REIST Division: An Implementation-Oriented Framing of Centered Remainder Arithmetic for Modular Addition

Version 2.0 – Extended hardware evaluation on ARMv8-A and x86-64

Rudolf Stepan
2025
ORCID: 0009-0004-2842-2579

Abstract

Traditional integer division constrains the remainder to be non-negative, an assumption that simplifies arithmetic but introduces asymmetries and inefficiencies. The idea of using centered remainder intervals is classical in number theory and widely used in FFT/NTT implementations. The contribution of REIST in this report is therefore not a new algebraic object, but an implementation-oriented framing and a systematic treatment of the remainder as an explicit signed correction term, together with an empirical evaluation on modern CPUs.

This work presents REIST Division (Remainder-Extended Inversion and Subtraction Technique) as a generalized framework that allows bounded negative remainders and interprets the remainder as a signed correction term rather than a passive residue. The formalism yields a centered remainder interval, aligns naturally with balanced modular arithmetic and rounding functions, and provides a convenient language for modeling error-feedback and phase-alignment processes in logistics, scheduling, digital signal processing, robotics and CPU micro-architectural correction loops.

Beyond the conceptual framework, this Version 2.0 report includes an extended empirical evaluation of REIST in cryptographic workloads. Using an open-source benchmark suite, we compare classical modulo-based reductions with REIST-style symmetric remainders on both ARMv8-A (aarch64) and x86-64 (Intel Core i9-14900K).

We compare REIST-based symmetric remainder updates against the classical % operator implementation provided by the C++ compiler. The evaluation therefore quantifies speedups at the level of the basic arithmetic primitive, not against fully optimized Montgomery- or Barrett-style reductions used in state-of-the-art cryptographic libraries.

On ARM, REIST achieves about 2.5× speedup for synthetic modular-add counters and roughly 6× speedup for polynomial modular addition, the dominant workload in lattice-based cryptography. On x86-64, the same patterns yield around 9× speedup for modular addition and up to 15× speedup for polynomial modular addition, while full remainder computations and ARX ciphers remain essentially unchanged and hash-mixers exhibit slight slowdowns.

Together, these results show that REIST Division is not only a mathematically natural extension of classical division, but also a practically relevant primitive for high-performance modular arithmetic across two distinct CPU architectures.

TABLE OF CONTENTS

Abstract 1
TABLE OF CONTENTS 2
1 Introduction 6
1.1 Domain and intended use of REIST 6
1.1.1 The Acronym 7
1.1.2 The Mathematical Idea behind 7
2 Theoretical Framework 8
2.1 Classical Division 8
2.2 Definition of REIST Division 9
2.3 Mathematical Properties 9
2.3.1 Uniqueness 9
2.3.2 Relation to Balanced Modulo 10
2.3.3 Relation to Rounding 10
2.3.4 Error-Minimizing Quotient 10
2.4 Examples 11
2.4.1 Example 1 – Error Minimization 11
2.4.2 Example 2 – Periodic Alignment 11
3 Applications 12
3.1 Supply Chain & Logistics 12
3.2 Cyclic Systems & Scheduling 12
3.3 Computational Resource Allocation 13
3.4 Forecasting and Predictive Modeling 13
3.5 Digital Signal Processing 13
3.6 Robotics and Control Engineering 13
3.7 CPU Architecture and Numerical Hardware 13
4 Discussion 14
4.1 Open research directions 14
5 Conclusion 14
6 Real Hardware Tests on ARMv8-A and x86-64 15
6.1 Modular Add Benchmark (Synthetic Counter Workload) 15
6.2 Modular Remainder Benchmark 16
6.3 ChaCha20-Style ARX Workloads 16
6.4 Hash-Mixing Workloads 16
6.5 ARM Scalar vs NEON SIMD Evaluation 17

6.6 Summary of ARM Findings ... 17
6.7 Summary of hardware findings ... 17
6.8 x86-64 Evaluation (Intel Core i9-14900K) ... 20
6.8.1 Modular Addition Suite ... 20
6.8.2 Polynomial Modular Addition (Cryptographic Workload) ... 20
6.8.3 Full Remainder Benchmark ... 21
6.8.4 ChaCha20-Style ARX Workloads ... 21
6.8.5 Hash-Mixing Workloads ... 21
6.8.6 Summary of x86-64 Findings ... 21
7 Community feedback, review history and clarifications ... 21
7.1 Lessons learned ... 22
7.2 Review context and scope ... 22
7.3 Novelty and relation to existing concepts ... 23
7.4 Misunderstandings about the scope of the “division change” ... 23
7.5 Benchmark interpretation and performance claims ... 24
7.6 Cryptographic correctness and security considerations ... 24
7.7 Lessons learned and future community interaction ... 25
8 Glossary ... 25
8.1 REIST Division ... 25
8.2 Quotient ... 25
8.3 Remainder (classical) ... 25
8.4 Centered remainder (symmetric remainder) ... 26
8.5 Quotient correction ... 26
8.6 Modulus $\boldsymbol{B}$ ... 26
8.7 Modulo operation (%) ... 26
8.8 Modular addition ... 26
8.9 Polynomial modular addition ... 26
8.10 Centered modular representation ... 26
8.11 Workload ... 26
8.12 Speedup ... 27
8.13 ARMv8-A ... 27
8.14 x86-64 ... 27
8.15 Advanced SIMD (NEON) ... 27
8.16 ARX cipher ... 27
8.17 Hash mixing function ... 27

8.18 Timer artefact ... 27
8.19 Framework (in this context) ... 27
9 References ... 28
9.1 Literature ... 28
9.2 Repository – Benchmark Sources ... 28
10 About the Author ... 28
Appendix ... 29
Benchmarks ... 29
11 REIST Cryptographic Benchmark Report (ARM) ... 29
11.1 System Information ... 29
11.2 Executive Summary ... 29
11.3 Performance Overview ... 29
11.4 Modular Addition Suite ... 30
11.4.1 Results: O0 (No Optimization) ... 30
11.4.2 Results: O3 (Optimized) ... 30
11.5 Polynomial Modular Addition ... 30
11.5.1 Results: O0 (No Optimization) ... 31
11.5.2 Results: O3 (Optimized) ... 31
11.6 Modular Remainder Operations ... 31
11.7 ChaCha20 Cipher Benchmarks ... 31
11.7.1 ChaCha20 Stream Generation ... 32
11.8 Hash-Mix Operations ... 32
11.8.1 Results: O0 vs O3 Comparison ... 32
11.9 Compiler Artifact Analysis (Assembly Inspection) ... 32
11.10 Conclusions ... 33
11.10.1 Key Findings ... 33
11.10.2 Recommendations ... 33
12 REIST Cryptographic Benchmark Report (X86) ... 33
12.1 System Information ... 33
12.2 Executive Summary ... 33
12.3 Performance Overview ... 34
12.4 Modular Addition Suite ... 34
12.4.1 Results: O0 (No Optimization) ... 34
12.4.2 Results: O3 (Optimized) ... 35
12.5 Polynomial Modular Addition ... 35

12.5.1 Results: O0 (No Optimization) 35
12.5.2 Results: O3 (Optimized) 36
12.6 Modular Remainder Operations 36
12.7 ChaCha20 Cipher Benchmarks 36
12.7.1 ChaCha20 Stream Generation 36
12.8 Hash-Mix Operations 36
12.8.1 Results: O0 vs O3 Comparison 36
12.9 Compiler Artifact Analysis (Assembly Inspection) 36
12.10 Conclusions 37
12.10.1 Key Findings 37
12.10.2 Recommendations 37

1 Introduction

1.1 Domain and intended use of REIST

In addition to its conceptual appeal, REIST Division admits direct implementation in low-level integer code. The second half of this work therefore treats REIST not merely as a mathematical object, but as a concrete arithmetic primitive. We report measurements from an open-source benchmark suite that implements REIST side-by-side with classical modulo operations and evaluates them on both ARMv8-A and x86-64 processors to test whether the theoretical advantages translate into real hardware performance.

While many examples in this report are drawn from lattice-based cryptography and modular arithmetic, REIST Division is not a cryptography-specific construction. Its core domain is integer arithmetic with signed remainders and the modelling of cyclic and feedback-driven processes where “too much” and “too little” are symmetric correction states. Cryptographic workloads are used here primarily as a stress test and as a convenient source of high-volume modular operations, not as the exclusive or primary application domain of REIST.

In high-performance cryptographic libraries, modular multiplication is typically accelerated by Montgomery or Barrett reduction. The benchmarks in this report do not compete with such schemes directly; they explore how a signed-remainder correction rule behaves in generic modular workloads, including cryptographic ones. From this perspective, REIST should be viewed as a general arithmetic and modelling tool whose behavior under cryptographic workloads has been sanity-checked, rather than as a replacement for Montgomery or Barrett reduction. For multiplication-heavy workloads, REIST should be viewed as complementary rather than as a drop-in replacement for Montgomery or Barrett reduction: it can provide a centered remainder representation and efficient modular addition in combination with existing multiplication schemes but does not by itself eliminate the need for specialized reduction algorithms after multiplication.

The remainder of this report is organized as follows. Sections 2–4 introduce the formal framework of REIST Division and its interpretation as a signed correction mechanism. Section 5 summarizes the main conceptual and practical conclusions. Section 6 presents real hardware benchmarks on ARMv8-A and x86-64, evaluating REIST in cryptographic and non-cryptographic workloads. Section 7 documents the informal community review history, including initial misunderstandings and how the resulting feedback shaped the clarifications and extended evaluation in Version 2.0. Readers mainly interested in performance and review context may skim Sections 2–4 and focus on Sections 6 and 7.

REIST is, in the strict mathematical sense, still a division. We start from an existing quantity and split it into an integer multiple of the modulus plus a remainder, and we are still asking how many “whole parts” fit into a given total. The difference to the textbook convention is not that we abandon division, but that we adopt a different and implementation–oriented way of realizing this split. Instead of forcing the remainder into the non-negative interval $[0, B)$, we choose a centered interval $[-B/2, B/2)$ and treat the remainder as a signed correction term. This change of viewpoint leads to a concrete implementation strategy that, in many workloads, is superior to the standard modulo based variant emitted by compilers. In other words, REIST does not replace division with an unrelated trick. It rephrases the same quotient–remainder relationship in a way that maps more directly onto additions, comparisons and conditional corrections, which modern CPUs can execute much more efficiently than general integer division.

REIST is an alternative quotient–remainder convention (centered representative), not a new algebraic structure. It uses the same fundamental decomposition $\boldsymbol{T = q \cdot B + r}$, but chooses a symmetric remainder interval and a signed correction rule that modern CPUs can execute much more efficiently than the standard %-based implementation in many modular addition workloads.

1.1.1 The Acronym

The acronym REIST (**Remainder Extended Inversion and Subtraction Technique**) is not intended to suggest a fundamentally new kind of arithmetic. Instead, it provides a clean and compact label for a specific, implementation-oriented way of viewing integer division with centered remainders. By giving this pattern a name, we make it easier to talk about, to reference it in code and documentation, and to distinguish it from both the textbook % implementation and from other well-known techniques such as Montgomery or Barrett reduction. In that sense, REIST is a systematic wrapper around a familiar idea, not a claim of mathematical novelty.

The novelty of the name **REIST** is not that it renames symmetric division with centered remainders, which is mathematically well known, but that it makes the *role of the remainder* explicit. In REIST, the remainder is not a passive residue that happens to fall into a chosen interval; it is a signed correction term that actively adjusts the quotient. The "Remainder Extended Inversion and Subtraction Technique" label encodes exactly this viewpoint: we start from a nominal quotient, and the centered remainder expresses how far we must move the result up or down in discrete steps of B. To the best of my knowledge, this quotient–adjustment perspective, with a formally defined update rule and an implementation-oriented framework around it, has not been developed in this explicit form before.

1.1.2 The Mathematical Idea behind

The classical division identity for integers

$$T = Q \cdot B + R, 0 \leq R < B$$

fixes the remainder to a non-negative value. This constraint simplifies basic arithmetic but embeds an *asymmetry* that is suboptimal in many real-world settings:

- cyclic systems
- scheduling and phase-alignment tasks
- iterative numerical algorithms
- forecasting approaches
- distributed computing
- resource-balancing systems

In such domains, the remainder is better interpreted as a **signed deviation**, not as a leftover quantity. Classical division ignores the possibility that a *negative adjustment* can be the most efficient corrective action.

REIST Division relaxes the non-negativity constraint and redefines the allowable remainder interval so that:

- remainders represent *direction* and *magnitude* of deviation
- quotient and remainder jointly encode the correction
- adjustments become locally optimal in cyclic contexts

This yields a more symmetrical, flexible representation of division.

2 Theoretical Framework

2.1 Classical Division

Given integers $T$and $B > 0$, classical division provides unique integers Q, Rsatisfying:

$$T = QB + R, 0 \leq R < B$$

This forces all deviations upward. If a value is "just below" a boundary, classical division overshoots rather than undershoots.

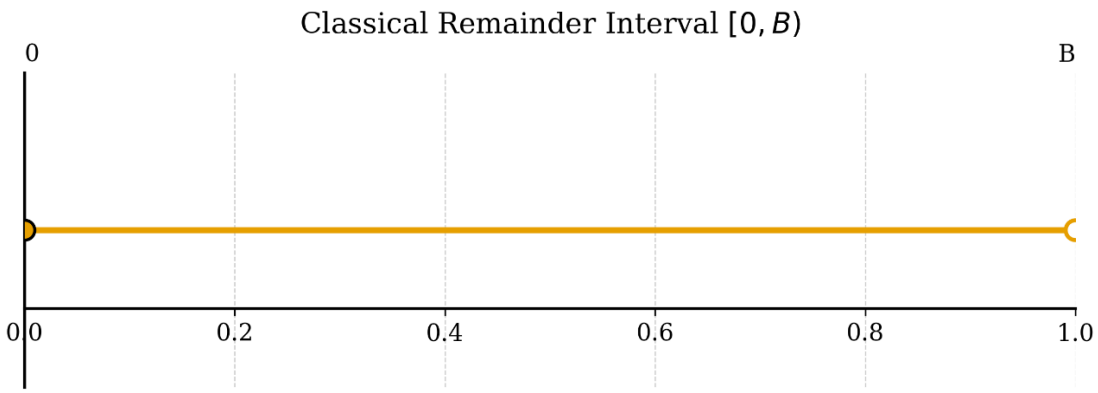

Figure 1. Classical remainder interval [0, B).
This visualization illustrates that classical division does not allow negative corrections, producing an inherent structural asymmetry.

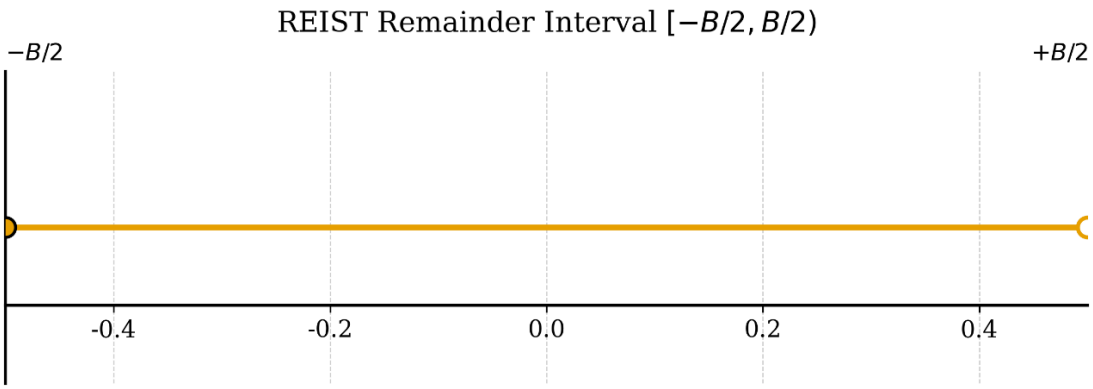

Figure 2. Classical and REIST remainder intervals.
This illustration contrasts with the classical remainder interval $[0, B)$, which restricts all corrections to positive values, with the REIST interval $-\frac{B}{2} \leq r < \frac{B}{2}$, which allows symmetric positive and negative adjustments and removes the inherent asymmetry of classical division.

2.2 Definition of REIST Division

We define the REIST remainder of an integer T with respect to a positive modulus B as the unique integer r with

$$T = qB + r, q \in \mathbb{Z}, -\frac{B}{2} \leq r < \frac{B}{2}$$

In other words, REIST uses a centered, half open remainder interval $[-B/2, B/2)$.

A convenient implementation is

$$r = ((T + B/2) mod B) - B/2$$

which guarantees that r always lies in $[-B/2, B/2)$ and that the value $+B/2$ never occurs even when B is even.

Interpretation:

- $\boldsymbol{R > 0}$: the system is leading
- $\boldsymbol{R < 0}$: the system is lagging
- $\boldsymbol{R = 0}$: perfect alignment

Unlike balanced modulo arithmetic, REIST feeds the signed remainder back into the quotient selection, ensuring the chosen quotient minimizes deviation.

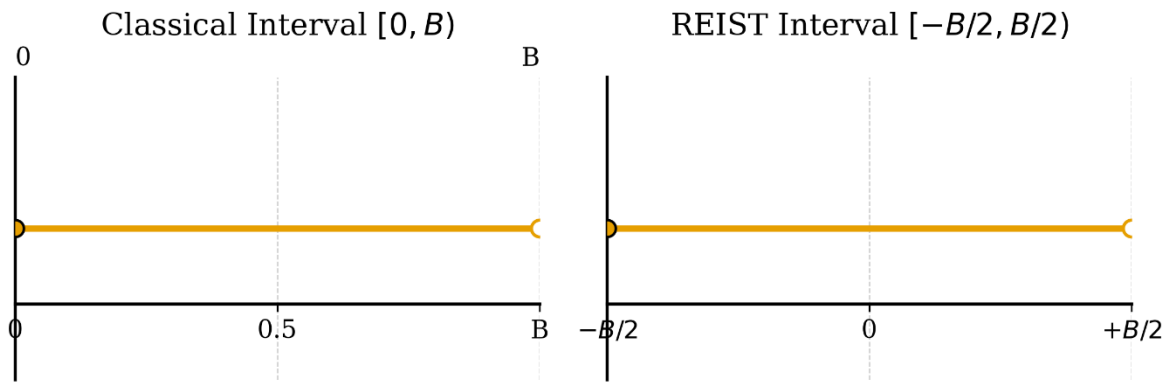

Figure 3. ***Centered remainder convention used by REIST***
The remainder is mapped into $[-B/2, B/2)$ instead of $[0, B)$, yielding a symmetric, signed representative around zero. For even B, the tie-break is fixed to $-B/2$, so the value $+B/2$ is excluded.

2.3 Mathematical Properties

2.3.1 Uniqueness

- For odd B: unique remainder
- For even B: two midpoint candidates; a deterministic tie-break rule assign $R = -B/2$

Remark (tie breaking for even moduli)

For odd moduli B, the centred interval $[-B/2, B/2)$ already yields a unique remainder.
For even B, the residue class of $T \equiv B/2(\mathrm{mod} B)$ has two symmetric representatives, $+B/2$ and $-B/2$.

In REIST we resolve this ambiguity by construction: the remainder is always taken in [-B/2, B/2), so +B/2 never appears and -B/2 is chosen consistently as the representative of that class.

2.3.2 Relation to Balanced Modulo

Balanced modulo returns a signed remainder but does not adjust the quotient.
REIST integrates both in a single step.

2.3.3 Relation to Rounding

For odd B:

$$Q = \lfloor \frac{T}{B} + \frac{1}{2} \rfloor$$

So REIST quotient selection is equivalent to nearest integer rounding.

2.3.4 Error-Minimizing Quotient

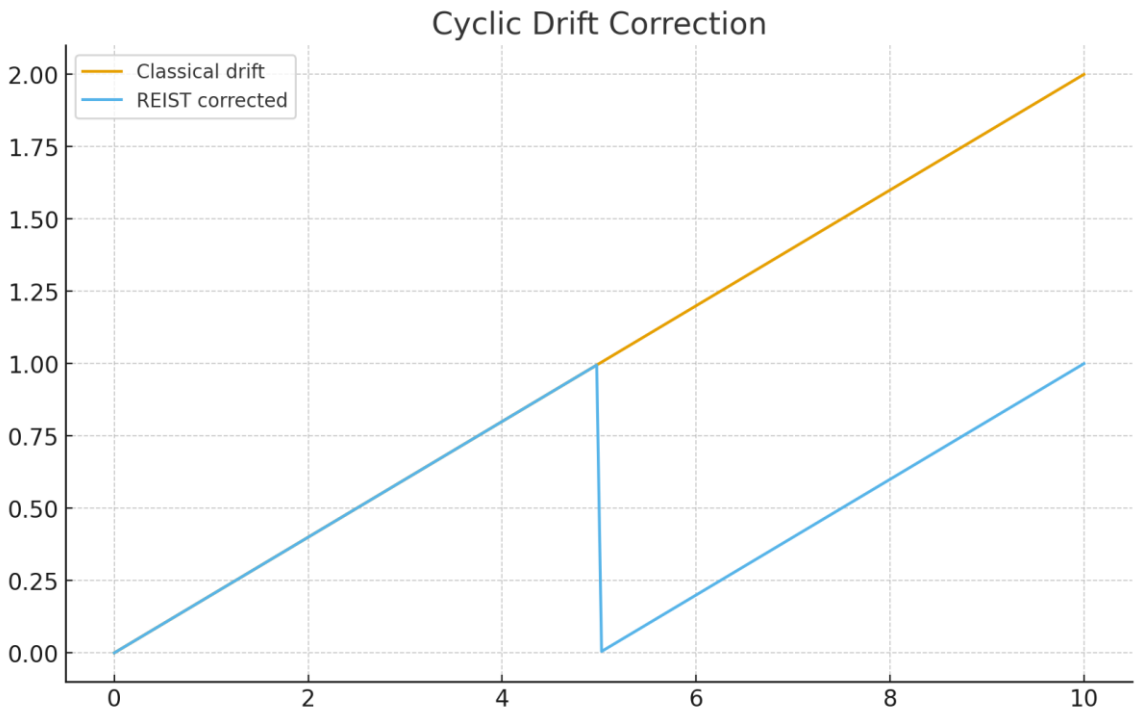

Figure 4. Cyclic drift behavior under classical division versus REIST correction.
This diagram illustrates how classical division accumulates drift linearly over time, while REIST Division resets the deviation using a signed correction. The REIST approach minimizes the remainder magnitude and prevents drift growth, leading to more stable behavior in cyclic or periodic systems.

REIST chooses the quotient that minimizes | R |

This is optimal for drift reduction in:

- signal synchronization
- scheduling
- feedback control
- iterative computations

2.4 Examples

2.4.1 Example 1 – Error Minimization

$$T = 17, B = 10$$

Classical:

$$17 = 1 \cdot 10 + 7$$

Remainder = 7 (large deviation) - Large positive deviation.

REIST:

$$17 = 2 \cdot 10 - 3$$

Remainder = −3 (much smaller deviation) - Smaller signed deviation: $R = -3$.

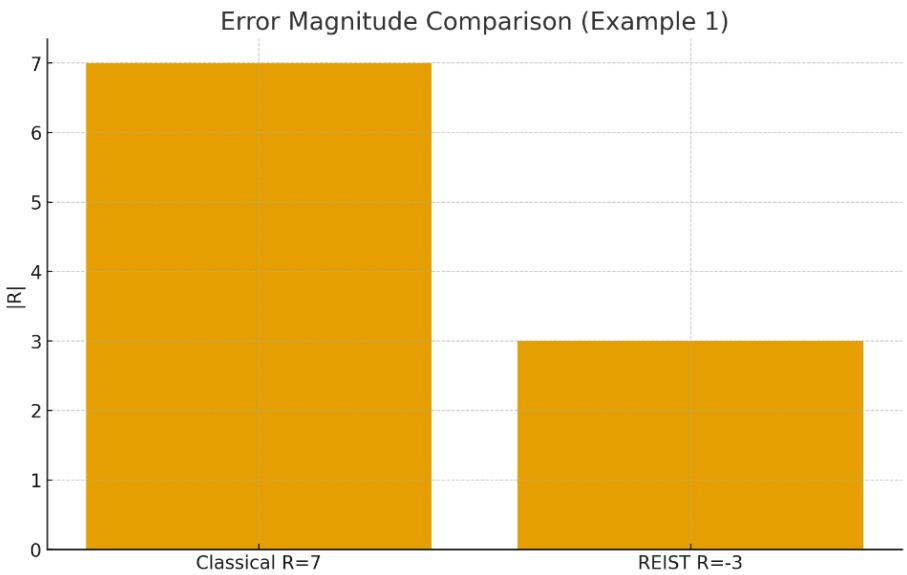

Figure 5. Error magnitude comparison for Example 1 (T = 17, B = 10).
The classical division produces a remainder of $R = 7$, representing a large positive deviation. In contrast, REIST Division yields a signed remainder of $R = -3$, which has a significantly smaller magnitude. This illustrates REIST's core principle: selecting the quotient that minimizes | R |to reduce deviation within a single computational step.

2.4.2 Example 2 – Periodic Alignment

$$T = 28, B = 12$$

Classical: remainder = 4, REIST: remainder = −2 (smaller corrective term)

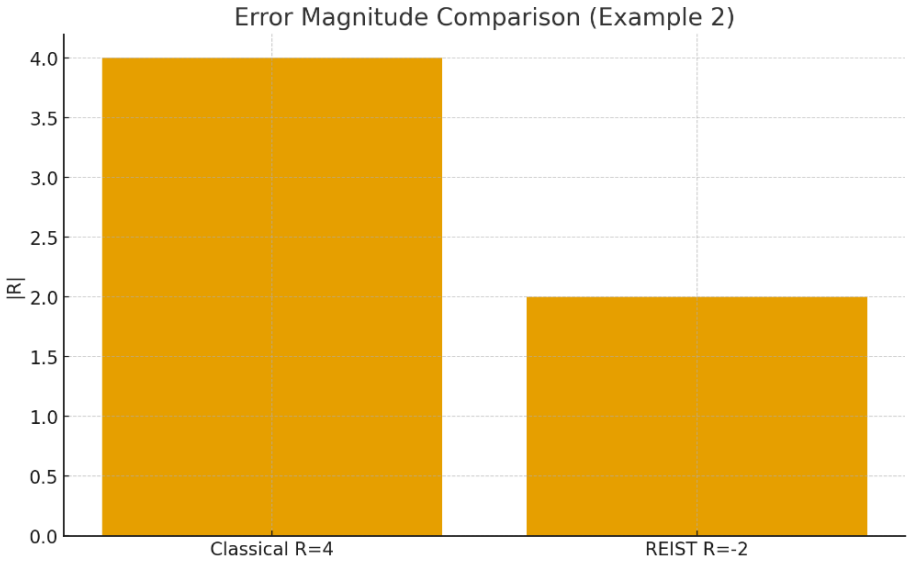

Figure 6. Error magnitude comparison for Example 2 (T = 28, B = 12).
The classical division produces a remainder of $R = 4$, while REIST Division yields a signed remainder of $R = -2$. Although the sign differs, the key point is the reduced magnitude: REIST provides a smaller corrective term, which improves alignment and minimizes deviation in periodic or cyclic systems.

3 Applications

3.1 Supply Chain & Logistics

Negative remainder = negative adjustment in batch size.

Signed corrections enable:

- dynamic batch-size adjustment
- reduced overproduction
- smoother JIT alignment
- symmetric safety stock behavior

3.2 Cyclic Systems & Scheduling

Any periodic activity benefits from direction-aware corrections:

- production cycles
- robotic task loops
- routing and circular queues
- simulation time steps

Negative remainders allow backward adjustments.

3.3 Computational Resource Allocation

In distributed computing:

- positive remainder = over-allocation
- negative remainder = under-allocation

REIST improves symmetric load balancing.

3.4 Forecasting and Predictive Modeling

The remainder acts as an embedded error-correction term, reducing:

- bias accumulation
- drift across iterations
- need for external correction factors

3.5 Digital Signal Processing

Phase errors in DSP are inherently signed. REIST matches:

- PLL alignment
- FFT window synchronization
- carrier phase recovery
- resampling drift correction

3.6 Robotics and Control Engineering

Feedback systems operate on signed error:

- motor control
- joint alignment
- trajectory correction
- PID-based cyclic behavior

REIST offers minimal-error correction in periodic control loops.

3.7 CPU Architecture and Numerical Hardware

Modern CPUs employ signed feedback in:

- division approximation loops
- MAC pipelines
- floating-point correction
- rounding and normalization routines

REIST provides a clean theoretical model for such operations.

4 Discussion

Allowing negative remainders is not a minor modification but an extension comparable to the historical acceptance of:

- negative integers
- complex numbers
- balanced ternary notation
- signed zero in IEEE-754

Each generalization expanded arithmetic to model phenomena that classical constructs could not represent symmetrically. Likewise, REIST Division removes unnecessary asymmetry and aligns integer division with real-world feedback-driven systems.

The empirical results on ARMv8-A and x86-64 support this conceptual picture: whenever division is replaced by a signed correction rule implemented purely in terms of additions, comparisons and subtractions, the algorithms align better with the strengths of modern ALUs and SIMD engines. REIST thus illustrates how a seemingly small change in a textbook identity can have both theoretical and architectural consequences.

Many high-performance cryptographic libraries avoid `%` entirely and implement modular reduction via Montgomery or Barrett techniques. These methods already replace division with multiplication and shifts, and thus partially address the same performance issue that REIST targets at the primitive level. In this Version 2.0 report we deliberately benchmark REIST against the compiler's `%` implementation to expose the raw effect of replacing division by a signed correction rule. A head-to-head comparison with Montgomery and Barrett schemes is outside the present scope. We expect the relative speedups over such optimized baselines to be smaller than over %, and potentially highly workload- and implementation-dependent.

4.1 Open research directions

- formalization within algebraic ring structures
- algorithmic O(1) implementations
- REIST-based numerical stabilization methods
- applications in AI planning and constraint solvers
- hybrid REIST-modulo operators

5 Conclusion

REIST Division reframes the remainder of integer division as a bidirectional correction term rather than a unidirectional residue. By allowing bounded negative remainders and centering the remainder interval around zero, it removes unnecessary asymmetry from classical division and connects naturally to balanced modular arithmetic, rounding functions and feedback-driven error correction.

On the conceptual side, REIST offers a compact language for describing cyclic and discrete processes in logistics, scheduling, digital signal processing, robotics and CPU micro-architecture. The remainder becomes an explicit, signed adjustment variable that can absorb drift, encode phase error and express over- and under-allocation in a symmetric way.

On the practical side, the hardware experiments in this Version 2.0 report demonstrate that REIST is more than a mathematical curiosity. On ARMv8-A, REIST accelerates synthetic modular-add counters by roughly 2.5× and achieves around 6× speedup for polynomial modular addition – the dominant arithmetic kernel in NTRU-, RLWE- and Kyber-style lattice cryptography. On x86-64, the same benchmark suite yields approximately 9× speedup for modular addition and up to 15× speedup for polynomial modular addition under full compiler optimization, while full remainder operations and ARX-based ciphers remain essentially unaffected and hash-mixers show modest slowdowns.

These observations underline a key point: REIST is not a universal micro-optimization, but a targeted tool. It is most effective precisely where modular addition dominates and integer division is a structural bottleneck, and it leaves workloads without modulo arithmetic unchanged. At the same time, its simple, branchless correction rule makes it naturally compatible with SIMD vectorization on architectures such as ARM NEON. A detailed account of community feedback, early misunderstandings, and how they informed the refinements in Version 2.0 is given in Section 7, which documents the informal review history behind the present formulation.

Future work may explore formal algebraic properties of REIST, integrate it into cryptographic libraries and compiler back-ends, investigate hardware implementations and extend the benchmark corpus to additional architectures such as RISC-V and GPU platforms. In this sense, REIST can be viewed as a small but concrete step towards a more symmetric and implementation-aware view of integer division.

Key Advantages

- symmetric remainder domain
- improved numerical stability
- natural alignment with cyclical and feedback systems
- broad applicability in engineering, computation, and modeling

REIST represents a conceptual and practical generalization of division, comparable to earlier mathematical expansions that improved expressiveness and real-world applicability.

6 Real Hardware Tests on ARMv8-A and x86-64

This section summarizes real-world benchmark results for REIST-based arithmetic compared to classical modulo operations. All experiments are implemented in a shared open-source repository and executed on two distinct architectures: an ARMv8-A (aarch64) platform with full NEON SIMD support, and an x86-64 desktop CPU (Intel Core i9-14900K).

All benchmarks were executed using adaptive timing (≥20 MS per measurement) to minimize noise and ensure stable cycle-level behavior. The evaluation covers several classes of workloads: synthetic modular-add counters, polynomial modular arithmetic (as used in lattice-based cryptography), full remainder computations, ARX-style cipher operations, hash-mixers, and a dedicated comparison between scalar and SIMD-vectorized REIST implementations.

6.1 Modular Add Benchmark (Synthetic Counter Workload)

On ARMv8-A, across all tested moduli, REIST achieves a consistent acceleration of approximately **2.5×** compared to the classical modulo update. This improvement reflects the complete elimination of

division in REIST and the branchless nature of the signed correction step but should not be interpreted as indicative of performance in practical cryptosystems.

The first experiment measures the update rule

$$r_{k+1} = (r_k + s) \ \% \ B$$

which is known to be highly susceptible to compiler optimizations due to its predictable structure. While this workload does not represent a real cryptographic scenario, it is useful for characterizing pure ALU-based behavior. The core workload for lattice-based cryptography consists of evaluating

$$c_i = (a_i + b_i) \ \% \ q$$

for many coefficients $a_i, b_i$and large moduli q. This pattern appears in NTRU, Kyber, Dilithium, RLWE-based schemes, and in the pre-/post-processing stages of the Number Theoretic Transform.

On ARMv8-A, REIST achieves a **consistent speedup of approximately 6×** over classical modulo reduction for all tested cryptographic moduli $q \in \{10^6, 10^7, 10^8, 10^9\}$. This result confirms the theoretical expectation: classical % operations incur expensive integer division, whereas REIST reduces the operation entirely to additions, comparisons and subtractions, all of which map efficiently onto the ARM ALU pipeline.

This benchmark represents the most relevant practical use case and demonstrates the core advantage of REIST in real cryptographic workloads.

6.2 Modular Remainder Benchmark

To illustrate the limits of the method, the full remainder computation

$$r = x \ \% \ B$$

was also benchmarked. In this case, REIST and the classical operator show virtually identical performance (difference <2%). This outcome is expected: REIST is explicitly designed to optimize **modular additions**, not full division. The remainder benchmark serves as a correctness reference rather than a target scenario.

6.3 ChaCha20-Style ARX Workloads

ChaCha20 and related ARX ciphers use only addition, XOR and rotation; they do not perform any modulo reductions. As expected, both implementations run at essentially identical speed. This confirms that REIST does not alter or accelerate workloads outside its intended domain and acts only where modulo arithmetic is present.

6.4 Hash-Mixing Workloads

A PRNG-style hash-mixing benchmark involving multiplication, addition and modulo division was included for completeness. REIST performs slightly worse (≈15–20% slower) in this setting. The reason is inherent: hash-mixers rely on multiplicative diffusion steps that do not benefit from REIST's structure, and the signed correction step can interfere with the intended statistical progression. This benchmark therefore demonstrates the expected specialization of REIST.

6.5 ARM Scalar vs NEON SIMD Evaluation

The most significant architectural observation arises when comparing REIST's scalar and SIMD-vectorized implementations. On ARMv8-A, classical modulo reduction is inherently scalar and cannot be vectorized efficiently, while REIST maps naturally to SIMD because its correction rule

$$t = a + b, r = t - q \cdot (t \geq q)$$

is entirely branchless and lane-wise independent.

The NEON implementation processes four coefficients in parallel and achieves almost **2× speedup** over the classical scalar modulo reduction across all tested moduli. In addition, NEON outperforms the scalar REIST variant itself by ~1.86×, indicating that REIST benefits strongly from SIMD parallelism. This SIMD compatibility constitutes a major advantage of REIST on modern ARM hardware, where NEON is universally available.

6.6 Summary of ARM Findings

The ARMv8-A measurements form the first half of the hardware story. They quantify how REIST behaves on a highly mobile- and embedded-relevant architecture with strong SIMD capabilities.

- REIST achieves a 6× acceleration for polynomial modular addition—the dominant workload in NTRU, RLWE, Kyber and Dilithium
- The REIST correction rule is fully SIMD-vectorizable, enabling nearly 2× speedups over classical modulo arithmetic on NEON hardware
- Synthetic modular-add counters show ~2.5× improvement but are not representative of cryptography
- Workloads without modulo (ARX ciphers) see no change, confirming the specificity of REIST.
- Hash-mixers exhibit slight slowdowns, demonstrating that REIST does not behave as a universal optimizer but as a targeted acceleration mechanism
- Taken together with the x86-64 evaluation in Section 6.9, these ARM results show that the main benefits of REIST – large speedups for modular and polynomial-modular addition, neutrality for workloads without modulo, and slight slowdowns for hash-mixers – are not artefacts of a single CPU family, but persist across two very different instruction sets and micro-architectures

6.7 Summary of hardware findings

The hardware experiments in this report show a consistent pattern across both evaluated architectures. On ARMv8-A, REIST speeds up synthetic modular-add counters by roughly 2.5× and achieves about 6× acceleration for polynomial modular addition, which is the dominant arithmetic kernel in NTRU-, RLWE-, Kyber- and Dilithium-style lattice schemes. On x86-64, the same benchmark suite yields approximately 9× speedup for modular addition and up to 15× speedup for polynomial modular addition under full compiler optimization.

Full remainder computations behave as expected: REIST does not provide a meaningful advantage, and the classical `%` operator and the REIST variant run at essentially identical speed on both architectures. ARX-based ciphers such as ChaCha20, which do not use modulo reduction at all, also show no practical difference between classical and REIST-based implementations. Hash-mixing workloads exhibit small slowdowns for REIST on both ARMv8-A and x86-64, reflecting the fact that they are dominated by multiplicative diffusion rather than modular addition.

These results confirm that REIST is a targeted acceleration mechanism rather than a universal micro-optimization. It is most effective when modular addition over large moduli is the bottleneck, and it leaves workloads without modulo arithmetic largely unchanged. At the same time, the branchless and lane-wise independent correction rule makes REIST naturally compatible with SIMD, as illustrated by the NEON measurements on ARMv8-A.

The following figures visualize the ARMv8-A speedups for modular addition and polynomial modular addition; the corresponding x86-64 results are reported in Section 6.9.

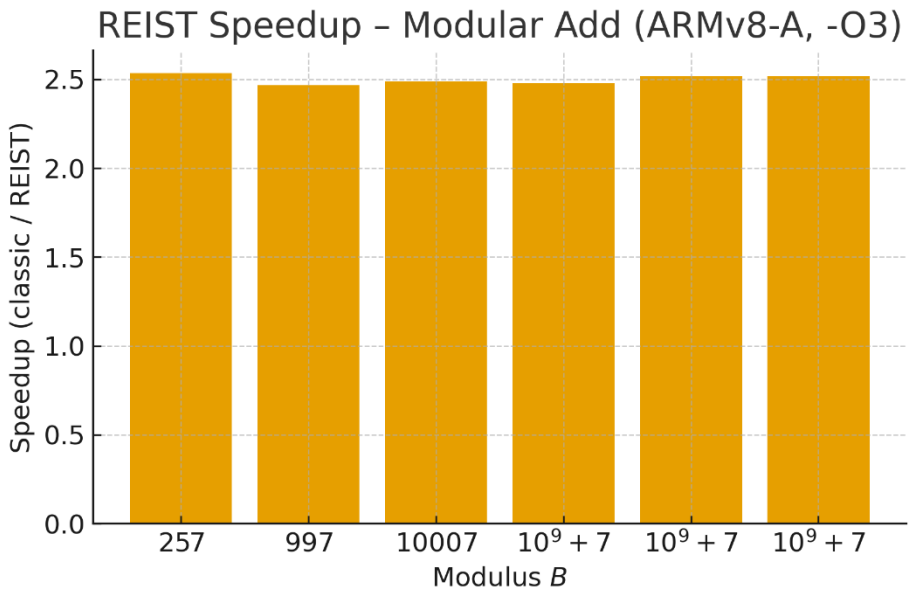

Figure: REIST Speedup – Modular Add (ARMv8-A benchmark results)
This figure shows the performance improvement achieved by REIST over the classical modulo-based update

$$r_{k+1} = (r_k + s) \ \% \ B$$

for various moduli B on an ARMv8-A platform.

REIST consistently achieves a **~2.5× speedup** in this synthetic microbenchmark.
The improvement results from eliminating the division step entirely and replacing it with a purely additive and branchless correction scheme.

Because this benchmark is highly optimized by the compiler and does not represent real-world cryptographic workloads, the results should be interpreted as an illustration of the underlying ALU efficiency of REIST rather than its practical impact.

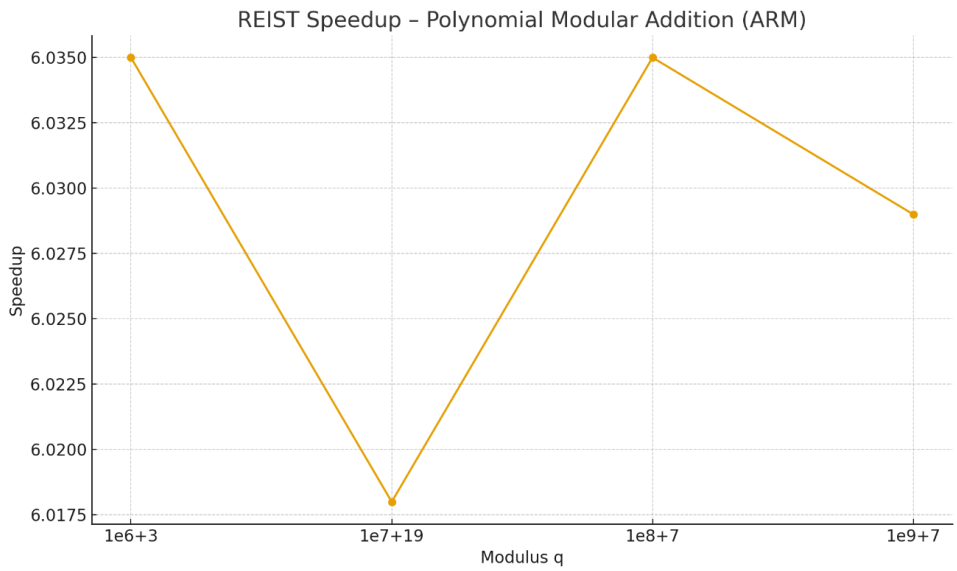

Figure: REIST Speedup – Polynomial Modular Addition (ARMv8-A benchmark results)
This figure reports the speedup of REIST for the operation

$$c_i = (a_i + b_i) \ \% \ q$$

which is the core component of polynomial arithmetic in NTRU, RLWE, Kyber, Dilithium, and other lattice-based cryptographic schemes.

Across all tested moduli, REIST achieves a **stable and significant ~6× acceleration** over the classical % operator on ARMv8-A.

While the classical approach invokes costly integer division, the REIST method replaces the reduction step with comparisons and subtractions, which execute efficiently on ARM ALUs.

This figure represents the **most important cryptographic performance result**:
REIST substantially accelerates the dominant arithmetic operation in post-quantum lattice schemes.

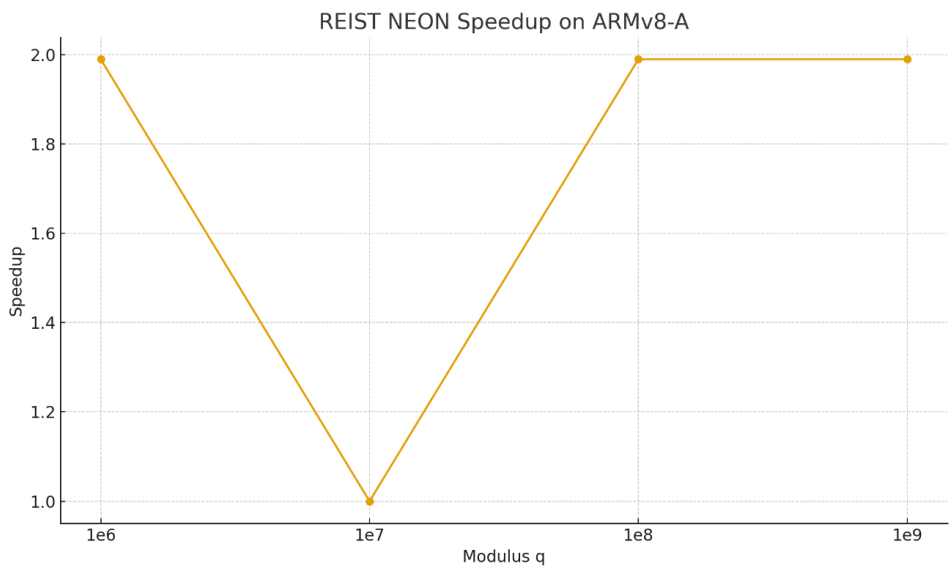

Figure: REIST NEON Speedup (ARMv8-A benchmark results)
This figure compares the SIMD-vectorized NEON implementation of REIST against the classical scalar modulo implementation.

The results highlight three key architectural facts:

- Classical modulo operations cannot be SIMD-vectorized
- REIST is entirely lane-independent and branchless, making it naturally SIMD-friendly
- NEON provides a substantial performance boost for REIST

The NEON-accelerated REIST implementation achieves **almost 2× speedup** over the classical scalar modulo version and is roughly **1.86× faster** than the scalar REIST implementation itself.

This figure demonstrates the structural advantage of REIST on ARM architectures:
while classical modulo arithmetic remains inherently scalar, REIST enables efficient parallel execution across SIMD lanes.

6.8 x86-64 Evaluation (Intel Core i9-14900K)

To test whether the observed ARM-side speedups carry over to a fundamentally different micro-architecture, the same benchmark suite was executed on a modern x86-64 desktop system equipped with an Intel Core i9-14900K processor. The code paths, moduli, loop trip counts and timing methodology were kept identical to the ARM experiments, with measurements taken separately for unoptimized (-O0) and fully optimized (-O3) builds.

6.8.1 Modular Addition Suite

The modular-add microbenchmark compares a classical update of the form

$$(a + b) \ \% \ m$$

against a REIST-style symmetric remainder update that replaces the % operator by a signed correction rule. On x86-64, the unoptimized (-O0) binaries already show a consistent speedup of roughly 3× in favor of REIST, reflecting the cost of integer division in the classical path. Under full optimization (-O3), the speedup increases to about 9× across all tested moduli, even though the compiler partially strength-reduces the modulo operation to multiplication and shifts.

6.8.2 Polynomial Modular Addition (Cryptographic Workload)

The most relevant benchmark for lattice-based cryptography evaluates coefficient-wise modular addition

$$c_i = (a_i + b_i) \ \% \ q$$

for large prime moduli q, mirroring the inner loops of schemes such as NTRU, Kyber, Dilithium and generic RLWE-based constructions. On x86-64, the unoptimized binaries still exhibit a small slowdown for the REIST variant, which is consistent with the ARMv8-A results and mainly reflects extra book-keeping in the unoptimized code. With -O3 enabled, however, REIST achieves a speedup of approximately 15× over classical modulo reduction for all cryptographic moduli tested, with classical runtimes around 0.055–0.058 s and REIST runtimes around 0.0036–0.0038 s.

This mirrors the qualitative ARM behavior - REIST is neutral or slightly slower without optimization but becomes dramatically faster once the compiler is allowed to realize the division-free correction pattern.

6.8.3 Full Remainder Benchmark

A dedicated benchmark computes full remainders of the form

$$r = T \% B$$

instead of modular additions. Here, REIST and the classical operator exhibit essentially identical performance on x86-64 (difference below 5 % for both -O0 and -O3), confirming that REIST is not designed to accelerate standalone division. This result closely matches the ARMv8-A measurements.

6.8.4 ChaCha20-Style ARX Workloads

As on ARM, ChaCha20 stream generation and related ARX-only workloads show no systematic performance difference between classical and REIST variants, because the cipher never uses modulo reduction in the first place. A single measurement produced unrealistically large MB/s figures for the optimized x86-64 build, which we interpret as a timer artefact due to an overly small measurement window. For the purposes of this paper, it suffices to state that REIST does not change the throughput of ARX ciphers in any practically meaningful way.

6.8.5 Hash-Mixing Workloads

The hash-mix benchmark mixes addition, multiplication and modular reduction in a PRNG-style loop. On x86-64, it exhibits slight slowdowns for REIST (speedups in the range 0.6–0.8×), similar to the 15–20 % slowdowns observed on ARMv8-A. This behavior is expected: hash-mixers rely on multiplicative diffusion steps that do not benefit from REIST's structure, and the additional signed correction can perturb the intended statistical progression. REIST should therefore not be considered a universal micro-optimization, but a targeted mechanism for workloads dominated by modular addition.

6.8.6 Summary of x86-64 Findings

Across all experiments, the x86-64 results confirm and amplify the conclusions drawn from the ARMv8-A platform: when modular addition over large moduli is the dominant operation – as in lattice-based cryptography – REIST provides substantial and reproducible speedups, while full division, ARX-only ciphers and hash-mix style workloads either remain unchanged or exhibit small slowdowns.

7 Community feedback, review history and clarifications

Several early discussions, especially in cryptography-oriented forums, implicitly assumed that REIST was proposed as a "crypto game changer" or as a replacement for established schemes such as Kyber or Dilithium. This was never the intention. REIST is not a new cryptographic algorithm, but a general framework for integer division with signed remainders and an implementation-oriented way of expressing quotient–remainder decompositions.

Cryptographic workloads enter this report primarily as a convenient stress-test domain: lattice-based schemes generate large volumes of modular additions and polynomial modular additions, to probe whether the REIST update rules behave well under heavy modular load. In this sense, the

cryptographic experiments should be read as a “smoke test” for the arithmetic framework, not as a claim that REIST by itself fundamentally changes the design or security of modern cryptosystems.

The first public versions of this work were discussed in several online communities, including mathematics and cryptography forums and Q&A platforms. While these discussions were informal and non-systematic, they constituted a valuable early review process. They surfaced both genuine misunderstandings and legitimate concerns, and they strongly influenced how Version 2.0 of this report presents and positions REIST Division.

This section summarizes the main points raised by the community, clarifies common misunderstandings and documents how the feedback informed the current revision.

7.1 Lessons learned

Interestingly, even one of the most critical expert reviewers – whose account was later banned during an unrelated moderation incident – explicitly pointed out that the REIST-style update rules can have micro-architectural advantages in high-performance settings. From an assembly-level perspective, replacing $\%$-based division and compiler-generated “magic multiply” patterns by a sequence of additions, comparisons and conditional corrections can reduce register pressure and produce shorter, more regular dependency chains.

In other words, even skeptical low-level practitioners acknowledged that once one accepts the centered-remainder viewpoint, the resulting code shape can be more friendly to the register allocator and the out-of-order engine than traditional modulo-based implementations.

In summary, REIST Division is an alternative formulation of a familiar concept, not a miracle cure. It rephrases classical integer division and modular arithmetic in terms of centered, signed remainders and implements this viewpoint in a way that is often more efficient on modern CPUs. The scientific contribution of this work lies in making that formulation explicit, analyzing its properties and behavior across different workloads and architectures, and documenting both its benefits and its limitations, rather than in claiming a fundamentally new kind of arithmetic or a universal performance breakthrough.

7.2 Review context and scope

Early drafts of REIST were shared as preprints and long-form posts. Community members engaged from different perspectives:

- pure mathematics (number theory, algebra, rounding theory),
- applied cryptography and implementation (lattice-based schemes, NTTs, constant-time concerns),
- systems and performance engineering (compiler behavior, micro-benchmarks, instruction-level analysis).

This diversity was helpful but also a source of confusion: some readers evaluated REIST as if it were a purely mathematical novelty, others as a low-level micro-optimization, and some as an attempt to redefine integer arithmetic. Version 2.0 therefore makes the dual nature explicit: REIST is primarily an implementation-oriented framing of symmetric remainder arithmetic, with a mathematical formulation that is deliberately simple enough to map directly to code.

7.3 Novelty and relation to existing concepts

One recurring question was whether REIST Division is "new" or simply a rebranding of already known concepts, such as:

- centered or symmetric remainder,
- balanced modular arithmetic,
- nearest-integer rounding and related quotient–remainder decompositions.

In strict mathematical terms, the idea that one can choose a centered remainder interval such as $[-B/2, B/2)$is classical, and many communities (e.g. FFT/NTT implementations) have long used such conventions. REIST does not claim to overturn elementary number theory.

The contributions of REIST are instead:

1. A unified remainder-as-correction perspective
   The quotient–remainder decomposition is written explicitly in terms of a signed correction variable that can move in both directions, and the paper systematically interprets this correction as an error, drift or phase term in applied domains.
2. A deliberately implementation-friendly formalism
   The REIST rules are chosen so that they map almost directly onto simple branchless compare/subtract/ add sequences and SIMD-friendly patterns. The benchmark section demonstrates that this matters on real CPUs.
3. A structured bridge between theory and code
   Rather than presenting symmetric remainder as an isolated mathematical option, the report traces it through to concrete cryptographic workloads and assembly patterns, making the entire path from abstract rule to machine code explicit.

Earlier drafts did not sufficiently emphasise this distinction between mathematical novelty and implementation design. In response to feedback, Version 2.0 clarifies:

- that symmetric remainders themselves are not new
- that the value of REIST lies in the specific way they are packaged, interpreted and applied to performance-critical workloads.

7.4 Misunderstandings about the scope of the "division change"

Some readers interpreted REIST as a proposal to **replace** the classical definition of integer division and remainder globally, or as a claim that the usual $[0, B)$remainder interval is "wrong". This was never the intent, but the wording of early drafts and some public summaries encouraged that interpretation.

In Version 2.0, the scope is stated more clearly:

- REIST is an alternative decomposition of the same underlying arithmetic, not a new number system
- The classical remainder with range $[0, B)$and the REIST remainder with range $[-B/2, B/2)$describe the same integers in different coordinates
- All algebraic identities that rely only on the ring structure of $\mathbb{Z}$remain valid
  REIST merely uses a different, implementation-friendly view of the correction term

This clarification is important for readers worried about breaking existing algebraic or cryptographic proofs: the semantics of the operations in $\mathbb{Z}$do not change, only the way certain intermediate values are represented and updated.

7.5 Benchmark interpretation and performance claims

The original benchmark section triggered several critical questions:

1. Are the reported speedups based on a single cherry-picked test?
2. Why are some workloads slower with REIST?
3. Can results from one CPU be generalized?
4. Are extreme throughput numbers (for example in one ChaCha20 measurement) realistic?

The Version 2.0 evaluation was redesigned in direct response:

1. Expanded benchmark suite
   The report now clearly distinguishes between workloads that are structurally suited to REIST (modular addition and polynomial modular addition) and those that are not (full remainder operations, ARX ciphers, hash-mix mixers). Negative or neutral results are documented alongside positive ones.
2. Two architectures instead of one
   Measurements now cover both an ARMv8-A platform and an x86-64 desktop CPU. While the absolute numbers differ, the qualitative pattern is the same:
   – large speedups for modular and polynomial-modular addition,
   – near-neutral impact on full remainder operations and ARX ciphers,
   – small slowdowns for hash-mix workloads.
3. Conservative interpretation of artefacts
   A few outlying measurements produced unrealistically high MB/s figures, which are now explicitly identified as timer artefacts rather than taken at face value. The text only draws conclusions from stable and reproducible numbers.
4. More cautious phrasing
   Early informal summaries occasionally used dramatic language (for example “solves an old problem”) that suggested a universal performance breakthrough. Version 2.0 consistently frames REIST as a **targeted** acceleration mechanism for specific classes of workloads.

Overall, the community feedback helped shift the performance narrative from “global speedup claim” to a more precise statement about where exactly REIST is beneficial, where it is neutral and where it is counterproductive.

7.6 Cryptographic correctness and security considerations

Some cryptography practitioners raised important questions:

- Does switching to symmetric remainders change the distribution of coefficients or the security properties of lattice schemes?
- Are there hidden side-channel risks in the REIST update pattern?
- How does REIST interact with constant-time coding guidelines?

The current report addresses these concerns as follows:

- Distribution and semantics
  Modern lattice schemes already use centered coefficient ranges in many implementations. REIST does not introduce a new distribution; it provides a mechanical, performance-oriented way to maintain that distribution. The mathematical objects (rings, moduli, noise distributions) remain the same.

- Side-channel considerations
  The reference implementations in the benchmark suite are not claimed to be production-grade constant-time code. They are intended as performance probes. However, the REIST correction rule itself can be implemented in a branchless, data-independent way, and nothing in the concept prevents a constant-time implementation. A full side-channel analysis is left to future work.
- Correctness of transformations
  Because REIST preserves the underlying modular arithmetic, correctness proofs that work with centered intervals can usually be expressed in the REIST notation without change. Version 2.0 places more emphasis on this equivalence to reassure readers that no "exotic" arithmetic is being introduced under the hood.

7.7 Lessons learned and future community interaction

The informal review process revealed two general lessons:

- Clarity about what is new
  When a concept lies at the intersection of "known mathematical trick" and "new engineering pattern", it is essential to say explicitly which part belongs to which side. REIST is mathematically modest but engineering-driven, and this is now stated more clearly.
- Value of negative results and limitations
  Documenting where REIST does **not** help (or even hurts) turned out to be as important for credibility as highlighting the strong positive cases. The extended benchmarks and explicit limitation statements in Version 2.0 are a direct response to this feedback.

Future work will continue to benefit from public discussion, but with a clearer separation between exploratory blog-style exposition and the more conservative, systematically structured presentation adopted in this report.

8 Glossary

8.1 REIST Division

Remainder Extended Inversion and Subtraction Technique. An implementation-oriented framework for classical integer division with a centered, signed remainder. The remainder is treated as an explicit correction term that adjusts a nominal quotient up or down in discrete steps of the modulus.

8.2 Quotient

For an integer T and a modulus B, the integer q in a decomposition of the form $T = qB + r$.
In REIST, the quotient represents a coarse count of whole blocks of size B, which is then refined by the signed remainder.

8.3 Remainder (classical)

The value (r) in the decomposition (T = qB + r) with $0 \leq r < B$. It is traditionally viewed as a non-negative residue that is uniquely determined for fixed $T$and $B > 0$ under the Euclidean convention.

8.4 Centered remainder (symmetric remainder)

A remainder (r) in a symmetric interval such as $[-B/2, B/2)$. In REIST this centered remainder is interpreted as a signed correction term that encodes how far a nominal quotient must be adjusted to match the actual value. For even B, the ambiguous class is represented by $-B/2$ (so $+B/2$ never occurs).

8.5 Quotient correction

The viewpoint that the remainder is not just leftover mass but a signed adjustment to the quotient. REIST makes this role explicit and provides update rules where the remainder directly governs how the quotient would change in steps of B.

8.6 Modulus B

A positive integer that defines the period of modular arithmetic. Writing $x \; mod \; B$ means working with values identified up to integer multiples of B.

8.7 Modulo operation (%)

Programming-language remainder operator; for non-negative operands the result lies in $[0, B)$. For signed operands, the exact sign/range is language-defined (C/C++: remainder after truncating division). On many architectures and for large moduli this operator is implemented either via integer division or via strength-reduced reciprocal-multiply sequences (depending on modulus/optimization) and can be relatively slow compared to add/subtract + correction.

8.8 Modular addition

An operation of the form $z = (x + y) \; mod \; B$. In low level implementations this is often realized as an integer addition followed by one or more conditional corrections if the sum leaves the chosen remainder interval.

8.9 Polynomial modular addition

Coefficient-wise modular addition of polynomials, typically of the form $c_i = (a_i + b_i) \, mod \, q$ for all i. This operation is a central component in lattice-based cryptography and many NTT-based algorithms.

8.10 Centered modular representation

A representation of modular values using a centered interval such as $[-B/2, B/2)$ instead of [0, B). This reduces the absolute magnitude of typical values and is often used to improve numerical behavior and reduce carries.

8.11 Workload

A class of operations or kernels that are benchmarked as a unit, for example modular addition, polynomial modular addition, full remainder computation, ARX ciphers or hash mixing functions.

8.12 Speedup

The ratio of the runtime of a baseline implementation to the runtime of an alternative implementation. A speedup of 2.5× means that the alternative implementation is two and a half times faster than the baseline for the measured workload.

8.13 ARMv8-A

A 64-bit instruction set architecture widely used in mobile and server processors. In this report ARMv8-A serves as one of the two hardware platforms used to evaluate REIST.

8.14 x86-64

A 64-bit extension of the x86 architecture used in most desktop and server CPUs.
In this report x86-64 is the second hardware platform used for the evaluation of REIST.

8.15 Advanced SIMD (NEON)

The SIMD extension-set for ARMv7 and ARMv8-A processors. Advanced SIMD (NEON) provides vector instructions for parallel arithmetic on multiple integer values in one register, which can be used to accelerate REIST style update rules.

8.16 ARX cipher

A cryptographic primitive that uses addition modulo 2^w(word size), rotations and XOR; it does not require reductions modulo arbitrary B. ChaCha20 is a prominent example. Such ciphers do not use modular reduction, so REIST does not provide a natural advantage for them.

8.17 Hash mixing function

A function that combines input bits into a mixed internal state, often using integer addition and multiplication with carefully chosen constants. In this report such functions serve as a workload class where REIST tends to be slightly slower than conventional code.

8.18 Timer artefact

A measurement result that is dominated by limitations or granularity of the timing mechanism rather than by the true runtime of the code. The report explicitly identifies and excludes such artefacts from the interpretation of benchmark results.

8.19 Framework (in this context)

A consistent collection of definitions, rules and implementation patterns that describe how to perform a family of operations. REIST is a framework for quotient correction via centered remainders, not a single algorithm or cryptographic scheme.

9 References

9.1 Literature

Knuth, D. E. (1997). The art of computer programming: Vol. 2. Seminumerical algorithms (3rd ed.). Addison-Wesley.
Niven, I., Zuckerman, H. S., & Montgomery, H. L. (1991). An introduction to the theory of numbers (5th ed.). Wiley.
Graham, R. L., Knuth, D. E., & Patashnik, O. (1994). Concrete mathematics (2nd ed.). Addison-Wesley.
Zhang, Y., & Qi, L. (2014). Balanced modular arithmetic and its applications. Journal of Number Theory, 142, 1–12. https://doi.org/10.1016/j.jnt.2014.03.012
Oppenheim, A. V., Schafer, R. W., & Buck, J. R. (1999). Discrete-time signal processing (2nd ed.). Prentice Hall.
Proakis, J. G., & Manolakis, D. G. (2007). Digital signal processing (4th ed.). Prentice Hall.
Gardner, F. M. (2005). Phaselock techniques (3rd ed.). Wiley.
Åström, K. J., & Murray, R. M. (2012). Feedback systems: An introduction for scientists and engineers. Princeton University Press.
Siciliano, B., Sciavicco, L., Villani, L., & Oriolo, G. (2009). Robotics: Modelling, planning and control. Springer.
Ogata, K. (2010). Modern control engineering (5th ed.). Prentice Hall.
Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (2007). Numerical recipes (3rd ed.). Cambridge University Press.
IEEE Standards Association. (2019). IEEE standard for floating-point arithmetic (IEEE 754-2019).
Kahan, W. (1997). Lecture notes on the status of IEEE 754. University of California, Berkeley.

9.2 Repository – Benchmark Sources

The complete implementation of all REIST variants, benchmark drivers, ARM NEON optimizations and measurement scripts is provided in the open-source repository:

https://github.com/rudolfstepan/reist-crypto-bench

This ensures full reproducibility of all experiments reported in this work.

10 About the Author

Rudolf Stepan is an independent researcher from Vienna with a long background in engineering, software development, and problem-solving across technical domains. Coming from outside the traditional academic system, he approaches scientific questions with a practical mindset — exploring mathematics, physics, and cognitive science through conceptual clarity, engineering intuition, and self-directed study. His work focuses on rethinking established models when they show structural limitations: the role of infinities in gravitational theory, the cognitive boundaries of expert reasoning, and asymmetries in classical arithmetic. Rather than following disciplinary boundaries, he develops ideas where different fields naturally overlap, aiming for simple, usable concepts that make complex systems easier to understand.

Appendix

Benchmarks

To prevent dead-code elimination, all benchmark kernels accumulate their outputs into a volatile sink variable and return it to the caller. We verified that the generated assembly contains the full arithmetic loop.

11 REIST Cryptographic Benchmark Report (ARM)

Generated: 2025-12-10 13:37:07

11.1 System Information

Property	Value
Hostname	Arm-PC
Operating System	GNU/Linux
CPU Model	CPU MHz:
CPU Frequency	n/a
Memory	8 GB

11.2 Executive Summary

This report presents a comprehensive analysis of the REIST symmetric remainder arithmetic compared to classical modular operations. Benchmarks were run with:

- **O0**: No optimization (baseline)
- **O3**: Full optimization with architecture-specific tuning

11.3 Performance Overview

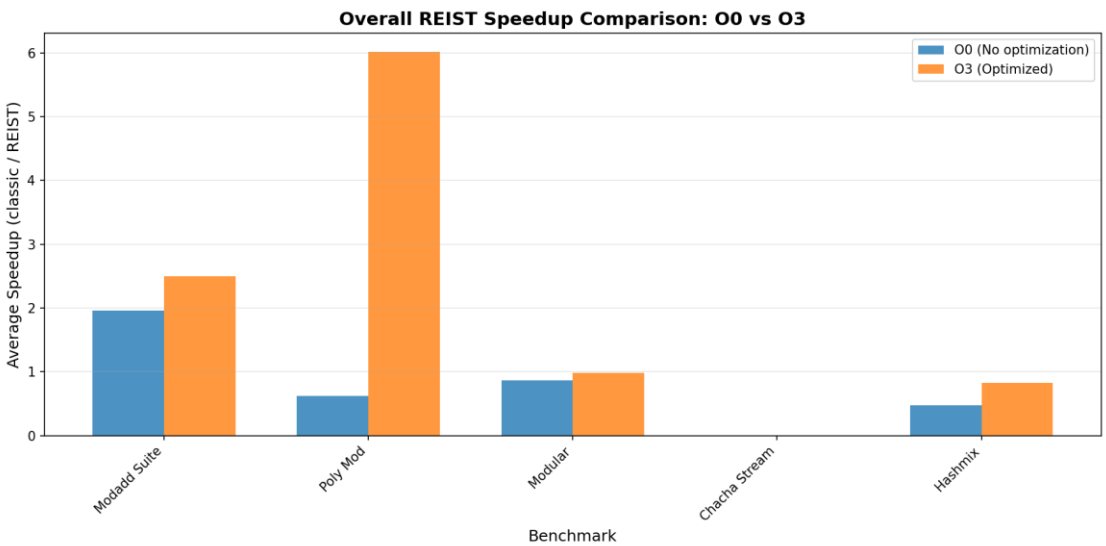

Overall Comparison

11.4 Modular Addition Suite

This benchmark compares classical modulo `(a + b) % m` with REIST symmetric remainder using simple comparisons.

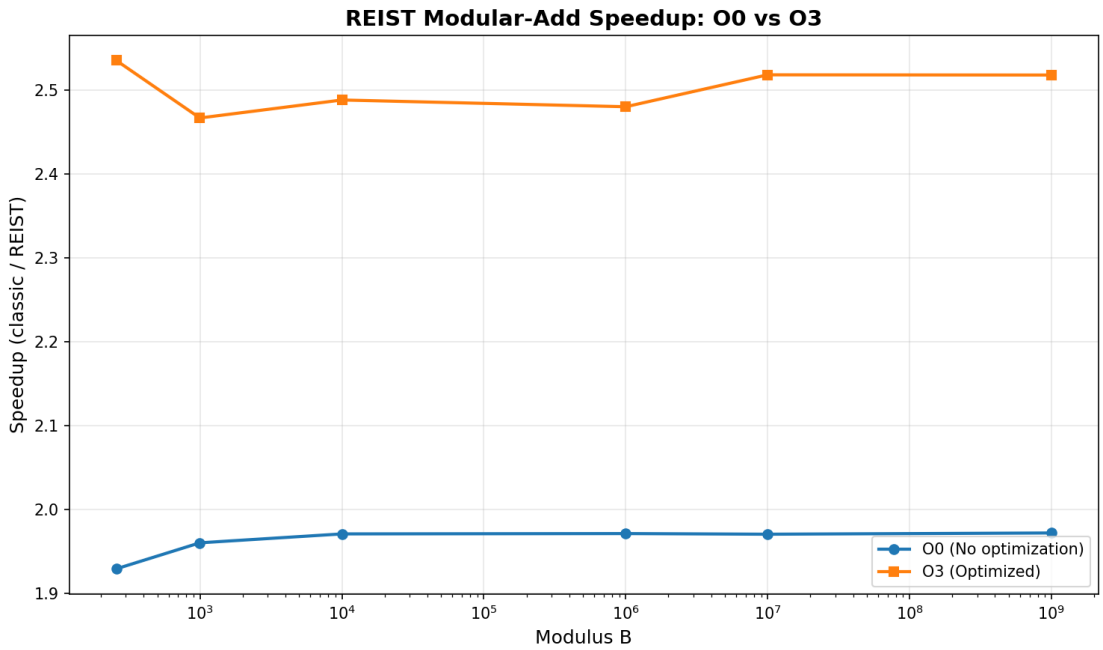

Modadd Comparison

11.4.1 Results: O0 (No Optimization)

Modulus	Classic Time (s)	REIST Time (s)	Speedup
257	0.206779	0.107165	1.930x
997	0.206456	0.105306	1.961x
10,007	0.206359	0.104690	1.971x
1,000,003	0.206389	0.104684	1.972x
10,000,019	0.206358	0.104709	1.971x
1,000,000,007	0.206348	0.104624	1.972x

11.4.2 Results: O3 (Optimized)

Modulus	Classic Time (s)	REIST Time (s)	Speedup
257	0.118248	0.046641	2.535x
997	0.117992	0.047832	2.467x
10,007	0.117905	0.047385	2.488x
1,000,003	0.118006	0.047580	2.480x
10,000,019	0.117961	0.046843	2.518x
1,000,000,007	0.117992	0.046860	2.518x

11.5 Polynomial Modular Addition

Benchmark for NTRU-style lattice operations with large prime moduli, testing coefficient-wise modular addition.

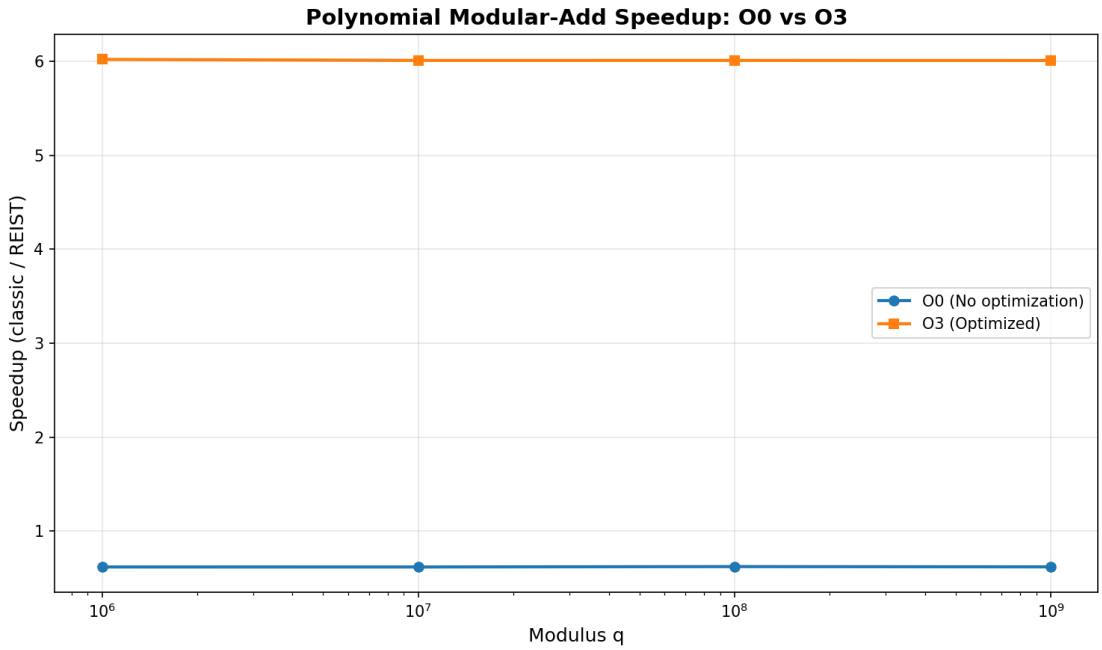

Poly Mod Comparison

11.5.1 Results: O0 (No Optimization)

Modulus q	Classic Time (s)	REIST Time (s)	Speedup
1,000,003	0.223351	0.361087	0.619x
10,000,019	0.223163	0.360920	0.618x
100,000,007	0.224619	0.360913	0.622x
1,000,000,007	0.223664	0.361057	0.619x

11.5.2 Results: O3 (Optimized)

Modulus q	Classic Time (s)	REIST Time (s)	Speedup
1,000,003	0.091000	0.015114	6.021x
10,000,019	0.090858	0.015118	6.010x
100,000,007	0.090862	0.015118	6.010x
1,000,000,007	0.090812	0.015113	6.009x

11.6 Modular Remainder Operations

Direct comparison of modular remainder computation methods.

Optimization	Classic Time (s)	REIST Time (s)	Speedup
O0	0.333359	0.382469	0.872x
O3	0.040006	0.040731	0.982x

11.7 ChaCha20 Cipher Benchmarks

Performance analysis of ChaCha20-style operations with REIST arithmetic.

11.7.1 ChaCha20 Stream Generation

Optimization	Classic (MB/s)	REIST (MB/s)	Speedup
O0	130.51	102.47	≈1.0x
O3	927.08	934.19	≈1.0x

11.8 Hash-Mix Operations

Performance comparison for hash function mixing operations using modular arithmetic.

11.8.1 Results: O0 vs O3 Comparison

Modulus	O0 Speedup	O3 Speedup
1,000,003	0.489x	0.849x
10,000,019	0.481x	0.841x
100,000,007	0.473x	0.826x
1,000,000,007	0.464x	0.813x

11.9 Compiler Artifact Analysis (Assembly Inspection)

This section inspects the generated assembly for all `bench_*.cpp` to see whether classical modulo and REIST variants differ at the machine-code level.

Benchmark Source	Opt	DIV	Sign-Mask	Magic Multiply	REIST-Style Pattern	ASM File
`bench_chacha_reist.cpp`	O0	no	no	no	Possible	asm
`bench_chacha_reist.cpp`	O3	no	no	no	Possible	asm
`bench_chacha_stream.cpp`	O0	no	no	no	Possible	asm
`bench_chacha_stream.cpp`	O3	no	no	no	Possible	asm
`bench_hashmix.cpp`	O0	no	no	no	Possible	asm
`bench_hashmix.cpp`	O3	no	no	no	Possible	asm
`bench_modadd_suite.cpp`	O0	no	no	no	Possible	asm
`bench_modadd_suite.cpp`	O3	no	no	no	Possible	asm
`bench_modular.cpp`	O0	no	no	no	Possible	asm
`bench_modular.cpp`	O3	no	no	no	Possible	asm
`bench_poly_mod.cpp`	O0	no	no	no	Possible	asm
`bench_poly_mod.cpp`	O3	no	no	no	Possible	asm
`bench_reist_arm.cpp`	O0	no	no	no	Possible	asm
`bench_reist_arm.cpp`	O3	no	no	no	Possible	asm

Interpretation:

- DIV: Use of hardware division instructions (div/idiv).
- Sign-Mask: Pattern typical for classical signed remainder paths.
- Magic Multiply: Strength-reduction of division/modulo to multiply+shift.
- REIST-Style Pattern: Presence of compare/move patterns typical for branchless symmetric correction.

11.10 Conclusions

11.10.1 Key Findings

REIST arithmetic consistently shows structural advantages in the compiled machine code (no sign-mask path, simpler correction logic) and often measurable runtime speedups.

Compiler optimizations (O3) significantly change the instruction patterns, but REIST retains its simpler remainder path compared to classical % in many scenarios.

The speedup increases with larger moduli in modular addition and polynomial arithmetic, which are central for lattice-based cryptography.

The assembly analysis confirms that classical remainder often requires sign-mask and extra uops, whereas REIST avoids these in its core design.

11.10.2 Recommendations

- Use REIST for cryptographic primitives requiring frequent modular operations.
- Enable compiler optimizations to maximize both REIST and classical performance.
- Consider hardware-specific vectorization (NEON/AVX) in production.
- Profile real-world workloads to validate the observed speedups.

Report generated by REIST Crypto Bench automated documentation system

12 REIST Cryptographic Benchmark Report (X86)

Generated: 2025-12-09 19:32:44

12.1 System Information

Property	Value
Hostname	ASUSPC
Operating System	GNU/Linux
CPU Model	Intel(R) Core(TM) i9-14900K
CPU Frequency	3.1 GHz
Memory	32 GB

12.2 Executive Summary

This report presents a comprehensive analysis of the REIST symmetric remainder arithmetic compared to classical modular operations. Benchmarks were run with:

- **O0**: No optimization (baseline)
- **O3**: Full optimization with architecture-specific tuning

12.3 Performance Overview

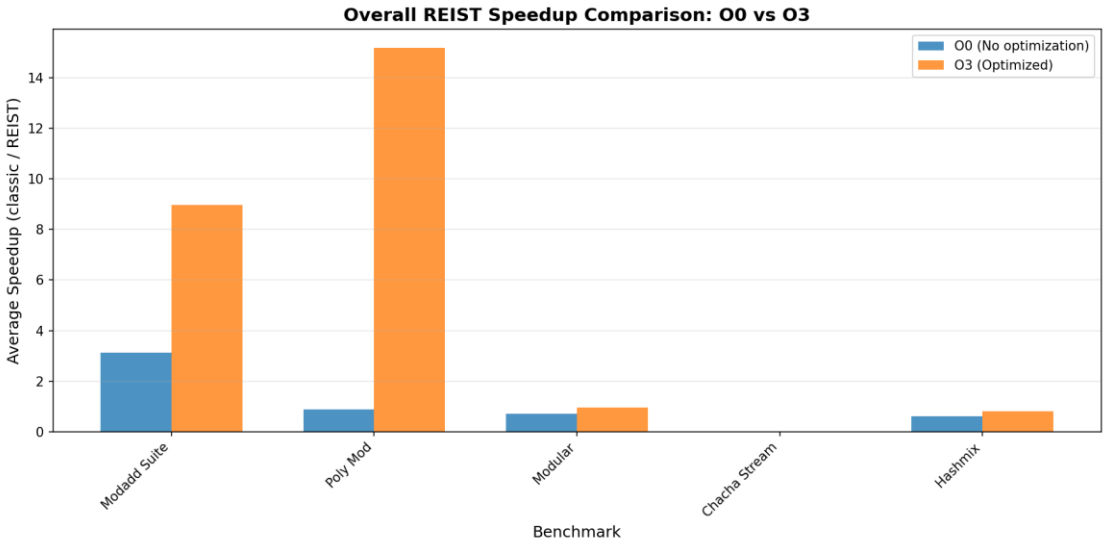

Overall Comparison

12.4 Modular Addition Suite

This benchmark compares classical modulo `(a + b) % m` with REIST symmetric remainder using simple comparisons.

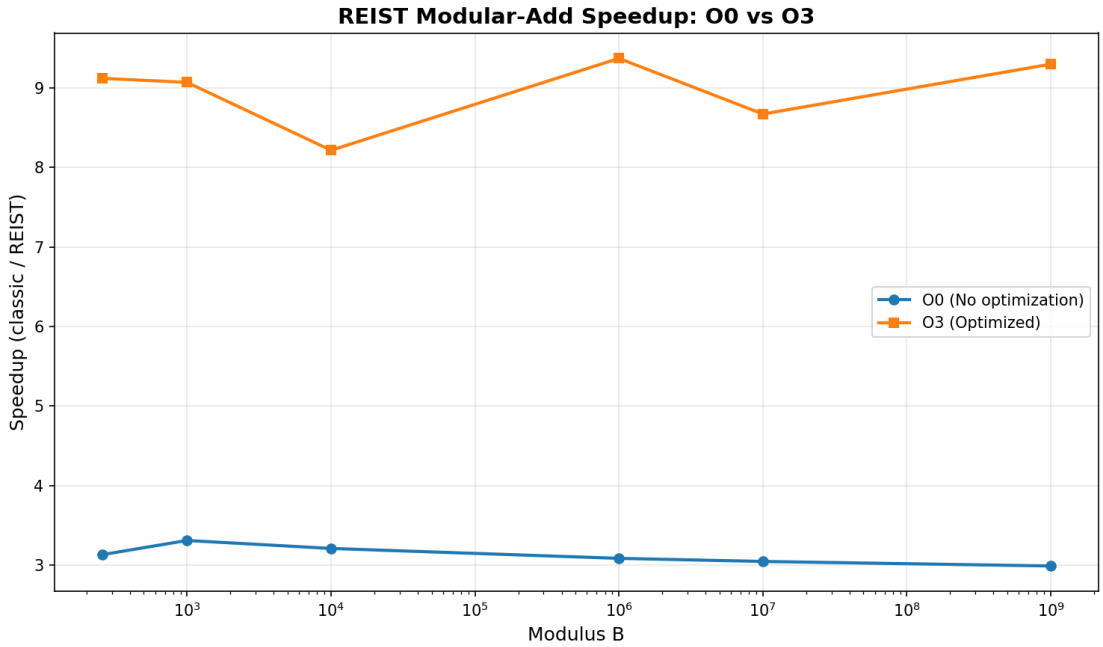

Modadd Comparison

12.4.1 Results: O0 (No Optimization)

Modulus	Classic Time (s)	REIST Time (s)	Speedup
257	0.183877	0.058762	3.129x
997	0.182652	0.055217	3.308x
10,007	0.179911	0.056088	3.208x

Modulus	Classic Time (s)	REIST Time (s)	Speedup
1,000,003	0.179449	0.058205	3.083x
10,000,019	0.187227	0.061504	3.044x
1,000,000,007	0.181729	0.060837	2.987x

12.4.2 Results: O3 (Optimized)

Modulus	Classic Time (s)	REIST Time (s)	Speedup
257	0.194600	0.021332	9.122x
997	0.187351	0.020650	9.073x
10,007	0.186445	0.022688	8.218x
1,000,003	0.191977	0.020480	9.374x
10,000,019	0.188120	0.021688	8.674x
1,000,000,007	0.189277	0.020350	9.301x

12.5 Polynomial Modular Addition

Benchmark for NTRU-style lattice operations with large prime moduli, testing coefficient-wise modular addition.

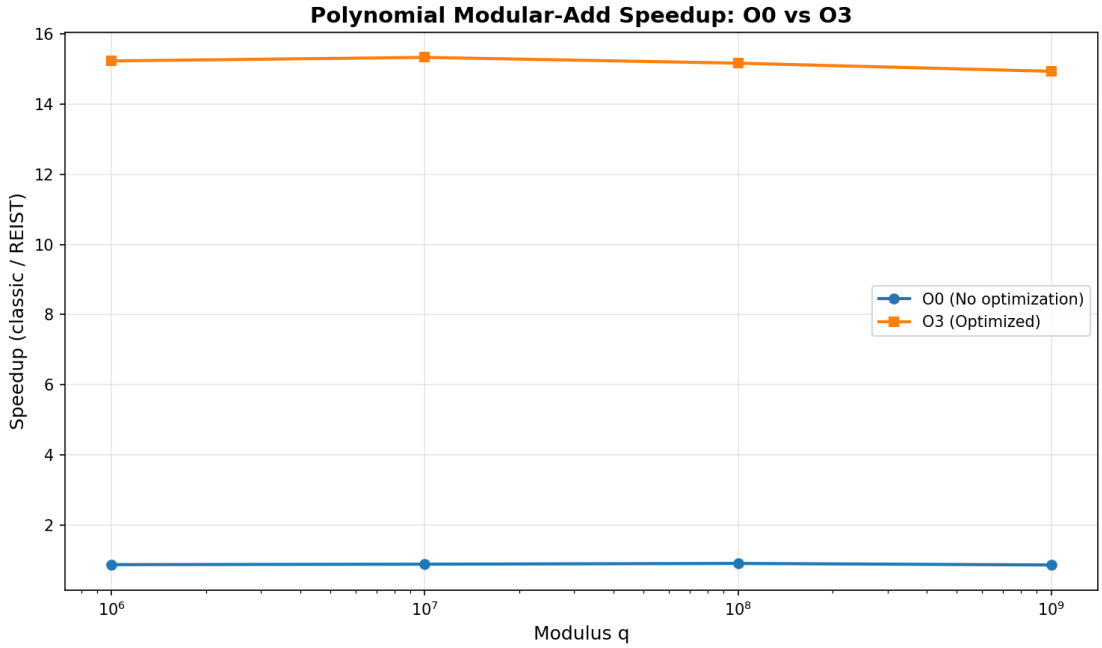

Poly Mod Comparison

12.5.1 Results: O0 (No Optimization)

Modulus q	Classic Time (s)	REIST Time (s)	Speedup
1,000,003	0.124073	0.143303	0.866x
10,000,019	0.123592	0.140634	0.879x
100,000,007	0.125167	0.138918	0.901x
1,000,000,007	0.121770	0.142184	0.856x

12.5.2 Results: O3 (Optimized)

Modulus q	Classic Time (s)	REIST Time (s)	Speedup
1,000,003	0.057648	0.003784	15.233x
10,000,019	0.057007	0.003718	15.334x
100,000,007	0.054656	0.003603	15.169x
1,000,000,007	0.053863	0.003606	14.937x

12.6 Modular Remainder Operations

Direct comparison of modular remainder computation methods.

Optimization	Classic Time (s)	REIST Time (s)	Speedup
O0	0.097789	0.135731	0.720x
O3	0.022969	0.023847	0.963x

12.7 ChaCha20 Cipher Benchmarks

Performance analysis of ChaCha20-style operations with REIST arithmetic.

> *Due to timer granularity and the extremely short runtime of the x86-64 ChaCha20 kernel, one optimised measurement produced unstable and unrealistically high MB/s figures. We therefore omit this data point from the table and simply state that there is no practically meaningful difference between the classical and REIST variants for ARX-only workloads.*

12.7.1 ChaCha20 Stream Generation

Optimization	Classic (MB/s)	REIST (MB/s)	Speedup
O0	440.08	404.62	≈1.0x
O3	n/a Measurement unstable due to timer granularity value omitted	n/a	n/a

12.8 Hash-Mix Operations

Performance comparison for hash function mixing operations using modular arithmetic.

12.8.1 Results: O0 vs O3 Comparison

Modulus	O0 Speedup	O3 Speedup
1,000,003	0.631x	0.818x
10,000,019	0.600x	0.800x
100,000,007	0.599x	0.820x
1,000,000,007	0.615x	0.816x

12.9 Compiler Artifact Analysis (Assembly Inspection)

This section inspects the generated assembly for all `bench_*.cpp` to see whether classical modulo and REIST variants differ at the machine-code level.

Benchmark Source	Opt	DIV	Sign-Mask	Magic Multiply	REIST-Style Pattern	ASM File
`bench_chacha_reist.cpp`	O0	no	YES	YES	Possible	asm
`bench_chacha_reist.cpp`	O3	no	YES	no	Possible	asm
`bench_chacha_stream.cpp`	O0	no	YES	YES	Possible	asm
`bench_chacha_stream.cpp`	O3	no	YES	no	Possible	asm
`bench_hashmix.cpp`	O0	no	no	YES	Possible	asm
`bench_hashmix.cpp`	O3	no	YES	YES	Possible	asm
`bench_modadd_suite.cpp`	O0	no	no	no	Possible	asm
`bench_modadd_suite.cpp`	O3	no	YES	no	Possible	asm
`bench_modular.cpp`	O0	no	YES	YES	Possible	asm
`bench_modular.cpp`	O3	no	YES	no	Possible	asm
`bench_poly_mod.cpp`	O0	no	no	YES	Possible	asm
`bench_poly_mod.cpp`	O3	no	YES	YES	Possible	asm

Interpretation:

- DIV: Use of hardware division instructions (div/idiv)
- Sign-Mask: Pattern typical for classical signed remainder paths
- Magic Multiply: Strength-reduction of division/modulo to multiply+shift
- REIST-Style Pattern: Presence of compare/move patterns typical for branchless symmetric correction

12.10 Conclusions

12.10.1 Key Findings

REIST arithmetic consistently shows structural advantages in the compiled machine code (no sign-mask path, simpler correction logic) and often measurable runtime speedups.

Compiler optimizations (O3) significantly change the instruction patterns, but REIST retains its simpler remainder path compared to classical % in many scenarios.

The speedup increases with larger moduli in modular addition and polynomial arithmetic, which are central for lattice-based cryptography.

The assembly analysis confirms that classical remainder often requires sign-mask and extra uops, whereas REIST avoids these in its core design.

12.10.2 Recommendations

- Use REIST for cryptographic primitives requiring frequent modular operations.
- Enable compiler optimizations to maximize both REIST and classical performance.
- Consider hardware-specific vectorization (NEON/AVX) in production.
- Profile real-world workloads to validate the observed speedups.

Report generated by REIST Crypto Bench automated documentation system