The dominant post-quantum signature paradigm is “Fiat–Shamir with aborts”[11], instantiated at scale by CRYSTALS-Dilithium, now standardized as ML-DSA (FIPS 204)[12]. ML-DSA signatures range from 2,420 to 4,595 bytes across security levels. The NTRU hash-and-sign signature Falcon (FN-DSA) produces still smaller signatures, at the cost of floating-point Gaussian samplers. Active research continues to compress lattice FS signatures: HAETAE attains signatures up to 39% smaller than Dilithium through

bimodal distributions and a hyperball uniform sampler[13], while the G+G scheme builds compact Fiat–Shamir signatures from convolved Gaussians with a tight QROM proof[14]. Surveys trace the steady evolution of the hash-and-sign and Fiat–Shamir families and their specialized variants[15]. All of these rest on structured-lattice (Module-LWE/SIS or NTRU) assumptions.

Knapsack public-key cryptography began with Merkle–Hellman and was famously broken by lattice reduction and Shamir-style attacks, because the underlying super-increasing subset-sum instances were low-density; modern lattice attacks recover Merkle–Hellman keys in under a second. These failures motivated the move from $\{0,1\}$ subset-sum to *compact*, matrix-defined knapsacks with large integer coordinates, which evade the low-density attack regime[16]. The hardness of subset-sum under lattice attacks remains an active topic: improvements to the Lagarias–Odlyzko density threshold and faster practical lattice reduction continually refine the secure parameter frontier, underscoring the need for *empirically grounded* parameter selection rather than worst-case heuristics.

In the attack direction, the security of compact knapsack reduces to instances of the closest vector problem (CVP), which is NP-hard[17]. In practice one reduces a basis with LLL [18] and then applies Babai's nearest-plane algorithm[19] to the kernel lattice of A . The base scheme's experiments take exactly this form, and our hardness oracle faithfully reproduces them. Side-channel and hybrid analyses of standardized lattice schemes[20] further illustrate that concrete security must be evaluated empirically and per-parameter-set.

Although individual schemes tune their parameters, the literature contains few *general* procedures that treat block-structure / bit-allocation selection for a knapsack-type signature as an explicit optimization driven by measured attack hardness. The base paper selects S and S' by hand and flags optimization as open. Table 1 situates the present work among representative studies.

Table 1. Representative studies in post-quantum / knapsack signatures and lattice attacks, and their relation to this work.

Study	Year	Focus	Relation to this work
G+G	2023	Convolved-Gaussian FS sig.	Alternative compaction route
Ryan–Heninger	2023	Faster lattice reduction	Refines secure thresholds
Darzi	2023	PQ standardization survey	Motivates assumption diversity
Rizos–Draziotis	2023	Compact knapsack sig.	Base scheme we optimize
HAETAE	2024	Shorter lattice FS sig.	Compaction in MLWE family
Survey	2024	Lattice-signature survey	Context / taxonomy
NIST FIPS 204	2024	ML-DSA standard	Standardized reference sizes
Zong	2024	Math. foundations of PQC	SVP/CVP geometry
Wang	2024	Hybrid attack on Kyber	Empirical-security argument
Lagarias–Odlyzko impr	2024	Subset-sum density	Hardness-threshold refinement
Merkle–Hellman attack	2025	LLL key recovery	Why “compact” is needed
This work	2025	Adaptive bit-allocation	Optimizes: –19.6% size

2. Materials and Methods

We recall the parts of needed below. For $\alpha > 0$ with $\alpha R \in \mathbb{Z}$, let $I_{\alpha R} = \{v \in \mathbb{Z} : 2^{\alpha R-1} \leq v \leq 2^{\alpha R} - 1\}$ be the set of (αR) -bit integers, and for $k | n$ define the structured solution space

$$S_{\alpha_1, \dots, \alpha_k}(n, R) = I_{\alpha_1 R}^{n/k} \times I_{\alpha_2 R}^{n/k} \times \dots \times I_{\alpha_k R}^{n/k} \subset \mathbb{Z}^n. \quad (1)$$

The public key is (A, b, S, S', t, H) with $A \leftarrow \mathcal{M}_{m \times n}(I_{R/8})$, a secret $x \in S$, $b^T = Ax^T$, and a hash $H : \{0,1\}^* \rightarrow \{0,1\}^t$. The t -fold parallel ID-scheme has the prover commit $R = AK$ for a nonce matrix K with columns in S' , receive a challenge $e \in \{0,1\}^t$, and respond with $S = K + xe^T$ (column-wise $s_i = k_i + e_i x$). The FS signature is $\sigma = (R, S)$, verified by recomputing $e = H(R, msg)$ and checking $As_i^T = r_i^T + e_i b^T$ together with membership $s_i \in S'$ if $e_i = 0$, else $s_i \in S$.

Completeness. For the response to lie in S when $e_i = 1$ one needs $x + k \in S$ with overwhelming probability. The authors prove (Cor. 4.2–4.3, Lemma 4.4) that if $\beta_i \leq \frac{1}{R \ln 2} \ln(2^{\alpha_i R-1} + 1)$, then, per coordinate of block i ,

$$\varepsilon_i = \frac{3 \cdot 2^{\beta_i R-1} - 1}{2^{\alpha_i R}}, \quad \Pr[x + k \in S] = \prod_{i=1}^k (1 - \varepsilon_i)^{n/k} \quad (2)$$

Unpredictable commitments. EUF-CMA via FS requires unpredictable commitments. The probability of a commitment collision is bounded combinatorially by

$$P \leq \frac{1}{2^{\beta_{\min} R}}, \quad \beta_{\min} = \min_i \beta_i, \quad (3)$$

so a larger $\beta_{\min} R$ yields stronger unpredictability.

Soundness and size. Running t parallel binary rounds gives forgery probability 2^{-t} . With the accounting of the authors' parameters.py—writing $A_i = \alpha_i R$ and $B_i = \beta_i R$ for the per-block bit-lengths and $w = R/8$ for the matrix-entry size—the signature occupies

$$\frac{mt}{k} \sum_{i=1}^k (w + A_i) + \frac{nt}{k} \sum_{i=1}^k (w + B_i) \text{ bits}. \quad (4)$$

For the base parameters this evaluates to 261,120 bits = 32,640 bytes $\approx$ 31.9 KB, reproducing the authors' figure exactly. The first term is the commitment R ($m \times t$ entries); the second is the response S ($n \times t$ entries), which dominates because $n \gg m$.

The base scheme fixes a single hand-chosen bit-profile. Our method instead treats the profile as a decision variable and selects it through a closed loop (Figure 1): a candidate profile is first screened by cheap analytic constraints (size, completeness, collision, key entropy), and only the survivors are sent to an *empirical hardness oracle*—the reproduced CVP and D&C attacks—whose measured success rate gates admissibility. Among all admissible profiles the optimizer retains the one of minimum signature size. The adaptivity lies in *where bits are spent*: a block's bit-length is increased only when the attack data show it is required for hardness, and decreased to the security floor everywhere else.




Figure 1. Conceptual model of the adaptive optimizer. Cheap analytic constraints pre-filter candidates; the surviving profiles are gated by the reproduced lattice attacks acting as an empirical hardness oracle.

We fix the structural parameters $n = 48, m = 4, R = 192, t = 80$ at the base values so that hardness is compared like-for-like (the kernel-lattice dimension $n = \text{rank}(A)$ and the matrix-entry size $w = R/8$ are then unchanged). A configuration is the pair of integer bit-profiles $\{A_i\}_{i=1}^k$ (secret space S) and $\{B_i\}_{i=1}^k$ (nonce space S'). Relative to a reference configuration (the base scheme) we require:

(C1) **Soundness:** $t = 80 \Rightarrow 2^{-80}$. (5)

(C2) **Completeness:** $\text{err}(\{A_i\}, \{B_i\}) \leq \text{err}_{\text{ref}}$ (*err from Eq. 2*). (6)

(C3) **Unpredictability:** $\min_i B_i \geq \kappa$ ($= 24$). (7)

(C4) **Hardness:** $\text{succ}(\hat{S}) \leq \text{succ}_{\text{ref}}$ ($= 50\%$), *empirical* (§4.4). (8)

(C5) **Key entropy:** $\log_2 |S| = \frac{n}{k} \sum_i A_i \geq \lambda$ ($= 256$). (9)

(C6) **Validity:** $B_i \leq A_i - 1 \quad \forall i$ (*Cor. 4.2 hypothesis*). (10)

2.1 The optimization problem

Problem 1 (Adaptive bit-allocation). Among all block bit-profiles $(\{A_i\}, \{B_i\})$ satisfying (C1)–(C6) [Eq. 5–10], minimize the signature size of Eq. (4).

A structural observation drives the solution. Constraints (C3) and (C6), together with the completeness budget, force $A_i \geq B_i + g$ for a gap $g \approx 24$ (matching the reference ε), hence $A_i \geq 48$ and $B_i \geq 24$ for every block. Substituting the minimal $B_i = 24$ into Eq. (4) fixes the dominant response term at $nt \cdot 48 = 184,320$ bits and leaves the total monotone increasing in $\sum_i A_i$:

$$\text{size} = 184,320 + mt \cdot w + \frac{mt}{k} \sum_{i=1}^k A_i \quad \text{bits}. \quad (11)$$

Thus the problem reduces to finding the bit-profile of smallest mean A_i whose secret space S remains hard (C4). The base profile $A = (48, 96)$ wastes bits in two ways: the nonce profile $B = (24, 48)$ over-provisions the second block (collision needs only $B_{\min} = 24$), and the secret block $A_2 = 96$ exceeds what hardness requires.

The empirical hardness oracle

Constraint (C4) is evaluated by the reproduced attacks of. The CVP attack computes an integer particular solution y of $Az^T = b^T$, an LLL-reduced basis of the kernel lattice $L = \{z : Az^T = 0\}$, a Babai nearest-plane vector $b^* = \text{CVP}(L, t)$ toward the structured target t , and reports the best of $\{y \pm jb^*\}_{j=0}^{\alpha}$ following the authors' success criterion, a timing coordinate counts as recovered once it lies in the union of the per-block intervals. The divide-and-conquer attack builds on the same idea; it simply works block by block and solves each one on its own. To keep the estimate conservative, we take $\text{succ}(\hat{S}) = \max(\text{CVP}, \text{D\&C})$ as the attacker's success measure.

Whichever of the two attacks does better sets the score. Satisfying condition (C4) therefore yields no more than a lower bound on the security level relative to the reference scheme: the attacker, after all, never settles for the weaker of the two.

Interpreting the 50% plateau. The success metric counts a coordinate as recovered whenever it lands in the union of the block ranges. With two well-separated blocks, this gives the metric a built-in floor at exactly 50%: the attacker can always read off the coordinates of the wider block, but the disjoint ranges leave no room to do better. The figure is therefore relative, not absolute. It says the construction is no easier to break than the base space under the published attack; it says nothing about the absolute probability of a single-shot forgery. That question comes down, in the end, to the full lattice problem and the entropy left in each block. Our gate enforces exactly this property: the scheme must match the base level under the stronger of the two reproduced attacks.

The adaptive optimizer

Algorithm 1 puts this into practice. It walks through the well-separated profiles and gives each block the shortest nonce length the constraints allow. A cheap analytic screen then weeds out most candidates before any attack runs. Only the survivors reach the attack evaluation, and even there an incumbent bound prunes anything that cannot beat the best profile found so far. This keeps calls to the expensive oracle rare.

Algorithm 1 — AdaptiveBitAllocation

Require: structural params (n, m, R, t); reference (succ_ref, err_ref, κ, λ);
 bounds A_min, A_max; trials T; seed p

Ensure: optimal profile ({A*}, {B*}) of minimum signature size

```

1 best ← Nil
2 for all well-separated profiles {As}, As ∈ [A_min, A_max], k | n do
3   Bs ← MinNonce(As, κ, err_ref) ▷ smallest Bs ≥ κ with εs ≤ floor
4   if ∃ i : Bs > As - 1 then continue ▷ (C6) validity
5   if err({As}, {Bs}) > err_ref then continue ▷ (C2) completeness
6   if mina Ba < κ or (n/κ)ΣaAa < λ then cont. ▷ (C3), (C5)
7   z ← Size(n, m, R, t, {As}, {Bs}) ▷ Eq. (4)
8   if best ≠ Nil and z ≥ best.z then continue ▷ size prune before oracle
9   s ← HardnessOracle(m, n, R, {As}; T, p) ▷ (C4) expensive
10  if s ≤ succ_ref and (best = Nil or z < best.z) then
11    best ← ({As}, {Bs}, z, s)
12  end if
13 end for
14 return best

```

Cost. Let P = #candidate profiles, P_t = #survivors of the analytic screen (P_t << P).
 Analytic screen: O(P·k) closed-form evaluations. Oracle calls (LLL + Babai per trial):
 O(P_t·T) – bounded further by the incumbent-size prune (line 8). The polynomial-time
 LLL reduction dominates runtime but is invoked only on the few smallest survivors.
 Output for the base reference: A* = (48, 64), B* = (24, 24) → 26,240 B (−19.6%).

Algorithm 1. AdaptiveBitAllocation. The analytic screen (lines 4–8) discards inadmissible or non-improving profiles before the costly hardness oracle (line 9) is ever called.

Outcome

Running Algorithm 1 with the base scheme as reference returns the two-block profile

$$S^* = S_{1/4,1/3}(48,192) \ (A = (48,64)), \ S'^* = S'_{1/8,1/8}(48,192) \ (B = (24,24)). \quad (12)$$

Both moves are justified *within the base security model*: $S'_{1/8,1/8}$ keeps $\beta_{\min}R = 24$ (identical commitment unpredictability) while removing the unnecessary 48-bit second nonce block; and $S_{1/4,1/3}$ keeps the two secret blocks separated by 16 bits, which Section 5 shows is sufficient to remain on the same 50% hardness plateau as the base $S_{1/4,1/2}$.

3. Results**Experimental setup and reproducibility**

All experiments run in Python 3.12 with numpy and the fpylll 0.6.4 lattice engine (the same fpylll core used by SageMath), with exact integer arithmetic and a Babai nearest-plane CVP implemented as in. Every reported number uses a fixed random seed and is regenerated by a single script. As a validation anchor we first reproduce the base paper's published attack tables.

Table 2. Validation of the reproduced CVP attack against the base paper. Success = mean fraction of coordinates recovered (union metric).

(a) Table 1 of, n = 50, R = 80, 20 trials			
m	S ₁ (ours / paper)	S _{1,1/2} (ours / paper)	S _{1,2,1/4} (ours / paper)
1	100.0 / 100	50.0 / 50	50.0 / 50
2	100.0 / 100	50.0 / 50	50.0 / 50
10	100.0 / 100	50.0 / 50	50.0 / 50
20	99.8 / 100	50.0 / 50	50.0 / 50
(b) Table 2 of, n = 54, R = 64, S _{1,1/2,1/4,1/8,1/16,1/32} , 30 trials			
m	ours	paper	abs. diff.
1	80.1	78.7	1.4
10	30.8	31.2	0.4
40	23.5	22.9	0.6

Table 2 shows exact agreement on Table 1 (for $m \leq 20$, away from the boundary regime $m \rightarrow n$ that the original flags as special) and agreement to within 1.4 percentage points on the six-block space of Table 2. We conclude that the oracle is a faithful reproduction of the published attack, suitable for gating new parameter choices.

Hardness frontier

Figure 2 sweeps the bit-separation $A_2 - A_1$ of a two-block secret space S ($A_1 = 48$ fixed) and measures CVP success at the real parameters $n = 48, R = 192, m = 4$. The behaviour is a sharp threshold: when the two blocks are well separated ($A_2 - A_1 \gtrsim 16$) their value ranges are disjoint and the attack saturates at the 50% security plateau—the same level as the base space $S_{1/4,1/2}$ at separation 48. As the blocks approach each other, their ranges begin to merge under the union metric and success climbs (51.5% at separation 8, 71.8% at separation 4). The proposed profile $S_{1/4,1/3}$ sits at separation 16—the smallest separation cleanly on the plateau. Operationally, the threshold is what licenses the size reduction: it identifies the point below which additional separation stops purchasing measurable hardness, so that any separation beyond 16 bits is expenditure without security return at these parameters.




Figure 2. Empirical hardness frontier. CVP-attack success as a function of the bit-separation between the two secret blocks ($n = 48, R = 192, m = 4$, 60 trials per point). Both the base profile and the proposed one sit on the 50% plateau; success climbs only once the blocks are too close to tell apart.

Signature size

With the nonce blocks at the collision floor and the secret profile on the plateau, the proposed signature comes to 209,920 bits, or 26,240 bytes (25.6 KB), against 261,120 bits, or 32,640 bytes (31.9 KB), for the base scheme. That is a 19.6% reduction (Figure 3).

Two terms account for the saving: the response term, which dominates and drops by 20% once the second nonce block is halved, and the commitment term, which falls by 16.7% as we tighten from 96 to 64 bits.




Figure 3. Signature-size breakdown. The response matrix S dominates the total; the proposed parameters cut it by 19.6% at the same security level

Security and correctness preservation

Table 3 puts the two schemes side by side. Soundness (2^{-80}), commitment unpredictability (2^{-24}), and the completeness error (2.15×10^{-6}) match by construction. The secret space is just as hard in practice: under CVP the two profiles sit at the same level (50.01% vs. 50.00%, within the noise), and under D&C the attack succeeds no more often against ours (10.50% vs. 10.83%, a difference within sampling error). Key entropy stays at 2^{2688} , comfortably above the 2^{256} floor. Figure 4 shows the comparison at a glance. Both attack comparisons should be read as statements of statistical parity: at 200 CVP and 100 D&C trials the differences reported here are smaller than one standard error, so the optimized profile is established as no weaker rather than strictly stronger (Section 4.6.6).

Completeness was further checked by direct Monte-Carlo simulation: over 2×10^5 random pairs $(x, k) \in S \times S'$ we observed zero failures of $x + k \in S$ (consistent with the predicted 2.15×10^{-6}), and the $e = 0$ branch $s = k \in S'$ never failed.

Table 3. Base vs. proposed profile at $t = 80, n = 48, m = 4, R = 192$. Lower attack success means a harder scheme. Success rates come from 200 CVP and 100 D&C trials with a fixed seed.

Metric	Base	Proposed	Change
Secret space S	$S_{1/4, 1/2}$	$S_{1/4, 1/3}$	profile
Nonce space S'	$S'_{1/8, 1/4}$	$S'_{1/8, 1/8}$	profile
Secret bit-profile (A_i)	(48, 96)	(48, 64)	—
Nonce bit-profile (B_i)	(24, 48)	(24, 24)	—
Signature size	32,640 B	26,240 B	−19.6%
commitment R	30,720 b	25,600 b	−16.7%
response S	230,400 b	184,320 b	−20.0%
Soundness	2^{-80}	2^{-80}	=
Commitment unpredictability	2^{-24}	2^{-24}	=
Completeness error	2.15×10^{-6}	2.15×10^{-6}	=
Key entropy $\log_2 S $	2^{3456}	2^{2688}	$\gg 2^{256}$
CVP-attack success	50.00%	50.01%	$\approx$
D&C-attack success	10.83%	10.50%	no worse




Figure 4. Security and correctness metrics, base vs. proposed. All are equal or better for the proposed scheme; only the signature size differs materially.

Size–security frontier

Figure 5 plots the candidate profiles by attack success and signature size. Anything left of the 50% gate is admissible. Among these, the starred profile is the smallest one that stays firmly on the plateau under both attacks. One trade-off is worth spelling out: a three-block profile ($S_{1/4,1/4,1/3}$, $S'_{1/8,1/8,1/8}$) gets down to 26,133 bytes (a 19.9% saving) and holds up better under CVP (46.4%), but its finer block structure leaves it weaker under D&C (14.8%, against 10.8% for the base). Since our gate assumes the strongest attacker, we put the conservative two-block solution forward as the main result. For deployments that want every metric to stay put, a strictly dominant alternative ($S_{1/4,3/8}$, $S'_{1/8,1/8}$) saves 19.1% while keeping completeness at exactly the base value and CVP at 50.0%.




Figure 5. Size-versus-security frontier of candidate profiles. Green: admissible; grey: rejected (too easy or completeness exceeded). Stars mark base and proposed.

Interpretation of the main findings

The 19.6% reduction reported in Section 4.3 is not produced by a single tuning knob, and separating its two mechanically distinct sources is what makes the result interpretable rather than merely favourable. The first source is the nonce space. Commitment unpredictability in the base scheme is governed by the smallest block exponent alone, through the quantity $\beta_{\min R} = 24$ bits; the second nonce block of the published parameter

set nevertheless carries 48 bits. Every bit above the 24-bit floor is therefore paid for twice — once in the commitment and again, with multiplicity, in the response matrix — while purchasing no additional unpredictability. Contracting $S^{1/8,1/4}$ to $S^{1/8,1/8}$ removes exactly this surplus: the commitment term falls by 16.7% and the dominant response term by 20.0%, with the unpredictability level left at the 2^{-24} value chosen by the base scheme. The second source is the secret space, and here the operative quantity is not the total bit budget of S but the separation between its blocks. Figure 2 pins CVP success to a 50% plateau for $A_2 - A_1 \gtrsim 16$ and shows it rising steeply only below that point (51.5% at separation 8, 71.8% at separation 4). The published profile spends a separation of 48 where 16 already saturates the plateau; in the precise sense of our measurements, the surplus 32 bits are security-neutral.

The most informative single entry of Table 3 is, in this light, a negative one. The key-space exponent falls from $\log_2 |S| = 3456$ to 2688 — a loss of 768 bits, more than a fifth of the secret space — while measured CVP success moves from 50.00% to 50.01%, that is, not at all within the resolution of the experiment. This dissociation between the cardinality of the secret space and its measured resistance to the reduction-based attack is precisely what the optimizer exploits, and it also indicates where the difficulty of the attack actually resides: not in how many candidates S contains, but in the geometry that the block structure imposes on the target vector relative to the kernel lattice of A . Once the blocks are separated widely enough that their coordinate ranges do not overlap under the union metric, the nearest-plane step recovers a comparable fraction of coordinates irrespective of how much absolute range each block spans. The practical reading for this family is that block geometry should drive parameter selection, while key entropy is best treated as a floor constraint — here satisfied with a very large margin, 2688 bits against a 256-bit requirement — rather than as a proxy for concrete security.

The frontier of Section 4.5 adds a third observation, namely that the optimum is not unique and that the ranking of candidates depends on which attack is taken as binding. The three-block candidate reaches 26,133 bytes (19.9%) and is measurably harder under CVP (46.4%), yet its finer block structure makes it easier under divide-and-conquer (14.8%, against 10.83% for the base profile). A plausible reading is that additional blocks give the CVP formulation a less favourable target geometry while simultaneously handing the divide-and-conquer attack more structure to split on; we report the association as an empirical regularity rather than as a proven mechanism. Because our admissibility gate is defined against the strongest attacker available to us, the two candidates are not comparable by size alone, and we deliberately advance the conservative two-block profile. That the strictly dominant alternative $S^{1/4,3/8}$ still saves 19.1% places the cost of this conservatism at roughly half a percentage point of signature size, which we regard as inexpensive insurance.

Relation to the research objectives

The first objective was to establish an independently reproducible baseline for both the attacks and the size accounting of. Table 2 meets this objective: the homogeneous and structured spaces of the original Table 1 are reproduced exactly for $m \leq 20$, and the six-block space of the original Table 2 to within 1.4 percentage points. This is not a presentational detail. Every size claim we make is conditional on the gate that admits a candidate profile, and the gate is only as trustworthy as the oracle behind it; without the reproduction step, a 19.6% saving would rest on an unvalidated attack implementation and would be uninterpretable.

The second objective was to replace hand-selection of the block bit-profile with an explicit optimization problem. Problem 1 discharges this by encoding soundness, completeness, commitment unpredictability, key entropy and empirical attack hardness as constraints over the bit-profile, and Algorithm 1 by searching that space under a hardness oracle. The results bear out the value of the formalization in a specific way: the two savings the optimizer found are of different kinds — one follows analytically from

the unpredictability constraint once it is written down, the other is visible only empirically through Figure 2 — and a purely analytic parameter search would have recovered the first but not the second.

The third objective was to answer, concretely, the open question left by Rizos and Draziotis, who conjectured that a more careful choice of parameters might shorten the signature. The answer is affirmative and quantified: 19.6% at $t = 80$, $n = 48$, $m = 4$, $R = 192$, with soundness, commitment unpredictability and the completeness error identical by construction and with measured attack hardness statistically indistinguishable from the baseline. It is equally informative that the best candidate anywhere in our search reaches only 19.9%. Within a fixed (n, m, t) and a two- or three-block structure, the slack in the published parameters therefore appears to be close to exhausted, and further substantial compaction in this family will likely require changing the ambient dimensions or the challenge encoding rather than the bit profile.

Comparison with previous studies

The natural point of comparison is the base scheme itself. Relative to the present parameter set is 19.6% smaller under an identical security model, an identical hardness assumption and an unchanged protocol; the improvement is therefore additive to, rather than in competition with, the original construction. It is worth stating what this comparison does not establish. We do not claim a stronger scheme, only a tighter one: every security quantity we can measure is equal to the baseline, and none is improved beyond the divide-and-conquer figure, which we discuss below as a difference within sampling error.

Against the structured-lattice Fiat–Shamir family the comparison is instructive precisely because it is unfavourable. ML-DSA signatures occupy between 2,420 and 4,595 bytes across security levels, so the optimized compact-knapsack signature remains between roughly six and eleven times larger, and Falcon (FN-DSA) is smaller still. Recent compaction work inside that family — HAETAE, which reports up to 39% shorter signatures than Dilithium through bimodal distributions and a hyperball sampler, and the G+G construction based on convolved Gaussians — achieves its savings by changing the distribution from which the response is drawn. That route is unavailable here: the compact-knapsack response is an integer vector constrained to a structured set, not a sampled Gaussian, and there is no rejection-sampling distribution to reshape. Our route, hardness-gated allocation of bits across blocks, is orthogonal to theirs and exploits a degree of freedom that the Gaussian-response schemes do not possess. The consequence for the reader is that this work should not be read as an attempt to close the size gap with ML-DSA. Its value lies in the diversity-of-assumptions argument: if a compact-knapsack signature is to be held in reserve against a structural break in Module-LWE/SIS, then the cost of holding it in reserve is what has been reduced, by a factor that materially affects whether the primitive fits inside existing transport budgets.

Against the older knapsack lineage the results are reassuring but should not be over-read. Merkle–Hellman and its successors[21] fell because low-density subset-sum instances collapse under lattice reduction, and modern implementations of those attacks recover keys in under a second. Our measurements are consistent with the standard view that the compact, matrix-defined formulation lies outside that regime: across the whole admissible region the attack does not break the scheme, it plateaus at 50% recovery under the union metric. The sweep of Figure 2 nevertheless shows how near the cliff edge lies — success reaches 71.8% once the blocks are only 4 bits apart — which is a direct argument for measuring rather than assuming the margin. This is the same lesson drawn from the continuing refinement of density thresholds for subset sum and from faster practical reduction such as iterated compression, both of which move the secure frontier without warning; a parameter set that was chosen by hand against a fixed estimate has no natural way to be revisited, whereas the pipeline reported here can be re-run against an improved oracle.

Finally, on parameter-selection practice: structured-lattice schemes are routinely parameterized against published cost estimators for the underlying reduction attacks, but we are not aware of a comparable estimator for compact knapsack, and the base scheme selects S and S' by hand. The contribution of the present work at this level is to substitute measurement for estimate within a stated scope — a weaker instrument than an analytic estimator, but one that is available now and that is falsifiable, since the gate is re-runnable by any reader.

Theoretical implications

Three implications follow, each stated at the level of confidence the evidence supports. First, and most directly, the experiments suggest that within the attack scope considered here the security-relevant descriptor of a structured secret space is the ordered set of gaps between its block exponents rather than the exponents themselves or their sum. This is a testable conjecture: profiles with equal gap vectors but different absolute exponents should exhibit equal attack success, and we can confirm it only at the two points our sweep visits at a fixed $A_1 = 48$. Establishing or refuting it over a wider grid would convert an empirical gate into a structural statement about the family.

Second, the plateau sits at exactly 50% rather than tapering, and the same value appears in the base paper's own results for structured spaces. A value this sharp invites an analytic explanation: a proof that the union-metric recovery rate of the nearest-plane attack against a well-separated two-block space is exactly one half would replace the empirical hardness oracle in Problem 1 with a closed-form constraint, and would remove the principal source of scope-dependence in our conclusions. We record the regularity; we do not prove it.

Third, the divergence between the CVP and divide-and-conquer rankings of the three-block candidate has a structural consequence for how parameter selection in this family should be posed. The admissible region is defined by a maximum over attack families, so it is not in general convex, and a candidate that improves monotonically along one attack axis may cross the gate along another. This is the formal justification for the discrete, gated search of Algorithm 1 in preference to a continuous relaxation of the bit allocation, and it will apply to any knapsack-type scheme whose security must be argued against a heterogeneous attack set. Independently of these three points, the fact that no protocol step is altered means the completeness, special-soundness, S-HVZK and EUF-CMA arguments of apply verbatim at the new parameters: the security argument is modular in the bit profile, and any subsequent analysis of the base scheme — a QROM treatment, for instance — would transfer to the optimized parameters without re-derivation.

Practical implications

For a deployment, the saving is 6,400 bytes on every signature. Where signatures travel in handshakes or certificate chains this scales with chain depth, and on links with small transport units it changes the fragment count directly: at a 1,280-byte payload the base signature spans 26 fragments and the optimized one 21, a reduction that matters most in exactly the constrained settings — low-power wide-area links, embedded provisioning, datagram-based transports — where a 31.9 KB signature is otherwise prohibitive. We are careful not to overstate the scope of the improvement: it is a reduction in transmitted and stored bytes only. Verification performs the same integer dot products over a smaller response matrix, and our size model does not measure runtime, so we make no claim of a speed improvement and none should be inferred.

For a practitioner choosing among our candidates, the recommendation is ordered by risk appetite. $S = S1/4, 1/3$ with $S' = S'1/8, 1/8$ is the default: the largest saving that stays on the hardness plateau under both attacks. Deployments that require every measured quantity to be at least as good as the published baseline should prefer the strictly dominant $S1/4, 3/8$, at 19.1%. The three-block profile is appropriate only where divide-and-conquer can be argued out of scope, which we do not recommend as a default posture.

Because none of these choices alters the protocol, migration between them is a parameter change rather than a redesign, and the security model, the assumption and the existing proofs are untouched.

The methodological implication may outlast the specific numbers. The pipeline — reproduce the published attacks, express the parameter constraints explicitly, then gate a discrete search on measured hardness — is inexpensive at the scale reported here (200 CVP and 100 divide-and-conquer trials per candidate profile) and is not specific to compact knapsack. Any post-quantum candidate whose concrete security is argued one parameter set at a time, and for which no cost estimator exists, is amenable to the same treatment; and because the whole pipeline is scripted and seeded, revisiting a parameter set after an advance in lattice reduction becomes a re-run rather than a new study.

Threats to validity and limitations

Our security-parity claims are statistical, and should be read as such. With 200 CVP trials the standard error at a success rate near 50% is approximately 3.5 percentage points, and with 100 divide-and-conquer trials the standard error near 11% is approximately 3.1 points. The CVP figures of 50.00% and 50.01% are therefore indistinguishable at this sample size, which supports parity but does not establish equality; and the divide-and-conquer figures of 10.83% and 10.50% differ by well under one standard error, so the honest reading is that the optimized profile is no worse under that attack rather than measurably harder. We have kept the wording of Table 3 conservative on this point. Larger trial counts, or a paired design that reuses instances across profiles, would tighten these intervals and are the obvious next experimental step.

The completeness check is subject to the same caution. Observing zero failures in 2×10^5 samples bounds the failure probability above by roughly 1.5×10^{-5} at 95% confidence by the rule of three. This is consistent with the analytic prediction of 2.15×10^{-6} but is an order of magnitude too coarse to confirm it; resolving the predicted rate directly would require a substantially larger run or an importance-sampling design.

Beyond sampling, four scope limitations bound the conclusions.

We have kept the comparison fair on purpose. (i) The improvement is a parameter optimization, not a new protocol. Completeness, soundness, S-HVZK, and the EUF-CMA proof of all carry over unchanged: S stays in the same geometric class, and every constraint is already met at the reference level. (ii) The matched commitment-unpredictability level is 2^{-24} , the value chosen by the base scheme; applications wanting a higher level can raise κ , which our model and optimizer handle directly at a quantified size cost. (iii) We fix n, m, t to make hardness comparable; jointly optimizing these, and richer (non-binary) challenge encodings, are natural extensions. (iv) Our hardness oracle uses the same attacks as the base paper; while it reproduces their published data faithfully, it inherits their scope, and concrete security against future improved reductions should be re-evaluated as the state of the art advances—exactly the kind of re-evaluation that our reproducible pipeline is designed to make routine.

4. Discussion

Two further limitations lie outside that list and should be stated explicitly. Our hardness oracle is built on a floating-point lattice-reduction engine, so its outputs inherit that engine's numerical behaviour; agreement with the published tables gives us confidence in the reproduction but does not eliminate the dependency. And the present work addresses signature size only: we provide no constant-time reference implementation and no side-channel evaluation, and the literature on implementation attacks against deployed lattice schemes indicates that such an evaluation is a prerequisite for any real deployment rather than an optional refinement. Both are identified as immediate future work.

5. Conclusion

We revisited the compact-knapsack Fiat–Shamir signature of Rizos and Draziotis and answered its open question on parameter selection. By reproducing the scheme's lattice attacks and size model faithfully, formalizing parameter choice as a constrained optimization, and driving it with an adaptive optimizer that allocates bits only where measured attack hardness demands them, we obtained a parameter set that is *19.6% smaller at identical security*. The two ingredients of the saving—an over-provisioned nonce space and an under-tightened secret space—are removed without touching the security model or introducing any new assumption, and every claim is backed by reproducible measurement, including a 2×10^5 -sample completeness check. Beyond the specific numbers, the main contribution is methodological. Empirical, hardness-gated parameter optimization is a recipe that transfers well: it applies to compact knapsack and, more broadly, to any knapsack-type post-quantum signature whose security has to be argued one parameter set at a time. Several directions remain open. These include jointly optimizing, moving to non-binary challenge spaces, building a constant-time reference implementation, and tracking the size–security frontier as lattice attacks grow stronger. Reproducibility. All the scripts that support this paper are included: the attack reproduction, the size model, the adaptive optimizer, the Monte-Carlo completeness test, and the figure-generation code. With the seeds fixed, they reproduce every value we report.

REFERENCES

- [1] National Institute of Standards and Technology, *Module-Lattice-Based Digital Signature Standard*, Federal Information Processing Standards (NIST FIPS) 204, U.S. Dept. Commerce, Gaithersburg, MD, USA, Aug. 13, 2024, doi: 10.6028/NIST.FIPS.204.
- [2] National Institute of Standards and Technology, “Announcing issuance of Federal Information Processing Standards (FIPS) FIPS 203, Module-Lattice-Based Key-Encapsulation Mechanism Standard, FIPS 204, Module-Lattice-Based Digital Signature Standard, and FIPS 205, Stateless Hash-Based Digital Signature Standard,” *Federal Register*, vol. 89, no. 157, Aug. 14, 2024.
- [3] S. Darzi, K. Ahmadi, S. Aghapour, A. A. Yavuz, and M. Mozaffari Kermani, “Envisioning the future of cyber security in post-quantum era: A survey on PQ standardization, applications, challenges and opportunities,” arXiv preprint arXiv:2310.12037, 2023.
- [4] C. Zong, “The mathematical foundation of post-quantum cryptography,” arXiv preprint arXiv:2404.19186, 2024.
- [5] J. Bi, L. Su, H. Peng, and L. Wang, “A simple and efficient attack on the Merkle–Hellman knapsack cryptosystem,” *PLoS One*, vol. 20, no. 5, Art. no. e0322726, 2025, doi: 10.1371/journal.pone.0322726.
- [6] *Improving the Lagarias–Odlyzko algorithm for average-case subset sum: A modular arithmetic approach*, arXiv preprint arXiv:2408.16108, 2024.
- [7] K. Ryan and N. Heninger, “Fast practical lattice reduction through iterated compression,” in *Advances in Cryptology—CRYPTO 2023*, Part I, Lecture Notes in Computer Science, vol. 14081, Cham, Switzerland: Springer, 2023, pp. 3–36.
- [8] G. S. Rizos and K. A. Draziotis, “Cryptographic primitives based on compact knapsack problem,” arXiv preprint arXiv:2303.08973, 2023.
- [9] A. Fiat and A. Shamir, “How to prove yourself: Practical solutions to identification and signature problems,” in *Advances in Cryptology—CRYPTO ’86*, Lecture Notes in Computer Science, vol. 263, Berlin, Germany: Springer, 1987, pp. 186–194, doi: 10.1007/3-540-47721-7_12.
- [10] M. Abdalla, J. H. An, M. Bellare, and C. Namprempe, “From identification to signatures via the Fiat–Shamir transform: Minimizing assumptions for security and forward-security,” in *Advances in Cryptology—EUROCRYPT 2002*, Lecture Notes in Computer Science, vol. 2332, Berlin, Germany: Springer, 2002, pp. 418–433, doi: 10.1007/3-540-46035-7_28.
- [11] V. Lyubashevsky, “Lattice signatures without trapdoors,” in *Advances in Cryptology—EUROCRYPT 2012*, Lecture Notes in Computer Science, vol. 7237, Berlin, Germany: Springer, 2012, pp. 738–755.

- [12] L. Ducas, E. Kiltz, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé, "CRYSTALS-Dilithium: A lattice-based digital signature scheme," *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, vol. 2018, no. 1, pp. 238–268, 2018, doi: 10.13154/TCHES.V2018.I1.238-268.
- [13] J. H. Cheon, H. Choe, J. Devevey, T. Güneysu, D. Hong, M. Krausz, G. Land, M. Möller, D. Stehlé, and M. Yi, "HAETAE: Shorter lattice-based Fiat-Shamir signatures," *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, vol. 2024, no. 3, pp. 25–75, 2024, doi: 10.46586/tches.v2024.i3.25-75.
- [14] J. Devevey, A. Passelègue, and D. Stehlé, "G+G: A Fiat-Shamir lattice signature based on convolved Gaussians," in *Advances in Cryptology – ASIACRYPT 2023*, Lecture Notes in Computer Science, vol. 14305, Singapore: Springer, 2023, pp. 37–64, doi: 10.1007/978-981-99-8739-9_2.
- [15] F. Liu, Z. Zheng, Z. Gong, K. Tian, Y. Zhang, Z. Hu, J. Li, and Q. Xu, "A survey on lattice-based digital signature," *Cybersecurity*, vol. 7, Art. no. 7, 2024, doi: 10.1186/s42400-023-00198-1.
- [16] K. A. Draziotis and A. Papadopoulou, "Improved attacks on knapsack problem with their variants and a knapsack type ID-scheme," *Adv. Math. Commun.*, vol. 12, no. 3, pp. 429–449, 2018, doi: 10.3934/amc.2018026.
- [17] D. Micciancio and O. Regev, "Lattice-based cryptography," in *Post-Quantum Cryptography*, D. J. Bernstein, J. Buchmann, and E. Dahmen, Eds. Berlin, Germany: Springer, 2009, pp. 147–191, doi: 10.1007/978-3-540-88702-7_5.
- [18] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász, "Factoring polynomials with rational coefficients," *Math. Ann.*, vol. 261, no. 4, pp. 515–534, 1982.
- [19] S. D. Galbraith, *Mathematics of Public Key Cryptography*. Cambridge, U.K.: Cambridge University Press, 2012, doi: 10.1017/CBO9781139012843.
- [20] K. Wang, D. Xu, and J. Tian, "An improved two-step attack on lattice-based cryptography: A case study of Kyber," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 44, no. 9, pp. 3643–3647, 2025, doi: 10.1109/TCAD.2025.3550443.
- [21] R. C. Merkle and M. E. Hellman, "Hiding information and signatures in trapdoor knapsacks," *IEEE Trans. Inf. Theory*, vol. 24, no. 5, pp. 525–530, Sep. 1978, doi: 10.1109/TIT.1978.1055927.