

Vectorising and distributing NTTs to count Goldbach partitions on Arm-based supercomputers

Ricardo Jesus¹ Tomás Oliveira e Silva² Michèle Weiland¹

¹EPCC, The University of Edinburgh

²IEETA/DETI, Universidade de Aveiro

May 2021

- Introduction
- Vectorising modular arithmetic loops with SVE
- Distributing and parallelising NTTs
- Preliminary results on counting Goldbach partitions
- Conclusions

- **Problem: count Goldbach partitions to large limits**
 - Evaluate for all even n below cutoff point

$$R(n) = \# \{(p, q) : n = p + q \wedge p, q \text{ prime}\} \quad (1)$$

- Can be done via a polynomial product \Rightarrow requires exact convolution
- **Number-theoretic transforms (NTTs) enable fast exact convolutions**
 - “DFTs with modular arithmetic”
 - Used extensively for bignum and polynomial arithmetic
 - Because we want to achieve very large limits, we need a distributed implementation
- NTTs rely on **modular arithmetic**
 - Tends to hinder vectorisation (lack of suitable instructions)
 - SVE supports the required operations

- Elementary operations: addition, subtraction, and multiplication

$$z_{\pm} = (x \pm y) \bmod N; \quad z_{\times} = (xy) \bmod N$$

- Additions and subtractions are trivial, **multiplications are harder**
- The Montgomery method is a well-known algorithm for doing fast modular multiplication
 - For $R > N$ and $\gcd(R, N) = 1$:

$$\text{Let } \tilde{x} = (xR) \bmod N, \quad N' = (-N^{-1}) \bmod R, \quad \text{and} \quad (2)$$

$$\text{REDC}(T) = [T + N ((TN') \bmod R)] / R \bmod N^{\dagger}. \quad \text{Then} \quad (3)$$

$$\boxed{\tilde{xy} = \text{REDC}(\tilde{x}\tilde{y})} \quad (4)$$

[†]If $0 \leq T < RN$, $\bmod N$ can be replaced by a conditional subtraction.


- Additions and subtractions are the same in Montgomery representation:

$$\widetilde{x \pm y} = (\tilde{x} \pm \tilde{y}) \bmod N$$

- For multiplications:

$$\widetilde{xy} = w - N [w \geq N]$$

- With $R = 2^{64}$ the high and low parts of intermediate products are manipulated almost independently



⇒ modular arithmetic loops can be efficiently vectorised with SVE

- A generalisation of the Discrete Fourier Transform (DFT) defined over a ring or field

$$X_k = \sum_{j=0}^{m-1} x_j g^{-jk} \bmod p \quad (5)$$

- **Operations are devoid of errors**
- Used to perform convolutions with guarantee of exact results
- Most FFT algorithms can be used with NTTs, as long as operations are recast to the new domain

- Tests on an A64FX processor
 - **SVE** version implemented with **ACLE** intrinsics
- Major compilers
 - Arm 21.0
 - Cray 10.0.1[†]
 - Fujitsu 4.3.1,
 - Gnu 10.2 & 11

```

inline svuint64_t
mul_sve(svbool_t pg, svuint64_t a, svuint64_t b)
{
    svbool_t pc;
    svuint64_t m, xh, xl, z;

    /* x = a*b */
    xl = svmul_x(pg, a, b);
    xh = svmulh_x(pg, a, b);

    /* m = (x*N') mod R */
    m = svmul_x(pg, xl, N_prime);





    /* c = (x mod R)+(m*N mod R) >= R */
    pc = svcmpgt(pg, xl, svmla_x(pg, xl, m, N));

    /* z = (x+m*N)/R + c */
    z = svadd_x(pg, xh, svmulh_x(pg, m, N));
    z = svadd_m(pc, z, (uint64_t)1);

    /* adjust result */
    pc = svcmpge(pg, z, N);
    return svsub_m(pc, z, N);
}

```


[†]Did not support SVE intrinsics.



- Hybrid MPI and OpenMP implementation
- Adaptation of Bailey's four-step algorithm¹
 0. Organise array of m points into a $H \times W$ distributed matrix
 1. Perform NTTs on the columns
 2. Apply twiddle factors
 3. Partially transpose the matrix*
 4. Perform NTTs on the rows*
- Each MPI process works on a portion of the global matrix
 - OpenMP threads do the work in parallel

¹David H. Bailey (1990). "FFTs in external or hierarchical memory".

- ① Each process starts with a slice of columns $H \times w$
 - Compute NTTs of the columns in parallel and apply twiddle factors
 - Master thread exchanges columns with other processes
- ② Once all blocks have been exchanged, rearrange data into slice of rows
 - Can be seen as transposing a $P \times w$ matrix, moving h elements at a time
 - Done in-place by following cycles and utilising tasks for parallelism
- ③ Perform NTTs on the rows in bulk



- NTTs on the columns done with Stockham algorithm¹
 - Requires an additional buffer, but **output is in-order**
- NTTs on the rows done in bulk with DIF/DIT algorithms¹
 - In-place algorithms (**no additional buffer required**)
 - Bit-scrambling unnecessary
 - Two possibilities for parallelisation
 - i. Subdivide the rows amongst the threads
 - ii. **Parallelise the butterfly loops**

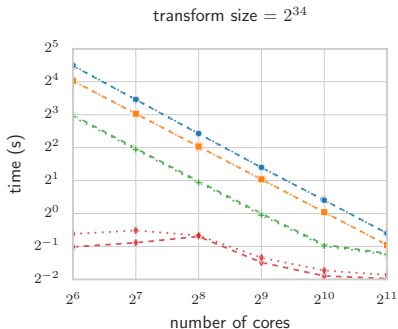
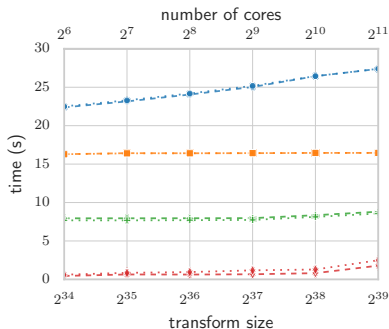
▪ Typical loop structure for DIF

```
for(ulong k = m/2; k > 0; k >>= 1) {
    for(ulong j = 0; j < k; j++)
        for(ulong i = j; i < m; i += 2*k) {
            /* ... */
        }
}
```

▪ DIF with butterfly loops collapsed

```
for(ulong k = m/2; k > 0; k >>= 1) {
    #pragma omp for schedule(...)
    for(ulong h = 0; h < m/2; h++) {
        ulong i = h/k, j = h&(k-1);
        /* ... */
    }
}
```



¹Richard Crandall & Carl Pomerance, “Fast Algorithms for Large-Integer Arithmetic”, in *Prime Numbers: A Computational Perspective*.




- Row transforms
- Column transforms
- + Others
- ◆ MPI Overhead
- - Forward transform
- ⋯ Inverse transform

†On Fulhame (HPE Apollo 70).

- Isambard 1 (Cray XC50) shows $\sim 10\%$ slowdown wrt Fulhame
- SVE delivers over 3.5x speedup on Isambard 2 (HPE Apollo 80)



- **Counts up to 2^{40}** in less than 10 minutes
- In the near future, our **goal is to reach at least 2^{45}**
- The analysis of these results will be the subject of another publication



- **SVE** shows great potential **for methods that utilise modular arithmetic**
 - Easily achieved speedups greater than 4
 - Competitive even for extremely short transform sizes
- Our **NTT methods scale well** on the Arm-based systems we tested
 - **MPI almost perfectly overlapped** with computation
 - Enables **fast, efficient, and exact convolutions**
 - Underlying ideas can be applied to FFTs
- These methods have been put into action to **compute the number of Goldbach partitions** of all even numbers up to 2^{40} (soon 2^{45})

Questions?