

Article

Discrete Ricci Curvature on Complex Network Topologies: An Application to Management Information System Networks

Emre Öztürk 

Department of Management Information Systems, Faculty of Economics and Administrative Sciences,
Hatay Mustafa Kemal University, Hatay 31060, Türkiye; emre.ozturk@mku.edu.tr

Abstract

Discrete Ricci curvature has been applied to biological, financial, and communication graphs, but typed management information system (MIS) transaction networks still lack MIS-specific analytical guarantees and a reproducible, operationally interpretable workflow. We combine Ollivier–Ricci optimal transport with a role-based access control representation of users, roles, services, endpoints, and databases. Closed-form results show that edges joining a hub to degree-one leaves have non-negative curvature, whereas canonical inter-hub bridges converge to curvature -1 when the idleness parameter is $\alpha = 0.5$. We also prove the existence of an optimal coupling on finite graphs and explain why the coupling itself need not be unique. On a reproducible synthetic MIS network, the curvature ranks authentication, shared-data, and cross-module links among the leading candidate bottlenecks and correlates with edge betweenness (Spearman $\rho = -0.78$). Across 40 randomized MIS instances, graph-instance-level edge-class differences are significant (Friedman $p < 10^{-32}$), and all five leading edges belong to predefined structural bridge classes in every run. On five public Internet backbones, curvature–betweenness correlations range from -0.53 to -0.91 . Comparisons with Forman curvature, local edge connectivity, algebraic connectivity loss, and a spectral embedding show a complementary geometric signal rather than a universal cut-edge or spectral surrogate. We additionally report the idleness parameter sensitivity, empirical runtime and process memory measurements, study limitations, and a conceptual temporal extension.

Keywords: discrete Ricci curvature; optimal transport; Wasserstein distance; complex networks; management information systems; computational geometry

MSC: 53C21; 05C12; 05C82; 90B15; 94C15



Academic Editor: Darren Narayan

Received: 18 June 2026

Revised: 26 July 2026

Accepted: 29 July 2026

Published: 3 August 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The robustness and structural integrity of complex networks are important in disciplines ranging from biology to management information systems (MISs). Traditional graph-theoretic metrics such as the degree and betweenness centrality identify important nodes and edges but do not directly encode the geometry of transport between probability neighborhoods. Enterprise architectures may exhibit hub-dominated, modular, or small-world organization [1–4].

Recent work adapted continuous geometric ideas to discrete settings. In differential geometry, Ricci curvature is a fundamental tensor on a Riemannian manifold that quantifies the deviation of the manifold’s geometry from being locally Euclidean. It governs volume growth, the convergence or divergence of geodesics, and the behavior of heat flow; its

synthetic reformulations on metric-measure spaces, via the curvature-dimension condition and optimal transport, underpin the discrete theory used here [5–7], building on the classical link between curvature and harmonic maps and heat flow [8]. Extending this continuous concept to graphs, which lack a classical tangent space, yields several notions of discrete curvature [9–12]. Discrete curvature has since been used as an indicator of robustness and systemic risk in domains ranging from financial markets to biological and communication networks [13–16].

Existing network-curvature studies establish the value of Ollivier–Ricci curvature for Internet topology, community detection, and systemic risk analysis [13,17,18]. The unresolved MIS-specific question is not merely whether curvature can be computed but which typed transaction links should be expected to be negatively curved, why that sign has an operational interpretation, and whether the signal generalizes beyond one generated graph. The present study therefore differs from generic curvature applications by combining closed-form analysis of canonical enterprise substructures with a typed role-based access control (RBAC) architecture, a reproducible vulnerability-detection algorithm, statistical replication, and external communication-network evaluation.

The contributions are fourfold:

- (i) We state the graph assumptions explicitly and prove that the finite optimal transport problem defining curvature is well posed, while distinguishing uniqueness of the optimal value from the generally false uniqueness of an optimal plan;
- (ii) We derive closed-form Ollivier–Ricci curvatures for degree-one hub spokes and inter-hub bridges and provide their geometric interpretation;
- (iii) We map users, roles, services, endpoints, and databases to a typed MIS graph and give reproducible pseudocode for curvature-based vulnerability ranking; and
- (iv) We evaluate robustness and parameter sensitivity, report runtime and memory measurements, compare graph-theoretic baselines, and analyze five public communication networks.

1.1. Continuous Motivation and Finite-Graph Assumptions

Let (M, g) be a smooth Riemannian manifold. For a point $x \in M$ and a unit tangent vector $v \in T_x M$, the Ricci curvature $\text{Ric}(v, v)$ averages the sectional curvatures of planes containing v . Heuristically, and rigorously in suitable small-scale limits, it describes whether small probability neighborhoods around nearby points are closer or farther apart than their centers. Ollivier’s construction transfers this comparison of neighborhoods to a metric space by measuring the transport cost between local probability measures [9].

Throughout this paper, unless stated otherwise, $G = (V, E)$ is a finite, simple, connected, undirected, unweighted graph with $n = |V| \geq 2$. We write $m = |E|$, $\mathcal{N}(x) = \{v \in V : \{x, v\} \in E\}$, and $d_x = |\mathcal{N}(x)|$. The function $d : V \times V \rightarrow \mathbb{N}_0$ is the shortest-path metric. Connectedness makes every distance finite and every degree positive, while the finiteness of V implies that (V, d) is complete and has a finite diameter. Sparsity is not required for the definition or existence results below; it affects only the computational cost. For disconnected empirical data, curvature is computed componentwise; our experiments retain the largest connected component.

1.2. Optimal Transport and Discrete Curvatures

For probability measures μ and ν on V , let $\Pi(\mu, \nu)$ denote the set of non-negative matrices $\pi : V \times V \rightarrow [0, 1]$ satisfying $\sum_v \pi(u, v) = \mu(u)$ and $\sum_u \pi(u, v) = \nu(v)$. The L^1 -Wasserstein distance is

$$W_1(\mu, \nu) = \min_{\pi \in \Pi(\mu, \nu)} \sum_{(u, v) \in V \times V} d(u, v) \pi(u, v). \quad (1)$$

Proposition 1 (Well-posedness of the finite transport problem). *For every pair of probability measures μ, ν on the finite graph G , the feasible set $\Pi(\mu, \nu)$ is non-empty, and the minimum in Equation (1) is attained. The number $W_1(\mu, \nu)$ is unique, although an optimal coupling π^* need not be unique. Consequently, the Ollivier–Ricci curvature defined below is unambiguous without imposing uniqueness of the transport plan.*

Proof. The product coupling $\pi(u, v) = \mu(u)\nu(v)$ proves the non-emptiness. The marginal constraints and non-negativity define a closed and bounded polytope in the finite-dimensional space $\mathbb{R}^{n^2}$; hence, $\Pi(\mu, \nu)$ is compact. The transport objective is linear and continuous, and thus it attains its minimum on this polytope. All minimizers have the same scalar objective value by definition, but several matrices may attain it. For example, on a four-cycle set-up with cyclic vertices 1, 2, 3, 4, every coupling between $\frac{1}{2}(\delta_1 + \delta_3)$ and $\frac{1}{2}(\delta_2 + \delta_4)$ has a unit cost. Therefore, distinct optimal plans exist. The curvature depends only on the unique optimal value and is therefore unambiguous. $\square$

Under the Kantorovich duality [19], we have

$$W_1(\mu, \nu) = \sup_{\|f\|_{\text{Lip}} \leq 1} \sum_{v \in V} f(v)(\mu(v) - \nu(v)), \quad (2)$$

where $\|f\|_{\text{Lip}} = \sup_{x \neq y} |f(x) - f(y)|/d(x, y)$. This dual form certifies the optimality of the explicit transport costs derived in Section 2. Computationally, the finite linear program can be solved directly or approximated via entropic regularization [20,21].

For $\alpha \in [0, 1]$, the α -lazy random walk measure at x is

$$m_x^\alpha(v) = \begin{cases} \alpha, & v = x, \\ (1 - \alpha)/d_x, & v \in \mathcal{N}(x), \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

For an edge $e = \{x, y\} \in E$, the Ollivier–Ricci curvature is

$$\kappa_\alpha(x, y) = 1 - \frac{W_1(m_x^\alpha, m_y^\alpha)}{d(x, y)}. \quad (4)$$

Because $d(x, y) = 1$ for an unweighted edge, $\kappa_\alpha(x, y) = 1 - W_1(m_x^\alpha, m_y^\alpha)$. The parameter α is the probability of remaining at the current vertex; Section 5.5 examines its empirical influence. When α is fixed, we write $\kappa(x, y)$ for $\kappa_\alpha(x, y)$ to simplify the notation.

For comparison, the augmented Forman–Ricci curvature of an unweighted edge is [22,23]

$$F(x, y) = 4 - d_x - d_y + 3\Delta_{xy}, \quad (5)$$

where $\Delta_{xy} = |\mathcal{N}(x) \cap \mathcal{N}(y)|$ is the number of triangles containing $\{x, y\}$. Forman curvature is combinatorial and inexpensive, whereas Ollivier–Ricci curvature compares probability neighborhoods through the graph metric. Their distinction is evaluated in Sections 5.3 and 5.7.

Table 1 summarizes the notation used throughout the manuscript.

Table 1. Notation used throughout the manuscript.

Symbol	Meaning
$G = (V, E)$	Finite, simple, connected, undirected, unweighted graph
n, m	Graph order $ V $ and size $ E $
$\mathcal{N}(x), d_x$	Open neighborhood and degree of vertex x
$d(x, y)$	Shortest-path metric on V
α	Lazy walk idleness parameter in $[0, 1]$
m_x^α	Probability measure supported on x and $\mathcal{N}(x)$
$\Pi(\mu, \nu)$	Set of couplings with marginals μ and ν
$W_1(\mu, \nu)$	L^1 -Wasserstein transport distance
$\kappa_\alpha(x, y)$	Ollivier–Ricci curvature of edge $\{x, y\}$
$\mathbf{F}(x, y)$	Augmented Forman–Ricci curvature
Δ_{xy}	Number of triangles incident to edge $\{x, y\}$
$\lambda_2(G)$	Algebraic connectivity (second Laplacian eigenvalue)
$d_{\max}$	Maximum vertex degree
$b_{xy}, b_{\max}$	$b_{xy} = d_x + d_y + 2$; $b_{\max}$ is its maximum over E
$\text{EBC}(e)$	Edge betweenness centrality
$\lambda_{\text{local}}(e)$	Local edge connectivity between the endpoints of e
$\Delta\lambda_2(e)$	Algebraic connectivity loss after deleting e
$d_{\text{LE}}(e)$	Endpoint distance in the normalized-Laplacian eigenmap
$\tau, E_{\text{critical}}(\tau)$	Operational score threshold and corresponding flagged edge set
$G_t, \Delta\kappa_t(e)$	Temporal graph and edge curvature change score

2. Mathematical Analysis of MIS Topologies

Management information system (MIS) transaction graphs can contain hub-and-spoke substructures when shared databases or authentication services connect many clients. Data exchange between functional modules may also be mediated by a small number of inter-hub links. We now derive closed-form curvature expressions for these two canonical structures and show that they have geometrically opposite signatures.

2.1. Curvature of Degree-One Hub Spokes

We first treat a degree-one client attached to an administrative hub. The result is local: the hub may have connections elsewhere in the graph, and thus the statement applies directly to user-to-role spokes in the empirical MIS network. Contrary to the intuition that every edge incident to a high-degree hub is a bottleneck, such degree-one hub-to-leaf edges are *non-negatively* curved.

Lemma 1 (Curvature of MIS hub-to-leaf edges). *Let $\{c, l\} \in E$, where l is a degree-one vertex and $d_c = N \geq 1$. The other neighbors of c and the rest of the graph are otherwise unrestricted. For $\alpha \in [\frac{1}{N+1}, 1]$, the discrete Ollivier–Ricci curvature of the spoke edge is*

$$\kappa(c, l) = \frac{2(1 - \alpha)}{N}. \quad (6)$$

In particular, $\kappa(c, l) \geq 0$, and as $N \rightarrow \infty$, the curvature tends toward zero (from above when $\alpha < 1$).

Proof. Write $\mathcal{N}(c) = \{l, z_1, \dots, z_{N-1}\}$. The lazy walk measures are

$$m_c^\alpha(v) = \begin{cases} \alpha, & v = c \\ \frac{1-\alpha}{N}, & v \in \mathcal{N}(c) \\ 0, & \text{otherwise,} \end{cases} \quad m_l^\alpha(v) = \begin{cases} \alpha, & v = l \\ 1 - \alpha, & v = c \\ 0, & \text{otherwise,} \end{cases}$$

since the unique neighbor of l is c . Put $q = (1 - \alpha)/N$, and let $\delta = m_l^\alpha - m_c^\alpha$. Then, we have

$$\delta(c) = 1 - 2\alpha, \quad \delta(l) = \alpha - q, \quad \delta(z_j) = -q \quad (1 \leq j \leq N - 1),$$

and $\sum_v \delta(v) = 0$. The hypothesis $\alpha \geq 1/(N + 1)$ is equivalent to $\alpha \geq q$. An admissible transport from m_l^α to m_c^α keeps the mass q at l , sends q from c to each z_j , keeps the remaining mass q at c , and sends the surplus $\alpha - q$ from l to c . Its cost is

$$(\alpha - q) + (N - 1)q = \alpha + (N - 2)q = 1 - \frac{2(1 - \alpha)}{N}.$$

For the matching lower bound, define $f(c) = 0$, $f(l) = 1$, and $f(z_j) = -1$. This assignment is 1-Lipschitz on $\{c\} \cup \mathcal{N}(c)$ because l is a leaf, and hence $d(l, z_j) = 2$; it has a 1-Lipschitz extension to all of V . Evaluating the Kantorovich dual in Equation (2) against δ gives $(\alpha - q) + (N - 1)q$, and thus the displayed transport is optimal. Applying Equation (4) yields

$$\kappa(c, l) = 1 - W_1(m_c^\alpha, m_l^\alpha) = \frac{2(1 - \alpha)}{N},$$

which completes the proof. $\square$

Interpretation. Lemma 1 shows that a spoke connecting a hub to a degree-one client is *not* a bottleneck in the optimal transport sense; its curvature is small but non-negative and vanishes only in the large hub limit. Geometrically, the endpoint probability neighborhoods overlap strongly: all non-idle mass at the leaf is already placed on the hub, and mass assigned to the hub's other neighbors moves through the same center. They therefore overlap rather than diverge. This corrects the common misconception that every edge incident to a highly connected hub is fragile; node-level load and edge-level geometric bottleneck status are distinct properties.

2.2. Curvature of Inter-Hub Bridge Edges

In the canonical structure considered next, negative curvature arises on an edge joining two otherwise weakly connected hubs. We model this using two stars joined by a single bridge.

Theorem 1 (Curvature of MIS inter-hub bridges). *Let G consist of two hubs A and B , each additionally adjacent to $k \geq 1$ private leaves, joined by a single bridge edge $\{A, B\}$ such that $d_A = d_B = k + 1$ and A, B share no common neighbor. For $\alpha \in [0, 1]$, the discrete Ollivier–Ricci curvature of the bridge is*

$$\kappa(A, B) = -\frac{2(1 - \alpha)(k - 1)}{k + 1}. \quad (7)$$

Hence, $\kappa(A, B) \leq 0$ for $k \geq 1$, and as the hubs grow ($k \rightarrow \infty$), the curvature tends toward $-2(1 - \alpha)$, whose value at $\alpha = \frac{1}{2}$ is -1 .

Proof. Write $p = \frac{1 - \alpha}{k + 1}$. With $d(A, B) = 1$, the lazy walk measures place the mass

$$m_A^\alpha : A \mapsto \alpha, \quad B \mapsto p, \quad a_i \mapsto p \quad (i \leq k); \quad m_B^\alpha : B \mapsto \alpha, \quad A \mapsto p, \quad b_i \mapsto p \quad (i \leq k),$$

where $\{a_i\}$ and $\{b_i\}$ are the private leaves of A and B , respectively. The signed difference $\delta = m_A^\alpha - m_B^\alpha$ satisfies $\delta(A) = \alpha - p$, $\delta(B) = p - \alpha$, $\delta(a_i) = p$, and $\delta(b_i) = -p$. Deleting the bridge $\{A, B\}$ leaves the total imbalance $\alpha + (k - 1)p$ on the A side, while each of the $2k$ leaf edges carries an imbalance p . Because this canonical graph is a tree, the transport cost equals the sum of the absolute component imbalances across its edges, and hence

$$W_1(m_A^\alpha, m_B^\alpha) = \alpha + (k - 1)p + 2kp = \alpha + (3k - 1)p,$$

which is also certified by the dual potential $f(a_i) = \frac{3}{2}$, $f(A) = \frac{1}{2}$, $f(B) = -\frac{1}{2}$, and $f(b_i) = -\frac{3}{2}$. This function is 1-Lipschitz, and its dual value is $\alpha + (3k - 1)p$. Substituting $p = (1 - \alpha)/(k + 1)$ yields

$$\kappa(A, B) = 1 - W_1 = (1 - \alpha) \left(1 - \frac{3k - 1}{k + 1} \right) = -\frac{2(1 - \alpha)(k - 1)}{k + 1},$$

as claimed. $\square$

Geometric interpretation and systemic-risk scope. Lemma 1 and Theorem 1 distinguish overlapping transport neighborhoods on degree-one spokes from diverging neighborhoods on the canonical inter-hub bridge. The latter curvature tends toward $-2(1 - \alpha)$ as the adjacent modules grow; at $\alpha = \frac{1}{2}$, the limit is -1 . When the modeled bridge is the unique inter-module connector, its removal fragments the graph. In a network with redundant routes, however, strong negative curvature denotes transport-neighborhood divergence and does not by itself prove that the edge is a graph-theoretic cut edge. This distinction guides the empirical interpretation in Section 5.

2.3. Computational Complexity of the Transport Matrix

For an edge $\{x, y\}$, the source and target measures have support sizes of, at most, $d_x + 1$ and $d_y + 1$. The transport linear program therefore has at most $(d_x + 1)(d_y + 1)$ variables and $d_x + d_y + 2$ marginal constraints, rather than a variable for every pair in $V \times V$. We define the support-size upper bound

$$b_{xy} = d_x + d_y + 2, \quad b_{\max} = \max_{\{x, y\} \in E} b_{xy}.$$

The per-edge cost matrix occupies $O(b_{xy}^2)$ memory. For the empirical scaling analysis, we use the conservative workload proxy $b_{xy}^3 \log b_{xy}$ and its graph-level surrogate $m b_{\max}^3 \log b_{\max}$. This is not asserted as a tight worst-case bound for the network-simplex solver. The m edge problems are mutually independent once the required distances are available and are therefore parallelizable.

There is an important implementation distinction. Because the measures are one-hop supported, only the distances between the two local supports are needed mathematically. The reproducibility code nevertheless uses the `all_pairs` option of `GraphRicciCurvature` for deterministic small-graph execution; this adds an $O(n^2)$ distance matrix to memory. A pairwise or dynamically cached implementation avoids that global matrix. Section 5.6 reports the measured runtime and whole-process peak resident memory as functions of n , m , $d_{\max}$, and $b_{\max}$; those benchmarks support practical feasibility but do not constitute an asymptotic-complexity proof.

3. Proposed MIS Network Architecture and Methodology

3.1. Mapping MIS to a Typed Metric Graph

A modern management information system (MIS) can be modeled as a typed transaction graph. Consistent with the assumptions in Section 1.1, $G = (V, E)$ is simple, connected, undirected, and unweighted, and its vertex set is partitioned into the functional layers of a role-based access control (RBAC) architecture [24,25]:

- **Users (U):** individual accounts that initiate transactions;
- **Roles (R):** RBAC roles that aggregate users and mediate authorization (e.g., SALESREP or FINANCEMGR);
- **Application servers (A) and API endpoints (E_p):** the service layer exposing business logic;

- **Database tables (D):** persistent data stores queried by endpoints;
- **Shared infrastructure:** a central *authentication server* (AUTH) and a small number of *core database tables* (CORE) accessed across the whole organization.

The edge set E encodes the corresponding transactions: user-to-role assignments, role-to-AUTH authentication, role-to-service authorization grants, application-to-endpoint calls, endpoint-to-table queries, and the shared-infrastructure links (endpoint-to-CORE and a few cross-module integration calls).

This typing induces a characteristic two-scale geometry. Each business module (the users, roles, servers, endpoints, and private tables of one department) forms a comparatively dense community, whereas inter-module communication is funneled through a handful of shared hubs—the authentication server and the core tables—and through sparse cross-module integration links. The resulting graph is hub-dominated and modular; we do not assume or infer a power-law degree distribution. The curvature analysis is designed to expose edges joining these modules through the transport-divergence mechanism illustrated by Theorem 1 and related to curvature-based community detection [17]. These predefined structural bridge classes are distinct from the degree-one user-to-role spokes treated by Lemma 1.

The experimental pipeline is centralized and offline; in an operational implementation, a trusted analyzer would receive a de-identified adjacency list extracted from transaction logs. The mathematical computation is nevertheless edge-decomposable. Business modules may compute their incident probability measures locally and send only support and distance information to a coordinator, or independent workers may process disjoint edge batches. Thus, the present architecture is centralized for reproducibility and not by theoretical necessity. The node functions are explicit: users initiate access, roles aggregate authorization, application servers and endpoints execute services, private tables store module data, and AUTH and CORE provide shared infrastructure.

3.2. Curvature-Based Vulnerability Detection Framework

The proposed method complements rather than replaces established centrality and connectivity diagnostics. For a chosen operational threshold $\tau \leq 0$, define $E_{\text{critical}}(\tau) = \{e \in E : \kappa_\alpha(e) \leq \tau\}$. Because no universal threshold transfers across organizations, the implementation returns the complete ranking; an administrator may set τ from a monitoring budget or select a lower quantile. Algorithm 1 gives the exact workflow.

Algorithm 1 Unregularized curvature-based vulnerability ranking.

Input: a finite connected graph $G = (V, E)$; idleness $\alpha \in [0, 1]$; optional threshold τ .

Output: edge scores $\{\kappa_\alpha(e)\}_{e \in E}$, a non-decreasing risk ranking, and $E_{\text{critical}}(\tau)$.

1. Validate the graph, remove self-loops and duplicate transactions, retain or separately process each connected component, and assign unit edge lengths.
2. For every vertex x , construct m_x^α from Equation (3).
3. For each edge $e = \{x, y\}$, form the cost matrix $C_{uv} = d(u, v)$ for $u \in \text{supp}(m_x^\alpha)$ and $v \in \text{supp}(m_y^\alpha)$.
4. Solve the finite linear program in Equation (1) with the unregularized OTD and network-simplex implementation, and set $\kappa_\alpha(e) = 1 - W_1(m_x^\alpha, m_y^\alpha)$.
5. Sort edges by increasing κ_α ; flag scores below τ (if supplied), and attach entity-layer labels for interpretation.

Termination and reproducibility: Each transport problem is a finite linear program. The network-simplex implementation returns an optimal feasible solution up to numerical precision; no curvature-flow convergence tolerance is used. The experiments record the master seed, generated instance seeds, package versions, α , graph preprocessing, and solver method.

4. Algorithmic Implementation and Simulation

To evaluate the proposed framework computationally, we developed a simulation environment in Python. Network generation and topological operations use `NetworkX`, while `GraphRicciCurvature` [18] solves the unregularized finite transport problem for each edge up to the solver's numerical precision.

4.1. A Generative Model for MIS Transaction Topologies

Rather than relying on a generic scale-free model, we generated the test network from a layered RBAC process that reproduced the typed structure in Section 3. The instance reported here comprises four business modules (SALES, HR, FINANCE, and INVENTORY); each module contained 8–10 users, two RBAC roles, two application servers, four API endpoints, and three private database tables. The edges were sampled as follows. Every user was assigned one primary role in its module and, independently with probability 0.25, a distinct secondary role; every role was authenticated through the shared server AUTH and granted access to its module's application servers. Applications call endpoints, and endpoints query private tables with fixed probabilities; each endpoint additionally accesses a shared core table with a probability of 0.45; and three cross-module integration links connect endpoints of different departments. The largest connected component is retained, yielding a network of 82 nodes and 135 edges with heterogeneous degrees and prominent RBAC, authentication, and core-data hubs. This is a controlled structural analogue of a transaction-log graph and not a claim that the generator reproduces every empirical enterprise distribution. The full generator is provided in the reproducibility package.

4.2. Curvature Calculation Algorithm

The Ollivier–Ricci curvature for the entire network was computed with $\alpha = 0.5$, a standard lazy walk choice that prevents periodic alternation on bipartite-like substructures. The unregularized Wasserstein problem was solved with `GraphRicciCurvature 0.5.3.2` using `method = OTD`, one worker, unit edge weights, and all-pairs shortest-path distances. The main instance used seed 42. The master seed, generated instance seeds, preprocessing rules, and dependency versions are included in the reproducibility package.

4.3. Extension to Dynamic and Temporal MIS Graphs

For event time t , let $G_t = (V_t, E_t)$ be the transaction snapshot, and define the curvature change $\Delta\kappa_t(e) = \kappa_{G_t}(e) - \kappa_{G_{t-1}}(e)$. An insertion, deletion, or weight change first invalidates measures at its incident vertices. Candidate edge scores in the affected neighborhoods can then be recomputed in parallel. Because a topological update may also alter shortest-path distances outside that local region, an exact incremental implementation must combine this local invalidation with a dynamic shortest-path check; if distances change globally, then all dependent transport matrices are refreshed. A streaming monitor can rank both the current score $\kappa_{G_t}(e)$ and an abrupt negative change $\Delta\kappa_t(e)$, optionally smoothing repeated snapshots to suppress transient events. Directed or weighted transaction graphs require the corresponding asymmetric measures or weighted ground metric and are natural extensions, but the present experiments retain the static, undirected, and unweighted setting so that the theoretical formulas and baselines are directly comparable.

5. Experimental Results and Discussion

5.1. Identification of Structural Bottlenecks via Negative Curvature

The discrete Ollivier–Ricci curvature was computed for all 135 transaction links of the generated MIS network. The curvature ranged over $[-0.66, +0.33]$ and partitioned the edges by structural role. Figure 1 visualizes the typed network and its edge-level

curvature pattern. Table 2 lists the five lowest-curvature candidate bottlenecks; four joined the central server AUTH to module roles, and one was a cross-module endpoint link. The authentication layer was therefore the dominant but not exclusive highest-ranked structural class in this realization.

Table 2. The five lowest-curvature candidate bottlenecks at $\alpha = 0.5$. Every listed edge belongs to a predefined structural bridge class; no degree-one user-to-role spoke appeared.

Rank	MIS Transaction Link	Edge Type	Curvature (κ)
1	AUTH–Role:SALES (r_1)	Role authentication	−0.659
2	AUTH–Role:FINANCE (r_1)	Role authentication	−0.650
3	AUTH–Role:HR (r_0)	Role authentication	−0.639
4	FINANCE: e_3 –INVENTORY: e_1	Cross-module	−0.633
5	AUTH–Role:FINANCE (r_0)	Role authentication	−0.625

Aggregating by edge type made the geometric contrast explicit (Figure 2). Among user-to-role edges, the 26 degree-one spokes covered by Lemma 1 were all positive (mean: +0.114; range: [0.091, 0.143]); the complete user-to-role class also included 18 edges incident to users with two roles and had a mean of +0.107. The module-internal edges scattered around zero (mean: −0.042). In contrast, the three predefined structural bridge classes had negative means: cross-module integration links (−0.458), shared core database accesses (−0.396), and role authentication links (−0.630), with the last being the most negative and most tightly concentrated one. This empirical separation is consistent with the transport neighborhood mechanism illustrated by Lemma 1 and Theorem 1; the generated structural classes are not asserted to satisfy the theorem’s symmetric two-star assumptions exactly. Operationally, disruption of authentication or shared-core links can affect several departments; actual fragmentation depends on route redundancy, and thus the curvature ranking is a screening signal rather than a cut edge certificate.




Figure 1. Layered role-based access control (RBAC) management information system (MIS) transaction network (82 nodes, 135 edges, $\alpha = 0.5$). Nodes are colored by entity layer (users, roles, application servers, API endpoints, database tables, plus the shared authentication server and core tables) and sized by degree; edges are colored from positive curvature (green) to negative curvature (red), with width proportional to negative-curvature magnitude. Authentication, core database, and cross-module links are prominent among the strongly negative edges joining business modules.




Figure 2. Distribution of discrete Ricci curvature by transaction edge type. The user-to-role class includes both degree-one spokes, which are covered by Lemma 1, and edges incident to users with two roles. Module-internal edges center near zero, whereas the three predefined structural bridge classes have negative means in this realization, with role authentication links being the most negative. Points are individual edges; boxes show the interquartile range and median.

5.2. Comparison with Edge Betweenness Centrality

To position the geometric approach against an established baseline, we computed the exact edge betweenness centrality (EBC) [26,27] for every link and compared its ranking with that of the discrete Ricci curvature. The Spearman rank correlation between κ and the EBC was $\rho = -0.78$, indicating that many of the most negatively curved edges also had a large share of shortest paths. We treated this single-graph coefficient descriptively because edges within one graph are structurally dependent; the graph-instance-level uncertainty analysis is reported in Section 5.4. As Figure 3 shows, the predefined structural bridge classes tended to populate the region of higher betweenness and more negative curvature, whereas many other edges had lower betweenness and a curvature closer to zero.




Figure 3. Discrete Ollivier–Ricci curvature versus edge betweenness centrality for all 135 links of the MIS network. Predefined structural bridge classes (red diamonds) clustered at higher betweenness and more negative curvature, whereas the other edges (blue) tended toward lower betweenness and curvature closer to zero. The descriptive rank association is negative (Spearman $\rho = -0.78$), while the two metrics remained conceptually distinct.

This association is informative but does not make the measures redundant. The EBC is a global shortest path-based quantity. Ollivier–Ricci curvature compares one-hop probability measures using the graph shortest path metric; once the required distances have been computed, the transport problem for each edge can be solved independently and in parallel. Its closed-form interpretation in Lemma 1 and Theorem 1 further distinguishes the two canonical neighborhood geometries. Monitoring negative curvature can therefore provide a geometrically interpretable candidate risk ranking that complements centrality analysis.

5.3. Comparison with Forman–Ricci Curvature

Equation (5) defines Forman–Ricci curvature as a computationally cheaper but strictly local alternative [28,29]. To make the trade-off concrete, we computed the augmented Forman curvature $F(e) = 4 - d_x - d_y + 3\Delta(e)$ for every edge of the MIS network and compared it with the Ollivier–Ricci values. The two measures were positively correlated (Spearman $\rho = 0.69$). Both ranked the role authentication links among the most negative, but for different reasons; Forman curvature was strongly aligned with the negative incident-degree sum (Spearman $\rho = 0.98$), whereas the corresponding association for Ollivier–Ricci curvature was weaker ($\rho = 0.72$).

The two curvatures consequently induced different rankings in this architecture. As Figure 4 shows, the complete user-to-role class—to which Ollivier–Ricci assigns a mean curvature of $+0.107$ —received a negative mean Forman curvature of -6.2 , comparable to the cross-module (-5.0) and core database (-5.9) classes. Because Forman curvature responds directly to endpoint degrees and triangles, it does not cleanly separate user-to-role edges from those two predefined structural classes here, although it ranks role-authentication links as highly negative. Ollivier–Ricci curvature supplies a different transport neighborhood signal; neither measure was treated as the ground truth, and the external comparisons below delimit their agreement.




Figure 4. Forman–Ricci versus Ollivier–Ricci curvature for every edge of the MIS network by type. In this realization, degree-sensitive Forman values for user-to-role edges (green) overlap those of the cross-module and core database classes (orange and red diamonds, respectively). Ollivier–Ricci curvature (horizontal axis) yielded a different separation between the complete user-to-role class and the predefined structural bridge classes.

5.4. Robustness Across Randomized Instances

To separate architectural effects from one random seed, we generated 40 independently randomized MIS instances while varying the per-module user counts over $[6, 12]$, the probability of assigning a distinct secondary role over $[0.15, 0.35]$, the core access probability over $[0.30, 0.55]$, and the number of cross-module links from two to four. The graphs contained 83.88 vertices on average ($SD = 4.23$). Curvature was first averaged within each edge class for every graph; Table 3 summarizes those 40 graph-level means.

Table 3. Graph-instance-level edge-class curvature across 40 independently randomized synthetic MIS instances at $\alpha = 0.5$. SD is the between-instance standard deviation. Confidence limits are percentile-bootstrap 95% intervals over graph instances (10,000 resamples).

Edge Class	Mean	SD	Bootstrap 95% CI	Runs
User-role	0.102	0.012	[0.098, 0.105]	40
Intra-module	−0.065	0.025	[−0.073, −0.058]	40
Cross-module	−0.473	0.080	[−0.497, −0.448]	40
Endpoint-core DB	−0.265	0.112	[−0.299, −0.230]	40
Role-AUTH	−0.631	0.013	[−0.635, −0.627]	40

The edge-class differences were significant (Friedman $\chi^2(4) = 156.88$, $p = 6.82 \times 10^{-33}$). Three one-sided paired Wilcoxon comparisons showed that the pooled structural bridge mean was lower than the intra-module mean, the intra-module mean was lower than the user-role mean, and the pooled structural bridge mean was lower than the degree-one user-spoke mean; all had $W = 0$ and a Holm-adjusted $p = 2.73 \times 10^{-12}$. The per-instance Spearman correlation between curvature and the edge betweenness averaged -0.715 ($SD = 0.045$; bootstrap 95% CI $[-0.728, -0.701]$). In all 40 instances, every one of the five most negatively curved edges belonged to a predefined structural bridge class. The degree-one user spokes remained non-negative, whereas a user assigned two roles could act as a small bridge and is not covered by Lemma 1. These inferential results quantify variability under the stated generator and parameter ranges; they do not estimate variability across operational enterprise networks.

5.5. Sensitivity to the Lazy Walk Parameter

The reference 82-vertex, 135-edge graph was recomputed at $\alpha \in \{0, 0.25, 0.5, 0.75, 0.9\}$. Table 4 reports the rank stability against the $\alpha = 0.5$ result. The fixed-size top decile contained 14 edges; equal scores were resolved by a reproducible lexicographic endpoint tie-breaker.

Table 4. Sensitivity of the bottleneck ranking to α on the fixed 82-vertex, 135-edge MIS realization. “Top-decile overlap” is relative to $\alpha = 0.5$; “top-five bridge precision” is the fraction belonging to the role-AUTH, endpoint-core database (DB), or cross-module classes.

α	Rank ρ vs. 0.5	Top-Decile Overlap	Top-Five Bridge Precision
0.00	0.840	1.000	1.000
0.25	0.999	1.000	1.000
0.50	1.000	1.000	1.000
0.75	0.999	1.000	1.000
0.90	0.999	1.000	1.000

Varying α changed the curvature magnitudes but preserved the selected critical edge set on this realization. For $\alpha \in \{0.25, 0.75, 0.9\}$, the full ranking was almost unchanged ($\rho > 0.998$); at the non-lazy endpoint $\alpha = 0$, lower-ranked edges were reordered, and the

agreement decreased to 0.840. Nevertheless, the fixed-size top-decile set and all five leading bridge classifications were unchanged. For $\alpha \geq 0.25$, the magnitudes of all edge-class means contracted toward zero as α increased because the measures retained increasing mass at their source vertices (Figure 5). This experiment supports stability of the selected bottlenecks on the tested graph and not parameter invariance for every MIS topology.




Figure 5. Mean Ollivier–Ricci curvature by transaction-edge class under five lazy walk parameters. For $\alpha \geq 0.25$, class-mean magnitudes contracted toward zero as α increased, while the ordering of the structural classes remained unchanged on the reference MIS graph.

5.6. Empirical Runtime and Memory Scaling

We benchmarked unregularized transport on five generated MIS sizes using three independently generated graphs per size. Measurements were obtained on Windows 11 with Python 3.12.13, a 14-core Intel Core i9-13900H processor, 15.7 GB of RAM, and one curvature worker. The runtime excluded graph generation. The peak resident set size (RSS) was sampled for the complete Python process and therefore included interpreter and library overhead. Table 5 reports the corresponding runtime and memory summaries.

Table 5. Empirical runtime and memory for unregularized Ollivier–Ricci computation at $\alpha = 0.5$. Values are medians over three graphs; parentheses contain the interquartile range (IQR), and RSS denotes resident set size. Here, $b_{\max}$ is the largest local support-size upper bound.

Configuration	n	m	$d_{\max}$	$b_{\max}$	Runtime, s (IQR)	Peak RSS, MB (IQR)
MIS-2M	43	71	11	19	0.0104 (0.0020)	157.55 (0.31)
MIS-4M	83	135	11	21	0.0192 (0.0050)	158.07 (0.14)
MIS-6M	123	198	12	25	0.0272 (0.0012)	158.86 (0.36)
MIS-8M	163	271	16	29	0.0426 (0.0024)	159.52 (0.54)
MIS-12M	242	400	24	37	0.0609 (0.0029)	162.17 (0.48)

The median computation time increased from 0.0104 to 0.0609 s as n grew from 43 to 242 and m grew from 71 to 400 (Figure 6). A descriptive log–log regression on the workload proxy $mb_{\max}^3 \log b_{\max}$ gave a slope of 0.448 ($R^2 = 0.984$, $p = 0.00087$). With only five aggregated sizes, this association neither estimates nor proves an asymptotic exponent. The peak RSS increased from 157.55 to 162.17 MB; the values are dominated by fixed Python and library overhead and should not be interpreted as an algorithm-only memory law.




Figure 6. Exact-curvature runtime and whole-process peak resident memory as the generated MIS topology grew. Points are medians, and error bars span the first to third quartiles over three independent graphs.

5.7. External Communication Networks and Additional Baselines

External evaluation used five public communication backbone topologies from the Internet Topology Zoo [30] distributed in the Repetita data collection [31]: Abilene, AARNet, GEANT 2012, UNINETT 2010, and Interoute. Paired directed arcs were collapsed into simple undirected edges, self-loops were removed, and the largest connected component was retained. Capacity and delay fields were not used, and thus this is an unweighted topological comparison. These networks provide real communication data, but they are not observed enterprise-MIS transaction logs.

For every external edge, we compared κ with the edge betweenness centrality (EBC), augmented Forman curvature F , local edge connectivity λ_{local} , algebraic connectivity loss $\Delta\lambda_2 = \lambda_2(G) - \lambda_2(G - e)$, and endpoint distance d_{LE} in a four-dimensional normalized-Laplacian eigenmap. Table 6 gives the within-network Spearman correlations.

Table 6. External comparison on five public communication network topologies at $\alpha = 0.5$. Entries are edge-level Spearman correlations with Ollivier–Ricci curvature. EBC denotes edge betweenness centrality.

Network	n	m	Negative κ (%)	$\rho(\kappa, \text{EBC})$	$\rho(\kappa, F)$	$\rho(\kappa, \lambda_{\text{local}})$	$\rho(\kappa, \Delta\lambda_2)$	$\rho(\kappa, d_{\text{LE}})$
Abilene	11	14	7.1	−0.909	0.807	0.283	−0.765	−0.259
AARNet	19	24	29.2	−0.759	0.668	−0.059	−0.427	−0.785
GEANT 2012	40	61	54.1	−0.527	0.781	−0.511	0.058	−0.574
UNINETT 2010	74	101	61.4	−0.672	0.744	−0.502	0.075	−0.201
Interoute	110	146	56.2	−0.740	0.735	−0.419	−0.323	−0.179

The curvature agreed directionally with the EBC in all five networks; ρ ranged from −0.909 to −0.527 (median −0.740), and thus the greater shortest-path traffic tended to accompany more negative curvature. Agreement with the Forman curvature was also consistent (0.668–0.807, median = 0.744). By contrast, the local edge connectivity, $\Delta\lambda_2$, and spectral embedding distance were network-dependent. In particular, $\rho(\kappa, \Delta\lambda_2)$ ranged from −0.765 to 0.075 and was near zero for GEANT 2012 and UNINETT 2010. The correlation patterns are summarized in Figure 7.

The most negative 15% of edges overlapped the same-sized high-EBC set by 0.409–1.000 (median: 0.500) but the overlap with high $\Delta\lambda_2$ edges had a median of 0.182. Moreover, the graph-theoretic cut-edge recall was zero for AARNet, GEANT 2012, and Interoute, 0.063 for UNINETT 2010, and undefined for bridge-free Abilene. A peripheral cut edge and an inter-community transport bottleneck are different concepts. A negative curvature should therefore be interpreted as transport neighborhood divergence and not as a universal

replacement for algebraic connectivity or an exact cut-edge test. This limitation also explains why the different baselines are complementary.




Figure 7. Spearman correlations between Ollivier–Ricci curvature and five comparator metrics on public communication backbones. Consistent edge betweenness centrality (EBC) and Forman agreement contrasts with heterogeneous local connectivity, algebraic connectivity loss, and spectral embedding comparisons.

6. Conclusions

This study developed a mathematically well-posed and reproducible Ollivier–Ricci framework for typed MIS transaction networks. The closed-form analysis separates degree-one hub spokes, which are non-negatively curved, from inter-hub bridges, which become strongly negative as the adjacent modules grow. The distinction follows from the overlap or divergence of the endpoint transport neighborhoods and is therefore more informative than the degree alone.

On the reference MIS realization, curvature correlated with the edge betweenness at $\rho = -0.78$ and ranked authentication, shared-data, and cross-module links among the most negative scores. Across 40 independently randomized MIS graphs, edge-class differences remained significant, the curvature–betweenness association averaged -0.715 , and the five leading edges belonged to predefined structural bridge classes in every run. On the reference graph, the selected critical set was stable over the five tested α values. Five public communication backbones provided an external check; the curvature–betweenness correlation remained negative in every network (-0.53 to -0.91), and the Forman agreement was consistently positive.

The expanded comparisons also delimited the claim. Associations with the local edge connectivity, algebraic connectivity loss, and spectral embedding distance varied by network, and negative curvature did not universally recover graph-theoretic cut edges. Ollivier–Ricci curvature is therefore a geometric transport divergence indicator that complements global spectral and cut-based criteria; it is not interchangeable with them. Other limitations are that the MIS data were generated rather than taken from a live enterprise, the external graphs were unweighted communication backbones, and the present theoretical results concern static, undirected graphs. Future work should evaluate privacy-preserving transaction logs, weighted and directed interactions, temporal change scores, and incremental shortest-path updates in operational cloud environments.

Funding: This research received no external funding.

Data Availability Statement: The MIS network is fully reproducible from the layered RBAC generator described in Section 4.1. Code and outputs for the Ollivier–Ricci, Forman–Ricci, betweenness, scaling, sensitivity, statistical, spectral, and external network experiments are publicly available at <https://github.com/emreozi/mis-ricci-curvature> (accessed on 28 July 2026) (revision snapshot: commit 72db3ba1fa0b0a7a6facdac61e05d0a8830186a7) and permanently archived as version v3.1 on Zenodo at <https://doi.org/10.5281/zenodo.21574868>. The same files are also provided in the reproducibility package accompanying this revision. The external topologies are public Internet Topology Zoo graphs distributed by Repetita at <https://github.com/svissicchio/Repetita> (accessed on 28 July 2026); the analyzed snapshot is commit 60e679c2f34d9b65b7f256ff6f6963938fa040f9.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Barabási, A.-L.; Albert, R. Emergence of scaling in random networks. *Science* **1999**, *286*, 509–512. [\[CrossRef\]](#) [\[PubMed\]](#)
2. Albert, R.; Barabási, A.-L. Statistical mechanics of complex networks. *Rev. Mod. Phys.* **2002**, *74*, 47–97. [\[CrossRef\]](#)
3. Watts, D.J.; Strogatz, S.H. Collective dynamics of ‘small-world’ networks. *Nature* **1998**, *393*, 440–442. [\[CrossRef\]](#) [\[PubMed\]](#)
4. Newman, M.E.J. The structure and function of complex networks. *SIAM Rev.* **2003**, *45*, 167–256. [\[CrossRef\]](#)
5. Bakry, D.; Émery, M. Diffusions hypercontractives. In *Séminaire de Probabilités XIX 1983/84*; Lecture Notes in Mathematics; Springer: Berlin/Heidelberg, Germany, 1985; Volume 1123, pp. 177–206. [\[CrossRef\]](#)
6. Sturm, K.-T. On the geometry of metric measure spaces. *Acta Math.* **2006**, *196*, 65–131. [\[CrossRef\]](#)
7. Lott, J.; Villani, C. Ricci curvature for metric-measure spaces via optimal transport. *Ann. Math.* **2009**, *169*, 903–991. [\[CrossRef\]](#)
8. Eells, J.; Sampson, J.H. Harmonic mappings of Riemannian manifolds. *Am. J. Math.* **1964**, *86*, 109–160. [\[CrossRef\]](#)
9. Ollivier, Y. Ricci curvature of Markov chains on metric spaces. *J. Funct. Anal.* **2009**, *256*, 810–864. [\[CrossRef\]](#)
10. Lin, Y.; Lu, L.; Yau, S.-T. Ricci curvature of graphs. *Tohoku Math. J.* **2011**, *63*, 605–627. [\[CrossRef\]](#)
11. Bauer, F.; Jost, J.; Liu, S. Ollivier–Ricci curvature and the spectrum of the normalized graph Laplace operator. *Math. Res. Lett.* **2012**, *19*, 1185–1205. [\[CrossRef\]](#)
12. Jost, J.; Liu, S. Ollivier’s Ricci curvature, local clustering and curvature–dimension inequalities on graphs. *Discret. Comput. Geom.* **2014**, *51*, 300–322. [\[CrossRef\]](#)
13. Ni, C.-C.; Lin, Y.; Gao, J.; Gu, X.D.; Saunderson, E. Ricci curvature of the Internet topology. In Proceedings of the 2015 IEEE Conference on Computer Communications (INFOCOM), Hong Kong, China, 26–30 April 2015; pp. 2758–2766. [\[CrossRef\]](#)
14. Sandhu, R.; Georgiou, T.; Reznik, E.; Zhu, L.; Kolesov, I.; Senbabaoglu, Y.; Tannenbaum, A. Graph curvature for differentiating cancer networks. *Sci. Rep.* **2015**, *5*, 12323. [\[CrossRef\]](#) [\[PubMed\]](#)
15. Sandhu, R.; Georgiou, T.; Tannenbaum, A. Ricci curvature: An economic indicator for market fragility and systemic risk. *Sci. Adv.* **2016**, *2*, e1501495. [\[CrossRef\]](#) [\[PubMed\]](#)
16. Pouryahya, M.; Oh, J.H.; Mathews, J.C.; Deasy, J.O.; Tannenbaum, A.R. Characterizing cancer drug response and biological correlates: A geometric network approach. *Sci. Rep.* **2018**, *8*, 6402. [\[CrossRef\]](#) [\[PubMed\]](#)
17. Sia, J.; Jonckheere, E.; Bogdan, P. Ollivier–Ricci curvature-based method to community detection in complex networks. *Sci. Rep.* **2019**, *9*, 9800. [\[CrossRef\]](#) [\[PubMed\]](#)
18. Ni, C.-C.; Lin, Y.-Y.; Luo, F.; Gao, J. Community detection on networks with Ricci flow. *Sci. Rep.* **2019**, *9*, 9984. [\[CrossRef\]](#) [\[PubMed\]](#)
19. Villani, C. *Optimal Transport: Old and New*; Grundlehren der mathematischen Wissenschaften; Springer: Berlin/Heidelberg, Germany, 2009; Volume 338. [\[CrossRef\]](#)
20. Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems (NeurIPS)*; Neural Information Processing Systems Foundation, Inc.: South Lake Tahoe, NV, USA, 2013; Volume 26, pp. 2292–2300. Available online: https://proceedings.neurips.cc/paper_files/paper/2013/hash/af21d0c97db2e27e13572cbf59eb343d-Abstract.html (accessed on 28 July 2026).
21. Peyré, G.; Cuturi, M. Computational optimal transport. *Found. Trends Mach. Learn.* **2019**, *11*, 355–607. [\[CrossRef\]](#)
22. Forman, R. Bochner’s method for cell complexes and combinatorial Ricci curvature. *Discret. Comput. Geom.* **2003**, *29*, 323–374. [\[CrossRef\]](#)
23. Sreejith, R.P.; Mohanraj, K.; Jost, J.; Saucan, E.; Samal, A. Forman curvature for complex networks. *J. Stat. Mech. Theory Exp.* **2016**, *2016*, 063206. [\[CrossRef\]](#)
24. Sandhu, R.S.; Coyne, E.J.; Feinstein, H.L.; Youman, C.E. Role-based access control models. *Computer* **1996**, *29*, 38–47. [\[CrossRef\]](#)
25. Ferraiolo, D.F.; Sandhu, R.; Gavrila, S.; Kuhn, D.R.; Chandramouli, R. Proposed NIST standard for role-based access control. *ACM Trans. Inf. Syst. Secur.* **2001**, *4*, 224–274. [\[CrossRef\]](#)
26. Freeman, L.C. A set of measures of centrality based on betweenness. *Sociometry* **1977**, *40*, 35–41. [\[CrossRef\]](#)

27. Brandes, U. A faster algorithm for betweenness centrality. *J. Math. Sociol.* **2001**, *25*, 163–177. [[CrossRef](#)]
28. Samal, A.; Sreejith, R.P.; Gu, J.; Liu, S.; Saucan, E.; Jost, J. Comparative analysis of two discretizations of Ricci curvature for complex networks. *Sci. Rep.* **2018**, *8*, 8650. [[CrossRef](#)] [[PubMed](#)]
29. Weber, M.; Saucan, E.; Jost, J. Characterizing complex networks with Forman-Ricci curvature and associated geometric flows. *J. Complex Netw.* **2017**, *5*, 527–550. [[CrossRef](#)]
30. Knight, S.; Nguyen, H.X.; Falkner, N.; Bowden, R.; Roughan, M. The Internet Topology Zoo. *IEEE J. Sel. Areas Commun.* **2011**, *29*, 1765–1775. [[CrossRef](#)]
31. Gay, S.; Schaus, P.; Vissicchio, S. REPETITA: Repeatable experiments for performance evaluation of traffic-engineering algorithms. *arXiv* **2017**, arXiv:1710.08665. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.