Matryoshka: Optimization of Dynamic Diverse Quantum Chemistry Systems via Elastic Parallelism Transformation

Tuowei Wang[*]
Microsoft Research
Tsinghua University
Beijing, China

Kun Li[†]
Microsoft Research
Beijing, China

Donglin Bai
Microsoft Research
Beijing, China

Fusong Ju
Microsoft Research
Beijing, China

Leo Xia
Microsoft Research
Beijing, China

Ju Ren
Tsinghua University
Beijing, China

Yaoxue Zhang
Tsinghua University
Beijing, China

Ting Cao
Microsoft Research
Beijing, China

Mao Yang
Microsoft Research
Beijing, China

Abstract

AI infrastructures, predominantly GPUs, have delivered remarkable performance gains for deep learning. Conversely, scientific computing, exemplified by quantum chemistry systems, suffers from dynamic diversity, where computational patterns are more diverse and vary dynamically, posing a significant challenge to sponge acceleration off GPUs.

In this paper, we propose Matryoshka, a novel elastically-parallel technique for the efficient execution of quantum chemistry system with dynamic diversity on GPU. Matryoshka capitalizes on Elastic Parallelism Transformation, a property prevalent in scientific systems yet underexplored for dynamic diversity, to elastically realign parallel patterns with GPU architecture. Structured around three transformation primitives (Permutation, Deconstruction, and Combination), Matryoshka encompasses three core components. The Block Constructor serves as the central orchestrator, which reformulates data structures accommodating dynamic inputs and constructs fine-grained GPU-efficient compute blocks. Within each compute block, the Graph Compiler operates offline, generating high-performance code with clear computational path through an automated compilation process. The Workload Allocator dynamically schedules workloads with varying operational intensities to threads online. It achieves highly efficient parallelism for compute-intensive operations and facilitates fusion with neighboring memory-intensive operations automatically. Extensive evaluation shows that Matryoshka effectively addresses dynamic diversity, yielding acceleration improvements of up to 13.86× (average 9.41×) over prevailing state-of-the-art approaches on 13 quantum chemistry systems.

1 Introduction

Deep learning, powered by AI infrastructures like GPUs, has sparked transformative revolutions in various AI systems [17, 21, 27, 28, 48, 56]. Concurrently, scientific computing assumes an equally vital role, propelling breakthrough research in various scientific domains [15, 18, 31]. However, an intriguing observation lies in the prevailing state where the numerical scientific computing systems predominantly revolve around case-by-case optimizations on CPUs [5, 16, 29, 37, 57]. A noticeable dichotomy emerges as deep learning and scientific computing appear to progress as parallel entities. This intriguing scenario prompts a compelling question: Why the absence of systematic efforts to integrate AI infrastructures like GPUs into scientific computing?

This paper aims to unravel this puzzle by delving into quantum chemistry systems within scientific computing. Quantum Chemistry (QC) [10, 23, 26] is a pivotal scientific discipline that investigates the quantum mechanical properties of atomic and molecular systems. The profound insights gleaned from computational quantum chemistry find extensive applications across diverse industries, including materials science [41, 42], pharmaceuticals [12], and energy production [20].

Notwithstanding its transformative potential, the formidable computational complexity of QC poses a substantial demand on computing resources, thereby impeding progress within the scientific domain. With the booming rise of AI techniques, accelerators, primarily GPUs, have been increasingly employed to expedite QC systems.

While GPU could deliver a promising performance, particularly in General Matrix Multiply (GEMM) operations on tensors predominant in AI systems, it is imperative to recognize that computational patterns in scientific systems
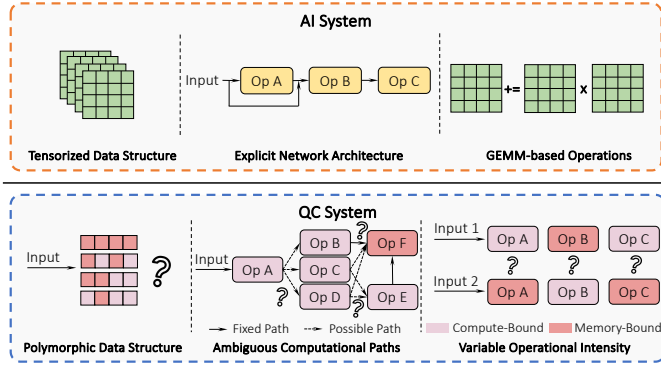

Figure 1. Comparison between AI system and QC system.

are significantly more *diverse*, i.e., a single system is built by various distinct operations with polymorphic data structures. Adding to this complexity, these diverse computational patterns are often varying at compile time or influenced by inputs at runtime, i.e., dynamic diversity. This introduces indiscernible and unpredictable performance bottlenecks in scientific systems, making optimization challenging without sufficient priors, as shown in Figure 1.

As a fundamental system in scientific computing, QC grapples with the significant challenge of dynamic diversity, evident across three primary misalignments against GPU architecture. These include polymorphic data structures causing divergent execution instructions, ambiguous computational paths contributing to a complex computing landscape at compile time, and variable operational intensity spanning compute-intensive to memory-intensive workloads at runtime.

In the presence of dynamic diversity, efficiently transitioning from CPU-centric to GPU-centric architectures for QC systems remains a substantial challenge, with the majority of current work still rooted in CPU-centric designs. Existing GPU-based QC systems typically perform a fixed optimization strategy on diverse operators known at compile time [8, 32–35, 50–53]. However, this static parallelism lacks adaptability to dynamic changes in tasks during runtime, leading to a noticeable decrease in GPU utilization. While fine-grained, operator-wise runtime solutions exist, their applicability is constrained to supporting a single module within the QC system. The concatenation of these operators to construct a comprehensive system introduces significant overhead for data layout conversion across diverse operators [50–53]. Furthermore, this approach may lead to performance conflicts when alternative techniques are applied to different operators. Beyond the performance impact, the re-design of operators for mapping to GPU SIMT architecture and the data layout conversion between operators can also incur accuracy loss, which hardly meets the stringent requirements of scientists from a physics perspective [8, 32–35].

In this paper, we present MATRYOSHKA, a novel elastically-parallel technique for the accurate and efficient execution of quantum chemistry system with dynamic diversity on GPU.

The cornerstone of MATRYOSHKA is Elastic Parallelism Transformation, a property prevalent in scientific systems yet underexplored for dynamic diversity. Most scientific operators have one or more dimensions (we call them EPT-axes), whose computation can be arbitrarily reordered without affecting the result. For instance, in a QC system, contracted basis functions operator ψ_a are fixed linear combinations of fundamental basis function ϕ_{ak}, given by $\psi_a(\mathbf{r}) = \sum_k^K D_{ak}\phi_{ak}(\mathbf{r})$. Leveraging EPT, MATRYOSHKA decouples ψ_a along an EPT-axis (k dimension) into multiple compute tiles. These tiles serve as parallelism-friendly units of the operator ψ_a at various stages, allowing them to be *permuted* in any order, *deconstructed* along the intrinsic EPT-axis, and *combined* to larger size for realigning an efficient parallelism on GPU elastically.

Through the manipulation of compute tiles along EPT-axes, MATRYOSHKA achieves the best of both worlds. It can achieve fine-grained, dense compute tile coverage of diverse operators without incurring layout transformation overhead, while promote highly-adapted parallelism with exhaustive GPU utilization.

A significant challenge to leverage Elastic Parallelism Transformation is how to capture and address the dynamic characteristics effectively. MATRYOSHKA tackles this challenge by the design of divergence-free Block Constructor, path-explicit Graph Compiler and thread-saturated Workload Allocator, which employs three EPT primitives (Permutation, Deconstruction, and Combination) to address dynamic diversity in QC systems.

The Block Constructor serves as the central orchestrator. It reformulates versatile data structures immune to dynamic inputs, and generates fine-grained blocks for efficient mapping to GPU SIMT architecture. The Graph Compiler, operating offline, transforms original scientific operators into computation graphs, applies domain-specific expert knowledge distilled from fine-tuned optimizations, and produces highly path-clear yet execution-efficient code through an automated compilation process. Meanwhile, the Workload Allocator combines compute tiles with varying operational intensity online. It rapidly evaluates workload schedule proposals, automatically allocating workloads to threads based on the strategic selection. This results in highly efficient parallelism for compute-intensive operations and facilitates fusion with neighboring memory-intensive operations.

We conduct a comprehensive evaluation of MATRYOSHKA on 13 representative QC systems (Chignolin, DNA, Crambin, etc.), demonstrating a remarkable speedup of up to 5.9x compared to state-of-the-art solutions. Notably, existing state-of-the-art approaches, while preserving original accuracy, have not surpassed the simulation scale of more than 1,000 atoms [32–36, 45–47]. MATRYOSHKA achieves a

Table 1. Symbol Table in § 2.2

Symbol	Meaning
ψ	Contracted basis function
ϕ	Fundamental basis function
a, **b**, **c**, **d**	Angular momentum of basis function
$\mathbf{a}_x$, $\mathbf{a}_y$, $\mathbf{a}_x$	Components of the angular momentum **a**
K, L, M, N	Degree of contraction
D_{ak}	Contraction coefficient
(**ab**|**cd**)	Contracted ERI
[**ab**|**cd**]	Fundamental ERI

breakthrough, maintaining original *ab initio* accuracy while simulating 11,259 atoms using a single GPU in one day (19.5 hrs).

Having traversed the content thus far, the rationale behind the initial posed question has been revealed: dynamic diversity. In response to this challenge, we introduce Matryoshka. Matryoshka signifies a paradigm shift, ushering in a new era-System4Science. This innovation lays the foundation for more efficient and scalable scientific computing on AI infrastructure in the context of quantum chemistry systems and beyond.

2 Background

2.1 Quantum Chemistry System

The cornerstone of quantum chemistry lies in the elucidation of atomic systems through the resolution of the Schrödinger equation [9]. However, the Schrödinger equation proves tractable solely for a finite set of elementary single-electron systems. To probe more intricate systems, the preponderance of computational quantum chemistry methodologies, including the widely employed Hartree-Fock (HF) and density-functional theory (DFT), resort to the self-consistent field (SCF) approximation. This approximation posits that each electron within the system traverses an average potential generated by all other electrons, forming the basis for constructing a QC system [38].

The resolution process of a QC system is completed through iterative steps. Commencing with an initial conjecture for the molecular orbitals (MOs), the process entails constructing the Fock matrix employing the extant MOs, which encapsulates contributions from kinetic energy, two-electron repulsion, and an effective potential. Subsequently, the electron density matrix is derived from the occupied MOs, and the total energy of the system is computed. Finally, an evaluation ensues to ascertain the convergence of electronic density and total energy to a stable solution. Should convergence remain elusive, the MOs are updated based on the prevailing electron density, and the entire process is reiterated until convergence is attained.

2.2 Electron Repulsion Integrals

The computation of two-electron repulsion integrals (ERIs) within a QC system constitutes a significant time consumption, accounting for approximately 95% of the overall computation time [39].

In quantum mechanics, a contracted basis function ψ characterizes an electron's spatial position, representing a linear combination of more fundamental basis functions ϕ. Each ϕ shares identical angular momentum, defined by the vector **a**. This vector consists of three integers (a_x, a_y, a_z), with their cumulative sum determining the angular momentum of the fundamental basis function. Mathematically, this relationship is expressed as:

$$\psi_a = \sum_{k}^{K} D_{ak}\phi_{ak}. \tag{1}$$

Here, the length of the linear combination, denoted as K, is termed the degree of contraction, and D_{ak} represents the contraction coefficients.

Angular momentum serves as a unique identifier for the contracted basis function. Assuming the angular momentum of the four involved contracted basis functions as **a**, **b**, **c**, **d**, the integral between two electrons can be denoted as (**ab**\|**cd**), which involves four contracted basis functions.

Given that the contracted basis function is a linear combination of fundamental basis functions, the integral (**ab**\|**cd**) can be expanded into the following form of basis function quadruple [22]:

$$(\mathbf{ab}|\mathbf{cd}) = \sum_{k}^{K} \sum_{l}^{L} \sum_{m}^{M} \sum_{n}^{N} D_{ak}D_{bl}D_{cm}D_{dn}[\mathbf{a}_k\mathbf{b}_l|\mathbf{c}_m\mathbf{d}_n] \tag{2}$$

To solve [**ab**\|**cd**], a recurrence relations algorithm is employed in ERI computations. The key idea is the calculation of [**ab**\|**cd**] with larger angular momentum can be derived from previously computed ones with smaller angular momentum. The algorithm involves two principles: Horizontal Recurrence Relation (HRR), which shifts angular momentum by transforming [**ab**\|**cd**] into [(**a** + **b**)**0**\|(**c** + **d**)**0**]; and Vertical Recurrence Relation (VRR), which reduces angular momentum at the **a** or **c**. Once all angular momenta at these 4 positions have been reduced to **0**, the integral can be computed analytically on [**00**\|**00**].

2.3 State-of-the-art

Recent efforts to accelerate research on large-scale QC systems with ab initio accuracy have pursued two primary approaches [6, 8, 32–36, 50–53]. One approach involves utilizing a greater number of CPU cores for distributed computing acceleration, a solution that consumes vast computational resources and is prohibitively expensive [36].The other focuses on efficiently porting QC systems to GPU architecture using High Performance Computing (HPC) techniques. However, as discussed earlier, the current approaches face significant

challenges due to dynamic diversity, which considerably limits the potential of GPUs[6, 8, 32–35, 50–53]. To the best of our knowledge, the industry-recognized state-of-the-art approaches maintaining original ab initio accuracy have yet to surpass the scale of simulating more than 1,000 atoms [8, 32–35]. Consequently, the design at the system level to efficiently support larger-scale QC systems with fewer computational resources emerges as a crucial and challenging endeavor.

3 Key Challenges and Insights

3.1 Dynamic Diversity

The existence of dynamic diversity presents a notable challenge, creating a significant misalignment between scientific systems and AI infrastructure with SIMT architecture. Over an extended period, the development and optimization of scientific systems have predominantly revolved around CPU architecture with serial computation logic. Consequently, these systems have not fully harnessed the advantages offered by cutting-edge AI infrastructure, such as GPUs renowned for their high performance. In QC systems, this limitation becomes particularly pronounced across three key dimensions.

Polymorphic data structures. Unlike the prevalent tensors with uniform data structures found in AI systems, a QC system is marked by a diverse array of data structures, including basis function, basis function pair, and basis function quadruple, to describe integrals of different input atoms. This diversity introduces a complex data organization and distinct execution instructions, leading to two key challenges.

Firstly, the involvement of multiple data structures escalates the memory cost, reaching up to the fourth power of basis functions. However, given the relatively constrained GPU memory compared to CPUs, accommodating large-scale QC systems becomes challenging. Secondly, the existence of distinct execution instructions for these diverse data structures often leads to warp divergence. This phenomenon emerges when threads within a warp deviate from a uniform execution path, resulting in only one thread being active in a warp. Consequently, this divergence triggers the serialization of execution and low throughput on the GPU.

Ambiguous computational paths. In AI systems, model structures are designed hierarchically, establishing explicit and well-defined computation paths through interconnected layers. In contrast, the computation paths in QC systems are highly diverse and non-unique, resulting in an intricate and continually-changing computational landscape for the entire system. This diversity introduces two significant challenges.

Firstly, the complex computation paths bring forth numerous high-precision floating-point operations, rapidly depleting limited register resources. This, in turn, leads to register spilling and subsequent performance degradation. Secondly, in existing work, all candidate computational paths are often manually coded and individually optimized, resulting in suboptimal efficiency and poor scalability when applied to
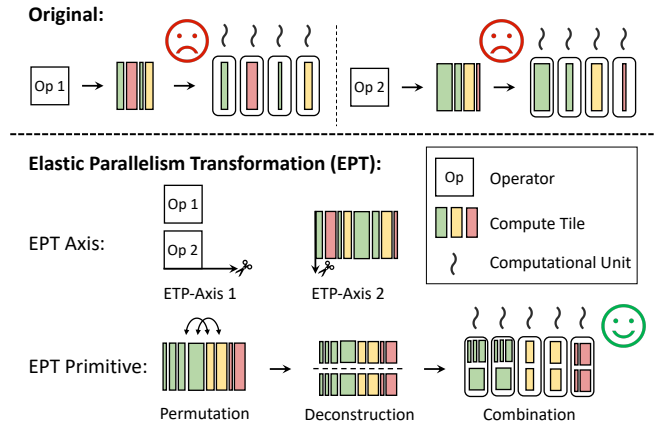

Figure 2. Illustration of Elastic Parallelism Transformation. Through the identification of EPT-axes and the application of EPT primitives, we can achieve fine-grained, dense compute tile coverage of diverse operators without incurring layout transformation overhead, while promote GPU utilization.

larger systems. Adding to the complexity, for a specific computational path, its cost is often dynamic. This dynamism arises because the computational workload is influenced by the input at each recursive entrance during runtime. This dynamic characteristic transforms each computation path into a highly serial and unpredictable process.

Variable operational intensity. GEMM operation serves as a fundamental and pivotal component within the domain of AI systems. Renowned for its high parallelizability, GEMM is particularly well-suited for hardware acceleration, notably on GPUs.

In stark contrast, the operations involved in QC systems deviate markedly from the prevalent homogeneity observed in AI systems. These operations display substantial variability across different types of atom, ranging from compute-intensive to memory-intensive. In particular, the locations where these operations are bound can dynamically shift. This implies that a specific operation may transition from being compute-intensive to memory-intensive based on varying inputs, resulting in notable idle periods that only become discernible at runtime.

3.2 Elastic Parallelism Transformation

As shown in Figure 2, Elastic Parallelism Transformation stands as the cornerstone insight employed by MATRYOSHKA to address dynamic diversity within QC systems, contributing to heightened GPU-efficient parallelism without incurring any data layout conversion costs.

In QC systems, each operator typically involves multiple reduction dimensions. MATRYOSHKA designates these dimensions as EPT-axes if and only if all computations on this axis are commutative and associative. The commutative and associative property ensures mathematically equivalent computation transformations on the original operators, allowing

for the random shuffling of compute tiles along the EPT-axis to achieve a more elastic parallel pattern in any order.

Along the EPT-axis, MATRYOSHKA introduces the design of compute tile—a fragment of an operator split in an adjustable size, which serves to compose a parallelism-friendly unit for efficient computation. For instance, in a QC system where contracted basis functions (ψ_a) are linear combinations of fundamental Gaussians (ϕ_{ak}), EPT recognizes that ψ_a has one EPT-axis (k dimensions), allowing the operator to be split into at most k independent ϕ_{ak} compute tiles along this axis.

Based on the introduced concepts of EPT-axis and compute tiles, EPT defines three pivotal transformation primitives. Each primitive works on a specific EPT-axis, orchestrating precise transformations within operators by manipulating nested compute tiles:

(1) **Permutation**: Along an EPT-axis, a portion or the entirety of compute tiles can rearrange their calculation order.

(2) **Deconstruction**: If a compute tile retains an EPT-axis, it can undergo further split into more sub-compute tiles.

(3) **Combination**: Sub-compute tiles with the same EPT-axis can also be nested into a larger compute tile.

Through the strategic application of these three primitives, dynamic diversity within QC systems can be effectively addressed by elastically aligning the nested compute tiles with the GPU architecture.

4 Design Overview

Figure 3 shows an overview of MATRYOSHKA, which consists of three core components: divergence-free Block Constructor, path-explicit Graph Compiler, and thread-saturated Workload Allocator.

The Block Constructor serves as the central orchestrator for the clustering of ERIs and the construction of ERI blocks with consistent execution instructions, which leverages the Permutation EPT primitive. It reformulates versatile data structures resilient to dynamic inputs and generates fine-grained ERI blocks tailored for efficient mapping onto GPU SIMT architecture (§ 5).

In an offline manner, the Graph Compiler produces highly path-clear yet execution-efficient code through an automated compilation process by utilizing the Deconstruction EPT primitive. The domain-specific expert knowledge is also distilled into the Compiler to guide an efficient path-search algorithm (§ 6).

Employing the Combination EPT primitive, the Workload Allocator dynamically schedules ERI workloads in response to their varying operational intensities online. Rapid evaluation of workload schedule proposals and automatic assignment of workloads to threads are key features, resulting
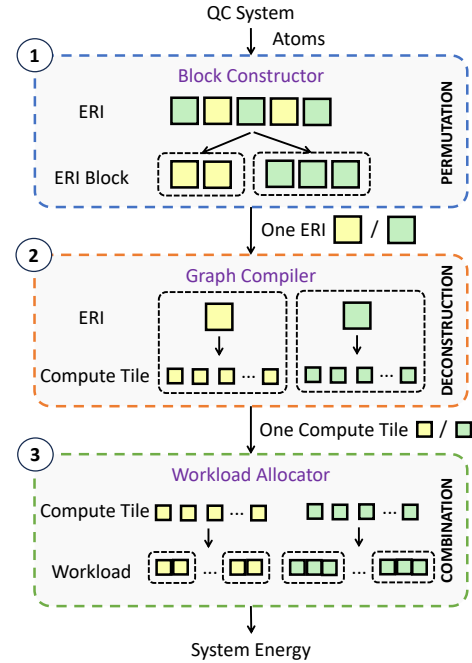

Figure 3. Overview of MATRYOSHKA.

in highly efficient parallelism for compute-intensive operations and facilitating fusion with adjacent memory-intensive operations (§ 7).

5 Block Constructor

In contrast to the uniform data structures, such as tensors, that facilitate batch processing in AI systems, QC systems present a series of distinctive challenges of Dynamic Diversity due to their reliance on polymorphic data structures.

The primary challenge is the accelerated depletion of computational and memory resources. As detailed in § 2.2, the ERI computation centers around a specific data structure—the basis function quadruple (**ab**|**cd**). In a QC system with N basis functions, the computation involves $O(N^4)$ ERIs of diverse types. As the size of QC systems scales up, the total number of ERIs reaches an overwhelming order of magnitude. Managing such a voluminous set is not only time-consuming but also places an immense demand on memory resources.

The second challenge is the issue of warp divergence during GPU computations. In QC systems, basis functions sharing the same angular momentum belong to the same class with consistent execution instructions. However, when a substantial number of basis functions are sent to the GPU for computation, threads within one warp may compute ERIs belonging to different classes. This divergence necessitates different threads executing different instructions, resulting in severe warp divergence.
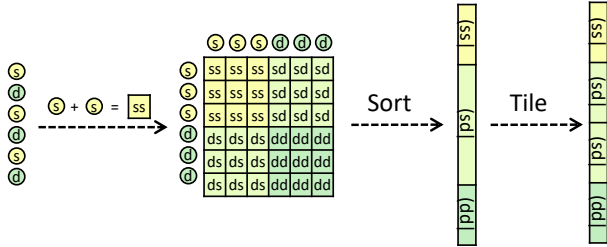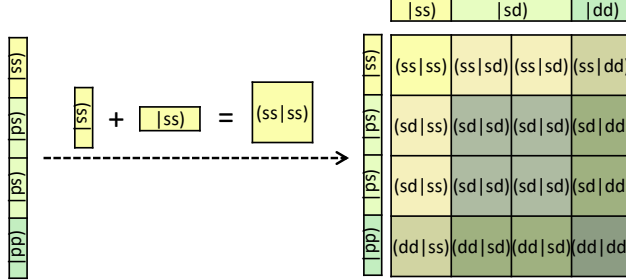
Stage 1: one → pair

Stage 2: pair → quadruple

Figure 4. Block Constructor's two-stage construction process. Circles, rectangles, and squares represent basis function individuals, pairs, and quadruples respectively.

Moreover, the distribution of these quadruples is inherently stochastic and undergoes dynamic changes with different QC systems. This dynamic nature implies that preprocessing these ERIs at compile time is unfeasible; instead, it must be conducted at runtime.

In Matryoshka we propose the Block Constructor which constructs basis function quadruple blocks in a streaming manner to handle the challenges of Dynamic Diversity caused by polymorphic data structures.

The Block Constructor is conceived based on two key insights, employing the initial EPT primitive—Permutation. Firstly, Matryoshka discerns that each basis function quadruple inherently possesses an EPT-axis of basis function pairs. Along this EPT-axis, $(\mathbf{ab}|\mathbf{cd})$ can be derived from the permutation of two pairs of basis functions, namely $(\mathbf{ab}|$ and $|\mathbf{cd})$. This realization indicates that there is no necessity to pre-construct all basis function quadruples; rather, constructing all basis function pairs and permuting them as required during computational proceedings suffices. This approach significantly reduces the memory cost of $(\mathbf{ab}|\mathbf{cd})$ from $O(N^4)$ to $O(N^2)$. Secondly, acknowledging the independence of each ERI computation within a basis function quadruple, Matryoshka identifies an additional EPT-axis of ERIs. Leveraging the Permutation primitive, the ERI computation can be permuted to group ERIs belonging to the same class together. This flexible adjustment allows for the batch processing of ERIs with consistent execution instructions within a single warp.

As shown in Figure 4, the Block Constructor takes two stages to complete the construction:

Stage 1: basis function → basis function pair. In this stage, the Block Constructor constructs an exhaustive pairing of all basis functions. Subsequently, these pairs are sorted in ascending order based on their angular momentum. To optimize data locality, we further segment these basis function pairs into multiple tiles. Notably, to guarantee that the resulting ERIs from the combination of these tiles fall within the same ERI class, the tiling process is exclusively performed within groups of basis function pairs belonging to the same class.

Stage 2: basis function pair → basis function quadruple. In this stage, the Block Constructor permutes the basis function pair tiles to form blocks of basis function quadruples. Should a tile contains M pairs, the resulting block will be the size of M^2, enabling the parallel computation of M^2 ERIs. These blocks of ERIs become the fundamental units of computation which share no data dependencies with each other. As a result, different ERI blocks can be assigned to different GPU streams for concurrent execution.

Through the design of the Block Constructor, Matryoshka successfully reduces the memory cost of $(\mathbf{ab}|\mathbf{cd})$ from $O(N^4)$ to $O(N^2)$, thereby freeing up a significant amount of memory space. Simultaneously, this design effectively alleviates the warp divergence problem, significantly enhancing the parallelism of ERI computations.

6 Graph Compiler

After enhancing the efficiency of batching various ERIs with the Block Constructor, this subsection will concentrate on the computation inside a single ERI. As introduced in § 2.2, recurrence relations algorithm is used to compute a single ERI $(\mathbf{ab}|\mathbf{cd})$ in QC systems. It first reduces the angular momentum on all four positions to $\mathbf{0}$ using these relations and then computes the ERI $(\mathbf{00}|\mathbf{00})$ analytically.

The dynamic diversity poses two main challenges in this procedure. First, the ERI computation involves diverse complex arithmetic operations with high-precision floating-point numbers, including division, modulo, exponentiation, etc. These operations necessitate a substantial allocation of registers per thread, often leading to severe register spilling. Second, a single ERI can give rise to diverse unclear computational paths during the recurrence process. This is because the order of applying the recurrence relations can change, yet all paths can yield the same result. Moreover, different computational paths have dynamic computational costs, which makes pinpointing the most efficient computational path a complex task. These computational costs are not determined until the recurrence process is actually executed.

In Matryoshka we introduce the Graph Compiler, which automatically generates an optimized kernel for each ERI. The Graph Compiler provides four stages to address the above challenges.
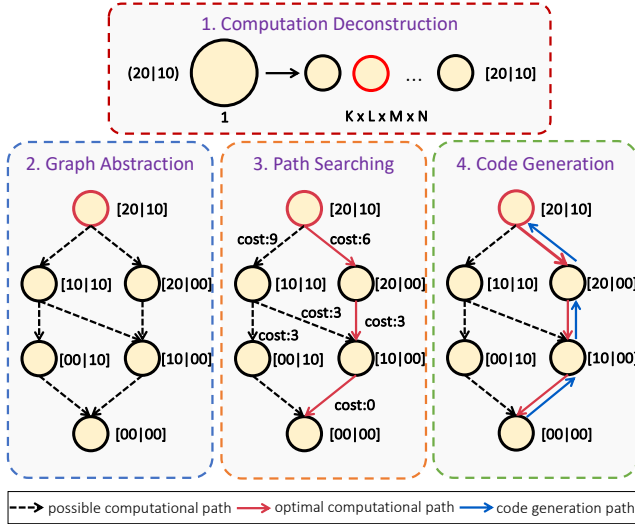
Figure 5. Graph Compiler's four-stage workflow on optimizing [**20**|**10**].

Stage 1: Computation Deconstruction. First, the Graph Compiler identifies the EPT axis existed in the dimension of (**ab**|**cd**) according to the principle of EPT. In Equation 2, (**ab**|**cd**) is calculated from the sum of all $[\mathbf{a}_k\mathbf{b}_l|\mathbf{c}_m\mathbf{d}_n]$. Since different $[\mathbf{a}_k\mathbf{b}_l|\mathbf{c}_m\mathbf{d}_n]$ share no data dependencies, they are naturally commutative and associative. Drawing on the Deconstruction EPT primitive, the Graph Compiler then deconstructs the computation of (**ab**|**cd**) into $K * L * M * N$ [**ab**|**cd**] (where K, L, M, and N represent the contraction degrees of the four contracted basis functions involved in the ERI computation). As a result, the basic computational workload for each thread is reduced by a factor of $K * L * M * N$, which markedly mitigates register spilling.

Stage 2: Graph Abstraction. The Graph Compiler first abstracts the recurrence process involved in computing each ERI into a *Directed Acyclic Graph* (DAG). Within this DAG, each node represents an intermediate result generated during the recursion process and each edge represents that one intermediate result can be derived from another. The direction of the edge reflects the direction of derivation. This abstraction simplifies the complex recurrence process into a format that is more accessible for computational handling, serving as the foundational basis for the following stages.

Stage 3: Path Searching. Drawing on the recurrence process of ERI computation, we observe that the computational cost of ERI mainly depends on two factors: (1) The length of the computational path; (2) The degree to reuse intermediate results. The first is because the longer the computational path is, the amount of computation naturally grows. The second is because thanks to the recurrence relations, many intermediate results in the previous recurrence steps actually can be reused in the following recurrence steps, which can greatly reduce the amount of computation.

Based on these observations, we conclude three guiding principles for the Graph Compiler to identify an optimized computational path:

(1) Strive for **0** angular momentum whenever possible.
(2) Prioritize the reuse of intermediate results from previous computations.
(3) Minimize the generation of new intermediate results.

The first principle is designed to terminate the recurrence process as soon as possible. The intent of the other two principles is to diminish computational efforts by preventing unnecessary recalculations and limiting the introduction of new calculations.

Based on these three principles, We design a greedy algorithm in Matryoshka to search for an optimized path on DAG. As Algorithm 1 shown, we express this algorithm as a function to determine the appropriate recurrence position at each recurrence step. From the Line 4, the algorithm iterates through each possible recurrence positions and greedily selects the one with the lowest computational cost. In Line 8, the computational cost is calculated based on the three guiding principles. Particularly, we use a hyperparameter λ to balance the first principle with the other two.

Algorithm 1 Greedy Path Searching

1: **function** FindOptimalPosition(*positions*, λ)
2:     $minC \leftarrow \infty$
3:     $optPos \leftarrow$ NULL
4:     **for** each *pos* in *positions* **do**
5:         $r \leftarrow$ count of reused results at *pos*
6:         $n \leftarrow$ count of new results at *pos*
7:         $a \leftarrow$ value of angular momentum at *pos*
8:         $cost \leftarrow (n - r) + \lambda * a$
9:         **if** $cost < minC$ **then**
10:             $minC \leftarrow cost$
11:             $optPos \leftarrow pos$
12:     **return** *optPos*

Stage 4: Code Generation. The Graph Compiler accepts the optimized computational path as input, interpreting it as a sequence of nodes arranged by topological order on a DAG. It then reverses this sequence and begins to generate computation codes starting from the base case of recurrence, which is [**00**|**00**]. After determining the value of ERI [**00**|**00**] analytically, the Graph Compiler proceeds to generate code for computing the nodes sequentially, working its way up to the desired target ERI [**ab**|**cd**].

It is worth noting that the Graph Compiler handles the entire task of kernel generation at compile time, resulting in no overhead during runtime. In addition to generating optimized kernels, the automation of the Graph Compiler equips Matryoshka with outstanding scalability for QC systems of any size. This eliminates the need for manual enumeration and case-by-case optimization across different ERI classes.
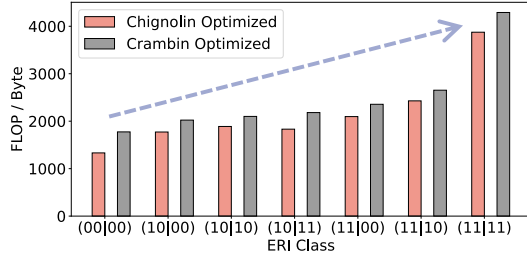
Figure 6. OP/B trends in Chignolin and Crambin. An upward trend is observed between OP/B and angular momentum.

7 Workload Allocator

While the Graph Compiler offers optimized kernels for [**ab**|**cd**] computations, it is noteworthy that different kernels for [**ab**|**cd**] exhibit varying operational intensities, spanning from compute-intensive to memory-intensive. In this context, we conduct an analysis of the compute-to-memory ratio (OP/B) for various [**ab**|**cd**] operations within Chignolin, a widely-used QC system composed of an artificial mini-protein. As depicted in Figure 6, an increase in angular momentum corresponds to a rise in the OP/B of [**ab**|**cd**]. This trend arises from the fact that in each ERI computation, the number of memory operations tends to remain stable, while the number of computational operations grows with increasing angular momentum. The multitude of ERI computations with diverse angular momenta leads to a frequent transition in operational intensity between memory-intensive and compute-intensive scenarios.

More complicated, however, is the dynamic nature of operational intensity for each kernel, where the bound position can vary with inputs. This dynamism implies a limited potential to pre-determine the bound bottleneck at compile time, often resulting in the generation of idle bubbles in the computing pipeline and leading to suboptimal GPU utilization at runtime.

As shown in Figure 7, MATRYOSHKA introduces the Workload Allocator—a key component designed to automatically allocate workloads to each thread efficiently, exploiting both the computational power and memory bandwidth of the GPU. Guided by the Combination EPT primitive, the Workload Allocator initially expands the [**ab**|**cd**] into larger dependency-free compute tiles with varying degrees of combination first. Then two allocation principles are presented to achieve an efficient overlapping of computation and memory access by managing these compute tiles.

(1) **For memory-intensive operations.** The Workload Allocator tends to allocate a larger workload per thread. Since computation is constrained by memory access, assigning more computational workloads can help mitigate idle bubbles caused by memory access latency.

(2) **For compute-intensive operations.** Adding more computational workload may not be advantageous, as it could already exceed each thread's computational capacity. In
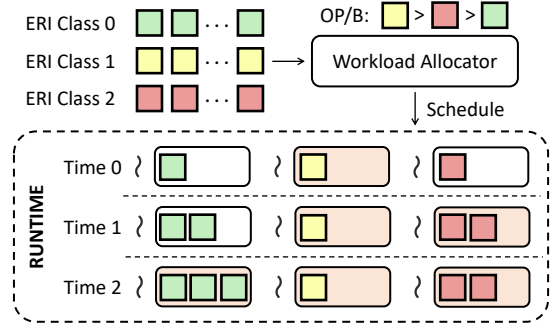

Figure 7. Workload Allocator matches threads to their ideal workload with varying operational intensity. At time 0, each thread is allocated a basic computational unit. Over time, the Workload Allocator progressively assigns additional tasks to the threads until they are saturated (orange-colored).

such cases, distributing the workload across more threads can help reduce computational latency. Simultaneously, the additional memory access associated with this approach can be mitigated by the underutilized memory bandwidth.

Based on the two allocation principles, we design an auto-tuning framework for efficiently scheduling workloads with varying operational intensities within the Workload Allocator. As shown in Algorithm 2, the workload of each thread is allocated as a [**ab**|**cd**] initially in Line 1-2. From the Line 3, the Workload Allocator systematically augments the workload for each ERI class by merging more basic computational units into a bigger one. The criterion for continuing this expansion is contingent upon observing improvements in the wall time of ERI calculations (Line 7-13). At last, the Workload Allocator converges, and each thread aligns with their suitable workload for diverse [**ab**|**cd**].

Algorithm 2 Auto-tuning Framework

1: **for** each *cls* in *ERI Classes* **do**
2:     *workload*[*cls*] ← [**ab**|**cd**]
3: *improved* ← true
4: **while** *improved* **do**
5:     *improved* ← false
6:     **for** each *cls* in *ERI Classes* **do**
7:         $t1$ ← TIME(*cls*)
8:         *workload*[*cls*] ← COMBINE(*workload*[*cls*])
9:         $t2$ ← TIME(*cls*)
10:        **if** $t2 < t1$ **then**
11:           *improved* ← true
12:        **else**
13:           *workload*[*cls*] ← REVERT(*workload*[*cls*])

Notably, the Workload Allocator seamlessly integrates with ongoing computations at runtime, introducing minimal additional overhead. Furthermore, the Workload Allocator efficiently executes reduction across diverse threads by enabling each thread to employ atomic operations directly. This strategic approach is based on the observation that update

positions from different threads exhibit relative sparsity, presenting an opportunity to minimize write conflicts among threads.

8 Evaluation

In this section, we conduct a comprehensive set of experiments to demonstrate the effectiveness of Matryoshka from various perspectives. Specifically, we begin by evaluating the correctness of Matryoshka with Elastic Parallelism Transformation against state-of-the-art approaches on five classical QC systems in § 8.2. Subsequently, in § 8.3 we assess each Matryoshka component individually: § 8.3.1 shows the performance breakdown of Matryoshka, § 8.3.2 examines the Block Constructor's role in handling warp divergence, § 8.3.3 validates the Graph Compiler's efficacy in mitigating register spilling and optimizing computational paths, and § 8.3.4 emphasizes the impact of the Workload Allocator in finely scheduling workloads to alleviate varying operational intensity. The scalability of Matryoshka is explored in § 8.4, and finally, in § 8.5, we evaluate the end-to-end simulation performance of Matryoshka on both A100 and A6000 GPUs across six larger representative QC systems. In summary, our results show that:

- The mathematical rigor of Matryoshka is firmly established through Elastic Parallelism Transformation. From a physicist's stringent viewpoint, the accuracy of Matryoshka can be deemed practically error-free (within 10^{-3} [24, 49]), with errors no more than 10^{-5} on five representative QC systems in comparison to state-of-the-art approaches.
- With Elastic Parallelism Transformation, Matryoshka addresses dynamic diversity effectively. Specifically, the Block Constructor shows a 4.7× improvement, the Graph Compiler a 2.3× improvement, and the Workload Allocator a 4.5× improvement across six QC systems on average.
- Matryoshka outperforms the previous state-of-art works and achieves up to 13.86×, 9.56×, and 4.82× speedup over Libint, PySCF, and QUICK respectively. Besides, it also achieves a breakthrough, maintaining original ab initio accuracy while simulating 11,259 atoms for 99 iterations using a single GPU in one day (19.5 hrs).

8.1 Experimental Setup

Platforms. We conduct our evaluation on two platforms. Platform A comprises an AMD EPYC 7V13 processor with 24 physical cores, and an Nvidia A100 80GB PCIe GPU. The A100 GPU could provide 1,935 GB/s memory bandwidth and 9.7 TFLOPs for FP64 operations. Platform B consists of an AMD EPYC 7742 processor with 64 physical cores, and four Nvidia A6000 48GB PCIe GPUs. Each A6000 GPU offers 768 GB/s memory bandwidth and 1.25 TFLOPs for FP64 operations. Both platforms employ GCC 9.4.0 and CUDA 12.2.

Table 2. Configuration for Matryoshka benchmarks.

Correctness		Performance		Scalability	
Name	Atoms	Name	Atoms	Name	Atoms
Water	3	Chinoglin	166	Water Cluster	300-11259
Benzene	12	DNA	566	GluAla Cluster	28-6658
Water-10	30	Crambin	642		
Methanol-7	42	Collagen	692		
C60	60	tRNA	1656		
		Pepsin	2797		

Benchmarks. We employ a wide array of representative QC systems to serve as benchmarks, as listed in Table 2. For correctness validation, we prioritize a diverse mix of systems, encompassing both organic and inorganic molecules. For performance evaluation, we pay particular attention to the computation complexity of these benchmarks. For scalability assessment, the size of the benchmarks is our primary concern. In all evaluations, we utilize the STO-3G basis set, which has relatively lower angular momentum for the sake of simplicity in presentation. It's important to note, however, that Matryoshka is compatible with any basis set.

State-of-the-arts. Due to the majority of work being closed-source or exhibiting significant differences in algorithm implementation, it is challenging to reproduce the results of related work on a one-to-one basis. Despite these obstacles, we have endeavored to collect and reproduce several industry-recognized state-of-the-art works, covering both CPU and GPU for the purpose of quantitative comparison.

We employ Libint [54] and PySCF [45–47] as representative examples of CPU-centric designs, which continues to dominate the majority of current work. Libint is one of the most efficient libraries for evaluating molecular integrals. PySCF is a widely utilized quantum chemistry Python package, featuring critical components that are highly optimized in C language. To showcase the emerging trend of GPU-centric work, we opt for QUICK, which stands as the current state-of-the-art for GPU implementations, supported by a series of studies [32–35].

8.2 Correctness Validation

We calculate the total energy for five classic QC systems using Matryoshka as well as other state-of-the-art approaches. In line with the common practice, we set the convergence threshold of electronic density as 10^{-6}.

Table 3 shows that the precision of Matryoshka aligns with other state-of-the-art approaches to within 10^{-5}, meeting the accuracy demands for most QC problems. For C60, a noticeable deviation in the result from QUICK is observed. However, the other approaches including Matryoshka maintain consistency.

We also visualize the Lowest Unoccupied Molecular Orbital (LUMO) of these QC systems using the results from Matryoshka, as shown in Figure 8. These visualizations are in agreement with the experimental data [7, 11], which reinforces the accuracy of Matryoshka.

Table 3. Comparison of precision with state-of-the-arts.

Molecules	Libint	PySCF	QUICK	Matryoshka
Water	**-74.9646977**	**-74.9646977**	**-74.9646977**	-74.9646977
Benzene	**-227.8909828**	**-227.8909828**	**-227.890982**7	-227.8909828
Water-10	**-749.6898790**	**-749.6898790**	**-749.6898790**	-749.6898793
Methanol-7	**-794.7735844**	**-794.7735844**	**-794.773584**3	-794.7735845
C60	**-2134.265202**4	**-2134.265202**02	**-2134.2**155554	-2134.2652033

* The underline highlights the identical digits among the state-of-the-art results.

(a) Water (b) Benzene (c) Water-10 (d) Methanol-7 (e) C60
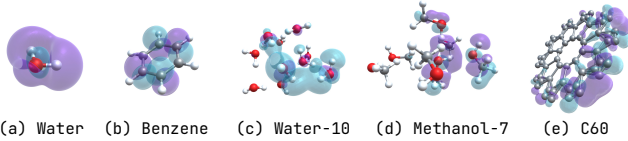
Figure 8. LUMO visualization for five classical QC systems.
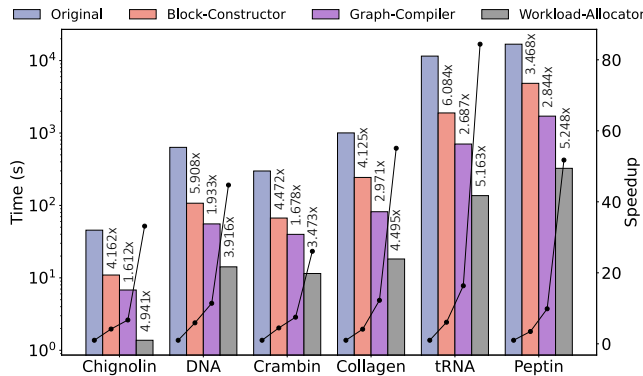
Figure 9. Performance breakdown of Matryoshka.

8.3 Breakdown Evaluation

8.3.1 Performance Breakdown. We first investigate how three core components of Matryoshka improve the performance by integrating them progressively. Figure 9 shows the performance breakdown of Matryoshka on six representative QC systems with great computational complexity. By clustering ERIs with consistent execution instructions, the Block Constructor mapping the ERI computation onto GPU SIMT architecture efficiently, thus achieving a 4.7× improvement on average. Next, the Graph Compiler deconstructs each ERI into multiple ERI compute tiles and generates the path-clear yet execution-efficient code for them, which brings a 2.3× improvement on average. Finally, the Workload Allocator is integrated which schedules ERI compute tiles to each thread dynamically, which provides a significant 4.5× improvement on average. Here Matryoshka has reached the 33.1×, 44.7×, 26.1×, 55.1×, 84.4×, 51.8× speedups accumulatively on six representative QC systems.

8.3.2 Block Constructor. The Block Constructor takes a two-fold approach to complete the construction, systematically organizing the data structures from individual basis functions to pairs, and ultimately to quadruples. This approach significantly reduces the memory cost of ERI computation from the $O(N^4)$ complexity of basis function quadruples to the more manageable $O(N^2)$ of pairs. We count the
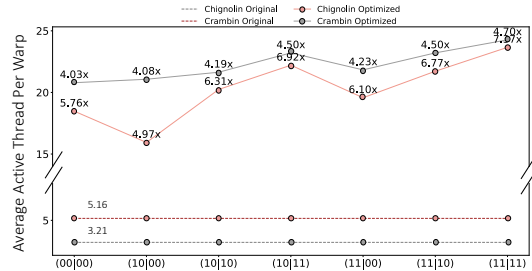
Figure 10. Comparative analysis of Average Active Threads per Warp in two representative QC systems. Matryoshka's values are classified by ERI classes, while the baseline values are depicted as horizontal lines, indicating the absence of any clustering process.

Table 4. The number of basis function pairs and quadruples in six representative QC systems.

QC System	Chignolin	DNA	Crambin	Collagen	tRNA	Peptin
Pair	24.0K	123.2K	156.8K	143.3K	381.6K	668.9K
Quadruple	577.1M	15.1G	24.5G	20.5G	145.6G	371.0G

number of basis function pairs and basis function quadruples in real-world QC systems respectively, as listed in Table 4. It shows that this nested construction slashes the memory demands by a factor of 10^3, transforming what was once an infeasible computational task into a feasible one.

We choose the metric *average active threads per warp* to study the effectiveness of the Block Constructor in reducing warp divergence, as shown in Figure 10. A higher number of average active threads per warp indicates less divergence within a single warp. Prior the integration of the Block Constructor, the average active threads per warp were as low as 3.21 and 5.16 for the Chignolin and Crambin respectively, indicating significant warp divergence issues. However, the Block Construct improves this situation markedly by leveraging the Permutation primitive of EPT, achieving increases of up to 7.37× and 4.70× across different ERI classes in the two QC systems.

8.3.3 Graph Compiler. The Graph Compiler is designed to address two key challenge inherent in a single ERI computation. The first challenge is the severe register spilling due to the complex arithmetic operations involved in ERI computations. To assess the effectiveness of the Graph Compiler in tackling this issue, we measure two specific metrics within the computations of Chignolin and Crambin.

The first metric we examine is the number of local memory requests, as shown in Figure 11a. This metric is a direct indicator of the severity of register spilling; when the register file is unable to hold all the data needed for computations, the excess data spills over into the local memory. The result shows that the Graph Compiler, by employing the Deconstruction primitive to partition the complete ERI into several compute tiles, efficiently reduces the local memory requests by up to 2.48× and 2.40× on two QC systems. This
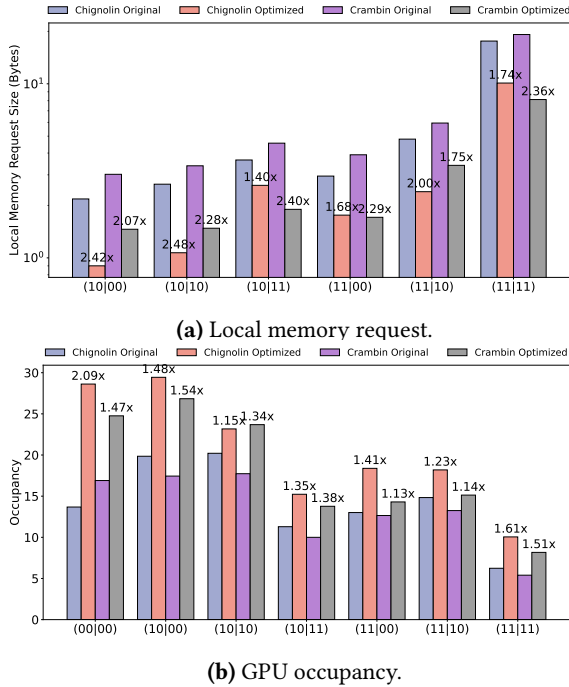
(a) Local memory request.

(b) GPU occupancy.

Figure 11. Comparative analysis of local memory request and GPU occupancy on two representative QC systems.

(a) Arithmetic Intensity.

(b) Compute Throughput.

Figure 12. Comparative analysis of compute throughput and arithmetic intensity on two representative QC systems.

substantial decrease in local memory requests implies that the Graph Compiler significantly alleviates the problem of register spilling, thereby enhancing the efficiency of the ERI computation process on these QC systems.

The second metric we select is the GPU Occupancy, an important metrics that reflects the GPU's ability to hide latencies. Register spilling can negatively impact GPU Occupancy as it increases the number of registers required per thread. Figure 11b shows that the Graph Compiler improves the occupancy of kernels that computes all ERI classes, from 1.23× to 2.09×, and from 1.13× to 1.54× on two QC systems respectively.

The second challenge is the ambiguous computational paths whose computational cost is dynamic. The Graph Compiler uses a greedy path searching algorithm to find an optimized path. Typically, it takes less than 10 seconds for finding the target computational path and generating the corresponding computational kernel. Take the Crambin as an example, the Graph Compiler uses just 2.57 seconds to complete the above tasks within a search space comprising approximately $O(10^5)$ potential computational paths. Notably, this kernel performs 1.42× faster than one derived from a randomly-generated computational path.

8.3.4 Workload Allocator Evaluation. Employing the Combination EPT primitive, the Workload Allocator schedules ERI compute tiles at runtime in response to their varying operational intensity. Specially, we measure two metrics to evaluate the effectiveness of the Workload Allocator in Chignolin and Crambin.
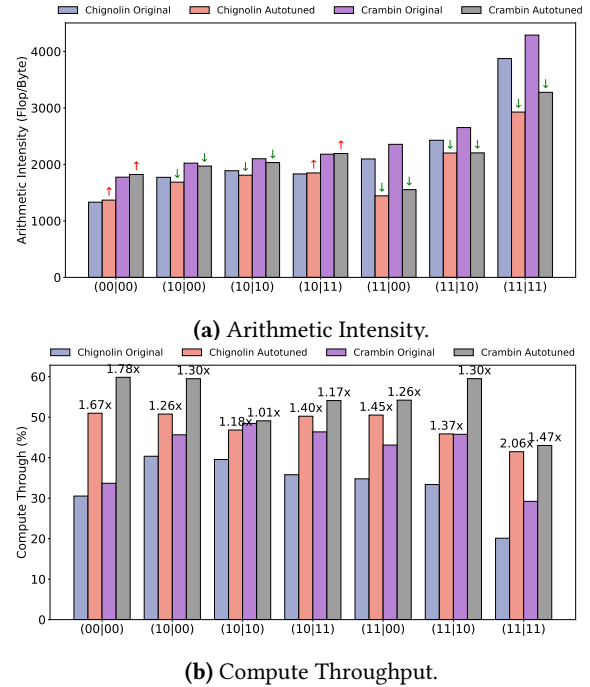
The first metrics is the arithmetic intensity, which measures the ratio of the FLOP to byte. A higher value indicates greater computational intensity for the kernel. As shown in Figure 12a, the Workload Allocator's tuning process adjusts the arithmetic intensity of different ERI classes to address their specific computational bottlenecks.

The other metrics is the compute throughput, which represents the system's computational performance, as shown in Figure 12b. The result shows that the Workload Allocator is effective; after tuning, compute throughput across all ERI classes has improved significantly, with increases of up to 2.06× and 1.78×, respectively. Additionally, since the auto-tuning algorithm is integrated with the computation process, there is no additional overhead measured.

8.4 Scalability Evaluation

Single GPU. We measure the execution time of MATRYOSHKA as it scales with two QC system sizes on single A100 GPU. Figure 13 shows that when the axes are scaled logarithmically, the execution time curve aligns closely with the ERI number curve. This result indicates that both execution time and ERI number grow exponentially with the atom count, while the performance of MATRYOSHKA remains stable despite problem size variations, ensuring consistent efficiency and adaptability across diverse scenarios. It is noteworthy that the largest scale of Water QC system accommodates up to 11,259 atoms and MATRYOSHKA finishes the simulation for
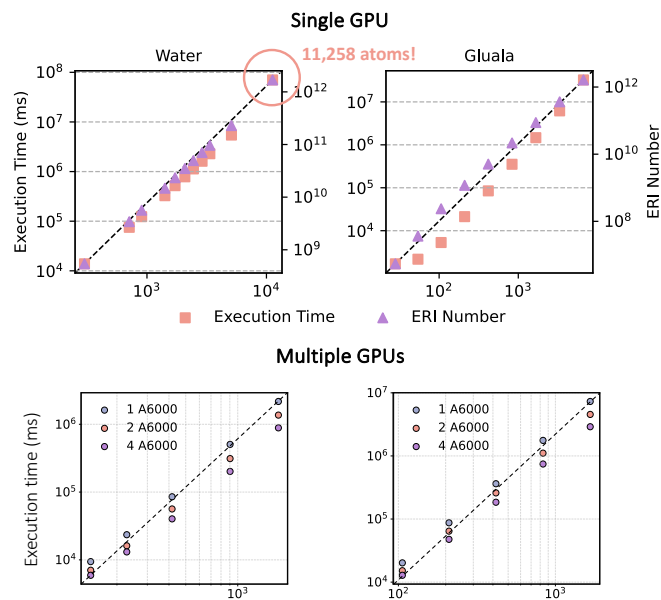
Figure 13. Scalability of Matryoshka on larger QC systems using both a single GPU and multiple GPUs.

Figure 14. Performance comparison of end-to-end execution time. Speedup is calculated relative to the execution time of Libint.

99 iterations in 19.5 hours, which establishes a novel benchmark especially when current studies are predominantly confined to systems with fewer than 1,000 atoms.

Multiple GPUs. We further evaluate the performance of Matryoshka in multi-GPU scenarios with four A6000 GPUs. Figure 13 presents the results of the weak scaling tests on Water and Gluala clusters with increasing molecular scales. With different problem sizes, the speedup of Matryoshka grows approximately proportional to the number of A6000 cards, which demonstrates that the performance of Matryoshka remains consistent regardless of problem size variations, even in multi-GPU scenarios.

8.5 End-to-End Evaluation

Finally, we compare the end-to-end execution time between Matryoshka and other state-of-the-art approaches. To eliminate the impact of different iteration counts, we limit the maximum number of iteration at 99. Figure 14 shows that by capitalizing on EPT, Matryoshka effectively tackles the challenge of dynamic diversity within QC systems and consistently surpasses the previous state-of-the-art approaches across all six representative QC systems on both A100 and A6000 GPUs. The performance of PySCF is slow and insufficient for producing results for large-sized molecules, such as tRNA and Peptin. Libint exhibits superior performance compared to PySCF, owing to more robust multi-thread support. Matryoshka achieves up to 13.35× and 13.86× speedup over Libint on A100 and A6000 GPUs, respectively. Leveraging GPU, QUICK surpasses CPU-based approaches in most benchmarks. However, it performs slower than Libint for the
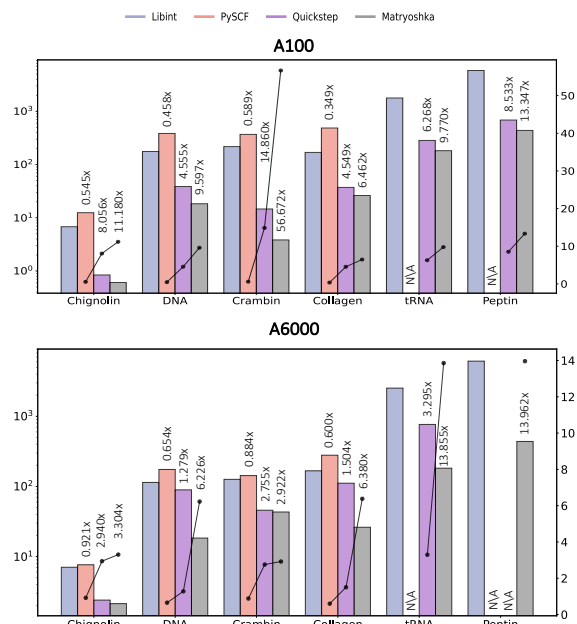
Crambin molecule, emphasizing multi-core CPUs' computational capability. QUICK fails to complete Pepsin calculations due to memory limitations. Matryoshka excels over QUICK across benchmarks, achieving up to 2.11× and 4.82× speedup on A100 and A6000 GPUs.

9 Related Work

CPU-centric. Traditional CPU-based approaches [4, 14, 25, 40] primarily focus on minimizing computation and can be divided into two strategies: reducing the number of iterations [4, 14] and limiting the number of integrals recalculated per iteration [14, 40]. Although these improvements significantly reduce computation, they do not consider computer hardware, resulting in suboptimal computational efficiency. Recently, several studies [1–3, 13, 19, 30, 36, 43, 44, 55] have concentrated on developing efficient implementations using distributed CPU nodes. Some works aim to design distributed parallel algorithms [1, 3, 19, 36, 55], while others address the prevalent load imbalance issues in the field [2, 13, 30, 43, 44]. However, the limitations in computational capability of CPUs hinder their further advancement.

GPU-centric. Yasudo [58] was among the first to implement GPU technology for the evaluation of ERIs, resulting in substantial speed improvements when compared to CPU-based computations. Following this, Ufimtsev and Martinez [50–53] conducted a series of studies that implemented a full Fock build on GPUs, wherein each thread is mapped to a different fundamental integral class. However, their works mainly depended on single-precision arithmetic, leading to

considerable computational errors. Asadchev et al. [6] introduced a Rys quadrature ERI method on GPUs that utilized double precision, but overlooked the data transfer costs between CPUs and GPUs. Miao et al. [32–35] and Barca et al. [8] employed the HGP method to reduce the computational expenses associated with the evaluation of ERIs, ultimately achieving remarkable speedups. Despite these advancements, these approaches did not account for the dynamic diversity inherent in QC systems, resulting in under-utilization.

10 Conclusion

This paper presents Matryoshka, a novel elastically-parallel technique for the efficient and accurate execution of quantum chemistry system with dynamic diversity on GPU. Leveraging Elastic Parallelism Transformation, Matryoshka realigns an efficient parallelism with GPU architecture with three key components: The Block Constructor formulates resilient data structures and constructs fine-grained, GPU-efficient compute blocks; the Graph Compiler generates code with clear computational path through an automated compilation process; the Workload Allocator achieves highly efficient parallelism for compute-intensive operations and facilitating automatic fusion with memory-intensive operations online.

References

[1] Yuri Alexeev, Ricky A. Kendall, and Mark S. Gordon. 2002. The distributed data SCF. *Computer Physics Communications* 143, 1 (2002), 69–82. https://doi.org/10.1016/S0010-4655(01)00439-8

[2] Yuri Alexeev, Sheri Mickelson, Sven Leyffer, Robert Jacob, and Anthony Craig. 2014. The Heuristic Static Load-Balancing Algorithm Applied to the Community Earth System Model. In *2014 IEEE International Parallel & Distributed Processing Symposium Workshops*. 1581–1590. https://doi.org/10.1109/IPDPSW.2014.177

[3] Yuri Alexeev, Michael W. Schmidt, Theresa L. Windus, and Mark S. Gordon. 2007. A parallel distributed data CPHF algorithm for analytic Hessians. *Journal of Computational Chemistry* 28, 10 (2007), 1685–1694. https://doi.org/10.1002/jcc.20633 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.20633

[4] J. Almlöf, K. Faegri Jr., and K. Korsell. 1982. Principles for a direct SCF approach to LICAO–MOab-initio calculations. *Journal of Computational Chemistry* 3, 3 (1982), 385–399. https://doi.org/10.1002/jcc.540030314 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.540030314

[5] Kazuto Ando, Rahul Bale, ChungGang Li, Satoshi Matsuoka, Keiji Onishi, and Makoto Tsubokura. 2021. Digital transformation of droplet/aerosol infection risk assessment realized on" Fugaku" for the fight against COVID-19. *arXiv preprint arXiv:2110.09769* (2021).

[6] Andrey Asadchev, Veerendra Allada, Jacob Felder, Brett M. Bode, Mark S. Gordon, and Theresa L. Windus. 2010. Uncontracted Rys Quadrature Implementation of up to G Functions on Graphical Processing Units. *Journal of Chemical Theory and Computation* 6, 3 (2010), 696–704. https://doi.org/10.1021/ct9005079 arXiv:https://doi.org/10.1021/ct9005079 PMID: 26613300.

[7] C.J. Ballhausen and H.B. Gray. 1965. *Molecular Orbital Theory: An Introductory Lecture Note and Reprint Volume*. W.A. Benjamin. https://books.google.com/books?id=HYw-AAAAIAAJ

[8] Giuseppe M. J. Barca, Jorge L. Galvez-Vallejo, David L. Poole, Alistair P. Rendell, and Mark S. Gordon. 2020. High-Performance, Graphics Processing Unit-Accelerated Fock Build Algorithm. *Journal of Chemical Theory and Computation* 16, 12 (2020), 7232–7238. https://doi.org/10.1021/acs.jctc.0c00768 arXiv:https://doi.org/10.1021/acs.jctc.0c00768 PMID: 33206515.

[9] Eric Cances, Mireille Defranceschi, Werner Kutzelnigg, Claude Le Bris, and Yvon Maday. 2003. Computational quantum chemistry: a primer. *Handbook of numerical analysis* 10 (2003), 3–270.

[10] Yudong Cao, Jonathan Romero, Jonathan P Olson, Matthias Degroote, Peter D Johnson, Mária Kieferová, Ian D Kivlichan, Tim Menke, Borja Peropadre, Nicolas PD Sawaya, et al. 2019. Quantum chemistry in the age of quantum computing. *Chemical reviews* 119, 19 (2019), 10856–10915.

[11] F.A. Carey and R.J. Sundberg. 2000. *Advanced Organic Chemistry: Part A: Structure and Mechanisms*. Kluwer Academic/Plenum Pub. https://books.google.com/books?id=z3jXKOYuqQAC

[12] Paolo Carloni, Ursula Rothlisberger, and Michele Parrinello. 2002. The Role and Perspective of Ab Initio Molecular Dynamics in the Study of Biological Systems. *Accounts of Chemical Research* 35, 6 (2002), 455–464. https://doi.org/10.1021/ar010018u arXiv:https://doi.org/10.1021/ar010018u PMID: 12069631.

[13] Edmond Chow, Xing Liu, Sanchit Misra, Marat Dukhan, Mikhail Smelyanskiy, Jeff Hammond, Yunfei Du, Xiang-Ke Liao, and Pradeep Dubey. 2015. Scaling up Hartree-Fock calculations on Tianhe-2. *International Journal of High Performance Computing Applications* 30 (07 2015). https://doi.org/10.1177/1094342015592960

[14] Dieter Cremer and JüRgen Gauss. 1986. An unconventional scf method for calculations on large molecules. *Journal of Computational Chemistry* 7, 3 (1986), 274–282. https://doi.org/10.1002/jcc.540070305 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.540070305

[15] Sambit Das, Bikash Kanungo, Vishal Subramanian, Gourab Panigrahi, Phani Motamarri, David Rogers, Paul Zimmerman, and Vikram Gavini. 2023. Large-Scale Materials Modeling at Quantum Accuracy: Ab Initio Simulations of Quasicrystals and Interacting Extended Defects in Metallic Alloys *(SC '23)*. Association for Computing Machinery, New York, NY, USA, Article 1, 12 pages. https://doi.org/10.1145/3581784.3627037

[16] Kaushik Datta, Mark Murphy, Vasily Volkov, Samuel Williams, Jonathan Carter, Leonid Oliker, David Patterson, John Shalf, and Katherine Yelick. 2008. Stencil computation optimization and auto-tuning on state-of-the-art multicore architectures. In *SC'08: Proceedings of the 2008 ACM/IEEE conference on Supercomputing*. IEEE, 1–12.

[17] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).

[18] L. Fedeli, A. Huebl, F. Boillod-Cerneux, T. Clark, K. Gott, C. Hillairet, S. Jaure, A. Leblanc, R. Lehe, A. Myers, C. Piechurski, M. Sato, N. Zaim, W. Zhang, J. Vay, and H. Vincenti. 2022. Pushing the Frontier in the Design of Laser-Based Electron Accelerators with Groundbreaking Mesh-Refined Particle-In-Cell Simulations on Exascale-Class Supercomputers. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Computer Society, Los Alamitos, CA, USA, 1–12. https://doi.org/10.1109/SC41404.2022.00008

[19] Graham D. Fletcher, Michael W. Schmidt, Brett M. Bode, and Mark S. Gordon. 2000. The Distributed Data Interface in GAMESS. *Computer Physics Communications* 128, 1 (2000), 190–200. https://doi.org/10.1016/S0010-4655(00)00073-4

[20] Annarita Giani and Zachary Eldredge. 2021. Quantum computing opportunities in renewable energy. *SN Computer Science* 2, 5 (2021), 393.

[21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[22] Martin Head-Gordon and John A Pople. 1988. A method for two-electron Gaussian integral and integral derivative evaluation using recurrence relations. *The Journal of chemical physics* 89, 9 (1988),

5777–5786.

[23] Trygve Helgaker, Wim Klopper, and David P Tew. 2008. Quantitative quantum chemistry. *Molecular Physics* 106, 16-18 (2008), 2107–2143.

[24] Trygve Helgaker, Wim Klopper, and David P Tew. 2008. Quantitative quantum chemistry. *Molecular Physics* 106, 16-18 (2008), 2107–2143.

[25] Marco Häser and Reinhart Ahlrichs. 1989. Improvements on the direct SCF method. *Journal of Computational Chemistry* 10, 1 (1989), 104–111. https://doi.org/10.1002/jcc.540100111 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.540100111

[26] Keith H Johnson. 1975. Quantum chemistry. *Annual review of physical chemistry* 26, 1 (1975), 39–57.

[27] John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. 2021. Highly accurate protein structure prediction with AlphaFold. *Nature* 596, 7873 (2021), 583–589.

[28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012).

[29] Kun Li, Liang Yuan, Yunquan Zhang, Yue Yue, and Hang Cao. 2022. An efficient vectorization scheme for stencil computation. In *2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 650–660.

[30] Xing Liu, Aftab Patel, and Edmond Chow. 2014. A New Scalable Parallel Algorithm for Fock Matrix Construction. In *2014 IEEE 28th International Parallel and Distributed Processing Symposium*. 902–914. https://doi.org/10.1109/IPDPS.2014.97

[31] Yong (Alexander) Liu, Xin (Lucy) Liu, Fang (Nancy) Li, Haohuan Fu, Yuling Yang, Jiawei Song, Pengpeng Zhao, Zhen Wang, Dajia Peng, Huarong Chen, Chu Guo, Heliang Huang, Wenzhao Wu, and Dexun Chen. 2021. Closing the "Quantum Supremacy" Gap: Achieving Real-Time Simulation of a Random Quantum Circuit Using a New Sunway Supercomputer. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (St. Louis, Missouri) *(SC '21)*. Association for Computing Machinery, New York, NY, USA, Article 3, 12 pages. https://doi.org/10.1145/3458817.3487399

[32] Madushanka Manathunga, Hasan Metin Aktulga, Andreas W. Götz, and Kenneth M. Jr. Merz. 2023. Quantum Mechanics/Molecular Mechanics Simulations on NVIDIA and AMD Graphics Processing Units. *Journal of Chemical Information and Modeling* 63, 3 (2023), 711–717. https://doi.org/10.1021/acs.jcim.2c01505 arXiv:https://doi.org/10.1021/acs.jcim.2c01505 PMID: 36720086.

[33] Madushanka Manathunga, Chi Jin, Vinícius Wilian D. Cruzeiro, Yipu Miao, Dawei Mu, Kamesh Arumugam, Kristopher Keipert, Hasan Metin Aktulga, Kenneth M. Jr. Merz, and Andreas W. Götz. 2021. Harnessing the Power of Multi-GPU Acceleration into the Quantum Interaction Computational Kernel Program. *Journal of Chemical Theory and Computation* 17, 7 (2021), 3955–3966. https://doi.org/10.1021/acs.jctc.1c00145 arXiv:https://doi.org/10.1021/acs.jctc.1c00145 PMID: 34062061.

[34] Yipu Miao and Kenneth M. Jr. Merz. 2013. Acceleration of Electron Repulsion Integral Evaluation on Graphics Processing Units via Use of Recurrence Relations. *Journal of Chemical Theory and Computation* 9, 2 (2013), 965–976. https://doi.org/10.1021/ct300754n arXiv:https://doi.org/10.1021/ct300754n PMID: 26588740.

[35] Yipu Miao and Kenneth M. Jr. Merz. 2015. Acceleration of High Angular Momentum Electron Repulsion Integrals and Integral Derivatives on Graphics Processing Units. *Journal of Chemical Theory and Computation* 11, 4 (2015), 1449–1462. https://doi.org/10.1021/ct500984t arXiv:https://doi.org/10.1021/ct500984t PMID: 26574356.

[36] Vladimir Mironov, Alexander Moskovsky, Michael D'Mello, and Yuri Alexeev. 2019. An efficient MPI/OpenMP parallelization of the Hartree–Fock–Roothaan method for the first generation of Intel® Xeon Phi™ processor architecture. *The International Journal of High Performance Computing Applications*

33, 1 (2019), 212–224. https://doi.org/10.1177/1094342017732628 arXiv:https://doi.org/10.1177/1094342017732628

[37] B Mohr, P Raghavan, H Fu, et al. 2017. 18.9-Pflops nonlinear earthquake simulation on Sunway TaihuLight. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–12.

[38] Robert G Parr and Weitao Yang. 1995. Density-functional theory of the electronic structure of molecules. *Annual review of physical chemistry* 46, 1 (1995), 701–728.

[39] John A Pople and Warren J Hehre. 1978. Computation of electron repulsion integrals involving contracted Gaussian basis functions. *J. Comput. Phys.* 27, 2 (1978), 161–168.

[40] Péter Pulay. 1980. Convergence acceleration of iterative sequences. the case of scf iteration. *Chemical Physics Letters* 73, 2 (1980), 393–398. https://doi.org/10.1016/0009-2614(80)80396-4

[41] Jakob Schiøtz, Francesco D Di Tolla, and Karsten W Jacobsen. 1998. Softening of nanocrystalline metals at very small grain sizes. *Nature* 391, 6667 (1998), 561–563.

[42] Jakob Schiøtz and Karsten W Jacobsen. 2003. A maximum in the strength of nanocrystalline copper. *Science* 301, 5638 (2003), 1357–1359.

[43] Hongzhang Shan, Brian Austin, Wibe De Jong, Leonid Oliker, N. J. Wright, and Edoardo Apra. 2014. Performance Tuning of Fock Matrix and Two-Electron Integral Calculations for NWChem on Leading HPC Platforms. In *High Performance Computing Systems. Performance Modeling, Benchmarking and Simulation*, Stephen A. Jarvis, Steven A. Wright, and Simon D. Hammond (Eds.). Springer International Publishing, Cham, 261–280.

[44] Hongzhang Shan, Samuel Williams, Wibe de Jong, and Leonid Oliker. 2015. Thread-Level Parallelization and Optimization of NWChem for the Intel MIC Architecture. In *Proceedings of the Sixth International Workshop on Programming Models and Applications for Multicores and Manycores* (San Francisco, California) *(PMAM '15)*. Association for Computing Machinery, New York, NY, USA, 58–67. https://doi.org/10.1145/2712386.2712391

[45] Qiming Sun. 2015. Libcint: An efficient general integral library for Gaussian basis functions. *Journal of Computational Chemistry* 36, 22 (2015), 1664–1671. https://doi.org/10.1002/jcc.23981 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.23981

[46] Qiming Sun, Timothy C. Berkelbach, Nick S. Blunt, George H. Booth, Sheng Guo, Zhendong Li, Junzi Liu, James D. McClain, Elvira R. Sayfutyarova, Sandeep Sharma, Sebastian Wouters, and Garnet Kin-Lic Chan. 2018. PySCF: the Python-based simulations of chemistry framework. *WIREs Computational Molecular Science* 8, 1 (2018), e1340. https://doi.org/10.1002/wcms.1340 arXiv:https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/wcms.1340

[47] Qiming Sun, Xing Zhang, Samragni Banerjee, Peng Bao, Marc Barbry, Nick S. Blunt, Nikolay A. Bogdanov, George H. Booth, Jia Chen, Zhi-Hao Cui, Janus J. Eriksen, Yang Gao, Sheng Guo, Jan Hermann, Matthew R. Hermes, Kevin Koh, Peter Koval, Susi Lehtola, Zhendong Li, Junzi Liu, Narbe Mardirossian, James D. McClain, Mario Motta, Bastien Mussard, Hung Q. Pham, Artem Pulkin, Wirawan Purwanto, Paul J. Robinson, Enrico Ronca, Elvira R. Sayfutyarova, Maximilian Scheurer, Henry F. Schurkus, James E. T. Smith, Chong Sun, Shi-Ning Sun, Shiv Upadhyay, Lucas K. Wagner, Xiao Wang, Alec White, James Daniel Whitfield, Mark J. Williamson, Sebastian Wouters, Jun Yang, Jason M. Yu, Tianyu Zhu, Timothy C. Berkelbach, Sandeep Sharma, Alexander Yu. Sokolov, and Garnet Kin-Lic Chan. 2020. Recent developments in the PySCF program package. *The Journal of Chemical Physics* 153, 2 (07 2020), 024109. https://doi.org/10.1063/5.0006074 arXiv:https://pubs.aip.org/aip/jcp/article-pdf/doi/10.1063/5.0006074/16722275/024109_1_online.pdf

[48] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric

Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).

[49] Ivan S Ufimtsev and Todd J Martinez. 2008. Quantum chemistry on graphical processing units. 1. Strategies for two-electron integral evaluation. *Journal of Chemical Theory and Computation* 4, 2 (2008), 222–231.

[50] Ivan S. Ufimtsev and Todd J. Martinez. 2009. Quantum Chemistry on Graphical Processing Units. 2. Direct Self-Consistent-Field Implementation. *Journal of Chemical Theory and Computation* 5, 4 (2009), 1004–1015. https://doi.org/10.1021/ct800526s arXiv:https://doi.org/10.1021/ct800526s PMID: 26609609.

[51] Ivan S. Ufimtsev and Todd J. Martinez. 2009. Quantum Chemistry on Graphical Processing Units. 3. Analytical Energy Gradients, Geometry Optimization, and First Principles Molecular Dynamics. *Journal of Chemical Theory and Computation* 5, 10 (2009), 2619–2628. https://doi.org/10.1021/ct9003004 arXiv:https://doi.org/10.1021/ct9003004 PMID: 26631777.

[52] Ivan S. Ufimtsev and Todd J. Martínez. 2008. Graphical Processing Units for Quantum Chemistry. *Computing in Science & Engineering* 10, 6 (Nov 2008), 26–34. https://doi.org/10.1109/MCSE.2008.148

[53] Ivan S. Ufimtsev and Todd J. Martínez. 2008. Quantum Chemistry on Graphical Processing Units. 1. Strategies for Two-Electron Integral Evaluation. *Journal of Chemical Theory and Computation* 4, 2 (2008), 222–231. https://doi.org/10.1021/ct700268q

arXiv:https://doi.org/10.1021/ct700268q PMID: 26620654.

[54] E. F. Valeev. 2022. Libint: A library for the evaluation of molecular integrals of many-body operators over Gaussian functions. http://libint.valeyev.net/. version 2.8.0.

[55] M. Valiev, E.J. Bylaska, N. Govind, K. Kowalski, T.P. Straatsma, H.J.J. Van Dam, D. Wang, J. Nieplocha, E. Apra, T.L. Windus, and W.A. de Jong. 2010. NWChem: A comprehensive and scalable open-source solution for large scale molecular simulations. *Computer Physics Communications* 181, 9 (2010), 1477–1489. https://doi.org/10.1016/j.cpc.2010.04.018

[56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).

[57] Qinglin Wang, Dongsheng Li, Xiandong Huang, Siqi Shen, Songzhu Mei, and Jie Liu. 2020. Optimizing FFT-based convolution on ARMv8 multi-core CPUs. In *European Conference on Parallel Processing*. Springer, 248–262.

[58] Koji Yasuda. 2008. Two-electron integral evaluation on the graphics processor unit. *Journal of Computational Chemistry* 29, 3 (2008), 334–342. https://doi.org/10.1002/jcc.20779 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.20779