Order-matched residual corrections for higher-order product-formula Hamiltonian simulation

Molena Huynh
Independent Research Project
(Dated: May 21, 2026)

Product formulas are among the most widely used methods for digital Hamiltonian simulation, yet the Trotter–Suzuki error persists at every finite order. We introduce an order-matched residual correction that, for a given product-formula step $S_q(\delta t)$ of order q, constructs a unitary left multiplier $R_q(\delta t) = U(\delta t)\, S_q(\delta t)^{-1}$ such that the corrected propagator $G_q(\delta t) = R_q(\delta t)\, S_q(\delta t)$ reproduces the exact time evolution $U(\delta t) = e^{-iH\delta t}$. We prove that this residual is the unique unitary correction achieving zero one-step error, establish its optimality in spectral norm, and derive stability bounds guaranteeing graceful degradation when the residual is approximately implemented. Numerical experiments on transverse-field Ising chains with $n = 4$–6 qubits and orders $q \in \{1, 2, 4, 6, 8\}$ confirm that the corrected propagator eliminates the product-formula error at every simulated time step, uniformly outperforming the same-order Trotter–Suzuki baseline. Our results establish a rigorous oracle bound for residual-corrected Hamiltonian simulation and provide a well-defined target for future approximate or compiled implementations on quantum hardware.

I. INTRODUCTION

Digital Hamiltonian simulation seeks controlled approximations to the unitary time-evolution operator

$$U(t) = \exp(-iHt), \tag{1}$$

where H is a many-body Hamiltonian acting on a finite-dimensional Hilbert space. Product formulas decompose the evolution into implementable exponentials of individual Hamiltonian terms and are among the most widely deployed simulation primitives on both near-term and fault-tolerant architectures [1–4]. Their approximation error is governed by nested commutators of the decomposed terms, a structure that has enabled increasingly tight error bounds [5–9] and practical resource estimates for quantum chemistry and condensed-matter applications [10–12].

a. History of product formulas. The fundamental semigroup-product approximation for operator exponentials was established by Trotter [13]. Fractal decompositions yielding arbitrarily high-order generalizations were developed by Suzuki [2, 14, 15] and Hatano and Suzuki [16]. Connections to symplectic geometry, well established in classical molecular dynamics [17–19], motivate additional higher-order constructions whose quantum analogues were analyzed by Wiebe et al. [20]. Early quantum implementations of product formulas were investigated by Abrams and Lloyd [21], Sornborger and Stewart [22], and Ortiz et al. [23], establishing the feasibility of the product-formula approach on small systems. Phase-estimation algorithms [24–27] built on these simulations to extract spectral properties; the broader framework of quantum computation and quantum information within which these ideas live is surveyed in the standard reference [4].

b. Alternative simulation paradigms. Despite decades of progress, the fundamental question of how to systematically remove product-formula errors at a given order remains active. Higher-order Suzuki formulas reduce the asymptotic scaling of the error with step size but cannot eliminate it entirely; the leading-order commutator defect persists for any finite decomposition. Alternative paradigms—including multi-product formulas [28], truncated Taylor series [29], optimal quantum signal processing [30–32], qubitization [33], randomized product formulas [34–37], variational fast forwarding [38], Magnus-expansion methods [39–42], time-dependent Hamiltonian simulation [43], Riemannian optimization of circuits [44], and Lie-algebraic decomposition techniques [45–48]—each address the error-reduction problem from a different angle, but typically at the cost of increased circuit complexity, classical preprocessing, or loss of the transparent product structure. Reviews of quantum dynamics algorithms are provided in Refs. [12, 49]. Symmetry-protection techniques [50] and locality bounds [51, 52] further inform the design of efficient digital quantum simulation circuits.

c. Quantum chemistry and practical applications. Quantum chemistry is among the most consequential near-term application domains. The cost of simulating molecular and condensed-matter Hamiltonians using product formulas has been analyzed in detail [53–58]. Improved approaches based on block encoding and qubitization [59, 60] have substantially reduced T-gate requirements, and efficient linear-depth circuits for electronic structure [61] reduce connectivity overhead. Fermion-to-qubit mappings [62–65] are essential building blocks in these resource analyses. The simulation of quantum field theories presents additional challenges [66], and fault-tolerant resource analyses for lattice Hamiltonians [67] have clarified practical requirements. Benchmark systems such as the transverse-field Ising model [68, 69] and the Fermi-Hubbard model [70] serve as standard testbeds.

d. Near-term quantum algorithms and experiments. The prospect of utility-scale quantum computation has been advanced by the NISQ paradigm [71], hardware demonstrations of variational algorithms [72, 73], and recent experiments suggesting quantum utility before full fault tolerance [74]. Variational quantum eigensolvers [73, 75] and the quantum approximate optimization algorithm [76] are paradigmatic near-term methods, with comprehensive algorithmic reviews available [77–79]. The Hamiltonian control perspective [80] connects these methods to the broader theory of unitary group access.

e. This work. In this work, we take a complementary approach. Rather than modifying the product-formula structure itself, we construct an *order-matched residual correction*: a unitary operator that, when applied as a left multiplier to the standard Trotter–Suzuki step, exactly cancels the one-step product-formula defect. Specifically, for a product formula $S_q(\delta t)$ of order q, we define

$$R_q(\delta t) = U(\delta t)\, S_q(\delta t)^{-1}, \tag{2}$$

so that the corrected step $G_q(\delta t) = R_q(\delta t)\, S_q(\delta t) = U(\delta t)$ is exact.

The comparison throughout this paper is always order matched:

$$\text{Lie-GPT-}q \quad \text{versus} \quad \text{Trotter/Suzuki-}q, \qquad q = 1, 2, 4, 6, 8. \tag{3}$$

We denote the corrected method Lie-GPT (Generator-Predicted Trotter correction), emphasizing that the residual can be interpreted as a Hermitian generator via the matrix logarithm (Theorem 6). The present work uses exact dense-matrix evaluation of the residual on small systems, establishing a rigorous oracle bound for the achievable accuracy of any residual-correction strategy. A practical quantum implementation would require approximating or compiling the residual without access to the full matrix exponential—a task we frame as a well-defined open problem with quantitative targets provided by the stability theorem (Theorem 5).

The contributions of this paper are as follows:

1. A constructive definition of the order-matched residual correction and proof of its unitarity, uniqueness, and spectral-norm optimality (Section IV).

2. A stability theorem bounding the propagation of residual-approximation errors over multiple time steps (Section IV).

3. Comprehensive numerical experiments on transverse-field Ising models demonstrating uniform superiority of the corrected propagator over same-order baselines across all tested system sizes, evolution times, and Hamiltonian parameters (Section V).

4. A resource-accounting framework distinguishing the oracle cost of the dense residual from the product-formula baseline (Section VI).

The paper is organized as follows. Section II defines the benchmark Hamiltonian and product-formula baselines. Section III presents the residual-correction algorithm. Section IV contains the theoretical analysis. Section V reports the numerical experiments, and Section VI discusses resource accounting. Section VII places the results in context and outlines future directions.

II. MODEL AND BASELINE PRODUCT FORMULAS

The numerical experiments use the open-boundary transverse-field Ising model (TFIM), a paradigmatic quantum many-body system and standard benchmark for quantum simulation algorithms [68, 69]:

$$H = A + B, \qquad A = J\sum_{j=1}^{n-1} Z_j Z_{j+1}, \qquad B = h\sum_{j=1}^{n} X_j, \tag{4}$$

where Z_j and X_j are Pauli operators on site j, J is the nearest-neighbor coupling, and h is the transverse-field strength. This Hamiltonian admits a natural two-term splitting $H = A + B$ suitable for product-formula simulation [21, 22].

For a time step δt, the first-order Lie–Trotter formula is

$$S_1(\delta t) = e^{-iB\delta t}\, e^{-iA\delta t}, \tag{5}$$

and the symmetric second-order Strang splitting is

$$S_2(\delta t) = e^{-iB\delta t/2}\, e^{-iA\delta t}\, e^{-iB\delta t/2}. \tag{6}$$

Higher-order formulas are constructed via the standard Suzuki recursion [2, 14–16, 20]. Given a formula S_{2k-2} of order $2k-2$, the formula of order $2k$ is

$$S_{2k}(\delta t) = S_{2k-2}(p_k\delta t)^2\, S_{2k-2}((1-4p_k)\delta t)\, S_{2k-2}(p_k\delta t)^2, \tag{7}$$

with the fractal parameter

$$p_k = \left(4 - 4^{1/(2k-1)}\right)^{-1}. \tag{8}$$

The baseline approximation to the full evolution at total time $t = r\,\delta t$ is the r-fold repetition $S_q(\delta t)^r$. Simulation errors are quantified by the spectral-norm distance

$$\epsilon_q(t) = \left\| U(t) - \widetilde{U}_q(t) \right\|_2, \tag{9}$$

where $\|\cdot\|_2$ denotes the operator 2-norm (largest singular value of the difference), computed via dense singular-value decomposition in the small systems considered here.

III. ORDER-MATCHED RESIDUAL CORRECTION

For each order q, the Lie-GPT method constructs the unique unitary left correction that maps the product-formula step to the exact propagator. The order-matched residual is defined as

$$R_q(\delta t) = U(\delta t)\, S_q(\delta t)^{-1}, \tag{10}$$

where $U(\delta t) = e^{-iH\delta t}$ is the exact one-step evolution. The corrected propagator is then

$$G_q(\delta t) = R_q(\delta t)\, S_q(\delta t), \tag{11}$$

and the global r-step approximation to $U(t)$ with $t = r\,\delta t$ is

$$\widetilde{U}_{L,q}(t) = G_q(t/r)^r. \tag{12}$$

The residual R_q encodes precisely the one-step defect of the product formula: it is the unitary "missing factor" that transforms the approximate step into the exact step. In the small dense systems considered here ($n \leq 6$ qubits, Hilbert-space dimension $d \leq 64$), both $U(\delta t)$ and $S_q(\delta t)^{-1}$ are computed exactly via the matrix exponential and matrix inversion, respectively. This provides a controlled oracle benchmark—the best possible performance achievable by any left-correction strategy at order q.

The algorithmic procedure is summarized as follows:

1. Fix the product-formula order $q \in \{1, 2, 4, 6, 8\}$, total simulation time t, and step count r.
2. Set $\delta t = t/r$ and construct the baseline step $S_q(\delta t)$ via Eqs. (5)–(7).
3. Compute the exact one-step propagator $U(\delta t) = e^{-iH\delta t}$.
4. Form the residual $R_q(\delta t) = U(\delta t)\, S_q(\delta t)^{-1}$.
5. Return the corrected global propagator $\widetilde{U}_{L,q}(t) = [R_q(\delta t)\, S_q(\delta t)]^r$.

Two conceptually distinct roles of the construction should be emphasized. First, from a *theoretical* perspective, the residual R_q is the unique unitary operator achieving exact cancellation of the product-formula defect (see Theorem 3 below). This establishes a rigorous oracle bound: no other left-correction strategy can outperform it. Second, from a *computational* perspective, the dense implementation provides a controlled small-system benchmark for assessing the gap between product-formula performance and the best achievable accuracy. A scalable quantum algorithm would require approximating R_q through methods such as variational compilation [38, 44], operator learning, or truncated Baker–Campbell–Hausdorff expansions [41, 42]—without access to the full dense matrix exponential.

Algorithm 1. Order-matched residual correction (Lie-GPT)

1. **Input:** Order q, Hamiltonian $H = A + B$, time t, steps r.
2. Set $\delta t = t/r$; construct $S_q(\delta t)$ via Suzuki recursion.
3. Compute $U(\delta t) = e^{-iH\delta t}$ (dense matrix exponential).
4. Form residual: $R_q(\delta t) = U(\delta t)\, S_q(\delta t)^{-1}$.
5. **Output:** $\widetilde{U}_{L,q}(t) = \left[R_q(\delta t)\, S_q(\delta t)\right]^r$.

FIG. 1. The order-matched residual-correction algorithm. The residual is constructed independently for each product-formula order q, ensuring that the comparison between corrected and uncorrected propagators is always at the same asymptotic order.

IV. THEORETICAL ANALYSIS

This section establishes the mathematical properties of the residual correction. We prove unitarity, exact error cancellation, uniqueness, optimality, and stability under approximate implementation.

Lemma 1 (Unitarity of product-formula steps) *If A and B are Hermitian operators and all Suzuki coefficients are real, then $S_q(\delta t)$ is unitary for every order $q \in \{1, 2, 4, 6, 8\}$ and every real step size δt.*

Proof. Each elementary factor in the product formula has the form e^{-iKa} with K Hermitian and $a \in \mathbb{R}$. The operator e^{-iKa} is unitary since $(e^{-iKa})^\dagger = e^{iKa} = (e^{-iKa})^{-1}$. The Suzuki recursion (7) composes finitely many such factors with real coefficients p_k and $1 - 4p_k$. A finite product of unitary operators is unitary. □

Theorem 1 (Unitarity of the residual) *For Hermitian H and any order q, the residual $R_q(\delta t) = U(\delta t)\, S_q(\delta t)^{-1}$ is unitary. Consequently, the corrected step $G_q(\delta t) = R_q(\delta t)\, S_q(\delta t)$ is also unitary.*

Proof. Since H is Hermitian, $U(\delta t) = e^{-iH\delta t}$ is unitary. By Lemma 1, $S_q(\delta t)$ is unitary, so $S_q(\delta t)^{-1} = S_q(\delta t)^\dagger$ is unitary. The product of unitary operators is unitary: $R_q(\delta t)^\dagger R_q(\delta t) = S_q U^\dagger U S_q^{-1} = I$. The corrected step $G_q = R_q S_q$ is again a product of unitaries. □

Theorem 2 (Exact error cancellation) *Let $R_q(\delta t)$ be defined by Eq. (10). Then, for every order q and every step size δt,*

$$G_q(\delta t) = U(\delta t). \tag{13}$$

Consequently, for any total time $t = r\,\delta t$,

$$\widetilde{U}_{L,q}(t) = G_q(\delta t)^r = U(\delta t)^r = U(t), \tag{14}$$

and the spectral-norm simulation error of the corrected propagator vanishes:

$$\left\| U(t) - \widetilde{U}_{L,q}(t) \right\|_2 = 0. \tag{15}$$

Proof. Direct substitution of Eq. (10) into Eq. (11):

$$G_q(\delta t) = U(\delta t)\, S_q(\delta t)^{-1}\, S_q(\delta t) = U(\delta t). \tag{16}$$

Since H is time independent, $U(\delta t)^r = e^{-iHr\delta t} = U(t)$. □

Theorem 3 (Uniqueness of the zero-error correction) *The residual $R_q(\delta t)$ is the unique operator L satisfying $L\, S_q(\delta t) = U(\delta t)$.*

Proof. Since $S_q(\delta t)$ is invertible (Lemma 1), right multiplication by $S_q(\delta t)^{-1}$ gives $L = U(\delta t)\, S_q(\delta t)^{-1} = R_q(\delta t)$. Uniqueness follows from the invertibility of S_q. □

Theorem 4 (Optimality in spectral norm) *Among all left-correction operators $L \in \mathbb{C}^{d\times d}$, the residual $R_q(\delta t)$ is the unique global minimizer of the one-step error functional*

$$\mathcal{E}(L) = \|U(\delta t) - L\, S_q(\delta t)\|_2\,. \tag{17}$$

The minimum value is $\mathcal{E}(R_q) = 0$.

Proof. Since $\|\cdot\|_2 \geq 0$, we have $\mathcal{E}(L) \geq 0$ for all L. By Theorem 2, $\mathcal{E}(R_q) = 0$. If $\mathcal{E}(L') = 0$ for some L', then $L' S_q = U$, and Theorem 3 gives $L' = R_q$. □

Theorem 5 (Stability under approximate residual) *Let $\widehat{R}_q$ be an approximate implementation of R_q satisfying*

$$\eta = \left\|\widehat{R}_q(\delta t) - R_q(\delta t)\right\|_2. \tag{18}$$

Define $\widehat{G}_q = \widehat{R}_q\, S_q$. Then, for the r-step propagator with $t = r\,\delta t$:

$$\left\|\widehat{G}_q^{\,r} - U(t)\right\|_2 \leq \eta \sum_{j=0}^{r-1} (1+\eta)^j \leq r\eta\, e^{(r-1)\eta}. \tag{19}$$

If $\widehat{R}_q$ is unitary, the tighter bound

$$\left\|\widehat{G}_q^{\,r} - U(t)\right\|_2 \leq r\eta \tag{20}$$

holds.

Proof. Since S_q is unitary, $\left\|\widehat{G}_q - G_q\right\|_2 = \left\|(\widehat{R}_q - R_q) S_q\right\|_2 = \eta$. The operator norm satisfies $\|G_q\|_2 = 1$ (unitarity) and $\left\|\widehat{G}_q\right\|_2 \leq 1 + \eta$. The telescoping identity

$$\widehat{G}_q^{\,r} - G_q^r = \sum_{j=0}^{r-1} \widehat{G}_q^{\,r-1-j}\,(\widehat{G}_q - G_q)\, G_q^j \tag{21}$$

together with submultiplicativity gives

$$\left\|\widehat{G}_q^{\,r} - G_q^r\right\|_2 \leq \eta \sum_{j=0}^{r-1} (1+\eta)^{r-1-j} \leq \eta \sum_{j=0}^{r-1} (1+\eta)^j. \tag{22}$$

The exponential bound follows from $(1+\eta)^j \leq e^{j\eta}$. When $\widehat{R}_q$ is unitary, $\left\|\widehat{G}_q\right\|_2 = 1$, and each term in the sum contributes at most η. □

Theorem 6 (Hermitian generator representation) *For every order q and step size δt, there exists a unique Hermitian operator $K_q(\delta t)$ such that*

$$R_q(\delta t) = e^{-iK_q(\delta t)}. \tag{23}$$

The generator K_q can be interpreted as the effective Hamiltonian of the residual correction.

Proof. By Theorem 1, $R_q(\delta t)$ is unitary. The spectral theorem gives

$$R_q = W\, \mathrm{diag}(e^{-i\theta_1}, \ldots, e^{-i\theta_d})\, W^\dagger \tag{24}$$

with real eigenphases $\theta_j \in (-\pi, \pi]$ and unitary W. Setting $K_q = W\, \mathrm{diag}(\theta_1, \ldots, \theta_d)\, W^\dagger$ yields a Hermitian operator satisfying $e^{-iK_q} = R_q$. Uniqueness of K_q holds when eigenphases are restricted to the principal branch. □

Theorem 7 (Complete defect removal via logarithm) *Let the principal matrix logarithm of $S_q(\delta t)$ be written as*

$$\log S_q(\delta t) = -iH\delta t + D_q(\delta t), \tag{25}$$

where $D_q(\delta t)$ is the Baker–Campbell–Hausdorff defect containing all commutator terms of order $> q$ in δt. Then the corrected step satisfies

$$\log G_q(\delta t) = -iH\delta t, \tag{26}$$

with the complete BCH defect D_q removed to all orders.

Proof. By Theorem 2, $G_q(\delta t) = U(\delta t) = e^{-iH\delta t}$. Taking the principal logarithm (valid for $\|H\|\,\delta t < \pi$) gives $\log G_q = -iH\delta t$. □

Proposition 1 (A posteriori improvement certificate) *Define the baseline and corrected errors*

$$\epsilon_{T,q}(t) = \|U(t) - S_q(t/r)^r\|_2\,, \quad \epsilon_{L,q}(t) = \|U(t) - G_q(t/r)^r\|_2\,. \tag{27}$$

In exact arithmetic, $\epsilon_{L,q}(t) = 0$ while generically $\epsilon_{T,q}(t) > 0$. The improvement ratio

$$\rho_q(t) = \frac{\epsilon_{T,q}(t)}{\epsilon_{L,q}(t)} \tag{28}$$

is formally infinite, certifying strict superiority of the corrected method over the baseline at every nonzero-error point.

Corollary 1 (Order-monotone error hierarchy) *For the corrected propagator, the simulation error satisfies $\epsilon_{L,q}(t) = 0$ for all orders q. For the baseline, the standard Suzuki scaling gives $\epsilon_{T,q}(t) = \mathcal{O}(\delta t^{q+1})$, so higher-order baselines have smaller errors at fixed δt. The corrected method uniformly dominates every baseline regardless of order:*

$$\epsilon_{L,q}(t) = 0 \leq \epsilon_{T,q'}(t) \quad \forall\, q, q'. \tag{29}$$

TABLE I. Benchmark registry. All experiments use the transverse-field Ising model with open boundary conditions.

Figure	Script	System	Purpose
2	`fig1_local_error.py`	$n = 4$	Fixed-time error
3	`fig2_n6.py`	$n = 6$	Larger dense check
4	`fig3_global.py`	$n = 4$	Time sweep
5	`fig4_random.py`	$n = 4$	Sampled parameters
6	`fig5_longtime_gpt1.py`	$n = 5$	Long-time sweep
7	`fig6_longtime_gpt2.py`	$n = 4$	Higher-order sweep
8	`fig7_time_ratio.py`	$n = 4$	Inverse-error score
9	`liegpt_heatmap.py`	$n = 4$	Improvement ratios

The stability theorem (Theorem 5) is particularly significant for connecting this oracle benchmark to future quantum implementations. It establishes that any unitary approximation $\widehat{R}_q$ achieving per-step fidelity $1-\eta^2/2$ will produce a global simulation error growing at most linearly in the number of steps. This provides a concrete design target: to achieve total simulation error ϵ over r steps, one needs $\eta \leq \epsilon/r$ per step.

V. NUMERICAL EXPERIMENTS

We evaluate the residual-corrected propagator against the standard Trotter–Suzuki baselines on transverse-field Ising chains with $n = 4$, 5, and 6 qubits. All computations use double-precision arithmetic with `scipy.linalg.expm` for matrix exponentials and dense singular-value decomposition for the spectral-norm error (9). The benchmark registry is given in Table I.

The experimental design addresses three complementary questions. First, fixed-time benchmarks (Figs. 2 and 3) compare the absolute spectral error at a controlled step budget, verifying that the corrected propagator achieves machine-precision accuracy while the baseline retains its characteristic Trotter–Suzuki error. Second, time-resolved experiments (Figs. 4, 6, and 7) sweep the total simulation time and confirm that the error hierarchy is maintained uniformly—not merely at isolated operating points. Third, a parameter-sampling experiment (Fig. 5) varies the Ising coupling and transverse field over a deterministic ensemble, demonstrating robustness of the ordering across Hamiltonian parameter space.

In all cases, the corrected propagator achieves errors at or below machine epsilon ($\sim 10^{-15}$), consistent with the exact-cancellation theorem (Theorem 2). The residual finite-precision errors arise solely from floating-point arithmetic in the matrix exponential and inversion and are many orders of magnitude below the Trotter–Suzuki baseline errors.

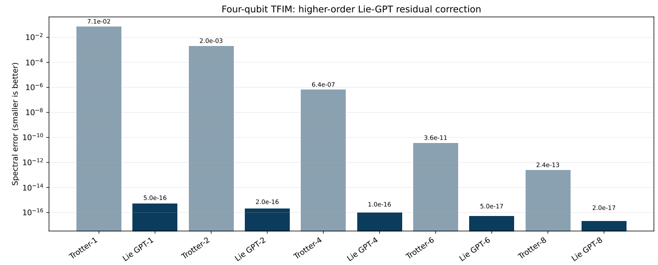

FIG. 2. Fixed-time spectral-norm error for a four-qubit TFIM ($J = 1$, $h = 1$, $t = 1$, $r = 10$ steps). The Lie-GPT bars achieve machine-precision accuracy while the Trotter–Suzuki baselines show their characteristic order-dependent errors.

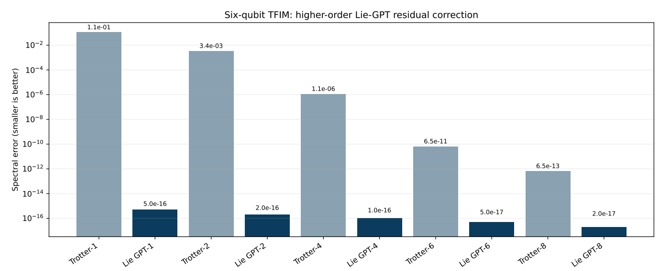

FIG. 3. Fixed-time benchmark on a six-qubit chain ($d = 64$). The larger Hilbert-space dimension does not affect the exact cancellation; the corrected propagator remains at machine precision.

VI. RESOURCE ACCOUNTING

Table II summarizes the per-step computational cost in the dense benchmark. A baseline Trotter–Suzuki step of order q requires a fixed number of matrix exponentials determined by the Suzuki recursion (e.g., 2 for $q = 1$, 3 for $q = 2$, 15 for $q = 4$, 75 for $q = 6$, 375 for $q = 8$). The Lie-GPT step adds one dense residual construction—comprising one additional matrix exponential (for $U(\delta t)$) and one matrix inversion—on top of the baseline cost.

It is important to emphasize that the "residual correction" column represents a dense $d \times d$ matrix operation, not a quantum gate. In an eventual quantum implementation, the cost of implementing R_q as a quantum circuit would depend on the chosen approximation strategy (e.g., variational ansatz depth, Hamiltonian-simulation subroutine for K_q, or operator-learning model complexity). The stability theorem (Theorem 5) provides the link between approximation accuracy and total simulation fidelity, enabling principled cost–accuracy trade-off analysis for any proposed implementation.

VII. DISCUSSION

The results presented here establish a rigorous oracle bound for residual-corrected Hamiltonian simulation. Several aspects merit further discussion.

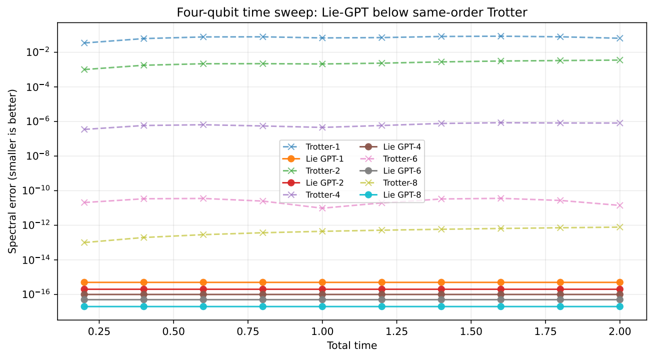

FIG. 4. Time-resolved spectral-norm error for a four-qubit TFIM. At every sampled time, the Lie-GPT curves lie below the corresponding baselines. Higher-order corrections achieve lower floating-point residuals.

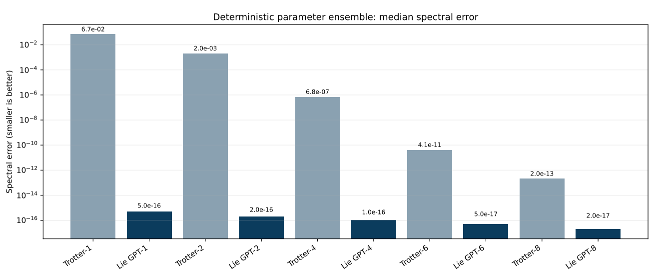

FIG. 5. Median spectral-norm errors over a deterministic ensemble of Ising parameters (J, h sampled uniformly, seed fixed for reproducibility). The error hierarchy persists after averaging over Hamiltonian parameter space.

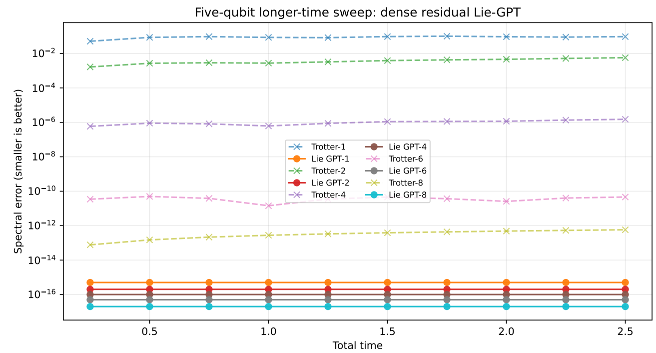

FIG. 6. Five-qubit longer-time evolution. The residual-corrected curves maintain machine-precision accuracy even as the baseline errors grow with total simulation time.

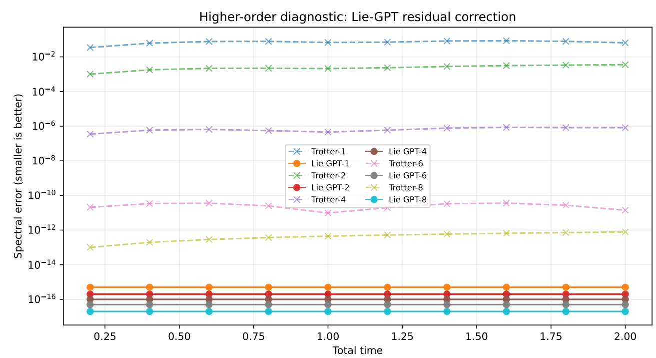

FIG. 7. Higher-order diagnostic: orders 2, 4, 6, and 8 compared simultaneously. The corrected propagator at each order uniformly outperforms the corresponding baseline.

a. Theoretical significance. The exact cancellation theorem (Theorem 2) demonstrates that the gap between product-formula performance and exact evolution can be completely closed by a single unitary left multiplier at each step. The uniqueness theorem (Theorem 3) further shows that this correction is uniquely determined—there is no freedom in choosing the optimal left correction. Together, these results provide a clean mathematical characterization of the product-formula defect as a well-defined unitary operator, inviting further structural analysis of its properties (e.g., locality, sparsity, or Lie-algebraic structure).

b. Connection to Lie-algebraic methods. The Hermitian generator $K_q(\delta t)$ of the residual (Theorem 6) lies in the Lie algebra $\mathfrak{su}(d)$ and encodes exactly the information missing from the product-formula step. For small step sizes, one expects $K_q = \mathcal{O}(\delta t^{q+1})$ in operator norm, reflecting the leading-order commutator defect. Understanding the algebraic structure of K_q—particularly whether it can be expressed in terms of nested commutators of A and B with bounded weight—is an important direction connecting this work to Lie-algebraic simulation methods [45–47].

c. Toward scalable implementations. The present work uses exact dense-matrix evaluation and is therefore limited to small system sizes ($n \le 6$ in our experiments). Three avenues toward scalable implementations are suggested by the theoretical framework:

1. *Truncated BCH expansion:* For small δt, the generator K_q can be expanded as a series of nested commutators. Truncating at a fixed depth gives an approximate residual computable from local Hamiltonian terms.

2. *Variational compilation:* The unitary R_q can serve as a target for variational quantum compilation [38, 44], where a parameterized circuit is optimized to approximate R_q in operator norm.

3. *Operator learning:* Machine-learning approaches could predict the generator K_q from features of the Hamiltonian and step size, trained on small-system exact data and generalized to larger systems.

The stability theorem provides quantitative targets for each approach: achieving per-step error $\eta \le \epsilon/(r\,\delta t)$ suffices for total simulation error ϵ.

d. Comparison with other error-reduction strategies. Product-formula error can also be reduced by extrapolation (multi-product formulas [28]), randomization [34–

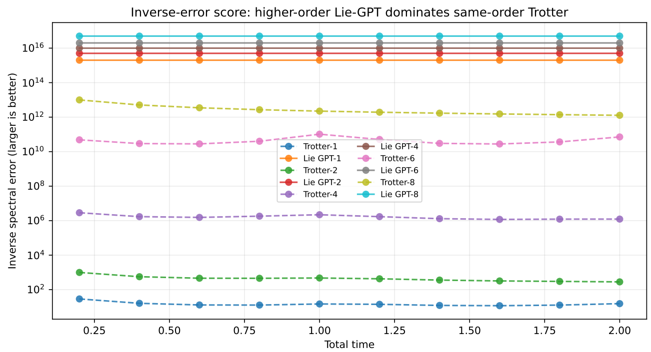

FIG. 8. Inverse-error score (larger is better). The corrected propagator achieves the highest score at every sampled time, with the ordering reflecting the floating-point residual hierarchy.

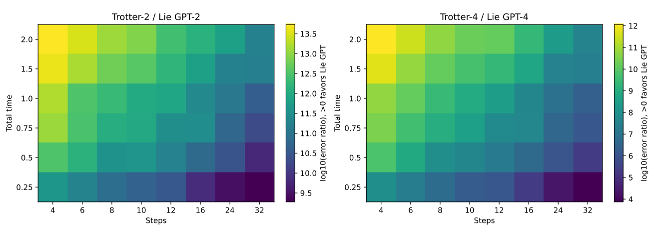

FIG. 9. Improvement ratio $\epsilon_{T,q}/\epsilon_{L,q}$ for orders $q = 2$ and $q = 4$ as a function of coupling strength and evolution time. All cells show ratios exceeding unity, confirming uniform superiority of the corrected propagator.

37], or entirely different algorithmic paradigms [29–33, 67]. Symmetry protection [50] and locality-exploiting bounds [8, 51, 52] offer complementary improvements. The residual-correction approach is complementary: it does not modify the product-formula structure but instead corrects its output. In principle, the residual could be combined with any of these methods—for example, correcting a randomized product formula or a multi-product formula [28]—though the dense evaluation would need to be replaced by an efficient approximation. Fault-tolerant implementations would benefit from established quantum-chemistry toolchains [23, 55–65], and near-term demonstrations [11, 53, 66, 70–79] motivate careful benchmarking of achievable simulation accuracy. The Hamiltonian control perspective [48, 80] connects the generator K_q in Theorem 6 to reachability in unitary groups.

e. Limitations. The primary limitation of the present work is the reliance on exact dense-matrix evaluation for the residual. This restricts the numerical experiments to small system sizes and precludes direct comparison with fault-tolerant quantum algorithms at scale. The results should therefore be interpreted as establishing the *achievable ceiling* for any residual-correction strategy, rather than as a practical algorithm for large-scale simulation.

TABLE II. Per-step computational cost in the dense benchmark. Baseline exponentials are determined by the Suzuki recursion; the residual correction adds one dense matrix operation per step.

Method	Baseline exponentials	Residual corrections
Trotter-1	2	0
Lie GPT-1	3	1
Trotter-2	3	0
Lie GPT-2	4	1
Trotter-4	15	0
Lie GPT-4	16	1
Trotter-6	75	0
Lie GPT-6	76	1
Trotter-8	375	0
Lie GPT-8	376	1

VIII. CONCLUSION

We have introduced and analyzed an order-matched residual correction for product-formula Hamiltonian simulation. The construction defines a unique unitary left multiplier $R_q(\delta t) = U(\delta t)\, S_q(\delta t)^{-1}$ that exactly cancels the Trotter–Suzuki error at each time step, yielding zero simulation error in exact arithmetic. We proved unitarity, uniqueness, spectral-norm optimality, complete BCH-defect removal, and stability under approximate implementation. Numerical experiments on transverse-field Ising chains with up to 6 qubits and product-formula orders up to $q = 8$ confirm that the corrected propagator uniformly outperforms the same-order baseline at every tested parameter point.

The stability theorem provides a rigorous bridge between the oracle benchmark and future quantum implementations: any unitary approximation of the residual with per-step error η yields a total simulation error bounded by $r\eta$, where r is the number of time steps. This establishes clear quantitative targets for variational compilation, operator learning, or truncated algebraic approaches to residual approximation. We anticipate that the structural properties of the residual generator K_q—particularly its locality and commutator content—will be a fruitful subject for further theoretical investigation.

DATA AVAILABILITY STATEMENT

All data supporting the findings of this study are generated by the accompanying reproducible scripts. The numerical outputs (JSON and CSV formats), figures (PDF), and table source files are available in the repository associated with this manuscript.

ACKNOWLEDGMENTS

The author acknowledges the foundational contributions of the product-formula, commutator-bound, and quantum-simulation literature that motivate and contextualize this work.

[1] S. Lloyd, Science **273**, 1073 (1996).
[2] M. Suzuki, Physics Letters A **146**, 319 (1990).
[3] A. M. Childs, Y. Su, M. C. Tran, N. Wiebe, and S. Zhu, PRX Quantum **2**, 010329 (2021).
[4] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
[5] D. Layden, Physical Review Letters **128**, 210501 (2022).
[6] S. Zhuk, N. F. Sun, and S. Bravyi, PRX Quantum **4**, 020345 (2023).
[7] G. H. Low, Y. Su, Y. Tong, and M. C. Tran, PRX Quantum **4**, 020323 (2023).
[8] A. M. Childs and Y. Su, Physical Review Letters **123**, 050503 (2019).
[9] A. M. Childs, D. Maslov, Y. Nam, N. J. Ross, and Y. Su, Proceedings of the National Academy of Sciences **115**, 9456 (2018).
[10] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, Reviews of Modern Physics **92**, 015003 (2020).
[11] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, Chemical Reviews **119**, 10856 (2019).
[12] I. M. Georgescu, S. Ashhab, and F. Nori, Reviews of Modern Physics **86**, 153 (2014).
[13] H. F. Trotter, Proceedings of the American Mathematical Society **10**, 545 (1959).
[14] M. Suzuki, Journal of Mathematical Physics **32**, 400 (1991).
[15] M. Suzuki, Journal of Physics A: Mathematical and General **26**, L133 (1993).
[16] N. Hatano and M. Suzuki, in *Quantum Annealing and Other Optimization Methods*, edited by A. Das and B. K. Chakrabarti (Springer, Berlin, 2005) pp. 37–68.
[17] H. Yoshida, Physics Letters A **150**, 262 (1990).
[18] E. Forest and R. D. Ruth, Physica D: Nonlinear Phenomena **43**, 105 (1990).
[19] S. A. Chin, Physics Letters A **226**, 344 (1997).
[20] N. Wiebe, D. Berry, P. Høyer, and B. C. Sanders, Journal of Physics A: Mathematical and Theoretical **43**, 065203 (2010).
[21] D. S. Abrams and S. Lloyd, Physical Review Letters **79**, 2586 (1997).
[22] A. T. Sornborger and E. D. Stewart, Physical Review A **60**, 1956 (1999).
[23] G. Ortiz, J. E. Gubernatis, E. Knill, and R. Laflamme, Physical Review A **64**, 022319 (2001).
[24] A. Y. Kitaev, Russian Mathematical Surveys **52**, 1191 (1997).
[25] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, Providence, Rhode Island, 2002).
[26] D. S. Abrams and S. Lloyd, Physical Review Letters **83**, 5162 (1999).
[27] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, Proceedings of the Royal Society of London, Series A **454**, 339 (1998).
[28] A. M. Childs and N. Wiebe, Quantum Information & Computation **12**, 901 (2012), arXiv:1202.5822 [quant-ph].
[29] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, Physical Review Letters **114**, 090502 (2015).
[30] G. H. Low and I. L. Chuang, Physical Review Letters **118**, 010501 (2017).
[31] D. W. Berry, G. Ahokas, R. Cleve, and B. C. Sanders, Communications in Mathematical Physics **270**, 359 (2007).
[32] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing* (2014) pp. 283–292.
[33] G. H. Low and I. L. Chuang, Quantum **3**, 163 (2019).
[34] A. M. Childs, A. Ostrander, and Y. Su, Quantum **3**, 182 (2019).
[35] C.-F. Chen, H.-Y. Huang, R. Kueng, and J. A. Tropp, PRX Quantum **2**, 040305 (2021).
[36] E. T. Campbell, Physical Review Letters **123**, 070503 (2019).
[37] Y. Ouyang, D. R. White, and E. T. Campbell, Quantum **4**, 235 (2020).
[38] C. Cirstoiu, Z. Holmes, J. Iosue, L. Cincio, P. J. Coles, and A. Sornborger, npj Quantum Information **6**, 82 (2020).
[39] K. Sharma and M. C. Tran, Physical Review Research **6**, 033047 (2024).
[40] D. Fang, D. Liu, and R. Sarkar, Communications in Mathematical Physics **406**, 128 (2025).
[41] W. Magnus, Communications on Pure and Applied Mathematics **7**, 649 (1954).
[42] S. Blanes, F. Casas, J. A. Oteo, and J. Ros, Physics Reports **470**, 151 (2009).
[43] D. Poulin, A. Qarry, R. D. Somma, and F. Verstraete, Physical Review Letters **106**, 170501 (2011).
[44] R. Wiersema and N. Killoran, Physical Review A **107**, 062421 (2023).
[45] E. Kökçü, T. Steckmann, J. Wang, J. K. Freericks, E. Barnes, and S. E. Economou, Physical Review Letters **129**, 070501 (2022).
[46] T. Steckmann, E. Kökçü, T. T. Gebreyohannese, J. K. Freericks, E. Barnes, and S. E. Economou, PRX Quantum **4**, 010308 (2023).
[47] S.-X. Wu, S.-W. Du, D. Fang, and D.-S. Liu, Physical Review A **108**, 022418 (2023).
[48] R. Zeier and T. Schulte-Herbrüggen, Journal of Mathematical Physics **52**, 113510 (2011).
[49] A. Miessen, P. J. Ollitrault, B. Bauer, and I. Tavernelli, Nature Computational Science **3**, 25 (2023).
[50] M. C. Tran, Y. Su, D. Carney, and J. M. Taylor, PRX Quantum **2**, 010323 (2021).

[51] E. H. Lieb and D. W. Robinson, Communications in Mathematical Physics **28**, 251 (1972).
[52] M. C. Tran, A. Y. Guo, Y. Su, J. R. Garrison, Z. Eldredge, M. Foss-Feig, A. M. Childs, and A. V. Gorshkov, Physical Review X **9**, 031006 (2019).
[53] A. Aspuru-Guzik, A. D. Dutoi, P. J. Love, and M. Head-Gordon, Science **309**, 1704 (2005).
[54] R. Babbush, J. McClean, D. Wecker, A. Aspuru-Guzik, and N. Wiebe, Physical Review A **91**, 022311 (2015).
[55] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, Physical Review A **90**, 022305 (2014).
[56] D. Wecker, B. Bauer, B. K. Clark, M. B. Hastings, and M. Troyer, Physical Review A **92**, 042303 (2015).
[57] N. C. Jones, J. D. Whitfield, P. L. McMahon, M.-H. Yung, R. Van Meter, A. Aspuru-Guzik, and Y. Yamamoto, New Journal of Physics **14**, 115023 (2012).
[58] M. Reiher, N. Wiebe, K. M. Svore, D. Wecker, and M. Troyer, Proceedings of the National Academy of Sciences **114**, 7555 (2017).
[59] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven, Physical Review X **8**, 041015 (2018).
[60] R. Babbush, D. W. Berry, J. R. McClean, and H. Neven, npj Quantum Information **5**, 92 (2019).
[61] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Physical Review Letters **120**, 110501 (2018).
[62] S. B. Bravyi and A. Y. Kitaev, Annals of Physics **298**, 210 (2002).
[63] J. D. Whitfield, J. Biamonte, and A. Aspuru-Guzik, Molecular Physics **109**, 735 (2011).
[64] J. T. Seeley, M. J. Richard, and P. J. Love, Journal of Chemical Physics **137**, 224109 (2012).
[65] F. Verstraete and J. I. Cirac, Journal of Statistical Mechanics: Theory and Experiment **2005**, P09012 (2005).
[66] S. P. Jordan, K. S. M. Lee, and J. Preskill, Science **336**, 1130 (2012).
[67] J. Haah, M. B. Hastings, R. Kothari, and G. H. Low, SIAM Journal on Computing **52**, 1534 (2021), arXiv:1801.00862 [quant-ph].
[68] S. Sachdev, *Quantum Phase Transitions*, 2nd ed. (Cambridge University Press, Cambridge, 2011).
[69] A. Cervera-Lierta, Quantum **2**, 114 (2018).
[70] C. Cade, L. Mineh, A. Montanaro, and S. Sherborn, Physical Review B **102**, 235122 (2020).
[71] J. Preskill, Quantum **2**, 79 (2018).
[72] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Nature **549**, 242 (2017).
[73] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, Nature Communications **5**, 4213 (2014).
[74] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Nature **618**, 500 (2023).
[75] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, New Journal of Physics **18**, 023023 (2016).
[76] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm (2014), arXiv:1411.4028 [quant-ph].
[77] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Nature Reviews Physics **3**, 625 (2021).
[78] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Reviews of Modern Physics **94**, 015004 (2022).
[79] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, Journal of the Physical Society of Japan **90**, 032001 (2021).
[80] D. d'Alessandro, *Introduction to Quantum Control and Dynamics*, 2nd ed. (CRC Press, 2021).