Phase-Weighted Path-Volume Geometry for Distributed Inverse QFT

Jonathan R. Landers

May 15, 2026

A standalone extension note connecting distributed iQFT routing costs with the path-volume framework of *Path-Length Distributions, Inverse Geometry, and Designed Neighborhood Spectra.*

Abstract

The distributed inverse Quantum Fourier Transform is usually analyzed by counting which remote controlled-phase interactions remain after approximation and pruning. A routing extension then asks how far those retained interactions travel through a physical QPU interconnect. This note adds a more geometric layer, applying the path-length distribution program from *Path-Length Distributions, Inverse Geometry, and Designed Neighborhood Spectra* to distributed iQFT routing [2]. The exponential phase law of the iQFT induces a metric on logical QPU blocks; thresholding weak phase interactions is equivalent to taking a ball graph in that induced metric. Once a placement into hardware is chosen, the retained interactions define a phase-weighted routing spectrum: a probability measure on physical route lengths. The ordinary routing cost is the first moment of this spectrum. A Gray-code hypercube placement becomes a spectral compactification result, placing all retained phase-weighted routing mass inside radius $\min\{D, \log_2 P\}$. Replacing shortest-route distance by route-volume distributions then yields an inverse-design problem: choose a physical QPU graph whose path-volume law matches the communication spectrum demanded by the distributed iQFT.

1 Two starting points

This note joins two related viewpoints. The first is architectural: the inverse Quantum Fourier Transform is a circuit-level object on a single processor, but in a distributed architecture it acquires geometry. If a controlled phase rotation acts between qubits on the same QPU, it is local. If it acts between qubits on different QPUs, it becomes a communication event. Approximate distributed iQFT constructions exploit the fact that phase angles decay exponentially with separation, so weak long-range rotations can be pruned. After pruning, one obtains a retained logical communication graph and must route its edges through a physical interconnect [1, 3].

The second viewpoint comes directly from the companion path-volume paper *Path-Length Distributions, Inverse Geometry, and Designed Neighborhood Spectra* [2]. In Euclidean space, the collection of one-waypoint paths between two endpoints has a specific path-volume law: equal-length sets are ellipses, and the density of available detours follows from area growth. More generally, a metric-measure space induces path-length spectra. Conversely, one may try to infer or design a geometry whose path-volume or neighborhood spectra have a prescribed shape.

Central idea

The distributed iQFT does not merely run on a physical graph. Its phase decay law induces a logical geometry. A hardware placement turns that logical geometry into a phase-weighted routing spectrum, and architecture design becomes an inverse problem over physical path-volume geometry.

In slogan form,

$$\text{iQFT phase decay} \implies \text{logical phase geometry} \implies \text{physical routing spectrum.}$$

2 Distributed iQFT as a logical interaction graph

Let there be P QPUs, each holding Q logical qubits, so $n = PQ$. Logical node $p \in \{0, \dots, P-1\}$ holds the block

$$B_p = \{pQ, pQ+1, \dots, (p+1)Q-1\}.$$

The inverse QFT contains controlled phase rotations whose magnitudes decay exponentially with qubit separation. At the node level, after replacing many qubit-level interactions by a block-level abstraction, this motivates a horizon graph.

Model: horizon-D distributed iQFT graph

Let

$$G_D = (V, E_D), \qquad V = \{0, \dots, P-1\},$$

where

$$(p,q) \in E_D \quad \Longleftrightarrow \quad 0 < |p-q| \le D.$$

Exact distributed iQFT corresponds to $D = P - 1$. Approximate distributed iQFT corresponds to a finite communication horizon $D \ll P$.

This graph records which QPU blocks must communicate after pruning. It does not yet say how those communications move through hardware.

Intuitive Interpretation

The horizon D is a geometric shadow of a quantum approximation. Weak long-range phase corrections are removed; the surviving phase corrections form a banded graph on the QPU blocks. The circuit has become a weighted interaction geometry.

3 The phase-induced metric

The horizon graph itself comes from the phase law. Let $a_{pq} \ge 0$ denote effective retained phase magnitude, communication demand, or phase-weighted interaction strength between logical QPU blocks p and q. The idealized iQFT scaling is

$$a_{pq} = a_0 b^{-|p-q|}, \qquad b > 1.$$

For the usual binary phase hierarchy, one may take $b = 2$, up to convention-dependent constants.

Definition: phase-induced distance

For positive interaction weights a_{pq}, define

$$d_\phi(p,q) = -\frac{1}{\log b} \log\left(\frac{a_{pq}}{a_0}\right).$$

When $a_{pq} = a_0 b^{-|p-q|}$, this gives $d_\phi(p,q) = |p-q|$.

Proposition 1: exponential phase decay induces a line metric

Suppose $a_{pq} = a_0 b^{-|p-q|}$, with $b > 1$. Then

$$d_\phi(p,q) = |p-q|.$$

Consequently, thresholding weak iQFT interactions is equivalent to taking a ball graph in the induced metric:

$$a_{pq} \ge \tau \quad \Longleftrightarrow \quad d_\phi(p,q) \le R_\tau, \qquad R_\tau = -\frac{1}{\log b} \log\left(\frac{\tau}{a_0}\right).$$

Proof. Substituting the exponential form gives

$$d_\phi(p,q) = -\frac{1}{\log b} \log(b^{-|p-q|}) = |p-q|.$$

The threshold equivalence follows by applying the monotone logarithm to $a_0 b^{-|p-q|} \ge \tau$. □

A robust version is useful when constants or block aggregation distort the ideal law.

Proposition 2: approximate exponential decay gives a quasi-line geometry

Suppose there are constants $0 < c_1 \le c_2$ and $\alpha > 0$ such that

$$c_1 e^{-\alpha|p-q|} \le \frac{a_{pq}}{a_0} \le c_2 e^{-\alpha|p-q|}.$$

Define $d_\phi(p,q) = -(1/\alpha)\log(a_{pq}/a_0)$. Then

$$|p-q| - \frac{\log c_2}{\alpha} \le d_\phi(p,q) \le |p-q| - \frac{\log c_1}{\alpha}.$$

Thus the induced phase geometry differs from the block-line metric by only an additive constant determined by multiplicative distortion in the phase weights.

Proof. Take logarithms of the two-sided bound and multiply by $-1/\alpha$, reversing the inequalities. □

Intuitive Interpretation

The iQFT already contains a geometry. Large phase corrections are close. Small phase corrections are far. Taking a negative logarithm turns multiplicative phase decay into additive distance. The familiar banded communication pattern is a metric ball graph, not merely an engineering cutoff.

4 From scalar routing cost to a routing spectrum

Let the physical QPU interconnect be

$$G_{\text{phys}} = (V_{\text{phys}}, E_{\text{phys}}), \qquad |V_{\text{phys}}| = P.$$

A placement is a bijection $\sigma : V \to V_{\text{phys}}$. Let

$$d_{\text{phys}}(u,v) = \text{dist}_{G_{\text{phys}}}(u,v)$$

be physical graph distance, and attach nonnegative weights w_{pq} to retained logical edges. These weights may represent remote gates, EPR-pair demand, latency demand, or phase-weighted communication mass. The usual routing functional is

$$C_D(\sigma) = \sum_{(p,q)\in E_D} w_{pq}\, d_{\text{phys}}(\sigma(p),\sigma(q)).$$

This is a scalar summary. The richer object is the distribution whose first moment it computes.

Definition: phase-weighted routing spectrum

Let $W = \sum_{(p,q)\in E_D} w_{pq}$. For a placement σ, define

$$\mu_\sigma = \frac{1}{W} \sum_{(p,q)\in E_D} w_{pq}\, \delta_{d_{\text{phys}}(\sigma(p),\sigma(q))}.$$

This is the phase-weighted distribution of physical route lengths required by the retained distributed iQFT interactions.

Proposition 3: routing cost is the first spectral moment

For any placement σ,

$$\frac{C_D(\sigma)}{W} = \int r\, d\mu_\sigma(r), \qquad C_D(\sigma) = W\, \mathbb{E}_{\mu_\sigma}[r].$$

Proof. Substitute the definition of μ_σ into the integral and multiply by W. □

Intuitive Interpretation

The scalar routing cost is not discarded. It becomes the mean of a more informative distribution. The spectrum μ_σ says how much phase-weighted communication occurs at radius one, radius two, radius three, and so on. A good architecture should not only lower the mean; it should control the tail.

5 Support and tail bounds

The routing spectrum immediately converts geometric support information into cost information.

Proposition 4: support bound

If $\operatorname{supp}(\mu_\sigma) \subseteq [0, R]$, then

$$C_D(\sigma) \le WR.$$

Proof. Since $r \le R$ on the support of μ_σ,

$$C_D(\sigma) = W \int r \, d\mu_\sigma(r) \le W \int R \, d\mu_\sigma(r) = WR.$$

□

Proposition 5: tail-mass routing bound

Assume the physical graph has diameter Δ. If $\mu_\sigma([R, \Delta]) \le \beta$, then

$$C_D(\sigma) \le W\big((1-\beta)R + \beta\Delta\big) = W\big(R + \beta(\Delta - R)\big).$$

Proof. At least $1-\beta$ of the mass lies at radius at most R, and at most β lies beyond R. Since no radius exceeds Δ, the stated expectation bound follows. □

The long-distance tail is not merely a visual feature of a distribution. It directly bounds how much weighted communication cost can leak into expensive routes.

6 Hypercube placement as spectral compactification

Assume $P = 2^m$. Let the physical graph be the m-dimensional hypercube

$$H_m = \{0, 1\}^m,$$

where two vertices are adjacent if their binary labels differ in exactly one coordinate. The graph distance is Hamming distance, so

$$\operatorname{diam}(H_m) = m = \log_2 P.$$

Place logical blocks by Gray code,

$$\sigma(p) = \operatorname{Gray}(p) = p \oplus (p \gg 1),$$

so consecutive Gray-code words differ in exactly one bit.

Proposition 6: Gray-code hypercube embedding of the phase geometry

Let $P = 2^m$, let $G_{\text{phys}} = H_m$, and place logical nodes by $\sigma(p) = \operatorname{Gray}(p)$. For every pair p, q,

$$d_{H_m}(\sigma(p), \sigma(q)) \le \min\{|p-q|, \log_2 P\}.$$

If $d_\phi(p, q) = |p - q|$, then

$$d_{H_m}(\sigma(p), \sigma(q)) \le \min\{d_\phi(p, q), \log_2 P\}.$$

Consequently, for every retained horizon-D edge,

$$d_{H_m}(\sigma(p), \sigma(q)) \le \min\{D, \log_2 P\}.$$

Proof. Every pair of vertices in H_m is separated by at most $m = \log_2 P$, giving the diameter bound. Gray-code ordering gives $d_{H_m}(\sigma(r), \sigma(r+1)) = 1$. If $q > p$, then $\sigma(p), \sigma(p+1), \dots, \sigma(q)$ is a path of length $q - p$, so $d_{H_m}(\sigma(p), \sigma(q)) \le |p - q|$. Combining the two bounds proves the claim; retained edges satisfy $|p - q| \le D$. □

Corollary 7: hypercube support bound for the routing spectrum

Under the assumptions of Proposition 6,

$$\operatorname{supp}(\mu_\sigma) \subseteq [0, \min\{D, \log_2 P\}],$$

and therefore

$$C_D(\sigma) \le \min\{D, \log_2 P\} \sum_{(p,q)\in E_D} w_{pq}.$$

Proof. Proposition 6 places every retained edge at radius no larger than $\min\{D, \log_2 P\}$. Proposition 4 then gives the cost bound. □

Intuitive Interpretation

The hypercube is not merely a graph with small diameter. In this setting it compactifies the iQFT phase geometry. Consecutive logical blocks remain adjacent through Gray code, while the worst possible retained communication is forced into a logarithmic-radius ball.

7 Path-volume geometry for physical routes

The previous spectrum places one atom at the shortest physical distance for each retained logical interaction. The link to *Path-Length Distributions, Inverse Geometry, and Designed Neighborhood Spectra* is that shortest distance is only the first geometric observable: the richer object counts not only the shortest route, but the available routes around it.

For each retained logical edge $e = (p, q) \in E_D$, let $\Gamma_e = \Gamma_{\sigma(p),\sigma(q)}$ be a family of admissible physical routes from $\sigma(p)$ to $\sigma(q)$. Let $c_a \ge 0$ be an edge cost, and define route length

$$L_c(\gamma) = \sum_{a\in\gamma} c_a.$$

Let ν_e be a counting measure or weighted route measure on Γ_e. Define

$$F_e(\ell) = \nu_e\{\gamma \in \Gamma_e : L_c(\gamma) \le \ell\}, \qquad \rho_e = dF_e.$$

Definition: phase-weighted iQFT path-volume spectrum

The global phase-weighted route-volume spectrum is

$$\rho_{\text{iQFT}} = \frac{1}{W} \sum_{e=(p,q)\in E_D} w_e \rho_e,$$

or equivalently

$$F_{\text{iQFT}}(\ell) = \frac{1}{W} \sum_{e=(p,q)\in E_D} w_e F_e(\ell).$$

The shortest-route spectrum μ_σ is recovered by replacing each route family with a single atom at shortest distance:

$$\rho_e = \delta_{d_{\text{phys}}(\sigma(p),\sigma(q))}.$$

Thus the earlier theory is the degenerate shortest-path case of the path-volume theory.

Intuitive Interpretation

For a fixed retained iQFT interaction, hardware may offer many possible routes. Some are shortest, some are slightly longer, and some are long detours. The path-volume spectrum asks how much phase-weighted routing capacity exists at each route length.

8 Inverse design of a QPU path-volume geometry

The previous definitions suggest a genuine inverse problem, in the same spirit as reconstructing or designing geometry from prescribed path-length spectra. Instead of choosing a familiar physical graph and then analyzing its cost, choose a physical graph whose route-volume law matches the structure demanded by the distributed iQFT.

Let $\rho^\star$ be a desired route-volume spectrum: compact support for hard latency control, an exponential law matching iQFT phase decay, a shifted lognormal law for multiplicative route distortions, or an anchor-wise target spectrum around each QPU block. Then an inverse architecture design problem is

$$(G^\star_{\mathrm{phys}}, \sigma^\star, c^\star) = \arg\min_{G,\sigma,c} D(\rho_{\mathrm{iQFT}}, \rho^\star) + \lambda\Omega(G,\sigma,c) + \eta\mathcal{L}_{\mathrm{hardware}}(G,c).$$

Here D may be Wasserstein distance, binned squared error, KL divergence when densities are available, or another discrepancy. The term Ω prevents degeneracy, while $\mathcal{L}_{\mathrm{hardware}}$ may encode degree constraints, congestion, fabrication cost, latency, fidelity, or platform limits.

An anchor-wise version mirrors designed neighborhood spectra. For each logical block p, define

$$W_p = \sum_{q:(p,q)\in E_D} w_{pq}, \qquad \mu_{\sigma,p} = \frac{1}{W_p}\sum_{q:(p,q)\in E_D} w_{pq}\,\delta_{d_{\mathrm{phys}}(\sigma(p),\sigma(q))}.$$

Then optimize

$$\sigma^\star = \arg\min_\sigma \sum_p D(\mu_{\sigma,p}, \mu_p^\star) + \lambda C_D(\sigma) + \Omega(\sigma).$$

Proposition 8: spectrum matching controls mean routing cost

Let μ_σ be the shortest-route spectrum and let $\mu^\star$ be a target probability measure on route lengths. If

$$W_1(\mu_\sigma, \mu^\star) \le \varepsilon,$$

where W_1 is Wasserstein-1 distance, then

$$\left|\frac{C_D(\sigma)}{W} - \int r\,d\mu^\star(r)\right| \le \varepsilon.$$

In particular,

$$C_D(\sigma) \le W\left(\int r\,d\mu^\star(r) + \varepsilon\right).$$

Proof. The function $f(r) = r$ is 1-Lipschitz on the nonnegative real line. By the Kantorovich-Rubinstein dual representation of W_1,

$$\left|\int r\,d\mu_\sigma(r) - \int r\,d\mu^\star(r)\right| \le W_1(\mu_\sigma, \mu^\star) \le \varepsilon.$$

Using Proposition 3 gives the stated cost bound. □

Intuitive Interpretation

This is the cleanest bridge from inverse geometry back to engineering cost. If the learned or designed architecture matches a target routing spectrum in W_1, then its normalized communication cost is close to the target mean. Matching the spectrum is not just aesthetic; it controls the routing objective.

9 Sparsified routing fabrics and stretch

The inverse-geometry viewpoint also suggests a sparsification principle: remove edges according to a controlled spectrum, then require replacement paths to bound shortest-path distortion. The same logic applies here.

Let G_{phys} be a full physical routing fabric and let $H \subseteq G_{\text{phys}}$ be a sparse physical routing fabric. Let d_H denote shortest-path distance inside H.

Proposition 9: phase-weighted stretch bound

Suppose that for every retained distributed-iQFT edge $(p, q) \in E_D$,

$$d_H(\sigma(p), \sigma(q)) \le \tau d_{G_{\text{phys}}}(\sigma(p), \sigma(q))$$

for some $\tau \ge 1$. Then

$$C_D^H(\sigma) \le \tau C_D^{G_{\text{phys}}}(\sigma).$$

Proof. Multiply the pointwise stretch inequality by w_{pq} and sum over $(p, q) \in E_D$. □

A tail-controlled version is also useful. Suppose a sparse fabric preserves all retained interactions up to radius R, and the remaining long-distance phase mass is at most β. If the sparse fabric has worst-case replacement stretch τ on the long tail, then the total penalty is concentrated on a controlled portion of the phase mass. Distributional tail control says how much important structure is exposed to pruning; stretch control says how much error that pruning can introduce.

10 What has actually been added?

The routing note already gives a practical theorem: hypercube geometry with Gray-code placement bounds retained iQFT communication by $\min\{D, \log_2 P\}$. The companion inverse-geometry note, *Path-Length Distributions, Inverse Geometry, and Designed Neighborhood Spectra*, gives a broader principle: geometries can be represented, inferred, or designed through path-length and neighborhood spectra. The present extension adds four things.

1. **An intrinsic iQFT geometry.** Exponential phase decay induces a phase metric d_ϕ, and the horizon graph is a ball graph in that metric.
2. **A spectral view of routing.** The scalar cost C_D is the first moment of a phase-weighted routing spectrum μ_σ.
3. **A compactification interpretation.** The Gray-code hypercube result says the retained phase-weighted spectrum has bounded support inside $[0, \min\{D, \log_2 P\}]$.
4. **An inverse-design problem.** Physical QPU geometry can be optimized so that its route-volume spectrum matches the phase-weighted communication spectrum demanded by the iQFT.

11 Limitations and next experiments

The model abstracts away effects that a full architecture paper would need to include: edge congestion, routing conflicts, entanglement generation latency, purification overhead, correlated noise, and platform-specific degree constraints. A natural congestion functional is

$$C_D^{\text{cong}}(\sigma) = \sum_{a \in E_{\text{phys}}} \Psi\left(\sum_{(p,q) \in E_D} w_{pq} \mathbf{1}\{a \in \gamma_{pq}\}\right),$$

where γ_{pq} is the chosen route and Ψ penalizes edge load. The route-volume spectrum could then be paired with a congestion spectrum.

The best next experiment is probably computational: fix P, D, and w_{pq}; compute μ_σ for several candidate physical graphs, such as a line, ring, grid, hypercube, expander, or sparse learned fabric; then compare their spectra. The picture should make the architecture choice visible: where is the phase mass, how heavy is the tail, and how much of the iQFT is forced into expensive routes?

12 One-sentence summary

The distributed iQFT induces a phase-weighted logical geometry; a hardware placement converts that geometry into a routing spectrum; and architecture design becomes the inverse problem of choosing a physical QPU path-volume geometry whose short routes carry the important phase mass.

References

[1] Jonathan R. Landers. *Routing Costs for Distributed Inverse QFT Architectures*. Short extension note, May 2026.

[2] J. R. Landers. *Path-Length Distributions, Inverse Geometry, and Designed Neighborhood Spectra*. Companion note, April 2026.

[3] F. Javier Cardama, Jorge Vazquez-Perez, Tomas F. Pena, and Andres Gomez. *Communication-Efficient Distributed Inverse Quantum Fourier Transform*. arXiv:2605.10710, 2026. https://arxiv.org/abs/2605.10710

[4] Richard Cleve and John Watrous. *Fast Parallel Circuits for the Quantum Fourier Transform*. arXiv:quant-ph/0006004, 2000. Later published in *SIAM Journal on Computing*, 45(4):1570–1595, 2016.

[5] D. Coppersmith. *An Approximate Fourier Transform Useful in Quantum Factoring*. arXiv:quant-ph/0201067, 2002. Originally IBM Research Report RC19642, 1994.

[6] Peter W. Shor. *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*. arXiv:quant-ph/9508027, 1995.