

# Lawrence Berkeley National Laboratory

## LBL Publications

### Title

Distributed memory, GPU accelerated Fock construction for hybrid, Gaussian basis density functional theory

### Permalink

<https://escholarship.org/uc/item/5h87r4f5>

### Journal

The Journal of Chemical Physics, 158(23)

### ISSN

0021-9606

### Authors

Williams-Young, David B

Asadchev, Andrey

Popovici, Doru Thom

et al.

### Publication Date

2023-06-21

### DOI

10.1063/5.0151070

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

## Distributed Memory, GPU Accelerated Fock Construction for Hybrid, Gaussian Basis Density Functional Theory

David B. Williams-Young,<sup>1, a)</sup> Andrey Asadchev,<sup>2</sup> Doru Thom Popovici,<sup>1</sup> David Clark,<sup>3</sup> Jonathan Waldrop,<sup>4</sup> Theresa Windus,<sup>4, 5</sup> Edward F. Valeev,<sup>2</sup> and Wibe A. de Jong<sup>1</sup>

<sup>1)</sup>*Applied Mathematics and Computational Research Division,  
Lawrence Berkeley National Laboratory, Berkeley, CA 94720*

<sup>2)</sup>*Department of Chemistry, Virginia Tech, Blacksburg,  
VA 24061*

<sup>3)</sup>*NVIDIA Corporation, Santa Clara, CA 95051*

<sup>4)</sup>*Chemical and Biological Sciences Division, Ames National Laboratory, Ames,  
IA 50011*

<sup>5)</sup>*Department of Chemistry, Iowa State University, Ames,  
IA 50011*

(Dated: 28 March 2023)

With the growing reliance of modern supercomputers on accelerator-based architectures such as GPUs, the development and optimization of electronic structure methods to exploit these massively parallel resources has become a recent priority. While significant strides have been made in the development of GPU accelerated, distributed memory algorithms for many-body (e.g. coupled-cluster) and spectral single-body (e.g. planewave, real-space and finite-element density functional theory [DFT]), the vast majority of GPU-accelerated Gaussian atomic orbital methods have focused on shared memory systems with only a handful of examples pursuing massive parallelism on distributed memory GPU architectures. In the present work, we present a set of distributed memory algorithms for the evaluation of the Coulomb and exact-exchange matrices for hybrid Kohn-Sham DFT with Gaussian basis sets via direct density-fitted (DF-J-Engine) and seminumerical (sn-K) methods, respectively. The absolute performance and strong scalability of the developed methods are demonstrated on systems ranging from a few hundred to over one thousand atoms using up to 128 NVIDIA A100 GPUs on the Perlmutter supercomputer.

---

<sup>a)</sup>dbwy@lbl.gov

## I. INTRODUCTION

Since its inception, quantum chemistry has relied on its ability to quickly adapt to an ever evolving landscape of computer architectures to enable the next generation of scientific applications. A particular emphasis has been placed on the leverage of distributed memory parallelism to enable *ab initio* simulations of large molecular systems on the world’s largest supercomputers<sup>1-4</sup>. The past two decades have been no different, with an enormous research effort having been afforded to targeting the latest introduction into the high-performance computing (HPC) ecosystem: graphics processing units (GPU)<sup>4-6</sup>. The use of GPUs in quantum chemistry is relatively long standing, with applications ranging from single-body methods such as Hartree-Fock (HF) and density functional (DFT) theories<sup>7-39</sup>, to many-body methods<sup>3</sup> such as coupled-cluster theory<sup>40-49</sup>, many-body perturbation theory<sup>50-56</sup>, and configuration interaction and complete active space theories<sup>57-61</sup> to name a few. Despite significant advances in the development of GPU-accelerated quantum chemistry software, these efforts are not yet as mature as their central processing unit (CPU) counterparts, and growing requirements for the desired scale and accuracy of *ab initio* simulations require further development in this area. As such, the pursuance of improved, GPU-accelerated quantum chemistry methods capable of leveraging the latest advances in modern HPC remains an active area of research.

With the growing reliance of modern exascale HPC systems on accelerator architectures,<sup>62,63</sup> recent years have seen the deployment of GPUs in distributed memory systems. This deployment has in turn been accompanied by its own unique set of optimization challenges. In particular, the increased local processing rate of GPU-accelerated compute nodes has exposed bottlenecks involving communication and load imbalance which were less apparent on the massively parallel CPU systems of years past<sup>63-65</sup>. As necessitated by the predominant availability of consumer-grade, gaming-oriented GPU hardware, early development of GPU-accelerated quantum chemistry methods focused on shared memory systems where host and device memory spaces, while disjoint, are accessible to one another without the need for communication across distributed memory networks. Due to their relative computational cost, many-body methods (e.g., coupled cluster theory<sup>46-49</sup>) and spectral single-body methods (e.g., planewave<sup>27-32</sup>, real-space<sup>33,34</sup>, finite-element<sup>35</sup>, and wavelet<sup>36</sup> discretizations of DFT) were among the first to be successfully ported to distributed memory GPU architec-

tures. The success of these developments has been, in large part, enabled by the prevalent use of GPU optimized libraries for common operations, such as (multi-)linear algebra<sup>43,66-69</sup> and fast Fourier transforms (FFT)<sup>70-72</sup>, in these applications. On the other hand, atomic orbital (AO) single-body methods, such as Gaussian, Slater, and numerical atomic orbital HF and DFT, present a unique challenge for GPU optimization due to their significant dependence on domain-specific kernels (e.g. AO integrals and highly-specialized numerical integration techniques) which exhibit less-regular compute patterns than those common operations aforementioned. For Gaussian basis DFT in particular, kernels involving manipulation of the AO electron repulsion integral (ERI) tensor to form the Coulomb and exchange matrices are among the most cumbersome for GPU architectures. As such, the vast majority of both shared<sup>7,8,10,13,16,20,21,25,73-76</sup> and distributed<sup>11,12,24,77</sup> memory Gaussian DFT GPU efforts have centered around optimization of these terms.

The primary innovations of the present work concern the evaluation of Coulomb and exact-exchange matrices. Although the naive (based on 2-body 4-center ERI) approach for evaluation of these contributions is simple and thus used by majority of implementations on distributed heterogeneous platforms,<sup>11,12,23</sup> their steep asymptotic scaling complicates their application to large molecular systems. For the Coulomb matrix, the use of 4-center integrals results in suboptimal  $\mathcal{O}(N^2)$  cost, whereas fast approaches with  $\mathcal{O}(N)$  for the Coulomb potential evaluation are well known<sup>78,79</sup>. More importantly, even if combined with the fast  $\mathcal{O}(N)$  treatment of the long-range contributions to the potential, the use of 4-center integrals for the near-field contributions is still prohibitively expensive. The solution is well known also: clever optimizations (like early integral digestion<sup>80-83</sup>) and/or numerical approximations like density fitting<sup>84,85</sup> can dramatically reduce the cost of Coulomb matrix evaluation such that even with naive  $\mathcal{O}(N^2)$  evaluation its cost is dwarfed by the cost of the exact exchange.

While with proper screening the exact exchange evaluation is  $\mathcal{O}(N)$  in system size, the use of 4-center integrals, again, results in suboptimally-high prefactor. Thus, there has been a recent resurgence of interest in the development of numerical methods for the evaluation of exact exchange for molecular systems. In these methods, one of the two ERI coordinate integrations is replaced by a numerical integration. The general concept for this approach has been around for over 30 years, beginning with the pseudospectral method on Friesner and co-workers<sup>86,87</sup> in the early 1980s. This class of techniques was revisited in the early

2000s as the chain-of-spheres exchange (COSX) method of Neese, *et al.*,<sup>88</sup> and more recently as the seminumerical exchange (sn-K) method of Laqua, *et al.*<sup>10</sup> The primary differences between these approaches are in their consideration of numerical sparsity, spatial locality, and in their attempts to reduce the required size of the numerical quadrature to attain a particular accuracy. We will adopt the sn-K nomenclature in the following. The sn-K algorithm is particularly attractive for GPU hardware as there is a natural expression of vectorization in the numerical quadrature which is superior to that of analytical integral methods. Significant progress has been made in the development of GPU sn-K algorithms for shared memory systems including multiple GPUs per node.<sup>10,13,14,89</sup> These methods have demonstrated a sizable performance improvement over analytical integral methods on GPU architectures, but have yet to be explored in the context of distributed memory parallelism.

In addition to the Coulomb and exact exchange terms, hybrid AO DFT methods also require the evaluation of the exchange-correlation (XC) potential matrix by numerical integration methods due to the non-linear nature of the XC energy functional. In contrast to AO ERI methods, numerical integration methods developed for molecular DFT are much simpler to port to GPU architectures<sup>7,17</sup> and are able to heavily utilize optimized BLAS libraries to attain near peak floating point performance on modern GPU hardware<sup>18</sup>. Recently, there have been a number of works which have addressed key optimization challenges pertaining to the distributed memory evaluation of the XC potential on GPU accelerated computing clusters<sup>17,24</sup>. In this work, we extend the parallel integration infrastructure developed for the XC potential by the authors<sup>17</sup> to treat the Gaussian basis sn-K method on distributed memory architectures.

The remainder of this work is organized as follows. In Sec. II we review the salient aspects of Gaussian basis DFT necessary to develop parallel algorithms for the evaluation of the Coulomb (Sec. IIB) and exact-exchange (Sec. IID) matrices. In Secs. IIE and IIF, we described the extension of the distributed memory integration procedure developed for the XC potential (Sec. IIC) to treat sn-K on GPU clusters. In Sec. III, we demonstrate the strong scaling performance of the developed methods for a range of systems and, finally, discuss future outlook for further development of distributed memory GPU algorithms based on the methods presented in this work in Sec. IV.

## II. THEORY AND IMPLEMENTATION

### A. The Hybrid Kohn-Sham Fock Matrix

Given a set of basis functions,  $\mathcal{B} = \{\phi_\mu : \mathbb{R}^3 \rightarrow \mathbb{R}\}_{\mu=1}^{N_b}$ , the hybrid Kohn-Sham (KS) Fock matrix,  $\mathbf{F} \in \mathbb{R}^{N_b \times N_b}$ , is given by<sup>90,91</sup>

$$\mathbf{F} = \mathbf{h} + \mathbf{J}[\mathbf{D}] - c_x \mathbf{K}[\mathbf{D}] + \mathbf{V}^{xc}[\mathbf{D}], \quad (1)$$

where  $\mathbf{h}$  is the core Hamiltonian which contains the free-particle (kinetic) Hamiltonian and the electron-nuclear interaction potential, and  $\mathbf{D}$  is the basis-discretized one-particle density matrix.  $\mathbf{J}$ ,  $\mathbf{K}$ , and  $\mathbf{V}^{xc}$  are the density-dependent (as denoted by  $[\cdot]$ ) Coulomb, exact-exchange, and exchange-correlation (XC) potential matrices, respectively, which describe the mean-field electron-electron interactions in the hybrid KS model as modulated by the exact-exchange parameter  $c_x \in \mathbb{R}^+$ . For real-valued basis functions and density matrices, these matrices take the forms<sup>91</sup>

$$J_{\mu\nu} = \sum_{\lambda\kappa} (\mu\nu|\lambda\kappa) D_{\lambda\kappa} \quad (2)$$

$$K_{\mu\nu} = \sum_{\lambda\kappa} (\mu\lambda|\nu\kappa) D_{\lambda\kappa} \quad (3)$$

$$V_{\mu\nu}^{xc} = \int_{\mathbb{R}^3} d^3\mathbf{r} \phi_\mu(\mathbf{r}) \frac{\delta E^{xc}[\rho(\mathbf{r})]}{\delta \rho(\mathbf{r})} \phi_\nu(\mathbf{r}) \quad (4)$$

where  $E^{xc}$  is the XC energy functional evaluated at the electron density,  $\rho : \mathbb{R}^3 \rightarrow \mathbb{R}$

$$\rho(\mathbf{r}) = \sum_{\mu\nu} D_{\mu\nu} \phi_\mu(\mathbf{r}) \phi_\nu(\mathbf{r}) \quad (5)$$

and

$$(\mu\lambda|\nu\kappa) = \iint_{\mathbb{R}^3} d^3\mathbf{r} d^3\mathbf{r}' \frac{\phi_\mu(\mathbf{r}) \phi_\lambda(\mathbf{r}) \phi_\nu(\mathbf{r}') \phi_\kappa(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \quad (6)$$

is the electron-repulsion integral (ERI) tensor.

In this work, we take  $\mathcal{B}$  to be comprised of contracted, atom-centered Gaussian-type orbitals (GTO), however, we note that the general principles presented here may be extended to other atom-centered ansätze, such as Slater-type orbitals (STO), as well. We denote the atomic center of  $\phi_\mu$  as  $\mathbf{R}_\mu$  in the following. GTO Fock matrix construction is dominated by the three density-dependent terms which must be evaluated for each new density in e.g., an SCF optimization or dynamics simulation. In the remainder of this section, we examine the scalable evaluation of these terms on distributed memory GPU architectures.

## B. The Coulomb Matrix via Density-Fitting J-Engine

Efficient evaluation of the Coulomb potential Eq. (2) takes advantage of the density fitting (DF; also known as the Resolution-of-the-Identity) approximation of the two-electron integrals:

$$(\mu\nu|\lambda\kappa) \approx \sum_{KL} (\mu\nu|K)(\mathbf{V}^{-1})_{KL}(L|\lambda\kappa) \quad (7)$$

where 3- and 2-center Coulomb integrals involving the (auxiliary) density-fitting basis  $\{\chi_K : \mathbb{R}^3 \rightarrow \mathbb{R}\}_{K=1}^{N_{\text{aux}}}$  were introduced:

$$(\mu\nu|K) = \iint_{\mathbb{R}^3} d^3\mathbf{r}d^3\mathbf{r}' \frac{\phi_\mu(\mathbf{r})\phi_\nu(\mathbf{r})\chi_K(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} \quad (8)$$

$$V_{LK} = \iint_{\mathbb{R}^3} d^3\mathbf{r}d^3\mathbf{r}' \frac{\chi_L(\mathbf{r})\chi_K(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}. \quad (9)$$

DF approximation of Eq. (2) leads to the following factorization of the Coulomb potential:

$$V_L = \sum_{\lambda\kappa} (L|\lambda\kappa)D_{\lambda\kappa}, \quad (10)$$

$$D_K = \sum_L (\mathbf{V}^{-1})_{KL}V_L, \quad (11)$$

$$J_{\mu\nu} \stackrel{\text{DF}}{\approx} \sum_K (\mu\nu|K)D_K, \quad (12)$$

with  $V_L$  and  $D_K$  representing the Coulomb potential of the exact density  $\rho$  and the *robust* approximation of  $\rho$  in the DF basis, respectively. DF-based evaluation of  $J$  costs  $\mathcal{O}(N^3)$  in the traditional approach where the Cholesky factorization of the positive-definite matrix  $\mathbf{V}$  is computed and stored (this allows to amortize its cost over the SCF iterations). In practice, however, the cost is largely controlled by the evaluation of  $\mathcal{O}(N^2)$  nonnegligible Coulomb 3-center AO integrals in Eqs. (10) and (12). While for 3-center AO integrals dedicated optimization of AO integral evaluation is possible,<sup>92-94</sup> including specific developments for the GPU architectures,<sup>95</sup> optimal evaluation of Eqs. (10) and (12) involves blurring the line between the integral evaluation and density contraction via several related ideas<sup>80-83</sup> that is often dubbed the *J-engine* approach. The original J-engine approaches were demonstrated in the context of Eq. (2); Kussmann et al.<sup>13</sup> recently illustrated the utility of the McMurchie-Davidson-based<sup>96</sup> J-engine<sup>83</sup> in the DF context. Here we only briefly recap the DF-based J-engine (“DF-J-engine”) formalism using the established notation<sup>97</sup> for Gaussian AO integrals.

An uncontracted primitive Cartesian Gaussian with exponent  $\zeta_a \in \mathbb{R}^+$  and non-negative integer Cartesian “quanta”  $\mathbf{a} \equiv \{a_x, a_y, a_z\}$  centered at  $\mathbf{R}_a \equiv \{A_x, A_y, A_z\}$  will be denoted by

$$\phi_{\mathbf{a}}(\mathbf{r}) \equiv x_A^{a_x} y_A^{a_y} z_A^{a_z} \exp(-\zeta_a r_A^2), \quad (13)$$

with  $\mathbf{r}_A \equiv \{x_A, y_A, z_A\}$ ,  $x_A \equiv x - A_x$ , etc.  $l_{\mathbf{a}} \equiv a_x + a_y + a_z \geq 0$  is colloquially referred to as the “angular momentum” of a Gaussian. Closely related to the Cartesian Gaussian is a Hermite Gaussian:

$$\Lambda_{\tilde{\mathbf{a}}}(\mathbf{r}) \equiv \left( \frac{\partial}{\partial x_A} \right)^{\tilde{a}_x} \left( \frac{\partial}{\partial y_A} \right)^{\tilde{a}_y} \left( \frac{\partial}{\partial z_A} \right)^{\tilde{a}_z} \exp(-\zeta_a r_A^2). \quad (14)$$

A primitive Cartesian Gaussian and a product of two primitive Cartesian Gaussians can be expressed as linear combinations of Hermite Gaussians,

$$\phi_{\mathbf{a}}(\mathbf{r}) = \sum_{\tilde{a}_x=0}^{\tilde{a}_x \leq a_x} E_{a_x}^{\tilde{a}_x} \sum_{\tilde{a}_y=0}^{\tilde{a}_y \leq a_y} E_{a_y}^{\tilde{a}_y} \sum_{\tilde{a}_z=0}^{\tilde{a}_z \leq a_z} E_{a_z}^{\tilde{a}_z} \Lambda_{\tilde{\mathbf{a}}}(\mathbf{r}) \equiv \sum_{\tilde{\mathbf{a}}} E_{\mathbf{a}}^{\tilde{\mathbf{a}}} \Lambda_{\tilde{\mathbf{a}}}(\mathbf{r}), \quad (15)$$

$$\phi_{\mathbf{a}}(\mathbf{r})\phi_{\mathbf{b}}(\mathbf{r}) = \sum_{\tilde{p}_x=0}^{\tilde{p}_x \leq a_x + b_x} (E_x)_{a_x b_x}^{\tilde{p}_x} \sum_{\tilde{p}_y=0}^{\tilde{p}_y \leq a_y + b_y} (E_y)_{a_y b_y}^{\tilde{p}_y} \sum_{\tilde{p}_z=0}^{\tilde{p}_z \leq a_z + b_z} (E_z)_{a_z b_z}^{\tilde{p}_z} \Lambda_{\tilde{\mathbf{p}}}(\mathbf{r}) \equiv \sum_{\tilde{\mathbf{p}}} E_{\mathbf{ab}}^{\tilde{\mathbf{p}}} \Lambda_{\tilde{\mathbf{p}}}(\mathbf{r}), \quad (16)$$

with  $\zeta_p \equiv \zeta_a + \zeta_b$ ,  $\mathbf{R}_p \equiv \frac{\zeta_a \mathbf{R}_a + \zeta_b \mathbf{R}_b}{\zeta_a + \zeta_b}$ , and  $E_{\mathbf{a}}^{\tilde{\mathbf{a}}} \equiv \prod_{i=x,y,z} E_{a_i}^{\tilde{a}_i}$ ,  $E_{\mathbf{ab}}^{\tilde{\mathbf{p}}} \equiv \prod_{i=x,y,z} (E_i)_{a_i b_i}^{\tilde{p}_i}$ . The Hermite-to-Cartesian transformation coefficients are evaluated straightforwardly by recursion.<sup>96</sup>

The use of Eqs. (15) and (16) allows to express a 3-center Coulomb integral over primitive Cartesian Gaussians as a linear combination,

$$(\mathbf{ab}|\mathbf{c}) = \sum_{\tilde{\mathbf{p}}, \tilde{\mathbf{c}}} E_{\mathbf{ab}}^{\tilde{\mathbf{p}}} E_{\mathbf{c}}^{\tilde{\mathbf{c}}} (\tilde{\mathbf{p}}|\tilde{\mathbf{c}}), \quad (17)$$

of the Coulomb integral between two primitive Hermite Gaussians:

$$(\tilde{\mathbf{p}}|\tilde{\mathbf{c}}) \equiv \iint_{\mathbb{R}^3} d^3\mathbf{r} d^3\mathbf{r}' \frac{\Lambda_{\tilde{\mathbf{p}}}(\mathbf{r}) \Lambda_{\tilde{\mathbf{c}}}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}. \quad (18)$$

The latter can be evaluated directly,

$$(\tilde{\mathbf{p}}|\tilde{\mathbf{c}}) \equiv (-1)^{l_{\tilde{\mathbf{c}}}} (\tilde{\mathbf{p}} + \tilde{\mathbf{c}})^{(0)}, \quad (19)$$



from the auxiliary integral,

$$(\tilde{\mathbf{r}})^{(m)} \equiv \left( \frac{\partial}{\partial x_R} \right)^{\tilde{r}_x} \left( \frac{\partial}{\partial y_R} \right)^{\tilde{r}_y} \left( \frac{\partial}{\partial z_R} \right)^{\tilde{r}_z} (\mathbf{0})^{(m)}, \quad (20)$$

with  $(\mathbf{0})^{(m)}$  related to the Boys function  $F_m(x)$ :

$$(\mathbf{0})^{(m)} \equiv (-2\rho)^m \frac{2\pi^{5/2}}{\zeta_p \zeta_c \sqrt{\zeta_p + \zeta_c}} F_m(\rho |\mathbf{R}_p - \mathbf{R}_c|^2), \quad (21)$$

$$F_m(x) \equiv \int_0^1 dy y^{2m} \exp(-xy^2), \quad (22)$$

$$\rho \equiv \frac{\zeta_p \zeta_c}{\zeta_p + \zeta_c}. \quad (23)$$

Auxiliary integrals are evaluated recursively,

$$(\tilde{\mathbf{r}} + \mathbf{1}_i)^{(m)} = \tilde{r}_i (\tilde{\mathbf{r}} - \mathbf{1}_i)^{(m+1)} + (P_i - C_i) (\tilde{\mathbf{r}})^{(m+1)}, \quad (24)$$

starting from  $(\mathbf{0})^{(m)}$ .

Efficient evaluation of Eqs. (10) and (12) involves contracting densities with Hermite-to-Cartesian transformation coefficients of the ket functions (Eqs. (16) and (15), respectively) to produce ‘‘Hermite’’ densities which can be stored and reused for every bra function in Eqs. (10) and (12), as illustrated here for a single primitive shell contribution to Eq. (10):

$$\sum_{\mathbf{ab}} (\mathbf{c}|\mathbf{ab}) D_{\mathbf{ab}} \stackrel{\text{Eq. (17)}}{=} \sum_{\mathbf{ab}} \left( \sum_{\tilde{\mathbf{c}}|\tilde{\mathbf{p}}} E_{\tilde{\mathbf{c}}}^{\tilde{\mathbf{c}}} (\tilde{\mathbf{c}}|\tilde{\mathbf{p}}) E_{\mathbf{ab}}^{\tilde{\mathbf{p}}} \right) D_{\mathbf{ab}} \quad (25)$$

$$= \sum_{\tilde{\mathbf{c}}} E_{\tilde{\mathbf{c}}}^{\tilde{\mathbf{c}}} \left( \sum_{\tilde{\mathbf{p}}} (\tilde{\mathbf{c}}|\tilde{\mathbf{p}}) D_{\tilde{\mathbf{p}}} \right), \quad (26)$$

where

$$D_{\tilde{\mathbf{p}}} \equiv \sum_{\mathbf{ab}} E_{\mathbf{ab}}^{\tilde{\mathbf{p}}} D_{\mathbf{ab}} \quad (27)$$

and the order of evaluation is indicated by the parentheses. Refactorization of Eq. (25) via Eq. (26) is the key idea of J-engine: it leads to great FLOP reduction which is easily rationalized as follows: instead of multiplying ‘‘matrix’’  $(\tilde{\mathbf{c}}|\tilde{\mathbf{p}})$  by ‘‘matrix’’  $E_{\mathbf{ab}}^{\tilde{\mathbf{p}}}$ , then evaluating the inner product with ‘‘vector’’  $D_{\mathbf{ab}}$ , in the J-engine factorization the inner product of ‘‘matrix’’  $(\tilde{\mathbf{c}}|\tilde{\mathbf{p}})$  with ‘‘vector’’  $D_{\tilde{\mathbf{p}}}$  is evaluated directly. This can also be viewed as early ‘‘digestion’’ of the integrals; more general framework for early digestion of the integrals beyond the J matrix evaluation has been considered by Gill et al.<sup>82</sup>

Equations (10) and (12) of the DF-J-engine approach were implemented in the open-source `LibintX` library<sup>95,98</sup> and Eq. (11) was implemented using the `TiledArray` library for distributed tensor contractions.<sup>66,99</sup> Evaluation of Eq. (10) in `LibintX` proceeds as follows:

1. **CPU**: construct a batch of **ab** shell pairs with same  $l_{\text{ket}} = l_{\mathbf{a}} + l_{\mathbf{b}}$  such that  $D_{\mathbf{ab}}$  is local, copy to mapped memory buffers;
2. **GPU**: launch a kernel to transform current batch of  $D_{\mathbf{ab}}$  to  $D_{\tilde{\mathbf{p}}}$  (Eq. (27));
3. **GPU**: launch kernels to evaluate Hermite integrals (Eq. (19)) and their contributions to the Coulomb potential (Eq. (26)) for all shells **c** (one kernel per  $l_{\mathbf{c}}$ );
4. repeat for the next **ab** batch.

Note after launching the GPU kernels, the CPU starts to work on preparing the density and metadata for the next batch, thus the GPU and CPU work overlaps. Note that no significant metadata is kept persistent between SCF iterations to allow GPU memory use by other application stages.

Evaluation of Eq. (12) in `LibintX` is only slightly more complicated:

1. **CPU**: construct a batch of **ab** shell pairs with same  $l_{\text{bra}} = l_{\mathbf{a}} + l_{\mathbf{b}}$  such that  $J_{\mathbf{ab}}$  is local;
2. **GPU**: launch a kernel to evaluate  $E_{\mathbf{ab}}^{\tilde{\mathbf{p}}}$ ;
3. **GPU**: launch kernels to evaluate Hermite density  $D_{\tilde{\mathbf{c}}} \equiv \sum_{\mathbf{c}} E_{\mathbf{c}}^{\tilde{\mathbf{c}}} D_{\mathbf{c}}$ , Hermite integrals (Eq. (19)), and their contributions to the Coulomb potential in the Hermite basis,

$$J_{\tilde{\mathbf{p}}} = \sum_{\tilde{\mathbf{c}}} (\tilde{\mathbf{p}}|\tilde{\mathbf{c}}) D_{\tilde{\mathbf{c}}}, \quad (28)$$

(this is analogous to the first step in Eq. (26), but adapted to the evaluation of Eq. (12)) for all shells **c** (one kernel per  $l_{\mathbf{c}}$ );

4. **GPU**: launch kernel to transform  $J_{\tilde{\mathbf{p}}}$  to  $J_{\mathbf{ab}}$  via

$$J_{\mathbf{ab}} = \sum_{\tilde{\mathbf{p}}} E_{\mathbf{ab}}^{\tilde{\mathbf{p}}} J_{\tilde{\mathbf{p}}} \quad (29)$$

5. **CPU**: store  $J_{\mathbf{ab}}$  shell-sets

6. repeat for the next **ab** batch;

CPU and GPU work is again overlapped, with the CPU thread starting to work on steps 1-3 of the next batch while the work on the previous batch is being completed; step 4 is scheduled only after completion of the previous batch’s step 5.

The distributed memory work distribution in `LibintX`’s DF-J-engine is statically determined by the distribution of the user-provided density matrix among ranks and by the expected distribution of the J matrix result. This design decision is motivated by the desire to simplify the API of the `LibintX` library and avoid mandating the density to be in any particular data structure or distribution manner; the user provides the density to the library in the form of a C++ lambda function (closure) that provides one block of the density matrix at a time (or fails to provide it if it is not located on the current node). Production computations in Sec. III provided density and stored the resulting J matrix as a block-sparse distributed array (`DistArray`) object of the `TiledArray` framework.<sup>67,99</sup> Tiling of the density matrix was determined by the k-means-based clustering of the atoms,<sup>100</sup> with tiles divided evenly among ranks (for  $p$  distributed ranks, the first  $n^2/p$  tile ordinals assigned to rank 0, the next  $n^2/p$  tile ordinals assigned to rank 1, etc.; note than due to sparsity actual tile counts per rank are lower and end up being problem dependent). In evaluation of Eq. (10), each rank evaluates all contributions that involve the tiles of the density matrix that are screened out and that reside on that rank; each rank produces contributions to all elements of vector  $V_L$ , thus these contributions are reduced across all ranks before evaluating Eq. (11). A similar lambda-based mechanism is used by `LibintX` to return the distributed J matrix to the user and distribute the work in evaluation of Eq. (12).

In the interest of keeping the focus on the distributed-memory parallelization, the kernel-level implementation details of `LibintX`’s DF-J-engine will be kept to a minimum here. Each GPU thread evaluates the entire set of “1-center” auxiliary integrals (Eq. (24)), converts them to “2-center” integrals via Eq. (19) and “digests” them via Eq. (26) (or its analog for Eq. (12)) for a single primitive shell. This allows the algorithm to completely eliminate thread divergences since every thread in a thread block performs exactly the same computation. This approach is radically different from the work distribution when computing integrals<sup>95</sup> in which multiple threads cooperatively evaluated shell components; this is possible due to the greatly reduced memory requirements of the DF-J-engine compared to the AO integral engine. A key feature of the kernel implementation is the use of compile-time

(templates and `constexpr`) C++ programming instead of custom code generation, similarly to how the 3-center integral evaluation (rather than the J-engine) was implemented earlier.<sup>95</sup>

### C. The Exchange-Correlation Potential Matrix

Unlike  $\mathbf{J}$ , which may be assembled directly from analytical integrals, the integrals involved in  $\mathbf{V}^{xc}$  (Eq. (4)) must be evaluated numerically due to the nonlinear nature of  $E^{xc}$  and its functional derivatives. As detailed elsewhere<sup>7,101–103</sup>, for atom centered bases,  $\mathbf{V}^{xc}$  may be efficiently evaluated via

$$V_{\mu\nu}^{xc} \approx \sum_{i \in \mathcal{Q}} \Phi_{\mu i} Z_{\nu i} + Z_{\mu i} \Phi_{\nu i} \quad (30)$$

$$\Phi_{\mu i} = \phi_{\mu}(\mathbf{r}_i) \quad (31)$$

where  $\Phi$  is the collocation matrix and  $\mathcal{Q} = \{(w_i, \mathbf{r}_i)\}_{i=1}^{N_g}$  is a numerical quadrature consisting of nodes,  $\mathbf{r}_i \in \mathbb{R}^3$ , and grid weights,  $w_i \in \mathbb{R}$ . For molecular calculations, due to the irregular nature of the integrands in the vicinity of nuclear charges,  $\mathcal{Q}$  is typically taken as the union of spherical quadratures ( $\mathbb{R} \times S^2$ ) originated at each atomic center coupled with a modified weight partitioning scheme to account for overlapping regions. We refer the reader to Refs. 104–112 for more comprehensive discussions regarding the construction of generic molecular quadratures. In this work, atomic grids are constructed according to the Mura-Knowles (MK) scheme<sup>107</sup> with Lebedev-Laikov<sup>113</sup> angular grids, and we use the molecular weight partitioning scheme of Stratmann, *et al*<sup>105</sup>. In addition, angular grids are radially pruned according to the scheme of Treutler, *et al*<sup>109</sup>.

The auxiliary matrix  $\mathbf{Z} \in \mathbb{R}^{N_b \times N_g}$  is method-dependent; the form of which depends on  $E^{xc}$ . Within the spin-restricted generalized-gradient approximation (GGA),  $\mathbf{Z}$  takes the form<sup>101,102</sup>

$$Z_{\mu i} = \frac{w_i}{2} (E_i^{\rho} \Phi_{\mu i} + 4E_i^{\gamma} (\nabla \rho(\mathbf{r}_i) \cdot \nabla \Phi_{\mu i})), \quad (32)$$

$$E_i^{\rho/\gamma} = \left. \frac{\partial \varepsilon(\rho(\mathbf{r}), \gamma(\mathbf{r}))}{\partial \rho/\gamma} \right|_{\mathbf{r}=\mathbf{r}_i} \quad (33)$$

$$\rho(\mathbf{r}_i) = \sum_{\mu} F_{\mu i} \Phi_{\mu i}, \quad \nabla \rho(\mathbf{r}_i) = 2 \sum_{\mu} F_{\mu i} \nabla \Phi_{\mu i}, \quad (34)$$

$$F_{\mu i} = \sum_{\nu} D_{\mu\nu} \Phi_{\nu i}, \quad (35)$$

where  $\varepsilon : \mathbb{R}^2 \rightarrow \mathbb{R}$  is the GGA energy density which depends on  $\rho$  and  $\gamma = \nabla\rho \cdot \nabla\rho$ . Similar expressions have been derived for spin-generalized and meta-GGA XC functionals<sup>13,101,103,114</sup>. This evaluation scheme for  $\mathbf{V}^{xc}$  is particularly attractive for GPU architectures due to the fact that the most compute-intensive operations can be implemented using level-3 (Eq. (30) via SYR2K and Eq. (35) via GEMM) and level-1 (Eq. (34) via DOT<sup>115</sup>) BLAS operations<sup>7,13,17</sup>. As such, these operations can leverage heavily optimized GPU BLAS libraries, leaving only a small number of DFT-specific kernels (e.g. Eqs. (31) to (33)) to be optimized by the chemistry-domain developer. We will examine specific details regarding the use of GPU BLAS libraries for this application in Sec. II F.

#### D. The Seminumerical Exact-Exchange Matrix

Although  $\mathbf{K}$  can be evaluated directly from Coulomb integrals as in Eq. (3), this is rarely the most efficient approach due to the steep rise of the computational cost of the 4-center integrals with the angular momenta. As detailed elsewhere<sup>10,86–88</sup>, the ERI tensor may be factorized on a quadrature grid via

$$(\mu\lambda|\nu\kappa) \approx \frac{1}{2} \sum_{i \in \mathcal{Q}} w_i A_{\mu\lambda i} \Phi_{\nu i} \Phi_{\kappa i} + (\mu\lambda \leftrightarrow \nu\kappa), \quad (36)$$

$$A_{\mu\lambda i} = \int_{\mathbb{R}^3} d^3\mathbf{r} \frac{\phi_{\mu}(\mathbf{r})\phi_{\lambda}(\mathbf{r})}{|\mathbf{r} - \mathbf{r}_i|}, \quad (37)$$

where  $\mathbf{A}$  is the 3-center Coulomb potential (3c-CP) integral tensor, and  $\Phi$  and  $\mathcal{Q}$  are the collocation matrix and quadrature discussed in Sec. II C. Inserting Eq. (36) into Eq. (3), we obtain a simple expansion for  $\mathbf{K}$ <sup>10,86–88</sup>,

$$G_{\mu i} = \sum_{\lambda} w_i A_{\mu\lambda i} F_{\lambda i}, \quad (38)$$

$$\tilde{K}_{\mu\nu} = \sum_i G_{\mu i} \Phi_{\nu i}, \quad (39)$$

$$K_{\mu\nu} \approx \frac{1}{2} \left( \tilde{K}_{\mu\nu} + \tilde{K}_{\nu\mu} \right), \quad (40)$$

where  $\mathbf{F}$  is same intermediate given in Eq. (35) for the evaluation of  $\mathbf{V}^{xc}$ . This method for  $\mathbf{K}$  assembly will be referred to as seminumerical exchange<sup>10</sup> (sn-K) in the following. Apart from the use of common intermediates, there exists a striking similarity between Eqs. (38) to (40) and the  $\mathbf{V}^{xc}$  assembly in Eqs. (30), (32), (34) and (35). As has been posited elsewhere<sup>10,88</sup>,

sn-K can reuse many of the same algorithmic primitives as those used for the evaluation of  $\mathbf{V}^{xc}$ . One might even be tempted to implement Eq. (38) via BLAS, but we (as have others<sup>10,88</sup>) have found this to be inefficient in practice, particularly on GPU architectures. We present the scheme by which we perform the combined evaluation of Eqs. (37) and (38) in Appendix A.

## E. Batching and Screening of Numerical Integrals

Each of the expressions in Eqs. (30) to (35), (38) and (39) is *local* with respect to the numerical quadrature, and thus may be partitioned along their grid indices ( $i$ ) with final results (e.g., Eqs. (30) and (39)) being recovered as a sum over locally integrated intermediates. This partitioning has clear ramifications in the context of parallelism which will be discussed in Sec. II F. Grid point locality also allows for the exploitation of the spatial sparsity of  $\mathcal{B}$  and various operators (e.g. Eq. (37)), which in turn takes the nominal  $\mathcal{O}(N_b^2 N_g)$  scaling of these numerical schemes to methods which scale (near-)linearly with respect to system size.<sup>10,88,102,105</sup>

As the sparsity profiles of spatially adjacent grid points are similar, it is canonical to *batch* these grid points into subsets  $\mathcal{Q}_j$  such that  $\mathcal{Q} = \bigcup_j \mathcal{Q}_j$  with  $\mathcal{Q}_j \cap \mathcal{Q}_k = \emptyset$  for  $j \neq k$ . Several approaches for quadrature grid batching have been suggested in various contexts.<sup>10,102,105</sup> In this work, we utilize an octree-based partition scheme which has been used in previous studies by the authors.<sup>17</sup> In general, the octree approach will yield more spatially local quadrature batches than other approaches<sup>10</sup>, but will also yield batches of irregular sizes which potentially leads to load imbalance between tasks. We discuss the implications of this irregularity in the following. Given a partitioning  $\{\mathcal{Q}_j\}_{j=1}^{N_{\text{batch}}}$ , one may construct sets of important quantities which are approximately considered non-negligible for a particular integrand (e.g.,  $\mathbf{K}$  or  $\mathbf{V}^{xc}$ ). The simplest of these sets, which is required for both numerical integrands considered in this work, pertains to negligible basis functions,

$$\mathcal{B}_j = \{\phi_\mu \mid \exists \mathbf{r}_i \in \mathcal{Q}_j \quad \text{s.t.} \quad |\phi_\mu(\mathbf{r}_i)| \geq \varepsilon_B\}, \quad (41)$$

where  $\varepsilon_B$  is a basis tolerance. For GTO bases, Eq. (41) may be quickly determined by encircling each basis function with a sphere of radius  $r_\mu^{\text{cut}}$  such that  $|\phi_\mu(\mathbf{r} - \mathbf{R}_\mu)| < \varepsilon_B \forall |\mathbf{r} - \mathbf{R}_\mu| > r_\mu^{\text{cut}}$ , and only keeping those elements of  $\mathcal{B}$  for which their non-negligible sphere

spatially overlaps  $\mathcal{Q}_j$ <sup>10,17,88,105</sup>. This process formally scales  $\mathcal{O}(N_b N_{\text{batch}})$  as each basis shell must be checked against each  $\mathcal{Q}_j$ . It is worth noting that in scenarios where the nuclei remain stationary for several Fock builds, e.g., in a self-consistent field optimization (SCF),  $\mathcal{B}_j$  need only be computed once and reused for subsequent iterations. This process of basis screening is the primary mechanism by which sparsity may be exploited in the numerical integration of  $\mathbf{V}^{xc105}$ . For example, given  $\mathcal{B}_j$ , the compute-intensive Eq. (35) becomes

$$F_{\mu i}^{xc,(j)} = \sum_{\nu \in \mathcal{B}_j} D_{\mu\nu}^{xc,(j)} \Phi_{\mu i}^{(j)}, \quad \mu \in \mathcal{B}_j, i \in \mathcal{Q}_j, \quad (42)$$

where  $\mathbf{D}^{xc,(j)}$ ,  $\mathbf{F}^{xc,(j)}$  and  $\Phi^{(j)}$  are contiguous batch local matrices pertaining to  $\mathcal{Q}_j$ .

The case is significantly different for the sn-K method as basis function and integral screening alone can only produce  $\mathcal{O}(N^2)$  scaling<sup>116</sup>. Several methods have been suggested for density-dependent screening in the sn-K method.<sup>10,86–88,117</sup> Due to its simple form and demonstrated ability to generate sufficient sparsity for near-linear scaling, we adopt the sn-Link screening approach of Laqua, *et al*<sup>10</sup> in this work, though we note that other schemes could be employed with only minor modifications of the overall algorithm design. In the sn-Link method, a list of basis pairs,  $\mathcal{V}_j$ , is selected for each  $\mathcal{Q}_j$  such that

$$\mathcal{V}_j = \{(\phi_\mu, \phi_\nu) \mid \varepsilon_{\mu\nu}^{E(j)} \geq \varepsilon_E \vee \varepsilon_{\mu\nu}^{K(j)} \geq \varepsilon_K\}, \quad (43)$$

where  $\varepsilon_K$  and  $\varepsilon_E$  are tolerances which screen shell-pair contributions to  $\mathbf{K}$  and  $\text{Tr}[\mathbf{KD}]$  (the exact-exchange energy), respectively, and

$$\varepsilon_{\mu\nu}^{E(j)} = \tilde{F}_\mu^{K,(j)} \tilde{F}_\nu^{K,(j)} \tilde{A}_{\mu\nu}, \quad (44)$$

$$\varepsilon_{\mu\nu}^{K(j)} = \max\left(\tilde{F}_\mu^{K,(j)}, \tilde{F}_\nu^{K,(j)}\right) \tilde{\phi}^{(j)} \tilde{A}_{\mu\nu}, \quad (45)$$

$$\tilde{F}_\mu^{K,(j)} = \sum_{\nu \in \mathcal{B}_j} |D_{\mu\nu}| \tilde{\Phi}_\nu^{(j)}, \quad (46)$$

$$\tilde{A}_{\mu\nu} = \max_{i \in \mathcal{Q}_j} |A_{\mu\nu i}|, \quad \Phi_\nu^{(j)} = \max_{i \in \mathcal{Q}_j} |\sqrt{w_i} \Phi_{\nu i}|, \quad \tilde{\phi}^{(j)} = \max_{i \in \mathcal{Q}_j} \sqrt{w_i} \sum_{\mu \in \mathcal{B}_j} |\Phi_{\mu i}|. \quad (47)$$

To avoid re-computation of the 3c-CP integrals in the screening procedure,  $\tilde{A}_{\mu\nu}$  is approximated with global upper-bound estimates from the appendix of Ref. 118 in practice<sup>10</sup>. We refer the reader to Laqua, *et al*.<sup>10</sup> for further details pertaining to the veracity of this screening procedure for sn-K. Given  $\mathcal{V}_j$ , block-sparse expressions similar to Eq. (42) can be derived

for sn-K, e.g.,

$$F_{\lambda i}^{K,(j)} = \sum_{\nu \in \mathcal{B}_j} D_{\lambda\nu}^{K,(j)} \Phi_{\nu i}^{(j)}, \quad G_{\mu i}^{K,(j)} = \sum_{(\cdot, \lambda) \in \mathcal{V}_j} w_i A_{\mu\lambda i} F_{\lambda i}^{K,(j)}, \quad \mu, \lambda \in \mathcal{V}_j, i \in \mathcal{Q}_j. \quad (48)$$

The  $(\cdot, \lambda)$ -sum pertaining to  $G_{\mu i}^{K,(j)}$  indicates only inclusion of terms for which  $\exists \phi_\mu$  such that  $(\phi_\mu, \phi_\lambda) \in \mathcal{V}_j$ . Both the  $\mathcal{B}$ -collision detection and the evaluation of performance critical  $\mathcal{V}_j$ -intermediates (e.g., Eq. (46) can be implemented as a single, large dimension-GEMM over tasks) are executed on the GPU in this work.

There are several key differences in screening processes for sn-K in comparison with the XC integration:

1. As this procedure is density-dependent,  $\mathcal{V}_j$  must be reevaluated each new density and its cost cannot generally be amortized over e.g., an SCF optimization.
2. While the matrix multiplication in Eq. (46) can be restricted over  $\nu \in \mathcal{B}_j$ , the  $\mu$ -index *must* be over the entire basis set  $\mathcal{B}$  in order to check for all possible non-trivial elements of  $\mathcal{V}_j$ . Therefore, Eq. (46) formally scales  $\mathcal{O}(N_b N_{\text{batch}})$  assuming  $\mathcal{B}_j$  can be screened to an amortized constant.<sup>105</sup> While this formally scales the same as  $\mathcal{B}$ -screening, it will be accompanied by a much larger prefactor (GEMM vs scalar collision detection).
3. Eq. (43) scales  $\mathcal{O}(N_b N_{\text{batch}})$  as each non-negligible shell-pair (which asymptotically scales  $\mathcal{O}(N_b)$ ) must be checked against every  $\mathcal{Q}_j$ .

We discuss the implications of this screening procedure on the overall performance and scalability of numerical integration procedures in Sec. III.

## F. Load Balancing and Parallel Numerical Integration

In addition to enabling the screening of basis functions and operator integrals, the locality of the numerical integration procedure considered in this work is also particularly advantageous for distributed memory implementations<sup>17,24,31</sup>. As the batch-local quantities discussed in the previous subsection are independent, they may be executed concurrently and assembled *a posteriori* via collective reduction operations to form final integrands. The general distributed memory scheme explored in this work is a simple three step process (Alg. 1) that is the same for both the sn-K and XC integrations<sup>17</sup>.



- 1: **Input:** Density matrix  $\mathbf{D}$ , basis set  $\mathcal{B}$ .
- 2: **Output:** Desired integrand  $\mathbf{X} \in \{\mathbf{V}^{xc}, \mathbf{K}\}$
- 3:  $\mathcal{Q}_{local} \leftarrow$  Generate balanced local quadrature batches ▷ Alg. 2
- 4: Preallocate fraction of device memory.
- 5: Send quadrature independent data (e.g.  $\mathbf{D}$ ) to the device ▷ cpu/gpu
- 6:  $\mathbf{X}_{local} \leftarrow \mathbf{0}$  ▷ gpu
- 7: **while**  $\mathcal{Q}_{local} \neq \emptyset$  **do**
- 8:      $\mathcal{Q}_{batch} \leftarrow$  Subset of  $\mathcal{Q}_{local}$  s.t.  $\mathbf{X}$ -intermediates saturate device memory. ▷ cpu
- 9:     Send  $\mathcal{Q}_{batch}$  data (e.g.  $\{\mathcal{B}_j\}_{local}, \{\mathcal{V}_j\}_{local}$ ) to device ▷ cpu/gpu
- 10:      $\mathcal{Q}_{local} \leftarrow \mathcal{Q}_{local} \setminus \mathcal{Q}_{batch}$  ▷ cpu
- 11:      $\mathbf{X}_{local} \leftarrow \mathbf{X}_{local} + \mathcal{Q}_{batch}$  contributions to  $\mathbf{X}$  ▷ gpu
- 12: **end while**
- 13:  $\mathbf{X} \leftarrow (\text{All})\text{reduce}(\mathbf{X}_{local})$  ▷ collective

Algorithm 1. General scheme for the parallel, GPU-accelerated evaluation of  $\mathbf{V}^{xc}$  or  $\mathbf{K}$

Due to the varying extent to which sparsity can be realized between differing spatial regions, there is a significant potential for load imbalance in parallel implementations of molecular integration which rely on grid partitioning.<sup>17,24,25</sup> The optimal task assignment problem for irregular work is NP-Hard<sup>119</sup>, but reasonable solutions can be obtained using heuristic driven methods. A recent distributed memory DFT application<sup>24</sup> has adopted a dynamic load balancing scheme to address this problem. At large processor counts, dynamic load balancing algorithms rely on the ability to overlap the costs associated with the generation, assignment, and communication of tasks with their execution on local processors. Given the fine-granularity and large number of tasks generated by the octree-batching algorithm discussed in the previous section, the communication costs associated with the memory requirement of each task (e.g.,  $\mathcal{Q}_j$ ,  $\mathcal{B}_j$  and  $\mathcal{V}_j$ ) far outweigh their associated computational work. Instead, we have developed a *static* load balancing procedure for molecular integration which allows for the *a priori* distribution of work without the need to carefully balance the overlap of local work and task communication. In this static load balancing scheme, a heuristic cost,  $W_j$ , is associated with each task and is subsequently assigned to a rank via a greedy algorithm illustrated in Alg. 2. For the XC integration, the work heuristic

is given by<sup>17</sup>

$$W_j = (|\mathcal{B}_j| \times (9 + 2|\mathcal{B}_j|) + N_A^2 + 3) \times |\mathcal{Q}_j|. \quad (49)$$

Algorithm 2 is a replicated procedure which allows for near optimal task assignment given work heuristics which properly rank the relative computational costs between different tasks. As such, while being communication-free, it constitutes a scaling bottleneck due to Amdahl’s law. Similar algorithms based on prefix-sums have been suggested for the distributed generation of tasks<sup>119</sup> However, we have found that because Alg. 2 is able to globally optimize task assignment (via the sort of all tasks on  $W_j$ ), it generally is able to obtain a more optimal task partitioning than those based on prefix sums. Due to the fact that  $\mathcal{B}_j$  in particular can be computed once-per-nuclear configuration, the scaling bottleneck associated with the replicated nature of Alg. 2 can be amortized over several Fock builds for the XC integration.

For the sn-K integration, the problem becomes slightly more difficult. While it would be possible to generalize Eq. (49) to account for  $|\mathcal{V}_j|$ , replication of Eqs. (43) and (46) on each processor for all tasks would be a considerable bottleneck due to the steep scaling and prefactors associated with sn-K screening discussed in the previous subsection. Given an initial work partitioning generated by e.g., Eq. (49),  $\mathcal{V}_j$  may be computed locally, and could in principle be re-balanced according to the distributed prefix-sum algorithms previously mentioned. Apart from generally obtaining less-optimal work partitions than Alg. 2, for hybrid DFT simulations which require the evaluation of both  $\mathbf{V}^{xc}$  and  $\mathbf{K}$ , this procedure would require rebalancing the work distribution at every SCF step, which would in turn incur significant communication overhead. Instead, we have chosen to reuse the same task distribution for both sn-K and the XC integration based on Alg. 2 and Eq. (49). We examine the veracity of this reuse to yield balanced work for both sn-K and the XC integration in Sec. III.

If screened sufficiently well, the dimensions of the packed batch-local sub-matrices in Eqs. (31) and (32) will be small which means that the GEMM operations in e.g., Eqs. (42) and (48) will be of low dimension. While this screening leads to a significant decrease in the required computational work, any particular GEMM operation (or other kernel invocation) will not constitute enough work to occupy the resources of a GPU. While it is possible to concurrently execute GPU kernels via streaming (as has been explored in other works for GPU-accelerated sn-K integration<sup>10</sup>), the large number of kernels invoked for large systems partitioned by the octree-method would incur significant kernel launch overhead. As has

```

1: Input: Quadrature batches  $\{\mathcal{Q}_j\}_{j=1}^{N_{\text{batch}}}$ 
2: Output: Local quadrature batches  $\mathcal{Q}_{\text{local}} \subset \{\mathcal{Q}_j\}$ 
3: world_size  $\leftarrow$  Number of ranks
4: world_rank  $\leftarrow$  Index of this rank
5:  $\mathcal{T} \leftarrow []$ 
6: for  $j \in [1, N_{\text{batch}}]$  do
7:    $\mathcal{T}_j = (\mathcal{Q}_j, W_j)$  ▷ Eq. (49)
8: end for
9:  $\mathcal{T} \leftarrow$  Sort  $\mathcal{T}$  by  $W_j$ 
10:  $\mathcal{G} \leftarrow \text{zeros}(\text{world\_size})$ 
11:  $\mathcal{Q}_{\text{local}} \leftarrow []$ 
12: for  $j \in [1, N_{\text{batch}}]$  do
13:    $I \leftarrow \arg \min_J \mathcal{G}_J$  ▷ global
14:    $\mathcal{G}_I \leftarrow \mathcal{G}_I + W_j$  ▷ global
15:   if  $I = \text{world\_rank}$  then
16:      $\mathcal{Q}_{\text{local}} \leftarrow \mathcal{Q}_{\text{local}} \cup \mathcal{Q}_j$  ▷ local
17:   end if
18: end for
19: return  $\mathcal{Q}_{\text{local}}$ 

```

Algorithm 2. Replicated load balancing algorithm for distributed memory molecular integration.

been demonstrated in previous studies regarding the XC integration<sup>17,18</sup>, the use of *batched* kernels to concurrently execute logically identical tasks can lead to large performance improvements over concurrent stream injection for fine task granularity. The challenge for DFT applications is in that the dimensions of  $\mathcal{B}_j$  and  $\mathcal{V}_j$  can vary drastically in different spatial regions. For the evaluation of batched GEMM operations (e.g., Eqs. (42) and (48) batched over  $j$  indices), standard solutions to the problem exist in community GPU linear algebra libraries in the form of variable-sized batched BLAS.<sup>120,121</sup>

For sn-K, the batched evaluation of Eq. (39) is a straight-forward application of variable-sized batched GEMM as was applied to the formation of Eq. (42) for the XC integration. Batching of Eq. (38) can be achieved hierarchically over three dimensions (1)  $\{\mathcal{Q}_j\}$ , (2)

$(\mu, \nu) \in \mathcal{V}_j$  within  $\mathcal{Q}_j$ , and at the lowest level (3)  $i \in \mathcal{Q}_j$ . The kernel-level details for (3) are given in Appendix A. For any batched kernel algorithm, the natural challenge is to decide which task to batch together to execute concurrently on the device. For the algorithms presented here, we make the simple choice of executing as many tasks as will fit into device memory (Alg. 1). While this choice does mitigate kernel launch overhead and allows for the overlap of host data manipulations with GPU activity, it also introduces significant pressure on the GPU scheduler to effectively execute large numbers of independent tasks with variable amounts of work. However, we have seen that for DFT applications that this device-saturation approach is capable of yielding excellent utilization of GPU resources on an array of modern GPU hardware.<sup>17,18</sup>

In the large processor limit, the reduction step poses a considerable bottleneck for large  $N_b$  as the message size for the `(All)reduce` operation scales  $\mathcal{O}(N_b^2)$ . While standard MPI collectives are often sufficient for small  $N_b$ , we have found that for NVIDIA hardware, the NVIDIA Collective Communication Library (NCCL)<sup>122</sup> all-reduce primitive can yield significant strong scaling improvements for large  $N_b$ . NCCL provides collective primitives with an API similar to that of MPI with the addition of a CUDA stream argument which allows for its injection into asynchronous workflows. The all-reduce algorithm and exact choice of communication routes are selected by NCCL at runtime based on the system topology allowing for optimal performance on a wide-range of hardware configurations and message sizes. We will examine the comparison of MPI and NCCL collectives for this application in Sec. III.

### III. NUMERICAL RESULTS

The methods presented in this work are made publicly available as a part of the open-source `GauXC` library for exascale Gaussian basis DFT (XC and sn-K)<sup>17,18,103</sup> and the `LibintX` library for Gaussian AO integrals (DF-J-engine).<sup>95</sup> Each of these libraries were developed to be modular, reusable GPU-based components under the NWChemEx Exascale Computing Project and are freely available on GitHub<sup>98,123</sup> for use in other electronic structure packages. As present, these methods have been integrated into the `NWChemEx`<sup>124</sup> and `MPQC`<sup>125</sup> computational chemistry software packages. In the present study, all numerical results were obtained using `MPQC`; however, as the focus of this work is on the parallel construction of

Molecule	$N_A$	$N_b$	$N_{\text{aux}}$
Taxol	113	1,032	3,599
Olestra	453	3,181	11,633
Crambin	642	5,559	19,500
Ubiquitin	1,231	10,292	36,419

TABLE I. Representative systems considered in this study.  $N_A$  is the number of atoms,  $N_b$  is given for Cartesian 6-31G(d) and  $N_{\text{aux}}$  for def2-tzvp-j.

Grid Name	$N_{\text{rad}}$	$N_{\text{ang}}$
<b>Grid1</b>	50	302
<b>Grid2</b>	30	110

TABLE II. Spherical atomic quadratures consisting of  $N_{\text{rad}}$  radial point and  $N_{\text{ang}}$  angular points used in this work. Angular resolution has been pruned according to the scheme of Ref. 109

hybrid DFT Fock matrices, the performance characterizations presented here are expected to be representative of any software integration involving these libraries. All numerical experiments were carried out on the GPU partition of the Perlmutter (PM) supercomputer at the National Energy Research Scientific Computing Center (NERSC). Each PM GPU node has 4 NVIDIA A100 GPUs (40GB HBM2e RAM) and 1 AMD EPYC 7763 CPU (64 cores @ 3.5GHz). The GPUs within a PM node are connected via NVLink and are connected to the CPU via PCI-e. Internode communication is facilitated through the HPE Slingshot 11 interconnect. All numerical experiments were configured with 1 MPI rank per GPU (4 ranks per node, 16 threads per rank) and the threads of each MPI rank were bound to respective NUMA domains (of which there are 4 per PM node). GPU Batched BLAS operations for the sn-K and XC integrations were provided by the MAGMA library<sup>120,121</sup>.

As have been utilized in other studies<sup>17,24</sup>, we have included experiments with 4 test molecules, the specifics of which can be found in Tab. I. The geometries for Taxol, Olestra and Ubiquitin were taken from Ref. 17 and the Crambin geometry was taken from Ref. 24. Each of these geometries can be found in the Supplemental Information. All experiments utilize the Pople 6-31G(d) basis set<sup>126–128</sup> with Cartesian  $d$ -functions for  $\mathcal{B}$ , def2-tzvp-j fitting basis for DF-J<sup>129</sup>, and the PBE0 hybrid XC functional<sup>130</sup>. GPU evaluation of the XC functional

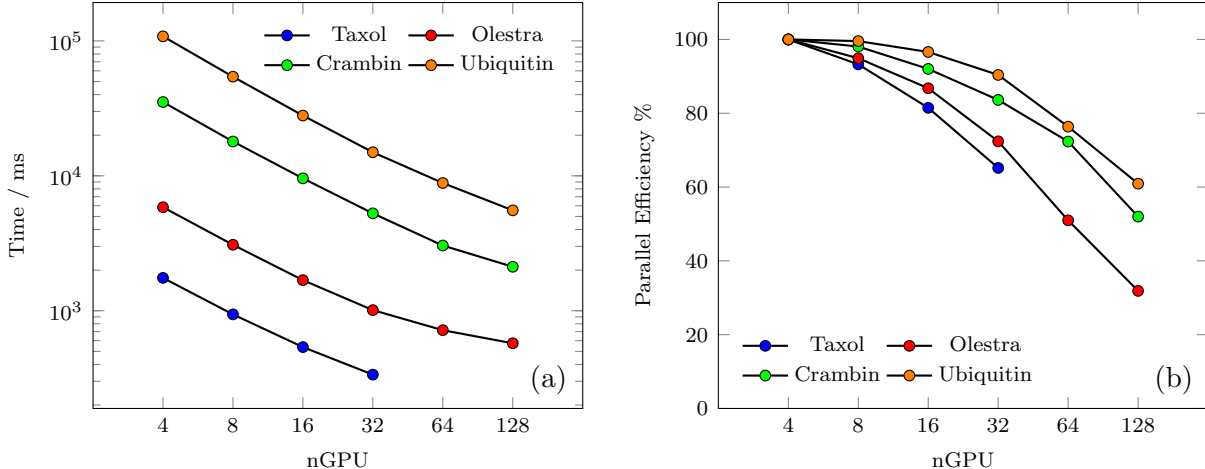


FIG. 1. Strong scaling of distributed memory Fock algorithm for **Grid1**. (a) Presents the overall timings and (b) illustrates the parallel efficiency relative to single-node performance




FIG. 2. Strong scaling of distributed memory Fock algorithm for **Grid2**. (a) Presents the overall timings and (b) illustrates the parallel efficiency relative to single-node performance

was performed using the open source **ExchCXX** library<sup>131</sup>. All calculations were performed with  $\varepsilon_B = \varepsilon_E = \varepsilon_K = 10^{-9}$  and with a maximum grid partition size of 4096. Evaluation of 3-center integrals in Eqs. (10) and (12) is screened using user-provided screener that is provided by the user via the **LibintX** API. All computational experiments used the default (Schwarz) screener of 3-center AO integrals which omitted evaluation of integral sets below the 64-bit floating point epsilon ( $\approx 2 * 10^{-16}$ ).

Figures 1 and 2 show the strong scaling performance and parallel efficiency (PE) of the overall Fock matrix construction for two different integration grids (summarized in Tab. II). The same grid is used for both XC and sn-K. PE is measured relative to single-node (4 GPU)

Molecule	DF-J	Grid1			Grid2		
		XC	sn-K	Fock	XC	sn-K	Fock
Taxol	63.2	69.4	107.5	240.0	95.1	256.9	423.6
Olestra	113.5	320.5	438.0	872.0	363.6	565.9	1,043.3
Crambin	177.8	968.6	1,512.6	2,659.0	1,043.4	2,226.4	3,448.8
Ubiquitin	378.8	2,268.0	3,674.0	6,320.8	2,443.1	5,653.1	8,465.6

TABLE III. Minimum time to solution for considered systems on 128 NVIDIA A100 GPUs. All times are given in milliseconds (ms)

performance and all results are presented with the NCCL reduction optimization discussed in Sec. II F. For the larger grid (**Grid1**), we see that excellent PE ( $> 80\%$ ) is achieved for all considered problems out to 16 GPUs and maintains a very reasonable PE ( $> 70\%$ ) out to 32 GPUs except for Taxol which exhibits 64%. This result is comparable with the achieved PE in other recent distribution memory Fock algorithms at similar GPU counts.<sup>23,24</sup> For Crambin,  $> 70\%$  PE is maintained out to 64 GPUs and achieves 52% PE at 128 GPUs. For the largest problem (Ubiquitin),  $>90\%$  PE is maintained out to 32 GPUs,  $> 75\%$  PE out to 64 GPUs and achieves 61% PE at 128 GPUs. For small problems (Taxol and Olestra) and the smaller grid (**Grid2**), strong scaling stagnation is encountered immediately. Apart from the self-evident increase in available local work exhibited by the larger test problems and grids, the smaller test problems expose bottlenecks (such as communication and load imbalance) in the relative performance and scalability of the individual Fock matrix components. The minimum time to solution for these problems, including individual component timings, are accounted in Tab. III.

Figures 3 and 4 illustrate the strong scaling of the **J**, **K** (sn-K) and  $\mathbf{V}^{xc}$  components relative to the total Fock formation for the Olestra and Ubiquitin test cases at various grid sizes. The textual annotations denote the wall-time percentage of each component relative to the overall Fock build. For the Ubiquitin test case (Fig. 3), sn-K dominates the overall Fock build by a large margin at all GPU counts, and the strong scaling of the Fock build is virtually identical to that of sn-K (as illustrated by the parallelism of the scaling plots). As such, the overall scaling behaviour of the Fock build is insensitive to the scaling behaviours of **J** and XC. The opposite case is illustrated for Olestra (Fig. 4), where the **J**-build and XC




FIG. 3. Strong scaling of individual Fock matrix components for Ubiquitin 6-31G(d)/def2-tzvp-j/PBE0 using (a) Grid1 and (b) Grid2 for the XC and sn-K integration. Annotations depict the relative contribution of each component to the overall Fock timings.




FIG. 4. Strong scaling of individual Fock matrix components for Olestra 6-31G(d)/def2-tzvp-j/PBE0 using (a) Grid1 and (b) Grid2 for the XC and sn-K integration. Annotations depict the relative contribution of each component to the overall Fock timings.

integration constitute a much more considerable portion of the overall Fock build than was exhibited for Ubiquitin, particularly at large GPU counts. As such, we can conclude that the scaling behaviours of **J** and XC have a much more measurable impact on the overall strong scaling of the Fock build for smaller systems.

The performance and scaling behaviour of the XC integration via the algorithms presented in this work (modulo the NCCL reduction) have been explored in other works.<sup>17,18</sup> In Fig. 5 we examine the scalability of the 3 primary components of the DF-J algorithm: DF-V (Eq. (10)),



DF-D (Eq. (11)) and DF-J (Eq. (12)). For Ubiquitin, the integral-driven DF-V and DF-J exhibit  $> 65\%$  PE for up to 32 GPUs where as the scaling of the TiledArray GEMV in DF-D stagnates immediately. Lack of scaling in the latter can be attributed to the lack of computational intensity exhibited by distributed GEMV. For Olestra, DF-D and DF-J exhibit immediate scaling stagnation while DF-V demonstrates  $> 50\%$  PE for up to 32 GPUs.

For the sn-K algorithm (Fig. 6), we examine the strong scaling of the local integration (`LocalWork`), shell-pair screening (`Screen`), and collective reduction (`Allreduce`) for both grid sizes. Host-to-device and device-to-host data movement timings are included in the `LocalWork` timings. For `Grid1` the overall sn-K integration is dominated by the local integration at all considered GPU counts. However, for `Grid2`, the slight deviation of the sn-K scaling from the `LocalWork` scaling can be attributed to the growing importance of the reduction and screening operations in the strong scaling limit. In Fig. 6(c), we see that for `Grid1`, the `LocalWork` timings exhibit near perfect strong scaling  $> 95\%$  out to 64 GPUs and degrades only to 88% at 128 GPUs. The `LocalWork` timings for `Grid2` and the `Screen` timings `Grid1` exhibit reasonable scalability ( $> 60\%$ ) out to 32 GPUs but quickly stagnate thereafter. These results indicate that the reuse of the XC scheduling heuristic discussed in Sec. IIF is viable for the *a priori* scheduling of sn-K integration tasks given enough distributable work.

While remaining communication free, it is important to note that the chosen work distribution scheme for sn-K requires the reevaluation of Eqs. (43) and (46) for  $\mathcal{Q}_{\text{local}}$  at each SCF step. Due to the fact that the scaling of each of these terms is linear in  $N_g$ , it is expected




FIG. 5. Component strong scaling of the distributed DF-J algorithm presented in Sec. IIB for (a) Ubiquitin and (b) Olestra 6-31G(d)/def2-tzvp-j. (c) Illustrates the parallel efficiencies of these components for both problems.




FIG. 6. Component strong scaling of the distributed sn-K algorithm presented in Alg. 2 for Ubiquitin 6-31G(d) using (a) Grid1 and (b) Grid2 numerical quadratures. (c) Illustrates the parallel efficiency of these components for both problems.




FIG. 7. Comparison of Cray MPI and NCCL Allreduce performance for Ubiquitin 6-31G(d). (a) Shows the strong scaling of Allreduce for these two libraries and (b) illustrates the speedup of NCCL over Cray MPI.

that the work associated with sn-K screening will decrease as the quadrature batches are distributed amongst independent ranks. However, in the large processor limit where there are very few tasks per rank, there remains a large prefactor associated with the formation of  $\mathcal{V}_j$  which scales  $\mathcal{O}(N_b)$ . As such the scaling stagnation of the Screen timings can be attributed to the  $\mathcal{O}(N_b)$  prefactor of the screening procedure overtaking the scalable distribution the quadrature batches in the large processor limit.

It is expected that the cost of the reduction operation will grow with growing processor counts as the connectivity of both Allreduce communication graphs grows logarithmically with the number of processors. To demonstrate the performance of the NCCL reduction

optimization, we compare the Ubiquitin XC/sn-K reduction with `MPI_Allreduce` from the Cray MPI library in Fig. 7. As we can see, NCCL improves the overall communication cost by about a factor of 4 at all GPU counts, except for 1 node at which the speedup is 15. The reason for this speed up can be attributed to the fact that the 4 GPUs on a PM are connected via NVLink, which allows for a very fast intranode reduction to take place prior to triggering internode communication over the Slingshot interconnect. This fast intranode reduction is clearly present in the 1 node case where no internode communication is triggered. When using the MPI reduction primitives, sn-K and XC exhibit a 20% reduction in parallel efficiency due to the growing relative communication cost in the strong scaling limit.

The consistent stagnation of strong scaling behaviour as problem size (e.g.  $N_g$  and/or  $N_b$ ) decreases can be attributed to two primary factors (1) communication overhead and (2) a lack of divisible work to be distributed amongst independent ranks leading to load imbalance in the large processor limit. The communication overhead (Fig. 7) represents a slow growing, nearly constant cost with processor count, which means that even in the case of perfect strong scaling for the local work of the Fock build, communication will eventually dominate the strong scaling behaviour. While these scaling inefficiencies are considerable in the large processor limit, It is important to contextualize the scaling behaviour of these smaller problems in reference to their overall wall time at this scale (Tab. III). For example, while the DF-J algorithm becomes a considerable wall time contribution for the smaller problems in the strong scaling limit, the overall wall time for this operation is  $<0.4s$  for all problems considered. Further, the 128 GPU timing for sn-K is under  $\leq 0.6s$  for Olesta and Taxol,  $<3s$  for Crambin and  $< 6s$  for Ubiquitin. Due to the small timing margins at this scale, it is unlikely that further algorithmic optimization would overcome inherent bottlenecks to significantly change the presented scaling results.

#### IV. CONCLUSIONS

In this work, we have presented a set of GPU accelerated, distributed memory algorithms for the evaluation of the performance-critical Coulomb and exact-exchange matrices for Gaussian basis DFT. For the Coulomb matrix, we developed a DF-based J-engine implementation capable of efficient execution on distributed memory heterogeneous platforms. This work extends our recent developments of algorithms for evaluation of 3-center inte-

grals on accelerated architectures.<sup>95</sup> For the exact-exchange matrix, we have developed a distributed memory sn-K algorithm which extends recently developed numerical integration methods for the Gaussian basis XC potential<sup>17</sup>. These methods were implemented in the open-source `LibintX` and `GauXC` libraries and are publicly available on GitHub<sup>98,123</sup>. We have demonstrated the absolute and strong scaling performance of the presented algorithms for a representative set of molecular systems out to 128 NVIDIA A100 GPUs on the Perlmutter supercomputer. For the largest problem considered (Ubiquitin), total Fock construction via our methods exhibited  $> 90\%$  parallel efficiency out to 32 GPUs and achieved 61% efficiency out to 128 GPUs for a total Fock matrix duration of approximately 8.5s. While degraded strong scalability was exhibited for the smaller problems considered due to communication overhead and load imbalance, the magnitude of their effects on overall scalability are magnified relative to low execution time of these operations and are small on an absolute scale.

While results of the present work indicate a promising future for distributed memory, GPU accelerated algorithms for Gaussian basis DFT, there remain several areas for exploration in future work. The efficient evaluation of the Coulomb potential described here can be combined straightforwardly with fast  $[\mathcal{O}(N)]$  approaches for evaluation of Coulomb potential, such as the fast multipole method (FMM) that can be applied in both non-periodic<sup>79</sup> and periodic<sup>132,133</sup> settings, as well as the Ewald summation<sup>134</sup> for periodic systems. While the seminumerical method was only applied to the problem of exact exchange in this work, similar approaches have been developed for treatment of correlated many-body methods as well<sup>135,136</sup>. Pursuance of many-body extensions to the presently developed distributed memory sn-K method will be the subject of future work by the authors. Finally, as has recently been demonstrated for the XC integration<sup>18</sup>, the modular nature `GauXC` library allows for rapid development of performance portable DFT methods by separating the implementation details of performance critical kernels from their inclusion in high-level workflows, thereby exposing the highest potential and flexibility for targeting current an emerging accelerator hardware with minimal developer effort. The pursuance of performance portable DF-J and sn-K algorithms will also be pursued by the authors in upcoming work.

## ACKNOWLEDGMENTS

The authors thank Norm M. Tubman for proofreading the draft of this manuscript and providing useful feedback. This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC award ASCR-ERCAP-M3946.

## Appendix A: Three-Center Coulomb Potential Integrals

In this section, we present our approach for the direct evaluation of the contributions arising from the 3c-CP integrals to the sn-K intermediate  $\mathbf{G}$  (Eq. (38)). As opposed to the McMurchie-Davidson approach taken for the DF- $J$ -engine in Sec. II B, we have adopted the Obara-Saika GTO integral recursions<sup>97,137</sup> for the evaluation of  $\mathbf{G}$ . We refer the reader for Sec. II B for notation regarding basis functions and integrals used in this section.

For a particular shell pair  $\phi_{\mathbf{a}}, \phi_{\mathbf{b}} \in \mathcal{B}$ , and grid point  $\mathbf{r}_i \in \mathcal{Q}$ , we define the auxiliary integrals,  $\Theta_{\mathbf{a}}^n$  and  $\Xi_{\mathbf{ab}}$  where

$$\Theta_{\mathbf{a}+1_q}^n = (\mathbf{R}_p - \mathbf{R}_{\mathbf{a}})_q \Theta_{\mathbf{a}}^n - (\mathbf{P} - \mathbf{r}_i)_q \Theta_{\mathbf{a}}^{n+1} + \frac{a_q}{2(\zeta_{\mathbf{a}} + \zeta_{\mathbf{b}})} \left( \Theta_{\mathbf{a}-1_q}^n - \Theta_{\mathbf{a}-1_q}^{n+1} \right), \quad (\text{A1})$$

$$\Xi_{\mathbf{a}(\mathbf{b}+1_q)} = \Xi_{(\mathbf{a}+1_q)\mathbf{b}} + (\mathbf{R}_{\mathbf{a}} - \mathbf{R}_{\mathbf{b}})_q \Xi_{\mathbf{ab}}, \quad (\text{A2})$$

and

$$\Theta_{\mathbf{0}}^n = F_n((\zeta_{\mathbf{a}} + \zeta_{\mathbf{b}})|\mathbf{R}_p - \mathbf{r}_i|^2), \quad \Xi_{\mathbf{a0}} = \Theta_{\mathbf{a}}^0, \quad A_{\mathbf{abi}} \equiv \Xi_{\mathbf{ab}}. \quad (\text{A3})$$

$F_n$  is the Boys function defined in Eq. (22). Will refer to Eqs. (A1) and (A2) as the vertical (VRR) and horizontal (HRR) OS recursions, respectively, in the following. In the canonical two-step OS algorithm, VRR intermediates are constructed and assembled into target integrals via the HRR. For brevity, the following discussion considers the case that  $\phi_{\mathbf{a}}$  and  $\phi_{\mathbf{b}}$  are primitive GTO functions. This may be generalized to contracted functions by contracting the  $\Theta$  intermediates with basis coefficients between the VRR and HRR<sup>137</sup>. There are two important aspects of these expressions for the 3c-CP

1. In contrast to their application to the evaluation of ERIs<sup>137</sup>, only the  $\Theta^0$  intermediates are required for the assembly of target integrals via the HRR. As such, for a particular  $\mathbf{a}$  and  $\mathbf{b}$ , the VRR recursion requires the evaluation of many intermediate quantities which do not contribute to the HRR.
2. Both the HRR and VRR are independent in the grid point, and are thus well suited to be parallelized in this dimension. On GPU architectures, each thread can independently perform the same VRR and HRR steps for different grid points (batch level parallelism) with minimal thread divergence. However, adopting this simple approach can be resource (e.g. register and shared memory) intensive, which can hinder the number of thread blocks that can be executed concurrently on a particular SM.




FIG. 8. A pre-order tree transversal approach for the evaluation of the VRR expressed in Eq. (A1). The nodes and edges of this tree represent intermediate integrals and their VRR connections, respectively. Grey tiles represent temporary terms which must be computed to form the  $\Theta$  values (red) required for the HRR. Only one path through the VRR tree transversal is shown for brevity.

In the following, we develop an efficient algorithm for the transversal of the VRR and HRR recursions for the sn-K method to minimize GPU resource requirements.

**VRR Implementation.** Figure 8 depicts our approach for implementing the VRR in the context of the sn-K method. First, we cast the recursion as a tree traversal, where the nodes of the tree represent the intermediate integrals and the edges between nodes represent the direct connections between the integrals via the VRR. The levels of the tree represent  $\Theta^n$  intermediates with the same total angular momentum,  $L$ , with the root of the tree representing the state  $L = 0$ . For a target  $\mathbf{a}$  and  $\mathbf{b}$ ,  $L \in [0, l_a + l_b + 1]$ , storage is kept to a minimum by performing a pre-order transversal of the VRR tree. This allows for the direct storage of  $\Theta^0$  intermediates without having to form all elements of a tree level simultaneously. For example, via the pre-order transversal depicted in Fig. 8, only 15  $\Theta$  values need to be held simultaneously in 30 single-precision GPU registers (15 x 2 FP64) to evaluate every target integral for  $l_a = l_b = 2$  as compared to 20  $\Theta$  values for a breath first search traversal. Additionally, 30 double precision intermediate values must be stored in shared memory, values that are needed for the HRR step. In general, our approach of the VRR step requires  $(l_a + l_b + 1)(l_a + l_b + 2)$  single precision GPU registers, and  $\sum_{i=1}^{l_b} (l_a + i + 1)(l_a + i + 2)$  single precision intermediate values in the GPU shared memory. The quantities must be

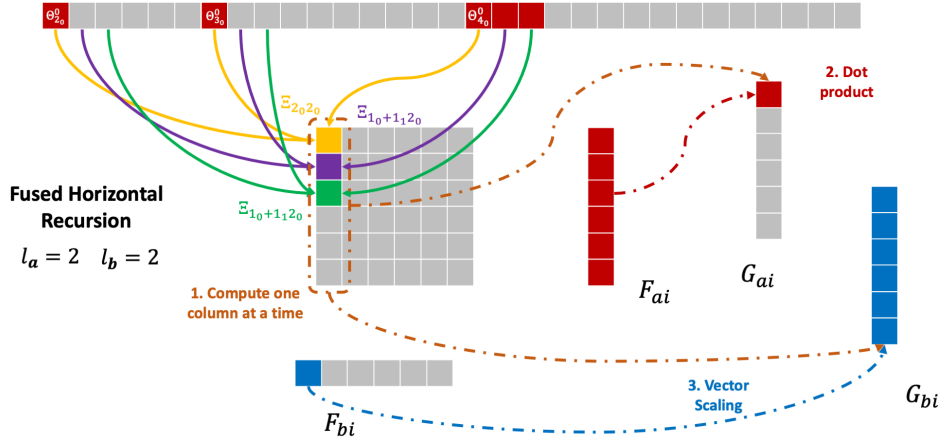


FIG. 9. The construction of the  $\Xi_{\mathbf{a}(\mathbf{b})}$  column by column and the subsequent merger with the contraction with the elements of  $F_{\mathbf{a}/\mathbf{b}i}$  to form  $\mathbf{G}$ . The full 3c-CP integral tensor is never fully materialized.

multiplied with the total number of threads within an execution block to get the total resources needed. For small angular momenta (e.g.,  $l_{\mathbf{a}/\mathbf{b}} < 2$ ), we do not utilize shared memory and keep all values in registers, including the intermediate values. This has shown to provide the best performance at the expense of an increased register count.

**Merged HRR Implementation.** Given the required  $\Theta^0$  intermediates from the VRR, Figure 9 depicts our approach for merging the HRR with the contraction with elements of  $F_{\mathbf{a}/\mathbf{b}i}$  to directly form its contributions to  $G_{\mathbf{a}/\mathbf{b}i}$  in Eq. (38). In this approach, the full 3c-CP integral tensor (which quadratically grows with angular momenta) need not be materialized in order to perform the contraction in Eq. (38). For example, for  $l_{\mathbf{a}} = l_{\mathbf{b}} = 2$ , the full 6x6 integral tensor occupies 72 single-precision (32 x 2 FP64) registers alone, which would constitute a significant bottleneck for resource utilization on modern GPU hardware, particularly when vectorizing over many grid points. In our approach, we only need to materialize a single column of the target integral at a time, which in the preceding example only requires 12 registers (a 6-fold reduction). Additionally, we need 36 single precision registers to store the values from  $F_{\mathbf{a}/\mathbf{b}i}$  and the results for  $G_{\mathbf{a}/\mathbf{b}i}$ . The values for  $F_{\mathbf{a}/\mathbf{b}i}$  can be pre-loaded to hide the latency of reading data from the GPU main memory. In general, for the HRR step our approach needs  $2 * (l_{\mathbf{a}} + 1)(l_{\mathbf{a}} + 2) + (l_{\mathbf{b}} + 1)(l_{\mathbf{a}} + 2)$  single precision GPU registers. This values must be multiplied with the total number of threads within a thread block to



obtain the final resource utilization.

For the GPU implementation, we have constructed device kernels for different angular momenta. More specifically we have constructed an in-house code generator that performs the pre-order traversal and fuses the horizontal recursion with the other contractions, and generates CUDA code. The current implementation exploits the independent execution across the grid points, each GPU thread executing the same VRR and HRR but on different grid points. For small angular momenta, we don't utilize shared memory and keep all intermediate values in registers. For larger values, we make use of shared memory, which is dynamically allocated given the angular momenta. Therefore the main strategy, given a shell pair of angular momenta  $l_{\mathbf{a}}$  and  $l_{\mathbf{b}}$ , is to first group the grid points for that specific shell pair, launch a specialized kernel on the GPU and perform the batched VRR and HRR computations, and finally obtain the contracted  $G_{\mathbf{a}/\mathbf{b}i}$  stored in device memory.

## REFERENCES

- <sup>1</sup>C. L. Janssen and I. M. Nielsen, *Parallel Computing in Quantum Chemistry* (2008).
- <sup>2</sup>W. A. De Jong, E. Bylaska, N. Govind, C. L. Janssen, K. Kowalski, T. Müller, I. M. Nielsen, H. J. van Dam, V. Veryazov, and R. Lindh, "Utilizing high performance computing for chemistry: parallel computational chemistry," *Physical Chemistry Chemical Physics* **12**, 6896–6920 (2010).
- <sup>3</sup>J. A. Calvin, C. Peng, V. Rishi, A. Kumar, and E. F. Valeev, "Many-body quantum chemistry on massively parallel computers," *Chemical Reviews* **121**, 1203–1231 (2021).
- <sup>4</sup>V. Gavini, S. Baroni, V. Blum, D. R. Bowler, A. Buccheri, J. R. Chelikowsky, S. Das, W. Dawson, P. Delugas, M. Dogan, *et al.*, "Roadmap on electronic structure codes in the exascale era," arXiv preprint arXiv:2209.12747 (2022).
- <sup>5</sup>M. S. Gordon, G. Barca, S. S. Leang, D. Poole, A. P. Rendell, J. L. Galvez Vallejo, and B. Westheimer, "Novel computer architectures and quantum chemistry," *The Journal of Physical Chemistry A* **124**, 4557–4582 (2020).
- <sup>6</sup>M. S. Gordon and T. L. Windus, "Editorial: Modern architectures and their impact on electronic structure theory," *Chemical Reviews* **120**, 9015–9020 (2020).
- <sup>7</sup>K. Yasuda, "Accelerating density functional calculations with graphics processing unit," *Journal of Chemical Theory and Computation* **4**, 1230–1236 (2008), PMID: 26631699.

- <sup>8</sup>J. Kalinowski, F. Wennmohs, and F. Neese, “Arbitrary angular momentum electron repulsion integrals with graphical processing units: Application to the resolution of identity hartree–fock method,” *Journal of Chemical Theory and Computation* **13**, 3160–3170 (2017), pMID: 28605592.
- <sup>9</sup>J. Kussmann and C. Ochsenfeld, “Employing opencl to accelerate ab initio calculations on graphics processing units,” *Journal of Chemical Theory and Computation* **13**, 2712–2716 (2017), pMID: 28561575.
- <sup>10</sup>H. Laqua, T. H. Thompson, J. Kussmann, and C. Ochsenfeld, “Highly efficient, linear-scaling seminumerical exact-exchange method for graphic processing units,” *Journal of Chemical Theory and Computation* **16**, 1456–1468 (2020), pMID: 32053375.
- <sup>11</sup>J. Kussmann and C. Ochsenfeld, “Preselective screening for linear-scaling exact exchange-gradient calculations for graphics processing units and general strong-scaling massively parallel calculations,” *Journal of Chemical Theory and Computation* **11**, 918–922 (2015).
- <sup>12</sup>J. Kussmann and C. Ochsenfeld, “Hybrid cpu/gpu integral engine for strong-scaling ab initio methods,” *Journal of Chemical Theory and Computation* **13**, 3153–3159 (2017).
- <sup>13</sup>J. Kussmann, H. Laqua, and C. Ochsenfeld, “Highly efficient resolution-of-identity density functional theory calculations on central and graphics processing units,” *Journal of Chemical Theory and Computation* **17**, 1512–1521 (2021).
- <sup>14</sup>H. Laqua, J. C. B. Dietschreit, J. Kussmann, and C. Ochsenfeld, “Accelerating hybrid density functional theory molecular dynamics simulations by seminumerical integration, resolution-of-the-identity approximation, and graphics processing units,” *Journal of Chemical Theory and Computation* **18**, 6010–6020 (2022).
- <sup>15</sup>I. S. Ufimtsev and T. J. Martinez, “Quantum chemistry on graphical processing units. 2. direct self-consistent-field implementation,” *Journal of Chemical Theory and Computation* **5**, 1004–1015 (2009), pMID: 26609609.
- <sup>16</sup>N. Luehr, A. Sisto, and T. J. Martínez, “Gaussian basis set hartree–fock, density functional theory, and beyond on gpus,” in *Electronic Structure Calculations on Graphics Processing Units* (John Wiley and Sons, Ltd, 2016) Chap. 4, pp. 67–100.
- <sup>17</sup>D. B. Williams-Young, W. A. de Jong, H. J. J. van Dam, and C. Yang, “On the efficient evaluation of the exchange correlation potential on graphics processing unit clusters,” *Frontiers in Chemistry* **8** (2020), 10.3389/fchem.2020.581058.
- <sup>18</sup>D. B. Williams-Young, A. Bagusetty, W. A. de Jong, D. Doerfler, H. J. van Dam, Álvaro

- 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* **108**, 102829 (2021).
- <sup>19</sup>H. Ahmed, D. B. Williams-Young, K. Z. Ibrahim, and C. Yang, “Performance modeling and tuning for dft calculations on heterogeneous architectures,” in *2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (2021) pp. 714–722.
- <sup>20</sup>A. Asadchev and M. S. Gordon, “New multithreaded hybrid cpu/gpu approach to hartree-fock,” *Journal of Chemical Theory and Computation* **8**, 4166–4176 (2012).
- <sup>21</sup>G. M. J. Barca, J. L. Galvez-Vallejo, D. L. Poole, A. P. Rendell, and M. S. Gordon, “High-performance, graphics processing unit-accelerated fock build algorithm,” *Journal of Chemical Theory and Computation* **16**, 7232–7238 (2020).
- <sup>22</sup>G. M. J. Barca, D. L. Poole, J. L. G. Vallejo, M. Alkan, C. Bertoni, A. P. Rendell, and M. S. Gordon, “Scaling the hartree-fock matrix build on summit,” in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis* (2020) pp. 1–14.
- <sup>23</sup>G. M. J. Barca, M. Alkan, J. L. Galvez-Vallejo, D. L. Poole, A. P. Rendell, and M. S. Gordon, “Faster self-consistent field (scf) calculations on gpu clusters,” *Journal of Chemical Theory and Computation* **17**, 7486–7503 (2021).
- <sup>24</sup>M. Manathunga, C. Jin, V. W. D. Cruzeiro, Y. Miao, D. Mu, K. Arumugam, K. Keipert, H. M. Aktulga, K. M. J. Merz, and A. W. Götz, “Harnessing the power of multi-gpu acceleration into the quantum interaction computational kernel program,” *Journal of Chemical Theory and Computation* **17**, 3955–3966 (2021).
- <sup>25</sup>M. Manathunga, Y. Miao, D. Mu, A. W. Götz, and K. M. J. Merz, “Parallel implementation of density functional theory methods in the quantum interaction computational kernel program,” *Journal of Chemical Theory and Computation* **16**, 4315–4326 (2020).
- <sup>26</sup>M. Manathunga, H. M. Aktulga, A. W. Götz, and K. M. Merz Jr, “Quantum mechanics molecular mechanics simulations on nvidia and amd graphics processing units,” *Journal of Chemical Information and Modeling* (2023).
- <sup>27</sup>S. Maintz, B. Eck, and R. Dronskowski, “Speeding up plane-wave electronic-structure calculations using graphics-processing units,” *Computer Physics Communications* **182**, 1421–1427 (2011).

- <sup>28</sup>L. Wang, Y. Wu, W. Jia, W. Gao, X. Chi, and L.-W. Wang, “Large scale plane wave pseudopotential density functional theory calculations on gpu clusters,” in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis* (2011) pp. 1–10.
- <sup>29</sup>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* **33**, 2581–2589 (2012).
- <sup>30</sup>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* **152**, 154105 (2020).
- <sup>31</sup>E. Aprà, E. J. Bylaska, W. A. de Jong, N. Govind, K. Kowalski, T. P. Straatsma, M. Vailiev, H. J. J. van Dam, Y. Alexeev, J. Anchell, V. Anisimov, F. W. Aquino, R. Atta-Fynn, J. Autschbach, N. P. Bauman, J. C. Becca, D. E. Bernholdt, K. Bhaskaran-Nair, S. Bogatko, P. Borowski, J. Boschen, J. Brabec, A. Bruner, E. Cauët, Y. Chen, G. N. Chuev, C. J. Cramer, J. Daily, M. J. O. Deegan, T. H. Dunning, M. Dupuis, K. G. Dyall, G. I. Fann, S. A. Fischer, A. Fonari, H. Früchtl, L. Gagliardi, J. Garza, N. Gawande, S. Ghosh, K. Glaesemann, A. W. Götz, J. Hammond, V. Helms, E. D. Hermes, K. Hirao, S. Hirata, M. Jacquelin, L. Jensen, B. G. Johnson, H. Jónsson, R. A. Kendall, M. Klemm, R. Kobayashi, V. Konkov, S. Krishnamoorthy, M. Krishnan, Z. Lin, R. D. Lins, R. J. Littlefield, A. J. Logsdail, K. Lopata, W. Ma, A. V. Marenich, J. Martin del Campo, D. Mejia-Rodriguez, J. E. Moore, J. M. Mullin, T. Nakajima, D. R. Nascimento, J. A. Nichols, P. J. Nichols, J. Nieplocha, A. Otero-de-la Roza, B. Palmer, A. Panyala, T. Pirojsirikul, B. Peng, R. Peverati, J. Pittner, L. Pollack, R. M. Richard, P. Sadayappan, G. C. Schatz, W. A. Shelton, D. W. Silverstein, D. M. A. Smith, T. A. Soares, D. Song, M. Swart, H. L. Taylor, G. S. Thomas, V. Tipparaju, D. G. Truhlar, K. Tsemekhman, T. Van Voorhis, A. Vázquez-Mayagoitia, P. Verma, O. Villa, A. Vishnu, K. D. Vogiatzis, D. Wang, J. H. Weare, M. J. Williamson, T. L. Windus, K. Woliński, A. T. Wong, Q. Wu, C. Yang, Q. Yu, M. Zacharias, Z. Zhang, Y. Zhao, and R. J. Harrison, “Nwchem: Past, present, and future,” *The Journal of Chemical Physics* **152**, 184102 (2020).
- <sup>32</sup>K. Wilkinson and C.-K. Skylaris, “Porting onetep to graphical processing unit-based coprocessors. 1. fft box operations,” *Journal of Computational Chemistry* **34**, 2446–2459 (2013), <https://onlinelibrary.wiley.com/doi/pdf/10.1002/jcc.23410>.

- <sup>33</sup>X. Andrade and A. Aspuru-Guzik, “Real-space density functional theory on graphical processing units: Computational approach and comparison to gaussian basis set methods,” *Journal of Chemical Theory and Computation* **9**, 4360–4373 (2013), PMID: 26589153.
- <sup>34</sup>S. Hakala, V. Havu, J. Enkovaara, and R. Nieminen, “Parallel electronic structure calculations using multiple graphics processing units (gpus),” in *Applied Parallel and Scientific Computing*, edited by P. Manninen and P. Öster (Springer Berlin Heidelberg, Berlin, Heidelberg, 2013) pp. 63–76.
- <sup>35</sup>S. Das, P. Motamarri, V. Subramanian, D. M. Rogers, and V. Gavini, “Dft-fe 1.0: A massively parallel hybrid cpu-gpu density functional theory code using finite-element discretization,” *Computer Physics Communications* **280**, 108473 (2022).
- <sup>36</sup>L. Genovese, M. Ospici, T. Deutsch, J.-F. Méhaut, A. Neelov, and S. Goedecker, “Density functional theory calculation on many-cores hybrid central processing unit-graphics processing unit architectures,” *The Journal of Chemical Physics* **131**, 034103 (2009).
- <sup>37</sup>H. van Schoot and L. Visscher, “Gpu acceleration for density functional theory with slater-type orbitals,” in *Electronic Structure Calculations on Graphics Processing Units* (John Wiley and Sons, Ltd, 2016) Chap. 5, pp. 101–114.
- <sup>38</sup>T. Yoshikawa, N. Komoto, Y. Nishimura, and H. Nakai, “Gpu-accelerated large-scale excited-state simulation based on divide-and-conquer time-dependent density-functional tight-binding,” *Journal of Computational Chemistry* **40**, 2778–2786 (2019).
- <sup>39</sup>W. P. Huhn, B. Lange, V. W. zhe Yu, M. Yoon, and V. Blum, “Gpu acceleration of all-electron electronic structure theory using localized numeric atom-centered basis functions,” *Computer Physics Communications* **254**, 107314 (2020).
- <sup>40</sup>A. E. DePrince III and J. R. Hammond, “Coupled cluster theory on graphics processing units i. the coupled cluster doubles method,” *Journal of chemical theory and computation* **7**, 1287–1295 (2011).
- <sup>41</sup>A. E. DePrince III, J. R. Hammond, and C. D. Sherrill, “Iterative coupled-cluster methods on graphics processing units,” in *Electronic Structure Calculations on Graphics Processing Units* (John Wiley & Sons, Ltd, 2016) Chap. 13, pp. 279–300.
- <sup>42</sup>A. Asadchev and M. S. Gordon, “Fast and flexible coupled cluster implementation,” *Journal of Chemical Theory and Computation* **9**, 3385–3392 (2013).
- <sup>43</sup>I. A. Kaliman and A. I. Krylov, “New algorithm for tensor contractions on multi-core cpus, gpus, and accelerators enables ccSD and eom-ccSD calculations with over 1000 basis

- functions on a single compute node,” *Journal of Computational Chemistry* **38**, 842–853 (2017).
- <sup>44</sup>B. S. Fales, E. R. Curtis, K. G. Johnson, D. Lahana, S. Seritan, Y. Wang, H. Weir, T. J. Martínez, and E. G. Hohenstein, “Performance of coupled-cluster singles and doubles on modern stream processing architectures,” *Journal of Chemical Theory and Computation* **16**, 4021–4028 (2020).
- <sup>45</sup>E. G. Hohenstein and T. J. Martínez, “Gpu acceleration of rank-reduced coupled-cluster singles and doubles,” *The Journal of Chemical Physics* **155**, 184110 (2021).
- <sup>46</sup>W. Ma, S. Krishnamoorthy, O. Villay, and K. Kowalski, “Acceleration of streamed tensor contraction expressions on gpgpu-based clusters,” in *2010 IEEE International Conference on Cluster Computing* (IEEE, 2010) pp. 207–216.
- <sup>47</sup>W. Ma, K. Bhaskaran-Nair, O. Villa, E. Aprà, A. Tumeo, S. Krishnamoorthy, and K. Kowalski, “Perturbative coupled-cluster methods on graphics processing units: Single- and multi-reference formulations,” in *Electronic Structure Calculations on Graphics Processing Units* (John Wiley & Sons, Ltd, 2016) Chap. 14, pp. 301–326, <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118670712.ch14>.
- <sup>48</sup>C. Peng, J. A. Calvin, and E. F. Valeev, “Coupled-cluster singles, doubles and perturbative triples with density fitting approximation for massively parallel heterogeneous platforms,” *International Journal of Quantum Chemistry* **119**, e25894 (2019), <https://onlinelibrary.wiley.com/doi/pdf/10.1002/qua.25894>.
- <sup>49</sup>J. V. Pototschnig, A. Papadopoulos, D. I. Lyakh, M. Repisky, L. Halbert, A. Severo Pereira Gomes, H. J. A. Jensen, and L. Visscher, “Implementation of relativistic coupled cluster theory for massively parallel gpu-accelerated computing architectures,” *Journal of Chemical Theory and Computation* **17**, 5509–5529 (2021).
- <sup>50</sup>R. Olivares-Amaya, M. A. Watson, R. G. Edgar, L. Vogt, Y. Shao, and A. Aspuru-Guzik, “Accelerating correlated quantum chemistry calculations using graphical processing units and a mixed precision matrix multiplication library,” *Journal of Chemical Theory and Computation* **6**, 135–144 (2010).
- <sup>51</sup>R. Olivares-Amaya, A. Jinich, M. A. Watson, and A. Aspuru-Guzik, “Gpu acceleration of second-order møller–plesset perturbation theory with resolution of identity,” in *Electronic Structure Calculations on Graphics Processing Units* (John Wiley & Sons, Ltd, 2016) Chap. 12, pp. 259–278,

<https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781118670712.ch12>.

- <sup>52</sup>C. Song and T. J. Martínez, “Atomic orbital-based sos-mp2 with tensor hypercontraction. i. gpu-based tensor construction and exploiting sparsity,” *The Journal of Chemical Physics* **144**, 174111 (2016).
- <sup>53</sup>S. A. Maurer, J. Kussmann, and C. Ochsenfeld, “Communication: A reduced scaling j-engine based reformulation of sos-mp2 using graphics processing units,” *The Journal of Chemical Physics* **141**, 051106 (2014).
- <sup>54</sup>D. Bykov and T. Kjaergaard, “The gpu-enabled divide-expand-consolidate ri-mp2 method (dec-ri-mp2),” *Journal of Computational Chemistry* **38**, 228–237 (2017).
- <sup>55</sup>G. M. J. Barca, S. C. McKenzie, N. J. Bloomfield, A. T. B. Gilbert, and P. M. W. Gill, “Q-mp2-os: Møller–plesset correlation energy by quadrature,” *Journal of Chemical Theory and Computation* **16**, 1568–1577 (2020).
- <sup>56</sup>G. M. Barca, J. L. G. Vallejo, D. L. Poole, M. Alkan, R. Stocks, A. P. Rendell, and M. S. Gordon, “Enabling large-scale correlated electronic structure calculations: scaling the ri-mp2 method on summit,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (2021) pp. 1–15.
- <sup>57</sup>E. G. Hohenstein, N. Luehr, I. S. Ufimtsev, and T. J. Martínez, “An atomic orbital-based formulation of the complete active space self-consistent field method on graphical processing units,” *The Journal of Chemical Physics* **142**, 224103 (2015).
- <sup>58</sup>B. S. Fales and T. J. Martínez, “Efficient treatment of large active spaces through multi-gpu parallel implementation of direct configuration interaction,” *Journal of Chemical Theory and Computation* **16**, 1586–1596 (2020).
- <sup>59</sup>J. W. Mullinax, E. Maradzike, L. N. Koulias, M. Mostafanejad, E. Epifanovsky, G. Gidofalvi, and A. E. DePrince III, “Heterogeneous cpu+ gpu algorithm for variational two-electron reduced-density matrix-driven complete active-space self-consistent field theory,” *Journal of chemical theory and computation* **15**, 6164–6178 (2019).
- <sup>60</sup>T. Straatsma, R. Broer, S. Faraji, R. Havenith, L. A. Suarez, R. Kathir, M. Wibowo, and C. De Graaf, “Gronor: Massively parallel and gpu-accelerated non-orthogonal configuration interaction for large molecular systems,” *The Journal of Chemical Physics* **152**, 064111 (2020).
- <sup>61</sup>P. Maris, C. Yang, D. Oryspayev, and B. Cook, “Accelerating an iterative eigensolver for nuclear structure configuration interaction calculations on gpus using openacc,” *Journal*

- of Computational Science **59**, 101554 (2022).
- <sup>62</sup>D. Kothe, S. Lee, and I. Qualters, “Exascale computing in the united states,” *Computing in Science & Engineering* **21**, 17–29 (2019).
- <sup>63</sup>F. Alexander, A. Almgren, J. Bell, A. Bhattacharjee, J. Chen, P. Colella, D. Daniel, J. DeSlippe, L. Diachin, E. Draeger, A. Dubey, T. Dunning, T. Evans, I. Foster, M. Francois, T. Germann, M. Gordon, S. Habib, M. Halappanavar, S. Hamilton, W. Hart, Z. (Henry) Huang, A. Hungerford, D. Kasen, P. R. C. Kent, T. Kolev, D. B. Kothe, A. Kronfeld, Y. Luo, P. Mackenzie, D. McCallen, B. Messer, S. Mniszewski, C. Oehmen, A. Perazzo, D. Perez, D. Richards, W. J. Rider, R. Rieben, K. Roche, A. Siegel, M. Sprague, C. Steefel, R. Stevens, M. Syamlal, M. Taylor, J. Turner, J.-L. Vay, A. F. Voter, T. L. Windus, and K. Yelick, “Exascale applications: skin in the game,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **378**, 20190056 (2020).
- <sup>64</sup>S. Ashby, P. Beckman, J. Chen, P. Colella, B. Collins, D. Crawford, J. Dongarra, D. Kothe, R. Lusk, P. Messina, *et al.*, “The opportunities and challenges of exascale computing,” Summary Report of the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee , 1–77 (2010).
- <sup>65</sup>S. Amarasinghe, M. Hall, R. Lethin, K. Pingali, D. Quinlan, V. Sarkar, J. Shalf, R. Lucas, K. Yelick, P. Balanji, *et al.*, “Exascale programming challenges,” in *Proceedings of the Workshop on Exascale Programming Challenges, Marina del Rey, CA, USA. US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR)* (2011).
- <sup>66</sup>J. A. Calvin, C. A. Lewis, and E. F. Valeev, “Scalable task-based algorithm for multiplication of block-rank-sparse matrices,” in *Proceedings of the 5th Workshop on Irregular Applications: Architectures and Algorithms* (2015) pp. 1–8.
- <sup>67</sup>J. A. Calvin and E. F. Valeev, “Task-based algorithm for matrix multiplication: A step towards block-sparse tensor computing,” arXiv preprint arXiv:1504.05046 (2015).
- <sup>68</sup>D. I. Lyakh, “An efficient tensor transpose algorithm for multicore cpu, intel xeon phi, and nvidia tesla gpu,” *Computer Physics Communications* **189**, 84–91 (2015).
- <sup>69</sup>J. Kim, A. Sukumaran-Rajam, V. Thumma, S. Krishnamoorthy, A. Panyala, L.-N. Pouchet, A. Rountev, and P. Sadayappan, “A code generator for high-performance tensor contractions on gpus,” in *2019 IEEE/ACM International Symposium on Code Generation*



*and Optimization (CGO)* (IEEE, 2019) pp. 85–95.

- <sup>70</sup>D. T. Popovici, A. Canning, Z. Zhao, L.-W. Wang, and J. Shalf, “A systematic approach to improving data locality across fourier transforms and linear algebra operations,” in *Proceedings of the ACM International Conference on Supercomputing* (2021) pp. 329–341.
- <sup>71</sup>A. Ayala, S. Tomov, A. Haidar, and J. Dongarra, “heffte: Highly efficient fft for exascale,” in *Computational Science–ICCS 2020: 20th International Conference, Amsterdam, The Netherlands, June 3–5, 2020, Proceedings, Part I* (Springer, 2020) pp. 262–275.
- <sup>72</sup>F. Franchetti, D. G. Spampinato, A. Kulkarni, D. T. Popovici, T. M. Low, M. Franusich, A. Canning, P. McCorquodale, B. Van Straalen, and P. Colella, “Fftx and spectralpack: A first look,” in *2018 IEEE 25th International Conference on High Performance Computing Workshops (HiPCW)* (IEEE, 2018) pp. 18–27.
- <sup>73</sup>N. Luehr, I. S. Ufimtsev, and T. J. Martinez, “Dynamic precision for electron repulsion integral evaluation on graphical processing units (gpu),” *Journal of Chemical Theory and Computation* **7**, 949–954 (2011).
- <sup>74</sup>A. Asadchev, V. Allada, J. Felder, B. M. Bode, M. S. Gordon, and T. L. Windus, “Uncontracted rys quadrature implementation of up to g functions on graphical processing units,” *Journal of Chemical Theory and Computation* **6**, 696–704 (2010), pMID: 26613300.
- <sup>75</sup>Y. Miao and K. M. Merz, “Acceleration of electron repulsion integral evaluation on graphics processing units via use of recurrence relations,” *Journal of Chemical Theory and Computation* **9**, 965–976 (2013), pMID: 26588740.
- <sup>76</sup>Y. Miao and K. M. J. Merz, “Acceleration of high angular momentum electron repulsion integrals and integral derivatives on graphics processing units,” *Journal of Chemical Theory and Computation* **11**, 1449–1462 (2015).
- <sup>77</sup>K. G. Johnson, S. Mirchandaney, E. Hoag, A. Heirich, A. Aiken, and T. J. Martínez, “Multinode multi-gpu two-electron integrals: Code generation using the regent language,” *Journal of Chemical Theory and Computation* **18**, 6522–6536 (2022).
- <sup>78</sup>L. Greengard and V. Rokhlin, “A fast algorithm for particle simulations,” *Journal of Computational Physics* **73**, 325–348 (1987).
- <sup>79</sup>C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, “The continuous fast multipole method,” *Chemical physics letters* **230**, 8–16 (1994).
- <sup>80</sup>G. Reza Ahmadi and J. Almlöf, “The Coulomb operator in a Gaussian product basis,”

- Chemical Physics Letters **246**, 364–370 (1995).
- <sup>81</sup>C. A. White and M. Head-Gordon, “A  $J$  matrix engine for density functional theory calculations,” The Journal of Chemical Physics **104**, 2620–2629 (1996).
- <sup>82</sup>T. R. Adams, R. D. Adamson, and P. M. W. Gill, “A tensor approach to two-electron matrix elements,” The Journal of Chemical Physics **107**, 124–131 (1997).
- <sup>83</sup>Y. Shao and M. Head-Gordon, “An improved  $J$  matrix engine for density functional theory calculations,” Chemical Physics Letters **323**, 425–433 (2000).
- <sup>84</sup>J. L. Whitten, “Coulombic potential energy integrals and approximations,” The Journal of chemical physics **58**, 4496–4501 (1973).
- <sup>85</sup>O. Vahtras, J. Almlöf, and M. Feyereisen, “Integral approximations for LCAO-SCF calculations,” Chemical Physics Letters **213**, 514–518 (1993).
- <sup>86</sup>R. A. Friesner, “Solution of self-consistent field electronic structure equations by a pseudospectral method,” Chemical Physics Letters **116**, 39–43 (1985).
- <sup>87</sup>M. N. Ringnalda, M. Belhadj, and R. A. Friesner, “Pseudospectral hartree–fock theory: Applications and algorithmic improvements,” The Journal of Chemical Physics **93**, 3397–3407 (1990).
- <sup>88</sup>F. Neese, F. Wennmohs, A. Hansen, and U. Becker, “Efficient, approximate and parallel hartree–fock and hybrid dft calculations. a ‘chain-of-spheres’ algorithm for the hartree–fock exchange,” Chemical Physics **356**, 98–109 (2009).
- <sup>89</sup>H. Laqua, J. Kussmann, and C. Ochsenfeld, “Accelerating seminumerical fock-exchange calculations using mixed single- and double-precision arithmetic,” The Journal of Chemical Physics **154**, 214116 (2021).
- <sup>90</sup>A. Szabo and N. S. Ostlund, *Modern quantum chemistry: introduction to advanced electronic structure theory* (Courier Corporation, 2012).
- <sup>91</sup>R. G. Parr and Y. Weitao, *Density-Functional Theory of Atoms and Molecules* (Oxford University Press, 1995).
- <sup>92</sup>R. Ahlrichs, “Efficient evaluation of three-center two-electron integrals over Gaussian functions,” Physical chemistry chemical physics : PCCP **6**, 5119 (2004).
- <sup>93</sup>D. S. Hollman, H. F. Schaefer, and E. F. Valeev, “A tight distance-dependent estimator for screening three-center Coulomb integrals over Gaussian basis functions,” The Journal of chemical physics **142**, 154106 (2015).
- <sup>94</sup>E. F. Valeev and T. Shiozaki, “Comment on “A tight distance-dependent estimator for

- screening three-center Coulomb integrals over Gaussian basis functions” [J. Chem. Phys. 142, 154106 (2015)],” The Journal of chemical physics **153**, 097101 (2020).
- <sup>95</sup>A. Asadchev and E. F. Valeev, “Memory-efficient recursive evaluation of 3-center gaussian integrals,” Journal of Chemical Theory and Computation **0**, null (2023), <https://doi.org/10.1021/acs.jctc.2c00995>.
- <sup>96</sup>L. E. McMurchie and E. R. Davidson, “One- and two-electron integrals over cartesian gaussian functions,” Journal of Computational Physics **26**, 218–231 (1978).
- <sup>97</sup>S. Obara and A. Saika, “Efficient recursive computation of molecular integrals over cartesian gaussian functions,” The Journal of chemical physics **84**, 3963–3974 (1986).
- <sup>98</sup>“LibintX,” <https://github.com/ValeevGroup/LibintX>.
- <sup>99</sup>J. Calvin, E. Valeev, C. Peng, D. Lewis, D. Williams-Young, K. Pierce, R. Richard, A. Asadchev, X. Wang, J. M. Waldrop, J. Zhang, M. Ambroise, and samslattery, “Valeev-group/tiledarray: 1.0.0,” (2020).
- <sup>100</sup>C. A. Lewis, J. A. Calvin, and E. F. Valeev, “Clustered low-rank tensor format: Introduction and application to fast construction of Hartree–Fock exchange,” Journal of Chemical Theory and Computation **12**, 5868–5880 (2016).
- <sup>101</sup>J. A. Pople, P. M. Gill, and B. G. Johnson, “Kohn–sham density-functional theory within a finite basis set,” Chemical physics letters **199**, 557–560 (1992).
- <sup>102</sup>A. M. Burow and M. Sierka, “Linear scaling hierarchical integration scheme for the exchange-correlation term in molecular and periodic systems,” Journal of Chemical Theory and Computation **7**, 3097–3104 (2011).
- <sup>103</sup>A. Petrone, D. B. Williams-Young, S. Sun, T. F. Stetina, and X. Li, “An efficient implementation of two-component relativistic density functional theory with torque-free auxiliary variables,” The European Physical Journal B **91**, 1–14 (2018).
- <sup>104</sup>A. Becke, “A multicenter numerical integration scheme for polyatomic molecules,” The Journal of Chemical Physics **88**, 2547–2553 (1988).
- <sup>105</sup>R. E. Stratmann, G. E. Scuseria, and M. J. Frisch, “Achieving linear scaling in exchange-correlation density functional quadratures,” Chemical Physics Letters **257**, 213–223 (1996).
- <sup>106</sup>H. Laqua, J. Kussmann, and C. Ochsenfeld, “An improved molecular partitioning scheme for numerical quadratures in density functional theory,” The Journal of Chemical Physics **149**, 204111 (2018).

- <sup>107</sup>M. E. Mura and P. J. Knowles, “Improved radial grids for quadrature in molecular density-functional calculations,” *The Journal of Chemical Physics* **104**, 9848–9858 (1996).
- <sup>108</sup>C. W. Murray, N. C. Handy, and G. J. Laming, “Quadrature schemes for integrals of density functional theory,” *Molecular Physics* **78**, 997–1014 (1993).
- <sup>109</sup>O. Treutler and R. Ahlrichs, “Efficient molecular numerical integration schemes,” *The Journal of Chemical Physics* **102**, 346–354 (1995).
- <sup>110</sup>P. M. W. Gill and S.-H. Chien, “Radial quadrature for multiexponential integrands,” *Journal of Computational Chemistry* **24**, 732–740 (2003).
- <sup>111</sup>S.-H. Chien and P. M. W. Gill, “Sg-0: A small standard grid for dft quadrature on large systems,” *Journal of Computational Chemistry* **27**, 730–739 (2006).
- <sup>112</sup>P. M. Gill, B. G. Johnson, and J. A. Pople, “A standard grid for density functional calculations,” *Chemical Physics Letters* **209**, 506 – 512 (1993).
- <sup>113</sup>V. Lebedev, “Quadratures on a sphere,” *USSR Computational Mathematics and Mathematical Physics* **16**, 10 – 24 (1976).
- <sup>114</sup>F. Egidi, S. Sun, J. J. Goings, G. Scalmani, M. J. Frisch, and X. Li, “Two-component noncollinear time-dependent spin density functional theory for excited state calculations,” *Journal of Chemical Theory and Computation* **13**, 2591–2603 (2017).
- <sup>115</sup>Evaluation of the density and its gradient actually proceeds as a batched-DOT as discussed in Sec. II E.
- <sup>116</sup>W. Kohn, “Analytic properties of bloch waves and wannier functions,” *Physical Review* **115**, 809–821 (1959).
- <sup>117</sup>B. Helmich-Paris, B. de Souza, F. Neese, and R. Izsák, “An improved chain of spheres for exchange algorithm,” *The Journal of Chemical Physics* **155**, 104109 (2021).
- <sup>118</sup>T. H. Thompson and C. Ochsenfeld, “Integral partition bounds for fast and effective screening of general one-, two-, and many-electron integrals,” *The Journal of Chemical Physics* **150**, 044101 (2019).
- <sup>119</sup>P. Sanders, K. Mehlhorn, M. Dietzfelbinger, and R. Dementiev, *Sequential and Parallel Algorithms and Data Structures* (Springer, 2019).
- <sup>120</sup>A. Haidar, T. Dong, P. Luszczek, S. Tomov, and J. Dongarra, “Batched matrix computations on hardware accelerators based on gpus,” *The International Journal of High Performance Computing Applications* **29**, 193–208 (2015).
- <sup>121</sup>A. Abdelfattah, A. Haidar, S. Tomov, and J. Dongarra, “Performance, design, and

- autotuning of batched gemm for gpus,” in *High Performance Computing*, edited by J. M. Kunkel, P. Balaji, and J. Dongarra (Springer International Publishing, 2016) pp. 21–38.
- <sup>122</sup>“NCCL,” <https://github.com/NVIDIA/ncc1>.
- <sup>123</sup>“GauXC,” <https://github.com/wavefunction91/GauXC>.
- <sup>124</sup>K. Kowalski, R. Bair, N. P. Bauman, J. S. Boschen, E. J. Bylaska, J. Daily, W. A. de Jong, T. J. Dunning, N. Govind, R. J. Harrison, M. Keçeli, K. Keipert, S. Krishnamoorthy, S. Kumar, E. Mutlu, B. Palmer, A. Panyala, B. Peng, R. M. Richard, T. P. Straatsma, P. Sushko, E. F. Valeev, M. Valiev, H. J. J. van Dam, J. M. Waldrop, D. B. Williams-Young, C. Yang, M. Zalewski, and T. L. Windus, “From nwchem to nwchemex: Evolving with the computational chemistry landscape,” *Chemical Reviews* **121**, 4962–4998 (2021).
- <sup>125</sup>C. Peng, C. A. Lewis, X. Wang, M. C. Clement, K. Pierce, V. Rishi, F. Pavošević, S. Slattery, J. Zhang, N. Teke, A. Kumar, C. Masteran, A. Asadchev, J. A. Calvin, and E. F. Valeev, “Massively Parallel Quantum Chemistry: A high-performance research platform for electronic structure,” *The Journal of chemical physics* **153**, 044120 (2020).
- <sup>126</sup>R. Ditchfield, W. J. Hehre, and J. A. Pople, “Self-consistent molecular-orbital methods. ix. an extended gaussian-type basis for molecular-orbital studies of organic molecules,” *J. Chem. Phys.* **54** (1971), 10.1063/1.1674902.
- <sup>127</sup>M. M. Francl, W. J. Pietro, W. J. Hehre, J. S. Binkley, M. S. Gordon, D. J. DeFrees, and J. A. Pople, “Self-consistent molecular orbital methods. xxiii. a polarization-type basis set for second-row elements,” *J. Chem. Phys.* **77** (1982), 10.1063/1.444267.
- <sup>128</sup>M. S. Gordon, J. S. Binkley, J. A. Pople, W. J. Pietro, and W. J. Hehre, “Self-consistent molecular-orbital methods. 22. small split-valence basis sets for second-row elements,” *J. Am. Chem. Soc.* **104** (1982), 10.1021/ja00374a017.
- <sup>129</sup>F. Weigend, M. Häser, H. Patzelt, and R. Ahlrichs, “Ri-mp2: optimized auxiliary basis sets and demonstration of efficiency,” *Chem. Phys. Lett.* **294**, 143–152 (1998).
- <sup>130</sup>C. Adamo and V. Barone, “Toward reliable density functional methods without adjustable parameters: The pbe0 model,” *The Journal of Chemical Physics* **110**, 6158–6170 (1999), <https://doi.org/10.1063/1.478522>.
- <sup>131</sup>“ExchCXX,” <https://github.com/wavefunction91/ExchCXX>.
- <sup>132</sup>M. Challacombe, C. White, and M. Head-Gordon, “Periodic boundary conditions and the fast multipole method,” *The Journal of chemical physics* **107**, 10131–10140 (1997).
- <sup>133</sup>K. N. Kudin and G. E. Scuseria, “A fast multipole method for periodic systems with

- arbitrary unit cell geometries,” *Chemical physics letters* **283**, 61–68 (1998).
- <sup>134</sup>P. P. Ewald, “Die Berechnung optischer und elektrostatischer Gitterpotentiale,” *Annalen der Physik* **369**, 253–287 (1921).
- <sup>135</sup>A. K. Dutta, F. Neese, and R. Izsák, “Accelerating the coupled-cluster singles and doubles method using the chain-of-sphere approximation,” *Molecular Physics* **116**, 1428–1434 (2018).
- <sup>136</sup>R. B. Murphy, M. D. Beachy, R. A. Friesner, and M. N. Ringnalda, “Pseudospectral localized molecular-orbital-plesset methods: Theory and calculation of conformational energies,” *The Journal of chemical physics* **103**, 1481–1490 (1995).
- <sup>137</sup>P. M. Gill, M. Head-Gordon, and J. A. Pople, “An efficient algorithm for the generation of two-electron repulsion integrals over gaussian basis functions,” *International Journal of Quantum Chemistry* **36**, 269–280 (1989).