Ordered Patch Theory

Appendix T-3: MERA Tensor Networks and the Informational Causal Cone

Anders Jarevåg

April 5, 2026 | DOI: 10.5281/zenodo.19300777

Original Task T-3: MERA Tensor Networks and the Causal Cone Problem: OPT proposes an Informational Causal Cone composed of sequential compression, but relies on a bespoke geometric description rather than standard quantum tensor formalisms. **Deliverable:** Formal mapping of OPT's Informational Causal Cone to the MERA tensor network structure.

Closure status: CONDITIONAL ISOMORPHISM (structural homomorphism confirmed; strict physical isomorphism conditionally upgraded via P-2). This appendix delivers the target structural mapping required by T-3. Three theorems establish a strong topological analogy: (T-3a) the OPT Stability Filter's iterative coarse-graining is structurally homomorphic to a MERA tensor network; (T-3b) the Informational Causal Cone of §3.3 corresponds in order-of-magnitude to the MERA causal cone; and (T-3c) the Forward Fan structurally maps to the un-renormalized boundary degrees of freedom. Mathematically elevating this purely stochastic structural homomorphism into the strict Hilbert-space isometries required for a true discrete Ryu-Takayanagi bound originally remained open, but is now conditionally resolved via the explicit computational basis embedding and Isometry Identification bridge postulates established sequentially in problem P-2.

---

§1. The Multi-Layer Compression Structure

The preprint's §3.3 defines the OPT observer by a single bottleneck optimisation (Eq. 4): a compressed state $Z_t \in \{1, \ldots, 2^B\}$ is selected from the full boundary state X_t to maximise predictive information at minimum description length. What §3.3 does not make explicit is that the path from $X_{\partial A}$ to Z_t naturally decomposes into a *cascade* of compression layers — each one discarding short-range correlations irrelevant to prediction at the next scale. This hierarchical structure is the OPT side of the MERA correspondence.

1.1 The L-Layer Bottleneck Cascade

Let $s \geq 2$ be a fixed coarse-graining factor and L the total number of compression layers. Define the cascade:

$$Z_t^{(0)} := X_{\partial_R A} \qquad \text{(layer 0: full Markov boundary, } H = B_0 \text{ bits)}$$

At each subsequent layer $\tau = 0, \ldots, L-1$:

$$Z_t^{(\tau+1)} = \arg\min_q \Big[I\Big(Z_t^{(\tau)} \,;\, Z_t^{(\tau+1)}\Big) - \beta_\tau \, I\Big(Z_t^{(\tau+1)} \,;\, X_{t+1:\infty}\Big)\Big]$$

$$\text{subject to: } I\Big(Z_t^{(\tau)} \,;\, Z_t^{(\tau+1)}\Big) \leq B_\tau, \qquad B_\tau = B_0 \cdot s^{-\tau}$$

The final state is $Z_t := Z_t^{(L)}$, with $B_L = B_0 \cdot s^{-L}$ bits. The cascade defines a Markov chain:

$$X_{t+1:\infty} \;\text{---}\; Z_t^{(0)} \;\text{---}\; Z_t^{(1)} \;\text{---}\; \cdots \;\text{---}\; Z_t^{(L)} = Z_t$$

By the data processing inequality, predictive information is monotone non-increasing:

$$I\Big(Z_t^{(\tau)} \,;\, X_{t+1:\infty}\Big) \geq I\Big(Z_t^{(\tau+1)} \,;\, X_{t+1:\infty}\Big)$$

Each layer loses a controlled quantity of predictive information — controlled by the distortion budget D_τ of that layer's bottleneck.

1.2 Decomposition into Disentangle-then-Coarsen

Each layer transition $Z^{(\tau)} \to Z^{(\tau+1)}$ decomposes into two canonical steps:

- **Disentanglement:** Apply a local reversible rearrangement modeled as a permutation mapping $U_\tau \in S_{|\mathcal{Z}|}$ to $Z^{(\tau)}$ that brings mutually irrelevant branches of the Forward Fan — branches sharing no predictive information about the future — into adjacent positions. This classical step is reversible; no information is lost.
- **Coarse-graining (bottleneck mapping):** Partition the states into groups of s and apply the classical stochastic bottleneck compression map $W_\tau : \mathcal{Z}^s \to \Delta(\mathcal{Z})$ at each group. The bond dimension is fixed as $\chi = 2^{B_0/N}$, where N is the number of boundary sites. To function formally as an exact discrete Hilbert-space tensor dimension rather than an effective continuous scale, the framework strictly mandates the diophantine constraint $2^{B_0/N} \in \mathbb{Z}^+$. This explicitly ensures the exact integer dimension χ yields a per-site entropy $\log \chi = B_0/N$ consistent geometrically with the capacity schedule $B_\tau = B_0 \cdot s^{-\tau}$. **Note:** The quantum target structures used in §2

are the MERA isometry $w_\tau : \mathbb{C}^\chi \to (\mathbb{C}^\chi)^{\otimes s}$ (whose adjoint $w_\tau^\dagger$ implements coarse-graining) and disentangler u_τ. The §1 maps $W_\tau : \mathcal{Z}^s \to \Delta(\mathcal{Z})$ and $U_\tau \in S_{|\mathcal{Z}|}$ are the classical OPT objects. The embedding that connects them is established in Appendix P-2.

The composition $W_\tau \circ U_\tau$ at each layer, stacked for $\tau = 0, \dots, L-1$, constitutes the full tensor network. We now show this is precisely MERA.

---

§2. MERA — Formal Definitions

We state the relevant definitions from Vidal (2008) [43] in a form suited to the OPT mapping.

2.1 Tensors

A MERA for a 1D chain of N boundary sites with local Hilbert space $\mathbb{C}^\chi$ consists of L layers. Each layer τ contains two classes of tensor:

- **Disentanglers u_τ:** unitary tensors $u_\tau : (\mathbb{C}^\chi)^{\otimes 2} \to (\mathbb{C}^\chi)^{\otimes 2}$ acting on adjacent pairs of sites. They remove short-range entanglement without changing the total Hilbert space dimension. Unitarity: $u^\dagger u = uu^\dagger = I$.
- **Isometries w_τ:** tensors $w_\tau : \mathbb{C}^\chi \to (\mathbb{C}^\chi)^{\otimes s}$ satisfying $w_\tau^\dagger w_\tau = I_{\mathbb{C}^\chi}$ (isometric: the map is an injection from the coarse-grained space into the fine-grained space). The adjoint $w_\tau^\dagger : (\mathbb{C}^\chi)^{\otimes s} \to \mathbb{C}^\chi$ implements the coarse-graining, mapping s fine-grained sites to 1 coarse site.

The full MERA maps the top state $|\psi_{\text{top}}\rangle \in \mathbb{C}^\chi$ (the bulk) to the boundary state $|\psi_{\text{boundary}}\rangle \in (\mathbb{C}^\chi)^{\otimes N}$ by applying the layers from bulk to boundary, each layer expanding the state space by factor s.

2.2 The MERA Causal Cone

The **causal cone** $\mathcal{C}(x)$ of a boundary site $x \in \{1, \dots, N\}$ is the minimal set of tensors in the network whose values can affect the reduced density matrix ρ_x of site x. It is computed bottom-up (from bulk toward boundary).

At the bulk layer (depth $\tau = L$ from boundary): $\mathcal{C}(x)$ contains the single top tensor. At each subsequent layer going toward the boundary, the causal cone expands by factor s at each isometry layer and by at most 2 at each disentangler layer. The width of $\mathcal{C}(x)$ at boundary depth τ from the top is:

$$w(\tau) = \mathcal{O}(s^\tau) \qquad \text{[grows exponentially from bulk toward boundary]}$$

For the critical MERA ($s = 2$), the causal cone width grows as 2^τ at depth τ, and after L layers reaches the full boundary width $N = s^L$.

2.3 Entanglement Entropy and the Minimal Cut

For a contiguous boundary region A of length $|A| = l$, the entanglement entropy $S(A)$ in a MERA state is bounded by the number of bonds cut by the minimal surface γ_A through the bulk of the tensor network:

$$S(A) \leq |\gamma_A| \cdot \log \chi$$

where $|\gamma_A|$ is the number of bonds in the minimal cut and χ is the bond dimension. For a scale-invariant MERA, $|\gamma_A| \sim \frac{c}{3} \log l$, recovering the CFT entanglement entropy $S(A) \sim \frac{c}{3} \log l$ with $c/3 = \log \chi$. This is the discrete analogue of the Ryu-Takayanagi formula in AdS/CFT.

---

§3. Theorem T-3a — Structural Homomorphism

Theorem T-3a (MERA–OPT Homomorphism). The OPT L-layer Information Bottleneck cascade $\{Z_t^{(\tau)}, \tau = 0, \ldots, L\}$ with boundary state $Z_t^{(0)} = X_{\partial_R A}$, bulk state $Z_t^{(L)} = Z_t$, layer capacity $B_\tau = B_0 \cdot s^{-\tau}$, and bond dimension $\chi = 2^{B_0/N}$, is structurally homomorphic to the layer topology of a MERA with L layers, scale factor s, and bond dimension χ, under the formal classical mapping: - (i) OPT coarse-graining $W_\tau \leftrightarrow$ MERA isometry adjoint $w_\tau^\dagger$ - (ii) OPT disentangler $U_\tau \leftrightarrow$ MERA disentangler u_τ

3.1 Proof — Isometry Identification

The OPT coarse-graining tensor at layer τ computes via the conditional distribution $q^*(z^{(\tau+1)} \mid z^{(\tau)})$ produced by the bottleneck optimization. While the overall information budget enforces an average macroscopic capacity ratio of $B_\tau / B_{\tau+1} = s$, the classical stochastic bottleneck does not natively force exact uniform fiber cardinality (a strict discrete preimage equivalently matching size s for every $z^{(\tau+1)}$ output). Formalizing this explicit step therefore restricts the architecture to the idealized tight mapping limit ($D \to 0$), conditionally assuming the parameters perfectly isolate uniform information structures.

However, q^* represents a classical stochastic probability matrix, not a complex quantum unitary matrix. Claiming the true Hilbert space isometry condition ($W_\tau W_\tau^\dagger = I_{\mathbb{C}^\chi}$) would constitute a category error. A true partial isometry requires an explicit embedding of these discrete states into a computational basis on $\mathbb{C}^\chi$. **Appendix P-2** (Conditional Quantum Correspondence) establishes this embedding: Theorem P-2.0 provides the computational basis identification, and Theorem P-2c proves that the optimal bottleneck map in the tight limit acts as a partial isometry within the QECC-protected subspace. Conditional on the local noise model of P-2, the structural homomorphism upgrades to a genuine tensor-network isomorphism within the code space. ■

3.2 Proof — Disentangler Identification

The purely classical disentangler U_τ is established as a local bijection (a state alphabet permutation from the symmetric group $S_{|\mathcal{Z}|}$) that rearranges $Z^{(\tau)}$ to minimize inter-group redundancies (identically: mutual information) before they are coarse-grained.

$$U_\tau = \arg\min_{U\in S_{|\mathcal{Z}|}} \sum_{j\neq k} I\Big(U(Z^{(\tau)})_{\text{group}\,j}\,;\,U(Z^{(\tau)})_{\text{group}\,k}\Big)$$

This matches the structural objective of the MERA disentangler: removing short-range entanglement (correlations between adjacent groups) before coarse-graining. True complex unitarity ($U^\dagger U = I$) is established by Theorem P-2.0 (Appendix P-2): under the computational basis embedding, the permutation $U_\tau \in S_{|\mathcal{Z}|}$ lifts uniquely to a unitary matrix in $U(\mathbb{C}^\chi)$ via the permutation representation.

Caveat (Permutation vs. General Unitary). Theorem P-2.0 lifts OPT's disentanglers into the *permutation subgroup* of $U(\mathbb{C}^\chi)$, not the full unitary group. Standard MERA disentanglers are general unitaries $u_\tau \in U(\mathbb{C}^\chi \otimes \mathbb{C}^\chi)$; the permutation subgroup is a strict subset ($|S_\chi| = \chi!$ vs. $\dim U(\chi) = \chi^2$ continuous parameters). The isomorphism established by P-2.0+P-2c is therefore to *permutation MERA* — a restricted sub-class. Extending to full MERA would require identifying an OPT-native mechanism that generates general unitaries rather than permutations. This gap does not affect the RT entropy bound (P-2d), which depends only on the isometry condition P-2c, not on the disentangler class. ■

MERA–OPT Isomorphism Dictionary

MERA component	OPT counterpart	Formal OPT definition
Boundary layer (UV)	Markov boundary $X_{\partial_R A}$	Full physical substrate states; $H = B_0$ bits (§3.4 preprint)
Bulk layer (IR)	Compressed state Z_t	Optimal bottleneck output; $H = B_L$ bits (preprint Eq. 4)
Isometry adjoint $w_\tau^\dagger$	Coarse-graining $W_\tau(\Delta\mathcal{Z})$	Classical stochastic bottleneck map at layer τ; reduces capacity $B_\tau \to B_{\tau+1}$
Disentangler $u_\tau(U(\mathbb{C}^\chi))$	Branch disentangler $U_\tau(S_{|\mathcal{Z}|})$	Classical permutation removing inter-group correlations before coarse-graining

MERA component	OPT counterpart	Formal OPT definition
Bond dimension χ	$\chi = 2^{B_0/N}$	Per-site channel capacity; $\log \chi = B_0/N$ bits per site, consistent with geometric schedule $B_\tau = B_0 s^{-\tau}$ (see §1.1).
Scale factor s	Coarse-graining ratio s	Compression factor per layer; $B_{\tau+1} = B_\tau / s$
Number of layers L	Compression depth L	$L = \log_s(B_0/B_L)$; depth of Stability Filter hierarchy
Top tensor	Present aperture Z_t	The $C_{\max}$ bottleneck; the NOW of the Informational Causal Cone

---

§4. Theorem T-3b — Causal Cone Identity

Theorem T-3b (Causal Cone Correspondence). Under the homomorphism of T-3a, the Informational Causal Cone of OPT (preprint §3.3) corresponds structurally (in order-of-magnitude scaling) to the MERA causal cone. The present aperture Z_t maps to the bulk top tensor; the settled Causal Record $\mathcal{R}_t$ corresponds to the past bulk states; the Forward Fan $\mathcal{F}_h(z_t)$ corresponds to the un-renormalized degrees of freedom at the MERA boundary h layers from the present.

4.1 Direction of the Correspondence

There is a subtlety of orientation that must be stated precisely. In MERA, the network runs from boundary (UV, fine-grained) to bulk (IR, coarse-grained). In OPT, the Informational Causal Cone runs from past (settled, compressed) through the present aperture to the future (Forward Fan, unresolved). The correspondence is:

MERA direction	OPT direction	Interpretation
Boundary → Bulk (UV→IR)	Substrate → Present Z_t	Compressing fine-grained boundary into the compressed causal state
Bulk → Boundary (IR→UV)	Present Z_t → Forward Fan	Expanding from the aperture into un-renormalized future branches

MERA direction	OPT direction	Interpretation
Causal cone of bulk point	Forward Fan $\mathcal{F}_h(z_t)$	Boundary states reachable from bulk point; width $\sim s^h$

4.2 Proof — Causal Cone Width = Forward Fan Capacity

In the MERA, the causal cone of the bulk state Z_t (at depth L from the boundary) expands as it moves toward the boundary: at depth τ layers from the top, the cone has width s^τ. This counts the number of boundary sites that can independently influence Z_t.

In OPT, the Forward Fan $\mathcal{F}_h(z_t)$ at depth h time-steps from the present aperture contains at most $2^{B\cdot h}$ distinguishable future states (preprint Eq. 5: $\log|\mathcal{F}_h| \leq Bh$). MERA layer depth corresponds to $\tau = h$. We observe an exponential vs linear bounding mismatch ($s^\tau \cdot B/L$ bits in MERA via scale expansion vs $B\tau$ in the Forward Fan via chronological accretion). The causal cone width and the OPT Forward Fan capacity agree robustly in order of magnitude, but find strict exact agreement only in the limit of a single-layer codec ($L = 1$). Furthermore, identifying the passive topology of MERA with the action-dependent Forward Fan implies we are operating exclusively within the passive observer limit ($a \equiv \text{const}$). ■

4.3 Proof — Causal Record = Past Bulk

The settled Causal Record $\mathcal{R}_t = (Z_0, Z_1, \ldots, Z_t)$ (preprint §3.3) consists of all past compressed states — the bulk states that have already been rendered into the settled past. In the MERA, these correspond to the sequence of past bulk states connected by the codec's temporal dynamics K_θ (preprint Eq. 6). The settled, low-entropy character of $\mathcal{R}_t$ corresponds to the fact that bulk states in MERA have low entanglement entropy by construction — they are the coarse-grained result of the disentangling procedure. ■

---

§5. Theorem T-3c — The Forward Fan as Boundary UV and the Discrete Ryu-Takayanagi Formula

Theorem T-3c (Forward Fan = Boundary UV; Discrete RT).

(a) The Forward Fan $\mathcal{F}_h(z_t)$ maps probabilistically to the set of un-renormalized degrees of freedom at the MERA boundary — the boundary UV layer of the MERA applied to the codec at time step $t + h$.

(b) Classical Data Processing Limit (Bulk Cut Bound): The predictive cut entropy evaluated correctly at the internal bulk minimal cut layer satisfies explicitly:

$$S_{\text{cut}}^{(\tau^*)}(A) \leq |\gamma_A| \log \chi$$

(c) Discrete Quantum RT Extension (Conditional on P-2d embedding):

$$\boxed{S_{\mathrm{vN}}(\rho_A) \leq |\gamma_A| \log \chi}$$

where γ_A is the minimal-cut surface in the MERA bulk and $\chi = 2^{B_0/N}$ is the bond dimension. This bound holds conditional on the P-2d isometry; it reduces to the classical bulk-cut bound of Part (b) when the quantum structure is unavailable.

5.1 Proof — Forward Fan as Boundary UV

The MERA boundary UV layer at time $t + h$ consists of all possible input states $X^{(t+h)}_{\partial_R A}$ — the fine-grained, un-coarse-grained boundary states that will be processed by the codec over the next h time-steps. By the cascade structure, these are exactly the states reachable from the present aperture $Z_t = Z_t^{(L)}$ by running the MERA in reverse (from bulk toward boundary) for h layers — i.e., by expanding the causal cone of Z_t for h steps.

The Forward Fan $\mathcal{F}_h(z_t)$ is defined in the preprint (§3.3) as:

$$\mathcal{F}_h(z_t) := \{z_{t+1:t+h} : p(z_{t+1:t+h} \mid z_t,\, a_{t:t+h-1}) > 0\}$$

These are precisely the sequences of bulk states reachable from Z_t within h MERA layers by operating the cascade probabilistically in the expanded direction. The identification requires the MERA be evaluated in both directions — boundary → bulk (past compression) and bulk → boundary (future expansion). The Forward Fan corresponds explicitly to the second direction, which is the exact support set of the causal cone expansion of the bulk state toward the boundary UV, under the time-reversal identification properly noted in §4.1. ■

5.2 Proof — Discrete Ryu-Takayanagi Mapped Bound

Let A and $\bar{A} = V \setminus A$ be a bipartition of the boundary. Let τ^* be the minimal layer at which the $A/\bar{A}$ interface is exactly severed in the tensor network (the minimal cut layer). At this layer, the local mutual information bottleneck capacity is strictly clamped by the capacity of those severed bonds:

$$S^{(\tau^*)}_{\mathrm{cut}}(A) \leq |\gamma_A| \cdot \log \chi \qquad \text{(Inter-group bulk bound)}$$

While this successfully establishes the discrete Ryu-Takayanagi capacity bound exactly at the bulk minimal cut layer, formally pushing this bound upward to limit the exterior boundary predictive cut entropy $S_{\mathrm{cut}}(A) = S^{(0)}_{\mathrm{cut}}(A)$ cannot be accomplished using the Data Processing Inequality (as the DPI mandates that entropy must monotonically decrease, not increase, as we compress downward: $S^{(0)}_{\mathrm{cut}} \geq S^{(\tau^*)}_{\mathrm{cut}}$).

The correct path to the full target discrete RT boundary bound ($S_{\mathrm{vN}}(\rho_A) \leq |\gamma_A| \log \chi$) requires bounding the **Schmidt rank** across the bipartition — a strategy that requires treating the network as constructing the boundary state via true linear isometries. This is now established in **Appendix P-2**: Theorem P-2d proves the discrete quantum Ryu-Takayanagi formula via the Schmidt decomposition of the MERA state across the minimal cut, conditional on the isometry condition of P-2c. ■ (conditional on P-2d isometry).

---

§6. The Epistemic Ladder — From Classical to Quantum RT

The three theorems above establish the MERA structure at the classical information-theoretic level. The Epistemic Ladder of §3.4 of the preprint describes the conditions under which each rung can be climbed.

Rung	Entropy law	Condition	Status
1. Classical Area Law	$S_{\mathrm{cut}} \leq \lvert\partial A\rvert \log q$	Locality + Markov screening (§3.4 preprint)	Proven (preprint Eq. 8)
2a. Classical bulk-cut	$S_{\mathrm{cut}}^{(\tau^*)} \leq \lvert\gamma_A\rvert \log \chi$	T-3a cascade + classical DPI	Proven (T-3c Part b)
2b. Discrete quantum RT	$S_{\mathrm{vN}}(\rho_A) \leq \lvert\gamma_A\rvert \log \chi$	2a + P-2 isometry embedding	Proven (P-2d, conditional)
3. Quantum RT	$S(A) = \frac{\mathrm{Area}(\gamma_A)}{4G_{\mathrm{OPT}}} + S_{\mathrm{bulk}}$	Rung 2b + continuum limit	Conditional on continuum limit
4. Full AdS/CFT	Exact bulk/boundary duality	Quantum RT + geometric reconstruction of bulk operators	Long-term (v3.0+)

The quantum RT formula requires replacing the classical predictive cut entropy $I(X_A; X_{V\setminus A})$ with the von Neumann entanglement entropy $S_{\mathrm{vN}}(\rho_A)$ of a density matrix ρ_A. This presupposes a Hilbert space structure for the state space of Z_t. The derivation of this structure — via the ADH quantum error correction argument (preprint P-2) — remains the next formal step. Once P-2 is closed, the bond dimension $\chi = 2^{B_0/N}$ becomes a quantum bond dimension, and the classical mutual information in the proof of T-3c is replaced by quantum mutual information, recovering the full quantum RT formula with the bulk correction term S_{bulk}.

---

§7. Emergent Bulk Geometry from Code Distance

The MERA bulk geometry is not a pre-existing container. Under the isomorphism of T-3a, it is the informational metric space of the codec: the geometry of compression distances.

7.1 Code Distance as Bulk Metric

Define the discrete integer **code distance** $d(z^{(\tau)}, z'^{(\tau)})$ between two states at layer τ of the cascade as the minimum number of disentangler-swaps required to connect them within the tensor network.

Under a proper thermodynamic or continuum limit ($N \to \infty, a \to 0$), one might approximate the bulk metric $g_{ij}^{\text{bulk}}(\tau)$ at continuous spatial layer scale τ as:

$$g_{ij}^{\text{bulk}}(\tau) \propto \lim_{a \to 0} \frac{d\Big(z_i^{(\tau)},\, z_j^{(\tau)}\Big)^2}{d\Big(z_i^{(0)},\, z_j^{(0)}\Big)^2}$$

This is a structural expectation, conditional on scale invariance of the cascade and the assumption that Permutation MERA is continuously approximable by a general MERA in the continuum limit — consistent with the known results of Swingle (2012) and Nozaki-Ryu-Takayanagi (2012), but not guaranteed for a discrete cascade with finitely many layers. Thus, under these continuum-limit conjectures, we expect that spacetime geometry would curve precisely where code distance diverges — i.e., where the predictive rate R_{req} approaches $C_{\max}$, consistent strategically with T-2's rate-distortion overflow identification.

7.2 Connection to T-2

T-2 established that gravitational curvature $G_{\mu\nu}$ is the metric derivative of rendering entropy S_{render}. The MERA structure now specifies the microscopic origin of S_{render}: it is the minimal-cut entropy $|\gamma_A| \log \chi$, and the Einstein tensor $G_{\mu\nu}$ is the response of this cut entropy to metric perturbations in the bulk geometry induced by code distance. The two appendices are therefore consistent: T-2 gives the macroscopic field equations; T-3 gives the microscopic tensor-network origin of the entropy functional they extremise.

---

§8. Closure Summary and Open Edges

T-3 Deliverables — Partially Resolved → Conditionally Upgraded (with P-2)

1. **T-3a (MERA isomorphism).** The OPT L-layer bottleneck cascade is structurally homomorphic to a MERA with layer factor s and depth L. With Appendix P-2 (Theorems P-2.0 and P-2c), this upgrades to a tensor-network isomorphism within the QECC-protected subspace, conditional on local noise. **Note:** The isomorphism is to *permutation MERA* (disentanglers in the permutation subgroup of $U(\mathbb{C}^\chi)$), not to general MERA with arbitrary

unitary disentanglers. This restriction does not affect the RT bound (P-2d) but limits the correspondence to a sub-class of MERA networks.

2. **T-3b (Causal cone correspondence).** The Informational Causal Cone scales with order-of-magnitude symmetry to the MERA causal cone structure within the passive-observer limit, though depth profiles differ. The Forward Fan corresponds to un-renormalized boundary data. (P-2's isometry result applies within the passive observer limit; the action-dependent $a_{t:t+h-1}$ terms in the Forward Fan definition require an open-systems extension not addressed by P-2.)
3. **T-3c (Discrete quantum RT).** The original DPI-based proof bounded the bulk but not the boundary entropy. With the isometry from P-2c, Theorem P-2d establishes the full boundary bound $S_{\mathrm{vN}}(\rho_A) \leq |\gamma_A| \log \chi$ via the Schmidt rank of the MERA state.
4. **Emergent bulk geometry.** The MERA bulk metric g_{ij}^{bulk} is induced from code distance in the cascade. Spacetime curves where code distance diverges, consistent with T-2's identification of $G_{\mu\nu}$ as the metric derivative of rendering entropy. (Continuum limit still needed.)
5. **Epistemic Ladder status.** Rung 2 (discrete quantum RT) is now proven via P-2d. Rung 3 (full quantum RT with bulk correction) requires a continuum limit not yet derived from OPT primitives.

Open edges enabled by this closure

- **P-2 (Born Rule / Hilbert Space)** now has its exact entry point: the bond dimension χ must be embedded as a quantum Hilbert space dimension. Once ADH error correction forces the logical qubit structure, the classical bond $\chi = 2^{B_0/N}$ upgrades to a quantum bond with von Neumann entropy, and the discrete RT of T-3c becomes the full quantum RT with bulk correction S_{bulk}.
- **P-3 (Asymmetric Holography):** the MERA bulk reconstruction and Fano's inequality now have a shared formal home. Fano's inequality (preprint §3.10) bounds the observer's ability to reconstruct the substrate from within the render — precisely the irreversibility of the MERA map (boundary → bulk is the codec; bulk → boundary inversion is impossible past the minimal cut depth τ^*).
- **T-5 (Constants Recovery):** the bond dimension $\chi = 2^{B_0/N}$ and coarse-graining factor s provide new constraints on the dimensionless constants. In particular, $s = 2$ and $L = \log_s(B_0/B_L)$ must be consistent with the Planck-scale identification $l_{\mathrm{codec}} = l_P$ from T-2, constraining the ratio B_0/B_L.
- **§8.3 preprint item 3 (MERA/causal set):** formally mapping the MERA boundary layers of the Forward Fan to the causal set framework to extract metric properties of perceived spacetime purely from codec

sequencing. The code-distance metric g_{ij}^{bulk} of §7 is the starting point.

This appendix is maintained as part of the OPT project repository alongside theoretical_roadmap.pdf. References: Vidal (2008) [43], Pastawski et al. (2015) [44], Almheiri-Dong-Harlow (2015) [42], Tishby et al. (1999) [28], Ryu-Takayanagi (2006).