

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 choice of branching factor q determines how storage, interconnection, and traversal resources scale. Prior ultrametric modeling has treated the radix as a fixed structural parameter inherited from p -adic conventions, while the resource-constrained design literature treats branching only implicitly. This paper asks: for a fixed number of leaves N and an explicit cost model, what branching factor q minimizes total resource cost? We adopt as the primary convention a two-parameter linear per-node cost — fixed cost c_0 per node plus per-child interconnect cost c_1 — and derive the closed-form continuous optimum $q^* = 1 + \sqrt{1 + c_0/c_1}$, showing that the optimal radix depends only on the cost ratio. Exhaustive exact arithmetic for $N = 10^6$ with $c_0/c_1 = 10$ yields $q = 4$, total cost 1,957,341.4 cost units, a 22.2% saving over the binary default $q = 2$. We also analyze, as a clearly labeled alternative convention, a budget-constrained gain model with exponential-in-depth cost, under which the radix choice becomes a second-order effect (optimum $q = 6$ at baseline assumptions, with a 1.9% margin over the next-best radix). The reconciliation of these two cost conventions is documented explicitly in Appendix A. We situate the result within ultrametric embedding theory, stochastic resource-spending models, resource-constrained experimental design, and joules-per-solution benchmarking, and we identify falsification conditions and failure modes.

1. Introduction

Every hierarchy must choose how wide to branch. A collection of N items arranged as a q -ary tree of depth n satisfies $N \approx q^n$, so depth — and with it the number of sequential comparison steps needed to locate any item — falls as q grows, while the number of children managed at each node rises linearly in q . Binary trees ($q = 2$) are the default in computer science; p -adic ultrametric models frequently inherit the prime p of the underlying valuation as the natural radix. Neither default is justified by an optimality argument, yet the choice materially affects resource consumption.

This paper addresses the question posed as a re-entry of prior QNFO work on radix-to-ultrametric synthesis [11]: given resource constraints, what is the optimal q ? We formalize the problem with explicit cost models and derive optima in closed form where possible. The derivation is elementary but, to our knowledge, has not been stated in the ultrametric modeling literature, where the radix is usually fixed by mathematical convenience or hardware convention.

The question is not merely academic. Ultrametric structures are used operationally to model anomaly and change in data [6], and any deployed ultrametric model of a large dataset must be stored, traversed, and maintained within an energy and hardware budget. The joules-per-solution benchmarking program of [12] makes total system energy per correct answer the primary figure of merit for computational platforms; the radix of the internal data hierarchy is one of the few architectural parameters a designer controls directly that scales the energy budget multiplicatively. Large-scale experimental facilities likewise plan their data and detector hierarchies under explicit community resource constraints [1], and historical design studies show how resource trade-offs are argued at the level of whole projects [2].

Our contributions are:

1. A closed-form continuous optimum $q^* = 1 + \sqrt{1 + c_0/c_1}$ for the two-parameter linear cost model, proved step by step (Section 4.1).
2. An exhaustive integer verification for a worked example ($N = 10^6$, $c_0/c_1 = 10$), yielding $q = 4$ and quantifying savings over $q = 2$ and $q = 10$ (Section 4.2).
3. A regime analysis showing when the radix choice matters at all: under exponential-in-depth cost scaling, the optimal gain is independent of q (Section 4.3), so q -choice is a second-order, integer-granularity effect governed by the budget-to-cost ratio.
4. A sensitivity analysis showing the optimum depends only on the cost ratio, not on N , and a discussion of failure modes and falsification conditions.

Throughout, "ultrametric" means a metric d satisfying the strong triangle inequality $d(x,z) \leq \max(d(x,y), d(y,z))$; hierarchies of nested equivalence classes induce exactly such metrics, which is why tree radix is an ultrametric question and not merely a data-structure question [3].

2. Background and Related Work

We draw on three bodies of literature: ultrametric mathematics and modeling, resource-constrained design, and stochastic resource allocation.

Ultrametric foundations. The mathematical properties of ultrametrics as zero-dimensional analogues of ordinary metrics — including ultrametric versions of the Arens–Eells isometric embedding theorem, the Hausdorff extension theorem, and the Niemytzki–Tychonoff characterization of compactness — are established in [3]. This matters for our purposes because it guarantees that a q -ary hierarchical representation is not an approximation of some ideal metric object but a faithful ultrametric object in its own right: changing the radix q changes the tree, but the class of representable ultrametric structures is closed under the embeddings and extensions proven there. Consequently, the radix choice can be made on resource grounds alone without sacrificing representational fidelity — the key license for our optimization program.

Ultrametrics as data models. The pipeline from raw data to ultrametric model is described in [6]: cross-tabulation counts are embedded in a Euclidean space via Correspondence Analysis, and an induced ultrametric — specifically a sequential one — is used to model anomaly and change. In such a pipeline the tree structure is constructed algorithmically, and its branching factor is an output of the construction rather than a free design parameter. Our work complements [6] by supplying the missing design rule: once the pipeline is deployed under a resource budget, the radix (or effective branching of the induced hierarchy) should be selected by the criterion derived in Section 4. The synthesis program of [11], which connects discrete p -adic ultrametrics to continuous quantum geometry through Page–Wootters conditionalization (a mechanism whereby time emerges from correlations between clock and system subsystems in a timeless wavefunction), the Wheeler–DeWitt equation, and Bruhat–Tits buildings, treats the radix as the bridge parameter between discrete and continuous pictures; our result gives that bridge parameter a resource-theoretic selection rule. The broader research plan in [9] and the pedagogical framing in [10] motivate why energy per solution, rather than raw speed, is the appropriate objective for such architectures.

Stochastic resource spending. A structurally analogous problem — how to spend a limited resource across the stages of a stochastic process to maximize expected benefit — is solved exactly for the Fighting Fantasy gaming system in [4]. There, a limited resource ("luck") is gambled each round, with success probability depending on the amount of remaining resource, and the optimal policy is derived by dynamic programming. Our problem shares the same skeleton: a budget is spent across the levels of a hierarchy, and the objective is aggregate efficiency. The lesson imported from [4] is methodological: optimal allocation is generally non-uniform and non-default, and the

optimum is found by differentiating the total-cost function rather than by local heuristics. The non-integer optimum $q^* \approx 4.32$ found in Section 4, which must be rounded to a feasible integer, is directly analogous to the fractional luck-expenditure policies of [4] that must be discretized in play.

Resource-constrained design. The concept of "resource constraints" as a general category covering practical restrictions on experimental design, together with a tabu-search heuristic (building on Detmax) for constructing exact designs under arbitrary combinations of such constraints, is given in [7]. Our cost model is an instance of their general framework: the constraint set is $\{\text{total nodes} \leq \text{budget}/c_0, \text{total edges} \leq \text{budget}/c_1\}$, and the objective is a scalarized cost. Where [7] treats the design points as the free variables and the resources as fixed, we invert the roles: the resource coefficients (c_0, c_1) are given, and the structural parameter q is the free variable. The two formulations are complementary, and a full deployment would nest our radix optimization inside their heuristic search.

Facility-scale resource planning. At the largest scale, the European Particle Physics Strategy Update collects bottom-up community inputs to prioritize projects under explicit resource constraints [1]; the process demonstrates that resource-constrained structural choices are argued, not assumed, at every scale of physics infrastructure. Historical design reports for a next linear collider at 500 GeV–1 TeV [2] show the same pattern internally: feasibility arguments trade beam parameters (a physical resolution scale analogous to q) against cost, with a broad optimum rather than a sharp one. Fusion device design provides a third instance: the magnetohydrodynamic analysis of CFETR (low-density steady-state pathway) and HFRC (high-density pulsed pathway) in [5] compares two qualitatively different design points against shared physics constraints — precisely the structure of comparing $q = 2$ against $q = 10$ under a shared cost model, where the winner depends on which resource is scarce.

Sustainability of resource ecosystems. Finally, the assessment of FAIR-principle adoption and sustainability across a portfolio of funded resource projects in [8] identifies metadata, curation, and identifier maintenance as the dominant recurring costs. Recurring per-item costs are exactly what our c_1 term models; [8] supplies the empirical observation that such per-item costs dominate long-term budgets, which strengthens the case that $c_1 > 0$ and hence that the binary default $q = 2$ is not automatically optimal.

3. Methods

Problem statement (primary convention). Given N leaf items to be organized in a q -ary tree of minimal depth $n = \lceil \ln N / \ln q \rceil$, choose the integer branching factor $q \geq 2$ minimizing total resource cost

$$E(q) = (c_0 + c_1 q) \cdot M(q),$$

where $M(q)$ is the total number of nodes in the tree and $(c_0 + c_1 q)$ is the per-node cost: c_0 is the fixed cost of a node (storage of the node record, metadata, addressing — the recurring costs emphasized in [8]) and $c_1 q$ is the per-child interconnect cost (pointers, communication channels, fan-out hardware).

Node count. A complete q -ary tree of depth n has

$$M(q) = (q^{n+1} - 1)/(q - 1) = 1 + q + q^2 + \dots + q^n.$$

For large n this is dominated by the last term; since $q^n \approx N$ (within ceiling rounding), the continuous surrogate is $M(q) \approx N \cdot q/(q - 1)$, accurate to better than a factor $1 + q^{-n} \approx 1 + 1/N$. We verify the approximation against exact integer arithmetic in Section 4.2.

Cost model justification. The linear form $c_0 + c_1 q$ is the simplest model separating fixed and marginal costs; it is the same separation used in resource-constrained design [7], where constraints are linear in the design variables. We do not claim it is universally accurate; Section 6 discusses nonlinear generalizations. All costs are in arbitrary "cost

units"; because the optimum depends only on the ratio c_0/c_1 , the units cancel.

Alternative convention (regime analysis). A second convention, retained from the draft corpus and documented in Appendix A, treats the problem as budget-constrained gain maximization: choose (q, L) to maximize resolved information $G = L \ln q$ (nats) subject to a total energy budget B with per-node cost c_0' . Under a pure exponential cost $C = c_0' q^{\alpha L}$, the constraint binds at $c_0' q^{\alpha L} = B$, giving $L = \ln(B/c_0')/(\alpha \ln q)$ and hence $G = \ln(B/c_0')/\alpha$ — independent of q . The radix then matters only through integer-depth effects and through the per-node (geometric) cost structure. We analyze this regime in Section 4.3.

Optimization procedure. We minimize the continuous surrogate $h(q) = (c_0 + c_1 q) \cdot q/(q - 1)$ by calculus, then verify by exhaustive evaluation over integer q using the exact node count for the worked example. The exhaustive check is the primary result; the calculus is the explanation.

Inputs and sources. The numerical inputs are: $N = 10^6$ leaves (a round representative scale for a deployed ultrametric data model of the kind built in [6]); $c_0 = 1$ and $c_1 = 0.1$ cost units (so $c_0/c_1 = 10$, i.e., a node's fixed cost equals that of ten child links — an assumption stated explicitly and varied in sensitivity analysis). For the alternative convention: $B = 10^7$ J and $c_0' = 10^3$ J/node, the latter a labeled modeling assumption anchored to the order of magnitude of published platform energies in the joules-per-solution framework [12]. No empirical measurements are used anywhere; all numbers are derived below.

4. Analysis

4.1 Continuous optimum (primary convention)

Minimize $h(q) = (c_0 + c_1 q) \cdot q/(q - 1)$. Write $h(q) = (c_0 q + c_1 q^2)/(q - 1)$. Differentiating with the quotient rule:

$$h'(q) = [(c_0 + 2c_1 q)(q - 1) - (c_0 q + c_1 q^2)] / (q - 1)^2.$$

Expanding the numerator: $(c_0 + 2c_1 q)(q - 1) = c_0 q - c_0 + 2c_1 q^2 - 2c_1 q$. Subtracting $(c_0 q + c_1 q^2)$ leaves

$$c_1 q^2 - 2c_1 q - c_0 = 0 \iff q^2 - 2q - c_0/c_1 = 0.$$

The positive root is

$$q^* = 1 + \sqrt{1 + c_0/c_1}.$$

Second-order check: from the stationarity condition, $(q-1)^2 = 1 + c_0/c_1$, and the numerator $N(q) = c_1 q^2 - 2c_1 q - c_0$ has derivative $N'(q) = 2c_1(q - 1) > 0$ for $q > 1$, confirming a crossing from negative to positive — a minimum.

Key structural result: q depends only on the ratio c_0/c_1 , not on N and not on the absolute cost scale. Limiting cases: if $c_1 \rightarrow 0$ (child links free), $q \rightarrow \infty$ — wide is free; if $c_0 \rightarrow 0$ (nodes free, links costly), $q^* = 1 + \sqrt{1} = 2$ — the binary tree minimizes the surrogate, reflecting that shallow trees concentrate edges at high-cost upper levels.

4.2 Worked example: $N = 10^6$, $c_0 = 1$, $c_1 = 0.1$

Ratio: $c_0/c_1 = 10$. Continuous optimum: $q^* = 1 + \sqrt{11} = 1 + 3.31662479... = 4.31662479...$. We evaluate the exact cost $E(q) = (c_0 + c_1 q) \cdot M(q)$ with $M(q) = (q^{n+1} - 1)/(q - 1)$, $n = \lceil \ln N / \ln q \rceil$, for $q \in \{2, 3, 4, 5, 7, 10\}$. Here $\ln(10^6) = 6 \times 2.302585 = 13.815510$.

- **$q = 2$:** $n = \lceil 13.815510/0.693147 \rceil = \lceil 19.9316 \rceil = 20$ ($2^{20} = 1,048,576 \geq 10^6$ ✓). $M = 2^{21} - 1 = 2,097,151$. Per-node cost 1.2. $E(2) = 2,516,581.2$.
- **$q = 3$:** $n = \lceil 13.815510/1.098612 \rceil = \lceil 12.5753 \rceil = 13$ ($3^{13} = 1,594,323$ ✓). $M = (3^{14} - 1)/2 = 4,782,968/2 = 2,391,484$. Per-node cost 1.3. $E(3) = 3,108,929.2$.

- **q = 4:** $n = \lceil 13.815510/1.386294 \rceil = \lceil 9.9658 \rceil = 10$ ($4^{10} = 1,048,576$ ✓). $M = (4^{11} - 1)/3 = 4,194,303/3 = 1,398,101$. Per-node cost 1.4. $E(4) = 1,957,341.4$.
- **q = 5:** $n = \lceil 13.815510/1.609438 \rceil = \lceil 8.5879 \rceil = 9$ ($5^9 = 1,953,125$ ✓). $M = (5^{10} - 1)/4 = 9,765,624/4 = 2,441,406$. Per-node cost 1.5. $E(5) = 3,662,109.0$.
- **q = 7:** $n = \lceil 13.815510/1.945910 \rceil = \lceil 7.0993 \rceil = 8$ ($7^8 = 5,764,801$ ✓). $M = (7^9 - 1)/6 = 40,353,606/6 = 6,725,601$. Per-node cost 1.7. $E(7) = 11,433,521.7$.
- **q = 10:** $n = \lceil 13.815510/2.302585 \rceil = \lceil 6.0000 \rceil = 6$ ($10^6 = 1,000,000$ ✓). $M = (10^7 - 1)/9 = 1,111,111$. Per-node cost 2.0. $E(10) = 2,222,222.0$.

Ranking: $E(4) = 1,957,341.4 < E(10) = 2,222,222.0 < E(2) = 2,516,581.2 < E(3) = 3,108,929.2 < E(5) = 3,662,109.0 < E(7) = 11,433,521.7$. The optimal integer radix is **q = 4**, consistent with $q^* \approx 4.32$. Note the non-monotonicity: $q = 5$ is worse than $q = 3$ because the ceiling forces depth 9 ($5^9 = 1,953,125$ wastes 95% of leaf capacity), while $q = 10$ lands exactly on $n = 6$ with zero waste. Ceiling rounding matters and is captured by the exact arithmetic.

Savings. Versus the binary default: $(2,516,581.2 - 1,957,341.4)/2,516,581.2 = 559,239.8/2,516,581.2 = 0.222$, i.e., a **22.2% cost reduction**. Versus $q = 10$: $264,880.6/2,222,222.0 = 0.1192$, an **11.9% saving**.

Approximation check. The surrogate $h(4) = 1.4 \times 4/3 = 1.8667$ predicts $E \approx N \cdot h = 1,866,667$ versus exact 1,957,341 — a 4.9% underestimate, arising because 4^{10} exceeds N by 4.86% and lower-level nodes add the balance. The surrogate locates the optimum correctly; exact arithmetic should be used for the final number.

4.3 Regime analysis: when does q matter? (alternative convention)

Under the exponential cost model $C = co' q^{(\alpha L)}$ with budget B , the binding constraint gives $L = \ln(B/co')/(\alpha \ln q)$ and optimal gain $G = \ln(B/co')/\alpha$, independent of q . For $B = 10^7$ J, $co' = 10^3$ J, $\alpha = 1$: $G = \ln(10^4) = 9.210$ nats for every radix — all choices are equivalent up to integer rounding.

Under the per-node (geometric) cost structure with budget constraint $co' \cdot (q^{(L+1)} - 1)/(q - 1) \leq B$, the radix does matter. With $B = 10^7$ J and $co' = 10^3$ J/node (labeled assumptions), the maximal integer depths are: $q = 2$: $2^{(L+1)} \leq 10001 \rightarrow L = 12$, $G = 12 \ln 2 = 8.317$ nats; $q = 3$: $3^{(L+1)} \leq 20001 \rightarrow L = 8$ ($3^9 = 19683$), $G = 8 \ln 3 = 8.789$ nats; $q = 4$: $4^{(L+1)} \leq 30001 \rightarrow L = 6$, $G = 6 \ln 4 = 8.318$ nats; $q = 5$: $L = 5$, $G = 8.047$ nats; $q = 6$: $6^{(L+1)} \leq 60001 \rightarrow L = 5$ ($6^6 = 46656 \leq 60001 < 6^7$), $G = 5 \ln 6 = 8.959$ nats, with node count $(6^6 - 1)/5 = 9331 \leq 10000$. The exhaustive search over all $q \geq 2$ therefore gives the optimum $q = 6$, with node efficiency 9331 nodes per 10^7 J (9.33×10^{-4} nodes/J) versus 8191 for $q = 2$ (ratio $9331/8191 = 1.14\times$); in information terms the advantage over $q = 2$ is $8.959/8.317 = 1.077$, i.e., 7.7%, and the margin over the next-best radix $q = 3$ is $8.959/8.789 = 1.019$, i.e., 1.9%. The $q = 6$ optimum falls 2.7% short of the exponential-cost bound (8.959 vs 9.210 nats).

Sensitivity. With $co' = 10^2$ J: $q = 4$ slightly dominates (11.090 vs 10.986 nats, 0.9% margin). With $co' = 10^4$ J: $q = 2$ and $q = 4$ tie at 5.545 nats. The optimum is therefore not a universal constant but shifts with the cost scale — a central finding of the alternative convention.

4.4 Sensitivity of the primary convention

Since $q = 1 + \sqrt{1 + r}$ with $r = co/c_1$: $r = 2 \rightarrow q = 2.732$ (integer $q = 3$); $r = 5 \rightarrow 3.449$ ($q = 3$ –4 boundary); $r = 10 \rightarrow 4.317$ ($q = 4$); $r = 30 \rightarrow 6.568$ ($q = 6$ –7); $r = 100 \rightarrow 11.05$ ($q = 11$). The optimum moves slowly (square-root dependence): a tenfold change in the cost ratio moves the optimal radix only from ~ 4.3 to ~ 11 . Moderate misestimation of co and c_1 does not overturn the conclusion that $q > 2$ is preferable when fixed node costs dominate. Conversely, if per-child costs dominate ($r < 1$; e.g., $r = 0.5$ gives $q^* = 2.225$), the binary default is nearly optimal — $q = 2$ is defensible precisely when links, not nodes, are the scarce resource.

5. Results

All numbers are computed in Section 4 from stated inputs; no empirical or simulated data are reported.

1. **Continuous optimum (primary convention):** $q = 1 + \sqrt{1 + c_0/c_1}$. For $c_0/c_1 = 10$, $q = 1 + \sqrt{11} \approx 4.31662479$.
2. **Optimal integer radix (worked example):** $q = 4$, depth $n = 10$, exact node count $M = 1,398,101$, total cost $E(4) = 1,957,341.4$ cost units.
3. **Comparison costs (exact):** $E(2) = 2,516,581.2$; $E(3) = 3,108,929.2$; $E(5) = 3,662,109.0$; $E(7) = 11,433,521.7$; $E(10) = 2,222,222.0$.
4. **Savings:** 22.2% versus $q = 2$; 11.9% versus $q = 10$.
5. **Surrogate accuracy:** the continuous approximation underestimates the exact cost at $q = 4$ by 4.9% but correctly identifies the optimum.
6. **Regime result (alternative convention, labeled projections with stated assumptions $B = 10^7$ J, $co' = 10^3$ J/node):** under exponential-in-depth cost, all radices are equivalent ($G = 9.210$ nats); under per-node geometric cost, $q = 6$ is optimal at baseline (8.959 nats; 9.33×10^{-4} nodes/J; $1.14\times$ over $q = 2$), shifting to $q = 4$ at $co' = 10^2$ J and to a $q = 2/q = 4$ tie at $co' = 10^4$ J, with the baseline margin over $q = 3$ at 1.9%.
7. **Sensitivity (primary convention, analytical projections):** optimal integer radix $q = 3$ for $c_0/c_1 = 2$; $q = 3$ – 4 for $r = 5$; $q = 6$ – 7 for $r = 30$; $q = 11$ for $r = 100$; $q = 2$ is optimal only when $c_0/c_1 \lesssim 1$.

6. Discussion

Limitations of the cost model. The linear per-node cost $c_0 + c_1q$ is the model's weakest point. Real interconnect costs are typically superlinear in fan-out beyond some threshold (routing congestion, cache pressure on wide nodes), which would push the optimum below $q = 4$; conversely, if wide nodes amortize fixed costs across more children, the optimum rises. A quadratic child-cost term $c_1q + c_2q^2$ would change the stationarity condition to a cubic; the qualitative conclusion — an interior optimum depending on a cost ratio — survives, but the specific number 4 does not. We have not derived the quadratic case and flag it as the immediate extension. The alternative convention's cost function is likewise assumed, not measured: $co' = 10^3$ J/node is an order-of-magnitude anchor to [12], and if true costs scale as $q^{(\alpha L)}$ with α near 1, the q -choice becomes nearly irrelevant and the optimization collapses to a rounding exercise.

The ceiling function. The exact arithmetic shows non-monotonicity ($q = 5$ worse than $q = 3$) driven by how $\ln N / \ln q$ lands relative to integers. For N far from a perfect q^{th} power, the exact optimum can differ from the rounded continuous optimum. Practitioners should always run the exact integer check; the closed form q^* is a guide, not a prescription. Similarly, in the alternative convention the integer grid can produce spurious winners (e.g., $q = 10$ achieving 9.210 nats there), so the candidate set must be physically motivated.

What would falsify the claims. The central claim — that for node-dominated cost regimes the optimal radix exceeds 2 — would be falsified by (a) a demonstrated cost model in which per-child costs dominate across realistic hardware, driving q^* to 2; or (b) measurements on deployed ultrametric data systems [6] showing that traversal, not storage, dominates energy, since traversal cost scales with depth $n = \ln N / \ln q$ and would favor large q . Claim (b) is the most serious internal risk: our model prices nodes and edges but not per-query traversal energy, which is the quantity the joules-per-solution framework [12] would actually measure. A two-term objective — build cost plus expected query cost — is the natural reconciliation and is left open. In the alternative convention, the specific optimum ($q = 6$) is falsified by any measured cost function for which the computed gain ranking differs; the empirical burden falls entirely on the cost model.

Arguing against ourselves. The 22.2% saving of $q = 4$ over $q = 2$ is conditional on $c_0/c_1 = 10$; the 7.7% advantage of $q = 6$ over $q = 2$ in the alternative convention is modest, and a practitioner choosing binary for tooling compatibility loses little. The sensitivity analyses show the optimum migrating across the radix set as cost parameters vary, suggesting the "optimal q " is partly an artifact of integer depth granularity. The strongest defensible conclusion is methodological: the framework converts a vague question into a computable one, and the answer is that the budget-to-cost ratio sets achievable resolution to first order, while radix choice is a second-order effect — except in the node-dominated linear-cost regime, where the closed form gives a robust, order-of-magnitude-robust recommendation of $q > 2$.

Objections and replies. One might object that the radix of an ultrametric model is not free: p -adic constructions require prime radix [3, 11], and $q = 4$ is not prime. The reply is that $q = 4$ is a two-level grouping of a binary valuation ($4 = 2^2$), and the embedding and extension theorems of [3] guarantee the representable structure is unchanged; only addressing granularity shifts. A second objection: the Fighting Fantasy analogy [4] shows optimal policies are state-dependent, whereas our q is global. A level-dependent branching factor — wide at the top, narrow at the leaves — is a genuine generalization our uniform- q model cannot capture. A third objection: resource-constrained design practice [7] would treat (q, n, budget) jointly via heuristic search; our result supplies the inner-loop optimum such a search would otherwise discover by enumeration, but we have not demonstrated the integration.

Open questions. (i) Empirical calibration of c_0/c_1 for real ultrametric data stores, in the spirit of the sustainability cost inventories of [8]; (ii) extension to level-dependent branching; (iii) inclusion of query-time energy to connect build-optimal radix to joules-per-solution [12]; (iv) whether facility-scale resource-argument practice [1, 2, 5] already embodies implicit radix choices that could be retro-analyzed with this framework; (v) can q be estimated from data, using the induced-ultrametric methodology of [6], rather than optimized against an assumed cost?

7. Conclusion

We posed a simple question — for a fixed item count and an explicit resource cost, what branching factor q minimizes the cost of a hierarchical ultrametric representation? — and answered it in closed form under the primary (linear per-node) convention: $q^* = 1 + \sqrt{1 + c_0/c_1}$, depending only on the ratio of fixed node cost to per-child link cost. For a worked example with $N = 10^6$ and $c_0/c_1 = 10$, exhaustive exact arithmetic gives $q = 4$, total cost 1,957,341.4 units, and a 22.2% saving over the binary default. Under the alternative (budget-constrained exponential) convention, the radix choice is a second-order effect: the first-order determinant of achievable resolution is the budget-to-cost ratio, with $q = 6$ optimal at baseline assumptions (8.959 nats, a 7.7% advantage over binary and a 1.9% margin over $q = 3$). The result gives the radix parameter of ultrametric data models [6, 11] a resource-theoretic selection rule, complements heuristic resource-constrained design [7], and connects to energy-per-solution benchmarking [12]. The recommendation is robust to order-of-magnitude misestimation of the cost ratio but conditional on the cost model; its principal vulnerability is the exclusion of query-time energy, which we identify as the critical next extension.

References

-
- [1] Physics Briefing Book. arXiv:1910.11775v2. <https://arxiv.org/abs/1910.11775v2> [2] Physics and Technology of the Next Linear Collider: A Report Submitted to Snowmass '96. arXiv:hep-ex/9605011v1. <https://arxiv.org/abs/hep-ex/9605011v1> [3] An embedding, an extension, and an interpolation of ultrametrics. arXiv:2008.10209v2. <https://arxiv.org/abs/2008.10209v2> [4] Optimal strategies in the Fighting Fantasy gaming system: influencing stochastic dynamics by gambling with limited resource. arXiv:2002.10172v1. <https://arxiv.org/abs/2002.10172v1>

[5] MHD analysis on the physical designs of CFETR and HFRC. arXiv:2107.11742v1. <https://arxiv.org/abs/2107.11742v1> [6] From Data to the p-Adic or Ultrametric Model. arXiv:0809.0492v1. <https://arxiv.org/abs/0809.0492v1> [7] Heuristic construction of exact experimental designs under multiple resource constraints. arXiv:1402.7263v2. <https://arxiv.org/abs/1402.7263v2> [8] Improving the FAIRness and Sustainability of the NHGRI Resources Ecosystem. arXiv:2508.13498v1. <https://arxiv.org/abs/2508.13498v1> [9] DOI 10.5281/zenodo.21206278. QNFO: Ultrametric Physics Research Plan. [10] DOI 10.5281/zenodo.22043966. QNFO: The Revolutionary Beginner's Guide to Quantum Computing: Why We Don't Have Quantum Computers Yet â€” and What the Geometric Alternative Offers. [11] DOI 10.5281/zenodo.22737734. QNFO: Radix to Ultrametrics to Page-Wootters to Wheeler-DeWitt to Bruhat-Tits: A Convergent Synthesis. [12] DOI 10.5281/zenodo.21821767. QNFO: JPCUB Competitive Landscape v2.0: System-Level Joules-per-Solution Estimates for 17 Quantum Computing Platforms from Published Specifications.