

MakoXC: Rearchitecting DFT Exchange–Correlation with Matrix-Aligned and Knowledge-Organized Sparsity

Haozhi Han^{1,4}, Fusong Ju², Jing Bai³, Ruge Zhang^{5,4}, Xiang Zhao⁵, Liang Yuan⁵, Yunquan Zhang⁵, Ting Cao⁴, Yunxin Liu⁴, Yifeng Chen¹, and Kun Li^{4†}

¹School of Computer Science, Peking University, Beijing, China

²Microsoft Research, Beijing, China ³Microsoft Research, Redmond, USA

⁴Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China

⁵Chinese Academy of Sciences, Beijing, China

Abstract—Density Functional Theory (DFT) is essential to materials science and drug discovery, yet exchange–correlation (XC) evaluation remains a major bottleneck. Linear-scaling methods exploit electronic nearsightedness but introduce irregular sparse workloads that underutilize modern AI accelerators. We present MakoXC, a matrix-aligned XC engine that restructures nearsightedness-induced sparsity into regular accelerator-friendly computation. MakoXC combines Matrix-Aligned Cells to organize interactions into dense clusters, Sparsity-Guided Activation to expose implicit sparsity for practical linear scaling, and a Kernel-Fused Pipeline to consolidate fragmented workloads into a compute-intensive execution path. MakoXC achieves average speedups of $67.8\times$ over standard XC evaluation and $4.7\times$ over state-of-the-art linear-scaling methods. Integrated into a production-grade commercial DFT package, MakoXC scales to ubiquitin (1,231 atoms, def2-SVP) on 64 GPUs, completing end-to-end DFT in under five minutes. MakoXC demonstrates how scientific workloads can achieve genuine low complexity while maximizing parallel efficiency on AI accelerators.

Index Terms—Density Functional Theory, AI Accelerators, Tensor Contraction, Sparsity, Linear-Scaling

I. INTRODUCTION

The rapid rise of artificial intelligence (AI) workloads has fundamentally reshaped the computing landscape [1]–[3]. Modern accelerators—particularly GPUs—are increasingly architected around AI-native operations, with specialized units such as Tensor Cores that execute fixed-shape matrix multiply–accumulate (MMA) tiles [4]–[6]. These units deliver peak throughput only when operands are dense, regular, and tile-aligned, thereby rewarding only computation that is uniform and highly structured [7]–[10].

As a cornerstone of electronic-structure theory [11], [12], Density Functional Theory (DFT) exerts widespread impact on domains ranging from materials science [13]–[15] to drug discovery [16]–[18]. Its accuracy and efficiency, however, rely critically on the evaluation of the exchange–correlation

(XC) contribution—the integration of energy densities and their derivatives over large, atom-centered grids [19]–[22]. In practically relevant large-scale DFT calculations, numerical XC evaluation dominates the overall cost and becomes the main computational bottleneck. [23]–[26]. With its conventional cubic-scaling complexity $\mathcal{O}(N^3)$, XC evaluation is both the most pressing challenge in practice and a compelling target for acceleration on modern hardware [23], [25], [27].

Attempts to accelerate XC evaluation on modern accelerators has evolved from dense formulations to linear-scaling methods. As shown in Table I, dense formulations align naturally with hardware and achieve high utilization, but only by inflating computation and rigidifying the cubic complexity [28]–[30]. Linear-scaling methods exploit the electronic nearsightedness to lower the asymptotic complexity of XC evaluation from $\mathcal{O}(N^3)$ to near-linear $\mathcal{O}(N)$, enabling substantial performance gains at large scale [31]–[35].

However, the numerical manifestation of linear-scaling methods is highly **irregular sparsity**, which introduces two-level bottlenecks: *at the algorithmic level*, the sparsity is exposed only in a fragmented form centered on individual atomic orbital (AO) values ϕ_μ , leaving the richer implicit sparsity induced by AO products $\phi_\mu\phi_\nu$ hidden and difficult to translate into usable algorithmic benefits [23]; *at the hardware level*, the irregular patterns lead to fragmented execution that misaligns with MMA tiling and fails to efficiently utilize modern AI accelerators [8], [36], [37]. As a result, linear-scaling XC evaluation still falls short of its theoretical promise in practice, leaving a substantial gap between algorithmic sparsity and efficient execution on modern AI accelerators.

In this paper, we propose **MakoXC**¹, a modular XC evaluation system that rearchitects electronic-nearsightedness-induced sparsity into matrix-aligned structures, enabling im-

[†]Corresponding author (likun@air.tsinghua.edu.cn).

Part of this work was done at Microsoft Research and subsequently supported by PerfXLab Technologies.

¹The name “MakoXC” reflects its principle—*Matrix-Aligned and Knowledge-Organized Sparsity*—and evokes the Mako shark, renowned for its speed, symbolizing the pursuit of peak performance on modern AI accelerators.

TABLE I: Comparison of Computational Structures and Execution Models of XC Evaluation in DenseXC [28], GPU4PySCF [38], and MakoXC.

XC Method	Scaling (Theo.)	Matrix Structure	Sparsity Exploitation	Execution Feature
DenseXC	cubic	dense	○	monolithic
GPU4PySCF	near-linear	irregular sparse	●	kernel-fragmented
MakoXC	linear	aligned sparse	●	pipeline-fused

licit sparsity exploitation and highly parallel execution on Tensor Core GPUs.

The central insight is that sparsity, long viewed as an obstacle, can instead be restructured into a bridge for algorithm–hardware co-design, connecting algorithmic efficiency with hardware utilization. Electronic nearsightedness induces global irregularity yet simultaneously creates local density. MakoXC exploits this duality through the principle of matrix-aligned and knowledge-organized sparsity: reorganizing irregular patterns into accelerator-aligned structures and unlocks implicit sparsity through domain knowledge and physical laws. Guided by this principle, MakoXC enables XC evaluation, for the first time, to unite practical linear complexity with high parallel efficiency on GPUs.

To realize this insight, MakoXC builds upon three tightly integrated techniques:

Matrix-Aligned Cells (MACs). MACs reorganize locally dense clusters hidden within global irregularity into fixed-shape, MMA-aligned cells by (1) clustering locality induced by electronic nearsightedness, (2) aligning computation with accelerator tiles, and (3) adapting to basis-set and grid resolution. This design transforms physically induced irregularity into accelerator-compatible units, preserving locality while exposing regular compute structures that enable subsequent implicit sparsity exploitation and efficient parallel execution.

Sparsity-Guided Activation (SGA). Building on MACs, SGA exposes the deeper implicit sparsity embedded in AO products and translates it into executable structured computation through a knowledge-organized two-stage activation scheme. Going beyond explicit sparsity pruning, it first applies value-based filtering to identify numerically significant cells, and then uses AO product–density correspondence to activate only the density cells associated with significant AO products. By fully exploiting both explicit and implicit sparsity induced by electronic nearsightedness, SGA enables structured computation with numerical correctness and practical linear complexity.

Kernel-Fused Pipeline (KFP). To translate structural advantages into real performance, KFP unifies XC evaluation into a fused on-chip pipeline that collapses the disjoint GEMM and reduction stages into a single tensor-core-centric dataflow, together with direct SRAM reuse and domain-knowledge-guided redundancy elimination, including symmetry folding and normcache gating. This fusion transforms XC evaluation from a fragmented process dominated by frequent small-kernel launches, scheduling overheads, and memory traffic into a compute-intensive, high-throughput pipeline, fully unleashing

Tensor Cores.

We evaluate MakoXC across three distinct molecular series and two widely used basis set configurations. Our results show that MakoXC delivers average speedups of $67.8\times$, $5.15\times$ and $4.7\times$ over DenseXC [28], GPU4PySCF [38]–[40] and GauXC [41]–[43], respectively, while exhibiting essentially perfect agreement with the theoretical linear-scaling trend. Notably, when integrated into a mature commercial DFT package, MakoXC enables scalable XC evaluation for ubiquitin (1,231 atoms, def2-SVP) on 64 GPUs, allowing the end-to-end DFT calculation to complete in under five minutes and extending practical DFT simulations to larger molecular systems.

The main contributions of our paper include:

- We introduce matrix-aligned cells, a hardware-native expression for electronic nearsightedness that unifies algorithmic sparsity with accelerator regularity.
- Built on the matrix-aligned foundation, we co-design sparsity-guided activation and kernel-fused pipeline to uncover implicit sparsity and unify fragmented XC workloads for practical linear scaling on modern accelerators.
- We implement these innovations as a modular XC engine compatible with mainstream DFT software and deployed in a production-grade commercial package, demonstrating broad integrability and industrial reliability.
- Beyond XC evaluation, this work advances HPC by showing that matrix-aligned formulation offers a viable design principle for translating irregular scientific sparsity into accelerator-native execution.

II. BACKGROUND AND CHALLENGES

A. Density Functional Theory and Exchange–Correlation Evaluation

DFT describes the ground-state properties of an interacting electron system entirely in terms of its electron density [11], [44], [45]. The ground-state energy can be written as a functional of ρ ,

$$E[\rho] = T_s[\rho] + E_{\text{ext}}[\rho] + E_J[\rho] + E_{\text{xc}}[\rho], \quad (1)$$

where T_s , E_{ext} , and E_J denote the kinetic, external, and Coulomb terms, and E_{xc} is the XC contribution. The XC functional $E_{\text{xc}}[\rho]$ embodies the many-body effects of electron exchange and correlation, whose inherent complexity precludes any closed analytic representation [22], [27], [46], [47]. As a result, E_{xc} must be approximated through density functionals evaluated by numerical integration [19], [48]. To this end, progress along Jacob’s ladder systematically enriches XC functional dependence on the density and its derivatives, thereby improving accuracy, but simultaneously imposes greater computational demands, making E_{xc} the principal bottleneck in practical DFT calculations [20].

Standard XC Evaluation. In practice, XC contributions are evaluated numerically on atom-centered quadrature grids in three steps. First, the electron density is constructed on the grid:

$$\rho(r_g) = \sum_{\mu\nu} D_{\mu\nu} \phi_\mu(r_g) \phi_\nu(r_g), \quad (2)$$

where $D_{\mu\nu}$ is the density matrix and ϕ_μ atomic orbital basis functions. Second, the XC functional is evaluated pointwise to produce energy densities and functional derivatives, a computationally light step [42], [49]. Third, the XC energy and Fock contributions are assembled by integrations:

$$V_{\mu\nu}^{\text{XC}} = \sum_g w_g \phi_\mu(r_g) v_{\text{XC}}(r_g) \phi_\nu(r_g). \quad (3)$$

Among these, Steps 1 and 3 dominate runtime, both requiring dense basis–grid integrations with complexity $O(N_{\text{grid}}N^2)$ (basis size N and number of grid points N_{grid}). This cubic-like scaling rapidly becomes prohibitive for large systems [28]–[30].

Linear-Scaling XC Evaluation. The key algorithmic innovation for overcoming this cost relies on electronic near-sightedness [31]: atomic orbitals (AOs) decay rapidly, so each grid point interacts only with nearby basis functions [50]–[52]. Linear-scaling implementations exploit this property by introducing shell-dependent cutoff radii $r_{\text{cut},M}$, allowing distant AO contributions to be discarded [50], [53], [54]. This reduces both density construction and potential assembly to near-linear scaling with system size [55]. However, the resulting AO-locality-driven sparsity is highly irregular in practice [41]. This irregular sparsity underpins all linear-scaling XC methods, yet simultaneously constitutes the central challenge we target in this work.

B. Memory Hierarchy and Tensor Cores on GPU

The NVIDIA GPU architecture is organized as a scalable array of multithreaded Streaming Multiprocessors (SMs), each hosting numerous CUDA Cores, specialized Tensor Cores, and a hierarchical memory system [56]. Execution within an SM proceeds in warps of 32 threads, each equipped with its own register file. Threads within the same thread block share an on-chip SRAM space that offers low-latency, programmer-managed access, while all threads can access off-chip HBM, which provides larger capacity but significantly higher latency. As summarized in Table II, this hierarchy reflects the fundamental trade-off that increasing memory capacity comes at the cost of longer access latency [57].

Among the compute units, Tensor Cores are specialized for high-throughput matrix multiply–accumulate, defined as $D_{m \times n} = A_{m \times k} \times B_{k \times n} + C_{m \times n}$, delivering several-fold higher throughput than conventional CUDA Cores but restricted to fixed tile dimensions (e.g., $m = 8$, $n = 8$, $k = 4$ in FP64 precision) [4]. While this design achieves exceptional efficiency for dense GEMM, the rigid tile geometry prevents effective utilization under the irregular and fragmented sparsity patterns of scientific workloads such as XC evaluation [8].

TABLE II: Memory Hierarchy on A100 Tensor Core GPU.

Memory Tier	Capacity	Latency (cycles)
Off-chip HBM (Global Memory)	80 GiB / GPU	~290
On-chip SRAM (Shared Memory)	164 KiB / SM	~22
On-chip SRAM (Registers, 32-bit)	64 Ki / SM	~1

C. Opportunity: Bridge Low Complexity to High Parallelism with Sparsity

XC evaluation exposes a paradox: dense methods saturate hardware parallelism but rigidify cubic complexity, whereas linear-scaling methods reduce complexity yet yield irregular sparsity misaligned with accelerators. Crucially, this sparsity is not random but physically induced—and this distinction is key: while globally irregular, it exhibits local density and deeper implicit sparsity that can be systematically reorganized into regular, tile-aligned computational units and further mined for the implicit sparsity obscured by the native formulation. This rearchitecting not only aligns linear-scaling XC evaluation with modern GPUs, but also drives it from nominal near-linear scaling toward practical linear scaling. Sparsity, long regarded as an obstacle, thus becomes the key lever that bridges low algorithmic complexity with high accelerator parallelism.

D. Challenges

In this subsection, we outline three key challenges:

Challenge #1: *How can nearsightedness-induced irregularity be reorganized into regular, accelerator-aligned units for efficient parallel execution?* In linear-scaling XC evaluation, the resulting workloads are fragmented and irregular, rendering them “anti-hardware” and unable to exploit modern accelerators. Existing remedies—such as enlarging batch sizes [36], [42], invoking CUTLASS device functions [23], or employing complex scheduling and memory control in libraries like DBCSR [58] and NTPoly [37]—only mask superficial overheads. The fundamental pathology of fragmented, irregular workloads persists.

Challenge #2: *How can nearsightedness-induced sparsity be fully exploited as an algorithmic advantage to realize practical linear scaling?* Existing methods exploit only explicit sparsity from Grid–AO mapping at the level of individual basis functions [23], [59], while overlooking the richer, implicit sparsity inherent in AO products (e.g., $\phi_\mu(r_g)\phi_\nu(r_g)$). In principle, this implicit sparsity could greatly reduce computation and memory traffic. Yet doing so in a physically valid manner gives rise to highly irregular workloads, which have thus far precluded efficient mapping onto GPUs. As a result, existing linear-scaling XC evaluation methods achieve linear scaling primarily in asymptotic theory, but have yet to realize true linear scaling in practical GPU execution.

Challenge #3: *How can fragmented, memory-bound workloads be fused into unified pipelines to efficiently harness the throughput of modern AI accelerators?* Existing linear-scaling methods remain constrained by irregular workloads fragmented into numerous small kernel invocations. Such fragmentation not only incurs substantial scheduling and launch overheads, but also produces workloads too small and uneven to sustain the high-throughput execution model on which modern AI accelerators rely. Consequently, even though each individual kernel can often attain reasonably good implementation efficiency with vendor-optimized libraries, overall performance is still dominated by HBM traffic and launch




Fig. 1: Overall Architecture of MakoXC.

III. SYSTEM DESIGN

As illustrated in Fig. 1, MakoXC is an algorithm–hardware co-designed system that maps domain-specific knowledge in XC evaluation onto the execution model of modern AI accelerators through three tightly coupled components. Matrix-Aligned Cells (§ III-A) form the structural foundation, which reorganizes domain-specific knowledge-induced irregularity into regular, tile-aligned computational units for highly parallel execution on modern AI accelerators. Built on this foundation, Sparsity-Guided Activation (§ III-B) serves as the activation stage, which identifies the physically effective cells and computations through domain-specific knowledge, translating deeper implicit sparsity into executable structured sparsity embedded into mainstream DFT software before the SCF loop. Built on the activated workload, Kernel-Fused Pipeline (§ III-C) serves as the execution pipeline, which consolidates fragmented operations into a unified, high-throughput pipeline integrated into mainstream DFT software within the SCF loop.

A. Matrix-Aligned Cells

In addressing the inherent tension between low algorithmic complexity and high parallel efficiency, the primary challenge lies in leveraging electronic nearsightedness—an essential source of near-linear algorithms—in a parallel-friendly manner. Traditional radius-based batching strategies can, in principle, achieve low complexity, but they induce highly irregular computations that misalign with the regular execution model of modern accelerators and incur substantial scheduling overhead.

To overcome this challenge, we propose Matrix-Aligned Cells (MAC), a mechanism that systematically maps the locality induced by electronic nearsightedness into regular computational structures. The core unit is the **MAC**, which encapsulates locality into dense micro-matrices tailored for efficient accelerator utilization. The design of MAC is guided by three key observations: (1) Local Clustering: Owing to electronic nearsightedness, the Grid-AO matrix is globally sparse but locally dense, forming clusters of nonzeros due to localized basis overlaps and neighboring grid regions. (2)




Fig. 2: Matrix-Aligned Cells.

Alignment Optimality: Modern accelerators (e.g., Tensor Core Units) execute matrix operations through fixed-size MMA tiles (e.g., 16×16). Aligning computation units to these tiles maximizes hardware efficiency. (3) Granularity Variability: In DFT calculations, the choice of basis set and grid level directly influences the size of dense clusters. Based on the above observations, the MAC embodies three essential properties: (1) Regularity: Each MAC is a fixed-shape micro-matrix whose dimensions align with multiples of MMA tiles (e.g., 16×16 , 32×32), ensuring compatibility with accelerator execution units. (2) Adaptivity: The MAC size is not fixed but subject to a tradeoff: larger MACs yield more efficient matrix operations, but may partially compromise sparsity. At the beginning of a DFT calculation, the optimal MAC size is automatically determined based on the chosen quadrature grid and basis set, with the constraint that it must not exceed the available on-chip memory. (3) Sparsity-Awareness: If a MAC exhibits extreme sparsity (e.g., for a 32×32 tile, fewer than 103 nonzero entries, i.e., $> 90\%$ sparsity), it is stored in CSR format and marked as sparse by setting the corresponding entry in the *is_dense* array to 0.

While individual MACs encapsulate the local manifestations of electronic nearsightedness, leveraging this locality at scale requires a global perspective. We therefore design a **MAC-based sparse storage format**, which organizes local structures into a globally consistent layout, thereby enabling efficient parallel execution of XC evaluation. As shown in Fig. 2, the format employs a CSR-like index at the MAC granularity: *RowPtr* (length $N_{\text{row}} + 1$) stores the prefix offsets of MACs per row, so that the interval $[RowPtr[i], RowPtr[i + 1])$ enumerates all MACs of row i ; *ColumnIdx* records the column index of each MAC. Each MAC is annotated with *isDense*, indicating whether it is stored densely or sparsely, and with a *normcache*, which holds a cached magnitude metric (in our case, the Frobenius norm) to support normcache gating (§ III-C). The actual data is stored in a contiguous MAC List, ensuring efficient access and parallel traversal.

B. Sparsity-Guided Activation

Building upon the MAC design, the next critical challenge is to fully exploit the electronic nearsightedness principle of linear-scaling DFT—maximizing the utilization of both




Fig. 3: (a) Sparsity-Guided Activation. (b) Correspondence Activation for Exploiting Implicit Sparsity.

explicit and implicit sparsity within the MAC paradigm. To address this, we propose Sparsity-Guided Activation (SGA), a unified scheme operating through knowledge-organized sparsity. By enforcing knowledge-driven and physically rigorous filtering of MACs—activating only those with significant contributions while discarding all others as completely sparse—SGA transforms electronic nearsightedness into actionable computational efficiency, thereby pushing the practical complexity of XC evaluation ever closer to the theoretical linear-scaling limit.

SGA is founded on a key observation we term *static sparsity*: the sparsity induced by electronic nearsightedness depends solely on the basis set and quadrature grid, and thus remains invariant throughout the SCF loop. Leveraging this property, SGA performs a one-time activation at the very beginning of the DFT computation, before the SCF loop. Thereafter, all subsequent XC evaluations are confined to the MACs activated in this initial SGA, enabling consistent acceleration across the entire SCF loop. As shown in Fig. 3(a), the design of SGA consists of two stages, namely Value-Based Filtering and Correspondence Activation.

1) *Value-Based Filtering*: Within the theory of electronic nearsightedness, the exponential decay of Gaussian-type orbitals (GTOs) implies that their values diminish rapidly with increasing distance from atomic centers. Exploiting this property, SGA introduces a magnitude threshold ε (default 10^{-12}) on each AO basis function μ evaluated at grid points r . For any MAC in the Grid-AO matrix, if all elements fall below ε , the MAC is discarded outright; if more than 90% of its elements fall below ε , the MAC is retained but stored in the CSR sparse format. The remaining MACs are deemed value-significant and are fully activated, preserved in the global MAC-based sparse storage format to sustain efficient downstream XC evaluation.

2) *Correspondence Activation*: For the Grid-AO matrix, SGA leverages Value-Based Filtering to uncover explicit sparsity. However, this approach, akin to traditional radius-based batching strategies, only exploits explicit sparsity at the level of individual AO basis functions and thus falls short of revealing the deeper implicit sparsity inherent in the computation.

To address this, SGA introduces Correspondence Activation, which operates at the MAC level to further exploit the implicit sparsity embedded in the theory of electronic nearsightedness. This method is founded on the following key observations: (1) **Richer combinatorial sparsity**. The evaluation of grid density involves AO products $\phi_\mu(r_g)\phi_\nu(r_g)$, whose product structure inherently contains numerous near-zero values. This implies abundant implicit sparsity that can be exploited to substantially reduce redundant memory accesses and computations. (2) **One-to-one correspondence**. As can be seen from electron density formula (Equation 2), each density matrix entry $D_{\mu\nu}$ has a strict one-to-one correspondence with the AO product $\phi_\mu(r_g)\phi_\nu(r_g)$. (3) **Multiplicative property**. If an AO product is zero, the result of multiplying it with $D_{\mu\nu}$ is necessarily zero, even if $D_{\mu\nu}$ itself is nonzero. Consequently, such $D_{\mu\nu}$ entries can be excluded from storage and computation. (4) **Static invariance**. Both the AO products $\phi_\mu(r_g)\phi_\nu(r_g)$ and the one-to-one correspondence remain invariant throughout the computation. Consequently, once a density matrix entry $D_{\mu\nu}$ is identified as having no contribution to the final result, this property holds throughout the entire DFT computation. Based on these observations, while the density matrix D is dynamically updated and typically dense, its fixed correspondence with AO products induces a form of static correspondence sparsity: the mapping between $D_{\mu\nu}$ and AO products is invariant, and near-zero AO products render the corresponding $D_{\mu\nu}$ entries persistently sparse across the entire computation. This property not only exposes deeper implicit sparsity under electronic nearsightedness but also enables pre-activation at the beginning of the DFT computation by exploiting its static invariance.

Therefore, Correspondence Activation requires only the prior identification of numerically significant (non-negligible) AO products and the localization of their corresponding positions within the density-matrix MACs, which are then activated; all remaining inactive MACs are deemed sparse and can be safely discarded. In its ideal design, the numerical significance of an AO product is determined by directly evaluating the magnitude of $p_{\mu\nu}(r) = \phi_\mu(r)\phi_\nu(r)$. However, explicitly computing all AO products at every grid point is prohibitively expensive. To address this challenge, we leverage the shellpair construction and its screening results—already pre-computed in earlier stages of DFT—as an efficient surrogate for AO-product evaluation, thereby enabling the identification of MACs to be activated in the density matrix.

Importantly, the connection between shell-pair construction and the magnitude of AO products is not heuristic or empirical, but is grounded in rigorous mathematical principles. We next formalize why shell-pair construction serves as a principled surrogate for AO-product significance, ensuring a conservative screening criterion for Correspondence Activation—one that guarantees no false negatives.

Theorem 1 (Gaussian Product Theorem). *For any two primitive Gaussian functions χ_μ, χ_ν with exponents α_μ, α_ν and centers R_μ, R_ν , the product can be rewritten as a single*

Gaussian:

$$\chi_\mu(r)\chi_\nu(r) = P_{\mu\nu}(r) e^{-(\alpha_\mu+\alpha_\nu)\|r-R_c\|^2} e^{-\gamma\|R_\mu-R_\nu\|^2}, \quad (4)$$

where

$$\gamma_{\mu\nu} = \frac{\alpha_\mu\alpha_\nu}{\alpha_\mu + \alpha_\nu}, \quad R_c = \frac{\alpha_\mu R_\mu + \alpha_\nu R_\nu}{\alpha_\mu + \alpha_\nu}, \quad (5)$$

and $P_{\mu\nu}$ is a polynomial whose degree equals the combined angular momenta of χ_μ and χ_ν .

Corollary 1. *The spatial magnitude of an AO product is concentrated near the line connecting the two shell centers, and its global amplitude is exponentially suppressed by the factor $e^{-\gamma\|R_\mu-R_\nu\|^2}$. Thus, when the shell centers are sufficiently far apart, $p_{\mu\nu}(r)$ is negligible for all grid points.*

Theorem 2 (Schwarz Bound). *The norm of an AO product satisfies*

$$\|p_{\mu\nu}\|_2^2 = (\mu\nu|\mu\nu) \leq \sqrt{(\mu\mu|\mu\mu)} \sqrt{(\nu\nu|\nu\nu)}. \quad (6)$$

Corollary 2. *If the self-overlap integral of an AO is small, then its product with any other AO is necessarily small in the global sense. Hence, AO product significance can be strictly bounded by orbital self-integrals.*

Proposition 1 (Shellpair Significance Criterion). *Let shellpair (P, Q) denote the set of all AO combinations from shells P and Q . If either the shell-center distance $\|R_P - R_Q\|$ is sufficiently large or the Schwarz bound falls below a threshold ε , then for any $\mu \in P, \nu \in Q$,*

$$|p_{\mu\nu}(r)| < \varepsilon, \quad \forall r. \quad (7)$$

Proof 1. *By Theorem 1, the maximum AO product amplitude is controlled by the exponential decay term $e^{-\gamma\|R_P-R_Q\|^2}$; when the shell-center distance is large, this term tends to zero, ensuring insignificance. By Theorem 2, the global norm of the AO product does not exceed the Schwarz bound; if this bound is below ε , then the AO product is uniformly negligible. Hence, the proposition follows. $\square$*

Taken together, the significance of AO products is rigorously determined by inter-shell distances and orbital self-integrals—the very quantities exploited in shellpair construction and screening. Thus, the screened shellpair set constitutes a conservative surrogate for AO-product significance: while some negligible products may survive due to the use of upper bounds (over-activation), no genuinely significant contributions are ever discarded (no false negatives). This conservative guarantee preserves numerical correctness while dramatically reducing the cost of explicit AO-product evaluation, thereby providing both the theoretical rigor and practical efficiency that justify shellpair-based Correspondence Activation.

As shown in Fig. 3(b), SGA employs the shell pairs pre-computed in the preceding stage to construct a set of unique MAC indices $\mathcal{S}$, which is automatically ordered and duplicate-free. Correspondence Activation is then realized by activating the MACs in the density matrix that correspond to the indices

Algorithm 1 Correspondence Activation on Density Matrix.

Require: List of screened shellpairs $\mathcal{P}$, Dim of MACs dim

Ensure: Activated MACs in Density Matrix $\mathcal{D}$

```

1:  $\mathcal{S} \leftarrow \emptyset$  ▷ set of unique MAC indices  $\mathcal{S} : (r, c)$ 
2: for all  $(x, y) \in \mathcal{P}$  do
3:   for  $i \leftarrow \text{sh2bf}[x]$  to  $\text{sh2bf}[x+1] - 1$  do
4:     for  $j \leftarrow \text{sh2bf}[y]$  to  $\text{sh2bf}[y+1] - 1$  do
5:        $r \leftarrow \lfloor i/dim \rfloor, c \leftarrow \lfloor j/dim \rfloor$ 
6:        $\mathcal{S} \leftarrow \mathcal{S} \cup \{(r, c)\}, \{(c, r)\}$  ▷ symmetric
7:     end for
8:   end for
9: end for
10:  $\mathcal{D} \leftarrow \text{ACTIVATION}(\mathcal{S})$ 

```

in $\mathcal{S}$. The $\mathcal{S}$ construction process is detailed in Algorithm 1: for each shell pair, the procedure invokes the `sh2bf` function to locate all contracted Gaussian basis functions within the pair and determine their corresponding positions (i, j) in the density matrix. These positions are then mapped to MAC indices (r_b, c_b) and, by exploiting the symmetry of the density matrix, the resulting indices are inserted into $\mathcal{S}$. Finally, at Line 10, the effective MACs of the density matrix D are activated based on the index set $\mathcal{S}$. Correspondence Activation leverages the AO–density correspondence, rigorously activating only the significant MACs by exploiting shellpair screening and its strict equivalence to AO product significance.

In summary, SGA fully capitalizes on the sparsity induced by electronic nearsightedness: it first eliminates redundant MACs through value-based filtering, and then rigorously activates only the truly effective density-matrix MACs. This two-stage design uncovers both explicit and implicit sparsity and, for the first time, translates implicit sparsity into a practically effective acceleration framework.

C. Kernel-Fused Pipeline

Before entering the SCF loop, MakoXC employs MAC to encode physical locality into regularized structural units and uses SGA to pre-filter and activate only those contributing to computation. To translate these structural advantages into hardware-level gains, it is imperative to design a kernel pipeline tailored to the MACs.

To this end, we propose the Kernel-Fused Pipeline (KFP)—a high-performance XC execution framework that systematically leverages the structural regularity of MACs to fuse critical operations, eliminate redundancy, and map execution efficiently onto modern accelerators, thereby unlocking fine-grained kernel-level optimizations and maximizing computational throughput. KFP is motivated by three key observations: (1) HBM Bottleneck: XC evaluation typically involves multiple GEMM and DOT operators; executing them independently makes HBM transfers the dominant performance bottleneck. (2) On-Chip Reuse: Each MAC is sufficiently small to fit within the on-chip SRAM of modern GPUs, enabling cross-operator data reuse directly in fast SRAM. (3) Shape Alignment: The consistent shape of MACs ensures that the input and output granularity across different operators is naturally aligned, eliminating additional data transformation overhead




Fig. 4: Kernel-Fused Pipeline.

and making computation fusion virtually cost-free. Building on these observations, KFP achieves operator-level fusion and on-chip reuse at the MAC granularity: by pipelining GEMM, DOT, and related operators within SRAM, it maximizes on-chip data reuse and drastically reduces HBM traffic. Meanwhile, the regularity of MAC shapes enables efficient operator fusion without requiring costly data reordering.

In this section, we take the kernel **COMPUTERHO** in XC evaluation as a representative example. The computation of grid density ρ involves multiplying the Grid-AO matrix with the density matrix, followed by a dot product with the Grid-AO matrix on a per-grid-row basis. For clarity, we denote the former as Φ_1 and the latter as Φ_2 , though they represent the same matrix: $\rho = (\Phi_1 \times D) \cdot \Phi_2$. Instead of issuing separate GEMM and DOT calls, KFP fuses them into a single pipeline to fully exploit on-chip data locality and minimize memory traffic. The overall workflow is shown in Fig. 4, with the corresponding pseudocode in Algorithm 2.

In **COMPUTERHO**, each thread block is assigned one row of Φ_2 and processes all valid MACs within that row (after SGA pruning). The first loop (Lines 5, DOT Loop) iterates over all valid MACs in the current row of Φ_2 , where the iterator mid points to one such MAC, denoted as ao' . Based on its position (row, col) in Φ_2 , the corresponding MAC acc —with which MAC ao' will be involved in a DOT computation—can be identified in the intermediate matrix $\Phi_1 D$, although this matrix is never explicitly materialized. The position (row, col) indicates that acc is generated by multiplying all valid MACs in row row of Φ_1 with those corresponding in column col of D . Accordingly, in Lines 12 and 15, the second and third loops (GEMM-row Loop & GEMM-col Loop) traverse all valid MACs in row row of Φ_1 and column col of D , respectively. At Line 19, pairs of candidate MACs are checked for alignment in the matrix multiplication; once matched, the multiplication is executed on Tensor Cores or CUDA Cores. If both a and d correspond to dense MACs, the computation is executed on Tensor Cores for maximum throughput, whereas sparse cases are executed on CUDA cores. After the second and third loops, complete, acc has already been computed and cached in SRAM. Without writing back to HBM, it can directly participate in the DOT operation with ao' on SRAM. In Line 11-36, denotes GEMM operations executed on SRAM, while denotes DOT operations executed on SRAM. Both

Algorithm 2 **COMPUTERHO** in Kernel-Fused Pipeline.

Require: Sparse Grid-AO Matrix AO , Sparse Density Matrix D , threshold ϵ

Ensure: Grid density values ρ

```

1:  $row = blockIdx.x$ ;
2:  $start = AO.row\_ptr[row]$ 
3:  $end = AO.row\_ptr[row + 1]$ 
4: // DOT Loop
5: for  $mid \leftarrow start$  to  $end - 1$  do
6:    $v_{ao2} = AO.norm[mid]$ 
7:    $col = AO.col\_index[mid]$ 
8:    $D\_index = D.row\_ptr[col]$ 
9:    $D\_end = D.row\_ptr[col + 1]$ 
10:   $d_{row} = D.col\_index[D\_index]$ 
11:  // GEMM-row Loop ▷ GEMM
12:  for  $A\_index \leftarrow start$  to  $mid$  do ▷ Symmetry Folding
13:     $a_{col} = AO.col\_index[A\_index]$ 
14:    // GEMM-col Loop
15:    while  $d_{row} < a_{col}$  and  $D\_index < D\_end$  do
16:       $D\_index++$ 
17:       $d_{row} = D.col\_index[D\_index]$ 
18:    end while
19:    if  $a_{col} = d_{row}$  then
20:       $v_{ao1} = AO.norm[A\_index]$ 
21:       $v_d = D.norm[D\_index]$ 
22:      if  $v_{ao1} * v_d * v_{ao2} < \epsilon$  then
23:        Continue ▷ Normcache Gating
24:      end if
25:       $scale \leftarrow (A\_index = mid) ? 0.5 : 1.0$ 
26:       $Load\ d = D.data[D\_index].T * scale$ 
27:       $Load\ ao = AO.data[A\_index]$ 
28:      if  $ao.isDense \& d.isDense$  then
29:        Tensor Core Compute:  $acc = acc + ao \otimes d$ 
30:      else
31:        CUDA Core Compute:  $acc = acc + ao \otimes d$ 
32:      end if
33:    end if
34:  end for
35:     $ao' \leftarrow Load\ AO.data[mid]$  ▷ DOT
36:     $out \leftarrow out + 2 * ao' * acc$ 
37: end for
38: StoreResults:  $\rho \leftarrow Reduce\_sum(out)$ 

```

are fused on-chip to maximize data reuse. Finally, at Line 38, all MACs in row row of Φ_2 are reduced to obtain the corresponding density ρ .

The above description demonstrates how KFP transforms traditionally disjoint GEMM and DOT invocations into a unified on-chip pipeline, substantially reducing both memory traffic and kernel launch overhead. Crucially, by enforcing a global fusion perspective, KFP also exposes redundant computations and data flows that contribute negligibly to numerical accuracy—patterns that are largely invisible when operations are performed in isolation.

More importantly, the fused execution framework of KFP naturally creates opportunities to eliminate such inefficiencies: within a unified pipeline, operators, data streams, and storage paths are systematically organized, enabling patterns that were previously scattered and elusive to emerge. Building on this insight, KFP further integrates two redundancy-elimination

mechanisms:

(1) Symmetry Folding: Because the density matrix D is symmetric ($D_{\mu\nu} = D_{\nu\mu}$), it is unnecessary to load and compute the entire matrix; evaluating only its upper triangular part suffices. In Algorithm 2 (Lines 12 and 25), KFP exploits this property by folding symmetric operations into one fused computation, thereby reducing memory loads and computation by nearly half.

(2) Normcache Gating: Although MACs activated by SGA are themselves numerically significant, their contributions after GEMM+DOT computations may still be negligible (e.g., $v_{ao1} \cdot v_d \cdot v_{ao2} < \varepsilon$). Therefore, as shown in Algorithm 2 (Line 22), KFP computes the product of the three MACs’ NormCaches before executing the triplet computation, compares it against the threshold ε , and filters out insignificant data streams on the fly.

By integrating Symmetry Folding and Normcache Gating into the fused pipeline, KFP retains only numerically significant computations while discarding redundant operations at the source, achieving both higher efficiency and stronger alignment with hardware execution.

KFP transforms XC evaluation from a memory-bound sequence of fragmented operators into a tightly fused on-chip pipeline, combining operator fusion with SRAM-level reuse to fully unleash accelerator performance. Beyond operator fusion, it further eliminates redundancy through symmetry folding and norm-cache gating, delivering a high-throughput execution pipeline for XC evaluation.

IV. EVALUATION

A. Evaluation Setup

Hardware. We evaluate performance on both single- and multi-GPU systems. The single-GPU setup consists of an AMD EPYC 7V13 CPU paired with an NVIDIA A100 Tensor Core GPU (Ampere architecture, Compute Capability 8.0). For multi-GPU evaluations, we employ a cluster where each node integrates 8 NVIDIA A100 GPUs and 96 AMD EPYC CPU cores, interconnected through 200 GB/s NVIDIA Mellanox HDR InfiniBand links.

Metrics. For standalone performance evaluation of MakoXC, we measure the average XC evaluation time within each SCF cycle [23]. For end-to-end evaluation, we instead use the average iteration time of each SCF cycle [60]. The average is taken over the first ten iterations, excluding the initial one to eliminate initialization overhead, thereby ensuring consistent and fair timing across software packages with different convergence algorithms.

State-of-the-art. Since many DFT packages are closed-source, reproducing prior results remains difficult. To ensure a fair comparison, we evaluated several widely recognized open-source state-of-the-art packages, GauXC (v1.0) [41], [42], and GPU4PySCF (v1.4.1) [38]–[40]. We also include DenseXC, our in-house implementation that ignores electronic near-sightedness, serving as a non-linear-scaling baseline. Notably, GauXC, which cannot independently perform SCF cycles, is integrated into our DFT framework and executed under the

same interface to ensure fairness; however, its heavy reliance on global memory often leads to out-of-memory failures on larger molecular systems.

Dataset. We have prepared three scalable classes of representative systems: (1) Polyglycine chains, a linear system scalable with chain length; (2) Boron-nitride sheets, two-dimensional materials scalable with sheet size; and (3) Water clusters, three-dimensional globular systems scalable with cluster size. These scalable benchmarks provide a systematic basis for assessing MakoXC’s performance across different molecular dimensionalities and complexities.

Aligned Parameter Settings. To ensure consistency and fairness in evaluating the performance of MakoXC, we adopted a unified set of parameter configurations across all experiments: (1) Basis sets: We primarily tested the widely used def2-SVP and def2-TZVP basis sets, obtained from the Basis Set Exchange [61]. (2) XC functional: All calculations employed the same B3LYP functional [62], [63]. (3) Integration grids: We adopted the same grid generation strategy as PySCF [64] and GPU4PySCF [38]. The grid level can be specified from 0 (very sparse) to 9 (very dense). Since different packages implement slightly different grid construction and grading schemes, we compared the corresponding grid levels by additionally reporting the number of grid points per atom. In all subsequent calculations, we fixed the grid level to 5, which corresponds to approximately 30,019 grid points per atom [65]–[67]. (4) Threshold (ε): For accuracy validation, we adopt the default threshold of 10^{-12} , whereas for performance evaluation, we set the threshold to 10^{-10} to ensure consistency with other software packages [41], [42].

B. Overall Performance

State-of-the-art Comparison. We evaluate the average XC evaluation time per SCF iteration and the corresponding speedup of MakoXC against all state-of-the-art software across diverse molecular systems and basis sets on a single NVIDIA A100 GPU. As shown in Fig. 5, MakoXC delivers substantial acceleration over all baselines. Under the def2-SVP basis set, MakoXC achieves an average speedup of $87.1\times$ over DenseXC, $3.6\times$ over GauXC, and $6.7\times$ over GPU4PySCF. When moving to the larger def2-TZVP basis set, MakoXC continues to sustain substantial gains of $48.4\times$ over DenseXC, $2.9\times$ over GauXC, and $3.6\times$ over GPU4PySCF, demonstrating that its performance advantage is sustained as system size and basis complexity increase. These results demonstrate MakoXC’s efficiency, robustness across basis sets, and ability to fully exploit modern accelerators.

Practical Linear Scaling Verification. Fig. 6 plots the execution time of DenseXC and MakoXC as the system size increases proportionally, together with reference cubic- and linear-scaling curves. The results clearly show that MakoXC is the first method to realize practical linear-scaling XC evaluation as molecular size grows, turning what was previously only a theoretical complexity advantage into an achieved system-level behavior. To the best of our knowledge, MakoXC is also the first method to combine theoretical linear complexity with




Fig. 5: Performance Comparison with State-of-the-art on a Single A100 GPU. (a) use def2-SVP, and (b) use def2-TZVP basis.




Fig. 6: Practical Linear Scaling Verification of MakoXC.




Fig. 8: Performance Breakdown of MakoXC.




Fig. 7: Accuracy Validation of MakoXC. Per-atom errors relative to Grid-Level=7 for all baseline across varying atomic grid densities with the threshold of 10^{-12} .

full alignment to modern hardware execution units, thereby fully unlocking their computational potential.

Accuracy Validation. We evaluate the numerical accuracy of all XC evaluation baselines using per-atom energy differences (Hartree/atom), taking the grid-level=7 result from PySCF as the reference across varying atomic grid densities. As a widely used and well-established quantum chemistry package, PySCF provides a reliable numerical reference for accuracy evaluation. To ensure a fair comparison, DenseXC,

GPU4PySCF, PySCF, and GauXC use the same grid generation scheme and screening threshold of 10^{-12} as MakoXC. The energy error ranges of all methods are shown in Fig. 7. Following the convention in quantum chemistry, XC energy errors within 1 mHartree/atom are regarded as chemically accurate [68]–[70]. Under this criterion, MakoXC exhibits excellent agreement with established software, thereby validating its correctness and reliability for practical quantum chemistry workloads.

C. Ablation Study

Performance Breakdown. Fig. 8 presents a performance breakdown of MakoXC across three representative molecular systems and multiple basis sets. First, applying only value-based filtering—exploiting the sparsity of the Grid-AO matrix while ignoring that of the Density Matrix—delivers an average $9.0\times$ speedup over the dense baseline. Second, employing the full SGA strategy, which additionally leverages Density Matrix sparsity, yields a further $1.71\times$ improvement. Finally, when combined with the high-performance KFP implementation,




Fig. 9: Runtime Overhead of SGA before SCF Loop.




Fig. 10: Kernel Performance Improvements in MakoXC.

MakoXC achieves an additional $4.1\times$ acceleration on average. These results demonstrate that each technique in MakoXC makes a substantial and complementary contribution to its overall performance improvement.

Overhead Analysis of SGA. To ensure that the SGA procedure itself does not incur additional performance overhead, we measured its execution time. As shown in Fig. 9, the XC SGA stage consistently accounts for less than 3% of the total runtime across both globular molecular systems and boron-nitride sheets, with its relative cost further diminishing as system size increases. In contrast, the performance benefits of SGA are substantial, as evidenced by the performance breakdown results. Similar observations hold for all other tested molecular systems. These results decisively establish SGA as a highly effective and scalable strategy for XC evaluation.

Breakdown of KFP. Fig. 10 illustrates a breakdown analysis of these kernel techniques using the COMPUTERHO kernel as a representative case, offering fine-grained insight into the performance improvements delivered by KFP. As shown, all three techniques consistently enhance kernel efficiency, with the fusion strategy yielding the most pronounced gains. Meanwhile, symmetry folding and normcache gating deliver increasing benefits as molecular system size grows. These results highlight the critical role of a high-performance kernel-fused pipeline design, which establishes a solid, hardware-aligned foundation for the overall MakoXC framework.

D. End-to-End Evaluation and Scalability

End-to-End Evaluation. We integrated MakoXC into our in-house DFT package, which provides a fully featured software stack together with highly optimized ERI kernels. This integration leads to substantial end-to-end performance advantages over state-of-the-art implementations. As shown in Fig. 11, the integrated system delivers an average $3.94\times$ speedup over GPU4PySCF under the def2-SVP basis set.

Scalability Evaluation. Since GPU4PySCF provides limited support for multi-GPU execution, scalability testing was




Fig. 11: End-to-End Performance Comparison. Speedup refers to the performance gain of the in-house DFT package with MakoXC over GPU4PySCF.




Fig. 12: Scalability of the In-house DFT Package with MakoXC.

performed only on the integrated system. XC integration is inherently amenable to multi-GPU parallelization, as each GPU can process a disjoint set of grid points with communication required only for the final reduction. Using ubiquitin (1,231 atoms) with the def2-SVP basis set, we scale the system from 1 to 8 nodes (1–64 GPUs). As shown in Fig. 12, the system sustains over 74% parallel efficiency even at 64 A100 GPUs across 8 nodes. At this scale, it completes a single-point energy calculation for ubiquitin in under five minutes, demonstrating both strong scalability and practical capability for large molecular simulations.

V. RELATED WORK

DFT is a cornerstone of electronic structure calculations, but traditional plane-wave codes such as VASP [71]–[75] and Quantum ESPRESSO [76]–[80] typically exhibit cubic scaling with system size, limiting their applicability to large-scale simulations [81], [82]. Linear-scaling DFT mitigates this limitation by exploiting electronic nearsightedness to approach $\mathcal{O}(N)$ complexity [11], [26], [33]–[35], [50], [51], [55]. Packages such as CP2K [83]–[85], SIESTA [86]–[89], and FHI-aims [90], [91] employ localized basis functions and sparsity-aware algorithms to reduce computational complexity. However, the resulting computations often exhibit irregular sparsity and data-dependent execution patterns that are difficult to map efficiently onto modern GPUs [50], [52], [92]–[94]. Recent work, Mako [95], demonstrates a matrix-centric alternative by rearchitecting irregular quantum-chemistry computations, particularly electron-repulsion integral (ERI) evaluation, into structured matrix-aligned kernels that can efficiently exploit modern AI accelerators. However, these techniques mainly target ERI evaluation, while exchange–correlation calculations

exhibit distinct geometry-dependent, dynamically evolving sparsity that requires new matrix-aligned techniques.

Among the major components of DFT, exchange-correlation (XC) evaluation is a performance-critical stage because it lies on the critical path of every SCF iteration and repeatedly processes large numbers of grid points, basis functions, and density-dependent quantities [20], [23], [26], [27], [46], [47]. Although decades of GPU acceleration efforts—from early GGA kernel offloading to systems such as QUICK, GauXC, and GPU4PySCF—have substantially improved XC performance, the irregular sparsity in linear-scaling XC still causes fragmented workloads, irregular memory access, and load imbalance, leaving a large gap between its theoretical efficiency and realized performance on modern AI accelerators [28]–[30], [36], [38]–[42], [96]. This gap arises because the sparsity exposed by nearsightedness is numerically beneficial but computationally irregular, making it difficult to organize into the regular, high-throughput execution patterns favored by modern accelerator architectures.

VI. CONCLUSION

MakoXC bridges linear-scaling XC evaluation with modern AI accelerators by uncovering implicit sparsity and reorganizing it into structured, matrix-aligned computation. It converts fragmented and irregular XC workloads into an efficient high-throughput execution path. This design enables practical linear scaling, significantly improves performance, and makes large-scale accurate DFT calculations practical.

REFERENCES

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017. [Online]. Available: <https://dl.acm.org/doi/10.5555/3295222.3295349>
- [2] J. Wang, J. Hsu, and Z. Qin, “A comprehensive analysis of nvidia’s technological innovations, market strategies, and future prospects,” *International Journal of Information Technologies and Systems Approach (IJITSA)*, vol. 17, no. 1, pp. 1–16, 2024. [Online]. Available: <https://doi.org/10.4018/IJITSA.344423>
- [3] A. Wael and A. Madi, “Accelerating artificial intelligence: The role of gpus in deep learning and computational advancements,” *East Journal of Engineering*, vol. 1, no. 1, pp. 31–46, 2025. [Online]. Available: <https://doi.org/10.63496/eje.Vol1.Iss1.34>
- [4] S. Markidis, S. W. Der Chien, E. Laure, I. B. Peng, and J. S. Vetter, “Nvidia tensor core programmability, performance & precision,” in *2018 IEEE international parallel and distributed processing symposium workshops (IPDPSW)*. IEEE, 2018, pp. 522–531. [Online]. Available: <https://doi.org/10.1109/IPDPSW.2018.00091>
- [5] G. Schieffer, D. A. De Medeiros, J. Faj, A. Marathe, and I. Peng, “On the rise of amd matrix cores: Performance, power efficiency, and programmability,” in *2024 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 2024, pp. 132–143. [Online]. Available: <https://doi.org/10.1109/ISPASS61541.2024.00022>
- [6] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers *et al.*, “In-datacenter performance analysis of a tensor processing unit,” in *Proceedings of the 44th annual international symposium on computer architecture*, 2017, pp. 1–12. [Online]. Available: <https://doi.org/10.1145/3140659.3080246>
- [7] Y. E. Wang, G.-Y. Wei, and D. Brooks, “Benchmarking tpu, gpu, and cpu platforms for deep learning,” *arXiv preprint arXiv:1907.10701*, 2019. [Online]. Available: <https://doi.org/10.48550/arXiv.1907.10701>
- [8] W. Sun, A. Li, T. Geng, S. Stuijk, and H. Corporaal, “Dissecting tensor cores via microbenchmarks: Latency, throughput and numeric behaviors,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 34, no. 1, pp. 246–261, 2022. [Online]. Available: <https://doi.org/10.1109/TPDS.2022.3217824>
- [9] V. Sharma, G. K. Gupta, and M. Gupta, “Performance benchmarking of gpu and tpu on google colab for convolutional neural network,” in *Applications of Artificial Intelligence in Engineering: Proceedings of First Global Conference on Artificial Intelligence and Applications (GCAIA 2020)*. Springer, 2021, pp. 639–646. [Online]. Available: https://doi.org/10.1007/978-981-33-4604-8_49
- [10] Y. Wang, D. Murnane, N. Choma, S. Farrell, P. Calafiura *et al.*, “Benchmarking gpu and tpu performance with graph neural networks,” *arXiv preprint arXiv:2210.12247*, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2210.12247>
- [11] W. Kohn, A. D. Becke, and R. G. Parr, “Density functional theory of electronic structure,” *The journal of physical chemistry*, vol. 100, no. 31, pp. 12974–12980, 1996. [Online]. Available: <https://doi.org/10.1021/jp9606691>
- [12] C. D. Sherrill, “Frontiers in electronic structure theory,” *The Journal of chemical physics*, vol. 132, no. 11, 2010. [Online]. Available: <https://doi.org/10.1063/1.3369628>
- [13] J. Hafner, C. Wolverton, and G. Ceder, “Toward computational materials design: the impact of density functional theory on materials research,” *MRS bulletin*, vol. 31, no. 9, pp. 659–668, 2006. [Online]. Available: <https://doi.org/10.1557/mrs2006.174>
- [14] A. E. Mattsson, P. A. Schultz, M. P. Desjarlais, T. R. Mattsson, and K. Leung, “Designing meaningful density functional theory calculations in materials science—a primer,” *Modelling and Simulation in Materials Science and Engineering*, vol. 13, no. 1, p. R1, 2004. [Online]. Available: <https://doi.org/10.1088/0965-0393/13/1/R01>
- [15] J. Neugebauer and T. Hickel, “Density functional theory in materials science,” *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 3, no. 5, pp. 438–448, 2013. [Online]. Available: <https://doi.org/10.1002/wcms.1125>
- [16] N. Ye, Z. Yang, and Y. Liu, “Applications of density functional theory in covid-19 drug modeling,” *Drug discovery today*, vol. 27, no. 5, pp. 1411–1419, 2022. [Online]. Available: <https://doi.org/10.1016/j.drudis.2021.12.017>
- [17] V. T. Sabe, T. Ntombela, L. A. Jhamba, G. E. Maguire, T. Govender, T. Naicker, and H. G. Kruger, “Current trends in computer aided drug design and a highlight of drugs discovered via computational techniques: A review,” *European Journal of Medicinal Chemistry*, vol. 224, p. 113705, 2021. [Online]. Available: <https://doi.org/10.1016/j.ejmech.2021.113705>
- [18] H. Tandon, T. Chakraborty, and V. Suhag, “A brief review on importance of dft in drug design,” *Res. Med. Eng. Sci.*, vol. 7, no. 4, pp. 791–795, 2019. [Online]. Available: <https://doi.org/10.31031/RMES.2019.07.00068>
- [19] D. N. Laikov, “Fast evaluation of density functional exchange-correlation terms using the expansion of the electron density in auxiliary basis sets,” *Chemical Physics Letters*, vol. 281, no. 1-3, pp. 151–156, 1997. [Online]. Available: [https://doi.org/10.1016/S0009-2614\(97\)01206-2](https://doi.org/10.1016/S0009-2614(97)01206-2)
- [20] J. P. Perdew and K. Schmidt, “Jacob’s ladder of density functional approximations for the exchange-correlation energy,” in *AIP conference proceedings*, vol. 577, no. 1. American Institute of Physics, 2001, pp. 1–20. [Online]. Available: <https://doi.org/10.1063/1.1390175>
- [21] J. M. Martin, C. W. Bauschlicher Jr, and A. Ricca, “On the integration accuracy in molecular density functional theory calculations using gaussian basis sets,” *Computer physics communications*, vol. 133, no. 2-3, pp. 189–201, 2001. [Online]. Available: [https://doi.org/10.1016/S0010-4655\(98\)00174-0](https://doi.org/10.1016/S0010-4655(98)00174-0)
- [22] S. E. Wheeler and K. Houk, “Integration grid errors for meta-gga-predicted reaction energies: Origin of grid errors for the m06 suite of functionals,” *Journal of Chemical Theory and Computation*, vol. 6, no. 2, pp. 395–404, 2010. [Online]. Available: <https://doi.org/10.1021/ct900639j>
- [23] R. Stocks and G. M. J. Barca, “Efficient algorithms for gpu accelerated evaluation of the dft exchange-correlation functional,” 2025. [Online]. Available: <https://doi.org/10.1021/acs.jctc.5c01229>
- [24] Y. Wang, D. Hait, K. G. Johnson, O. J. Fajen, J. H. Zhang, R. D. Guerrero, and T. J. Martínez, “Extending gpu-accelerated gaussian integrals in the terachem software package to f type orbitals: Implementation and applications,” *The Journal of Chemical Physics*,

- vol. 161, no. 17, 2024. [Online]. Available: <https://doi.org/10.1063/5.0233523>
- [25] H. Laqua, J. Kussmann, and C. Ochsenfeld, "An improved molecular partitioning scheme for numerical quadratures in density functional theory," *The Journal of Chemical Physics*, vol. 149, no. 20, 2018. [Online]. Available: <https://doi.org/10.1063/1.5049435>
- [26] A. M. Burrow and M. Sierka, "Linear scaling hierarchical integration scheme for the exchange-correlation term in molecular and periodic systems," *Journal of Chemical Theory and Computation*, vol. 7, no. 10, pp. 3097–3104, 2011. [Online]. Available: <https://doi.org/10.1021/ct200412r>
- [27] R. J. Bartlett, I. Grabowski, S. Hirata, and S. Ivanov, "The exchange-correlation potential in ab initio density functional theory," *The Journal of chemical physics*, vol. 122, no. 3, p. 034104, 2005. [Online]. Available: <https://doi.org/10.1063/1.1809605>
- [28] K. Yasuda, "Accelerating density functional calculations with graphics processing unit," *Journal of Chemical Theory and Computation*, vol. 4, no. 8, pp. 1230–1236, 2008. [Online]. Available: <https://doi.org/10.1021/ct8001046>
- [29] I. S. Ufimtsev and T. J. Martinez, "Quantum chemistry on graphical processing units. 3. analytical energy gradients, geometry optimization, and first principles molecular dynamics," *Journal of Chemical Theory and Computation*, vol. 5, no. 10, pp. 2619–2628, 2009. [Online]. Available: <https://doi.org/10.1021/ct9003004>
- [30] S. Seritan, C. Bannwarth, B. S. Fales, E. G. Hohenstein, S. I. Kokkila-Schumacher, N. Luehr, J. W. Snyder, C. Song, A. V. Titov, I. S. Ufimtsev *et al.*, "Terachem: Accelerating electronic structure and ab initio molecular dynamics with graphical processing units," *The Journal of chemical physics*, vol. 152, no. 22, 2020. [Online]. Available: <https://doi.org/10.1063/5.0007615>
- [31] E. Prodan and W. Kohn, "Nearsightedness of electronic matter," *Proceedings of the National Academy of Sciences*, vol. 102, no. 33, pp. 11 635–11 638, 2005. [Online]. Available: <https://doi.org/10.1073/pnas.0505436102>
- [32] R. F. Bader, "Nearsightedness of electronic matter as seen by a physicist and a chemist," *The Journal of Physical Chemistry A*, vol. 112, no. 51, pp. 13 717–13 728, 2008. [Online]. Available: <https://doi.org/10.1021/jp806282j>
- [33] J. C. Prentice, J. Aarons, J. C. Womack, A. E. Allen, L. Andrinopoulos, L. Anton, R. A. Bell, A. Bhandari, G. A. Bramley, R. J. Charlton *et al.*, "The onetep linear-scaling density functional theory program," *The Journal of chemical physics*, vol. 152, no. 17, 2020. [Online]. Available: <https://doi.org/10.1063/5.0004445>
- [34] E. Hernández, M. Gillan, and C. Goringe, "Linear-scaling density-functional-theory technique: The density-matrix approach," *Physical Review B*, vol. 53, no. 11, p. 7147, 1996. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.53.7147>
- [35] G. E. Scuseria, "Linear scaling density functional calculations with gaussian orbitals," *The Journal of Physical Chemistry A*, vol. 103, no. 25, pp. 4782–4790, 1999. [Online]. Available: <https://doi.org/10.1021/jp990629s>
- [36] M. Manathunga, Y. Miao, D. Mu, A. W. Gotz, and K. M. Merz Jr, "Parallel implementation of density functional theory methods in the quantum interaction computational kernel program," *Journal of Chemical Theory and Computation*, vol. 16, no. 7, pp. 4315–4326, 2020. [Online]. Available: <https://doi.org/10.1021/acs.jctc.0c00290>
- [37] W. Dawson and T. Nakajima, "Massively parallel sparse matrix function calculations with ntpoly," *Computer Physics Communications*, vol. 225, pp. 154–165, 2018. [Online]. Available: <https://doi.org/10.1016/j.cpc.2017.12.010>
- [38] X. Wu, Q. Sun, Z. Pu, T. Zheng, W. Ma, W. Yan, Y. Xia, Z. Wu, M. Huo, X. Li *et al.*, "Enhancing gpu-acceleration in the python-based simulations of chemistry frameworks," *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 15, no. 2, p. e70008, 2025. [Online]. Available: <https://doi.org/10.48550/arXiv.2404.09452>
- [39] R. Li, Q. Sun, X. Zhang, and G. K.-L. Chan, "Introducing gpu acceleration into the python-based simulations of chemistry framework," *The Journal of Physical Chemistry A*, vol. 129, no. 5, pp. 1459–1468, 2025. [Online]. Available: <https://doi.org/10.1021/acs.jpca.4c05876>
- [40] Z. Pu and Q. Sun, "Enhancing pyscf-based quantum chemistry simulations with modern hardware, algorithms, and python tools," *arXiv preprint arXiv:2506.06661*, 2025. [Online]. Available: <https://doi.org/10.48550/arXiv.2506.06661>
- [41] D. B. Williams-Young, W. A. De Jong, H. J. Van Dam, and C. Yang, "On the efficient evaluation of the exchange correlation potential on graphics processing unit clusters," *Frontiers in chemistry*, vol. 8, p. 581058, 2020. [Online]. Available: <https://doi.org/10.3389/fchem.2020.581058>
- [42] D. B. Williams-Young, A. Bagusetty, W. A. de Jong, D. Doerfler, H. J. van Dam, Á. Vázquez-Mayagoitia, T. L. Windus, and C. Yang, "Achieving performance portability in gaussian basis set density functional theory on accelerator based architectures in nwchemex," *Parallel Computing*, vol. 108, p. 102829, 2021. [Online]. Available: <https://doi.org/10.1016/j.parco.2021.102829>
- [43] D. B. Williams-Young, A. Asadchev, D. T. Popovici, D. Clark, J. Waldrop, T. L. Windus, E. F. Valeev, and W. A. de Jong, "Distributed memory, gpu accelerated fock construction for hybrid, gaussian basis density functional theory," *The Journal of Chemical Physics*, vol. 158, no. 23, p. 234104, 06 2023. [Online]. Available: <https://doi.org/10.1063/5.0151070>
- [44] L. J. Bartolotti and K. Flurichick, "An introduction to density functional theory," *Reviews in computational chemistry*, pp. 187–216, 1996.
- [45] F. M. Bickelhaupt and E. J. Baerends, "Kohn-sham density functional theory: predicting and understanding chemistry," *Reviews in computational chemistry*, pp. 1–86, 2000. [Online]. Available: <https://doi.org/10.1002/9780470125922.ch1>
- [46] A. M. Köster, R. Flores-Moreno, and J. U. Reveles, "Efficient and reliable numerical integration of exchange-correlation energies and potentials," *The Journal of chemical physics*, vol. 121, no. 2, pp. 681–690, 2004. [Online]. Available: <https://doi.org/10.1063/1.1759323>
- [47] L. Balbás, J. L. Martins, and J. M. Soler, "Evaluation of exchange-correlation energy, potential, and stress," *Physical Review B*, vol. 64, no. 16, p. 165110, 2001. [Online]. Available: <https://doi.org/10.1103/PhysRevB.64.165110>
- [48] M. Challacombe, "Linear scaling computation of the fock matrix. v. hierarchical cubature for numerical integration of the exchange-correlation matrix," *The Journal of Chemical Physics*, vol. 113, no. 22, pp. 10 037–10 043, 2000. [Online]. Available: <https://doi.org/10.1063/1.1316012>
- [49] S. Lehtola, C. Steigemann, M. J. Oliveira, and M. A. Marques, "Recent developments in libxc—a comprehensive library of functionals for density functional theory," *SoftwareX*, vol. 7, pp. 1–5, 2018. [Online]. Available: <https://doi.org/10.1016/j.softx.2017.11.002>
- [50] J. Junquera, O. Paz, D. Sánchez-Portal, and E. Artacho, "Numerical atomic orbitals for linear-scaling calculations," *Physical Review B*, vol. 64, no. 23, p. 235111, 2001. [Online]. Available: <https://doi.org/10.1103/PhysRevB.64.235111>
- [51] K. N. Kudin and G. E. Scuseria, "Linear-scaling density-functional theory with gaussian orbitals and periodic boundary conditions: Efficient evaluation of energy and forces via the fast multipole method," *Physical Review B*, vol. 61, no. 24, p. 16440, 2000. [Online]. Available: <https://doi.org/10.1103/PhysRevB.61.16440>
- [52] J. Andzel and E. Wimmer, "Density functional gaussian-type-orbital approach to molecular geometries, vibrations, and reaction energies," *The Journal of chemical physics*, vol. 96, no. 2, pp. 1280–1303, 1992. [Online]. Available: <https://doi.org/10.1063/1.462165>
- [53] D. R. Bowler and T. Miyazaki, "O(n) methods in electronic structure calculations," *Reports on Progress in Physics*, vol. 75, no. 3, p. 036503, 2012. [Online]. Available: <https://doi.org/10.1088/0034-4885/75/3/036503>
- [54] M. Sierka, A. Hogeckamp, and R. Ahlrichs, "Fast evaluation of the coulomb potential for electron densities using multipole accelerated resolution of identity approximation," *The Journal of chemical physics*, vol. 118, no. 20, pp. 9136–9148, 2003. [Online]. Available: <https://doi.org/10.1063/1.1567253>
- [55] C. O'Rourke and D. R. Bowler, "Linear scaling density matrix real time tddft: Propagator unitarity and matrix truncation," *The Journal of chemical physics*, vol. 143, no. 10, 2015. [Online]. Available: <https://doi.org/10.1063/1.4919128>
- [56] Z. Jia, M. Maggioni, B. Staiger, and D. P. Scarpazza, "Dissecting the nvidia volta gpu architecture via microbenchmarking," *arXiv preprint arXiv:1804.06826*, 2018. [Online]. Available: <https://doi.org/10.48550/arXiv.1804.06826>
- [57] X. Mei and X. Chu, "Dissecting gpu memory hierarchy through microbenchmarking," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 1, pp. 72–86, 2016. [Online]. Available: <https://doi.org/10.1109/TPDS.2016.2549523>

- [58] I. Sivkov, P. Seewald, A. Lazzaro, and J. Hutter, "Dbscr: A blocked sparse tensor algebra library," *arXiv preprint arXiv:1910.13555*, 2019. [Online]. Available: <https://doi.org/10.48550/arXiv.1910.13555>
- [59] E. Besalú and R. Carbó-Dorca, "The general gaussian product theorem," *Journal of mathematical chemistry*, vol. 49, no. 8, pp. 1769–1784, 2011. [Online]. Available: <https://doi.org/10.1007/s10910-011-9857-9>
- [60] J. L. Galvez Vallejo, G. M. Barca, and M. S. Gordon, "High-performance gpu-accelerated evaluation of electron repulsion integrals," *Molecular Physics*, vol. 121, no. 9–10, p. e2112987, 2023. [Online]. Available: <https://doi.org/10.1080/00268976.2022.2112987>
- [61] B. S. Exchange, "Basis Set," <https://www.basissetexchange.org/>, 2025.
- [62] A. D. Becke, "A new mixing of hartree-fock and local density-functional theories," *Journal of chemical physics*, vol. 98, no. 2, pp. 1372–1377, 1993. [Online]. Available: <https://doi.org/10.1063/1.464304>
- [63] P. J. Stephens, F. J. Devlin, C. F. Chabalowski, and M. J. Frisch, "Ab initio calculation of vibrational absorption and circular dichroism spectra using density functional force fields," *The Journal of physical chemistry*, vol. 98, no. 45, pp. 11 623–11 627, 1994. [Online]. Available: <https://doi.org/10.1021/j100096a001>
- [64] Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. D. McClain, E. R. Sayfutyarova, S. Sharma *et al.*, "Pyscf: the python-based simulations of chemistry framework," *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 8, no. 1, p. e1340, 2018. [Online]. Available: <https://doi.org/10.48550/arXiv.1701.08223>
- [65] V. I. Lebedev, "Quadratures on a sphere," *USSR Computational Mathematics and Mathematical Physics*, vol. 16, no. 2, pp. 10–24, 1976. [Online]. Available: [https://doi.org/10.1016/0041-5553\(76\)90100-2](https://doi.org/10.1016/0041-5553(76)90100-2)
- [66] P. M. Gill and S.-H. Chien, "Radial quadrature for multiexponential integrands," *Journal of computational chemistry*, vol. 24, no. 6, pp. 732–740, 2003. [Online]. Available: <https://doi.org/10.1002/jcc.10211>
- [67] M. E. Mura and P. J. Knowles, "Improved radial grids for quadrature in molecular density-functional calculations," *Journal of Chemical Physics*, vol. 104, no. 24, pp. 9848–9858, 1996. [Online]. Available: <https://doi.org/10.1063/1.471749>
- [68] F. Ju, X. Wei, L. Huang, A. J. Jenkins, L. Xia, J. Zhang, J. Zhu, H. Yang, B. Shao, P. Dai *et al.*, "Acceleration without disruption: Dft software as a service," *Journal of Chemical Theory and Computation*, vol. 20, no. 24, pp. 10838–10851, 2024. [Online]. Available: <https://doi.org/10.1021/acs.jctc.4c00940>
- [69] N. Luehr, I. S. Ufimtsev, and T. J. Martínez, "Dynamic precision for electron repulsion integral evaluation on graphical processing units (gpus)," *Journal of Chemical Theory and Computation*, vol. 7, no. 4, pp. 949–954, 2011, PMID: 26606344. [Online]. Available: <https://doi.org/10.1021/ct100701w>
- [70] S. Tsuji, Y. Ito, H. Fujii, N. Yokogawa, K. Suzuki, K. Nakano, and A. Kasagi, "Dynamic screening of two-electron repulsion integrals in gpu parallelization," in *2024 Twelfth International Symposium on Computing and Networking Workshops (CANDARW)*. IEEE, 2024, pp. 211–217. [Online]. Available: <https://doi.org/10.1109/CANDARW64572.2024.00041>
- [71] M. Hacene, A. Anciaux-Sedrakian, X. Rozanska, D. Klahr, T. Guignon, and P. Fleurat-Lessard, "Accelerating vasp electronic structure calculations using graphic processing units," *Journal of computational chemistry*, vol. 33, no. 32, pp. 2581–2589, 2012. [Online]. Available: <https://doi.org/10.1002/jcc.23096>
- [72] A. Maniopoulou, E. R. Davidson, R. Grau-Crespo, A. Walsh, I. J. Bush, C. R. A. Catlow, and S. M. Woodley, "Introducing k-point parallelism into vasp," *Computer Physics Communications*, vol. 183, no. 8, pp. 1696–1701, 2012. [Online]. Available: <https://doi.org/10.1016/j.cpc.2012.03.009>
- [73] F. Wende, M. Marsman, J. Kim, F. Vasilev, Z. Zhao, and T. Steinke, "Openmp in vasp: Threading and simd," *International Journal of Quantum Chemistry*, vol. 119, no. 12, p. e25851, 2019. [Online]. Available: <https://doi.org/10.1002/qua.25851>
- [74] V. Stegailov, G. Smirnov, and V. Vechev, "Vasp hits the memory wall: Processors efficiency comparison," *Concurrency and Computation: Practice and Experience*, vol. 31, no. 19, p. e5136, 2019. [Online]. Available: <https://doi.org/10.1002/cpe.5136>
- [75] J. Hafner, "Ab-initio simulations of materials using vasp: Density-functional theory and beyond," *Journal of computational chemistry*, vol. 29, no. 13, pp. 2044–2078, 2008. [Online]. Available: <https://doi.org/10.1002/jcc.21057>
- [76] T. A. Barnes, T. Kurth, P. Carrier, N. Wichmann, D. Prendergast, P. R. Kent, and J. Deslippe, "Improved treatment of exact exchange in quantum espresso," *Computer Physics Communications*, vol. 214, pp. 52–58, 2017. [Online]. Available: <https://doi.org/10.1016/j.cpc.2017.01.008>
- [77] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo *et al.*, "Quantum espresso: a modular and open-source software project for quantum simulations of materials," *Journal of physics: Condensed matter*, vol. 21, no. 39, p. 395502, 2009. [Online]. Available: <https://doi.org/10.1088/0953-8984/21/39/395502>
- [78] P. Giannozzi, O. Baseggio, P. Bonfà, D. Brunato, R. Car, I. Carnimeo, C. Cavazzoni, S. De Gironcoli, P. Delugas, F. Ferrari Ruffino *et al.*, "Quantum espresso toward the exascale," *The Journal of chemical physics*, vol. 152, no. 15, 2020. [Online]. Available: <https://doi.org/10.1063/5.0005082>
- [79] S. Scandolo, P. Giannozzi, C. Cavazzoni, S. de Gironcoli, A. Pasquarello, and S. Baroni, "First-principles codes for computational crystallography in the quantum-espresso package," *Zeitschrift für Kristallographie-Crystalline Materials*, vol. 220, no. 5–6, pp. 574–579, 2005. [Online]. Available: <https://doi.org/10.1524/zkri.220.5.574.65062>
- [80] I. Carnimeo, F. Affinito, S. Baroni, O. Baseggio, L. Bellentani, R. Bertossa, P. D. Delugas, F. F. Ruffino, S. Orlandini, F. Spiga *et al.*, "Quantum espresso: One further step toward the exascale," *Journal of Chemical Theory and Computation*, vol. 19, no. 20, pp. 6992–7006, 2023. [Online]. Available: <https://doi.org/10.1021/acs.jctc.3c00249>
- [81] J. Feng, L. Wan, J. Li, S. Jiao, X. Cui, W. Hu, and J. Yang, "Massively parallel implementation of iterative eigensolvers in large-scale plane-wave density functional theory," *Computer Physics Communications*, vol. 299, p. 109135, 2024. [Online]. Available: <https://doi.org/10.1016/j.cpc.2024.109135>
- [82] W. Jia, Z. Cao, L. Wang, J. Fu, X. Chi, W. Gao, and L.-W. Wang, "The analysis of a plane wave pseudopotential density functional theory code on a gpu machine," *Computer Physics Communications*, vol. 184, no. 1, pp. 9–18, 2013. [Online]. Available: <https://doi.org/10.1016/j.cpc.2012.08.002>
- [83] J. Hutter, M. Iannuzzi, F. Schiffmann, and J. VandeVondele, "cp2k: atomistic simulations of condensed matter systems," *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 4, no. 1, pp. 15–25, 2014. [Online]. Available: <https://doi.org/10.1002/wcms.1159>
- [84] T. D. Kühne, M. Iannuzzi, M. Del Ben, V. V. Rybkin, P. Seewald, F. Stein, T. Laino, R. Z. Khaliullin, O. Schütt, F. Schiffmann *et al.*, "Cp2k: An electronic structure and molecular dynamics software package-quickstep: Efficient and accurate electronic structure calculations," *The Journal of Chemical Physics*, vol. 152, no. 19, 2020. [Online]. Available: <https://doi.org/10.1063/5.0007045>
- [85] D. Yokelson, N. V. Tkachenko, R. Robey, Y. W. Li, and P. A. Dub, "Performance analysis of cp2k code for ab initio molecular dynamics on cpus and gpus," *Journal of chemical information and modeling*, vol. 62, no. 10, pp. 2378–2386, 2022. [Online]. Available: <https://doi.org/10.1021/acs.jcim.1c01538>
- [86] A. García, N. Papior, A. Akhtar, E. Artacho, V. Blum, E. Bosoni, P. Brandimarte, M. Brandbyge, J. I. Cerdá, F. Corsetti *et al.*, "Siesta: Recent developments and applications," *The Journal of chemical physics*, vol. 152, no. 20, 2020. [Online]. Available: <https://doi.org/10.1063/5.0005077>
- [87] E. Artacho, E. Anglada, O. Diéguez, J. D. Gale, A. García, J. Junquera, R. M. Martin, P. Ordejón, J. M. Pruneda, D. Sánchez-Portal *et al.*, "The siesta method: developments and applicability," *Journal of Physics: Condensed Matter*, vol. 20, no. 6, p. 064208, 2008. [Online]. Available: <https://doi.org/10.1088/0953-8984/20/6/064208>
- [88] F. Corsetti, "Performance analysis of electronic structure codes on hpc systems: a case study of siesta," *PLoS One*, vol. 9, no. 4, p. e95390, 2014. [Online]. Available: <https://doi.org/10.1371/journal.pone.0095390>
- [89] D. Sánchez-Portal, P. Ordejón, and E. Canadell, "Computing the properties of materials from first principles with siesta," in *Principles and applications of density functional theory in inorganic chemistry II*. Springer, 2004, pp. 103–170. [Online]. Available: <https://doi.org/10.1007/b97943>
- [90] V. Blum, M. Rossi, S. Kokott, and M. Scheffler, "The fhi-aims code: All-electron, ab initio materials simulations towards the exascale," *arXiv preprint arXiv:2208.12335*, 2022. [Online]. Available: <https://doi.org/10.48550/arXiv.2208.12335>
- [91] J. W. Abbott, C. M. Acosta, A. Akkoush, A. Ambrosetti, V. Atalla, A. Bagrets, J. Behler, D. Berger, B. Bieniek, J. Björk *et al.*,

- “Roadmap on advancements of the fhi-aims software package,” *arXiv preprint arXiv:2505.00125*, 2025. [Online]. Available: <https://doi.org/10.48550/arXiv.2505.00125>
- [92] P. Lin, X. Ren, X. Liu, and L. He, “Ab initio electronic structure calculations based on numerical atomic orbitals: Basic formalisms and recent progresses,” *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 14, no. 1, p. e1687, 2024. [Online]. Available: <https://doi.org/10.1002/wcms.1687>
- [93] A. L. Magalhaes, “Gaussian-type orbitals versus slater-type orbitals: a comparison,” *Journal of Chemical Education*, vol. 91, no. 12, pp. 2124–2127, 2014. [Online]. Available: <https://doi.org/10.1021/ed500437a>
- [94] B. Dunlap and N. Rosch, “The gaussian-type orbitals density-functional approach to finite systems,” in *Advances in Quantum Chemistry*. Elsevier, 1990, vol. 21, pp. 317–339. [Online]. Available: [https://doi.org/10.1016/S0065-3276\(08\)60603-6](https://doi.org/10.1016/S0065-3276(08)60603-6)
- [95] H. Han, K. Li, F. Ju, Q. Li, H. An, Y. Chen, Y. Zhang, T. Cao, and M. Yang, “Matrix is all you need: Rearchitecting quantum chemistry to scale on ai accelerators,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 2126–2142. [Online]. Available: <https://doi.org/10.1145/3712285.3759829>
- [96] M. Manathunga, H. M. Aktulga, A. W. Gotz, and K. M. Merz Jr, “Quantum mechanics/molecular mechanics simulations on nvidia and amd graphics processing units,” *Journal of Chemical Information and Modeling*, vol. 63, no. 3, pp. 711–717, 2023. [Online]. Available: <https://doi.org/10.1021/acs.jcim.2c01505>