

Vectorized Falcon-Sign Implementations using SSE2, AVX2, AVX-512F, NEON, and RVV

Jipeng Zhang¹ Jiaheng Zhang¹

¹National University of Singapore, Singapore
jp-zhang@outlook.com

Paper: <https://eprint.iacr.org/2025/1867>

Artifact: <https://github.com/Ji-Peng/VecFalcon>

Slides: https://ji-peng.github.io/uploads/tches2026/VecFalcon_slides.pdf

IACR TCHES 2026-1

2026-01-08

Outline

1 Contributions

2 Motivations

3 Background

4 Vectorized BaseSampler

5 Vectorized FFT (RISC-V)

6 Results

Contributions

This paper focuses on optimizing **Falcon Signature Generation**.

- **Performance Profiling:** Identified BaseSampler ($> 30\%$) and FFT-related subroutines (on RISC-V) as bottlenecks.
- **Vectorized BaseSampler:**
 - Implemented across **6 ISAs**: SSE2, AVX2, AVX-512F, NEON, RVV, RV64IM.
 - Achieved up to **8.4 \times** (AVX2) and **7.7 \times** (RVV) speedup for the sampler.
- **Vectorized FFT on RVV:**
 - Novel **4+5 layer merging** strategy using strided load/store instructions.
 - Achieved **4.7 \times** speedup on RVV.
- Signature generation speedups of **23%** (AVX2), **36%** (AVX-512F), and **59%** (RV64GCVB).

FALCON (FN-DSA)

- Selected by NIST for standardization (FIPS 206). **Note:** The FIPS 206 standard document was **not yet published** at the time of this work.
- Fast verification, but **slow signature generation**.

Bottleneck 1: Discrete Gaussian Sampling

- BaseSampler accounts for > 30% of signing time.
- Non-vectorizable in reference code due to sequential UniformBits calls for KAT compatibility.

Bottleneck 2: FFT on RISC-V

- FFT-related ops take $\approx 38\%$ time on SpacemIT X60 (RV64GCVB).
- Existing optimizations lack efficient deep layer merging strategies.

Falcon Signature Generation

- Involves Fast Fourier Sampling (ffSampling) and Discrete Gaussian Sampling (SamplerZ).
- SamplerZ calls BaseSampler to sample integers z_0 from distribution χ .

Target Platforms

Architecture	CPU	ISA
x86-64	Intel i7-11700K	SSE2, AVX2, AVX-512F
ARMv8-A	Cortex-A72	NEON
RISC-V	SpacemiT X60	RV64GC, RVV (v1.0), Bit-manip

Background: FALCON & BaseSampler

Algorithm 1: Sign($m, sk, \lfloor \beta^2 \rfloor$)

Input : A message m , a secret key sk , and a bound $\lfloor \beta^2 \rfloor$

Output : A signature sig of message m

```
1:  $r \leftarrow \{0, 1\}^{320}$  uniformly
2:  $c \leftarrow \text{HashToPoint}(r \| m, q, n)$ 
3:  $t \leftarrow (-\frac{1}{q} \text{FFT}(c) \odot \text{FFT}(F), \frac{1}{q} \text{FFT}(c) \odot \text{FFT}(f))$ 
4: do
5:   do
6:      $z \leftarrow \text{ffSampling}_n(t, T)$ 
7:      $s = (t - z)\hat{B}$ 
8:     while  $\|s\|^2 > \lfloor \beta^2 \rfloor$ 
9:      $(s_1, s_2) \leftarrow \text{iFFT}(s)$ 
10:     $s \leftarrow \text{Compress}(s_2, 8 \cdot \text{sbytelen} - 328)$ 
11: while  $s = \perp$ 
12: return sig = ( $r, s$ )
```

Algorithm 2: ffSampling_n(t, T)

Input : $t = (t_0, t_1) \in \text{FFT}(\mathbb{Q}[x]/(x^n + 1))^2$; T

Output : $z = (z_0, z_1) \in \text{FFT}(\mathbb{Z}[x]/(x^n + 1))^2$

```
1: if  $n = 1$  then
2:    $\sigma' \leftarrow T.\text{value}$ 
3:    $z_0 \leftarrow \text{SamplerZ}(t_0, \sigma')$ 
4:    $z_1 \leftarrow \text{SamplerZ}(t_1, \sigma')$ 
5:   return  $z = (z_0, z_1)$ 
6:  $(\ell, T_0, T_1) \leftarrow (T.\text{value}, T.\text{leftchild}, T.\text{rightchild})$ 
7:  $t_1 \leftarrow \text{splitfft}(t_1)$ 
8:  $z_1 \leftarrow \text{ffSampling}_{n/2}(t_1, T_1)$ 
9:  $z_1 \leftarrow \text{mergefft}(z_1)$ 
10:  $t'_0 \leftarrow t_0 + (t_1 - z_1) \odot \ell$ 
11:  $t_0 \leftarrow \text{splitfft}(t'_0)$ 
12:  $z_0 \leftarrow \text{ffSampling}_{n/2}(t_0, T_0)$ 
13:  $z_0 \leftarrow \text{mergefft}(z_0)$ 
14: return  $z = (z_0, z_1)$ 
```

Background: FALCON & BaseSampler

Algorithm 3: SamplerZ(μ, σ')

Input : $\mu, \sigma' \in \mathcal{R}; \sigma' \in [\sigma_{\min}, \sigma_{\max}]$
Output : $z \in \mathbb{Z}$ close to $D_{\mathbb{Z}, \mu, \sigma'}$

- 1: $r \leftarrow \mu - |\mu|$
- 2: $ccs \leftarrow \sigma_{\min}/\sigma'$
- 3: **while** (1) **do**
- 4: | $z_0 \leftarrow \text{BaseSampler}()$
- 5: | $b \leftarrow \text{UniformBits}(8) \& 0x1$
- 6: | $z \leftarrow b + (2 \cdot b - 1)z_0$
- 7: | $x \leftarrow \frac{(z-r)^2}{2\sigma'^2} - \frac{z_0^2}{2\sigma_{\max}^2}$
- 8: | **if** $\text{BerExp}(x, ccs) = 1$ **then**
- 9: | | **return** $z + \lfloor \mu \rfloor$

Algorithm 4: ApproxExp(x, ccs)

Input : $x \in [0, \ln(2)]; ccs \in [0, 1];$
Output : A precomputed array C

Output : An integer
 $\approx 2^{63} \cdot ccs \cdot \exp(-x)$

- 1: $y \leftarrow C[0]$
- 2: $z \leftarrow \lfloor 2^{63} \cdot x \rfloor$
- 3: **for** $i = 1$ **to** 12 **do**
- 4: | $y \leftarrow C[i] - (z \cdot y) \gg 63$
- 5: $z \leftarrow \lfloor 2^{63} \cdot ccs \rfloor$
- 6: $y \leftarrow (z \cdot y) \gg 63$
- 7: **return** y

Algorithm 5: BaseSampler()

Output : $z_0 \in \{0, \dots, 18\}; z_0 \sim \chi$

- 1: $u \leftarrow \text{UniformBits}(72)$
- 2: $z_0 \leftarrow 0$
- 3: **for** $i = 0$ **to** 17 **do**
- 4: | $z_0 \leftarrow z_0 + \llbracket u < \text{RCDT}[i] \rrbracket$
- 5: **return** z_0

Algorithm 6: BerExp(x, ccs)

Input : Floating point values
 $x, ccs \geq 0$
Output : A single bit, equal to 1
with probability
 $\approx ccs \cdot \exp(-x)$

- 1: $s \leftarrow \lfloor x / \ln(2) \rfloor$
- 2: $r \leftarrow x - s \cdot \ln(2)$
- 3: $s \leftarrow \min(s, 63)$
- 4: $z \leftarrow (2 \cdot \text{ApproxExp}(r, ccs) - 1) \gg s$
- 5: $i \leftarrow 64$
- 6: **do**
- 7: | $i \leftarrow i - 8$
- 8: | $w \leftarrow \text{UniformBits}(8) - ((z \gg i) \& 0xFF)$
- 9: **while** (($w = 0$) **and** ($i > 0$))
- 10: **return** $\llbracket w < 0 \rrbracket$

Performance Profiling

Profiling Methodology

- **Baseline:** C-FN-DSA project (SHAKE256X4 variant). FALCON-1024's signature generation.
- **Tool:** gperftools.
- **Platforms:** Intel i7-11700K (AVX2) & SpacemIT X60 (RV64GCVB).

Key Observations

- **BaseSampler** is a consistent bottleneck ($> 30\%$) across architectures.
- **FFT-related** operations are expensive specifically on RISC-V ($\approx 38\%$).
- *Note: SHA-3 is already optimized for AVX2; BerExp is left for future work.*

Table: Breakdown of Execution Time

Component	AVX2	RISC-V
BaseSampler	30.2%	30.3%
FFT-related	15.1%	37.6%
SHA-3	22.6%	20.6%
BerExp	31.2%	14.3%

Optimization Targets

Based on this data, we focus on:

- ① **BaseSampler**
- ② **FFT-related ops (RISC-V only)**

Vectorized BaseSampler: Previous Methods on x86

```
void gaussian0_ref(sampler_state *ss, int32_t *z_bimodal,
                   int32_t *z_square)
{
    /* Get a random 72-bit value, into three 24-bit limbs (v0..v2). */
    uint64_t lo = prng_next_u64(&ss->pc);
    uint32_t hi = prng_next_u8(&ss->pc);
    uint32_t v0 = (uint32_t)lo & 0xFFFFFFF;
    uint32_t v1 = (uint32_t)(lo >> 24) & 0xFFFFFFF;
    uint32_t v2 = (uint32_t)(lo >> 48) | (hi << 16);
    /* Sampled value is z such that v0..v2 is lower than the first
     * z elements of the table. */
    int32_t z = 0;
    for (size_t i = 0; i < (sizeof GAUSS0) / sizeof(GAUSS0[0]); i++) {
        uint32_t cc;
        cc = (v0 - GAUSS0[i][2]) >> 31;
        cc = (v1 - GAUSS0[i][1] - cc) >> 31;
        cc = (v2 - GAUSS0[i][0] - cc) >> 31;
        z += (int32_t)cc;
    }
    // Get a random bit b to turn the sampling into a bimodal distribution.
    int32_t b = prng_next_u8(&ss->pc) & 1;
    *z_bimodal = b + ((b << 1) - 1) * z;
    *z_square = z * z;
}
```

C impl. in C-FN-DSA project (72 bits = 3×24 bits; 59 cycles without PRNG overhead)

```
void gaussian0_avx2_ref(sampler_state *ss, int32_t *z_bimodal,
                        int32_t *z_square)
{
    ...
    lo = prng_next_u64(&ss->pc); hi = prng_next_u8(&ss->pc);
    hi = (hi << 7) | (unsigned)(lo >> 57); lo = 0xFFFFFFFFFFFFFF;
    xhi = _mm256_broadcastw_epi16(_mm_cvtsi32_si128(hi));
    rhi = _mm256_loadu_si256(Brhi15.ymm[0]);
    gthi = _mm256_cmplt_epi16(rhi, xhi); eqhi = _mm256_cmpeq_epi16(rhi, xhi);
    t = _mm256_casts128_si256(gthi, 15);
    zt = _mm_setzero_si128();
    t = _mm_hadd_epi16(t, zt); t = _mm_hadd_epi16(t, zt);
    t = _mm_hadd_epi16(t, zt); r = _mm_cvtsi128_si32(t);
    xlo = _mm256_broadcastb_epi64(_mm_cvtsi64_si128(*(<int64_t>*)&lo));
    gtl0 = _mm256_cmplt_epi64(_mm256_loadu_si256(&rlo57.ymm[0]), xlo);
    gtl01 = _mm256_cmplt_epi64(_mm256_loadu_si256(&rlo57.ymm[1]), xlo);
    gtl02 = _mm256_cmplt_epi64(_mm256_loadu_si256(&rlo57.ymm[2]), xlo);
    gtl03 = _mm256_cmplt_epi64(_mm256_loadu_si256(&rlo57.ymm[3]), xlo);
    gtl04 = _mm256_cmplt_epi64(_mm256_loadu_si256(&rlo57.ymm[4]), xlo);
    gtl0 = _mm256_and_si256(
        gtl0, _mm256_cvtepi16_epi64(_mm256_casts128_si128(eqhi)));
    gtl01 = _mm256_and_si256(
        gtl01, _mm256_cvtepi16_epi64(
            _mm256_casts1256_si128(_mm256_bsrl1_epi128(eqhi, 8))));
    eqm = _mm256_permute4x64_ep164(eqhi, 0xFF);
    gtl02 = _mm256_and_si256(gtl02, eqm); gtl03 = _mm256_and_si256(gtl03, eqm);
    gtl04 = _mm256_and_si256(gtl04, eqm); gtl00 = _mm256_or_si256(gtl00, gtl01);
    gtl0 = _mm256_add_epi64(_mm256_and_epi64(gtl00, gtl02),
                           _mm256_and_epi64(gtl03, gtl04));
    t = _mm_add_epi64(_mm256_casts1256_si128(gtl00),
                      _mm256_extracti128_si256(gtl00, 1));
    t = _mm_add_epi64(t, _mm_srl1_si128(t, 8));
    r -= _mm_cvtsi128_si32(t);
    int32_t b = prng_next_u8(&ss->pc) & 1;
    *z_bimodal = b + ((b << 1) - 1) * r; *z_square = r * r;
}
```

AVX2 impl. in NIST submission (72 bits = 15 + 57 bits; 44 cycles without PRNG overhead)

Vectorized BaseSampler: Previous Methods on ARMv8-A NEON

```
void gaussian0_ref(sampler_state *ss, int32_t *z_bimodal,
                   int32_t *z_square)
{
    /* Get a random 72-bit value, into three 24-bit limbs (v0..v2). */
    uint64_t lo = prng_next_u64(&ss->pc);
    uint32_t hi = prng_next_u32(&ss->pc);
    uint32_t v0 = (uint32_t)lo & 0xFFFFFFF;
    uint32_t v1 = (uint32_t)(lo >> 24) & 0xFFFFFFF;
    uint32_t v2 = (uint32_t)(lo >> 48) | (hi << 16);
    /* Sampled value is z such that v0..v2 is lower than the first
     * z elements of the table. */
    int32_t z = 0;
    for (size_t i = 0; i < (sizeof GAUSS0) / sizeof(GAUSS0[0]); i++) {
        uint32_t cc;
        cc = (v0 - GAUSS0[i][2]) >> 31;
        cc = (v1 - GAUSS0[i][1] - cc) >> 31;
        cc = (v2 - GAUSS0[i][0] - cc) >> 31;
        z += (int32_t)cc;
    }
    // Get a random bit b to turn the sampling into a bimodal distribution.
    int32_t b = prng_next_u8(&ss->pc) & 1;
    *z_bimodal = b + ((b << 1) - 1) * z;
    *z_square = z * z;
}
```

C impl. in C-FN-DSA (72 bits = 3×24 bits;
54 cycles without PRNG overhead)

```
void gaussian0_NG23(sampler_state *ss, int32_t *z_bimodal,
                     int32_t *z_square)
{
    ...
    lo = prng_next_u64(&ss->pc); hi = prng_next_u8(&ss->pc);
    v0 = (uint32_t)lo & 0xFFFFFFF; v1 = (uint32_t)(lo >> 24) & 0xFFFFFFF;
    v2 = (uint32_t)(lo >> 48) | (hi << 16); x0 = vdupq_n_u32(v0);
    x1 = vdupq_n_u32(v1); x2 = vdupq_n_u32(v2);
    w = vld3q_u32(&dist[0]); cc0 = vsubq_u32(x0, w.val[2]);
    cc1 = vsubq_u32(x1, w.val[1]); cc2 = vsubq_u32(x2, w.val[0]);
    cc1 = (uint32x4_t)vsraq_n_s32((int32x4_t)cc1, (int32x4_t)cc0, 31);
    cc2 = (uint32x4_t)vsraq_n_s32((int32x4_t)cc2, (int32x4_t)cc1, 31);
    zz = vshrq_n_u32(cc2, 31); w = vld3q_u32(&dist[12]);
    cc0 = vsubq_u32(x0, w.val[2]); cc1 = vsubq_u32(x1, w.val[1]);
    cc2 = vsubq_u32(x2, w.val[0]);
    cc1 = (uint32x4_t)vsraq_n_s32((int32x4_t)cc1, (int32x4_t)cc0, 31);
    cc2 = (uint32x4_t)vsraq_n_s32((int32x4_t)cc2, (int32x4_t)cc1, 31);
    zz = vsraq_n_u32(cc2, 31); w = vld3q_u32(&dist[24]);
    cc0 = vsubq_u32(x0, w.val[2]); cc1 = vsubq_u32(x1, w.val[1]);
    cc2 = vsubq_u32(x2, w.val[0]);
    cc1 = (uint32x4_t)vsraq_n_s32((int32x4_t)cc1, (int32x4_t)cc0, 31);
    cc2 = (uint32x4_t)vsraq_n_s32((int32x4_t)cc2, (int32x4_t)cc1, 31);
    zz = vsraq_n_u32(cc2, 31); w = vld3q_u32(&dist[36]);
    cc0 = vsubq_u32(x0, w.val[2]); cc1 = vsubq_u32(x1, w.val[1]);
    cc2 = vsubq_u32(x2, w.val[0]);
    cc1 = (uint32x4_t)vsraq_n_s32((int32x4_t)cc1, (int32x4_t)cc0, 31);
    cc2 = (uint32x4_t)vsraq_n_s32((int32x4_t)cc2, (int32x4_t)cc1, 31);
    zz = vsraq_n_u32(cc2, 31); wh = vld3_u32(&dist[48]);
    cc0h = vsub_u32(vget_low_u32(x0), wh.val[2]);
    cc1h = vsub_u32(vget_low_u32(x1), wh.val[1]);
    cc2h = vsub_u32(vget_low_u32(x2), wh.val[0]);
    cc1h = (uint32x2_t)vsraq_n_s32((int32x2_t)cc1h, (int32x2_t)cc0h, 31);
    cc2h = (uint32x2_t)vsraq_n_s32((int32x2_t)cc2h, (int32x2_t)cc1h, 31);
    zzh = vshrq_n_u32(cc2h, 31); z = vaddv_u32(zzh);
    int32_t b = prng_next_u8(&ss->pc) & 1;
    *z_bimodal = b + ((b << 1) - 1) * z;
    *z_square = z * z;
}
```

NEON impl. in [NG23] (72 bits = 3×24 bits; 59 cycles without PRNG overhead)

Challenge: Strict KAT compatibility enforces sequential PRNG usage.

Our Solution:

- ① **Relax KAT Compatibility:** Does *not* affect interoperability with verification.
- ② **Batch Processing:** Generate N samples at once (e.g., $N = 16$ for AVX2).
- ③ **Modular Design:**
 - Main computation loop vectorized.
 - Bimodal transformation and squaring integrated.

The Trade-off

- Vectorization requires **parallel** PRNG usage.
- This breaks strict KAT compatibility (which mandates a sequential PRNG order).

Why Interoperability is Preserved?

- Our modification does *not* alter the specific Gaussian distribution of the output vector s (line 7 of Alg. 1).
- It preserves the rejection sampling condition (line 8 of Alg. 1). The condition $\|s\|^2 \leq \lfloor \beta^2 \rfloor$ remains strictly satisfied.

* Note: The C-FN-DSA project employs a similar trade-off in its SHAKE256X4 variant.

Algorithm 1: Sign($m, sk, \lfloor \beta^2 \rfloor$)

Input : A message m , a secret key sk , and a bound $\lfloor \beta^2 \rfloor$
Output : A signature sig of message m

```
1:  $r \leftarrow \{0, 1\}^{320}$  uniformly
2:  $c \leftarrow \text{HashToPoint}(r \| m, q, n)$ 
3:  $t \leftarrow (-\frac{1}{q}\text{FFT}(c) \odot \text{FFT}(F), \frac{1}{q}\text{FFT}(c) \odot \text{FFT}(f))$ 
4: do
5:   do
6:      $z \leftarrow \text{ffSampling}_n(t, T)$ 
7:      $s = (t - z)\bar{B}$ 
8:     while  $\|s\|^2 > \lfloor \beta^2 \rfloor$ 
9:      $(s_1, s_2) \leftarrow \text{iFFT}(s)$ 
10:     $s \leftarrow \text{Compress}(s_2, 8 \cdot \text{sbytelen} - 328)$ 
11: while  $s = \perp$ 
12: return  $sig = (r, s)$ 
```

Figure: Alg. 1

Implementation on x86: Algorithm Overview

Algorithm 7: Vectorized BaseSampler

Output: N independent pairs (z, z_0^2)

1 *Constants:*

2 | $N_s = 4$ for SSE2

3 | $N_s = 8$ for AVX2

4 | $N_s = 16$ for AVX-512F

5 | M is an integer such that $N = M \cdot N_s$

6 | RCDT__N_s : RCDT in vectorized form

7 *Variable declarations:*

8 | $\text{prn_24x3_}_N_s \text{ prn}[M]$

9 | $\text{ALIGNED_INT32}(N) b$

10 | $\mathbf{z}_0[0], \dots, \mathbf{z}_0[M-1] \leftarrow 0$

11 // Prepare random numbers

12 for $j \leftarrow 0$ to $M-1$ do

13 | for $i \leftarrow 0$ to N_s-1 do

14 | for $k \leftarrow 0$ to 2 do

15 | | $\text{prn}[j].i32[k][i] \leftarrow$
 | | UniformBits(24)

16 $bs \leftarrow \text{UniformBits}(N)$

17 for $i \leftarrow 0$ to $N-1$ do

18 | $b.i32[i] \leftarrow (bs \gg i) \& 1$

19 // Main computation loop

```
20 for  $i \leftarrow 0$  to  $17$  do
21   |  $t_l \leftarrow \text{RCDT\_}_N_s[i][0]$  // low 24-bit
22   |  $t_m \leftarrow \text{RCDT\_}_N_s[i][1]$  // middle
23   |  $t_h \leftarrow \text{RCDT\_}_N_s[i][2]$  // high
24   | for  $k \leftarrow 0$  to  $M-1$  do
25    |   |  $c \leftarrow \text{VSUB}(\text{prn}[k].\mathbf{v}[0], t_l)$ 
26    |   |  $c \leftarrow \text{VSRLI}(c, 31)$ 
27    |   |  $c \leftarrow \text{VSUB}(\text{prn}[k].\mathbf{v}[1], c)$ 
28    |   |  $c \leftarrow \text{VSUB}(c, t_m)$ 
29    |   |  $c \leftarrow \text{VSRLI}(c, 31)$ 
30    |   |  $c \leftarrow \text{VSUB}(\text{prn}[k].\mathbf{v}[2], c)$ 
31    |   |  $c \leftarrow \text{VSUB}(c, t_h)$ 
32    |   |  $c \leftarrow \text{VSRLI}(c, 31)$ 
33    |   |  $\mathbf{z}_0[k] \leftarrow \text{VADD}(\mathbf{z}_0[k], c)$ 
34 // Bimodal and squaring
35 for  $k \leftarrow 0$  to  $M-1$  do
36   |  $t_b \leftarrow b.\mathbf{v}[k]$ 
37   |  $t_1 \leftarrow \text{VADD}(t_b, t_b)$ 
38   |  $t_1 \leftarrow \text{VSUB}(t_1, 1)$ 
39   |  $t_2 \leftarrow \text{VMULLO}(t_1, \mathbf{z}_0[k])$ 
40   |  $\mathbf{z}[k] \leftarrow \text{VADD}(t_2, t_b)$ 
41   |  $\mathbf{z}_0[k]^2 \leftarrow \text{VMULLO}(\mathbf{z}_0[k], \mathbf{z}_0[k])$ 
42 return  $(\mathbf{z}[k], \mathbf{z}_0[k]^2), k = 0, \dots, M-1$ 
```

Comparison Optimization

- Reference: Requires subtraction with borrow on 3×24 -bit integers.
- SIMD: Use `vpcmpgtd` (Compare Packed Signed Integers Greater Than).

Instruction Scheduling

- **SSE2/AVX2:** Optimized using `vpcmpgtd` (Latency 1, CPI 0.5 on Rocket Lake).
- **AVX-512F:** Avoided `vpcmpgtd` due to higher latency (3 cycles).
- **Unrolling:** Fully unrolled inner loops (**for** $k \leftarrow 0$ **to** $M - 1$) to enable instruction interleaving and reduce pipeline stalls.

Implementation on RISC-V Vector (RVV): Algorithm Overview

Algorithm 8: Vectorized BaseSampler

Output: N independent pairs (z, z_0^2)

1 *Constants:*

2 | $N_s = 4$ for NEON

3 | $N_s = 8$ for RVV with VLEN=256

4 | $N = M \cdot N_s$

5 | RCDT_Ns

6 *Variable declarations:*

7 | prn_24x3_Ns prn[M]

8 | ALIGNED_INT32(N) b

9 | $z_0[0], \dots, z_0[M-1] \leftarrow 0$

10 // Prepare random numbers

11 for $j \leftarrow 0$ to $M-1$ do

12 | for $i \leftarrow 0$ to N_s-1 do

13 | for $k \leftarrow 0$ to 2 do

14 | | prn[j].i32[k][i] \leftarrow
 | | UniformBits(24)

15 $bs \leftarrow$ UniformBits(N)

16 for $i \leftarrow 0$ to $N-1$ do

17 | $b.i32[i] \leftarrow (bs \gg i) \& 1$

```
18 // Main computation loop
19 for  $k \leftarrow 0$  to  $M-1$  do
20   | for  $i \leftarrow 0$  to 17 do
21    |   |  $t_l \leftarrow$  RCDT_Ns[i][0]
22    |   |  $t_m \leftarrow$  RCDT_Ns[i][1]
23    |   |  $t_h \leftarrow$  RCDT_Ns[i][2]
24    |   |  $c \leftarrow$  VSUB(prn[k].v[0],  $t_l$ )
25    |   |  $c \leftarrow$  VSRLI( $c$ , 31)
26    |   |  $c \leftarrow$  VSUB(prn[k].v[1],  $c$ )
27    |   |  $c \leftarrow$  VSUB( $c$ ,  $t_m$ )
28    |   |  $c \leftarrow$  VSRLI( $c$ , 31)
29    |   |  $c \leftarrow$  VSUB(prn[k].v[2],  $c$ )
30    |   |  $c \leftarrow$  VSUB( $c$ ,  $t_h$ )
31    |   |  $c \leftarrow$  VSRLI( $c$ , 31)
32    |   |  $z_0[k] \leftarrow$  VADD( $z_0[k]$ ,  $c$ )
33    |   | // Bimodal and squaring
34    |   |  $t_b \leftarrow b.v[k]$ 
35    |   |  $t_1 \leftarrow$  VADD( $t_b$ ,  $t_b$ )
36    |   |  $t_1 \leftarrow$  VSUB( $t_1$ , 1)
37    |   |  $t_2 \leftarrow$  VMULLO( $t_1$ ,  $z_0[k]$ )
38    |   |  $z[k] \leftarrow$  VADD( $t_2$ ,  $t_b$ )
39    |   |  $z_0[k]^2 \leftarrow$  VMULLO( $z_0[k]$ ,  $z_0[k]$ )
40 return ( $z[k]$ ,  $z_0[k]^2$ ),  $k = 0, \dots, M-1$ 
```

Constraints

- RVV comparison instructions output to mask registers, not usable directly in arithmetic.
- vmsbc (subtract with borrow) has dependency chains through mask register v0.

Optimization Strategy

- **Hybrid approach:** Pack 3 iterations into one macro.
- Use direct subtraction (vsub, vsrl) for 2 iterations + borrow method for 1 iteration.
- **Load-once-use-many:** Keep RCDT table in vector registers ($N = 64$).
 - $18 \times 3 = 54$ 24-bit integer segments, where 12 segments are zero-valued.
 - 24 scalar + 18 vector registers can hold 42 non-zero segments.
 - vx-type instructions: vsub.vx v2, v1, t0.

Overview

- The implementation logic is **nearly identical to the RVV version**.
- Batch size: $N = 64$ (similar to RVV).

Register Allocation: Load-once-use-many

- *Challenge:* NEON lacks RVV's .vx feature (cannot use scalar registers as operands).
- *Strategy:*
 - Persistently store **20 RCDT table entries** in **20 vector registers**.
 - Load the remaining entries from memory on demand during execution.

Comparison Optimization

- Implemented using cmgt (Compare Greater Than).

Strategy

- Designed for RISC-V processors **without RVV support**.
- Adopts the **Batch Processing** strategy ($N = 64$).
- Partition 72-bit integer into **High 8-bits + Low 64-bits**.

Load-once-use-many for RCDT table:

- **High 8-bits:** All 18 entries fit into 12-bit signed immediates (e.g., encoded directly in addi).
- **Low 64-bits:**
 - **2 entries** are $< 2^{11}$, encoded via immediates.
 - **remaining 16 entries** are stored in 16 registers.

```
16 for  $j \leftarrow 0$  to  $N$  do
17   for  $i \leftarrow 0$  to 17 do
18     // Set to 1 if less than
19      $c \leftarrow$ 
20     SLTU( $prn[j].[0]$ ,  $RCDT[i][0]$ )
21      $c \leftarrow prn[j].[1] - c$ 
22      $c \leftarrow c - RCDT[i][1]$ 
23      $c \leftarrow c \gg 63$ 
24      $z_0[j] \leftarrow z_0[j] + c$ 
```

Figure: Part of Alg. 9

Motivation: FFT is a major bottleneck on RISC-V ($\approx 38\%$).

Key Innovation: Strided Load/Store

- RVV supports `v1se64.v` (Vector Load Strided Element).
- Minimal overhead compared to contiguous load (CPI 4 vs 3).
- Allows flexible coefficient loading for deep layer merging.

Layer Merging Strategies

- FALCON-512: **4+4** merging (vs NEON's 2+2+4 in [NG23]).
- FALCON-1024: **4+5** merging (vs NEON's 1+2+2+4 in [NG23]).
- Directly construct coefficient arrangements in registers, minimizing memory access.

Results: BaseSampler Performance

Speedup compared to Reference Implementation (C-FN-DSA project)

Instruction Set	Ref Cycles	Our Cycles	Speedup
SSE2	59	14	4.2×
AVX2	59	7	8.4×
AVX-512F	59	6	9.8×
NEON	54	30	1.8×
RVV	192	25	7.7×
RV64IM	192	51	3.8×

Note: The comparison excludes PRNG overhead to highlight sampler efficiency.

Results: FFT/iFFT on SpacemIT X60

Speedup compared to Reference Implementation (C-FN-DSA project)

Algorithm	Impl.	Strategy	Cycles	Speedup
FFT-1024	Ref	-	80,524	1.0×
	Our RV64D	3+3+3	27,115	3.0×
	Our RVV	4+5	17,181	4.7×
iFFT-1024	Ref	-	76,652	1.0×
	Our RV64D	3+3+3	27,074	2.8×
	Our RVV	5+4	17,974	4.3×

Results: FALCON-512 Signature Generation

Comparison with Reference Implementation (C-FN-DSA project)

Platform	Instruction Set	Ref (k)	Our (k)	Speedup
x86-64	SSE2	631	556	1.13×
	AVX2	543	441	1.23×
	AVX-512F	536	393	1.36×
ARMv8	NEON	1,230	1,053	1.17×
RISC-V	RV64GCB	2,535	1,867	1.36×
	RV64GCVB	2,530	1,590	1.59×

- **x86-64:** Significant gains from vectorized BaseSampler & 8-way Keccak (AVX-512).
- **RISC-V:** Massive speedup (1.59×) driven by RVV BaseSampler + FFT.
- Note: Cycle counts in thousands (k).

Conclusion

- We addressed the main bottlenecks in FALCON signature generation: BaseSampler and FFT.
- **Vectorized BaseSampler:** Implemented across 6 ISAs, yielding massive speedups (up to $9.8\times$).
- **RVV FFT:** Leveraged strided loads for 4+5 layer merging, achieving $> 4\times$ speedup.
- **Final Result:** Significant performance gains (up to $1.59\times$) for FALCON- $\{512, 1024\}$ signature generation across x86, ARM, and RISC-V.

Paper: <https://eprint.iacr.org/2025/1867>

Artifact: <https://github.com/Ji-Peng/VecFalcon>

Slides: https://ji-peng.github.io/uploads/tches2026/VecFalcon_slides.pdf

Thank You!