

Optimal Radix q under Resource Constraints: A Reconciled Derivation Framework for Ultrametric Hierarchies

Optimal Radix q under Resource Constraints: A Reconciled Derivation Framework for Ultrametric Hierarchies

Abstract

Hierarchical (ultrametric) representations organize N items in a q -ary tree of depth $n = \lceil \ln N / \ln q \rceil$, and the branching factor q determines how storage, interconnection, and traversal resources scale. Prior ultrametric modeling has largely treated the radix as a fixed structural parameter inherited from the problem statement rather than as a design variable to be optimized. This paper treats q as the single decision variable in a constrained resource-allocation problem. We formalize three resource functionals — pointer storage $S(q) = qN/(q-1)$, traversal cost $T(q) = q \lceil \ln N / \ln q \rceil$, and interconnection degree $D(q) = q$ — and minimize a weighted composite $F_{\alpha,\beta}(q)$. We show analytically that the continuous optimum of $q/\ln q$ lies at $q = e \approx 2.718$, that the best integer radix at equal weights for $N = 10^6$ is $q = 4$, with $q = 3$ optimal only in the traversal-dominated regime where the traversal weight β exceeds the derived threshold $\beta^* = 2.286$. We verify this with full arithmetic for $N = 10^6$: under equal weighting and unpadded storage the composite scores are $F_{1,1}(2) = 4.8956$, $F_{1,1}(3) = 4.3227$, $F_{1,1}(4) = 4.2289$, and $F_{1,1}(5) = 4.5075$, so the integer optimum is $q = 4$; $q = 3$ takes over only when the traversal weight exceeds $\beta^* = 2.286$. Under padded storage, where dummy leaves occupy real slots, $q = 4$ remains optimal with $F_{1,1}^{\text{pad}}(4) = 4.2937$. We situate the derivation within a synthesis literature that argues, in several adjacent fields, that optimization must respect cost metrics beyond the simplest count, and we state the falsification conditions under which the $q = 3$ recommendation would fail.

1. Introduction

Ultrametric hierarchies — trees in which the distance between two leaves is determined by the depth of their least common ancestor — are a canonical way of organizing N items when similarity is naturally nested. The branching factor q of such a tree is not cosmetic. It fixes the depth $n = \lceil \ln N / \ln q \rceil$, the number of pointers that must be stored per internal node, the fan-out of the interconnection network, and the number of comparisons performed at each level of a lookup. Yet in much of the modeling tradition the radix arrives as an inherited constant: base 2 because binary trees are the default, base p because a p -adic valuation is mathematically convenient, or whatever radix the downstream consumer of the hierarchy happens to expect.

The question addressed here, posed as a re-entry of prior QNFO work on radix-to-ultrametric synthesis [11], is: **given resource constraints, what is the optimal q ?** The word "re-entry" is used deliberately. In the aerospace literature, re-entry problems are characterized by severe uncertainty and by the need to extrapolate a small number of measured parameters (such as a ballistic coefficient) across a regime where direct observation is impossible; the GOCE case study is the canonical example, using accurate GPS and attitude measurements to compute a precise reference orbit for three weeks of decay and to extrapolate the ballistic coefficient evolution of the object [1]. Our problem has the same shape in miniature: a handful of measured or assumed resource weights must be extrapolated into a design regime where the optimum is sensitive to them, and the honest deliverable is not a single number but a derivation with stated sensitivities and falsification conditions.

The contribution is threefold. First, we give a closed-form derivation of the continuous optimum of the traversal-storage trade-off, showing that it is governed by the minimum of $q/\ln q$ at $q = e$. Second, we evaluate the integer radices $q \in \{2, 3, 4, 5\}$ with complete arithmetic for $N = 10^6$ and show that $q = 4$ is optimal at equal weights, with $q = 3$ becoming optimal only when the traversal weight exceeds the derived threshold $\beta^* = 2.286$. Third, we embed the derivation in a methodological context drawn from the synthesis literature: several adjacent fields have independently concluded that optimizing against a single scalar metric (gate count, alignment proxy, simplest layout score) is inadequate, and that the metric set itself must be widened before optimization is meaningful [2], [3], [4], [7]. Our composite functional $F_{\alpha,\beta}(q)$ is the ultrametric-hierarchy instance of that general lesson.

2. Background and Related Work

We discuss all twelve works in the supplied bibliography, in order. Because the corpus is deliberately cross-disciplinary, each entry is related to our argument only through what its own supplied summary states; where a summary is thin, we say so.

[1] Radar-based Re-Entry Predictions with very limited tracking capabilities: the GOCE case study (arXiv:1805.09171v1). The entry states that the re-entry predictions of GOCE have been deeply investigated because of the large amount of data, mainly radar and GPS, available until re-entry, and that accurate GPS and attitude measurements were used to compute a precise reference orbit for the three weeks of decay and to extrapolate the ballistic coefficient evolution of the object. This is the methodological template for our "re-entry" framing: a well-instrumented reference phase followed by extrapolation of a small parameter set into the regime of interest. Our sensitivity analysis in Section 4 plays the role of the ballistic-coefficient extrapolation: we fix a reference weighting and propagate its uncertainty into the design recommendation.

[2] Synthesis of Reversible Functions Beyond Gate Count and Quantum Cost (arXiv:1004.4609v1). The entry states that many synthesis approaches for reversible and quantum logic generate circuits with respect to simple metrics such as gate count or quantum cost, whereas physically realizing reversible and quantum hardware requires additional constraints, and that the paper describes cost metrics beyond gate count and quantum cost. This is the closest methodological analogue to our problem in the corpus: the radix q is our "simple metric" trap. A hierarchy optimized for depth alone (which drives $q \rightarrow N$) or for storage alone (which drives $q \rightarrow 2$) is exactly the kind of single-metric optimization the entry warns against; hence our composite functional.

[3] InstructLayout: Instruction-Driven 2D and 3D Layout Synthesis with Semantic Graph Prior (arXiv:2407.07580v3). The entry states that existing layout synthesis methods implicitly model object joint distributions and express object relations, hindering generation's controllability, and that InstructLayout integrates a semantic graph prior and a layout decoder to improve controllability and fidelity. The relevance is structural: a semantic graph prior is an explicit hierarchical organization imposed on a generation problem, and our ultrametric tree is the same kind of explicit structural prior. The entry's observation that implicit relational modeling hinders controllability supports our insistence that the radix be an explicit, declared design variable rather than an implicit default.

[4] Synth-by-Reg (SbR): Contrastive learning for synthesis-based registration of paired images (arXiv:2107.14449v3). The entry states that nonlinear inter-modality registration is challenging due to the lack of objective functions that are good proxies for alignment, and proposes a registration loss for weakly supervised image translation that does not require perfectly aligned training data. This is a warning about objective-function quality that applies directly to us: our composite $F_{\alpha,\beta}(q)$ is only as good as its proxy quality for real system cost. We address this by reporting the full Pareto frontier over weights rather than a single point estimate.

[5] Improving the FAIRness and Sustainability of the NHGRI Resources Ecosystem (arXiv:2508.13498v1).

The entry states that in 2024, NHGRI-funded genomic resource projects completed a Self-Assessment Tool and interviews to evaluate their application of FAIR (Findable, Accessible, Interoperable, Reusable) principles and sustainability, that key challenges were identified in metadata tools, data curation, variant identifiers, and data processing, and that the community was engaged through webinars and discussion. We use this as the resource-ecosystem analogue of our problem: a large collection of resources whose access structure (the FAIR dimensions) must be balanced against curation and processing costs, with the assessment instrument playing the role of our resource functionals.

[6] IVOA Recommendation: Resource Metadata for the Virtual Observatory Version 1.12 (arXiv:1110.0514v1).

The entry states that an essential capability of the Virtual Observatory is a means for describing what data and computational facilities are available where, and once identified, how to use them, and that metadata about data collections and data services is required so that VO users can easily find information of interest. A registry of resources is naturally organized as a hierarchy, and the question "how should the registry branch?" is precisely our question in an operational setting; the entry supplies the requirement (findability through metadata) that our depth term $T(q)$ operationalizes.

[7] Text-to-Speech Synthesis Techniques for MIDI-to-Audio Synthesis (arXiv:2104.12292v6).

The entry states that speech synthesis and music audio generation from symbolic input differ in many aspects but share some similarities, and that the study investigates text-to-speech techniques for piano MIDI-to-audio synthesis, using Tacotron and neural source-filter waveform models as basic components. The transfer-of-technique pattern — taking a machinery set built for one domain and re-targeting it to a structurally similar domain — is the pattern we follow when we import the multi-metric synthesis lesson of [2] into ultrametric hierarchy design.

[8] Uncertainty Quantification Study of a Re-entry Breakup (arXiv:2607.03212v1).

The entry states that the uncertainty associated with breakup events during atmospheric re-entry is severe, that limited attempts to gain knowledge of this environment have included breakup recorder-type sensor capsules designed to escape the demising debris cloud and survive in order to transmit data, and that the work models a breakup recorder undergoing this process as a rigid body. Severe uncertainty in a destructive-transition regime is the honest analogy for our weight-sensitivity analysis: when the resource weights α, β are themselves uncertain, the optimum can "break up" from one integer radix to another, and our job is to map where that transition occurs rather than to pretend it does not.

[9] QNFO: Optimal Radix q under Resource Constraints: A Reconciled Derivation Framework for Ultrametric Hierarchies.

The entry states that hierarchical (ultrametric) representations organize N items in a q -ary tree of depth $n = \lceil \ln N / \ln q \rceil$, that the choice of branching factor q determines how storage, interconnection, and traversal resources scale, and that prior ultrametric modeling has treated the radix as a fixed structural parameter. This is the direct predecessor of the present paper and supplies our problem statement, our depth formula, and our three-resource taxonomy (storage, interconnection, traversal). The present paper is the derivation that the predecessor's framework calls for.

[10] QNFO: FACTORING, Adelic Complexity, and the Silent-Radix Principle.

The supplied summary for this entry is empty; it gives no further detail beyond the title. We therefore relate it to our argument only through the title: an adelic treatment of factoring naturally involves a product over all primes p , and a "silent-radix" principle suggests that the radix choice can be made implicit or invisible in some formulations. Whether such an implicit-radix formulation would dissolve or merely relocate the optimization question of this paper is an open question we return to in Section 6; no substantive claim from this entry is used anywhere in our derivations.

[11] **QNFO: HM-LWE v3.0**. The entry states only "Hidden-modulus LWE: novel hardness assumption." The summary is thin, but the title-level concept is usable: a hidden-modulus assumption is one in which the radix-like quantity (the modulus) is not fixed in advance but is part of the hidden structure. This inverts our problem — there the radix is unknown and must be recovered; here it is known and must be chosen — and the inversion is a useful check that our optimization is well-posed only when the radix is a genuine free parameter.

[12] **QNFO: The Trapped-Ion Ultrametric Testbed: A Falsifiability Register for Testing p -Adic Structure in Quantum Dynamics**. The entry states that sixteen published records from a single research program, spanning December 2025 to August 2026, are organized into one testable claim: that trapped-ion quantum simulators are the first near-term platform on which ultrametric (p -adic) structure in quantum dynamics can be accepted or rejected by measurement. This supplies the physical instantiation and, crucially, the falsifiability culture our derivation adopts: if ultrametric structure in such a testbed is organized at a prime radix p fixed by the platform, then our free- q optimization is constrained, and the constraint is experimentally decidable.

3. Methods

3.1 Setup

Let N be the number of items to organize, and let the ultrametric representation be a rooted q -ary tree with all leaves at depth n . The depth is

$$n(q) = \left\lceil \frac{\ln N}{\ln q} \right\rceil.$$

We take the number of leaves as exactly N (padding to a full q -ary tree with at most $q^n - N$ dummy leaves; the padding overhead is analyzed in Section 4.4). Internal nodes at level k (root at $k = 0$) number N/q^{n-k} ; summing over levels, the total number of internal nodes is

$$I(q) = \sum_{k=0}^{n-1} \frac{N}{q^{n-k}} = N \cdot \frac{q^{-1}(1 - q^{-n})}{1 - q^{-1}} = \frac{N(1 - q^{-n})}{q - 1} \leq \frac{N}{q - 1}.$$

3.2 Resource functionals

Following the three-resource taxonomy of [9], we define:

1. **Storage**. Each internal node stores q child pointers, so total pointer storage is

$$S(q) = q \cdot I(q) = \frac{qN(1 - q^{-n})}{q - 1}.$$

The bound $S(q) \leq qN/(q - 1)$ follows in one line: since $n \geq 1$ we have $0 < q^{-n} \leq 1/q < 1$, so $0 < 1 - q^{-n} < 1$, and therefore

$$\frac{S(q)}{N} = \frac{q(1 - q^{-n})}{q - 1} < \frac{q}{q - 1}.$$

1. **Traversal**. A lookup descends n levels, comparing the query against up to q children at each level, giving an upper bound

$$T(q) = q \cdot n(q) = q \left\lceil \frac{\ln N}{\ln q} \right\rceil.$$

1. **Interconnection.** Each internal node has degree q (children) plus one parent link, so the maximum node degree is

$$D(q) = q + 1.$$

3.3 Composite objective

We minimize the weighted, normalized composite

$$F_{\alpha,\beta}(q) = \alpha \frac{S(q)}{N} + \beta \frac{T(q)}{T_{\text{ref}}},$$

where $T_{\text{ref}} = \ln N$ is the depth of an idealized unbounded-fanout lookup, so that both terms are dimensionless and $O(1)$ -to- $O(\log)$ scaled. The interconnection term $D(q)$ is monotone increasing in q and therefore acts in the same direction as the storage term; we fold it into the storage weight and note this explicitly as a modeling choice.

3.4 Continuous relaxation

Dropping the ceiling, minimize $f(q) = q / \ln q$ for $q > 1$:

$$f'(q) = \frac{\ln q - 1}{(\ln q)^2} = 0 \implies \ln q = 1 \implies q = e \approx 2.718.$$

Since $f''(e) = 1/e > 0$, this is the unique minimum, and $f(e) = e$. The continuous optimum of the traversal term therefore sits between the integer radices 2 and 3, and the ceiling function plus the storage term decide the winner.

4. Analysis

Every input number is stated with its source; every arithmetic step is shown.

4.1 Input numbers

- $N = 10^6$ (chosen reference workload; a round number in the range of registry-scale collections such as those motivating [5], [6]).
- $\ln 2 = 0.693147$, $\ln 3 = 1.098612$, $\ln 4 = 1.386294$, $\ln 5 = 1.609438$, $\ln 10^6 = 6 \ln 10 = 13.815511$ (standard natural-logarithm values, carried to six decimal places).
- Reference weights $\alpha = 1$, $\beta = 1$ (equal weighting; the sensitivity analysis in Section 4.5 varies these).

4.2 Depth and traversal for each integer radix

$\ln N / \ln q$ for each candidate:

- $q = 2$: $13.815511 / 0.693147 = 19.931569 \Rightarrow n(2) = \lceil 19.931569 \rceil = 20$.
- $q = 3$: $13.815511 / 1.098612 = 12.575452 \Rightarrow n(3) = 13$.
- $q = 4$: $13.815511 / 1.386294 = 9.965784 \Rightarrow n(4) = 10$.
- $q = 5$: $13.815511 / 1.609438 = 8.585568 \Rightarrow n(5) = 9$.

Traversal costs $T(q) = q n(q)$:

$$T(2) = 2 \times 20 = 40, \quad T(3) = 3 \times 13 = 39, \quad T(4) = 4 \times 10 = 40, \quad T(5) = 5 \times 9 = 45.$$

4.3 Storage and composite score

Storage per leaf, $S(q)/N = q(1 - q^{-n})/(q - 1)$:

- $q = 2$: $q^{-n} = 2^{-20} = 9.54 \times 10^{-7}$, so $S(2)/N = 2(1 - 9.54 \times 10^{-7})/1 = 2.0000$ (to four decimals).
- $q = 3$: $3^{-13} = 1/1594323 = 6.27 \times 10^{-7}$, so $S(3)/N = 3(1 - 6.27 \times 10^{-7})/2 = 1.5000$.
- $q = 4$: $4^{-10} = 9.54 \times 10^{-7}$, so $S(4)/N = 4(1 - 9.54 \times 10^{-7})/3 = 1.3333$.
- $q = 5$: $5^{-9} = 5.12 \times 10^{-7}$, so $S(5)/N = 5(1 - 5.12 \times 10^{-7})/4 = 1.2500$.

With $T_{\text{ref}} = \ln N = 13.815511$ and $\alpha = \beta = 1$:

$$F_{1,1}(q) = \frac{S(q)}{N} + \frac{T(q)}{13.815511}.$$

- $F_{1,1}(2) = 2.0000 + 40/13.815511 = 2.0000 + 2.8956 = 4.8956$.
- $F_{1,1}(3) = 1.5000 + 39/13.815511 = 1.5000 + 2.8227 = 4.3227$.
- $F_{1,1}(4) = 1.3333 + 40/13.815511 = 1.3333 + 2.8956 = 4.2289$.
- $F_{1,1}(5) = 1.2500 + 45/13.815511 = 1.2500 + 3.2575 = 4.5075$.

Result: $q = 4$ wins under equal weighting, with $F_{1,1}(4) = 4.2289 < F_{1,1}(3) = 4.3227 < F_{1,1}(2) = 4.8956 < F_{1,1}(5) = 4.5075$. This is an important correction to the naive intuition from the continuous relaxation: the storage term, which favors large q (fewer internal nodes per leaf), combines with the ceiling-quantized traversal term to push the integer optimum above e .

4.4 Padding overhead

A full q -ary tree of depth n has q^n leaf slots; the padding fraction is $(q^n - N)/N$:

- $q = 2$: $(2^{20} - 10^6)/10^6 = (1048576 - 1000000)/1000000 = 0.0486$, i.e. 4.86%.
- $q = 3$: $(3^{13} - 10^6)/10^6 = (1594323 - 1000000)/1000000 = 0.5943$, i.e. 59.43%.
- $q = 4$: $(4^{10} - 10^6)/10^6 = (1048576 - 1000000)/1000000 = 0.0486$, i.e. 4.86%.
- $q = 5$: $(5^9 - 10^6)/10^6 = (1953125 - 1000000)/1000000 = 0.9531$, i.e. 95.31%.

If dummy leaves must be physically stored, the effective storage becomes $S_{\text{eff}}(q)/N = q \cdot q^n/(N(q - 1))$:

- $S_{\text{eff}}(2)/N = 2 \times 1048576/(10^6 \times 1) = 2.0972$.
- $S_{\text{eff}}(3)/N = 3 \times 1594323/(10^6 \times 2) = 2.3915$.
- $S_{\text{eff}}(4)/N = 4 \times 1048576/(10^6 \times 3) = 1.3981$.
- $S_{\text{eff}}(5)/N = 5 \times 1953125/(10^6 \times 4) = 2.4414$.

Recomputing the composite with padded storage:

- $F_{1,1}^{\text{pad}}(2) = 2.0972 + 2.8956 = 4.9928$.
- $F_{1,1}^{\text{pad}}(3) = 2.3915 + 2.8227 = 5.2142$.
- $F_{1,1}^{\text{pad}}(4) = 1.3981 + 2.8956 = 4.2937$.
- $F_{1,1}^{\text{pad}}(5) = 2.4414 + 3.2575 = 5.6989$.

Under padded storage, $q = 4$ wins by a wider margin, and $q = 3$ falls to last-but-one. The reason is the unlucky fit of $3^{13} = 1594323$ to $N = 10^6$: the radix-3 tree must waste 59.43% of its leaf slots.

4.5 Weight sensitivity (the "breakup map")

Following the severe-uncertainty framing of [8], we map where the integer optimum changes as the traversal weight β varies with $\alpha = 1$, using unpadded storage. The optimum switches between $q = 3$ and $q = 4$ when $F(3) = F(4)$:

$$1.5000 + \beta \cdot 2.8227 = 1.3333 + \beta \cdot 2.8956 \implies 0.1667 = \beta(2.8956 - 2.8227) = 0.0729 \beta \\ \implies \beta^* = 0.1667/0.0729 = 2.286.$$

So for $\beta < 2.286$ (traversal weighted less than about $2.3\times$ storage), $q = 4$ is optimal; for $\beta > 2.286$, $q = 3$ takes over. Checking $q = 2$ against $q = 3$: $F(2) - F(3) = 0.5000 + \beta(2.8956 - 2.8227) = 0.5000 + 0.0729\beta > 0$ for all $\beta > 0$, so $q = 2$ is dominated by $q = 3$ for all $\beta > 0$. Comparing $q = 2$ against the other candidates: $F(2) - F(4) = (2.0000 - 1.3333) + \beta(2.8956 - 2.8956) = 0.6667 > 0$ for all β , so $q = 2$ is also dominated by $q = 4$; and $F(2) - F(5) = (2.0000 - 1.2500) + \beta(2.8956 - 3.2575) = 0.7500 - 0.3619\beta$, which is positive only for $\beta < 0.7500/0.3619 = 2.072$. Hence $q = 2$ is optimal for no value of β under this model for $N = 10^6$: it is beaten by $q = 3$ and $q = 4$ everywhere, and by $q = 5$ for $\beta < 2.072$, while for $\beta \geq 2.072$ it is beaten by $q = 3$ (since $2.072 < 2.286$).

4.6 Asymptotic check

For large N the ceiling becomes negligible and $T(q)/\ln N \rightarrow q/\ln q$, minimized at $q = e$ with value $e \approx 2.718$, while $S(q)/N \rightarrow q/(q-1)$, decreasing in q . The composite's continuous optimum then solves

$$\alpha \frac{d}{dq} \frac{q}{q-1} + \beta \frac{d}{dq} \frac{q}{\ln q} = 0 \implies \frac{-\alpha}{(q-1)^2} + \beta \frac{\ln q - 1}{(\ln q)^2} = 0.$$

At $\alpha = \beta = 1$: try $q = 4$: $-1/9 + (1.386294 - 1)/1.921812 = -0.1111 + 0.2010 = +0.0899 > 0$; try $q = 3$: $-1/4 + 0.098612/1.206949 = -0.2500 + 0.0817 = -0.1683 < 0$. The root lies between 3 and 4; interpolating linearly: $q \approx 3 + 0.1683/(0.1683 + 0.0899) = 3 + 0.6518 = 3.65$. The asymptotic continuous optimum is therefore $q \approx 3.65$, consistent with the integer winner $q = 4$ found exactly above.

5. Results

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

- Continuous optimum of the traversal term:** $q = e \approx 2.718$, with $f(e) = e$ (Section 3.4).
- Integer-radix comparison for $N = 10^6$, equal weights, unpadded storage:** $F_{1,1}(2) = 4.8956$, $F_{1,1}(3) = 4.3227$, $F_{1,1}(4) = 4.2289$, $F_{1,1}(5) = 4.5075$; **optimum $q = 4$** (Section 4.3).
- Padding overhead at $N = 10^6$:** 4.86% ($q = 2$), 59.43% ($q = 3$), 4.86% ($q = 4$), 95.31% ($q = 5$); with padded storage the composite scores are 4.9928, 5.2142, 4.2937, 5.6989, and $q = 4$ **remains optimal with a wider margin** (Section 4.4).
- Weight-sensitivity threshold:** the $q = 3/q = 4$ optimum switch occurs at $\beta^* = 2.286$ (with $\alpha = 1$, unpadded storage); $q = 2$ is dominated for all $\beta > 0$ (Section 4.5).
- Asymptotic continuous optimum of the full composite ($\alpha = \beta = 1$):** $q \approx 3.65$, bracketed by the exact evaluations at $q = 3$ and $q = 4$ shown in Section 4.6.

6. **Projection (labeled as such):** for workloads whose N makes q^n land close to N at $q = 3$ (padding below a few percent), the $q = 3$ vs $q = 4$ ranking can invert relative to Result 2; the inversion condition is exactly the equality in Section 4.5 with the padded storage values substituted, and should be recomputed per workload rather than assumed.

6. Discussion

Limitations. The model counts pointers and comparisons but not cache behavior, memory bandwidth, or the cost of arithmetic at each node; a wide-node radix-4 tree may be faster or slower in practice than the comparison count suggests. The normalization $T_{\text{ref}} = \ln N$ is a convention; any other reference changes β^* linearly, though not the ranking at fixed weights. The padding analysis assumes dummy leaves occupy real storage; in sparse implementations they may not, in which case the unpadded ranking (still $q = 4$ at equal weights) applies. The interconnection term was folded into storage by monotonicity; if degree- q wiring costs superlinearly in q , large radices are penalized further, strengthening the case against $q \geq 5$ but potentially reopening $q = 3$.

Failure modes and self-critique. The headline result — $q = 4$ at $N = 10^6$ — is workload-specific: it depends on the accidental near-fit of 2^{20} and 4^{10} (both 1048576) to 10^6 , and on the unlucky fit of 3^{13} . A different N can flip the ranking; Result 6 states this as a projection, not a finding. The claim that $q = 2$ is never optimal holds only under this functional family; a model with per-comparison costs that grow with node width (e.g., cache-line misses proportional to q) would add a term γq per level, changing $T(q)$ to $(q + \gamma q)n(q)$ and possibly rehabilitating $q = 2$. We have not modeled update costs (rebalancing), which for radix trees are typically $O(n)$ and mildly q -dependent; this is an open modeling gap.

What would falsify the claims. Three falsifiers: (i) a measured workload class in which the effective per-level cost is superlinear in q , for which the $\beta^* = 2.286$ threshold would shift and could drop below 1, making $q = 2$ optimal after all; (ii) a workload N for which the padded-storage ranking inverts in favor of $q = 3$, which would show that the $q = 4$ result is an artifact of the near-fit of $4^{10} = 1048576$ to 10^6 ; (iii) an empirical demonstration that the traversal cost per level is well approximated by $q/\ln q$ even at small q , which would push the effective optimum toward $q = e$ and make $q = 3$ the integer winner at equal weights. Any one of these, established with measured costs, would overturn the headline recommendation.

7. Conclusion

We treated the radix q of an ultrametric hierarchy as a free design variable and minimized a weighted composite of pointer storage and traversal cost. The continuous relaxation of the traversal term is minimized at $q = e \approx 2.718$, but the ceiling-quantized traversal term combined with the storage term shifts the integer optimum: for $N = 10^6$ at equal weights the optimum is $q = 4$ ($F_{1,1}(4) = 4.2289$), robust to padding, with $q = 3$ optimal only when the traversal weight exceeds $\beta^* = 2.286$ and $q = 2$ dominated by $q = 3$ for all $\beta > 0$. The asymptotic continuous optimum of the full composite is $q \approx 3.65$. The recommendation is workload-specific and stated with an explicit falsification map rather than as a universal constant; the central methodological lesson, imported from the multi-metric synthesis literature, is that single-metric radix choices are inadequate and the weight sensitivity must be reported alongside the optimum.

References

- [1] Radar-based Re-Entry Predictions with very limited tracking capabilities: the GOCE case study. arXiv:1805.09171v1. <https://arxiv.org/abs/1805.09171v1> [2] Synthesis of Reversible Functions Beyond Gate Count and Quantum Cost. arXiv:1004.4609v1. <https://arxiv.org/abs/1004.4609v1> [3] InstructLayout: Instruction-Driven

2D and 3D Layout Synthesis with Semantic Graph Prior. arXiv:2407.07580v3. <https://arxiv.org/abs/2407.07580v3>

[4] Synth-by-Reg (SbR): Contrastive learning for synthesis-based registration of paired images. arXiv:2107.14449v3. <https://arxiv.org/abs/2107.14449v3>

[5] Improving the FAIRness and Sustainability of the NHGRI Resources Ecosystem. arXiv:2508.13498v1. <https://arxiv.org/abs/2508.13498v1>

[6] IVOA Recommendation: Resource Metadata for the Virtual Observatory Version 1.12. arXiv:1110.0514v1. <https://arxiv.org/abs/1110.0514v1>

[7] Text-to-Speech Synthesis Techniques for MIDI-to-Audio Synthesis. arXiv:2104.12292v6. <https://arxiv.org/abs/2104.12292v6>

[8] Uncertainty Quantification Study of a Re-entry Breakup. arXiv:2607.03212v1. <https://arxiv.org/abs/2607.03212v1>

[9] QNFO: Optimal Radix q under Resource Constraints: A Reconciled Derivation Framework for Ultrametric Hierarchies

[10] QNFO: FACTORING, Adelic Complexity, and the Silent-Radix Principle

[11] QNFO: HM-LWE v3.0

[12] QNFO: The Trapped-Ion Ultrametric Testbed: A Falsifiability Register for Testing p -Adic Structure in Quantum Dynamics