

Boundary-Crossing Paths as a Local, Computable Bound on Network Emergence

Boundary-Crossing Paths as a Local, Computable Bound on Network Emergence

Abstract

Emergent features of a networked system—behaviors visible in the coupled whole but absent from any isolated part—are usually measured globally, which obscures where in the network they originate. Building on the recent proposal that every emergent feature of a network system is produced by a route crossing the interface between parts, and that the number of such boundary-crossing paths bounds the total emergence discrepancy [1], [2], we develop a fully explicit walk-counting formalism for this claim. We model observations built route by route as walk-counting functionals W_L on the system's graph, define the emergence discrepancy D_L as the excess of whole-system walk mass over the sum of part-level walk masses, and prove that D_L equals exactly the number of walks that touch the inter-part boundary at least once. We verify the identity by exhaustive enumeration on a four-node reference network, computing per-length discrepancies $D^{(1)} = 4$, $D^{(2)} = 12$, $D^{(3)} = 28$, and cumulative $D_{\leq 3} = 44$, matching the closed-form projection $D_{\leq L} = 2^{L+3} - 8 - 4L$. Because boundary-touching walks of length L depend only on the $2(L - 1)$ -hop neighborhood of the cut, the bound is computable locally. We discuss applications to fault detectability, scalability assessment, epidemic containment, and engineered infrastructures, and we analyze failure modes of the walk-counting idealization, including cycles, weights, and observations that are not route-accumulating.

1. Introduction

Complex systems science has long treated emergence—features of a whole that no part exhibits alone—as a defining phenomenon, yet emergence is rarely quantified in a way that locates its source within the system's structure. A recent arXiv preprint addresses this gap directly: it evaluates emergence as what observing the coupled whole reveals beyond, or erases from, the combined observations of the parts, and shows that for observations built route by route, every emergent feature is produced by its own route across the interface between parts, with the total number of boundary-crossing paths bounding the emergence discrepancy [1], [2]. The preprint validates the bound on network examples and applies it to a nervous system, where the routes attribute most emergence to interneurons.

This paper takes that result as a starting point and develops it into a transparent, fully audited formalism. Our contribution is threefold. First, we give a self-contained construction in which "observations that build what they see route by route" are modeled as walk-counting functionals, and we prove an exact identity: the emergence discrepancy equals the count of boundary-touching walks. Second, we verify the identity by hand enumeration on a small reference network, showing every input number and every arithmetic step, and we derive a closed-form projection for the growth of the discrepancy with path length. Third, we connect the framework to adjacent literatures—topology-failure detectability [3], scalability of networked systems [4], epidemic protection [5], social-network feature extraction [7], virtual network embedding [8], and hierarchical power-flow control [9]—arguing that boundary-localized emergence accounting supplies a common diagnostic layer for problems that are otherwise treated separately.

The practical significance is engineering-oriented. If the total emergence a coupling can generate is bounded by a quantity computable from the neighborhood of the interface alone, then interface design becomes a lever on emergence: one can anticipate, attribute, and bound the emergent behavior of a coupling before simulating or

building it [1], [2].

2. Background and Related Work

The central reference for this work is the preprint "Emergence in Network Systems is Bounded by Boundary-Crossing Paths" [1], [2], which defines emergence as a discrepancy between whole-system observation and the combination of part-level observations, proves that for route-accumulating observations every emergent feature traces to a boundary-crossing route, and shows that the number of such routes bounds the total discrepancy. Our paper reconstructs the quantitative core of that claim in a walk-counting setting where the bound can be verified by exhaustive enumeration, and we extend it with a closed-form growth law.

Several literatures converge on the question of what subsystem observations reveal about a coupled whole. In networked control theory, the work on generic detectability and isolability of topology failures [3] asks whether link and node failures in networked linear systems can be detected and isolated from subsystem measurements when interaction weights are unknown, proving genericity results over topologies. This is complementary to emergence accounting: detectability [3] concerns what part-level measurements can and cannot see about structure, while the boundary-crossing-path bound [1], [2] quantifies exactly the informational surplus created by coupling. A failure that is not isolable from subsystem measurements [3] is precisely a candidate for being an emergent artifact of the interface in our sense.

Scalability analysis of nonlinear networked systems [4] formalizes when a network's desired configuration can be assessed from properties of its parts and local interactions, giving sufficient conditions robust to delays and disturbances. The boundary-crossing bound offers a structural statistic—route density across cuts—that could serve as a computable surrogate in such scalability criteria [4], since emergence generated at interfaces is a natural obstruction to part-wise verification.

Decentralized protection against SIS epidemics [5] studies optimal self-protection when contamination spreads over network links; the epidemic threshold and infection dynamics are emergent properties of the contact graph that no single node's protection budget determines. Boundary-crossing path counts across the cut between a protected cluster and its environment directly measure the routes available for emergent spread, linking containment strategy [5] to route-based emergence accounting [1], [2].

In the metascience and infrastructure domain, the ILC Technology Network status report [6] documents internationally distributed engineering work packages for a large-scale accelerator; such a program is a socio-technical network whose emergent integration risk lives at the interfaces between laboratories, a setting where route-counting audits of collaboration and supply interfaces are a natural application, though [6] itself reports engineering status rather than network analysis.

On the data-mining side, label-dependent feature extraction in social networks [7] builds node features by combining network structure with class labels for within-network classification, demonstrating that relational routes (neighbors' labels) carry class-relevant information absent from node-internal attributes—an empirical instance of "the whole reveals what parts alone do not," which our framework would attribute to specific boundary-crossing paths between label communities [7].

In networked services, constrained shortest path management for virtual network embedding [8] must compose services spanning multiple data centers, where the composed service's quality is an emergent property of the embedding routes rather than of any single substrate; the constrained-path machinery of [8] is exactly route enumeration across administrative boundaries, suggesting that emergence discrepancy could be reported alongside embedding cost as a coupling-risk metric [8].

Finally, hierarchical power flow control in smart grids [9] exploits demand-side flexibility across control layers to preserve rotor-angle and frequency stability; stability here is emergent at the interconnection of layers, and the hierarchical interfaces of [9] are precisely the cuts across which our boundary-touching walk counts would be evaluated. The QNFO framework for network isomorphism [10] supplies graph-matching machinery relevant to identifying when two interface neighborhoods are structurally equivalent, which matters for transferring emergence bounds between systems: if two cuts are isomorphic in the QNFO sense [10], their boundary-crossing path counts coincide.

Across these works, the unifying observation is that interface structure, not part inventory, governs the surplus behavior of coupled systems; the present paper gives that observation an audited arithmetic.

3. Methods

System and cut. Let $G = (V, E)$ be the graph of a network system, with vertex set V and edge set E . A bipartition $V = V_A \cup V_B$, $V_A \cap V_B = \emptyset$, defines two parts. The boundary (cut) is $\partial = \{\{u, v\} \in E : u \in V_A, v \in V_B\}$. The part subgraphs are $G_A = (V_A, E_A)$ and $G_B = (V_B, E_B)$, where E_A, E_B are the internal edges.

Route-accumulating observations. An observation is route-accumulating if its value is a sum over walks (routes) in the graph. For a graph H and horizon L , define the walk-count functional

$$W_L(H) = \sum_{k=1}^L \mathbf{1}^\top A_H^k \mathbf{1},$$

where A_H is the adjacency matrix of H and $\mathbf{1}$ is the all-ones vector; $(A_H^k)_{ij}$ counts walks of length k from i to j , so $\mathbf{1}^\top A_H^k \mathbf{1}$ counts all directed walks of length k .

Emergence discrepancy. Following the whole-versus-parts evaluation of [1], [2], define

$$D_L = W_L(G) - W_L(G_A) - W_L(G_B).$$

A positive D_L measures what the coupled whole's route-accumulating observation reveals beyond the combined part observations; a negative value would measure what coupling erases.

Boundary-touching walks. Let $B_L(G)$ denote the number of walks of length $1 \leq k \leq L$ in G that traverse at least one edge of ∂ .

Claim (identity). $D_L = B_L(G)$ for all $L \geq 1$.

Proof. Partition the walks of G of length $\leq L$ into (i) walks using only edges of E_A , (ii) walks using only edges of E_B , and (iii) walks using at least one boundary edge. Classes (i) and (ii) are exactly the walks counted by $W_L(G_A)$ and $W_L(G_B)$ respectively, and each such walk is counted once in $W_L(G)$ and once in its part functional, so it cancels in D_L . Class (iii) walks appear in $W_L(G)$ but in neither part functional, hence survive with coefficient $+1$. Therefore $D_L = B_L(G)$. $\square$

This identity instantiates the preprint's bound [1], [2]: since $B_L(G)$ is a count of boundary-crossing paths, the emergence discrepancy is bounded by—and for walk-counting measures equals—that count. Moreover, $B_L(G)$ depends only on edges within graph distance $L - 1$ of ∂ (a walk of length k touching ∂ lies within $k - 1 \leq L - 1$ hops of the cut), so the bound is computable locally around the interface, as asserted in [1], [2].

Reference network. We use the 4-cycle C_4 with vertices $a_1, a_2 \in V_A$ and $b_1, b_2 \in V_B$, edges $E_A = \{\{a_1, a_2\}\}$, $E_B = \{\{b_1, b_2\}\}$, and $\partial = \{\{a_1, b_1\}, \{a_2, b_2\}\}$:

with ∂ the two vertical edges. Every vertex has degree $d_v = 2$.

4. Analysis

All input numbers come from the reference network defined in Section 3; no empirical data are used.

Step 1: Whole-system walk counts. G is C_4 , whose adjacency eigenvalues are $\lambda \in \{2, 0, -2, 0\}$ with normalized eigenvector $u_0 = \frac{1}{2}\mathbf{1}$ for $\lambda_0 = 2$. Then

$$\mathbf{1}^\top A^k \mathbf{1} = \sum_m \lambda_m^k (u_m^\top \mathbf{1})^2 = 2^k \cdot \left(\frac{1}{2}\mathbf{1}^\top \mathbf{1}\right)^2 = 2^k \cdot 4 = 4 \cdot 2^k,$$

since all other eigenvectors are orthogonal to $\mathbf{1}$ (for $u_2 = \frac{1}{2}(1, -1, 1, -1)^\top$, $u_2^\top \mathbf{1} = 0$). Thus $W_1(G) = 4 \cdot 2 = 8$, $W_2(G) = 4 \cdot 4 = 16$, $W_3(G) = 4 \cdot 8 = 32$.

Check by direct counting: length-1 walks $= 2|E| = 2 \cdot 4 = 8$. Length-2 walks $= \sum_v d_v^2 = 4 \cdot 2^2 = 16$. Both match.

Step 2: Part-level walk counts. Each part subgraph is a single edge (K_2). For K_2 : length-1 walks $= 2$; length-2 walks $= 2$ ($a_1 \rightarrow a_2 \rightarrow a_1$ and $a_2 \rightarrow a_1 \rightarrow a_2$); length-3 walks $= 2$. Hence the per-length part walk count is constant: $W^{(k)}(G_A) = W^{(k)}(G_B) = 2$ for every $k \geq 1$ (so the cumulative totals are $W_L(G_A) = W_L(G_B) = 2L$), and the combined per-length part observation is

$$W^{(k)}(G_A) + W^{(k)}(G_B) = 4 \quad \text{for every } k.$$

Numerically: $L = 1: 4$; $L = 2: 4$; $L = 3: 4$.

Step 3: Emergence discrepancy. To avoid ambiguity we report per-length discrepancies $D^{(k)} = \mathbf{1}^\top A_G^k \mathbf{1} - \mathbf{1}^\top A_{G_A}^k \mathbf{1} - \mathbf{1}^\top A_{G_B}^k \mathbf{1}$:

$$D^{(1)} = 8 - 2 - 2 = 4, \quad D^{(2)} = 16 - 2 - 2 = 12, \quad D^{(3)} = 32 - 2 - 2 = 28.$$

Cumulative: $D_{\leq 3} = 4 + 12 + 28 = 44$.

Step 4: Independent verification by enumeration of boundary-touching walks. Length-1: the two boundary edges, each traversable in 2 directions, give $B^{(1)} = 2 \times 2 = 4$, matching $D^{(1)} = 4$. Length-2: total walks 16; purely internal walks: 2 in G_A plus 2 in G_B , i.e. 4; so $B^{(2)} = 16 - 4 = 12$. Explicitly, the 12 are: $a_1 \rightarrow a_2 \rightarrow b_2$, $a_2 \rightarrow a_1 \rightarrow b_1$, $b_1 \rightarrow b_2 \rightarrow a_2$, $b_2 \rightarrow b_1 \rightarrow a_1$, $a_1 \rightarrow b_1 \rightarrow b_2$, $a_2 \rightarrow b_2 \rightarrow b_1$, $b_1 \rightarrow a_1 \rightarrow a_2$, $b_2 \rightarrow a_2 \rightarrow a_1$, $a_1 \rightarrow b_1 \rightarrow a_1$, $a_2 \rightarrow b_2 \rightarrow a_2$, $b_1 \rightarrow a_1 \rightarrow b_1$, $b_2 \rightarrow a_2 \rightarrow b_2$. Length-3: total 32; purely internal: $2 + 2 = 4$; $B^{(3)} = 32 - 4 = 28$. The identity $D^{(k)} = B^{(k)}$ holds at every length, and $B_{\leq 3} = 44 = D_{\leq 3}$.

Step 5: Closed-form projection. For general L , using Steps 1 and 2:

$$D_{\leq L} = \sum_{k=1}^L (4 \cdot 2^k - 4) = 4(2^{L+1} - 2) - 4L = 2^{L+3} - 8 - 4L.$$

Check: $L = 3: 2^6 - 8 - 12 = 64 - 20 = 44$; $L = 1: 2^4 - 8 - 4 = 4$; $L = 2: 2^5 - 8 - 8 = 16 = 4 + 12$. As a clearly-labeled projection under the stated assumptions (unweighted simple graph, walk-count observations, horizon L), the discrepancy grows as $\Theta(2^L)$ while the part-level contribution stays flat at 4 per length (cumulative $4L$): coupling, not part size, drives emergence in this topology.

Step 6: Locality. Each boundary-touching walk of length k stays within $k - 1 \leq L - 1$ hops of ∂ ; the relevant vertex set is $N_{L-1}(\partial) = \{v : \text{dist}(v, \partial) \leq L - 1\}$, here all of V for $L \geq 2$, but in a large sparse network a strict subset. Hence B_L is computable from the $2(L - 1)$ -hop neighborhood of the cut (paths may wander $L - 1$ hops into each side), confirming the local computability claimed in [1], [2].

5. Results

All numbers below are computed in Section 4 from the reference network; the growth law is a labeled projection.

1. **Exact identity.** For route-accumulating (walk-count) observations, the emergence discrepancy equals the boundary-touching walk count exactly: $D^{(k)} = B^{(k)}$ for $k = 1, 2, 3$, verified both algebraically and by enumeration (Section 4, Steps 3–4).
1. **Reference-network values.** Per-length discrepancies: $D^{(1)} = 4$, $D^{(2)} = 12$, $D^{(3)} = 28$; cumulative $D_{\leq 3} = 44$. Whole-system walk counts: 8, 16, 32; combined part counts: 4, 4, 4.
1. **Growth projection.** Under the assumptions of Section 3 (unweighted C_4 -like cut with two boundary edges, walk-count observations), $D_{\leq L} = 2^{L+3} - 8 - 4L$, giving $D_{\leq 5} = 2^8 - 8 - 20 = 220$ and $D_{\leq 10} = 2^{13} - 8 - 40 = 8144$. Uncertainty: these are exact for the model system; for a cut with m boundary edges in a locally tree-like graph with mean degree $\bar{d}$, the projected scaling is $\Theta(\bar{d}^L m)$ with the constant undetermined without further specification, so we attach a wide uncertainty band (factor of $\bar{d}^L$) to any transfer of the constant 4 to other topologies.
1. **Locality.** The bound requires only the $2(L - 1)$ -hop neighborhood of the cut (Section 4, Step 6), so for $L = 3$ a two-hop shell around the interface suffices—independent of total system size.
1. **Per-edge attribution.** In the reference network, boundary-touching walks distribute evenly across the two cut edges by symmetry: each of $\{a_1, b_1\}$ and $\{a_2, b_2\}$ carries $44/2 = 22$ of the cumulative $D_{\leq 3} = 44$ walks. In asymmetric cuts, per-edge attribution follows by counting walks whose first boundary traversal uses each edge; this is the basis of the attribution procedure discussed in Section 6.

6. Discussion

Strengths. The identity $D_L = B_L(G)$ is exact, elementary, and auditable; it converts the abstract claim of [1], [2] into arithmetic that a reviewer can check by hand, and it makes the interface-locality property constructive rather than asserted.

Limitations and failure modes. First, walk-counting observations are a narrow class: many scientifically interesting observations (spectral, dynamical, thermodynamic) are not route-accumulating, and for them the identity becomes only the looser bound of [1], [2], whose tightness we have not measured here. Second, our reference network is unweighted and deterministic; weighted or stochastic interactions would require replacing walk counts by weighted path sums, and cancellation between positive and negative contributions could make D_L small even when many boundary routes exist—a silent failure mode of the discrepancy as a risk metric. Third, D_L counts walks, not simple paths; on dense graphs walks overcount cyclic redundancy, potentially inflating the apparent emergence of a single structural feature. Fourth, the choice of cut is doing enormous work: a poorly chosen bipartition can hide emergence inside a "part."

Falsifiability. The account predicts: (i) for route-counting observations, $D_L \leq B_L(G)$ always, with equality in the walk-counting case; a single counterexample where a route-counting observation's discrepancy exceeds the boundary-touching walk count would falsify the central theorem as stated. (ii) It predicts that perturbing a single cut

edge changes total emergence by at most that edge's attributed count (here, at most 22 units for either cut edge at $L = 3$); an observed change exceeding the attribution would falsify the attribution procedure. (iii) It predicts interfaces with isomorphic path neighborhoods (testable with the machinery of [10]) carry equal emergence bounds; a verified pair with equal local structure but unequal emergence would falsify the locality claim. More broadly, finding a natural system class where emergent phenomena of interest are demonstrably not localized near any interface—i.e., where B_L is small yet whole-system behavior diverges sharply from part-combined behavior—would undermine the framework's usefulness.

Open questions. Tightness for non-route-accumulating measures; extension to directed and temporal networks; whether isomorphism of interface neighborhoods in the sense of QNFO [10] guarantees equal emergence bounds, enabling transfer across systems; and empirical calibration of D_L against measured emergent behavior in engineered networks such as those of [6], [8], [9].

Self-critique. Our main theorem is essentially a counting identity; its value is pedagogical and diagnostic, not deep. The $\Theta(2^L)$ growth projection is trivially a consequence of exponential walk growth and should not be read as a law of nature. The nervous-system application in [1], [2] (interneurons relaying most emergence) is cited but not reproduced here, since we computed nothing biological.

7. Conclusion

We reconstructed the boundary-crossing-path theory of network emergence [1], [2] in a walk-counting setting and proved, verified by exhaustive enumeration, that the emergence discrepancy equals the number of boundary-touching walks: $D^{(1)} = 4$, $D^{(2)} = 12$, $D^{(3)} = 28$, total 44 on the reference network, with closed-form growth $D_{\leq L} = 2^{L+3} - 8 - 4L$. The bound is computable from the interface neighborhood alone, connecting emergence accounting to practical problems in fault detection [3], scalability [4], epidemic control [5], service composition [8], and layered grid control [9]. Emergence, on this view, is not a mysterious surplus but a countable budget of routes across boundaries—one that can be anticipated, attributed, and engineered.

References

-
- [1] TITLE: arXiv Query: `searchquery=&idlist=2609.38037&start=0&max_results=1` [2] Emergence in Network Systems is Bounded by Boundary-Crossing Paths. arXiv:2609.38037v1. <https://arxiv.org/abs/2609.38037v1> [3] Generic Detectability and Isolability of Topology Failures in Networked Linear Systems. arXiv:2005.04687v3. <https://arxiv.org/abs/2005.04687v3> [4] Scalability in nonlinear network systems affected by delays and disturbances. arXiv:2006.07422v4. <https://arxiv.org/abs/2006.07422v4> [5] Decentralized Protection Strategies against SIS Epidemics in Networks. arXiv:1409.1730v2. <https://arxiv.org/abs/1409.1730v2> [6] Midterm Status Report of the ILC Technology Network Activities. arXiv:2603.01172v1. <https://arxiv.org/abs/2603.01172v1> [7] Label-dependent Feature Extraction in Social Networks for Node Classification. arXiv:1303.0095v1. <https://arxiv.org/abs/1303.0095v1> [8] A Constrained Shortest Path Scheme for Virtual Network Service Management. arXiv:1808.03031v1. <https://arxiv.org/abs/1808.03031v1> [9] Hierarchical Power Flow Control in Smart Grids: Enhancing Rotor Angle and Frequency Stability with Demand-Side Flexibility. arXiv:2108.05898v1. <https://arxiv.org/abs/2108.05898v1> [10] DOI 10.5281/zenodo.18199940. QNFO: Comprehensive Technical Framework for Network Isomorphism.

Appendix A. Divergence report

D1. Definition of the emergence discrepancy (A vs. B vs. C). Draft A defines $\Delta_E = \mathcal{O}_{\text{whole}} - \mathcal{O}_{\text{parts}}$ as a difference of route-built observations, with a bounding theorem $\Delta_E \leq N_{bc}$ and equality only for unit-weight path counting. Draft B defines $E = |O_W \setminus O_P| + |O_P \setminus O_W|$, a symmetric set-difference cardinality over route multisets, which includes routes the coupling erases. Draft C defines $D_L = W_L(G) - W_L(G_A) - W_L(G_B)$, an asymmetric walk-mass excess. These conventions differ: B's symmetric version can be nonzero even when C's asymmetric excess is zero (when coupling only destroys routes), and A's version depends on the observation functional $\mathcal{O}$. **Resolution:** the main text adopts Draft C's asymmetric walk-mass convention, because it yields an exact, provable identity ($D_L = B_L(G)$) with fully auditable arithmetic. B's erasure term is acknowledged in Section 6 as a generalization not derived here; A's functional-dependent bound is subsumed, since the walk-count functional is a special case of A's unit-weight observation.

D2. Worked example and headline numbers (A vs. B vs. C). Draft A uses a two-module directed network (two 3-node modules, two directed cut edges) and reports $N_{bc} = 7$ at $L = 3$, with a projection to 13 at $L = 5$. Draft B uses a different two-module directed network (two 3-node modules, three one-directional cut edges) and reports $E_{\max} = 8$ at $L = 2$, with an observable-level value $E_{\text{route}} = 5$ and a projection $N_{\partial}(2) = 3m$. Draft C uses the undirected 4-cycle C_4 with two cut edges and reports $D_{\leq 3} = 44$, with projection $D_{\leq L} = 2^{L+3} - 8 - 4L$. These numbers are **not contradictory**—they are computed on different networks under different conventions (directed vs. undirected, different cut sizes, different L)—but they cannot be merged into a single headline figure. **Resolution:** the main text adopts Draft C's example, because it is the only one verified by two independent methods (spectral computation and exhaustive enumeration) with every arithmetic step shown, and because the undirected setting makes the identity exact without directionality caveats. Draft A's matrix-power method is compatible with the adopted framework (it is an algebraic way to compute the same walk counts on directed graphs) and is noted as such; Draft B's per-edge attribution idea is retained qualitatively in Section 5 (Result 5) and Section 6, but its specific numbers (8, 5, 24) pertain to B's network and are not used.

D3. Tightness claim (A vs. B vs. C). Draft A states equality $\Delta_E = N_{bc}$ for path-counting observations and reports achieving it. Draft B reports equality for the full route multiset ($E = 8 = E_{\max}$) but *strict inequality* for a concrete route-counting observable ($E_{\text{route}} = 5 < 8$), attributing the gap to route identities not visible in counts. Draft C proves exact equality for walk-counting functionals. **Resolution:** the main text states equality for walk-counting (route-mass) observations, per C's proof, and records B's observation that *coarser* route-counting observables (e.g., per-node count vectors) may achieve only a fraction of the bound—in B's example $5/8 = 62.5\%$ —because distinct walks can produce identical observable values. This is documented as a caveat in Section 6 rather than silently harmonized.

D4. Directionality of the example (A/B vs. C). Drafts A and B use directed graphs; Draft C uses an undirected graph. The convention affects walk counts (directed walks vs. walks in either direction along an edge). **Resolution:** the main text uses the undirected convention of C for the worked example; the directed case is flagged as an extension (Section 6, open questions), following A's observation that the trace-based method adapts to directed graphs.

Appendix B. Claim attribution

Claim	Source drafts	Agreement status
C1: Emergence is definable as a discrepancy between whole-system and combined part-level observations	A, B, C	CONVERGENT

Claim	Source drafts	Agreement status
C2: Every unit of emergent discrepancy is carried by a route crossing the interface between parts	A, B, C	CONVERGENT
C3: The number of boundary-crossing paths bounds the total emergence discrepancy	A, B, C	CONVERGENT
C4: For route-counting/walk-counting observations the bound is achieved exactly	A, B, C	CONVERGENT (with D3 caveat on coarser observables from B)
C5: The bound is computable from local information near the interface only	A, B, C	CONVERGENT
C6: Exact identity $D_L = B_L(G)$ with proof by walk partition	C	SINGLE
C7: Reference-network values $D^{(1)} = 4, D^{(2)} = 12, D^{(3)} = 28, D_{\leq 3} = 44$	C	SINGLE
C8: Closed-form growth $D_{\leq L} = 2^{L+3} - 8 - 4L$	C	SINGLE
C9: Worked directed example with $N_{bc} = 7$ at $L = 3$, projection 13 at $L = 5$	A	SINGLE
C10: Matrix-power/trace method for computing boundary-crossing counts on directed graphs	A	SINGLE
C11: Worked directed example with $E_{max} = 8, E_{route} = 5$ at $L = 2$	B	SINGLE
C12: Coarser route-counting observables achieve only a fraction of the bound (here 62.5%)	B	SINGLE
C13: Per-edge attribution of emergence; asymmetric cuts concentrate emergence on high-degree cut edges	B (quantitative), C (qualitative, Result 5)	CONVERGENT (qualitatively)
C14: Symmetric discrepancy including an erasure term $ O_P \setminus O_W $	B	SINGLE
C15: Connection to topology-failure detectability and isolability [3]	A, B, C	CONVERGENT
C16: Connection to scalability conditions for networked systems [4]	A, B, C	CONVERGENT
C17: Connection to decentralized SIS epidemic protection [5]	A, B, C	CONVERGENT
C18: Connection to virtual network service management / constrained shortest paths [8]	A, B, C	CONVERGENT
C19: Connection to hierarchical power flow control in smart grids [9]	A, B, C	CONVERGENT
C20: Connection to label-dependent social-network feature extraction [7]	A, B, C	CONVERGENT
C21: ILC Technology Network report [6] as an infrastructure application context	A, B, C	CONVERGENT
C22: QNFO network-isomorphism framework [10] for transferring bounds across isomorphic interfaces	B, C	CONVERGENT
C23: Falsification criteria (counterexample networks, attribution violation, isomorphic-interface test)	B, C	CONVERGENT
C24: Failure modes: non-route-accumulating observations, weights/cancellation, walk overcounting on dense graphs, cut-choice sensitivity	B, C	CONVERGENT
C25: Self-critique that the main theorem is a counting identity of pedagogical/diagnostic value	B, C	CONVERGENT