arXiv:2611.09822v3 [quant-ph] 16 Nov 2026

Discovering formally verified quantum programs with AI

Sergey Shirobokov[1,*,†], Ben Jaderberg[1,*,†], Katerina Hristova[‡,*], George Tsoukalas[1,*], Johannes Bausch[1,*], Eric Wieser[1], Francisco J.H. Heras[1], Nicholas C. Rubin[2], Guang Hao Low[2], Sébastien Racanière[1], Nenad Tomašev[1], Lei M. Zhang[1], James D. Watson[2], Tanuj Khattar[2], Aws Albarghouthi[‡], Stavan Jain[‡], Harun Khan[1], Leopoldo Sarra[1], Andrew W. Senior[1], Michael Broughton[2], Hussain Masoom[1], Sergio Boixo[2], Ryan Babbush[2], Swarat Chaudhuri[1], and Pushmeet Kohli[1]

[1]Google DeepMind [2]Google Quantum AI

November 17, 2026

The timeline by which quantum computers can solve breakthrough problems in science depends on improved quantum algorithms as much as on better hardware.[1,2] Such algorithmic advances are typically achieved through many years of expert human collaboration.[3,4] In mathematics and computer science, large language models have accelerated progress through automated discovery loops.[5,6] However, quantum algorithms can only be evaluated by classical computers on restricted or small instances.[7,8] Here we demonstrate automated discovery of quantum programs beyond these regimes including order-of-magnitude improvements on scientific applications. To achieve this, we build a quantum program library in the Lean theorem prover in which Gemini 4 agents synthesize novel programs against formal problem specifications and verify their correctness and cost. Agent teams collaborate, sharing results and knowledge through wikis and discussing ideas at organized conferences. Our results improve on prior baselines for parallel search[9] (up to 2.57×), quantum arithmetic[10] (290×), and simulations relevant to fusion, photochemistry and drug-metabolism by 12×, 8× and 2.5× respectively. Combining human and artificial intelligence at scale, under the rigor of mathematical proof, lets researchers design algorithms that today's classical and quantum computers cannot test, shortening the path to practical quantum computation.

Quantum computers promise to tackle problems beyond classical reach across cryptanalysis,[11] optimization,[12] chemistry and drug discovery,[13–15] materials science,[16,17] and high-energy and plasma physics.[18–20] Many of the most promising applications are estimated to require a device with thousands of error-corrected (logical) qubits executing 10^9 to 10^{15} logical operations,[1,21] well beyond the roughly one million operations projected for early fault-tolerant processors.[2] Beyond improvements in hardware and quantum error correction to build better logical qubits from noisy physical qubits,[22] closing this gap relies on cumulative algorithmic resource reductions that lower qubit and gate counts.[23–27] Historically, such reductions take sustained scientific effort over many years.[3,4]

Artificial intelligence (AI) offers a promising route to accelerate these improvements. Traditional machine learning has already been successfully applied to e.g. quantum hardware control,[28–30] error-correction decoding,[31]

*Equal contributions.
†Corresponding authors: shirobokov@google.com, benjaderberg@google.com.
‡These authors performed this work as research consultants on assignment at Google.

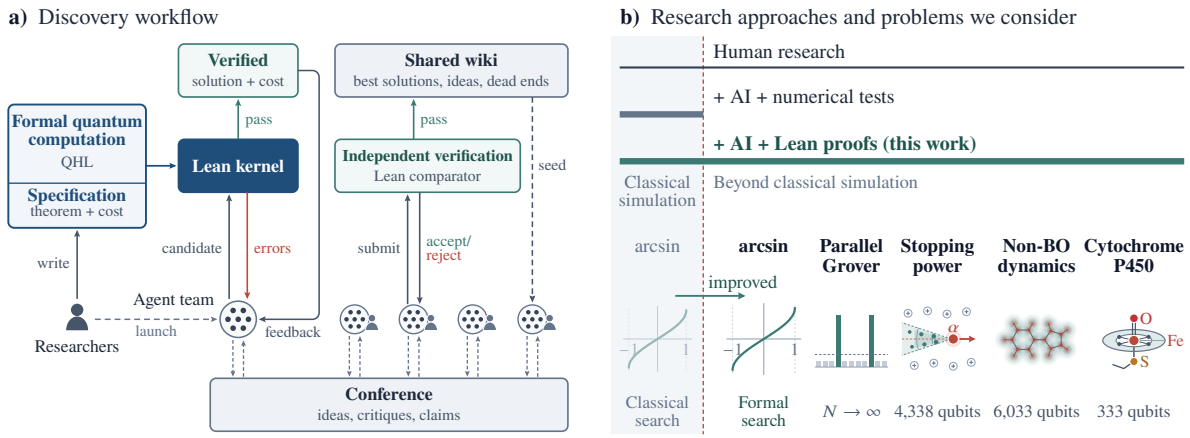

Figure 1: **Formally verified multi-agent quantum program discovery. a)** Researchers write a Lean quantum Hoare logic specification of correctness and cost, then launch their agent team, which iterates on error feedback from the Lean kernel. Teams publish independently re-verified solutions, partial results and dead ends to a shared wiki, and exchange ideas and critiques in scheduled conferences. **b)** Bars indicate which problem each approach can reach. AI search checked by numerical tests can only produce solutions that can be simulated classically (dashed line), which for arcsin means a restricted gate set.[7]

and small-scale circuit synthesis.[32] While large language models (LLMs) are rapidly advancing mathematics and classical algorithms via execution-based search and automated reasoning[5,33–39] as well as formal theorem proving for fixed conjectures,[6,40] their application to quantum computing lags behind due to the challenge of *verifying* general correctness without fault-tolerant hardware to test on. High-throughput validation of LLM-generated quantum programs has thus been restricted to circuit regimes that can be classically simulated,[7,8] limiting the design space and scale of discovery (Fig. 1b), whereas concurrent LLM-assisted formalization has focused on verifying known algorithms (Appendix A). Bringing AI discovery to general quantum algorithms instead requires pairing formal models of quantum computation with mathematical proofs of correctness and cost, with access to tools such as numerical heuristics where they are effective. Formal reasoning is particularly critical when targeting asymptotic speedups and large-scale applications spanning thousands of qubits—just like scientists have reasoned about quantum algorithms over the past decades.

In this work, we introduce a Lean[41] verification library based on quantum Hoare logic (QHL),[42] engineer a multi-agent research workflow built around it, and prove new resource reductions across asymptotic algorithms and large-scale scientific simulations. Our library pairs a language for quantum programs and classical control flow with formal proofs of correctness and resource cost, chosen per specification. For fault-tolerant circuits we use the spacetime cost $Q \cdot \mathcal{T}$, the number Q of qubits times the number $\mathcal{T}$ of T gates, which are costly in error-corrected quantum computers.[43] Researchers launch heterogeneous teams of LLM agents that iterate on algorithms, and exchange concepts, partial results, and dead ends through a shared wiki and scheduled cross-team virtual conferences, while an independent verifier audits every proof against a frozen specification written in our library (Fig. 1a). Across five problems (Fig. 1b)—computing arcsin, parallel Grover search, cytochrome P450 energy levels, fusion alpha-particle stopping power and non-Born–Oppenheimer molecular dynamics—this workflow discovers new programs and circuits with machine-checked proofs of correctness and cost, reducing resource costs by 290× for arcsin (combining classical and formal search), and by up to 12× for our scientific applications compared to the strongest prior baselines (Fig. 3). Together, these advances shift researchers effort from manual circuit design, optimization, and resource tracking to problem specification and scientific interpretation.

Specification `StoppingPower.lean`

```
abbrev systemQubits : ℕ := 3 * (218 * 6 + 8)
def stoppingPowerAncillas : ℕ := sorry
noncomputable def stoppingPowerProgram (c : Configuration)
    (j : Fin 10) ⋯ : DynamicCircuit Var n := sorry

theorem stoppingPower_correctness ⋯ :
  ∃ W : Unitary (2 ^ (systemQubits + stoppingPowerAncillas))
    (∀ spinUp, ApproximatesTimeEvolution (hamiltonian c)
      (j.val + 1 : ℝ) 0.01 (fermionicProjector spinUp) W) ∧
    RealizesUnitaryFromBasisOutcomesAndZeroWorkspace
      rSys gridStates rAnc hdisj W
      (stoppingPowerProgram c j ⋯).compile := by sorry

def stoppingPowerConstCost : ℕ := sorry
theorem stoppingPower_const_cost ⋯ :
  (∀ j, (stoppingPowerProgram c j ⋯).count .TOFF = 0) ∧
  (systemQubits + stoppingPowerAncillas) *
    (50 * ∑ j, (stoppingPowerProgram c j ⋯).count .T)
      ≤ stoppingPowerConstCost := by sorry
```

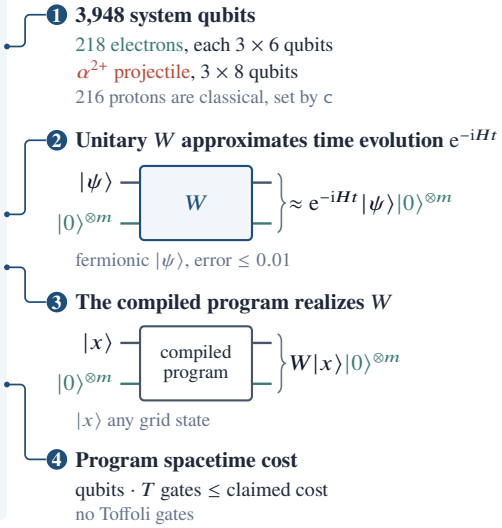

Figure 2: **Lean specifications for discovery tasks.** An abridged specification for electronic stopping power in warm-dense hydrogen,[48] where the goal is to construct a quantum program that simulates time evolution under the system Hamiltonian H. The agents fill in each `sorry` statement, which is Lean's placeholder for an unfilled definition or proof: the ancilla qubit count, the circuit and its claimed cost, and the proofs of the correctness and cost theorems. (**1**) The system qubits store the momentum of every quantum particle. (**2**) A unitary matrix W on the system and ancilla qubits (extra qubits used to aid computation) must simulate time evolution $e^{-iH(c)t}$ for every physically allowed electron state, up to a specified error, and return the ancillas to $|0\rangle$. (**3**) The compiled program must act as W on grid states whenever the ancillas start in $|0\rangle$. (**4**) The cost is a function of the logical spacetime volume $Q \cdot \mathcal{T}$ multiplied by the repeated runs of the stopping power routine: 50 repeats × 10 time steps × $Q \cdot \mathcal{T}$ (Sec. 2.3).

1. Formally Verified Multi-Agent Discovery

1.1. Formal Verification of Quantum Algorithms

We implement our verification and resource-counting library[44] in the Lean theorem prover,[41] building on Mathlib[45] and Physlib's Lean-QuantumInfo.[46,47] These formalize core objects of quantum information, but reasoning about quantum programs also requires a programming language to express them and a program logic to prove them correct. Our library fills this gap by combining a classical-quantum programming language, a classical-quantum Hoare logic to reason about program correctness[42] and support for resource counting. Besides classical operations such as assignment, conditional branching and unbounded loops, the formalized language supports core quantum operations: unitary transformations, mid-program qubit measurement, and state reset. We include further details in Appendix A.

Correctness theorems are formulated as so-called *Hoare triples* of (precondition, program, postcondition), which state that all executions of a program starting in a state satisfying a specified precondition end in a state satisfying a chosen postcondition. Triples are proved backwards from the postcondition by computing the weakest preconditions for each step of the program.[42] Alongside correctness, our library proves bounds on both expected and worst-case resource costs across branches and loops, including T-gate count and number of qubits required to execute the algorithm. To pose a program discovery task, we write a *problem specification* as a Lean file (Fig. 2) that states the target correctness and cost theorems while leaving placeholders (`sorry`) to be filled in with a candidate program, its cost bound, and their proofs.

1.2. Multi-Agent Discovery Architecture

To begin the search on a problem, researchers write the specification and launch independent teams of LLM agents with high-level optimization prompts rather than problem-specific ideas (Fig. 1a). The agent teams then explore candidate algorithms and prove their correctness and cost against the independent verifier. Each researcher chooses their own team setup, from a single agent iterating on a candidate to hierarchical teams targeting individual subroutines, using tools such as analytical derivations, simulation, and numerical optimization (Appendix C). Unlike multi-agent science systems with a fixed architecture[49–52] or single search loop,[5,36,40,53,54] we prescribe no internal team design and fix only the shared specification, verifier, and workflow. All large-scale experimental runs of autonomous discovery were performed using Gemini 4 Carbon models.

Agent teams can share progress through two channels. A wiki stores each problem's best verified solution, reusable Lean lemmas, and reviewed notes on promising ideas and dead ends.[55] Newly launched teams can start from the wiki's best solution. Teams also meet in scheduled forum conferences to critique proposals and discuss unverified conjectures, which count as results only once verified.

Because agents can freely modify files in their workspace, they may exploit loopholes by weakening theorem statements, altering definitions, or adding axioms.[56] To prevent this, an external verifier checks every submission against the frozen specification using the Lean Comparator,[57] confirming that all target theorems, underlying definitions, and allowed axioms are unchanged (Appendix B). Researchers thus review only the specification before a run and can trust that any found and verified solution is a valid algorithm.

2. Results

We demonstrate automated quantum program discovery on five applications (Fig. 3): parallel unstructured search (Sec. 2.1), fault-tolerant arithmetic (Sec. 2.2), and three physical and chemical simulations (Secs. 2.3–2.5).

2.1. Parallel Grover Search

Unstructured search, the problem of finding one or all of k marked items in a list of N candidates, provided one of the earliest provable quantum advantages, with Grover's algorithm achieving a quadratic speedup over classical search.[9] Marked items are identified by an *oracle* subroutine assumed to dominate runtime, so cost is measured solely in oracle queries. Allowing up to m parallel queries per round on a shared quantum register (rather than m independent, uncoupled devices),[9,63] we require algorithms to succeed with certainty and score them by their expected oracle rounds for large N, evaluated at powers of two up to 1,024 for both m and k (Fig. 3d). We compare against the strongest published protocol in each regime,[62–66] including Zalka's partition combined with weakly measured search at $k = 1$, evaluated under the same metric (Appendix E).

Our Lean-verified circuits match this baseline on the $m = 1$ and $k = 1$ boundaries for finding a single element and outperform it at all multi-target parallel settings ($m, k \geq 2$), reducing expected rounds by up to 1.22×; for finding all elements, they outperform it at all multi-target settings ($k \geq 2$), by up to 2.57×. Part of this gain comes from combining published ideas via classical control, such as pairing domain partitioning or parallel restarts with weak measurements, which collapse coherence only partially at the price of only partial success information. The largest find-all gains come from a new coherent mechanism: instead of splitting the list into m blocks with unknown target counts, the circuit holds a superposition over all sets of m candidates, and each parallel query imprints a sign set by the parity of how many remaining targets a set contains. Between queries, a fixed, target-independent walk on the Johnson graph of these sets steers amplitude toward sets containing every remaining target, which a final query records at once (Appendix E.3).

a) Summary of results using absolute cost

Problem	Application	Baseline $\mathcal{Q} \cdot \mathcal{T}$	This work $\mathcal{Q} \cdot \mathcal{T}$	Reduction
Fixed-point arcsin[10]	State preparation	$> 1.08 \times 10^8$	3.70×10^5	$> 290\times$
α + hydrogen stopping power[48]	Fusion	5.42×10^{15}	4.59×10^{14}	$11.8\times$
$C_{23}H_{20}N_3O$ non-BO dynamics[58,59]	Photochemistry	4.98×10^{15}	6.06×10^{14}	$8.22\times$
Cytochrome P450 ground and low-lying state energies[60,61]	Drug metabolism	1.09×10^{12}	4.37×10^{11}	$2.49\times$

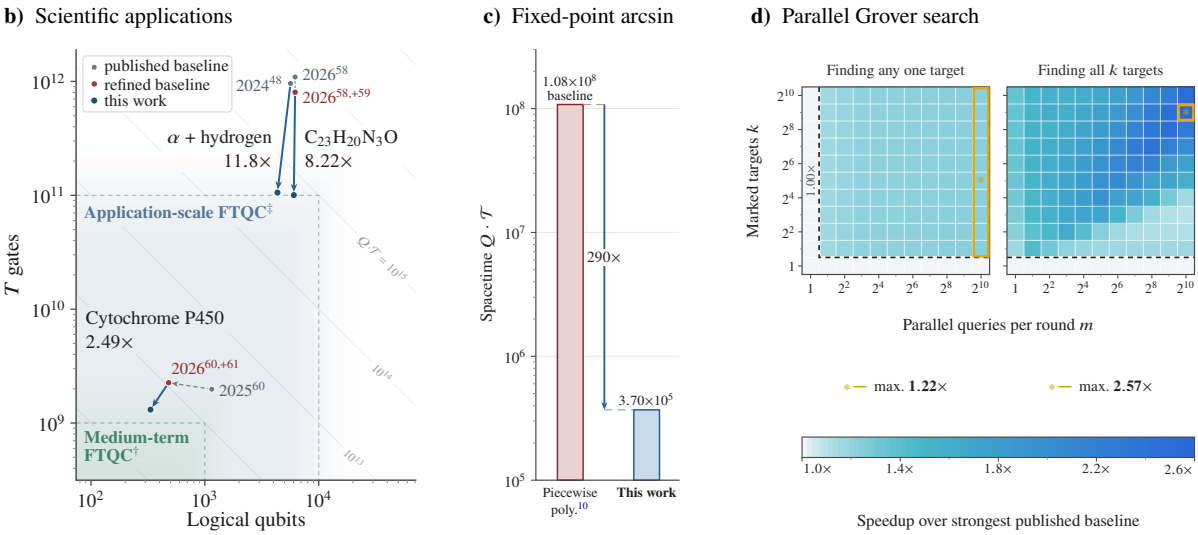

Figure 3: **Summary of verified results.** **a)** Quantum algorithm cost for baselines and solutions from this work across application domains. **b)** Logical-qubit and T-gate improvements on the three scientific applications, with feathered hardware horizons (dashed lines at geometric midpoints) for medium-term fault-tolerant quantum computing (FTQC†; teraquop planar architectures[7,61]) and application-scale FTQC‡ (fault-tolerant application-scale quantum computers, FASQ[1,2,61]). Grey points denote published baselines, red points denote refined baselines incorporating subsequent methodological improvements from the literature, and blue points denote this work. For cytochrome P450, the published ground-state estimate is extended to all eigenstates within 0.1 Ha of E_0, raising its cost by 0.85%. **c)** Cost for 25-bit fixed-point arcsin ($\varepsilon \leq 1.25 \times 10^{-7}$) against piecewise-polynomial evaluation.[10] **d)** Speedup in expected number of rounds for large N of our verified parallel Grover circuits over the strongest published baseline[62–66] for powers of two $m, k \leq 2^{10}$ when finding any one or all of k marked targets with m parallel queries per round (dashed contours: 1.00× boundary; $*$ markers: peak speedups of 1.22× and 2.57×).

2.2. Fixed-Point arcsin

Computing the function $\arcsin(y)$ on a quantum superposition of inputs is an important subroutine for loading numbers into a wavefunction, used in many algorithms including for solving linear systems of equations. We specifically consider the case where the input y is represented by 25-classical bits, and the output must be accurate to an error $\leq 1.25 \times 10^{-7}$ across the domain $y \in [-1, 1]$. For this setting, the best-known published baseline implementation is a piecewise-polynomial construction,[10] which requires a spacetime cost exceeding 1.08×10^8 for an error of $\sim 10^{-6}$.

Unlike the other examples in this work, arcsin is known to have quantum implementations that can be efficiently simulated classically by using a restricted gate set called *kickmix*.[7] Therefore, we run automated discovery in two steps. First, the agents only consider kickmix circuits, exploiting that solutions can be verified in a matter of seconds. This leads to fast improvements, overall reducing the spacetime cost to 4.15×10^5 (4,152 T gates on 100

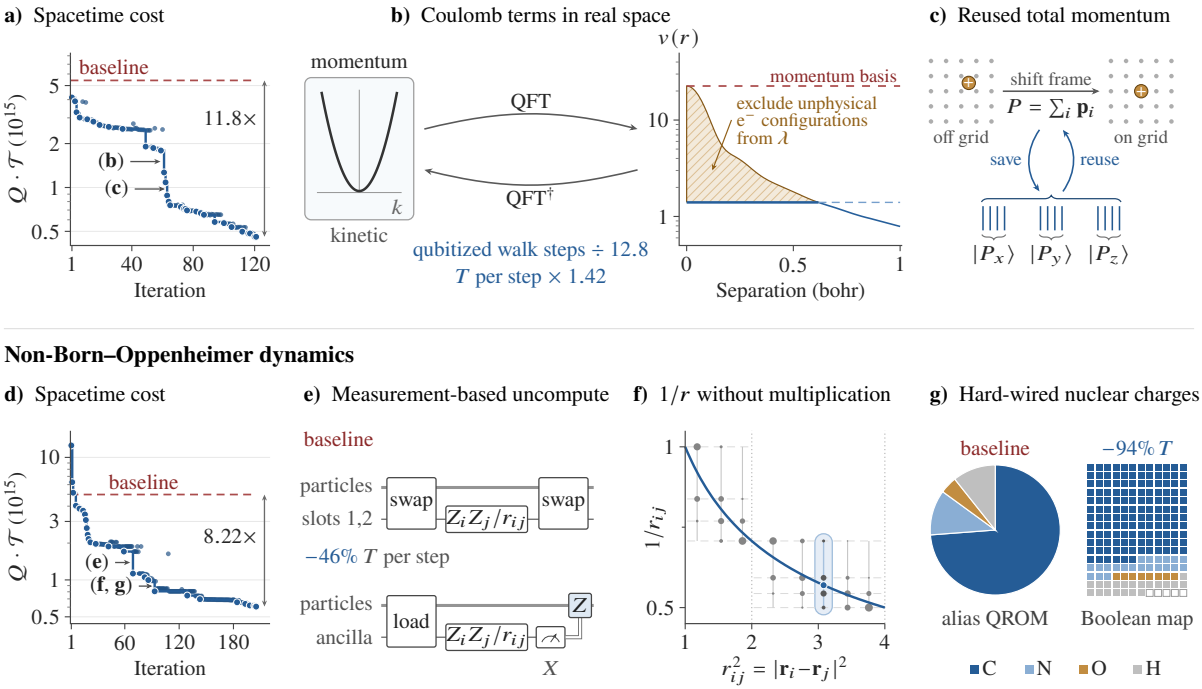

Figure 4: **Where program improvements come from. a, d)** Cost of verified solutions. Arrows mark where the mechanisms in the other panels were introduced. Interestingly, the drop at iteration 49 in **a)** is not present in the final solution. **b)** A quantum Fourier transform (QFT) moves the electrons onto a real-space grid and back. On the grid, the pair potential $v(r)$ is split into long-range (blue) and short-range (hatched) parts. The long-range part costs about 30 times less per pair than in the momentum basis (red dashed). For the short-range part, a subroutine bounds the interaction strength of configurations that violate the Pauli exclusion principle, further reducing its cost. **c)** Protons sit off the grid, so to apply an electron–proton interaction the circuit shifts all electrons to put the proton on a grid point, a phase set by the sum of all 218 electron momenta. Keeping this sum in a 42-qubit register between walk steps halves the summations per step. **e)** To apply the electrostatic interaction Z_iZ_j/r_{ij}, rather than swapping the coordinates of the interacting pair into fixed slots and back, a copy is loaded into ancilla qubits. This can be cleared cheaply by measurement without destroying the originals. **f)** From the loaded coordinates, r_{ij}^2 is calculated cheaply and rescaled within each interval to $\tilde{r}_{ij}^2 \in [0, 1)$. Through random sampling, one of four classically precomputed coefficients is chosen, with probabilities set by $\tilde{r}_{ij}^2$ (dot area). In expectation, this gives $1/r_{ij}$ as a Bernstein polynomial,[67] as shown for one value of r_{ij}^2 (shaded). **g)** Each nucleus has integer charge $Z_n \in \{1, 6, 7, 8\}$. In the electron–nucleus attraction, rather than preparing these weights via alias sampling,[24] a uniform superposition is created and reversible logic assigns Z_n of them to each nucleus. Applying a unit-charge interaction on every branch then gives each nucleus its full charge Z_n.

qubits). Subsequently, we use our formal verification library to lift the kickmix restriction and further optimize the algorithm, using a new quantum gate set that makes the solution hard to verify outside Lean. From this, the agents discover a solution in which a certain component, the coefficient lookup, can be implemented using 33% fewer gates, leading to a further overall 11% lower cost of 3.70×10^5 (3,698 T gates on 100 qubits), 290× below the baseline.[10]

2.3. Stopping Power in Inertial-Confinement Fusion

In inertial-confinement fusion, hydrogen fuel ignites only if the energetic α particles it produces deposit enough of their energy back into the fuel to sustain the reaction.[68,69] Modeling this requires knowing the fuel's stopping power, the rate at which it absorbs the kinetic energy of an α particle passing through it. No classical method for these dynamics is both efficient and systematically improvable.[48,70–73]

We consider an α particle and 218 electrons moving among 216 fixed protons. A candidate circuit must simulate the time evolution of the system e^{-iHt} with error $\varepsilon = 10^{-2}$ on every state antisymmetric under exchange of same-spin electrons. The full stopping-power calculation repeats this evolution 500 times.

Our baseline is constructed from Ref. 48: we take the 5,650 logical qubits and parameters in its Table 2 and evaluate its equation (12), giving a spacetime cost of 5.42×10^{15} per run, averaged over the 500 runs (Appendix G). This builds e^{-iHt} from a sequence of walk steps, each encoding H/λ, where λ is the sum of the absolute weights when H is written as a linear combination of unitaries (LCU). The number of walk steps[25] scales as λt.

Running our discovery workflow, we find a solution with 4,338 logical qubits and 11.8× lower spacetime cost (Fig. 4a). To understand the main improvement, we note that over 90% of the baseline's λ comes from electron–proton and electron–electron interactions, and that λ must account for every state the qubits can be in, including unphysical ones. Our main innovation is to quantum Fourier transform all electrons to a real-space grid to apply these interactions (Fig. 4b). Importantly, in real space the solution can exploit Pauli exclusion, which allows at most two opposite spin electrons on any grid point. For the short-range part of the potential, a subroutine kicks back a phase of i in one LCU branch and $-i$ in the other for each electron on a nearby grid point. The branches interfere to apply the correct interaction, up to a constant, for two or fewer electrons, while keeping it no larger for three or more, so these grid points no longer set λ. In the final circuit, this cuts the walk steps by 12.8×, at 1.42× the T gates per step.

The largest other improvement comes from keeping the total electron momentum, needed to apply the electron–proton interaction, in a 42-qubit register between walk steps. This halves how often it is recomputed and cuts the T gates per step by 10% (Fig. 4c). More details are given in Appendix G.

2.4. Non-Born–Oppenheimer Dynamics of $C_{23}H_{20}N_3O$

Light-driven chemical processes such as solar energy conversion,[74,75] photocatalysis[76] and biological photoreception[77] proceed through excited-state pathways where electronic and nuclear motion occur on comparable timescales. This invalidates the Born–Oppenheimer approximation,[78–80] meaning the nuclei must also be included in the quantum wavefunction. Methods that treat the nuclei classically struggle to simulate quantum tunneling and superpositions of electronic states,[81,82] while exact classical simulation scales exponentially with system size.[14,78]

Following Reaction 5 of Ref. 58, we simulate time evolution of $C_{23}H_{20}N_3O$, modeled on molecules proposed for photo-activated, proton-coupled electron transfer.[75,83] Unlike the fixed protons of Sec. 2.3, all 47 nuclei are treated quantum mechanically alongside the 187 electrons (5,616 system qubits). Here a candidate circuit must implement time evolution e^{-iHt} for $t = 1$ fs, to within error $\varepsilon = 10^{-2}$, for every initial state.

Improving the published estimate[58] by the authors' subsequent work[59] gives a baseline cost of $4.98\times10^{15}\ Q\cdot\mathcal{T}$ (Appendix H). In this work we find a solution that costs only $6.06\times10^{14}\ Q\cdot\mathcal{T}$, a reduction of 8.22× (Fig. 4d).

Within each walk step, three changes lead to the most significant improvements. First, rather than swapping interacting particles into fixed sets of qubits and back,[58] our solution loads their coordinates into ancilla qubits, which can be cleared by cheap measurement[23,27] (Fig. 4e). Second, instead of evaluating $1/r$ with a multiplication step,[58] we obtain $1/r$ in expectation from a piecewise Bernstein polynomial in r^2,[67] using only comparisons and a table lookup (Fig. 4f). Finally, a nucleus of charge Z_n attracts an electron just as Z_n unit charges at the same

point would. The solution exploits this by creating an equal superposition and assigning Z_n of the branches to each nucleus via classical Boolean logic,[27] rather than preparing unequal weights by alias sampling[24,58] (Fig. 4g). More details are given in Appendix H.

2.5. Cytochrome P450 Compound I

Cytochrome P450 enzymes break down most human medicines, so understanding how they react with a drug candidate helps predict how quickly the drug is cleared from the body and whether it produces toxic byproducts.[84–86] These reactions are driven by a short-lived intermediate of the enzyme, Compound I, whose reactivity depends on several closely-spaced spin states. Approximate classical methods such as density functional theory disagree on the energy ordering of these states,[87] making it a central target for quantum phase estimation, a quantum algorithm to obtain eigenstate energies.[3,88]

We consider a model of Compound I with 63 electrons in 58 orbitals (active space X[88]) with the fixed double-factorized tensor-hypercontraction (DFTHC) Hamiltonian.[60] For every eigenstate within $\Delta = 0.1$ Ha of the ground state, the circuit must estimate the energy with root-mean-square error at most $\varepsilon = 1$ mHa, including all approximation errors.

For this problem, the best published spacetime cost is 2.26×10^{12}.[60] To be conservative, we further integrate recent compilation improvements,[61,89] which roughly halve the cost. Finally, since the published estimate is for the ground-state energy only, we adjust it for our task of predicting all energies within Δ of the ground state, giving a cumulative baseline of 1.09×10^{12} on 481 logical qubits (Appendix I).

For this problem, we find a solution that achieves a spacetime cost of 4.37×10^{11}, a reduction of 2.49×. To understand the main improvement, consider the baseline. To apply a Hamiltonian term, each walk step performs 57 rotations in orbital space, to define new coordinates in which the term acts on one orbital.[4] The chosen term is stored as a binary number, which is used to read the rotation angles needed into a 247-qubit register from a quantum read-only memory.[24] Our circuit instead gives each term its own qubit, and sets only the qubit of the chosen term to 1 (a one-hot register). Each angle can then be read from these qubits with only cheap (Clifford) gates once the register is built, with no table lookup. Prior work[89] found this narrowly worse than the baseline because it needs more qubits, and suggested mixing one-hot and binary encodings. Our solution uses this both to look up the term weights and to supply the rotation angles, accounting for over 95% of the T-gate saving per step and most of the qubit saving. More details are given in Appendix I.

2.6. Collaboration Across Agent Teams

We now discuss the dynamics of the agentic runs to better understand the role of multiple agent teams in search. Our best solutions emerged from collaboration across teams rather than any single team working alone. In non-Born–Oppenheimer dynamics, for example, the leading record changed teams eight times after Team 1's initial run, and the final circuit combines techniques from three agent teams, superseding intermediate improvements from other teams (Figs. 4e–g, 5a). Whereas the wiki transferred verified code, conferences exchanged unverified proposals and critique: the largest single drop (−34%) began as a Team 1 conference proposal that Team 2 refined and proved (Fig. 4e), and a second key improvement grew from a similar exchange (Fig. 4f).

Such cross-team contributions were sustained across all three scientific applications. To estimate each channel's share once both were active, we classify every verified gain by where its new idea originated (Fig. 5b). Averaged across the three scientific applications, 61% of the total cost reduction came from within-team ideas, 15% from cross-team wiki entries or Lean files, and 24% from conferences. The cross-team share ranges from 13% in stopping power to 63% in cytochrome P450. Credit is often hard to separate cleanly, as one proof often combines ideas from a team's own search, the wiki, and conferences.

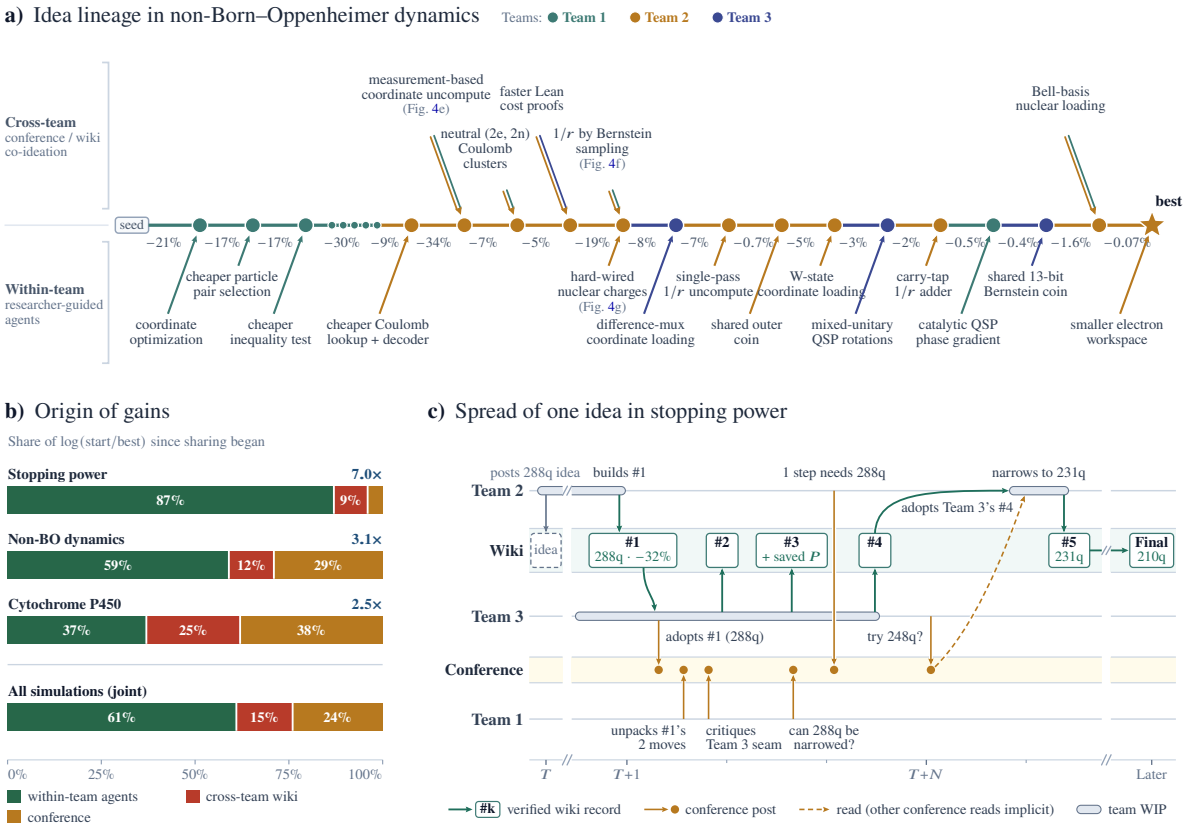

Figure 5: **Cross-team collaboration and idea lineage. a)** Lineage of verified circuits for non-Born–Oppenheimer dynamics across Teams 1–3 (of seven participating teams) from the baseline (seed) to the final circuit (star), with the $Q \cdot \mathcal{T}$ reduction of each step. Within-team discoveries enter from below; cross-team conference and wiki ideas enter from above. **b)** Share of the logarithmic $Q \cdot \mathcal{T}$ reduction from within-team ideas, cross-team wiki entries, and conference exchanges once both channels were active on each scientific application, with reduction factors over this window shown at right and the unweighted average at bottom. **c)** Cross-team evolution of one stopping-power idea across wiki records #1–#5 and conference critiques, narrowing the shared ancilla register from 288 to 231 qubits; later records narrowed it further to 210 qubits.

Figure 5c traces one such cross-team exchange in stopping power, where Teams 2 and 3 refined a shared ancilla register across five successive wiki records with conference input from Team 1, narrowing it from 288 to 231 qubits; later records narrowed it further to 210 qubits. Beyond complete proofs, dead-end notes and partial proofs also transferred across teams: a note proving why a proposed shortcut fails stopped two proposals from other teams, while helper Lean modules proved by one team were imported by a second and reused in a third team's stopping-power circuit.

Machine-checked verification is also what makes open collaboration safe at this scale, letting teams build on one another's Lean proofs without manual code review. Flawed candidates and claims fell into two categories (Appendix B): broken proofs or inaccurate script-based estimates (such as candidates that bypassed local Lean checks), which our external verifier and conference critique filtered out, and loopholes in the specifications themselves (such as exploiting how Toffoli gates were counted without improving the circuit), which researchers and teams closed by revising the specification.

3. Discussion and Outlook

In this work, we show that combining formal verification with LLM agents can push the frontier of quantum algorithm development. We demonstrate this across fundamental algorithms, fault-tolerant subroutines and full-scale scientific applications, achieving up to a 2.57× reduction in expected oracle rounds for parallel Grover search, a 290× spacetime reduction for fixed-point arcsin, and 11.8×, 8.22× and 2.49× reductions on simulations relevant to fusion, photochemistry and drug metabolism. Where classical simulation can check candidate circuits, as with arcsin, numerical testing with formal proving on top is effective. For asymptotic bounds and general algorithms on hundreds or thousands of qubits, formal proof is essential.

Beyond these results, our experiments point to a new research workflow. Rather than replacing human researchers, formal AI discovery shifts where human effort is applied. Researchers and domain experts choose the problems and carefully prepare faithful formal specifications. External verification then lets independent agent teams build on one another's work without manual review, and carry out the search and formal proofs that drive down costs. Finally, researchers interpret the verified solutions to guide future research directions.

Our end-to-end simulations still demand far more resources than projected for the first fault-tolerant quantum computers[2] (Fig. 3b; Appendix D), and reaching early hardware will require further orders-of-magnitude improvements. Part of the challenge is for us to choose specifications well, as agents cannot optimize beyond them. For example, to match the baseline, our non-Born–Oppenheimer dynamics specification requires correctness on *all* input states, even unphysical ones that break exchange symmetry. Similarly, fixing the factorization in cytochrome P450 excludes the representation advances behind recent gains.[4,60,88]

Bringing algorithms onto real hardware also means extending formal specifications to error-correcting codes, qubit routing and physical layout,[43,61] so human–agent teams can optimize algorithms for a given code or co-design both. Furthermore, qubits are expected to be a more limited resource than gates for early machines, suggesting optimizing each factor unequally.[90] Other hardware designs, where T gates can run in parallel, may instead favor minimizing qubits times T depth.[43] Formalizing these architectural constraints allows us to optimize algorithms and fault-tolerant layouts together before large-scale quantum processors exist. Over time, collaborative proof search can thus turn what were once isolated, multi-year derivations into cumulative resource reductions, closing the gap to useful quantum computation from above while physical hardware scales from below.

Code availability

Our Lean library, together with the problem specifications and verified solutions discussed in this manuscript, is available at `https://github.com/google-deepmind/formal-quantum-computation`.[44]

References

1. Beverland, M. E. *et al. Assessing requirements to scale to practical quantum advantage* 2022. arXiv: 2211.07629 [quant-ph].
2. Preskill, J. Beyond NISQ: The Megaquop Machine. *ACM Transactions on Quantum Computing* **6,** 1–7. arXiv: 2502.17368 [quant-ph] (2025).
3. Reiher, M., Wiebe, N., Svore, K. M., Wecker, D. & Troyer, M. Elucidating reaction mechanisms on quantum computers. *Proceedings of the National Academy of Sciences* **114,** 7555–7560 (2017).
4. Lee, J. *et al.* Even More Efficient Quantum Computations of Chemistry Through Tensor Hypercontraction. *PRX Quantum* **2,** 030305 (2021).
5. Novikov, A. *et al. AlphaEvolve: A coding agent for scientific and algorithmic discovery* 2025. arXiv: 2506.13131 [cs.AI].

6. Hubert, T. *et al.* Olympiad-level formal mathematical reasoning with reinforcement learning. *Nature* **651,** 607–613 (2026).
7. Babbush, R. *et al.* Securing Elliptic Curve Cryptocurrencies against Quantum Vulnerabilities: Resource Estimates and Mitigations. *PRX Quantum* **7,** 031001. arXiv: 2603.28846 [quant-ph] (2026).
8. Crichigno, M. & Flammia, S. *Automating quantum circuit optimization with AI agents* Oct. 2026. arXiv: 2610.31352 [quant-ph].
9. Grover, L. K. *A fast quantum mechanical algorithm for database search* in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (1996), 212–219.
10. Häner, T., Roetteler, M. & Svore, K. M. *Optimizing Quantum Circuits for Arithmetic* 2018. arXiv: 1805.12445 [quant-ph].
11. Shor, P. W. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Journal on Computing* **26,** 1484–1509 (1997).
12. Jordan, S. P. *et al.* Optimization by decoded quantum interferometry. *Nature* **646,** 831–836 (2025).
13. Aspuru-Guzik, A., Dutoi, A. D., Love, P. J. & Head-Gordon, M. Simulated Quantum Computation of Molecular Energies. *Science* **309,** 1704–1707 (2005).
14. Kassal, I., Jordan, S. P., Love, P. J., Mohseni, M. & Aspuru-Guzik, A. Polynomial-time quantum algorithm for the simulation of chemical dynamics. *Proceedings of the National Academy of Sciences* **105,** 18681–18686 (2008).
15. Santagati, R. *et al.* Drug design on quantum computers. *Nature Physics* **20,** 549–557 (2024).
16. Bauer, B., Bravyi, S., Motta, M. & Chan, G. K.-L. Quantum Algorithms for Quantum Chemistry and Quantum Materials Science. *Chemical Reviews* **120,** 12685–12717 (2020).
17. Delgado, A. *et al.* Simulating key properties of lithium-ion batteries with a fault-tolerant quantum computer. *Physical Review A* **106,** 032428 (2022).
18. Jordan, S. P., Lee, K. S. M. & Preskill, J. Quantum Algorithms for Quantum Field Theories. *Science* **336,** 1130–1133 (2012).
19. Bauer, C. W. *et al.* Quantum Simulation for High-Energy Physics. *PRX Quantum* **4,** 027001 (2023).
20. Joseph, I. *et al.* Quantum computing for fusion energy science applications. *Physics of Plasmas* **30,** 010501 (2023).
21. Dalzell, A. M. *et al. Quantum Algorithms: A Survey of Applications and End-to-end Complexities* arXiv: 2310.03011 [quant-ph] (Cambridge University Press, 2025).
22. Acharya, R. *et al.* Quantum error correction below the surface code threshold. *Nature* **638,** 920–926 (2025).
23. Gidney, C. Halving the cost of quantum addition. *Quantum* **2,** 74 (2018).
24. Babbush, R. *et al.* Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity. *Physical Review X* **8,** 041015 (2018).
25. Low, G. H. & Chuang, I. L. Hamiltonian Simulation by Qubitization. *Quantum* **3,** 163 (2019).
26. Gilyén, A., Su, Y., Low, G. H. & Wiebe, N. *Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics* in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (2019), 193–204.
27. Su, Y., Berry, D. W., Wiebe, N., Rubin, N. & Babbush, R. Fault-Tolerant Quantum Simulations of Chemistry in First Quantization. *PRX Quantum* **2,** 040332 (2021).
28. Fösel, T., Tighineanu, P., Weiss, T. & Marquardt, F. Reinforcement Learning with Neural Networks for Quantum Feedback. *Physical Review X* **8,** 031084. arXiv: 1802.05267 [quant-ph] (2018).
29. Sivak, V. V. *et al.* Real-time quantum error correction beyond break-even. *Nature* **616,** 50–55. arXiv: 2211.09116 [quant-ph] (2023).

30. Krenn, M., Landgraf, J., Foesel, T. & Marquardt, F. Artificial intelligence and machine learning for quantum technologies. *Physical Review A* **107,** 010101. arXiv: 2208.03836 [quant-ph] (2023).

31. Bausch, J. *et al.* Learning high-accuracy error decoding for quantum processors. *Nature* **635,** 834–840 (2024).

32. Ruiz, F. J. R. *et al.* Quantum circuit optimization with AlphaTensor. *Nature Machine Intelligence* **7,** 374–385. arXiv: 2402.14396 [quant-ph] (2025).

33. Fawzi, A. *et al.* Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* **610,** 47–53 (2022).

34. Mankowitz, D. J. *et al.* Faster sorting algorithms discovered using deep reinforcement learning. *Nature* **618,** 257–263 (2023).

35. Romera-Paredes, B. *et al.* Mathematical discoveries from program search with large language models. *Nature* **625,** 468–475 (2024).

36. Lange, R. T., Imajuku, Y. & Cetin, E. *ShinkaEvolve: Towards Open-Ended And Sample-Efficient Program Evolution* in *International Conference on Learning Representations (ICLR)* (2026). arXiv: 2509.19349 [cs.CL].

37. Alman, J. & Vassilevska Williams, V. *Truly Subquadratic 3SUM and Truly Subcubic APSP via Triangles in Sparse Lopsided Graphs* 2026. arXiv: 2610.06783 [cs.DS].

38. OpenAI. *The Exponential Time Hypothesis is False* https://github.com/openai/math/tree/main/preprints/The-Exponential-Time-Hypothesis-is-False-October-25-2026. Preprint, dated October 25, 2026; AI-generated, not peer-reviewed. Oct. 2026.

39. Anthropic & Claude. *A $2^{o(n)}$ Time Algorithm for 3-SAT* https://github.com/anthropics/formal-math/tree/main/eth. Paper and Lean 4 formalization (Lean v4.33.1, Mathlib v4.33.1); AI-generated, not peer-reviewed. Oct. 2026.

40. Achim, T. *et al. Aristotle: IMO-level Automated Theorem Proving* 2025. arXiv: 2510.01346 [cs.AI].

41. De Moura, L. & Ullrich, S. *The Lean 4 Theorem Prover and Programming Language* in *Automated Deduction – CADE 28* **12699** (Springer, 2021), 625–635.

42. Feng, Y. & Ying, M. Quantum Hoare Logic with Classical Variables. *ACM Transactions on Quantum Computing* **2,** 1–43 (2021).

43. Litinski, D. A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery. *Quantum* **3,** 128 (2019).

44. Google DeepMind. *Formal quantum computation in Lean 4* https://github.com/google-deepmind/formal-quantum-computation. 2026.

45. The mathlib Community. *The Lean mathematical library* in *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs* (2020), 367–381.

46. The Physlib Community. *Physlib: The Lean Physics Library* https://github.com/leanprover-community/physlib. 2024.

47. Meiburg, A. *Quantum Information in Lean* https://github.com/Timeroot/Lean-QuantumInfo. 2024.

48. Rubin, N. C. *et al.* Quantum computation of stopping power for inertial fusion target design. *Proceedings of the National Academy of Sciences* **121,** e2317772121 (2024).

49. Swanson, K., Wu, W., Bulaong, N. L., Pak, J. E. & Zou, J. The Virtual Lab of AI agents designs new SARS-CoV-2 nanobodies. *Nature* **646,** 716–723 (2025).

50. Gottweis, J., Weng, W.-H., Daryin, A., *et al.* Accelerating scientific discovery with Co-Scientist. *Nature* **655,** 487–496 (2026).

51. Lu, C. *et al.* Towards end-to-end automation of AI research. *Nature* **651,** 914–919 (2026).

52. Ghareeb, A. E. *et al.* A multi-agent system for automating scientific discovery. *Nature* **655,** 497–505 (2026).

53. Yuksekgonul, M. *et al. Learning to Discover at Test Time* in *Proceedings of the 43rd International Conference on Machine Learning (ICML)* (2026). arXiv: 2601.16175 [cs.LG].

54. Tsoukalas, G. *et al.* Advancing mathematics research with AI-driven formal proof search. *Science* **394,** 234–239. arXiv: 2605.22763 (2026).

55. Qu, A. *et al. CORAL: Towards Autonomous Multi-Agent Evolution for Open-Ended Discovery* in *Proceedings of the Conference on Language Modeling (COLM)* (2026). arXiv: 2604.01658 [cs.AI].

56. Zhong, Z., Raghunathan, A. & Carlini, N. *ImpossibleBench: Measuring LLMs' Propensity of Exploiting Test Cases* 2025. arXiv: 2510.20270.

57. Lean FRO. *comparator: A trustworthy judge for Lean proofs from untrusted sources* https://github.com/leanprover/comparator. Accessed Nov. 2026. 2025.

58. Pocrnic, M., Loaiza, I., Arrazola, J. M., Wiebe, N. & Motlagh, D. *Efficient Simulation of Pre-Born-Oppenheimer Dynamics on a Quantum Computer* 2026. arXiv: 2602.11272 [quant-ph].

59. Pocrnic, M. & Motlagh, D. *Improved constant factors for qubitized Hamiltonian simulation* 2026. arXiv: 2608.02734 [quant-ph].

60. Low, G. H. *et al.* Fast Quantum Simulation of Electronic Structure by Spectral Amplification. *Physical Review X* **15,** 041016 (2025).

61. Low, G. H. *et al. A denser planar surface code* 2026. arXiv: 2605.30455 [quant-ph].

62. Boyer, M., Brassard, G., Høyer, P. & Tapp, A. Tight Bounds on Quantum Searching. *Fortschritte der Physik* **46,** 493–505 (1998).

63. Zalka, C. Grover's quantum searching algorithm is optimal. *Physical Review A* **60,** 2746–2751 (1999).

64. Gingrich, R. M., Williams, C. P. & Cerf, N. J. Generalized quantum search with parallelism. *Physical Review A* **61,** 052313 (2000).

65. Grover, L. K. & Radhakrishnan, J. *Quantum search for multiple items using parallel queries* 2004. arXiv: quant-ph/0407217.

66. Fang, W. & Heunen, C. *Weakly Measured Loops for Quantum Amplitude Amplification: Oracle Savings and Adaptive Search with Unknown Target Probability* 2026. arXiv: 2609.39621 [quant-ph].

67. Goyal, V. & Sigman, K. On simulating a class of Bernstein polynomials. *ACM Transactions on Modeling and Computer Simulation* **22,** 1–5 (2012).

68. Hentschel, T. W. *et al.* Improving dynamic collision frequencies: Impacts on dynamic structure factors and stopping powers in warm dense matter. *Physics of Plasmas* **30,** 062703 (2023).

69. Hurricane, O. A. *et al.* Physics principles of inertial confinement fusion and U.S. program overview. *Reviews of Modern Physics* **95,** 025005 (2023).

70. Grabowski, P. E., Surh, M. P., Richards, D. F., Graziani, F. R. & Murillo, M. S. Molecular Dynamics Simulations of Classical Stopping Power. *Physical Review Letters* **111,** 215002 (2013).

71. Correa, A. A. Calculating electronic stopping power in materials from first principles. *Computational Materials Science* **150,** 291–303 (2018).

72. Ding, Y. H., White, A. J., Hu, S. X., Certik, O. & Collins, L. A. *Ab Initio* Studies on the Stopping Power of Warm Dense Matter with Time-Dependent Orbital-Free Density Functional Theory. *Physical Review Letters* **121,** 145001 (2018).

73. Grabowski, P. E. *et al.* Review of the first charged-particle transport coefficient comparison workshop. *High Energy Density Physics* **37,** 100905 (2020).

74. Scholes, G. D., Fleming, G. R., Olaya-Castro, A. & van Grondelle, R. Lessons from nature about solar light harvesting. *Nature Chemistry* **3,** 763–774 (2011).

75. Lennox, J. C., Kurtz, D. A., Huang, T. & Dempsey, J. L. Excited-State Proton-Coupled Electron Transfer: Different Avenues for Promoting Proton/Electron Movement with Solar Photons. *ACS Energy Letters* **2,** 1246–1256 (2017).

76. Gentry, E. C. & Knowles, R. R. Synthetic Applications of Proton-Coupled Electron Transfer. *Accounts of Chemical Research* **49,** 1546–1556 (2016).

77. Polli, D. *et al.* Conical intersection dynamics of the primary photoisomerization event in vision. *Nature* **467,** 440–443 (2010).

78. Worth, G. A. & Cederbaum, L. S. Beyond Born–Oppenheimer: Molecular Dynamics Through a Conical Intersection. *Annual Review of Physical Chemistry* **55,** 127–158 (2004).

79. Mai, S. & González, L. Molecular Photochemistry: Recent Developments in Theory. *Angewandte Chemie International Edition* **59,** 16832–16846 (2020).

80. Hammes-Schiffer, S. & Stuchebrukhov, A. A. Theory of Coupled Electron and Proton Transfer Reactions. *Chemical Reviews* **110,** 6939–6960 (2010).

81. Subotnik, J. E. *et al.* Understanding the Surface Hopping View of Electronic Transitions and Decoherence. *Annual Review of Physical Chemistry* **67,** 387–417 (2016).

82. Crespo-Otero, R. & Barbatti, M. Recent Advances and Perspectives on Nonadiabatic Mixed Quantum–Classical Dynamics. *Chemical Reviews* **118,** 7026–7068 (2018).

83. Ren, G., Meng, Q., Zhao, J. & Chu, T. Molecular Design for Electron-Driven Double-Proton Transfer: A New Scenario for Excited-State Proton-Coupled Electron Transfer. *The Journal of Physical Chemistry A* **122,** 9191–9198 (2018).

84. Denisov, I. G., Makris, T. M., Sligar, S. G. & Schlichting, I. Structure and Chemistry of Cytochrome P450. *Chemical Reviews* **105,** 2253–2278 (2005).

85. Guengerich, F. P. Cytochrome P450 and Chemical Toxicology. *Chemical Research in Toxicology* **21,** 70–83 (2008).

86. Rittle, J. & Green, M. T. Cytochrome P450 Compound I: Capture, Characterization, and C–H Bond Activation Kinetics. *Science* **330,** 933–937 (2010).

87. Shaik, S. *et al.* P450 Enzymes: Their Structure, Reactivity, and Selectivity—Modeled by QM/MM Calculations. *Chemical Reviews* **110,** 949–1017 (2010).

88. Goings, J. J. *et al.* Reliably assessing the electronic structure of cytochrome P450 on today's classical computers and tomorrow's quantum computers. *Proceedings of the National Academy of Sciences* **119,** e2203533119 (2022).

89. Caesura, A. *et al.* Faster Quantum Chemistry Simulations on a Quantum Computer with Improved Tensor Factorization and Active Volume Compilation. *PRX Quantum* **6,** 030337 (2025).

90. Gidney, C. *How to factor 2048 bit RSA integers with less than a million noisy qubits* 2025. arXiv: 2505.15917 [quant-ph].

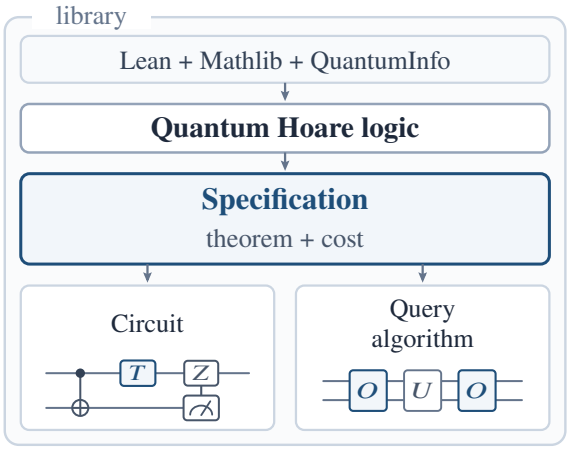

Figure S1: **Verification library and specification stack.** Built on `Mathlib`,[45] `Physlib`'s `Lean-QuantumInfo`[46,47] in Lean, our quantum Hoare logic[42,97] formalizes classical–quantum states, weakest preconditions and resource costs, supporting problem specifications (correctness theorem and cost bound) across dynamic Clifford+T circuits and black-box query algorithms.

A. Formal Quantum Computation Library

We formalize our verification and resource-counting framework for quantum algorithms in Lean,[41] an interactive theorem proving language based on dependent type theory. We chose Lean for its already established rich library ecosystem for broad mathematics, physics, and quantum information theory.[45–47] Moreover, Lean has seen significant traction in recent months for verification of AI-generated proofs. In particular, formalizing quantum algorithms in Lean allows us to model a broader semantics of quantum programming, complementing fast fuzz-testing evaluators like kickmix.[7] Prior formal frameworks for quantum program verification have been developed across several interactive theorem provers. In Rocq (formerly Coq), SQIR[91] verifies unitary quantum circuits, while CoqQ[92] formalizes quantum Hoare logic[93] for pure quantum `while` programs. In Isabelle/HOL, QHLProver[94] provides an earlier mechanization of quantum Hoare logic for imperative quantum algorithms. Concurrent work has also explored quantum verification in Lean,[95] most notably QEDIA,[96] which formalizes classical-quantum Hoare logic[97] for partial and total correctness. Ref. 98 also separates human-reviewed specifications from machine-checked proofs to reduce the human "trust cost" of AI-generated quantum circuits, a division our workflow mirrors. Beyond correctness verification, our framework includes formal support for measuring the worst-case and expected resource costs of algorithms, which is critical for quantum algorithm discovery.

The core of our Lean library is formed by a formalization of a classical-quantum imperative `while` language with an associated classical-quantum Hoare logic.[42] Supporting classical variables and classical control flow is crucial in practice: realistic quantum algorithms routinely interleave quantum operations with classical arithmetic, loops, and measurement-dependent branching. Handling classical states natively avoids the state-space overhead of encoding classical data into quantum registers.

Correctness in this work is expressed through total Hoare triples $\{\!\{\Theta\}\!\}\ P\ \{\!\{\Psi\}\!\}$. These assert that for every initial (classical-quantum) state Δ satisfying the *pre*condition Θ, all executions of the program P terminate in a (classical-quantum) state satisfying the *post*condition Ψ. In the classical case, pre– and postconditions are predicates on the state. In the quantum case, a state no longer simply satisfies or violates an assertion, the condition is quantitative and takes value in $[0, 1]$. We call this the *expectation* of satisfaction. A triple then

asserts that, for every initial state Δ, the expectation of Δ satisfying an assertion Θ is bounded above by the expectation of the resulting state Δ' (obtained after running the program P on Δ) to satisfy Ψ.

Verification is performed via reasoning about *weakest preconditions*. Starting from a given postcondition Ψ, we compute the weakest precondition for this to hold by inductively reasoning about the program backwards. The full semantics of weakest preconditions is given in Figure S3. We follow the source closely but not fully faithfully: we have made certain modifications both to fit our verification purposes and to address ambiguities in the paper. We next outline the main ingredients of the logic and where our formalization differs from Ref. 42.

Variable types. Programs operate on two kinds of variables: classical and quantum. In Ref. 42 classical variables can have two types, boolean and integer. We focus on integer-valued classical variables. This is without loss of generality, since boolean values can be encoded as integers. More precisely, given a type Var of classical variables, the type of classical states (stores) is:

```
abbrev Store (Var : Type) : Type := Var → ℤ
```

For quantum variables, Feng and Ying[42] consider a countably infinite set qVar of quantum variables $q_1, q_2, \ldots$ each associated with a Hilbert space of dimension $d_{q_1}, d_{q_2}, \ldots$ respectively. We restrict our set-up to a single quantum variable q with associated Hilbert space $\mathcal{H}(d)$ of dimension d, where d is a positive natural number. In practice, $\mathcal{H}(d)$ can be identified with $\mathbb{C}^d$. This loses no expressiveness if we limit ourselves to finitely many quantum variables: if these are of dimensions $d_1, \ldots, d_n$ can be combined into one register of dimension $d = \prod_{i=1}^{n} d_i$. Although, this requires lifting operations to the full space explicitly. Supporting multiple named quantum variables is left for future work.

Classical-quantum states. Let $\mathcal{L}(d)$ be the set of Hermitian operators (matrices) on $\mathcal{H}(d)$ and $\mathcal{D}(d) = \{\rho \in \mathcal{L}(d) \mid 0 \preceq \rho,\ \mathrm{Tr}(\rho) \leq 1\}$ the set of partial density operators, ordered by the Löwner order. A *classical-quantum state* Δ is a function $\Delta : \texttt{Store Var} \to \mathcal{D}(d)$ with countable support whose total trace $\sum_\sigma \mathrm{tr}(\Delta(\sigma))$ is at most 1. In particular, any deficit $1 - \|\Delta\|$ represents nontermination. In other words, for a classical store σ, $\mathrm{tr}(\Delta(\sigma))$ is the probability of being in the classical branch σ and $\Delta(\sigma)$, suitably normalized, is the quantum state in that branch. We represent this slightly more generally as:

```
def CQState (State : Type*) (d : ℕ) : Type :=
  {Δ : State → 𝒟 d // ∃ c, HasSum (fun σ ↦ (Δ σ).trace) c ∧ c ≤ 1}
```

and our formalization shows that `CQState (Store Var) d` is equivalent to Definition 3.1 of Ref. 42; countable support need not be separately imposed, since a suitable family of non-negative reals has countable support (covered by `HasSum`). The formalization and proofs of this can be found in `CQState.lean`.[44]

Classical-quantum assertions. We define the set of unit-interval observables[42] (quantum predicates) as $\mathcal{O}(d) = \{M \in \mathcal{L}(d) \mid 0 \preceq M \preceq I\}$, where 0 and I are the zero and identity matrix respectively. A *classical-quantum assertion* is a map $\Theta : \texttt{Store Var} \to \mathcal{O}(d)$. This approach is inspired by classical probabilistic program logics, mirroring the use of by real-valued expectations there.[99] We formalize this (again slightly more generally) as follows, where `Observable d` is our formalization of $\mathcal{O}(d)$:

```
abbrev CQAssertion (State : Type*) (d : ℕ) : Type :=
  State → Observable d
```

As an example, Figure S2 shows how `CQAssertion` is used to express the clean-ancilla property for subroutines that initialize workspace qubits in $|0\rangle$ and uncompute them upon completion.

Similarly to `CQState`, `CQAssertion` has a partial order: we can lift the Löwner order on observables by declaring $\Theta_1 \preceq \Theta_2$, if $\Theta_1(\sigma) \preceq \Theta_2(\sigma)$ for all classical stores σ. Note that since we work with a single quantum

```
def RealizesCleanSubroutine
    {Var : Type} [DecidableEq Var] {n k m : ℕ}
    (rTarget : QRegister k n) (rAnc : QRegister m n)
    (_hdisj : rTarget.DisjointWires rAnc)
    (U : Unitary (2 ^ k))
    (P : QProg Var (2 ^ n)) : Prop :=
  ∀ (M : Observable (2 ^ n)),
    let post : CQAssertion (Store Var) (2 ^ n) :=
      .restrictToZeroAncilla (fun _ ↦ M) rAnc
    {{ post.qAssign (rTarget.embedUnitary U)
        |>.restrictToZeroAncilla rAnc }}
    P
    {{ post }}
```

Figure S2: A formal definition of the clean-ancilla property. Here, a program `P` implements the unitary U on register `rTarget` and returns the disjoint ancilla register `rAnc` to $|0\cdots0\rangle$: for every predicate M, from any input with clean ancillas, `P` must end with clean ancillas and satisfy M at least as often as applying U would. Since M is arbitrary, `P` acts exactly as U on all such inputs.

variable, in contrast to Ref. 42, we need not impose further conditions on the quantum variables used when defining the order. `CQAssertion` is a pointed ω-CPO with respect to the Löwner order: it has a least element, and every increasing chain $\Theta_0 \preceq \Theta_1 \preceq \cdots$ has a supremum. This structure is crucial as it allows us to define the weakest precondition on a `while` loop as the supremum of the weakest preconditions of its finite unrollings and to prove the soundness of the `while` loop rule. In most logics, assertions form a complete lattice, a stronger structure in which every set of assertions has a supremum and an infimum. This matters in practice: existing verification generator frameworks such as Lean's `vcgen` tactic[100] and Loom[101,102] require this. The observables on a Hilbert space of dimension at least 2,[103] do not even form a lattice under the Löwner order, so these frameworks are not presently applicable. Therefore, we implemented the quantum Hoare logic and weakest precondition calculus from scratch.

At first glance, the formalization of `CQAssertion (Store Var)` above differs from Definition 3.5 of Ref. 42, which includes the two conditions on Θ : `Store Var` $\to \mathcal{O}(d)$:

1. The image set Θ(`Store Var`) is countable;
2. For each $M \in \Theta$(`Store Var`), the preimage $\Theta^{-1}(M)$ is definable by a classical assertion.

We do not need to impose condition (1) separately: when the set of classical variables `Var` is finite, `Store Var` = `Var` $\to \mathbb{Z}$ is countable, so the image of Θ : `Store Var` $\to \mathcal{O}(d)$ is automatically countable. We prove this in `CQAssertion.equivCountable`. Condition (2) holds automatically since we represent classical assertions as predicates `Store Var` $\to$ `Prop`. However, if the classical assertions are first-order logic formulas over `Var` as is stated in the beginning of Section 3.2 of Ref. 42, cq-assertions are no longer closed under suprema of increasing chains in the Löwner order and therefore, not an ω-CPO. We encounter the same roadblock to the ω-CPO structure if we explicitly impose condition (1) on the cq-assertions while allowing countably infinitely many classical variables, as opposed to obtaining it naturally by restricting to finitely many classical variables. The full formalization can be found in `CQAssertion.lean`.[44]

Expectation. To reason about program correctness, we need a quantifiable notion of *how much* a cq-state satisfies a cq-assertion. In the quantum case, this is no longer a boolean predicate; instead, we consider the expectation of a cq-state Δ satisfying a cq-assertion Θ:

$$\mathbb{E}[\Delta \models \Theta] \triangleq \sum_{\sigma} \operatorname{tr}(\Theta(\sigma)\,\Delta(\sigma)) \in [0,1]. \tag{1}$$

This can be thought of as the expected satisfaction of Θ in state Δ, averaging over all possible classical branches σ. Note that since we have one quantum variable, our formula is a simplified version of Definition 3.7 in Ref. 42.

We formalize Equation 1 as follows:

```
def expect (Δ : CQState State d) (θ : CQAssertion State d) : ℝ :=
  ∑' σ, ⟪(θ σ).val, Δ σ⟫_ℝ
notation "𝔼[" Δ " ⊨ " θ "]" => expect Δ θ
```

This definition and Equation 1 are equivalent: the inner product of the Hermitian matrices $\Theta(\sigma)$ and $\Delta(\sigma)$ is the real value of the trace of their product as defined in Equation 1. The prime around the sum stands for Mathlib's `tsum`, or topological sum, the limit of finite partial sums of possibly infinite index type. Finally, `expect` $\Delta\ \Theta$ takes values in the unit-interval, even though its return type is the real numbers. This follows from the trace of positive semi-definite matrices being non-negative combined with lemmas we prove in our library establishing the term inside the sum is bounded above by 1 (see theorems `expect_mono`, `expect_top`, `CQState.trace_le_one` in `Expectation.lean`[44]).

Syntax. Programs are expressed in the following variant of the classical-quantum imperative `while`-language of Feng and Ying[42]:

$$\begin{aligned} P \ ::= \ & \texttt{skip} \mid \texttt{abort} \mid x \coloneqq e \mid \texttt{qset} \mid \texttt{apply}\ U \\ & \mid x \coloneqq \texttt{measure}\ \mathcal{M} \mid P_1 \,;\, P_2 \\ & \mid \texttt{qif}\ b\ \texttt{then}\ P_1\ \texttt{else}\ P_2 \mid \texttt{qwhile}\ b\ \texttt{loop}\ P \end{aligned} \tag{2}$$

Our formalization diverges from Ref. 42 in the following ways: we omit random assignment $x :=_{\$} g$, we work with a single quantum system rather than several quantum variables, and we represent expressions e and guards b directly as Lean functions on the classical store. We embed this language in Lean as an inductive datatype `QProg`, with the commands as constructors.

A `QProg` program is parameterized by a finite type `Var` of integer-valued classical variables x and on a single d-dimensional quantum system (d an arbitrary positive natural number, with $d = 2^n$ for n qubits). Classical expressions e : `Store Var` $\to \mathbb{Z}$ and decidable guards b : `Store Var` $\to$ `Prop` are shallowly embedded as functions of the classical store σ, avoiding the need to encode boolean guards as nonzero integers and aligning branch conditions directly with Lean's propositional tactics. The command `qset` resets the quantum system to $|0\rangle$, `apply` U applies a unitary operator U, and $x \coloneqq$ `measure` $\mathcal{M}$ performs an i-outcome quantum measurement with Kraus operators $\mathcal{M} = \{M_o\}_{o<i}$, $\sum_o M_o^\dagger M_o \preceq I$, storing the outcome o in x. The commands `qif` and `qwhile` respectively branch and loop on classical guards. As an example, simple applications of `qwhile` provide support for repeat-until-success loops. Denotational semantics for `QProg` are given in Figure S3.

Weakest Preconditions. Program verification is performed backwards using the weakest precondition (wp) transformer (Figure S3), which takes a program and a target postcondition and computes the weakest assertion that must hold before execution:

```
def wp : QProg Var d → CQAssertion (Store Var) d → CQAssertion (Store Var) d
```

This formalizes the predicate transformer of Table 3 in Ref. 42. Since we are interested in terminating classical-quantum programs, we focus our attention on total correctness. Therefore, we formalize the weakest precondition calculus from Ref. 42 and omit the weakest liberal preconditions, which are only needed for partial correctness.

Particularly, given a classical-quantum program P and a postcondition Ψ, `wp` $P\ \Psi$ computes the "weakest" cq-assertion whose expectation in any initial state Δ equals the expectation of Ψ in the state reached after executing P. Since cq-assertions form an ω-CPO under the Löwner order, the *weakest* precondition is the largest element with respect to the Löwner order for which the desired postcondition still holds.

The wp rules on basic commands are described in Figure S3. They are all faithfully formalized to the source, apart from the syntax and variable-type choices in `QProg`, and can be found in `WP.lean`.[44]

We use the weakest precondition calculus to prove total-correctness Hoare triples $\{\!\{\Theta\}\!\}\ P\ \{\!\{\Psi\}\!\}$, formalized as propositions:

```
def Triple (Θ : CQAssertion (Store Var) d) (c : QProg Var d)
    (Ψ : CQAssertion (Store Var) d) : Prop :=
  ∀ Δ : CQState (Store Var) d, 𝔼[Δ ⊨ Θ] ≤ 𝔼[(interp c).toStateFun Δ ⊨ Ψ]
```

Here `(interp` c`).toStateFun` Δ formalizes the output state $[\![c]\!](\Delta)$ of Figure S3. The central property connecting the forward denotational semantics and the backward wp transformer is exact expectation duality (Lemma 4.14 in Ref. 42):

```
theorem expect_wp (c : QProg Var d) (Δ : CQState (Store Var) d)
    (Ψ : CQAssertion (Store Var) d) :
    𝔼[Δ ⊨ wp c Ψ] = 𝔼[(interp c).toStateFun Δ ⊨ Ψ]
```

Exact expectation duality (`expect_wp`) immediately implies *soundness*: that `wp` $c\ \Psi$ is itself a valid precondition, and that it is indeed the weakest such precondition. These are formalized in `wp_soundness` and `wp_weakest` in `WP.lean`.

Resource Counting and Cost Models. Alongside correctness, our library supports two ways of calculating the resource cost of a program: syntactic gate counting over structured circuits, and semantic expected-cost accounting over program executions with branching and unbounded loops. At the circuit level, straight-line circuits are represented as the inductive type `Circuit` parametrized by the number of wires n. This is built sequentially from nine primitive unitary gates `PrimitiveGate` (Clifford+T, SWAP and TOFF gates). The count of each gate type is obtained simply by summing its occurrences along the sequence. Because each gate type is tracked separately, the cost assigned to each gate is tunable in a specification (for instance, counting only T gates while treating Cliffords as free). Circuits can also include mid-circuit single-qubit measurements and classical branching, packaged in the type `DynamicCircuit` which appears in the problems described in this work. For these, we set the resource counting to have measurements add 0 to gate counts, while the cost of a classical branch is the maximum of its two branch counts, giving a worst-case bound. Through the typing of the output circuit, we also keep track of the number of qubits used to include in resource counts. All of this can be found in the folder `Circuits`, with basic definitions in `Basic.lean`, and the resource counting framework in `ResourceCounting.lean`.

To calculate exact expected resource usage directly on `QProg`, a user-supplied cost model (`CostModel`) assigns nonnegative integer weights w to each primitive command (`assignmentCost`, `resetCost`, `measurementCost`, and `unitaryCost`). Unitary and measurement weights may depend on the operator applied. For instance, one can charge 1 for a designated oracle unitary and 0 for all other operations.

Given w and an initial cq-state Δ with total trace $\|\Delta\| = \sum_\sigma \mathrm{tr}(\Delta(\sigma))$, the branch-weighted expected cost $\mathrm{cost}_w(P,\Delta)$ (`expectedCost`) is defined by structural recursion on P alongside the denotational semantics $[\![P]\!](\Delta)$:

$$
\begin{aligned}
\mathrm{cost}_w(\texttt{skip},\Delta) = \mathrm{cost}_w(\texttt{abort},\Delta) = 0, \qquad \mathrm{cost}_w(c,\Delta) &= \|\Delta\|\cdot w(c),\\
\mathrm{cost}_w(P_1\ ;\ P_2,\Delta) &= \mathrm{cost}_w(P_1,\Delta) + \mathrm{cost}_w(P_2,[\![P_1]\!](\Delta)),\\
\mathrm{cost}_w(\texttt{qif}\ b\ \texttt{then}\ P_1\ \texttt{else}\ P_2,\Delta) &= \mathrm{cost}_w(P_1,\Delta|_b) + \mathrm{cost}_w(P_2,\Delta|_{\neg b}),\\
\mathrm{cost}_w(\texttt{qwhile}\ b\ \texttt{loop}\ P,\Delta) = \bigsqcup_{n\ge 0} C_n(\Delta), \quad C_0(\Delta) = 0,\ C_{n+1}(\Delta) &= \mathrm{cost}_w(P,\Delta|_b) + C_n([\![P]\!](\Delta|_b)),
\end{aligned}
\tag{3}
$$

Program P	Denotational semantics $[\![P]\!](\Delta)$	Weakest precondition $\mathrm{wp}(P, \Psi)$
`skip`	Δ	Ψ
`abort`	0	0
$x \coloneqq e$	$\sigma' \mapsto \sum_{\sigma:\, \sigma[x \mapsto e(\sigma)] = \sigma'} \Delta(\sigma)$	$\sigma \mapsto \Psi(\sigma[x \mapsto e(\sigma)])$
`qset`	$\sigma \mapsto \mathrm{Tr}(\Delta(\sigma))\, \lvert 0\rangle\langle 0\rvert$	$\sigma \mapsto \langle 0\rvert \Psi(\sigma) \lvert 0\rangle\, I$
`apply` U	$\sigma \mapsto U\, \Delta(\sigma)\, U^\dagger$	$\sigma \mapsto U^\dagger\, \Psi(\sigma)\, U$
$x \coloneqq$ `measure` $\mathcal{M}$	$\sigma' \mapsto \sum_o M_o\, [\![x \coloneqq o]\!](\Delta)\, M_o^\dagger$	$\sigma \mapsto \sum_o M_o^\dagger\, \Psi(\sigma[x \mapsto o])\, M_o$
P_1 ; P_2	$[\![P_2]\!]([\![P_1]\!](\Delta))$	$\mathrm{wp}(P_1, \mathrm{wp}(P_2, \Psi))$
`qif` b `then` P_1 `else` P_2	$[\![P_1]\!](\Delta\lvert_b) + [\![P_2]\!](\Delta\lvert_{\neg b})$	$\mathrm{wp}(P_1, \Psi)\lvert_b + \mathrm{wp}(P_2, \Psi)\lvert_{\neg b}$
`qwhile` b `loop` P	$\bigsqcup_{n \geq 0} W_n(\Delta)$	$\bigsqcup_{n \geq 0} \Psi_n$

where $W_{n+1}(\Delta) = \Delta\lvert_{\neg b} + W_n([\![P]\!](\Delta\lvert_b))$, $\qquad \Psi_{n+1} = \mathrm{wp}(P, \Psi_n)\lvert_b + \Psi\lvert_{\neg b}$.

Figure S3: Inductive definitions of the denotational semantics $[\![P]\!](\Delta)$ and weakest precondition $\mathrm{wp}(P, \Psi)$. For a guard b, $\Delta\lvert_b$ retains $\Delta(\sigma)$ where $b(\sigma) \neq 0$ and sets all other branches to 0 (similarly for $\Psi\lvert_b$ and $\neg b$). Loop unrollings start from $W_0(\Delta) = 0$ and $\Psi_0 = 0$.

for primitive commands c, taking values in $[0, \infty]$ so that loop suprema always exist. The total trace, which could also be seen as the probability mass of entering the current branch of the program, is formalized as `stateMass`. Our library provides compositional rules to upper-bound $\mathrm{cost}_w(P, \Delta)$ across sequences, branches, and loops, bounding a `qwhile` loop by its per-iteration cost times the expected number of iterations, which can itself be bounded via `wp`. These definitions and compositional rules on the costs (formalized as `CostBound`) are mechanized in `QuantumHoareLogic/ResourceCounting`.

Each problem specification specializes this framework along two axes. First, the counted resource is chosen to match the computational model, such as gate counts & spacetime, or black-box oracle calls (charging weight 1 per oracle application and 0 for all other operations, as in parallel Grover search; Appendix E). Second, costs can be measured either as concrete numerical bounds at a fixed problem size (like the `arcsin` problem) or as asymptotic bounds. In many cases, these asymptotic bounds are enforced to be of the form $c \cdot N^a$ as the input size N grows, ordered lexicographically on (a, c), which gives a straightforward ordering on discovered solutions.

B. Verification Pipeline

Each problem is defined by a specification file pinned to a frozen repository commit, declaring the target correctness theorem and resource-cost function with typed `sorry` placeholders for the candidate construction, workspace parameters, and proofs (Sec. 1.1). Because agents may game local checks without discovering or improving the algorithm, for instance by weakening theorem statements, modifying helper definitions, adding axioms, disabling local kernel checks, or misreporting costs,[56,104] candidate grading is isolated from the agent environment.

Every submission is replayed in a clean sandbox by an external verifier using the Lean Comparator[57] against the frozen specification. A submission passes only if: (i) all target theorem statements and every transitive definition they depend on match the specification; (ii) the compiled proof terms depend solely on standard Lean axioms (`propext`, `Quot.sound`, and `Classical.choice`), with no `sorry` or custom axioms; and (iii) the proved resource theorem reduces to the claimed closed-form cost. The same check runs automatically before any solution is merged into the shared wiki. Numerical estimates and agent-reported figures are used only internally by teams

to prioritize search directions, whereas researchers need only audit the physical model and cost conventions of the specification itself, revising the pinned specification if an unintended modeling loophole is found.

C. Multi-Agent Orchestration

Agent teams. Each researcher configures and operates an independent team of LLM agents, choosing the agent framework, models, prompts, and division of labor. Setups ranged from a single agent iterating on a candidate construction to multi-agent hierarchical organizations that split subroutine design, error analysis, and Lean proof engineering across specialized workers. In this study, researchers put their domain knowledge, discussion, and design into the problem specification itself rather than seeding domain knowledge or steering the research while underway, though the workflow naturally supports active human seeding and steering as well. During exploration, agents may invoke any external tools, including natural-language derivations, classical simulation, numerical optimization, and remote compute clusters.

Shared wiki. Teams exchange persistent artifacts through a version-controlled repository with one directory per problem.[55,105] Each entry contains: (i) the current best verified solution, which any new run can adopt as its starting point; (ii) reusable partial Lean files verified to compile against the shared library; and (iii) reviewed notes classifying explored idea families as verified, promising, or dead ends, alongside short technical lessons. Whereas solutions and partial Lean files are machine-checked, idea notes are reviewed summaries rather than formal proofs (see Fig. S4). A candidate solution is merged only if it passes automatic re-verification against the frozen specification and clears a minimum relative cost-improvement threshold (5% by default) so that teams do not spend effort adapting to frequent minor updates. Because verified Lean solutions grow large across successive iterations, problems with established solutions also accept updates that substantially shorten the Lean code (by at least 20%) while strictly reducing cost, keeping files within agent context limits and reducing Lean compilation time. Recording negative results with concrete obstructions (such as numerical counterexamples or Lean lemmas) prevents later runs across all teams from repeating known dead ends.

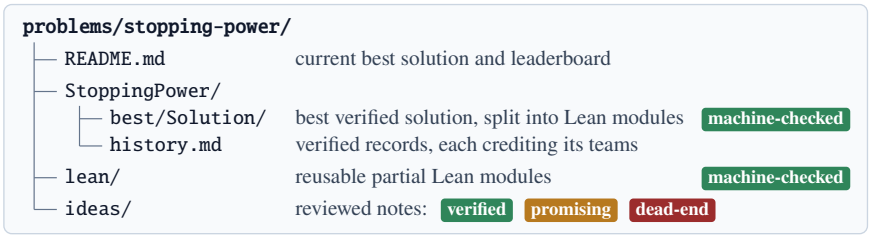

Figure S4: **Layout of one problem directory in the shared wiki (stopping power).** The best solution and the reusable partial Lean modules are checked by the external verifier against the frozen specification before merging, while idea notes are reviewed summaries labeled verified, promising, or dead-end. Each verified record in `history.md` lists the teams credited for it.

Conferences. Because many useful exchanges discuss unverified conjectures, teams convene periodically every four hours on a shared forum open across participating researchers and agent frameworks. During each session, agents pose questions, critique claims, and propose candidate constructions for other teams to test (see Fig. S5). Discussion rules encourage direct replies between teams as opposed to standalone status reports, preventing the one-way broadcasting typical of unstructured agent boards.[106,107] To uphold the correctness of intermediate work,

statements and cost projections posted in a conference carry no weight until independently verified and published to the wiki.

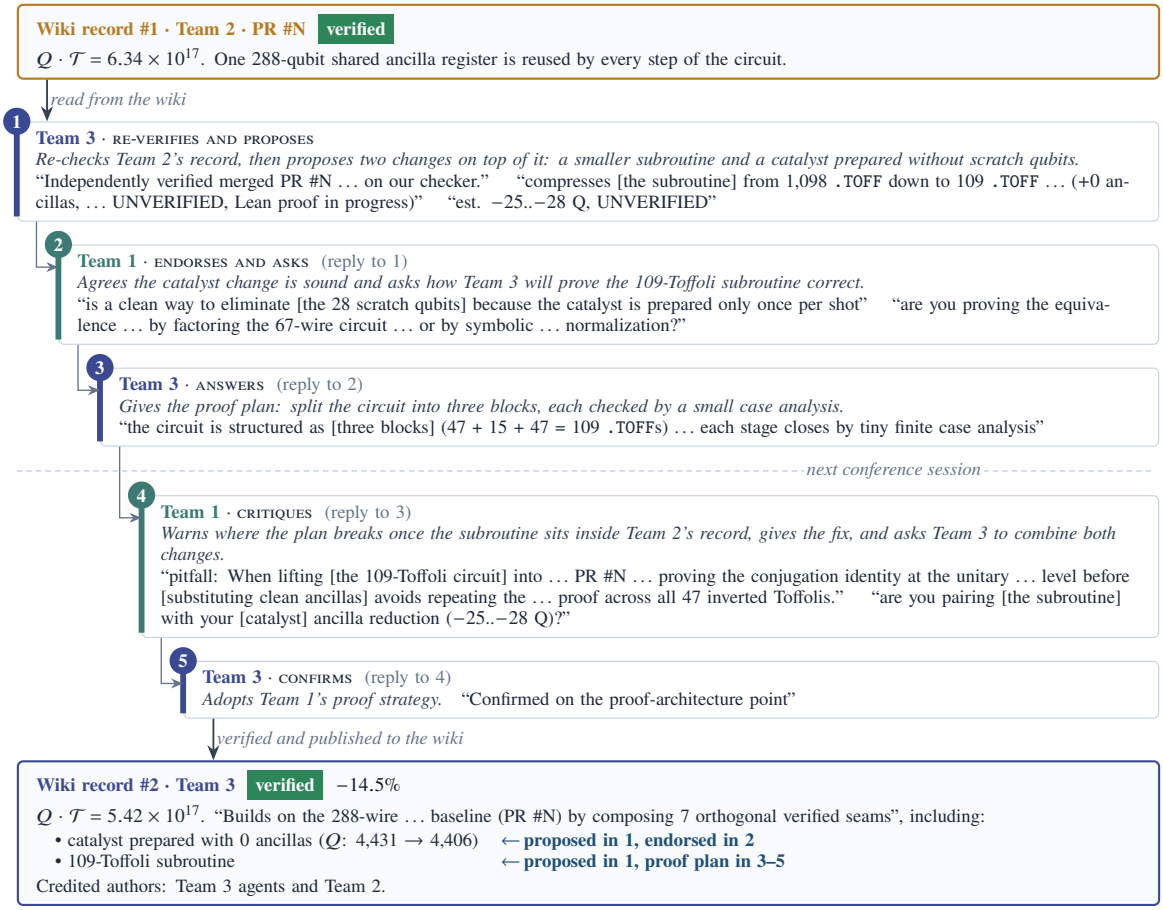

Figure S5: **A cross-team conference exchange between two verified records in the stopping power problem (records #1 and #2 of Fig. 5c).** Each numbered post replies directly to the one above it, across two conference sessions. Colored bars mark the team, the italic line is an analysis summary, and quoted text is verbatim from the conference; ". . ." marks omitted text, and words in square brackets, such as [the subroutine], replace Lean code snippets used in the original posts.

D. Fault-Tolerant Cost Metric and Hardware Horizons

D.1. Logical spacetime volume and architectural assumptions

Throughout the fault-tolerant benchmarks in this work, candidate circuits and literature baselines are compared at matched system encodings and error budgets by their logical spacetime volume $Q \cdot \mathcal{T}$, defined as total peak logical qubits Q (including clean workspace ancillas) multiplied by total T-gate count $\mathcal{T}$, pricing clean-ancilla Toffoli gates at $4\mathcal{T}$ via measurement-based temporary-AND uncomputation.[23,108,109] This metric models compilation onto two-dimensional planar surface-code architectures using lattice surgery.[43,110,111] Because Clifford gates can be tracked in the classical Pauli frame or absorbed into multi-qubit Pauli product measurements, fault-tolerant overhead is dominated by preparing non-Clifford magic states through distillation or cultivation.[43,112–114] Under Pauli-based or serial-factory execution—where magic states are consumed at a fixed logical cycle rate—

Table S1: **Fault-tolerant capability regimes and hardware horizons.** Summary of the capability regimes referenced in Fig. 3b, listing plotted dashed-line midpoints $(Q, \mathcal{T})$, nominal upper-bound ceilings $(Q_{\max}, \mathcal{T}_{\max})$ (rendered in Fig. 3b with 14-step normal-quantile feathering spanning $Q \approx 5.9\times10^2$–1.7×10^3, $\mathcal{T} \approx 3.5\times10^8$–$2.9\times10^9$ for teraquop planar and $Q \approx 5.0\times10^3$–2.0×10^4, $\mathcal{T} \approx 3.5\times10^{10}$–$2.9\times10^{11}$ for FASQ), nominal spacetime volume $Q\cdot\mathcal{T}$, and representative benchmarks and references.

Capability regime	**Midpoint** $(Q, \mathcal{T})$	Q **ceiling**	$\mathcal{T}$ **ceiling**	$Q\cdot\mathcal{T}$	**Benchmarks & references**
Medium-term FTQC (teraquop)	$(10^3,\ 10^9)$	5×10^2–2×10^3	3×10^8–3×10^9	10^{12}	Cytochrome P450; compact catalysis & ECC[7,60,61,88,89,114,115,119]
Application-scale FTQC (FASQ)	$(10^4,\ 10^{11})$	5×10^3–2×10^4	3×10^{10}–3×10^{11}	10^{15}	Stopping power, $C_{23}H_{20}N_3O$ dynamics; catalysis, materials & RSA[1–3,27,48,58,116,120,121,123]

execution time tracks $\mathcal{T}$ while memory footprint tracks Q, making $Q\cdot\mathcal{T}$ proportional to $Q_{\text{logical}}\cdot D_{\text{cycles}}$, the topological spacetime volume.[1,2,43,89,115,116] Optimizing $Q\cdot\mathcal{T}$ also regularizes space–time tradeoffs against degenerate extremes—neither permitting unbounded clean-ancilla QROAM replication to suppress T-gate count[4,24,117] nor forcing deep ancilla-starved recomputation to save a handful of qubits[10,118]—and aligns directly with the $(Q, \text{Toffoli}/\mathcal{T})$ metrics reported across standard resource-estimation benchmarks.[3,27,48,58,60,88,119,120]

Physical device cost departs from $Q\cdot\mathcal{T}$ in two principal architectural regimes. First, provisioning parallel magic-state factories compresses execution depth from $\mathcal{T}$ toward T-depth at the expense of dedicated factory footprint and routing overhead.[1,43,116] Second, architectures incorporating non-local connectivity, high-rate quantum low-density parity-check (qLDPC) memory, or Active Volume compilation decouple idle memory storage from active reaction layers, lowering the spatial cost of inactive data qubits relative to active workspace.[61,89,90,121,122] Because every verified scientific simulation in Fig. 3b reduces both Q and $\mathcal{T}$ relative to its baseline, our improvements hold under any monotone weighting of logical space and non-Clifford-gate count.

D.2. Fault-tolerant capability regimes and hardware horizons

To place end-to-end logical resource counts in hardware context, Fig. 3b overlays two fault-tolerant capability regimes (Table S1) as upper-bound resource envelopes ($Q \le Q_{\max}$, $\mathcal{T} \le \mathcal{T}_{\max}$), with dashed lines at reference midpoints and feathered boundaries reflecting spread across estimates.

Medium-term FTQC (teraquop planar). Gidney *et al.*[115] (Sec. 2.1) define the teraquop scale at 10^{12} reliable logical operations ($\sim 10^3$ logical qubits over $\sim 10^9$ steps at per-step error $\sim 10^{-12}$), and Low *et al.*[61] quantify how denser planar surface-code layouts and magic-state cultivation set the corresponding qubit and T-gate ceilings across planar code distances. This regime sits at midpoint $Q = 10^3$, $\mathcal{T} = 10^9$, and $Q\cdot\mathcal{T} = 10^{12}$, with feathered ceiling $Q \in [5\times10^2,\ 2\times10^3]$ and $\mathcal{T} \in [3\times10^8,\ 3\times10^9]$.

Application-scale FTQC (FASQ / petaquop). Preskill[2] defines the fault-tolerant application-scale quantum (FASQ) regime for full-scale scientific and cryptographic workloads, and Beverland *et al.*[1] (Table I) establish the logical-qubit and T-gate budgets required across materials, chemistry, dynamics, and factoring. This regime sits at midpoint $Q = 10^4$, $\mathcal{T} = 10^{11}$, and $Q\cdot\mathcal{T} = 10^{15}$, with feathered ceiling $Q \in [5\times10^3,\ 2\times10^4]$ and $\mathcal{T} \in [3\times10^{10},\ 3\times10^{11}]$.

E. Parallel Grover Search

E.1. Problem specification and oracle model

Given a search space $\{0,\dots,N-1\}$ of size $N = 2^n$ ($n \ge 1$) and an unknown marked subset $S \subseteq \{0,\dots,N-1\}$ of known size $|S| = k \le N$, we study two exact (zero-error) search problems across $(m, k) \in \{1, 2, 4, \dots, 1024\}^2$

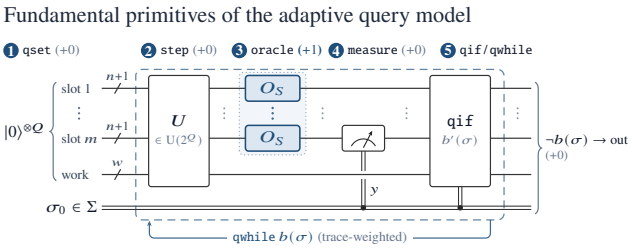

Figure S6: **Parallel Grover query model.** Workspace allocation (`qset`, +0 query cost) initializes m query registers of $n+1$ qubits and w ancilla qubits in $|0\rangle^{\otimes Q}$ together with a classical state $\sigma \in \Sigma$. Unitary evolution (`step`, +0) applies any S-independent unitary $U \in \mathrm{U}(2^Q)$ across all Q qubits. Parallel oracle querying (`oracle`, +1) applies $O_S^{\otimes m}$ (Eq. (5)) simultaneously across m disjoint $(n+1)$-qubit subregisters. Projective measurement (`measure`, +0) measures a subregister into classical variable y in σ while leaving unmeasured wires coherent. Classical control (`qif`/`qwhile`; Appendix A) branches and loops on predicates of σ, weighting each oracle round by the active-branch probability $\sum_{\sigma:\, b(\sigma)} \mathrm{Tr}(\Delta_t(\sigma))$ until $\neg b(\sigma)$ outputs the result.

(121 settings): *Find-Any*, which must terminate with probability 1 and output an index $x \in S$, and *Find-All*, which must terminate with probability 1 and output all k elements of S (so algorithms cannot trade error probability for fewer queries). We allow up to m parallel queries per round[63,124] (Fig. S6). Access to S is given by the $(n+1)$-qubit reversible marking oracle

$$O_S|b\rangle|x\rangle = |b \oplus [x \in S]\rangle|x\rangle, \qquad b \in \{0,1\},\ x \in \{0,\dots,N-1\}, \tag{4}$$

where b is a single-qubit target bit and x is the n-qubit index. Starting from $|0\rangle^{\otimes Q}$ on an unconstrained Q-qubit workspace, an algorithm interleaves zero-cost S-independent unitaries, projective measurements, and classical branching or loops with parallel m-query rounds

$$O_S^{\otimes m} = \bigotimes_{j=1}^{m} O_S^{(j)} \tag{5}$$

on any m disjoint $(n+1)$-qubit registers, weighted in expectation by the probability of being in an active loop branch. To query fewer than m registers in a round, the remaining target qubits are prepared in $|+\rangle$, on which O_S acts as the identity. Our cost metric is the expected number of parallel query rounds (query depth).

Depending on how the target qubit b is prepared (and optionally measured), a single call to O_S implements three operations. Preparing $|b\rangle = |-\rangle$ applies the phase flip $(I - 2\Pi_S)$ for $\Pi_S = \sum_{x\in S} |x\rangle\langle x|$, while preparing $|b\rangle = |0\rangle$ and measuring b tests whether $x \in S$ (and once a subset $F \subseteq S$ has been found, unmarking F via $|b\rangle|x\rangle \mapsto |b \oplus [x \in F]\rangle|x\rangle$ is independent of S and costs zero queries). More generally, preparing $|\chi_\gamma\rangle = \cos\gamma\,|0\rangle - \sin\gamma\,|1\rangle$ for $\gamma \in (0, \pi/4)$ leaves $|\chi_\gamma\rangle$ unchanged for $x \notin S$ while mapping $|\chi_\gamma\rangle \mapsto -\sin(2\gamma)|\chi_\gamma\rangle + \cos(2\gamma)|\chi_\gamma^\perp\rangle$ for $x \in S$, where $|\chi_\gamma^\perp\rangle = \sin\gamma\,|0\rangle + \cos\gamma\,|1\rangle$. Measuring b in $\{|\chi_\gamma\rangle, |\chi_\gamma^\perp\rangle\}$ therefore applies, with $\kappa = \cos^2(2\gamma) \in (0,1)$,

$$M_1 = \sqrt{\kappa}\,\Pi_S \ (\text{outcome } |\chi_\gamma^\perp\rangle), \qquad M_0 = (I - \Pi_S) - \sqrt{1-\kappa}\,\Pi_S \ (\text{outcome } |\chi_\gamma\rangle): \tag{6}$$

a click ($|\chi_\gamma^\perp\rangle$) projects onto S with no false positives, while a no-click ($|\chi_\gamma\rangle$) only slightly damps the marked amplitude and still flips its sign by -1 for the next Grover diffusion step.

To compare algorithms independent of the domain size $N = 2^n$, let $\mathbb{E}[D_n(S)]$ be a program's expected query depth on marked set S, and let $\overline{D}_{n,m,k} = \binom{N}{k}^{-1} \sum_{|S|=k} \mathbb{E}[D_n(S)]$ be its average over all k-element subsets (which equals the worst-case expected depth over $|S| = k$, since indices can be randomly permuted at zero query cost). We normalize $\overline{D}_{n,m,k}$ by the fixed reference scales

$$D_{\mathrm{std}}^{\mathrm{any}}(N,m,k) = \frac{\pi}{4}\sqrt{\frac{N}{mk}}, \quad D_{\mathrm{std}}^{\mathrm{all}}(N,m,k) = \sum_{r=1}^{k} \frac{\pi}{4}\sqrt{\frac{N}{mr}} = \frac{\pi}{4}\,H_k^{(1/2)}\sqrt{\frac{N}{m}}, \quad H_k^{(1/2)} = \sum_{r=1}^{k} \frac{1}{\sqrt{r}}, \tag{7}$$

and certify in Lean a rational constant $R(m,k)$ satisfying $\overline{D}_{n,m,k} \le R(m,k)\, D_{\mathrm{std}}(2^n,m,k)$ for all sufficiently large n. Because D_{std} is only a fixed reference scale rather than a lower bound, all speedups are computed as ratios against published baselines in the same metric, so D_{std} cancels out.

E.2. Published prior baselines

Every quantum search protocol obeys $\Theta(\sqrt{N})$ scaling in domain size,[62,63,125] with parallel round lower bounds $\Omega(\sqrt{N/(mk)})$ for Find-Any (Zalka's bound[63] for $k=1$; the general case follows from the parallel adversary bound[124]) and $\Omega(\sqrt{Nk/(m\min(m,k))})$ for Find-All,[65] so protocols at fixed (m,k) compete on their asymptotic leading constant $r(m,k) = \lim_{N\to\infty} \overline{D}_{n,m,k}/D_{\mathrm{std}}(N,m,k)$. Prior work spans exact and fixed-point amplitude amplification[9,126–129]; early-stopped projective restarts with domain partitioning, independent parallel runs, or sequential strike-out[62–64,130]; and damped or weakly measured loops.[66,131,132] We compare against the strongest published protocol in each (m,k) regime, together with the direct combination of Zalka's partition[63] and weakly measured search[66] on the $k=1$ boundary (Fig. S7). Parallel quantum-walk and learning-graph algorithms[124] target multi-element certificates and do not tighten these baselines.

Boundary regime ($m=1$ or $k=1$; BBHT, ZK, WG). On a single register ($m=1$) with k marked items, each Grover iteration rotates the state toward S by $2\theta = 2\arcsin\sqrt{k/N} \approx 2\sqrt{k/N}$, so j iterations reach angle $\phi = (2j+1)\theta \approx 2j\sqrt{k/N}$ and find a marked item upon measurement with probability $\sin^2\phi$. Standard Grover runs to $\phi = \pi/2$ ($r=1$), where $\sin^2\phi$ flattens out and the final queries buy little success probability. Boyer, Brassard, Høyer and Tapp[62] (**BBHT**) instead stop early at $\phi < \pi/2$ and restart from scratch on failure: each run costs $j \approx \frac{\phi}{2}\sqrt{N/k}$ queries and succeeds after $1/\sin^2\phi$ expected runs, minimizing expected queries $\frac{\phi}{2\sin^2\phi}\sqrt{N/k}$ at

$$r_{\mathrm{BBHT}} = \frac{2\phi^*}{\pi\sin^2\phi^*} \approx 0.878567, \qquad \tan\phi^* = 2\phi^* \;\; (\phi^* \approx 1.1656,\; \sin^2\phi^* \approx 0.845). \tag{8}$$

Across the $m=1$ and $k=1$ boundaries of Find-Any and Find-All (which coincide when $k=1$), r_{BBHT} sets the projective-restart baseline. For $k=1$ and $m\ge 2$ in both tasks, Zalka[63] (**ZK**) splits $\{0,\dots,N-1\}$ into m disjoint blocks of size N/m so the unique block containing the target runs at ϕ^* ($r_{\mathrm{ZK}} = r_{\mathrm{BBHT}}$). For $m=1$ and $k\ge 2$ in Find-All, repeating early-stopped search k times while unmarking each found item at zero query cost[62,130] achieves r_{BBHT} at every stage.

On the single-register boundary ($m=1, k\ge 1$ for Find-Any, and $(1,1)$ for Find-All), early-stopped projective search is superseded by Fang and Heunen's weakly measured loop[66] (**WG**), which avoids discarding the amplified state when a check misses. Instead, **WG** first runs standard Grover halfway to $\alpha^* = \pi/4$ (50% marked weight, costing $\frac{1}{2}D^{\mathrm{any}}_{\mathrm{std}}$), and then alternates calibrated weak checks (Eq. (6), $\kappa \approx 8\sqrt{k/N}$) with single zero-query diffusion steps. Whenever a check does not click, its slight damping and -1 phase flip are offset by the diffusion step to return the state back to $\pi/4$ without losing superposition; waiting for a click adds $\frac{1}{\pi}D^{\mathrm{any}}_{\mathrm{std}}$ expected queries, giving

$$r_{\mathrm{WG}} = \frac{1}{2} + \frac{1}{\pi} \approx 0.818310. \tag{9}$$

For $k=1$ and $m\ge 2$ in both tasks, running **WG** inside each of Zalka's m disjoint blocks of size N/m (**ZK+WG**) is a direct substitution—the $m-1$ empty blocks never click while the unique target block evolves at scale $\sqrt{N/m}$—upgrading the $k=1$ boundary from $r_{\mathrm{ZK}} \approx 0.879$ to $r_{\mathrm{WG}} \approx 0.818$.

Interior Find-Any ($m,k\ge 2$; GW). When both $m\ge 2$ and $k\ge 2$, naive domain splitting breaks down because each of the m blocks receives a random, unknown number of targets (and most blocks are empty when $k<m$), causing multi-target blocks to rotate faster and overshoot $\pi/2$ under any fixed schedule. Gingrich, Williams and

a Published baselines and verified bounds

		Find-Any		Find-All	
Reference	**Evaluation in metric**	**Regime**	**Rel. rounds**	**Regime**	**Rel. rounds**
Boyer *et al.*[62] (**BBHT**, Eq. (8))	Early-stop restart + strike-out	—	—	m=1, $k \geq 2^\star$	0.879
Zalka[63] (**ZK**, Eq. (8))	m-block partition + early stop	$m \geq 2$, k=1	0.879	$m \geq 2$, k=1	0.879
Gingrich *et al.*[64] (**GW**, Eq. (10))	Unpartitioned m-slot early stop	$m, k \geq 2^\star$	0.942–0.997	—	—
Grover & Radhakrishnan[65] (**GR0**)	Two-pass m-block staircase (t_1^*)	—	—	$m, k \geq 2^\star$	0.060–0.955
Fang & Heunen[66] (**WG**; +**ZK**, Eq. (9))	Calibrated weak loop (+**ZK** at k=1)	m=1 or k=1*	0.818	k=1*	0.818
This work (Lean-verified circuits)		all	0.818	all	0.024–0.818
Speedup vs. baseline envelope (Fig. 3d)		$m, k \geq 2$	up to **1.22**×	$k \geq 2$	up to **2.57**×

b Finding any one target

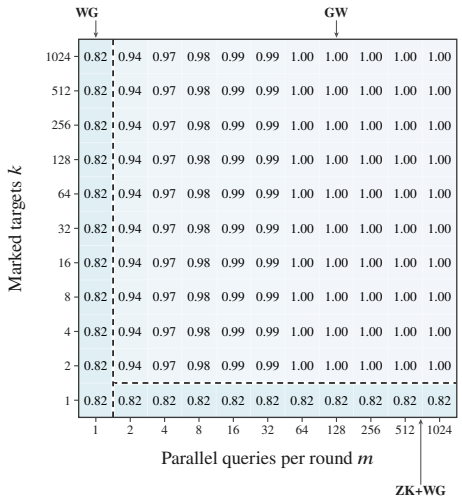

c Finding all k targets

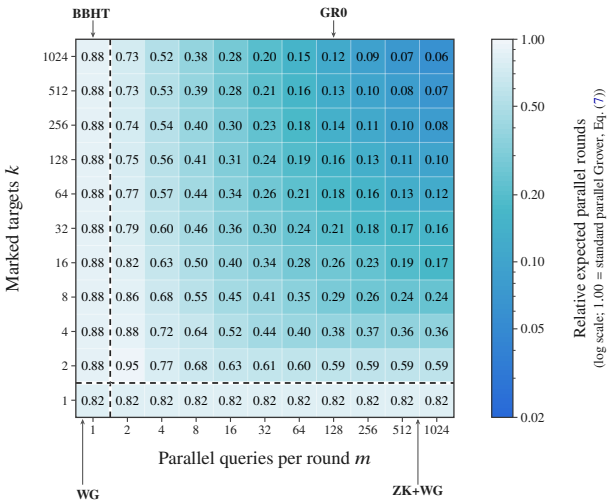

Figure S7: **Published parallel Grover baselines and verified bounds across** $m, k \in \{1, 2, 4, \ldots, 1024\}$**. a** Each row evaluates a prior baseline protocol in our zero-error asymptotic metric (* marks the winning baseline in each regime); the bottom rows summarize our Lean-verified circuits and speedups over the baseline envelope. Values report asymptotic expected parallel rounds as $N = 2^n \to \infty$ relative to D_{std} (Eq. (7)). **b, c** Winning baseline envelopes for finding any one (**b**) or all k (**c**) targets, on a logarithmic colour scale, with dashed contours marking regime boundaries.

Cerf[64] (**GW**) avoid this unknown-load problem by running all m registers independently over the full domain $\{0,\dots,N-1\}$. Stopping all m registers synchronously after $j \approx \frac{\phi}{2}\sqrt{N/k}$ rounds finds a target unless all m independent measurements miss (probability $\cos^{2m}\phi$); minimizing expected rounds $\frac{\phi}{2(1-\cos^{2m}\phi)}\sqrt{N/k}$ selects $\phi^*(m) < \phi^*$ with

$$r_{\mathrm{GW}}(m) = \frac{2\sqrt{m}\,\phi^*(m)}{\pi\big(1-\cos^{2m}\phi^*(m)\big)}, \qquad 1-\cos^{2m}\phi^*(m)\big(1+2m\phi^*(m)\tan\phi^*(m)\big) = 0. \tag{10}$$

Because m unentangled copies combine success probabilities classically rather than coherently, $r_{\mathrm{GW}}(m)$ is independent of k and degrades from 0.942 at $m = 2$ to 0.997 at $m = 1024$, setting the published baseline on all 100 interior cells (Fig. S7b).

Interior Find-All ($m, k \geq 2$; GR0). To let m registers find distinct targets in parallel at $\sqrt{N/m}$ block scale, Grover and Radhakrishnan[65] randomly partition $\{0,\dots,N-1\}$ into m blocks of size N/m. Each block receives an unknown target count c, which they clear via a *descending staircase* of exact amplitude-amplification stages[128] at assumed counts $s = t, t-1, \dots, 1$ (total depth $H_t^{(1/2)}\frac{\pi}{4}\sqrt{N/m}$), testing and unmarking hits after each stage. Counting *down* guarantees emptying every block with $c \leq t$: even if stages $s > c$ miss, once s reaches a block's remaining count c, the assumed and actual counts match, finding a target with certainty and stepping $c \to c-1$ in lockstep with s (while mismatched stages $s \neq c$ still hit with probability $\sin^2(\frac{\pi}{2}\sqrt{c/s})$). Because Theorem 3 of ref. 65 leaves tail cleanup at $\Theta(\cdot)$ order, we evaluate a concrete two-pass schedule (**GR0**): Pass 1 runs a staircase from the optimal integer cap $1 \leq t_1^*(m,k) < k$ down to 1, and Pass 2 clears all $R_1 = k - |\mathrm{found}|$ remaining targets via a second staircase from R_1 down to 1 (Fig. S7c).

E.3. Algorithmic mechanisms in verified circuits

Beyond the $k = 1$ boundary (**ZK+WG**), our interior baselines are single-paper protocols; Fig. S8a separates *automatic recombinations*, which join unchanged published subroutines by new classical control, from *new coherent mechanisms*, which alter the quantum dynamics or measurements within a search stage, listing a few of the many mechanisms found per class. Since mechanisms interact, Fig. S7a reports only joint speedups, while Fig. S8b,c details the verified relative round costs and winning circuit regimes across (m, k).

Unlike distributed baselines such as **GW** and **GR0** that run m uncoupled registers, our shared-workspace model allows inter-slot entanglement: coupling all m slots via odd-parity fanout in Find-Any extends the boundary constant r_{WG} to every interior $(m, k \geq 2)$, while Find-All entangles slot and storage registers. In our strongest Find-All mechanism ($k \lesssim m$; Fig. S9), three ideas avoid the bin-load variance of partitioned search: (i) preparing a uniform superposition over m-element subsets $|E\rangle$ so one $|-\rangle$-marked call to $O_S^{\otimes m}$ (unmarking archived hits F) supplies all target information via the hit-parity phase $(-1)^{|E\cap(S\setminus F)|}$; (ii) alternating that query with a fixed, S-independent reflection on the Johnson-scheme eigenspaces of m-subsets[124]—analogous to Grover's uniform diffusion step and compiled via quantum singular value transformation (QSVT),[26] plus a terminal phase kick, driving the hit count $w = |E \cap (S \setminus F)|$ coherently from 0 to $r = |S \setminus F|$ in under $\sqrt{N/m}$ rounds; and (iii) batch readout, where one final query writes all $w \in [w_{\mathrm{lo}}, r]$ hits into $|F\rangle$ at once, retrying or clearing any remainder via exact loops. Although S-independent unitaries carry zero query cost in the black-box specification, our verified circuits use $O((m+k)n)$ qubits and each unitary step compiles into $\mathrm{poly}(n, m, k)$ elementary gates (within error $2^{-\Omega(n)}$ for the QSVT-compiled diffusion steps).

a Selected mechanisms in verified circuits

Mechanism	Builds on	Regime	Description
Automatic recombinations			
Weak-loop pipelining	62, 64, 66, 132	Find-All, $k \gtrsim m$	Full-domain slots run weakly measured loops; a click strikes out its target and restarts only that slot, while the others keep their progress.
Sibling sub-block re-search	62, 63, 65	Find-All, $k \approx \sqrt{m}$	A target missed by a partitioned sweep most likely shares a block with a found one; only those blocks are re-searched, split across all slots.
Re-partitioned sweeps	62, 65, 66	Find-All, $k < m$	Under-rotated sweeps over fresh (bit-swapped) partitions stop once all targets are found; a weakly measured loop finds the last one.
New coherent mechanisms			
Odd-parity joint reflection	63, 66	Find-Any	Conjugating $O_S^{\otimes m}$ by marker-qubit gates gives a reflection about slot tuples holding an odd number of targets.
Tensor-power conveyor	9	Find-All, $k \gg m$	A zero-query rotation repeatedly moves found amplitude from each query register into $d-1$ storage registers.
Collective Johnson walk	26, 63, 124	Find-All, $m \geq 4$	Alternating odd-parity m-subset queries with S-independent Johnson-graph reflections (compiled via QSVT) drives the hit count from 0 to k in one pass.
Slice zoom	63, 65	Find-All, $k \leq 8$	Restricting the search to target-independent rotated slices of the domain zooms in on the few marked items, so later rounds search a much smaller effective list.

b Finding any one target

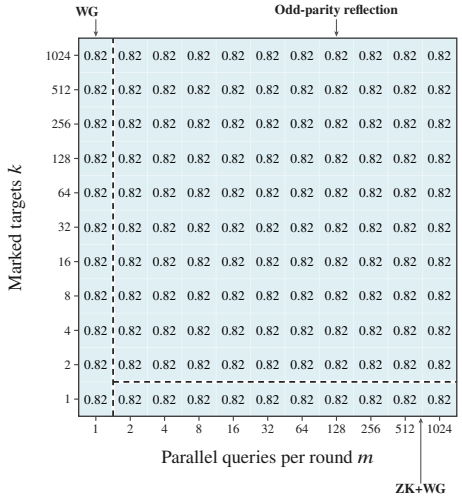

c Finding all k targets

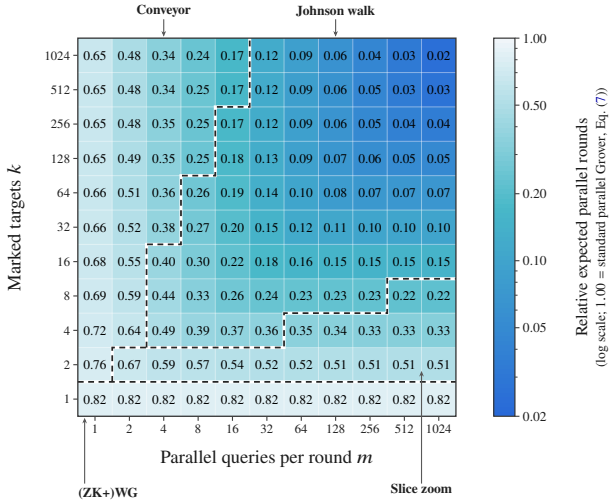

Figure S8: **Verified parallel Grover mechanisms and relative round costs across $m, k \in \{1, 2, 4, \ldots, 1024\}$.** **a** Selected algorithmic mechanisms in our Lean-verified circuits (not exhaustive). **b, c** Verified asymptotic expected parallel rounds relative to D_{std} (Eq. (7)) for finding any one (**b**) or all k (**c**) targets, on a logarithmic colour scale, with dashed contours and labels marking the winning mechanism in each region (cf. panel **a**; **WG** and **ZK+WG** denote ties with the baseline).

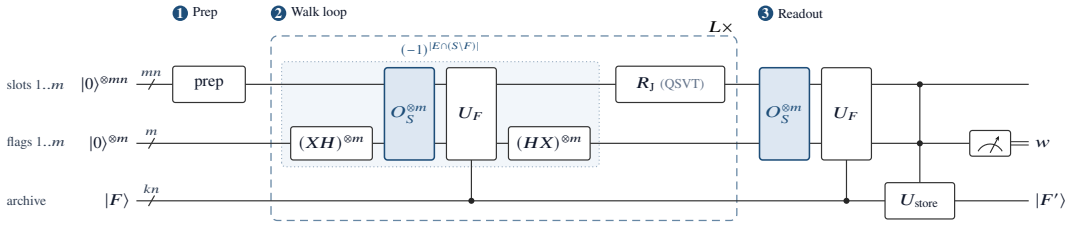

Figure S9: **High-level circuit for the collective Johnson-graph walk.** Registers hold m address slots ($|0\rangle^{\otimes mn}$), m flag qubits ($|0\rangle^{\otimes m}$), and a kn-qubit archive $|F\rangle$ of found targets. (1) prep prepares the uniform superposition $\binom{N}{m}^{-1/2}\sum_{|E|=m}|E\rangle$ over m-subsets $E\subseteq\{0,\ldots,N-1\}$. (2) For $L<\sqrt{N/m}$ rounds ($L\times$), $(XH)^{\otimes m}$ (X then H) and $(HX)^{\otimes m}$ switch the flags between $|0\rangle^{\otimes m}$ and $|-\rangle^{\otimes m}$; $O_S^{\otimes m}$ is the parallel marking oracle; U_F unmarks slots already in F, so the dotted block applies $(-1)^{|E\cap(S\setminus F)|}$; and R_J is an S-independent Johnson-graph reflection compiled via QSVT. (3) After one call to $O_S^{\otimes m}$ and U_F, U_{store} records new hits into $|F'\rangle$ and measuring the flags outputs the hit count $w=|E\cap(S\setminus F)|$.

F. Fixed-Point arcsin

F.1. Problem and specification

Many fault-tolerant quantum algorithms require converting amplitudes held in a quantum register into single-qubit rotation angles, including probability-distribution loading and state preparation,[133] eigenvalue inversion in quantum linear-system solvers,[134] and quantum rejection and Metropolis sampling.[135,136] Implementing these rotations via phase kickback or controlled Pauli gates first evaluates $\arcsin(y)$ reversibly on a fixed-point register, a subroutine that often dominates the non-Clifford gate count.[10,137]

We consider reversible evaluation $|y\rangle|b\rangle|0\rangle^{\otimes m}\mapsto|y\rangle|(b+z(y)) \bmod 2^k\rangle|0\rangle^{\otimes m}$ on $k=25$ signed two's-complement registers, where y is the input register, b is the output register, and $z(y)$ approximates $\arcsin(y)$. Because $\frac{\mathrm{d}}{\mathrm{d}y}\arcsin(y)=(1-y^2)^{-1/2}$ diverges as $|y|\to 1$, uniform polynomial approximations require high degree near the endpoints, while iterative shift-and-add methods accumulate deep chains of reversible multiplications.[10,138] Our signed two's complement representation (1 sign, 1 integer and 23 fractional bits) matches the `float32` mantissa size. With 23 fractional bits, the smallest representable value is $2^{-23}\simeq 1.19\times10^{-7}$. Allowing for some small rounding error, we require a maximum absolute error $|z(y)-\arcsin(y)|\le 1.25\times10^{-7}$ across all $2^{24}+1=16{,}777{,}217$ inputs $y\in[-1,1]$. Using m clean ancillas, a candidate program in Clifford+T gates and mid-circuit measurements must apply this exact action across all 2^{50} basis states (y,b) and return all ancillas to $|0\rangle^{\otimes m}$.

F.2. Baseline

We consider as baseline the algorithm by Häner et al.,[10] which uses a piecewise-polynomial approach.[1] We compare against their $(n,m,d)=(35,3,3)$ fixed-point two's-complement configuration, whose maximum error $\sim 10^{-6}$ is driven by the error in the square root computation (Fig. 6 of ref. 10), and larger than our target bound of 1.25×10^{-7}, thus providing a lower bound on the spacetime cost. As shown in Fig. S10, this baseline splits the computation of arcsin on the interval $[-1,1]$ between a degree-3 polynomial evaluation on $[-0.5,0.5]$ and a Newton iteration for $\sqrt{z}=z\cdot 1/\sqrt{z}$.[2] As their Appendix D describes, the former can be implemented using 140 qubits and requires 17,304 Toffoli gates and the same amount for reversing the computation, for a total of 34,608 Toffoli gates (138,432 T gates). The latter can be implemented using 245 qubits and requires 73,442 Toffoli

[1] This approximation is stronger than the quantum CORDIC paper.[138] While their paper does not report results for our input type, we ran the code from their github repository and got a maximum error of 9×10^{-3}.

[2] Our baseline requires $\sqrt{z}=z/\sqrt{z}$. A recent circuit[139] computes $1/\sqrt{x}$ to 2^{-24} relative error in 7,735 Toffoli gates and 260 qubits. Using this here would slightly increase the total number of qubits but significantly reduce the number of gates. The number of gates for the square-root could be cut by a factor of up to 10, and the baseline[10] cost by up to 3×.

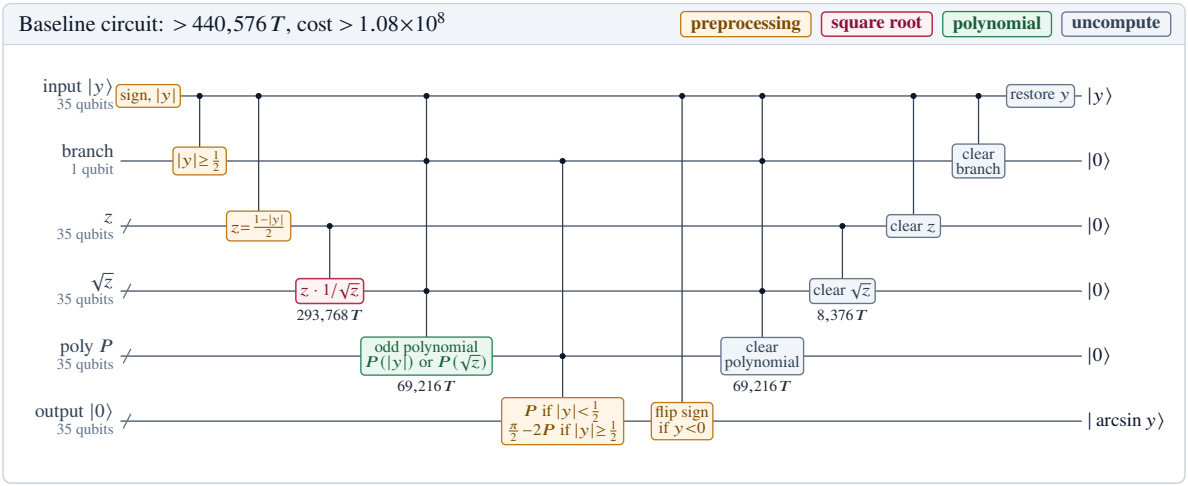

Figure S10: Baseline circuit to implement the arcsin function,[10] using $n = 35$, $m = 3$, $d = 3$. After restricting to the positive range, the approximation splits the calculation of arcsin between a degree-3 polynomial in the interval $[0, 0.5]$ and a Newton iteration for the square root function in $[0.5, 1]$. Number of T gates and number of qubits calculated as in their Appendix, for reversible computation.

gates to implement it as the product of z with $1/\sqrt{z}$, as described in their Appendix C.2, and an additional 2,094 Toffoli gates to also reverse the product operation, being $1/\sqrt{z}$ already a reversible operation, for a total of 75,536 Toffolis (302,144 T gates).

We calculate a lower bound of the full circuit spacetime cost of 1.08×10^8 by using the maximum number of required qubits in the two subcircuits, 245, and the sum of the T gates required for them, 440,576 T gates. We thus ignore persistent registers and the smaller branch-selection and combination blocks, which would further increase the cost.

F.3. Program Discovery

While our library lets us represent arbitrary gates and formally verify the cost of algorithms, the *kickmix* model[7] restricts circuits to classical reversible gates (X, CNOT, Toffoli and CSWAP), Pauli-X measurement-based uncomputation,[23] and diagonal Z, CZ and CCZ phase corrections so that every basis input stays a single signed basis state $(-1)^{\phi}|x\rangle$ along each measurement branch, enabling fast classical fuzz-testing without statevector simulation. The fact that agentic systems can produce compact arithmetic quantum circuits by using the feedback from kickmix testing is already known,[118,139] and has also been applied to a related $\arcsin(c/X)$ subroutine, but on a restricted input range rather than $[-1, 1]$. We exploited this speed by first running our multi-agent system to search for circuits using kickmix. The result was a circuit of spacetime cost 498,168 (102 qubits, 4,884 T gates).

However, the kickmix gate set is limited, and we expect that the full Clifford+T gate set could allow for a more efficient implementation. To obtain more general algorithms that are not limited by the kickmix restrictions, we need a different way to verify the circuit, and this is where our formal library helps. Starting from the algorithm obtained with kickmix, we instructed agents to start using our quantum algorithm specification and verification framework to further optimize the algorithm. For a while, agents were able to continue optimizing the circuit up to a cost of 415,200 (100 qubits, 4,152 T gates, Fig. S11a and Table S2), which continued using a restricted set of gates, and *could* still be translated to kickmix. Subsequently, the agents truly took advantage of our library and produced circuits that *could not* be expressed and efficiently simulated with kickmix (Fig. S11b,c). This drove the cost down by a further 11% to 369,800 (100 qubits, 3,698 T gates), 290× below the baseline[10] (Fig. 3a,c).

Table S2: **Summary for fixed-point** $\arcsin(y)$ **on** $[-1, 1]$ **at** $k = 25$**.** The baseline evaluates the $(n, m, d) = (35, 3, 3)$ configuration,[10] assuming $4\,T$ per Toffoli.[23,108] We derive a lower bound on its spacetime cost by summing the two dominant reversible submodules on the 245-qubit square-root workspace while neglecting persistent registers and the steps that combine the two parts. In the non-kickmix Lean 4 circuit, the 3,698 T gates comprise 776 four-T temporary-AND gates ($3{,}104\,T$) plus 594 single-qubit $T/T^\dagger$ gates across 41 inner-product gadgets.

Method	ε	**Toffoli**	$\mathcal{T}$	Q	**Spacetime cost** $(Q \cdot \mathcal{T})$
Baseline[10] (lower bound)	$\sim 10^{-6}$	110,144	440,576	245	1.08×10^8
– arcsin on $[-0.5, 0.5]$		34,608	138,432	—	—
– $1/\sqrt{z}$		75,536	302,144	—	—
Kickmix	1.25×10^{-7}	1,038	4,152	100	415,200
Non-kickmix (Lean 4)	1.25×10^{-7}	776	3,698	100	**369,800**

In other words, for this problem, our agents using kickmix can do an efficient search in the restricted subspace without the overhead of a formal proof. Afterwards, our agents using formal verification lift the gate-set restriction and allow further optimization that would not be easily testable classically. Every circuit that can be simulated with kickmix can also be proven in Lean, and seeding a Lean discovery agent from a kickmix solution further improves the solution.

F.4. Interpretation

Fig. S11 provides an interpretation of the algorithms that we obtained, and shows where the optimization using the Lean verifier obtained an improvement.

Since arcsin is an odd function, i.e. $z(-y) = -z(y)$, we can perform the calculation only for positive inputs, and flip the output register according to the sign. Subsequently, we also split off the linear part, $z(|y|) = |y| + r(|y|)$ and approximate the remainder $r(|y|) \approx \arcsin(|y|) - |y|$ using a separate quadratic polynomial on 440 intervals of $[0, 1]$. Each interval is chosen to have a width roughly proportional to its distance to 1, so on each interval the square root is smooth, a second order polynomial fits well, and the fit error shrinks when getting closer to 1. The intervals get narrower toward $|y| = 1$, where $\arcsin(|y|) \approx \frac{\pi}{2} - \sqrt{2(1 - |y|)}$ grows steeply. The closest input to 1 on the fixed-point grid is 2^{-23} away, so finitely many intervals cover all inputs. The point $|y| = 1$ is handled separately. Subsequently, the circuit looks for the interval that contains $|y|$ and stores its index in a 9-bit register s. It does so by using a barrel shift, i.e. it counts the leading zeros of $1 - |y|$, and shifts the magnitude bits left using controlled swaps to form the local offset t. A lookup block then computes the interval's coefficients A, B, C for s, and squarer and adder circuits add $A + Bt + Ct^2$ to b. As an optimization, the 19 high bits of A are added to b and erased so that the remaining coefficient bits can be computed on the same qubits. Similarly, C and t^2 are erased midway so their qubits can be reused for the rest of $+Bt$. Measurements with Z and CZ corrections erase the coefficients and all other scratch qubits at zero T cost. Finally, the circuit is made reversible by undoing the shift and the sign conversion, restoring y and returning all ancillas to $|0\rangle$.

As shown in Fig. S11, a large amount of the total kickmix T gates are spent on the coefficient lookup phase. Specifically, each coefficient bit is a fixed function of the interval number s, given by a table with one entry per interval. Any such table can be written as an XOR of products of bits of s, e.g. $s_1 s_4 \oplus s_2 s_3 s_7 \oplus s_5$. The kickmix circuit first builds a few products that many coefficient bits share, and then builds each coefficient bit by XOR-ing in the products it needs using a $4\,T$ AND circuit for each product.

The non-kickmix circuit replaces most of these bitwise product blocks with 41 inner-product circuits that need roughly half as many T gates. Indeed, each block puts its coefficient qubit in an equal superposition of 0 and 1 states, writes the products as a relative phase between the two branches using T and $T^\dagger$ gates, and then makes the branches interfere to get the appropriate bit. However, each inner-product circuit leaves a sign ± 1 that depends on the specific s, which is cancelled across each lookup pass with Z and CZ gates at zero T cost. Because each inner

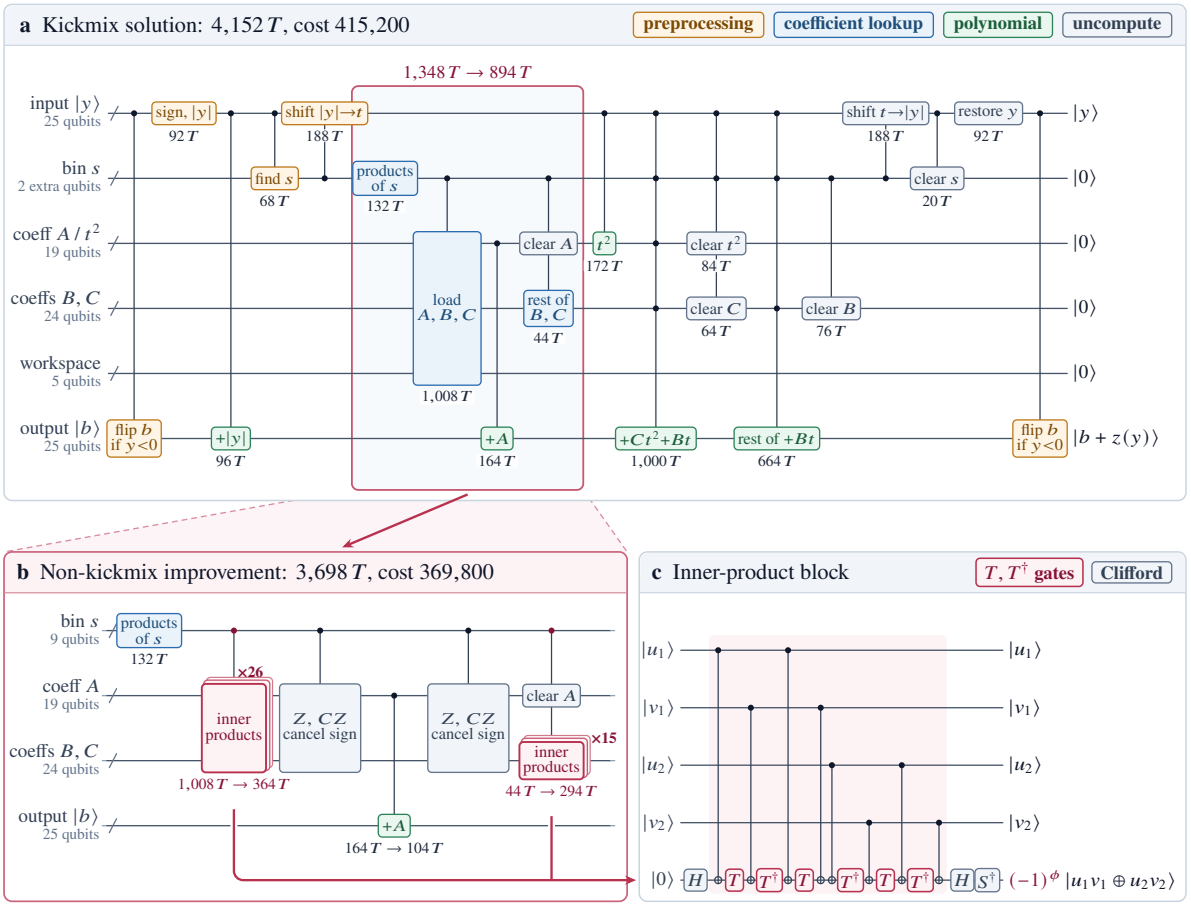

Figure S11: Sketch of the algorithm to implement the fixed-point arcsin module. (a) Full kickmix circuit. It uses 50 data qubits ($|y\rangle$ and $|b\rangle$, 25 each) and 50 ancilla qubits reused across stages, and a total of 4,152 T gates. The bin index s has 9 bits but needs only 2 extra ancillas. After "shift $|y| \to t$", 6 bits of s are simply the top 6 bits of t, which identify the bin within its group. A 7th bit replaces the $\frac{1}{2}$ bit of $|y|$, which can be recomputed from s and is erased. The remaining 2 bits use one ancilla each: one is stored directly, and the other replaces the top bit of $|y\rangle$, whose value is moved into an ancilla. t^2 is computed into the ancillas freed by clearing A. At peak, the ancillas hold A (19), B and C (24), 5 lookup work qubits, 1 bit of s, and the moved bit of $|y\rangle$. (b) Improved block in the non-kickmix circuit. Algorithm verification with our formal quantum computation library allows further optimization of the circuit. Specifically, the coefficient lookup block can be replaced with a more efficient block, saving overall 454 T gates. Each coefficient bit is computed by one inner-product block taking a variable number $r \in \{2, \ldots, 7\}$ of product pairs and using $2r + 2$ T gates. We have 26 blocks totaling 364 T gates before $+A$, and 15 blocks totaling 230 T gates on the freed qubits after clearing A. Each inner-product block leaves an extra sign phase that is then removed before returning the result. The whole circuit requires now 3,698 T gates instead of 4,152 T gates. (c) Example of the optimized inner-product block for two product pairs, using 6 T gates, which could not be expressed in the kickmix framework.

product writes directly into its target coefficient qubit and needs no extra scratch qubits, the circuit can postpone more coefficient bits of B and C until after A has been added to b and cleared, leaving more spare qubits for that adder and reducing the lookup-and-add stage from 1,348 to 894 T gates. Together, these changes lower the full circuit from 4,152 to 3,698 T gates on the same 100 qubits.

G. Stopping Power in Inertial-confinement Fusion

G.1. Hamiltonian and specification

Following Eqs. (1)–(7) and Tables 1–2 of ref. 48, the system consists of $\eta = 218$ electrons ($\zeta_e = -1$, mass 1) and one α-particle projectile ($\zeta_{\text{proj}} = +2$, mass $M_{\text{proj}} = 7{,}294.29954171\, m_e$, matching 3.9726 proton masses in Table 2 of ref. 48) in a periodic cubic cell of volume $\Omega = 2{,}419.68282\ a_0^3$ containing 216 fixed protons at arbitrary positions $R_\ell \in \mathbb{R}^3$. In plane-wave first quantization, each electron momentum coordinate is stored in $n_p = 6$ signed-magnitude qubits and each projectile coordinate in $n_n = 8$ signed-magnitude qubits,[27,48] giving $Q_{\text{sys}} = 3(\eta n_p + n_n) = 3{,}948$ system qubits. The valid grid subspace excludes the signed-magnitude negative zero $|1, 0\cdots0\rangle$ from every coordinate, leaving integer momenta in $G = [-31, 31]^3$ for each electron (63^3 plane waves, covering the 53^3 classical time-dependent density-functional grid of Table 1 in ref. 48) and $\tilde{G} = [-127, 127]^3$ for the projectile, at wavevectors $k_\nu = 2\pi\nu/\Omega^{1/3}$. The projectile's momenta are taken relative to its initial mean momentum k_{proj}, with $k_{\text{proj}}\Omega^{1/3}/(2\pi) \in \mathbb{Z}^3$ for compatibility with the cell's periodicity and $\|k_{\text{proj}}\|_2 \le 29{,}376$ a.u. ($v_\alpha \approx 4$ a.u.).

On valid grid states, the discretized Hamiltonian in atomic units is

$$\begin{aligned} H = &\sum_{i=1}^{\eta} \sum_{p\in G} \frac{\|k_p\|_2^2}{2} |p\rangle\langle p|_i + \sum_{p\in\tilde{G}} \frac{\|k_p - k_{\text{proj}}\|_2^2}{2M_{\text{proj}}} |p\rangle\langle p|_{\text{proj}} + \sum_a \sum_{\substack{p,q\in G_a \\ p\neq q}} \zeta_a S(p-q) K(p-q) |p\rangle\langle q|_a \\ &+ \frac{1}{2} \sum_{a\neq b} \sum_{\nu\neq 0} \sum_{\substack{p\in G_a,\, q\in G_b \\ p+\nu\in G_a,\, q-\nu\in G_b}} \zeta_a \zeta_b K(\nu) |p+\nu\rangle\langle p|_a |q-\nu\rangle\langle q|_b, \end{aligned} \tag{11}$$

and acts as zero outside the valid grid subspace. Here a and b range over the η electrons ($G_a = G$, $\zeta_a = -1$) and the projectile ($G_{\text{proj}} = \tilde{G}$, $\zeta_{\text{proj}} = +2$), $K(\nu) = 4\pi/(\Omega\,\|k_\nu\|_2^2)$, and $S(\nu) = \sum_{\ell=1}^{216} e^{-ik_\nu\cdot R_\ell}$.

For every proton configuration $\{R_\ell\}$, initial projectile momentum k_{proj}, evolution time $t \in \{1, \ldots, 10\}$ a.u., and partition of the 218 electron registers into spin-up and spin-down subsets, a candidate circuit over Clifford+T and mid-circuit measurement using m ancilla qubits must realize a single unitary W that matches e^{-iHt} up to a global phase φ, satisfying

$$\left\| W\left(|\psi\rangle \otimes |0\rangle^{\otimes m}\right) - e^{i\varphi}\left(e^{-iHt}|\psi\rangle\right) \otimes |0\rangle^{\otimes m} \right\|_2 \le \varepsilon \tag{12}$$

at $\varepsilon = 10^{-2}$ for every normalized grid state $|\psi\rangle$ antisymmetric under exchange of same-spin electrons. All ancillas begin in $|0\rangle^{\otimes m}$ in each run, rather than carrying pre-prepared resource states (such as phase-gradient states) across runs. Every circuit is compiled into Clifford+T gates, mid-circuit measurements and classical feed-forward, and checked at that gate level by the correctness proof. Where later gates are conditioned on measurement outcomes, we count the T gates of the most expensive outcome branch.

To extract the stopping power from the projectile's rate of energy loss, the protocol repeats time evolution 50 times at each of $t = 1, \ldots, 10$ a.u. (500 runs) to estimate the mean projectile kinetic energy, omitting state preparation and readout as their cost is negligible. In this work, our Lean specification asks for the algorithm to prove correctness and cost for all 500 circuits. However, in Sec. 2.3 and Fig. 4 we report the average T count over the circuits, since this more accurately reflects the size of circuit a fault-tolerant quantum computer would need

to coherently execute. We apply this to both the baseline and our solution, so the factor of improvement reported is equal regardless. However, the baseline costing convention[48] (total over all circuits) can be obtained by simply multiplying our arithmetic mean by 500.

G.2. Baseline

Table 2 of ref. 48 gives a momentum-space LCU 1-norm of λ = 1,744,784.42 Ha, a block-encoding cost of $C_{\mathrm{BE}} = 2.498 \times 10^4$ Toffolis per walk step, and Q_{base} = 5,650 logical qubits (3,948 system and 1,702 ancilla qubits). Each electron–proton or electron–electron pair contributes $\sum_{\nu \in [-62,62]^3 \setminus \{0\}} (\pi \Omega^{1/3} \|\nu\|_2^2)^{-1} = 22.53$ Ha to λ, giving 1.06×10^6 Ha from electron–proton attraction and 5.33×10^5 Ha from electron–electron repulsion (91.4% of λ).

Generalized QSP[140] implements e^{-iHt} as a polynomial in a qubitized walk operator[25] that block-encodes H/λ, with one walk step per query. Evaluating the query count[141] in Eq. (12) of ref. 48,

$$N_{\mathrm{steps}}(t) = \lambda t + \frac{3^{2/3}}{2} (\lambda t)^{1/3} \log_2^{2/3}(1/\varepsilon), \tag{13}$$

at $\varepsilon = 10^{-2}$ and summing over the 50 runs at each of $t = 1, \ldots, 10$ a.u. gives 4.80×10^9 walk steps, or 1.20×10^{14} Toffolis over all 500 runs. This is an average of 9.60×10^6 walk steps and $N_{\mathrm{Toff}}^{\mathrm{base}} = 2.40 \times 10^{11}$ Toffolis per run, so $\mathcal{T}_{\mathrm{base}} = 4 N_{\mathrm{Toff}}^{\mathrm{base}} = 9.59 \times 10^{11}$ T gates per run and $Q_{\mathrm{base}} \cdot \mathcal{T}_{\mathrm{base}} = 5.42 \times 10^{15}$ per run (2.71×10^{18} over 500 runs), the number presented in Fig. 3a. The same study[48] also quotes an eighth-order product-formula estimate (1.40×10^{15} Toffolis on 6,170 qubits over 500 runs), which we do not use as a baseline since it needs 11.7× more Toffolis than the QSP estimate above.

G.3. Main improvements

Section 2.3 summarizes the two changes that account for most of the reduction (Fig. 4b,c). Here we describe each one alongside the baseline construction it replaces, and Appendix G.4 gives the resulting resource counts.

The baseline applies the electron–proton and electron–electron terms in momentum space, where each pair contributes 22.53 Ha to λ (Appendix G.2). Our solution instead Fourier transforms all electrons to a real-space grid each step to apply the electron–proton, electron–projectile and electron–electron interactions, while the kinetic and projectile–proton terms stay in momentum space. This basis change alone halves the contribution of each pair to λ, to 11.3 Ha, but that reduction is canceled out by the extra cost of applying the walk step in real space. However, a further reduction follows by being in the position basis, where the solution can exploit Pauli exclusion. A subroutine checks every electron against each grid point near a given electron or proton, and each electron found there kicks back a phase of i in one LCU branch and $-i$ in the other. Averaging the two branches gives $\cos(\pi n/2)$ for n electrons on the point, which equals $1 - n$ for $n \leq 2$ and has magnitude at most one for any n. The Lean solution also proves that states antisymmetric under exchange of same-spin electrons have at most two electrons on any grid point, allowing it to satisfy the specification while not faithfully applying the $n > 2$ interaction terms. Therefore, the circuit reproduces the interaction up to a constant, while each such term adds its potential to λ once rather than once for every electron that could occupy the point. At separations beyond $\approx 0.6\ a_0$, where the potential is below 1.4 Ha, the interaction is too weak for the savings to outweigh the cost of the subroutine. The pair potential is therefore range-separated at 1.4 Ha, with the subroutine applied only to the short-range part and the long-range part summed over all pairs, each contributing 0.74 Ha (Fig. 4b).

The second most significant improvement found, in terms of cost reduction, relates to the electron–proton interaction. Protons do not sit on grid points, so to apply this term the circuit shifts all electrons by $-R_\ell$ to put each proton on a grid point, and then shifts them back. In the plane-wave basis this shift is a phase set by the total

electron momentum P. Earlier real-space solutions computed and uncomputed P around each shift, summing all 218 electron momenta four times per walk step, which made up about a fifth of the T gates per step. Only the protons and the projectile change P, and both act in the real-space step, so the final circuit keeps P in a 42-qubit register between walk steps, clearing it before the real-space step and rebuilding it after. This halves the sums and cuts the T gates per step by 10% (Fig. 4c).

G.4. Resource counts

Our solution uses $Q = 4{,}338$ logical qubits: 3,948 system qubits and 390 clean ancillas, against the baseline's 1,702. The ancillas are a 42-qubit register that keeps the total electron momentum between walk steps, a 48-qubit phase-gradient register prepared in each run, 87 control and data qubits for the block encoding, and 213 workspace qubits that every subroutine borrows and returns to $|0\rangle$. The improvements above lower λ to 1.356×10^5 Ha. Each run of length t is split into t segments of 1 a.u., each implemented by generalized QSP[140,141] of degree 135,867, giving an average of 7.47×10^5 walk steps per run.

Each walk step uses roughly 141,500 T gates. One-time initialization adds 59,600 T gates per run. Averaged over $t = 1,\ldots,10$ a.u., this gives 1.06×10^{11} T gates and $Q\cdot\mathcal{T} = 4.59\times10^{14}$ per run, or 229,357,900,876,653,900 over all 500 runs. Compared with the baseline, this is 12.8× fewer walk steps and 1.30× fewer qubits at 1.42× the T gates per step, an 11.8× lower cost (Fig. 4a).

H. Non-Born–Oppenheimer Dynamics of $C_{23}H_{20}N_3O$

H.1. Hamiltonian and specification

Following Reaction 5 in Eqs. (2)–(10) and (23)–(25) and Tables I and IV of ref. 58, the system consists of $\eta = 234$ particles of composition $C_{23}H_{20}N_3O$: 47 nuclei (23 ^{12}C, 20 ^{1}H, 3 ^{14}N, and 1 ^{16}O, with charges $\zeta_i = Z_i \in \{6,1,7,8\}$) followed by 187 electrons ($\zeta_i = -1$). Each Cartesian coordinate is stored as a signed $n_g = 8$-bit integer $q_{i,w} \in [-128,127]$ on $Q_{\text{sys}} = 3\eta n_g = 5{,}616$ system qubits, representing position $r_{i,w} = q_{i,w}\Delta$ in a cubic box of width $L = 44$ Å (83.15 a_0) with grid spacing $\Delta = L/(2^{n_g}-1) \approx 0.326\ a_0$. Electron masses are $m_i = 1$, and nuclear masses m_i (not stated in the baseline[58]) are set in electron-mass units from the ^{1}H, ^{12}C, ^{14}N, and ^{16}O isotopic atomic masses minus $Z_i m_e$.

The discretized non-Born–Oppenheimer Hamiltonian is $H = T + V$, with

$$T = \frac{2\pi^2}{L^2}\sum_{i=1}^{\eta}\frac{1}{m_i}\,\mathrm{QFT}_i^{\dagger}\left(\sum_{k\in G}\|k\|_2^2\ |k\rangle\langle k|_i\right)\mathrm{QFT}_i, \qquad V = \frac{1}{2}\sum_{i\neq j}\frac{\zeta_i\zeta_j}{\Delta\max\left(\left\|q_i-q_j\right\|_2,\tilde{\Gamma}_{ij}\right)}, \tag{14}$$

where $G = [-128,127]^3$, QFT_i is the quantum Fourier transform on the three coordinates of particle i, and the squared Coulomb saturation distance is $\tilde{\Gamma}_{ij}^2 = 8$ when both i and j are nuclei ($n_\Gamma = 3$ in Table IV of ref. 58) and $\tilde{\Gamma}_{ij}^2 = 1$ otherwise, capping electron interactions at $1/\Delta$ and nuclear repulsion at $1/(\sqrt{8}\Delta)$.

A candidate circuit over Clifford+T and mid-circuit measurement must realize a single clean-ancilla unitary W satisfying Eq. (12) at $\varepsilon = 10^{-2}$ for $t = 1$ fs on every state of the system register. Matching the baseline,[58] initial-state preparation and observable readout are not costed, since in end-to-end fault-tolerant simulations of photochemical dynamics, real-time Hamiltonian evolution dominates initial-state preparation by orders of magnitude.[142] All ancillas begin in $|0\rangle^{\otimes m}$, so resource states such as phase-gradient states must be prepared within the circuit. Every Toffoli is compiled into Clifford+T gates and checked at that gate level by the correctness proof, and where later gates are conditioned on measurement outcomes we count the T gates of the most expensive branch.

The baseline instead bounds only the block-encoding error $\left\|A - e^{i\varphi}e^{-iHt}\right\|_{\mathrm{op}} \le \epsilon_{\mathrm{BE}}$ (Appendix C of ref. 58), where $A = (I \otimes \langle 0^m|)\, W\, (I \otimes |0^m\rangle)$ is the part of the circuit that starts and ends with the ancillas in $|0^m\rangle$. This leaves out the ancilla error, the part $|\psi_\perp\rangle$ of the output in which the ancillas are not returned to $|0^m\rangle$. The two parts are orthogonal, so the error in Eq. (12) splits as

$$\left\|W(|\psi\rangle \otimes |0^m\rangle) - e^{i\varphi}(e^{-iHt}|\psi\rangle) \otimes |0^m\rangle\right\|_2^2 = \left\|(A - e^{i\varphi}e^{-iHt})|\psi\rangle\right\|_2^2 + \||\psi_\perp\rangle\|_2^2, \tag{15}$$

and any circuit meeting Eq. (12) at ε also meets the block-encoding criterion at $\epsilon_{\mathrm{BE}} = \varepsilon$. For the baseline QSP circuit, the polynomial approximation error bounds the ancilla error[25,140,141] by $\||\psi_\perp\rangle\|_2^2 \le 2\epsilon_{\mathrm{exp}}$, so with the error split in Table IV of ref. 58 ($\epsilon_{\mathrm{exp}} = \epsilon_{\mathrm{BE}}/12.95$) it is only guaranteed to meet Eq. (12) at $\sqrt{\epsilon_{\mathrm{BE}}^2 + 2\epsilon_{\mathrm{exp}}} \approx 0.041$, four times our target.

H.2. Baseline

Table I of ref. 58 reports $N_{\mathrm{Toff}}^{\mathrm{pub}} = 2.73 \times 10^{11}$ Toffolis per femtosecond on $Q_{\mathrm{pub}} = 6{,}198$ qubits at $\epsilon_{\mathrm{BE}} = 10^{-2}$, a spacetime cost of 6.77×10^{15}. Each walk step uses about 9.0×10^4 T gates, of which 52% go to the swap network, 25% to preparing the LCU weights and 22% to evaluating $1/r$ (Fig. 5a of ref. 58). These give the baseline values of 46.8k, 22.5k and 19.8k T gates per step in Sec. 2.4 and Fig. 4e–g.

The baseline costs the time evolution as the number of walk steps times the cost of one step, $N_{\mathrm{Toff}} = D_{\mathrm{old}}(\tau, \epsilon_{\mathrm{exp}}) \cdot \mathcal{T}_H$ (Eqs. (C1) and (C4)–(C5) of ref. 58), where the number of steps is the degree of the generalized QSP polynomial.[140,141] Two of the same authors tighten this degree bound, removing the factor $e/2 \approx 1.36$ in front of τ (Theorem 1 of ref. 59, via Kapteyn's inequality),

$$\begin{aligned} D_{\mathrm{old}}(\tau, \epsilon_{\mathrm{exp}}) &= \left\lceil \frac{e}{2}\,\tau + \ln\frac{2c}{\epsilon_{\mathrm{exp}}} \right\rceil, \\ d_{\mathrm{Thm\,1}}(\tau, \epsilon_{\mathrm{exp}}) &= \left\lceil \tau \cosh\left(\left[\frac{3}{\tau} \ln \frac{2}{\epsilon_{\mathrm{exp}}(1 - e^{-\theta_0})} \right]^{1/3} \right) \right\rceil, \end{aligned} \tag{16}$$

where $\tau := \alpha_H t$, $c = 4(\sqrt{2\pi}\, e^{1/13})^{-1} \approx 1.4776$ and $\theta_0 = \left[\frac{3}{\tau}\ln(2/\epsilon_{\mathrm{exp}})\right]^{1/3}$. Theorem 2 there further tightens the subleading term via Watson's inequality. That work[59] does not apply this bound to the system studied here, but it improves the same QSP construction, so we include it so that our solutions are compared against the strongest available baseline.

The tighter bound only reduces the number of walk steps. The error split $t\epsilon_H + \epsilon_{\mathrm{exp}} \le \epsilon_{\mathrm{BE}}$ fixes ϵ_H independently of the polynomial degree, so the T-gate count $\mathcal{T}_H$ per step and the qubit count $Q_{\mathrm{pub}} = 6{,}198$ are unchanged. At the reported $\alpha_H = 2.16 \times 10^5$ Ha ($\tau \approx 8.93 \times 10^6$ at $t = 41.34$ a.u.) and $\epsilon_{\mathrm{exp}} = \epsilon_{\mathrm{BE}}/12.95$, Theorem 2 lowers the number of walk steps from $D_{\mathrm{old}} = 12{,}136{,}779$ to $d_{\mathrm{Thm\,2}} = 8{,}930{,}416$ (8,930,877 with Theorem 1). Rescaling the published Toffoli count by this ratio gives our baseline,

$$\begin{aligned} N_{\mathrm{Toff}}^{\mathrm{base}} &= N_{\mathrm{Toff}}^{\mathrm{pub}}\, \frac{d_{\mathrm{Thm\,2}}}{D_{\mathrm{old}}} \approx 2.01 \times 10^{11}, \\ \mathcal{T}_{\mathrm{base}} &= 4 N_{\mathrm{Toff}}^{\mathrm{base}} \approx 8.04 \times 10^{11}, \\ Q_{\mathrm{pub}} \cdot \mathcal{T}_{\mathrm{base}} &\approx 4.98 \times 10^{15}, \end{aligned} \tag{17}$$

the number presented in Fig. 3a. At this large τ the subleading terms barely matter: using Theorem 1 instead, or simply dividing the published Toffoli count by $e/2$, gives the same cost to three significant figures.

H.3. Main improvements

Section 2.4 summarizes three changes to the walk step that account for almost all of the reduction in T gates per step relative to the baseline (Fig. 4e–g). Here we describe each one alongside the baseline construction it replaces, and Appendix H.4 gives the resulting T counts.

First, to apply a two-body interaction, the baseline swaps the coordinates of the selected pair into the two particle registers on which its interaction circuit acts, applies the interaction there, and swaps them back afterward.[58] Instead, our solution copies the selected coordinates into an ancilla register, leaving the particle register untouched. Because the ancilla holds only a copy, it can be cleared using mid-circuit Pauli-X measurements and classically controlled CZ corrections,[23,117] as done for momentum bits copied in plane-wave kinetic terms[27] (Fig. 4e).

Second, the potential V in Eq. (14) requires computing $1/r = (r^2)^{-1/2}$ for $r = \|q_i - q_j\|_2$. Because the coordinates are stored as integers, $r^2 = \sum_w (q_{i,w} - q_{j,w})^2$ can be obtained cheaply using integer subtraction and squaring. To get to $1/r$ without a square root, the baseline prepares an ancilla register in uniform superposition over integers $m < M$, assigning $+1$ to branches with $m^2 r^2 < M^2$ and an alternating sign $(-1)^m$ to the rest. The alternating terms cancel, so the average over branches approximates $1/r$, but the test requires multiplying m^2 by r^2, which dominates its cost.[58] In this work, the range of r^2 is split into coarse intervals with rescaled position $\tilde{r}^2 \in [0, 1)$, and $1/r$ on each interval is approximated by a cubic Bernstein polynomial $\sum_{k=0}^{3} \binom{3}{k} (\tilde{r}^2)^k (1-\tilde{r}^2)^{3-k} w_k$ (512 precomputed coefficients w_k in total). Rather than multiplying powers of $\tilde{r}^2$, the circuit samples this polynomial directly.[67,143] It compares $\tilde{r}^2$ against three ancilla registers in uniform superposition and counts how many are smaller, which gives $k \in \{0, 1, 2, 3\}$ with exactly the binomial weights above. Looking up w_k then gives the polynomial in expectation (Fig. 4f).

Finally, in the electron–nucleus terms of V, each nucleus i is weighted by its integer charge Z_i. Rather than preparing these weights via coherent alias sampling[24] with amplitude amplification,[58] a uniform superposition of basis states is prepared and a reversible classical circuit assigns Z_i states to each nucleus i, as proposed for fixed nuclei[27] and for quantum nuclei.[142] Our solution applies this to every weight of the linear combination (Fig. 4g).

H.4. Resource counts

The verified circuit found in this work uses $Q = 6{,}033$ logical qubits, consisting of 5,616 system qubits and 417 clean ancillas. Evaluating the 1-norm at $L = 44$ Å together with potential shifting gives block-encoding scale $\alpha = 8.045 \times 10^4$ Ha and generalized QSP degree $d = 3{,}326{,}847$.

Each uncontrolled walk step uses 30,116 T gates (30,132 T on the two boundary controlled steps): 1,256 T for reversible nuclear-charge and pair-weight preparation, 24,194 T to copy selected particle coordinates into ancillas and 908 T for X-basis measurement uncomputation (25,102 T total routing, against 46.8k T for the baseline swap network), 3,060 T for the 512-entry Bernstein $1/r$ evaluation, 446 T for the QFT kinetic term, and 252 T for the LCU reflection, plus 76 T per step for QSP rotations and 245 T of one-time overhead. The Lean solution fixes both the 3,326,848 generalized QSP rotations[140,141] and the Clifford+T words that approximate them by existence proofs, the words through a derandomized form of unitary mixing over a finite table,[144,145] so it certifies their T count and error without listing them. The rotations are computable in time polynomial in d[146] but have not been computed at this degree. The same holds for the baseline,[58,59] which uses the same doubled construction[141] and a larger degree polynomial which would require greater classical computation. This totals $\mathcal{T} = 3{,}326{,}847 \times (30{,}116 + 76) + 2 \times 30{,}132 + 245 = 100{,}444{,}225{,}133 \approx 1.00 \times 10^{11}$ T gates and $Q \cdot \mathcal{T} = 605{,}980{,}010{,}227{,}389 \approx 6.06 \times 10^{14}$ (8.22× below baseline).

I. Cytochrome P450 Compound I

I.1. Hamiltonian and specification

The system is active space X of Compound I, with $\eta = 63$ electrons in $N = 58$ spatial orbitals.[88] Under the Jordan–Wigner transformation each spin orbital (p, σ) is stored on qubit $2p + \sigma$, giving $Q_{\text{sys}} = 2N = 116$ system qubits.

GPU-accelerated density matrix renormalization group (DMRG) calculations now converge this 58-orbital active space classically,[147] yet leave the physical spin-state ordering unresolved because dynamic correlation from omitted orbitals exceeds the sub-millihartree doublet–quartet gap.[88,147] Active space X nonetheless serves as the standard benchmark across fault-tolerant resource estimates,[60,88,89] and circuit improvements to its DFTHC walk structure[60] carry over to the larger active spaces needed to capture dynamic correlation and resolve the spin ordering.[147]

The Hamiltonian is the $(R, B, C) = (9, 29, 14)$ double-factorized tensor-hypercontraction (DFTHC) factorization of Table V in ref. 60, with the released tensors[148] rounded to $(b_{\text{coeff}}, b_{\text{rot}}) = (10, 15)$ bits:

$$H_{\text{DFTHC}} = \sum_{\sigma,p,q} h'_{pq}\, a^\dagger_{p\sigma} a_{q\sigma} + \frac{1}{2} \sum_{r=0}^{R-1} \sum_{c=0}^{C-1} \left(W^{(rc)} I + \sum_{b=0}^{B-1} w_b^{(rc)} n_{u_b^{(r)}} \right)^2, \tag{18}$$

where $n_u = \sum_\sigma (\sum_p u_p a_{p\sigma})^\dagger (\sum_q u_q a_{q\sigma})$ for unit orbital vectors $u_b^{(r)} \in \mathbb{R}^{58}$ ($M = RB = 261$) and $W^{(rc)} = w_B^{(rc)} - \sum_{b<B} w_b^{(rc)}$, as follows from expanding $O_{\text{SF},rc}$ in the DFTHC construction.[60] As in the released dataset,[148] Eq. (18) omits the scalar $-\frac{1}{2}\sum_{r,c}(W^{(rc)})^2$, and adding the classical Hartree–Fock shift $E_{\text{HF}} - E_{\text{HF}}^{\text{DFTHC}}$ recovers the active-space total energy (Eqs. 42–43 of ref. 60). We define the problem to not allow for improvement in the DFTHC factors,[148] which reproduce the active-space coupled-cluster correlation energy to 0.28 mHa, to separate quantum circuit optimization from classical tensor fitting.

We cost one run of phase estimation on the 116 system qubits, started in a 63-electron eigenstate $|\psi\rangle$ of H_{DFTHC} with energy E. Assuming an eigenstate as input loses no generality for resource estimation: phase estimation on any state $|\phi\rangle$ gives, with probability $|\langle\psi_j|\phi\rangle|^2$, the same outcomes[149] as on the eigenstate $|\psi_j\rangle$. In practice $|\phi\rangle$ only approximates the ground state $|\psi_0\rangle$, so finding E_0 takes on average $1/|\langle\psi_0|\phi\rangle|^2$ runs, each also paying for trial-state preparation.[150] Like prior resource estimates,[60,88,89] we do not count this overhead. The circuit uses Clifford+T gates and mid-circuit measurements on the system and m ancillas that start in $|0\rangle$. The specification applies to this gate-level circuit, so its error bound includes rotation synthesis and approximate state preparation as well as phase estimation. At the end, a b-qubit phase register among the ancillas is measured, and its outcome $y \in \{0, \ldots, 2^b - 1\}$ is converted into the energy estimate

$$\hat{E}(y) = E_{\text{off}} + s\,|\cos(\pi y / 2^b)|, \tag{19}$$

where b, E_{off} and s are fixed numbers chosen with the circuit. This form covers both standard qubitization and sum-of-squares spectral amplification. For a walk eigenphase φ, the energy is $E = \lambda \cos\varphi$ in the first[4,24,25,89] and $E = E_{\text{SOS}} + \Lambda(1 + \cos\varphi)$ in the second.[60,151] Since $\cos\varphi \le 0$ at the energies measured here, both read $E = E_{\text{off}} + s\,|\cos\varphi|$, with $(E_{\text{off}}, s) = (0, -\lambda)$ and $(E_{\text{SOS}} + \Lambda, -\Lambda)$ respectively. When phase estimation applies the walk forward or backward depending on the phase register, $2\pi y/2^b$ estimates $\pm 2\varphi$, so Eq. (19) estimates E.

If the specification only covered the ground state, a circuit with no gates and $s = 0$ would pass by setting E_{off} within ε of the ground-state energy E_0. The specification therefore requires the same circuit and decoder to work for every eigenstate up to $E_{\text{top}} = E_0 + \Delta$, with $\Delta = 0.1$ Ha and $\varepsilon = 10^{-3}$ Ha. For each 63-electron eigenstate $|\psi\rangle$ with energy E, the outcome distribution $p(y \mid \psi)$ must satisfy

$$\sum_y p(y \mid \psi) \left(\min(\hat{E}(y), E_{\text{top}}) - \min(E, E_{\text{top}}) \right)^2 \le \varepsilon^2. \tag{20}$$

This is the root-mean-square error behind the phase-estimation step counts of Eqs. (21)–(23) in ref. 24. For an eigenstate inside the window, it bounds the root-mean-square error of the estimate by ε, with estimates above E_{top} counted as E_{top}. For an eigenstate above the window, it only requires that estimates do not fall more than ε below E_{top}. A circuit that ignores $|\psi\rangle$ gives the same estimates for every eigenstate, and in root mean square these would have to lie within ε of E_0 and at most ε below $E_0 + \Delta$ at once. This is possible only if $\Delta \le 2\varepsilon = 2$ mHa, far below $\Delta = 0.1$ Ha. The window also spans the gap of about 75 mHa between the quartet and sextet states found by DMRG (Fig. 4 of ref. 147). The tolerance $\varepsilon = 1$ mHa matches prior resource estimates on this active space, although separating the doublet and quartet states would require tighter tolerances.

All ancillas begin in $|0\rangle^{\otimes m}$, so resource states such as phase-gradient states must be prepared within the circuit. Every Toffoli is compiled into Clifford+T gates and checked at that gate level by the correctness proof, and where later gates are conditioned on measurement outcomes we count the T gates of the most expensive branch.

I.2. Baseline

Our baseline starts from the sum-of-squares spectral amplification estimate in Table V of ref. 60 (51,510 walk steps at 9,535 Toffolis/step on 1,150 qubits, or a spacetime cost of 2.26×10^{12}), where $H_{\text{DFTHC}} - E_{\text{SOS}} = H_{\text{sqrt}}^{\dagger} H_{\text{sqrt}}$ is block-encoded with norm Λ and energy-window gap $E_{\text{gap}} = E_{\text{ub}} + \Delta - E_{\text{SOS}}$ using the phase-estimation step count[24,60]

$$\lambda_{\text{eff}} = \sqrt{E_{\text{gap}}(2\Lambda - E_{\text{gap}})}, \qquad N_{\text{steps}} = \left\lceil \frac{\pi \lambda_{\text{eff}}}{2\varepsilon} \right\rceil. \tag{21}$$

Adding Δ to E_{gap} in that estimate[60] raises its step count from 51,510 to 51,948 and its cost to 2.28×10^{12} (+0.85%). The published construction[60] forms the one-body term h' from the exact two-electron tensor g, while the squares in Eq. (18) use the factorized tensor $\tilde{g}$. The Hamiltonian it block-encodes therefore differs from the $(h, \tilde{g})$ Hamiltonian by a one-body term that lowers the low-lying energies by about 0.6 Ha. Its E_{gap} subtracts the sum-of-squares bound of the block-encoded Hamiltonian from the sextet Hartree–Fock energy of the $(h, \tilde{g})$ Hamiltonian,[148] so it is about 0.6 Ha larger than when both energies come from one Hamiltonian, which overestimates the cost. We form h' from $\tilde{g}$, so that Eq. (18) is the $(h, \tilde{g})$ Hamiltonian up to a constant, whose CCSD(T) correlation energy is within 0.28 mHa of that of (h, g). This lowers Λ from 97.4395 to 97.3864 Ha ($E_{\text{SOS}} = -134.8874$ Ha), and with both energies evaluated on this Hamiltonian $E_{\text{gap}} = 5.189$ Ha including Δ (49,269 steps). Bounding $E_0 \le E_{\text{ub}} = -129.8443$ Ha with the quartet ($S = 3/2$) ROHF determinant instead lowers E_{gap} to 5.143 Ha, giving $\lambda_{\text{eff}} = 31.23$ Ha and $N_{\text{steps}} = 49{,}056$.

Adding a coherent 16-qubit phase-estimation register and storing each of the 57 Givens angles in 13 rather than 15 bits gives 1,052 qubits and a cost of 1.97×10^{12} before space–time tradeoffs. Thirteen bits suffice because the angles $\theta_p = 2\pi k_p / 2^{15}$ enter the orbital vectors[4] as $\cos 2\theta_p$ and $\sin 2\theta_p$, and the overall sign of each vector does not affect n_u, so fixing that sign puts every angle index of H_{DFTHC} in $0 \le k_p < 2^{13}$. As suggested in prior work,[89] we use both these observations and integrate them into our baseline. We then apply recent compilation options,[61] recomputing the $(G, r, c, \sigma, b = B)$ index registers from an 8-bit outer index x_o with minimum-Toffoli comparators[109] and loading the 57 Givens angles in three batches of 19 (5 QROM lookups per SELECT) while reusing idle angle wires as clean-QROAM[117] workspace. These options come from later work than the spectral-amplification study,[60] but we include them so that our solutions are compared against the strongest available baseline. On our 13-bit instance we also pad C rather than R in x_o, skip the two always-zero addend bits in the 15-bit phase-gradient adders,[23] and include 32 Toffolis/step to control the walk reflections from the phase-estimation register, omitting approximate state preparation and rotation synthesis so H_{DFTHC} remains exact. Minimizing $Q \cdot \mathcal{T}$ over QROAM block sizes and angle batch sizes gives $Q_{\text{base}} = 481$ logical qubits and $C_{\text{BE}} = 11{,}517$ Toffolis/step (5.65×10^8 Toffolis, $\mathcal{T}_{\text{base}} = 4 C_{\text{BE}} N_{\text{steps}} = 2.26 \times 10^9$ T gates), for a baseline cost of

$$Q_{\text{base}} \cdot \mathcal{T}_{\text{base}} = 1.087 \times 10^{12} \approx 1.09 \times 10^{12}, \tag{22}$$

the number presented in Fig. 3a. At a fixed number of walk steps, the 13-bit angles and these compilation options[61] together roughly halve the published cost,[60] from 1,150 qubits and 9,535 Toffolis per step to 481 qubits and 11,517, including the phase-estimation register. Fig. 3b shows the published estimate,[60] extended to the energy window, in gray and this baseline in red.

I.3. Main improvements

Section 2.5 summarizes the main change in our solution, which replaces the angle lookups of the baseline with a one-hot register for the chosen orbital vector, and uses the same trick in the inner state preparation. Here we describe both alongside the baseline construction they replace, together with a smaller change to the spin swaps that the main text leaves out. Appendix I.4 gives the resulting T counts.

Each walk step block-encodes Eq. (18) with a state preparation and two SELECT calls.[60] The outer stage of the state preparation picks a square (r, c) or a one-body term, and the inner stage picks b, which fixes the orbital vector $u_b^{(r)}$. SELECT rotates this vector into a single orbital with 57 Givens rotations, acts on that orbital and rotates back.[4]

In the baseline, each SELECT call loads the 13-bit Givens angles of the chosen vector into a 247-qubit register, 19 at a time with 5 QROM lookups per call, and each rotation adds its angle into the phase-gradient register (Appendix I.2). Our solution instead builds a 145-qubit one-hot register in each SELECT call, for 296 T. Each of its qubits stands for two orbital vectors, so the register covers 290 of them: the 261 vectors $u_b^{(r)}$ and 29 of the 58 eigenvectors of h'. The other 29 eigenvectors are handled by the second call of the step. A further qubit s records which vector of the pair was chosen. Each angle bit then enters the phase-gradient adder of its rotation as

$$(\alpha \oplus M_1)(\beta \oplus M_0) \oplus M_{01}, \qquad \beta = q_c s, \quad \alpha = q_c \oplus \beta, \tag{23}$$

where q_c is the control qubit of the rotation and M_0, M_1 and M_{01} are CNOT parities of the one-hot qubits whose first vector, second vector or both vectors have that bit set. At most one one-hot qubit is set, so $M_0 M_1 = M_{01}$, and Eq. (23) is 0 for $q_c = 0$ and the bit of the chosen vector for $q_c = 1$. The baseline already uses one AND gate per angle bit to control the rotation. Eq. (23) also needs one AND gate per bit, plus one per rotation for β. With the one-hot register, the rotations and their angles cost 11,818 T per call rather than 18,844 T, and the angle register is no longer needed.

The inner stage uses the same idea. It keeps a 126-qubit one-hot register for (r, c) and a 5-bit binary register for b, so each alias-sampling lookup[24] costs 390 T, against 1,968 T for the QROAM of the baseline. Together with the one-hot SELECT, this accounts for 96% of the T gates saved per walk step (Table S3).

Our solution also includes a second, smaller improvement, which halves the cost of the spin swaps. Rather than performing the basis rotations once for each spin, the baseline applies them to the spin-up qubits only and, controlled by a qubit in superposition, swaps the spin-down qubits into their place before and back after, with four blocks of 58 controlled swaps per step.[60] Since swapping two qubits only flips the sign of their antisymmetric state $(|01\rangle - |10\rangle)/\sqrt{2}$, each block can instead be implemented as a single controlled sign flip, applied when an odd number of the 58 orbitals are in that state. After a Clifford change of basis to Bell pairs, the circuit computes this parity with 58 Toffolis, or 232 T, as many as one block of swaps, and erases it with X-basis measurements for free. Since the system register is untouched between consecutive SELECT calls, one parity serves two swap blocks.

I.4. Resource counts

The best verified circuit uses $Q = 333$ logical qubits and $\mathcal{T} = 1{,}310{,}839{,}337 \approx 1.31 \times 10^9$ T gates, so $Q \cdot \mathcal{T} = 436{,}509{,}499{,}221 \approx 4.37 \times 10^{11}$, 2.49× below the baseline (Fig. 3a,b). This comes from 1.67× fewer T

Part of a walk step	Baseline	This work
Rotations and their angles	37,688	23,636
Index decoding	264	440
Spin swaps	928	464
Outer state preparation	1,032	598
Inner state preparation	5,888	2,264
Total	45,800	27,402

Table S3: T gates per walk step for the baseline and for a step inside the first eight blocks (45,440 steps) of the best verified circuit. Reflections and phase-estimation control are counted separately in the text.

gates per walk step (Appendix I.3), 1.44× fewer qubits and 3% fewer walk steps. The qubits are the 116 system qubits, a 16-qubit phase-estimation register, an 11-qubit phase-gradient register, 17 and 24 qubits for the outer and inner state-preparation registers, and 149 qubits of workspace shared by the subroutines. Most of the saving against the 481 qubits of the baseline comes from dropping its 247-qubit angle register.

The circuit runs $M = 47{,}487$ walk steps. We bound $E_0 \le E_{\mathrm{ub}} = -130.1526$ Ha with a 3,000-determinant quartet wavefunction, whose energy is checked in Lean. This is 0.31 Ha below the quartet ROHF bound of the baseline and gives $E_{\mathrm{gap}} = 4.835$ Ha and $\lambda_{\mathrm{eff}} = 30.30$ Ha, 3% below the 31.23 Ha of the baseline. A slightly tighter error bound for phase estimation than Eq. (21) saves another 0.2%. The state preparation is approximate, so the block-encoded Hamiltonian differs from H_{DFTHC} by at most 4.0×10^{-7} Ha on the 63-electron sector. This comes from the synthesized rotations in the state preparation, and the correctness proof adds it to the phase-estimation error to show that the total stays within ε.

The walk steps run in 16 blocks totaling 47,360 steps, followed by 127 single steps. As qubits of the phase-estimation register are prepared stage by stage, fewer remain free as extra workspace for SELECT: 8 blocks (45,440 steps in total) borrow 11 unprepared qubits for 160 workspace qubits at the 27,402 T per step of Table S3, 1 block of 1,024 steps borrows 10 qubits (28,298 T per step), and 7 blocks of 128 steps borrow 7 qubits (31,010 T per step). On top of Table S3, the baseline spends 268 T per step on reflections and phase-estimation control, for 46,068 T in total. Each of our blocks uses 84 T for each reflection between its steps and 232 T for an ordinary first swap block. The 127 single steps cost 37,642 T each (including 418 T per inner lookup), because all phase-estimation qubits are then in use and SELECT has to fit in 149 qubits of workspace. The reflections between blocks and between single steps, together with phase-estimation control, take 13,316 T, with 31 reflections controlled by the phase-estimation register. Preparing the initial state of the phase-estimation register takes 142,248 T, the inverse QFT 12,603 T and the phase-gradient state 1,036 T. Adding these up gives the $\mathcal{T} = 1{,}310{,}839{,}337$ quoted above.

Appendix References

1. Beverland, M. E. *et al. Assessing requirements to scale to practical quantum advantage* 2022. arXiv: 2211.07629 [quant-ph].
2. Preskill, J. Beyond NISQ: The Megaquop Machine. *ACM Transactions on Quantum Computing* **6,** 1–7. arXiv: 2502.17368 [quant-ph] (2025).
3. Reiher, M., Wiebe, N., Svore, K. M., Wecker, D. & Troyer, M. Elucidating reaction mechanisms on quantum computers. *Proceedings of the National Academy of Sciences* **114,** 7555–7560 (2017).
4. Lee, J. *et al.* Even More Efficient Quantum Computations of Chemistry Through Tensor Hypercontraction. *PRX Quantum* **2,** 030305 (2021).
7. Babbush, R. *et al.* Securing Elliptic Curve Cryptocurrencies against Quantum Vulnerabilities: Resource Estimates and Mitigations. *PRX Quantum* **7,** 031001. arXiv: 2603.28846 [quant-ph] (2026).

9. Grover, L. K. *A fast quantum mechanical algorithm for database search* in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (1996), 212–219.
10. Häner, T., Roetteler, M. & Svore, K. M. *Optimizing Quantum Circuits for Arithmetic* 2018. arXiv: 1805.12445 [quant-ph].
23. Gidney, C. Halving the cost of quantum addition. *Quantum* **2,** 74 (2018).
24. Babbush, R. *et al.* Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity. *Physical Review X* **8,** 041015 (2018).
25. Low, G. H. & Chuang, I. L. Hamiltonian Simulation by Qubitization. *Quantum* **3,** 163 (2019).
26. Gilyén, A., Su, Y., Low, G. H. & Wiebe, N. *Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics* in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (2019), 193–204.
27. Su, Y., Berry, D. W., Wiebe, N., Rubin, N. & Babbush, R. Fault-Tolerant Quantum Simulations of Chemistry in First Quantization. *PRX Quantum* **2,** 040332 (2021).
41. De Moura, L. & Ullrich, S. *The Lean 4 Theorem Prover and Programming Language* in *Automated Deduction – CADE 28* **12699** (Springer, 2021), 625–635.
42. Feng, Y. & Ying, M. Quantum Hoare Logic with Classical Variables. *ACM Transactions on Quantum Computing* **2,** 1–43 (2021).
43. Litinski, D. A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery. *Quantum* **3,** 128 (2019).
44. Google DeepMind. *Formal quantum computation in Lean 4* https://github.com/google-deepmind/formal-quantum-computation. 2026.
45. The mathlib Community. *The Lean mathematical library* in *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs* (2020), 367–381.
46. The Physlib Community. *Physlib: The Lean Physics Library* https://github.com/leanprover-community/physlib. 2024.
47. Meiburg, A. *Quantum Information in Lean* https://github.com/Timeroot/Lean-QuantumInfo. 2024.
48. Rubin, N. C. *et al.* Quantum computation of stopping power for inertial fusion target design. *Proceedings of the National Academy of Sciences* **121,** e2317772121 (2024).
55. Qu, A. *et al. CORAL: Towards Autonomous Multi-Agent Evolution for Open-Ended Discovery* in *Proceedings of the Conference on Language Modeling (COLM)* (2026). arXiv: 2604.01658 [cs.AI].
56. Zhong, Z., Raghunathan, A. & Carlini, N. *ImpossibleBench: Measuring LLMs' Propensity of Exploiting Test Cases* 2025. arXiv: 2510.20270.
57. Lean FRO. *comparator: A trustworthy judge for Lean proofs from untrusted sources* https://github.com/leanprover/comparator. Accessed Nov. 2026. 2025.
58. Pocrnic, M., Loaiza, I., Arrazola, J. M., Wiebe, N. & Motlagh, D. *Efficient Simulation of Pre-Born-Oppenheimer Dynamics on a Quantum Computer* 2026. arXiv: 2602.11272 [quant-ph].
59. Pocrnic, M. & Motlagh, D. *Improved constant factors for qubitized Hamiltonian simulation* 2026. arXiv: 2608.02734 [quant-ph].
60. Low, G. H. *et al.* Fast Quantum Simulation of Electronic Structure by Spectral Amplification. *Physical Review X* **15,** 041016 (2025).
61. Low, G. H. *et al. A denser planar surface code* 2026. arXiv: 2605.30455 [quant-ph].
62. Boyer, M., Brassard, G., Høyer, P. & Tapp, A. Tight Bounds on Quantum Searching. *Fortschritte der Physik* **46,** 493–505 (1998).
63. Zalka, C. Grover's quantum searching algorithm is optimal. *Physical Review A* **60,** 2746–2751 (1999).

64. Gingrich, R. M., Williams, C. P. & Cerf, N. J. Generalized quantum search with parallelism. *Physical Review A* **61,** 052313 (2000).

65. Grover, L. K. & Radhakrishnan, J. *Quantum search for multiple items using parallel queries* 2004. arXiv: quant-ph/0407217.

66. Fang, W. & Heunen, C. *Weakly Measured Loops for Quantum Amplitude Amplification: Oracle Savings and Adaptive Search with Unknown Target Probability* 2026. arXiv: 2609.39621 [quant-ph].

67. Goyal, V. & Sigman, K. On simulating a class of Bernstein polynomials. *ACM Transactions on Modeling and Computer Simulation* **22,** 1–5 (2012).

88. Goings, J. J. *et al.* Reliably assessing the electronic structure of cytochrome P450 on today's classical computers and tomorrow's quantum computers. *Proceedings of the National Academy of Sciences* **119,** e2203533119 (2022).

89. Caesura, A. *et al.* Faster Quantum Chemistry Simulations on a Quantum Computer with Improved Tensor Factorization and Active Volume Compilation. *PRX Quantum* **6,** 030337 (2025).

90. Gidney, C. *How to factor 2048 bit RSA integers with less than a million noisy qubits* 2025. arXiv: 2505.15917 [quant-ph].

91. Hietala, K., Rand, R., Hung, S.-H., Wu, X. & Hicks, M. A verified optimizer for quantum circuits. *Proceedings of the ACM on Programming Languages* **5,** 1–29 (2021).

92. Zhou, L., Barthe, G., Strub, P.-Y., Liu, J. & Ying, M. CoqQ: Foundational Verification of Quantum Programs. *Proceedings of the ACM on Programming Languages* **7,** 833–865 (2023).

93. Ying, M. Floyd–Hoare logic for quantum programs. *ACM Transactions on Programming Languages and Systems* **33,** 1–49 (2011).

94. Liu, J. *et al. Formal Verification of Quantum Algorithms Using Quantum Hoare Logic* in *Computer Aided Verification – CAV 2019* **11562** (Springer, 2019), 187–207.

95. Zhang, L., Jing, M., Zhao, Y., Liu, J. & Wang, X. *Lean-QuantumAlg: A Framework for AI Agents to Produce Quantum Algorithms that Humans Can Trust* Oct. 2026. arXiv: 2610.30216 [quant-ph].

96. Ren, Y. *et al. QEDIA: Quantum Program Verification in Lean* Oct. 2026. arXiv: 2610.17708 [quant-ph].

97. Ying, M. A practical quantum Hoare logic with classical variables, I. *Information and Computation* **309,** 105417. ISSN: 0890-5401 (2026).

98. Endo, K. *Reducing Trust Cost in Quantum Application Development Using AI* Oct. 2026. arXiv: 2610.14665 [quant-ph].

99. Morgan, C., McIver, A. & Seidel, K. Probabilistic Predicate Transformers. *ACM Transactions on Programming Languages and Systems* **18,** 325–353 (May 1996).

100. The Lean Developers. *The Lean Language Reference (Lean v4.35.0), Chapter 17: The* `vcgen` *Tactic* https://lean-lang.org/doc/reference/4.35.0/The--vcgen--tactic/. Lean FRO. Accessed: 2026-11-04. 2026.

101. Gladshtein, V., Pîrlea, G., Zhao, Q., Kurin, V. & Sergey, I. Foundational Multi-Modal Program Verifiers. *Proceedings of the ACM on Programming Languages* **10,** 2233–2264 (2026).

102. VERSE Lab. *Loom: A Framework for Foundational Multi-Modal Program Verifiers in Lean* https://github.com/verse-lab/loom. GitHub repository. Accessed: 2026-11-04. 2026.

103. Gudder, S. P. & Greechie, R. Effect Algebra Counterexamples. *Mathematica Slovaca* **46,** 317–325 (1996).

104. Baker, B. *et al. Monitoring Reasoning Models for Misbehavior and the Risks of Promoting Obfuscation* 2025. arXiv: 2503.11926 [cs.AI].

105. Schmidgall, S. & Moor, M. *AgentRxiv: Towards Collaborative Autonomous Research* 2025. arXiv: 2503.18102.

106. Jiang, Y., Zhang, Y., Shen, X., Backes, M. & Zhang, Y. *"Humans welcome to observe": A First Look at the Agent Social Network Moltbook* 2026. arXiv: 2602.10127.

107. Chen, Y., Niu, G., Cheng, J., Han, B. & Sugiyama, M. *When and Why Does Multi-Agent Debate Fail and Does It Really Underperform?* in *Proceedings of the 43rd International Conference on Machine Learning (ICML)* (2026). arXiv: 2510.20963 [cs.LG].

108. Jones, C. Low-overhead constructions for the fault-tolerant Toffoli gate. *Physical Review A* **87,** 022328 (2013).

109. Khattar, T. & Gidney, C. Rise of conditionally clean ancillae for efficient quantum circuit constructions. *Quantum* **9,** 1752 (2025).

110. Fowler, A. G., Mariantoni, M., Martinis, J. M. & Cleland, A. N. Surface codes: Towards practical large-scale quantum computation. *Physical Review A* **86,** 032324 (2012).

111. Horsman, D., Fowler, A. G., Devitt, S. & Van Meter, R. Surface code quantum computing by lattice surgery. *New Journal of Physics* **14,** 123011 (2012).

112. Bravyi, S. & Kitaev, A. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A* **71,** 022316 (2005).

113. Gidney, C. & Fowler, A. G. Efficient magic state factories with a catalyzed $|CCZ\rangle$ to $2|T\rangle$ transformation. *Quantum* **3,** 135 (2019).

114. Gidney, C., Shutty, N. & Jones, C. *Magic state cultivation: growing T states as cheap as CNOT gates* 2024. arXiv: 2409.17595 [quant-ph].

115. Gidney, C., Newman, M., Fowler, A. & Broughton, M. A Fault-Tolerant Honeycomb Memory. *Quantum* **5,** 605 (2021).

116. Kim, I. H. *et al.* Fault-tolerant resource estimate for quantum chemical simulations: Case study on Li-ion battery electrolyte molecules. *Physical Review Research* **4,** 023019 (2022).

117. Berry, D. W., Gidney, C., Motta, M., McClean, J. R. & Babbush, R. Qubitization of Arbitrary Basis Quantum Chemistry Leveraging Sparsity and Low Rank Factorization. *Quantum* **3,** 208 (2019).

118. Gu, A., Kemeny, A. H., Yelin, S. F. & Lukin, M. D. *Discovering state-of-the-art quantum circuits with large language models* Oct. 2026. arXiv: 2610.22367 [quant-ph].

119. Von Burg, V. *et al.* Quantum computing enhanced computational catalysis. *Physical Review Research* **3,** 033055 (2021).

120. Rubin, N. C. *et al.* Fault-Tolerant Quantum Simulation of Materials Using Bloch Orbitals. *PRX Quantum* **4,** 040303 (2023).

121. Litinski, D. & Nickerson, N. *Active volume: An architecture for efficient fault-tolerant quantum computers with limited non-local connections* 2022. arXiv: 2211.15465 [quant-ph].

122. Bravyi, S. *et al.* High-threshold and low-overhead fault-tolerant quantum memory. *Nature* **627,** 778–782 (2024).

123. Gidney, C. & Ekerå, M. How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum* **5,** 433 (2021).

124. Jeffery, S., Magniez, F. & de Wolf, R. Optimal Parallel Quantum Query Algorithms. *Algorithmica* **79,** 509–529. ISSN: 1432-0541 (2016).

125. Bennett, C. H., Bernstein, E., Brassard, G. & Vazirani, U. Strengths and Weaknesses of Quantum Computing. *SIAM Journal on Computing* **26,** 1510–1523 (1997).

126. Høyer, P. Arbitrary phases in quantum amplitude amplification. *Physical Review A* **62,** 052304 (2000).

127. Long, G. L. Grover algorithm with zero theoretical failure rate. *Physical Review A* **64,** 022307 (2001).

128. Brassard, G., Høyer, P., Mosca, M. & Tapp, A. in *Quantum Computation and Information* 53–74 (American Mathematical Society, 2002). arXiv: quant-ph/0005055.

129. Yoder, T. J., Low, G. H. & Chuang, I. L. Fixed-Point Quantum Search with an Optimal Number of Queries. *Physical Review Letters* **113,** 210501 (2014).

130. Aldenbro, D. & Skalski, A. *Comparing Grover's algorithm to linear search when finding multiple queried elements* Bachelor's thesis (KTH Royal Institute of Technology, 2024).

131. Mizel, A. Critically Damped Quantum Search. *Physical Review Letters* **102,** 150501 (2009).

132. Andrés-Martínez, P. & Heunen, C. Weakly measured while loops: peeking at quantum states. *Quantum Science and Technology* **7,** 025007 (2022).

133. Grover, L. & Rudolph, T. *Creating superpositions that correspond to efficiently integrable probability distributions* 2002. arXiv: quant-ph/0208112.

134. Harrow, A. W., Hassidim, A. & Lloyd, S. Quantum Algorithm for Linear Systems of Equations. *Physical Review Letters* **103,** 150502 (2009).

135. Temme, K., Osborne, T. J., Vollbrecht, K. G., Poulin, D. & Verstraete, F. Quantum Metropolis sampling. *Nature* **471,** 87–90 (2011).

136. Ozols, M., Roetteler, M. & Roland, J. Quantum rejection sampling. *ACM Transactions on Computation Theory* **5,** 1–33 (2013).

137. Sanders, Y. R., Low, G. H., Scherer, A. & Berry, D. W. Black-Box Quantum State Preparation without Arithmetic. *Physical Review Letters* **122,** 020502 (2019).

138. Burge, I., Barbeau, M. & Garcia-Alfaro, J. *Quantum CORDIC – Arcsine on a Budget* 2024. arXiv: 2411.14434 [quant-ph].

139. Khattar, T. *et al. Leveraging verifiable structure in quantum circuits to find more efficient compilations with agentic AI* 2026. arXiv: 2610.23832 [quant-ph].

140. Motlagh, D. & Wiebe, N. Generalized Quantum Signal Processing. *PRX Quantum* **5,** 020368. arXiv: 2308.01501 [quant-ph] (2024).

141. Berry, D. W., Motlagh, D., Pantaleoni, G. & Wiebe, N. Doubling the efficiency of Hamiltonian simulation via generalized quantum signal processing. *Physical Review A* **110,** 012612 (2024).

142. Eklund, E. C. *et al. End-to-End Simulation of Chemical Dynamics on a Quantum Computer* 2026. arXiv: 2603.19007 [quant-ph].

143. Nacu, Ş. & Peres, Y. Fast simulation of new coins from old. *The Annals of Applied Probability* **15,** 93–115 (2005).

144. Campbell, E. Shorter gate sequences for quantum computing by mixing unitaries. *Physical Review A* **95,** 042306 (2017).

145. Hastings, M. B. Turning gate synthesis errors into incoherent errors. *Quantum Information and Computation* **17,** 488–494. arXiv: 1612.01011 (2017).

146. Haah, J. Product Decomposition of Periodic Functions in Quantum Signal Processing. *Quantum* **3,** 190. arXiv: 1806.10236 (2019).

147. Menczer, A. *et al.* Parallel Implementation of the Density Matrix Renormalization Group Method Achieving a Quarter petaFLOPS Performance on a Single DGX-H100 GPU Node. *Journal of Chemical Theory and Computation* **20,** 8397–8404 (2024).

148. Rubin, N. *Fast quantum simulation of electronic structure by spectral amplification* Zenodo. 2025.

149. Abrams, D. S. & Lloyd, S. Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors. *Physical Review Letters* **83,** 5162–5165 (1999).

150. Lee, S. *et al.* Evaluating the evidence for exponential quantum advantage in ground-state quantum chemistry. *Nature Communications* **14,** 1952 (2023).

151. King, R., Low, G. H., Babbush, R., Somma, R. D. & Rubin, N. C. Quantum Simulation with Sum-of-Squares Spectral Amplification. *Physical Review Letters* **136,** 110601 (2026).