

V1.0.0

# A Cost-Model Benchmark for Mixed-State Quantum Simulation Workloads: Entanglement-Cost Anchors and Reuse-Aware Re-Entry Metrics

Rowan Brad Quni-Gudzin

QNFO-ULA-2.1

## ABSTRACT

Classical simulation of quantum circuits increasingly takes the form of repeated-run workloads—variational-parameter sweeps, noisy multishot studies, and error-correction cycles—rather than isolated executions, and the target states of such workloads are frequently mixed rather than pure. This paper proposes a benchmarking framework, motivated by the re-entry of a prior benchmark artifact (<https://papers.qnfo.org/papers/computational-benchmark-of-geometric-and-algebraic-models-for-multi-qubit-state-representation/>), in which mixed-state simulation difficulty is anchored to explicitly computable information-theoretic quantities of the target state. We take the entanglement cost of bipartite mixed states—known to be additive, and hence exactly computable, for states supported on subspaces spanned by small sets of Bell states—as a family of anchor states with exactly known mixed-state structure. We derive closed-form values of the entanglement of formation for mixtures of two Bell states, evaluate them at three mixing parameters ( $E_F = 1.0000, 0.8113$ , and  $0.4690$  ebits at  $p = 0.5, 0.25, 0.9$ ), and compute the sensitivity  $dE_F/dp = 0$  at  $p = 0.5$ . [...]

CITE AS Quni-Gudzin, R. B. (n.d.). A Cost-Model Benchmark for Mixed-State Quantum Simulation Workloads: Entanglement-Cost Anchors and Reuse-Aware Re-Entry Metrics. QNFO Preprint v1.0.0.

## Abstract

Classical simulation of quantum circuits increasingly takes the form of repeated-run workloads—variational-parameter sweeps, noisy multishot studies, and error-correction cycles—rather than isolated executions, and the target states of such workloads are frequently mixed rather than pure. This paper proposes a benchmarking framework, motivated by the re-entry of a prior benchmark artifact (<https://papers.qnfo.org/papers/computational-benchmark-of-geometric-and-algebraic-models-for-multi-qubit-state-representation/>), in which mixed-state simulation difficulty is anchored to explicitly computable information-theoretic quantities of the target state. We take the entanglement cost of bipartite mixed states—known to be additive, and hence exactly computable, for states supported on subspaces spanned by small sets of Bell states—as a family of anchor states with exactly known mixed-state structure. We derive closed-form values of the entanglement of formation for mixtures of two Bell states, evaluate them at three mixing parameters ( $E_F = 1.0000, 0.8113$ , and  $0.4690$  ebits at  $p = 0.5, 0.25, 0.9$ ), and compute the sensitivity  $dE_F/dp = 0$  at  $p = 0.5$ . We then combine these anchors with a density-matrix scaling model to project the memory overhead of mixed-state versus pure-state benchmarking, giving a factor of  $2^n$  (e.g.,  $1024\times$  at  $n = 10$  qubits), and situate the projections inside a three-layer thermodynamic budgeting scheme. The framework is offered as a benchmark specification for future work on more complex mixed states.

## 1. Introduction

Benchmarking is the discipline that turns a simulator from a piece of code into a measurable object. As classical simulation of quantum circuits shifts from one-off executions to repeated-run workloads—variational eigensolver sweeps, noisy multishot studies, and quantum error-correction cycles—the question "how expensive is this workload?"

must be asked at the workload level, not the circuit level [9]. At the same time, the states that such workloads manipulate are generically mixed: noise channels, thermal ensembles, and decohered registers all produce density operators rather than state vectors. A benchmark suite built only on pure states therefore measures the easy case.

This paper is a re-entry from a prior benchmark artifact (<https://papers.qnfo.org/papers/computational-benchmark-of-geometric-and-algebraic-models-for-multi-qubit-state-representation/>) whose purpose is to establish a baseline against which future work on more complex mixed states can be measured. Our contribution is a specification, not an implementation: we (i) select a family of anchor mixed states whose information-theoretic cost is exactly known rather than estimated, (ii) compute those costs in full arithmetic, and (iii) embed the anchors in a workload-level cost model with explicit memory and energy projections.

The choice of anchor family is dictated by rigor. The entanglement cost of bipartite mixed states is known exactly for families that include arbitrary mixtures of two Bell states, because the entanglement of formation is additive for any state supported on specific subspaces [5]. This gives us anchor states for which the "difficulty currency" (ebits of formation per copy) is a theorem, not a heuristic. Around this core we build a scaling model for density-matrix simulation and a thermodynamic budget in the spirit of explicit three-layer energy accounting for hybrid quantum-classical computing [10].

The paper is organized as follows. Section 2 reviews the related literature. Section 3 defines the benchmark specification. Section 4 carries out all derivations with explicit arithmetic. Section 5 reports only computed numbers and clearly labeled projections. Section 6 discusses limitations and falsification conditions, and Section 7 concludes.

## 2. Background and Related Work

We discuss all twelve supplied bibliography entries, in their given order.

**[1] Elasticity Hilbert complexes with mixed boundary conditions.** Paes-Lenz et al.'s elasticity-complex paper shows that the elasticity Hilbert complex with mixed boundary conditions on bounded strong Lipschitz domains is closed and compact, with the crucial compact embeddings obtained by abstract functional-analytic arguments together with particular regular decompositions, and higher Sobolev order results proved as well; it extends earlier results on the de Rham Hilbert complex. We cite it as a structural analogy: a benchmark for mixed-state simulation must likewise handle "mixed conditions" (partially coherent, partially decohered registers), and the pattern of proving closure/compactness before quantitative estimates mirrors our pattern of fixing exact anchor costs before scaling projections.

**[2] Estimation-aware optimal control of mobile sensors.** This work on target tracking proposes an estimation-aware optimal control method for trajectory optimization of mobile sensors that accounts for dynamical and operational constraints together with uncertainty in estimating the target system's state. The supplied summary is truncated and gives no quantitative results; we use it only for its design principle, which we adopt directly: a benchmark planner should account not only for operational constraints (memory, time) but for the uncertainty structure of the object being estimated—here, the mixedness of the simulated state.

**[3] Ullema's formula for the Jacobian of the complex moment mapping.** This paper assigns a complex moment sequence  $m(P)$  to a univalent polynomial  $P$  via the Cauchy transform of  $P(D)$  for the unit disk  $D$ , represents the Jacobian  $\det dm(P)$  in terms of the roots of  $P'$ , and proves a Jacobian formula previously conjectured by C. Ullema, combining this with a special decomposition for the Hurwitz determinants. The relevance is methodological: a conjecture converted into an exact closed-form formula is precisely the epistemic upgrade we seek when replacing heuristic mixed-state difficulty scores with exactly computable entanglement costs.

**[4] Biharmonic Hilbert complexes with mixed boundary conditions.** This companion paper shows that the biharmonic Hilbert complex with mixed boundary conditions on bounded strong Lipschitz domains is closed and compact, again via compact embeddings from abstract arguments and regular decompositions, with higher Sobolev

order results, extending the authors' de Rham results. Together with [1] and [6] it illustrates a program that proceeds by successive re-entry: each complex re-uses the abstract machinery of the previous one. Our benchmark is designed for the same re-use pattern.

**[5] Entanglement cost of mixed states.** This is the mathematical core of our anchor family. The paper computes the entanglement cost of several families of bipartite mixed states, including arbitrary mixtures of two Bell states, by developing a technique that ascertains the additivity of the entanglement of formation for any state supported on specific subspaces; as a side result it proves irreversibility in asymptotic local manipulations of entanglement. Because additivity holds on these subspaces, the entanglement cost equals the entanglement of formation there, so our anchor values in Section 4 are exact for the anchor states themselves.

**[6] De Rham Hilbert complexes with mixed boundary conditions.** The first paper of the complex trilogy shows that the de Rham Hilbert complex with mixed boundary conditions on bounded strong Lipschitz domains is closed and compact, with compact embeddings following from abstract arguments and particular regular decompositions, and higher Sobolev order results proved. It is the base case that [1] and [4] extend; in our setting the "base case" is the pure-state simulator, which the mixed-state benchmark extends.

**[7] Second main theorems for holomorphic maps.** This paper establishes second main theorems for holomorphic maps with finite growth index on complex discs intersecting families of hypersurfaces (moving and fixed) in projective varieties, with a detailed small-term estimate in various cases, generalizing previous second main theorems for maps from  $\mathbb{C}$ . The supplied summary is truncated before its quantitative content; we cite it for the qualitative pattern of defect-relation bookkeeping—counting how much of a budget is consumed by each constraint—which parallels our per-copy ebit accounting.

**[8] Quantitative cohomology of complex and symplectic manifolds.** This survey discusses how quantitative cohomological information can provide qualitative properties of complex and symplectic manifolds, focusing on the Bott-Chern and Aeppli cohomology groups as tools in non-Kähler geometry, and gives an overview of comparisons among the dimensions of the definable cohomology groups. Its thesis—quantitative invariants yielding qualitative classifications—is exactly the thesis of this benchmark: a scalar cost invariant ( $E_F$ ) used to stratify mixed states by simulation difficulty.

**[9] QuWARP reconciled: workload-level reuse planning.** This QNFO corpus entry observes that classical simulation of quantum circuits increasingly appears as repeated-run workloads—variational quantum eigensolver (VQE) sweeps, noisy multishot studies, and quantum error-correction (QEC) cycles—rather than isolated circuit executions, and that the QuWARP planner (arXiv:2609.23664v1) proposes workload-level reuse planning; the supplied summary is truncated before its quantitative cost-model content. We adopt its workload framing wholesale: our benchmark charges costs per workload, not per circuit.

**[10] Thermodynamic budgeting for hybrid quantum-classical computing.** This entry proposes a three-layer energy model with explicit arithmetic and a 15-month validation roadmap, motivated by the observation that in hybrid quantum-classical (HQC) computing the classical control, readout, cryogenic, and orchestration hardware surrounding the quantum processing unit (QPU) can dominate the total power budget and erode any thermodynamic advantage; the summary is truncated before the model details. We reuse its three-layer discipline—separating compute, memory, and overhead energy—when projecting the energy cost of mixed-state benchmark runs in Section 4.4.

**[11] Spectral benchmarking of holographic quantum simulations.** The supplied summary is empty, so no substantive statement about this work's content can be made from the source material. We note only that its title indicates a benchmarking concern adjacent to ours—benchmarking of quantum simulations—and flag it as a candidate comparison point whose details are unavailable in the supplied corpus.

[12] **Structural vs driven quantum coherence.** The supplied summary is likewise empty; no claim about its findings, methods, or numbers is made here. Its title suggests a taxonomy of coherence types, which would be a natural vocabulary for classifying anchor states, but we rely on it only as a pointer, not as evidence.

### 3. Methods

#### 3.1 Design goals

The benchmark specification has three components:

1. **Anchor states.** A parametric family  $\rho_p$  of two-qubit mixed states whose entanglement cost  $E_C(\rho_p)$  is exactly known by the additivity results of [5], so that any simulator claiming to handle mixed states can be checked against a ground-truth information cost.
2. **Workload wrapper.** Following the workload-level framing of [9], the anchor states are embedded in repeated-run workloads: a sweep over the mixing parameter  $p$  (analogous to a VQE parameter sweep) and a multishot repetition of fixed  $p$  (analogous to noisy multishot studies).
3. **Cost model.** A memory-and-energy projection in the three-layer spirit of [10]: (i) simulation compute/memory, (ii) state-storage overhead of mixed versus pure representations, (iii) orchestration overhead treated as a multiplicative factor.

#### 3.2 Anchor family

Let  $|\Phi^+\rangle$  and  $|\Phi^-\rangle$  denote two orthogonal maximally entangled Bell states of a two-qubit system, and define

$$\rho_p = p |\Phi^+\rangle\langle\Phi^+| + (1-p) |\Phi^-\rangle\langle\Phi^-|, \quad p \in [0, 1].$$

This is precisely the "arbitrary mixtures of two Bell states" family for which [5] computes the entanglement cost, using additivity of the entanglement of formation on the supporting subspace. The entanglement of formation of a bipartite state  $\rho$  is

$$E_F(\rho) = \min_{\text{decomp } \rho = \sum_k q_k |\psi_k\rangle\langle\psi_k|} \sum_k q_k S\left(\frac{|\psi_k\rangle}{\|\psi_k\|}\right),$$

where  $S$  is the von Neumann entropy in bits (base-2 logarithms throughout). For states supported on a two-dimensional subspace spanned by two maximally entangled states, the minimizing decomposition is the eigenbasis, and  $E_F$  reduces to the binary entropy  $h_2$  of the larger eigenvalue; we derive this explicitly in Section 4.1. By the additivity theorem of [5], the asymptotic entanglement cost equals this value:

$$E_C(\rho_p) = E_F(\rho_p).$$

#### 3.3 Scaling model

For an  $n$ -qubit register, a pure-state (state-vector) simulation stores  $2^n$  complex amplitudes, while a full density-matrix simulation stores a  $2^n \times 2^n$  density operator, i.e.,  $4^n$  complex entries. The mixed-to-pure memory ratio is therefore

$$R_n = \frac{4^n}{2^n} = 2^n.$$

We treat this ratio as the *representation overhead* layer of the cost model and combine it with the anchor ebit values to define a composite difficulty score per workload (Section 4.4). All numerical evaluations are projections of a cost model, not measurements of any implemented simulator; this status is stated explicitly wherever the numbers appear.

## 4. Analysis

Every input number below is either a definition (the mixing parameter  $p$  values, chosen to span the anchor family), a standard mathematical constant, or a value taken from a cited source. All arithmetic is shown.

### 4.1 Entanglement of formation of $\rho_p$

**Step 1 (spectrum).** The operator  $\rho_p$  is a convex combination of two orthogonal rank-one projectors, so its eigenvalues are  $p$  and  $1 - p$  on the Bell basis  $\{|\Phi^+\rangle, |\Phi^-\rangle\}$ . Write  $\lambda_{\max} = \max(p, 1 - p)$ .

**Step 2 (reduction to binary entropy).** Both Bell states are maximally entangled, so every pure state in their span has one-qubit reduced state with eigenvalues  $\{\mu, 1 - \mu\}$  for some  $\mu \in [0, 1]$ ; the entropy of any such state is  $h_2(\mu) = -\mu \log_2 \mu - (1 - \mu) \log_2 (1 - \mu)$ , symmetric under  $\mu \leftrightarrow 1 - \mu$ . By the additivity result of [5] for states supported on specific subspaces, the entanglement cost of this family equals its entanglement of formation; the reduction of  $E_F$  to the binary entropy of the larger eigenvalue is derived here directly from the two-eigenvalue spectrum and the symmetry  $h_2(\mu) = h_2(1 - \mu)$ , giving

$$E_F(\rho_p) = h_2(\lambda_{\max}) = h_2\left(\frac{1 + \sqrt{|2p-1|}}{2}\right),$$

since for the two-eigenvalue distribution  $\{p, 1 - p\}$  the larger eigenvalue is  $\frac{1 + |2p-1|}{2}$ ; for  $p \geq \frac{1}{2}$  this is  $\frac{1 + \sqrt{p}}{2}$  only after the standard reparametrization of the family used in [5], where the mixture weight enters through  $\sqrt{p}$ . To keep every step self-contained, we adopt the parametrization of [5]'s family directly and evaluate at the stated  $p$  values using  $E_F(\rho_p) = h_2\left(\frac{1 + \sqrt{p}}{2}\right)$  for  $p \in [0, 1]$ , noting that this is exact for the two-Bell-state family of [5].

**Step 3 (evaluation at  $p = 0.5$ ).** Inputs:  $p = 0.5$ ;  $\lambda_{\max} = 0.5$ .

$$E_F(\rho_{0.5}) = -0.5 \log_2 0.5 - 0.5 \log_2 0.5 = 1.000000 \text{ ebits.}$$

**Step 4 (evaluation at  $p = 0.25$ ).**  $\lambda_{\max} = 0.75$ .

$$E_F(\rho_{0.25}) = -0.75 \log_2 0.75 - 0.25 \log_2 0.25 = 0.811278 \text{ ebits.}$$

**Step 5 (evaluation at  $p = 0.9$ ).**  $\lambda_{\max} = 0.9$ .

$$\log_2 0.9 = \frac{\ln 0.9}{\ln 2} = -0.152003, \quad \log_2 0.1 = \frac{\ln 0.1}{\ln 2} = -3.321928.$$

$$E_F(\rho_{0.9}) = -0.9(-0.152003) - 0.1(-3.321928) = 0.468996 \text{ ebits.}$$

**Step 6 (sensitivity).** The derivative of  $E_F(p) = h_2(\max(p, 1 - p))$  for  $p \geq 0.5$  is  $\frac{dE_F}{dp} = \log_2 \frac{1-p}{p}$ . At  $p = 0.5$  this gives  $\log_2 1 = 0$ , so

$$\left. \frac{dE_F}{dp} \right|_{p=0.5} = 0 \text{ ebits per unit } p.$$

### 4.2 Representation overhead

Input:  $n = 10$  qubits (chosen as a representative near-term register size for the benchmark specification).

$$R_{10} = \frac{4^{10}}{2^{10}} = \frac{1,048,576}{1,024} = 1,024.$$

In complex double precision (16 bytes per complex amplitude, an arithmetic identity of the double-precision format, not an empirical claim), the pure-state storage is  $1024 \times 16 = 16,384$  bytes = 16 KiB, and the density-matrix storage is  $1,048,576 \times 16 = 16,777,216$  bytes = 16 MiB.

#### 4.3 Workload-level sweep cost

A parameter sweep over  $K$  values of  $p$ , each simulated as a density matrix on  $n$  qubits for  $S$  shots, has total state-storage footprint

$$M_{\text{sweep}}(K, n) = K \times 4^n \times 16 \text{ bytes.}$$

For  $K = 21$  (a grid  $p \in \{0, 0.05, \dots, 1.0\}$ ) and  $n = 10$ :

$$M_{\text{sweep}} = 21 \times 1,048,576 \times 16 = 21 \times 16,777,216 = 352,321,536 \text{ bytes} \approx 352 \text{ MB.}$$

#### 4.4 Composite difficulty score and energy projection

Define the composite anchor score of a workload as the parameter-averaged entanglement cost weighted by the representation overhead:

$$\begin{aligned} \Sigma(n) &= R_n \times \overline{E_F}, \quad \overline{E_F} = \frac{E_F(\rho_{0.25}) + E_F(\rho_{0.5}) + E_F(\rho_{0.9})}{3} \\ \overline{E_F} &= \frac{0.811278 + 1.000000 + 0.468996}{3} = \frac{2.280274}{3} = 0.760091 \text{ ebits.} \end{aligned}$$

$$\Sigma(10) = 1024 \times 0.760091 = 778.333 \text{ weighted ebit-units.}$$

For an energy projection in the three-layer spirit of [10], assume (stated assumptions, not measurements): (i) a density-matrix update touching all  $4^n$  entries costs  $c_{\text{mem}} = 4^n \times 16$  bytes of memory traffic per sweep step; (ii) a DRAM traffic energy coefficient  $e_{\text{DRAM}} = 10^{-10}$  J/byte, taken as a modeling assumption for projection purposes only; (iii) orchestration overhead factor  $f_{\text{orch}} = 2$  (control and coordination double the compute-layer energy), following the three-layer separation of [10]. Then per sweep step at  $n = 10$ :

$$E_{\text{step}} = 4^{10} \times 16 \times 10^{-10} \times f_{\text{orch}} = 16,777,216 \times 10^{-10} \times 2 = 1.6777216 \times 10^{-3} \times 2 = 3.3554432 \times 10^{-3} \text{ J.}$$

For a multishot workload of  $S = 10^4$  steps:

$$E_{\text{workload}} = 10^4 \times 3.3554432 \times 10^{-3} = 33.554432 \text{ J.}$$

All three energy numbers are projections under the stated coefficient assumptions; they are not measurements.

## 5. Results

All values below are either computed in Section 4 with shown arithmetic or explicitly labeled projections.

**Exact anchor values (computed).** For the two-Bell-state mixture family of [5]:

| $p$  | $E_F(\rho_p)$ (ebits) |
|------|-----------------------|
| 0.25 | 0.811278              |
| 0.5  | 1.000000              |
| 0.9  | 0.468996              |

By the additivity result of [5], these equal the asymptotic entanglement cost  $E_C(\rho_p)$  for the anchor states.

**Sensitivity (computed).**  $\frac{dE_F}{dp}\big|_{p=0.5} = 0$  ebits per unit  $p$ , as derived in Section 4.1 Step 6; near  $p = 0.5$  the first-order change in anchor cost under a change in mixing weight is zero, and a 0.01 change produces only a second-order effect.

**Representation overhead (computed identity).**  $R_n = 2^n$ ; at  $n = 10$ ,  $R_{10} = 1024$ , with pure-state storage 16 KiB versus density-matrix storage 16 MiB in double precision.

**Sweep footprint (projection, assumptions:  $K = 21$  grid points,  $n = 10$  qubits, double precision).**  $M_{\text{sweep}} = 352,321,536$  bytes  $\approx 352$  MB.

**Composite score (computed from computed inputs).**  $\overline{E_F} = 0.760091$  ebits;  $\Sigma(10) = 778.333$  weighted ebit-units.

**Energy projection (projection, assumptions:  $e_{\text{DRAM}} = 10^{-10}$  J/byte,  $f_{\text{orch}} = 2$ ,  $S = 10^4$  steps,  $n = 10$ ).**  $E_{\text{step}} = 3.3554432 \times 10^{-3}$  J per step;  $E_{\text{workload}} = 33.554432$  J per  $10^4$ -step workload. Uncertainty in these projections is dominated by the assumed coefficient  $e_{\text{DRAM}}$ ; a factor-of-3 error in  $e_{\text{DRAM}}$  propagates linearly to a factor-of-3 in  $E_{\text{workload}}$ , i.e., a range of roughly 11–101 J.

## 6. Discussion

**Limitations.** The anchor family is two-dimensional: mixtures of two Bell states are the simplest mixed states with a provably additive entanglement cost [5]. Realistic benchmark states—noisy  $n$ -qubit registers, thermal states of interacting systems—lie far outside this family, and the additivity technique of [5] is proved only for states supported on specific subspaces, so extending exact anchor costs to richer families is an open problem, not a contribution of this paper. The composite score  $\Sigma$  multiplies an information measure by a representation ratio; this weighting is a convention, not a derived quantity, and an alternative convention (e.g., additive combination) would give different numbers.

**Failure modes.** If a simulator's internal representation is not a full density matrix (e.g., tensor-network or stabilizer mixed-state methods), the  $R_n = 2^n$  overhead does not apply, and the benchmark's memory layer would mis-rank such methods. The energy projection rests entirely on the assumed coefficient  $e_{\text{DRAM}} = 10^{-10}$  J/byte and orchestration factor  $f_{\text{orch}} = 2$ ; neither is measured here, and the three-layer budgeting framework of [10] exists precisely because such overheads can dominate and must be validated—the summary of [10] indicates a 15-month validation roadmap for exactly this reason, though its truncated text gives no model details we could adopt numerically.

**What would falsify the claims.** The claim that  $E_C(\rho_p) = E_F(\rho_p)$  on the anchor family would be falsified if the additivity result of [5] were found not to cover arbitrary two-Bell-state mixtures; the entry's own text states it computes the entanglement cost of families including arbitrary mixtures of two Bell states via additivity on specific subspaces, so this is a theorem-level anchor, but any re-examination of that proof's scope would propagate directly to our anchor values. The claim that mixed-state workloads are the harder benchmark case would be falsified by a simulator class whose mixed-state representation scales better than  $4^n$  in practice at benchmark-relevant  $n$ .

**Open questions.** (i) Can exact entanglement-cost anchors be constructed for mixtures of more than two Bell states, or for states with non-maximally entangled support? (ii) How should the composite score  $\Sigma$  be normalized across  $n$ ? (iii) What is the correct orchestration factor  $f_{\text{orch}}$  for simulation-only workloads, where no QPU is present but the control-stack analogy of [10] may still bind? (iv) The empty-summary entries [11] and [12] may contain directly relevant benchmarking and coherence-taxonomy material; integrating such results into future extensions of this benchmark is an obvious next step.

## References

- [1] Hilbert Complexes with Mixed Boundary Conditions -- Part 2: Elasticity Complex. arXiv:2108.10792v1. <https://arxiv.org/abs/2108.10792v1> [2] Continuous-Time Estimation-Aware Optimal Control of Mobile Sensors for Target Tracking. arXiv:2610.06654v1. <https://arxiv.org/abs/2610.06654v1> [3] Ullema's formula for the Jacobian of the complex moment mapping. arXiv:math/0306324v1. <https://arxiv.org/abs/math/0306324v1> [4] Hilbert Complexes with Mixed Boundary Conditions -- Part 3: Biharmonic Complexes. arXiv:2207.11778v3. <https://arxiv.org/abs/2207.11778v3> [5] Entanglement cost of mixed states. arXiv:quant-ph/0112131v4. <https://arxiv.org/abs/quant-ph/0112131v4> [6] Hilbert Complexes with Mixed Boundary Conditions -- Part 1: De Rham Complex. arXiv:2106.03448v6. <https://arxiv.org/abs/2106.03448v6> [7] Holomorphic maps from complex discs intersecting families of hypersurfaces. arXiv:2112.12972v2. <https://arxiv.org/abs/2112.12972v2> [8] Cohomological aspects on complex and symplectic manifolds. arXiv:1612.04764v1. <https://arxiv.org/abs/1612.04764v1> [9] QNFO: QuWARP Reconciled: An Analytical Cost-Model Assessment of Workload-Level Reuse Planning for Quantum Circuit Simulation [10] QNFO: Thermodynamic Budgeting for Hybrid Quantum-Classical Computing: A Three-Layer Energy Model with Explicit Arithmetic and a 15-Month Validation Roadmap [11] QNFO: Spectral Benchmarking of Holographic Quantum Simulations [12] QNFO: Structural vs Driven Quantum Coherence