LLNL-JRNL-609929

LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Tensor hypercontraction: A universal technique for the resolution of matrix elements of local, finite-range, N-body potentials in many-body quantum problems

R. Parrish, E. Hoenstein, N. Schunck, D. Sherrill, T. Martinez

January 11, 2013

Physical Review Letters

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

Tensor hypercontraction: A universal technique for the resolution of matrix elements of local, finite-range N-body potentials in many-body quantum problems

Robert M. Parrish,[1] Edward G. Hohenstein,[2,3] Nicolas F. Schunck,[4,*] C. David Sherrill,[1,†] and Todd J. Martínez[2,3,‡]

[1] *Center for Computational Molecular Science and Technology, School of Chemistry and Biochemistry, School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0400, United States*

[2] *Department of Chemistry and the PULSE Institute, Stanford University, Stanford, CA 94305*

[3] *SLAC National Accelerator Laboratory, Menlo Park, CA 94025*

[4] *Lawrence Livermore National Laboratory, Livermore, CA 94551*

(Dated: January 10, 2013)

Configuration-space matrix elements of N-body potentials arise naturally and ubiquitously in the Ritz-Galerkin solution of many-body quantum problems. For the common specialization of local, finite-range potentials, we develop the Tensor HyperContraction (THC) method, which provides a quantized renormalization of the coordinate-space form of the N-body potential, allowing for a highly separable tensor factorization of the configuration-space matrix elements. This representation allows for substantial computational savings in chemical, atomic and nuclear physics simulations, particularly with respect to difficult "exchange-like" contractions.

PACS numbers:

The physics of many-body quantum systems is often captured by local, finite-range N-body potentials $\hat{V}(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_N)$. Key examples include the Coulomb potential of atomic and molecular physics, the Yukawa potential of particle physics, and the effective Skyrme and Gogny pseudo-potentials encountered in nuclear structure and nuclear astrophysics [1].

In some finite one-particle Ritz-Galerkin basis set $\{\psi_i(\boldsymbol{r})\}$, the configuration-space representation of $\hat{V}$ is the spatial integral tensor,

$$\langle i \ldots n|\hat{V}|i' \ldots n'\rangle \equiv \int \mathrm{d}\boldsymbol{r}_1 \ldots \int \mathrm{d}\boldsymbol{r}_N \psi_i^*(\boldsymbol{r}_1) \ldots \psi_n^*(\boldsymbol{r}_N)\hat{V}(\boldsymbol{r}_1, \ldots, \boldsymbol{r}_N)\psi_{i'}(\boldsymbol{r}_1) \ldots \psi_{n'}(\boldsymbol{r}_N). \quad (1)$$

The generation, manipulation, and storage of this tensor is a major hurdle in many-body quantum simulations.

Tensor HyperContraction (THC) is a general technique developed to obtain computational savings via physically-motivated tensor decomposition of the configuration-space integral tensor (1). THC was first introduced by some of us in the context of electronic structure theory, where we have used several variants to provide *approximate* resolution of the ubiquitous electron repulsion integral (ERI) tensor, which encapsulates the Coulomb potential between electrons [2–4].

In this letter, we show that an *exact* THC decomposition (X-THC) is possible for specific choices of basis sets that are widely used, especially in nuclear and atomic physics. Aside from providing significant gains for problems involving these bases, this finding also provides the rationale for the accuracy of approximate variants of THC in more general bases.

Below, we first demonstrate the key features of the X-THC representation through the pedagogical, but entirely representative example of a 1-dimensional, 2-body problem in Cartesian coordinates using Hermite functions. The D-dimensional, N-body generalization of X-THC is then presented. Finally, the approximate LS-THC method for non-polynomial basis sets is briefly explained, and generalized to N-body potentials.

X-THC Example - For a canonical example, consider the case of a 1-dimensional ($D = 1$) problem in Cartesian coordinates, involving a finite basis of $M + 1$ Hermite functions $\{\psi_i(x)\}$ with a local 2-body ($N = 2$) potential $\hat{V} \equiv \hat{V}(x_1, x_2)$. The matrix elements are,

$$\langle ij|\hat{V}|i'j'\rangle \equiv \iint \mathrm{d}x_1 \; \mathrm{d}x_2 \; \psi_i^*(x_1)\psi_j^*(x_2)\hat{V}(x_1, x_2)\psi_{i'}(x_1)\psi_{j'}(x_2). \quad (2)$$

The first stage in the X-THC procedure is to note that all $(M+1)^2$ products $\psi_i^*(x_1)\psi_{i'}(x_1)$ are exactly spanned by an orthonormal "auxiliary" basis $\{\chi_A(x_1)\}$ consisting of $2M + 1$ Hermite functions with a slightly modified spatial range, $\chi_A(x_1) \equiv \psi_A(\sqrt{2}x_1)$. Using the Einstein summation convention here and throughout,

$$\psi_i^*(x_1)\psi_{i'}(x_1) = [ii'|A]\chi_A^*(x_1), \quad (3)$$

where,

$$[ii'|A] \equiv \int_{\mathbb{R}} \mathrm{d}x_1 \; \psi_i^*(x_1)\psi_{i'}(x_1)\chi_A(x_1). \quad (4)$$

This step is analogous to the popular Density Fitting (DF) procedure of electronic structure theory [5–8], though here it is exact due to the closure properties of the polynomial-based Hermite functions. The integrals are now given as,

$$\langle ij|\hat{V}|i'j'\rangle = [ii'|A][jj'|B]G^{AB}, \quad (5)$$

where,

$$G^{AB} \equiv \iint_{\mathbb{R}^2} \mathrm{d}x_1 \; \mathrm{d}x_2 \; \chi_A^*(x_1)\chi_B^*(x_2)\hat{V}(x_1,x_2). \quad (6)$$

Thus, the fourth-order integral tensor is expressed as a product of second- and third-order tensors. Even though we have compressed the fourth-order tensor, this representation still precludes scaling reduction in "exchange-like" terms. A canonical example of such a term is the pairing field in the Hartree-Fock-Bogoliubov theory. This is analogous to the exchange term of Hartree-Fock, but easier to generalize for N-body operators,

$$\Delta_{ij} \equiv \langle ij|\hat{V}|i'j'\rangle \kappa_{i'j'} = [ii'|A][jj'|B]G^{AB}\kappa_{i'j'}, \quad (7)$$

where κ is the pairing tensor. Despite the factorization, computing this term still scales as $\mathcal{O}(M^4) = \mathcal{O}(M^{2ND})$.

The critical step in the THC procedure is to resolve the three-index overlap integral $[ii'|A]$ to "unpin" the indices i and i' across some additional linear-scaling index P. That is, we seek a PARAFAC-like decomposition of the form $[ii'|A] = X_i^{*P} X_{i'}^P Y_A^P$ [9]. Thanks to the choice of a polynomial basis, the overlap integral is exactly integrated by the $2M+1$-node Gaussian quadrature (in this case, Gauss-Hermite) defined by the nodes and weights $\{< x_P, w_P >\}$ [10]. Therefore, the quadrature grid index provides the natural PARAFC decomposition of the overlap integral,

$$[ii'|A] = w_P \psi_i^*(x_P)\psi_{i'}(x_P)\chi_A(x_P). \quad (8)$$

This resolution is reminiscent of the discrete variable representation [11–13] or pseudospectral [14] techniques of chemical physics. Defining $X_i^{*P} \equiv \psi_i^*(x_P)$ and $Y_A^P \equiv w_P\chi_A(x_P)$, we can form the intermediate object,

$$Z^{PQ} \equiv Y_A^P Y_B^Q G^{AB}. \quad (9)$$

The full integral (2) is thus expressed as,

$$\langle ij|\hat{V}|i'j'\rangle = X_i^{*P} X_j^{*Q} Z^{PQ} X_{i'}^P X_{j'}^Q. \quad (10)$$

This X-THC representation of the integral tensor is the key for the exact $\mathcal{O}(M^3) = \mathcal{O}(M^{ND+1})$ treatment of the pairing term, via several intermediate summations, indicated here by brackets for clarity,

$$\begin{aligned}\Delta_{ij} &= X_i^{*P} X_j^{*Q} Z^{PQ} X_{i'}^P X_{j'}^Q \kappa_{i'j'} \\ &= X_i^{*P}\left[X_j^{*Q}\left[Z^{PQ}\left[X_{i'}^P\left[X_{j'}^Q\kappa_{i'j'}\right]\right]\right]\right]. \quad (11)\end{aligned}$$

Interpretation - The Z operator appears here as a mathematical intermediate, but actually is a rich physical object: it is a quantized renormalization of the coordinate-space representation of the potential operator $\hat{V}$. To see this, we first consider the continuous, renormalized potential operator $\bar{V}$, defined as,

$$\bar{V}(x_1,x_2) \equiv \chi_A(x_1)\chi_B(x_2)G^{AB}. \quad (12)$$

This operator is not equivalent to the original in physical space, i.e., $\hat{V}(x_1,x_2) \neq \bar{V}(x_1,x_2)$, yet the matrix elements of both operators are identical, i.e., $\langle ij|\hat{V}|i'j'\rangle = \langle ij|\bar{V}|i'j'\rangle$. The renormalized operator is simply the raw operator $\hat{V}$ with all components outside of the finite product space $\psi_i^*(x_1)\psi_{i'}(x_1) \Leftrightarrow \chi_A(x_1)$ projected out in each coordinate. This projection is serendipitous: the coordinate-space integrand involving $\bar{V}$ and the products of basis functions is exactly resolved by the Gaussian quadrature for the auxiliary basis, while the corresponding integrand for $\hat{V}$ is not exact under any finite quadrature. Applying the Gaussian quadrature, we quantize the renormalized operator $\bar{V}$ to produce the operator $\tilde{V}$, adding quadrature weights to account for the spatial contribution of each point,

$$\tilde{V}(x_1,x_2) \equiv w_P w_Q \delta(x_1 - x_P)\delta(x_2 - x_Q)\bar{V}(x_1,x_2). \quad (13)$$

As with $\bar{V}$, the matrix elements of $\tilde{V}$ are identical to those of $\hat{V}$. Integrating $\tilde{V}$ instead of $\hat{V}$ naturally exposes the X-THC factorization,

$$\begin{aligned}&\langle ij|\hat{V}|i'j'\rangle \overset{\{\psi_i\}}{=} \langle ij|\tilde{V}|i'j'\rangle \\ &= \iint \mathrm{d}x_1 \; \mathrm{d}x_2 \; \psi_i^*(x_1)\psi_j^*(x_2)\tilde{V}(x_1,x_2)\psi_{i'}(x_1)\psi_{j'}(x_2) \\ &\qquad = X_i^{*P} X_j^{*Q} Z^{PQ} X_{i'}^P X_{j'}^Q. \quad (14)\end{aligned}$$

Here, the elements Z^{PQ} are simply the quantized values of the renormalized potential, with the weights rolled in, i.e., $Z^{PQ} = w_P w_Q \bar{V}(x_P, x_Q)$. An example involving a Gaussian potential in Hermite functions is shown in Figure 1. The renormalized potential (right) clearly shows the effects of projection from the raw potential (left). The locations of the quantization to Z^{PQ} are shown by the small crosses on the right.

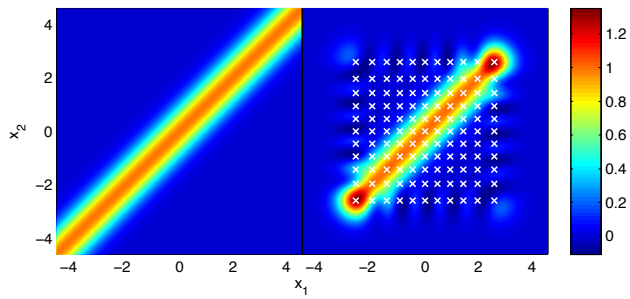

FIG. 1: (color online) Example of the X-THC process for a 1-dimensional, 2-body Gaussian potential $\hat{V}(x_1,x_2) = \exp(-x_{12}^2)$ in Hermite functions $\{\psi_i(x)\}$ up to $M = 5$. Left panel: raw $\hat{V}(x_1,x_2)$. Right panel: renormalized, quantizable $\bar{V}(x_1,x_2)$.

Generalized X-THC - The generalization of the 1-dimensional, 2-body, Hermite function example above to N-body potentials in D-dimensions and other choices of

polynomial direct-product bases is quite straightforward. We present the main results here, and leave a full derivation for the supplemental material.

For X-THC to hold, the one-particle basis must be of the D-dimensional direct-product polynomial type, i.e., $\psi_i(\boldsymbol{r}) \equiv \prod_{\mu=1}^{D} P_{i_\mu}(r_\mu) v_\mu(r_\mu)$. In each dimension μ, P_{i_μ} is a polynomial of up to degree i_μ, and v_μ is an arbitrary weight function (analogous to the Gaussian term in the Hermite functions above). Such basis sets are widely used in atomic and nuclear many-body physics in Cartesian, cylindrical, spherical, and other coordinate systems. The presence of a basis founded on polynomials in each dimension automatically guarantees closure: For the $M_\mu + 1$ functions in the dimension μ, the span $< \psi^*_{i_\mu}(r_\mu)\psi_{i'_\mu}(r_\mu) >$ lies wholly inside a $2M_\mu + 1$-function auxiliary basis, defined by a set of polynomials orthogonal with respect to the weight $|v_\mu(r_\mu)|^2$. Additionally, all quadratic products of auxiliary functions are exactly integrated by a $2M_\mu + 1$-node Gaussian quadrature $\{< r_{P_\mu}, w_{P_\mu} >\}$.

These properties allow for the X-THC factorization,

$$\langle i \dots n|\hat{V}|i' \dots n'\rangle = X_i^{*P} \dots X_n^{*W} Z^{P\dots W} X_{i'}^{P} \dots X_{n'}^{W}, \tag{15}$$

with each X_i^{*P} being the direct product of the D underlying $X_{i_\mu}^{*P_\mu}$. $Z^{P\dots W}$ is the straightforward generalization of (9) above, though with N-body auxiliary integrals $G^{A\dots N}$. A full definition of these factors is presented in the supplemental material.

Within the X-THC representation, the entirely representative generalized pairing term, $\Delta_{i\dots n} \equiv \langle i \dots n|\hat{V}|i' \dots n'\rangle \kappa_{i'\dots n'}$, now scales as $\mathcal{O}(M_\mu^{ND+1})$, rather than $\mathcal{O}(M_\mu^{2ND})$, with no approximation or restriction on the form of the local, finite-range potential $\hat{V}$.

It is worth noting that some alternative techniques to reduce the cost of treating exchange-like terms involve approximating the potential to be direct-product separable over N_w terms, e.g., by approximating the Coulomb operator as a sum of separable Gaussians [15, 16]. This reduces the cost of forming the generalized pairing tensor to $\mathcal{O}(M_\mu^{ND+N})$. X-THC can either avoid having to approximate the potential in this w-separable manner, for example if the G^{AB} factor can be computed analytically, or can additionally capitalize on the knowledge that the potential is w-separable to produce a $\mathcal{O}(M_\mu^{ND+1})$ implementation with lower prefactor and memory requirements. A summary of the scaling reductions afforded with various technologies and local potentials is shown in Table I.

Practical Demonstration - To illustrate the numerical equivalence and practical utility of the X-THC approach, a hybrid MATLAB/C++ code was developed to produce generalized pairing fields for D-dimensional, N-body forces in Hermite functions. A complete description of the code is presented in the supplemental material.

TABLE I: Comparison of computational scalings for the pairing term, using a variety of technologies and types of local potentials.

Technology	General Local	w-Separable Local
Conventional/DF	$\mathcal{O}(M_\mu^{2ND})$	$\mathcal{O}(N_w M_\mu^{ND+N})$
X-THC	$\mathcal{O}(M_\mu^{ND+1})$	$\mathcal{O}(N_w M_\mu^{ND+1})$

First and foremost, comparison of the generalized pairing fields returned by conventional and X-THC technologies strongly indicates that X-THC is exact within the machine precision (as expected mathematically). The next consideration is timing data; the representative example of $N = 2$ and $D = 1, 2, 3$ is depicted in Figure 2. For a general local potential, X-THC is clearly several orders of magnitude faster than conventional for the largest M_μ studied here. For a w-separable potential, the scaling advantage of X-THC over conventional is smaller, and the X-THC prefactor becomes visible, particularly for large D and small N. However, w-separable X-THC always scales lower than conventional, and can therefore provide some savings for sufficiently large M_μ. Moreover, X-THC with a conventional potential is generally of the same order of expense as either technology with a w-separable potential. Thus X-THC allows for the use of a wholly general local potential with the same tractability as a w-separable potential.

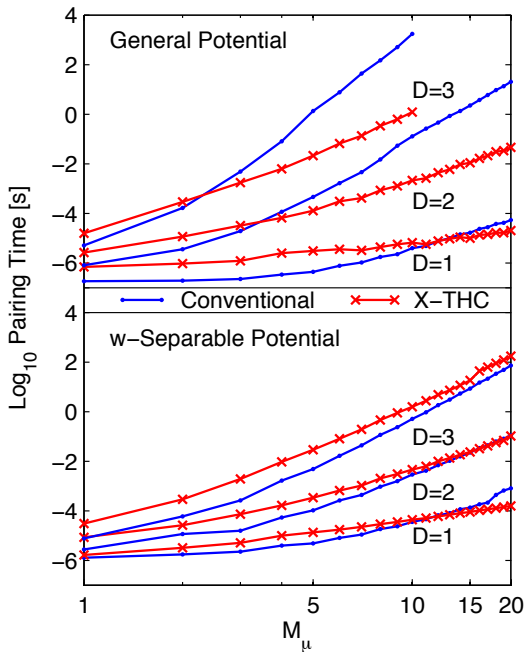

FIG. 2: (color online) Wall times for pairing tensor formation as a function of M_μ for $N = 2$ (log-log scale).

LS-THC - In some cases, more exotic bases than polynomial types are preferred. A canonical example is the

molecular physics problem, for which local atom-centered Gaussian or Slater functions are used for the basis. In these cases, the number of basis functions M is typically chosen to be proportional to the number of particles in the system, so the concepts of dimensionality or direct-product separability have little meaning.

For this choice of basis, no exact linear-span closure exists. In practice however, the locality of the basis functions and the redundancies of the remaining basis function pairs implies that a linear-span closure is *approximately* obtained. One might imagine that we could simply substitute an approximate linear-scaling quadrature grid and auxiliary basis into X-THC. However, we have found that this "pseudospectral THC" (PS-THC) approximation requires far too many points to be practically useful, in the context of the molecular ERI tensor.

A straightforward "least-squares THC" (LS-THC) modification is to choose the X operator to come from an approximate linear-scaling quadrature grid, and then to choose the Z operator to minimize the vector 2-norm of the residual in the desired integral tensor. A decomposition of the form of (15) results, where now the tuned Z largely corrects for deficiencies in the zeroth-order quadrature grid. For an N-body potential, this method yields the remakably analytical formula for the Z-operator,

$$Z^{P\ldots W} = S^{-1}_{PP'} \ldots S^{-1}_{WW'} X^{P'}_{ii'} \ldots X^{W'}_{nn'} \langle i \ldots n | \hat{V} | i' \ldots n' \rangle, \quad (16)$$

where the joint collocation is $X^P_{ii'} \equiv X^{*P}_i X^P_{i'}$ and the physical-space metric matrix is $S_{PP'} \equiv X^P_{ii'} X^{P'}_{ii'}$. In perturbation theory and coupled cluster applications in electronic structure, we have found that the LS-THC representation is remarkably accurate [3, 4].

Summary and Outlook - In this Letter, we have demonstrated that a quantized renormalization of the coordinate-space form of any local, spatial, finite-range N-body potential can be used to produce a powerful tensor factorization of the configuration-space integrals, a form denoted as Tensor HyperContraction or THC. The key advantage of THC is that it allows almost certainly optimal scaling reductions in the computation of exchange-like terms.

The determination of THC factors is a separate matter: we have shown that for polynomial direct-product bases, the X-THC factorization is exact, simple to form, and based on the $2M_\mu + 1$-node Gaussian quadratures related to the orthonormal polynomials in the basis. For either a non-direct-product basis or a direct-product basis built on non-polynomial functions, LS-THC is established as a a practical approximation, and is based on determining the approximate quantized physical-space representation of the potential by minimizing the 2-norm of the residual in the integrals, under the constraint of a preselected linear-scaling collocation grid.

In molecular physics, LS-THC separability of the ERI tensor provides an accurate and efficient treatment of many difficult terms in correlated methods. In nuclear physics, the potential for THC may be even brighter: the vast majority of basis sets are of the direct-product polynomial form, and can immediately benefit from application of X-THC, using the full span of local N-body potentials.

Acknowledgments - R.M.P. is supported by a DOE Computational Science Graduate Fellowship (Grant DE-FG02-97ER25308), particularly including a practicum rotation at Lawrence Livermore National Laboratory. This material is based on work supported by the National Science Foundation through grants to C.D.S. (Grant No. CHE-1011360) and T.J.M. (Grant No. CHE-1047577), by the Department of Defense through a grant to T.J.M. (Office of the Director of Defense Research and Engineering), and under the auspices of the US Department of Energy by the Lawrence Livermore National Laboratory (Contract DE-AC52-07NA27344). The Center for Computational Molecular Science and Technology is funded through an NSF CRIF award (Grant No. CHE-0946869) and by Georgia Tech.

* Electronic address: schunck1@llnl.gov
† Electronic address: sherrill@gatech.edu
‡ Electronic address: toddjmartinez@gmail.com

[1] M. Bender, P.-H. Heenen, and P.-G. Reinhard, Rev. Mod. Phys. **75**, 121 (2003).
[2] E. G. Hohenstein, R. M. Parrish, and T. J. Martínez, J. Chem. Phys. **137**, 044103 (2012).
[3] R. M. Parrish, E. G. Hohenstein, T. J. Martínez, and C. D. Sherrill, J. Chem. Phys. **137**, 224106 (2012).
[4] E. G. Hohenstein, R. M. Parrish, C. D. Sherill, and T. J. Martínez, J. Chem. Phys. **137**, 221101 (2012).
[5] J. L. Whitten, J. Chem. Phys. **58**, 4496 (1973).
[6] D. Gogny, Nucl. Phys. A **237**, 399 (1975).
[7] B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin, Int. J. Quantum Chem. Symp. **11**, 81 (1977).
[8] O. Vahtras, J. Almlöf, and M. W. Feyereisen, Chem. Phys. Lett. **213**, 514 (1993).
[9] J. D. Carroll and J. J. Chang, Psychometrika **31**, 279 (1970).
[10] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, Dover Publications, New York, 1964.
[11] D. O. Harris, G. G. Engerholm, and W. D. Gwinn, J. Chem. Phys. **43**, 1515 (1965).
[12] A. S. Dickinson and P. R. Certain, J. Chem. Phys. **49**, 4209 (1968).
[13] J. Lill, G. Parker, and J. Light, Chem. Phys. Lett. **89**, 483 (1982).
[14] R. A. Friesner, Chem. Phys. Lett. **116**, 39 (1985).
[15] J. Dobaczewski et al., Comp. Phys. Comm. **180**, 2361 (2009).
[16] L. M. Robledo, Phys. Rev. C **81**, 044312 (2010).

Supplemental material for: Tensor hypercontraction: A universal technique for the resolution of matrix elements of local, finite-range N-body potentials in many-body quantum problems

Robert M. Parrish,[1] Edward G. Hohenstein,[2,3] Nicolas F. Schunck,[4,*] C. David Sherrill,[1,†] and Todd J. Martínez[2,3,‡]

[1] *Center for Computational Molecular Science and Technology, School of Chemistry and Biochemistry, School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0400, United States*

[2] *Department of Chemistry and the PULSE Institute, Stanford University, Stanford, CA 94305*

[3] *SLAC National Accelerator Laboratory, Menlo Park, CA 94025*

[4] *Lawrence Livermore National Laboratory, Livermore, CA 94551*

INTRODUCTION

Below, we first discuss the high-level mathematics of the THC representation, with particular regard to conceptual understanding of the decomposition. We then provide detailed derivations of X-THC and LS-THC for N-body local potentials. We also provide a physical interpretation for the THC Z operator as a quantized renormalization of the coordinate-space potential, through a graphical example of 1-dimensional, 2-body Hermite functions in a Gaussian potential. Finally, we detail the implementation of and results from a MATLAB/C++ code developed to provide accuracy and timings data for X-THC treatment of D-dimensional, N-body local potentials in Hermite functions.

A rather large number of indices appear in this work, so we will summarize them here for clarity. The same notation is used in the primary manuscript.

Particle Number: The particle number is denoted implicitly by the presence of ellipses in relevant equations. This index ranges from 1 to N. We will often work in the special case of a 2-body potential, without loss of generality. This choice is often made for intuitive reasons: the presence of an easily-written chemists' notation in the 2-body case aids markedly in the visualization of the density fitting component of the THC procedure.

Dimension: The dimension index is denoted by μ and runs from 1 to D for each particle. In direct-product bases (e.g., X-THC), dimensionality plays a major role. In non-direct-product bases (e.g., LS-THC), dimensionality has little meaning, so we omit any reference to dimension in LS-THC (all scalings for the latter case can be found by substituting $D = 1$ into the scalings for the former).

Primary Basis: The primary single-particle basis is denoted by the indices i to n (bra), and i' to n' (ket). In a direct-product basis, i is a composite index corresponding to the underlying direct-product of 1-dimensional primary basis functions, e.g., $|i\rangle \equiv |i_x\rangle|i_y\rangle|i_z\rangle$. In a polynomial direct-product basis, the 1-dimensional primary basis functions in dimension μ range from 0 to M_μ (the zero is a consequence of the polynomial definition of the basis). In a non-direct-product basis, the full basis functions range from 1 to M.

Auxiliary Basis: Auxiliary basis indices are denoted by the indices $A, B, \ldots$. In a direct-product basis, A is a composite index corresponding to the underlying direct-product of 1-dimensional auxiliary basis functions, e.g., $|A\rangle \equiv |A_x\rangle|A_y\rangle|A_z\rangle$. In a polynomial direct-product basis, the 1-dimensional auxiliary basis functions in dimension μ range from 0 to $2M_\mu$ (the zero is a consequence of the polynomial definition of the basis). In a non-direct-product basis, the full auxiliary basis functions range from 1 to M_A, where increasing M_A increases the fidelity of the approximate density fitting procedure.

Quadrature Grid: Quadrature grid indices are denoted by the indices $P, Q, \ldots$. In a direct-product basis, P is a composite index corresponding to the union of the underlying one-dimensional quadratures, e.g., $\boldsymbol{r}_P \equiv (x_{P_x}, y_{P_y}, z_{P_z})$. In a polynomial direct-product basis, the exact quadrature in dimension μ dimension μ range from 0 to $2M_\mu$ (the zero is a consequence of the polynomial definition of the basis). In a non-direct-product basis, the LS-THC quadrature grid ranges from 1 to M_P, where increasing M_P increases the fidelity of the approximate LS-THC procedure.

For cases where two classes of indices are required to resolve a tensor element, a double subscript is used, e.g., the P-th grid point in the μ-th dimension is denoted r_{P_μ}. Also note that we use the generalized Einstein convention in this work: a repeated index on the right side of an equation is contracted over if it appears twice, or hypercontracted over if it appears more than twice, so long as the same index is not present on the left side of the same equation.

Also note that in the general discussions below, we will

try to distinguish between the case of a closed direct-product polynomial basis (for which all decompositions are exact) and the case of generic basis (for which all decompositions are approximate, but rigorously completable).

MATHEMATICAL DEVELOPMENT

Local Configuration-Space Integrals

A local, finite-range N-body potential is a special case of the most generic non-local, finite-range N-body potential, and has the coordinate-space form,

$$\hat{V}(\boldsymbol{r}_1,\ldots,\boldsymbol{r}_N) = \hat{V}(\boldsymbol{r}_1,\boldsymbol{r}'_1,\ldots,\boldsymbol{r}_N,\boldsymbol{r}'_N) \delta(\boldsymbol{r}_1-\boldsymbol{r}'_1)\ldots\delta(\boldsymbol{r}_N-\boldsymbol{r}'_N). \quad (1)$$

Here and throughout, the ellipses denote that the formalism holds for an arbitrary number of particles. The coordinate $\boldsymbol{r}$ has D components $(r_1,\ldots,r_\mu,\ldots,r_D)$. In this work, $\boldsymbol{r}$ is pedagogically understood to refer to spatial coordinates in the one-particle Hilbert space. However, there is certainly no restriction to this: the coordinate $\boldsymbol{r}$ could refer to a momentum-space or any other orthogonal representation of the one-particle Hilbert space.

We are considering the general configuration-space integral tensor for a local N-body potential in D dimensions,

$$\langle i\ldots n|\hat{V}|i'\ldots n'\rangle = \int \mathrm{d}\boldsymbol{r}_1\ldots\int \mathrm{d}\boldsymbol{r}_N \psi_i^*(\boldsymbol{r}_1)\ldots\psi_n^*(\boldsymbol{r}_N)V(\boldsymbol{r}_1,\ldots,\boldsymbol{r}_N)\psi_{i'}(\boldsymbol{r}_1)\ldots\psi_{n'}(\boldsymbol{r}_N). \quad (2)$$

The functions $\psi_i(\boldsymbol{r})$ are the finite Ritz-Galerkin basis for the one-particle Hilbert space $L_2(\mathbb{R}^D)$. If the potential were truly non-local, the coordinates in the ket would carry primes. When formulated in configuration, or spectral, space, all approaches to the quantum many-body problem involve contracting the integral tensor with an auxiliary object such as the one-body density matrix (Hartree-Fock), pairing tensor (Hartree-Fock-Bogoliubov), amplitudes (coupled-cluster), etc.

A characteristics of the configuration space representation of $\hat{V}$ is that the number of matrix elements (2) scales as $\mathcal{O}(M_\mu^{2ND})$ independently of the nature of the underlying potential (local, separable, etc.). The minimal cost of contracting the integral tensor with another tensor therefore scales as $\mathcal{O}(M_\mu^{2ND})$, in the specific case that all indices in the second tensor already lie within the integral tensor. By construction, the generic configuration space representation does not allow to take advantage of the specific properties of the potential. By contrast, in the field formulation of the quantum many-body problem, all contractions are expressed as integrals in $2D$ dimensions: the specific properties of the potential will be naturally exposed, and substantial computational simplifications may thus occur. For example, assuming a local potential reduces the dimensionality of the integrals from $2D$ to D.

Ideally, one would like to retain the formal simplicity (and abundant usage of linear algebra) of the configuration-space formulation of the quantum many-body problem while being able to take advantage of the spatial properties of the potential. The Tensor Hyper-Contraction (THC) representation of the integral tensor is a particularly efficient method to achieve this goal.

THC Stage 1: Density Fitting

This step is easily seen by considering the specific case of a 2-body potential, without loss of generality,

$$\langle ij|\hat{V}|i'j'\rangle = \iint \mathrm{d}\boldsymbol{r}_1\mathrm{d}\boldsymbol{r}_2\ \psi_i^*(\boldsymbol{r}_1)\psi_j^*(\boldsymbol{r}_2) \hat{V}(\boldsymbol{r}_1,\boldsymbol{r}_2)\psi_{i'}(\boldsymbol{r}_1)\psi_{j'}(\boldsymbol{r}_2). \quad (3)$$

Because the potential is 2-body, we can invoke the chemists' notation for pedagogical purposes [2],

$$\langle ij|\hat{V}|i'j'\rangle \equiv [ii'|\hat{V}|jj'] = \iint \mathrm{d}\boldsymbol{r}_1\mathrm{d}\boldsymbol{r}_2\ \psi_i^*(\boldsymbol{r}_1)\psi_{i'}(\boldsymbol{r}_1) \hat{V}(\boldsymbol{r}_1,\boldsymbol{r}_2)\psi_j^*(\boldsymbol{r}_2)\psi_{j'}(\boldsymbol{r}_2). \quad (4)$$

It is obvious that we can replace the bra and ket of the chemists' integral with the corresponding local products,

$$[ii'|\hat{V}|jj'] = \iint \mathrm{d}\boldsymbol{r}_1\mathrm{d}\boldsymbol{r}_2\ \rho_{ii'}(\boldsymbol{r}_1)\hat{V}(\boldsymbol{r}_1,\boldsymbol{r}_2)\rho_{jj'}(\boldsymbol{r}_2). \quad (5)$$

This substitution is, as-yet, purely formal, and offers no computational advantages. However, we may introduce an auxiliary basis $\{\chi_A(\boldsymbol{r})\}$ of the Hilbert space and expand the local products on that basis,

$$\rho_{ii'}(\boldsymbol{r}) = d^A_{ii'}\chi^*_A(\boldsymbol{r}), \quad (6)$$

leading to,

$$[ii'|\hat{V}|jj'] = d^A_{ii'}G^{AB}d^B_{jj'}, \quad (7)$$

with

$$G^{AB} = \iint \mathrm{d}\boldsymbol{r}_1\mathrm{d}\boldsymbol{r}_2\ \chi^*_A(\boldsymbol{r}_1)\hat{V}(\boldsymbol{r}_1,\boldsymbol{r}_2)\chi^*_B(\boldsymbol{r}_2). \quad (8)$$

This is the well-known Density Fitting (DF) procedure from electronic structure theory (so-named because the local products $\rho_{ii'}$ can be viewed as generalized probability densities). Technically, exact equality is only achieved in two special cases: (1) if the underlying finite primary basis has a special structure admitting strict closure in a linear-scaling auxiliary basis (we will exploit this in X-THC) or (2) in a generic basis, if the complete basis limit

is reached, which is obviously only conceptually useful. In the latter case however (corresponding to the standard DF procedure in molecular physics, and to LS-THC), the rigorous information growth established above provides for the existence of a δ-ϵ approximation, in which the error decays exponentially with respect to the linear size of the (well-chosen) auxiliary basis, and arbitrary gains in accuracy can be formally guaranteed with only a linear increase in the number of auxiliary functions. When the number of expansion coefficients in (6) is of the order of M_μ (for each dimension μ), DF successfully identifies the underlying rank structure of the integral tensor, effectively compressing the rank-$4D$ integrals $\langle ij|\hat{V}|i'j'\rangle$ to the rank-$2D$ auxiliary integrals G^{AB}.

The DF coefficients $d^A_{ii'}$ are determined in a number of ways.

- If a closed primary basis is used, an exact, linear-scaling, orthonormal auxiliary basis exists, and the DF coefficients are simply the projection of the local products onto this auxiliary basis, i.e., the overlap integrals,

$$d^A_{ii'} = [ii'|A] = \int \mathrm{d}\boldsymbol{r}_1 \; \psi^*_i(\boldsymbol{r}_1)\psi_{i'}(\boldsymbol{r}_1)\chi_A(\boldsymbol{r}_1). \tag{9}$$

  Note that these overlap integrals are rigorously direct-product separable by dimension if the underlying primary basis is direct-product separable.

- If a non-closed primary basis is used, the DF coefficients are typically determined by least-squares fitting against some positive definite 2-body operator $\hat{O}$,

$$d^A_{ii'} = \iint \mathrm{d}\boldsymbol{r}_1\mathrm{d}\boldsymbol{r}_2 \; \psi^*_i(\boldsymbol{r}_1)\psi_{i'}(\boldsymbol{r}_1)\hat{O}(\boldsymbol{r}_1,\boldsymbol{r}_2)\chi_B(\boldsymbol{r}_2) \times \left[\iint \mathrm{d}\boldsymbol{r}_1\mathrm{d}\boldsymbol{r}_2 \; \chi^*_B(\boldsymbol{r}_1)\hat{O}(\boldsymbol{r}_1,\boldsymbol{r}_2)\chi_A(\boldsymbol{r}_2)\right]^{-1}. \tag{10}$$

The DF representation often provides substantial savings in integral generation and/or storage, and also can be used to produce scaling gains for certain "Coulomb-like terms" such as, e.g., the direct 2-body term in Hartree-Fock theory, involving the contraction of the integral tensor with an auxiliary object,

$$J_{ii'} = \langle ij|\hat{V}|ij'\rangle\rho_{j'j} = [ii'|\hat{V}|jj']D_{jj'} = d^A_{ii'}\left[G^{AB}\left[d^B_{jj'}D_{jj'}\right]\right], \tag{11}$$

where ρ is the one-body density matrix. For $N = 2$, the scaling is reduced from $\mathcal{O}(M_\mu^{4D})$ to $\mathcal{O}(M_\mu^{2D+1})$. However, this term is somewhat serendipitous: the key contraction index is the compound index jj', which occurs in one piece in the factor $d^B_{jj'}$. For more generic contractions, such scaling reductions are not realized by the DF factorization, due to the ungainly "pinning" of the indices i and i' in the density fitting coefficients: The common thread between all "exchange-like" terms which cannot be factored by density fitting is that two or more of the contraction indices occur on the same side of the physicists' integral, implying that two or more of the density fitting coefficient tensors are required to remove the contraction indices. A canonical and entirely representative example is the pairing term of Hartree-Fock-Bogoliubov theory,

$$\Delta_{ij} = \langle ij|\hat{V}|i'j'\rangle\kappa_{i'j'} = d^A_{ii'}d^B_{jj'}G^{AB}\kappa_{i'j'}, \tag{12}$$

where κ is the pairing tensor. The best possible DF factorization is,

$$\Delta_{ij} = d^A_{ii'}\left[G^{AB}\left[d^B_{jj'}\kappa_{i'j'}\right]\right]. \tag{13}$$

Not only does it still scale as $\mathcal{O}(M_\mu^{4D})$, it now has also a markedly higher prefactor than conventionally. In fact, it is often better computationally to collapse the density-fitted integrals to re-form the conventional integrals before performing exchange-like contractions. In this case, the DF procedure only helps reduce integral generation and/or storage costs.

THC Stage 2: Unpinning

The key realization of THC is that the indices i and i' must be "unpinned" from each other (and from A) in the DF coefficients by direct-product expansion over some other linear-scaling decomposition index P. That is, we need a PARAFAC decomposition of the form,

$$d^A_{ii'} = X^{*P}_i X^P_i Y^P_A \tag{14}$$

The index P appears three times on the right: this is not a conventional two-factor contraction, but rather a higher-order "hypercontraction." If such a decomposition exists in linear rank P, the Y operators may be contracted with the auxiliary integrals G to form the intermediate Z operator,

$$Z^{PQ} = Y^P_A G^{AB} Y^Q_B. \tag{15}$$

This yields the THC representation,

$$\langle ij|\hat{V}|i'j'\rangle = X^{*P}_i X^{*Q}_j Z^{PQ} X^P_{i'} X^Q_{j'}. \tag{16}$$

The pairing term is now factored naturally as,

$$\Delta_{ij} = \langle ij|\hat{V}|i'j'\rangle\kappa_{i'j'} = X^{*P}_i\left[X^{*Q}_j\left[Z^{PQ}\left[X^P_{i'}\left[X^Q_{j'}\kappa_{i'j'}\right]\right]\right]\right]. \tag{17}$$

This factorization scales as $\mathcal{O}(M_\mu^{ND+1})$ for each bracket. The pre-factor comes from the number of brackets to consider, which itself is a function of the number of particles.

If the underlying basis is of the direct-product type, the direct-product separability of the factors $X_i^P = \prod_\mu^D X_{i_\mu}^{P_\mu}$ can be exploited to perform this term in $\mathcal{O}(M_\mu^{N+1})$ (with an increase in the pre-factor cost). If the basis is not of the direct-product type (implying no concept of dimensionality), this term can be performed in $\mathcal{O}(M^{N+1})$.

This finding is, as-yet, only conceptual, as we have not yet provided a physical justification or practical algorithm for the existence of the linear-scaling PARAFAC of (14). The former is quite simple. Whether the primary basis has exact closure or not, the DF coefficients can be taken to depend on the three-center overlap integrals,

$$[ii'|A] = \int \mathrm{d}\boldsymbol{r}_1 \ \psi_i^*(\boldsymbol{r}_1)\psi_{i'}(\boldsymbol{r}_1)\chi_A(\boldsymbol{r}_1). \tag{18}$$

For a closed primary basis, this is precisely the DF coefficient tensor $d_{ii'}^A = [ii'|A]$, which provides an exact decomposition. For a non-closed primary basis, this is part of the N-sided overlap-metric DF procedure, where $d_{ii'}^A \equiv [ii'|B][B|A]^{-1}$, which provides an approximate decomposition. While the number of elements in the overlap integral above is naïvely $\mathcal{O}(M_\mu^{3D})$, the local product $\psi_i^*(\boldsymbol{r}_1)\psi_{i'}(\boldsymbol{r}_1)\chi_A(\boldsymbol{r}_1)$ lives in $L_2(\mathbb{R}^D)$, and thus contains strictly $\mathcal{O}(M_\mu^D)$ nonredundant combinations of elements. Thus, a linear-scaling PARAFAC decomposition must exist.

Having established the existence of a PARAFAC of the form of (14), it remains to actually compute the decomposition by some practical algorithm. One approach is to directly apply an iterative PARAFAC algorithm to the three-center overlap integrals; this was the approach of our demonstration-level PARAFAC THC (PF-THC) procedure. However, it turns out that there exist two far more practical recipes: X-THC for closed polynomial bases and LS-THC for generic bases. Key features of both variants are the use of quadrature grids to provide the PARAFAC separability, and the presence of an analytical (i.e., non-iterative) factorization.

It is fairly obvious that a linear-scaling spatial quadrature $\{< \boldsymbol{r}_P, w_P >\}$ of the overlap integrals provides an (at-least approximate) PARAFAC decomposition,

$$[ii'|A] \overset{?}{\approx} \underbrace{\psi_i^*(\boldsymbol{r}_P)}_{X_i^{*P}} \underbrace{\psi_{i'}(\boldsymbol{r}_P)}_{X_{i'}^P} \underbrace{\chi_A(\boldsymbol{r}_P)w_P}_{Y_A^P}. \tag{19}$$

The key is that the integrand of the three-index overlap integral is naturally direct-product separable, and the invocation of quadrature explicitly exposes this over the index P. For a closed, polynomial direct-product basis with $M_\mu+1$ primary basis functions in dimension μ, there always exists an exact DF representation with $2M_\mu + 1$ auxiliary functions, and the resultant three-index overlap is thus exactly integrated by a Gaussian quadrature with $2M_\mu + 1$ nodes. Therefore, the ? symbol in the above is the equality. This is the eXact THC (X-THC) procedure discussed below. For non-closed bases, no exact quadrature exists, but we may certainly apply an approximate linear-scaling quadrature recipe. This is the idea behind the Least-Square THC (LS-THC) discussed further.

Full N-body D-dimensional X-THC Derivation

A direct-product basis founded on polynomials has the form,

$$\psi_i(\boldsymbol{r}) \equiv \prod_{\mu=1}^D P_{i_\mu}(r_\mu)v_\mu(r_\mu). \tag{20}$$

In each dimension μ, P_{i_μ} is a polynomial of up to degree i_μ, and v_μ is a weight function, typically chosen to bring the polynomial into $L_2(\mathbb{D})$, where $\mathbb{D}$ is the domain of the problem, and also to provide qualitative conformation to some *a priori* knowledge of the future shape of the wavefunction. The i_μ index runs from 0 to M_μ, so the polynomials range up to a maximum degree of M_μ. Note that these polynomials do not have to be orthogonal, though they are often defined to be so. The product of polynomials being itself a polynomial of up to order $2M_\mu$, the local product $\psi_{i_\mu}^*(r_\mu)\psi_{i'_\mu}(r_\mu)$ in dimension μ lies inside the span of the $2M_\mu + 1$ auxiliary functions,

$$\chi_{A_\mu}(r_\mu) = \tilde{P}_{A_\mu}(r_\mu)|v_\mu(r_\mu)|^2, \tag{21}$$

where $\tilde{P}_{A_\mu}$ is a polynomial of up to order $2M_\mu$, often different from P_{A_μ}. We will choose these auxiliary functions to be orthonormal for convenience (this avoids DF metric matrices).

The auxiliary functions $\chi_{A_\mu}(r_\mu)$ yield an exact DF representation for this basis,

$$\langle i \dots n|\hat{V}|i' \dots n'\rangle = [ii'|A] \dots [nn'|N]G^{A\dots N}. \tag{22}$$

The overlap integrals are separable in coordinates,

$$[ii'|A] = \prod_{\mu=1}^D [i_\mu i'_\mu|A_\mu]. \tag{23}$$

However, in general, the auxiliary potential integrals are *not* separable,

$$G^{A\dots N} \equiv \int \mathrm{d}\boldsymbol{r}_1 \ \dots \int \mathrm{d}\boldsymbol{r}_N \ \chi_A^*(\boldsymbol{r}_1)\dots\chi_N^*(\boldsymbol{r}_N)\hat{V}(\boldsymbol{r}_1,\dots,\boldsymbol{r}_N). \tag{24}$$

To produce the THC representation, it remains to find an exact quadrature for the three-index overlap integrals. Owing to the closure in the product $i_\mu i'_\mu$, any quadrature which can exactly integrate the auxiliary overlap metric $\langle B_\mu|A_\mu\rangle = \delta_{B_\mu A_\mu}$ can exactly integrate the three-index overlap integrals $[i_\mu i'_\mu|A_\mu]$, as the spans are identical in

both cases. A quadrature which can exactly integrate all quadratic products of functions based on orthogonal polynomials of up to degree $2M_\mu$ is *precisely* the definition of the $2M_\mu + 1$-node Gaussian quadrature. Regardless of the choice of weight $v_\mu(r_\mu)$, the nodes and roots of this Gaussian quadrature can always be determined efficiently by the Golub-Welsch algorithm. In the full direct-product basis, we will use the collapsed notation for the collocations of this Gaussian quadrature,

$$X_i^P = \prod_{\mu=1}^{D} X_{i_\mu}^{P_\mu}, \quad Y_A^P = \prod_{\mu=1}^{D} Y_{A_\mu}^{P_\mu}, \tag{25}$$

to save space. However, in real implementation, the direct-product separability of these two quantities is very important to achieve near-optimal scaling in formation of the X-THC factorization and subsequent utilization.

With these definitions, the X-THC Z operator reads

$$Z^{P\ldots W} = Y_A^P \ldots Y_N^W G^{A\ldots N}, \tag{26}$$

which yields the the full $\boldsymbol{N}$-body $\boldsymbol{D}$-dimensional X-THC,

$$\langle i \ldots n|\hat{V}|i' \ldots n\rangle = X_i^{*P} \ldots X_n^{*W} Z^{P\ldots W} X_{i'}^P \ldots X_{n'}^W. \tag{27}$$

The Z operator is not, in general, direct-product separable, but the factors X and Y are. In forming Z, the contraction of the Y factors with the auxiliary potential integrals would naïvely scale as $\mathcal{O}(M_\mu^{ND+D})$. However, for each Y, we can perform the transformation in one dimension at a time (e.g., replacing A_x with P_x, etc), reducing the formal scaling to $\mathcal{O}(M_\mu^{ND+1})$. Similarly, the separability of the X and Y factors is critically important to reduce the scaling of the generalized pairing term,

$$\Delta_{i\ldots n} = \langle i \ldots n|\hat{V}|i' \ldots n'\rangle \kappa_{i'\ldots n'}. \tag{28}$$

This contraction of a rank-ND tensor with the integral tensor is in fact the worst possible scenario, as far as the DF representation is concerned, since the compound contraction index involves *all* N DF coefficient tensors,

$$\Delta_{i\ldots n} = d_{ii'}^A \ldots d_{nn'}^N G^{A\ldots N} \kappa_{i'\ldots n'}. \tag{29}$$

In general, this term will always scale as $\mathcal{O}(M_\mu^{2ND})$ with both conventional and DF approaches, though the computational pre-factor is markedly higher in the latter case. Its THC factorization now reads

$$\begin{aligned} &\Delta_{i\ldots n} = \langle i \ldots n|\hat{V}|i' \ldots n'\rangle \kappa_{i'\ldots n'} \\ &= X_i^{*P} \ldots \left[X_n^{*W} \left[Z^{P\ldots W} \left[X_{i'}^P \ldots \left[X_{n'}^W \kappa_{i'\ldots n'}\right]\right]\right]\right]. \end{aligned} \tag{30}$$

Any contraction involving X here would naïvely scale as $\mathcal{O}(M_\mu^{ND+D})$. However, for each X, we can perform the transformation in one dimension at a time (e.g., replacing i'_x with P_x, etc), reducing the formal scaling to $\mathcal{O}(M_\mu^{ND+1})$.

Full LS-THC Derivation

In a finite non-direct-product non-polynomial basis with M functions, in exact arithmetic, the number of exactly non-redundant local products $\psi_i^*(\boldsymbol{r})\psi_{i'}(\boldsymbol{r})$ could be as high as M^2. Therefore, to produce an exact application of THC as in X-THC, a quadratric-scaling auxiliary basis *and* a quadratic-scaling quadrature grid would be required, leading to no computational savings. However, as discussed below (2), the asymptotic growth of the non-redundant combinations of local products must be linear. In practice, this manifests in exponential decay of the singular values of the overlap metric between local products,

$$S_{ii',jj'} \equiv \int \mathrm{d}\boldsymbol{r}_1 \; \rho_{ii'}(\vec{r}_1)\rho_{jj'}(\vec{r}_2). \tag{31}$$

Since the decay is exponential, arbitrary relative accuracy can always be achieved in a linear number of auxiliary functions, providing for the existence of the efficient and well-adopted approximate DF procedure in electronic structure theory.

To achieve a good THC factorization in these cases, it is critical to ensure that the projection on coordinate space brought about by the finite quadrature will not introduce an alias. we can insure this by requiring that a dealiased Z operator be *chosen* such that the 2-norm of the residual tensor is minimized. It turns out that an analytic form exists for such a Z, and that a separate DF step can be used to massively accelerate its formation. The results for the 2-body case have already been established in our series in the molecular physics series, so we will give the essential steps for the derivation in the N-body case here.

We assert,

$$\begin{aligned} &\langle i \ldots n|\hat{V}|i' \ldots n'\rangle \\ &\qquad \approx X_i^{*P} \ldots X_n^{*W} Z^{P\ldots W} X_{i'}^P \ldots X_{n'}^W. \end{aligned} \tag{32}$$

Equivalently (for shorthand), we define the joint collocation $X_{ii'}^P \equiv X_i^{*P} X_{i'}^P$, so the above is,

$$\langle i \ldots n|\hat{V}|i' \ldots n'\rangle \approx X_{ii'}^P \ldots X_{nn'}^W Z^{P\ldots W}. \tag{33}$$

Defining the residual tensor,

$$\epsilon_{i\ldots n}^{i'\ldots n'} = \langle i \ldots n|\hat{V}|i' \ldots n'\rangle - X_{ii'}^P \ldots X_{nn'}^W Z^{P\ldots W}. \tag{34}$$

The fidelity metric is the 2-norm of the residual tensor (squared to achieve a quadratic form, scaled for convenience),

$$O(Z^{P\ldots W}) = \frac{1}{2}\epsilon_{i\ldots n}^{i'\ldots n'}\epsilon_{i\ldots n}^{i'\ldots n'}. \tag{35}$$

We seek to minimize O under an unconstrained Z. The weak form of the minimization is the stationary condition,

$$\frac{\partial O}{\partial Z^{P'\ldots W'}} = 0, \; \forall \; P' \ldots W'. \tag{36}$$

Evaluating the derivative, this leads to,

$$X_{ii'}^{P'} \dots X_{nn'}^{W'} \epsilon_{i\dots n}^{i'\dots n'} = 0, \ \forall \ P' \dots W'. \tag{37}$$

This asserts that the projection of the residual onto the fitting subspace is zero, which is the canonical finding of least-squares. Expanding the residual,

$$X_{ii'}^{P'} \dots X_{nn'}^{W'} \Big[\langle i \dots n | \hat{V} | i' \dots n' \rangle - X_{ii'}^{P} \dots X_{nn'}^{W} Z^{P\dots W} \Big] = 0, \quad \forall \ P' \dots W'. \tag{38}$$

Moving the first term to the RHS (note the extra primes to distingish the summations in the Einstein convention),

$$X_{ii'}^{P} X_{ii'}^{P'} \dots X_{nn'}^{W} X_{nn'}^{W'} Z^{P\dots W} = X_{i'''i''''}^{P'} \dots X_{n'''n''''}^{W'} \langle i''' \dots n''' | \hat{V} | i'''' \dots n'''' \rangle, \quad \forall \ P' \dots W'. \tag{39}$$

Defining the grid metric matrices (symmetric positive definite if the collocations are full rank),

$$S_{PP'} = X_{ii'}^{P} X_{ii'}^{P'} = \Big[X_{i}^{*P} X_{i}^{*P'} \Big] \Big[X_{i'}^{P} X_{i'}^{P'} \Big]. \tag{40}$$

leads to

$$S_{PP'} \dots S_{WW'} Z^{P\dots W} = X_{ii'}^{P'} \dots X_{nn'}^{W'} \langle i \dots n | \hat{V} | i' \dots n' \rangle, \quad \forall \ P' \dots W'. \tag{41}$$

Finally, applying the inverses (or pseudoinverses, if rank problems arise in the collocation matrices),

$$Z^{P\dots W} = S_{PP'}^{-1} \dots S_{WW'}^{-1} X_{ii'}^{P'} \dots X_{nn'}^{W'} \langle i \dots n | \hat{V} | i' \dots n' \rangle. \tag{42}$$

The rate-limiting step here is the projection of the exact integrals onto the grid, forming the matrix $E^{P'\dots W'}$,

$$E^{P'\dots W'} \equiv X_{ii'}^{P'} \dots X_{nn'}^{W'} \langle i \dots n | \hat{V} | i' \dots n' \rangle. \tag{43}$$

This naïvely scales as $\mathcal{O}(M^{2ND+1})$, which is not acceptable for almost any application. However, the knowledge that the operator is local automatically provides for the existence of an accurate approximate DF representation, with any choice of fitting metric,

$$\langle i \dots n | \hat{V} | i' \dots n' \rangle \approx d_{ii'}^{A} \dots d_{nn'}^{N} G^{A\dots N}. \tag{44}$$

Here the choice of DF metric is reflected in the selection of DF coefficients $d_{ii'}^{A}$. We may substitute Eq.(44) into Eq.(43). By forming the intermediate L objects, via Coulomb-like contractions.

$$L_{A}^{P} = d_{ii'}^{A} X_{ii'}^{P'}, \tag{45}$$

the E matrix may be formed as,

$$E^{P'\dots W'} \approx L_{A}^{P'} \dots L_{N}^{W'} G^{A\dots N}. \tag{46}$$

In general, this scales as $\mathcal{O}(M^{N+1})$. However, it is worth noting that the formation of the L object scales as $\mathcal{O}(M^4)$, and is technically rate-limiting for the case $N = 2$.

Z-Operator Visualization

Recall that the X-THC Z-operators are,

$$Z^{P\dots W} = Y_{A}^{P} \dots Y_{N}^{W} G^{A\dots N} = w_{P} \dots w_{W} X_{A}^{P} \dots X_{N}^{W} G^{A\dots N}, \tag{47}$$

with the collocation $X_{A}^{P} = \chi_{A}(\boldsymbol{r}_{P})$. These are extremely interesting physical objects, as they contain the quantized renormalization of the physical-space N-body potential,

$$\hat{V}(\boldsymbol{r}_1, \dots, \boldsymbol{r}_N) \overset{\{\psi_i\}}{=} \delta(\boldsymbol{r}_1 - \boldsymbol{r}_P) \dots \delta(\boldsymbol{r}_N - \boldsymbol{r}_W) Z^{P\dots W}. \tag{48}$$

The quadrature weights in the Z-operators correct for the length scales and number of quadrature points involved, and thus are related to the quantization of the integral rather than the value of the renormalized potential. The object corresponding directly to the quantized, renormalized N-body potential is the Z-operator with the quadrature weights removed, which we denote here by $\bar{V}^{P\dots W}$,

$$\bar{V}^{P\dots W} = X_{A}^{P} \dots X_{N}^{W} G^{A\dots N}. \tag{49}$$

In fact, we can roll back one step further by generalizing the collocation locations, and extract a continuous, renormalized potential $\tilde{V}$ for the primary basis at hand,

$$\tilde{V}(\boldsymbol{r}_1, \dots, \boldsymbol{r}_N) \equiv \chi_{A}(\boldsymbol{r}_1) \dots \chi_{N}(\boldsymbol{r}_N) G^{A\dots N}. \tag{50}$$

It is easily verified that (due to the closure of the X-THC case), the spectral integral tensors for the operators $\hat{V}$ and $\tilde{V}$ are identical,

$$\langle i \dots n | \hat{V} | i' \dots n' \rangle = \langle i \dots n | \tilde{V} | i' \dots n' \rangle. \tag{51}$$

However, it is also obvious that, for a generic V, the spatial forms of these operators are not equivalent,

$$\hat{V}(\boldsymbol{r}_1, \dots, \boldsymbol{r}_N) \neq \tilde{V}(\boldsymbol{r}_1, \dots, \boldsymbol{r}_N). \tag{52}$$

In fact, $\hat{V}$ still lies in the full space, but $\tilde{V}$ lies in the finite subspace spanned by the products ii', etc. Thus, no general finite spatial quantization of $\hat{V}$ will provide an exact quadrature for the integral tensor. In contrast, because we know that $\tilde{V}$ lives in ii' (and ii' does too),

we can quantize the potential $\tilde{V}$ in an exact manner by applying the Gaussian quadrature.

To provide some physical insight into the X-THC renormalization and quantization processes, we consider the case of a 1-dimensional, 2-body Gaussian force,

$$\hat{V}(x_1, x_2) = \alpha_w \exp(-\beta_w x_{12}^2). \tag{53}$$

For sake of simplicity, we consider $\alpha_w = 1$ and $\beta_w = 1$. We take the basis to be Hermite functions of up to order M with a spatial range of $b_x = 1$.

Plots for $M = 5$ and $M = 9$ are shown in Figures 1 and 2, respectively. In each case, the top panel shows the Gaussian potential $\hat{V}$, which is obviously central and invariant to the basis used. The middle panel shows the renormalized operator $\tilde{V}$, which is overall qualitatively similar to $\hat{V}$, but differs in several quantitative details. Most obviously, the projection through ii' (via A) attenuates the potential for $x_1 \approx x_2 \gg 0$ and $x_1 \approx x_2 \ll 0$, as the finite basis ii' has no support for these regions of physical space. More subtly, incompleteness of ii' induces small differences (alias) between the potentials throughout physical space, with the differences becoming more marked near the periphery. Note that exact integration (e.g., an infinite quadrature grid) of either potential with the basis functions would return the same integral tensor, for this choice of basis. However, because we know that $\tilde{V}$ lies wholly in ii', we can invoke the *finite* Gaussian quadrature to quantize and integrate this potential exactly. This is denoted by the grid of dots on the middle panel. However, we cannot apply this same finite grid to $\hat{V}$: the alias, depicted in the bottom panel of each figure, is nonzero at the quadrature nodes.

In the limit that the primary basis is complete, $\tilde{V}$ approaches $\hat{V}$, as the finite primary basis spans the full Hilbert space, and there are no alias components to be projected. This is visible by comparing the $M = 5$ and $M = 9$ difference plots (note that the $M = 9$ quadrature spans a larger spatial region than $M = 5$). In the central region, from roughly $x = -2$ to $x = 2$, the $M = 9$ $\tilde{V}$ is much closer to $\hat{V}$ than in the $M = 5$ case; this corresponds to the additional completeness of the primary basis in this region for $M = 9$. The boundary regions of the quadratures are also interesting; the larger differences at the edges are similar, regardless of M. At the edges, the $M = 5$ and $M = 9$ primary basis sets have a qualitatively identical approximate span, albeit at a markedly different maximum spatial value.

These results indicate that X-THC can be interpreted as a two phase process: (1) production of a basis-specific potential $\tilde{V}$ by projection of $\hat{V}$ through A to clip components outside of ii' (renormalization) and then (2) alias-free invocation of the finite quadrature grid P to provide separability (quantization).

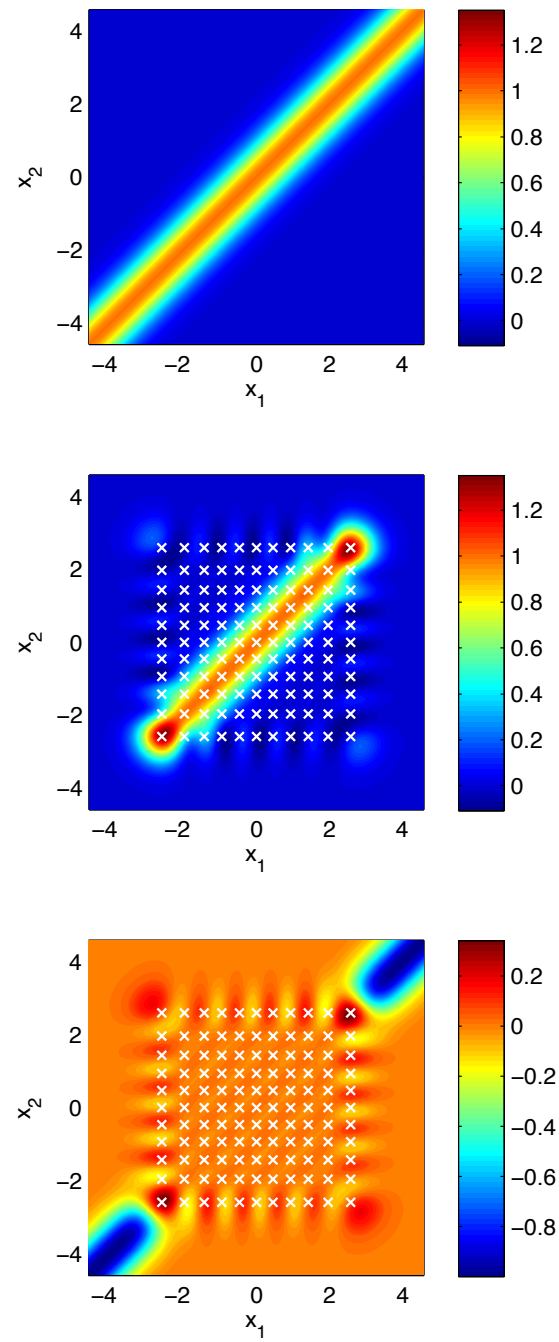

FIG. 1: Visualization of the original potential $\hat{V}$ (top), renormalized potential $\tilde{V}$ (middle), and difference (alias) $\tilde{V} - \hat{V}$ (bottom) for $M = 5$.

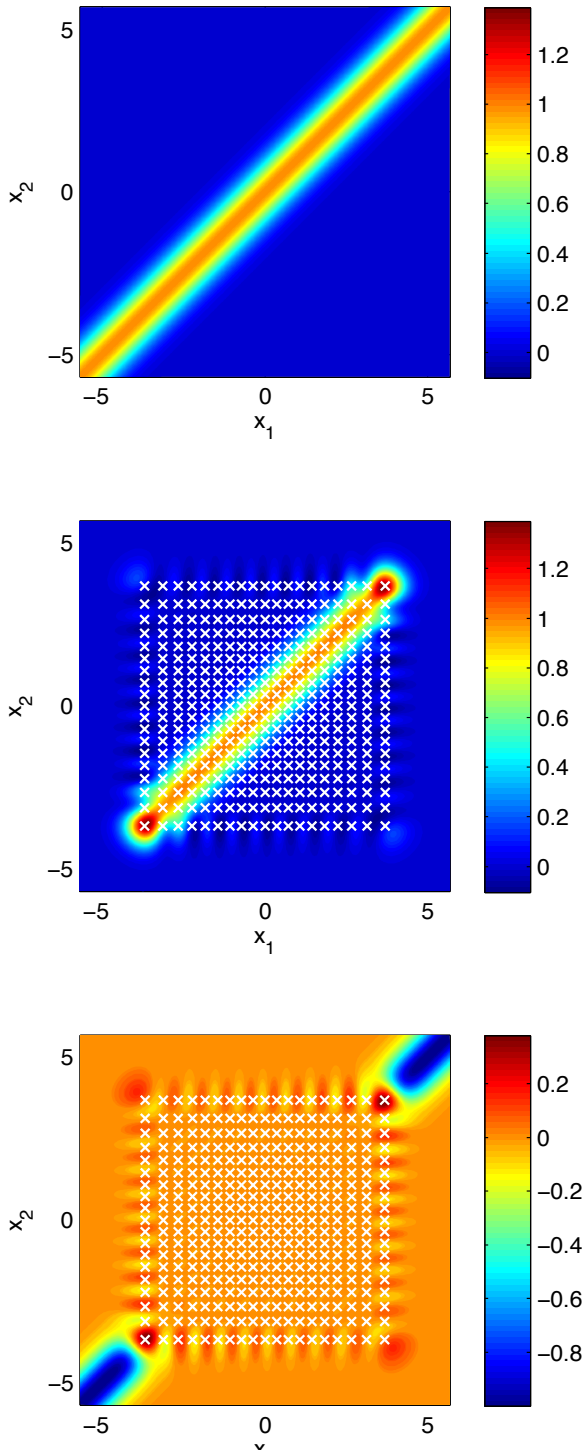

FIG. 2: Visualization of the original potential $\hat{V}$ (top), renormalized potential $\tilde{V}$ (middle), and difference (alias) $\tilde{V} - \hat{V}$ (bottom) for $M = 9$.

DEMONSTRATION CODE

Overview

To support the mathematical demonstraction of exact X-THC resolution of the potential integral tensor, and to provide a practical example of the scaling gains provided by X-THC, a mixed MATLAB/C++ code was developed for the case of D-dimensional Hermite functions with N-body w-contracted Gaussian forces. In this code, the required potential integrals (in the primary or auxiliary basis) and possible X-THC factors are generated in MATLAB and written to disk, for the given basis size M_μ, dimensionality D, number of bodies N, and number of w-contraction points N_w. For each integral technology and w-separable vs. w-nonseparable case, the generalized pairing field is computed for a randomly generated pairing tensor in a standalone C++ code for the particular w-separability and integrals technology case. In each C++ code, the integrals and factors are read in, the w indices contracted over first if simulating a nonseparable force, and then the generalized pairing tensor is computed according to the algorithms discussed below.

MATLAB was chosen for the integral and factor generation routines due to ease of implementation, and particular strength in treatment of arbitrary rank tensors. As this portion of the total procedure would typically be performed as a single-use overhead step (e.g., before HFB iterations or before application of the integrals in correlated methods), we do not include this step in the timings for the generalized pairing tensor, and thus there is no penalty for using the interpreted and rather memory-naïve MATLAB language for this stage. For the heavy linear algebra work of pairing tensor formation, C++ was selected for its "close to the metal" properties, particularly including explicit control of memory allocation and ability to swap pointers without performing explicit deep copy operations. Wherever possible (and for absolutely all contraction operations), BLAS calls are used, with the algorithms designed so as to allow for permutation of memory to be hidden in BLAS3 or BLAS2 operations as much as possible. The algorithms were formulated so as to rely preferentially on the BLAS3 DGEMM operation, followed by the BLAS2 DGEMV, followed by various BLAS1 operations. Every effort was expended to produce well-optimized algorithms, with the same amount of optimization present in both the X-THC and conventional methods. With these considerations, we believe that the timings reported are entirely representative of a practical application of the various integrals technologies, and show a wholly fair comparison of X-THC and conventional methods.

The C++ codes are compiled with the Intel `icpc 12.0.1` compiler, using `-O3` optimization. The BLAS calls are handled with Intel's very efficient `MKL 7.0.1`

library, with threading disabled. Performance measurements are performed using the `PAPI 5.0.1` library, which features a wall timer accurate to ~ 1 microsecond. Accumulated averaging was performed to ensure that all pairing kernels ran for at least 1 second of wall time, which ameliorates startup and noise costs for small problem sizes. All timings were produced on a single-socket node featuring a quad-core 3.4 GHz Intel i7 Processor (Sandy Bridge) with 32 GB of DDR3 and 8 MB of L3 cache. All timings are for wall times using a single thread.

Note that we do not show results for density-fitted algorithms in this study, though we have coded generalized pairing routines for this technology. All accuracy results are of essentially the same quality as X-THC (e.g., numerically exact to within a small pre-factor of the machine epsilon). For w-separable forces with $N > 1$, the optimal pathway is to contract the DF intermediates to the conventional integrals, and then form the pairing tensor in the conventional manner. As a result, the DF timings results for w-separable potentials are essentially indistinguishable from the conventional case. For non-w-separable potentials, the DF algorithms require the same $\mathcal{O}(M_\mu^{2ND})$ scaling as the conventional case, but the prefactors are many orders of magnitude larger, due to the larger auxiliary basis sizes involved. Moreover, in the non-w-separable case, formation of the conventional integrals from the DF integrals exhibits a higher scaling of $\mathcal{O}(M_\mu^{2ND+1})$, and is therefore not a viable alternative. In any case, the conventional algorithm always outperforms the DF algorithm for the generalized pairing tensor, so we have elected to not include the DF results here. This failure of DF methods for "exchange-like" contractions is well known, and was the primary motivation for the development of the THC representation.

Basis/Potential Choice

Hermite Function Primary Basis

The primary basis functions chosen for this demonstration are direct products of generalized Hermite functions,

$$\psi_i(\boldsymbol{r}) = \prod_{\mu}^{D} \psi_{i_\mu}(r_\mu). \tag{54}$$

Below, we will drop the μ labels and work in 1D ($\boldsymbol{r} \equiv x$) unless otherwise noted. The 1D generalized Hermite function is,

$$\psi_i(x) = (b_\mu)^{1/2} \underbrace{(\sqrt{\pi}2^i i!)^{-1/2} H_i(z) \exp(-z^2/2)}_{\phi_i(z)}, \tag{55}$$

where H_i is the i-th Hermite polynomial and the nondimensional coordinate is,

$$z = b_\mu x. \tag{56}$$

i runs from 0 to M_μ. Note that we reserve $\phi_i(z)$ for the true non-dimensional Hermite function, where $b_\mu = 1$.

The auxiliary functions for this problem are also generalized Hermite functions,

$$\chi_A(x) = (\sqrt{2}b_\mu)^{1/2} (\sqrt{\pi}2^A A!)^{-1/2} H_A(z') \exp(-z'^2/2), \tag{57}$$

where now,

$$z' = \sqrt{2} b_\mu x, \tag{58}$$

and A runs from 0 to $2M_\mu$. The THC quadrature for this problem is thus the $2M_\mu + 1$-node Gauss-Hermite quadrature with a spatial length scale of $\sqrt{2}b_\mu$.

Gaussian Potential

For flexibility, we use for potential a linear combination of Gaussians. Given the set $\{< \alpha_w, \beta_w >\}$, the form of the potential is,

$$\hat{V}(x_1, \ldots, x_N) = \sum_{w}^{N_w} \sum_{\eta=1}^{N-1} \sum_{\xi=\eta+1}^{N} \alpha_w^{1/D} \exp(-\beta_w x_{\eta\xi}^2). \tag{59}$$

For a 2-body potential in 1 dimension, with $N_w = 1$ this reduces to the usual Gaussian force,

$$\hat{V}(x_1, x_2) = \alpha_w^{1/D} \exp(-\beta_w x_{12}^2). \tag{60}$$

For higher N, this is simply a sum of two-body Gaussian forces over all possible pairs of two-body coordinates, $x_{\eta\xi} = x_\eta - x_\xi$.

We use a set of $\{< \alpha_w, \beta_w >\}$ with $N_w = 8$ which approximates the Coulomb operator $1/r_{12}$ for all computations shown in this work. This allows us to show separate accuracy and timings results for the w-separable and non-w-separable cases, depending on whether we choose to sum over the w index first or last.

Potential Integrals (MATLAB)

All of the D-dimensional N-body potential integrals above can be constructed from primitive 1-dimensional 2-body potential integrals. For conventional integral technology, four-index integrals in the primary basis are required. For X-THC technology, two-center integrals in the auxiliary basis are required. Integrals of this type are discussed extensively in [1], where Moshinski transformations and length-scale transformations are used to provide analytical conventional integrals (the generalization to auxiliary integrals is straightforward).

Conventional Integrals

Following [1], the 1-dimensional 2-dimensional primary integrals are computed as,

$$\langle ij|\hat{V}^w|i'j'\rangle = \alpha_w^{1/D} \iint \mathrm{d}x_1\mathrm{d}x_2\ \psi_i(x_1)\psi_j(x_2)\times \exp(-\beta_w x_{12}^2)\psi_{i'}(x_1)\psi_{j'}(x_2) = \alpha_w^{1/D} D_{mn}^{Aa} D_{m'n'}^{Ab} D_{aa'}(\eta_w) D_{ba'}(\eta_w), \quad (61)$$

where all summations run from 0 to $2M_\mu$. D_{mn}^{Aa} are the Moshinski transformation coefficients, and $D_{aa'}(\eta_w)$ are the length-scale transformation coefficients (see below). The length-scale change parameter is

$$\eta_w = \frac{1}{\sqrt{1+2\beta_w/b_\mu^2}}. \quad (62)$$

Auxiliary Integrals

By extension, the 1-dimensional primary integrals are computed as,

$$G_w^{AB} = \alpha_w^{1/D} \iint \mathrm{d}x_1\mathrm{d}x_2\ \chi_A(x_1)\exp(-\beta_w x_{12}^2)\chi_B(x_2) = \frac{1}{\sqrt{2}b_\mu}\frac{\alpha_w^{1/D}}{\sqrt[4]{1+2\beta_w/b_\mu^2}} D_{AB}^{Nn} D_{nn'}(\eta_w) I_N I_{n'}, \quad (63)$$

where again all summations run from 0 to $2M_\mu$. The I_N quantities are primitive integrals over single Hermite functions of unit length,

$$I_N = \int \mathrm{d}z\ \psi_N(z) = \begin{cases} \sqrt[4]{\pi}\sqrt{2}\frac{(N-1)!!}{\sqrt{N!}} & N \text{ even} \\ 0 & N \text{ odd} \end{cases} \quad (64)$$

Moshinski Transformation Coefficients D_{mn}^{Aa}

The Moshinski transformation coefficients relate products of Hermite functions in Eulerian coordinates x_1 and x_2 to corresponding Hermite functions in the Lagrangian coordinates X and x, where,

$$X = \frac{1}{\sqrt{2}}[x_1+x_2], \quad x = \frac{1}{\sqrt{2}}[x_1-x_2], \quad (65)$$

The transformation is

$$\psi_{n_1}(x_1)\psi_{n_2}(x_2) = D_{n_1,n_2}^{N,n}\psi_N(X)\psi_n(x). \quad (66)$$

As the Gaussian potential is central (i.e., depends only on x, not X), invoking the Moshinksi transformation reduces the two-coordinate potential integrals to a separable product of one-coordinate integrals in X and x. If n_1 and n_2 each range from 0 to M_μ, N and n each range from 0 to $2M_\mu$. The well-known selection rule is $n_1+n_2 = N+n$. A simple, explicit formula for these coefficients is [1]

$$D_{n_1,n_2}^{N,n} = \delta_{n_1+n_2,N+n}\left(\frac{n_1!n_2!}{N!n!}\right)\left(\frac{1}{\sqrt{2}}\right)^{N+n} \times \sum_{i,j=0}^{i<N,j<n,i+j=n_1} \binom{N}{i}\binom{n}{j}(-1)^j. \quad (67)$$

However, this formula is unstable for large values of $N+n$ due to the summation over alternating quantities which are both large. We have therefore derived stable recurrence relations for the Moshinksi coefficients. The recurrence relation in the first coordinate is,

$$D_{n_1+1,n_2}^{N,n} = \frac{1}{\sqrt{2}}\left[\sqrt{\frac{N}{n_1+1}}D_{n_1,n_2}^{N-1,n} + \sqrt{\frac{n}{n_1+1}}D_{n_1,n_2}^{N,n-1}.\right] \quad (68)$$

And the corresponding recurrence relation in the other coordinate is,

$$D_{n_1,n_2+1}^{N,n} = \frac{1}{\sqrt{2}}\left[\sqrt{\frac{N}{n_2+1}}D_{n_1,n_2}^{N-1,n} - \sqrt{\frac{n}{n_2+1}}D_{n_1,n_2}^{N,n-1}.\right] \quad (69)$$

The seed is of course,

$$D_{0,0}^{0,0} = 1. \quad (70)$$

For implementation purposes, any Moshinski transformation coefficient with a "negative" index can be taken to be zero.

Length-Scale Transformation Coefficients $D_{nn'}(\eta_w)$

The length-scale transformation coefficients provide the correspondence between Hermite polynomials of different length-scales,

$$H_n(z) = (\eta)^{-1/2}D_{nn'}(\eta)H_{n'}(z'=z/\eta) \quad (71)$$

Note that the transformation automatically renormalizes the polynomials. n and n' both run from 0 to M_μ.

A closed-form expression of these coefficients is [1]

$$D_{nn'}(\eta) = F_{n,n'}(-1)^{\frac{n-n'}{2}}\left(\frac{n!}{n'!}\right)^{1/2}\frac{\eta^{n'+1/2}(1-\eta^2)^{\frac{n-n'}{2}}}{2^{\frac{n-n'}{2}}\left(\frac{n-n'}{2}\right)!}, \quad (72)$$

on the condition that $n > n'$, and that the parity agrees,

$$F_{n,n'} = \frac{1}{2}[1+(-1)^{n+n'}] = \begin{cases} 1, & n+n' \text{ even} \\ 0, & n+n' \text{ odd} \end{cases} \quad (73)$$

Because $\eta \leq 1$, this formula appears to be universally stable. However, the numerators and denominators above both involve divergent values, so a log-space

formalism is required for explicit evaluation to prevent overflow. This may lose a few digits of precision, so we have elected to use a recurrence relation for these coefficients, which is easily derived. The recurrence relation is,

$$D_{n+1,n'} = \eta\sqrt{\frac{n'+1}{n+1}}D_{n,n'+1} + \eta\sqrt{\frac{n'}{n+1}}D_{n,n'-1} - \sqrt{\frac{n}{n+1}}D_{n-1,n'}. \quad (74)$$

The seed is,

$$D_{0,0}(\eta) = \sqrt{\eta}. \quad (75)$$

For implementation purposes, any length-scale transformation coefficient with a "negative" index can be taken to be zero.

D-Dimensional, N-Body Integrals

The generalization to N-body integrals as described above is carried out at the 1-dimensional stage, and is currently hard-coded for only the 3-body case (this could easily be extended, but the point is already demonstrated by 3-body). For primary-basis integrals, the required integral is simply,

$$\langle ijk|\hat{V}^w|i'j'k'\rangle = \langle ij|\hat{V}^w|j'j'\rangle\delta_{kk'} + \langle ik|\hat{V}^w|i'k'\rangle\delta_{jj'} + \langle jk|\hat{V}^w|j'k'\rangle\delta_{ii'}. \quad (76)$$

For auxiliary basis integrals, there is no delta function in the third coordinate, but rather a normalized primitive Hermite integral, i.e., the quantity I_N defined in the auxiliary potential integrals. Thus, the auxiliary potential integral is,

$$G_w^{ABC} = G_w^{AB}I_C + G_w^{AC}I_B + G_w^{BC}I_A. \quad (77)$$

In practice, the generalization to 3-body integrals is performed in MATLAB on the 1-dimensional integrals, before integrals are written to disk.

The generalization to non-w-separable D-dimensional potentials is carried out by summing over w, e.g., in the 2-dimensional, 2-body case,

$$\langle ij|\hat{V}|i'j'\rangle = \sum_w^{N_w}\langle i_x j_x|\hat{V}^w|i'_x j'_x\rangle\langle i_y j_y|\hat{V}^w|i'_y j'_y\rangle, \quad (78)$$

or,

$$G^{AB} = \sum_w^{N_w} G_w^{A_x B_x} G_w^{A_y B_y}. \quad (79)$$

In practice, the production of the nonseparable Z factors or conventional integrals is carried out in blocks in C++, to save memory. The formation of these integrals is not counted in the pairing timings, as infinite memory is assumed. The overhead for this is many orders of magnitude larger for the conventional case than the X-THC case, due to the larger size of the rank-$2ND$ conventional integral tensor.

X-THC Factors (MATLAB)

To complete the X-THC factorization, the Gauss-Hermite quadrature nodes/weights and collocation matrices X and Y are required.

Quadratures

The THC grid for this problem is the $2M_\mu + 1$ node Gauss-Hermite quadrature with the spatial range parameter $\sqrt{2}b_\mu$.

First, the non-dimensional quadrature is generated. The orthonormal-basis position operator $X_{PQ} = \langle P|\hat{x}|Q\rangle$ in the auxiliary Hermite functions is,

$$X_{PQ} = \langle P|\hat{x}|Q\rangle = \sqrt{\frac{P+1}{2}}\left(\delta_{P+1,Q} + \delta_{P,Q+1}\right). \quad (80)$$

This operator is symmetric tridiagonal, with zero diagonal. The eigen-decomposition is formed,

$$X_{PQ} = Q_{PP'}x_{P'}Q_{QP'}. \quad (81)$$

The eigenvalues are the nondimensional quadrature nodes. The weights are determined by first generating the moment vector,

$$v_P = \sqrt[4]{\pi}\delta_{P0}. \quad (82)$$

and applying the diagonalizing transformation,

$$v_{P'} = Q_{PP'}v_P \quad (83)$$

The full weights are then given by,

$$w_{P'} = v_{P'}^2\exp(x_{P'}^2). \quad (84)$$

The nodes and weights are then transformed to the problem domain, by,

$$x_P = \frac{x_{P'}}{\sqrt{2}b_\mu}, \quad w_P = \frac{w_{P'}}{\sqrt{2}b_\mu}. \quad (85)$$

Collocation

The X and Y X-THC factors require collocations at the quadrature nodes. This is accomplished by efficient, stable recurrence relations for the Hermite functions.

The non-dimensional coordinate is first computed,

$$z = b_\mu x. \tag{86}$$

The first two non-dimensional Hermite functions are computed explicitly,

$$\psi_0(z) = \pi^{-1/4} \exp(-z^2/2), \tag{87}$$

and

$$\psi_1(z) = \sqrt{2} z \pi^{-1/4} \exp(-z^2/2) = \sqrt{2} z \psi_0(z). \tag{88}$$

The recurrence relation is then applied iteratively,

$$\begin{aligned}\psi_{i+1}(z) &= \sqrt{\frac{2}{i+1}} \left[z\psi_i(z) - \sqrt{\frac{i}{2}} \psi_{i-1}(z) \right] \\ &= \sqrt{\frac{2}{i+1}} z\psi_i(z) - \sqrt{\frac{i}{i+1}} \psi_{i-1}(z).\end{aligned} \tag{89}$$

And finally the length-scale normalization is added,

$$\psi_i(z) = (b_\mu)^{1/2} \psi_i(z). \tag{90}$$

The X factor is,

$$X_i^P = \psi_i(x_P) = (b_\mu)^{1/2} \psi_i(z = b_\mu x_P), \tag{91}$$

and the Y factor is,

$$Y_A^P = w_P \chi_A(x_P) = w_P (\sqrt{2} b_\mu)^{1/2} \psi_A(z' = \sqrt{2} b_\mu x_P). \tag{92}$$

If desired, the DF 3-index overlap integrals can be generated exactly (within numerical precision) from the X-THC factors,

$$[ii'|A] = X_i^P X_{i'}^P Y_A^P. \tag{93}$$

The X-THC Z operators are immediately formed from the Y factors and corresponding G auxiliary potential integrals,

$$Z_w^{PQ} = Y_A^P Y_B^Q G_w^{AB}. \tag{94}$$

Generalized Pairing Algorithms (C++/BLAS)

Algorithms for the generalized pairing tensor with conventional or X-THC integrals, and general or w-separable potentials are described below, and depicted in Algorithms 1-4. In these algorithms, we use ellipses to denote arbitrary rank in D or N, and we follow the convention that dimensions are striped as the slow superindex, and particles are striped as the fast superindex. The tensors used in these algorithms are stored in practice as a collapsed single-dimensional array (a `double*`), regularly striped so that the left-most index is the fastest dimension and the right-most index is the slowest dimension (Fortran order, which allows for convenient application of BLAS operations). For key contraction operations, we group sets of neighboring indices to form the three superindices i (result row or fast index), j (result column or slow index) and k (contraction index) for use in GEMM,

$$C_{ij} = A_{ik} B_{kj}. \tag{95}$$

Note the use of color to emphasize the choice of superindices. By changing the order of A and B and altering the GEMM transposition arguments, we can take i, j, and k in any order in the factor tensors A or B (e.g., the contraction index could actually be the row dimension in A). However, striding is not permitted with BLAS3 operations, e.g., a tensor contraction of the form,

$$C_{ijk} = A_{il} B_{jlk}. \tag{96}$$

would require explicit transposition to B_{ljk} or B_{jkl} to be able to group the compound jk index. Our algorithms are designed to eliminate such explicit transposition as much as possible. However, no transposition-free algorithm exists for X-THC w-separable potentials in arbitrary N, see the discussion below.

Note that our discussions and codes all assume isotropic basis sets for simplicity (e.g., $M_x = M_y$), but X-THC is certainly not restricted to this. In general, X-THC can handle differing values of M_x and M_y, and even different classes of basis functions in each dimension (e.g. cylindrical coordinates).

Conventional Integrals, General Potential

This algorithm is quite simple, but extraordinarily expensive (this is why non-w-separable forces are rarely used in high D or N cases). The physicists' integral tensor is used in a simple matrix-vector product with the generalized pairing tensor, which is carried out by GEMV in $M_\mu^{2ND} \equiv \mathcal{O}(M_\mu^{2ND})$ operations.

In practice, such an algorithm will almost certainly be orders of magnitude more expensive than reported here. In our demonstration, we simulate infinite memory by forming blocks of the integrals prior to GEMM (i.e., an integral direct procedure). The integral generation costs $\mathcal{O}(M_\mu^{2ND})$ (in this case by contracting out the w index) with a much larger prefactor than the GEMV itself. In a fully generic potential, the integrals would have to be generated explicitly, at possibly even higher cost. This integral formation is not counted in the wall times reported here, but would increase the practical wall time considerably in the usual case that the rank-$2ND$ integral tensor does not fit in core memory.

X-THC Integrals, General Potential

This algorithm works by cyclically transforming from configuration space i'_μ to the quantized coordinate space P_μ by GEMM, scaling the coordinate-space pairing tensor by Z, and then performing another cyclic transformation to bring P_μ back to i_μ via GEMM. The cyclic transformations are written so that the fast index in the current buffer is contracted off, and the replacement index is placed at the slow index of the result buffer. The pointers for the result and current buffer are then swapped, and the next contraction index automatically appears in the fast index of the current buffer, without any need for explicit transposition.

Formally, $\mathcal{O}(M_\mu^{ND+1})$ operations are required for the cyclic permutations. The prefactor arises from the $\sim 2M_\mu$ size of the quadrature index, and from the fact that two cyclic permutations are required.

This algorithm currently requires essentially three buffers of size $2^{ND}M_\mu^{ND}$ (T, U, and Z). An efficient blocked or disk-based algorithm could be developed if this becomes a bottleneck, with considerable performance gains expected over conventional disk-based or integral-direct algorithms with general potentials.

Conventional Integrals, w-Separable Potential

This algorithm works by cyclically applying the integrals for dimension μ via GEMM, for each w point. Explicit transposition is avoided by placing the replacement index $(i\ldots n)_\mu$ as the slow index of the result, allowing $(i'\ldots n')_{\mu+1}$ to be exposed as the new fast index. The formal scaling of this algorithm is $\mathcal{O}(N_w M_\mu^{ND+N})$ operations, with remarkably small external overhead. The memory requirement is essentially 2 M_μ^{ND} buffers (T and U). The small overhead, combined with the efficiency of the long $(i'\ldots n')_\mu$ contraction index and small memory footprint explains much of the current success of w-separable potentials.

X-THC Integrals, w-Separable Potential

This algorithm works by, for each w point and dimension μ, cyclically forward transforming from i'_μ to P_μ, applying the $Z_w^{(P\ldots W)_\mu}$ operation, and then cyclically backtransforming from P_μ to i_μ. The formal prefactor is $\mathcal{O}(N_w M_\mu^{ND+1})$ FMA operations. The required memory is essentially 2 $2^N M_\mu^{ND}$ buffers T and U, a factor of 2^N more than the conventional w-separable algorithm.

Unfortunately, an explicit transposition is required for this algorithm for general N. The genesis of this requirement is that the cyclic permutation is not carried through all ND coordinates, but only N coordinates at a time.

In the 2-body case, a specialized algorithm can be applied to avoid the explicit transposition. The transformation in each dimension reads,

$$\begin{aligned} U_{j'LP} &\leftarrow T_{i'j'L} X_{i'}^P \\ U_{QLP} &\leftarrow T_{j'LP} X_{j'}^Q \\ U_{QLP} &\leftarrow T_{QLP} Z_w^{PQ} \\ U_{QLi} &\leftarrow T_{QLP} X_i^P \\ U_{Lij} &\leftarrow T_{QLi} X_j^Q \end{aligned} \tag{97}$$

Here L is the compound index corresponding to the other dimensions, and pointer swap is assumed between each step.

One additional modification can attenuate the prefactor somewhat, at the cost of doubling the buffer space. The forward transformation of the first dimension from $\kappa_{(i'j')_1 L}$ to $U_{L(PQ)_1}$ is the same for all w points, as κ is w-independent. Moreover, the back transformation of the last dimension from $U_{L(PQ)_D}$ to $\Delta_{L(ij)_D}$ does not depend on the individual w points, but only on their sum. Thus the buffers $U_{L(PQ)_1}$ and $U_{L(PQ)_D}$ can be used in a prelude/epilogue construct. In the limit that $N_w \to \infty$, the savings are 100%, 50%, and 33%, for D = 1, 2,, and 3, respectively. In the limit that $N_w = 1$ or $D \to \infty$, the savings approach 0%, so the deployment of this modification would depend greatly on the context and memory capacity.

Computational Results

Timings and accuracy results for generalized pairing tensor formation are shown in Figures 3 and 4, respectively, for various integral technologies, dimensions, number of bodies, and problem sizes. To help crystallize the information in the timing data, asymptotic scaling and predicted crossover metrics are presented in Tables I and II, respectively.

The accuracy results of Figure 4 depict the relative maximum residual D in the pairing tensor, defined as,

$$D_{\text{Method}} = \frac{\left\|\Delta_{i\ldots n}^{\text{Method}} - \Delta_{i\ldots n}^{\text{Reference}}\right\|_\infty}{\left\|\Delta_{i\ldots n}^{\text{Reference}}\right\|_\infty}. \tag{98}$$

Here the w-separable conventional integral technology was selected as a reference for the accuracy.[3] It is apparent that all methods (conventional and THC) are exact to within a reasonable growth factor against the machine epsilon. In fact, the worst relative maximum residual seen here is less than 10^{-12} (compared to the double-precision machine epsilon of 2.2×10^{-16}), and does not seem to be grow markedly with respect to problem

size. This provides strong numerical evidence that X-THC is a lossless compression of the potential integral tensor. Further agreement could doubtless be obtained if the Moshinksi coefficients were evaluated with higher precision.[4]

The timing results of Figure 3 are quite clean as timings go, with very smooth increase with respect to problem size, indicating that the accumulation procedure (smaller M_μ) or sheer size of the problem (larger M_μ) are sufficient to eliminate the bulk of the noise that often plagues timing studies. On a log-log scale, All timing curves exhibit a slight upward concavity which quickly tends toward linearity for larger M_μ as the rate-limiting DGEMM-based steps become dominant relative to the lower-scaling operations. Two noticable "jumps" exist: the last few points of the 2-body, 1-dimensional conventional separable case and the last few points of the 2-body, 3-dimensional THC separable case. These jumps are likely due to working set size limits exceeding some discrete performance threshold in hardware, for instance, cache or memory bank limits, respectively. From these plots, power-law regression of the form,

$$t_{\text{Pairing}} = \alpha M_\mu^\beta, \tag{99}$$

is performed, using points selected above an M_μ which appears to be visually free of noise in each case. The β from these regressions are depicted in Table I. The predicted THC crossover points from these regressions (the critical M_μ at which X-THC becomes practically superior to conventional integrals) are shown in Table II. For all cases in which an explicit crossover occurs, the predicted crossover point from the power-law model is within 1 M_μ of the observed value. Note that the observed asymptotic scaling is often less than the theoretical value. There are two geneses for this: residual contributions from lower scaling operations (which drags the scaling down at the cost of prefactor) and better GEMM efficiency for larger matrix sizes (which makes GEMM appear to scale better than cubic as the matrix size increases). However, the relative scaling relationships between all integral technologies is retained.

From this point forward, it is useful to consider general and w-separable potentials separately.

For all cases of N and D in general local potentials, X-THC is markedly more efficient than conventional. In all such cases, X-THC crosses over conventional (often at very small M_μ for $D > 1$), and is several orders of magnitude faster for the largest cases shown with $D > 1$. This is strong evidence for the immediate application of X-THC to problems involving general potentials, with possibly magnficent tractability gains realizable.

For w-separable cases, X-THC always exhibits lower asymptotic scaling than conventional, and provides crossovers and sometimes significant speedups for $D = 1$ and $D = 2$. However, the narrower asymptotic separation between conventional and X-THC ($\beta_{\text{Conv-Sep}} - \beta_{\text{THC-Sep}} = N-1$ vs. $\beta_{\text{Conv-Gen}} - \beta_{\text{THC-Gen}} = ND-1$) exposes the prefactor of X-THC, particularly for larger D. This prefactor has two geneses: the formal FMA prefactor due to successive substitution of primary basis indices for larger quadrature indices and the lower efficiency of X-THC DGEMM operations compared to conventional DGEMM operations. For a 3-body, 3-dimensional w-separable potential, the crossover seems likely to occur just outside of the memory-limited problem size explicitly shown in Figure 3. In fact, the predicted crossover point for this case is $M_\mu = 10.3$. The 2-body, 3-dimensional w-separable potential is somewhat more sinister: a "jump" in the timings curve caused by some hardware boundary causes the expected crossover point to increase to a practically unreachable value of $M_\mu = 1789$.

The inability of X-THC to provide a practical crossover for the 2-body, 3-dimensional w-separable potential is an indication that this technique is not a panacea. However, we point out that the utilization of a w-separable potential throughout the literature seems to stem from the lack of an X-THC representation for a general potential: the approximation of the potential as w-separable was required to provide a tractable numerical numerical recipe. In this light, X-THC treatment of general potentials provides a practical alternative pathway that avoids approximation of the potential as w-separable. For all of the cases studied here, the general X-THC curve is within roughly an order of magnitude of the conventional separable curve (in some cases even faster!). Moreover, X-THC does provide some gains in w-separable potentials, particularly for larger N. On more formal grounds, the demonstrated lossless asymptotic reduction of a generalized pairing term in any local potential from $\mathcal{O}(M_\mu^{2ND})$ to $\mathcal{O}(M_\mu^{ND+1})$ is a useful result in and of itself, considering that simply storing the pairing tensor requires $\mathcal{O}(M_\mu^{ND})$.

* Electronic address: schunck1@llnl.gov
† Electronic address: sherrill@gatech.edu
‡ Electronic address: toddjmartinez@gmail.com

[1] L. M. Robledo, Phys. Rev. C **81**, 044312 (2010).
[2] A tensor generalization of the chemists' notation still exists for $N > 2$, indeed, it is required for THC to function. However, it is very difficult to write down (as you would need N generalizations of the bra- or ket-like objects $[ii'|$ and $|jj']$). A possible generalization is $[ii'|\ldots[nn'|\hat{V}$, which makes sense physically, but just plain looks awkward. This is in contrast to the physicists' notation, which is trivial to write as a compound rank-2 object $\langle i\ldots n|\hat{V}|i'\ldots n'\rangle$, but perhaps harder to visualize.
[3] Note that there is some ambiguity here, as some roundoff error is intrinsic to the w-separable conventional reference itself. A particularly marked source of roundoff error is the Moshinski relations for the potentials. Since both the w-separable and general conventional pairing routines

share the same underlying Moshinski relations, superior agreement between these two conventional methods and the THC methods does not necessarily imply that one is more accurate than the other, against a hypothetical exact precision result. In any case, the agreement between all methods is sufficient that the point is moot.

[4] This is the most numerically susceptible portion of the procedure, even when using recurrence relations.

Algorithm 1 Generalized pairing algorithm: conventional integrals, general potential.

procedure PAIRING_CONV_GEN($\langle (i\dots n)_1 \dots (i\dots n)_D | \hat{V} | (i'\dots n')_1 \dots (i'\dots n')_D \rangle$, $\kappa_{(i'\dots n')_1\dots(i'\dots n')_D}$)
  Allocate $\Delta_{(i\dots n)_1\dots(i\dots n)_D}$ ▷ Target (Typically Preallocated)
  $\Delta_{(i\dots n)_1\dots(i\dots n)_D} = \langle (i\dots n)_1 \dots (i\dots n)_D | \hat{V} | (i'\dots n')_1 \dots (i'\dots n')_D \rangle \kappa_{(i'\dots n')_1\dots(i'\dots n')_D}$ ▷ GEMV, $\mathcal{O}(M_\mu^{2ND})$
  return $\Delta_{(i\dots n)_1\dots(i\dots n)_D}$
end procedure

Algorithm 2 Generalized pairing algorithm: X-THC integrals, general potential.

procedure PAIRING_THC_GEN($X_{i_\mu}^{P_\mu}$, $Z^{(P\dots W)_1\dots(P\dots W)_D}$, $\kappa_{(i'\dots n')_1\dots(i'\dots n')_D}$)
  Allocate $\Delta_{(i'\dots n')_1\dots(i'\dots n')_D}$ ▷ Target (Typically Preallocated)
  Allocate $T_{(P\dots W)_1\dots(P\dots W)_D}$ ▷ Scratch Array (Typically Preallocated)
  Allocate $U_{(P\dots W)_1\dots(P\dots W)_D}$ ▷ Scratch Array (Typically Preallocated)
  $T_{(i'\dots n')_1\dots(i'\dots n')_D} = \kappa_{(i'\dots n')_1\dots(i'\dots n')_D}$ ▷ Deep Copy
  for all $\mu \in [1, D]$ **do** ▷ Start Cyclic Permutation in μ and η
    for all $\eta \in [1, N]$ **do**
      $U_{(\dots n')_\mu\dots(P\dots W)_{\mu-1}(P)_\mu} = X_{i'_\mu}{}^{P_\mu} T_{(i')_\mu(\dots n')_\mu\dots(P\dots W)_{\mu-1}}$ ▷ GEMM, $\mathcal{O}(M_\mu^{ND+1})$
      swap(T, U) ▷ Pointer Swap
    end for
  end for ▷ End Cyclic Permutation in μ and η
  $T_{(P\dots W)_1\dots(P\dots W)_D} * = Z^{(P\dots W)_1\dots(P\dots W)_D}$ ▷ Hadamard Product, $\mathcal{O}(M_\mu^{ND})$
  for all $\mu \in [1, D]$ **do** ▷ Start Cyclic Permutation in μ and η
    for all $\eta \in [1, N]$ **do**
      $U_{(\dots W)_\mu\dots(i\dots n)_{\mu-1}(i)_\mu} = X_{i_\mu}{}^{P_\mu} T_{(P)_\mu(\dots W')_\mu\dots(i\dots n)_{\mu-1}}$ ▷ GEMM, $\mathcal{O}(M_\mu^{ND+1})$
      swap(T, U) ▷ Pointer Swap
    end for
  end for ▷ End Cyclic Permutation in μ and η
  $\Delta_{(i\dots n)_1\dots(i\dots n)_D} = T_{(i\dots n)_1\dots(i\dots n)_D}$ ▷ Deep Copy
  return $\Delta_{(i\dots n)_1\dots(i\dots n)_D}$
end procedure

Algorithm 3 Generalized pairing algorithm: conventional integrals, w-separable potential.

procedure PAIRING_CONV_SEP($\langle i\dots n | \hat{V} | i'\dots n' \rangle_\mu^w$, $\kappa_{(i'\dots n')_1\dots(i'\dots n')_D}$)
  Allocate $\Delta_{(i\dots n)_1\dots(i\dots n)_D} = 0$ ▷ Target (Typically Preallocated)
  Allocate $T_{(i'\dots n')_1\dots(i'\dots n')_D}$ ▷ Scratch Array (Typically Preallocated)
  Allocate $U_{(i'\dots n')_1\dots(i'\dots n')_D}$ ▷ Scratch Array (Typically Preallocated)
  for all $w \in [1, N_w]$ **do**
    $T^w_{(i'\dots n')_1\dots(i'\dots n')_D} = \kappa_{(i'\dots n')_1\dots(i'\dots n')_D}$ ▷ Deep Copy
    for all $\mu \in [1, D]$ **do** ▷ Start Cyclic Permutation in μ
      $U^w_{\dots(i\dots n)_{\mu-1}(i\dots n)_\mu} = \langle i\dots n | \hat{V} | i'\dots n' \rangle_\mu^w T^w_{(i'\dots n')_\mu\dots(i\dots n)_{\mu-1}}$ ▷ GEMM, $\mathcal{O}(N_w M_\mu^{ND+N})$
      swap(T^w, U^w) ▷ Pointer Swap
    end for ▷ End Cyclic Permutation in μ
    $\Delta_{(i\dots n)_1\dots(i\dots n)_D} + = T^w_{(i\dots n)_1\dots(i\dots n)_D}$ ▷ Contribution from w
  end for
  return $\Delta_{(i\dots n)_1\dots(i\dots n)_D}$
end procedure

Algorithm 4 Generalized pairing algorithm: X-THC integrals, w-separable potential.

procedure PAIRING_THC_SEP($X_{i_\mu}^{P_\mu}$, $Z_w^{(P\ldots W)_\mu}$, $\kappa_{(i'\ldots n')_1\ldots(i'\ldots n')_D}$)
  Allocate $\Delta_{(i'\ldots n')_1\ldots(i'\ldots n')_D}$ ▷ Target (Typically Preallocated)
  Allocate $T_{(P\ldots W)_1\ldots(i\ldots n)_D}$ ▷ Scratch Array (Typically Preallocated)
  Allocate $U_{(P\ldots W)_1\ldots(i\ldots n)_D}$ ▷ Scratch Array (Typically Preallocated)
  for all $w \in [1, N_w]$ **do**
    $T_{(i'\ldots n')_1\ldots(i'\ldots n')_D} = \kappa_{(i'\ldots n')_1\ldots(i'\ldots n')_D}$ ▷ Deep Copy
    for all $\mu \in [1, D]$ **do** ▷ Start Cyclic Permutation in μ
      for all $\eta \in [1, N]$ **do**
        $U_{(\ldots n')_\mu\ldots(i\ldots n)_{\mu-1}(P)_\mu} = X_{i'_\mu}{}^{P_\mu} T_{(i')_\mu(\ldots n')_\mu\ldots(i\ldots n)_{\mu-1}}$ ▷ GEMM, $\mathcal{O}(N_w M_\mu^{ND+1})$
        swap(T, U) ▷ Pointer Swap
      end for
      $T_{\ldots(i\ldots n)_{\mu-1}(P\ldots W)_\mu} *= Z_w^{(P\ldots W)_\mu}$ ▷ SCAL, $\mathcal{O}(N_w M_\mu^{ND})$
      for all $\eta \in [N, 1]$ **do**
        $U_{(n)_\mu\ldots(i\ldots n)_{\mu-1}(P\ldots)_\mu} = X_{n_\mu}{}^{W_\mu} T_{\ldots(i\ldots n)_{\mu-1}(P\ldots)_\mu(W)_\mu}$ ▷ GEMM, $\mathcal{O}(N_w M_\mu^{ND+1})$
        swap(T, U) ▷ Pointer Swap
      end for
      $U_{\ldots(i\ldots n)_{\mu-1}(i\ldots n)_\mu} = T_{(i\ldots n)_\mu\ldots(i\ldots n)_{\mu-1}}$ ▷ Explicit Transposition
      swap(T, U) ▷ Pointer Swap
    end for ▷ End Cyclic Permutation in μ
    $\Delta_{(i\ldots n)_1\ldots(i\ldots n)_D} += T_{(i\ldots n)_1\ldots(i\ldots n)_D}$ ▷ Contribution from w
  end for
  return $\Delta_{(i\ldots n)_1\ldots(i\ldots n)_D}$
end procedure

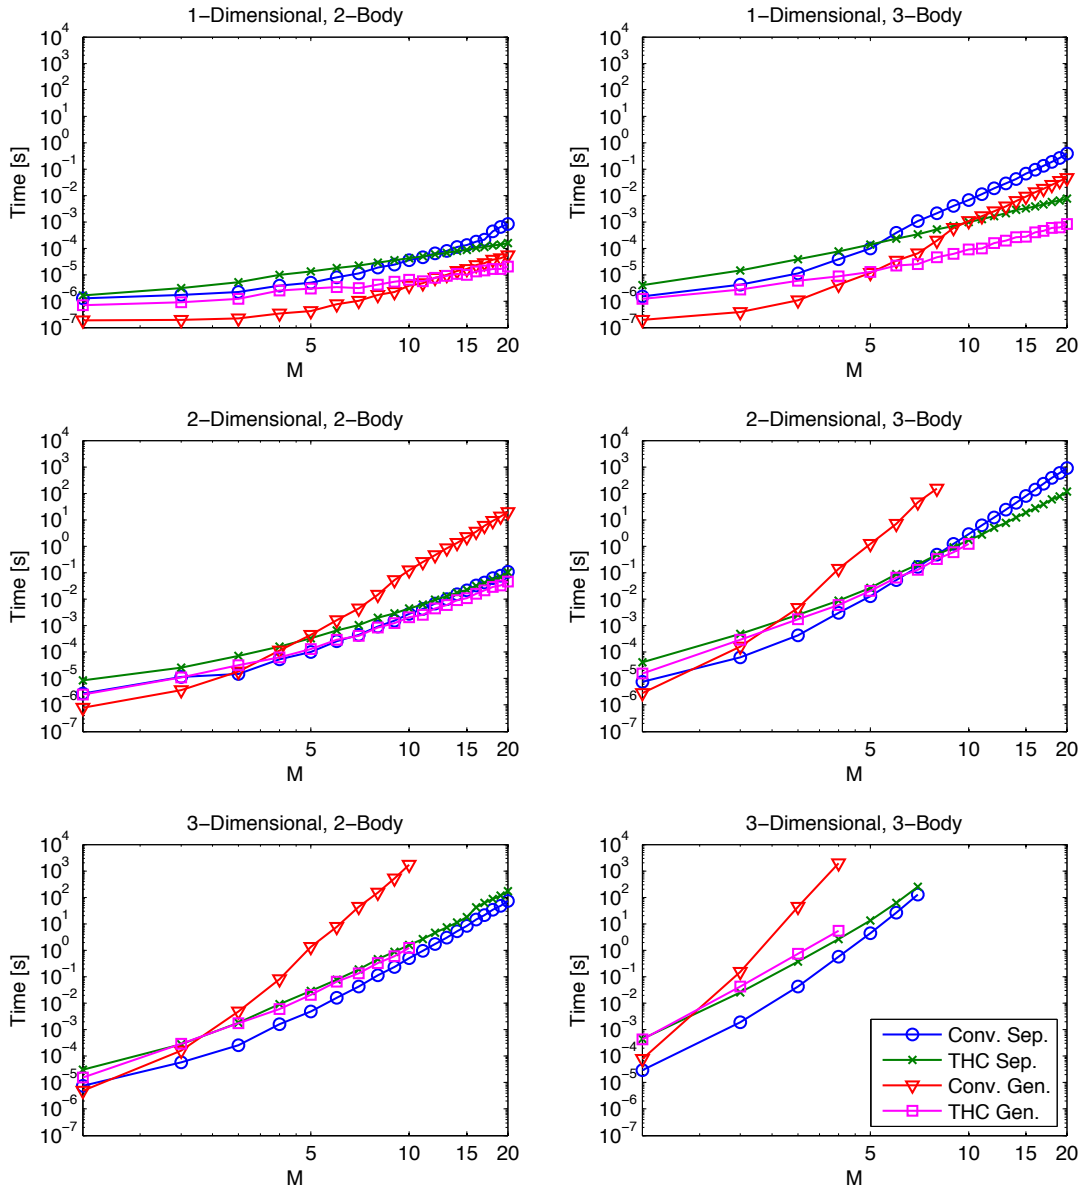

FIG. 3: Timings for generalized pairing tensor formation vs. problem size for various integral technologies, numbers of bodies N, and number of dimensions D.

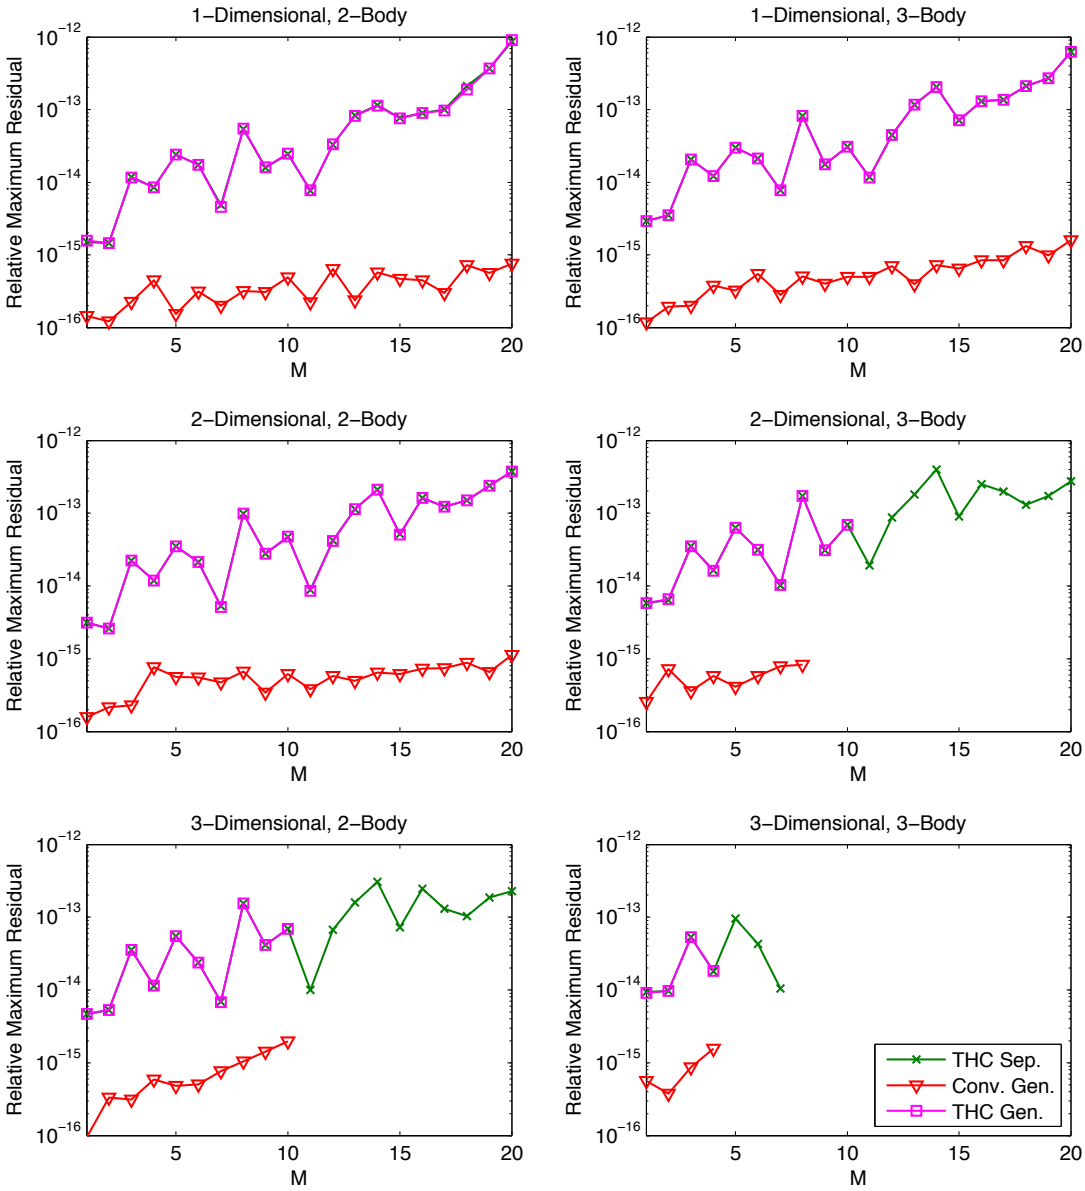

FIG. 4: Relative maximum residual for generalized pairing tensor formation vs. problem size for various integral technologies, numbers of bodies N, and number of dimensions D.

TABLE I: Asymptotic scalings of generalized pairing tensor formation with various integral technologies. Shown are observed scalings based on power-law fit to timing data and corresponding theoretical scaling (in parentheses).

D	N	$\beta_{\mathrm{Conv-Sep}}$	$\beta_{\mathrm{THC-Sep}}$	$\beta_{\mathrm{Conv-Gen}}$	$\beta_{\mathrm{THC-Gen}}$
1	2	4.5 (4)	1.9 (3)	3.9 (4)	1.8 (3)
1	3	5.8 (6)	3.0 (4)	5.5 (6)	3.2 (4)
2	2	5.3 (6)	4.4 (5)	7.3 (8)	4.6 (5)
2	3	8.3 (9)	6.1 (7)	10.5 (12)	5.5 (7)
3	2	7.2 (8)	7.0 (7)	10.3 (12)	5.8 (7)
3	3	8.9 (12)	7.3 (10)	13.7 (18)	7.0 (10)

TABLE II: Predicted THC vs. conventional crossover points for generalized pairing tensor formation based on power-law fit to timing data.

D	N	M_μ^{Gen}	M_μ^{Sep}
1	2	1.224E+01	1.172E+01
1	3	3.429E+00	5.170E+00
2	2	2.194E+00	1.657E+01
2	3	2.304E+00	7.928E+00
3	2	2.024E+00	1.789E+03
3	3	1.650E+00	1.035E+01

Prepared by LLNL under Contract DE-AC52-07NA27344.