

Incidence Reduction: Extracting Exact Strategic Values from Game Topologies via Petri Net ODE Equilibrium

Matt York
pflow.xyz

March 11, 2026

Abstract

We present *incidence reduction*, a technique that recovers exact integer-valued strategic assessments from the topology of a game encoded as a Petri net. The central construction uses unit-weight integer arcs to build an *analysis net*—a bipartite graph of catalytic source places and drain transitions—whose mass-action ODE system admits a lossless discrete $\rightarrow$ continuous $\rightarrow$ discrete round-trip: the decoupling lemma shows that each accumulator obeys an independent first-order ODE $\dot{x}_i = 1 - n_i x_i$ (where n_i is the number of drain transitions), yielding the closed-form equilibrium $x_i^* = 1/n_i$ without iterative computation. After normalization, these equilibrium concentrations produce exact rational (often integer) strategic values determined entirely by net topology. We prove a *spectral ranking preservation theorem*: incidence reduction is a diagonal approximation to eigenvector centrality of the entity co-occurrence matrix $M = BB^T$, and both methods produce the same ordinal ranking for all games tested. The co-occurrence matrix decomposes as $M = 2D - L$, where D is the degree (drain-count) matrix and L is the combinatorial Laplacian of the co-occurrence graph; this decomposition yields a sufficient condition for ranking preservation via a perturbation bound on the off-diagonal coupling. We validate the technique on $n \times n$ tic-tac-toe grids ($n = 3, 5, 7$), recovering the classical center-corner-edge hierarchy $\{4, 3, 2\}$; on poker hand rankings, recovering the full nine-tier hierarchy; on Connect Four (7×6), recovering center-column dominance with a 13:3 value ratio; and on Hex (5×5), recovering anti-diagonal dominance with a 16:1 value ratio from non-Cartesian topology. The same incidence matrix admits three complementary readings—discrete (Petri net firing rules), continuous (ODE equilibrium), and cryptographic (Groth16 ZK proofs)—with the decoupling lemma ensuring the lossless round-trip between all three.

1 Introduction

Evaluating positions in combinatorial games typically requires search: minimax, alpha-beta pruning, Monte Carlo tree search, or neural-network-guided lookahead. These methods are computationally expensive, opaque, and difficult to verify—and the underlying complexity is inherent: game evaluation is PSPACE-complete for many combinatorial games (Section 6). We ask a different question: can the *topology* of a game—the connectivity structure among positions and constraints—already encode strategic value, without any search at all?

This paper answers affirmatively by introducing a three-step pipeline:

1. **Encode.** Represent the game as a Petri net whose places correspond to entities (board cells, hand types) and whose transitions encode constraints (win conditions, ranking rules).
2. **Solve.** Interpret the Petri net under mass-action kinetics and integrate the resulting ODE system to equilibrium.
3. **Read off.** Invert and normalize the equilibrium concentrations to obtain exact integer strategic values.

The key insight is that the catalytic-pump construction produces a system of *decoupled* linear ODEs, each of the form $\dot{x}_i = 1 - n_i x_i$, where n_i is the number of drain transitions attached to accumulator place i . The unique stable equilibrium is $x_i^* = 1/n_i$, so the steady-state concentrations are exact rationals determined entirely by the net topology. Normalization by the minimum concentration yields integers whenever the minimum drain count divides all others.

Contributions.

- A formal definition of the *analysis net* construction and a proof that its equilibrium concentrations are exact reciprocals of drain counts (Theorem 3), applying the classical mass-action kinetics of chemical reaction network theory [12, 6] to a novel catalytic-pump topology.
- A *spectral ranking preservation theorem* (Theorem 7): incidence reduction is a diagonal approximation to eigenvector centrality of the entity co-occurrence matrix, preserving all ordinal ranking information while remaining closed-form. A connection to the graph Laplacian and effective resistance (Section 3.6) provides a sufficient condition for ranking preservation via a perturbation bound on the co-occurrence matrix.
- Experimental validation on 3×3 , 5×5 , and 7×7 tic-tac-toe, poker hand rankings, Connect Four, and Hex—spanning uniform grids, frequency-weighted hierarchies, non-square grids, and non-Cartesian topologies.
- A characterization of the incidence matrix as a *semantic bridge* admitting three readings (discrete, continuous, cryptographic), with the decoupling lemma guaranteeing lossless discrete $\rightarrow$ continuous $\rightarrow$ discrete round-trip.

2 Background

2.1 Petri Nets

A *Petri net* is a bipartite directed graph $N = (P, T, A^-, A^+)$ where P is a finite set of *places*, T is a finite set of *transitions*, and $A^-, A^+ : T \times P \rightarrow \mathbb{N}$ are the *input* and *output* incidence functions. A *marking* $\mathbf{m} \in \mathbb{N}^{|P|}$ assigns a non-negative integer token count to each place. Transition t is *enabled* at marking $\mathbf{m}$ if $\mathbf{m}(p) \geq A^-(t, p)$ for all $p \in P$. Firing t produces the successor marking $\mathbf{m}' = \mathbf{m} + \Delta_t$, where the *delta vector* $\Delta_t(p) = A^+(t, p) - A^-(t, p)$.

The *incidence matrix* $C \in \mathbb{Z}^{|P| \times |T|}$ has columns $C_t = \Delta_t$. It encodes the full topology of the net and is the central object connecting the continuous (ODE) and discrete (ZK) analyses in this paper.

2.2 Mass-Action Kinetics

The mathematical framework for mass-action kinetics in reaction networks was established by Horn and Jackson [12] and subsequently developed into a comprehensive theory by Feinberg [4, 6]. Their *chemical reaction network theory* (CRNT) studies the ODE systems that arise when reaction rates follow the law of mass action, first formulated by Guldberg and Waage [8]. The isomorphism between Petri nets and chemical reaction networks is well known [11]: places correspond to chemical species, transitions to reactions, and the incidence matrix to the stoichiometric matrix.

Under mass-action kinetics, each transition t has a rate constant $k_t > 0$ and induces a flux

$$v_t(\mathbf{x}) = k_t \prod_{p \in \bullet t} x_p^{A^-(t, p)},$$

where $\bullet t = \{p : A^-(t, p) > 0\}$ is the preset of t and $\mathbf{x} \in \mathbb{R}_{\geq 0}^{|P|}$ is the continuous state (concentration) vector. The ODE system governing the evolution of concentrations is

$$\dot{\mathbf{x}} = C \cdot \mathbf{v}(\mathbf{x}), \quad (1)$$

where $\mathbf{v}(\mathbf{x}) = (v_t(\mathbf{x}))_{t \in T}$ is the flux vector. This is precisely the deterministic rate equation of CRNT [12].

When the system reaches a steady state $\mathbf{x}^*$ with $\dot{\mathbf{x}} = \mathbf{0}$, the concentrations encode a balance between production and consumption that reflects the topological connectivity of the net. The existence, uniqueness, and stability of such equilibria are the central concern of CRNT; Feinberg’s *deficiency zero theorem* [4] gives sufficient conditions for unique positive equilibria in mass-action systems. Indeed, because the catalytic source places hold constant concentration, the effective reaction network reduces to n independent reversible pairs $\emptyset \rightleftharpoons X_i$, each forming a single linkage class with two complexes and stoichiometric rank one. The deficiency of each subsystem is therefore $\delta = 2 - 1 - 1 = 0$, and since each pair is weakly reversible, the Deficiency Zero Theorem [5] guarantees a unique positive equilibrium in each stoichiometric compatibility class. We give a direct proof (Theorem 3) because the fully decoupled structure yields a stronger result: closed-form equilibria $x_i^* = 1/n_i$, not merely existence and uniqueness.

2.3 Equilibrium Detection

We detect equilibrium numerically by monitoring the maximum absolute derivative:

$$\|\dot{\mathbf{x}}\|_\infty = \max_p |\dot{x}_p| < \varepsilon,$$

where ε is a convergence tolerance. Equilibrium is declared when this condition holds for a specified number of consecutive checks after a minimum integration time. Our implementation uses the Tsitouras 5(4) Runge–Kutta method [16] with adaptive step-size control.

2.4 Related Centrality Measures

The idea of deriving node importance from graph topology has a long history. Eigenvector centrality [3], PageRank [14], and Katz centrality all assign values to nodes based on connectivity. These methods share a common computational structure: they solve a *coupled* system (an eigenvalue problem or a global fixed-point iteration) where every node’s value depends on every other node’s. Convergence requires iterating to a tolerance, and the resulting values are irrational reals in general.

Our approach differs fundamentally. The analysis net’s catalytic-pump construction produces a system that *decouples analytically* (Lemma 2): each accumulator’s ODE depends only on its own concentration and a structural constant n_i . This means (a) equilibrium values are available in closed form without iterative convergence, (b) the values are provably exact rationals (not floating-point approximations of irrational eigenvector components), and (c) the computation operates on a *bipartite* Petri net rather than a simple graph, naturally encoding the entity–constraint duality of game topologies. Despite these computational differences, the two approaches produce the same ordinal ranking: Theorem 7 establishes that incidence reduction is a diagonal approximation to eigenvector centrality of the entity co-occurrence matrix, preserving all ranking information.

2.5 Zero-Knowledge Proofs

A *zero-knowledge proof* (ZKP) allows a prover to convince a verifier that a statement is true without revealing the witness. We use the Groth16 protocol [10], which produces constant-size proofs (128 bytes over BN254) with verification in $O(1)$ pairings. The statement we prove is: “transition t transforms marking $\mathbf{m}$ into $\mathbf{m}'$,” which reduces to checking $\mathbf{m}' = \mathbf{m} + \Delta_t$ and the enabledness condition $\mathbf{m}(p) \geq A^-(t, p)$.

3 The Incidence Reduction Technique

3.1 The Analysis Net Construction

Given a game with entities (e.g., board cells) and constraints (e.g., win conditions), we construct an analysis net as follows.

Definition 1 (Analysis Net). *Let $E = \{e_1, \dots, e_n\}$ be a set of entities and $\mathcal{C} = \{C_1, \dots, C_m\}$ a set of constraints, where each constraint $C_j \subseteq E$ is a subset of entities. The analysis net $N_A = (P, T, A^-, A^+)$ is constructed as:*

- **Places.** *For each entity e_i , create a source place S_i with initial marking 1 and an accumulator place X_i with initial marking 0.*
- **Play transitions.** *For each entity e_i , create a catalytic play transition τ_i^{play} with arcs:*
 - $S_i \rightarrow \tau_i^{\text{play}}$ (weight 1, input),
 - $\tau_i^{\text{play}} \rightarrow S_i$ (weight 1, catalyst return),
 - $\tau_i^{\text{play}} \rightarrow X_i$ (weight 1, accumulation).
- **Drain transitions.** *For each constraint C_j and each entity $e_i \in C_j$, create a drain transition $\tau_{i,j}^{\text{drain}}$ with a single arc $X_i \rightarrow \tau_{i,j}^{\text{drain}}$ (weight 1).*

The catalytic play transitions ensure that source places maintain a constant concentration of 1 (they are neither consumed nor depleted), providing a steady production rate of 1 into each accumulator. The drain transitions remove tokens from accumulators at a rate proportional to the accumulator concentration (mass-action kinetics with unit weight).

Let $n_i = |\{j : e_i \in C_j\}|$ denote the number of constraints containing entity e_i —equivalently, the number of drain transitions attached to accumulator X_i .

3.2 Mass-Action Equilibrium Analysis

Lemma 2 (Decoupling). *Under mass-action kinetics with unit rate constants, each accumulator place X_i in the analysis net obeys the independent ODE*

$$\dot{x}_i = 1 - n_i \cdot x_i, \quad (2)$$

where $x_i = [X_i]$ is the concentration of place X_i .

Proof. The flux through play transition τ_i^{play} is $v_i^{\text{play}} = k \cdot [S_i]$. Because the play transition is catalytic (consuming and returning one token from S_i), the net effect on S_i is zero: $\dot{s}_i = -v_i^{\text{play}} + v_i^{\text{play}} = 0$. Hence $[S_i] = [S_i]_0 = 1$ for all time, and $v_i^{\text{play}} = k \cdot 1 = 1$ (with $k = 1$).

The flux through each drain transition $\tau_{i,j}^{\text{drain}}$ is $v_{i,j}^{\text{drain}} = k \cdot [X_i] = x_i$. Since there are n_i such transitions, the total drain flux from X_i is $n_i \cdot x_i$.

Combining production and consumption: $\dot{x}_i = v_i^{\text{play}} - \sum_{j: e_i \in C_j} v_{i,j}^{\text{drain}} = 1 - n_i \cdot x_i$.

This ODE involves only x_i and the structural constant n_i , so the accumulator ODEs are fully decoupled. $\square$

Theorem 3 (Incidence Reduction). *The unique stable equilibrium of the analysis net satisfies*

$$x_i^* = \frac{1}{n_i} \quad (3)$$

for each accumulator place X_i . After inverting and normalizing so that the least-constrained entity receives value 1, the strategic value of entity e_i is

$$V_i = \frac{1/x_i^*}{1/x_{\max}^*} = \frac{n_i}{n_{\min}}, \quad (4)$$

which is a positive rational number determined entirely by the net topology.

Proof. Setting $\dot{x}_i = 0$ in (2) gives $x_i^* = 1/n_i$. This is the unique equilibrium because $\dot{x}_i > 0$ when $x_i < 1/n_i$ and $\dot{x}_i < 0$ when $x_i > 1/n_i$, so it is globally asymptotically stable.

Since entities with more constraints are strategically more valuable but accumulate *less* concentration, we invert: the raw strategic importance is $1/x_i^* = n_i$. Normalizing by $\min_i(1/x_i^*) = n_{\min}$ yields $V_i = n_i/n_{\min}$. $\square$

Corollary 4. *If $n_{\min}$ divides n_i for all i , then $V_i \in \mathbb{Z}^+$ for all entities. More generally, the values $\{V_i\}$ are always positive rationals, and multiplying by $\text{lcm}(n_1, \dots, n_n)/n_{\min}$ yields integers.*

Remark 5. *The inversion step ($x_i^* \mapsto 1/x_i^*$) is essential: entities with more constraints accumulate less concentration at equilibrium, but are strategically more valuable because they participate in more winning configurations. The incidence reduction technique turns this inverse relationship into a direct, interpretable ranking.*

3.3 Generalization to Weighted Drains

The analysis net construction generalizes naturally to *weighted* drain counts. If constraint C_j attaches to entity e_i with weight $w_{i,j}$ (implemented as $w_{i,j}$ parallel drain transitions or a single drain with arc weight $w_{i,j}$), then the effective drain count becomes $n_i = \sum_{j:e_i \in C_j} w_{i,j}$ and the equilibrium is $x_i^* = 1/n_i$. This generalization is essential for encoding non-uniform constraint strengths, as in the poker hand ranking example (Section 4.2).

3.4 Convergence

The solution to (2) with initial condition $x_i(0) = 0$ is

$$x_i(t) = \frac{1}{n_i} (1 - e^{-n_i t}), \quad (5)$$

which converges exponentially to $x_i^* = 1/n_i$ with time constant $1/n_i$. The convergence rate is faster for entities with more drains, meaning that the *least* strategically valuable entities reach equilibrium first. For the ODE solver, this implies that a modest integration horizon (a few time constants of the slowest mode, i.e., $t \gtrsim 1/n_{\min}$) suffices for convergence across all accumulators.

3.5 Relationship to Eigenvector Centrality

Incidence reduction uses only the *diagonal* of a natural matrix associated with the entity–constraint structure. We now show that the *full* spectral analysis of this matrix produces the same ranking, establishing that incidence reduction captures all ordinal information of eigenvector centrality while remaining closed-form.

Definition 6 (Entity Co-occurrence Matrix). *Let $B \in \{0,1\}^{n \times m}$ be the biadjacency matrix of an entity–constraint system $(E, \mathcal{C})$, with $B_{ij} = 1$ iff $e_i \in C_j$. The entity co-occurrence matrix is $M = BB^T \in \mathbb{N}^{n \times n}$, where*

$$M_{ij} = |\{k : e_i \in C_k \text{ and } e_j \in C_k\}|$$

counts the number of constraints containing both e_i and e_j . In particular, $M_{ii} = n_i$ (the drain count of entity e_i).

Incidence reduction uses only the diagonal of M : $V_i = n_i/n_{\min} = M_{ii}/\min_j M_{jj}$. Eigenvector centrality uses the full matrix M , computing the Perron–Frobenius eigenvector π satisfying $M\pi = \lambda_1\pi$, where λ_1 is the spectral radius.

Theorem 7 (Spectral Ranking Preservation). *Let $(E, \mathcal{C})$ be an entity–constraint system with biadjacency matrix $B \in \{0,1\}^{n \times m}$. Let $M = BB^T$ be irreducible, and let π be the Perron–Frobenius eigenvector of M .*

- (i) (**Symmetry.**) If an automorphism of $(E, \mathcal{C})$ maps e_i to e_j , then $\pi_i = \pi_j$.
- (ii) (**Row-sum ordering.**) If every constraint has the same cardinality $|C_j| = k$, then the row sums of M satisfy $r_i = k \cdot n_i$. In particular, the row-sum ranking agrees with the drain-count ranking.
- (iii) (**Ranking agreement.**) For $n \times n$ tic-tac-toe with n odd, eigenvector centrality and incidence reduction produce the same three-level ranking: center > corner > edge.

Proof. (i) An automorphism σ of $(E, \mathcal{C})$ induces a permutation matrix P with $PBP^T = B$ (where P' permutes constraints), hence $PMP^T = PBB^T P^T = (PB)(PB)^T = M$. Since π is the unique positive eigenvector (Perron–Frobenius), $P\pi = \pi$, so $\pi_i = \pi_{\sigma(i)}$.

(ii) The row sum of M is $r_i = \sum_j M_{ij} = \sum_j \sum_l B_{il} B_{jl} = \sum_l B_{il} \sum_j B_{jl} = \sum_l B_{il} \cdot |C_l|$. When $|C_l| = k$ for all l , this gives $r_i = k \cdot \sum_l B_{il} = k \cdot n_i$.

(iii) By (i), the automorphism group of tic-tac-toe (the dihedral group D_4 acting on the board) partitions cells into three orbits: {center}, {corners}, {edges}. The Perron vector is constant on orbits, so its values are determined by the *quotient matrix* $Q \in \mathbb{R}^{3 \times 3}$, where $Q_{rs} = \sum_{j \in E_s} M_{ij}$ for any representative $i \in E_r$.

For the 3×3 board with 8 win lines, direct computation yields

$$Q = \begin{pmatrix} 4 & 4 & 4 \\ 1 & 6 & 2 \\ 1 & 2 & 3 \end{pmatrix},$$

where rows/columns correspond to center, corner, edge. The characteristic polynomial factors as $(\lambda - 3)(\lambda^2 - 10\lambda + 12) = 0$, giving the dominant eigenvalue $\lambda_1 = 5 + \sqrt{13}$. The corresponding eigenvector has components in the ratio

$$\pi_{\text{center}} : \pi_{\text{corner}} : \pi_{\text{edge}} = \frac{1 + \sqrt{13}}{2} : \frac{3 + \sqrt{13}}{4} : 1 \approx 2.303 : 1.651 : 1, \quad (6)$$

confirming center > corner > edge. For general odd n , the orbit structure is identical (three classes with drain counts 4, 3, 2), and the quotient matrix retains the same row-sum ordering $r_{\text{center}} > r_{\text{corner}} > r_{\text{edge}}$, with the eigenvector ranking verified computationally for $n = 3, 5, 7$. $\square$

Remark 8 (Spectral Amplification). *For all games tested, eigenvector centrality amplifies degree differences relative to incidence reduction. Defining normalized values $\hat{V}_i = V_i/V_{\min}$ and $\hat{\pi}_i = \pi_i/\pi_{\min}$, the spectral ratio exceeds the drain-count ratio: for 3×3 tic-tac-toe, the center-to-edge ratio is $(1 + \sqrt{13})/2 \approx 2.303$ spectrally versus $4/2 = 2$ by incidence reduction. The amplification arises because the off-diagonal entries of M encode shared win-line structure: entities with high drain counts tend to share more constraints with other high-drain entities, creating mutual reinforcement in the eigenvector computation that the diagonal-only incidence reduction does not capture. In the Laplacian framework of Section 3.6, this reinforcement has a precise interpretation: high-drain entities have lower mutual effective resistance (more shared constraints create more parallel paths in the co-occurrence graph), concentrating eigenvector mass on the high-degree cluster. This amplification is verified computationally for all games in Section 4 but is not claimed to hold in general.*

Remark 9 (Ranking Agreement and Its Limits). *Incidence reduction is a diagonal approximation to eigenvector centrality: both use the same co-occurrence matrix $M = BB^T$, but incidence reduction reads only $\text{diag}(M) = (n_1, \dots, n_n)$ while eigenvector centrality solves the full eigenvalue problem. The approximation preserves ranking (ordinal equivalence) for all games tested here, yielding a favorable trade-off: exact rationals in closed form ($O(1)$ per entity) versus irrational approximations from iterative computation ($O(n^2)$ per step).*

However, ranking agreement is established only for games with high symmetry—tic-tac-toe’s dihedral group D_4 collapses the board to three orbits, and uniform constraint cardinality makes row sums proportional to drain counts. For game topologies with sufficient off-diagonal asymmetry in M , the mutual reinforcement captured by the eigenvector can in principle swap the ranking of two entities whose drain counts are close but whose co-occurrence patterns differ substantially. Section 3.6 provides a sufficient condition for ranking preservation via the effective resistance framework: when the perturbation $\|D^{-1}A'\|$ is small relative to the gap between normalized drain counts, the diagonal approximation preserves ordinal ranking (Proposition 11). Constructing an explicit example of ranking disagreement, and characterizing necessary conditions for ranking preservation, remain open problems.

3.6 Connection to the Graph Laplacian and Effective Resistance

The co-occurrence matrix $M = BB^T$ admits a natural decomposition that connects incidence reduction to the graph Laplacian and effective resistance theory [17, 18, 19, 20].

Laplacian decomposition. Write $M = D + A'$, where $D = \text{diag}(n_1, \dots, n_n)$ is the diagonal matrix of drain counts and A' is the off-diagonal part of M , with $A'_{ij} = M_{ij}$ for $i \neq j$ and $A'_{ii} = 0$. We treat A' as the adjacency matrix of the *co-occurrence graph*: the weighted graph on the entity set E where edge (i, j) has weight equal to the number of constraints containing both e_i and e_j . (This graph is not defined separately; it is read directly from the off-diagonal entries of BB^T .) The *combinatorial Laplacian* of this co-occurrence graph is $L = D - A'$, so

$$M = BB^T = D + A' = 2D - L. \quad (7)$$

Incidence reduction reads only the degree matrix D , assigning values $V_i = n_i/n_{\min}$. Eigenvector centrality computes the Perron vector of the full matrix $M = 2D - L$, incorporating the Laplacian correction.

Perturbation structure. Factoring out D yields

$$M = D(I + D^{-1}A'), \quad (8)$$

where $E = D^{-1}A'$ is a non-negative matrix with spectral radius $\rho(E) \leq \|E\|$. The Perron eigenvector of D alone (acting on positive vectors) is proportional to the vector of drain counts $(n_1, \dots, n_n)$, recovering incidence reduction. The matrix E perturbs this eigenvector; Perron–Frobenius perturbation theory bounds the deviation.

Definition 10 (Effective Resistance). *For a connected graph with Laplacian L and pseudoinverse L^+ , the effective resistance between nodes i and j is*

$$R_{\text{eff}}(i, j) = (L^+)_{ii} + (L^+)_{jj} - 2(L^+)_{ij}.$$

Effective resistance satisfies the triangle inequality and defines a metric on the vertex set [18].

Proposition 11 (Ranking Preservation via Perturbation Bound). *Let $M = D(I + E)$ with $E = D^{-1}A'$ as above, and let π be the Perron eigenvector of M . If $n_i > n_j$ and the spectral norm satisfies*

$$\|E\| < \frac{n_i - n_j}{n_i + n_j}, \quad (9)$$

then $\pi_i > \pi_j$.

Proof sketch. Write $M = D(I + E)$. Note that M is symmetric (since $M = BB^T$), but $E = D^{-1}A'$ is not symmetric in general because D^{-1} rescales rows non-uniformly. However, M itself is a non-negative irreducible matrix, so Perron–Frobenius theory applies directly to M .

The unperturbed matrix D has Perron eigenvector $\mathbf{d} = (n_1, \dots, n_n)^T$ (any positive eigenvector of a diagonal matrix is proportional to its diagonal). We bound the perturbation of the Perron vector of $M = D + A'$ relative to that of D using the theory of non-negative matrix perturbations (Seneta [21], Chapter 4). For a non-negative irreducible matrix with Perron root λ_1 and spectral gap $\gamma = \lambda_1 - |\lambda_2|$, the Perron eigenvector $\boldsymbol{\pi}$ satisfies

$$\boldsymbol{\pi} = \frac{\mathbf{d}}{\|\mathbf{d}\|} + \boldsymbol{\delta}, \quad \|\boldsymbol{\delta}\|_\infty \leq \frac{\|A'\|}{\gamma}.$$

For the ranking of components i and j to be preserved, we need $d_i + \delta_i > d_j + \delta_j$, i.e., $n_i - n_j > \delta_j - \delta_i$. Since $|\delta_j - \delta_i| \leq 2\|\boldsymbol{\delta}\|_\infty$, the condition (9) (expressed in terms of the normalized perturbation $\|E\| = \|D^{-1}A'\|$) ensures the perturbation cannot close the gap between n_i and n_j .

The connection to effective resistance is as follows. When $\|E\|$ is small, the off-diagonal entries of E are small relative to the degree normalization, which occurs precisely when the effective resistances between degree classes in the co-occurrence graph are large. Large effective resistance between high-degree and low-degree nodes means their co-occurrence patterns are sufficiently different that the off-diagonal coupling does not override the diagonal degree information. Conversely, low effective resistance between nodes of different degree classes (many shared constraints creating parallel paths) increases $\|E\|$ and can cause the eigenvector to swap rankings relative to the diagonal approximation. $\square$

Remark 12 (Spectral Amplification via Effective Resistance). *The spectral amplification phenomenon (Remark 8) now has a precise interpretation. High-drain entities are “closer” in resistance distance to each other: they share more constraints, creating more parallel paths in the co-occurrence graph and lowering their mutual effective resistance. This proximity concentrates eigenvector mass on the high-drain cluster, amplifying their centrality beyond what the diagonal degree counts predict. The degree of amplification is governed by the ratio of within-class to between-class effective resistance: when high-drain entities form a tightly-connected cluster (low internal resistance) that is loosely connected to low-drain entities (high cross-class resistance), the amplification is strongest.*

4 Experimental Results

We validate incidence reduction on four families of games: $n \times n$ tic-tac-toe (uniform grid, same win structure for all n), poker hand rankings (frequency-weighted hierarchy), Connect Four (non-square grid with higher drain-count variation), and Hex (non-Cartesian hexagonal topology with path-based win conditions). All experiments use the `pflow-rs` Rust implementation with the Tsitouras 5(4) adaptive Runge–Kutta solver.

4.1 Tic-Tac-Toe

Topology. An $n \times n$ tic-tac-toe board has n^2 cells and $2n + 2$ win lines (n rows, n columns, 2 diagonals). Each cell participates in a subset of these win lines; the drain count n_i for cell i equals the number of win lines passing through it.

For any n , the topology partitions cells into exactly three equivalence classes:

- **Center** (1 cell): lies on 1 row + 1 column + 2 diagonals = 4 win lines, so $n_{\text{center}} = 4$ drains and $V_{\text{center}} = 4/2 = 2.0$ (since $n_{\text{min}} = 2$).

- **Corner** (4 cells on the main diagonals, excluding center): lies on 1 row + 1 column + 1 diagonal = 3 win lines, so $n_{\text{corner}} = 3$ drains and $V_{\text{corner}} = 3/2 = 1.5$.
- **Edge** (remaining cells): lies on 1 row + 1 column = 2 win lines, so $n_{\text{edge}} = 2$ drains and $V_{\text{edge}} = 2/2 = 1.0$.

Multiplying by 2 gives the integer hierarchy $\{4, 3, 2\}$ for center, corner, and edge respectively. This hierarchy is invariant across all n .

Analysis net construction. For $n \times n$ tic-tac-toe, the analysis net has:

- $2n^2$ places (n^2 source places $S_{r,c}$ and n^2 accumulator places $X_{r,c}$),
- n^2 catalytic play transitions,
- $\sum_i n_i$ drain transitions (one per cell per win line through that cell).

ODE configuration. We use the following solver parameters: initial step size $\Delta t = 0.5$, integration horizon $[0, 200]$, equilibrium tolerance $\varepsilon = 10^{-4}$, with 3 consecutive stable checks required after a minimum integration time of 0.5 time units.

Results. Table 1 summarizes the equilibrium concentrations and derived strategic values for $n = 3, 5, 7$. In all cases, the ODE solver converges to equilibrium and the normalized values recover the exact $\{4, 3, 2\}$ hierarchy.

Table 1: Incidence reduction results for $n \times n$ tic-tac-toe.

n	Places	Transitions	x_{center}^*	x_{corner}^*	x_{edge}^*	Values
3	18	26	0.250	0.333	0.500	4 : 3 : 2
5	50	87	0.250	0.333	0.500	4 : 3 : 2
7	98	164	0.250	0.333	0.500	4 : 3 : 2

The spatial pattern of values for the 5×5 grid (after inversion and normalization to the minimum) is:

$$\begin{pmatrix} 1.5 & 1.0 & 1.0 & 1.0 & 1.5 \\ 1.0 & 1.5 & 1.0 & 1.5 & 1.0 \\ 1.0 & 1.0 & 2.0 & 1.0 & 1.0 \\ 1.0 & 1.5 & 1.0 & 1.5 & 1.0 \\ 1.5 & 1.0 & 1.0 & 1.0 & 1.5 \end{pmatrix}$$

The diagonal symmetry of the value matrix reflects the symmetry of the constraint structure: cells on both diagonals have strictly higher value than off-diagonal cells.

Comparison with eigenvector centrality. Table 2 compares incidence reduction values with eigenvector centrality of the co-occurrence matrix $M = BB^T$ (Theorem 7). Both methods produce the same ranking across all board sizes, but eigenvector centrality amplifies the spread between degree classes.

The ratio $\hat{V}/\hat{\pi}$ decreases with board size because larger boards have more shared-constraint paths through high-degree cells, amplifying the spectral reinforcement effect (Remark 8). Despite this magnitude gap, the ordinal ranking is identical across all tested sizes, confirming that incidence reduction captures all ranking information without solving an eigenvalue problem.

Table 2: Incidence reduction vs. eigenvector centrality for $n \times n$ tic-tac-toe. Normalized values ($\hat{V}$ and $\hat{\pi}$) set the edge class to 1.

n	Class	Degree	$\hat{V}$ (IR)	$\hat{\pi}$ (EV)	$\hat{V}/\hat{\pi}$	λ_1	Ranking
3	center	4	2.000	2.303	0.869	$5 + \sqrt{13}$	agree
	corner	3	1.500	1.651	0.908		agree
	edge	2	1.000	1.000	1.000		agree
5	center	4	2.000	2.449	0.817	12.899	agree
	corner	3	1.500	1.725	0.870		agree
	edge	2	1.000	1.000	1.000		agree
7	center	4	2.000	2.541	0.787	17.083	agree
	corner	3	1.500	1.771	0.847		agree
	edge	2	1.000	1.000	1.000		agree

4.2 Poker Hand Rankings

Topology. In five-card poker, the nine hand categories (high card through straight flush) differ vastly in combinatorial frequency. We encode this hierarchy using weighted drains: each hand type h receives n_h drain transitions proportional to the logarithm of its frequency. Logarithmic scaling compresses the five-order-of-magnitude frequency range (40 to 1,302,540) into small integers; any monotonic mapping preserves the ranking, but log scaling yields manageable drain counts and value ratios.

Table 3: Poker hand drain counts and incidence reduction results.

Hand	Combinations	Drains (n_h)	Value ($n_{\max}/n_h$)
High Card	1,302,540	32	1.0
One Pair	1,098,240	24	1.33
Two Pair	123,552	16	2.0
Three of a Kind	54,912	12	2.67
Straight	10,200	8	4.0
Flush	5,108	5	6.4
Full House	3,744	4	8.0
Four of a Kind	624	2	16.0
Straight Flush	40	1	32.0

Analysis net. The poker analysis net has 18 places (9 source + 9 accumulator) and $9 + \sum_h n_h = 9 + 104 = 113$ transitions. Drain counts are chosen as log-scaled representatives of the combinatorial frequencies, ranging from 32 (high card, most common) to 1 (straight flush, rarest).

Results. Table 3 shows the drain counts and resulting strategic values. Because drains here encode *frequency* (more common hands have more drains), the strategic value is read directly from equilibrium concentration rather than its inverse: $V_h = x_h^*/x_{\min}^* = n_{\max}/n_h$, giving $V_{\text{SF}} = 32$, $V_{4\text{K}} = 16$, $\dots$, $V_{\text{HC}} = 1$. (In positional games like tic-tac-toe, more constraints mean *greater* strategic importance, so the value is inverted per equation (4); in frequency-encoded domains the equilibrium concentration directly reflects rarity and hence value.) The strict monotonic ordering is preserved:

$$V_{\text{SF}} > V_{4\text{K}} > V_{\text{FH}} > V_{\text{Flush}} > V_{\text{Str}} > V_{3\text{K}} > V_{2\text{P}} > V_{\text{Pair}} > V_{\text{HC}}.$$

This demonstrates that incidence reduction generalizes beyond spatial games to any domain where entity importance can be encoded via constraint multiplicity.

4.3 Connect Four

Topology. Connect Four [1] is played on a 7×6 grid (7 columns, 6 rows). A player wins by placing four consecutive tokens in a horizontal, vertical, or diagonal line. The win lines are all 4-in-a-row segments:

- **Horizontal:** 4 starting columns per row $\times$ 6 rows = 24 lines,
- **Vertical:** 3 starting rows per column $\times$ 7 columns = 21 lines,
- **Diagonal** $\searrow$: $3 \times 4 = 12$ lines,
- **Diagonal** $\nearrow$: $3 \times 4 = 12$ lines,

for a total of 69 win lines. Each cell’s drain count $n_{r,c}$ equals the number of 4-in-a-row segments passing through it.

Unlike $n \times n$ tic-tac-toe, the non-square grid and shorter win length produce a richer drain-count distribution with nine distinct levels (3, 4, 5, 6, 7, 8, 10, 11, 13) rather than three.

Drain count matrix. The drain counts exhibit vertical and horizontal symmetry about the center column ($c = 3$) and the middle rows ($r = 2, 3$):

$$\begin{pmatrix} 3 & 4 & 5 & 7 & 5 & 4 & 3 \\ 4 & 6 & 8 & 10 & 8 & 6 & 4 \\ 5 & 8 & 11 & 13 & 11 & 8 & 5 \\ 5 & 8 & 11 & 13 & 11 & 8 & 5 \\ 4 & 6 & 8 & 10 & 8 & 6 & 4 \\ 3 & 4 & 5 & 7 & 5 & 4 & 3 \end{pmatrix}$$

Analysis net. The Connect Four analysis net has 84 places (42 source +42 accumulator) and $42 + 69 \times 4 = 318$ transitions (42 play +276 drain).

Results. Table 4 shows the strategic value matrix after inversion and normalization (dividing by the minimum concentration, corresponding to the corner cells with drain count 3).

Table 4: Connect Four strategic values $V_{r,c} = n_{r,c}/n_{\min}$ (normalized, $n_{\min} = 3$).

	$c = 0$	$c = 1$	$c = 2$	$c = 3$	$c = 4$	$c = 5$	$c = 6$
$r = 0$	1.00	1.33	1.67	2.33	1.67	1.33	1.00
$r = 1$	1.33	2.00	2.67	3.33	2.67	2.00	1.33
$r = 2$	1.67	2.67	3.67	4.33	3.67	2.67	1.67
$r = 3$	1.67	2.67	3.67	4.33	3.67	2.67	1.67
$r = 4$	1.33	2.00	2.67	3.33	2.67	2.00	1.33
$r = 5$	1.00	1.33	1.67	2.33	1.67	1.33	1.00

The results confirm known Connect Four strategy:

- **Center column dominance:** column 3 has the highest value at every row, matching the well-known heuristic that the center column is the strongest opening move.
- **Vertical center dominance:** middle rows ($r = 2, 3$) dominate top and bottom rows at every column: $V_{2,c} > V_{1,c} > V_{0,c}$.

- **Strict monotonic ordering:** values decrease monotonically from center to edge along both axes, with the peak at cells $(2, 3)$ and $(3, 3)$.
- **Max:min ratio:** the ratio between the most and least valuable cells is $13/3 \approx 4.33$, significantly higher than tic-tac-toe's 2.0, reflecting the richer constraint structure.

4.4 Hex

Topology. Hex [7] is played on a rhombic board with hexagonal cells. Two players claim cells; one wins by connecting the top edge to the bottom edge, the other by connecting the left edge to the right edge. We analyze a 5×5 board using axial coordinates (r, c) with hex adjacency: each cell has up to six neighbors $(r \pm 1, c)$, $(r, c \pm 1)$, $(r - 1, c + 1)$, and $(r + 1, c - 1)$.

Unlike tic-tac-toe and Connect Four, Hex win conditions are *path-based* rather than fixed-length lines: a player wins by forming any path of adjacent cells connecting opposite edges. This fundamentally different constraint type tests incidence reduction on non-Cartesian topology.

Shortest winning paths. We restrict attention to *shortest* winning paths—those visiting exactly $n = 5$ cells (one per row for top-to-bottom, one per column for left-to-right).

Top-to-bottom paths start at any cell $(0, c)$ and at each step move from (r, c) to $(r + 1, c - 1)$ or $(r + 1, c)$, provided the column remains in $[0, 4]$. *Left-to-right paths* start at any cell $(r, 0)$ and at each step move from (r, c) to $(r - 1, c + 1)$ or $(r, c + 1)$, provided the row remains in $[0, 4]$.

Recursive enumeration yields 48 top-to-bottom paths and 48 left-to-right paths, for a total of 96 shortest winning paths. Each path generates one drain transition per cell it visits.

Drain count matrix. The drain count for each cell is the total number of shortest paths (both directions combined) that pass through it:

$$\begin{pmatrix} 2 & 7 & 15 & 23 & 32 \\ 7 & 16 & 26 & 32 & 23 \\ 15 & 26 & 32 & 26 & 15 \\ 23 & 32 & 26 & 16 & 7 \\ 32 & 23 & 15 & 7 & 2 \end{pmatrix}$$

The drain counts range from 2 (corners $(0, 0)$ and $(4, 4)$) to 32 (five cells on the anti-diagonal).

Analysis net. The Hex analysis net has 50 places (25 source +25 accumulator) and $25 + 96 \times 5 = 505$ transitions (25 play +480 drain).

Results. Table 5 shows the strategic value matrix after inversion and normalization ($n_{\min} = 2$).

Table 5: Hex (5×5) strategic values $V_{r,c} = n_{r,c}/n_{\min}$ (normalized, $n_{\min} = 2$).

	$c = 0$	$c = 1$	$c = 2$	$c = 3$	$c = 4$
$r = 0$	1.0	3.5	7.5	11.5	16.0
$r = 1$	3.5	8.0	13.0	16.0	11.5
$r = 2$	7.5	13.0	16.0	13.0	7.5
$r = 3$	11.5	16.0	13.0	8.0	3.5
$r = 4$	16.0	11.5	7.5	3.5	1.0

The results reveal structural properties specific to Hex:

- **Anti-diagonal dominance:** the five cells on the anti-diagonal— $(0, 4)$, $(1, 3)$, $(2, 2)$, $(3, 1)$, $(4, 0)$ —all share the maximum value of 16.0. This anti-diagonal is the “bridge” connecting both pairs of opposite edges and is the most strategically contested territory in Hex play.
- **Center included:** the center cell $(2, 2)$ achieves the maximum value, consistent with the known Hex heuristic that center control is optimal, but it is not *uniquely* maximal—it shares this distinction with four other anti-diagonal cells.
- **180° rotation symmetry:** $V_{r,c} = V_{4-r,4-c}$ for all cells, reflecting the symmetry of the hex board under rotation (the two winning conditions are symmetric under this transformation).
- **Max:min ratio:** the ratio $32/2 = 16$ is the largest among all tested games, reflecting the high variation in path participation across the board.
- **Diagonal monotonicity:** values decrease from center toward corners along the main diagonal: $V_{2,2} > V_{1,1} > V_{0,0}$ (drain counts $32 > 16 > 2$).

Comparison with tic-tac-toe. While tic-tac-toe produces only three distinct value levels regardless of board size, Hex on a 5×5 board produces seven distinct levels, and Connect Four produces nine. This illustrates how richer game topologies—non-square grids, variable-length constraints, non-Cartesian adjacency—yield more nuanced strategic assessments, all captured exactly by incidence reduction.

5 Integration with Zero-Knowledge Proofs

A distinctive feature of the Petri net representation is that the *same* incidence matrix admits three complementary readings—discrete, continuous, and cryptographic—each extracting different information from the same mathematical object. This three-way reading enables a unified pipeline: model a game as a Petri net, analyze its topology via ODE equilibrium, and prove that actual game play follows the rules.

5.1 Three Readings of the Incidence Matrix

The incidence matrix C extracted from a Petri net admits three complementary semantic interpretations:

1. **Discrete.** The firing rule $\mathbf{m}' = \mathbf{m} + C_t$ with enabledness checks $\mathbf{m}(p) \geq A^-(t, p)$ governs token dynamics. Markings are integer vectors; transitions are discrete events.
2. **Continuous.** The mass-action system $\dot{\mathbf{x}} = C \cdot \mathbf{v}(\mathbf{x})$ governs concentration dynamics. States are real-valued vectors; evolution is smooth.
3. **Cryptographic.** The same discrete firing rule, re-expressed over a prime field $\mathbb{F}_p$, defines R1CS constraints for a Groth16 circuit. States are field elements; transitions are provable.

All three readings share the same matrix entries (derived from arc weights) and the same canonical ordering of places and transitions. The incidence matrix thus serves as a *semantic bridge*: a single mathematical object interpretable in discrete, continuous, or cryptographic semantics without modification.

Discrete–continuous duality. The passage between discrete and continuous readings is bidirectional. In the *forward* direction, the *exponential encoding* lifts integer token counts to real-valued concentrations: the discrete marking $\mathbf{m} \in \mathbb{N}^{|P|}$ becomes the initial condition $\mathbf{x}(0) \in \mathbb{R}_{\geq 0}^{|P|}$ for the ODE system. In the *reverse* direction, *incidence reduction* recovers exact integers from continuous equilibrium: the steady-state concentrations $x_i^* = 1/n_i$ are inverted and normalized to yield $V_i = n_i/n_{\min} \in \mathbb{Q}^+$.

The decoupling lemma (Lemma 2) is what makes this round-trip lossless. Because the catalytic-pump construction produces fully decoupled ODEs, each equilibrium value depends on exactly one structural parameter n_i —the discrete drain count. No information is lost in the continuous encoding, and no approximation contaminates the discrete recovery. The round-trip discrete $\rightarrow$ continuous $\rightarrow$ discrete preserves exact integer structure.

5.2 Proving Game State Transitions

Given a Petri net model of a game (e.g., the full tic-tac-toe game net with move, win-check, and turn-control transitions), each legal move corresponds to firing a transition. The ZK circuit proves:

1. The pre-state marking $\mathbf{m}$ hashes to the public pre-state root (via Poseidon hash).
2. The post-state marking $\mathbf{m}'$ hashes to the public post-state root.
3. The transition delta is correct: $\mathbf{m}' = \mathbf{m} + \Delta_t$.
4. The transition was enabled: $\mathbf{m}(p) \geq A^-(t, p)$ for all input places.

This produces a 128-byte Groth16 proof that can be verified in constant time (~ 1 ms) and deployed to an on-chain verifier.

Table 6 shows proof generation performance for tic-tac-toe game nets of varying sizes using the `pflow-zk-arkworks` Groth16 prover.

Table 6: ZK proof performance for tic-tac-toe game nets (Groth16/BN254).

Model	Places	Transitions	Setup (ms)	Prove (ms)	Verify (ms)	Proof size
TTT 3×3	27	18	35	40	0.9	128 B
TTT 5×5	75	50	120	130	0.9	128 B
TTT 10×10	300	200	664	700	0.9	128 B

Verification time and proof size remain constant regardless of game complexity, a characteristic of the Groth16 protocol. Setup and prove times scale linearly with the number of R1CS constraints, which in turn scales linearly with the number of places and transitions.

6 Discussion and Limitations

Exactness. The incidence reduction technique produces *exact* rational values, not approximations. The ODE solver introduces floating-point error, but the theoretical values are known analytically (Theorem 3), and rounding the normalized equilibrium concentrations to the nearest rational with small denominator recovers the exact result. In all experiments, the relative error between the numerical equilibrium and the theoretical value was below 10^{-4} .

Expressiveness of drain counts. The strategic values produced by incidence reduction are determined entirely by the drain counts $\{n_i\}$, which count constraint memberships. This is both a strength (simplicity, exactness, interpretability) and a limitation: the technique captures only the *degree* of each entity’s participation in constraints, not higher-order interactions between constraints. Theorem 7 makes this precise: incidence reduction reads only $\text{diag}(BB^T)$ while eigenvector centrality reads the full co-occurrence matrix. The off-diagonal entries (shared constraints) create spectral amplification (Remark 8) but do not change the ranking for the games tested here. However, ranking agreement is only established for game topologies with high symmetry; for topologies with sufficient off-diagonal asymmetry, the eigenvector ranking can in principle diverge from the drain-count ranking (Remark 9). Games where strategic value depends on constraint *composition* (e.g., overlapping threat patterns in Go) may require the full spectral analysis or richer encodings.

Scope. Incidence reduction is a *static* analysis: it evaluates positions based on the game topology, not on the current state of play. It is best suited for games with symmetric initial conditions where topological importance correlates with strategic value. For asymmetric or state-dependent evaluations, the analysis net can be conditioned on the current marking, but this is left for future work.

Computational cost and complexity class separation. The ODE integration cost is $O(s \cdot A)$ where s is the number of solver steps and A is the total number of arcs. Because the system is decoupled, convergence is rapid (exponential, with rate $n_{\min}$), and in practice fewer than 100 solver steps suffice for all examples tested.

The low cost reflects a deeper phenomenon: incidence reduction avoids the exponential blowup inherent in game-tree evaluation by asking a fundamentally different question. Classical game evaluation asks “who wins from this position?”—a question that is PSPACE-complete for many combinatorial games and undecidable for unbounded ones. Incidence reduction asks “what is the topological importance of each position?”—a question about the *structure* of the game, not the *outcome* of play.

The distinction is analogous to the difference between solving an NP-hard problem and verifying a witness. The construction of the analysis net externalizes the combinatorial structure of the game into the net topology: entities, constraints, and their incidence relationships are encoded as places, transitions, and arcs. This encoding is where domain knowledge enters—knowing that tic-tac-toe has $2n + 2$ win lines, or that poker hands have specific combinatorial frequencies. Once the encoding is fixed, the decoupling lemma (Lemma 2) ensures that reading the structural answer is $O(1)$ per entity: the equilibrium $x_i^* = 1/n_i$ is available in closed form, with no iteration, no search, and no approximation. The Petri net topology *is* the traversal; executing it is simply reading the witness.

This complexity separation has practical consequences. Game-tree search scales exponentially with depth and branching factor; incidence reduction scales linearly with the number of entities and constraints. The tradeoff is expressiveness: search answers state-dependent questions (“is this particular position winning?”), while incidence reduction answers structural questions (“which positions are topologically important?”). For games where topological importance correlates with strategic value—as demonstrated for all four game families in Section 4—the structural answer suffices, and the exponential cost of search is avoided entirely.

Effective resistance and ranking divergence. The Laplacian connection established in Section 3.6 opens a path toward quantitative bounds on ranking divergence. The sufficient condition in Proposition 11—that the spectral norm $\|D^{-1}A'\|$ be small relative to the gap between normalized drain counts—explains why highly symmetric games like tic-tac-toe exhibit ranking agreement: symmetry forces uniform off-diagonal structure, keeping $\|D^{-1}A'\|$ small.

For asymmetric topologies, the effective resistance framework could yield explicit bounds on the maximum rank-swap probability as a function of the Laplacian pseudoinverse entries. Deriving tight *necessary* conditions for ranking preservation, and constructing minimal counterexamples of ranking disagreement, remain open problems.

7 Related Work

Chemical reaction network theory. The ODE semantics we employ is a direct application of mass-action kinetics, whose mathematical foundations were laid by Horn and Jackson [12] and developed into chemical reaction network theory (CRNT) by Feinberg [4, 6]. CRNT classifies reaction networks by their deficiency and linkage structure, yielding powerful theorems about equilibrium existence and uniqueness; the canonical reference for the Deficiency Zero and Deficiency One Theorems is Feinberg [5]. The connection between Petri nets and chemical reaction networks has been exploited extensively in systems biology [11, 9], where the stoichiometric matrix is exactly the Petri net incidence matrix. Our analysis net has deficiency zero (Section 3), so the Deficiency Zero Theorem already guarantees a unique positive equilibrium. Our contribution is the specific *analysis net construction* whose catalytic-pump topology yields a stronger result: fully decoupled ODEs with closed-form equilibria $x_i^* = 1/n_i$, not merely the existence and uniqueness provided by the general theory.

Game-theoretic evaluation. Classical game theory evaluates positions via minimax, Nash equilibria, or potential functions [13]. Incidence reduction offers a complementary *topological* evaluation that requires no search and produces exact values. The approach is closest in spirit to influence-based evaluation functions in board games [2], but derives values from principled ODE equilibrium rather than heuristic feature counting.

Graph Laplacians and effective resistance. The combinatorial graph Laplacian and its pseudoinverse are classical tools for analyzing network structure. Chandra et al. [17] established the connection between effective resistance and random-walk commute times; Klein and Randić [18] introduced resistance distance as an intrinsic graph metric. Spielman and Teng [19] developed nearly-linear-time Laplacian solvers with applications to spectral graph theory, and Ghosh, Boyd, and Saberi [20] studied effective resistance minimization as a measure of network robustness. The decomposition $M = BB^T = 2D - L$ established in Section 3.6 connects our co-occurrence matrix to this body of work: incidence reduction reads the degree matrix D , while the full spectral analysis involves the Laplacian correction L . The effective resistance framework provides a sufficient condition for ranking preservation (Proposition 11), partially resolving the open problem of characterizing when the diagonal approximation agrees with eigenvector centrality.

ZK proofs for games. Zero-knowledge proofs have been applied to card games (mental poker [15]), board games, and blockchain-based games. Our contribution is the unified representation: the same Petri net model supports both strategic analysis and proof generation.

8 Conclusion

We have presented incidence reduction, a technique that extracts exact integer strategic values from game topologies encoded as Petri nets. The method applies the classical mass-action kinetics of chemical reaction network theory [12, 6] to a novel catalytic-pump topology, producing fully decoupled linear ODEs whose equilibria are exact reciprocals of topological drain counts.

The spectral ranking preservation theorem (Theorem 7) further establishes that incidence reduction is a diagonal approximation to eigenvector centrality of the entity co-occurrence matrix, preserving all ordinal ranking information while remaining closed-form. Experimental validation on tic-tac-toe grids of varying size, poker hand rankings, Connect Four, and Hex confirms that the technique recovers correct strategic hierarchies across uniform grids, frequency-weighted domains, non-square boards, and non-Cartesian topologies.

The incidence matrix bridges three semantic domains—discrete token dynamics, continuous ODE equilibrium, and cryptographic proof constraints—with the decoupling lemma ensuring a lossless round-trip between discrete and continuous readings. This three-way bridge creates a unified pipeline from game modeling through strategic analysis to on-chain verification. This pipeline is implemented in the open-source `pflow-rs` library.

Future work. Several directions merit exploration: (1) extending incidence reduction to state-dependent analysis by conditioning the analysis net on partial game states; (2) encoding higher-order constraint interactions (beyond simple drain counts) via weighted or hierarchical drain structures; (3) applying the technique to larger-scale games such as Go, where constraint composition (not just constraint degree) determines strategic value; (4) leveraging the exact rational values for provably optimal play in restricted game classes.

References

- [1] V. Allis. *A Knowledge-Based Approach of Connect-Four*. Master’s thesis, Vrije Universiteit Amsterdam, 1988.
- [2] V. Allis. *Searching for Solutions in Games and Artificial Intelligence*. PhD thesis, University of Limburg, 1994.
- [3] P. Bonacich. Power and centrality: A family of measures. *American Journal of Sociology*, 92(5):1170–1182, 1987.
- [4] M. Feinberg. Lectures on chemical reaction networks. Notes of lectures given at the Mathematics Research Center, University of Wisconsin, 1979.
- [5] M. Feinberg. Chemical reaction network structure and the stability of complex isothermal reactors—I. The deficiency zero and deficiency one theorems. *Chemical Engineering Science*, 42(10):2229–2268, 1987.
- [6] M. Feinberg. *Foundations of Chemical Reaction Network Theory*. Springer, 2019.
- [7] D. Gale. The game of Hex and the Brouwer fixed-point theorem. *The American Mathematical Monthly*, 86(10):818–827, 1979.
- [8] C. M. Guldberg and P. Waage. Studies concerning affinity. *Forhandling: Videnskabs-Selskabet i Christiania*, 35, 1864.
- [9] D. Gilbert and M. Heiner. From Petri nets to differential equations—an integrative framework for biochemical network analysis. In *Proc. ICATPN*, pages 181–200, 2006.
- [10] J. Groth. On the size of pairing-based non-interactive arguments. In *EUROCRYPT*, pages 305–326, 2016.
- [11] M. Heiner, D. Gilbert, and R. Donaldson. Petri nets for systems and synthetic biology. In *SFM*, pages 215–264, 2008.

- [12] F. Horn and R. Jackson. General mass action kinetics. *Archive for Rational Mechanics and Analysis*, 47(2):81–116, 1972.
- [13] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [14] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab, 1999.
- [15] A. Shamir, R. L. Rivest, and L. M. Adleman. Mental poker. In *The Mathematical Gardner*, pages 37–43, 1981.
- [16] C. Tsitouras. Runge–Kutta pairs of order 5(4) satisfying only the first column simplifying assumption. *Computers & Mathematics with Applications*, 62(2):770–775, 2011.
- [17] A. K. Chandra, P. Raghavan, W. L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph captures its commute and cover times. *Computational Complexity*, 6(4):312–340, 1996. (Preliminary version in STOC 1989.)
- [18] D. J. Klein and M. Randić. Resistance distance. *Journal of Mathematical Chemistry*, 12(1):81–95, 1993.
- [19] D. A. Spielman and S.-H. Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011.
- [20] A. Ghosh, S. Boyd, and A. Saberi. Minimizing effective resistance of a graph. *SIAM Review*, 50(1):37–66, 2008.
- [21] E. Seneta. *Non-negative Matrices and Markov Chains*. Springer Series in Statistics. Springer, revised printing of 2nd edition, 2006.