

The Exponential Arena: A Reconciled Quantitative Account of Hilbert-Space Growth with Qubit Number

The Exponential Arena: A Reconciled Quantitative Account of Hilbert-Space Growth with Qubit Number

Abstract

A recurring question in quantum information is what "space" the state of a quantum computer inhabits, and how that space grows as the number of qubits n increases. This paper treats the question quantitatively, formalizing the arena of an n -qubit register as the complex projective Hilbert space $\mathbb{CP}^{d-1}$ with $d = 2^n$, and deriving the scaling of its real dimension, the classical memory required to store one state vector, the per-gate cost of exact state-vector simulation, the entanglement entropy available across bipartitions, and the sparsity of classically simulable structured subsets. All headline numbers are computed explicitly from stated inputs: a general $n = 50$ register requires $2^{54} \approx 1.80 \times 10^{16}$ bytes (16 PiB) per stored state vector; a single gate touches $2^{50} \approx 1.13 \times 10^{15}$ amplitudes; the central cut of a 100-qubit maximally entangled state carries 50 bits of entropy, and reproducing it with a matrix-product state requires bond dimension 2^{50} and $\approx 2.5 \times 10^{32}$ parameters — worse than the raw state vector. Stabilizer states, by contrast, number 8.79×10^{19} for $n = 10$ and form a measure-zero subset of a 2046-real-dimensional manifold. We situate these scalings against twelve works spanning algebraic geometry, representation theory, number theory, plasma physics, and quantum-information methodology, and we argue that the exponential is a physical resource, with consequences for simulation strategy, benchmarking, and hardware-roadmap claims. Limitations, failure modes, and falsification criteria are discussed.

1. Introduction

When a physicist says that a quantum system of n qubits lives in a "large space," the statement is usually made in passing. Yet the precise sense in which that space grows — exponentially in n , as the Hilbert-space dimension $d = 2^n$ — is the single most consequential quantitative fact separating quantum from classical computation, and it deserves the same explicit bookkeeping that thermodynamics gives to phase-space volumes.

The research question posed by this paper is: *what is "space" as the number of qubits increases?* Following the explicit-bookkeeping methodology of the reconciled assessment in [9], we answer in five registers:

1. **Dimensional:** the manifold of pure states is $\mathbb{CP}^{d-1}$, of real dimension $2d - 2 = 2^{n+1} - 2$.
2. **Descriptive:** the number of complex amplitudes needed to specify a general state is $d = 2^n$, with direct memory consequences.
3. **Computational:** exact simulation of gate dynamics costs $\Theta(d) = \Theta(2^n)$ amplitude touches per gate.
4. **Entropic:** the entanglement accessible across a bipartition scales as $\min(m, n - m)$ bits, which is the operational content of the exponential.
5. **Structural:** the classically simulable fragments (stabilizer states, bounded-entanglement tensor networks) are countable or measure-zero subsets of the arena, and compression succeeds exactly when entanglement is sub-volume-law.

Our contribution is not a new theorem but a reconciled, fully explicit quantitative assessment: every number in Section 4 is derived from stated inputs with shown arithmetic, and every projection is labeled as such. This discipline matters because the literature is full of informal appeals to " 10^{100} -dimensional Hilbert spaces" that are

rarely checked against actual memory, time, and entropy budgets of concrete register sizes.

Throughout, a *qubit* is a two-level quantum system; *Haar-random* means drawn from the uniform measure on the unit sphere of the d -dimensional complex Hilbert space; a *stabilizer state* is a pure state stabilized by a maximal abelian subgroup of the n -qubit Pauli group; an *amplitude touch* is one update of one complex amplitude.

2. Background and Related Work

Because the question "what is the space of n qubits?" sits at a crossroads of geometry, computation, and physics, the relevant literature is heterogeneous. We discuss the twelve works of the bibliography in its exact numbering and indicate how each bears on the argument.

[1] Algebraic theta functions and Eisenstein-Kronecker numbers (arXiv:0709.0640v1). This overview applies Mumford's algebraic theta machinery to the algebraic and p -adic properties of Eisenstein-Kronecker numbers. Its relevance is structural: theta functions are canonical coordinate systems on high-dimensional complex tori, a rigorous model for how "coordinates on a large complex space" can be handled with far fewer parameters than the naive dimension suggests — precisely the tension between $d = 2^n$ amplitudes and compressed parametrizations studied in Section 4.

[2] Geometric Satake correspondence for affine Kac-Moody Lie algebras of type A (arXiv:1812.11710v1). This expository article explains the geometric Satake correspondence, in which representation-theoretic data of an affine group is encoded in the geometry of Grassmannians and quiver varieties. It is the archetype for our claim that the exponential Hilbert space is best attacked through its geometry (symmetries, entanglement structure) rather than through raw coordinate count.

[3] SPACE: the SPectroscopic All-sky Cosmic Explorer (arXiv:0804.4433v1). SPACE is a proposed near-IR spectroscopic mission aiming at redshifts for more than half a billion galaxies to map the Universe over the past 10 billion years. We use it as a data-volume benchmark: in Section 4.5 we compare the memory footprint of a single $n = 50$ state vector against the projected petabyte-scale data volume of such a survey, giving the exponential an operational feel.

[4] A geometric interpretation of Milnor's triple linking numbers (arXiv:math/0110001v3). This work interprets Milnor's triple linking numbers via intersection patterns of Seifert surfaces. Topological invariants of this kind extract low-dimensional structure from configurations whose naive description is high-dimensional; analogously, entanglement invariants extract tractable summaries of 2^n -dimensional states, a point quantified in Sections 4.4 and 4.7.

[5] Simultaneous approximation by conjugate algebraic numbers in fields of transcendence degree one (arXiv:math/0404268v2). This paper approximates several transcendental numbers simultaneously by conjugate algebraic numbers of bounded degree. Its bearing is diophantine: the amplitudes of a generic state are transcendental-like real parameters, and any classical encoding must approximate them to finite precision. The cost of that approximation enters our memory budget in Section 4.2, where we adopt 16 bytes per complex amplitude.

[6] Polynomial-time computing over quadratic maps I: sampling in real algebraic sets (arXiv:cs/0403008v3). This work gives procedures computing sampling points of real algebraic sets defined by quadratic maps in $(dn)^{O(k)}$ arithmetic operations. It shows that even over real closed fields, tractability is governed by the dimension of the *target* set, not the ambient space — exactly the "structured subspace" philosophy we quantify: stabilizer states and bounded-bond-dimension tensor networks are real-algebraic subsets of the full state space.

[7] **Odd, spoof perfect factorizations** (arXiv:2006.10697v1). This paper investigates Diophantine equations related to perfect numbers, generalizing Descartes' 1638 "spoof" factorization $3^2 \cdot 7^2 \cdot 11^2 \cdot 13^2 \cdot 22021$ and Voight's $3^4 \cdot 7^2 \cdot 11^2 \cdot 19^2 \cdot (-127)$. It is a cautionary tale about near-miss structure: factorizations that look perfect but fail one condition. We invoke it as an analogy for pseudo-structure in state space — families that appear to cover the arena but in fact form sparse, measure-zero subsets, made quantitative via the stabilizer count in Section 4.6.

[8] **Whistler instability stimulated by the suprathermal electrons present in space plasmas** (arXiv:1910.01506v1). This plasma-physics paper studies how self-generated whistler instabilities regulate the electron temperature anisotropy $A = T_\perp/T_\parallel$ in collisionless space plasmas. It is the classical many-body counterpoint: a plasma of N particles has a phase space also exponentially large in N , yet kinetic instabilities reduce its effective description to a few moments. This is the physical precedent for the compressed descriptions of quantum registers discussed in Section 6.

[9] **QNFO: Algorithmic Graph-Search Design of Heralded Linear Optical Circuits for Multipartite Entanglement** (<https://papers.qnfo.org/papers/algorithmic-graph-search-design-of-heralded-linear-optical-circuits-for-multipart/>). This reconciled quantitative assessment treats the synthesis of linear-optical circuits generating prescribed multipartite entangled states as a graph-search problem. It is our closest neighbor: its target states (GHZ- and cluster-type photonic states) are exactly the highly structured corners of $\mathbb{CP}^{2^n-1}$ whose existence shows that useful states do not sample the arena uniformly, and its methodology — explicit inputs, shown arithmetic, labeled projections — is the template for this paper.

[10] **QNFO: Bruhat-Tits Tree as a Unifying Geometric Object** (<https://papers.qnfo.org/papers/ballistic-transport-on-the-bruhattits-tree/>). This entry develops the Bruhat-Tits tree as a unifying geometric object across number theory and physics. Trees are the combinatorial skeleton behind tensor-network geometries (MERA, tree tensor networks), whose bond dimensions control how much of the exponential arena a compressed description captures; we use this framing in Section 4.7.

[11] **QNFO: Geometric Unity of Computation** (<https://papers.qnfo.org/papers/geometric-unity-of-computation/>). This entry argues for a geometric unification of computational models. We engage with its spirit by treating classical simulation cost, state-space dimension, and entanglement entropy as geometric quantities of one and the same arena, rather than as incidental complexity measures.

[12] **QNFO: Spectral Benchmarking of Holographic Quantum Simulations** (<https://papers.qnfo.org/papers/spectral-benchmarking-of-holographic-quantum-simulations/>). This entry develops spectral benchmarking for holography-inspired quantum simulations. Benchmarking is where the exponential bites hardest in practice: verifying an n -qubit device against classical prediction is possible only while 2^n remains simulable, and Section 4.3's cost figures make that boundary explicit.

We note honestly that the bibliography contains no standard quantum-information textbook references; the derivations in Section 4 are therefore self-contained and use only definitions stated in Section 3.

3. Methods

3.1 The arena

Definition 3.1 (Register and arena). An n -qubit register has Hilbert space

$$\mathcal{H}_n = (\mathbb{C}^2)^{\otimes n}, \quad d_n = \dim \mathcal{H}_n = 2^n.$$

The pure-state arena is the complex projective space

$$\mathcal{P}_n = \mathbb{CP}^{d_n-1} = \mathbb{CP}^{2^n-1},$$

the quotient of the unit sphere in $\mathcal{H}_n$ by the phase $|\psi\rangle \sim e^{i\phi}|\psi\rangle$.

Definition 3.2 (Real dimension). $\mathbb{CP}^{d-1}$ has real dimension $2d - 2$, hence

$$\dim_{\mathbb{R}} \mathcal{P}_n = 2^{n+1} - 2.$$

Definition 3.3 (Amplitude count and memory). A general pure state $|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle$ has $d_n = 2^n$ complex amplitudes. Storing each amplitude in double precision costs $2 \times 64 = 128$ bits = 16 bytes, so the memory for one state vector is

$$M_n = 16 \cdot 2^n \text{ bytes.}$$

Definition 3.4 (Gate cost). Applying a single- or two-qubit gate to a dense state vector requires updating all d_n amplitudes (a two-qubit gate acts in blocks of 4). The per-gate cost is

$$W_n = 2^n \text{ amplitude touches.}$$

Definition 3.5 (Entanglement entropy across a cut). For a bipartition with $m \leq n - m$ qubits on the smaller side, the maximum von Neumann entropy of either reduced state is

$$S_{\max}(m, n) = m \text{ bits} = m \ln 2 \text{ nats,}$$

achieved by any maximally entangled state across the cut.

Definition 3.6 (Typical entanglement, Page-type). For a Haar-random pure state on $k \otimes (n - k)$ qubits with $k \leq n/2$, the average entanglement entropy in bits is, to leading order,

$$S_{\text{typ}}(n, k) = k - \frac{2^{2k-n-1}}{\ln 2},$$

valid up to exponentially small corrections; this is the standard Page-type leading term, where the correction $d_1/(2d_2 \ln 2)$ arises from the finite ratio of subsystem dimensions $d_1 \leq d_2$.

Definition 3.7 (Stabilizer count). The number of pure n -qubit stabilizer states is

$$N_{\text{stab}}(n) = 2^n \prod_{k=1}^n (2^k + 1),$$

the standard exact count obtained by counting maximal commuting subgroups of the symplectic space $\mathbb{F}_2^{2n}$.

Definition 3.8 (MPS cost). A matrix-product state (MPS) with maximum bond dimension χ on n qubits requires $2n\chi^2$ complex parameters (n tensors of shape $\chi \times \chi \times 2$, up to boundary factors), i.e. $32 n \chi^2$ bytes at double precision. The entanglement across any cut obeys $S \leq \log_2 \chi$.

3.2 Method of assessment

Following the reconciled-assessment methodology of [9], every quantitative claim below is either (i) derived in Section 4 from the definitions above with all arithmetic shown, or (ii) explicitly labeled a *projection* with stated assumptions and uncertainty bounds. No empirical measurements are reported; this is an analytical paper.

4. Analysis

All inputs are stated here with sources: qubit counts n are illustrative register sizes; the constants 64 bits per double-precision float, 8 bits per byte, $\ln 2 = 0.693147$, $\log_{10} 2 = 0.301030$, 1 year = 3.156×10^7 s, 1 day = 8.64×10^4 s are standard; the cosmological atom count $\sim 10^{80}$ is an order-of-magnitude input.

4.1 Dimension of the arena

Input: $d_n = 2^n$ (Definition 3.1). Real dimension (Definition 3.2):

$$\dim_{\mathbb{R}} \mathcal{P}_n = 2^{n+1} - 2.$$

Checks: $n = 1$: $2^2 - 2 = 2$ (the Bloch sphere — correct). $n = 2$: $2^3 - 2 = 6$ (real dimension of $\mathbb{CP}^3$ — correct). $n = 10$: $2^{11} - 2 = 2046$. $n = 50$: $2^{51} - 2 = 2,251,799,813,685,246 \approx 2.25 \times 10^{15}$.

Hilbert-space dimensions: $D_{10} = 1024$; $D_{50} = 2^{50} = 10^{50 \times 0.301030} = 10^{15.0515} \approx 1.1259 \times 10^{15}$; $D_{100} = 2^{100} = 10^{30.1030} \approx 1.2677 \times 10^{30}$; $D_{300} = 2^{300} = 10^{90.3090} \approx 2.037 \times 10^{90}$, exceeding the baryon count of the observable universe ($\sim 10^{80}$) by ten orders of magnitude.

Mixed-state parameter count: a $D_n \times D_n$ Hermitian trace-one matrix has $D_n^2 - 1 = 4^n - 1$ real parameters. $P_{10} = 2^{20} - 1 = 1,048,575$; $P_{100} = 2^{200} - 1 = 10^{60.2060} - 1 \approx 1.6069 \times 10^{60}$. The mixed-state manifold grows as the *square* of the pure-state dimension in the exponent.

4.2 Memory for one state vector

Inputs: 16 bytes per complex amplitude (Definition 3.3); $2^{50} = 1,125,899,906,842,624$.

$$M_{50} = 16 \times 1,125,899,906,842,624 = 2^{54} = 18,014,398,509,481,984 \text{ bytes.}$$

Since 1 PiB = 2^{50} bytes, $M_{50} = 2^{54}/2^{50} = 2^4 = 16$ PiB $\approx 1.80 \times 10^{16}$ bytes. For $n = 60$: $M_{60} = 16 \cdot 2^{60} = 2^{64}$ bytes = 16 EiB $\approx 1.845 \times 10^{19}$ bytes. For $n = 100$: with $2^{100} \approx 1.2677 \times 10^{30}$,

$$M_{100} = 16 \times 1.2677 \times 10^{30} = 2.028 \times 10^{31} \text{ bytes.}$$

Scale comparison: if all $\sim 8 \times 10^9$ humans each stored 10^{15} bytes, the aggregate is 8×10^{24} bytes; M_{100} exceeds this by $2.028 \times 10^{31}/8 \times 10^{24} = 2.54 \times 10^6$.

4.3 Per-gate simulation cost

Input: $W_n = 2^n$ amplitude touches (Definition 3.4). For $n = 50$: $W_{50} = 1.1259 \times 10^{15}$ touches per gate. At a nominal throughput of $R = 10^{10}$ amplitude-updates per second per node (a stated assumption for a modern many-core node performing one complex multiply-add per touch), the wall-clock time per gate is the *projection*

$$t_{50} = \frac{1.1259 \times 10^{15}}{10^{10}} = 1.126 \times 10^5 \text{ s; } \quad \frac{1.126 \times 10^5}{8.64 \times 10^4} = 1.303 \text{ days} \approx 31.3 \text{ hours per gate.}$$

For $n = 40$: $W_{40} = 2^{40} = 1.0486 \times 10^{12}$, so $t_{40} = 1.0486 \times 10^{12}/10^{10} = 104.86 \text{ s} \approx 1.75$ minutes per gate — the practical classical frontier for dense state-vector methods with many gates. The uncertainty of these projections is the uncertainty in R , plausibly a factor of 10 in either direction on commodity hardware.

4.4 Entanglement entropy across the central cut

Input: $S_{\max}(m, n) = m$ bits with $m = \lfloor n/2 \rfloor$ (Definition 3.5). For $n = 100$, $m = 50$:

$$S_{\max} = 50 \text{ bits} = 50 \times 0.693147 = 34.6574 \text{ nats}.$$

The central cut of a 100-qubit maximally entangled state carries 50 bits — indistinguishable, across that cut, from 50 perfectly correlated Bell pairs. The reduced state on the smaller side has $(2^{50})^2 - 1 = 2^{100} - 1 \approx 1.2677 \times 10^{30}$ real parameters, consistent with Section 4.2.

Typical (Page-type) values, as a complementary statistic (Definition 3.6): for $n = 100$, $k = 50$, the correction is $2^{2 \cdot 50 - 100 - 1} / \ln 2 = 2^{-1} / 0.693147 = 0.5 / 0.693147 = 0.7213$, so $S_{\text{typ}}(100, 50) = 50 - 0.7213 = 49.28$ bits, i.e. $49.28/50 = 98.6\%$ of the maximum. For $n = 100$, $k = 10$: correction $= 2^{-81} / \ln 2 = 4.136 \times 10^{-25} / 0.693147 = 5.97 \times 10^{-25}$, so $S_{\text{typ}}(100, 10) = 10 - 5.97 \times 10^{-25} \approx 10$ bits to double precision: a 10-qubit subsystem of a Haar-random 100-qubit state is essentially maximally entangled.

4.5 Data-volume comparison against a real survey benchmark

Inputs: $M_{50} = 1.80 \times 10^{16}$ bytes (Section 4.2); the SPACE mission [3] targets redshifts for more than half a billion galaxies, a projected data volume at the petabyte scale; we take as a stated assumption $V_{\text{survey}} = 10^{15}$ bytes (1 PB). Ratio:

$$\frac{M_{50}}{V_{\text{survey}}} = \frac{1.80 \times 10^{16}}{10^{15}} = 18.0.$$

A single $n = 50$ state vector occupies ≈ 18 petabytes, i.e. roughly 18 times the assumed data volume of an all-sky spectroscopic survey of half a billion galaxies. This is a labeled comparison under the stated assumption for V_{survey} ; if the survey volume is 10^{16} bytes the ratio is 1.8, so the honest statement is that one $n = 50$ state vector is of the same order as, to within a factor of ~ 20 , a flagship survey dataset.

4.6 Stabilizer states: an exactly countable sparse subset

Input: $N_{\text{stab}}(n) = 2^n \prod_{k=1}^n (2^k + 1)$ (Definition 3.7). For $n = 10$ the factors $(2^k + 1)$ for $k = 1, \dots, 10$ are 3, 5, 9, 17, 33, 65, 129, 257, 513, 1025. Cumulative product, step by step:

$$3 \times 5 = 15; \quad 15 \times 9 = 135; \quad 135 \times 17 = 2295; \quad 2295 \times 33 = 75,735;$$

$$75,735 \times 65 = 4,922,775; \quad 4,922,775 \times 129 = 635,037,975;$$

$$635,037,975 \times 257 = 163,204,759,575; \quad 163,204,759,575 \times 513 = 83,724,041,661,975;$$

$$83,724,041,661,975 \times 1025 = 85,817,142,703,524,375.$$

Multiplying by $2^{10} = 1024$:

$$N_{\text{stab}}(10) = 85,817,142,703,524,375 \times 1024 = 87,876,754,128,408,960,000 \approx 8.79 \times 10^{19}.$$

Compare with the size of the arena: $\mathcal{P}_{10} = \mathbb{CP}^{1023}$ has real dimension 2046 (Section 4.1), so the stabilizer states form a countable set of points on a 2046-real-dimensional continuous manifold — a measure-zero subset. The count 8.79×10^{19} is large in absolute terms but occupies zero volume fraction of the arena; this is the quantitative content of the pseudo-structure caution drawn from [7].

4.7 Matrix-product states and the central cut of $n = 100$

Inputs: $n = 100$; central-cut entropy $S = 50$ bits (Section 4.4); MPS bound $S \leq \log_2 \chi$ (Definition 3.8), so reproducing 50 bits requires $\chi \geq 2^{50}$. Take $\chi = 2^{50} = 1.1259 \times 10^{15}$. Parameter count (Definition 3.8):

$$P_{\text{MPS}} = 2n\chi^2 = 2 \times 100 \times (2^{50})^2 = 200 \times 2^{100}.$$

With $2^{100} = 10^{30.1030} \approx 1.2677 \times 10^{30}$ (Section 4.1):

$$P_{\text{MPS}} = 200 \times 1.2677 \times 10^{30} = 2.535 \times 10^{32} \approx 2.5 \times 10^{32} \text{ complex parameters.}$$

At 16 bytes per complex parameter this is $2.535 \times 10^{32} \times 16 = 4.06 \times 10^{33}$ bytes, exceeding the raw dense storage $M_{100} = 2.028 \times 10^{31}$ bytes (Section 4.2) by a factor $4.06 \times 10^{33} / 2.028 \times 10^{31} = 200$. Conclusion: for a state that is maximally entangled across the central cut, the MPS description is strictly worse than the raw state vector — compression succeeds only when the entanglement spectrum across every cut is bounded, i.e. χ grows polynomially rather than as $2^{n/2}$.

5. Results

All numbers below are computed in Section 4 with shown arithmetic; none are empirical measurements.

1. **Arena dimension** (4.1): $\dim_{\mathbb{R}} \mathcal{P}_n = 2^{n+1} - 2$; for $n = 10$, 2046; for $n = 50$, $\approx 2.25 \times 10^{15}$. Mixed states: $4^n - 1$ real parameters; $P_{100} \approx 1.6069 \times 10^{60}$.
2. **Memory** (4.2): $M_{50} = 2^{54} = 1.80 \times 10^{16}$ bytes = 16 PiB; $M_{60} = 2^{64}$ bytes = 16 EiB; $M_{100} = 2.028 \times 10^{31}$ bytes, exceeding an 8×10^{24} -byte global archive by 2.54×10^6 .
3. **Per-gate cost** (4.3, projections at assumed $R = 10^{10}$ updates/s, uncertainty a factor of 10): $t_{50} = 1.303$ days per gate; $t_{40} = 104.86$ s per gate.
4. **Entanglement** (4.4): central cut of $n = 100$: $S_{\max} = 50$ bits = 34.66 nats; Page-type typical value $S_{\text{typ}}(100, 50) = 49.28$ bits (98.6% of maximum); $S_{\text{typ}}(100, 10) \approx 10$ bits.
5. **Survey comparison** (4.5, projection under $V_{\text{survey}} = 10^{15}$ bytes): $M_{50}/V_{\text{survey}} = 18.0$.
6. **Stabilizer sparsity** (4.6): $N_{\text{stab}}(10) = 87,876,754,128,408,960,000 \approx 8.79 \times 10^{19}$, a countable, measure-zero subset of the 2046-dimensional arena.
7. **MPS cost** (4.7): reproducing the $n = 100$ central cut needs $\chi = 2^{50}$ and 2.535×10^{32} parameters $\approx 4.06 \times 10^{33}$ bytes, a factor 200 worse than dense storage.

6. Discussion

Interpretation. The exponential is not an artifact of notation: it appears simultaneously in manifold dimension ($2^{n+1} - 2$), memory ($16 \cdot 2^n$ bytes), time (2^n touches per gate), and entropy ($\min(m, n - m)$ bits). Yet the simulable fragments — stabilizer states (Section 4.6), low-bond-dimension MPS — live on measure-zero or polynomially parametrized subsets, echoing the theme of [6] that tractability is governed by the dimension of the target set, and the kinetic-moment reduction of the exponentially large plasma phase space in [8]. Useful quantum states, as in the structured optical targets of [9], do not sample the arena uniformly.

Limitations. (i) The per-gate timings are projections resting on an assumed throughput $R = 10^{10} \text{ s}^{-1}$; real node performance varies by about a factor of 10 in either direction, and algorithmic improvements (vectorization, gate fusion) change constants, not the 2^n scaling. (ii) The survey comparison in Section 4.5 assumes $V_{\text{survey}} = 10^{15}$ bytes; the ratio is only order-of-magnitude. (iii) The MPS parameter formula $2n\chi^2$ ignores boundary factors and assumes a uniform bond dimension; open-boundary MPS at the first and last cuts need $\chi = 1$ there, an $O(1)$ correction. (iv) The Page-type formula (Definition 3.6) is a leading-order expression valid when d_1/d_2 is not close to 1; at the balanced cut $k = 50$ it carries an $O(1)$ correction that we have computed explicitly, but higher-order terms are neglected. (v) The bibliography contains no standard quantum-information textbook references; all definitions are therefore stated and derived self-contained in Sections 3-4, which limits external cross-checking.

Failure modes and falsification. The central claims would be falsified if: (a) a classical algorithm stored a generic $n = 50$ state vector in asymptotically less than $c \cdot 2^n$ bits with c bounded — this would contradict the information-theoretic lower bound implicit in Definition 3.3; (b) an MPS with polynomially bounded χ reproduced a state with volume-law entanglement across all cuts — this would contradict the bound $S \leq \log_2 \chi$; (c) the stabilizer count formula failed an exact enumeration at small n , e.g. $N_{\text{stab}}(1) = 2 \times 3 = 6$, which matches the six single-qubit stabilizer states ($\pm x, \pm y, \pm z$ eigenstates). The analogy drawn from spoof perfect factorizations [7] is heuristic, not a theorem: we claim sparsity of structured subsets, not that every apparently structured family is sparse.

Open questions. Where exactly is the simulability frontier for realistic (non-Haar) circuits of depth D ? How do the geometric frameworks of [2], [10], [11] sharpen the boundary between compressible and incompressible regions of $\mathcal{P}_n$? And at what n does spectral benchmarking of the kind proposed in [12] become impossible on any classical infrastructure — our Section 4.3 figures put the dense-vector frontier near $n = 40\text{-}50$ on single nodes, leaving multi-node and tensor-network methods as the open margin.

7. Conclusion

We have given a fully explicit quantitative account of the exponential growth of the n -qubit arena: real dimension $2^{n+1} - 2$, memory $16 \cdot 2^n$ bytes, per-gate cost 2^n touches, central-cut entropy $\lfloor n/2 \rfloor$ bits, stabilizer count $2^n \prod_{k=1}^n (2^k + 1)$, and MPS parameter cost $2n\chi^2$ with $\chi \geq 2^S$. Every headline number is derived from stated inputs with shown arithmetic, and every timing is a labeled projection. The exponential is a physical resource: it is what quantum hardware buys and what classical simulation pays, and the boundary between the two is set by entanglement structure, not by raw dimension.

References

- [1] Algebraic theta functions and Eisenstein-Kronecker numbers. arXiv:0709.0640v1. <https://arxiv.org/abs/0709.0640v1> [2] Geometric Satake correspondence for affine Kac-Moody Lie algebras of type A . arXiv:1812.11710v1. <https://arxiv.org/abs/1812.11710v1> [3] SPACE: the SPectroscopic All-sky Cosmic Explorer. arXiv:0804.4433v1. <https://arxiv.org/abs/0804.4433v1> [4] A geometric interpretation of Milnor's triple linking numbers. arXiv:math/0110001v3. <https://arxiv.org/abs/math/0110001v3> [5] Simultaneous approximation by conjugate algebraic numbers in fields of transcendence degree one. arXiv:math/0404268v2. <https://arxiv.org/abs/math/0404268v2> [6] Polynomial-time computing over quadratic maps I: sampling in real algebraic sets. arXiv:cs/0403008v3. <https://arxiv.org/abs/cs/0403008v3> [7] Odd, spoof perfect factorizations. arXiv:2006.10697v1. <https://arxiv.org/abs/2006.10697v1> [8] Whistler instability stimulated by the suprathermal electrons present in space plasmas. arXiv:1910.01506v1. <https://arxiv.org/abs/1910.01506v1> [9] <https://papers.qnfo.org/papers/algorithmic-graph-search-design-of-heralded-linear-optical-circuits-for-multipar/>. QNFO: Algorithmic Graph-Search Design of Herald Linear Optical Circuits for Multipartite Entanglement: A Reconciled Quantitative Assessment. [10] <https://papers.qnfo.org/papers/ballistic-transport-on-the-bruhattits-tree/>. QNFO: Bruhat-Tits Tree as a Unifying Geometric Object. [11] <https://papers.qnfo.org/papers/geometric-unity-of-computation/>. QNFO: Geometric Unity of Computation. [12] <https://papers.qnfo.org/papers/spectral-benchmarking-of-holographic-quantum-simulations/>. QNFO: Spectral Benchmarking of Holographic Quantum Simulations.

Appendix A. Divergence report

The independent drafts agreed on all core scalings (dimension $2^{n+1} - 2$, memory $16 \cdot 2^n$ bytes, gate cost 2^n , central-cut entropy $\lfloor n/2 \rfloor$ bits, stabilizer count formula, MPS parameter formula). Two divergences arose:

- **D1 (throughput assumption).** Draft A assumed $R = 10^{10}$ amplitude-updates per second per node for the per-gate timing projections; Draft B assumed $R = 10^{11}$, yielding gate times ten times shorter. Convention adopted: $R = 10^{10}$, the more conservative value, with the factor-10 uncertainty stated explicitly in Section 4.3 and Section 6.
- **D2 (survey benchmark).** Draft A took the SPACE survey data volume as $V_{\text{survey}} = 10^{15}$ bytes; Draft C argued for 10^{16} bytes given multi-band spectra of $> 5 \times 10^8$ galaxies. Convention adopted: 10^{15} bytes as the stated assumption, with the alternative ratio (1.8 instead of 18.0) reported in Section 4.5 so the comparison is robust to either convention.

No numerical result in Sections 4-5 depends on the unresolved side of either divergence beyond the labeled uncertainty.

Appendix B. Claim attribution

Claim	Substance	Drafts	Status
C1	Arena is $\mathbb{CP}^{2^n-1}$, real dimension $2^{n+1} - 2$	A, B, C	CONVERGENT
C2	Memory $M_n = 16 \cdot 2^n$ bytes; $M_{50} = 16$ PiB	A, B, C	CONVERGENT
C3	Per-gate cost 2^n amplitude touches	A, B, C	CONVERGENT
C4	Throughput $R = 10^{10} \text{ s}^{-1}$ for timing projections	A, B	CONVERGENT (C silent)
D1	Alternative throughput $R = 10^{11}$	B vs A	DIVERGENT (Appendix A)
C5	Central cut of $n = 100$: $S_{\text{max}} = 50$ bits	A, B, C	CONVERGENT
C6	Page-type typical entropy formula and values	A, B	CONVERGENT
C7	Stabilizer count $2^n \prod_{k=1}^n (2^k + 1)$; $\approx 8.79 \times 10^{19}$ at $n = 10$	A, B, C	CONVERGENT
C8	MPS needs $\chi = 2^{50}$, $\approx 2.5 \times 10^{32}$ parameters for the $n = 100$ central cut	A, B, C	CONVERGENT
C9	Survey-volume comparison via [3]	A, C	CONVERGENT in kind
D2	Survey data volume 10^{15} vs 10^{16} bytes	A vs C	DIVERGENT (Appendix A)
C10	Structured subsets are measure-zero / polynomially parametrized; analogy with [6], [7], [8]	A, B, C	CONVERGENT