

SuperFermion: A Leap Forward in Quantum Circuit Simulation Efficiency

Statevector learning matrix product states dynamic circuits and specialized execution modes

Md. Hasibul Hassan Himal

Catstate hasibul265686@gmail.com ORCID: 0009-0007-3613-5408

Muhebbullah Tamim

Catstate tamim@catstate.dev ORCID: 0009-0004-3790-1408

Preprint

Abstract

Efficient quantum simulation requires execution paths that match circuit structure and requested outputs. This study evaluates SuperFermion using three recorded benchmark suites covering statevectors, matrix product states (MPS), dynamic circuits and specialized execution modes. Six statevector configurations share 252 conditions from 4 to 20 qubits, including five-update variational learning from 4 to 12 qubits. The selected corpus contains 1,512 simulator jobs and 140 controls. SuperFermion has the lowest reconstructed primary-compute total, 27.36 min, compared with 34.30 min for Qiskit. Median paired process-memory savings are approximately 16.6% against Qiskit, 44.1% against PennyLane and 34.6-34.7% against Cirq; Qulacs memory is similar and slightly lower. The MPS suite contributes 182 selected unique jobs, including 141 successful observations. Structured circuits reach 1,000 qubits, and native MPS teleportation reaches 120 qubits with passing target-bit checks. At 120 qubits and 1,000 shots, SuperFermion records 0.318 s versus 1.015 s for Aer, while attributed interleaved checks report a narrower ratio near 1.6. The later suite adds 24 conditions across tableau sampling, compiled adjoint gradients and dense-core density matrices. Stabilizer and Clifford labels share one engine. SuperFermion is faster than every available external comparator in 10 of 11 prepared sampling conditions but only four cold conditions. All ten prepared gradients are faster than Qulacs, with maximum saved component-wise discrepancy $2.8172\text{e-}15$ and one independent eight-qubit finite-difference anchor. The dense-core density-matrix route is faster than Aer at eight qubits and slower at twelve; the ten-qubit Aer diagonal-check flag leaves that comparison unranked. MPS truncation counterexamples, process-memory exceptions, preparation costs, unequal concurrency, different output contracts and incomplete repeated-run records constrain generalization. The evidence supports workload-specific execution benefits and preparation reuse; the three suites retain separate timing definitions and correctness scopes.

Keywords: quantum circuit simulation; SuperFermion; quantum machine learning; matrix product states; dynamic circuits; stabilizer sampling; Clifford circuits; adjoint gradients; density matrices

1 Introduction

1.1 Motivation and research questions

Classical simulation supports quantum algorithm development by making circuit behavior, numerical error and execution cost inspectable. Quantum kernels, variational learning and approximate

optimization repeatedly evaluate related circuits, so lower execution cost can make broader experiments practical. This study evaluates SuperFermion as a quantum simulation framework across quantum machine learning, objective evaluation, circuit sweeps, matrix product states, dynamic circuits, Clifford sampling, adjoint gradients and density-matrix evaluation. Its contribution is an assessment of software execution efficiency under recorded conditions.

Quantum feature maps represent observations as states whose overlaps define a kernel for a classical estimator. Variational methods instead adjust circuit parameters using an objective built from observable expectations. These alternatives impose different demands on simulation and readout [1, 2, 3, 4]. Returning a complete statevector and returning a scalar observable therefore require separate analysis, even when the logical circuit is the same.

The study asks how SuperFermion compares with external frameworks on matching workloads, how computation and process memory change with qubit count and output contract, and how bounded-bond simulation extends the tested circuit sizes. The six statevector configurations are SuperFermion, Qulacs, Qiskit, PennyLane, Cirq f32 and Cirq f64. The MPS suite uses SuperFermion, Qiskit Aer, PennyLane, Cirq with quimb, and Qulacs as an explicitly not-offered adapter path.

The statevector analysis includes 1,512 simulator records on 252 identical conditions per configuration, plus 105 classical controls and 35 analytical controls. The replacement MPS archive contributes 182 selected unique jobs, of which 141 complete successfully. The later execution-mode suite adds 24 conditions and 81 selected framework observations. Together, the experiments describe measured behavior across exact-state outputs, native readout, truncation diagnostics and measurement-driven execution. Cases where another framework is faster, uses less recorded memory or has lower truncation error remain part of the comparison.

The specialized paths address different computational regimes. A stabilizer tableau exploits the Clifford gate class without requiring low entanglement, whereas an MPS exploits a bounded auxiliary dimension and can support general rotations with local truncation. Reusable dense-state adjoint handles target repeated variational gradients. We therefore assess representation, preparation and requested output separately, rather than pooling incompatible timing contracts. [5, 6, 7, 8]

1.2 Related work and contribution

Qulacs targets efficient quantum circuit simulation, while Qiskit Aer, PennyLane, Cirq and qsim provide distinct simulation and programming interfaces [9, 10, 4, 11, 12]. Kernel and variational studies motivate evaluation beyond simple gate throughput because end-to-end tasks can also depend on readout, repeated state construction and dense matrix operations. The present comparison describes the configured adapters in the supplied studies rather than every execution path available in those frameworks.

Matrix product states represent a state through tensors linked by bounded auxiliary dimensions. Their favorable storage depends on the circuit’s entanglement structure; increasing the maximum bond can increase both accuracy and cost [13, 6]. Large-qubit results must consequently be read together with the measured circuits, bond limits, diagnostics and requested outputs.

The contribution is a source-traceable evaluation of SuperFermion across a broad workload grid. Separate experiment-and-qubit tables preserve feature maps and output contracts, while matched scatter plots and all five framework heatmaps expose both advantages and counterexamples. The integrated paper contains 132 tables and 42 figures. The later methods report adds mathemati-

cal descriptions of canonical MPS repair, two-site splitting, symbolic tableau measurements and prepared gradients, while the execution-mode suite adds measured sampling, gradient and density-matrix results. [14, 8] No simulator measurements were generated during this manuscript revision.

2 Materials and methods

2.1 Quantum representations and output contracts

For a feature map U , a classical observation x determines a state and a fidelity kernel:

$$|\phi(x)\rangle = U(x)|0\rangle^{\otimes n}, \quad K(x, y) = |\langle\phi(x) | \phi(y)\rangle|^2. \quad (1)$$

Kernel classification and regression use these overlaps in a classical estimator. Quantum kernel principal component analysis uses a kernel representation for dimensionality analysis. Reported primary runtimes can include more than state preparation, and the saved summaries do not isolate every matrix operation or estimator step. Variational classification, regression and objective evaluation instead use expectation values of an observable O :

$$f(x, \theta) = \langle\psi(x, \theta)|O|\psi(x, \theta)\rangle. \quad (2)$$

The statevector track exports a complete state and carries out its associated readout. The native-readout track requests an observable through the framework’s native interface. These routes have different materialized outputs and can use different execution primitives. Condition keys retain the track so that a scalar-output observation is not silently pooled with a state-export observation.

2.2 SuperFermion execution design

SuperFermion exposes circuit construction, a native intermediate representation and state or observable execution. The design considered here prepares reusable execution structure and offers statevector and native-readout routes, with batching where the study requests it [15]. Figure 1 describes that organization. The records do not provide a component-level ablation, so they cannot assign a measured improvement to a particular kernel, cache or dispatch change.




Figure 1: Conceptual organization of prepared circuit execution and its two output contracts. This schematic describes the software design; the saved measurements do not isolate each component’s contribution.

A native implementation can reduce repeated dispatch, avoid unnecessary state export and retain prepared work across related evaluations. These are useful engineering mechanisms, but their individual effects must be established through controlled profiling. The recorded task costs are the evidence for this paper; the schematic is not a measurement of each mechanism.

Figure 2 distinguishes the representations used by the recorded paths. The stabilizer and Clifford labels share the same Pauli-frame engine; the gradient handle and density-matrix core are separate dense-state paths. The later report describes MPS mechanisms but does not establish which native source revision produced every earlier campaign record. [14, 8]

Clifford circuits Binary stabilizer tableau Affine measurement plan 64-shot sampling lanes	Bounded entanglement MPS canonical tensors Local two-site split Sequential conditionals	Variational circuits Dense statevector Compiled adjoint handle Reverse gradient sweep
---	--	--

Dense density matrices use a separate native-core path

Figure 2: Representations associated with the reported execution paths. Stabilizer and Clifford mode labels share one tableau engine; the adjoint handle uses a different statevector path. MPS mechanisms are described in the methods report and are not benchmarked by the new mode suite. The density-matrix observations exercise a separate native-core path.

2.3 Evidence and workload selection

The statevector evidence comes from a completed campaign overview and its job-by-job report [16, 17]. The included six configurations share exactly the same 252 condition keys. Every included statevector job is recorded as successful. The comparison subset contains 1,652 rows when its 140 controls are included. The earlier incomplete report is retained as source evidence but does not contribute to these estimates.

The three suites remain separate evidence strata. The statevector campaign supplies matched task records and controls, the MPS archive supplies bounded-bond and dynamic-circuit observations, and the later executed notebook supplies 24 specialized conditions. Duplicate copies within the execution-mode archive are reconciled as the same dataset. No combined compute total or global framework ranking is formed across these different workloads. [16, 18, 8]

The classification datasets are breast_cancer, iris, wine, digits_01 and moons; regression uses diabetes and friedman1. Feature selection uses breast_cancer and wine. MaxCut covers ring and Erdos-Renyi graph families, VQE uses an open-chain transverse-field Ising objective, and gradient and circuit tasks use synthetic inputs. The final reports specify a full 80/20 held-out split and seed 42. Complete sample cardinalities, preprocessing transformations and leakage audits are not independently available from the displayed records, so no additional dataset sizes are assumed.

Table 1: Selected statevector study protocol

Item	Recorded setting or evidence boundary
Comparison set	Six configurations; 252 identical conditions per configuration
Qubit grid	4, 8, 12, 16 and 20; VQC/VQR limited to 4, 8 and 12
Learning split	Full 80/20 held-out split; seed 42
Training budget	VQC/VQR: five updates
Repeated measurements	One fresh-process repetition per condition
Job timeout	1,200 s
State-bank cap	512 MiB exported-bank limit; not a total process-RSS cap
Selected coverage	1,512 simulator jobs plus 105 classical and 35 analytical controls
Implementation metadata	Full package lock, pinned source patch and detailed adapter settings unavailable
Learning metadata	Row-level predictive outcomes and complete preprocessing audit unavailable

These are reported study settings. Descriptive comparisons retain the supplied output contracts and precision differences.

Table 2: Matched workload coverage per simulator lane

Experiment	Inputs and configured outputs	Qubits	Conditions
QSVM	5 datasets; angle and ZZ; statevector	4/8/12/16/20	50
QKRR	2 datasets; angle and ZZ; statevector	4/8/12/16/20	20
QKPCA	5 datasets; ZZ; statevector	4/8/12/16/20	25
VQC	5 datasets; angle; two output tracks; 5 updates	4/8/12	30
VQR	2 datasets; angle; two output tracks; 5 updates	4/8/12	12
QAOA feature selection	2 datasets; two output tracks	4/8/12/16/20	20
QAOA MaxCut	2 graph families; two output tracks	4/8/12/16/20	20
VQE	Open-chain TFIM; two output tracks	4/8/12/16/20	10
Gradient	Synthetic objective; two output tracks	4/8/12/16/20	10
Circuit	7 families plus 4 batch/depth sweep settings	4/8/12/16/20	55
Total	201 statevector and 51 native-readout conditions	Measured grid	252

Every included framework has the same 252 condition keys. The six configurations contribute 1,512 simulator jobs; 105 classical controls and 35 analytical controls bring the analyzed statevector corpus to 1,652 rows.

Cirq f32 is the lower-precision configuration identified by the overview, while Cirq f64 is retained as a separate comparator. Detailed simulator flags and fusion settings are not independently exposed in the final statevector reports. The five external entries therefore describe the study’s configured execution paths. Analytical-angle and classical estimators remain controls rather than general simulator competitors.

2.4 Timing memory and reconstruction

Statevector time is the reported primary compute stage, converted to seconds for analysis. It differs from process wall time, which also includes work outside that stage. Worker peak RSS measures the entire fresh process, including interpreter, imports, preparation and simulation. Lower process memory may therefore reflect a lighter execution environment as well as numerical allocation.

Each record is keyed by framework, task, dataset, variant, map, output track, qubits, batch and depth. The reconstruction verifies uniqueness and the equality of the six 252-condition sets. Mil-

liseconds are converted to seconds and GiB to MiB using a factor of 1,024. Original display strings and source line numbers remain in the analysis CSV. Matched time and memory comparisons use:

$$S_{b,c} = \frac{T_{b,c}}{T_{\text{SF},c}}, \quad \Delta R_{b,c} = 100 \left(1 - \frac{R_{\text{SF},c}}{R_{b,c}} \right). \quad (3)$$

The SF subscript denotes SuperFermion. A time ratio above one favors SuperFermion, and a positive memory saving means its process peak is smaller. A median of condition-level ratios describes the typical retained pair; a ratio of compute sums describes the aggregate cost of a specified grid and gives more weight to expensive conditions. The main heatmaps use the latter definition.

Displayed zero latencies are censored at report precision and are never treated as zero computational work. They are omitted from logarithmic time scatters and from positive-pair medians. Display-rounding intervals are reconstructed separately. A condition has a precision-separated leader only when one upper rounding bound is below every other configuration’s lower bound. The 35 unresolved conditions are not assigned a winner. These bounds concern display rounding and are not statistical confidence intervals.

Detailed tables summarize each experiment at each measured qubit count. Feature maps and output tracks stay separate. Dataset or circuit medians are cohort summaries, not repeated-run estimates. VQE and gradient tracks with one condition retain the individual observation. QKPCA and circuit tables also show maxima so that a cohort median does not conceal the expensive cases.

2.5 Matrix product state and measurement protocol

All MPS and mid-circuit results in this revision come from the replacement archive [18]. Its append-only JSONL contains 282 rows, which reduce to 264 unique jobs under the supplied reducer key. Selecting the five included configurations yields 182 unique jobs. Their final statuses, numeric values and trial medians agree with the supplied CSV. No observations from the earlier MPS archive enter the current estimates, tables or figures.

The MPS state representation and bounded-bond storage are:

$$|\psi\rangle = \sum_{s_1, \dots, s_n} A_1^{s_1} A_2^{s_2} \cdots A_n^{s_n} |s_1 \cdots s_n\rangle, \quad M_{\text{MPS}} = O(n\chi^2). \quad (4)$$

Here χ is the maximum auxiliary bond dimension. This storage expression is useful for circuits whose required bonds remain bounded; it does not establish efficient simulation of arbitrary states. The ordinary MPS ladder and stress paths use one untimed warm call followed by three timed calls and retain their median. The saved adapter source constructs circuit or intermediate representation before those timers, but execution and observable handling remain adapter-specific.

Table 3: MPS and dynamic circuit measurement protocol

Item	Recorded setting or limitation
Included records	182 unique jobs; 141 OK, 12 budget, three error and 26 not offered
Circuit families	GHZ, brick4, brick12, QFT, teleportation, measurement-based wire and repetition code
Structured size grid	50 to 1,000 qubits at maximum bond 64; stress at 200, 500 and 1,000
Accuracy anchors	GHZ/brick4 at 12 and 16; TVD at 16; bond sweeps from 2 to 128
Ordinary timing	One warm call and three timed calls; median recorded
Validation paths	Dense reference and TVD may use single-pass or extraction work; eight successful rows have no time
Teleportation	12 qubits: 500 shots; 24-qubit boundary: 500 shots and a common 90 s budget
Other dynamic cases	17-qubit wire: 200 shots; repetition code: four patterns with 200 shots per pattern
Native MPS teleportation	100 shots at 12; 1,000 at 30, 60 and 120 qubits
Memory metric	Process RSS snapshot/high-water value in MiB, potentially including sampling and earlier chunk work
Sampling	Separate sample-time field; excluded from primary MPS compute time
Failure treatment	Budget, error, not-offered and absent-job outcomes remain distinct

MPS and statevector time totals are not pooled. Dynamic paths use a single timed replay per recorded condition.

Table 4: MPS software versions and execution paths

Framework	Version	Recorded path
SuperFermion	0.1.15	Native MPS state and dynamic measurement interfaces
Qiskit Aer	0.17.2	matrix_product_state plus conditional execution
PennyLane	0.45.0	default.tensor for MPS; default.qubit trajectories for MCM
Cirq quimb	Cirq 1.6.1 / quimb 1.13.0	MPS simulator; separate trajectory measurement path
Qulacs	0.6.14	No offered MPS/feed-forward path in this adapter

The SuperFermion release label is updated to 0.1.15 for this preprint revision; the benchmark values are retained from the archived runs and were not remeasured for the release update. The probe and row identifiers remain in the source archive. The 40 completed SuperFermion rows include 38 with compiled identifier 99304153f6... and two refreshed 12-qubit teleportation/repetition-code rows with f576d0eb7a.... The timeout row lacks an identifier. The records do not establish a single-revision implementation effect.

Sampling is measured separately and may occur before the RSS snapshot. Ordinary MPS time therefore excludes sampling, while the associated process peak can include it. Chunks can contain multiple jobs, so a peak may also include earlier work in that process. The field suffix mb is misleading: the recorder divides bytes by 2 to the power 20, making both memory fields MiB. These observations do not isolate incremental tensor allocation.

Dynamic correctness checks use deterministic target outputs. Teleportation starts from the specified one-bit input, the measurement-based wire checks its final target, and the repetition-code worker tests no error and one bit flip on each of three data bits. The repetition-code wall covers all four constructions and replays, with 200 shots per pattern and 800 in total. Other dynamic walls cover the single specified replay. Matching shots within a case does not make different circuit cases equivalent workloads.

Accuracy is assessed using the Z0 Zlast correlation, dense-versus-MPS small-circuit anchors, and total variation distance at 16 qubits:

$$\text{TVD}(p, q) = \frac{1}{2} \sum_z |p(z) - q(z)|. \quad (5)$$

Discarded weight and the reported fidelity lower bound are engine diagnostics at larger sizes. Their definitions and aggregation policies vary across engines. A matching correlation or a favorable diagnostic does not independently certify equality of complete states or output distributions. Fixed maximum bond is not a matched-accuracy constraint.

2.6 Correctness and controls

The statevector overview reports a completed correctness gate, backend checks and gradient tests, including a median per-job maximum gradient error of 4.996e-10 and a largest reported value of 1.448e-7 [16]. Detailed correctness files are not supplied, so those statements remain attributed summaries. Successful job status alone does not replace a numerical correctness audit.

The analytical-angle control evaluates a separable RY-map kernel directly:

$$K_{\text{angle}}(x, y) = \prod_{j=1}^n \cos^2\left(\frac{\theta_j(x) - \theta_j(y)}{2}\right). \quad (6)$$

Its $O(n)$ arithmetic per data pair can avoid state simulation when the specified map has that structure, although constructing a dense Gram matrix remains quadratic in sample count. It is not applicable to every entangling map or variational circuit. Classical controls include RBF SVC, logistic regression, random-forest classification, ridge regression, RBF kernel ridge regression and random-forest regression. Their nominal qubit label only indexes a paired study setting.

2.7 MPS canonical structure and local repair

In the open-chain MPS representation above, each site tensor has a left bond, a right bond and a binary physical index. Both boundary bond dimensions equal one, and internal dimensions are bounded by χ . The $O(n \chi^2)$ storage order counts complex entries and is not a measured process-memory result.

The methods report describes a default bond cap of 64 and each site as two physical blocks stacked in a matrix. Left QR sweeps absorb the triangular factor into the next site; right LQ sweeps establish an isometric suffix. The corresponding canonical conditions are

$$\sum_s (A_i^s)^\dagger A_i^s = I, \quad \sum_s B_i^s (B_i^s)^\dagger = I. \quad (7)$$

The report describes a dirty-window frontier for right canonicity and a cached left measurement environment. Rewriting a local tensor invalidates the affected frontier, so subsequent measurement repairs the needed interval instead of sweeping the entire chain unconditionally. These mechanisms reduce repeated preparation work when their invariants hold; the mode suite does not contain a component ablation measuring their individual contribution. [14, 6]

2.8 Two site gates and truncation

A two-qubit gate first contracts the shared bond and acts on the physical indices. With l and r denoting the outer bonds and a the shared bond, the merged matrix can be written as

$$M_{(l,o_1),(o_2,r)} = \sum_{a,i_1,i_2} g_{o_1 o_2, i_1 i_2} (A_i^{i_1})_{la} (A_{i+1}^{i_2})_{ar}. \quad (8)$$

The methods report describes an exact structured split when a block-diagonal gate admits a factorization whose rank fits the bond cap. That path avoids a full singular-value decomposition for the qualifying merge. The general default uses a thin SVD and a relative pruning threshold of $1e-12$ times the merged Frobenius norm. An exact retained singular subspace gives the rank- k approximation; an orthonormal retained basis also permits the projected-mass formulation

$$M_k = M V_k V_k^\dagger, \quad \varepsilon = \frac{\|M\|_F^2 - \|M V_k\|_F^2}{\|M\|_F^2}, \quad V_k^\dagger V_k = I. \quad (9)$$

The report describes selecting retained columns using measured norms of M v when reported singular values and vectors are inconsistent. This is a numerical safeguard described by that source, not an independently reproduced native-code finding. It also distinguishes the default SVD route from a legacy positional QR cut. The internal values summarized below are source-comment observations, kept separate from the recorded execution-mode benchmark. [14]

Table 5: Report attributed internal MPS observations

Condition	Reported legacy	Reported default	Evidence scope
Brick12, 200 qubits	2.18 s; reported discard metric about $5.2e-4$	1.04 s; accumulated discarded mass about $6.7e-14$	Source-comment timing; telemetry labels differ
Brick12, 26 qubits; cap 64	Reported error $4.2e-6$	Reported error $3.6e-14$	The report does not define the error metric here
Same 26-qubit condition	Accumulated discarded mass $6.7e-2$	Accumulated discarded mass $1.1e-16$	Telemetry is not certified global fidelity

Values are attributed to internal measurements in native source comments summarized by the supplied methods report. The 200-qubit legacy discard metric has an unspecified aggregation, while the default is stated as an accumulated sum. No call-level record, complete recipe or native source revision is included. These observations do not enter framework rankings, measured RAM comparisons or the 24-condition suite statistics. [14]

Truncation telemetry records per-step discarded mass, event counts and its accumulated sum. The field named `fidelity_lower_bound` is described in the methods report as a diagnostic rather than a certified global fidelity bound. Repeated truncation, renormalization and circuit structure can change the relationship between local discarded mass and final-state error. Quantitative accuracy claims consequently require independent state or probability comparisons.

2.9 MPS routing observables and sampling

Lazy routing tracks a logical-to-site permutation. Bringing distant sites together still requires adjacent tensor operations, but leaving the resulting permutation in place can avoid unconditional

swap-back work. The bookkeeping update is constant time; that statement does not make a tensor SWAP itself constant time.

Sequential sampling uses the probability chain rule. In a gauge with a right-isometric tail, a prefix vector $\mathbf{v}$ can be propagated through the chosen physical block. Normalized local branch probabilities and the complete sample distribution are

$$p(s_i | s_{<i}) = \frac{\|v_{i-1} A_i^{s_i}\|^2}{\sum_{b=0}^1 \|v_{i-1} A_i^b\|^2}, \quad p(s) = \prod_{i=1}^n p(s_i | s_{<i}). \quad (10)$$

Under the stated canonical structure and fixed physical dimension, per-sample contraction work is linear in chain length and quadratic in bond dimension, without constructing a dense statevector. Measurement projects and renormalizes the chosen branch; reset can then apply X when the observed bit is one. A Pauli-product expectation can likewise be evaluated with a sequence of transfer matrices,

$$E_i = \sum_{s,t} (P_i)_{st} (A_i^s)^\dagger E_{i-1} A_i^t, \quad E_0 = 1, \quad \psi^\dagger (P_1 \otimes \cdots \otimes P_n) \psi = E_n. \quad (11)$$

Bounded-bond storage is useful only when the retained bond dimensions describe the state accurately. A Clifford circuit can create extensive entanglement and remain tractable for a stabilizer tableau, while a fixed-bond MPS can require truncation. Conversely, the Clifford class by itself does not guarantee that a particular tableau implementation wins every timing comparison. The new mode suite contains no shared-workload head-to-head MPS-versus-tableau measurement. [5, 6, 8]

2.10 Pauli frame tableau and symbolic measurements

A pure stabilizer state can be represented with stabilizer and destabilizer rows containing binary X and Z indicators and phase bits. The described implementation packs row data into 64-bit words. Tableau storage is quadratic in qubit count; elementary H , S and controlled- X updates touch selected bit positions in each row. This storage order should not be confused with a universal measured gate-time complexity or an automatic 64-fold speedup. [14, 5]

For H on qubit q , the X and Z bits exchange and the phase toggles when both were set. For S , the Z bit is XORed with X and the same phase correction applies. The controlled- X update combines the two target bit positions with a Boolean phase term. The report describes fusing the operation stream row by row and using word-level phase products to avoid repeated per-bit traversal.

The sampling preparation performs the symbolic stabilizer measurement procedure once. Independent random branches introduce unbiased binary variables; determined outcomes become affine parity functions of those variables. In vector notation,

$$\mathbf{y} = A\mathbf{r} \oplus \mathbf{c}, \quad \mathbf{r} \sim \text{Uniform}(\mathbb{F}_2^k). \quad (12)$$

The described sampler evaluates these parity expressions across 64 independent shot lanes packed into a word, then transposes them into output keys. Sparse dependency lists reduce parity work when the measurement structure permits it. A column index accelerates finding active tableau rows, and repeated strings are materialized once per distinct count key. GHZ conditions have two

reported keys, whereas the tested brickwork and random-Clifford conditions record 2,000 distinct keys; output structure is therefore part of the measured work. [14]

The tableau path applies to Clifford circuits. Non-Clifford gates require another execution path; a general variational rotation is not represented by this Pauli-frame sampler. Stim is a specialized stabilizer comparator, and its recorded sample-array contract differs from the SuperFermion counts-map contract. [19, 8]

2.11 Compiled adjoint gradients

The gradient experiment uses RY and RZ on every qubit in each layer, followed by a controlled-X ladder. It has $P = 2nL$ parameters and the full Z-product observable. A compiled handle separates circuit construction from repeated gradient evaluation. For a parameterized gate in a circuit with G gates, a reverse sweep uses the forward state and the backward observable state,

$$\frac{\partial f}{\partial \theta_j} = 2 \operatorname{Re} \left(\lambda_j^\dagger \frac{\partial U_j}{\partial \theta_j} \psi_{j-1} \right), \quad \lambda_j = U_{j+1}^\dagger \cdots U_G^\dagger O \psi_G. \quad (13)$$

This form avoids independent full-circuit replay for every parameter and can reuse a constant number of dense statevectors. Its memory still grows exponentially with qubit count. The compiled-handle experiment measures this dense-state gradient route, not the Clifford tableau engine. [7, 8]

2.12 Dense density matrix checks

For the noiseless unitary circuits in this suite, the expected density matrix is a pure-state outer product. Its diagonal agrees with statevector probabilities, and its trace and purity are one:

$$\rho = \psi \psi^\dagger, \quad \rho_{jj} = \psi_j \overline{\psi_j}, \quad \operatorname{Tr} \rho = 1, \quad \operatorname{Tr}(\rho^2) = 1. \quad (14)$$

These checks are useful but do not identify a complete density matrix. For example, the two states below have the same diagonal, trace and purity but different coherences,

$$\rho_+ = \frac{1}{2} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, \quad \rho_- = \frac{1}{2} \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}. \quad (15)$$

A complete density-matrix validation would compare off-diagonal entries or suitable informationally complete observables, with explicit tolerance. No noisy-channel workload is measured by these three pure-state conditions. Dense density-matrix storage grows as four to the power of the qubit count in complex entries; this is a representation count, not a measured RAM estimate.

2.13 Execution mode benchmark protocol

The executed notebook contains no saved error outputs. The analysis reconciles all 24 unique JSONL condition keys with the four supplied CSVs and confirms that the nested results ZIP is a duplicate copy of the same recorded data. The selected comparison has 81 framework observations after retaining the primary SuperFermion execution path and external comparators. The archive

provides aggregate window statistics rather than complete timed-call arrays, so repeated-run uncertainty cannot be reconstructed. [8]

Table 6: Recorded execution mode protocol

Experiment	Recorded grid	Timing and output contract
Stabilizer	6 conditions; 64 to 1,024 qubits	GHZ, depth-eight brickwork and random Clifford; 2,000 shots
Large Clifford	5 conditions; 1,024 and 4,096 qubits	Same Pauli-frame path in the larger regime; 2,000 shots
Gradient	10 conditions; 6 to 20 qubits	One layer at every measured size; two and four layers at 8 qubits; $P = 2 \times n \times L$
Density matrix	3 conditions; 8, 10 and 12 qubits	Noiseless depth-eight brickwork; separate native dense-core path
Concurrency	SuperFermion: Rayon limit 4	Competitors use package defaults; equal-concurrency comparison was not recorded
Random seeds	Brickwork 11; random Clifford 7	Gradient parameters use seed 7 with draws between 0.1 and 1.2
Sampling repetitions	Two reversed comparison windows	SF and Stim: 7 timed calls/window; Aer and PennyLane: 2 cold and 3 warm; Cirq: 1
Gradient repetitions	Two forward comparison cycles	7 timed gradient calls/window; SF handle compile cost reported separately
Density repetitions	Two SF windows, then two Aer windows	SF: 7 calls/window; Aer: 5 after an excluded warm call
Timing boundaries	Native process startup excluded	Recorded cold contracts differ; SF warm counts versus Stim sample array
Resource guards	40 s first-call criterion	Cirq capped above 64 qubits; stabilizer Aer above 256; budget checked after the call
RAM	No RSS or allocation measurements	No measured memory ranking is inferred from this suite

Timings are means of two window medians. Adjoint and density-matrix order are not the reversed tableau order. Quoted repeat counts describe the saved adapter; per-call values are absent.

Tableau windows run SuperFermion, Stim, Aer, PennyLane and Cirq in forward and then reverse order. Gradient windows instead use two forward comparison cycles. Density-matrix windows are grouped as two SuperFermion calls followed by two Aer calls in the notebook code. This distinction matters because a general description of all sections as ABBA-interleaved would be inaccurate.

SuperFermion cold sampling includes native tableau construction, measurement-plan preparation and first execution inside its timed region. Stim full timing includes sampler compilation and sampling, with Python circuit assembly already outside that region. Aer, PennyLane and Cirq have their own fresh-construction boundaries. SuperFermion warm timing produces a count map; Stim warm timing produces a sample array. The additional Aer reuse measurement is retained in each detailed sampling table but is separate from the archive’s warm headline.

The saved software versions are Stim 1.16.0, Qiskit 2.1.2 with Aer 0.17.2, PennyLane 0.45.0, Cirq 1.6.1 and Qulacs 0.6.14. The native example has the recorded short binary identifier c0c9679b215e and configuration identifier 3b88281923. These identify the measured implementation but do not replace a full source manifest. The Rust execution paths in the suite are engine-level measurements through a release example; the supplied README states that those paths are not yet exposed by the shipped Python package.

The Qulacs gradient adapter compensates for its rotation convention by setting negative angles and

negating the returned backpropagation derivative. Gradient parity refers to those compensated derivatives. Construction of the SuperFermion handle remains separate from its timed evaluations; the recorded Qulacs function includes parameter updates, state allocation, circuit application and backpropagation.

The density-matrix observations call the native core’s QuantumDAG.simulate_dm implementation through the example. The separate density-matrix planner mode is held out and is not claimed as implemented by these timings. The supplied methods report traces implementation mechanisms to native source files, while the ZIP ships the benchmark adapter and executed outputs rather than those complete native implementations. Mechanism descriptions are consequently attributed to the report; no component ablation is inferred.

3 Results

3.1 Overall computation and process memory

SuperFermion has the lowest reconstructed successful primary-compute total among the six included configurations: 27.36 min over the common 252 conditions. Qiskit totals 34.30 min, Qulacs 70.04 min, Cirq f32 72.38 min, PennyLane 132.99 min and Cirq f64 138.52 min. The Qiskit-to-SuperFermion total ratio is approximately 1.254. These totals summarize the selected workload grid and omit startup, imports, retries and discarded attempts.

Table 7: Reconstructed performance on 252 matched conditions per configuration

Framework	Compute median (ms)	Compute total (min)	Median pair ratio (n)	RAM median (MiB)	Paired RAM saving (%)
SuperFermion	29.70	27.36	-	150.60	-
Qulacs	100.95	70.04	1.083 (238)	149.55	-0.37
Qiskit	76.30	34.30	2.291 (242)	184.70	16.56
PennyLane	163.95	132.99	9.216 (242)	268.95	44.15
Cirq f32	212.85	72.38	14.405 (242)	228.60	34.63
Cirq f64	548.10	138.52	34.190 (242)	229.35	34.68

Pair ratios compare positive displayed times with SuperFermion; n is the retained pair count. Totals weight expensive conditions more strongly than paired medians. Memory is whole-worker peak RSS; negative saving means SuperFermion’s peak is higher. No repeated-run confidence intervals are available.




Figure 3: Six included simulator configurations evaluated on the same 252 conditions. Compute totals sum the reported primary stage and exclude startup, imports and retries. Memory is the median whole-worker peak RSS. Cirq f32 is a mixed-precision comparison.

The typical positive paired Qulacs time ratio is 1.083 across 238 retained pairs, whereas its ratio of compute totals is approximately 2.560. These quantities answer different questions. Omitting displayed-zero pairs also changes the median’s population, so it is not substituted for an unrounded full-cohort statistic. The overview’s 1.42 s median SuperFermion process wall is much larger than its 29.7 ms primary-compute median, indicating that the measured stage does not describe the full cost of every fresh invocation.

At display precision, 217 conditions have a leader separated from all other configurations’ rounding intervals. SuperFermion leads 131, Qulacs 80, Qiskit four and Cirq f32 two; 35 conditions remain unresolved. This rule is deliberately based on the available rounded records. Precision separation does not establish statistical significance or repeated-run stability.

Median SuperFermion worker peak is 150.60 MiB, compared with 149.55 MiB for Qulacs, 184.70 MiB for Qiskit, 268.95 MiB for PennyLane, 228.60 MiB for Cirq f32 and 229.35 MiB for Cirq f64. Median condition-paired SuperFermion RAM savings are 16.56%, 44.15%, 34.63% and 34.68% against the latter four configurations. Against Qulacs the median saving is negative, approximately -0.37%. The small native-framework difference supplies no evidence for a universal lower-memory claim.

3.2 Framework comparisons by task and qubit count

All five framework heatmaps appear in this main results section. Each cell divides the comparator’s compute-time sum by SuperFermion’s sum over the same task and qubit cohort. Values above one favor SuperFermion; values below one favor the comparator. The common logarithmic color range is 0.5 to 512, and every planned value lies inside it. Of the 250 cells, 230 are planned and 20 remain explicitly unplanned because VQC/VQR stop at 12 qubits.




Figure 4: Qulacs/SuperFermion ratios of summed primary compute time for matching task and qubit cohorts. Values above one favor SuperFermion; values below one favor Qulacs. Cells pool matching datasets, maps and output tracks. All five heatmaps share the 0.5 to 512 logarithmic scale. Unplanned VQC/VQR cells retain their status. These are workload-total ratios, not medians of individual speedups.

Qulacs has lower QSVM totals at 4, 8 and 16 qubits, with ratios of approximately 0.851, 0.844 and 0.937. Its 12-qubit QSVM total is about 29 times the SuperFermion total, and a marked 12-qubit increase also appears in several other configured tasks. This pattern identifies an adapter/configuration question for controlled profiling. The records alone do not identify the execution primitive responsible [20].




Figure 5: Qiskit/SuperFermion ratios of summed primary compute time for matching task and qubit cohorts. Values above one favor SuperFermion; values below one favor Qiskit. Cells pool matching datasets, maps and output tracks. All five heatmaps share the 0.5 to 512 logarithmic scale. Unplanned VQC/VQR cells retain their status. These are workload-total ratios, not medians of individual speedups.

Qiskit is competitive in the larger kernel cohorts. At 16 qubits its QSVM, QKRR and QKPCA total ratios are approximately 0.981, 0.988 and 0.926, respectively; its gradient total ratio is about 0.862. The Qiskit heatmap nevertheless shows larger totals for many smaller objective and circuit cohorts. The outcome depends on both task and size, rather than a single framework-wide ordering.




Figure 6: PennyLane/SuperFermion ratios of summed primary compute time for matching task and qubit cohorts. Values above one favor SuperFermion; values below one favor PennyLane. Cells pool matching datasets, maps and output tracks. All five heatmaps share the 0.5 to 512 logarithmic scale. Unplanned VQC/VQR cells retain their status. These are workload-total ratios, not medians of individual speedups.

Every planned PennyLane task/qubit compute-sum cell is above one in this selected grid, with ratios ranging approximately from 1.35 to 48.8. This describes the saved adapter and output contracts. It does not compare every differentiable execution or gradient path that PennyLane provides, and it does not imply a difference in predictive quality.




Figure 7: Cirq f32/SuperFermion ratios of summed primary compute time for matching task and qubit cohorts. Values above one favor SuperFermion; values below one favor Cirq f32. Cells pool matching datasets, maps and output tracks. All five heatmaps share the 0.5 to 512 logarithmic scale. Unplanned VQC/VQR cells retain their status. These are workload-total ratios, not medians of individual speedups.

The Cirq f32 heatmap is explicitly a mixed-precision comparison. SuperFermion has lower totals in most displayed cells, while the 20-qubit VQE cohort is close to parity and slightly favors Cirq f32 at a ratio of approximately 0.995. A close one-run difference is not interpreted as a stable advantage.




Figure 8: Cirq f64/SuperFermion ratios of summed primary compute time for matching task and qubit cohorts. Values above one favor SuperFermion; values below one favor Cirq f64. Cells pool matching datasets, maps and output tracks. All five heatmaps share the 0.5 to 512 logarithmic scale. Unplanned VQC/VQR cells retain their status. These are workload-total ratios, not medians of individual speedsups.

The separate Cirq f64 view preserves its higher-precision configuration. All planned compute-sum ratios are above one in the saved grid, ranging approximately from 1.71 to 109. The full matched-condition distributions remain available in Appendix D, where statevector and native-readout markers also expose differences concealed by task totals.

3.3 Quantum kernels and mathematical structure

Kernel costs grow sharply in several larger cohorts, making state export, retained banks and Gram-matrix work consequential. For the 20-qubit breast-cancer ZZ QSVM condition, SuperFermion takes 186.5 s, compared with 218.3 s for Qulacs and 238.5 s for Qiskit. PennyLane takes 996.0 s, Cirq f32 339.1 s and Cirq f64 911.7 s. These are matched task observations rather than a predictive-accuracy ranking.

The separable angle control changes the computational problem. For breast cancer at 20 qubits it takes 64.2 ms, compared with 212.6 s for SuperFermion and 218.2 s for Qulacs. At four qubits, however, the control’s 5.6 ms observation is slower than SuperFermion’s 4.9 ms. Figure 9 therefore records both the larger-scale benefit of verified structure and the importance of fixed overhead at small sizes.




Figure 9: Breast-cancer angle-map QSVM observations. The analytical control evaluates the separable kernel formula and is kept outside general-simulator rankings. At 20 qubits it takes 64.2 ms, versus 212.6 s for SuperFermion and 218.2 s for Qulacs. Memory includes work outside the kernel stage.

An execution system can use such an identity only after verifying that the requested feature map is separable. General ZZ and variational circuits still require the appropriate state or observable route. The control remains outside simulator rankings because its scope is a specific closed-form kernel.

3.4 Variational learning and objective evaluation

VQC and VQR contain five updates at 4, 8 and 12 qubits. On breast cancer at 12 qubits, SuperFermion VQC records 588.7 ms for native readout and 804.5 ms for statevector output, with worker peaks of 148.3 and 204.8 MiB. On diabetes at the same size, VQR records 442.7 ms with 146.9 MiB for native readout and 586.3 ms with 191.1 MiB for statevector output. The limited update budget supports execution comparisons but does not establish convergence or competitive learning quality.

QAOA feature selection and MaxCut retain five measured qubit sizes and both output tracks. Their objective-evaluation costs provide useful examples of avoiding unnecessary state export [21]. Table 8 also retains VQE and gradient counterexamples: at 20 qubits SuperFermion VQE native readout takes 0.944 s, while Qiskit takes 0.5401 s and Cirq f32 0.2972 s. The corresponding gradient takes 0.8974 s for SuperFermion and 0.8688 s for Cirq f32.

Table 8: Individual 20-qubit native-readout times for selected objectives

Framework	QAOA FS breast cancer (s)	MaxCut ring (s)	VQE TFIM (s)	Gradient synthetic (s)
SuperFermion	0.2604	0.1483	0.944	0.8974
Qulacs	7.73	1.66	1.63	5.77
Qiskit	1.52	0.3959	0.5401	2.36
PennyLane	24.99	5.42	5.73	14.92
Cirq f32	1.19	0.1845	0.2972	0.8688
Cirq f64	59.33	5.84	12.92	29.84

Every numeric entry is one reported primary compute observation. The objective, qubit count, output track, batch, and depth are matched across frameworks within each column. These examples are intentionally shown alongside cases where another lane is faster; they are not repeated-run speedup estimates.

Circuit tests include angle, GHZ, Pauli, phase, QFT, random and ZZ families, together with batch and depth sweeps. The 20-qubit QFT condition at batch one and depth two takes 118.7 ms for SuperFermion, 83.7 ms for Qulacs, 77.8 ms for Qiskit and 82.3 ms for Cirq f32. This counterexample motivates circuit-specific profiling rather than treating prepared execution as a uniform gain. The qubit tables and scaling figures in Appendix A preserve each workload’s relevant output contract.

3.5 Memory scaling and task quality

$$M_{\text{state}}(n) = 16 \, 2^n \text{ bytes}, \quad M_{\text{bank}} \approx B \, M_{\text{state}}(n). \quad (16)$$

For complex128 storage, one dense state at 20 qubits occupies 16 MiB and one at 24 qubits would occupy 256 MiB. The latter is a storage calculation, not a completed 24-qubit learning result. Retaining B states multiplies that allocation before accounting for estimator workspaces, buffers and interpreter state. The reported exported-bank medians are 0.1, 1.6, 25.0, 360.0 and 512.0 MiB over 4 to 20 qubits.




Figure 10: Theoretical dense-state storage, attributed exported-bank medians and reconstructed QSVM worker medians measure different objects. The 512 MiB cap limits the exported bank; it does not cap total RSS or remove exponential dense-state storage.

The 512 MiB exported-bank cap constrains a materialized output bank; it does not constrain total process memory. SuperFermion’s maximum recorded worker peak is 936.4 MiB, while Cirq f64 reaches 2,355.2 MiB on a 16-qubit diabetes QKRR condition. Process allocation can vary with execution strategy, retained objects and workload. Lower process peaks do not imply a reduction in the exponential storage complexity of arbitrary dense states.

Table 9: Task-quality summaries attributed to the supplied overview

Task	Metric	Reported median	Interpretation
QSVM	Accuracy	0.7778	Mixed datasets, maps, qubits, and lanes
Classical classification	Accuracy	0.9625	Different estimator cohort; not a paired accuracy effect
QKRR	R2	0.01056	Limited predictive fit in the reported grid
Classical regression	R2	0.4571	Context control; hyperparameter parity not established
VQC	Accuracy	0.4167	Five-update budget; no convergence claim
VQR	R2	0.004448	Five-update budget; no convergence claim
QKPCA	Variance fraction	0.1184	Representation metric, not classification accuracy
QAOA feature selection	Downstream accuracy	0.9327	Objective quality is separate from simulator speed
MaxCut	Approximation ratio	0.7285	No general optimality guarantee
VQE	Energy gap	12.14	Reference and units require original task metadata
Gradient	Maximum error	4.996e-10	Median of reported per-job maximum gradient errors

These are attributed overview summaries of its original heterogeneous task cohort; row-level learning outcomes are unavailable for recalculation over the six included configurations. They do not rank predictive quality. The largest reported gradient error is 1.448e-7.

Task quality is separate from execution cost. The overview’s heterogeneous QSVM and classical-classification summaries are not paired learning-accuracy estimates. Its low VQC accuracy and near-zero VQR R2 under five updates do not support a convergence claim. Row-level predictive outputs are unavailable for recalculation over the six included configurations, so the attributed summaries are contextual evidence rather than a revised quality leaderboard.

3.6 Matrix product state coverage and scaling

Table 10: Replacement MPS study outcomes by included configuration

Framework	OK	Budget	Error	Not offered	Total
SuperFermion	40	1	0	0	41
Qiskit Aer	40	1	0	0	41
PennyLane	31	5	0	0	36
Cirq quimb	30	5	3	0	38
Qulacs	0	0	0	26	26

Outcomes are deduplicated by recorded job key. A budget outcome has no assigned runtime. Not offered describes the recorded adapter path, not general framework impossibility.




Figure 11: Outcomes after retry deduplication and selection of the included MPS configurations: 182 jobs, comprising 141 OK, 12 budget, three error and 26 not-offered results. Coverage differs across adapters. Counts describe the planned interface, not general framework success or speed.

The selected MPS grid contains 141 successful observations, 12 budget outcomes, three errors and 26 not-offered entries. SuperFermion and Aer each complete 40 of 41 recorded jobs. PennyLane completes 31 of 36, while Cirq quimb completes 30 of 38 with three QFT construction errors. Qulacs has 26 not-offered rows. These differing interfaces and grids prevent success counts from functioning as a speed ranking.

At maximum bond 64, SuperFermion completes the 1,000-qubit GHZ case in 0.012040 s and the four-layer brick case in 0.020053 s. Aer records 0.031483 s and 0.288328 s for those same cases. The structured circuit and requested correlation matter to interpretation: these results extend the measured low-entanglement workloads, rather than demonstrating efficient general dense-state simulation at that size.

For the twelve-layer brick stress case, SuperFermion records 0.455114, 1.140850 and 2.359950 s at 200, 500 and 1,000 qubits. Aer records 0.813775, 1.705730 and 2.879600 s. At 1,000 qubits the recorded correlation differs by approximately $1.1\text{e-}12$. Memory supplies a counterexample: SuperFermion’s recorded process peak is 473.8 MiB, versus 296.7 MiB for Aer. This prevents interpreting the faster recorded replay as a simultaneous lower-memory result.

3.7 Truncation diagnostics and QFT boundaries

At 16 qubits and maximum bond two, SuperFermion TVD is approximately 0.09054 while Aer reaches 0.04076. At bond four and above their reported TVDs are approximately $8.0\text{e-}16$ and $4.9\text{e-}11$. SuperFermion therefore has a larger error at the tightest recorded bond. At 200 qubits, its reported discarded-weight sequence is non-monotonic: 0.2418, 0.9330 and 0.5001 at bonds 8, 16 and 32, followed by approximately $6.65\text{e-}14$ at bond 64. These diagnostics warrant numerical-policy investigation and are not global fidelity certificates.

At 200 qubits, the saved QFT path takes 0.586371 s for SuperFermion and 2.081645 s for Aer. PennyLane completes its 20-qubit QFT but exceeds the recorded 50-qubit budget. Cirq quimb succeeds at eight qubits and has three construction-error jobs at 20, 30 and 50 qubits. The errors

concern the pinned construction and dependencies; they do not establish a general size limit for Cirq. Matching one correlation also does not validate every QFT output amplitude.

The eight successful dense-reference rows without recorded time remain N/A in time tables. Cirq’s largest dense-versus-MPS correlation difference is approximately $6.0\text{e-}8$, while the SuperFermion, Aer and PennyLane differences are below $4.1\text{e-}12$. The precision and extraction paths are not uniform. Sixteen-qubit TVD rows include dense-reference/extraction work for SuperFermion and Aer, while the other paths have different validation contracts, so their times are not ranked as equivalent outputs.

3.8 Measurement feed forward and execution limits

Table 11: Completed dynamic circuit observations with matched case specifications

Case	Framework	Shots	Wall (s)	RAM (MiB)	Correct
TP 12	SuperFermion	500	0.067684	45.5	Yes
TP 12	Qiskit Aer	500	0.789876	76.7	Yes
TP 12	PennyLane	500	6.304416	334.2	Yes
TP 12	Cirq quimb	500	3.088467	202.5	Yes
Rep 12	SuperFermion	200 x 4	0.076387	47.5	Yes
Rep 12	Qiskit Aer	200 x 4	0.894468	77.5	Yes
Rep 12	PennyLane	200 x 4	6.537493	334.5	Yes
Rep 12	Cirq quimb	200 x 4	5.352905	202.7	Yes
Wire 17	SuperFermion	200	2.014718	52.8	Yes
Wire 17	Qiskit Aer	200	12.330271	84.6	Yes
Wire 17	PennyLane	200	69.025288	342.5	Yes
Wire 17	Cirq quimb	200	4.104582	208.6	Yes
MPS TP 120	SuperFermion	1000	0.317963	46.0	Yes
MPS TP 120	Qiskit Aer	1000	1.015076	79.6	Yes

TP is ordinary teleportation; Wire is the measurement-based circuit; MPS TP is native MPS teleportation. Rep tests four error patterns at 200 shots each. Compare within a case. Correct means all tested target-bit outcomes satisfy the specified invariant, not validation for arbitrary inputs.




Figure 12: Dynamic replay and process memory from the replacement archive. TP uses 500 shots; Wire 17 uses 200; Rep 12 tests four patterns with 200 shots each. Native MPS teleportation uses 100 shots at 12 qubits and 1,000 at 30, 60 and 120. Counts match within each case. Successful target-bit checks pass for the tested inputs. The common 24-qubit budget boundary has no assigned runtime.

The refreshed 12-qubit teleportation test uses 500 shots and records 0.067684 s for SuperFermion, 0.789876 s for Aer, 3.088467 s for Cirq and 6.304416 s for PennyLane. The four-pattern repetition-code test records 0.076387, 0.894468, 5.352905 and 6.537493 s, respectively, with all tested correction outputs passing. At the common 24-qubit ordinary-teleportation specification, all four offered paths exceed the 90 s budget. Those threshold outcomes are neither assigned zero time nor interpreted as an exclusive limitation of SuperFermion.




Figure 13: Native MPS teleportation with matched shots within each qubit count: 100 shots at 12 qubits and 1,000 at 30, 60 and 120. Curves connect tested conditions and do not isolate qubit scaling from the shot-count change. One timed replay is recorded per condition; both paths pass the target-bit check.

Native MPS teleportation completes at 12, 30, 60 and 120 qubits. At 120 qubits and 1,000 shots, SuperFermion takes 0.317963 s with 46.0 MiB recorded peak RSS, compared with 1.015076 s and 79.6 MiB for Aer. The time ratio is approximately 3.19 in these recorded rows, and both target-bit checks pass. At 12 qubits the matched anchor uses only 100 shots, so the connected scaling plot does not isolate qubit count from shot count.

Measurement-window variation remains material. The supplied narrative reports interleaved stress medians of about 2.31 and 2.37 s at 1,000 qubits, much closer than the grid observations, and a native MPS teleportation ratio near 1.6. It also reports a reversal in the SuperFermion-versus-Cirq measurement-based-wire ordering between windows [18]. The supplementary call-level files referenced by that narrative are not present in the replacement archive. These report-level summaries are kept as contextual checks, not imported into the primary numeric tables or regenerated plots.

3.9 Stabilizer and large Clifford sampling

Table 12: Sampling outcomes and available comparator standings

Regime	Conditions	Faster cold	Faster warm	Complete field
Stabilizer	6	2 / 6	5 / 6	1 / 6
Large Clifford	5	2 / 5	5 / 5	0 / 5
Combined	11	4 / 11	10 / 11	1 / 11

Faster means below every external comparator with a recorded timing in that condition. Complete field requires all four external sampling comparators to have timings. Missing capped or budget jobs do not count as wins or losses. The selected sampling observations comprise 39 OK, 15 capped and one budget outcome.

SuperFermion is faster than every available external sampling comparator in ten of eleven prepared

warm conditions. The exception is 64-qubit brickwork: SuperFermion takes 0.522775 ms while Stim takes 0.23665 ms, a comparator-to-SuperFermion ratio of about 0.453. Cold standings are narrower, with four wins across the eleven recorded conditions. Only the 64-qubit GHZ condition has a complete four-comparator field, so the large-regime observations do not establish a win against every framework at every size.




Figure 14: Recorded comparator-to-SuperFermion time ratios for 2,000 shots. Values above one favor SuperFermion; k means one thousand. Both panels use one logarithmic color scale and include every reviewed condition. Gray cells distinguish capped jobs and the recorded budget outcome. Ratios compare the saved output and preparation contracts, with unequal concurrency settings; they are not significance tests.

At 1,024 and 4,096 qubits, GHZ prepared counting takes 0.8224 and 2.652575 ms for SuperFermion, versus Stim's 10.75125 and 76.52965 ms. The corresponding recorded ratios are 13.07 and 28.85. Large brickwork has ratios 3.07 and 5.40, and 4,096-qubit random Clifford has ratio 6.68. These values concern the saved prepared-output contracts and concurrency configuration, not a universal simulator or public-API speedup.




Figure 15: Cold and prepared warm sampling times for all five included tableau-framework configurations. Discrete conditions use unconnected markers. Missing capped or budget measurements have no plotted time. Cold compilation boundaries and warm output contracts are adapter-specific, as defined in the protocol.

Cold large-Clifford costs retain the preparation counterexamples. At 4,096 qubits, SuperFermion brickwork takes 2,106.4573 ms compared with Stim’s 231.2324 ms. For random Clifford, the cold values are 9,396.8310 and 4,413.2164 ms. SuperFermion’s reported random-Clifford measurement-plan cost is 8,021.0467 ms, much larger than its 14.0754 ms prepared counting cost. This identifies preparation reuse as a practical distinction without assigning causality to a particular kernel.

3.10 Prepared gradient evaluation

All ten saved SuperFermion prepared-gradient timings are below the corresponding Qulacs timings. The maximum reported component-wise gradient discrepancy is $2.8172\text{e-}15$. The independent central finite-difference anchor occurs only at eight qubits and one layer, with maximum error $3.9037\text{e-}11$; it is not a finite-difference validation of the entire size grid.




Figure 16: Prepared gradient timings and separately recorded SuperFermion handle-compilation cost. The one-layer grid spans 6 to 20 qubits; the eight-qubit grid has 16, 32 and 64 parameters. Curves connect tested conditions only. Qulacs includes per-call state preparation and backpropagation; SuperFermion construction is outside gradient timing.

At 20 qubits and one layer, SuperFermion records 923.79665 ms for the gradient, Qulacs records 1,264.0098 ms, and the separate SuperFermion compilation cost is 363.89395 ms. The prepared gradient ratio is about 1.368. The largest ratio occurs at 12 qubits, where the reported Qulacs timing rises to 181.6728 ms while SuperFermion records 0.6909 ms. The resulting ratio is about 262.95; the irregular comparator trend needs targeted profiling before a general scaling interpretation is justified.

Repeated-use cost can be examined with a transparent arithmetic estimate. Let C be the saved SuperFermion compile time, g its prepared gradient time, q the Qulacs per-call time, and K the number of repeated evaluations:

$$C + Kg < Kq \iff K > \frac{C}{q - g}, \quad q > g. \quad (17)$$

Using that estimate, four calls amortize the recorded compile cost at ten qubits, five at sixteen qubits, and two at twenty qubits. The calculation excludes a separate Qulacs setup cost because that quantity was not recorded; it is not a measured cold end-to-end comparison. The archive's `cold_ms` field is not substituted for $C + g$, since the saved timing windows and boundaries differ.




Figure 17: Each left-panel ratio is a saved Qulacs gradient time divided by the SuperFermion prepared-handle time. The right panel is a calculation, not a cold-run measurement: it is the smallest integer K for which $C + K g < K q$, using the recorded SF compile cost C and per-call times g and q . Qulacs setup cost is not added because it was not measured separately.

3.11 Density matrix evaluation and numerical checks

At eight qubits, the native SuperFermion density-matrix engine takes 28.424 ms and Aer takes 49.21935 ms, giving an Aer-to-SuperFermion ratio of about 1.732. At twelve qubits, SuperFermion takes 14,239.7466 ms while Aer takes 9,192.7935 ms; SuperFermion is approximately 1.549 times slower. Both saved diagonal flags pass at those sizes, but the limited checks do not certify complete matrix equality.

The same native worker records statevector anchors of 0.5535, 1.45895 and 2.7277 ms at eight, ten and twelve qubits. Those anchors return a different object and remain contextual observations, not framework speed comparisons. They illustrate why a pure-state application should request a dense density matrix only when that output is required.




Figure 18: Recorded density-matrix times for noiseless depth-eight circuits. The hollow Aer marker and gray ratio marker identify its 10-qubit diagonal-check flag; that point is not used as a validated performance ranking. At 8 and 12 qubits both saved diagonal flags pass. These are dense-core observations, not a result for the held-out density-matrix planner mode.

The ten-qubit row records 815.93745 ms for SuperFermion and 659.0791 ms for Aer. Aer’s diagonal flag is false, even though its trace and purity are near one. The adapter rounds the diagonal and reference probabilities to eight decimal places and then uses exact array equality. The archive does not save Aer’s unrounded diagonal discrepancy or density matrix, so this flag cannot be diagnosed as a particular implementation error or dismissed as harmless rounding. The time remains visible as a recorded observation, with validated relative-performance interpretation withheld.

Table 13: Scope of the recorded numerical checks

Check	Saved observation	Interpretation
Gradient parity	10 / 10 conditions; maximum 2.8172e-15	Component-wise SF/Qulacs agreement in the retained second window
Finite differences	8 qubits, 1 layer; 3.9037e-11	One independent gradient anchor, not all sizes
GHZ sample support	All recorded GHZ checks pass	Samples lie in the two GHZ branches; frequencies and all output states are not certified
SF density diagonal	Flags pass at 8, 10 and 12 qubits	Native worker probability check; thresholds differ from the Aer recipe
Aer density diagonal	Pass at 8 and 12; false at 10	Eight-decimal array comparison; missing unrounded error prevents diagnosis
Density trace and purity	All saved deviations below 4.4e-14	Reported scalar checks; complete coherences and Hermiticity are not independently verified

These values are audit scopes, not repeated-run confidence bounds. Full brickwork/random sampling distributions and complete density matrices are not supplied for independent reconstruction.

4 Discussion

4.1 Contribution to quantum computing software

The practical contribution is lower recorded execution cost for many of the tested quantum research workloads. SuperFermion’s smaller primary-compute total can shorten repeated feature-map comparisons, objective evaluation and circuit exploration under these configurations. Native observable requests can also avoid exporting a full state when the application needs only a scalar result. The qubit-specific tables identify where those routes help and where a different framework is competitive.

The memory contribution is narrower. SuperFermion has lower whole-worker peaks than the configured Qiskit, PennyLane and Cirq paths across the median matched conditions, while Qulacs memory is similar and slightly lower. The MPS stress result additionally favors Aer on recorded peak RSS. A low-memory claim should therefore state the metric, workload and comparator rather than describe a universal framework property.

Bounded-bond MPS and native measurement/feed-forward broaden the tested execution scope. Structured circuits reach 1,000 qubits, and native MPS teleportation reaches 120 qubits with passing target-bit checks. This supports an interface that exposes dense states, compact tensor representations and native readout with explicit diagnostics. It does not extend the learning experiments to those sizes or certify efficient simulation for arbitrary entanglement.

The later observations extend this contribution to preparation reuse. The same tableau path supplies 2,000-shot Clifford samples through 4,096 qubits, while a separately compiled handle evaluates ten dense-state gradient conditions. SuperFermion leads the available external field in 10 of 11 prepared sampling conditions and all ten recorded Qulacs gradient comparisons. Cold construction and the density-matrix counterexample show why these observations require explicit output and preparation contracts. [8]

A useful community release would pair these execution paths with documented output contracts, explicit precision settings, correctness examples and a complete pinned implementation. Publishing the native source patch and exact adapters would allow independent studies to test whether the observations persist across other workloads and repeated executions. Software efficiency is the supported research contribution; improved predictive quality or a general quantum advantage requires separate evidence.

4.2 Engineering priorities

The kernel data suggest controlled bank lifetimes, blocked Gram-matrix construction and reuse of prepared feature maps. The angle-control example also motivates verified analytical dispatch for separable representations. Such a route should check applicability explicitly and fall back to the requested general simulation when the identity does not apply.

For variational and optimization objectives, output-aware APIs should request the smallest result that satisfies the application. The later compiled-handle gradient experiment demonstrates a recorded benefit from separating construction and repeated evaluation, but it does not compare every analytic, adjoint, automatic-differentiation or parameter-shift route. The arithmetic compile-amortization analysis can guide repeated-use profiling, provided its estimate remains distinct from a measured cold end-to-end comparison.

For Clifford sampling, preparation and output conversion need separate profiling. The random-Clifford measurement plan can dominate cold cost even when subsequent packed-shot counting is short. Equivalent count or sample-array outputs and balanced concurrency are needed to assess the effect of the affine plan independently. For density matrices, a follow-up should retain unrounded diagonal differences and complete off-diagonal checks before investigating performance, then extend validation to noisy channels.

For MPS execution, the low-bond truncation sequence calls for numerical-policy tests using full-distribution distances where a dense reference is possible. Validation should then distinguish engine diagnostics from independently checked state or probability errors. Dynamic correctness should expand beyond the tested deterministic inputs to superpositions and additional feed-forward histories. Repeated matched-shot comparisons would help separate stable execution behavior from measurement-window variation.

Language conversion by itself is not evaluated here. Profiling should locate repeated preparation, dispatch, export, allocation and primitive costs before an implementation change is chosen. The saved counterexamples identify concrete targets for such investigation, while the missing component ablation prevents assigning causality to a particular optimization mechanism.

4.3 Limitations and validity

The main statistical limitation is repetition. The statevector grid has one fresh-process measurement per condition, and most MPS trials are warm calls within a process rather than independent fresh repetitions. Workload medians and display-rounding intervals do not quantify repeated-run variation. Close ratios and precision-separated leaders remain descriptive observations. The later mode suite reports means of two window medians, with different ordering in the sampling, gradient and density-matrix sections. It does not save complete per-call timing arrays, so those aggregates cannot supply a reconstructed uncertainty interval.

The MPS grid combines saved measurement windows and compiled SuperFermion revisions. The narrative’s interleaved checks indicate that some large nominal margins narrow and the measurement-based-wire direction can reverse. Their call-level supplements are absent from the replacement archive. Primary tables retain the exact grid values and provenance rather than treating those report summaries as another independently reconstructible dataset.

Output fairness is also limited. Dense-state export, native observables, tensor diagnostics, exact-reference extraction and shot replay materialize different results. Cirq f32 uses mixed precision, and fixed MPS bond does not match approximation error across engines. A comparison of configured adapters consequently does not rank the complete capability of each framework. The mode suite additionally contrasts SuperFermion counts maps with Stim sample arrays, uses adapter-specific cold boundaries and fixes a four-thread Rayon limit for native paths while competitors use their defaults. These engine-level release-example paths are not identified as equivalent shipped Python interfaces.

Measurement scope restricts memory interpretation. Fresh-worker RSS includes imports and setup. MPS peaks can include sampling or earlier work in a shared chunk. Those metrics are useful operational observations but cannot identify incremental tensor allocation. Reported compute totals also exclude some costs present in process or notebook walls. The later execution-mode suite records no RAM field; its representation storage orders do not extend the process-RSS comparison.

Implementation reproducibility remains incomplete. The final statevector overview lists correctness

artifacts and an implementation patch without providing their complete contents. The replacement MPS notebook exposes adapter code, records and probes, but the native source patch and supplementary interleaved calls are not supplied. Independent reruns with a full environment and source manifest remain necessary before a general performance conclusion. The later adapter and executed notebook are available, while the complete native MPS, tableau, adjoint and density-matrix implementations cited by the methods report are not included. Source-comment MPS timings remain attributed context, with no component ablation or mapping to all earlier compiled revisions.

Correctness coverage also differs across regimes. The later gradient parity is computed from the retained second-window vectors, and independent finite differences cover only eight qubits and one layer. GHZ support checks do not validate full brickwork or random-Clifford distributions. Trace, purity and diagonal checks cannot identify complete density matrices; the ten-qubit Aer diagonal flag therefore remains explicit and its performance comparison is unranked. [8]

The tested scope is finite. No completed 24-qubit learning results or VQC/VQR training above 12 qubits are supplied. The 24-qubit dynamic test is a shared budget boundary. Five training updates and aggregate task-quality summaries do not establish convergence, clinical efficacy or a predictive advantage. Those questions require their own experiments and row-level outcomes.

5 Conclusion

SuperFermion records the lowest reconstructed primary-compute total among six included statevector configurations on 252 matched conditions. The analysis preserves separate output contracts, all five main-paper framework heatmaps, and detailed comparisons at each measured qubit count. SuperFermion is the precision-separated leader in 131 conditions; 35 conditions remain unresolved at display precision. Qulacs, Qiskit and Cirq retain specific advantages, and Qulacs process memory is similar to and slightly lower than SuperFermion’s.

The replacement MPS results extend the evaluation to bounded-bond structured circuits and dynamic measurement. They include 141 successful observations among 182 selected jobs, 1,000-qubit ladder and stress cases, and correct native MPS teleportation through 120 qubits. Low-bond accuracy counterexamples, a higher SuperFermion stress peak than Aer, symmetric ordinary-teleportation budget limits and measurement-window variation remain explicit.

The later 24-condition suite adds mathematical methods and specialized execution evidence. The shared stabilizer/Clifford engine is faster than every available external comparator in 10 of 11 prepared sampling conditions, with only four corresponding cold wins. All ten prepared adjoint gradients are below Qulacs, while the density-matrix path is faster at eight qubits and slower at twelve. The flagged ten-qubit comparison remains unranked. These results extend the observed scope without replacing or pooling the earlier statevector and MPS measurements.

Across the three suites, the supported conclusion is practical execution efficiency for stated workloads and output contracts. Preparation reuse and representation selection are promising engineering mechanisms; their individual effects, public-interface parity and broad performance stability require complete native provenance, balanced repeated measurements and stronger distribution or matrix validation.

6 Data and code availability

The accompanying source archive contains the supplied statevector reports, the MPS ZIP and its executed notebook, JSONL/CSV records, capability probe, narrative and original plots. It also contains the unchanged later methods report and execution-mode ZIP, together with its executed notebook, benchmark adapter, raw JSONL, four CSVs and original figure exports. Analysis scripts verify matching statevector conditions, reconcile MPS jobs and mode records with the supplied CSVs, regenerate figures and reconstruct tables. Each of the 182 selected MPS jobs occurs once in its detailed condition tables; all 24 specialized conditions have separate comparison tables. Source hashes, unit conversions, missing-data outcomes, numerical flags and arithmetic gradient-amortization estimates remain inspectable.

One shared manuscript source builds the integrated LaTeX, PDF and Word text, table values, captions and editable Word equations without rerunning a quantum simulator. Raw source material is preserved for traceability and the nested mode results copy is not counted as an additional study. The manuscript retains the stated comparison sets and separates source-comment mechanisms from benchmark observations. Complete native implementation sources, the original statevector correctness artifacts and MPS supplementary interleaved calls remain requirements for independent reproduction.

References

- [1] Maria Schuld and Nathan Killoran. Quantum Machine Learning in Feature Hilbert Spaces. *Physical Review Letters*, 122:040504, 2019. <https://doi.org/10.1103/PhysRevLett.122.040504>.
- [2] Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567:209–212, 2019. <https://doi.org/10.1038/s41586-019-0980-2>.
- [3] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O’Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5:4213, 2014. <https://doi.org/10.1038/ncomms5213>.
- [4] Ville Bergholm et al. PennyLane: Automatic differentiation of hybrid quantum-classical computations, 2018. arXiv:1811.04968 <https://arxiv.org/abs/1811.04968>.
- [5] Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70:052328, 2004. <https://doi.org/10.1103/PhysRevA.70.052328>.
- [6] Ulrich Schollwoeck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326:96–192, 2011. <https://doi.org/10.1016/j.aop.2010.09.012>.
- [7] Tyson Jones and Julien Gacon. Efficient calculation of gradients in classical simulations of variational quantum algorithms, 2020. arXiv preprint 2009.02823 <https://arxiv.org/abs/2009.02823>.
- [8] SuperFermion execution modes benchmark suite. Stabilizer adjoint density matrix and Clifford benchmark results, 2026. User-supplied executed notebook and CSV/JSONL records, 12 October 2026.
- [9] Yasunari Suzuki et al. Qulacs: a fast and versatile quantum circuit simulator for research purpose. *Quantum*, 5:559, 2021. <https://doi.org/10.22331/q-2021-10-06-559>.
- [10] Qiskit Aer developers. Qiskit Aer documentation, 2026. Software documentation, accessed 8 October 2026 <https://qiskit.github.io/qiskit-aer/>.
- [11] Cirq developers. Cirq: Python framework for quantum circuits, 2026. Software documentation, accessed 8 October 2026 <https://quantumai.google/cirq>.
- [12] Google Quantum AI. qsim: quantum circuit simulation, 2026. Software documentation, accessed 8 October 2026 <https://quantumai.google/qsim>.
- [13] Guifre Vidal. Efficient Classical Simulation of Slightly Entangled Quantum Computations. *Physical Review Letters*, 91:147902, 2003. <https://doi.org/10.1103/PhysRevLett.91.147902>.
- [14] SuperFermion methods report. SF optimization methods at the math level MPS engine versus cat Clifford engine, 2026. User-supplied code-oriented mathematical methods report.
- [15] SuperFermion Team. SuperFermion: quantum circuit simulation library, 2026. Source repository, accessed 8 October 2026 <https://github.com/Catstate101/superfermion>.
- [16] SuperFermion benchmark campaign. SuperFermion 24Q VQC/VQR Campaign - Outputs and Benchmarks Report, 2026. Supplied campaign report, 7 October 2026; CAMPAIGN_OUTPUTS_REPORT.md.
- [17] SuperFermion benchmark campaign. Campaign job-by-job report - 4 to 20 qubits, 2026. Supplied job report, 7 October 2026; report4to20q.md.
- [18] SuperFermion benchmark suite. Matrix product state and mid-circuit measurement benchmark results, 2026. User-supplied executed notebook, CSV/JSONL records, capability probe and narrative report; October 2026.
- [19] Craig Gidney. Stim: a fast stabilizer circuit simulator. *Quantum*, 5:497, 2021. <https://doi.org/10.22331/q-2021-07-06-497>.
- [20] Qulacs developers. Qulacs documentation: installation, threading, and advanced simulation, 2026. Software documentation, accessed 8 October 2026 https://docs.qulacs.org/en/latest/intro/1_install.html.
- [21] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A Quantum Approximate Optimization Algorithm, 2014. arXiv:1411.4028 <https://arxiv.org/abs/1411.4028>.

A Quantum experiment comparisons by qubit count

Each table is an individual experiment-and-qubit comparison. Six simulator configurations share the same condition keys; QSVM and QKRR also retain the analytical-angle control. Time is the primary compute stage and RAM is whole-worker peak RSS. Cohort medians do not represent repeated-run uncertainty. Source display strings and complete condition records are retained in the analysis data.

Feature maps remain separate for QSVM/QKRR, and output tracks remain separate for variational and objective tasks. QKPCA and circuit tables show median and maximum values. N/A denotes an absent evaluation; a reported latency rounded to zero remains a censored value. No 24-qubit learning tables are constructed because completed observations are not supplied.

A.1 Quantum support vector machine

The measured qubit grid is 4, 8, 12, 16, 20. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 19: Compute-time and whole-worker memory scaling for QSVM. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred.

Table 14: Quantum support vector machine at 4 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	2.9 ms	147.9	4.4 ms	148.2
Qulacs	2.5 ms	147.3	3.3 ms	147.0
Qiskit	17.5 ms	181.3	20.3 ms	181.8
PennyLane	24.8 ms	265.0	45.6 ms	264.7
Cirq f32	69.1 ms	225.0	140 ms	225.6
Cirq f64	133 ms	225.0	226 ms	225.3
Analytical angle	5.6 ms	147.2	N/A	N/A

Each framework uses the same 5 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 15: Quantum support vector machine at 8 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	10.3 ms	148.2	14.3 ms	148.7
Qulacs	9.1 ms	147.5	11.2 ms	147.6
Qiskit	28.1 ms	183.0	33.2 ms	184.0
PennyLane	52 ms	265.6	99.7 ms	267.1
Cirq f32	123 ms	226.4	283 ms	226.5
Cirq f64	261 ms	228.9	507 ms	230.3
Analytical angle	11 ms	152.7	N/A	N/A

Each framework uses the same 5 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 16: Quantum support vector machine at 12 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	108 ms	169.6	122 ms	170.9
Qulacs	111 ms	169.2	7.17 s	171.1
Qiskit	138 ms	215.3	159 ms	216.8
PennyLane	197 ms	299.3	314 ms	301.7
Cirq f32	282 ms	273.0	539 ms	282.8
Cirq f64	542 ms	320.0	998 ms	335.1
Analytical angle	15.8 ms	157.4	N/A	N/A

Each framework uses the same 5 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 17: Quantum support vector machine at 16 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	1.88 s	517.3	2.22 s	520.2
Qulacs	1.83 s	515.3	2.02 s	516.1
Qiskit	1.91 s	564.9	2.06 s	567.2
PennyLane	2.37 s	650.2	3.21 s	652.3
Cirq f32	2.28 s	707.9	2.88 s	742.1
Cirq f64	3.05 s	1,044.5	4.25 s	1,105.9
Analytical angle	20.6 ms	162.7	N/A	N/A

Each framework uses the same 5 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 18: Quantum support vector machine at 20 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	75.19 s	932.4	75.13 s	920.7
Qulacs	82.47 s	929.5	86.45 s	930.0
Qiskit	90.37 s	943.6	95.27 s	945.6
PennyLane	224.8 s	1,044.5	404.1 s	1,044.5
Cirq f32	160.4 s	1,187.8	139.5 s	1,249.3
Cirq f64	307.5 s	1,884.2	367.9 s	2,017.3
Analytical angle	22.2 ms	337.9	N/A	N/A

Each framework uses the same 5 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.2 Quantum kernel ridge regression

The measured qubit grid is 4, 8, 12, 16, 20. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 20: Compute-time and whole-worker memory scaling for QKRR. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred.

Table 19: Quantum kernel ridge regression at 4 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	3.4 ms	150.8	4.75 ms	151.0
Qulacs	3.35 ms	149.8	3.85 ms	150.1
Qiskit	24.7 ms	183.2	18.1 ms	183.1
PennyLane	31.1 ms	268.1	54.8 ms	268.6
Cirq f32	80.7 ms	228.2	169 ms	229.2
Cirq f64	156 ms	228.8	271 ms	229.3
Analytical angle	8 ms	148.8	N/A	N/A

Each framework uses the same 2 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 20: Quantum kernel ridge regression at 8 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	13 ms	151.2	17.3 ms	151.2
Qulacs	12.2 ms	150.1	13.8 ms	150.2
Qiskit	37 ms	185.2	37.7 ms	185.6
PennyLane	60.2 ms	268.6	116 ms	271.1
Cirq f32	142 ms	228.6	325 ms	229.9
Cirq f64	304 ms	230.9	594 ms	232.8
Analytical angle	15.9 ms	155.9	N/A	N/A

Each framework uses the same 2 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 21: Quantum kernel ridge regression at 12 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	144 ms	174.7	161 ms	175.3
Qulacs	159 ms	173.6	8.555 s	175.5
Qiskit	190 ms	219.4	216 ms	220.1
PennyLane	245 ms	305.6	374 ms	307.9
Cirq f32	351 ms	280.6	636 ms	292.6
Cirq f64	645 ms	341.1	1.16 s	353.4
Analytical angle	22.5 ms	163.2	N/A	N/A

Each framework uses the same 2 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 22: Quantum kernel ridge regression at 16 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	2.505 s	578.7	2.79 s	582.1
Qulacs	2.48 s	577.2	2.67 s	578.0
Qiskit	2.54 s	852.8	2.69 s	853.0
PennyLane	3.07 s	938.6	4.065 s	942.4
Cirq f32	3.055 s	1,126.4	3.805 s	1,239.0
Cirq f64	3.9 s	2,073.6	5.18 s	2,252.8
Analytical angle	26.3 ms	170.1	N/A	N/A

Each framework uses the same 2 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 23: Quantum kernel ridge regression at 20 qubits

Framework	Angle time	Angle RAM (MiB)	ZZ time	ZZ RAM (MiB)
SuperFermion	87.78 s	935.9	110 s	921.7
Qulacs	96.71 s	930.8	131.2 s	931.0
Qiskit	114.3 s	944.7	135.9 s	945.7
PennyLane	251.6 s	1,044.5	708.2 s	1,044.5
Cirq f32	179.9 s	1,187.8	429.6 s	1,249.3
Cirq f64	344.8 s	1,935.4	582.1 s	2,017.3
Analytical angle	29.5 ms	337.9	N/A	N/A

Each framework uses the same 2 datasets. Map columns are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.3 Quantum kernel principal component analysis

The measured qubit grid is 4, 8, 12, 16, 20. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 21: Compute-time and whole-worker memory scaling for QKPCA. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred.

Table 24: Quantum kernel principal component analysis at 4 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	4.3 ms	6.7 ms	150.6	158.1
Qulacs	3.3 ms	5.7 ms	149.6	156.9
Qiskit	17.7 ms	35.7 ms	181.2	192.4
PennyLane	45.6 ms	70.6 ms	266.9	275.2
Cirq f32	140 ms	224 ms	228.3	236.6
Cirq f64	230 ms	364 ms	228.2	237.1

Each framework uses the same 5 datasets. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 25: Quantum kernel principal component analysis at 8 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	14.4 ms	27.5 ms	151.0	158.8
Qulacs	11 ms	22.9 ms	150.1	157.6
Qiskit	31.5 ms	52.9 ms	183.6	193.6
PennyLane	97.1 ms	156 ms	269.1	279.0
Cirq f32	272 ms	442 ms	229.3	240.4
Cirq f64	492 ms	804 ms	229.7	240.0

Each framework uses the same 5 datasets. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 26: Quantum kernel principal component analysis at 12 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	124 ms	289 ms	170.7	186.5
Qulacs	7.34 s	11.74 s	171.0	186.8
Qiskit	157 ms	314 ms	216.1	239.1
PennyLane	305 ms	583 ms	301.9	324.0
Cirq f32	527 ms	920 ms	282.2	318.0
Cirq f64	1 s	1.65 s	334.4	399.9

Each framework uses the same 5 datasets. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 27: Quantum kernel principal component analysis at 16 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	2.14 s	4.33 s	520.1	557.6
Qulacs	1.93 s	3.59 s	516.4	556.9
Qiskit	2.01 s	3.79 s	566.6	819.5
PennyLane	3.28 s	9.72 s	652.3	909.2
Cirq f32	2.83 s	6.96 s	742.6	1,187.8
Cirq f64	4.22 s	12.59 s	1,105.9	2,150.4

Each framework uses the same 5 datasets. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 28: Quantum kernel principal component analysis at 20 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	74.71 s	185.9 s	920.1	922.0
Qulacs	88.95 s	273.3 s	930.8	932.6
Qiskit	95.14 s	247.1 s	944.9	947.1
PennyLane	406.9 s	1117 s	1,044.5	1,044.5
Cirq f32	136.1 s	461.1 s	1,249.3	1,249.3
Cirq f64	363.8 s	949.4 s	2,017.3	2,017.3

Each framework uses the same 5 datasets. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.4 Variational quantum classification

The measured qubit grid is 4, 8, 12. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 22: Compute-time and whole-worker memory scaling for VQC. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred. Training uses five updates at 4, 8 and 12 qubits.

Table 29: Variational quantum classification at 4 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	65.5 ms	145.0	66.8 ms	144.6
Qulacs	59.2 ms	144.0	8.73 s	146.5
Qiskit	280 ms	183.6	242 ms	183.3
PennyLane	898 ms	262.7	973 ms	263.4
Cirq f32	2.66 s	223.8	2.5 s	222.6
Cirq f64	4.91 s	222.9	6.14 s	221.6

Each framework uses the same 5 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 30: Variational quantum classification at 8 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	174 ms	148.1	145 ms	145.6
Qulacs	138 ms	146.9	9.59 s	147.0
Qiskit	383 ms	186.9	310 ms	185.0
PennyLane	2.31 s	269.9	1.89 s	265.1
Cirq f32	6.17 s	229.6	4.62 s	223.6
Cirq f64	12.64 s	231.6	12.56 s	223.5

Each framework uses the same 5 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 31: Variational quantum classification at 12 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	553 ms	182.5	416 ms	147.0
Qulacs	138.5 s	183.0	151.2 s	147.3
Qiskit	904 ms	216.2	670 ms	185.7
PennyLane	5.05 s	306.8	5.15 s	267.0
Cirq f32	10 s	281.9	11.97 s	226.1
Cirq f64	23.83 s	327.6	29.3 s	224.7

Each framework uses the same 5 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.5 Variational quantum regression

The measured qubit grid is 4, 8, 12. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 23: Compute-time and whole-worker memory scaling for VQR. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred. Training uses five updates at 4, 8 and 12 qubits.

Table 32: Variational quantum regression at 4 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	59.8 ms	145.2	61 ms	144.8
Qulacs	53.7 ms	143.8	8.165 s	146.6
Qiskit	248 ms	185.0	228 ms	185.8
PennyLane	843 ms	263.1	921 ms	262.7
Cirq f32	2.43 s	223.7	2.31 s	222.3
Cirq f64	4.535 s	223.2	5.72 s	222.0

Each framework uses the same 2 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 33: Variational quantum regression at 8 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	155 ms	149.2	141 ms	145.9
Qulacs	115 ms	146.9	8.22 s	146.6
Qiskit	335 ms	188.6	268 ms	186.8
PennyLane	1.735 s	271.0	1.775 s	265.1
Cirq f32	4.665 s	230.2	4.31 s	224.1
Cirq f64	9.53 s	233.3	11.1 s	223.4

Each framework uses the same 2 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 34: Variational quantum regression at 12 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	569 ms	188.9	421 ms	146.8
Qulacs	130.1 s	189.0	141 s	147.2
Qiskit	771 ms	223.6	558 ms	187.2
PennyLane	4.55 s	313.0	4.735 s	268.2
Cirq f32	9.185 s	291.2	10.84 s	226.9
Cirq f64	21.1 s	344.0	25.31 s	225.8

Each framework uses the same 2 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.6 QAOA feature selection

The measured qubit grid is 4, 8, 12, 16, 20. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 24: Compute-time and whole-worker memory scaling for QAOA FS. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred.

Table 35: QAOA feature selection at 4 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	0.75 ms	145.0	0.6 ms	144.9
Qulacs	0.55 ms	144.0	33.5 ms	145.6
Qiskit	12.2 ms	171.6	11.7 ms	171.6
PennyLane	24.7 ms	261.2	24.8 ms	261.3
Cirq f32	21.1 ms	220.9	23.5 ms	221.1
Cirq f64	33.5 ms	220.2	42.3 ms	220.6

Each framework uses the same 2 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 36: QAOA feature selection at 8 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	1.65 ms	144.8	1.2 ms	145.0
Qulacs	1.25 ms	143.8	35.7 ms	145.4
Qiskit	27 ms	171.4	27.2 ms	171.8
PennyLane	65.5 ms	261.1	73.6 ms	261.4
Cirq f32	51 ms	221.2	67.2 ms	221.0
Cirq f64	89.9 ms	220.2	119 ms	220.4

Each framework uses the same 2 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 37: QAOA feature selection at 12 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	7.2 ms	145.8	7 ms	145.6
Qulacs	2.155 s	146.0	2.155 s	146.4
Qiskit	87.2 ms	172.6	93.2 ms	172.7
PennyLane	155 ms	261.6	155 ms	262.0
Cirq f32	108 ms	222.2	158 ms	223.2
Cirq f64	193 ms	221.9	267 ms	221.3

Each framework uses the same 2 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 38: QAOA feature selection at 16 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	39.2 ms	155.0	33.8 ms	156.0
Qulacs	41.1 ms	154.9	295 ms	154.9
Qiskit	153 ms	181.2	186 ms	181.8
PennyLane	432 ms	270.9	672 ms	271.7
Cirq f32	204 ms	230.6	257 ms	230.9
Cirq f64	496 ms	229.7	2.035 s	229.7

Each framework uses the same 2 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 39: QAOA feature selection at 20 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	359 ms	254.1	255 ms	262.7
Qulacs	657 ms	281.0	5.775 s	281.1
Qiskit	656 ms	277.8	1.145 s	278.4
PennyLane	6.71 s	397.8	18.15 s	398.1
Cirq f32	1.045 s	327.2	849 ms	327.6
Cirq f64	11.05 s	327.7	37.98 s	328.4

Each framework uses the same 2 datasets. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.7 QAOA MaxCut

The measured qubit grid is 4, 8, 12, 16, 20. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 25: Compute-time and whole-worker memory scaling for MaxCut. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred.

Table 40: QAOA MaxCut at 4 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	0.6 ms	144.1	0.55 ms	144.2
Qulacs	0.55 ms	143.2	33.5 ms	144.9
Qiskit	11.1 ms	170.6	10.5 ms	170.8
PennyLane	21.2 ms	260.5	23.5 ms	260.5
Cirq f32	15.8 ms	220.3	17.4 ms	220.2
Cirq f64	30.9 ms	219.2	30.8 ms	219.6

Each framework uses the same 2 graph families. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 41: QAOA MaxCut at 8 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	0.95 ms	144.2	0.8 ms	144.3
Qulacs	0.9 ms	143.2	34.7 ms	145.1
Qiskit	17.8 ms	171.1	19.4 ms	170.9
PennyLane	40.8 ms	260.8	41.7 ms	260.6
Cirq f32	35 ms	220.7	38.8 ms	220.4
Cirq f64	57.2 ms	219.8	73.6 ms	219.7

Each framework uses the same 2 graph families. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 42: QAOA MaxCut at 12 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	4.75 ms	144.9	4 ms	144.9
Qulacs	843 ms	146.1	913 ms	145.8
Qiskit	42.6 ms	171.8	46 ms	171.4
PennyLane	85.3 ms	261.2	80.5 ms	261.1
Cirq f32	56.9 ms	221.1	82.9 ms	222.2
Cirq f64	108 ms	221.0	144 ms	221.1

Each framework uses the same 2 graph families. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 43: QAOA MaxCut at 16 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	38.9 ms	156.9	31.6 ms	157.9
Qulacs	28.9 ms	155.8	170 ms	155.8
Qiskit	108 ms	181.9	116 ms	181.7
PennyLane	265 ms	271.9	403 ms	272.0
Cirq f32	126 ms	231.2	145 ms	231.1
Cirq f64	378 ms	230.9	974 ms	230.2

Each framework uses the same 2 graph families. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 44: QAOA MaxCut at 20 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	316 ms	256.1	201 ms	264.9
Qulacs	484 ms	284.5	3.38 s	284.8
Qiskit	538 ms	281.1	788 ms	281.4
PennyLane	4.635 s	400.9	11.09 s	401.3
Cirq f32	913 ms	330.1	473 ms	331.0
Cirq f64	8.44 s	331.1	21.91 s	331.4

Each framework uses the same 2 graph families. Output tracks are summarized separately. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.8 Variational quantum eigensolver

The measured qubit grid is 4, 8, 12, 16, 20. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 26: Compute-time and whole-worker memory scaling for VQE. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred.

Table 45: Variational quantum eigensolver at 4 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	1.8 ms	144.0	0.9 ms	144.5
Qulacs	1.8 ms	143.5	34.9 ms	145.3
Qiskit	13.6 ms	170.8	11.5 ms	170.8
PennyLane	27.7 ms	260.6	30.6 ms	260.8
Cirq f32	24 ms	220.7	26 ms	221.2
Cirq f64	33.2 ms	220.4	42.2 ms	220.1

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 46: Variational quantum eigensolver at 8 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	3.7 ms	144.1	1 ms	144.4
Qulacs	4.2 ms	143.7	33.5 ms	145.7
Qiskit	22.7 ms	171.6	19.3 ms	171.3
PennyLane	47.4 ms	260.7	45.5 ms	261.3
Cirq f32	41.6 ms	220.7	44.7 ms	221.0
Cirq f64	68.4 ms	220.2	76.7 ms	220.3

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 47: Variational quantum eigensolver at 12 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	16.4 ms	144.7	3.7 ms	145.1
Qulacs	823 ms	146.0	805 ms	145.8
Qiskit	51.9 ms	171.5	38.9 ms	171.3
PennyLane	87.8 ms	260.5	61.3 ms	261.4
Cirq f32	72.8 ms	221.5	76.1 ms	222.0
Cirq f64	130 ms	220.7	125 ms	220.6

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 48: Variational quantum eigensolver at 16 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	447 ms	154.5	62.2 ms	154.7
Qulacs	435 ms	154.3	83.4 ms	153.7
Qiskit	531 ms	180.6	101 ms	179.9
PennyLane	635 ms	270.8	265 ms	270.5
Cirq f32	526 ms	229.8	103 ms	230.2
Cirq f64	699 ms	229.7	712 ms	230.0

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 49: Variational quantum eigensolver at 20 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	15.47 s	260.6	944 ms	259.8
Qulacs	15.89 s	273.2	1.63 s	272.8
Qiskit	15.9 s	269.1	540 ms	268.8
PennyLane	18.49 s	389.6	5.73 s	389.5
Cirq f32	16.04 s	318.6	297 ms	318.4
Cirq f64	20.92 s	318.7	12.92 s	334.4

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.9 Gradient evaluation

The measured qubit grid is 4, 8, 12, 16, 20. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 27: Compute-time and whole-worker memory scaling for Gradient. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred.

Table 50: Gradient evaluation at 4 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	0.1 ms	144.3	0.3 ms	144.0
Qulacs	0.1 ms	143.2	32.6 ms	145.4
Qiskit	1.6 ms	172.8	1.9 ms	173.0
PennyLane	3.7 ms	260.9	4.2 ms	260.3
Cirq f32	10 ms	220.5	9.7 ms	220.6
Cirq f64	19.1 ms	220.2	24.6 ms	220.0

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 51: Gradient evaluation at 8 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	0.8 ms	144.8	1.1 ms	143.7
Qulacs	0.5 ms	142.7	67.7 ms	145.9
Qiskit	2.8 ms	173.6	2.8 ms	174.1
PennyLane	12.5 ms	260.8	14.9 ms	260.7
Cirq f32	34.6 ms	221.7	34.2 ms	220.5
Cirq f64	75.3 ms	221.1	84.3 ms	220.6

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 52: Gradient evaluation at 12 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	5.3 ms	150.5	4.1 ms	144.6
Qulacs	1.16 s	151.6	1.31 s	145.7
Qiskit	9.3 ms	180.4	6.2 ms	175.8
PennyLane	43.9 ms	270.1	35.6 ms	262.1
Cirq f32	74.6 ms	232.1	90.6 ms	222.8
Cirq f64	167 ms	236.8	212 ms	221.2

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 53: Gradient evaluation at 16 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	126 ms	275.8	98.1 ms	153.6
Qulacs	86.4 ms	274.5	193 ms	152.3
Qiskit	113 ms	304.2	80.8 ms	194.6
PennyLane	366 ms	392.5	434 ms	271.0
Cirq f32	284 ms	416.4	182 ms	229.8
Cirq f64	770 ms	554.1	1.56 s	229.6

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 54: Gradient evaluation at 20 qubits

Framework	Statevector time	Statevector RAM (MiB)	Native readout time	Native readout RAM (MiB)
SuperFermion	1.43 s	684.4	897 ms	245.1
Qulacs	1.75 s	680.3	5.77 s	240.9
Qiskit	2.53 s	694.9	2.36 s	431.0
PennyLane	9.98 s	797.5	14.92 s	356.6
Cirq f32	3.31 s	999.3	869 ms	302.5
Cirq f64	13.87 s	1,577.0	29.84 s	327.3

Each framework uses the same 1 condition per output track. Output tracks are summarized separately. The single-condition columns are individual measurements, not repeated-run medians. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

A.10 Circuit simulation and sweeps

The measured qubit grid is 4, 8, 12, 16, 20. The scaling view retains the output or feature-map distinctions used in the following tables.




Figure 28: Compute-time and whole-worker memory scaling for Circuit. Points summarize matched dataset or circuit cohorts; feature maps or output contracts remain separate. Curves connect measured sizes. One fresh-process repetition is available per condition, so no confidence bands are inferred. Open downward triangles denote cohort medians rounded to zero at the 0.05 ms display bound.

Table 55: Circuit simulation at 4 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	<0.05 ms	0.2 ms	144.0	144.6
Qulacs	<0.05 ms	0.1 ms	143.1	143.5
Qiskit	0.9 ms	1.2 ms	171.5	172.2
PennyLane	1.4 ms	3 ms	260.4	261.0
Cirq f32	2 ms	4.7 ms	220.2	220.7
Cirq f64	2.8 ms	7 ms	219.6	220.3

Each framework uses the same 11 matched circuit and sweep configurations. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 56: Circuit simulation at 8 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	0.1 ms	0.2 ms	144.0	144.8
Qulacs	0.1 ms	0.1 ms	143.3	143.6
Qiskit	1.4 ms	2.2 ms	171.4	172.3
PennyLane	2.5 ms	6.7 ms	260.6	260.8
Cirq f32	3.3 ms	8 ms	220.4	221.0
Cirq f64	7.4 ms	15.1 ms	219.7	220.2

Each framework uses the same 11 matched circuit and sweep configurations. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 57: Circuit simulation at 12 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	0.3 ms	0.9 ms	144.7	145.3
Qulacs	24.7 ms	311 ms	145.2	145.9
Qiskit	1.6 ms	6.4 ms	172.4	173.2
PennyLane	3.6 ms	17.7 ms	261.1	262.0
Cirq f32	5.4 ms	12 ms	221.8	223.1
Cirq f64	13.7 ms	23.7 ms	221.5	224.6

Each framework uses the same 11 matched circuit and sweep configurations. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 58: Circuit simulation at 16 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	9.2 ms	16.3 ms	155.4	158.4
Qulacs	5.2 ms	10.1 ms	152.7	154.8
Qiskit	10.8 ms	15.6 ms	182.1	189.2
PennyLane	22.6 ms	48.3 ms	270.2	278.1
Cirq f32	21.5 ms	33.1 ms	231.3	246.7
Cirq f64	40.6 ms	75.7 ms	229.9	269.8

Each framework uses the same 11 matched circuit and sweep configurations. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

Table 59: Circuit simulation at 20 qubits

Framework	Median time	Maximum time	Median RAM (MiB)	Maximum RAM (MiB)
SuperFermion	92.2 ms	119 ms	263.7	309.9
Qulacs	83.7 ms	138 ms	256.7	304.7
Qiskit	77.8 ms	180 ms	272.0	428.8
PennyLane	518 ms	1.28 s	374.2	533.1
Cirq f32	83.3 ms	326 ms	321.0	606.1
Cirq f64	619 ms	1.03 s	320.9	926.9

Each framework uses the same 11 matched circuit and sweep configurations. Medians describe the listed workload cohort, not repeated-run uncertainty. RAM is whole-worker peak RSS. N/A means this evaluation is not defined for that lane.

B Classical controls by nominal qubit setting

Classical controls use no quantum state allocation. The nominal qubit label indexes the paired experiment configuration. Results summarize the specified estimator cohort rather than an accuracy-controlled predictive comparison.




Figure 29: Classical estimators over five classification or two regression datasets. The nominal qubit axis identifies paired study settings; these estimators allocate no quantum qubits. Time and RAM do not measure predictive accuracy.

Table 60: Classical controls for the nominal 4 qubit setting

Estimator	Datasets	Median time	Median RAM (MiB)	Maximum RAM (MiB)
rbf svc	5	3.6 ms	142.3	142.6
logistic	5	4.1 ms	142.4	143.2
rf classifier	5	60.3 ms	143.0	143.8
ridge	2	1.5 ms	141.9	141.9
rbf krr	2	7.2 ms	146.5	147.0
rf regressor	2	94.8 ms	145.0	145.0

Classical controls execute no quantum circuit. The nominal qubit setting indexes the paired study configuration; it is not a physical qubit allocation. Classification estimators cover five datasets and regression estimators cover two.

Table 61: Classical controls for the nominal 8 qubit setting

Estimator	Datasets	Median time	Median RAM (MiB)	Maximum RAM (MiB)
rbf svc	5	4.1 ms	142.3	142.8
logistic	5	4.2 ms	142.5	143.1
rf classifier	5	60.9 ms	143.0	144.3
ridge	2	1.8 ms	142.2	142.4
rbf krr	2	7.05 ms	146.6	147.1
rf regressor	2	138 ms	144.8	144.8

Classical controls execute no quantum circuit. The nominal qubit setting indexes the paired study configuration; it is not a physical qubit allocation. Classification estimators cover five datasets and regression estimators cover two.

Table 62: Classical controls for the nominal 12 qubit setting

Estimator	Datasets	Median time	Median RAM (MiB)	Maximum RAM (MiB)
rbf svc	5	4.6 ms	142.8	142.9
logistic	5	3.9 ms	142.8	142.9
rf classifier	5	65.3 ms	143.2	143.8
ridge	2	1.7 ms	142.4	142.4
rbf krr	2	7.55 ms	146.4	146.6
rf regressor	2	163 ms	145.1	145.2

Classical controls execute no quantum circuit. The nominal qubit setting indexes the paired study configuration; it is not a physical qubit allocation. Classification estimators cover five datasets and regression estimators cover two.

Table 63: Classical controls for the nominal 16 qubit setting

Estimator	Datasets	Median time	Median RAM (MiB)	Maximum RAM (MiB)
rbf svc	5	4.9 ms	142.5	143.3
logistic	5	3.9 ms	142.8	143.2
rf classifier	5	65.4 ms	143.1	143.8
ridge	2	1.6 ms	141.9	142.0
rbf krr	2	7.35 ms	146.8	147.6
rf regressor	2	161 ms	145.1	145.2

Classical controls execute no quantum circuit. The nominal qubit setting indexes the paired study configuration; it is not a physical qubit allocation. Classification estimators cover five datasets and regression estimators cover two.

Table 64: Classical controls for the nominal 20 qubit setting

Estimator	Datasets	Median time	Median RAM (MiB)	Maximum RAM (MiB)
rbf svc	5	4.8 ms	142.6	143.0
logistic	5	4.5 ms	142.8	143.0
rf classifier	5	83.4 ms	143.4	144.0
ridge	2	1.55 ms	142.1	142.2
rbf krr	2	6.6 ms	146.8	147.0
rf regressor	2	154 ms	145.2	145.3

Classical controls execute no quantum circuit. The nominal qubit setting indexes the paired study configuration; it is not a physical qubit allocation. Classification estimators cover five datasets and regression estimators cover two.

C Matrix product state and dynamic circuit comparisons

Every time value and status in this appendix is reconstructed from the replacement MPS archive. Time is seconds and recorded process memory is MiB. Missing values and unoffered cases remain explicit. Eight successful dense-reference rows lack recorded timing. A single Z0 Zlast value is a correlation check, while TVD and discarded weight have distinct definitions. Each of the 182 selected jobs occurs in exactly one table.

C.1 Small circuit correctness anchors




Figure 30: Dense and MPS correctness-anchor costs. Successful dense-reference rows without recorded timing remain missing from time panels, while their measured RSS is retained. These paths have different validation and extraction contracts and are not ranked as equivalent simulation outputs.

Table 65: GHZ correctness anchor at 12 qubits using dense

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	N/A	44.4	1	OK
Qiskit Aer	0.002834	75.1	1	OK
PennyLane	0.003626	333.4	1	OK
Cirq quimb	N/A	200.8	1	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 66: GHZ correctness anchor at 12 qubits using mps

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.000145	52.8	1	OK
Qiskit Aer	0.001606	84.6	1	OK
PennyLane	0.048847	404.7	1	OK
Cirq quimb	0.050912	276	1	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 67: GHZ correctness anchor at 16 qubits using dense

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	N/A	46.9	1	OK
Qiskit Aer	0.004229	76.2	1	OK
PennyLane	0.028864	335.5	1	OK
Cirq quimb	N/A	201.1	1	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 68: GHZ correctness anchor at 16 qubits using mps

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.000176	48.7	1	OK
Qiskit Aer	0.001612	74.9	1	OK
PennyLane	0.033939	404.4	1	OK
Cirq quimb	0.073207	276.8	1	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 69: Brick4 correctness anchor at 12 qubits using dense

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	N/A	46.7	-0.0330144	OK
Qiskit Aer	0.006051	74.7	-0.0330144	OK
PennyLane	0.022606	333	-0.0330144	OK
Cirq quimb	N/A	200.9	-0.0330144	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 70: Brick4 correctness anchor at 12 qubits using mps

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.000668	47.6	-0.0330144	OK
Qiskit Aer	0.003687	75.4	-0.0330144	OK
PennyLane	0.074584	405.7	-0.0330144	OK
Cirq quimb	0.083455	277	-0.0330144	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 71: Brick4 correctness anchor at 16 qubits using dense

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	N/A	48.6	-0.157829	OK
Qiskit Aer	0.017519	76.3	-0.157829	OK
PennyLane	0.183592	336.8	-0.157829	OK
Cirq quimb	N/A	201.6	-0.15783	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 72: Brick4 correctness anchor at 16 qubits using mps

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.000975	49.1	-0.157829	OK
Qiskit Aer	0.006275	76.2	-0.157829	OK
PennyLane	0.085308	406.6	-0.157829	OK
Cirq quimb	0.109924	276.6	-0.157829	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

C.2 GHZ matrix product state ladder




Figure 31: Recorded adapter time and process peak RSS for matching circuit conditions. Only successful observations are plotted; absent, not-offered, error and budget outcomes remain missing. Sampling time is excluded from the execution metric but can contribute to recorded RSS. Output and validation contracts vary across adapters.

Table 73: GHZ at 50 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.000513	45.3	1	OK
Qiskit Aer	0.003371	78.1	1	OK
PennyLane	0.091111	405.8	1	OK
Cirq quimb	0.331065	288.6	1	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 74: GHZ at 100 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.00088	45.5	1	OK
Qiskit Aer	0.005211	79.8	1	OK
PennyLane	0.139522	412	1	OK
Cirq quimb	0.695074	354.2	1	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 75: GHZ at 200 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.001788	45.7	1	OK
Qiskit Aer	0.005984	83.8	1	OK
PennyLane	0.45368	428.8	1	OK
Cirq quimb	N/A	N/A	N/A	OB
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 76: GHZ at 500 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.005105	47.1	1	OK
Qiskit Aer	0.01704	96.3	1	OK
PennyLane	1.15688	514.3	1	OK
Cirq quimb	2.41519	294	1	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 77: GHZ at 1000 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.01204	48.3	1	OK
Qiskit Aer	0.031483	117.3	1	OK
PennyLane	3.73314	825.1	1	OK
Cirq quimb	5.44048	331.3	1	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

C.3 Four layer brick circuit ladder




Figure 32: Recorded adapter time and process peak RSS for matching circuit conditions. Only successful observations are plotted; absent, not-offered, error and budget outcomes remain missing. Sampling time is excluded from the execution metric but can contribute to recorded RSS. Output and validation contracts vary across adapters.

Table 78: Brick4 at 50 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.001412	49.5	-0.362315	OK
Qiskit Aer	0.01736	79.1	-0.362315	OK
PennyLane	0.219291	407.6	-0.362315	OK
Cirq quimb	0.428856	289.3	-0.362315	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 79: Brick4 at 100 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.001773	50.8	-0.192091	OK
Qiskit Aer	0.053491	81.9	-0.192091	OK
PennyLane	0.478772	413.5	-0.192091	OK
Cirq quimb	0.863467	354.5	-0.192091	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 80: Brick4 at 200 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.003783	52.7	0.840186	OK
Qiskit Aer	0.046151	87.6	0.840186	OK
PennyLane	0.972763	433.4	0.840186	OK
Cirq quimb	N/A	N/A	N/A	OB
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 81: Brick4 at 500 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.008581	58	-0.0857184	OK
Qiskit Aer	0.100235	103.9	-0.0857184	OK
PennyLane	3.4381	525.7	-0.0857184	OK
Cirq quimb	3.23684	288	-0.0857184	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 82: Brick4 at 1000 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.020053	64.8	0.0596901	OK
Qiskit Aer	0.288328	130.9	0.0596901	OK
PennyLane	9.00519	845.2	0.0596901	OK
Cirq quimb	6.86412	304.8	0.0596901	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

C.4 Twelve layer brick circuit stress




Figure 33: Recorded adapter time and process peak RSS for matching circuit conditions. Only successful observations are plotted; absent, not-offered, error and budget outcomes remain missing. Sampling time is excluded from the execution metric but can contribute to recorded RSS. Output and validation contracts vary across adapters.

Table 83: Brick12 stress at 200 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.455114	135.6	0.133363	OK
Qiskit Aer	0.813775	126.3	0.133363	OK
PennyLane	71.8163	500.1	0.133363	OK
Cirq quimb	N/A	N/A	N/A	OB
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 84: Brick12 stress at 500 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	1.14085	260.5	-0.00368592	OK
Qiskit Aer	1.70573	190.3	-0.00368592	OK
PennyLane	N/A	N/A	N/A	OB
Cirq quimb	35.8766	582.1	-0.00368592	OK
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 85: Brick12 stress at 1000 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	2.35995	473.8	-0.235859	OK
Qiskit Aer	2.8796	296.7	-0.235859	OK
PennyLane	N/A	N/A	N/A	OB
Cirq quimb	N/A	N/A	N/A	OB
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

C.5 Quantum Fourier transform conditions




Figure 34: Recorded adapter time and process peak RSS for matching circuit conditions. Only successful observations are plotted; absent, not-offered, error and budget outcomes remain missing. Sampling time is excluded from the execution metric but can contribute to recorded RSS. Output and validation contracts vary across adapters.

Table 86: QFT at 8 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	N/A	N/A	N/A	NT
Qiskit Aer	N/A	N/A	N/A	NT
PennyLane	N/A	N/A	N/A	NT
Cirq quimb	0.041238	276.5	0	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 87: QFT at 20 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	N/A	N/A	N/A	NT
Qiskit Aer	N/A	N/A	N/A	NT
PennyLane	35.457	414.4	-9.02695e-18	OK
Cirq quimb	N/A	N/A	N/A	E
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 88: QFT at 30 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	N/A	N/A	N/A	NT
Qiskit Aer	N/A	N/A	N/A	NT
PennyLane	N/A	N/A	N/A	NT
Cirq quimb	N/A	N/A	N/A	E
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 89: QFT at 50 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.049639	40.8	0	OK
Qiskit Aer	0.124498	90.4	-3.69779e-32	OK
PennyLane	N/A	N/A	N/A	OB
Cirq quimb	N/A	N/A	N/A	E
Qulacs	N/A	N/A	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 90: QFT at 100 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.152423	49.1	0	OK
Qiskit Aer	0.455453	85.7	-1.47911e-31	OK
PennyLane	N/A	N/A	N/A	NT
Cirq quimb	N/A	N/A	N/A	NT
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 91: QFT at 200 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Z0 Zlast	Outcome
SuperFermion	0.586371	62.3	0	OK
Qiskit Aer	2.08164	116.4	-1.72563e-31	OK
PennyLane	N/A	N/A	N/A	NT
Cirq quimb	N/A	N/A	N/A	NT
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

C.6 Bond sweep at 16 qubits




Figure 35: Recorded adapter time and process peak RSS for matching circuit conditions. Only successful observations are plotted; absent, not-offered, error and budget outcomes remain missing. Sampling time is excluded from the execution metric but can contribute to recorded RSS. Output and validation contracts vary across adapters.

Table 92: Brick4 at 16 qubits and maximum bond 2

Framework	Time (s)	RAM (MiB)	TVD	Outcome
SuperFermion	0.013897	58.1	0.0905422	OK
Qiskit Aer	0.103846	91.9	0.0407567	OK
PennyLane	0.085862	406.6	N/A	OK
Cirq quimb	0.037756	329.1	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 93: Brick4 at 16 qubits and maximum bond 4

Framework	Time (s)	RAM (MiB)	TVD	Outcome
SuperFermion	0.014847	55.3	7.99819e-16	OK
Qiskit Aer	0.117088	90.4	4.86381e-11	OK
PennyLane	0.083513	405.6	N/A	OK
Cirq quimb	0.040404	278.6	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 94: Brick4 at 16 qubits and maximum bond 8

Framework	Time (s)	RAM (MiB)	TVD	Outcome
SuperFermion	0.015421	56.3	7.99819e-16	OK
Qiskit Aer	0.118846	90.9	4.86381e-11	OK
PennyLane	0.08711	406	N/A	OK
Cirq quimb	0.057603	278.9	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 95: Brick4 at 16 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	TVD	Outcome
SuperFermion	0.015597	55.7	7.99819e-16	OK
Qiskit Aer	0.120473	90.4	4.86381e-11	OK
PennyLane	0.089649	405.9	N/A	OK
Cirq quimb	0.103294	276.6	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

C.7 Bond sweep at 200 qubits




Figure 36: Recorded adapter time and process peak RSS for matching circuit conditions. Only successful observations are plotted; absent, not-offered, error and budget outcomes remain missing. Sampling time is excluded from the execution metric but can contribute to recorded RSS. Output and validation contracts vary across adapters.

Table 96: Brick12 at 200 qubits and maximum bond 8

Framework	Time (s)	RAM (MiB)	Discarded weight	Outcome
SuperFermion	0.030895	52.2	0.241805	OK
Qiskit Aer	0.167375	83	N/A	OK
PennyLane	3.21919	439.2	N/A	OK
Cirq quimb	3.70493	285.6	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 97: Brick12 at 200 qubits and maximum bond 16

Framework	Time (s)	RAM (MiB)	Discarded weight	Outcome
SuperFermion	0.08271	57.1	0.933032	OK
Qiskit Aer	0.332962	86.5	N/A	OK
PennyLane	3.35149	442.9	N/A	OK
Cirq quimb	4.96813	289.3	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 98: Brick12 at 200 qubits and maximum bond 32

Framework	Time (s)	RAM (MiB)	Discarded weight	Outcome
SuperFermion	0.253603	68.7	0.500085	OK
Qiskit Aer	0.571216	91	N/A	OK
PennyLane	5.71271	454.4	N/A	OK
Cirq quimb	7.91128	299.4	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 99: Brick12 at 200 qubits and maximum bond 64

Framework	Time (s)	RAM (MiB)	Discarded weight	Outcome
SuperFermion	0.445186	104.8	6.65024e-14	OK
Qiskit Aer	0.646639	91.1	N/A	OK
PennyLane	26.5579	500.3	N/A	OK
Cirq quimb	14.8467	376.7	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

Table 100: Brick12 at 200 qubits and maximum bond 128

Framework	Time (s)	RAM (MiB)	Discarded weight	Outcome
SuperFermion	0.441834	104.9	6.65024e-14	OK
Qiskit Aer	0.853833	90.5	N/A	OK
PennyLane	N/A	N/A	N/A	OB
Cirq quimb	17.6958	374.6	N/A	OK
Qulacs	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. A single correlation does not establish full-state equality.

C.8 Truncation diagnostics

TVD at 16 qubits uses an exact probability reference. The 200-qubit discarded-weight sequence is an engine diagnostic rather than an independently measured full-state error. Missing fields are not assigned zero, and the non-monotonic sequence is preserved as recorded.




Figure 37: TVD at 16 qubits compares full output probabilities with a dense reference. At 200 qubits, discarded weight is an engine diagnostic with no dense reference in this sweep. Missing fields are not zero. The non-monotonic low-bond SuperFermion sequence is retained.

C.9 Dynamic measurement conditions

Figures 12 and 13 appear in the main results. The following tables preserve each case, qubit count, shot specification, target-bit check and outcome. Repetition-code shots are per error pattern; its wall covers all four patterns. Budget and not-offered entries receive no invented runtime.

Table 101: Ordinary teleportation at 12 qubits

Framework	Wall (s)	RAM (MiB)	Shots	Correct	Outcome
SuperFermion	0.067684	45.5	500	Yes	OK
Qiskit Aer	0.789876	76.7	500	Yes	OK
PennyLane	6.30442	334.2	500	Yes	OK
Cirq quimb	3.08847	202.5	500	Yes	OK
Qulacs	N/A	N/A	500	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. Target-bit checks cover only the tested inputs.

Table 102: Ordinary teleportation at 24 qubits

Framework	Wall (s)	RAM (MiB)	Shots	Correct	Outcome
SuperFermion	N/A	N/A	500	N/A	OB
Qiskit Aer	N/A	N/A	500	N/A	OB
PennyLane	N/A	N/A	500	N/A	OB
Cirq quimb	N/A	N/A	500	N/A	OB
Qulacs	N/A	N/A	500	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. Target-bit checks cover only the tested inputs.

Table 103: Measurement based wire at 17 qubits

Framework	Wall (s)	RAM (MiB)	Shots	Correct	Outcome
SuperFermion	2.01472	52.8	200	Yes	OK
Qiskit Aer	12.3303	84.6	200	Yes	OK
PennyLane	69.0253	342.5	200	Yes	OK
Cirq quimb	4.10458	208.6	200	Yes	OK
Qulacs	N/A	N/A	200	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. Target-bit checks cover only the tested inputs.

Table 104: Repetition code at 12 qubits

Framework	Wall (s)	RAM (MiB)	Shots	Correct	Outcome
SuperFermion	0.076387	47.5	200/pattern	Yes	OK
Qiskit Aer	0.894468	77.5	200/pattern	Yes	OK
PennyLane	6.53749	334.5	200/pattern	Yes	OK
Cirq quimb	5.3529	202.7	200/pattern	Yes	OK
Qulacs	N/A	N/A	200/pattern	N/A	NO

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. Target-bit checks cover only the tested inputs. Four error patterns are replayed with 200 shots per pattern (800 total per completed row).

Table 105: Native MPS teleportation at 12 qubits

Framework	Wall (s)	RAM (MiB)	Shots	Correct	Outcome
SuperFermion	0.00698	38.4	100	Yes	OK
Qiskit Aer	0.037715	76.7	100	Yes	OK
PennyLane	N/A	N/A	N/A	N/A	NT
Cirq quimb	N/A	N/A	N/A	N/A	NT
Qulacs	N/A	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. Target-bit checks cover only the tested inputs.

Table 106: Native MPS teleportation at 30 qubits

Framework	Wall (s)	RAM (MiB)	Shots	Correct	Outcome
SuperFermion	0.09374	40.5	1000	Yes	OK
Qiskit Aer	0.225771	77.1	1000	Yes	OK
PennyLane	N/A	N/A	N/A	N/A	NT
Cirq quimb	N/A	N/A	N/A	N/A	NT
Qulacs	N/A	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. Target-bit checks cover only the tested inputs.

Table 107: Native MPS teleportation at 60 qubits

Framework	Wall (s)	RAM (MiB)	Shots	Correct	Outcome
SuperFermion	0.236671	38.8	1000	Yes	OK
Qiskit Aer	0.368817	90.4	1000	Yes	OK
PennyLane	N/A	N/A	N/A	N/A	NT
Cirq quimb	N/A	N/A	N/A	N/A	NT
Qulacs	N/A	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. Target-bit checks cover only the tested inputs.

Table 108: Native MPS teleportation at 120 qubits

Framework	Wall (s)	RAM (MiB)	Shots	Correct	Outcome
SuperFermion	0.317963	46	1000	Yes	OK
Qiskit Aer	1.01508	79.6	1000	Yes	OK
PennyLane	N/A	N/A	N/A	N/A	NT
Cirq quimb	N/A	N/A	N/A	N/A	NT
Qulacs	N/A	N/A	N/A	N/A	NT

Time is recorded adapter execution in seconds; RAM is process peak RSS in MiB. OK: completed; OB: budget limit; E: adapter error; NO: not offered; NT: no recorded job. N/A is unrecorded, never zero. Target-bit checks cover only the tested inputs.

D Matched condition time and memory distributions

All task/qubit heatmaps are in the main paper. This appendix gives the complementary condition-level time/RSS scatter view for each external configuration against SuperFermion. Markers distinguish statevector and native readout. The horizontal reference is SuperFermion on both panels, and the equality line indicates which recorded value is lower. Time panels omit pairs censored by a displayed zero; memory panels retain all 252 pairs.

D.1 Qulacs




Figure 38: Matched Qulacs and SuperFermion observations. The logarithmic time panel retains 238 of 252 pairs; 14 displayed-zero pairs are omitted. All 252 memory pairs are shown. Points above the equality lines favor SuperFermion. Markers distinguish output contracts. RSS is whole-worker memory; there is one fresh-process repetition per condition.

D.2 Qiskit




Figure 39: Matched Qiskit and SuperFermion observations. The logarithmic time panel retains 242 of 252 pairs; 10 displayed-zero pairs are omitted. All 252 memory pairs are shown. Points above the equality lines favor SuperFermion. Markers distinguish output contracts. RSS is whole-worker memory; there is one fresh-process repetition per condition.

D.3 PennyLane




Figure 40: Matched PennyLane and SuperFermion observations. The logarithmic time panel retains 242 of 252 pairs; 10 displayed-zero pairs are omitted. All 252 memory pairs are shown. Points above the equality lines favor SuperFermion. Markers distinguish output contracts. RSS is whole-worker memory; there is one fresh-process repetition per condition.

D.4 Cirq f32




Figure 41: Matched Cirq f32 and SuperFermion observations. The logarithmic time panel retains 242 of 252 pairs; 10 displayed-zero pairs are omitted. All 252 memory pairs are shown. Points above the equality lines favor SuperFermion. Markers distinguish output contracts. RSS is whole-worker memory; there is one fresh-process repetition per condition.

Cirq f32 is a mixed-precision comparison. These plots preserve the study’s configured outputs rather than asserting precision-matched framework capability.

D.5 Cirq f64




Figure 42: Matched Cirq f64 and SuperFermion observations. The logarithmic time panel retains 242 of 252 pairs; 10 displayed-zero pairs are omitted. All 252 memory pairs are shown. Points above the equality lines favor SuperFermion. Markers distinguish output contracts. RSS is whole-worker memory; there is one fresh-process repetition per condition.

E Execution mode comparisons by condition

Each table preserves one experiment, qubit count and circuit or depth setting. Framework rows are the selected SuperFermion path and the external comparators actually measured or guarded by that section. Capped and budget results remain distinct from measured times. The density-matrix flag remains explicit, and compilation stays separate from prepared gradient timing. No RAM measurements are present for these conditions.

E.1 Stabilizer sampling

Table 109: GHZ sampling at 64 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	0.225525	0.065425	N/A	OK
Stim	0.42625	0.3435	N/A	OK
Qiskit Aer	224.86	232.911	249.169	OK
PennyLane	16.311	11.986	N/A	OK
Cirq	27694.1	25525.7	N/A	OK

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 110: GHZ sampling at 256 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	1.2739	0.203625	N/A	OK
Stim	2.23025	2.05825	N/A	OK
Qiskit Aer	4661.02	4817.3	4808.93	OK
PennyLane	52.7585	52.024	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 111: Brickwork sampling at 64 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	2.02715	0.522775	N/A	OK
Stim	0.3666	0.23665	N/A	OK
Qiskit Aer	294.473	270.567	318.788	OK
PennyLane	55.5716	54.2614	N/A	OK
Cirq	N/A	N/A	N/A	Budget

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 112: Brickwork sampling at 256 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	9.04725	2.00845	N/A	OK
Stim	2.64795	2.1142	N/A	OK
Qiskit Aer	3625.6	3419.34	3309.93	OK
PennyLane	381.935	237.748	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 113: Random Clifford sampling at 256 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	11.4707	1.51153	N/A	OK
Stim	3.6537	2.1122	N/A	OK
Qiskit Aer	7534.08	7432.47	7261.06	OK
PennyLane	180.892	125.289	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 114: Random Clifford sampling at 1,024 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	210.049	3.98597	N/A	OK
Stim	54.9199	10.5591	N/A	OK
Qiskit Aer	N/A	N/A	N/A	Capped
PennyLane	746.521	661.044	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

E.2 Large Clifford sampling

Table 115: GHZ sampling at 1,024 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	8.67817	0.8224	N/A	OK
Stim	13.8272	10.7512	N/A	OK
Qiskit Aer	N/A	N/A	N/A	Capped
PennyLane	355.95	345.806	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 116: GHZ sampling at 4,096 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	116.224	2.65258	N/A	OK
Stim	149.574	76.5296	N/A	OK
Qiskit Aer	N/A	N/A	N/A	Capped
PennyLane	3524.18	3181.03	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 117: Brickwork sampling at 1,024 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	119.211	4.6761	N/A	OK
Stim	20.3697	14.3595	N/A	OK
Qiskit Aer	N/A	N/A	N/A	Capped
PennyLane	1189.89	1267.28	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 118: Brickwork sampling at 4,096 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	2106.46	17.2773	N/A	OK
Stim	231.232	93.3314	N/A	OK
Qiskit Aer	N/A	N/A	N/A	Capped
PennyLane	7230.9	6917.17	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

Table 119: Random Clifford sampling at 4,096 qubits

Framework	Cold (ms)	Warm (ms)	Extra reuse (ms)	Outcome
SuperFermion	9396.83	14.0754	N/A	OK
Stim	4413.22	94.0489	N/A	OK
Qiskit Aer	N/A	N/A	N/A	Capped
PennyLane	9546.97	9599.09	N/A	OK
Cirq	N/A	N/A	N/A	Capped

2,000 shots per recorded window. Timings are means of two window medians; cold and warm contracts vary by adapter. Extra reuse is the separately recorded Aer simulator-reuse observation. Capped and budget rows have no assigned time. RAM was not measured.

E.3 Adjoint gradients

Table 120: Adjoint gradient at 6 qubits and 1 layer

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	0.0184	0.04255	4.44089e-16
Qulacs	0.0739	N/A	Reference

P = 12 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 121: Adjoint gradient at 8 qubits and 1 layer

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	0.04165	0.0563	4.44089e-16
Qulacs	0.2203	N/A	Reference

P = 16 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 122: Adjoint gradient at 10 qubits and 1 layer

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	0.25205	0.20635	3.88578e-16
Qulacs	0.31955	N/A	Reference

P = 20 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 123: Adjoint gradient at 12 qubits and 1 layer

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	0.6909	0.69245	1.50535e-16
Qulacs	181.673	N/A	Reference

P = 24 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 124: Adjoint gradient at 14 qubits and 1 layer

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	5.30115	3.6993	2.498e-16
Qulacs	102.818	N/A	Reference

P = 28 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 125: Adjoint gradient at 16 qubits and 1 layer

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	19.8242	15.8798	2.77556e-16
Qulacs	23.1683	N/A	Reference

P = 32 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 126: Adjoint gradient at 18 qubits and 1 layer

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	191.298	63.9346	1.29063e-15
Qulacs	244.965	N/A	Reference

P = 36 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 127: Adjoint gradient at 20 qubits and 1 layer

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	923.797	363.894	2.81719e-15
Qulacs	1264.01	N/A	Reference

P = 40 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 128: Adjoint gradient at 8 qubits and 2 layers

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	0.0855	0.164	1.66533e-16
Qulacs	0.2176	N/A	Reference

P = 32 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

Table 129: Adjoint gradient at 8 qubits and 4 layers

Framework	Gradient (ms)	Separate compile (ms)	Maximum gradient difference
SuperFermion	0.1485	0.08495	1.39916e-16
Qulacs	0.36975	N/A	Reference

P = 64 parameters; observable is the full Z product. SuperFermion uses a compiled handle; construction is outside gradient timing. Qulacs timing includes parameter setters, state preparation, execution and backpropagation. The parity field is the saved component-wise maximum from the retained second comparison window. RAM was not measured.

E.4 Density matrix evaluation

Table 130: Density matrix evaluation at 8 qubits

Framework	Time (ms)	Trace deviation	Purity deviation	Diagonal check
SuperFermion	28.424	5.77316e-15	1.11022e-16	Pass
Qiskit Aer	49.2193	1.59872e-14	4.28546e-14	Pass

Noiseless depth-eight brickwork. SuperFermion is the native dense-core density-matrix engine, not the held-out density-matrix planner mode. Aer compares diagonals after rounding to eight decimals; the saved SuperFermion flag comes from its native worker. Trace and purity do not establish complete density-matrix equality. The flagged 10-qubit comparison is unranked. RAM was not measured. The separately recorded SuperFermion statevector anchor is 0.5535 ms and has a different output contract.

Table 131: Density matrix evaluation at 10 qubits

Framework	Time (ms)	Trace deviation	Purity deviation	Diagonal check
SuperFermion	815.937	3.55271e-15	1.11022e-16	Pass
Qiskit Aer	659.079	1.9984e-14	1.15463e-14	Flagged

Noiseless depth-eight brickwork. SuperFermion is the native dense-core density-matrix engine, not the held-out density-matrix planner mode. Aer compares diagonals after rounding to eight decimals; the saved SuperFermion flag comes from its native worker. Trace and purity do not establish complete density-matrix equality. The flagged 10-qubit comparison is unranked. RAM was not measured. The separately recorded SuperFermion statevector anchor is 1.45895 ms and has a different output contract.

Table 132: Density matrix evaluation at 12 qubits

Framework	Time (ms)	Trace deviation	Purity deviation	Diagonal check
SuperFermion	14239.7	5.32907e-15	1.11022e-16	Pass
Qiskit Aer	9192.79	2.44249e-14	3.55271e-15	Pass

Noiseless depth-eight brickwork. SuperFermion is the native dense-core density-matrix engine, not the held-out density-matrix planner mode. Aer compares diagonals after rounding to eight decimals; the saved SuperFermion flag comes from its native worker. Trace and purity do not establish complete density-matrix equality. The flagged 10-qubit comparison is unranked. RAM was not measured. The separately recorded SuperFermion statevector anchor is 2.7277 ms and has a different output contract.