

1 **A SIMPLE GPU-ACCELERATED SOLVER FOR THE**
2 **SCHRÖDINGER OPERATOR WITH APPLICATIONS TO GROUND**
3 **STATES AND HAMILTONIAN SIMULATION**

4 XINYU LIU* AND XIANGXIONG ZHANG†

5 **Abstract.** We extend the tensor-product direct solver from the Laplacian to the Schrödinger
6 operator $-\Delta + V$. When the potential V_1 is separable, the operator $-\Delta + V_1$ is inverted or exponenti-
7 ated at cost $O(N^{1+1/d})$ in d dimensions via per-axis eigendecomposition. On a single NVIDIA A100
8 GPU, this costs less than one second for 10^9 degrees of freedom in 3D. For non-separable potentials
9 $V = V_1 + V_2$, the same solver provides a preconditioner $(-\Delta + V_1)^{-1}$ for the preconditioned conju-
10 gate gradient (PCG) method and a propagator for operator-splitting time integrators. For bounded
11 V_2 , we prove that the preconditioned operator has a bounded condition number and a clustered
12 spectrum with at most finitely many outlier eigenvalues, independently of the mesh size, and also
13 independently of the domain size when V_1 is a confining potential. This explains the mesh- and
14 domain-independent PCG iteration counts observed in practice. We apply this method to ground
15 state computation via inverse iteration for linear problems and via the a_u gradient flow for Gross-
16 Pitaevskii energy in 3D, and also Hamiltonian simulation via the approximated qHOP and Magnus-2
17 splitting methods from 3D to 9D on a single NVIDIA GH200 GPU.

18 **Key words.** tensor-product spectral method, GPU computing, Schrödinger equation, Gross-
19 Pitaevskii equation, Strang splitting, Yoshida splitting

20 **AMS subject classifications.** 65M70, 65N35, 65F15, 81-08, 35Q55, 81Q05

21 **1. Introduction.** It is well known that the discrete Laplacian on Cartesian
22 meshes has a tensor-product structure, which can be used as a simple and fast direct
23 solver since the 1960s [29, 18, 33, 25, 28]. For a general potential V , this structure is
24 lost for the Schrödinger operator $-\Delta + V$. In this paper, we extend the tensor-product
25 solver from the Laplacian to the Schrödinger operator $-\Delta + V$ as follows. For a sep-
26 arable potential $V_1(\mathbf{x}) = \sum_k f_k(x_k)$, the operator $-\Delta + V_1$ can still be inverted or
27 exponentiated via per-axis eigendecomposition at cost $O(N^{1+1/d})$ with $O(N)$ mem-
28 ory. For non-separable potentials $V = V_1 + V_2$, the tensor-product solver provides a
29 preconditioner $(-\Delta + V_1)^{-1}$ for the preconditioned conjugate gradient (PCG) method
30 and a propagator for operator-splitting time integrators. For bounded V_2 , we prove
31 that the preconditioned operator $(-\Delta + V_1)^{-1}(-\Delta + V_1 + V_2)$ has a bounded condition
32 number and a clustered spectrum with at most finitely many outliers, independently
33 of the mesh size (Theorem 3.4), and also independently of the domain size when V_1 is
34 a confining potential (Theorem 3.5). This can explain mesh- and domain-independent
35 PCG iteration counts observed in numerical tests for suitable potentials. Similar to
36 the simple GPU acceleration of the direct inversion of the Laplacian in [28], we show
37 that the Schrödinger operator $-\Delta + V$ can be easily inverted on GPU (Sections 2–3),
38 e.g., it takes less than 1 second for inverting $-\Delta + V_1$ for one billion DoFs on one
39 Nvidia GPU such as A100 and GH200.

40 Fast inversion of the Schrödinger operator has several applications. We first
41 consider ground state computation (Section 4): for separable potentials via shifted
42 inverse iteration, and for non-separable potentials via shifted inverse iteration with
43 PCG, preconditioned by $(-\Delta + V_1)^{-1}$. We also demonstrate that the defocusing
44 Gross–Pitaevskii ground state via the a_u flow can be easily implemented in 3D on a

*Department of Mathematics, The Ohio State University, 231 West 18th Avenue, Columbus, OH 43210 (liu.12165@osu.edu).

†Corresponding author. Department of Mathematics, Purdue University, 150 N. University Street, West Lafayette, IN 47907 (zhan1966@purdue.edu).

45 single GPU.

46 Another application is Hamiltonian simulation. If (λ_1, u_1) is the ground state
 47 eigenpair of $H = -\Delta + V$, then $\psi(x, t) = e^{-i\lambda_1 t} u_1(x)$ is an exact solution of the
 48 time-dependent Schrödinger equation $i\partial_t \psi = H\psi$. We use this stationary solution as a
 49 reference for testing operator-splitting methods (Sections 5–6).

50 In the quantum computing literature, two recent splitting methods have been
 51 considered for simulating $i\partial_t \psi = (A+B)\psi$ in the interaction picture, where A has large
 52 spectral norm but can be fast-forwarded (e.g., $A = -\Delta$ via QFT). The *quantum Highly*
 53 *Oscillatory Protocol* (qHOP) [1] approximates the first-order Magnus expansion and
 54 achieves $O(\Delta t^2)$ superconvergence with an error preconstant independent of problem
 55 size n [1]. The *Magnus-2* algorithm [12] uses the second-order Magnus expansion
 56 and achieves $O(\Delta t^4)$ superconvergence [12, 7]. On a quantum computer, the matrix
 57 exponential of a sum of interaction-picture Hamiltonians is implemented via linear
 58 combination of unitaries (LCU). Classically, this cannot be computed directly at large
 59 n . We implement an approximate version using a product formula where each factor
 60 is computed exactly via tensor-product propagations and a pointwise multiplication
 61 (equation (5.2) and Remark 5.1 in Section 5.2). We test these methods for multi-body
 62 Coulomb Hamiltonians in 4D, 6D, and 9D (Section 6). The approximated versions of
 63 qHOP and Magnus-2 are tested in multiple dimensions for two purposes. First, these
 64 tests can serve as a verification of their theoretical superconvergence orders in high
 65 dimensions. Since the approximated versions introduce additional product-formula
 66 errors beyond the original splitting error, observing the predicted convergence orders
 67 despite these additional errors provides supporting numerical evidence for the original
 68 schemes. Second, the same product-formula implementation yields a family of classical
 69 splitting methods with parameter M : $M = 1$ recovers the classical Strang and Yoshida
 70 splittings, while $M \geq 3$ reduces the error constant at the same convergence order,
 71 demonstrating that qHOP and Magnus-2 can serve as improved classical integrators
 72 compared to standard schemes. Furthermore, as a classical numerical scheme, the
 73 splitting $A = -\Delta + V_1$, $B = V_2$ gives smaller errors than $A = -\Delta$, $B = V_1 + V_2$
 74 (Tables 7 and 8 in Section 5).

75 In related work, the tensor-product approach to Schrödinger equations has been
 76 explored by Caliri et al. [8], whose μ -mode integrator exploits Kronecker structure
 77 for 3D problems at $\sim 2 \times 10^6$ DoFs. GPU acceleration for Schrödinger simulation has
 78 also been explored in [16]. The split-step spectral method was introduced in [5]. We
 79 refer to [26, 24] for reviews of semiclassical methods and to [23, 36, 21] for splitting
 80 error analysis and classical Magnus integrators. On the quantum computing side, gate
 81 complexity for real-space simulation of Schrödinger equations was established in [10],
 82 Trotter error bounds for interacting electrons were given in [35, 17], and quantum
 83 resource estimates for 2D Coulomb simulation are given in [15].

84 All results in this paper can be reproduced via Python (JAX) code available
 85 at <https://github.com/zhan1966/schrodinger>. GPU cards are available to many re-
 86 searchers through programs such as the NSF-funded ACCESS (Advanced Cyberinfras-
 87 tructure Coordination Ecosystem: Services & Support) [6], which provides allocations
 88 on GPU clusters including NVIDIA A100 and GH200 nodes. The direct inversion of
 89 the tensor-structured Laplacian on A100 was demonstrated in [28, 9, 19]. This paper
 90 extends the same approach to the Schrödinger operator $-\Delta + V$. All computations
 91 in this paper require only a single GPU.

92 The rest of the paper is organized as follows. Section 2 describes the tensor-
 93 product spectral solver. Section 3 develops the PCG preconditioner and analyzes its
 94 mesh- and domain-independent condition number and spectral clustering. Section 4

95 presents ground state computation via shifted inverse iteration and Gross–Pitaevskii
 96 gradient flows. Sections 5 and 6 present Hamiltonian simulation results in 3D through
 97 9D. Some concluding remarks are given in Section 7. Appendix A presents the Hermite
 98 spectral discretization on unbounded domains.

99 **2. Direct solver for $-\Delta + V$ with a separable potential.** The idea in this
 100 paper applies to any numerical method with a tensor-product structure on rectangular
 101 domains. For simplicity, throughout the paper we focus on the Q^k spectral element
 102 method (SEM), which is the high order finite element method with Q^k basis and
 103 $(k + 1)$ -point Gauss-Lobatto quadrature. For $k = 1$, it is the classical second order
 104 finite difference method. For $k \geq 2$, Q^k SEM can be proven to be a $(k + 2)$ -th order
 105 finite difference scheme in ℓ^2 -norm over quadrature point values [27]. As an example
 106 of another suitable method, we also include a Hermite spectral method on unbounded
 107 domains in Appendix A. See also [28] for using Q^k SEM in a Poisson solver.

108 **2.1. A direct solver via eigendecomposition for separable potentials.**
 109 We consider the elliptic PDE

$$110 \quad (2.1) \quad -\Delta u(\mathbf{x}) + V(\mathbf{x})u(\mathbf{x}) = b(\mathbf{x}), \quad \mathbf{x} \in \Omega,$$

111 on a rectangular domain $\Omega = [-L, L]^d$ (or on \mathbb{R}^d via Hermite functions in Appen-
 112 dix A). For simplicity, we consider the homogeneous Dirichlet boundary condition,
 113 and extensions of the scheme to homogeneous Neumann and periodic boundary con-
 114 ditions are straightforward. With the Q^k spectral element method using Gauss-
 115 Lobatto quadrature [27], let $K_d = M_d^{-1}S_d$ denote the 1D discrete Laplacian in di-
 116 rection $d \in \{x, y, z\}$, with M_d the mass matrix, S_d the stiffness matrix, and I_d the
 117 identity matrix. For Q^1 elements on a uniform mesh, K_d is the well-known tridiagonal
 118 $(-1, 2, -1)$ matrix from second order finite difference. Let n be the number of interior
 119 grid points in each direction. For any d -dimensional array X of size $n \times \dots \times n$, let
 120 $\text{vec}(X)$ denote the vector of size n^d obtained by reshaping all entries of X into a
 121 column vector. Then in 3D, the discretized system takes the form (see [28, 9]):

$$122 \quad (2.2) \quad [I_z \otimes I_y \otimes K_x + I_z \otimes K_y \otimes I_x + K_z \otimes I_y \otimes I_x + \text{diag}(\text{vec}(V))] \text{vec}(U) = \text{vec}(B),$$

123 where V is the array of potential values $V(\mathbf{x}_i)$ at the grid points. When the po-
 124 tential is separable, $V(\mathbf{x}) = f_x(x) + f_y(y) + f_z(z)$, the potential matrix decom-
 125 poses as $\text{diag}(\text{vec}(V)) = I_z \otimes I_y \otimes F_x + I_z \otimes F_y \otimes I_x + F_z \otimes I_y \otimes I_x$ with $F_d =$
 126 $\text{diag}(f_d(x_1), \dots, f_d(x_n))$, and the system (2.2) reduces to the Kronecker form

$$127 \quad (2.3) \quad (I_z \otimes I_y \otimes H_x + I_z \otimes H_y \otimes I_x + H_z \otimes I_y \otimes I_x) \text{vec}(U) = \text{vec}(B),$$

128 where $H_d = K_d + F_d = M_d^{-1}S_d + F_d$ for each axis $d \in \{x, y, z\}$. Each H_d can be
 129 diagonalized by

$$130 \quad (2.4) \quad H_d = M_d^{-1/2}(M_d^{-1/2}S_dM_d^{-1/2} + F_d)M_d^{1/2} = T_d\Lambda_dT_d^{-1},$$

131 where $T_d = M_d^{-1/2}Q_d$, $T_d^{-1} = Q_d^T M_d^{1/2}$, and $Q_d\Lambda_dQ_d^T$ is the eigendecomposition of
 132 the symmetric matrix $M_d^{-1/2}S_dM_d^{-1/2} + F_d$.

133 The solution is then computed in four steps:

- 134 1. Off-line setup (one-time): compute the eigendecomposition $H_d = T_d\Lambda_dT_d^{-1}$
 135 for each axis d . Since H_d is an $n \times n$ matrix, this costs $O(n^3) = O(N)$.
- 136 2. Forward transform: apply $T_z^{-1} \otimes T_y^{-1} \otimes T_x^{-1}$ to B .

137 3. Solve the diagonal system: divide by $(\Lambda_x)_{ii} + (\Lambda_y)_{jj} + (\Lambda_z)_{kk}$.

138 4. Backward transform: apply $T_z \otimes T_y \otimes T_x$.

139 Steps 2 and 4 contain the dominant computational cost $\mathcal{O}(N^{4/3})$ for $N = n^3$. The
140 memory cost is $\mathcal{O}(N)$.

141 The same factorization applies to the time-dependent Schrödinger propagator
142 $e^{-i(-\Delta+V_1)\Delta t}$. Since $-\Delta+V_1 = T \Lambda T^{-1}$, where $T = T_z \otimes T_y \otimes T_x$ and Λ is the diagonal
143 tensor of eigenvalue sums $\Lambda_{ijk} = (\Lambda_x)_{ii} + (\Lambda_y)_{jj} + (\Lambda_z)_{kk}$, the matrix exponential is

$$144 \quad e^{-i(-\Delta+V_1)\Delta t} \psi = T (e^{-i\Lambda\Delta t} \odot (T^{-1} \psi)),$$

145 with \odot denoting pointwise multiplication. The only change from the direct inversion
146 is that the eigenvalue division in step 3 is replaced by multiplication by $e^{-i\Lambda_{ijk}\Delta t}$.
147 Steps 1, 2 and 4 are the same. Thus, once the eigendecomposition is computed, both
148 the direct solve $(-\Delta + V_1)^{-1}b$ and the propagator $e^{-i(-\Delta+V_1)\Delta t}\psi$ cost $\mathcal{O}(N^{1+1/d})$.

149 **2.2. Numerical tests.** We demonstrate the solver on the equation

$$150 \quad (2.5) \quad (-\Delta + V) u = f, \quad \mathbf{x} \in [-1, 1]^d.$$

151 For 3D, we use the non-isotropic separable potential

$$152 \quad (2.6) \quad V(\mathbf{x}) = 1600 \left(\sin^2\left(\frac{\pi}{4}x\right) + \sin^2\left(\frac{\pi}{4}y\right) + \sin^2\left(\frac{\pi}{4}z\right) \right) + (x^2 + 2y^2 + 3z^2),$$

153 exact solution $u^* = \sin(\pi x) \sin(2\pi y) \sin(3\pi z)$, and $f = (-\Delta + V) u^*$. For 4D and
154 6D tests, we use the harmonic trapping potential $V(\mathbf{x}) = \sum_{k=1}^d x_k^2$, and the exact
155 solution $u^* = \prod_{k=1}^d \sin(k\pi x_k)$.

156 All computations are implemented in Python with JAX, similar to the Python
157 implementation in [28]. Table 1 reports *setup* time for offline eigendecomposition and
158 eigenvalue array construction, and *solve* time for applying the direct solver, which
159 would be the per-iteration cost in an iterative method such as PCG in later sections.

160 In Table 1, we also report different precision formats on a GH200 GPU (96 GB
161 HBM3). FP32 (single precision) stores each number in 32 bits, using half the memory
162 of FP64. TF32 (TensorFloat-32) uses 32-bit storage but performs matrix multiplica-
163 tions on GPU tensor cores at reduced precision for fast matrix operations. See [28] for
164 more comparison of TF32 with FP32. BF16 (bfloat16) uses only 16 bits per number,
165 further reducing memory and enabling larger grids on the same device, but with only
166 ~ 3 -digit precision.

167 **3. Efficient Inversion of $-\Delta + V$.** When $V = V_1 + V_2$ contains a separable V_1
168 and a non-separable V_2 , we use PCG with the preconditioner $(-\Delta + V_1)^{-1}$.

169 **3.1. Test potentials.** We consider two physically motivated potentials.

170 (a) *Harmonic plus quartic potential* [3]:

$$171 \quad (3.1) \quad V(x, y, z) = 2(1-\alpha)(\gamma_x x^2 + \gamma_y y^2) + \frac{\kappa}{2}(x^2 + y^2)^2 + \gamma_z z^2.$$

172 We consider $\gamma_x = \gamma_y = 1$, $\gamma_z = 3$, $\alpha = 1.4$, $\kappa = 0.3$, and we split $V = V_1 + V_2$ with
173 separable part $V_1 = \gamma_x x^2 + \gamma_y y^2 + \gamma_z z^2$ and non-separable remainder

$$174 \quad V_2 = (2(1-\alpha)-1)(\gamma_x x^2 + \gamma_y y^2) + \frac{\kappa}{2}(x^2 + y^2)^2 = -1.8(x^2 + y^2) + 0.15(x^2 + y^2)^2.$$

175 The quartic coupling $(x^2 + y^2)^2$ makes V_2 genuinely non-separable.

TABLE 1

Direct inversion of $(-\Delta + V)$ via Q^k SEM on Nvidia GH200 96 GB. 3D: potential (2.6), 4D/6D: $V = \sum x_k^2$. At matched n , solve times (in seconds) are identical across Q^k . In FP64, all three dimensions reach $\sim 10^9$ DoFs: $1099^3 \approx 1.3 \times 10^9$, $179^4 \approx 1.0 \times 10^9$, $33^6 \approx 1.3 \times 10^9$.

Precision	Dim	Q^k	N_{cell}	n	DoFs	Setup (s)	Solve (s)	ℓ^2 rel. err
FP64	3D	Q^1	1100	1099	1099^3	5.2	0.508	3.66×10^{-6}
		Q^2	550	1099	1099^3	5.1	0.512	1.93×10^{-9}
		Q^{10}	110	1099	1099^3	6.4	0.510	$7.66 \times 10^{-13\dagger}$
	4D	Q^{10}	18	179	179^4	6.2	0.195	$1.29 \times 10^{-13\dagger}$
	6D	Q^2	17	33	33^6	5.9	0.258	$2.95 \times 10^{-2§}$
FP32	3D	Q^2	500	999	999^3	5.4	0.270	7.17×10^{-7}
		Q^{10}	15	149	149^4	5.8	0.058	4.99×10^{-7}
		Q^{10}	20	199	199^4	5.8	0.188	$6.20 \times 10^{-7\dagger}$
	6D	Q^{10}	3	29	29^6	5.7	0.052	4.75×10^{-4}
TF32	3D	Q^2	500	999	999^3	6.6	0.133	5.13×10^{-4}
BF16	3D	Q^2	500	999	999^3	5.4	0.088	1.65×10^{-2}
		Q^2	800	1599	1599^3	5.6	0.517	1.64×10^{-2}

\dagger FP64 rounding floor. \ddagger FP32 rounding floor. \S Mesh too coarse for Q^2 to resolve $\prod_{k=1}^6 \sin(k\pi x_k)$.

176 (b) Harmonic plus stirrer potential [2, 22, 4]:

$$177 \quad (3.2) \quad V(x, y, z) = \underbrace{\gamma_x^2 x^2 + \gamma_y^2 y^2 + \gamma_z^2 z^2}_{V_1} + \underbrace{2w_0 e^{-\delta((x-r_0)^2 + y^2)}}_{V_2},$$

178 with $\gamma_x = 1$, $\gamma_y = 1$, $\gamma_z = 2$, $w_0 = 4$, $\delta = 1$, $r_0 = 1$. This potential includes a
179 harmonic trap with a localized Gaussian stirrer.

180 **3.2. PCG for $-\Delta + V_1 + V_2$ (non-separable case).** When $V = V_1 + V_2$
181 contains a non-separable component V_2 , the system $(-\Delta + V_1 + V_2)u = f$ cannot be
182 solved directly by the tensor-product solver. We consider the preconditioned conjugate
183 gradient (PCG) method. Antoine et al. [3] showed that PCG with the combined
184 preconditioner $P_C = V^{-1/2}(-\Delta)^{-1}V^{-1/2}$ achieves grid-independent iteration counts.
185 In this subsection, we show that the preconditioner $P = (-\Delta + V_1)^{-1}$ is stronger. It
186 is applied by the tensor-product direct solver with a cost $O(N^{4/3})$ in 3D.

187 **EXAMPLE 3.1 (PCG on bounded domain).** We consider the quartic (3.1) and
188 stirrer (3.2) potentials on $[-8, 8]^3$ with Q^6 SEM, Dirichlet BC, a random right-hand
189 side f , FP32 arithmetic, on a GH200 GPU. In all FP32 PCG tests reported here, the
190 offline eigendecomposition of $-\Delta + V_1$ is computed in FP64 and only then cast to lower
191 precision for the online PCG iteration. Figure 1 shows the convergence at two grid
192 sizes (197^3 and 701^3 DoFs). Both preconditioners exhibit grid-independent iteration
193 counts. The proposed preconditioner converges in 33–38 iterations (quartic) and 5–
194 6 iterations (stirrer), compared to 214–255 iterations with $(-\Delta)^{-1}$ for the quartic
195 potential and ~ 375 iterations with $P_C = V^{-\frac{1}{2}}(-\Delta)^{-1}V^{-\frac{1}{2}}$ for the stirrer potential.

196 Figure 1 also shows that the performance of the preconditioner $(-\Delta + V_1)^{-1}$ is
197 independent of both the domain size L and the mesh size h for the stirrer potential,
198 while for the quartic potential, the iteration count grows with L .

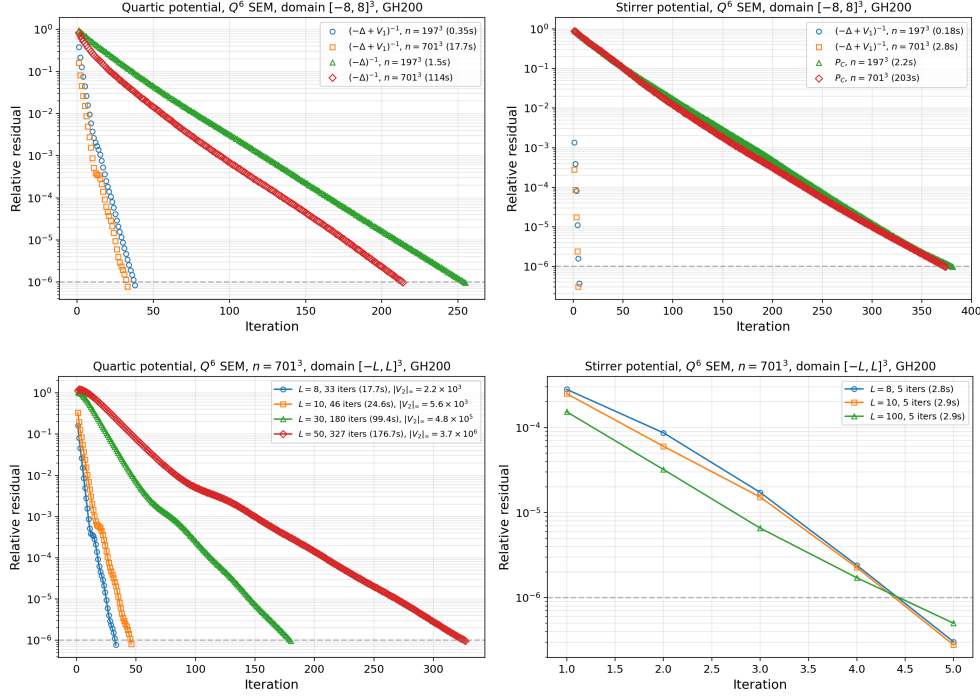


FIG. 1. PCG convergence (relative residual vs. iteration) for $(-\Delta + V_1 + V_2)u = f$, Q^6 SEM, FP32, GH200. Top row: quartic potential (3.1) (left) and stirrer potential (3.2) (right), both at $n = 701^3$, comparing the proposed preconditioner $(-\Delta + V_1)^{-1}$ with the $(-\Delta)^{-1}$ or P_C . Bottom row: domain dependence of $(-\Delta + V_1)^{-1}$ for the quartic potential (left, $L = 8, 10, 30, 50$) and the stirrer potential (right, $L = 8, 10, 100$). Total GPU time for each case is shown in the legend.

199 In the rest of this section, we explain this behavior by analyzing the spectrum of
 200 the preconditioned operator, which involves standard spectral theory for conforming
 201 finite element discretizations (such as Q^k SEM) under two assumptions on V_1 and V_2 .
 202 **Assumption I** (bounded domain): $V_1 \in L^\infty(\Omega)$ with $V_1 \geq 0$, and $V_2 \in L^\infty(\Omega)$.
 203 **Assumption II** (confining potential): $V_1 \geq 0$ with $V_1(x) \rightarrow \infty$ as $|x| \rightarrow \infty$ and
 204 $V_1 \in L^\infty_{\text{loc}}(\mathbb{R}^d)$, and $V_2 \in L^\infty(\mathbb{R}^d)$.
 205 Assumption II is stronger since it implies Assumption I on any bounded domain Ω .

206 **DEFINITION 3.1** (A -inner product). For a bounded domain $\Omega \subset \mathbb{R}^d$, under Assumption I, define the inner product for $H_0^1(\Omega)$:
 207

$$208 \quad \langle u, v \rangle_A = \int_{\Omega} (\nabla u \cdot \nabla v + V_1 u v) \, dx.$$

209 Since $V_1 \geq 0$ and $V_1 \in L^\infty(\Omega)$, the Poincaré inequality gives $(1 + C_P^2)^{-1} \|u\|_{H^1}^2 \leq$
 210 $\|u\|_A^2 \leq \max(1, \|V_1\|_{L^\infty}) \|u\|_{H^1}^2$, so $\|\cdot\|_A$ is equivalent to the standard H^1 norm.

211 **DEFINITION 3.2** (Perturbation operators and Riesz map). Under the setting of
 212 Definition 3.1, let $X_h \subset H_0^1(\Omega)$ be a conforming finite element subspace with mesh
 213 parameter $h > 0$. Define:

214 (a) The Riesz map $R : L^2(\Omega) \rightarrow H_0^1(\Omega)$ by $\langle Rf, v \rangle_A = (f, v)_{L^2}$ for all $v \in H_0^1(\Omega)$.

215 (b) The continuous perturbation operator $K : H_0^1(\Omega) \rightarrow H_0^1(\Omega)$ by

$$216 \quad (3.3) \quad \langle Ku, v \rangle_A = \int_{\Omega} V_2 u v \, d\mathbf{x}, \quad \forall v \in H_0^1(\Omega).$$

217 (c) $K_h : X_h \rightarrow X_h$ by $\langle K_h u_h, v_h \rangle_A = \int_{\Omega} V_2 u_h v_h \, d\mathbf{x}$ for all $v_h \in X_h$.

218 (d) The projection $\Pi_h : H_0^1(\Omega) \rightarrow X_h$ by $\langle \Pi_h u, v_h \rangle_A = \langle u, v_h \rangle_A$ for all $v_h \in X_h$.

219 (e) The Galerkin discretizations $-\Delta_h + V_{1,h}$ and $-\Delta_h + V_{1,h} + V_{2,h}$ on X_h :

$$220 \quad ((-\Delta_h + V_{1,h}) u_h, v_h)_{L^2} := \langle u_h, v_h \rangle_A,$$

221

$$222 \quad ((-\Delta_h + V_{1,h} + V_{2,h}) u_h, v_h)_{L^2} := \langle u_h, v_h \rangle_A + \int_{\Omega} V_2 u_h v_h \, d\mathbf{x},$$

223 for all $u_h, v_h \in X_h$. Here $V_{2,h}$ denotes the restriction of V_2 to X_h via the L^2
224 inner product, i.e., $(V_{2,h} u_h, v_h)_{L^2} = \int_{\Omega} V_2 u_h v_h \, d\mathbf{x}$.

225 The operators in Definition 3.2 correspond directly to the components of the
226 preconditioned system: R represents $(-\Delta + V_1)^{-1}$, K represents $(-\Delta + V_1)^{-1}V_2$, and
227 K_h is the discrete version of K . We now establish their key properties.

228 LEMMA 3.3. Under the setting of Definition 3.2:

229 (a) $K = R \circ (V_2 \cdot) \circ \iota$, where $\iota : H_0^1(\Omega) \hookrightarrow L^2(\Omega)$ is the natural embedding. In
230 particular, K is compact and self-adjoint on $(H_0^1(\Omega), \langle \cdot, \cdot \rangle_A)$.

231 (b) $K_h u_h = \Pi_h(K u_h)$ for all $u_h \in X_h$.

232 (c) $K = (-\Delta + V_1)^{-1}V_2$ and $K_h = (-\Delta_h + V_{1,h})^{-1}V_{2,h}$.

233 *Proof.* (a) First, the multiplication $u \mapsto V_2 u$ is bounded on L^2 , since $V_2 \in L^\infty(\Omega)$.
234 Second, R is a bounded operator, since it represents $(-\Delta + V_1)^{-1}$. The embedding ι
235 is compact by the Rellich–Kondrachov theorem, so $K = R \circ (V_2 \cdot) \circ \iota$ is compact. For
236 self-adjointness, (3.3) gives $\langle Ku, v \rangle_A = \int V_2 uv$ and $\langle Kv, u \rangle_A = \int V_2 vu$. Since $\langle \cdot, \cdot \rangle_A$
237 is symmetric, $\langle Ku, v \rangle_A = \langle u, Kv \rangle_A$.

238 (b) For any $u_h, v_h \in X_h$,

$$239 \quad \langle K_h u_h, v_h \rangle_A = \int_{\Omega} V_2 u_h v_h = \langle K u_h, v_h \rangle_A = \langle \Pi_h(K u_h), v_h \rangle_A,$$

240 so $K_h u_h = \Pi_h(K u_h)$ for all $u_h \in X_h$, i.e., $K_h = \Pi_h K|_{X_h}$.

241 (c) By definition, $\langle Ku, v \rangle_A = \int V_2 uv = (V_2 u, v)_{L^2} = \langle R(V_2 u), v \rangle_A$, so $Ku = R(V_2 u)$,
242 i.e., $K = (-\Delta + V_1)^{-1}V_2$. Similarly, $K_h u_h$ is the unique element of X_h satisfying
243 $\langle K_h u_h, v_h \rangle_A = \int V_2 u_h v_h$ for all $v_h \in X_h$, which is precisely $(-\Delta_h + V_{1,h})^{-1}(V_{2,h} u_h)$. \square

244 THEOREM 3.4 (Spectral convergence of the preconditioned operator). For a
245 bounded domain $\Omega \subset \mathbb{R}^d$, under Assumption I, let $T = (-\Delta + V_1)^{-1}(-\Delta + V_1 + V_2) =$
246 $I + K$ and $T_h = (-\Delta_h + V_{1,h})^{-1}(-\Delta_h + V_{1,h} + V_{2,h}) = I_h + K_h$, where I_h denotes
247 the identity on X_h , be the continuous and discrete preconditioned operators, and let
248 $\tilde{K}_h = \Pi_h K \Pi_h : H_0^1(\Omega) \rightarrow H_0^1(\Omega)$. Then:

249 (i) **Norm convergence.** $\|\tilde{K}_h - K\| \rightarrow 0$ as $h \rightarrow 0$ in the operator norm on
250 $(H_0^1(\Omega), \|\cdot\|_A)$.

251 (ii) **Spectral convergence and clustering.** For every $\epsilon > 0$, T_h has at most
252 $M(\epsilon)$ eigenvalues outside $(1-\epsilon, 1+\epsilon)$, where $M(\epsilon)$ is independent of h . Each
253 nonzero eigenvalue μ of K with multiplicity m is approximated by exactly m
254 eigenvalues of K_h converging to μ as $h \rightarrow 0$.

255 (iii) **Uniform condition number.** If $-\Delta + V_1 + V_2 > 0$, then $\kappa(T_h) \leq \kappa(T)$ for
 256 every h , and $\kappa(T_h) \rightarrow \kappa(T)$ as $h \rightarrow 0$.

257 *Proof.* Since Π_h maps $H_0^1(\Omega)$ onto X_h , $\tilde{K}_h = \Pi_h K \Pi_h$ maps X_h^\perp to zero and
 258 restricts to K_h on X_h (by Lemma 3.3(b), $\Pi_h K u_h = K_h u_h$ for $u_h \in X_h$). Therefore
 259 the nonzero eigenvalues of \tilde{K}_h are exactly those of K_h : if $\tilde{K}_h f = \mu f$ with $\mu \neq 0$, then
 260 $f = \mu^{-1} \tilde{K}_h f \in X_h$, so f is an eigenfunction of K_h .

261 *Part (i): Norm convergence.* Write $K - \tilde{K}_h = (I - \Pi_h)K + \Pi_h K (I - \Pi_h)$. Since K
 262 and Π_h are both self-adjoint in $\langle \cdot, \cdot \rangle_A$, the adjoint of $\Pi_h K (I - \Pi_h)$ is $(I - \Pi_h)K \Pi_h$.
 263 Because $\|B\| = \|B^*\|$ for any bounded operator B and its adjoint B^* on a Hilbert
 264 space, $\|\Pi_h K (I - \Pi_h)\| = \|(I - \Pi_h)K \Pi_h\| \leq \|(I - \Pi_h)K\| \|\Pi_h\| = \|(I - \Pi_h)K\|$, thus

$$265 \quad (3.4) \quad \|\tilde{K}_h - K\| \leq 2 \|(I - \Pi_h)K\|.$$

266 It remains to show $\|(I - \Pi_h)K\| \rightarrow 0$. Since K is compact, the image of the closed
 267 unit ball $\bar{B}_1 = \{f \in H_0^1(\Omega) : \|f\|_A \leq 1\}$ under K is precompact in $(H_0^1(\Omega), \|\cdot\|_A)$.
 268 Consider any fixed $\epsilon > 0$, and cover $K(\bar{B}_1)$ by finitely many ϵ -balls centered at
 269 $g_1, \dots, g_M \in H_0^1(\Omega)$. Since $C_c^\infty(\Omega)$ is dense in $H_0^1(\Omega)$ and conforming finite element
 270 spaces approximate smooth functions as $h \rightarrow 0$ (by polynomial interpolation), $\bigcup_h X_h$
 271 is dense in $H_0^1(\Omega)$, so $(I - \Pi_h)g_j \rightarrow 0$ in $\|\cdot\|_A$ for each j , thus there exists h_0 such
 272 that $\|(I - \Pi_h)g_j\|_A < \epsilon$ for all j and all $h < h_0$. For any f with $\|f\|_A \leq 1$, choose g_j
 273 with $\|Kf - g_j\|_A < \epsilon$. Since $\|I - \Pi_h\| \leq 1$ (orthogonal projection),

$$274 \quad \|(I - \Pi_h)Kf\|_A \leq \underbrace{\|(I - \Pi_h)(Kf - g_j)\|_A}_{\leq \epsilon} + \underbrace{\|(I - \Pi_h)g_j\|_A}_{< \epsilon} < 2\epsilon.$$

275 Hence $\|(I - \Pi_h)K\| \leq 2\epsilon$ for $h < h_0$. Since ϵ was arbitrary, $\|(I - \Pi_h)K\| \rightarrow 0$, and
 276 by (3.4), $\|\tilde{K}_h - K\| \rightarrow 0$.

277 *Part (ii): Clustering.* Since Π_h is an orthogonal projection in $\langle \cdot, \cdot \rangle_A$, its operator
 278 norm is $\|\Pi_h\| = 1$, so the singular values s_n satisfy $s_n(\tilde{K}_h) = s_n(\Pi_h K \Pi_h) \leq s_n(K)$.
 279 For self-adjoint operators, singular values s equal absolute values of eigenvalues μ , so
 280 $|\mu_n(\tilde{K}_h)| \leq |\mu_n(K)|$ for each n , where eigenvalues are ordered by decreasing absolute
 281 value. Since K is compact and self-adjoint, its eigenvalues accumulate only at 0 [31,
 282 Theorem VI.16], so for every $\epsilon > 0$, at most $M(\epsilon) < \infty$ of them satisfy $|\mu_n(K)| \geq \epsilon$.
 283 Since $|\mu_n(\tilde{K}_h)| \leq |\mu_n(K)|$, the same bound holds for \tilde{K}_h , independently of h . Hence
 284 $T_h = I_h + \tilde{K}_h$ has at most $M(\epsilon)$ eigenvalues outside $(1-\epsilon, 1+\epsilon)$.

285 *Convergence.* Since $\tilde{K}_h \rightarrow K$ in operator norm and K and \tilde{K}_h are compact
 286 and self-adjoint, Osborn's spectral approximation theory [30] applies directly. By the
 287 spectral convergence results in [30, Section 3], for each nonzero eigenvalue μ of K
 288 with multiplicity m , there exist exactly m eigenvalues of \tilde{K}_h converging to μ . Since
 289 the nonzero eigenvalues of \tilde{K}_h are exactly those of K_h , the same convergence holds
 290 for the eigenvalues of K_h .

291 *Part (iii): Uniform condition number.* For any $u_h \in X_h$ with $\|u_h\|_A = 1$,

$$292 \quad \langle K_h u_h, u_h \rangle_A = \int_{\Omega} V_2 |u_h|^2 \, d\mathbf{x} = \langle K u_h, u_h \rangle_A,$$

293 since $u_h \in X_h \subset H_0^1(\Omega)$. By the Courant–Fischer min-max principle, $\mu_{\max}(K_h) =$
 294 $\max_{u_h \in X_h, \|u_h\|_A=1} \langle K u_h, u_h \rangle_A$ and $\mu_{\max}(K) = \sup_{u \in H_0^1, \|u\|_A=1} \langle K u, u \rangle_A$. Since $X_h \subset$

295 $H_0^1(\Omega)$, the maximum over the smaller set X_h cannot exceed the supremum over
 296 $H_0^1(\Omega)$. The same reasoning applies to μ_{\min} . Therefore:

$$297 \quad \mu_{\min}(K) \leq \mu_{\min}(K_h) \quad \text{and} \quad \mu_{\max}(K_h) \leq \mu_{\max}(K).$$

298 Since $T_h = I_h + K_h$ and $T = I + K$, this gives $\mu_{\min}(T) \leq \mu_{\min}(T_h)$ and $\mu_{\max}(T_h) \leq$
 299 $\mu_{\max}(T)$. If $T > 0$, then $\mu_{\min}(T) > 0$, so

$$300 \quad \kappa(T_h) = \frac{\mu_{\max}(T_h)}{\mu_{\min}(T_h)} \leq \frac{\mu_{\max}(T)}{\mu_{\min}(T)} = \kappa(T) \quad \text{for every } h.$$

301 Since $\mu_{\max}(K)$ and $\mu_{\min}(K)$ (if nonzero) are isolated eigenvalues of K , Part (ii) gives
 302 eigenvalues of K_h converging to them. Combined with $\mu_{\max}(K_h) \leq \mu_{\max}(K)$, this
 303 gives $\mu_{\max}(K_h) \rightarrow \mu_{\max}(K)$. Similarly $\mu_{\min}(K_h) \rightarrow \mu_{\min}(K)$. Hence $\kappa(T_h) \rightarrow \kappa(T)$. \square

304 **THEOREM 3.5** (Domain-independent clustering for confining potentials). *Under*
 305 *Assumption II, consider the family of domains $\Omega_L = [-L, L]^d$ with conforming finite*
 306 *element subspaces $X_h \subset H_0^1(\Omega_L)$, and the preconditioned operator $T_{h,L}$ defined as*
 307 *in Lemma 3.3, then for every $\epsilon > 0$, there exist at most $M(\epsilon)$ eigenvalues of $T_{h,L}$*
 308 *outside $(1-\epsilon, 1+\epsilon)$, where $M(\epsilon) < \infty$ is independent of both h and L . Moreover, if*
 309 *$-\Delta + V_1 + V_2 > 0$ on \mathbb{R}^d , then $\kappa(T_{h,L})$ is bounded independently of both h and L .*

310 *Proof.* Since V_1 is confining, the Schrödinger operator $-\Delta + V_1$ on \mathbb{R}^d has compact
 311 resolvent and therefore purely discrete spectrum $\Lambda_1 \leq \Lambda_2 \leq \dots \rightarrow \infty$ [32,
 312 Theorem XIII.16]. By the min-max principle, the n -th eigenvalue of $-\Delta + V_1$ on Ω_L
 313 satisfies

$$314 \quad \lambda_n(-\Delta_L + V_1) = \min_{\substack{\dim S=n \\ S \subset H_0^1(\Omega_L)}} \max_{u \in S, \|u\|_{L^2}=1} \|u\|_A^2.$$

315 Restricting to $H_0^1(\Omega_L) \subset H^1(\mathbb{R}^d)$ reduces the set of candidate subspaces for the outer
 316 min, so the eigenvalues can only increase: $\lambda_n(-\Delta_L + V_1) \geq \Lambda_n$ for all n and L . The
 317 operator norm of the mapping $(-\Delta_L + V_1)^{-1} : L^2(\Omega_L) \rightarrow L^2(\Omega_L)$ satisfies

$$318 \quad (3.5) \quad \|(-\Delta_L + V_1)^{-1}\|_{L^2 \rightarrow L^2} = \frac{1}{\lambda_1(-\Delta_L + V_1)} \leq \frac{1}{\Lambda_1}, \quad \text{independently of } L.$$

319 By contrast, for $\Omega_L = [-L, L]^d$, $\|(-\Delta_L)^{-1}\|_{L^2 \rightarrow L^2} = 1/\lambda_1(-\Delta_L) \sim L^2/\pi^2 \rightarrow \infty$.

320 Recall that Lemma 3.3 gives $K_L = R \circ (V_2 \cdot) \circ \iota$, where $V_2 \cdot$ is bounded with
 321 norm $\|V_2\|_{L^\infty}$. By the definition of R with test function Rf , $\|Rf\|_A^2 = (f, Rf)_{L^2} \leq$
 322 $\|f\|_{L^2} \|Rf\|_{L^2} \leq \|f\|_{L^2} \cdot \Lambda_1^{-1} \|f\|_{L^2}$ by (3.5), so $\|Rf\|_A \leq \Lambda_1^{-1/2} \|f\|_{L^2}$. Thus R is
 323 bounded as a map $L^2(\Omega_L) \rightarrow (H_0^1(\Omega_L), \|\cdot\|_A)$ with norm $\leq 1/\sqrt{\Lambda_1}$.

324 The singular values of the compact embedding $\iota : (H_0^1(\Omega_L), \|\cdot\|_A) \hookrightarrow L^2(\Omega_L)$
 325 are determined as follows. The adjoint $\iota^* : L^2(\Omega_L) \rightarrow H_0^1(\Omega_L)$ satisfies $\langle \iota^* f, v \rangle_A =$
 326 $(f, v)_{L^2}$ for all $v \in H_0^1(\Omega_L)$, so $\iota^* f$ solves $(-\Delta + V_1)w = f$ weakly, i.e., $\iota^* = (-\Delta_L +$
 327 $V_1)^{-1}$ as a map $L^2 \rightarrow H_0^1$. Therefore $\iota^* \iota : H_0^1 \rightarrow H_0^1$ has eigenvalues $1/\lambda_n(-\Delta_L + V_1)$,
 328 and the singular values of ι are $s_n(\iota) = 1/\sqrt{\lambda_n(-\Delta_L + V_1)} \leq 1/\sqrt{\Lambda_n}$. Since K_L is
 329 self-adjoint, its singular values equal the absolute values of its eigenvalues, and the
 330 multiplicative singular value inequality gives

$$331 \quad |\mu_n(K_L)| \leq \frac{\|V_2\|_{L^\infty(\mathbb{R}^d)}}{\sqrt{\Lambda_1 \Lambda_n}}.$$

332 Therefore $|\mu_n(K_L)| \geq \epsilon$ requires $\Lambda_n \leq \|V_2\|_{L^\infty}^2 / (\Lambda_1 \epsilon^2)$, which holds for at most $M(\epsilon)$
 333 values of n , independently of L . This bound transfers to the discrete operators.

334 Define $\tilde{K}_{h,L} = \Pi_h K_L \Pi_h$ as in Theorem 3.4. Since $\|\Pi_h\| = 1$ (orthogonal projection),
 335 the singular value submultiplicativity gives $s_n(\Pi_h K_L \Pi_h) \leq s_n(K_L)$. For self-adjoint
 336 operators $|\mu_n| = s_n$, so $|\mu_n(K_{h,L})| \leq |\mu_n(K_L)|$ for each n . Hence $K_{h,L}$ also has at
 337 most $M(\epsilon)$ eigenvalues outside $[-\epsilon, \epsilon]$, independently of both h and L .

338 For the condition number bound, Theorem 3.4(iii) gives $\kappa(T_{h,L}) \leq \kappa(T_L)$ for each
 339 L . It remains to show $\kappa(T_L)$ is bounded independently of L . Let $K_{\mathbb{R}^d} = (-\Delta + V_1)^{-1}V_2$
 340 on \mathbb{R}^d (well-defined since V_1 is confining). Any $u \in H_0^1(\Omega_L)$, extended by zero to
 341 $\tilde{u} \in H^1(\mathbb{R}^d)$, satisfies $\|\tilde{u}\|_A = \|u\|_A$ and $\langle K_{\mathbb{R}^d} \tilde{u}, \tilde{u} \rangle_A = \int V_2 |\tilde{u}|^2 = \int_{\Omega_L} V_2 |u|^2 =$
 342 $\langle K_L u, u \rangle_A$. Thus the Rayleigh quotients of K_L over $H_0^1(\Omega_L)$ are a subset of those
 343 of $K_{\mathbb{R}^d}$ over $H^1(\mathbb{R}^d)$, giving $\mu_{\min}(K_L) \geq \mu_{\min}(K_{\mathbb{R}^d})$ and $\mu_{\max}(K_L) \leq \mu_{\max}(K_{\mathbb{R}^d})$.
 344 Therefore $\kappa(T_L) \leq \kappa(T_{\mathbb{R}^d})$ for all L , and $\kappa(T_{h,L}) \leq \kappa(T_{\mathbb{R}^d})$ for all h and L . \square

345 **REMARK 3.6** (Grid-independent PCG convergence). *For the stirrer potential (3.2),*
 346 V_1 *is confining and* $V_2 \in L^\infty(\mathbb{R}^d)$, *so Theorem 3.5 guarantees that the eigenvalues of*
 347 $T_{h,L}$ *cluster near 1 with at most* $M(\epsilon)$ *outliers, and* $\kappa(T_{h,L})$ *is bounded, independently*
 348 *of both* h *and* L , *which explains mesh- and domain-independent PCG iteration counts.*
 349 *In particular, CG effectively ignores finitely many outlier eigenvalues [37], which can*
 350 *explain the fast convergence (5–6 iterations) observed in Example 3.1 (Figure 1, bot-*
 351 *tom right), as well as the 5 iterations on* \mathbb{R}^3 *with the Hermite spectral method (Figure 9*
 352 *in Appendix A). For the quartic potential (3.1),* $V_2 \notin L^\infty(\mathbb{R}^d)$, *so only Theorem 3.4*
 353 *applies. For each fixed* L , $\kappa(T_h) \leq \kappa(T)$ *ensures mesh-independent iteration counts*
 354 *(Figure 1, top left), but* $\kappa(T_L)$ *grows with* L , *and iterations increase as* L *increases*
 355 *(Figure 1, bottom left).*

356 4. Ground State Computation.

357 **4.1. Linear ground state via shifted inverse iteration.** When $V = V_1$ is
 358 fully separable, the tensor-product solver computes $(-\Delta + V_1)^{-1}b$ directly, which can
 359 be used in shifted inverse iteration (power method for $(-\Delta + V_1 - \sigma)^{-1}$) for computing
 360 the ground state eigenvector of $(-\Delta + V_1)u = \lambda u$. We use the potential in [9]:

$$361 \quad (4.1) \quad V_1(x, y, z) = \sum_{k=1}^3 (x_k^2 + 100 \sin^2(\pi x_k/4)).$$

TABLE 2

Shifted inverse iteration for the ground state of $(-\Delta + V_1)u = \lambda u$ on $[-8, 8]^3$, Q^{10} SEM, FP64. Potential (4.1). Shift $\sigma = 0.9 \lambda_{\min}(-\Delta + V_1)$. Convergence criterion $|\Delta\lambda|/|\lambda| < 10^{-12}$. Ground state eigenvalue $\lambda_1 = 23.2878438176$.

GPU	DoFs	Setup (s)	JIT (s) [†]	Iters / time (s)	Peak mem (GB)
A100	599 ³	6.2	1.0	9 / 1.25	11.2
A100	999 ³	6.4	4.3	9 / 7.61	52.1
GH200	999 ³	9.3	0.9	9 / 3.50	52.1
GH200	1099 ³	9.2	1.2	9 / 5.03	69.4

[†] JAX/XLA compilation time. At $n=999$: GH200 is $2.2\times$ faster per solve than A100 (0.39s vs. 0.85s). 1.33×10^9 DoFs converge in 5s of iteration time on a single GH200.

362

363 Table 2 shows that the shifted inverse iteration converges in 9 iterations regardless
 364 of grid size, producing the ground state of $(-\Delta + V_1)u = \lambda u$. This eigenfunction will
 365 serve as a natural initial guess for the nonlinear GPE gradient flow below.

366 **4.2. GPE ground state via Sobolev gradient flow.** We consider the ground
 367 state of the defocusing ($\beta > 0$) Gross–Pitaevskii eigenvalue problem

$$368 \quad (4.2) \quad -\Delta u + V(\mathbf{x})u + \beta|u|^2u = \lambda u, \quad \mathbf{x} \in \Omega \subset \mathbb{R}^3, \quad \|u\|_{L^2} = 1,$$

369 with potential $V \geq 0$ and the parameter $\beta > 0$. The ground state minimizes

$$370 \quad (4.3) \quad E(u) = \frac{1}{2} \int_{\Omega} (|\nabla u|^2 + V|u|^2) d\mathbf{x} + \frac{\beta}{4} \int_{\Omega} |u|^4 d\mathbf{x}.$$

371 For $\beta > 0$, E has a unique positive ground state [20, 9], thus all computations in this
 372 section use real arithmetic. The ground state can be computed by the normalized
 373 gradient flow (GFDN) [4] or by projected Sobolev gradient flows [20, 11, 9].

374 **4.3. Riemannian gradient flows.** The Sobolev gradient flows can also be per-
 375 ceived as Riemannian gradient descent methods. We consider two Riemannian gradi-
 376 ent flows on the L^2 unit sphere $\mathcal{S} = \{u : \|u\|_{L^2} = 1\}$. Both take the form

$$377 \quad (4.4) \quad u^{n+1} = \frac{u^n - \tau \nabla_g^{\mathcal{R}} E(u^n)}{\|u^n - \tau \nabla_g^{\mathcal{R}} E(u^n)\|_{L^2}},$$

378 where $\nabla_g^{\mathcal{R}} E$ is the Riemannian gradient with respect to the metric g , and $\tau > 0$ is
 379 the step size. They differ only in the choice of g . We consider just two of them.

380 **Modified H^1 flow** [9]. The metric is

$$381 \quad (4.5) \quad g_{H^1}(w, z) = (\nabla w, \nabla z) + \alpha(w, z), \quad \alpha > 0.$$

382 Each iteration requires solving $(-\Delta + \alpha I)w = (-\Delta + V + \beta u^2)u$ and a projection
 383 step, amounting to two direct solves of $(-\Delta + \alpha I)^{-1}$ per iteration, a separable tensor-
 384 product operation at cost $O(N^{4/3})$ in 3D.

385 **a_u flow** [20]. The metric depends on the current iterate u :

$$386 \quad (4.6) \quad g_{a_u}(w, z) = (\nabla w, \nabla z) + (w, Vz) + \beta(w, u^2 z).$$

387 Each iteration requires solving $(-\Delta + V + \beta u^2)^{-1}$. We solve this by PCG with
 388 preconditioner $(-\Delta + V)^{-1}$, which is a separable tensor-product solve. We refer
 389 to [20, 9] for more details about a_u flow. The a_u metric adapts to the current iterate.
 390 It reduces the iteration count, but each step is more expensive because of the PCG
 391 solve. The tensor-product solver enters both methods: (i) shifted inverse iteration
 392 provides the eigenfunction of the linearized problem $(-\Delta + V_1)u = \lambda u$ as an initial
 393 guess (Section 4.1), (ii) the modified H^1 flow requires $(-\Delta + \alpha I)^{-1}$ (direct solve),
 394 and (iii) the a_u flow uses $(-\Delta + V)^{-1}$ as a PCG preconditioner.

395 **4.4. Numerical results.** When FP32 is used below for the PCG solves in the
 396 a_u flow, the offline eigendecompositions per axis are still computed in FP64. We
 397 consider the ground state of (4.2) for a potential $V = V_1$ with V_1 given in (4.1).

398 **EXAMPLE 4.1** (GPE ground state with eigenfunction initial guess). *Since the*
 399 *GPE ground state approaches the linear ground state as $\beta \rightarrow 0$, the eigenfunction*
 400 *from Table 2 provides a good initial guess for moderate β . Table 3 compares the*
 401 *modified H^1 flow [9] with two initial guesses: a constant function $u_0 \equiv c$ (normalized*
 402 *s.t. $\|u_0\|_{L^2} = 1$) and the linear ground state from shifted inverse iteration. Both use*
 403 *Q^{20} SEM with 599^3 DoFs, $\alpha = 20$, $\tau = 0.1$, tested for $\beta = 10$ and 100 on an NVIDIA*

404 GH200 96 GB GPU. For moderate β , the nonlinear ground state remains close to the
 405 linear eigenfunction, so the eigenfunction initial guess starts the flow near the solution.
 406 Figure 2 illustrates the progression: at $\beta = 100$ the lattice structure is already emerging
 407 but the solution remains concentrated near the center, while at $\beta = 1600$ the ground
 408 state has spread into a broad profile far from the linear eigenfunction. As shown in
 409 Table 3, for $\beta = 10, 100$, modified H^1 flow with both initializations converges to the
 410 same energy and eigenvalue, and the eigenfunction initial guess saves 35–71% of total
 411 time with 39–77% fewer iterations. Figure 3 illustrates that both H^1 flow and a_u
 412 flow for $\beta = 100$ can benefit from the eigenfunction initial guess. But for $\beta = 1600$,
 413 as shown in Table 4 ($N = 599^3$) and Figure 4 ($N = 999^3$), the eigenfunction is no
 414 longer a good initial guess. Both Figure 3 and Figure 4 also show that a_u flow can be
 415 efficiently implemented by the PCG with the proposed preconditioner $(-\Delta + V_1)^{-1}$.




FIG. 2. Cross-sections of u at $x=0, y=0, z=0$, plotted on $[-6, 6]^3$. Q^{20} SEM, 199^3 DoFs, computed on $[-8, 8]^3$. Each panel uses its own color scale. Left: eigenfunction initial guess (linear ground state, $\beta=0$). Center: GPE ground state, $\beta=100$. Right: GPE ground state, $\beta=1600$.

TABLE 3

Modified H^1 flow for the GPE ground state on $[-8, 8]^3$, Q^{20} SEM, FP64, $\alpha = 20$, step size $\tau = 0.1$, 599^3 DoFs, GH200. Potential (4.1). Convergence criterion: $|E_k - E^*|/|E^*| < 10^{-14}$. The total computation time includes both off-line and iteration time.

β	Init	H^1 iters	Total (s)	Energy	Eigenvalue
10	constant	745	122.0	14.1965761916	32.4916917439
	eigenfunction	175	34.8	14.1965761916	32.4916917439
100	constant	660	111.4	20.6824463703	47.7831207152
	eigenfunction	405	71.9	20.6824463703	47.7831207152

TABLE 4

Comparison of a_u flow and modified H^1 flow for $\beta = 1600$, Q^{20} SEM, $599^3 \approx 2.1 \times 10^8$ DoFs, FP64, $\alpha=20$, step size $\tau=0.1$ for modified H^1 and $\tau=1$ for a_u , GH200.

Method	Initialization	Iters	Linear solves	Time (s)	Energy
a_u flow	constant	66	790	72.8	33.80227900550
a_u flow	eigenfunction	147	1846	171.9	33.80227900550
Modified H^1	constant	272	544	42.6	33.80227900551
Modified H^1	eigenfunction	not converging			

Convergence criterion: $|E_k - E^*|/|E^*| < 10^{-12}$, where $E^* = 33.80227900547423$.




FIG. 3. Comparison of a_u flow and modified H^1 flow for $\beta = 100$, Q^{20} SEM, $599^3 \approx 2.1 \times 10^8$ DoFs, GH200. Left: relative energy error vs. iteration. Right: relative energy error vs. GPU time.

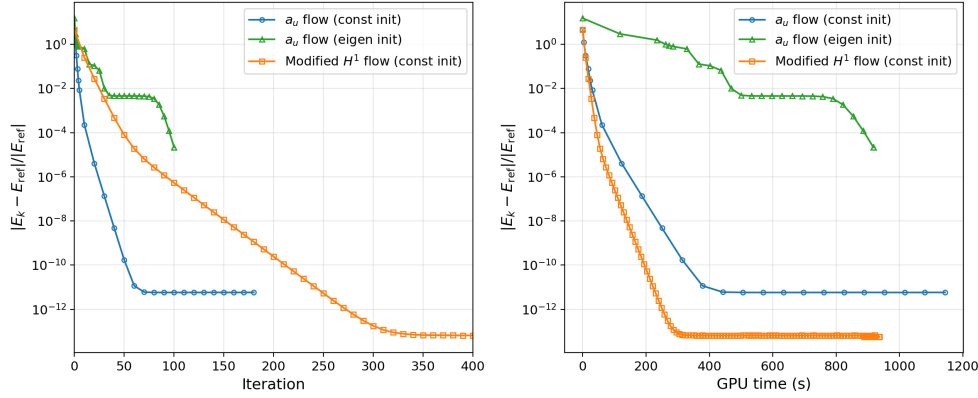


FIG. 4. Comparison of a_u flow and modified H^1 flow for $\beta = 1600$, Q^{20} SEM, $999^3 \approx 10^9$ DoFs, constant initial guess, GH200. Left: relative energy error vs. iteration. Right: relative energy error vs. GPU time. The reference energy $E^* = 33.80227900547$ is from [9].

416 **5. Hamiltonian Simulation in 3D.** We consider the Schrödinger equation

417 (5.1)
$$i \partial_t \psi = H \psi, \quad H = -\Delta + V, \quad \psi(\cdot, 0) = \psi_0,$$

418 on $\Omega \subset \mathbb{R}^3$ with either periodic or Dirichlet boundary conditions. When V is sep-
 419 arable, the tensor-product solver gives the exact propagator $e^{-iH\Delta t}$. When V is
 420 non-separable, it provides the propagator for each operator-splitting sub-step.

421 **5.1. Splitting methods and classical implementation.** We split $H = A +$
 422 B with $A = -\Delta$ (kinetic) and $B = V$ (potential). The standard Strang splitting
 423 $e^{-iA\Delta t/2} e^{-iB\Delta t} e^{-iA\Delta t/2}$ is second-order in Δt , but its operator-norm error grows
 424 with the spectral radius of $[A, B]$, which may scale with N . Two recent quantum
 425 algorithms consider the interaction-picture Hamiltonian $H_I(s) = e^{iAs} B e^{-iAs}$ and
 426 apply the Magnus expansion to the time-ordered evolution:

427 • The $qHOP$ algorithm [1] uses the first-order Magnus truncation, achieving

428 $O(\Delta t^2)$ with an error constant independent of N .

- 429 • The *Magnus-2* algorithm [12] adds a commutator correction (second-order
- 430 Magnus), achieving $O(\Delta t^4)$ superconvergence [7, 14].

431 Both methods require a matrix exponential of a sum of interaction-picture Hamilto-
 432 nians, which on a quantum computer is implemented via LCU [1]. Classically, e^{iAs} is
 433 computed by the tensor-product propagator and $e^{-iB\Delta t}$ is pointwise, so each factor
 434 $\exp(-i w_k \Delta t H_I(s_k))$ can be evaluated exactly:

$$435 \quad e^{-i w_k \Delta t H_I(s_k)} = e^{iAs_k} e^{-i w_k \Delta t B} e^{-iAs_k},$$

436 which leads to two families of methods, implemented via tensor-product propagations.

437 **$O(\Delta t^2)$ approximated qHOP via product formula.** The qHOP approxima-
 438 tion [1] replaces the time-ordered exponential by the first-order Magnus truncation in
 439 the interaction picture (see [12, Eq. (2)]):

$$440 \quad e^{-iH\Delta t} = e^{-iA\Delta t} \mathcal{T} \exp\left(-i \int_0^{\Delta t} H_I(s) ds\right) \approx e^{-iA\Delta t} \exp\left(-i \Delta t \sum_{k=1}^M w_k H_I(s_k)\right),$$

441 where $s_k \in [0, \Delta t]$ are Gauss–Legendre nodes and w_k ($\sum_{k=1}^M w_k = 1$) are the corre-
 442 sponding normalized Gauss–Legendre weights. We implement this classically via the
 443 product approximation

$$444 \quad (5.2) \quad \exp\left(-i \Delta t \sum_{k=1}^M w_k H_I(s_k)\right) \approx \prod_{k=1}^M e^{-i w_k \Delta t H_I(s_k)}.$$

445 Since $H_I(s_k) = e^{iAs_k} B e^{-iAs_k}$, each factor in the product is computed *exactly* as
 446 $e^{-i w_k \Delta t H_I(s_k)} = e^{iAs_k} e^{-i w_k \Delta t B} e^{-iAs_k}$ (two tensor-product propagations plus one
 447 pointwise multiplication). Combining (5.2) with $e^{-iH\Delta t} = e^{-iA\Delta t} \mathcal{T} \exp(\dots)$ gives
 448 the full time step $e^{-iH\Delta t} \approx e^{-iA\Delta t} \prod_{k=1}^M e^{-i w_k \Delta t H_I(s_k)}$, where the approximation
 449 comes from the quadrature and product formula (5.2).

450 **REMARK 5.1** (Order of the product formula). *The product formula (5.2) re-*
 451 *places $\exp(\text{sum})$ by a product of exponentials, which is in general a first-order Lie-*
 452 *Trotter approximation with error governed by the commutators $[A_j, A_k]$ where $A_k =$*
 453 *$-i w_k \Delta t H_I(s_k)$. In Lie–Trotter splitting of two different operators A and B , the com-*
 454 *mutator $[A, B]$ is a fixed operator, giving first-order global error $O(\Delta t)$. Here, how-*
 455 *ever, all the operators $H_I(s_k)$ have $s_k \in [0, \Delta t]$, so as $\Delta t \rightarrow 0$, $H_I(s_k) \rightarrow H_I(0) = B$*
 456 *for all k , and the commutators $[H_I(s_j), H_I(s_k)] \rightarrow [B, B] = 0$. More precisely,*
 457 *smoothness of $H_I(s)$ in s gives $[H_I(s_j), H_I(s_k)] = O(|s_j - s_k|) = O(\Delta t)$, so each*
 458 *$[A_j, A_k] = w_j w_k \Delta t^2 [H_I(s_j), H_I(s_k)] = O(\Delta t^3)$. The product formula error is there-*
 459 *fore $O(\Delta t^3)$ per step and $O(\Delta t^2)$ globally, matching the qHOP approximation order.*

460 **$O(\Delta t^4)$ approximated Magnus-2 via Yoshida composition.** The Magnus-2
 461 algorithm [12] replaces the first-order Magnus truncation with the second-order one:

$$462 \quad e^{-iH\Delta t} \approx e^{-iA\Delta t} \exp(\Omega_1 + \Omega_2),$$

463 where $\Omega_1 = -i \Delta t \sum_{k=1}^M w_k H_I(s_k)$ is the qHOP exponent and

$$464 \quad \Omega_2 = -\frac{1}{2} \int_0^{\Delta t} \int_0^{s_1} [H_I(s_1), H_I(s_2)] ds_2 ds_1$$

465 is the commutator correction, achieving $O(\Delta t^4)$ superconvergence. On a quantum
 466 computer, $e^{\Omega_1 + \Omega_2}$ is implemented via LCU. Classically, it is impractical to compute
 467 $e^{\Omega_1 + \Omega_2}$ directly. Instead, we apply the fourth-order Yoshida splitting [38] to the
 468 approximated qHOP:

$$469 \quad (5.3) \quad e^{-iH\Delta t} \approx P(\gamma_1\Delta t) P(\gamma_2\Delta t) P(\gamma_1\Delta t), \quad \gamma_1 = \frac{1}{2 - 2^{1/3}}, \quad \gamma_2 = \frac{-2^{1/3}}{2 - 2^{1/3}},$$

470 where $P(h)$ denotes one approximated qHOP step with step size h , i.e., $P(h) =$
 471 $e^{-iAh} \prod_{k=1}^M e^{-i w_k h H_I(s_k)}$, and $\gamma_1 + \gamma_2 + \gamma_1 = 1$. Yoshida cancels the $O(h^3)$ leading
 472 error ($2\gamma_1^3 + \gamma_2^3 = 0$), yielding $O(\Delta t^4)$ global accuracy, the same order as the original
 473 Magnus-2 method. Although (5.3) is derived by applying Yoshida composition to
 474 the approximated qHOP rather than by directly approximating $e^{\Omega_1 + \Omega_2}$, it achieves
 475 the same $O(\Delta t^4)$ rate as the original Magnus-2, and can therefore be regarded as an
 476 approximation to Magnus-2.

477 **REMARK 5.2** (Classical splitting as limiting cases). *qHOP with $M=1$ (midpoint*
 478 *rule) recovers the Strang splitting [1]. At the next order, our approximated Magnus-*
 479 *2 with $M=1$ reduces to the plain Yoshida method. Thus Strang and Yoshida can*
 480 *be perceived as the $M=1$ special cases of such approximated versions of qHOP and*
 481 *Magnus-2, respectively, while larger M reduces the error constant.*

482 **REMARK 5.3** (Cost per step). *For $M \geq 2$, the approximated qHOP in (5.2) has a*
 483 *direct count of $(2M+1)$ tensor-product propagations per step, i.e., number of multiply-*
 484 *ing $e^{\pm iAs_k}$ per step. For $M = 1$, the optimized Strang form $e^{-iV\Delta t/2} e^{-iA\Delta t} e^{-iV\Delta t/2}$*
 485 *needs only 1 tensor-product propagation per step. A better implementation for $M > 1$*
 486 *is to merge two adjacent factors, e.g., for $M = 3$, the merged form of the qHOP is*

$$487 \quad e^{-iA(\Delta t - s_3)} e^{-i w_3 \Delta t B} e^{-iA(s_3 - s_2)} e^{-i w_2 \Delta t B} e^{-iA(s_2 - s_1)} e^{-i w_1 \Delta t B} e^{-iAs_1},$$

488 *which has only $(M+1)$ A -propagations per step. By merging e^{-iAs_1} at step n with*
 489 *$e^{-iA(\Delta t - s_M)}$ at step $n+1$, the cost would be M A -propagations per time step. Ap-*
 490 *plying the same merging strategy to each of the three Yoshida sub-steps and across*
 491 *consecutive time steps, the approximated Magnus-2 costs $3M$ A -propagations per step*
 492 *asymptotically for $M \geq 2$, and 3 A -propagations plus 3 pointwise B -multiplications*
 493 *per step for $M = 1$ (plain Yoshida).*

494 **5.2. Exact propagator as reference for testing splitting methods.** When
 495 $V = V_1$ is separable, the eigendecomposition $H = T \Lambda T^{-1}$ gives the exact propagator

$$496 \quad (5.4) \quad e^{-iH\Delta t} \psi = T (e^{-i\Lambda\Delta t} \odot (T^{-1} \psi)),$$

497 at cost $O(N^{4/3})$ in 3D. We test with the separable potential (4.1) from Section 4.1 on
 498 $[-8, 8]^3$, Dirichlet BC, Q^k SEM, and the exact propagator (5.4) as reference.

499 Most existing numerical tests of these methods are in 1D with $N \leq 1024$, using
 500 dense matrix exponentials (`expm`). Our tensor-product solver provides the exact prop-
 501 agator as a reference solution at $\sim 10^8$ DoFs, enabling large-scale 3D validation of the
 502 splitting convergence rates. In the rest of this section, “qHOP” and “Magnus-2” refer
 503 to the approximated versions implemented via the product formula (5.2) and Yoshida
 504 composition (5.3), respectively.

505 **EXAMPLE 5.1** (Order of approximated qHOP). *Table 5 reports qHOP with $M =$*
 506 *1 (Strang), 3, 5, and 7 on a GH200 GPU. The discretization uses Q^{10} SEM with*
 507 *499^3 DoFs, complex128 arithmetic, and $T = 0.1$. The initial condition is $\psi_0(\mathbf{x}) =$*

508 $\prod_{k=1}^3 \sin(\pi(x_k+L)/2L)$. The observed convergence is $O(\Delta t^2)$ for all M . Relative to
 509 Strang, the error constant is smaller by ~ 5 ($M=3$), ~ 13 ($M=5$), and ~ 24 ($M=7$).
 510 See Remark 5.3 for cost comparison.

TABLE 5
 $qHOP$ splitting error $\|\psi_{\text{split}}(T) - \psi_{\text{exact}}(T)\|_2$ on GH200, Q^{10} SEM, 499³ DoFs, complex128,
 $T=0.1$. $M=1$ is the Strang splitting. t (s) is the GPU time.

Δt	$M = 1$ (Strang)			$M = 3$			$M = 5$			$M = 7$		
	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)
0.1	5.85e-1	—	0.1	1.07e-1	—	0.4	4.25e-2	—	0.5	2.28e-2	—	0.7
0.01	5.22e-3	2.05	0.9	1.02e-3	2.02	3.0	4.11e-4	2.01	4.8	2.21e-4	2.01	6.7
0.005	1.30e-3	2.01	1.9	2.55e-4	2.00	5.8	1.03e-4	2.00	9.6	5.52e-5	2.00	13
0.001	5.21e-5	2.00	9.5	1.02e-5	2.00	29	4.11e-6	2.00	48	2.21e-6	2.00	67

511 **5.3. Non-separable potential: manufactured exact solution.** When $V =$
 512 $V_1 + V_2$ is not fully separable, we use the PCG solver from Section 3.2 to compute the
 513 ground state (λ_1, u_1) of $H = -\Delta + V_1 + V_2$. The function

$$514 \quad (5.5) \quad \psi_{\text{exact}}(\mathbf{x}, t) = e^{-i\lambda_1 t} u_1(\mathbf{x})$$

515 is a stationary solution of (5.1), against which we verify splitting convergence orders.
 516 We test with the stirrer potential (3.2) on $[-8, 8]^3$, Dirichlet BC, Q^{20} SEM, and we
 517 also consider splitting $H = A + B$ with $A = -\Delta + V_1$ and $B = V_2$.

518 **5.3.1. Multi-level generation of the reference solution.** The ground state
 519 (λ_1, u_1) of $H = A + B$ is computed by shifted inverse iteration: each step solves
 520 $(H - \sigma I)w = u$ via PCG with the tensor-product preconditioner $A^{-1} = (-\Delta + V_1)^{-1}$
 521 from Section 3.2. The shift $\sigma = 0.9 \lambda_{\min}(A)$ keeps $H - \sigma I$ positive definite (since
 522 $V_2 \geq 0$) and close enough to λ_1 for rapid convergence. The initial guess is the ground
 523 state eigenfunction of A , obtained from the tensor-product eigendecomposition.

524 To efficiently reach large grid sizes, we use a multi-level strategy: solve on a
 525 coarse grid, interpolate the eigenvector to the next finer grid via tensor-product linear
 526 interpolation, and continue the inverse iteration. Each refinement inherits a good
 527 initial guess, reducing both the number of inverse iterations and the PCG iterations
 528 per step. Table 6 demonstrates this for the stirrer potential (3.2) on an A100 GPU. At
 529 the finest level (499³ DoFs), only 11 inverse iterations with 21 PCG iterations each are
 530 needed, a total of 231 applications of $(-\Delta + V_1)^{-1}$. The entire multi-level computation
 531 takes about one minute, producing a ground state accurate to $|\Delta\lambda|/|\lambda| < 10^{-13}$.

TABLE 6
 Multi-level shifted inverse iteration for the ground state of $H = -\Delta + V_1 + V_2$, stirrer po-
 tential (3.2), Q^{20} SEM, FP64, A100. Preconditioner: $(-\Delta + V_1)^{-1}$, PCG tolerance 10^{-12} , shift
 $\sigma = 0.9 \lambda_{\min}(A)$. Ground state eigenvalue $\lambda_1 = 5.286155366963$.

n (DoFs)	Setup (s)	Interpolation (s)	Inverse iters	PCG/iter	Total A^{-1}
99 ³	6.1	—	30	75	2242
199 ³	2.6	0.3	21	24	504
499 ³	2.6	0.5	11	21	231

Initial guess at 99³: ground state of $A = -\Delta + V_1$. **Total wall time: ~ 62 s on A100.**

5.3.2. Testing splitting methods for Hamiltonian simulations.

EXAMPLE 5.2 (Order of approximated qHOP for non-separable stirrer potential).
 We test qHOP with $M = 1$ (Strang), 3, 5, and 7, 499³ DoFs, complex128, $T = 0.1$, on a GH200 GPU. The initial condition is $\psi_0 = u_1$ (ground state eigenvector, $\lambda_1 = 5.286155366963$). Table 7 compares two splittings: $A = -\Delta$, $B = V_1 + V_2$ (top half), and $A = -\Delta + V_1$, $B = V_2$ (bottom half), where A includes the separable part of the potential. Both show $O(\Delta t^2)$ convergence for all M , but the splitting $A = -\Delta + V_1$ gives $\sim 1.5\text{--}3\times$ smaller errors because $\|V_2\| \ll \|V_1 + V_2\|$. For qHOP as a quantum algorithm, $A = -\Delta$ is required since A must be the operator that can be fast-forwarded on a quantum computer. For classical ODE/PDE solvers, including the separable potential in A is advantageous.

TABLE 7
 qHOP splitting error for the non-separable stirrer potential (3.2), Q²⁰ SEM, 499³ DoFs, complex128, $T=0.1$, GH200. t (s) is the GPU time. Reference: manufactured solution (5.5).

Δt	$M = 1$ (Strang)			$M = 3$			$M = 5$			$M = 7$		
	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)
<i>Splitting $A = -\Delta$, $B = V_1 + V_2$</i>												
0.1	8.04e-3	–	0.7	3.93e-4	–	1.0	1.58e-4	–	1.1	8.49e-5	–	1.4
0.01	7.73e-5	2.02	1.6	3.90e-6	2.00	3.5	1.57e-6	2.00	5.6	8.44e-7	2.00	7.8
0.005	1.93e-5	2.00	2.6	9.74e-7	2.00	6.5	3.93e-7	2.00	11	2.11e-7	2.00	15
0.001	7.73e-7	2.00	11	3.94e-8	1.99	30	1.64e-8	1.97	51	9.46e-9	1.93	72
<i>Splitting $A = -\Delta + V_1$, $B = V_2$</i>												
0.1	4.75e-3	–	0.9	1.37e-4	–	1.0	5.50e-5	–	1.3	2.96e-5	–	1.5
0.01	4.52e-5	2.02	1.7	1.35e-6	2.01	3.7	5.47e-7	2.00	5.6	2.95e-7	2.00	7.7
0.005	1.13e-5	2.00	2.6	3.40e-7	1.99	6.6	1.38e-7	1.99	10	7.50e-8	1.98	15
0.001	4.52e-7	2.00	10	1.61e-8	1.90	30	8.99e-9	1.70	49	7.19e-9	1.46	71

EXAMPLE 5.3 (Order of approximated Magnus-2 for non-separable stirrer potential). We use the same setup, but with $T = 1$. Table 8 compares the two splittings as in Table 7. Both show $O(\Delta t^4)$ convergence for $M \geq 3$. The splitting $A = -\Delta + V_1$, $B = V_2$ again gives smaller errors, with the advantage growing with M .

6. Multi-Body Hamiltonian Simulation (4D, 6D, 9D). In this section we consider multi-particle systems with pairwise interactions in 4D, 6D, and 9D. As in Section 5.3, the reference solution is $\psi_{\text{exact}} = e^{-i\lambda_1 t} u_1$, with the ground state computed by PCG using $(-\Delta + V_{\text{trap}})^{-1}$ as preconditioner.

6.1. Two particles in 2D with Coulomb interaction (4D). We consider two quantum particles in a 2D harmonic trap with Coulomb interaction. The Coulomb singularity $1/|\mathbf{x}_1 - \mathbf{x}_2|$ at $\mathbf{x}_1 = \mathbf{x}_2$ needs be regularized unless special transformation is used. One approach is the cell-averaged potential of [15], which integrates $1/|\mathbf{x}_1 - \mathbf{x}_2|$ against the basis functions. We use the simpler soft Coulomb potential $c/\sqrt{|\mathbf{x}_1 - \mathbf{x}_2|^2 + \delta^2}$ with a small parameter $\delta > 0$, which is smooth for any $\delta > 0$, so we expect $O(\Delta t)$ for Lie–Trotter and $O(\Delta t^2)$ for Strang, although Lie–Trotter is only $O(\Delta t^{1/4})$ for the exact (unregularized) Coulomb potential [13]. The equation is

$$(6.1) \quad i \partial_t \psi = (-\Delta_{\mathbf{x}_1} - \Delta_{\mathbf{x}_2} + V_{\text{trap}}(\mathbf{x}_1) + V_{\text{trap}}(\mathbf{x}_2) + \frac{c}{\sqrt{|\mathbf{x}_1 - \mathbf{x}_2|^2 + \delta^2}}) \psi, \quad \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^2,$$

TABLE 8

Magnus-2 via Yoshida splitting error for the non-separable stirrer potential (3.2), Q^{20} SEM, 499³ DoFs, complex128, $T=1$, GH200. t (s) is the GPU time. Top: $A = -\Delta$, $B = V_1 + V_2$. Bottom: $A = -\Delta + V_1$, $B = V_2$, which gives $\sim 1.5\text{--}3\times$ smaller errors. Reference: manufactured solution via multi-level inverse iteration.

Δt	$M = 1$ (Yoshida)			$M = 3$			$M = 5$			$M = 7$		
	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)
<i>Splitting $A = -\Delta$, $B = V_1 + V_2$</i>												
0.2	8.63e-2	–	2.2	5.36e-3	–	5.1	4.61e-4	–	8.2	1.02e-4	–	11
0.125	1.30e-2	4.03	3.1	5.44e-4	4.87	7.7	2.93e-5	5.86	13	1.32e-5	4.35	18
0.1	5.34e-3	3.99	3.7	1.33e-4	6.31	9.5	1.10e-5	4.39	16	5.02e-6	4.33	22
0.05	3.39e-4	3.98	6.9	3.25e-6	5.35	19	6.77e-7	4.02	30	3.00e-7	4.06	44
<i>Splitting $A = -\Delta + V_1$, $B = V_2$</i>												
0.2	2.48e-2	–	2.2	5.06e-3	–	5.1	4.21e-4	–	8.1	7.04e-5	–	11
0.125	3.75e-3	4.02	3.1	5.15e-4	4.86	7.8	1.46e-5	7.15	13	7.29e-6	4.82	17
0.1	1.50e-3	4.11	3.8	1.20e-4	6.53	9.5	4.18e-6	5.60	16	2.37e-6	5.04	21
0.05	1.04e-4	3.85	6.5	8.50e-7	7.14	18	2.37e-7	4.14	30	1.29e-7	4.20	42

560 where $V_{\text{trap}}(\mathbf{x}) = |\mathbf{x}|^2$ is a harmonic trap.

561 **6.1.1. Preconditioner comparison.** To select the best preconditioner for PCG, █
562 we solve $(-\Delta + V_1 + V_2)u = f$ with a random right-hand side f on the 49^4 grid
563 with $\delta = 0.01$, comparing three preconditioners in Table 9. The tensor-product pre-
564 preconditioner $(-\Delta + V_1)^{-1}$ converges in 24 iterations, far fewer than the combined
565 preconditioner $P_C = (V_1 + V_2)^{-1/2}(-\Delta)^{-1}(V_1 + V_2)^{-1/2}$ (443 iterations). The variant
566 $V_2^{-1/2}(-\Delta + V_1)^{-1}V_2^{-1/2}$ does not converge within 500 iterations. We therefore use
567 $(-\Delta + V_1)^{-1}$ throughout.

TABLE 9

PCG iteration count for solving $(-\Delta + V_1 + V_2)u = f$ with three preconditioners. Q^{10} SEM, 49⁴ DoFs, $\delta = 0.01$, FP64, A100. Tolerance 10^{-12} , max 500 iterations.

Preconditioner	Iters	Time (s)	Status
$(-\Delta + V_1)^{-1}$	24	0.5	converged
$(V_1 + V_2)^{-1/2}(-\Delta)^{-1}(V_1 + V_2)^{-1/2}$	443	2.9	converged
$V_2^{-1/2}(-\Delta + V_1)^{-1}V_2^{-1/2}$	500	3.0	stalled at 5×10^{-10}

568 **6.1.2. Ground state computation.** Table 10 shows the ground state computa-
569 tion cost for several values of δ , with $c = 1$, $V_{\text{trap}}(\mathbf{x}) = |\mathbf{x}|^2$, domain $[-8, 8]^2 \times [-8, 8]^2$.
570 PCG is warm-started with the previous inverse iteration solution, which reduces iter-
571 ations after the first few steps.

572 Figure 5 shows slices of the 4D ground state wavefunction for $\delta = 0.01$ and
573 0.0001, interpolated onto a fine grid via cell-by-cell Q^{10} polynomial reconstruction.
574 As δ decreases, the Coulomb singularity sharpens: the 1D slice narrows, and the
575 inter-particle correlation plot develops a more pronounced dip along $x_{1a} = x_{2a}$.

576 **6.1.3. Splitting convergence.** As in Section 5.2, “qHOP” and “Magnus-2”
577 refer to the approximated versions (5.2) and (5.3). All 4D splitting tests use Q^{10}
578 SEM, 99⁴ DoFs, complex128, $T = 0.1$, GH200.

TABLE 10

Ground state of $H = -\Delta + V_{\text{trap}} + V_{\text{Coulomb}}$ via shifted inverse iteration with PCG (preconditioner $(-\Delta + V_{\text{trap}})^{-1}$, max 500 iterations, warm-started with the previous iteration's solution). Q^{10} SEM, $99^4 \approx 9.6 \times 10^7$ DoFs, GH200. Multi-level: coarse 49^4 solution interpolated as initial guess; shift $\sigma = \lambda_{\min}(-\Delta + V_{\text{trap}}) - 10^{-4} = 3.9999$. PCG tolerance 10^{-9} .

δ	V_C^{\max}	Inv. iters	PCG/iter	Total solves	Time (s)	λ_1
0.1	10	22	1–18	121	8	5.060514417326
0.01	100	24	1–32	204	12	5.221829172892
0.001	1,000	12	1–500 [†]	4,286	208	5.400105827385
0.0001	10,000	12	1–500 [†]	5,816	274	5.444872472429

[†] PCG reaches max 500 iterations only in the first few inverse iterations, which does not affect the final accuracy since PCG converges quickly in the last few inverse iterations.

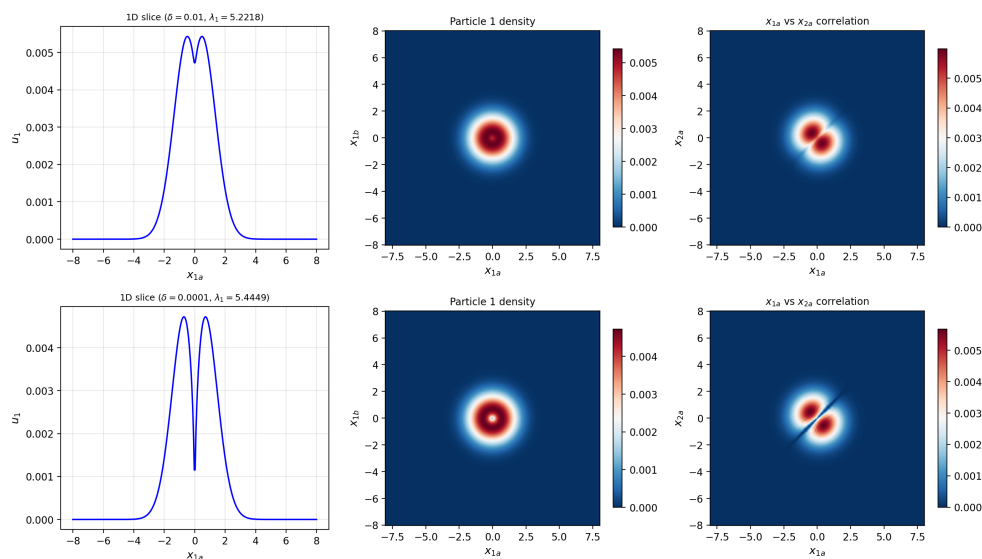


FIG. 5. Slices of the 4D Coulomb ground state for $\delta = 0.01$ (top) and $\delta = 0.0001$ (bottom), Q^{10} SEM, 99^4 DoFs. Top: 1D slice along x_{1a} . Middle: particle 1 density in the (x_{1a}, x_{1b}) plane. Bottom: x_{1a} vs. x_{2a} inter-particle correlation. The Coulomb repulsion strengthens as $\delta \rightarrow 0$. Plotted on a 300×300 fine grid via cell-by-cell Q^{10} polynomial interpolation.

579 EXAMPLE 6.1 (Order of approximated qHOP convergence in 4D, $\delta = 0.001$).
580 Table 11 reports qHOP with $M = 1$ (Strang), 3, 5, 7 for the Coulomb system with
581 $\delta = 0.001$ ($V_C^{\max} = 1000$). All methods show $O(\Delta t^2)$ convergence in the asymptotic
582 regime ($\Delta t \leq 0.001$). Larger M reduces the error constant: at $\Delta t = 0.0005$, $M=3, 5,$
583 7 reduce the error by $\sim 14\times, 35\times, 65\times$ relative to Strang.

584 EXAMPLE 6.2 (Order of approximated Magnus-2 in 4D, $\delta = 0.1$). Table 12
585 shows the convergence of Magnus-2 via Yoshida with $M = 1$ (plain Yoshida), 3, 5,
586 7 for the Coulomb system with $\delta = 0.1$ ($V_C^{\max} = 10$). For $M = 7$, the error reaches
587 2.5×10^{-10} at $\Delta t = 0.001$. Larger M reduces the error constant and enters the
588 high-order asymptotic regime at larger Δt .

TABLE 11

qHOP splitting error $\|\psi_{\text{split}}(T) - \psi_{\text{exact}}(T)\|_2$ on GH200, Q^{10} SEM, 99^4 DoFs, complex128, $T=0.1$, Coulomb $\delta = 0.001$. t (s) is the GPU time.

Δt	$M = 1$ (Strang)			$M = 3$			$M = 5$			$M = 7$		
	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)
0.005	3.32e-1	–	1.8	4.93e-2	–	4.2	4.27e-2	–	5.9	2.96e-2	–	8.0
0.002	5.21e-2	2.02	3.3	2.51e-2	0.74	8.5	1.26e-3	3.84	13	5.18e-4	4.42	18
0.001	3.75e-2	0.47	5.9	1.33e-3	4.24	16	2.34e-4	2.43	26	1.25e-4	2.05	36
0.0005	1.98e-3	4.24	11	1.41e-4	3.24	31	5.64e-5	2.05	51	3.03e-5	2.04	71

TABLE 12

Magnus-2 via Yoshida splitting error, 99^4 DoFs, complex128, $T=0.1$, Coulomb $\delta = 0.1$. t (s) is the GPU time on GH200.

Δt	$M = 1$ (Yoshida)			$M = 3$			$M = 5$			$M = 7$		
	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)
0.005	2.54e-3	–	3.7	7.49e-4	–	10	2.83e-4	–	16	1.51e-4	–	22
0.002	5.82e-5	4.12	8.4	2.38e-5	3.76	24	5.36e-6	4.33	39	5.52e-7	6.12	53
0.001	8.39e-6	2.79	16	1.61e-6	3.89	46	3.74e-8	7.16	77	2.54e-10	11.09	106

589 **6.2. Two particles in 3D with Coulomb interaction (6D).** We extend the
590 4D test to two particles in 3D ($\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^3$). The Hamiltonian is

$$591 \quad (6.2) \quad H = -\Delta_{\mathbf{x}_1} - \Delta_{\mathbf{x}_2} + V_{\text{trap}}(\mathbf{x}_1) + V_{\text{trap}}(\mathbf{x}_2) + \frac{c}{\sqrt{|\mathbf{x}_1 - \mathbf{x}_2|^2 + \delta^2}},$$

592 discretized on $[-L, L]^3 \times [-L, L]^3$ with Q^{10} SEM, $29^6 \approx 5.9 \times 10^8$ DoFs. The splitting
593 is $A = -\Delta$ (6D Laplacian), $B = V_{\text{trap}} + V_{\text{Coulomb}}$, and the reference solution is
594 $\psi_{\text{exact}}(\mathbf{x}, t) = e^{-i\lambda_1 t} u_1(\mathbf{x})$ where (λ_1, u_1) is the ground state of H .

595 **6.2.1. Ground state.** In these TF32 runs, the offline eigendecomposition of
596 $-\Delta + V_{\text{trap}}$ is still computed in FP64 and then cast to lower precision for PCG solves.
597 See Table 13 for the performance. The solution is plotted in Figure 6.

TABLE 13

Ground state of the 6D Coulomb Hamiltonian (6.2), Q^{16} SEM, $N_{\text{cell}} = 2$, $n = 31$, $31^6 \approx 8.9 \times 10^8$ DoFs, TF32, $L = 5$, GH200. Shifted inverse iteration with $\sigma = 0$, PCG tolerance 10^{-4} .

δ	V_C^{max}	Inv. iters	Total PCG solves	Time (s)	λ_1
0.01	100	12	31	17	6.792689
0.001	1000	12	99	34	6.844492
0.0001	10000	12	331	94	6.861799

598 **6.2.2. Splitting convergence.**

599 **EXAMPLE 6.3** (Order of approximated qHOP in 6D, $\delta = 0.01$). *Table 14 reports*
600 *qHOP with $M = 1$ (Strang), 3, 5, 7 for the 6D Coulomb system with $\delta = 0.01$.*
601 *All methods show $O(\Delta t^2)$ convergence. Larger M reduces the error constant by $\sim 7 \times$*
602 *($M=3$), $15 \times$ ($M=5$), $18 \times$ ($M=7$) relative to Strang. The splitting tests use complex64*
603 *arithmetic to fit the 29^6 arrays in GPU memory, and the reference ground state is*
604 *computed in FP64.*

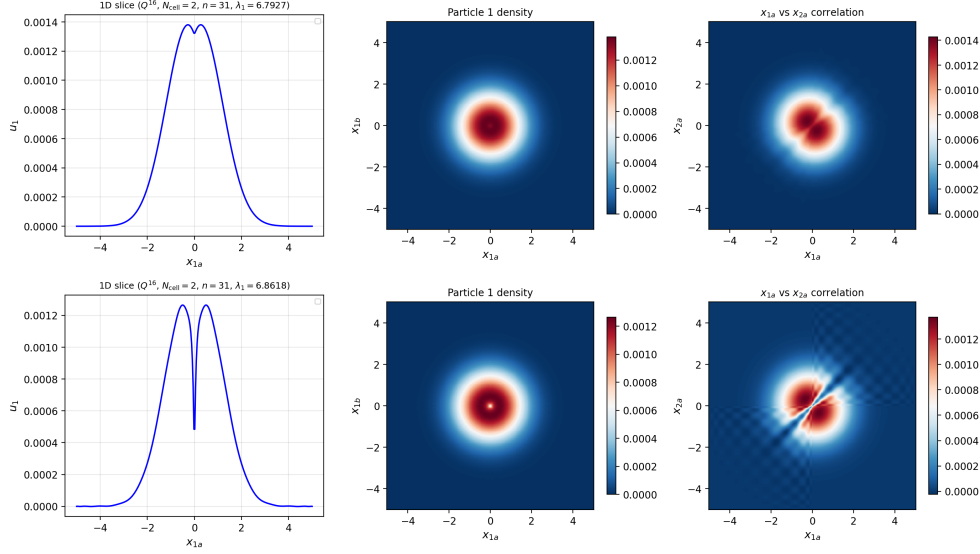


FIG. 6. Slices of the 6D Coulomb ground state at $31^6 \approx 8.9 \times 10^8$ DoFs, Q^{16} SEM, $N_{\text{cell}} = 2$, TF32, $L = 5$, GH200. Top: $\delta = 0.01$. Bottom: $\delta = 0.0001$. Left: 1D slice along x_{1a} . Center: particle 1 density. Right: inter-particle correlation x_{1a} vs. x_{2a} . Plotted on a 300×300 fine grid via cell-by-cell Q^{16} polynomial interpolation.

TABLE 14

qHOP splitting error for the 6D Coulomb system (6.2), $\delta = 0.01$, Q^{10} SEM, 29^6 DoFs, complex64, $T=0.1$, GH200. t (s) is the GPU time.

Δt	$M = 1$ (Strang)			$M = 3$			$M = 5$			$M = 7$		
	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)
0.1	1.98e-1	–	1.5	1.29e-1	–	2.2	1.75e-2	–	3.5	8.34e-3	–	3.6
0.02	9.76e-3	1.87	2.9	1.22e-3	2.90	5.8	4.87e-4	2.23	8.7	2.61e-4	2.15	12
0.005	5.51e-4	2.07	8.0	7.43e-5	2.02	19	3.00e-5	2.01	31	1.62e-5	2.01	43

605 EXAMPLE 6.4 (Order of approximated Magnus-2 in 6D, $\delta = 0.01$). Table 15
606 reports Magnus-2 via Yoshida with $M = 1, 3, 5, 7$. At large Δt , larger M reduces the
607 error: at $\Delta t = 0.01$, the errors range from 7.93×10^{-4} ($M=1$) to 3.28×10^{-6} ($M=7$).
608 At smaller Δt , the errors plateau at $\sim 2.5 \times 10^{-6}$ due to the complex64 arithmetic
609 floor: with 20–50 time steps, single-precision rounding errors accumulate to $O(10^{-6})$.

610 **6.3. Three particles in 3D with Coulomb interaction (9D).** We consider
611 three particles in 3D ($\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3 \in \mathbb{R}^3$) in a harmonic trap with pairwise soft Coulomb
612 interaction:

$$613 \quad (6.3) \quad H = \sum_{j=1}^3 (-\Delta_{\mathbf{x}_j} + V_{\text{trap}}(\mathbf{x}_j)) + \sum_{1 \leq j < k \leq 3} \frac{c}{\sqrt{|\mathbf{x}_j - \mathbf{x}_k|^2 + \delta^2}},$$

614 where $V_{\text{trap}}(\mathbf{x}) = |\mathbf{x}|^2$. Discretized on $[-L, L]^3 \times [-L, L]^3 \times [-L, L]^3$ with Q^k SEM,
615 this gives n^9 DoFs.

TABLE 15

Magnus-2 via Yoshida splitting error for the 6D Coulomb system (6.2), $\delta = 0.01$, Q^{10} SEM, 29^6 DoFs, complex64, $T=0.1$, GH200. t (s) is the GPU time.

Δt	$M = 1$ (Yoshida)			$M = 3$			$M = 5$			$M = 7$		
	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)
0.1	2.43e-1	–	2.3	1.78e-1	–	4.4	2.52e-1	–	6.0	2.72e-2	–	7.7
0.01	7.93e-4	2.49	12	1.55e-5	4.06	28	3.29e-6	4.88	46	1.99e-6	4.14	65
0.005	5.22e-5	3.93	22	2.34e-6 [†]	2.73	55	1.66e-6 [†]	0.99	91	2.77e-6 [†]	–0.48	128

[†]Error dominated by complex64 rounding accumulation.

616 **6.3.1. Ground state.** The offline eigendecomposition used in the separable pre-
617 conditioner is computed in FP64, even though the online inverse iteration and PCG
618 solves are carried out in FP32. The ground state is computed by shifted inverse
619 iteration with $\sigma = 0$ and warm-start PCG (preconditioner $(\sum_j(-\Delta_{\mathbf{x}_j} + V_{\text{trap}}))^{-1}$,
620 tolerance 10^{-4} with max 500 iterations). Warm-start uses the previous outer iteration's
621 PCG solution as the initial guess for the next solve. Table 16 shows the ground
622 state cost for $L = 3$, $c = 1$, Q^5 SEM, $N_{\text{cell}} = 2$, $n = 9$ ($9^9 \approx 3.9 \times 10^8$ DoFs) on
623 GH200. PCG iterations increase with smaller δ , but all cases converge within 18
624 inverse iterations.

TABLE 16

Ground state of the 9D Coulomb Hamiltonian (6.3), 3 particles in 3D, Q^5 SEM, $N_{\text{cell}} = 2$, $9^9 \approx 3.9 \times 10^8$ DoFs, FP32, $\sigma = 0$, $L = 3$, GH200.

δ	V_C^{max}	Inv. iters	PCG solves	Time (s)	λ_1
0.01	100	18	110	490	11.7561
0.001	1000	18	349	588	11.9268
0.0001	10000	18	708	714	11.9502

625 Figure 7 shows slices of the ground state for two values of δ . As δ decreases,
626 the Coulomb repulsion sharpens and the inter-particle anti-correlation pattern along
627 $x_{1a} = x_{2a}$ becomes more pronounced.

628 **6.3.2. Splitting convergence.** The splitting $A = \sum_j(-\Delta_{\mathbf{x}_j} + V_{\text{trap}})$, $B =$
629 $V_{12} + V_{13} + V_{23}$ is tested with $\delta = 0.1$, $T = 0.1$, complex64 arrays, using Q^5 SEM with
630 $N_{\text{cell}} = 2$ ($n = 9$, $9^9 \approx 3.9 \times 10^8$ DoFs) on GH200. Errors of qHOP and Magnus-2 via
631 Yoshida for $M = 1, 3, 5, 7$ are shown in Figure 8.

632 **7. Concluding remarks.** We extended the tensor-product solver of [28] from
633 the Laplacian to $-\Delta + V$ on a single GPU, handling up to 10^9 DoFs in under one
634 second. For non-separable potentials $V = V_1 + V_2$, the preconditioner $(-\Delta + V_1)^{-1}$
635 yields a preconditioned operator with a bounded condition number and spectrum
636 clustering near 1, independently of the mesh size (Theorem 3.4), and also indepen-
637 dently of the domain size when V_1 is confining and V_2 is bounded (Theorem 3.5).
638 We applied this framework to ground state computation via shifted inverse iteration,
639 PCG, and Gross–Pitaevskii gradient flows, and used the resulting eigenpairs as exact
640 stationary solutions to validate the approximated versions of qHOP and Magnus-2
641 splitting methods at $\sim 10^8$ DoFs in 3D through 9D. The product-formula implemen-
642 tation showed the expected convergence orders and also provides a family of classical




FIG. 7. Slices of the 9D Coulomb ground state, 3 particles in 3D, Q^5 SEM, 9^9 DoFs, FP32, $L = 3$, GH200. Top: $\delta = 0.01$. Bottom: $\delta = 0.0001$. Left: 1D slice along x_{1a} . Center: particle 1 density in the (x_{1a}, x_{1b}) plane. Right: x_{1a} vs. x_{2a} inter-particle correlation. Plotted on a 300×300 fine grid via cell-by-cell Q^5 polynomial interpolation.

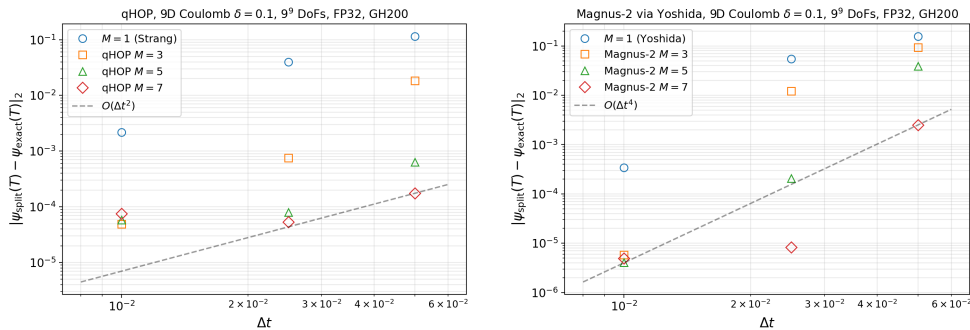


FIG. 8. Splitting convergence for 9D Coulomb, $\delta = 0.1$, $9^9 \approx 3.9 \times 10^8$ DoFs, complex64, FP64 transforms, $T = 0.1$, GH200. Left: qHOP ($O(\Delta t^2)$). Right: Magnus-2 via Yoshida ($O(\Delta t^4)$). Larger M reduces the error constant, and all methods hit the complex64 floor at $\sim 5 \times 10^{-6}$.

643 splitting methods generalizing second order Strang and fourth order Yoshida split-
644 tings.

645 **Acknowledgments.** The authors are grateful to Prof. Antoine Levitt at Labo-
646 ratoire de Mathématiques d’Orsay, Université Paris-Saclay and Prof. Jianfeng Lu at
647 Duke University for discussions on the Schrödinger operator. This work was partially
648 supported by NSF DMS-2208518. This work used DeltaAI at the National Center for
649 Supercomputing Applications (NCSA) through allocation MTH260013 from the Ad-
650 vanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS)
651 program, which is supported by U.S. National Science Foundation grants #2138259,
652 #2138286, #2138307, #2137603, and #2138296.

653 **Declaration of AI-assisted technologies.** All original ideas, as well as the
 654 direct solver and PCG codes, are attributed to the authors. During the preparation
 655 of this work, the authors used Anthropic’s Claude Code to assist with mathematical
 656 discussions, further numerical implementation, and drafting of the manuscript. After
 657 using this tool, the authors carefully reviewed and edited the content as needed and
 658 take full responsibility for the final publication.

659 **Appendix A. Hermite spectral method on unbounded domains.**

660 The tensor-product solver of Section 2 applies to *any* discretization that produces
 661 a separable 1D operator $K = T\Lambda T^{-1}$ per axis. To illustrate this generality, we present
 662 the Hermite spectral method, which replaces the SEM on $[-L, L]^d$ with Hermite
 663 functions on \mathbb{R}^d , eliminating domain truncation artifacts for quantum systems in
 664 confining potentials. The n Hermite-Gauss nodes $\{x_j\}$ are eigenvalues of the $n \times n$
 665 symmetric tridiagonal matrix J with $J_{k,k} = 0$ and $J_{k,k+1} = J_{k+1,k} = \sqrt{k/2}$ for
 666 $k = 1, \dots, n-1$. The Hermite function differentiation matrix at these nodes can be
 667 written as

668
$$D_{ij} = \frac{\psi_{n-1}(x_i)}{\psi_{n-1}(x_j)(x_i - x_j)}, \quad i \neq j, \quad D_{ii} = 0,$$

669 where ψ_{n-1} is the normalized Hermite function of degree $n-1$. See [34, Eq. (7.93)].
 670 Given a separable potential $f(x)$, the 1D operator is $K = -D^2 + \text{diag}(f(x_j))$, with
 671 eigendecomposition $K = T\Lambda T^{-1}$. The d -dimensional tensor-product solver is identical
 672 to SEM. No boundary conditions are needed, since the Hermite functions decay as
 673 $e^{-x^2/2}$. The method is limited to $n \lesssim 745$ in FP64 due to underflow in the Hermite
 674 function recurrence at the outermost nodes. Table 17 shows the accuracy of $(-\Delta +$
 675 $V)^{-1}$ on \mathbb{R}^3 with the potential (2.6). Figure 9 shows that the proposed preconditioner
 676 $(-\Delta + V_1)^{-1}$ converges in 5 PCG iterations on \mathbb{R}^3 , while the baseline P_C does not
 677 converge within 500 iterations. Table 18 reports the Magnus-2 convergence using
 678 the Hermite spectral method with separable potential (4.1), 499^3 DoFs, $T = 1$. For
 679 $M \geq 3$, the observed convergence is $O(\Delta t^4)$.

TABLE 17


Accuracy for using $(-\Delta + V)^{-1}$ to solve $(-\Delta + V)u = f$ with potential (2.6) on \mathbb{R}^3 via Hermite spectral method, FP64, A100. The exact solution is $u^* = \frac{\sin(\frac{\pi}{2}(x+1))}{1+x^2} \cdot \frac{\sin(\pi(y+1))}{1+y^2} \cdot \frac{\sin(\frac{3\pi}{2}(z+1))}{1+z^2}$.

n	DoFs	Setup (s)	Solve (s)	Weighted L^2 err	ℓ^2 rel. err
199	199^3	2.7	0.002	1.53×10^{-5}	6.31×10^{-5}
399	399^3	3.2	0.024	4.05×10^{-6}	2.20×10^{-5}
599	599^3	3.2	0.117	1.87×10^{-6}	1.20×10^{-5}

TABLE 18

Approximated Magnus-2, Hermite spectral on \mathbb{R}^3 , 499^3 DoFs, complex128, $T=1$, GH200. t (s) is the GPU time.

Δt	Yoshida ($M = 1$)			Magnus-2 $M = 3$			Magnus-2 $M = 5$			Magnus-2 $M = 7$		
	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)	Error	Rate	t (s)
0.2	1.34e+0	–	1.7	1.41e+0	–	11	3.05e-1	–	16	1.11e-1	–	23
0.125	1.21e+0	0.22	2.7	4.85e-1	2.27	17	3.59e-2	4.55	26	8.85e-3	5.38	36
0.1	1.52e+0	-1.02	3.4	1.11e-1	6.61	21	1.26e-2	4.69	33	3.51e-3	4.14	45
0.05	5.48e-1	1.47	6.7	5.17e-3	4.42	42	7.60e-4	4.05	65	2.22e-4	3.98	91

FIG. 9. PCG convergence with Hermite spectral method on \mathbb{R}^3 . Stirrer potential, FP32, GH200.

680

REFERENCES

- 681 [1] D. AN, D. FANG, AND L. LIN, *Time-dependent Hamiltonian simulation of highly oscillatory*
682 *dynamics and superconvergence for Schrödinger equation*, Quantum, 6 (2022), p. 690.
683 [2] X. ANTOINE AND R. DUBOSCQ, *Modeling and computation of Bose–Einstein condensates: sta-*
684 *tionary states, nucleation, dynamics, stochasticity*, in Nonlinear Optical and Atomic Sys-
685 *tems: At the Interface of Physics and Mathematics*, Springer International Publishing,
686 2015, pp. 49–145.
687 [3] X. ANTOINE, A. LEVITT, AND Q. TANG, *Efficient spectral computation of the stationary states of*
688 *rotating Bose–Einstein condensates by preconditioned nonlinear conjugate gradient meth-*
689 *ods*, J. Comput. Phys., 343 (2017), pp. 92–109.
690 [4] W. BAO AND Q. DU, *Computing the ground state solution of Bose–Einstein condensates by a*
691 *normalized gradient flow*, SIAM J. Sci. Comput., 25 (2004), pp. 1674–1697.
692 [5] W. BAO, S. JIN, AND P. A. MARKOWICH, *On time-splitting spectral approximations for the*
693 *Schrödinger equation in the semiclassical regime*, J. Comput. Phys., 175 (2002), pp. 487–
694 524.
695 [6] T. J. BOERNER, S. DEEMS, T. R. FURLANI, S. L. KNUTH, AND J. TOWNS, *ACCESS: Advanc-*
696 *ing innovation: NSF’s advanced cyberinfrastructure coordination ecosystem: Services &*
697 *support*, in Practice and Experience in Advanced Research Computing 2023: Computing
698 for the Common Good, 2023, pp. 173–176.
699 [7] Y. BORNS-WEIL, D. FANG, AND J. ZHANG, *Discrete superconvergence analysis for quantum*
700 *Magnus algorithms of unbounded Hamiltonian simulation*, Commun. Math. Phys., 407
701 (2026), p. 29.
702 [8] M. CALIARI, F. CASSINI, L. EINKEMMER, A. OSTERMANN, AND F. ZIVCOVICH, *A μ -mode inte-*
703 *grator for solving evolution equations in Kronecker form*, J. Comput. Phys., 455 (2022),
704 p. 110989.
705 [9] Z. CHEN, J. LU, Y. LU, AND X. ZHANG, *Fully discretized Sobolev gradient flow for the Gross–*
706 *Pitaevskii eigenvalue problem*, Math. Comp., 94 (2025), pp. 2723–2760.
707 [10] A. M. CHILDS, J. LENG, T. LI, J.-P. LIU, AND C. ZHANG, *Quantum simulation of real-space*
708 *dynamics*, Quantum, 6 (2022), p. 860.
709 [11] I. DANAILA AND B. PROTAS, *Computation of ground states of the Gross–Pitaevskii functional*
710 *via Riemannian optimization*, SIAM J. Sci. Comput., 39 (2017), pp. B1102–B1129.
711 [12] D. FANG, D. LIU, AND R. SARKAR, *Time-dependent Hamiltonian simulation via Magnus ex-*
712 *pansion: Algorithm and superconvergence*, Commun. Math. Phys., 406 (2025), p. 128.
713 [13] D. FANG, X. WU, AND A. SOFFER, *On the Trotter error in many-body quantum dynamics with*
714 *Coulomb potentials*, Commun. Math. Phys., 407 (2026), p. 83.
715 [14] D. FANG AND J. ZHANG, *Superconvergence of high-order Magnus quantum algorithms*, 2025.
716 arXiv:2509.22897.
717 [15] X. FENG, H. H. S. CHAN, AND D. P. TEW, *Improved grid-based simulation of Coulombic*
718 *dynamics*, 2026. arXiv:2603.02954.
719 [16] I. K. GAINULLIN AND M. A. SONKIN, *High-performance parallel solver for 3D time-dependent*
720 *Schrödinger equation for large-scale nanosystems*, Comput. Phys. Commun., 188 (2015),
721 pp. 68–75.
722 [17] A. HAHN, P. HARTUNG, D. BURGARTH, P. FACCHI, AND K. YUASA, *Lower bounds for the*
723 *Trotter error*, Phys. Rev. A, 111 (2025), p. 022417.

- 724 [18] D. B. HAIDVOGEL AND T. ZANG, *The accurate solution of Poisson's equation by expansion in*
725 *Chebyshev polynomials*, J. Comput. Phys., 30 (1979), pp. 167–180.
- 726 [19] W. HAO, S. LEE, AND X. ZHANG, *An efficient quasi-Newton method with tensor product im-*
727 *plementation for solving quasi-linear elliptic equations and systems*, J. Sci. Comput., 103
728 (2025), p. 89.
- 729 [20] P. HENNING AND D. PETERSEIM, *Sobolev gradient flow for the Gross–Pitaevskii eigenvalue*
730 *problem: Global convergence and computational efficiency*, SIAM J. Numer. Anal., 58
731 (2020), pp. 1744–1772.
- 732 [21] M. HOCHBRUCK AND C. LUBICH, *On Magnus integrators for time-dependent Schrödinger equa-*
733 *tions*, SIAM J. Numer. Anal., 41 (2003), pp. 945–963.
- 734 [22] B. JACKSON, J. F. MCCANN, AND C. S. ADAMS, *Vortex formation in dilute inhomogeneous*
735 *Bose–Einstein condensates*, Phys. Rev. Lett., 80 (1998), pp. 3903–3906.
- 736 [23] T. JAHNKE AND C. LUBICH, *Error bounds for exponential operator splittings*, BIT Numer.
737 Math., 40 (2000), pp. 735–744.
- 738 [24] S. JIN, P. A. MARKOWICH, AND C. SPARBER, *Mathematical and computational methods for*
739 *semiclassical Schrödinger equations*, Acta Numer., 20 (2011), pp. 121–209.
- 740 [25] Y.-Y. KWAN AND J. SHEN, *An efficient direct parallel spectral-element solver for separable*
741 *elliptic problems*, J. Comput. Phys., 225 (2007), pp. 1721–1735.
- 742 [26] C. LASSER AND C. LUBICH, *Computing quantum dynamics in the semiclassical regime*, Acta
743 Numer., 29 (2020), pp. 229–401.
- 744 [27] H. LI, D. APPELÖ, AND X. ZHANG, *Accuracy of spectral element method for wave, parabolic,*
745 *and Schrödinger equations*, SIAM J. Numer. Anal., 60 (2022), pp. 339–363.
- 746 [28] X. LIU, J. SHEN, AND X. ZHANG, *A simple GPU implementation of spectral-element methods for*
747 *solving 3D Poisson type equations on rectangular domains and its applications*, Commun.
748 Comput. Phys., 36 (2024), pp. 1157–1185.
- 749 [29] R. E. LYNCH, J. R. RICE, AND D. H. THOMAS, *Direct solution of partial difference equations*
750 *by tensor product methods*, Numer. Math., 6 (1964), pp. 185–199.
- 751 [30] J. E. OSBORN, *Spectral approximation for compact operators*, Math. Comp., 29 (1975), pp. 712–
752 725.
- 753 [31] M. REED AND B. SIMON, *Methods of Modern Mathematical Physics, Vol. I: Functional Analy-*
754 *sis*, Academic Press, New York, 1972.
- 755 [32] ———, *Methods of Modern Mathematical Physics, Vol. IV: Analysis of Operators*, Academic
756 Press, New York, 1978.
- 757 [33] J. SHEN, *Efficient spectral-Galerkin method I. Direct solvers of second- and fourth-order equa-*
758 *tions using Legendre polynomials*, SIAM J. Sci. Comput., 15 (1994), pp. 1489–1505.
- 759 [34] J. SHEN, T. TANG, AND L.-L. WANG, *Spectral Methods: Algorithms, Analysis and Applications*,
760 vol. 41 of Springer Series in Computational Mathematics, Springer, Heidelberg, 2011.
- 761 [35] Y. SU, H.-Y. HUANG, AND E. T. CAMPBELL, *Nearly tight Trotterization of interacting electrons*,
762 Quantum, 5 (2021), p. 495.
- 763 [36] M. THALHAMMER, *High-order exponential operator splitting methods for time-dependent*
764 *Schrödinger equations*, SIAM J. Numer. Anal., 46 (2008), pp. 2022–2038.
- 765 [37] A. VAN DER SLUIS AND H. A. VAN DER VORST, *The rate of convergence of conjugate gradients*,
766 Numer. Math., 48 (1986), pp. 543–560.
- 767 [38] H. YOSHIDA, *Construction of higher order symplectic integrators*, Phys. Lett. A, 150 (1990),
768 pp. 262–268.