

Adaptive Precision Sparse Matrix-Vector Product and its application to Krylov Solvers

Roméo Molina, Stef Graillat, Fabienne Jézéquel, Théo Mary

► To cite this version:

Roméo Molina, Stef Graillat, Fabienne Jézéquel, Théo Mary. Adaptive Precision Sparse Matrix-Vector Product and its application to Krylov Solvers. 13èmes Rencontres Arithmétique de l'Informatique Mathématique (RAIM 2022), Nov 2022, Nantes, France. hal-03931125

HAL Id: hal-03931125 https://hal.science/hal-03931125

Submitted on 9 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

RAIM 2022

3rd November 2022

Adaptive Precision Sparse Matrix–Vector Product

and its Application to Krylov Solvers

Roméo Molina

LIP6, Sorbonne Université Service Online, Département Informatique, IJCLab

Joint work with

Stef Graillat, Fabienne Jézéquel, and Theo Mary

Today's floating-point landscape

Bits					
		Signif. (t) Exp.	Range	$u = 2^{-t}$
bfloat16	В	8	8	10 ^{±38}	$4 imes 10^{-3}$
fp16	Η	11	5	$10^{\pm 5}$	$5 imes 10^{-4}$
fp32	S	24	8	10 ^{±38}	$6 imes 10^{-8}$
fp64	D	53	11	$10^{\pm 308}$	$1 imes 10^{-16}$
fp128	Q	113	15	$10^{\pm 4932}$	$1 imes 10^{-34}$

• Low precision increasingly supported by hardware

• Great benefits:

- Reduced storage, data movement, and communications
- Reduced energy consumption (5× with fp16, 9× with bfloat16)
- $\circ~$ Increased speed on emerging hardware (16 \times on A100 from fp32 to fp16/bfloat16)

• Some limitations too:

- Low accuracy (large *u*)
- Narrow range

Mix several precisions in the same code with the goal of

- Getting the performance benefits of low precisions
- While preserving the accuracy and stability of high precision

Terminology varies: Mixed precision, Multiprecision, Adaptive precision, Variable precision, Transprecision, Dynamic precision, ...

Mix several precisions in the same code with the goal of

- Getting the performance benefits of low precisions
- While preserving the accuracy and stability of high precision

Terminology varies: Mixed precision, Multiprecision, Adaptive precision, Variable precision, Transprecision, Dynamic precision, ...

How to select the right precision for the right variable/operation

- **Precision tuning:** autotuning based on the source code, my thesis area: CADNA / PROMISE...
 - ▲ Does not need any understanding of what the code does
 - Does not have any understanding of what the code does
- This work: another point of view, exploit as much as possible the knowledge we have about the code

Adaptive precision algorithms

- Given an algorithm and a prescribed accuracy ε , adaptively select the minimal precision for each computation
- \Rightarrow Why does it make sense to make the precision vary?

Adaptive precision algorithms

- Given an algorithm and a prescribed accuracy ε , adaptively select the minimal precision for each computation
- \Rightarrow Why does it make sense to make the precision vary?
 - Because not all computations are equally "important" ! Example:



and small elements produce small errors :

 $|\operatorname{\mathsf{fl}}(\operatorname{\mathsf{a}}\operatorname{\operatorname{op}}\operatorname{\mathsf{b}}) - \operatorname{\mathsf{a}}\operatorname{\operatorname{op}}\operatorname{\mathsf{b}}| \le u|\operatorname{\mathsf{a}}\operatorname{\operatorname{op}}\operatorname{\mathsf{b}}|, \quad \operatorname{\operatorname{op}} \in \{+, -, *, \div\}$

 $\Rightarrow \mbox{ Opportunity for mixed precision: adapt the precisions to the data at hand by storing and computing "less important" (usually smaller) data in lower precision$

- Pushing adaptive precision to the extreme: can we benefit from storing **each variable** in a (possibly) different precision?
- Example: Ax = b with adaptive precision for each A_{ij}
 - Is it worth it?

Need to have elements of widely different magnitudes

• Is it practical?

Probably not for compute-bound applications, but could it work for memory-bound ones?

 \Rightarrow Natural candidate: sparse matrices

Sparse matrix-vector product (SpMV)

 $y = Ax, A \in \mathbb{R}^{m \times n}$ for i = 1: m do $y_i = 0$ for $j \in nnz_i(A)$ do $y_i = y_i + a_{ij}x_j$ end for
end for

• Standard error analysis for y = Ax performed in a uniform precision ε gives,

$$|\widehat{y}_i - y_i| \le n_i \varepsilon \sum_{j \in nnz_i(A)} |a_{ij}x_j|$$

• **Idea:** store elements of A in a precision inversely proportional to their magnitude (smaller elements in lower precision)

Adaptive precision SpMV

for i = 1: m do $y_i = 0$ for k = 1: p do $y_i^{(k)} = 0$ for $j \in nnz_i(A)$ do if $a_{ij}x_j \in B_{ik}$ then $y_i^{(k)} = y_i^{(k)} + a_{ij}x_j$ at precision u_k end if end for $y_i = y_i + y_i^{(k)}$ end for end for end for

- Split row *i* of *A* into *p* buckets *B_{ik}* and sum elements of *B_{ik}* in precision *u_k*
- Error analysis: $|\widehat{y}_i^{(k)} y_i^{(k)}| \le n_i^{(k)} u_k \sum_{a_{ij} \times_j \in B_{ik}} |a_{ij} \times_j|$

Building the buckets

•
$$|\widehat{y}_{i}^{(k)} - y_{i}^{(k)}| \le n_{i}^{(k)} u_{k} \sum_{a_{ij} \times_{j} \in B_{ik}} |a_{ij} \times_{j}|$$

 \Rightarrow Build the buckets such that $u_k \sum_{a_{ij} \times_j \in B_{ik}} |a_{ij} \times_j| \approx \varepsilon \sum_j |a_{ij} \times_j|$

- By setting B_{ik} to the interval $(\varepsilon \beta_i / u_{k+1}, \varepsilon \beta_i / u_k]$, we obtain $|\widehat{y}_i^{(k)} y_i^{(k)}| \le n_i^{(k)} \varepsilon \beta_i$ and so $|\widehat{y}_i y_i| \le n_i \varepsilon \beta_i$
- Two possible choices for β_i :

 $\circ \ \ \beta_i = \sum_j |a_{ij}x_j| \Rightarrow \text{guarantees } O(\varepsilon) \text{ componentwise error:}$ $|\widehat{y_i} - y_i| \le n\epsilon \sum_j |a_{ij}x_j| \quad \forall i \in \{1, ..., n\} \\ \circ \ \ \beta_i = \|A\| \|x\| \Rightarrow \text{guarantees } O(\varepsilon) \text{ normwise error:} \\ ||\widehat{y_i} - y_i|| \le n\epsilon \|A\| \|x\|$

• Smallest elements are simply dropped











• 33 matrices coming from SuiteSparse collection and industrial partners with at most 166M non-zeros

- 33 matrices coming from SuiteSparse collection and industrial partners with at most 166M non-zeros
- 3 different accuracy targets:

$$\epsilon = 2^{-24}$$

$$\epsilon = 2^{-37}$$

$$\epsilon = 2^{-53}$$

SpMV experimental settings

Various sets of precision formats:

- 2 precisions: fp32, fp64
- 3 precisions: bfloat16, fp32, fp64
- 7 precisions: bfloat16, "fp24", fp32, fp64, "fp40", "fp48", "fp56"

	Bits		
	Mantissa	Exponent	
bfloat16	8	8	
"fp24"	8	8	
fp32	24	8	
"fp40"	29	11	
"fp48"	37	11	
"fp56"	45	11	
fp64	53	11	

Maintaining componentwise accuracy



Adaptive methods preserve an accuracy close to the accuracy of uniform methods.

Maintaining normwise accuracy



Adaptive methods preserve an accuracy close to the accuracy of uniform methods.

Theoretical storage gains targetting $\epsilon = 2^{-53}$ accuracy (fp64)



Up to 88% of storage reduction

Actual time gains targetting $\epsilon = 2^{-53}$ accuracy (fp64)



Up to 85% of time reduction

Theoretical storage gains targetting $\epsilon = 2^{-24}$ accuracy (fp32)



Up to 97% of storage reduction

Actual time gains targetting $\epsilon = 2^{-24}$ accuracy (fp32)



Up to 88% of time reduction



We are able to target any kind of accuracy with only natively supported precisions.



We are able to target any kind of accuracy with only natively supported precisions.

Performance of GMRES rely on SpMV

 $\begin{aligned} r &= b - Ax_0 \\ \beta &= \|r\|_2 \\ q_1 &= r/\beta \\ \text{for } k &= 1, 2, \dots \text{ do} \\ y &= Aq_k \\ \text{for } j &= 1: k \text{ do} \\ h_{jk} &= q_j^T y \\ y &= y - h_{jk}q_j \\ \text{end for} \\ h_{k+1,k} &= \|y\|_2 \\ q_{k+1} &= y/h_{k+1,k} \\ \text{Solve the least squares problem } \min_{c_k} \|Hc_k - \beta e_1\|_2 \\ x_k &= x_0 + Q_k c_k \\ \text{end for} \end{aligned}$

How does the adaptive method affect the convergence?

GMRES-IR experimental settings

- Adaptive SpMV achieve larger speedups for lower accuracy targets
- GMRES-based iterative refinement particularly attractive

1: for
$$i = 1, 2, ...$$
 do
2: $r_i = b - Ax_{i-1}$
3: Solve $Ad_i = r_i$ by GMRES
4: $x_i = x_{i-1} + d_i$
5: end for

- SpMV line 2 in CW adaptive, target $\epsilon = 2^{-53}$
- GMRES line 3 experimented with multiple variants
- Use of simple Jacobi preconditioner
- Fixed outside threshold $\epsilon_{out} = 2^{-53}$
- Varying inside threshold $\epsilon_{in} = 2^{-24}$
- Fixed inside restart to 80 iterations



- With reasonable accuracy targets, adaptive SpMV, does not affect the confergence scheme
- Choose the best compromise between the iteration cost and the number of iterations

Application to GMRES: convergence scheme experiments



- For a given ϵ_{in} value, NW variants achieve a lower cost but a slower convergence than CW ones
- $\epsilon_{in} = 2^{-24}$ leads to the best NW variant, which converges in 1040 iterations with an SpMV cost of 36% of the fp32 uniform one
- $\epsilon_{in} = 2^{-20}$ leads to the best CW variant, which converges in 320 iterations with a corresponding SpMV cost of 56%
- 18/19 Both options should be considered.



- Surprising behavior
- Consistently reproduced and occurs for several other matrices in our set
- Aggressive dropping of small coefficients might lead to a "nicer" matrix for which GMRES can converge quickly.

Conclusion: take-home messages

• Adaptive precision SpMV algorithm

- Ruilds computing buckets according to the elements magnitude
- $\circ~$ Targets any accuracy
- Matrix-dependent gains
- Application to Krylov solvers
 - $\circ\;$ Reasonable accuracy targets preserves convergence scheme

More info



Thank you! Any questions?