

Error Analysis of Matrix Multiplication with Narrow Range Floating-Point Arithmetic

Mantas Mikaitis

School of Computer Science, University of Leeds, Leeds, UK

Work with Theo Mary, CNRS, Paris, France

The Chinese Academy of Sciences

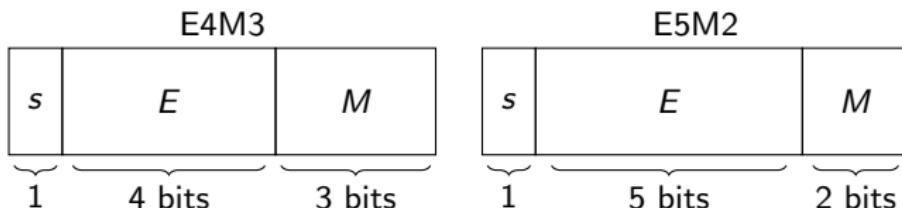
Workshop on Approximate computing in Numerical Linear Algebra

Beijing, China

24 April, 2025



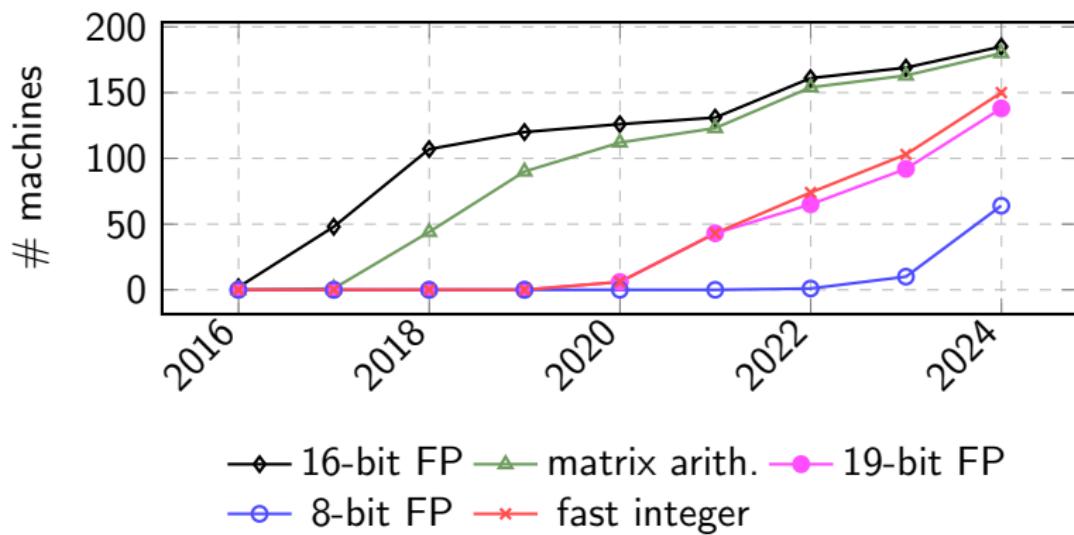
8-bit floating-point formats



Differences among hardware architectures and standards:

- OCP standard (<https://bit.ly/3G9EsyG>). E5M2 similar to IEEE 754 formats. E4M3: no infinities, one signed NaN.
- IEEE p3109 (interim report <https://bit.ly/42gPwcy>). No -0 . One unsigned NaN.
- NVIDIA (PTX ISA 8.7 <https://bit.ly/3RNElve>) follows the OCP spec.
- AMD does not fully follow the OCP spec. for NaN and inf.
- **Standardisation work ongoing.**

8-bit floating point on the TOP500



Devices counted: P100, V100, A100, H100, MI210, MI250X, MI300X, Intel Data Center GPU, from <https://www.top500.org>.

With NVIDIA Blackwell 4/6-bit FP will appear.

GPUs in the TOP500

98% of the top 1 Frontier's power comes from GPUs.
If you don't use GPUs, go home!

J. Dongarra presenting in Manchester, 2022

(Not a direct quote)

4/6/8/16-bit floating point formats have narrow ranges

Format	precision	min pos.	max pos.	u
binary64 (double)	53	2^{-1022}	$\sim 1.798 \times 10^{308}$	2^{-53}
binary32 (single)	24	2^{-126}	$\sim 3.403 \times 10^{38}$	2^{-24}
tf32 (19-bit)	11	2^{-126}	$\sim 3.401 \times 10^{38}$	2^{-11}
bfloating16	8	2^{-126}	$\sim 3.389 \times 10^{38}$	2^{-8}
binary16	11	2^{-14}	65504	2^{-11}
fp8-E4M3	4	2^{-6}	448	2^{-4}
fp8-E5M2	3	2^{-14}	57344	2^{-3}
fp6-E2M3	4	2^0	7.5	2^{-4}
fp6-E3M2	3	2^{-2}	28	2^{-3}
fp4-E2M1	2	2^0	6	2^{-2}

Mixed-precision matrix multipliers

Formats with narrow ranges are available in matrix multiply operation.

$$D = C + A \times B,$$
$$\underbrace{\begin{bmatrix} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{bmatrix}}_{\text{binary16 or binary32}} = \underbrace{\begin{bmatrix} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{bmatrix}}_{\text{binary16 or binary32}} + \underbrace{\begin{bmatrix} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{bmatrix}}_{\text{8-bit FP}} \times \underbrace{\begin{bmatrix} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{bmatrix}}_{\text{8-bit FP}}$$

Hardware level differences

- Example above is 4×4 , but dimensions differ across architectures.
- Internal dot product precision, rounding, and subnormal support.

Some of these will not affect our model, but we use round-to-nearest and parameterize subnormal support.

Mixed-precision matrix multipliers

Architecture	Input format	Accumulation format
NVIDIA PTX ISA	fp8-E5M2	binary32
	fp8-E4M3	binary32
	binary16	binary16
	binary16	binary32
	bfloat16	binary32
	19-bit FP	binary32
AMD MI300 ISA	fp8-E5M2	binary32
	fp8-E4M3	binary32
	binary16	binary32
	bfloat16	binary32
	19-bit FP	binary32

Matrix Multiply-Accumulate (MMA)

Model 1

The following model describes a mixed-precision MMA operation to compute $C = AB$, assuming round-to-nearest ties-to-even is used. We have two FP formats:

- *Input format* with precision t , unit roundoff $u = 2^{-t}$, exponent in $[e_{\min}, e_{\max}]$, range of normalized values $\pm[f_{\min}, f_{\max}]$. The maximum distance between any number in $[-f_{\min}, f_{\min}]$ and the nearest FP number is

$$g_{\min} = \begin{cases} f_{\min}/2 & \text{if subnormals are not available} \\ uf_{\min} & \text{with subnormals (gradual underflow)} \end{cases}$$

- *Accumulation format* with $T \geq t$, $U = 2^{-T}$, exponent in $[E_{\min}, E_{\max}] \supseteq [e_{\min}, e_{\max}]$, and range of norm. numbers $\pm[F_{\min}, F_{\max}]$. The maximum distance between any number in $[-F_{\min}, F_{\min}]$ and the nearest FP number is

$$G_{\min} = \begin{cases} F_{\min}/2 & \text{if subnormals are not available} \\ UF_{\min} & \text{with subnormals (gradual underflow)} \end{cases}$$

Models of worst-case rounding errors

Rounding error model based on [Demmel, 1984]

Take $x, y \in \mathbb{R}$. Assuming no overflows occur, the rounding operator to the *input format* is described as

$$\text{fl}(x) = x(1 + \delta) + \eta, \quad |\delta| \leq u, \quad |\eta| \leq g_{\min}, \quad \eta\delta = 0,$$

and arithmetic operations in the *accumulation format* as

$$\text{FL}(x \text{ op } y) = (x \text{ op } y)(1 + \delta) + \eta, \quad |\delta| \leq U, \quad |\eta| \leq G_{\min}, \quad \eta\delta = 0,$$

with $\text{op} \in \{+, -, \times, \div\}$.

Here $\eta\delta = 0$ accounts for only one type of error, rounding or overflow.

Part 1: Basic single-word algorithm

Single-word algorithm

- ① Scale input matrices A and B .
- ② Round input matrices to the *input format*.
- ③ Multiply scaled and rounded A and B in the *accumulation format*.
- ④ Scale the output matrix.

$$C = \Lambda^{-1} \left(\text{fl}(\Lambda A) \text{fl}(BM) \right) M^{-1}$$

- Λ and M are nonsingular diagonal matrices with diagonal coefficients λ_i and μ_i respectively.
- Scale coefficients λ_i and μ_i are powers of two.

Single-word algorithm

Let θ be the maximum value we can afford in the scaled A and B .

Scaling by powers of two means the maximum entry per row of A or column of B is in $(\theta/2, \theta]$.

We should maximise θ to reduce number of underflows, but at the same time remove possibility of overflow.

Choose:

$$\theta = \min(f_{\max}, \sqrt{F_{\max}/n}).$$

which avoids overflow in the input and in the accumulation of n products.

Single-word algorithm: an example

- Take $A \in \mathbb{R}^{4 \times 4}$ and $B \in \mathbb{R}^{4 \times 4}$.
- Set fp8-E4M3 as the *input format* with $f_{\max} = 448$.
- Set binary16 as the *accumulation format* with $F_{\max} = 65504$.
- No subnormal floating-point numbers.
- This gives $\min(448, \sqrt{65504/4}) = \min(448, 127.9687) \approx 127 = \theta$.

Scaling factors

In this case before rounding matrices to the *input format* we need to scale them such that 127 is the maximum value that appears.

- 127 is lower than $f_{\max} = 448$ - no *input format* overflows.
- $127 \times 127 = 16129$ and if we accumulate four such products we get $64616 < F_{\max} = 65504$. No *accumulation format* overflows.

Single-word algorithm: an example

Take

$$A = \begin{bmatrix} 500 & 1 & 1 & 2^{-6} \\ 128 & 128 & 128 & 128 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \end{bmatrix}.$$

We have

$$AB = \begin{bmatrix} 502.015625 & 64258 & 502.015625 & 502.015625 \\ 512 & 65536 & 512 & 512 \\ 4 & 512 & 4 & 4 \\ 4 & 512 & 4 & 4 \end{bmatrix}.$$

Overflows in the above example if no scaling is applied

(Input) $500 > f_{\max} = 448$ and (output) $65536 > F_{\max} = 65504$.

Single-word algorithm: an example

$$C = \Lambda^{-1} (\text{fl}(\Lambda A) \text{fl}(BM)) M^{-1}, \quad \theta = 127$$

Step 1: Scale A and B .

$$\Lambda A = \begin{bmatrix} 2^{-2} & 0 & 0 & 0 \\ 0 & 2^{-1} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 500 & 1 & 1 & 2^{-6} \\ 128 & 128 & 128 & 128 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 2^{-8} \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

$$BM = \begin{bmatrix} 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2^{-1} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix}$$

How the scale coefficients are calculated

For example, take the first row of A . The largest value is 500 and we need to get it below $\theta = 127$. $\lambda_1 = 2^{\lfloor \log_2(127/500) \rfloor} = 2^{-2}$.

Single-word algorithm: an example

$$C = \Lambda^{-1} \left(\text{fl}(\Lambda A) \text{fl}(BM) \right) M^{-1}$$

Step 2: Round to the *input format* fp8-E4M3 ($f_{\min} = 2^{-6}$).

$$\begin{aligned}\text{fl}(\Lambda A) &= \text{fl} \left(\begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 2^{-8} \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \right) = \begin{bmatrix} 125 & 2^{-2} & 2^{-2} & \mathbf{0} \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \\ \text{fl}(BM) &= \text{fl} \left(\begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix} \right) = \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix}\end{aligned}$$

Underflow in the above example

Notice that since subnormals are off, numbers $\leq f_{\min}/2$ will round to zero, causing underflow. This happened to $\Lambda A(1, 4) = 2^{-8}$, which resulted from scaling the first row of A , where originally $A(1, 4) = 2^{-6}$.

Single-word algorithm: an example

$$C = \Lambda^{-1} \left(\text{fl}(\Lambda A) \text{fl}(BM) \right) M^{-1}$$

Step 3: Perform matrix multiply in the *accumulation format* binary16 ($T = 11$, $F_{\max} = 65504$).

$$\begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 0 \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 125.5 & 8032 & 125.5 & 125.5 \\ 256 & 16384 & 256 & 256 \\ 4 & 256 & 4 & 4 \\ 4 & 256 & 4 & 4 \end{bmatrix}$$

Single-word algorithm: an example

$$C = \Lambda^{-1} \left(\text{fl}(\Lambda A) \text{fl}(BM) \right) M^{-1}$$

Step 4: Undo the scaling.

$$\begin{bmatrix} 4 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 125.5 & 8032 & 125.5 & 125.5 \\ 256 & 16384 & 256 & 256 \\ 4 & 256 & 4 & 4 \\ 4 & 256 & 4 & 4 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} =$$

$$\begin{bmatrix} 502 & 64256 & 502 & 502 \\ 512 & 65536 & 512 & 512 \\ 4 & 512 & 4 & 4 \\ 4 & 512 & 4 & 4 \end{bmatrix}$$

Single-word algorithm: an example

$$C = \Lambda^{-1} \left(\text{fl}(\Lambda A) \text{fl}(BM) \right) M^{-1}$$

Comparison. Our result computed with mixed-precision MMA:

$$AB \approx \begin{bmatrix} 502 & 64256 & 502 & 502 \\ 512 & 65536 & 512 & 512 \\ 4 & 512 & 4 & 4 \\ 4 & 512 & 4 & 4 \end{bmatrix}$$

And the exact result

$$AB = \begin{bmatrix} 502.015625 & 64258 & 502.015625 & 502.015625 \\ 512 & 65536 & 512 & 512 \\ 4 & 512 & 4 & 4 \\ 4 & 512 & 4 & 4 \end{bmatrix}$$

Single-word algorithm: the new error bound

The previous bound of [Blanchard et. al., 2020] was developed without considering range limitations (8-bit FP was not available then):

$$\|\hat{C} - AB\|_\infty \lesssim (2u + nU)\|A\|_\infty\|B\|_\infty.$$

Our analysis adds two extra terms for two types of underflow:

$$\|\hat{C} - AB\|_\infty \lesssim \left(2u + nU + 4n^2\theta^{-1}g_{\min} + 4n^2\theta^{-2}G_{\min}\right)\|A\|_\infty\|B\|_\infty.$$

Our example

Input format: fp8-E4M3, accumulation: binary16, no subnormals.

- $2u = 2 \times 2^{-4} = 0.125$
- $nU = 4 \times 2^{-11} = 2^{-9}$
- $4n^2\theta^{-1}g_{\min} = 4 \times 16 \times (1/127) \times 2^{-6}/2 = 0.00394$
- $4n^2\theta^{-2}G_{\min} = 4 \times 16 \times 1/127^2 \times 2^{-14}/2 \approx 0$

Part 2: Multi-word algorithm

Double-word algorithm: example

Again, for a step-by-step example, take

$$A = \begin{bmatrix} 500 & 1 & 1 & 2^{-6} \\ 128 & 128 & 128 & 128 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \end{bmatrix}.$$

Double-word algorithm: an example

Step 1: Scale A and B (same as before).

$$\Lambda A = \begin{bmatrix} 2^{-2} & 0 & 0 & 0 \\ 0 & 2^{-1} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 500 & 1 & 1 & 2^{-6} \\ 128 & 128 & 128 & 128 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} = \begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 2^{-8} \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

$$BM = \begin{bmatrix} 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \\ 1 & 128 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2^{-1} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix}$$

Double-word algorithm: an example

Step 2: Round to the *input format*, in **double-word representation**.

We will round each ΛA and BM to two fp8-E4M3 matrices instead of one.

Compute the first word (first of the two matrices):

$$A^{(0)} = \text{fl}(\Lambda A) = \text{fl} \left(\begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 2^{-8} \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \right) = \begin{bmatrix} 125 & 2^{-2} & 2^{-2} & \mathbf{0} \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix}$$

$$B^{(0)} = \text{fl}(BM) = \text{fl} \left(\begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix} \right) = \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix}$$

Double-word algorithm: an example

Step 2: Round to the *input format* fp8-E4M3, in **double-word representation**.

Compute the second word (rounding/underflow error in the first step):

$$A^{(1)} = \text{fl}((\Lambda A - A^{(0)})/u^1) = \\ \text{fl} \left(\left(\begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 2^{-8} \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} - \begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 0 \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \right) ./2^{-4} \right) = \begin{bmatrix} 0 & 0 & 0 & 2^{-4} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

Since $B^{(0)} = BM$, $B^{(1)} = \text{zeros}(4, 4)$.

Extra scaling

Notice the division by $u^1 = 2^{-4}$ before rounding, which is done to reduce underflows in the input format. In general, the multi-word split is

$$A^{(i)} = \text{fl} \left(\left(\Lambda A - \sum_{k=0}^{i-1} u^k A^{(k)} \right) / u^i \right).$$

Double-word algorithm: an example

Step 3: Perform matrix products and add them in the *accumulation format* binary16.

p-word case

After splitting ΛA and BM into $A^{(0)}, \dots, A^{(p-1)}$ and $B^{(0)}, \dots, B^{(p-1)}$, approximate matrix multiply by $p(p + 1)/2$ products

$$C \approx \Lambda^{-1} \left(\sum_{i+j < p} u^{i+j} A^{(i)} B^{(j)} \right) M^{-1}.$$

In our double-word case

$$A^{(0)}B^{(0)} + uA^{(1)}B^{(0)} =$$

$$\begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 0 \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 2^{-4} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix}$$

Double-word algorithm: an example

$$A^{(0)}B^{(0)} + uA^{(1)}B^{(0)} =$$

$$\begin{bmatrix} 125 & 2^{-2} & 2^{-2} & 0 \\ 64 & 64 & 64 & 64 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 2^{-4} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \\ 1 & 64 & 1 & 1 \end{bmatrix} =$$

$$\begin{bmatrix} 125.5 & 8032 & 125.5 & 125.5 \\ 256 & 16384 & 256 & 256 \\ 4 & 256 & 4 & 4 \\ 4 & 256 & 4 & 4 \end{bmatrix} + \begin{bmatrix} 2^{-8} & 0.25 & 2^{-8} & 2^{-8} \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} =$$

$$\begin{bmatrix} 125.50390625 & 8032.25 & 125.50390625 & 125.50390625 \\ 256 & 16384 & 256 & 256 \\ 4 & 256 & 4 & 4 \\ 4 & 256 & 4 & 4 \end{bmatrix}$$

Double-word algorithm: an example

$$C \approx \Lambda^{-1} \left(\sum_{i+j < p} u^{i+j} A^{(i)} B^{(j)} \right) M^{-1}.$$

Step 4: Undo the scaling.

$$\begin{bmatrix} 4 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 125.50390625 & 8032.25 & 125.50390625 & 125.50390625 \\ 256 & 16384 & 256 & 256 \\ 4 & 256 & 4 & 4 \\ 4 & 256 & 4 & 4 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} =$$

$$\begin{bmatrix} 502.015625 & 64258 & 502.015625 & 502.015625 \\ 512 & 65536 & 512 & 512 \\ 4 & 512 & 4 & 4 \\ 4 & 512 & 4 & 4 \end{bmatrix} = AB.$$

Double-word algorithm: the new error bound

The previous bound of [Fasi et. al., 2022] without the range limitations (no 8-bit FP available at the time):

$$\|\widehat{C} - AB\|_\infty \lesssim \left((p+1)u^p + (n+p^2)U \right) \|A\|_\infty \|B\|_\infty.$$

Our analysis recovered the rounding error terms and added two terms for the underflows:

$$\begin{aligned} \|\widehat{C} - AB\|_\infty &\lesssim \left((p+1)u^p + 4nu^{p-1}\theta^{-1}g_{\min} \right. \\ &\quad \left. + (n+p^2)U + 2p(p+1)n^2\theta^{-2}G_{\min} \right) \|A\|_\infty \|B\|_\infty. \end{aligned}$$

Error analysis: summary

Single-word algorithm:

$$\|\widehat{C} - AB\|_\infty \lesssim \left(2u + nU + 4n^2\theta^{-1}g_{\min} + 4n^2\theta^{-2}G_{\min} \right) \|A\|_\infty \|B\|_\infty.$$

p -word algorithm:

$$\begin{aligned} \|\widehat{C} - AB\|_\infty &\lesssim \left((p+1)u^p + 4nu^{p-1}\theta^{-1}g_{\min} \right. \\ &\quad \left. + (n+p^2)U + 2p(p+1)n^2\theta^{-2}G_{\min} \right) \|A\|_\infty \|B\|_\infty. \end{aligned}$$

Part 3: Numerical experiments

Simulating mixed precision with CPFloat in MATLAB/Octave

Paper in ACM TOMS provides details. [Fasi and Mikaitis, 2023]

<https://github.com/north-numerical-computing/cpfloat>

Example use:

```
>> input.format = 'q43';
>> input.subnormal = 0;
>> accum.format = 'binary16';
>> accum.subnormal = 0;
>> cpfloat(pi, input)
ans =
    3.250000000000000
>> cpfloat(cpfloat(
    cpfloat(pi, input)*cpfloat(pi, input), accum)+0.5, accum)
ans =
    11.062500000000000
```

Numerical experiments

We generate $A \in \mathbb{R}^{10 \times n}$ and $B \in \mathbb{R}^{n \times 10}$ and vary n .

Elements in $[-10^{10}, -10^{-10}] \cup [10^{-10}, 10^{10}]$.

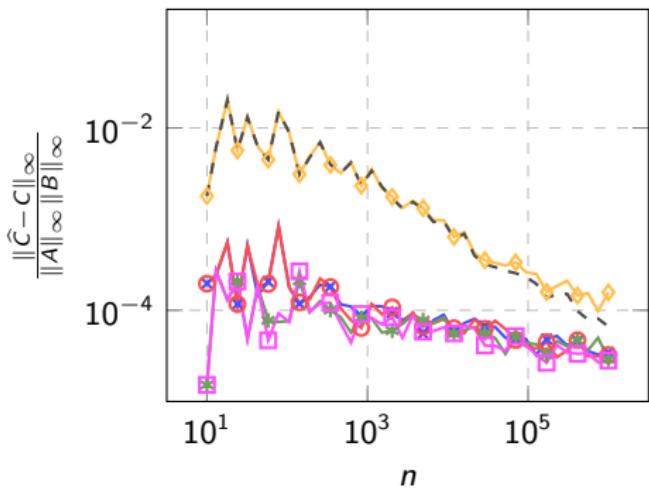
Measure the accuracy with $\frac{\|\hat{C} - C\|_\infty}{\|A\|_\infty \|B\|_\infty}$ where C is computed in binary64.

We check with subnormals on/off to detect any improvements due to *gradual underflow*.

We also plot the variants of MMA without any range (exponent) limitations, which CPFloat makes easy to do.

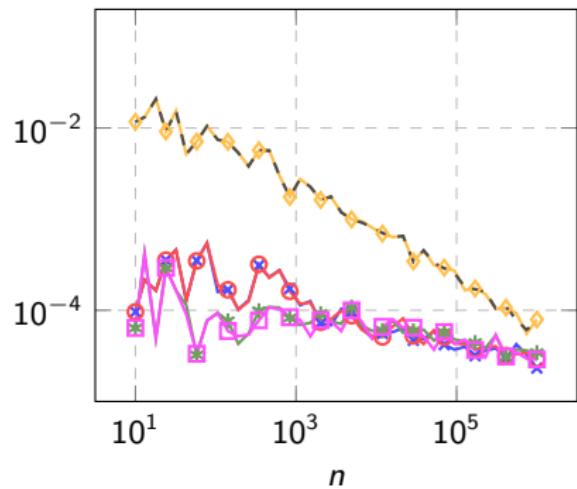
Numerical experiments I

fp8-E4M3 input
binary16 accumulation
subnormals off



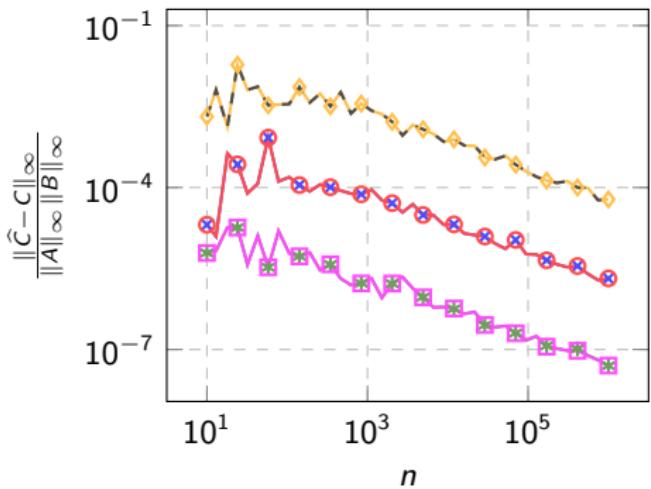
- ◇— $p = 1$ --- $p = 1$ unbounded range
- *— $p = 2$ —○— $p = 2$ unbounded range
- *— $p = 3$ —□— $p = 3$ unbounded range

fp8-E4M3 input
binary16 accumulation
subnormals on

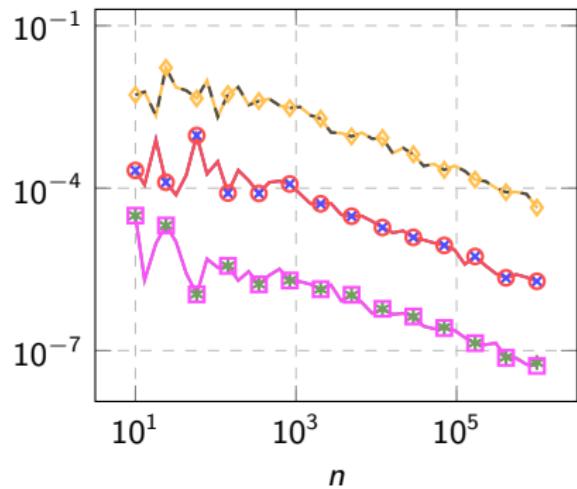


Numerical experiments II

fp8-E4M3 input
binary32 accumulation
subnormals off

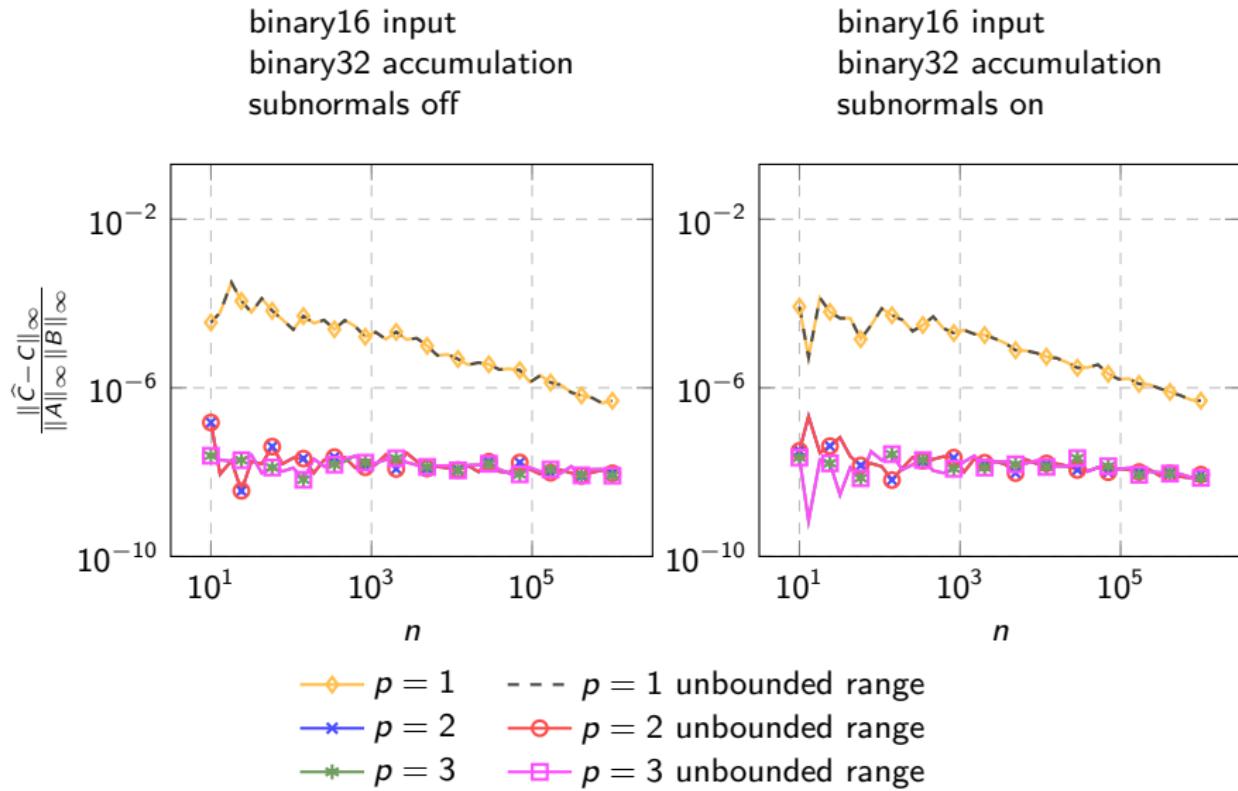


fp8-E4M3 input
binary32 accumulation
subnormals on



- ◇— $p = 1$ --- $p = 1$ unbounded range
- *— $p = 2$ —○— $p = 2$ unbounded range
- *— $p = 3$ —□— $p = 3$ unbounded range

Numerical experiments III



Summary

- Underflows in narrow-range FP formats not a problem, provided three types of scaling are used.
- Shown algorithms can be used to obtain binary32 accuracy in high performance.
- If higher accuracy is needed, MMA can still be used in conjunction with binary64—see the next talk.

SIAM SISC paper

T. Mary, and M. Mikaitis. *Error Analysis of Matrix Multiplication with Narrow Range Floating-Point Arithmetic*. **Preprint, v2. Accepted for SIAM SISC**. Mar. 2025.



<https://bit.ly/42dqpkn>.

Slides and more info at <http://mmikaitis.github.io>

References I



J. Demmel

Underflow and the reliability of numerical software

SIAM J. Sci. Comput., 5:4. Dec. 1984.



P. Blanchard, N. J. Higham, F. Lopez, T. Mary, and S. Pranesh

Mixed precision block fused multiply-add: Error analysis and application to GPU tensor cores

SIAM J. Sci. Comput., 42:3. Jan. 2020.



M. Fasi, N. J. Higham, F. Lopez, T. Mary, and M. Mikaitis

Matrix Multiplication in Multiword Arithmetic: Error Analysis and Application to GPU Tensor Cores

SIAM J. Sci. Comput., 45:1. Feb. 2023.



M. Fasi and M. Mikaitis

CPFloat: A C library for Simulating Low-Precision Arithmetic

ACM Trans. Math. Software, 49. 2023