

Accelerating Density Fitting with Adaptive-precision and 8-bit Integer on AI Accelerators

Hua Huang^{*1}, Wenkai Shao², and Jeff Hammond³

¹NVIDIA, Santa Clara, USA 94086

²Individual researcher, Anqing, China 246000

³NVIDIA, Helsinki, Finland 00180

*Email: huah@nvidia.com

Abstract

The emergence of artificial intelligence (AI) accelerators like NVIDIA Tensor Cores offers new opportunities to speed up tensor-heavy scientific computations. However, applying them to quantum chemistry is challenging due to strict accuracy demands and irregular data patterns. We propose an adaptive precision algorithm to accelerate the density fitting (DF) method with Gaussian basis sets on AI accelerators using 8-bit integer (INT8) arithmetics. Implemented in the GPU-accelerated PySCF package, the algorithm is tested on more than twenty molecular systems with different NVIDIA GPUs. Compared to the standard FP64 code, our algorithm is up to 204% faster on a RTX 4090 gaming GPU and up to 364% faster on a RTX 6000 Ada workstation GPU without compromising the converged energy. This work demonstrates a practical approach to use AI hardware for reliable quantum chemistry simulations.

Keywords

density fitting, mixed-precision, matrix multiplication, AI accelerator

Introduction

The recent surge in artificial intelligence (AI) applications has driven the development of hardware accelerators with extremely high throughput for low-precision matrix operations. Notably, NVIDIA Tensor Cores¹⁻³ and Google TPUs⁴ offer order-of-magnitude improvements in dense general matrix multiplication (GEMM) performance using half- and single-precision and even lower precision arithmetic. While these accelerators have transformed deep learning, their impact on scientific computing is limited. One major barrier is that scientists typically consider double precision (FP64), which is also supported on Tensor Cores on some GPUs, as the standard for scientific computing. Moreover, although GEMM is a common computational primitive in both AI and scientific computing, it does not encompass the spectrum of mathematical operations needed in scientific computing.

To address these challenges, researchers have proposed a variety of single- and mixed-precision algorithms for scientific computing. These include general purpose methods, such as mixed-precision solvers for systems of linear equations⁵⁻⁷, as well as techniques that emulate a higher precision GEMM using multiple lower precision GEMMs⁸⁻¹². In addition, domain-specific low- and mixed-precision algorithms have also been developed. For instance, molecular dynamics simulations often perform force calculations in single precision while retaining double precision for critical steps like integration and energy accumulation. However, due to the complexity and irregularity of their computational patterns, most of these algorithms cannot fully exploit the GEMM units designed for AI. Given that GEMM units can deliver low-precision performance that is several times, or even an order of magnitude, higher than that of double precision, developing tailored algorithms that effectively leverage them could result in substantial performance gains.

Researchers have proposed different mixed-precision algorithms to accelerate quantum chemistry computations with different types of basis. For the plane wave basis, a very recent work by Hou *et al.* proposed using single precision in Hartree-Fock exchange computation with low-rank approximation techniques¹³. For real-space basis sets, Das *et al.* employed single precision for some matrix multiplications in Cholesky QR decomposition and the Rayleigh-Ritz procedure in the DFT-FE package¹⁴. For Gaussian basis, several early works¹⁵⁻¹⁷ and a recent study by Laqua *et al.*¹⁸ investigated computing four-center and three-center electron repulsion integrals (ERIs) on CUDA cores using mixed double and single precision. However, none of these methods leverage specialized GEMM units. Habib *et al.* proposed a mixed-precision matrix factorization of the inverse overlap matrix with Tensor Cores¹⁹, but did not analyze the proportion of matrix decomposition time in one iteration or the speed up of the entire DFT calculation using this method.

In this paper, we present an adaptive precision density fitting (DF) approach for accelerating DFT computation with Gaussian basis on Tensor Cores with 8-bit signed integer (INT8) GEMMs. The contraction of large tensors forms the bulk of the computation in DF²⁰⁻²⁷, making it particularly well suited for acceleration using Tensor Cores. We propose an adaptive precision strategy for DF that converges to the same energies, using the same number or no more than two extra self-consistent field (SCF) iterations. Our method is implemented and validated in the GPU-enabled PySCF package²⁸⁻³¹. It can be readily integrated into other quantum chemistry, offering a promising pathway to accelerate quantum chemical simulations without compromising accuracy. Our contributions are summarized as follows:

- We propose an adaptive precision algorithm utilizing INT8 GEMM on Tensor Core for DF calculation.
- Our implementation in PySCF achieves up to 364% full-DFT speedup on NVIDIA RTX 6000 Ada, confirming the practical benefits of precision-aware Tensor Core usage.
- We demonstrate robust convergence behavior and numerical stability across a wide range of molecules and basis sets of our method.

Background

Density Fitting

Density fitting (DF), also called resolution-of-the-identity (RI), is a method to approximate the ERI tensor. DF uses a set of N_{aux} auxiliary basis functions $\{\hat{\psi}_p\}$ to fit the ERI tensor:

$$(ij|lk) \approx \sum_{p,q} (ij|p) \hat{J}_{pq}^{-1} (q|kl). \quad (1)$$

In practice, we compute and store

$$B_{ij}^q = \sum_p (ij|p) L_{pq}^{-1}, \quad (2)$$

where $LL^T = \hat{J}$ is the Cholesky decomposition of \hat{J} . Using B_{ij}^q and the density matrix D , the Coulomb matrix J is computed in two steps:

$$Y^q = \sum_{ij} B_{ij}^q D_{ij}, \quad (3)$$

$$J_{ij} = \sum_q B_{ij}^q Y^q. \quad (4)$$

The exchange matrix K is also computed in two steps. Let C_{rs} denote the expansion coefficients for the s -th occupied molecular orbital. Then, assuming symmetry of D , we compute

$$W_{is}^q = \sum_r B_{ir}^q C_{rs}, \quad (5)$$

$$K_{ij} = \sum_q \sum_s W_{is}^q W_{js}^q. \quad (6)$$

Schwarz screening is applied in density fitting to reduce the cost of certain computations³². Shell pair MN will be neglected if

$$\sqrt{(MN|MN)\max_P(P|P)} < \tau,$$

where $\max_P(P|P)$ is the maximum value of $(P|P)$ on all auxiliary basis shells. The sparsity pattern of B_{ij}^q is the same as the sparsity pattern of $(ij|p)$ after screening, since the contraction in Formula 2 is in the dimension of the auxiliary basis set.

Single- and Mixed-Precision Algorithms for Quantum Chemistry

The rapid development of general-purpose GPUs (GPGPUs) in the past 15 years has prompted research on how to utilize the computing, especially in single precision, capabilities of GPGPUs to accelerate quantum chemical calculations. In this section, we briefly review the efforts to accelerate quantum chemistry calculations with single- and mixed-precision.

Multiple single precision and mixed-precision algorithms have been proposed for Gaussian basis sets. One of the earliest works by Ufimtsev *et al.*¹⁵ proposed computing ERIs on GPU using mixed double- and single precision. Depending on domain-specific limits, “important” ERIs are calculated in double precision while others are calculated in single precision. The authors claim that this approach gives a nearly $2\times$ speed-up without compromising the accuracy as a double precision calculation. However, these early algorithms for computing ERI on GPUs can only handle smaller basis sets and are therefore less practical. More recently, Laqua *et al.* also proposed computing “significant” three-center-one-electron (3c1e) integrals using double precision and less “significant” 3c1e integrals in single precision, and these 3c1e integrals are used exclusively in the semi-numerical integration method¹⁸ developed by the same authors. Performance improvement and convergence behaviors in this work are comparable to the work of Ufimtsev *et al.*

For real space basis, Das *et al.* contributed a DFT package DFT-FE using finite element basis with mixed-precision acceleration¹⁴. The key part of DFT-FE is a Chebyshev polynomial filtered subspace iteration (CheFSI)³³ for solving electron orbital eigenvalues and eigenvectors. In their implementation, multiple GEMMs within the Cholesky QR decomposition and Rayleigh-Ritz procedures were carefully selected for single precision calculations. However, even for an extremely large system with 6164 Mg atoms, single precision GEMMs take about 35% of the overall running time. In most practical cases, those single precision GEMMs will provide only a very limited improvement in performance. Dawson *et al.*³⁴ demonstrated the use of mixed precision in a density matrix purification solver, including both FP32 and FP16 steps to achieve FP64 results.

For DFT using the plane wave basis, a very recent work by Hou *et al.*¹³ proposed using single precision in the interpolative separable density fitting (ISDF)^{35–37} and the adaptively compressed exchange (ACE)^{38,39} operator. Fast Fourier transforms (FFTs), the computational bottleneck, are also accelerated using single precision. This approach achieves a nearly $2\times$ speedup for multiple test cases.

Multiple works have also investigated using single precision for post-HF calculations. Olivares-Amaya *et al.* explored the use of single- and mixed-precision GEMM for tensor contractions in the second-order Møller-Plesset perturbation theory (MP2)⁴⁰. DePrince *et al.* studied using single precision GEMMs for tensor contractions in coupled cluster theory with single and double excitations (CCSD)⁴¹. They also suggested doing one or more final iterations in double precision after the iterations in single precision to reduce the error of the converged energy to the required level. The work of Pokhilko *et al.* further suggests that using single precision for correlation energy calculation is sufficient for many-body methods in a wide range of practical applications⁴².

The newly released cuEST⁴³ library supports emulated FP64 GEMM in density fitting, which was developed independently from the current effort. A user can use the default emulation level in cuEST or manually specify one emulation level. cuEST also supports the “float float” approximation⁴⁴ for FP64 in the VV10 functional⁴⁵ evaluation, which is computationally intensive but does not involve matrix multiplication.

Floating-point Arithmetic and Low-precision Emulation

A floating-point value consists of three parts: a sign bit, an exponent part, and a mantissa (fraction) part. Then a normalized value x in a floating-point value set is represented as follows:

$$x = (-1)^s \times 1.m_1m_2 \cdots m_{(l_m-1)} \times 2^e, \quad (7)$$

where s is the sign bit, e is the exponent part (a signed integer), and $1.m_1m_2 \cdots m_{(l_m-1)}$ is the mantissa part in binary number, l_m is the number of bits of the mantissa part. The first bit in the mantissa part is not explicitly stored. Different floating point data types have different numbers of bits of both the mantissa part and the exponent part, for example, FP64 has 11 bits exponent and 52 (+1 hidden) bits mantissa, and single precision (FP32) has 8 bits exponent and 23+1 bits mantissa.

Floating-point emulation methods use multiple low-precision floating-point or integer numbers to represent one high-precision floating-point number and use multiple low-precision floating-point or integer arithmetic operations to simulate one high-precision floating-point arithmetic operations. In this work, we focus on emulating FP64 dense general matrix multiplication (GEMM) using low-precision GEMMs.

The Ozaki scheme^{8,46} is one of the earliest floating-point GEMM emulation schemes that can achieve any accuracy. Mukunoki *et al.* proposed a method to compute the Ozaki scheme on Tensor Cores using half-precision (FP16, 5 bits exponent and 10+1 bits mantissa) GEMMs⁹. This FP16 emulation approach represents one FP64 value x^{fp64} as the scaled sum of multiple FP16 values (“splits”) \hat{x}_i^{fp16} :

$$x^{\text{fp64}} = \sum_{i=1}^s \left(\hat{x}_i^{\text{fp16}} \times 2^{-e_i} \right), \quad (8)$$

where e_i are positive integers, and the FP64 GEMM is emulated by $s(s+1)/2$ FP16 GEMMs. One major disadvantage of this approach is the mantissa waste bits in FP16. The work of Ootomo *et al.*¹¹ pointed out that each FP16 split can only capture $\alpha < 11$ bits of mantissa from the original FP64 input. The “bits per one split” (BPS) value α is computed as

$$\alpha = \min(11, \lfloor (24 - \log_2 k) / 2 \rfloor), \quad (9)$$

where k is the size of the inner product of matrix multiplication. In other words, for a larger k value, the FP16 emulation approach needs to increase s to maintain accuracy, which leads to a rapid decrease in performance. It is possible to split the original GEMM in the inner product dimension to reduce the k value in each GEMM and avoid wasting too many mantissa bits, but a smaller problem size also harms the GEMM performance.

In the past two years, new research has focused on emulating FP64 GEMM using 8-bit signed integer (INT8) GEMMs. Ootomo *et al.* first employed the Ozaki scheme with INT8 GEMMs¹¹ and their implementation ozIMMU is released on GitHub⁴⁷. ozIMMU handles the exponent part and the mantissa part separately and directly emulate the addition and multiplication between mantissa parts using INT8. Ozaki *et al.* further proposed a new method¹² called Ozaki scheme II, which scales and converts FP64 inputs to INT8 using Chinese Remainder Theorem, then performs INT8 GEMMs. Their implementation GEMMul8 is also released on GitHub⁴⁸. These INT8 emulation algorithms significantly outperform the FP16 emulation algorithm.

Accelerate DF with Adaptive Precision and INT8 GEMMs on Tensor Cores

We implemented an adaptive precision scheme that utilizes INT8-emulated GEMMs to accelerate DF calculations. In summary, the algorithm is:

- Compute the Coulomb matrix J using Formulas (3) and (4) in FP64;
- Compute the Hartree-Fock exchange matrix K using Formulas (5) and (6) and INT8-emulated FP64 GEMM with adaptive precision determined by the relative change magnitude of SCF energy;
- Fall back to direct FP64 GEMM when SCF is almost converge and INT8-emulated FP64 GEMM is no longer faster than direct FP64 GEMM since the cost of emulated FP64 GEMM increases with the emulation accuracy.

The details of the algorithm will soon be discussed in Section .

Justification of Accelerating only DF K Matrix Build

We choose DF instead of the direct method because the computations in DF are better suitable for Tensor Cores. Although tensor contractions are also one of the performance bottlenecks in the direct methods, the sizes of contracted tensors in the direct method are very small. Let M, N, P, Q denote shell indices. Then blocks of the Coulomb and exchange matrices are

$$J_{MN} = \sum_{PQ} (MN|PQ)D_{PQ} \quad \text{and} \quad K_{MN} = \sum_{PQ} (MP|NQ)D_{PQ}, \quad (10)$$

respectively, where $(MN|PQ)$ denotes a shell quartet of ERI. The sizes of each $(MN|PQ)$ tensor depend on the selected basis and the shell indices M, N, P , and Q ; in most cases the size of $(MN|PQ)$ smaller than 21^4 . These tensor contractions are computed as matrix-vector multiplication after reordering and fusion of tensor indices, while the matrix is smaller than 441×441 in most cases. Such matrix-vector multiplications with small matrices are not suitable for Tensor Cores. Furthermore, exploiting the 8-way symmetry property leads to irregular memory access patterns to $(MN|PQ)$ tensor and J, K, D matrices. Some earlier work by Huang *et al.*^{49,50} has discussed these challenges and proposed several strategies to mitigate them on CPUs, but these techniques cannot be applied to Tensor Cores. In contrast, density fitting only involves four large tensor contractions, which are much more suitable for GPUs and Tensor Cores.

The second question is why only the exchange matrix K is computed using the adaptive precision approach. The first and most important reason is that building the K matrix requires significantly more floating point operations (FLOPs) and has a much longer computation time compared to building the J matrix. The second reason is that the constructions of the J and V matrices cannot be expressed as GEMMs. We now proceed to estimate the FLOP counts for constructing both the J and K matrices.

After applying Schwarz screening to the $(ij|p)$ pairs of the 3c2e integrals $(ij|p)$, only a subset of the $(ij|p)$ pairs survives. Let N_{scr} denote the number of pairs $(ij|p)$ that remain after screening. $N_{scr} \leq N_{bf}^2$ always holds, and normally $N_{scr} = \Theta(N_{bf})$ for large N_{bf} . For each pair of indexes (q, s) , the computation in Formula (5) is to multiply a sparse matrix B_{ir} with a dense vector C_r . We can then compute the number of FLOPs required in each step:

- Formulas (3) and (4): $2N_{scr}N_{aux}$,
- Formula (5): $2N_{scr}N_{aux}N_{occ}$,
- Formula (6): $2N_{bf}^2N_{aux}N_{occ}$.

Thus, the J matrix build requires a total of $4N_{scr}N_{aux}$ FLOPs, while the K matrix build requires $2(N_{scr} + N_{bf}^2)N_{aux}N_{occ}$ FLOPs, which is at least N_{occ} times more than the FLOPs needed for the J matrix build. This significant difference in computational cost motivates us to keep the J matrix computation in FP64, while accelerating the K matrix computation using the adaptive precision scheme.

Adaptive Precision DF K Matrix Build with INT8 GEMM

We use the newly released emulated FP64 GEMM in cuBLAS 13.0 update 2 for computing the K matrix. The emulated FP64 GEMM implementation in cuBLAS uses a fixed point emulation and follows the Ozaki scheme, and it allows the user to set the number of mantissa bits m^{emu} to be used for fixed emulation⁵¹. The implementation uses $s^{\text{emu}} = \lceil (m^{\text{emu}} + 1)/8 \rceil$ splits of INT8 to represent the input. For details about this algorithm, we recommend that readers refer to a recent work of Schwarz *et al.*⁵².

Consequently, we use an *adaptive precision selection* scheme in the SCF iterations. Since the SCF procedure is iterative, different initial guesses of the solution can often converge to the same final result with the same or slightly different numbers of iterations. If the early iterations are computed with reduced precision, it can be considered as using a slightly different initial guess, and the overall convergence may remain unchanged. Let E_i^{total} denote the total energy after the i -th SCF iteration, and let

$$\Delta E_i = E_i^{\text{total}} - E_{i-1}^{\text{total}}, \quad \Delta E_i^{\text{rel}} = |\Delta E_i / E_i^{\text{total}}|. \quad (11)$$

$|\Delta E_i|$ is usually used as a criterion to determine if the SCF iteration has converged. We base our precision selection on ΔE_i^{rel} , the relative change in total energy between two adjacent SCF iterations, since relative

errors are more meaningful in numerical analysis and methods. We select m_{i+1}^{emu} for the $(i+1)$ -th SCF iteration based on Table 1 and hardware specifications. For example, we return to the standard FP64 calculation for $\Delta E_i^{\text{rel}} < 2 \times 10^{-6}$ (emulation level ≤ 3) on H100 GPU since the H100 Tensor Cores support FP64 and the standard FP64 calculation is faster than the INT8 emulation using 10 or more INT8 GEMMs.

ΔE_i^{rel}	Emu. Level	m_{i+1}^{emu}	Emu. Relative Accuracy	# INT8 Splits
$\geq 2 \times 10^{-6}$	4	23	1.2×10^{-7}	3
$\in [5 \times 10^{-9}, 2 \times 10^{-6})$	3	31	4.7×10^{-10}	4
$\in [2 \times 10^{-11}, 5 \times 10^{-9})$	2	39	1.8×10^{-12}	5
$\in [7 \times 10^{-14}, 2 \times 10^{-11})$	1	47	7.1×10^{-15}	6
$< 7 \times 10^{-14}$	0	(FP64)	(2.2×10^{-16})	(0)

Table 1: Adaptive precision selection scheme based on SCF total energy convergence

The algorithm starts with low precision ($m_1^{\text{emu}} = 23$) and gradually transitions to higher precision as the total energy approaches convergence. As SCF iterations converge, ΔE_{i+1} could be 10 times smaller than ΔE_i if the computation is performed in FP64. Meanwhile, m_{i+1}^{emu} used in the $(i+1)$ -th SCF iteration is determined by ΔE_i^{rel} instead of $\Delta E_{i+1}^{\text{rel}}$. Therefore, we set ΔE_i^{rel} thresholds to promote emulation accuracy approximately 10 times higher than the current relative accuracy of INT8 emulation. In addition, $|\Delta E_i|$ in standard FP64 SCF calculations does not always decrease monotonically. To stabilize the computation, we disallow reverting to a lower precision once a higher precision has been used.

Implementation Details

We implement our adaptive precision DF algorithm in the PySCF package²⁸, which leverages the CuPy library⁵³ for GPU acceleration.

The original PySCF GPU implementation computes and stores the 3D tensor B_{ij}^q in a symmetric sparse form (only non-zero B_{ij}^q values with $i \geq j$ are stored) to reduce memory usage. Then, it partitions the computations along the q dimension. Using MATLAB colon notation, the computation of Formulas (5) and (6) in PySCF can be rewritten as:

$$W_{is}^{q1:q2} = \sum_r B_{ir}^{q1:q2} C_{rs}, \quad (12)$$

$$K_{ij} = K_{ij} + \sum_{q=q1}^{q2} \sum_s W_{is}^q W_{js}^q. \quad (13)$$

Each tensor block $B_{ir}^{q1:q2}$ is first symmetrized and restored to a dense tensor, then $W_{is}^{q1:q2}$ is calculated using the dense tensor. In addition, $\sum_{q=q1}^{q2} \sum_s W_{is}^q W_{js}^q$ is computed using one GEMM instead of one symmetric rank-k update (SYRK). We follow these practices in our implementation for three reasons. First, leveraging sparsity requires repacking the corresponding rows or columns of $C(r, s)$ to match the sparsity pattern of B_{ij}^q , which introduces nontrivial data movement and additional memory overhead. Second, the original tensor contraction needs to be computed as many GEMMs with different sizes of matrices, but these sizes are often not large enough to fully saturate the GPU in a single call. One potential optimization is to use grouped GEMM kernels from the CUTLASS library⁵⁴, which support batched GEMMs with varying input dimensions. However, this feature is currently only available on NVIDIA H100 and newer Blackwell GPUs. Considering these factors, we opt not to utilize the sparsity in B_{ij}^q . Lastly, the INT8 emulation code is also not optimized for SYRK yet. Therefore, we follow the practices in the PySCF GPU implementation. Future implementations in C/C++-based quantum chemistry libraries may explore sparsity-aware and symmetry-aware optimizations more effectively.

To enable DFT calculations with density fitting on GPUs for larger molecular systems, we adjust several GPU memory usage thresholds in PySCF. In particular, we allow up to 80% of the available GPU memory to be used to store the sparse B_{ij}^q . We pad the dimensions of $B_{ir}^{q1:q2}$ and C_{rs} to multipliers of 32, allowing cuBLAS to select and use GEMM kernels optimized for Tensor Cores.

Numerical Experiments

Experiment Setup

We evaluate our mixed-precision density fitting algorithm across multiple basis sets and molecular systems on a range of NVIDIA GPUs. Specifically, we use three types of GPUs in our experiments: a consumer-grade RTX 4090, a workstation-class RTX A6000, and a server-grade flagship H100. Table 2 summarizes their key specifications. These GPUs are installed on different host systems; however, all systems share the same operating system and software environment, detailed as follows:

- Ubuntu 24.04 LTS with kernel 6.8.0;
- NVIDIA GPU driver 580.105.8 with CUDA toolkit 13.1;
- Python 3.12 + PySCF 2.8.0 + cupy-cuda13x 13.6.0 + cutensor-cu13 2.3.1.

We employ two high-accuracy basis sets in our tests: the split-valence Pople basis set 6-311G(d,p)⁵⁵ and the Karlsruhe basis set def2-TZVPP^{56,57}. The corresponding auxiliary basis sets are cc-pVTZ-JKFIT and def2-TZVPP-JKFIT, respectively. For all DFT calculations, we use the widely adopted hybrid exchange-correlation functional B3LYP^{58,59}.

The test set includes diverse molecular systems representing different dimensionalities and complexities: 1D alkane chains, 2D graphene-like molecules, 3D water clusters, and 3D organic molecules. Table 3 provides a comprehensive summary of the tested molecular systems, their associated basis sets, and reference DFT results computed using the standard FP64 implementation. All reference results are obtained with DF on a single NVIDIA H100 GPU with 80 GB of HBM3 memory.

	RTX 4090	RTX 6000 Ada	H100
Architecture	Ada Lovelance		Hopper
Compute Capability	8.9		9.0
Mem. Size (GB)	24	48	80
Mem. Bandwidth	1008 GB/s	960 GB/s	3.35 TB/s
# CUDA cores	16384	18176	16896
# Tensor Cores	512	568	528
Peak FP64 (TFLOPS)	1.3	1.4	33.5/67*
Peak FP32 (TFLOPS)	82.6	91.1	67
Peak TC INT8 (TOPS)	660.6*	728.5*	1979*

Table 2: Key specifications of the GPUs used in numerical experiments. The asterisk indicates that the Tensor Cores of this GPU support GEMM of this data type, and the corresponding values are the performance of Tensor Cores.

Single SCF Iteration Performance Improvement

We first study the performance improvement of a single SCF iteration. Figure 1 shows the running time breakdowns (in milliseconds) of the second SCF iteration on different GPUs. For each GPU, we choose the largest molecular system it can run in the test set. The time for constructing the J matrix is not listed separately since it is too small (smaller than 50 ms in all three test cases). For all molecular systems and GPUs tested, the construction of the K matrix is the most time-consuming part. On RTX 4090 and RTX 6000 Ada, switching from FP64 to INT8 emulation significantly accelerates the construction of the K matrix, since these GPUs have weak FP64 vector performance and no FP64 support on their Tensor Cores. In addition, the performance of different levels of INT8 emulation is different: level 4 (3xINT8 splits) only needs about half the time of level 1 (6xINT8 splits). On H100, only level 4 INT8 emulation is slightly faster than the standard FP64 implementation, since H100 Tensor Cores support FP64. These results also show the necessity of the proposed adaptive precision algorithm instead of simply using INT8 emulated FP64 GEMM as a drop-in replacement of the native FP64 GEMM.

Molecule	Atoms	N_{bf}	N_{aux}	N_{occ}	# SCF Iters.	Converged Energy
Alkane-62	C20H42	612	2840	81	9	-787.67278936
		1208	2256	81	9	-787.76535664
Alkane-122	C40H82	1212	5620	161	9	-1574.13589317
		2388	4476	161	9	-1574.32029172
Alkane-182	C60H122	1812	8400	241	9	-2360.59899665
		3568	6696	241	9	-2360.87522643
Graphene-36	C24H12	504	2256	78	9	-922.07474766
		912	2016	78	9	-922.19791221
Graphene-72	C54H18	1080	4806	171	10	-2069.31091985
		1926	4374	171	10	-2069.58764198
(H2O)20	(H2O)20	600	2780	100	12	-1529.16143540
		1180	2260	100	12	-1529.43152890
(H2O)47	(H2O)47	1410	6533	235	10	-3593.63712813
		2773	5311	235	11	-3594.23830243
Tamoxifen	C26H29N1O1	678	3082	100	11	-1138.22247481
		1274	2626	100	11	-1138.37146192
Sphingomyelin	C25H51N2O5P1	908	4179	135	13	-1810.65324984
		1748	3460	135	13	-1810.90326016

Table 3: Molecular systems used for numerical experiments and reference converged energy. For each molecule, the first and the second rows correspond to 6-311G(d,p) and def2-TZVPP basis sets, respectively. N_{bf} and N_{aux} are the numbers of basis functions of the original and DF auxiliary basis sets, respectively. N_{occ} is the number of occupied orbitals. The reference converged energy are computed using standard FP64 calculations with DF.

DFT Calculation Convergence

In this section, we evaluate the convergence of the proposed adaptive precision algorithm.

Table 4 reports the number of SCF iterations required to reach $|\Delta E| \leq 10^{-7}$ on different GPUs. We note that different quantum chemistry packages use different default thresholds: 10^{-9} in PySCF, 10^{-7} in NWChem⁶⁰, and 10^{-6} in both Psi4⁶¹ and ORCA⁶². We checked all converged energies of the adaptive precision algorithm; the absolute errors between these energies and the reference converged energies computed by the standard FP64 algorithm are less than 10^{-7} . For RTX 4090 and RTX 6000 Ada, the adaptive precision algorithm uses the same number of SCF iterations as the standard FP64 algorithm to converge; a few systems require one extra SCF iteration. For H100, the adaptive precision algorithm has the same convergence behavior as the standard FP64 algorithm. Since the adaptive precision algorithm returns to FP64 on H100 when $\Delta E_i^{\text{rel}} < 2 \times 10^{-6}$ instead of $\Delta E_i^{\text{rel}} < 7 \times 10^{-14}$, the convergence of H100 is slightly better.

Figure 2 demonstrates both the effectiveness of the adaptive mixed-precision scheme and the rationale behind its switching thresholds. In the left panel, we plot the SCF energy error relative to the converged reference value as a function of SCF iteration. All methods show similar early-stage error reduction, but the fixed level-4 run reaches a clear error floor at approximately 10^{-5} Ha, after which additional iterations provide little improvement. By contrast, the adaptive schedule continues to reduce the error to the 10^{-9} Ha range, close to the high-accuracy reference result. The right panel shows that the adaptive precision level decreases from 4 to 2 as the relative energy change ΔE_i^{rel} decreases. This threshold-driven schedule avoids the low-precision error floor while retaining reduced precision in the early SCF regime, where it is most efficient.

To summary, if a molecular system has good convergence using the standard FP64 DF algorithm, our proposed adaptive precision algorithm should also have good convergence.




Figure 1: Running time breakdowns (in milliseconds) of the second SCF iteration on different GPUs and molecular systems: (a) Alkane-122/6-311G(d,p) on a RTX 4090 (left); (b) Alkane-182/6-311G(d,p) on a RTX 6000 Ada (middle); and (c) Alkane-182/def2-TZVPP on a H100 (right).

	6-311G(d,p)		def2-TZVPP	
	Standard FP64	Adaptive Precision	Standard FP64	Adaptive Precision
Alkane-62	9	<i>10 / 10 / 9</i>	9	<i>10 / 10 / 9</i>
Alkane-122	9	<i>10 / 10 / 9</i>	9	- / <i>10 / 9</i>
Alkane-182	9	- / <i>10 / 9</i>	9	9 / - / -
Graphene-36	9	9 / 9 / 9	9	9 / 9 / 9
Graphene-72	10	10 / <i>11 / 10</i>	10	10 / - / -
(H2O)20	12	12 / 12 / 12	12	12 / 12 / 12
(H2O)47	10	- / <i>11 / 10</i>	11	11 / - / -
Tamoxifen	11	11 / 11 / 11	11	11 / 11 / 11
Sphingomyelin	13	13 / 13 / 13	13	13 / 13 / -

Table 4: Number of SCF iterations needed for convergence on different GPUs. Values in *italic* indicate a higher iteration count than that of the standard FP64 algorithm. For adaptive precision, the reported values are the numbers on RTX 4090 / RTX 6000 Ada / H100, respectively. A short dash indicates the GPU does not have enough memory to run this test case.

DFT Calculation Performance Improvement

Finally, we analyze the overall performance improvement of a DFT calculation using adaptive precision algorithms. All computational steps are included, including the integrals $(ij|p)$ and the construction of the B_{ij}^q tensor in Formula (2). Figure 3 presents the speedup of using the adaptive precision algorithm over the standard FP64 algorithm.

On all GPUs, the adaptive precision algorithm has a larger speedup over the standard FP64 algorithm on a larger molecular system, since the time required to construct the K matrix accounts for a larger proportion of the total running time of a larger molecular system. On RTX 4090 and RTX 6000 Ada, the adaptive precision algorithm has a speedup of up to 3.04 and 4.64, respectively. For molecular systems that the adaptive precision algorithm needs one more SCF iteration to converge, for example, Alkane-122/6-311G(d,p), the adaptive precision algorithm still has a very large speedup. These two GPUs have the same architecture and similar numbers of Tensor Cores, while RTX 6000 Ada has twice the memory size of RTX 4090. Therefore, RTX 6000 Ada can run larger systems and have a larger acceleration on larger systems. On H100, even if its Tensor Cores are designed to accelerate both traditional scientific workloads and modern AI tasks with FP64 and low-precision data types, the adaptive precision algorithm can still reach a speedup of

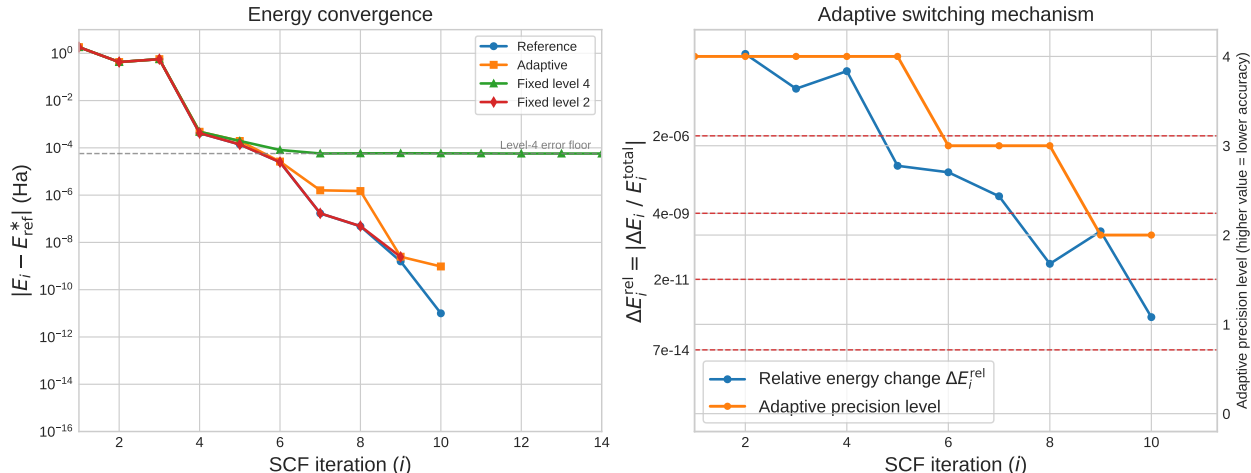


Figure 2: SCF energy convergence with different mixed-precision schemes on an RTX 6000 Ada GPU for molecular system Alkane-122 / 6-311G(d,p) (left) and the adaptive precision switching mechanism (right).

1.37. Even if the overhead of INT8 emulation cancels out the performance improvement of adaptive precision algorithm on small molecular systems, only 3% overall performance loss at most on a H100. To summary, using the adaptive precision algorithm is beneficial in most cases.

We also compared the performance between density fitting and the direct method on the same GPUs. Figure 4 and Figure 5 present the speedup of using the standard FP64 density fitting method over the FP64 direct method and the speedup of using the adaptive precision density fitting method over the FP64 direct method, respectively. The results show that density fitting method with and without adaptive precision has a large advantage over the direct method for most test cases. For Alkane-122 and Alkane-182, direct method is faster than density fitting method as the direct method can better leverage the sparsity created by screening on long alkane molecules.

Implementing the Adaptive Precision Algorithm with Other Libraries

We also implemented our adaptive precision algorithm with the ozIMMU library⁴⁷ and the GEMMu8 library⁴⁸ and tested their performance. One common and significant issue of our ozIMMU-based and GEMMu8-based implementations is the inaccurate estimation of the emulation level to achieve desired accuracy levels. Unlike cuBLAS, ozIMMU and GEMMu8 do not support emulating FP64 GEMM to a specific accuracy level. We have to estimate the number of INT8 slices to use in ozIMMU and GEMMu8 based on the results in their paper and adjust Table 1, and we might slightly over-estimate the number of slices in some cases. As a result, our ozIMMU-based and GEMMu8-based implementations can still achieve a good speedup on some large systems (for example, 6-311G(d,p)+Alkane-182) compared to the standard FP64 GEMM implementation, but no speedup on moderate or small size systems (detailed results omitted). These results also show the advantage of the new emulated FP64 GEMM APIs in cuBLAS.

Conclusion and Future Work

In this work, we propose an adaptive precision algorithm to accelerate the construction of exchange matrix in the widely used density fitting method. Our algorithm exploits the newly proposed INT8 emulation of FP64 GEMM and the property of iterative methods to fully utilize computational capacities of AI accelerators. We implemented our algorithm in the GPU version of PySCF for evaluation. Numerical experiment results show that our algorithm can deliver up to 204%, 365%, and 37% speedup on a gaming GPU, a workstation GPU, and a flagship GPU while maintaining a high accuracy of DFT calculation.

Experiment results also point out some issues that need further research. The first topic is numerical stability and error consideration. Although the numerical stability of INT8 emulation has been studied in




Figure 3: Speedup of DFT calculations when using the adaptive precision method over using the standard FP64 algorithm. All computational steps, including the integrals $(ij|p)$ and the construction of the B_{ij}^q tensor in Formula (2), are included. Missing data points indicate that the GPU did not have sufficient memory to complete the corresponding calculation.




Figure 4: Speedup of DFT calculations when using FP64 density fitting over using the FP64 direct method. Missing data points indicate that the GPU did not have sufficient memory to complete the corresponding DF calculation.

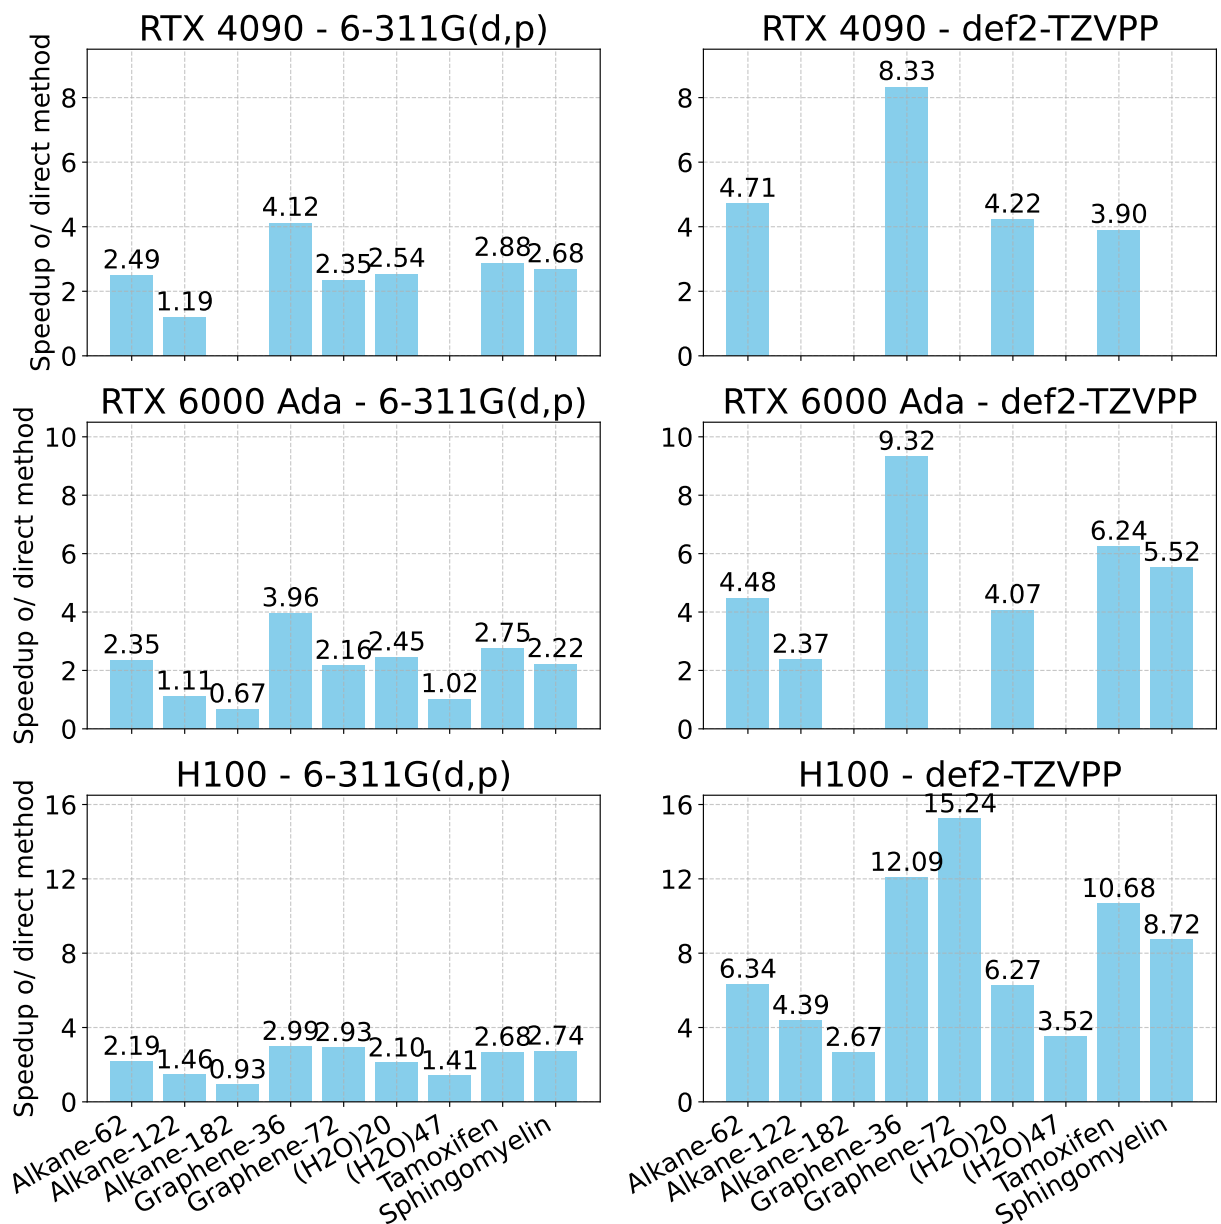


Figure 5: Speedup of DFT calculations when using density fitting with adaptive precision method over using the FP64 direct method. Missing data points indicate that the GPU did not have sufficient memory to complete the corresponding DF calculation.

earlier works, our strategies of switching between emulation accuracies are empirical and have no formal error bounds or convergence guarantees. We admit that it is very hard to perform formal numerical error analysis for complicated non-linear iterative algorithms. In the future, we will try to further refine our algorithm by linking precision and emulation level to residual norms or density matrix changes.

The second topic is the use of single precision (FP32) to construct the exchange-correlation (XC) matrix. Figure fig. 1 shows that once the K matrix construction is significantly accelerated, building the XC matrix could be a new performance bottleneck. PySCF and many other quantum chemistry packages use the LibXC library⁶³ to evaluate the values of different XC functionals on a set of points for constructing the V matrix. LibXC supports FP32, but it needs to be compiled separately. This means that the quantum chemistry library will need to handle linking and invoking two LibXC libraries. In addition, the selection of XC functional evaluation points also has a large impact on the accuracy of DFT. When using FP32, the number of evaluation points should be reduced compared to using FP64 to achieve a balance between accuracy and computational cost.

Acknowledgements

The authors thank Minseok Lee for the discussion of utilizing Tensor Cores in scientific computing and Hiroyuki Ootomo for the meaningful discussion of FP64 GEMM emulation using low-precision GEMM operations.

References

- (1) NVIDIA NVIDIA Tesla V100 GPU Architecture, <https://images.nvidia.com/content/volta-architecture/pdf/volta-architecture-whitepaper.pdf>, 2017.
- (2) NVIDIA NVIDIA A100 Tensor Core GPU Architecture, <https://images.nvidia.com/aem-dam/en-zz/Solutions/data-center/nvidia-ampere-architecture-whitepaper.pdf>, 2020.
- (3) NVIDIA NVIDIA H100 Tensor Core GPU Architecture, <https://resources.nvidia.com/en-us-tensor-core>, 2022.
- (4) Jouppi, N. P.; Young, C.; Patil, N.; Patterson, D.; Agrawal, G.; Bajwa, R.; Bates, S.; Bhatia, S.; Boden, N.; Borchers, A.; Boyle, R.; Cantin, P.-l.; Chao, C.; Clark, C.; Coriell, J.; Daley, M.; Dau, M.; Dean, J.; Gelb, B.; Ghaemmaghami, T. V.; Gottipati, R.; Gulland, W.; Hagmann, R.; Ho, C. R.; Hogberg, D.; Hu, J.; Hundt, R.; Hurt, D.; Ibarz, J.; Jaffey, A.; Jaworski, A.; Kaplan, A.; Khaitan, H.; Killebrew, D.; Koch, A.; Kumar, N.; Lacy, S.; Laudon, J.; Law, J.; Le, D.; Leary, C.; Liu, Z.; Lucke, K.; Lundin, A.; MacKean, G.; Maggiore, A.; Mahony, M.; Miller, K.; Nagarajan, R.; Narayanaswami, R.; Ni, R.; Nix, K.; Norrie, T.; Omernick, M.; Penukonda, N.; Phelps, A.; Ross, J.; Ross, M.; Salek, A.; Samadiani, E.; Severn, C.; Sizikov, G.; Snelham, M.; Souter, J.; Steinberg, D.; Swing, A.; Tan, M.; Thorson, G.; Tian, B.; Toma, H.; Tuttle, E.; Vasudevan, V.; Walter, R.; Wang, W.; Wilcox, E.; Yoon, D. H. In *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*, ACM: New York, NY, USA, 2017, pp 1–12.
- (5) Abdelfattah, A.; Anzt, H.; Ayala, A.; Boman, E.; Carson, E.; Cayrols, S.; Cojean, T.; Dongarra, J.; Falgout, R.; Gates, M., et al. *Advances in Mixed Precision Algorithms: 2021 Edition*; tech. rep.; Lawrence Livermore National Lab. (LLNL), Livermore, CA (United States), 2021.
- (6) Higham, N. J.; Mary, T. Mixed precision algorithms in numerical linear algebra. *Acta Numerica* **2022**, *31*, 347–414.
- (7) Kashi, A.; Lu, H.; Brewer, W.; Rogers, D.; Matheson, M.; Shankar, M.; Wang, F. Mixed-precision numerics in scientific applications: survey and perspectives, 2025.
- (8) Ozaki, K.; Ogita, T.; Oishi, S.; Rump, S. M. Error-free transformations of matrix multiplication by using fast routines of matrix multiplication and its applications. *Numer. Algorithms* **2012**, *59*, 95–118.
- (9) Mukunoki, D.; Ozaki, K.; Ogita, T.; Imamura, T. In *High Performance Computing*, ed. by Sadayappan, P.; Chamberlain, B. L.; Juckeland, G.; Ltaief, H., Springer International Publishing: 2020, pp 230–248.

- (10) Ootomo, H.; Yokota, R. Recovering single precision accuracy from Tensor Cores while surpassing the FP32 theoretical peak performance. *The International Journal of High Performance Computing Applications* **2022**, *36*, 475–491.
- (11) Ootomo, H.; Ozaki, K.; Yokota, R. DGEMM on integer matrix multiplication unit. *The International Journal of High Performance Computing Applications* **2024**, *38*, 297–313.
- (12) Ozaki, K.; Uchino, Y.; Imamura, T. Ozaki Scheme II: A GEMM-oriented emulation of floating-point matrix multiplication using an integer modular technique, 2025.
- (13) Hou, B.; Chen, S.; Qin, X.; Hu, W.; Yang, J. Dual-Grid and Mixed-Precision Methods for Accelerating Plane-Wave Hybrid Functional Electronic Structure Calculations. *Journal of Chemical Theory and Computation* **2025**, *21*, 787–802.
- (14) Das, S.; Motamarri, P.; Gavini, V.; Turcksin, B.; Li, Y. W.; Leback, B. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ACM: Denver, Colorado, 2019, pp 1–11.
- (15) Ufimtsev, I. S.; Martínez, T. J. Quantum Chemistry on Graphical Processing Units. 1. Strategies for Two-Electron Integral Evaluation. *Journal of Chemical Theory and Computation* **2008**, *4*, 222–231.
- (16) Asadchev, A.; Allada, V.; Felder, J.; Bode, B. M.; Gordon, M. S.; Windus, T. L. Uncontracted Rys Quadrature Implementation of up to G Functions on Graphical Processing Units. *Journal of Chemical Theory and Computation* **2010**, *6*, 696–704.
- (17) Miao, Y.; Merz, K. M. J. Acceleration of Electron Repulsion Integral Evaluation on Graphics Processing Units via Use of Recurrence Relations. *Journal of Chemical Theory and Computation* **2013**, *9*, 965–976.
- (18) Laqua, H.; Kussmann, J.; Ochsenfeld, C. Accelerating seminumerical Fock-exchange calculations using mixed single- and double-precision arithmetic. *The Journal of Chemical Physics* **2021**, *154*, 214116.
- (19) Habib, A.; Finkelstein, J.; Niklasson, A. M. N. Efficient Mixed-Precision Matrix Factorization of the Inverse Overlap Matrix in Electronic Structure Calculations with AI-Hardware and GPUs. *Journal of Chemical Theory and Computation* **2024**, *20*, 7102–7112.
- (20) Whitten, J. L. Coulombic potential energy integrals and approximations. *The Journal of Chemical Physics* **1973**, *58*, 4496–4501.
- (21) Eichkorn, K.; Treutler, O.; Öhm, H.; Häser, M.; Ahlrichs, R. Auxiliary basis sets to approximate Coulomb potentials. *Chemical Physics Letters* **1995**, *240*, 283–290.
- (22) Weigend, F. A fully direct RI-HF algorithm: Implementation, optimised auxiliary basis sets, demonstration of accuracy and efficiency. *Physical Chemistry Chemical Physics* **2002**, *4*, 4285–4291.
- (23) Sodt, A.; Subotnik, J. E.; Head-Gordon, M. Linear scaling density fitting. *The Journal of Chemical Physics* **2006**, *125*, 194109.
- (24) Aquilante, F.; Lindh, R.; Bondo Pedersen, T. Unbiased auxiliary basis sets for accurate two-electron integral approximations. *The Journal of Chemical Physics* **2007**, *127*, 114107.
- (25) Aquilante, F.; Gagliardi, L.; Pedersen, T. B.; Lindh, R. Atomic Cholesky decompositions: A route to unbiased auxiliary basis sets for density fitting approximation with tunable accuracy and efficiency. *The Journal of Chemical Physics* **2009**, *130*, 154107.
- (26) Reine, S.; Tellgren, E.; Krapp, A.; Kjærgaard, T.; Helgaker, T.; Jansik, B.; Høst, S.; Salek, P. Variational and robust density fitting of four-center two-electron integrals in local metrics. *The Journal of Chemical Physics* **2008**, *129*, 104101.
- (27) Merlot, P.; Kjærgaard, T.; Helgaker, T.; Lindh, R.; Aquilante, F.; Reine, S.; Pedersen, T. B. Attractive electron–electron interactions within robust local fitting approximations. *Journal of Computational Chemistry* **2013**, *34*, 1486–1496.
- (28) Sun, Q.; Berkelbach, T. C.; Blunt, N. S.; Booth, G. H.; Guo, S.; Li, Z.; Liu, J.; McClain, J. D.; Sayfutyarova, E. R.; Sharma, S.; Wouters, S.; Chan, G. K.-L. PySCF: the Python-based simulations of chemistry framework. *WIREs Computational Molecular Science* **2018**, *8*, e1340.

- (29) Sun, Q.; Zhang, X.; Banerjee, S.; Bao, P.; Barbry, M.; Blunt, N. S.; Bogdanov, N. A.; Booth, G. H.; Chen, J.; Cui, Z.-H.; Eriksen, J. J.; Gao, Y.; Guo, S.; Hermann, J.; Hermes, M. R.; Koh, K.; Koval, P.; Lehtola, S.; Li, Z.; Liu, J.; Mardirossian, N.; McClain, J. D.; Motta, M.; Mussard, B.; Pham, H. Q.; Pulkin, A.; Purwanto, W.; Robinson, P. J.; Ronca, E.; Sayfutyarova, E. R.; Scheurer, M.; Schurkus, H. F.; Smith, J. E. T.; Sun, C.; Sun, S.-N.; Upadhyay, S.; Wagner, L. K.; Wang, X.; White, A.; Whitfield, J. D.; Williamson, M. J.; Wouters, S.; Yang, J.; Yu, J. M.; Zhu, T.; Berkelbach, T. C.; Sharma, S.; Sokolov, A. Y.; Chan, G. K.-L. Recent developments in the PySCF program package. *The Journal of Chemical Physics* **2020**, *153*, 024109.
- (30) Li, R.; Sun, Q.; Zhang, X.; Chan, G. K.-L. Introducing GPU-acceleration into the Python-based Simulations of Chemistry Framework, 2024.
- (31) Wu, X.; Sun, Q.; Pu, Z.; Zheng, T.; Ma, W.; Yan, W.; Yu, X.; Wu, Z.; Huo, M.; Li, X.; Ren, W.; Gong, S.; Zhang, Y.; Gao, W. Enhancing GPU-acceleration in the Python-based Simulations of Chemistry Framework, 2024.
- (32) Werner, H.-J.; Manby, F. R.; Knowles, P. J. Fast linear scaling second-order Møller-Plesset perturbation theory (MP2) using local and density fitting approximations. *The Journal of Chemical Physics* **2003**, *118*, 8149–8160.
- (33) Zhou, Y.; Saad, Y.; Tiago, M. L.; Chelikowsky, J. R. Self-consistent-field calculations using Chebyshev-filtered subspace iteration. *Journal of Computational Physics* **2006**, *219*, 172–184.
- (34) Dawson, W.; Ozaki, K.; Domke, J.; Nakajima, T. Reducing Numerical Precision Requirements in Quantum Chemistry Calculations. *Journal of Chemical Theory and Computation* **2024**, *20*, 10826–10837.
- (35) Lu, J.; Ying, L. Compression of the electron repulsion integral tensor in tensor hypercontraction format with cubic scaling cost. *Journal of Computational Physics* **2015**, *302*, 329–335.
- (36) Hu, W.; Lin, L.; Yang, C. Interpolative Separable Density Fitting Decomposition for Accelerating Hybrid Density Functional Calculations with Applications to Defects in Silicon. *Journal of Chemical Theory and Computation* **2017**, *13*, 5420–5431.
- (37) Dong, K.; Hu, W.; Lin, L. Interpolative Separable Density Fitting through Centroidal Voronoi Tessellation with Applications to Hybrid Functional Electronic Structure Calculations. *Journal of Chemical Theory and Computation* **2018**, *14*, 1311–1320.
- (38) Lin, L. Adaptively Compressed Exchange Operator. *Journal of Chemical Theory and Computation* **2016**, *12*, 2242–2249.
- (39) Hu, W.; Lin, L.; Banerjee, A. S.; Vecharynski, E.; Yang, C. Adaptively Compressed Exchange Operator for Large-Scale Hybrid Density Functional Calculations with Applications to the Adsorption of Water on Silicene. *Journal of Chemical Theory and Computation* **2017**, *13*, 1188–1198.
- (40) Olivares-Amaya, R.; Watson, M. A.; Edgar, R. G.; Vogt, L.; Shao, Y.; Aspuru-Guzik, A. Accelerating Correlated Quantum Chemistry Calculations Using Graphical Processing Units and a Mixed Precision Matrix Multiplication Library. *Journal of Chemical Theory and Computation* **2010**, *6*, 135–144.
- (41) DePrince, A. E. I.; Hammond, J. R. Coupled Cluster Theory on Graphics Processing Units I. The Coupled Cluster Doubles Method. *Journal of Chemical Theory and Computation* **2011**, *7*, 1287–1295.
- (42) Pokhilko, P.; Epifanovsky, E.; Krylov, A. I. Double Precision Is Not Needed for Many-Body Calculations: Emergent Conventional Wisdom. *Journal of Chemical Theory and Computation* **2018**, *14*, 4088–4096.
- (43) NVIDIA cuEST: Accelerating Quantum Chemistry on NVIDIA GPUs, <https://docs.nvidia.com/cuda/cuest/>, 2026.
- (44) Thall, A. In *ACM SIGGRAPH 2006 Research Posters*, ACM: New York, NY, USA, 2006, 52–es.
- (45) Vydrov, O. A.; Van Voorhis, T. Nonlocal van der Waals density functional: The simpler the better. *The Journal of Chemical Physics* **2010**, *133*, 244103.
- (46) Ozaki, K.; Ogita, T.; Oishi, S.; Rump, S. M. Generalization of error-free transformation for matrix multiplication and its application. *Nonlinear Theory and Its Applications, IEICE* **2013**, *4*, 2–11.

- (47) Ootomo, H. ozIMMU - DGEMM on Int8 Tensor Core, <https://github.com/enpls0/ozIMMU>, Accessed: 2025-07-29, 2025.
- (48) Ozaki, K.; Uchino, Y.; Imamura, T. GEMMu8 - GEMM emulation using int8 matrix engines based on the Ozaki Scheme II, <https://github.com/RIKEN-RCCS/GEMMu8>, Accessed: 2025-07-31, 2025.
- (49) Huang, H.; Chow, E. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE: Dallas, TX, USA, 2018, pp 529–542.
- (50) Huang, H.; Sherrill, C. D.; Chow, E. Techniques for high-performance construction of Fock matrices. *The Journal of Chemical Physics* **2020**, *152*, 024122.
- (51) NVIDIA cuBLAS Release 13.2 Documentation, <https://docs.nvidia.com/cuda/cublas/>, 2026.
- (52) Schwarz, A.; Anders, A.; Brower, C.; Bayraktar, H.; Gunnels, J.; Clark, K.; Xu, R. G.; Rodriguez, S.; Cayrols, S.; Tabaszewski, P.; Podlozhnyuk, V. In *Proceedings of the Supercomputing Asia and International Conference on High Performance Computing in Asia Pacific Region*, Association for Computing Machinery: New York, NY, USA, 2026, pp 91–101.
- (53) Okuta, R.; Unno, Y.; Nishino, D.; Hido, S.; Loomis, C. In *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- (54) Thakkar, V.; Ramani, P.; Cecka, C.; Shivam, A.; Lu, H.; Yan, E.; Kosaian, J.; Hoemmen, M.; Wu, H.; Kerr, A.; Nicely, M.; Merrill, D.; Blasig, D.; Qiao, F.; Majcher, P.; Springer, P.; Hohnerbach, M.; Wang, J.; Gupta, M. CUTLASS, 2023.
- (55) Krishnan, R.; Binkley, J. S.; Seeger, R.; Pople, J. A. Self-consistent molecular orbital methods. XX. A basis set for correlated wave functions. *The Journal of Chemical Physics* **1980**, *72*, 650–654.
- (56) Weigend, F.; Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Physical Chemistry Chemical Physics* **2005**, *7*, 3297–3305.
- (57) Weigend, F. Accurate Coulomb-fitting basis sets for H to Rn. *Physical Chemistry Chemical Physics* **2006**, *8*, 1057–1065.
- (58) Becke, A. D. Density-functional thermochemistry. III. The role of exact exchange. *The Journal of Chemical Physics* **1993**, *98*, 5648–5652.
- (59) Stephens, P. J.; Devlin, F. J.; Chabalowski, C. F.; Frisch, M. J. Ab Initio Calculation of Vibrational Absorption and Circular Dichroism Spectra Using Density Functional Force Fields. *The Journal of Physical Chemistry* **1994**, *98*, 11623–11627.
- (60) Aprà, E. et al. NWChem: Past, present, and future. *The Journal of Chemical Physics* **2020**, *152*, 184102.
- (61) Smith, D. G. A.; Burns, L. A.; Simmonett, A. C.; Parrish, R. M.; Schieber, M. C.; Galvelis, R.; Kraus, P.; Kruse, H.; Di Remigio, R.; Alenaizan, A.; James, A. M.; Lehtola, S.; Misiewicz, J. P.; Scheurer, M.; Shaw, R. A.; Schriber, J. B.; Xie, Y.; Glick, Z. L.; Sirianni, D. A.; O’Brien, J. S.; Waldrop, J. M.; Kumar, A.; Hohenstein, E. G.; Pritchard, B. P.; Brooks, B. R.; Schaefer Henry F., I.; Sokolov, A. Y.; Patkowski, K.; DePrince A. Eugene, I.; Bozkaya, U.; King, R. A.; Evangelista, F. A.; Turney, J. M.; Crawford, T. D.; Sherrill, C. D. PSI4 1.4: Open-source software for high-throughput quantum chemistry. *The Journal of Chemical Physics* **2020**, *152*, 184108.
- (62) Neese, F.; Wennmohs, F.; Becker, U.; Riplinger, C. The ORCA quantum chemistry program package. *The Journal of Chemical Physics* **2020**, *152*, 224108.
- (63) Lehtola, S.; Steigemann, C.; Oliveira, M. J.; Marques, M. A. Recent developments in libxc - A comprehensive library of functionals for density functional theory. *SoftwareX* **2018**, *7*, 1–5.

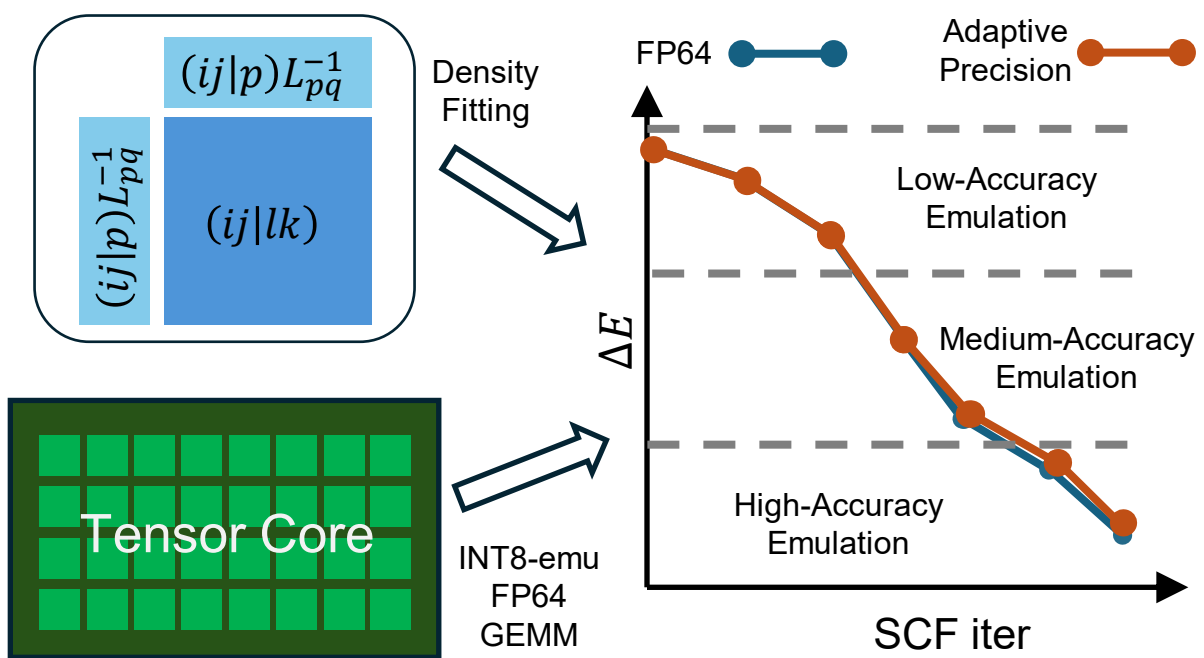


Figure 6: TOC Graphic.