

The Price of the Digital Abstraction: Entropy and Encoding Overhead When Discrete Error Models Meet Continuous Quantum Hardware

The Price of the Digital Abstraction: Entropy and Encoding Overhead When Discrete Error Models Meet Continuous Quantum Hardware

Abstract

A working assumption of much quantum error correction (QEC) is that noise can be modeled as a discrete set of Pauli errors acting on two-level systems. We ask what it costs, in information-theoretic terms, to impose this digital abstraction on hardware whose native dynamics are continuous. Using the framing of the QNFO corpus — in particular the observation that a transmon is better described as a qudit of dimension $d \approx 12$ than as a qubit, and that active QEC continuously consumes physical resources — we derive four concrete quantities with fully shown arithmetic. First, the classical register entropy of digitizing one bosonic mode at qudit dimension $d = 12$ is $\log_2 12 \approx 3.585$ bits, and the base truncation entropy $S_{\text{Base}} = \ln(d/2) \approx 1.792$ nats ≈ 2.585 bits. Second, binary encoding of a $d = 12$ qudit requires $m = 4$ qubits, an encoding efficiency of $\eta \approx 0.896$ and a dimension overhead of $16/12 \approx 1.333$. Third, Pauli twirling of an amplitude-damping channel with damping rate $\gamma = 10^{-3}$ yields a discrete surrogate whose environment entropy is 1.291×10^{-2} bits per use, overstating the true environment entropy of 1.141×10^{-2} bits by a factor ≈ 1.131 at matched total error probability. Fourth, for a cold mode with assumed mean occupation $\bar{n} = 0.1$, the discarded Hilbert-space weight is $\approx 3.19 \times 10^{-13}$, i.e., negligible. These are analytic results for idealized channels, not simulations; we state their assumptions and failure modes explicitly and argue that they quantify a map-territory mismatch rather than a physical law.

1. Introduction

Quantum error correction (QEC) as usually taught and deployed assumes a lattice of qubits, a discrete error set $\{I, X, Y, Z\}$, and a threshold theorem guaranteeing that if physical error rates lie below a threshold, arbitrarily long computation is possible with polylogarithmic overhead. This package of assumptions is an engineering convenience. Physical devices — superconducting transmons chief among them — are anharmonic oscillators with a ladder of energy levels, coupled to environments through continuous processes such as amplitude damping (energy relaxation) and dephasing. Forcing the continuous onto the discrete must be paid for somewhere: in discarded Hilbert-space weight, in extra encoding qubits, in entropy bookkeeping that no longer matches the physical environment.

This paper formalizes a slice of that cost. We take as our starting point the QNFO research conjecture that the digital/discrete paradigm in quantum computing is an engineering convenience rather than a physical universal, and that imposing number-theoretic, Pauli-error abstractions on natively bosonic systems incurs a quantifiable overhead — leakage, discretization-induced entropy, and error-correction cost — that is itself a measurable consequence of the mismatch between the map (the discrete model) and the territory (the physical system). The related QNFO analysis of the resource burden of active QEC emphasizes that QEC protects quantum information only by continuously consuming physical resources — qubits, measurement bandwidth, classical processing, and energy [12]; we ask a complementary, information-axis question: how much of this running burden is attributable not to noise itself, but to the *discretization step* through which noise is modeled and corrected?

Our contributions are deliberately modest and fully derived:

1. **Register entropy of digitization.** For one bosonic mode truncated to a qudit of dimension d , the classical side information needed to label the digitized outcome is $\log_2 d$ bits per mode per use. For the transmon-relevant value $d \approx 12$ suggested in the QNFO Project Rosetta analysis [10], this is ≈ 3.585 bits, and the associated base truncation entropy is $S_{\text{Base}} = \ln(d/2) \approx 1.792$ nats.
1. **Binary encoding overhead.** Encoding a d -dimensional qudit into qubits requires $m = \lceil \log_2 d \rceil$ qubits, wasting $2^m - d$ levels. For $d = 12$: $m = 4$, waste = 4 levels, efficiency $\eta \approx 0.896$, dimension overhead ≈ 1.333 .
1. **Entropy distortion under Pauli twirling.** Replacing the continuous-error amplitude-damping channel $\mathcal{L}_\gamma$ by its Pauli-twirled discrete surrogate changes the entropy attributed to the environment. For $\gamma = 10^{-3}$ we compute both entropies exactly and find the surrogate *overstates* the environment entropy by $\Delta H = 1.500 \times 10^{-3}$ bits per use, a ratio of ≈ 1.131 , even though the total error probability is matched exactly.
1. **A negative result on leakage.** For a cold mode with assumed mean occupation $\bar{n} = 0.1$ and cutoff $d = 12$, the discarded population is $\approx 3.20 \times 10^{-13}$: at this operating point the overhead of digitization is bookkeeping overhead, not physical leakage.

We emphasize scope: these are analytic results for idealized single-mode and single-qubit channels. We perform no simulations; where we extrapolate to architectures we do so as clearly labeled projections with stated assumptions. The broader claim — that quantum advantage might be better sought in analog Hamiltonian evolution or natively discrete fermionic systems — is stated as a hypothesis, not a result.

2. Background and Related Work

The supplied bibliography is dominated by works on digitization in non-quantum domains; we cite them for what their entries actually support, namely the general observation that representing physical or social reality digitally carries costs, standards, and failure modes. We then draw on the QNFO corpus for the quantum-specific framing. We note plainly that none of the eight arXiv entries below addresses quantum hardware; their summaries are used only as thematic context, and where a summary is thin we say so.

[1] presents the digital twin model as emerging from the interaction of cyber-physical systems with sensors, Industry 4.0, the Internet of Things, machine learning, and artificial intelligence, and states that the model has been used in many areas as a cost-reduction method. The relevance to our argument is structural: a digital twin is a discrete computational surrogate of a continuous physical asset, and its value proposition is explicitly framed as cost — the same cost axis along which we quantify QEC's digital overhead. The supplied summary gives no quantitative content we can import.

[2] addresses fashion digital archives from a humanities perspective, arguing that analyzing everyday fashion trends requires a digital archive of images over an extended period. It illustrates that digitization entails selection and encoding choices — which images, at what fidelity — analogous to our truncation choice of d levels. The summary supplies no further technical detail.

[3] surveys XML-based approaches to representing compound digital assets, including METS, IMS Content Packaging, and XFDU, within the MPEG-21 Digital Item Declaration framework. It is a standards document for discrete representation of assets; we cite it as evidence that even in classical computing, the choice of discrete representation format is consequential enough to standardize. The summary is truncated and yields no quantitative claims.

[4] analyzes the role of AI and machine learning in digital forensics and incident response, describing their integration as transformative but nascent. The analogy is limited; we cite it only to note that the supplied summary does not connect to quantum topics, and we draw no technical claim from it.

[5] proposes data lakes for digital humanities, noting that traditional humanities data bear various formats and need substantial transformations — encoding and tagging, stemming, lemmatization — to be managed. This is the closest classical analogue to our problem: raw, quasi-continuous, unstructured source data must be transformed into a discrete encoded form before analysis, and the transformation itself is a cost center. The summary gives no numbers.

[6] discusses digital democracy, arguing that despite increasingly powerful digital technology, digital democracy has yet to realize its potential for deliberative transformation. We cite it for its theme that a digital abstraction of a deliberative social process underdelivers relative to the underlying territory — a social, not physical, instance of map-territory mismatch. The summary supports no quantitative statement.

[7] presents EviPlant, a solution for creating, manipulating, and distributing digital forensic challenges, noting that realistic challenge corpora require wear-and-tear, background noise, and actual digital traces, and that creating such challenges is arduous. The point of contact is that realistic noise is hard to synthesize in discrete form — precisely the difficulty we formalize when a discrete Pauli error set must stand in for continuous noise. The summary is truncated beyond this.

[8] reports on MediVerify, Sri Lanka's national-scale online medicines information platform, recording over 1.49 million queries in its first year, in the context of large-scale medicine search behaviour in low- and middle-income countries. We restate its summary but make no derivable claim from it: its subject matter — national-scale digital health search — bears no derivable relation to quantum error models, and we do not use its query volume in any computation (see Appendix A for the divergence on this point).

Turning to the QNFO corpus: [9] frames quantum computing as often presented through a small set of recurring claims — that continuous-variable systems are naturally advantaged, that more modes means more power, that quantum advantage has been demonstrated at device level, that statistics could in principle be non-standard, and that complex amplitudes are essential. Our paper interrogates the first of these recurring claims from the opposite direction: rather than assuming continuous-variable advantage, we ask what it costs to deny continuity by digitizing. The summary supplied for [9] is truncated mid-list, so we use only the claims it explicitly names.

[10] is the most substantive quantum-side source. Its summary states that version 2.0 corrects an earlier error: the transmon spectrum IS discrete (Planck quantization), and binary truncation is a radix error, not a physics error; the transmon is a qudit of $d \approx 12$, not a qubit; and it proposes $S_{\text{Base}} = \ln(d/2) \approx 1.79$ nats. We adopt the qudit picture and the S_{Base} quantity, recompute both exactly in Section 4, and note that the correction cuts both ways for our thesis: if the spectrum is genuinely discrete, part of the "continuity" we invoke is itself an approximation, and the overhead we quantify is radix overhead, not a fight against physics. This nuance is developed in Section 6.

[11], "Lifecycle of a Fault-Tolerant Quantum Computer," has an empty summary in the supplied bibliography. We therefore cannot state anything about its content; we note only that its title indicates topical overlap with our accounting frame in Section 3.4, and we make no claim deriving from it.

[12] analyzes the resource burden of active quantum error correction, stating that QEC protects quantum information only by continuously consuming physical resources — qubits, measurement bandwidth, classical processing, and energy — and that while the theoretical foundations of QEC are mature, the resource burden of running a code, rather than constructing one, is the object of study. This is the direct conceptual ancestor of our overhead accounting: we complement its resource-axis analysis with an information-axis analysis (entropy and dimension), and its summary supports our framing that QEC cost is a running cost, not a one-time construction cost. The supplied summary is truncated and contains no numerical thresholds, so we import no numbers from it.

3. Methods

3.1 Model of digitization

Let $\mathcal{H}_{\text{bos}}$ be the Hilbert space of one bosonic mode, spanned by number states $\{|k\rangle\}_{k=0}^{\infty}$. Digitization consists of choosing a cutoff d and the map

$$\Pi_d = \sum_{k=0}^{d-1} |k\rangle\langle k|,$$

discarding the orthogonal complement. The digitized system is a qudit of dimension d ; the classical control stack must additionally carry a register of $\log_2 d$ bits per mode to address levels, syndrome symbols, and control outcomes.

3.2 Entropy quantities

We track three entropies:

- $H_{\text{reg}}(d) = \log_2 d$: the classical register entropy per digitized mode (bits).
- $S_{\text{Base}}(d) = \ln(d/2)$: the base truncation entropy proposed in [10] (nats).
- H_{env} : the environment entropy per channel use, computed exactly for the physical channel and for its discrete surrogate.

For a diagonal two-outcome distribution with probabilities $(p, 1-p)$ we use the binary entropy $h_2(p) = -p \log_2 p - (1-p) \log_2 (1-p)$.

3.3 Channels

The physical continuous-error channel is amplitude damping with damping rate γ :

$$\mathcal{L}_{\gamma}(\rho) = K_0 \rho K_0^{\dagger} + K_1 \rho K_1^{\dagger}, \quad K_0 = |0\rangle\langle 0| + \sqrt{1-\gamma} |1\rangle\langle 1|, \quad K_1 = \sqrt{\gamma} |0\rangle\langle 1|.$$

Its discrete surrogate is the Pauli channel obtained by twirling $\mathcal{L}_{\gamma}$ over the single-qubit Pauli group, which yields probabilities p_I, p_X, p_Y, p_Z derived in Section 4.3. Twirling is the standard operation by which continuous noise is forced into the discrete Pauli formalism required by stabilizer codes; our object of study is precisely the entropy accounting of this step.

3.4 Accounting frame

Following the framing of [12], we treat QEC cost as a running cost: at each correction cycle the machine consumes qubit-time, measurement bandwidth, classical processing, and energy. Our entropy and dimension quantities are per-cycle, per-mode accounting units that can in principle be added to such a resource budget. We do not attempt a full fault-tolerance lifecycle analysis; [11] would be the natural reference but its supplied summary is empty, so we proceed without it.

4. Analysis

Every input number below is stated with its source, and every arithmetic step is shown.

4.1 Register entropy and base truncation entropy

Input 1: $d = 12$, the transmon qudit dimension stated in [10].

Register entropy:

$$H_{\text{reg}}(12) = \log_2 12 = \frac{\ln 12}{\ln 2} = \frac{2.484907}{0.693147} = 3.584963 \text{ bits.}$$

Base truncation entropy, per the definition in [10]:

$$S_{\text{Base}}(12) = \ln \left(\frac{12}{2} \right) = \ln 6 = 1.791759 \text{ nats.}$$

Converting to bits: $1.791759 \times \log_2 e = 1.791759 \times 1.442695 = 2.584963$ bits. (Check: $\log_2 6 = \log_2 12 - 1 = 3.584963 - 1 = 2.584963$; consistent.)

Interpretation: per digitized mode per use, the digital abstraction requires ≈ 3.585 bits of classical bookkeeping that the un-digitized labeling does not; this is a hard floor of the abstraction, independent of code choice.

4.2 Binary encoding overhead

Input 2: $d = 12$ (from [10], as above).

Number of qubits needed:

$$m = \lceil \log_2 12 \rceil = \lceil 3.584963 \rceil = 4.$$

Encoded Hilbert dimension: $2^m = 2^4 = 16$. Wasted levels: $2^m - d = 16 - 12 = 4$. Dimension overhead ratio:

$$\frac{2^m}{d} = \frac{16}{12} = \frac{4}{3} \approx 1.333333.$$

Encoding efficiency:

$$\eta = \frac{\log_2 d}{m} = \frac{3.584963}{4} = 0.896241.$$

So binary digitization of one transmon qudit idles a fraction $1 - \eta = 1 - 0.896241 = 0.1037594 \approx 10.4\%$ of the encoded Hilbert space and inflates the state space by a factor $4/3$.

4.3 Entropy distortion under Pauli twirling

Input 3: $\gamma = 10^{-3}$, a representative single-qubit relaxation probability per cycle, chosen by us as an illustrative operating point (assumption, not a measurement).

Step 1 — twirled Pauli probabilities. Twirling $\mathcal{L}_\gamma$ over the Pauli group gives the Pauli channel equivalent on average:

$$p_X = p_Y = \frac{\gamma}{4}, \quad p_Z = \frac{\gamma}{2}, \quad p_I = 1 - \gamma.$$

For $\gamma = 10^{-3}$: $p_X = p_Y = 2.5 \times 10^{-4}$, $p_Z = 5 \times 10^{-4}$, $p_I = 0.999$.

Step 2 — environment entropy of the physical channel. In amplitude damping, the environment starts in $|0\rangle$ and ends in $|1\rangle$ with probability γ (it absorbed the excitation) and in $|0\rangle$ with probability $1 - \gamma$. Hence

$$H_{\text{env}}^{\text{phys}} = h_2(\gamma) = -\gamma \log_2 \gamma - (1 - \gamma) \log_2 (1 - \gamma).$$

Compute each term for $\gamma = 10^{-3}$:

$$\begin{aligned}
-\gamma \log_2 \gamma &= -10^{-3} \times \log_2(10^{-3}) = 10^{-3} \times 9.965784 = 9.965784 \times 10^{-3}, \\
-(1 - \gamma) \log_2(1 - \gamma) &= -0.999 \times \log_2 0.999 = 0.999 \times 1.443418 \times 10^{-3} = 1.441975 \times 10^{-3}, \\
(\text{using } \log_2 0.999 &= \ln 0.999 / \ln 2 = (-1.000500 \times 10^{-3}) / 0.693147 = -1.443418 \times 10^{-3}).
\end{aligned}$$

$$H_{\text{env}}^{\text{phys}} = 9.965784 \times 10^{-3} + 1.441975 \times 10^{-3} = 1.140776 \times 10^{-2} \text{ bits.}$$

Step 3 — environment entropy of the twirled surrogate. The twirled Pauli channel can be realized by an environment that carries a classical label with distribution (p_I, p_X, p_Y, p_Z) :

$$H_{\text{env}}^{\text{twirl}} = - \sum_i p_i \log_2 p_i.$$

Terms:

$$\begin{aligned}
-(1 - \gamma) \log_2(1 - \gamma) &= 1.441975 \times 10^{-3} \quad (\text{as above}), \\
\log_2 p_X &= \log_2(2.5 \times 10^{-4}) = \log_2 2.5 + \log_2 10^{-4} = 1.321928 - 13.287712 = -11.965784, \\
-2 p_X \log_2 p_X &= 2 \times 2.5 \times 10^{-4} \times 11.965784 = 5.982892 \times 10^{-3}, \\
\log_2 p_Z &= \log_2(5 \times 10^{-4}) = \log_2 5 - 13.287712 = 2.321928 - 13.287712 = -10.965784, \\
-p_Z \log_2 p_Z &= 5 \times 10^{-4} \times 10.965784 = 5.482892 \times 10^{-3}.
\end{aligned}$$

Summing:

$$H_{\text{env}}^{\text{twirl}} = 1.441975 \times 10^{-3} + 5.982892 \times 10^{-3} + 5.482892 \times 10^{-3} = 1.290776 \times 10^{-2} \text{ bits.}$$

Self-check. The twirled channel's error total is $p_X + p_Y + p_Z = \gamma/4 + \gamma/4 + \gamma/2 = \gamma$, matching the physical excitation-loss probability — but the entropy bookkeeping differs. The twirled entropy (1.291×10^{-2} bits) *exceeds* the physical entropy (1.141×10^{-2} bits). The distortion is therefore an *overstatement* by the surrogate:

$$\Delta H = H_{\text{env}}^{\text{twirl}} - H_{\text{env}}^{\text{phys}} = 1.290776 \times 10^{-2} - 1.140776 \times 10^{-2} = 1.500000 \times 10^{-3} \text{ bits,}$$

$$\frac{H_{\text{env}}^{\text{twirl}}}{H_{\text{env}}^{\text{phys}}} = \frac{1.290776 \times 10^{-2}}{1.140776 \times 10^{-2}} = 1.131489.$$

Interpretation: at $\gamma = 10^{-3}$, replacing the continuous damping channel by its discrete Pauli surrogate misattributes $\approx 1.5 \times 10^{-3}$ bits per use of entropy to the environment — a $\approx 13.1\%$ distortion — even though the total error probability γ is matched exactly. The distortion arises because the twirl redistributes probability into three error species, flattening the label distribution. This is a concrete, fully derived instance of discretization-induced entropy error: the map (Pauli channel) and the territory (damping channel) agree on error weight but disagree on entropy.

4.4 Scaling of the distortion

For small γ , $h_2(\gamma) \approx \gamma \log_2(1/\gamma) + \gamma / \ln 2$, while the twirled entropy is dominated by the three error-species terms. The ratio is a function of γ alone; at the computed operating point it is **1.131489**. We do not extrapolate beyond $\gamma = 10^{-3}$ without recomputation, and we flag the $\gamma \rightarrow 0$ limit as an open question (Section 6).

4.5 Discarded Hilbert-space weight (leakage proxy)

For a thermal bosonic mode with mean occupation $\bar{n}$, the population outside a cutoff d is

$$p_{\text{disc}}(\bar{n}, d) = \left(\frac{\bar{n}}{1 + \bar{n}} \right)^d.$$

Input 4: $\bar{n} = 0.1$ (assumed cold mode; assumption, not a measurement), $d = 12$ (from [10]).

$$\frac{\bar{n}}{1 + \bar{n}} = \frac{0.1}{1.1} = 0.0909091, \quad \ln(0.0909091) = -2.397895,$$

$$p_{\text{disc}} = e^{12 \times (-2.397895)} = e^{-28.774743} \approx 3.20 \times 10^{-13}.$$

At this operating point the discarded weight is negligible; the truncation cost is dominated by the register entropy H_{reg} and the dimension overhead $4/3$, not by leakage. We report this negative result honestly: for cold modes, the overhead of digitization is bookkeeping overhead, not physical leakage.

5. Results

All numbers below are computed in Section 4 with shown arithmetic; none are simulated or measured.

1. **Register entropy of digitization** (Input: $d = 12$ from [10]): $H_{\text{reg}}(12) = \log_2 12 = 3.584963$ bits per mode per use.
1. **Base truncation entropy** (Inputs: $d = 12$ from [10]; definition from [10]): $S_{\text{Base}}(12) = \ln 6 = 1.791759$ nats = 2.584963 bits.
1. **Binary encoding overhead** (Input: $d = 12$): $m = 4$ qubits; wasted levels = 4; dimension overhead = $16/12 = 1.333333$; encoding efficiency $\eta = 0.896241$; idle fraction = 0.1037594.
1. **Entropy distortion under Pauli twirling** (Input: $\gamma = 10^{-3}$, assumed operating point): $H_{\text{env}}^{\text{phys}} = 1.140776 \times 10^{-2}$ bits; $H_{\text{env}}^{\text{twirl}} = 1.290776 \times 10^{-2}$ bits; distortion $\Delta H = 1.500 \times 10^{-3}$ bits per use; ratio = 1.131489 (the discrete surrogate overstates environment entropy by $\approx 13.1\%$ at matched total error probability).
1. **Discarded weight for a cold mode** (Inputs: $\bar{n} = 0.1$ assumed, $d = 12$): $p_{\text{disc}} \approx 3.19 \times 10^{-13}$ — negligible; the digitization overhead at this operating point is bookkeeping, not leakage.

Labeled projection (not a computed result). If a fault-tolerant machine runs N_{cycles} correction cycles on M digitized modes, the cumulative register bookkeeping is $N_{\text{cycles}} \times M \times H_{\text{reg}}(d)$ bits. For illustration, $M = 1000$ modes and $N_{\text{cycles}} = 10^9$ cycles would require $10^9 \times 1000 \times 3.584963 = 3.584963 \times 10^{12}$ bits of cumulative register entropy accounting. This is a projection under stated assumptions (M , N_{cycles} chosen by us; per [12], QEC cost is a running cost, so cumulative quantities are the relevant axis), with uncertainty dominated entirely by the assumed M and N_{cycles} , which span many orders of magnitude across proposed architectures.

6. Discussion

What we have and have not shown. We have derived exact entropy and dimension quantities for idealized single-mode and single-qubit channels. We have not simulated a Josephson circuit, nor compared discrete Pauli channels against continuous drift in a dynamical simulation; that remains future work. Our results quantify the *bookkeeping* face of the digital abstraction; whether it translates into a *fundamental* lower bound on QEC overhead as a function of Hilbert-space continuity is conjecture.

The strongest argument against our own thesis. Entry [10] itself corrects the premise we build on: the transmon spectrum IS discrete by Planck quantization, and binary truncation is a radix error, not a physics error. If the territory is already discrete (a qudit of $d \approx 12$), then the "continuity" we invoke is partly an artifact of the harmonic-oscillator idealization, and our overhead is an encoding inefficiency ($\eta = 0.896$, overhead $4/3$) that better radix choices — qudit encodings, ternary logic, or d -ary digitization — could eliminate. On this reading, the digital abstraction is not fighting physics; it is merely using the wrong radix. Our entropy-distortion result (Result 4) survives this critique in weakened form: even on a genuinely discrete qudit, the *choice* of discrete error algebra (Pauli on qubits) misattributes entropy, but a d -ary error algebra might not.

Failure modes of the derivation. (i) The twirling equivalence holds only on average and for input states diagonal in the right basis; for coherent superpositions the twirled and physical channels differ in more than entropy, and our single entropy metric may mislead. (ii) The environment entropy of the twirled channel assumes the label is carried by a classical environment; a quantum environment realizing the same channel could have different entropy. (iii) $\gamma = 10^{-3}$ and $\bar{n} = 0.1$ are assumed operating points, not measurements; the distortion ratio is γ -dependent and we have computed it at exactly one point. (iv) The register entropy H_{reg} is a lower bound on bookkeeping; real control stacks carry far more.

What would falsify our claims. A demonstration that Pauli-twirled and physical damping channels have identical environment entropy under a defensible environment model would falsify Result 4. A construction achieving binary encoding of a $d = 12$ qudit with efficiency approaching 1 and no discarded weight would falsify Result 3. Measurement of substantial discarded population at cold operating points would falsify Result 5. Conversely, if the distortion ratio $H_{\text{env}}^{\text{twirl}}/H_{\text{env}}^{\text{phys}}$ were shown to be an artifact of our choice of classical-label environment, the entropy-distortion claim would need to be restated in terms of a different operational quantity, such as syndrome statistics or logical error rates under a specific code.

Open questions. How does the distortion ratio behave as $\gamma \rightarrow 0$ and at experimentally measured damping rates? Does the entropy misattribution translate into a measurable penalty in syndrome-processing bandwidth or energy, connecting our information-axis result to the resource-axis analysis of [12]? Can d -ary error algebras on native qudits recover the misattributed entropy, and at what cost in code structure? We leave these open; the supplied bibliography contains no data bearing on them.

7. Conclusion

We have quantified, with fully shown arithmetic, four faces of the cost of imposing a digital abstraction on bosonic quantum hardware: a register-entropy floor of $\log_2 12 \approx 3.585$ bits per digitized mode per use; a base truncation entropy of $\ln 6 \approx 1.792$ nats ≈ 2.585 bits as proposed in [10]; a binary-encoding inefficiency of $\eta \approx 0.896$ with dimension overhead $4/3$; and a Pauli-twirling entropy distortion of 1.500×10^{-3} bits per use (ratio ≈ 1.131) at $\gamma = 10^{-3}$, in which the discrete surrogate overstates environment entropy despite matching total error probability exactly. A negative result bounds discarded leakage weight at $\approx 3.20 \times 10^{-13}$ for cold modes. These quantities support the conjecture that the digital abstraction carries a measurable, computable price — while the Rosetta correction in [10] tempers the claim: the price may be a radix mismatch rather than a fight against physics. Distinguishing these readings is, in our view, the most valuable next experiment.

References

-
- [1] Digital Twin As A Cost Reduction Method. arXiv:2107.14109v1. <https://arxiv.org/abs/2107.14109v1> [2] A Novel Approach to Analyze Fashion Digital Archive from Humanities. arXiv:2107.08351v2. <https://arxiv.org/abs/2107.08351v2> [3] Representing Digital Assets using MPEG-21 Digital Item Declaration.

arXiv:cs/0508065v1. <https://arxiv.org/abs/cs/0508065v1> [4] A Comprehensive Analysis of the Role of Artificial Intelligence and Machine Learning in Modern Digital Forensics and Incident Response. arXiv:2309.07064v2. <https://arxiv.org/abs/2309.07064v2> [5] Data Lakes for Digital Humanities. arXiv:2012.02454v1. <https://arxiv.org/abs/2012.02454v1> [6] Digital Democracy: Episode IV -- A New Hope, How a Corporation for Public Software Could Transform Digital Engagement for Government and Civil Society. arXiv:1910.08604v3. <https://arxiv.org/abs/1910.08604v3> [7] EviPlant: An efficient digital forensic challenge creation, manipulation and distribution solution. arXiv:1704.08990v1. <https://arxiv.org/abs/1704.08990v1> [8] Digital Medicines Information at National Scale: Search Behaviour and System Performance of MediVerify Across 1.5 Million Queries. arXiv:2609.28946v1. <https://arxiv.org/abs/2609.28946v1> [9] QNFO: Locale Framework Applied to Quantum Computing Innovations & Practical Applications [10] QNFO: Project Rosetta: The Approximation Entropy & The Fractal Limits of Digital Physics — v2.0 [11] QNFO: Lifecycle of a Fault-Tolerant Quantum Computer [12] QNFO: The Resource Burden of Active Quantum Error Correction: A Scaling Analysis of Qubit, Bandwidth, and Energy Overheads