

Numerically exact configuration interaction at quadrillion-determinant scale

Received: 26 May 2025

Accepted: 24 October 2025

Published online: 10 December 2025

Agam Shayit^{1,6}, Can Liao^{2,6}, Shiv Upadhyay^{2,6}, Hang Hu³, Tianyuan Zhang², A. Eugene DePrince III⁴, Chao Yang⁵ & Xiaosong Li²✉

The combinatorial growth of configuration interaction (CI) has long limited this formally exact quantum chemistry method to only the smallest molecules. Here, we report a numerically exact CI calculation exceeding one quadrillion (10^{15}) determinants, made possible by a lossless categorical compression strategy within the small-tensor-product distributed active space (STP-DAS) framework. This approach overcomes the traditional memory bottlenecks of CI by a numerically exact compression of the wavefunction representation and reformulating the most computationally demanding matrix–vector operations. Using this method, we performed a fully relativistic CI calculation of the ground state of HBrTe with over 10^{15} complex-valued determinants in just 34.5 h on 1000 computing nodes—the largest CI calculation ever reported. We further achieved fast computation for systems with hundreds of billions of determinants on only a few compute nodes. Extensive benchmarks confirm that the method retains full numerical exactness while cutting memory and computational cost by orders of magnitude. Compared to previous state-of-the-art CI calculations, this work achieves a 1000 times increase in CI space, a 10^6 -fold increase in floating-point operations performed, and a 10^6 -fold improvement in computational speed.

Full configuration interaction (FCI) offers the most complete and accurate description of a molecule's electronic structure within a given basis set, providing the exact spectral solution to the non-relativistic electronic Schrödinger equation^{1–6}. Due to its variational nature, FCI is particularly well-suited for treating relativistic effects, such as spin–orbit and spin–spin couplings beyond perturbation theory^{7–23}, which are fundamentally rooted in the electronic Dirac equation.

At its core, solving the FCI problem reduces to diagonalizing a large many-electron Hamiltonian matrix. This matrix is Hermitian, sparse, and typically diagonally dominant, making it well-suited to iterative diagonalization techniques that can efficiently converge on a few eigenstates without requiring full storage or construction of the entire matrix. However, as the number of determinants grows factorially with system size, reflecting the combinatorial nature of Slater

determinant enumeration within the full Hilbert space, even iterative methods become intractable beyond a certain threshold.

The CI wavefunction is often expressed as a linear combination of Slater determinants, typically generated by excitations from a mean-field self-consistent field (SCF) ground-state reference. In the relativistic regime, this framework must be reformulated using complex-valued 2- or 4-spinor wavefunctions, as required by the Dirac formalism^{24,25}.

Figure 1 illustrates the historical progression of CI implementations, highlighting major breakthroughs in the achievable scale of determinants. Prior to this work, over a span of 35 years (1990–2024), the field advanced from handling billions to trillions of determinants, driven largely by advances in computer hardware technologies. Although CI is amenable to large-scale parallel processing schemes^{26–31},

¹Department of Physics, University of Washington, Seattle, WA, USA. ²Department of Chemistry, University of Washington, Seattle, WA, USA. ³Molecular Engineering and Sciences Institute, University of Washington, Seattle, WA, USA. ⁴Department of Chemistry and Biochemistry, Florida State University, Tallahassee, FL, USA. ⁵Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. ⁶These authors contributed equally: Agam Shayit, Can Liao, Shiv Upadhyay. ✉e-mail: xsli@uw.edu




Fig. 1 | The evolution of the state-of-the-art for CI calculations over time^{26–31}.

Each historical point is colored according to the nature of the key development of the respective work, which we classify as either intrinsic algorithmic developments (in purple) or optimizations on the HPC hardware of the time (in blue). * Only one CI iteration was performed. Source data are provided as a Source data file.

the explosive growth in memory requirements has historically restricted its applicability to only the smallest chemical systems. Relativistic CI is even more limited, due to the intrinsically larger spinor configuration space associated with complex-valued 2- or 4-component wavefunctions. Simply put, enabling CI for practical quantum chemistry applications demands alternative theoretical frameworks and data representations that can circumvent the brute-force enumeration of the CI space.

Many CI-based wavefunction methods aim to approximate the FCI solution. These methods use different types of approximations, which affect the accuracy of the resulting wavefunction. The two main approaches are complete active space CI (CASI) and selected CI (SCI). While both CASI and SCI methods effectively truncate the Hilbert space of the system to a subspace of significant determinants, this significance is determined differently and at different stages of the computation.

In the CASI method^{9,10,29,30,32–38}, it is assumed that only a subspace of the full Hilbert space of the system contains meaningful correlation, and the FCI wavefunction is approximated as the CI wavefunction in the truncated space (the so-called active space). The truncation often leads to an underestimation of dynamic correlation. Applying more computationally demanding methods such as multiconfigurational self-consistent field (MCSCF)^{6–8,12,13,16,39–48}, multireference configuration interaction (MRCI)^{6,19,48–54}, and many-body perturbation theory (MRPT2, CASPT2, NEVPT2, MC-PDFT)^{13,14,22,55–61} is typically required to achieve qualitative and quantitative agreement with experiment.

SCI-based methods estimate the importance of each configuration in the total wavefunction based on a predefined significance criterion, which depends on the chemical problem of interest^{62–78}. Once the most significant determinants are identified, the Hamiltonian is constructed and diagonalized within this reduced space to approximate the FCI wavefunction. As in CASI, SCI approaches are often combined with perturbation theory to recover contributions from neglected configurations and improve quantitative accuracy^{64,65,67–69}.

Even with advances in dimensionality reduction, CI remains fundamentally constrained by memory limitations. State-of-the-art implementations still require explicit storage of either Hamiltonian matrix elements or excitation lists to support on-the-fly matrix-vector operations, commonly referred to as the σ -build^{32,79–81}. For large CI spaces, storing the full Hamiltonian matrix and performing direct diagonalization is clearly impractical. In on-the-fly CI algorithms, the one-electron excitation list, which encodes the allowed excitations between determinants for efficient Hamiltonian construction, scales as $n_e \times (n_h + 1) \times N$, where N is the number of determinants, and n_e and n_h

denote the numbers of electrons and holes (unoccupied orbitals), respectively. Since N increases factorially with system size, the associated memory requirements grow rapidly, making conventional CI calculations infeasible for anything beyond the smallest systems.

For a CI problem involving N determinants, the size of each CI expansion vector scales linearly with N . For instance, a relativistic CI problem with one quadrillion (10^{15} , 100 orbitals, 88 electrons) determinants would require ~16 petabytes (PB) of memory just to store a single CI vector composed of complex-valued double-precision coefficients. While this memory footprint alone makes such problems challenging to tackle, the memory required to store the excitation list can easily scale to the exabyte (EB) regime. This poses a fundamental barrier to scalability, even before considering computational cost.

A recently introduced CI matrix-vector product algorithm³¹ leverages the exact factorization of the active space into small tensor products of distributed active spaces—an approach known as the small-tensor product distributed active space (STP-DAS) framework, illustrated in Fig. 2A. The STP-DAS algorithm reformulates the large CI matrix-vector product into a sequence of small tensor products, each embedded within a distributed active space, computed on-the-fly using string-based methods. This advance exploits the mathematical condition governing the phase relationship between the global address and the local DAS address of any CI matrix element, enabling the use of only small local determinant address strings in the CI matrix-vector build and overcoming the memory bottleneck associated with storing the full excitation list. This formulation enables extensive reuse of Hamiltonian excitation lists, leading to a dramatic reduction in memory demands. For a CI problem involving one quadrillion (10^{15}) determinants, the STP-DAS framework reduces the excitation list memory requirement from 12 exabytes (EB) to 25 gigabytes (GB), an 8-orders-of-magnitude reduction! In addition, by evenly distributing the computation of small tensor products, the STP-DAS algorithm achieves excellent load balance with minimal node-to-node communication overhead, ensuring strong scalability across both single-node and large-scale parallel architectures.

Since the excitation list typically dominates the storage requirements in CI calculations, the STP-DAS framework overcomes a long-standing memory bottleneck and enables CI computations that were previously deemed intractable. By effectively eliminating the memory footprint of the excitation lists, the storage of CI vector coefficients now emerges as the bottleneck in many large-scale CI problems.

Revisiting the earlier CI example of 10^{15} determinants: even after eliminating the memory bottleneck associated with storing excitation lists, the same calculation still demands 16 PB of memory to store the numerically exact coefficients of a single CI expansion vector in an iterative solver. Given that practical CI calculations typically require multiple subspace vectors for convergence, it becomes clear that storing these expansion vectors now represents the dominant memory bottleneck in large CI calculations. To address this challenge, we leverage the locally compressed nature of the STP-DAS framework to efficiently compute ultra-large-scale CI problems involving up to a quadrillion (10^{15}) determinants. This approach yields deterministic, numerically exact solutions and effectively shifts CI calculations from being memory-bound to compute-bound.

Results

Before presenting methodological details and performance benchmarks, we highlight the largest CASI calculation to date for the ground state of HBrTe, enabled by the compression-compatible STP-DAS framework to be introduced herein. HBrTe is a substituted form of a hydrogen chalcogenide where one of the hydrogens was substituted with a bromine atom to decrease the symmetry to the molecule. We performed a relativistic exact two-component^{21,22,82–93} CASI (X2C-CASI) calculation (100 2-spinor orbitals, 88 electrons, complex-valued 1.05×10^{15} 2-spinor determinants) of the ground state of HBrTe




Fig. 2 | Categorical compression within the small-tensor-product distributed active space (STP-DAS) framework. **A** The STP-DAS framework decomposition of a complete active space configuration interaction (CASCI) calculation into a direct sum of categorical excitations. The large excitation lists can be factored into much smaller categorical excitation lists. Purple sections within active spaces represent electron-occupied orbitals. **B** The exact two-component full configuration interaction (X2C-FCI) ground state energy of the Mg^{2+} ion within the cc-pVNZ-DK^{142,143} ($N = 2, 3, 4$) basis sets, along with the extrapolated complete basis set limit. Source data are provided as a Source data file. **C** Average execution time (in seconds) of the compression-compatible STP-DAS algorithm per σ -build of a thallium hydride (TH) test case versus the node count (5 iterations, 1 message passing interface (MPI) process per node, 40 symmetric multiprocessing (SMP) threads per MPI process). Here, **H** is the Hamiltonian, **C** is the CI vector, and σ is their product. The dashed lines illustrate the ideal strong scaling behavior of each CASCI calculation. Source data are provided as a Source data file. **D** The representation of the subspace

expansion vector in a traditional configuration interaction (CI) picture, the decomposed subspace vector in the STP-DAS framework, and the numerically exact compression of the subspace expansion vector in the categorically compression-compatible STP-DAS representation. The color of CI coefficients indicates their configuration category, while their brightness symbolizes their magnitude. White indicates a magnitude of zero. **E** A schematic representation of the lossless, compression-compatible, STP-DAS σ -build algorithm. The Hamiltonian matrix is represented as a heatmap, where brighter elements have larger magnitudes. The color of the vector elements indicates the configuration category of the corresponding CI coefficients. Note that the σ -build preserves categorical compression. **F** An illustration comparing the traditional Davidson preconditioner with the compression-compatible preconditioner to generate successive subspace expansion vectors. The compression-compatible preconditioner appends the subspace with the same effective search direction as the traditional Davidson preconditioner without compromising its compression.

using the compression-compatible STP-DAS framework. We also performed a calculation with the same number of determinants for the ground state of a magnesium atom (see Section S3 of the Supplementary Information). The calculation ran on the National Energy Research Scientific Computing Center's Perlmutter high-performance supercomputer with a total of 1000 nodes (AMD EPYC 7763 Milan, 128,000 compute cores, 512 GB of RAM per node, 200 GB/s NIC, 2 message passing interface (MPI) processes per node, and 64 symmetric multiprocessing (SMP) threads per MPI process).

The excitation list is generated without assuming any symmetry of the target state. Consequently, the calculation is formally performed in a quadrillion-determinant space, with all determinants explicitly included. Because the CI coefficients are complex, the memory footprint is twice that of an analogous non-relativistic calculation.

Moreover, the complex arithmetic makes the computational cost (in FLOP count) equivalent to that of a non-relativistic calculation with more than twice as many determinants.

Table 1 summarizes the results of the HBrTe calculation. The ground-state energy converged in 9 iterations to microhartree precision ($<10^{-6}$ a.u.) with a total runtime of 34.5 h. In each iteration, an additional CI expansion vector was introduced to accelerate convergence, while the compression algorithm dynamically adapted to the expanding vector space, leading to a gradual increase in computational cost. A total of 9 expansion vectors were involved in the σ -build. Despite the enormous configuration space of over one quadrillion (1.05×10^{15}) complex-valued 2-spinor determinants, the average σ -build time per vector remained just 3.8 h.

Table 1 | Details of the convergence of the X2C-CASCI calculation (100 2-spinor orbitals, 88 electrons, 10^{15} 2-spinor determinants) of the ground state of HBrTe (x2c-TZVPall⁹⁶)

Iteration	Duration (s)	Active Energy (E_h)	ΔE (E_h)	$\max\{ \Delta C_i \}$	$\ r\ $ (E_h)
1	122	-9812.080102321897	–	–	2.62×10^{-1}
2	111	-9812.129587025540	4.95×10^{-2}	5.45×10^{-2}	8.73×10^{-2}
3	262	-9812.133341658482	3.75×10^{-3}	1.36×10^{-2}	3.23×10^{-2}
4	710	-9812.133735698319	3.94×10^{-4}	2.26×10^{-3}	1.40×10^{-2}
5	1536	-9812.133784486688	4.88×10^{-5}	1.08×10^{-3}	7.14×10^{-3}
6	3293	-9812.133790409140	5.92×10^{-6}	4.97×10^{-4}	4.09×10^{-3}
7	6432	-9812.133791245860	8.37×10^{-7}	1.31×10^{-4}	2.54×10^{-3}
8	11,565	-9812.133791408300	1.62×10^{-7}	9.17×10^{-5}	1.61×10^{-3}
9	20,384	-9812.133791451150	4.28×10^{-8}	4.25×10^{-5}	9.71×10^{-4}

Each row corresponds to a Davidson iteration. The fourth and fifth columns track the changes in the computed Rayleigh-Ritz eigenpair across iterations, and the sixth column tracks the residual norm, $\|r\| = \|H\mathbf{C} - E\mathbf{C}\|$, where H is the Hamiltonian, $\mathbf{C}$ is the CI vector, σ is their product, and E is the corresponding energy. The preconditioning dropping threshold was taken to be $\varepsilon = 1.0 \times 10^{-5}$. The inactive energy is 417.105787 a.u.

The ground-state energy of the HBrTe molecule from this CI calculation is -9395.028004 a.u. Leveraging the gap theorem^{94,95}, we determine that our $\binom{100}{88}$ X2C-CASCI result lies within 10×10^{-6} a.u. of the true x2c-TZVPall⁹⁶ ground state energy within that active space, well below any chemically meaningful threshold. A detailed analysis is provided in “Methods”.

This work represents a 3-orders-of-magnitude increase in CI space and a 6-orders-of-magnitude increase in FLOP count, which is estimated using $\mathcal{O}(N^2)$, compared to the previous state-of-the-art in CI calculations^{30,31}. Compared to previous state-of-the-art CI calculations, this work also achieves a 6-orders-of-magnitude speedup in time-to-completion, as measured in core seconds per exaFLOP ($\frac{\text{core-second}}{\text{exaFLOP}}$, see Section S2 in the Supplementary Information for analysis). This ultra-large-scale CI calculation is enabled by the STP-DAS-based numerically exact categorical compression scheme, which reduces the memory required to store the 9 CI expansion vectors from 134 PB to less than 500 TB while maintaining a good load balance³¹, making the computation feasible on most existing supercomputing infrastructures.

We now describe the algorithmic developments that enable such CI calculations to be performed on existing supercomputing resources. Detailed algorithms, parallel implementation strategies, and error bound analyses are provided in the Supplementary Information. The central concept of the STP-DAS framework is the systematic partitioning of the full CI orbital space into a collection of distributed active spaces. Within each active space, configurations are further classified into categorical subspaces, rigorously defined by distinct electron occupation patterns, as illustrated in Fig. 2A³¹. The STP-DAS framework reformulates the CI σ -build as a sum of small-tensor products, each uniquely addressed via a global tensor looping structure. The major memory bottleneck associated with storing the excitation list is eliminated by allowing categorical subspaces to share compact, local excitation lists.

In the largest CASCI calculation (10^{15} determinants) presented here, employing 13 distributed active spaces, the STP-DAS approach reduces the excitation list memory requirement from 12×10^9 GB to just 25 GB. However, storing all nine CI expansion vectors for a system with 10^{15} determinants would require -134 PB of memory. On a high-performance computing system such as Perlmutter, this translates to more than 275,000 nodes, each equipped with 512 GB of memory. A straightforward element-wise sparsity treatment, however, does not meet the STP-DAS condition. To overcome this limitation, the following section introduces a categorical compression scheme that achieves the necessary memory reduction, making it possible to carry out large relativistic CI calculations on a medium-sized computing cluster.

With the STP-DAS framework, the memory bottleneck associated with storing CI expansion vectors is effectively eliminated by applying numerically exact categorical compression. The STP-DAS CI expansion vectors take the form

$$\mathbf{C} = \bigoplus_{\mathcal{B}} \mathbf{C}^{\mathcal{B}}, \quad (1)$$

where $\mathcal{B}$ is a category, defined by a unique electron occupation pattern within the distributed active spaces³¹.

The compression scheme, categorical compression (see Fig. 2D), stores the $\mathbf{C}^{\mathcal{B}}$ vectors as compressed sparse column (CSC) vectors. In this format, each categorical expansion vector $\mathbf{C}^{\mathcal{B}}$ is represented by two equally sized arrays: one, $V_{\text{CSC}}^{\mathcal{B}} \equiv \{C_{K^{\mathcal{B}}} : C_{K^{\mathcal{B}}} \neq 0\}$, stores the values of the nonzero coefficients, while the other, $A_{\text{CSC}}^{\mathcal{B}} \equiv \{K^{\mathcal{B}} : C_{K^{\mathcal{B}}} \neq 0\}$, stores their corresponding local addresses. The tensor-loop structure in the STP-DAS σ -build algorithm is reformulated in terms of categorically compressed local addresses together with their corresponding global phase factors (see Fig. 2E).

In contrast to element-wise compression, categorical compression can eliminate all configurations within a category, i.e., skip an entire category at once. This approach is better able to preserve the vectorized structure compared to element-wise compression. The major difficulty of any sparse matrix-vector product algorithm is the lack of a priori knowledge of the location of the nonzero elements. This difficulty leads to a major bottleneck rooted in the nonuniform accesses to memory. The categorical compression localizes nonuniform memory accesses within a category, which is always orders-of-magnitudes smaller than the size of the full CI space. This localized memory access pattern is the central strength of the categorical compression scheme.

Within each category, element-wise compression is still applicable to maximize sparsity. Most importantly, a category-based compression scheme naturally supports distributed small tensor products, taking advantage of the reduced memory footprint and improved parallel load balance of the STP-DAS framework.

In summary, the numerically exact categorical compression introduced here allows the STP-DAS σ -build algorithm to bypass entire categories of determinants while preserving both the vectorized structure and the local addressing scheme of STP-DAS. Because the compression is fully lossless, the omitted determinants have no impact on the resulting Ritz eigenvalue-eigenvector pair.

The final hurdle in reducing CI memory demands lies in the iterative solver. In CI, the Hamiltonian operator is diagonalized iteratively within the full Hilbert space of the system. As a result, the error in the computed Ritz value directly reflects the missing correlation energy in the associated approximate wavefunction. This enables the

Ritz residual, defined as $\mathbf{r} \equiv \|\mathbf{H}\mathbf{C} - E\mathbf{C}\|$, to be computed and appended as an additional CI expansion vector. The norm of the residual provides rigorous bounds on the missing correlation energy relative to the true eigenpair (eigenvector and eigenvalue) of the Hamiltonian in the chosen basis^{94,95,97–103}. A widely used approach that leverages this principle is the Davidson iterative solver¹⁰⁴. Other related methods that use residual norm as the convergence criterion include the locally optimal block preconditioned conjugate gradient (LOBPCG) method¹⁰⁵, the Jacobi-Davidson¹⁰⁶ method and generalized preconditioned locally harmonic residual (GPLHR) method^{107,108}.

By applying numerical or convergence thresholds at various stages of the Davidson method¹⁰⁴, one can exploit the sparsity of newly generated CI expansion vectors. With sufficiently tight thresholds, these vectors can span the part of the Hilbert space required to accurately represent the desired wavefunction and drive the iterative diagonalization to any desired level of precision. The Davidson method¹⁰⁴ utilized the Davidson preconditioner, which generates the i th component of the next trial expansion vector according to

$$t_i \leftarrow \begin{cases} 0, & \text{if } |\lambda - H_{ii}| \text{ is small} \\ \frac{r_i}{\lambda - H_{ii}}, & \text{else} \end{cases} \quad (2)$$

Here, λ is the Ritz value of the current iteration, H_{ii} is the i th element of the diagonal of the Hamiltonian $\mathbf{H}$, and r_i is the i th component of the current residual. Among the various preconditioners employed in the Davidson method^{97,98,103,109–111}, compression-compatible preconditioners, which discard terms t_i below a numerical threshold ε , have been shown to achieve convergence to the exact same results as the traditional Davidson preconditioner^{97–99,101,112}.

In this work, we apply the compression-compatible categorical preconditioner

$$\begin{aligned} s_i &\leftarrow \begin{cases} \frac{r_i}{\lambda - H_{ii}}, & \text{if } |\lambda - H_{ii}| \geq 10^{-12} \\ 0, & \text{else} \end{cases} \\ t_i &\leftarrow \begin{cases} s_i, & \text{if } |s_i| \geq \varepsilon \|\mathbf{s}\| \\ 0, & \text{else} \end{cases} \end{aligned} \quad (3)$$

using the nonzero residual elements r_i . We do this on-the-fly to avoid explicitly storing the prohibitively large diagonal of $\mathbf{H}$ (a dense vector of size N_{dets}). Figure 2F illustrates the expansion of the CI vector space enabled by the compression-compatible categorical preconditioner used in the Davidson method implemented here. Eq. (3) closely resembles the preconditioner proposed in ref. 112, with the key distinction being the inclusion of the Davidson-preconditioned residual norm $\|\mathbf{s}\|$ in the dropping criterion. This facilitates dynamic threshold adjustment: as the iterations progress and $\|\mathbf{s}\|$ decreases, the criterion becomes more stringent. More importantly, the factor of $\|\mathbf{s}\|$ ensures that numerical thresholding is applied to the generated expansion vectors relative to the total norm of their exact (traditional Davidson) counterparts, rather than some fixed cutoff on the absolute values of their entries. This results in a very accurate CI space expansion scheme at the cost of computing and contracting a significant number of Hamiltonian matrix elements (to evaluate $\mathbf{s}$ exactly), which would be intractable without the STP-DAS framework.

Algorithms and pseudocodes of the compression-compatible STP-DAS method are presented in “Methods”, along with discussions on load balancing and parallel implementation. The convergence behavior of the STP-DAS framework equipped with the compression-compatible preconditioner defined in Eq. (3) was evaluated across three systems with varying degrees of electron correlation: the magnesium atom, diatomic nitrogen, and a model carbon nanotube. The results are provided in Section S1 of the Supplementary Information.

The analysis reveals that overly aggressive thresholding can cause the Davidson procedure to stagnate, thereby preventing convergence

to the correct electronic wavefunction. When thresholds are too loose, newly generated trial vectors quickly become linearly dependent on the existing subspace vectors, signaling that the span of the modified subspace has saturated before achieving convergence. However, when the preconditioning threshold satisfies

$$\varepsilon \lesssim \frac{10}{\sqrt{N_{\text{dets}}}}, \quad (4)$$

the resulting energies agree with their exact values to better than $10^{-7} E_h$, and the residual norms become correspondingly small, demonstrating successful and reliable convergence.

The CI wavefunction of highly correlated systems is comprised of a large number of determinants with small CI coefficients. Because the Hilbert space of the problem is never truncated, and no determinants are discarded from the wavefunction itself. As a result, the compression-compatible preconditioner easily facilitates convergence to the exact wavefunction, provided that the preconditioning threshold ε is sufficiently small for the true eigenvector to be accurately represented in the subspace spanned by the modified expansion vectors.

With the capability to perform large CI calculations, energetic extrapolation to the correlation limit becomes feasible for many-electron systems. Figure 2B illustrates the correlation-consistent extrapolation for the Mg^{2+} ion using double-zeta (DZ, 36 orbitals), triple-zeta (TZ, 68 orbitals), and quadruple-zeta (QZ, 118 orbitals) basis sets, involving 2.54×10^8 , 2.91×10^{11} , and 9.75×10^{13} 2-spinor determinants, respectively. The complete basis set limit of -199.16704295 a.u. was obtained using a mixed Gaussian extrapolation scheme tailored for correlation-consistent basis sets¹¹³.

To demonstrate the strong scaling behavior of the compression-compatible STP-DAS σ -build algorithm, we performed relativistic X2C-CASCI calculations on the thallium hydride (TIH) molecule using active spaces of 40 and 41 2-spinor orbitals and 24 electrons, $\begin{pmatrix} 40 \\ 24 \end{pmatrix}$ and $\begin{pmatrix} 41 \\ 24 \end{pmatrix}$, corresponding to 63 and 152 billion 2-spinor determinants, respectively. Five distributed active spaces (DASs) were employed. These calculations were executed on the University of Washington's Hyak HPC system, a small-sized cluster where each node is equipped with two Intel Xeon 6230 Gold CPUs and a single 100 GB/s network interface card. As shown in Fig. 2C, even with just 5 compute nodes, relativistic X2C-CASCI calculations involving tens to hundreds of billions of determinants require only 4–6 min per σ -build on average. Increasing to 30 nodes further reduces the cost to just over 1 min. Past 30 nodes, the communication time dominates the runtime and the calculations no longer scale. This benchmark demonstrates that billion- and even trillion-determinant CI calculations are now feasible on a small-scale computing cluster.

Additionally, we performed X2C-CASCI calculations for the ground states of two highly correlated systems, showcasing the applicability of the compression-compatible STP-DAS framework to highly correlated systems. These systems were chosen from opposite ends of the correlation spectrum from strongly statically correlated to strongly dynamically correlated. Square Rb_4 , a relativistic analog of H_4 ^{114–126}, displays strong static correlation and Xe_2 is a dynamically correlated noble gas dimer^{127–135}.

Tables 2 and 3 summarize the results of the Rb_4 and Xe_2 calculations. The Rb_4 calculation (50 2-spinor orbitals, 28 electrons, 8.9×10^{13} 2-spinor determinants) ran on the National Energy Research Scientific Computing Center's Perlmutter high-performance supercomputer with a total of 100 nodes (AMD EPYC 7763 Milan, 12,800 compute cores, 512 GB of RAM per node, 200 GB/s NIC, 1 MPI processes per node, and 128 SMP threads per MPI process) and took 6 iterations and

Table 2 | Details of the convergence of the X2C-CASCI calculation (50 2-spinor orbitals, 28 electrons, 8.9×10^{13} 2-spinor determinants) of the ground state of Rb_4 (cc-pvtz-x2c¹⁴¹)

Iteration	Active Energy (E_h)	ΔE (E_h)	$\max\{ \Delta C_i \}$	$\ r\ $ (E_h)
1	-90.038089795466	–	–	4.48×10^{-2}
2	-90.048919589524	1.08×10^{-2}	1.21×10^{-1}	1.74×10^{-2}
3	-90.049940164869	1.02×10^{-3}	1.64×10^{-2}	4.95×10^{-3}
4	-90.050002638482	6.25×10^{-5}	2.53×10^{-3}	1.20×10^{-3}
5	-90.050006376495	3.74×10^{-6}	1.37×10^{-3}	2.96×10^{-4}
6	-90.050006593144	2.17×10^{-7}	1.55×10^{-4}	7.36×10^{-5}

Each row corresponds to a Davidson iteration. The third and fourth columns track the changes in the computed Rayleigh-Ritz eigenpair across iterations, and the fifth column tracks the residual norm, $\|r\| = \|H\mathbf{C} - E\mathbf{C}\|$, where H is the Hamiltonian, $\mathbf{C}$ is the CI vector, σ is their product, and E is the corresponding energy. The preconditioning dropping threshold was taken to be $\epsilon = 1.0 \times 10^{-6}$. The inactive energy is -11826.102719 a.u.

Table 3 | Details of the convergence of the X2C-CASCI calculation (60 2-spinor orbitals, 12 electrons, 1.4×10^{12} 2-spinor determinants) of the ground state of Xe_2 (x2c-TZVPall-2c⁹⁶)

Iteration	Active Energy (E_h)	ΔE (E_h)	$\max\{ \Delta C_i \}$	$\ r\ $ (E_h)
1	-21.059579374232	–	–	4.76×10^{-1}
2	-21.194173478231	1.35×10^{-1}	3.90×10^{-2}	1.86×10^{-1}
3	-21.206728805229	1.26×10^{-2}	8.89×10^{-3}	5.62×10^{-2}
4	-21.207569776081	8.41×10^{-4}	8.74×10^{-4}	1.75×10^{-2}
5	-21.207634517148	6.47×10^{-5}	4.51×10^{-4}	7.88×10^{-3}
6	-21.207644283591	9.77×10^{-6}	1.19×10^{-4}	4.48×10^{-3}
7	-21.207646647756	2.36×10^{-6}	3.90×10^{-5}	2.93×10^{-3}

Each row corresponds to a Davidson iteration. The third and fourth columns track the changes in the computed Rayleigh-Ritz eigenpair across iterations, and the fifth column tracks the residual norm, $\|r\| = \|H\mathbf{C} - E\mathbf{C}\|$, where H is the Hamiltonian, $\mathbf{C}$ is the CI vector, σ is their product, and E is the corresponding energy. The preconditioning dropping threshold was taken to be $\epsilon = 1.0 \times 10^{-5}$. The inactive energy is -14868.439049 a.u.

11.8 h to converge. The Xe_2 calculation (60 2-spinor orbitals, 12 electrons, 1.4×10^{12} 2-spinor determinants) ran on the same platform with 256 nodes and took 7 iterations and 36.1 h to converge.

The ground-state energies of the Rb_4 and Xe_2 molecules from these calculations are -11916.152725 a.u. and -14889.646696 a.u., accordingly. The gap theorem^{94,95} guarantees that these X2C-CASCI results lie within 0.53×10^{-6} a.u. and 17.56×10^{-6} a.u. of the true ground state energies within the corresponding active spaces and basis sets (see “Methods”).

In summary, by combining compression-compatible preconditioners with compression-compatible categorical CI vectors, the STP-DAS framework drastically reduces the memory footprint of both the excitation lists and the CI expansion vectors. In the largest relativistic CASCI calculation presented here, spanning 10^{15} determinants across 13 distributed active spaces, the STP-DAS approach reduces the memory required for the excitation list from 12×10^9 GB to just 25 GB, and for 9 CI expansion vectors from 134 PB to less than 500 TB. These reductions make quadrillion-determinant calculations tractable on current supercomputing architectures. While most of the community may not have access to the hundreds of compute nodes required for such runs, this work also demonstrates the practical feasibility of trillion-determinant calculations on just a few nodes and even on a laptop.

Discussion

In this work, we conducted a relativistic configuration interaction (CI) calculation for the ground state of HBrTe in a quadrillion-determinantal space. This calculation was enabled by numerically

exact categorical compression within the STP-DAS framework, which effectively eliminates the memory bottlenecks associated with storing both excitation lists and CI expansion vectors. Compared to previous state-of-the-art CI calculations, this work represents a 3-orders-of-magnitude increase in CI space, a 6-orders-of-magnitude increase in FLOP count, and a 6-orders-of-magnitude improvement in computational speed.

We introduced a categorically compressed representation of the CI expansion vectors and reformulated the STP-DAS σ -build algorithm to take advantage of this structure. By expressing the global expansion vector as a direct sum of compressed local components, the algorithm efficiently skips all coefficients that do not contribute to the categorical σ -vector. This approach is further enabled by a compression-compatible preconditioner, which generates compressed expansion directions within the Davidson procedure.

The resulting categorically compressed STP-DAS σ -build algorithm demonstrates excellent strong scaling behavior and yields dramatic reductions in both runtime and memory footprint. These benefits extend seamlessly to both relativistic (two- and four-component) and non-relativistic CI calculations. To highlight this capability, we computed the $\binom{100}{88}$ X2C-CASCI ground-state energy of HBrTe

using over one quadrillion (10^{15}) complex-valued 2-spinor determinants. The categorically compressed STP-DAS approach spans 10^{15} determinants across 13 distributed active spaces, reducing the memory required for the excitation list from 12×10^9 GB to only 25 GB, and for nine CI expansion vectors from 134 PB to under 500 TB. It converges the ground-state wavefunction of HBrTe in just nine iterations over a 34.5-h runtime. This achievement represents the largest CI calculation reported to date. Additionally, we achieved σ -build times of just 5 minutes for systems with ~150 billion complex-valued 2-spinor determinants using only a few compute nodes. The capability to perform large CI calculations makes basis set extrapolations to the complete basis set limit and computations on highly correlated molecular systems readily achievable with CI.

The integration of categorical compression with STP-DAS marks a paradigm shift in tackling large-scale CI problems. As quantum chemistry continues to push the limits of system complexity, the ability to carry out quadrillion-determinant calculations within tractable resource bounds establishes a powerful foundation for studying highly correlated, multireference, relativistic systems. While access to hundreds of compute nodes for quadrillion-determinant calculations may remain out of reach for most of the community, this work demonstrates the practical feasibility of trillion-determinant calculations on a small cluster.

For transition-metal, rare-earth, and heavy-element complexes, such large-scale CI calculations enable predictive simulations of electronic structure properties (bond order, covalency, polarization, etc.), spectroscopic observables (UV/Vis, X-ray, etc.), and reaction pathways, with the full orbital space consisting of both metal and ligand orbitals, treated on an equal footing.

The ability to simulate a full CI space of 100 orbitals on a classical computer not only challenges current notions of quantum supremacy, but also establishes a robust platform for developing and benchmarking quantum algorithms aimed at achieving chemical accuracy.

Methods

Lossless σ -build using the categorical compression of small tensor products

The categorical σ -build algorithm within STP-DAS³¹ can be reformulated to exploit the categorical compression of the expansion vectors. The compact nature of the categorical representation enables fast and memory-efficient computation of σ -vectors. The categorical σ -build

algorithm implements the evaluation of

$$\sigma_{L^A} = {}^{1e}\sigma_{L^A} + {}^{2e}\sigma_{L^A}, \quad (5)$$

$${}^{1e}\sigma_{L^A} = \sum_B \sum_{\mathbb{K}_\mu^B \oplus \mathbb{K}_\nu^B} \sum_{pq} P_{\mu\nu} \delta_{\mathbb{X}_\mu^A \mathbb{X}_\nu^B} h'_{pq} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{K}_\mu^B \oplus \mathbb{K}_\nu^B \rangle C_{K^B}, \quad (6)$$

$${}^{2e}\sigma_{L^A} = \frac{1}{2} \sum_{CB} \sum_{\mathbb{J}_\mu^C \oplus \mathbb{J}_\nu^C} \sum_{\mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C} \sum_{\mathbb{K}_\mu^B \oplus \mathbb{K}_\nu^B} \sum_{pqrs} P_{\mu\nu} P_{\kappa\lambda} \delta_{\mathbb{X}_\mu^A \mathbb{X}_\nu^B} \delta_{\mathbb{X}_\kappa^A \mathbb{X}_\lambda^B} g_{pqrs} \langle \mathbb{L}_\mu^A \oplus \mathbb{L}_\nu^A | \hat{E}_{pq} | \mathbb{J}_\mu^C \oplus \mathbb{J}_\nu^C \rangle \langle \mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C | \hat{E}_{rs} | \mathbb{K}_\mu^B \oplus \mathbb{K}_\nu^B \rangle C_{K^B}, \quad (7)$$

where $p \in \mathbb{X}_\mu^A$, $q \in \mathbb{X}_\nu^C$, $r \in \mathbb{X}_\kappa^C$, $s \in \mathbb{X}_\lambda^B$, $\langle \mathbb{X}_\mu^A \oplus \mathbb{X}_\nu^A | \hat{E}_{pq} | \mathbb{X}_\mu^B \oplus \mathbb{X}_\nu^B \rangle$ are categorical one-electron excitation lists, $P_{\mu\nu}$ are global phase factors, and h'_{pq} (g_{pqrs}) are one (two) body Hamiltonian elements. See ref. 31 for algorithmic details.

Equations (5) to (7) define the categorical σ -vector in terms of local STP-DAS one-electron excitation lists and the categorical CI expansion vector. Notably, when the expansion coefficients C_{K^B} are categorically compressed, the resulting σ -vector coefficients σ_{L^A} are also categorically compressed. In such cases, the categorical σ -build reduces to a contraction between a categorically compressed expansion vector and categorically compressed STP-DAS one-electron excitation lists. Thus, Eqs. (5) to (7) yield a categorically compressed σ -vector, in a manner directly analogous to the compression-preserving behavior of sparse matrix-sparse vector products (SpMSpV). Importantly, this compression preservation is general and independent of the specific storage format used for the categorically compressed representations.

The categorically compressed representation can be implemented in various forms, with the choice of storage format guided primarily by computational efficiency. Since the categorical σ -build algorithm often involves reading numerous expansion coefficients with increasing local addresses during contraction, it is natural to adopt a compressed sparse column (CSC) format for storing the categorical expansion coefficients.

Eigenvalue bound analysis

We wish to apply the gap theorem^{94,95,136} to bound the error in the computed X2C-CASCI ground state energy:

$$|\delta E| \leq \frac{\|\mathbf{r}\|^2}{\gamma_0}, \quad (8)$$

where $\mathbf{r}$ is the Ritz residual of the computed ground state and the gap $\gamma_0 \equiv E_1 - \tilde{E}_0$ is the difference between E_1 , the (unknown) exact energy of the first excited state, and $\tilde{E}_0$, the computed Ritz value of the ground state. Because γ_0 is unknown, one can estimate its order-of-magnitude using approximate methods or use experimental values to compute a surrogate for the true gap. One can also obtain an exact lower bound on the gap by including the posterior error bound of the first excited state^{136–138} in the Davidson calculation¹⁰⁰:

$$\gamma_0 = E_1 - \tilde{E}_0 \geq (\tilde{E}_1 - \|\mathbf{r}_1\|) - \tilde{E}_0 \equiv \gamma_0^-, \quad (9)$$

where $\mathbf{r}_1$ is the residual associated with the Ritz value $\tilde{E}_1$.

Using the gap theorem^{94,95,136}, we can place an exact bound on the error in the computed X2C-CASCI ground-state energy. This requires an estimate of the energy gap between the ground and first excited states of the X2C-CASCI Hamiltonian. To obtain this, we performed an X2C-CISD calculation for the two lowest-lying states and determined a

gap of ~ 0.095 a.u. for HBrTe. Based on this estimate, the gap theorem bounds the error in our X2C-CASCI ground-state energy for HBrTe to within 10 microhartree, which is well below any chemically meaningful threshold.

Compression-compatible STP-DAS algorithm

Algorithm 1 shows the categorically compressed STP-DAS algorithm, in which only nonzero elements contribute to the categorical σ -vector. Its advantage over the traditional STP σ -build algorithm³¹ is twofold: in the outermost loop, where we skip entire categories whose expansion vector vanishes (see lines 2–3), and in the inner loop of line 7, where we only process excitations $\langle \mathbb{J}_\kappa^C \oplus \mathbb{J}_\lambda^C | \hat{E}_{rs} | \mathbb{K}_\mu^B \oplus \mathbb{K}_\nu^B \rangle$ for which $C_{K^B} \neq 0$.

Algorithm 1: Two-electron σ -build using categorical compression. Bold text represents algorithmic logic and typewritten text represents comments.

Data: the categorically compressed expansion vector $\mathbf{C} = \bigoplus_B (V_{\text{CSC}}^B, A_{\text{CSC}}^B)$ and integrals g_{pqrs}

Result: the categorically compressed contraction ${}^{2e}\sigma$

```

1 for categories A, B, C and DAS indices μ, ν, κ, λ for which δμνAκλB δκλCμνB = 1 parallel do
2   if |ACSCB| = 0 then
3     continue to next iteration
4   for JμνκλC do
5     // Obtain the sorted DAS one-electron excitation list of JμνκλC
6     A ← {KκB ⊕ KλB : ∃ r ∈ XκC, ∃ s ∈ XλB s.t. ⟨JκC ⊕ JλC | Ers | KκB ⊕ KλB⟩ ≠ 0}
7     // Intersect ACSCB with A to contract nonzero expansion coefficients
8     i, j ← 0
9     while i < |A| and j < |ACSCB| do
10      if A[i] < ACSCB[j] then
11        // Advance i to the lowest i' for which Ai ≥ ACSCB[j]
12        exponentialSearch(A, i, ACSCB[j])
13      else if A[i] > ACSCB[j] then
14        // Advance j to the lowest j' for which ACSCB[j'] ≥ A[i]
15        exponentialSearch(ACSCB, j, A[i])
16      else
17        // Process the σ contribution at the address LA into a
18        // thread-safe hash table belonging to category A
19        for p ∈ XμA, q ∈ XνC do
20          2eσLA ← 2eσLA + Pμν Pκλ gpqrs ⟨LμA ⊕ LνA | Epq | JμC ⊕ JνC⟩
21          × ⟨JκC ⊕ JλC | Ers | KκB ⊕ KλB⟩ VCSCB[j]
22        i ← i + 1
23        j ← j + 1
24 Subroutine: exponentialSearch(A, i, x)
25   Result: An index i < j < |A| such that A[j] ≥ x or j = |A| if nonexistent
26   δ ← 1
27   while j < |A| and A[j] < x do
28     j ← j + δ
29     δ ← 2δ
30   jmin ← j - δ/2
31   jmax ← min{j, |A| - 1}
32   // Return the lower bound of x in the range A[jmin...jmax]
33   return binarySearch(A[jmin...jmax], x)

```

There is a potential workload imbalance associated with Algorithm 1: because the collection of categorical expansion vectors is distributed among computing nodes, the contraction workload of a given node is proportional to the number of nonzero categorical expansion coefficients it has. We alleviated some of the resulting computational delay by implementing passive one-sided MPI communication of categorical expansion vectors using remote memory access (RMA). This allows idle nodes to contract more categorical expansion vectors with their excitation lists without waiting for the corresponding busy nodes to broadcast them. Ultimately, overcoming this load-balancing issue requires dynamically redistributing categorical expansion vectors according to their sparsity, which changes during the iterative diagonalization.

As illustrated in Algorithm 1, the computational cost, both in memory and runtime, of the compression-compatible STP-DAS σ -build procedure increases with the density of the expansion vectors. Therefore, it is essential to maintain maximal compression in these vectors. To achieve this, we replace the traditional Davidson preconditioner with a compression-compatible alternative for generating new trial expansion vectors. This modification alters only the subspace expansion strategy in the Davidson algorithm, while preserving exact

treatment of the full determinantal space. As a result, the computed matrix-vector product $\mathbf{HC}$ and the corresponding residual norm $\|\mathbf{r}\| = \|\mathbf{HC} - \lambda\mathbf{C}\|$ remain exact, unlike in selected CI and other truncated approaches, where both the Hamiltonian and CI vectors are explicitly approximated.

The compression-compatible STP-DAS σ -build algorithm significantly reduces the overall workload associated with the σ build. The reduction in workload can be nonuniform: the contraction workload associated with a determinant $\mathbb{J}_{\mu\nu\kappa\lambda}^C$ is proportional to the number of local addresses containing both nonzero categorical expansion vectors elements and Hamiltonian matrix elements. To improve the load-balance, we implemented dynamic SMP thread-level parallelism in the outermost loop (line 1 in Algorithm 1) instead of in the loop over determinants $\{\mathbb{J}_{\mu\nu\kappa\lambda}^C\}$ (line 4 in Algorithm 1). Under dynamic parallelism, some threads execute many light contractions, while others execute fewer heavy contractions, resulting in a more uniform distribution of contraction workload. Such dynamic parallelism is ineffective in the loop over determinants $\{\mathbb{J}_{\mu\nu\kappa\lambda}^C\}$ due to the small tensor product nature of the STP-DAS framework.

Data availability

All datasets discussed in this work are included in the Supplementary Information/Source data files. The atomic coordinates for the HBrTe, Rb_4 , Xe_2 , Mg, N_2 , and carbon nanotube systems are available as Supplementary Data 1. Source data are provided with this paper.

Code availability

The compression-compatible STP-DAS framework is implemented in the development version of the Chronus Quantum software package³⁹. The Chronus Quantum software package is open-source software, distributed under the GNU General Public License v2 (GPLv2). All calculations were performed using executables compiled with GCC version 13.2. The STP-DAS implementation is currently accessible within the developers' community. It will be included in the next official Chronus public release, following extensive testing across multiple platforms to ensure robustness, scalability, and reproducibility. In the meantime, we created an MPI-compatible, portable, and containerized version of Chronus Quantum containing the compression-compatible STP-DAS implementation. This version is freely available, along with the relevant input files, in Figshare⁴⁰.

References

- Shavitt, I. In *Methods of Electronic Structure Theory* (ed. Schafer, H. F. I.) 189–275 (Plenum, 1977).
- Shavitt, I. The history and evolution of configuration interaction. *Mol. Phys.* **94**, 3–17 (1998).
- Sherrill, C. D. & Schaefer, H. F. The configuration interaction method: advances in highly correlated approaches. *Adv. Quantum Chem.* **34**, 143–269 (1999).
- Čársky, P. In *Encyclopedia of Computational Chemistry* (ed. Schleyer, P. V. R.) (John Wiley & Sons, Ltd, 2002).
- Karwowski, J. A. & Shavitt, I. In *Handbook of Molecular Physics and Quantum Chemistry* (ed. Wilson, S.) (John Wiley & Sons, Ltd, 2003).
- Szalay, P. G., Müller, T., Gidofalvi, G., Lischka, H. & Shepard, R. Multiconfiguration self-consistent field and multireference configuration interaction methods and applications. *Chem. Rev.* **112**, 108–181 (2012).
- Jørgen Aa. Jensen, H., Dyall, K. G., Saue, T. & Fægri, K. Relativistic four-component multiconfigurational self-consistent-field theory for molecules: formalism. *J. Chem. Phys.* **104**, 4083–4097 (1996).
- Thyssen, J., Fleig, T. & Jensen, H. J. A. A direct relativistic four-component multiconfiguration self-consistent-field method for molecules. *J. Chem. Phys.* **129**, 034109 (2008).
- Knecht, S., Jensen, H. J. A. & Fleig, T. Large-scale parallel configuration interaction. I. Nonrelativistic and scalar-relativistic general active space implementation with application to $(\text{Rb}-\text{Ba})^+$. *J. Chem. Phys.* **128**, 014108 (2008).
- Knecht, S., Jensen, H. J. A. & Fleig, T. Large-scale parallel configuration interaction. II. Two- and four-component double-group general active space implementation with application to BiH. *J. Chem. Phys.* **132**, 014108 (2010).
- Knecht, S., Legeza, O. & Reiher, M. Communication: four-component density matrix renormalization group. *J. Chem. Phys.* **140**, 041101 (2014).
- Bates, J. E. & Shiozaki, T. Fully relativistic complete active space self-consistent field for large molecules: quasi-second-order minimax optimization. *J. Chem. Phys.* **142**, 044112 (2015).
- Shiozaki, T. & Mizukami, W. Relativistic internally contracted multireference electron correlation methods. *J. Chem. Theory Comput.* **11**, 4733–4739 (2015).
- Vlaisavljevich, B. & Shiozaki, T. Nuclear energy gradients for internally contracted complex active space second-order perturbation theory: multistate extensions. *J. Chem. Theory Comput.* **12**, 3781–3787 (2016).
- Almoukhalalati, A., Knecht, S., Jensen, H. J. A., Dyall, K. G. & Saue, T. Electron correlation within the relativistic no-pair approximation. *J. Chem. Phys.* **145**, 074104 (2016).
- Reynolds, R. D., Yanai, T. & Shiozaki, T. Large-scale relativistic complete active space self-consistent field with robust convergence. *J. Chem. Phys.* **149**, 014106 (2018).
- Battaglia, S., Keller, S. & Knecht, S. Efficient relativistic density-matrix renormalization group implementation in a matrix-product formulation. *J. Chem. Theory Comput.* **14**, 2353–2369 (2018).
- Knecht, S., Jensen, H. J. A. & Saue, T. Relativistic quantum chemical calculations show that the uranium molecule U_2 has a quadruple bond. *Nat. Chem.* **11**, 40–44 (2019).
- Hu, H. et al. Relativistic two-component multireference configuration interaction method with tunable correlation space. *J. Chem. Theory Comput.* **16**, 2975–2984 (2020).
- Jenkins, A. J., Hu, H., Lu, L., Frisch, M. J. & Li, X. Two-component multireference restricted active space configuration interaction for the computation of L-edge X-ray absorption spectra. *J. Chem. Theory Comput.* **18**, 141–150 (2022).
- Sharma, P. et al. Exact-two-component multiconfiguration pair-density functional theory. *J. Chem. Theory Comput.* **18**, 2947–2954 (2022).
- Lu, L., Hu, H., Jenkins, A. J. & Li, X. Exact-two-component relativistic multireference second-order perturbation theory. *J. Chem. Theory Comput.* **18**, 2983–2992 (2022).
- Hoyer, C. E. et al. Correlated Dirac–Coulomb–Breit multiconfigurational self-consistent-field methods. *J. Chem. Phys.* **158**, 044101 (2023).
- Dyall, K. G. & Fægri, Jr, K. *Introduction to Relativistic Quantum Chemistry* (Oxford University Press, 2007).
- Reiher, M. & Wolf, A. *Relativistic Quantum Chemistry*, 2nd ed. (Wiley-VCH, 2015).
- Olsen, J., Jørgensen, P. & Simons, J. Passing the one-billion limit in full configuration-interaction (FCI) calculations. *Chem. Phys. Lett.* **169**, 463–472 (1990).
- Ansaloni, R., Bendazzoli, G. L., Evangelisti, S. & Rossi, E. A parallel full-CI algorithm. *Comp. Phys. Comm.* **128**, 496–515 (2000).
- Gan, Z. & Harrison, R. Calibrating quantum chemistry: a multi-teraflop, parallel-vector, full-configuration interaction program for the Cray-X1. In *SC '05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing 22–22* (2005).

29. Vogiatzis, K. D., Ma, D., Olsen, J., Gagliardi, L. & De Jong, W. A. Pushing configuration-interaction to the limit: towards massively parallel MCSCF calculations. *J. Chem. Phys.* **147**, 184111 (2017).
30. Gao, H., Imamura, S., Kasagi, A. & Yoshida, E. Distributed implementation of full configuration interaction for one trillion determinants. *J. Chem. Theory Comput.* **20**, 1185–1192 (2024).
31. Hu, H. et al. Small tensor product distributed active space (STPDAS) framework for relativistic and non-relativistic multi-configuration calculations: scaling from 10^9 on a laptop to 10^{12} determinants on a supercomputer. *Chem. Phys. Rev.* **5**, 041404 (2024).
32. Knowles, P. & Handy, N. A new determinant-based full configuration interaction method. *Chem. Phys. Lett.* **111**, 315–321 (1984).
33. Roos, B. O. The complete active space self-consistent field method and its applications in electronic structure calculations. *Adv. Chem. Phys.* **69**, 399–445 (1987).
34. Olsen, J., Roos, B. O., Jørgensen, P. & Jensen, H. J. A. Determinant based configuration interaction algorithms for complete and restricted configuration interaction spaces. *J. Chem. Phys.* **89**, 2185–2192 (1988).
35. Klene, M., Robb, M. A., Frisch, M. J. & Celani, P. Parallel implementation of the CI-vector evaluation in full CI/CAS-SCF. *J. Chem. Phys.* **113**, 5653–5665 (2000).
36. Fleig, T., Olsen, J. & Marian, C. M. The generalized active space concept for the relativistic treatment of electron correlation. I. Kramers-restricted two-component configuration interaction. *J. Chem. Phys.* **114**, 4775–4790 (2001).
37. Fleig, T., Olsen, J. & Visscher, L. The generalized active space concept for the relativistic treatment of electron correlation. II. Large-scale configuration interaction implementation based on relativistic 2- and 4-spinors and its application. *J. Chem. Phys.* **119**, 2963–2971 (2003).
38. Fleig, T., Jensen, H. J. A., Olsen, J. & Visscher, L. The generalized active space concept for the relativistic treatment of electron correlation. III. Large-scale configuration interaction and multi-configuration self-consistent-field four-component methods with application to UO_2 . *J. Chem. Phys.* **124**, 104106 (2006).
39. Klene, M., Robb, M. A., Blancafort, L. & Frisch, M. J. A new efficient approach to the direct restricted active space self-consistent field method. *J. Chem. Phys.* **119**, 713–728 (2003).
40. Ma, D., Li Manni, G. & Gagliardi, L. The generalized active space concept in multiconfigurational self-consistent field methods. *J. Chem. Phys.* **135**, 044128 (2011).
41. Vogiatzis, K. D., Li Manni, G., Stoneburner, S. J., Ma, D. & Gagliardi, L. Systematic expansion of active spaces beyond the CASSCF limit: a GASSCF/SplitGAS benchmark study. *J. Chem. Theory Comput.* **11**, 3010–3021 (2015).
42. Weser, O., Guther, K., Ghanem, K. & Li Manni, G. Stochastic generalized active space self-consistent field: theory and application. *J. Chem. Theory Comput.* **18**, 251–272 (2022).
43. Malmqvist, P. Å. Calculation of transition density matrices by nonunitary orbital transformations. *Int. J. Quant. Chem.* **30**, 479–494 (1986).
44. Werner, H.-J. & Meyer, W. A quadratically convergent MCSCF method for the simultaneous optimization of several states. *J. Chem. Phys.* **74**, 5794–5801 (1981).
45. Ganyushin, D. & Neese, F. A fully variational spin-orbit coupled complete active space self-consistent field approach: application to electron paramagnetic resonance g-tensors. *J. Chem. Phys.* **138**, 104113 (2013).
46. Malmqvist, P. Å & Roos, B. O. The CASSCF state interaction method. *Chem. Phys. Lett.* **155**, 189–194 (1989).
47. Zhang, B., Vandezande, J. E., Reynolds, R. D. & Schaefer, H. F. Spin-orbit coupling via four-component multireference methods: benchmarking on p-block elements and tentative recommendations. *J. Chem. Theory Comput.* **14**, 1235–1246 (2018).
48. Lischka, H. et al. Columbus—a program system for advanced multireference theory calculations. *WIREs Comput. Mol. Sci.* **1**, 191–199 (2011).
49. Liu, B. Ab initio potential energy surface for linear H_3 . *J. Chem. Phys.* **58**, 1925–1937 (1973).
50. Roos, B. O., Taylor, P. R. & Sigbahn, P. E. M. A complete active space SCF method (CASSCF) using a density matrix formulated super-CI approach. *Chem. Phys.* **48**, 157–173 (1980).
51. Lischka, H., Shepard, R., Brown, F. B. & Shavitt, I. New implementation of the graphical unitary group approach for multireference direct configuration interaction calculations. *Int. J. Quant. Chem.* **20**, 91–100 (1981).
52. Werner, H.-J. & Knowles, P. J. An efficient internally contracted multiconfiguration-reference configuration interaction method. *J. Chem. Phys.* **89**, 5803–5814 (1988).
53. Malmqvist, P.-Å & Roos, B. O. The restricted active space self-consistent-field method, implemented with a split graph unitary group approach. *J. Phys. Chem.* **94**, 5477–5482 (1990).
54. Hirao, K. (ed.). *Recent Advances in Multireference Methods* (World Scientific, 1999).
55. Andersson, K., Malmqvist, P. A., Roos, B. O., Sadlej, A. J. & Wolinski, K. Second-order perturbation theory with a CASSCF reference function. *J. Chem. Phys.* **94**, 5483–5488 (1990).
56. Andersson, K., Malmqvist, P. & Roos, B. O. Second-order perturbation theory with a complete active space self-consistent field reference function. *J. Chem. Phys.* **96**, 1218–1226 (1992).
57. Roos, B. O. et al. Multiconfigurational perturbation theory: applications in electronic spectroscopy. *Adv. Chem. Phys.* **93**, 219–331 (1996).
58. Celani, P. & Werner, H.-J. Multireference perturbation theory for large restricted and selected active space reference wave functions. *J. Chem. Phys.* **112**, 5546–5557 (2000).
59. Angeli, C., Cimraglia, R., Evangelisti, S., Leininger, T. & Malrieu, J. P. Introduction of n -electron valence states for multireference perturbation theory. *J. Chem. Phys.* **114**, 10252–10264 (2001).
60. Ma, D., Li Manni, G., Olsen, J. & Gagliardi, L. Second-order perturbation theory for generalized active space self-consistent-field wave functions. *J. Chem. Theory Comput.* **12**, 3208–3213 (2016).
61. Li Manni, G. et al. Multiconfiguration pair-density functional theory. *J. Chem. Theory Comput.* **10**, 3669–3680 (2014).
62. Bender, C. F. & Davidson, E. R. Studies in configuration interaction: the first-row diatomic hydrides. *Phys. Rev.* **183**, 23–30 (1969).
63. Whitten, J. L. & Hackmeyer, M. Configuration interaction studies of ground and excited states of polyatomic molecules. I. The CI formulation and studies of formaldehyde. *J. Chem. Phys.* **51**, 5584–5596 (1969).
64. Huron, B., Malrieu, J. P. & Rancurel, P. Iterative perturbation calculations of ground and excited state energies from multiconfigurational zeroth-order wavefunctions. *J. Chem. Phys.* **58**, 5745–5759 (1973).
65. Evangelisti, S., Daudey, J.-P. & Malrieu, J.-P. Convergence of an improved CIPSI algorithm. *Chem. Phys.* **75**, 91–102 (1983).
66. Bytautas, L. & Ruedenberg, K. A priori identification of configurational deadwood. *Chem. Phys.* **356**, 64–75 (2009).
67. Holmes, A. A., Tubman, N. M. & Umrigar, C. J. Heat-bath configuration interaction: an efficient selected configuration interaction algorithm inspired by heat-bath sampling. *J. Chem. Theory Comput.* **12**, 3674–3680 (2016).
68. Schriber, J. B. & Evangelista, F. A. Communication: an adaptive configuration interaction approach for strongly correlated electrons with tunable accuracy. *J. Chem. Phys.* **144**, 161106 (2016).
69. Tubman, N. M., Lee, J., Takeshita, T. Y., Head-Gordon, M. & Whaley, K. B. A deterministic alternative to the full configuration

- interaction quantum Monte Carlo method. *J. Chem. Phys.* **145**, 044112 (2016).
70. Liu, W. & Hoffmann, M. R. iCI: iterative CI toward full CI. *J. Chem. Theory Comput.* **12**, 1169–1178 (2016).
71. Sharma, S., Holmes, A. A., Jeanmairet, G., Alavi, A. & Umrigar, C. J. Semistochastic heat-bath configuration interaction method: selected configuration interaction with semistochastic perturbation theory. *J. Chem. Theory Comput.* **13**, 1595–1604 (2017).
72. Holmes, A. A., Umrigar, C. J. & Sharma, S. Excited states using semistochastic heat-bath configuration interaction. *J. Chem. Phys.* **147**, 164111 (2017).
73. Anderson, J. S., Heidar-Zadeh, F. & Ayers, P. W. Breaking the curse of dimension for the electronic Schrödinger equation with functional analysis. *Comput. Theor. Chem.* **1142**, 66–77 (2018).
74. Li, J., Otten, M., Holmes, A. A., Sharma, S. & Umrigar, C. J. Fast semistochastic heat-bath configuration interaction. *J. Chem. Phys.* **149**, 214110 (2018).
75. Greene, S. M., Webber, R. J., Weare, J. & Berkelbach, T. C. Beyond walkers in stochastic quantum chemistry: reducing error using fast randomized iteration. *J. Chem. Theory Comput.* **15**, 4834–4850 (2019).
76. Zhang, N., Liu, W. & Hoffmann, M. R. Iterative configuration interaction with selection. *J. Chem. Theory Comput.* **16**, 2296–2316 (2020).
77. Goings, J. J., Hu, H., Yang, C. & Li, X. Reinforcement learning configuration interaction. *J. Chem. Theory Comput.* **17**, 5482–5491 (2021).
78. Zhang, N., Liu, W. & Hoffmann, M. R. Further development of iCIPT2 for strongly correlated electrons. *J. Chem. Theory Comput.* **17**, 949–964 (2021).
79. Siegbahn, P. E. A new direct CI method for large CI expansions in a small orbital space. *Chem. Phys. Lett.* **109**, 417–423 (1984).
80. Olsen, J., Roos, B. O. & Jørgensen, P. Determinant based configuration interaction algorithms for complete and restricted configuration interaction spaces. *J. Chem. Phys.* **89**, 2185–2192 (1988).
81. Zarrabian, S., Sarma, C. & Paldus, J. Vectorizable approach to molecular CI problems using determinantal basis. *Chem. Phys. Lett.* **155**, 183–188 (1989).
82. Dyall, K. G. Interfacing relativistic and nonrelativistic methods. I. Normalized elimination of the small component in the modified Dirac equation. *J. Chem. Phys.* **106**, 9618–9626 (1997).
83. Dyall, K. G. Interfacing relativistic and nonrelativistic methods. II. Investigation of a low-order approximation. *J. Chem. Phys.* **109**, 4201–4208 (1998).
84. Kutzlenigg, W. & Liu, W. Quasirelativistic theory equivalent to fully relativistic theory. *J. Chem. Phys.* **123**, 241102 (2005).
85. Liu, W. & Peng, D. Infinite-order quasirelativistic density functional method based on the exact matrix quasirelativistic theory. *J. Chem. Phys.* **125**, 044102 (2006).
86. Peng, D., Liu, W., Xiao, Y. & Cheng, L. Making four- and two-component relativistic density functional methods fully equivalent based on the idea of from atoms to molecule. *J. Chem. Phys.* **127**, 104106 (2007).
87. Ilias, M. & Saue, T. An infinite-order relativistic hamiltonian by a simple one-step transformation. *J. Chem. Phys.* **126**, 064102 (2007).
88. Liu, W. & Peng, D. Exact two-component hamiltonians revisited. *J. Chem. Phys.* **131**, 031104 (2009).
89. Peng, D., Middendorff, N., Weigend, F. & Reiher, M. An efficient implementation of two-component relativistic exact-decoupling methods for large molecules. *J. Chem. Phys.* **138**, 184105 (2013).
90. Konecny, L. et al. Acceleration of relativistic electron dynamics by means of X2C transformation: application to the calculation of nonlinear optical properties. *J. Chem. Theory Comput.* **12**, 5823–5833 (2016).
91. Egidi, F. et al. Two-component non-collinear time-dependent spin density functional theory for excited state calculations. *J. Chem. Theory Comput.* **13**, 2591–2603 (2017).
92. Liu, J. & Cheng, L. Relativistic coupled-cluster and equation-of-motion coupled-cluster methods. *WIREs Comput. Mol. Sci.* **11**, 1536 (2021).
93. Hoyer, C. E., Hu, H., Lu, L., Knecht, S. & Li, X. Relativistic Kramers-unrestricted exact-two-component density matrix renormalization group. *J. Phys. Chem. A* **126**, 5011–5020 (2022).
94. Davis, C. & Kahan, W. The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.* **7**, 1–46 (1970).
95. Parlett, B. N. *The Symmetric Eigenvalue Problem* (Society for Industrial and Applied Mathematics, 1998).
96. Pollak, P. & Weigend, F. Segmented contracted error-consistent basis sets of double- and triple- ζ valence quality for one- and two-component relativistic all-electron calculations. *J. Chem. Theory Comput.* **13**, 3696–3705 (2017).
97. Knowles, P. J. Very large full configuration interaction calculations. *Chem. Phys. Lett.* **155**, 513–517 (1989).
98. Knowles, P. J. & Handy, N. C. A determinant based full configuration interaction program. *Comp. Phys. Comm.* **54**, 75–83 (1989).
99. Mitrushenkov, A. O. Passing the several billions limit in FCI calculations on a mini-computer. *Chem. Phys. Lett.* **217**, 559–565 (1994).
100. Zhou, Y., Shepard, R. & Minkoff, M. Computing eigenvalue bounds for iterative subspace matrix methods. *Comp. Phys. Comm.* **167**, 90–102 (2005).
101. Rolik, Z., Szabados, A. & Surján, P. R. A sparse matrix based full-configuration interaction algorithm. *J. Chem. Phys.* **128**, 144101 (2008).
102. Ritz, W. Über eine neue Methode zur Lösung gewisser Variationsprobleme der mathematischen Physik. *J. Reine Angew. Math.* **135**, 1–61 (1909).
103. Cotton, S. J. A truncated Davidson method for the efficient “chemically accurate” calculation of full configuration interaction wavefunctions without any large matrix diagonalization. *J. Chem. Phys.* **157**, 224105 (2022).
104. Davidson, E. R. The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices. *J. Chem. Phys.* **17**, 87–94 (1975).
105. Knyazev, A. V. Toward the optimal preconditioned eigensolver: locally optimal block preconditioned conjugate gradient method. *SIAM J. Sci. Comp.* **23**, 517–541 (2001).
106. Sleijpen, G. L. & Van der Vorst, H. A. A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.* **17**, 401–425 (1996).
107. Vecharynski, E., Yang, C. & Xue, F. Generalized preconditioned locally harmonic residual method for non-Hermitian eigenproblems. *SIAM J. Sci. Comp.* **38**, A500–A527 (2016).
108. Kasper, J. M., Williams-Young, D. B., Vecharynski, E., Yang, C. & Li, X. A well-tempered hybrid method for solving challenging time-dependent density functional theory (TDDFT) systems. *J. Chem. Theory Comput.* **14**, 2034–2041 (2018).
109. Liang, W., Fischer, S. A., Frisch, M. J. & Li, X. Energy-specific linear response TDHF/TDDFT for calculating high-energy excited states. *J. Chem. Theory Comput.* **7**, 3540–3547 (2011).
110. Peng, B., Lestrang, P. J., Goings, J. J., Caricato, M. & Li, X. Energy-specific equation-of-motion coupled-cluster methods for high-energy excited states: application to K-edge X-ray absorption spectroscopy. *J. Chem. Theory Comput.* **11**, 4146–4153 (2015).
111. Zhou, Z. & Parker, S. M. Converging time-dependent density functional theory calculations in five iterations with minimal auxiliary preconditioning. *J. Chem. Theory Comput.* **20**, 6738–6746 (2024).

112. Hernandez, T. M., Van Beeumen, R., Caprio, M. A. & Yang, C. A greedy algorithm for computing eigenvalues of a symmetric matrix with localized eigenvectors. *Numer. Linear Algebra Appl.* **28**, e2341 (2021).
113. Peterson, K. A., Woon, D. E. & Dunning Jr, T. H. Benchmark calculations with correlated molecular wave functions. IV. The classical barrier height of the $\text{H} + \text{H}_2 \rightarrow \text{H}_2 + \text{H}$ reaction. *J. Chem. Phys.* **100**, 7410–7415 (1994).
114. Silver, D. M. & Stevens, R. M. Reaction paths on the H_4 potential energy surface. *J. Chem. Phys.* **59**, 3378–3394 (1973).
115. Hernández, M. I. & Clary, D. C. Four-center reactions: a quantal model for H_4 . *J. Chem. Phys.* **104**, 8413–8423 (1996).
116. Contant, D., Casula, M. & Hellgren, M. Assessing many-body methods on the potential energy surface of the $(\text{H}_2)_2$ hydrogen dimer. *J. Chem. Phys.* **161**, 074106 (2024).
117. Schätzle, Z., Hermann, J. & Noé, F. Convergence to the fixed-node limit in deep variational Monte Carlo. *J. Chem. Phys.* **154**, 124108 (2021).
118. Nakashima, H. & Nakatsuji, H. Solving the Schrödinger equation of a planar model H_4 molecule. *Chem. Phys. Lett.* **815**, 140359 (2023).
119. Nakano, M. et al. Full configuration interaction calculations of the second hyperpolarizabilities of the H_4 model compound: summation-over-states analysis and interplay with diradical characters. *J. Chem. Phys.* **136**, 024315 (2012).
120. Kedžuch, S., Demel, O., Pittner, J., Ten-no, S. & Noga, J. Multi-reference F12 coupled cluster theory: the Brillouin-Wigner approach with single and double excitations. *Chem. Phys. Lett.* **511**, 418–423 (2011).
121. Senjean, B., Yalouz, S., Nakatani, N. & Fromager, E. Reduced density matrix functional theory from an ab initio seniority-zero wave function: exact and approximate formulations along adiabatic connection paths. *Phys. Rev. A* **106**, 032203 (2022).
122. Scuseria, G. E., Jiménez-Hoyos, C. A., Henderson, T. M., Samanta, K. & Ellis, J. K. Projected quasiparticle theory for molecular electronic structure. *J. Chem. Phys.* **135**, 124108 (2011).
123. Van Voorhis, T. & Head-Gordon, M. Benchmark variational coupled cluster doubles results. *J. Chem. Phys.* **113**, 8873–8879 (2000).
124. Genovese, C., Meninno, A. & Sorella, S. Assessing the accuracy of the Jastrow antisymmetrized geminal power in the H_4 model system. *J. Chem. Phys.* **150**, 084102 (2019).
125. Baek, U. et al. Say NO to optimization: a nonorthogonal quantum eigensolver. *Phys. Rev. X Quantum* **4**, 030307 (2023).
126. Sand, A. M. & Mazziotti, D. A. Parametric two-electron reduced-density-matrix method with application to diradical rectangular H_4 . *Comput. Theor. Chem* **1003**, 44–49 (2013).
127. Kullie, O. & Saue, T. Range-separated density functional theory: a 4-component relativistic study of the rare gas dimers He_2 , Ne_2 , Ar_2 , Kr_2 , Xe_2 , Rn_2 and Uuo_2 . *Chem. Phys.* **395**, 54–62 (2012).
128. Burda, J. V., Zahradnik, R. & Hobza, P. Dimers of rare gas atoms: CCSD(T), CCSDT and FCI calculations on the $(\text{He})_2$ dimer, CCSD(T) and CCSDT calculations on the $(\text{Ne})_2$ dimer, and CCSD(T) all-electron and pseudopotential calculations on the dimers from $(\text{Ne})_2$ through $(\text{Xe})_2$. *Mol. Phys.* **89**, 425–432 (1996).
129. Zhang, Y., Pan, W. & Yang, W. Describing van der Waals Interaction in diatomic molecules with generalized gradient approximations: the role of the exchange functional. *J. Chem. Phys.* **107**, 7921–7925 (1997).
130. Laschuk, E. F., Martins, M. M. & Evangelisti, S. Ab initio potentials for weakly interacting systems: homonuclear rare gas dimers. *Int. J. Quant. Chem.* **95**, 303–312 (2003).
131. Slaviček, P. et al. State-of-the-art correlated ab initio potential energy curves for heavy rare gas dimers: Ar_2 , Kr_2 , and Xe_2 . *J. Chem. Phys.* **119**, 2102–2119 (2003).
132. Tao, J. & Perdew, J. P. Test of a nonempirical density functional: short-range part of the van der Waals interaction in rare-gas dimers. *J. Chem. Phys.* **122**, 114102 (2005).
133. Goll, E., Werner, H.-J. & Stoll, H. A short-range gradient-corrected density functional in long-range coupled-cluster calculations for rare gas dimers. *Phys. Chem. Chem. Phys.* **7**, 3917–3923 (2005).
134. Kannemann, F. O. & Becke, A. D. Van der Waals interactions in density-functional theory: rare-gas diatomics. *J. Chem. Theory Comput.* **5**, 719–727 (2009).
135. Zhu, W., Toulouse, J., Savin, A. & Ángyán, J. G. Range-separated density-functional theory with random phase approximation applied to noncovalent intermolecular interactions. *J. Chem. Phys.* **132**, 244108 (2010).
136. Zhu, P., Argentati, M. E. & Knyazev, A. V. Bounds for the Rayleigh quotient and the spectrum of self-adjoint operators. *SIAM J. Matrix Anal. Appl.* **34**, 244–256 (2013).
137. Yosida, K. *Functional Analysis* (Springer, 1995).
138. Saad, Y. *Numerical Methods for Large Eigenvalue Problems* (Soc. for Industr. and Appl. Math., 2011).
139. Williams-Young, D. B. et al. The Chronus Quantum (ChronusQ) software package. *WIREs Comput. Mol. Sci.* **10**, e1436 (2020).
140. Shayit, A. et al. Numerically exact configuration interaction at quadrillion-determinant scale. Chronus quantum preview. Figshare. Software. <https://doi.org/10.6084/m9.figshare.30411583.v2> (2025).
141. Hill, J. G. & Peterson, K. A. Gaussian basis sets for use in correlated molecular calculations. XI. Pseudopotential-based and all-electron relativistic basis sets for alkali metal (K–Fr) and alkaline earth (Ca–Ra) elements. *J. Chem. Phys.* **147**, 244106 (2017).
142. de Jong, W. A., Harrison, R. J. & Dixon, D. A. Parallel Douglas–Kroll energy and gradients in NWChem: estimating scalar relativistic effects using Douglas–Kroll contracted basis set. *J. Chem. Phys.* **114**, 48–53 (2001).
143. Prascher, B. P., Woon, D. E., Peterson, K. A., Dunning Jr, T. H. & Wilson, A. K. Gaussian basis sets for use in correlated molecular calculations. VII. Valence, core-valence, and scalar relativistic basis sets for Li, Be, Na, and Mg. *Theor. Chem. Acc.* **128**, 69–82 (2010).

Acknowledgements

The development of variational relativistic multi-reference methods is supported by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, in the Computational and Theoretical Chemistry program (Grant No. DE-SC0006863 to X.L.). The development of the Chronus Quantum computational software is supported by the Office of Advanced Cyberinfrastructure, National Science Foundation (Grants No. OAC-2103717 and OAC-2103705). X.L., C.Y., and A.E.D. acknowledge the support to develop reduced scaling computational methods from the Scientific Discovery through Advanced Computing (SciDAC) program sponsored by the Offices of Advanced Scientific Computing Research (ASCR) and Basic Energy Sciences (BES) of the U.S. Department of Energy (Grant No. DE-SC0022263). This project used the resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC awards CSGB-ERCAP 0032454 and DDR-ERCAP 0034806.

Author contributions

A. Shayit, S. Updhyay, H. Hu, and X. Li conceived the project and developed the computer code. A.E. DePrince and A. Shayit acquired computing resources. C. Liao, A. Shayit, and S. Updhyay carried out the benchmark calculations and analyzed the data. X. Li and T. Zhang developed the software infrastructure. A.E. DePrince, C. Yang, and X. Li acquired research funding. X. Li, A. Shayit, S. Updhyay, and C. Liao wrote

the manuscript with input from all authors. All authors discussed the results and approved the final version of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-025-65967-7>.

Correspondence and requests for materials should be addressed to Xiaosong Li.

Peer review information *Nature Communications* thanks Laura Gagliardi, Rahul Maitra and the other anonymous, reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025