Indirect Reciprocity and Strategic Agents I

Raúl Landa
Networks and Services Research Laboratory
Department of Electronic and Electrical Engineering
University College London

Overlay Networks

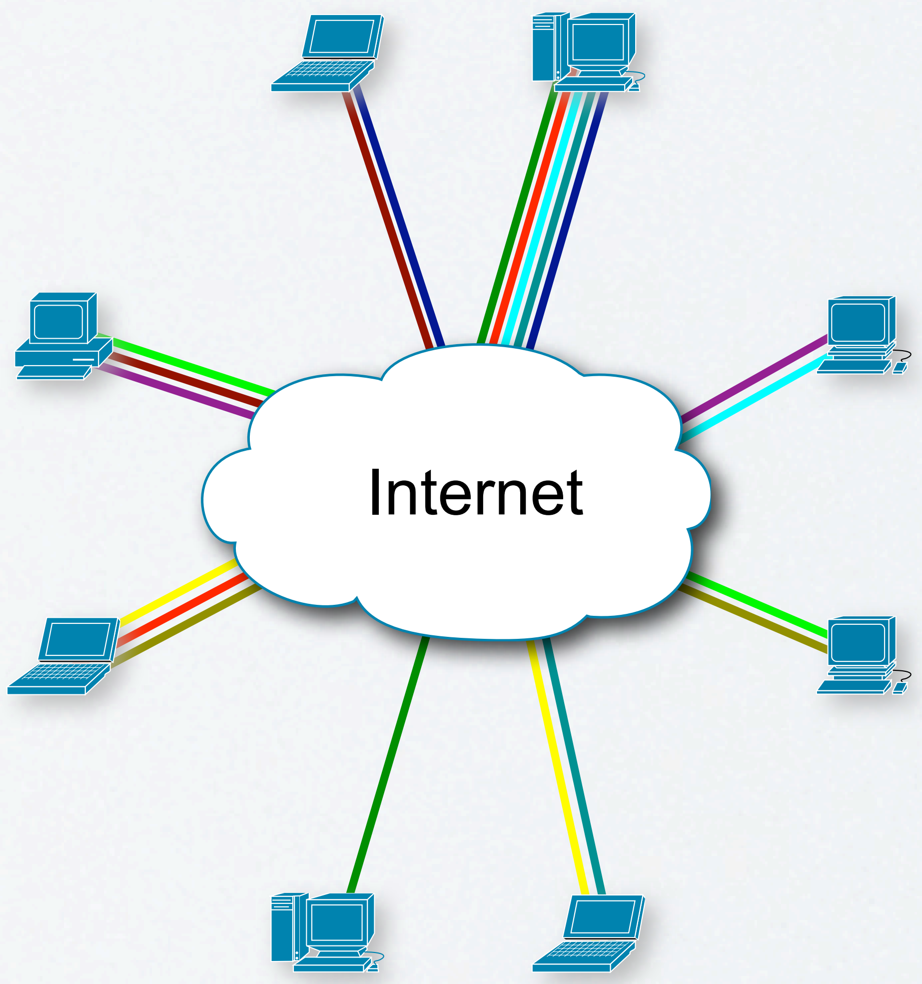

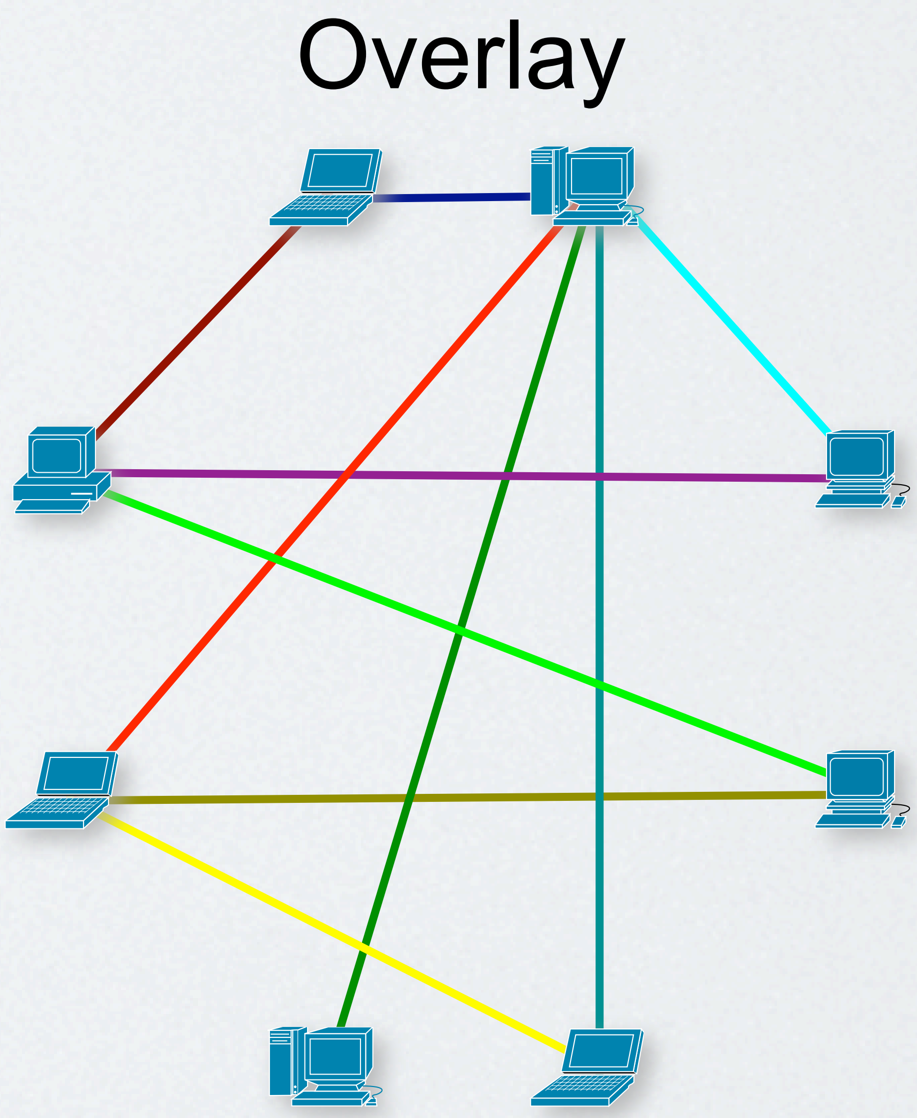

The Tragedy of the Commons

- Every participant is called a *peer*, and it has both client and server roles.
- Peers are assumed to be self interested
- If there is no incentive for contribution, there is a tendency to *freeload*
- A solution for this is *reciprocity*

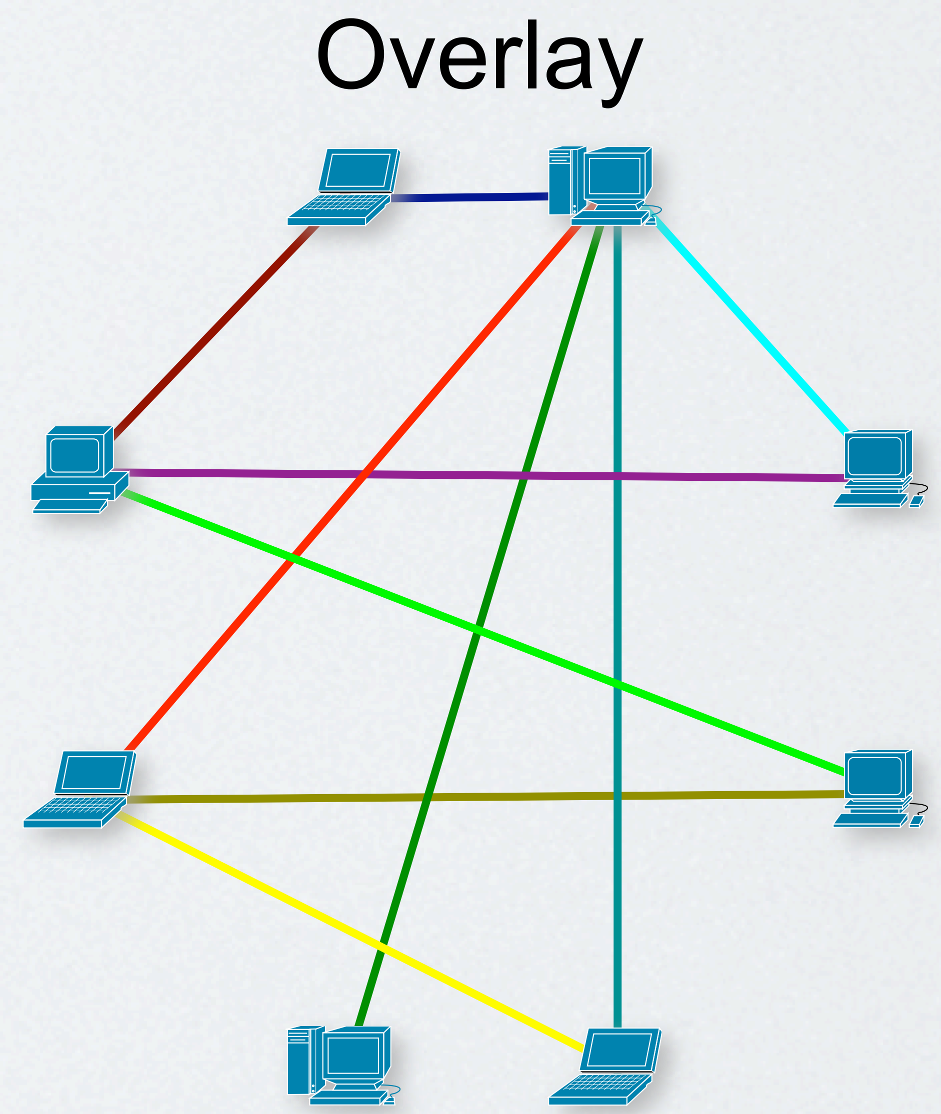

Direct Reciprocity

- Tit-for-Tat

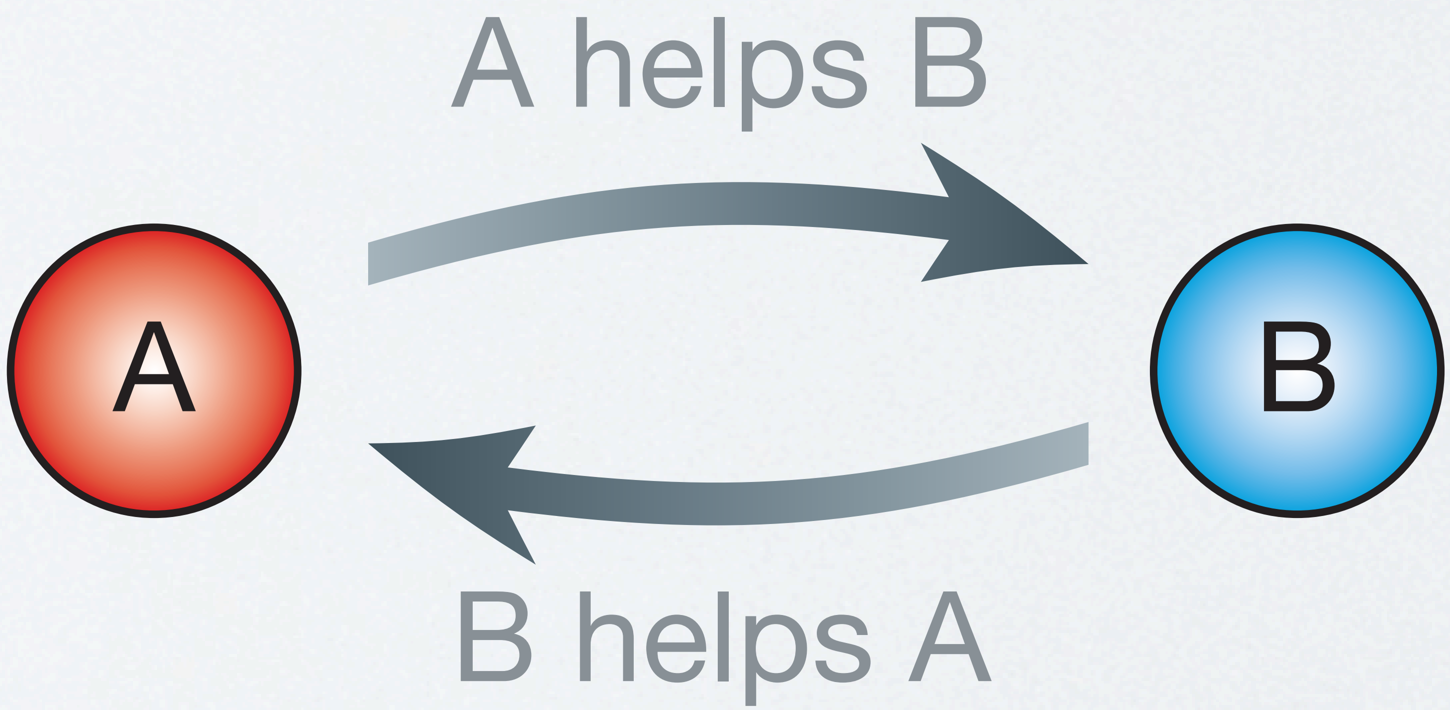

Nature 437, 1291-1298 (2005)

Indirect Reciprocity

- Tit-for-Tit-for-Tat

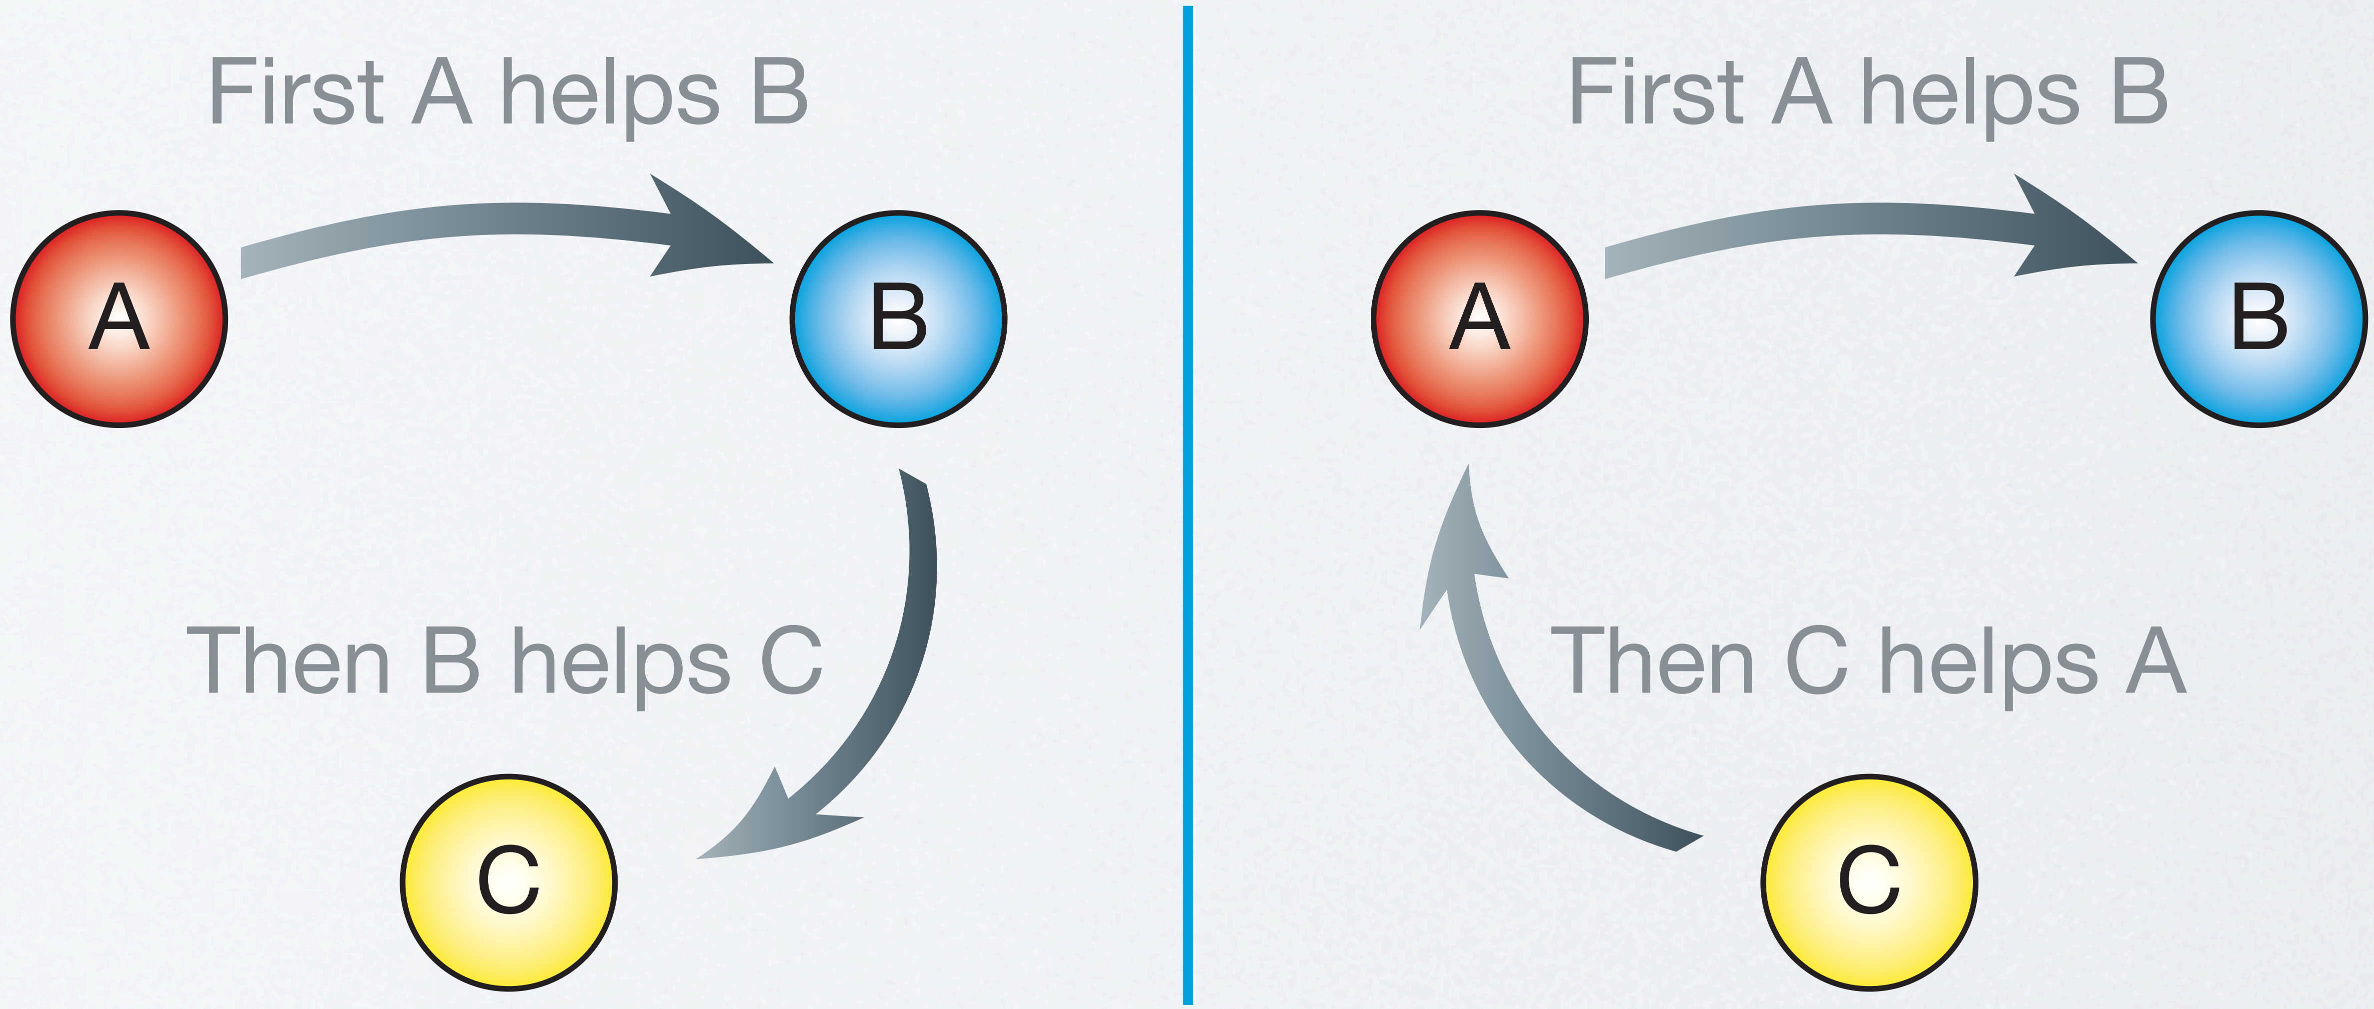

Nature 437, 1291-1298 (2005)

Our Contribution

- An analytic technique for the geometric analysis of contribution flows

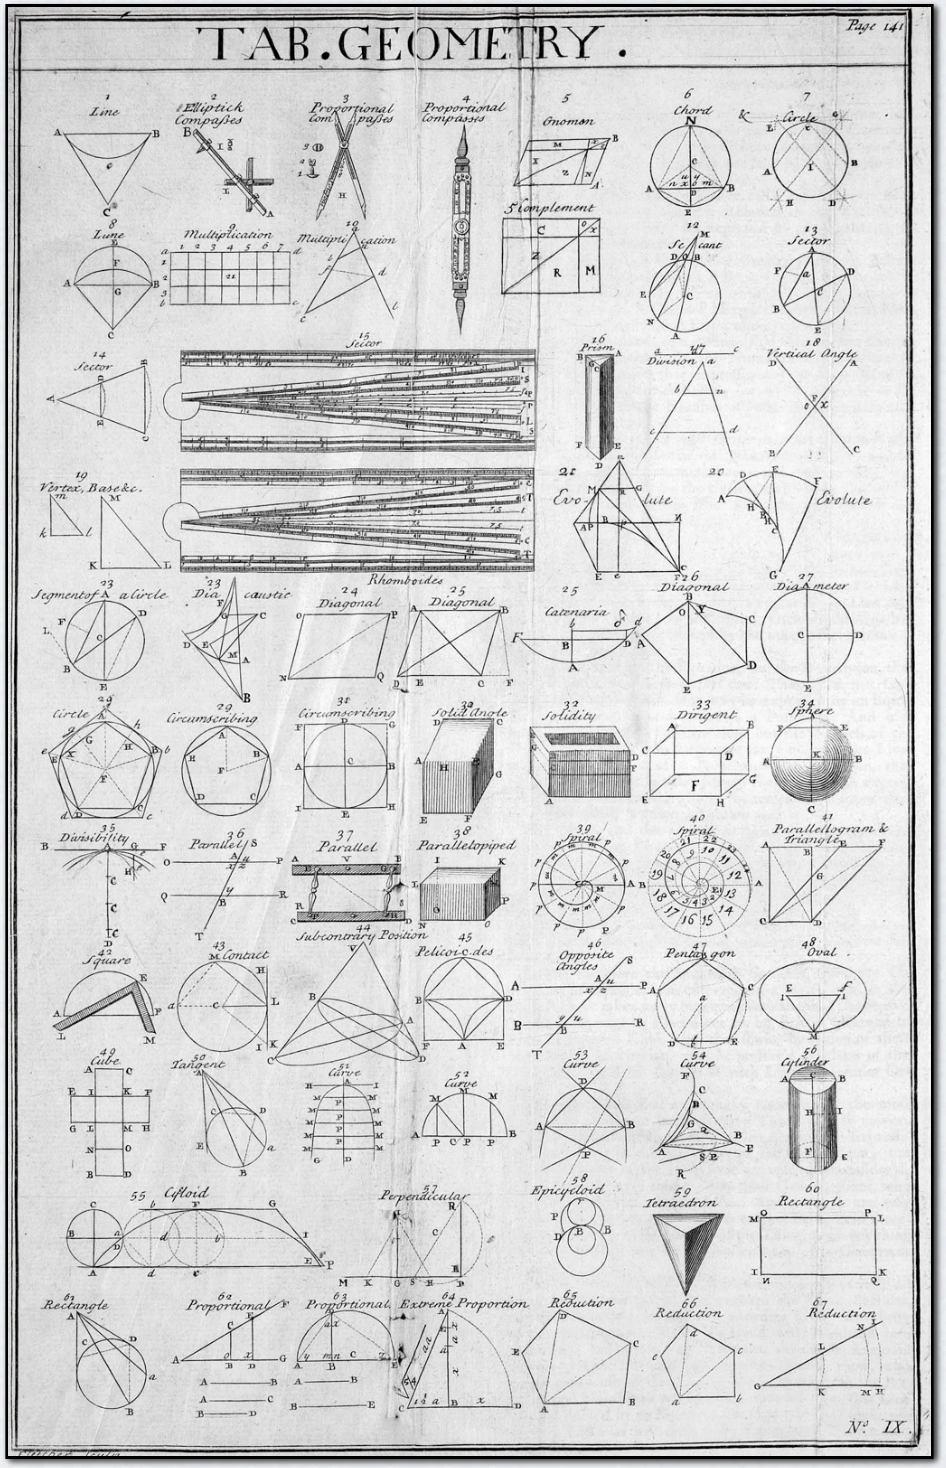

The Geometry of Indirect Reciprocity

- We consider the **contribution topology**, where a link is created between two nodes if one gives a contribution to the other.

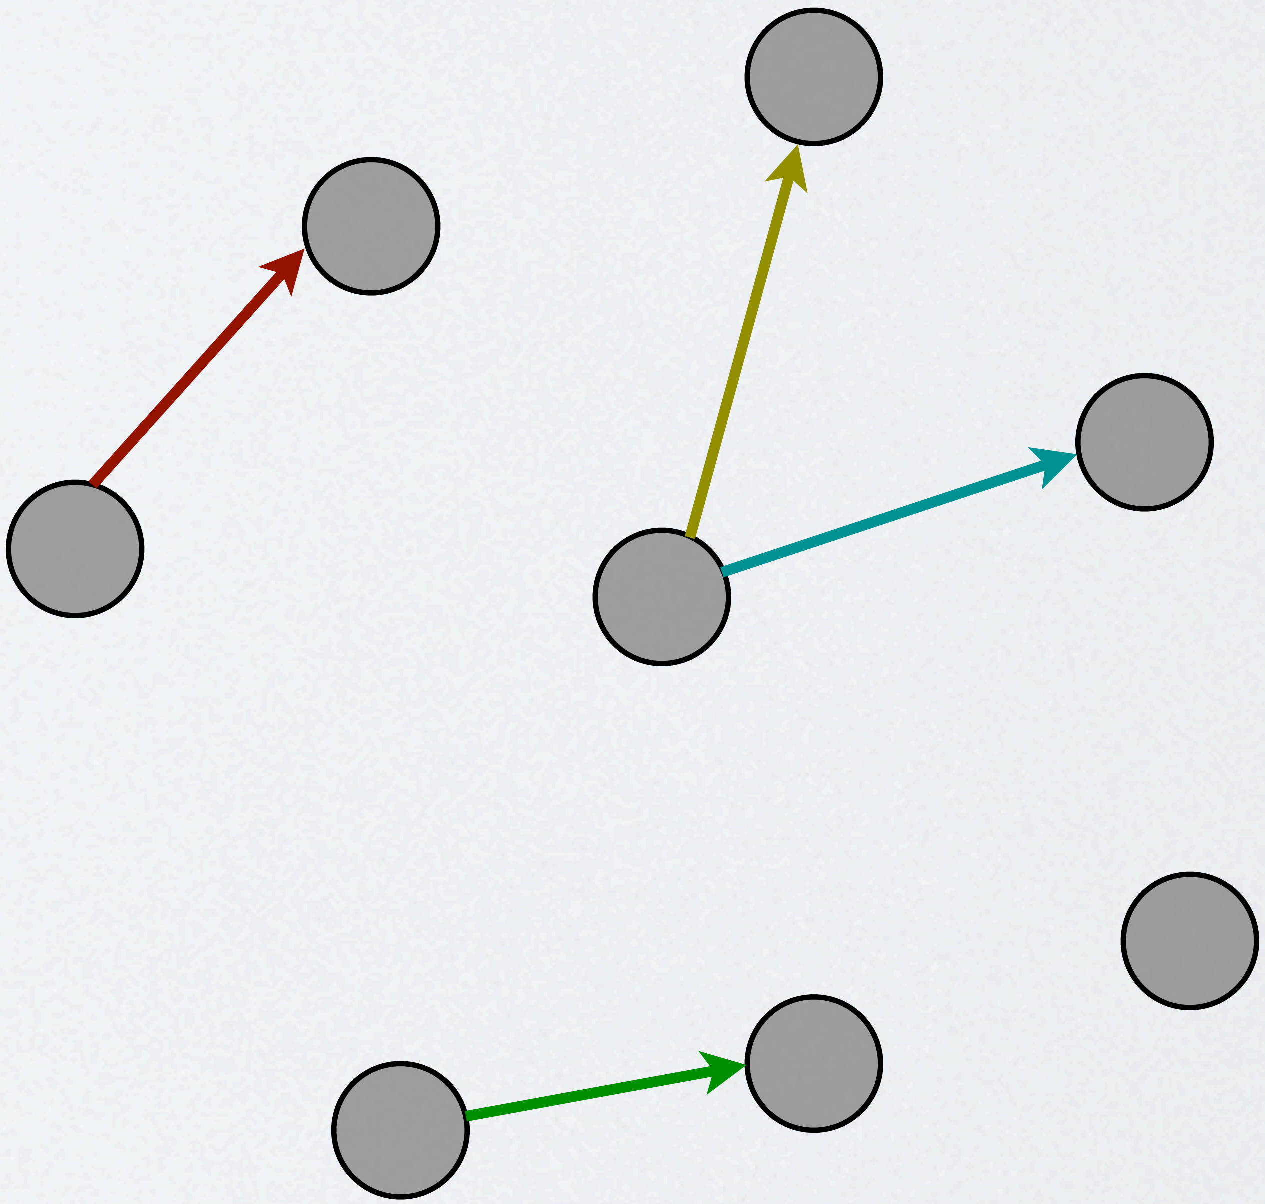

The Geometry of Indirect Reciprocity

- Reciprocity creates loops in the contribution topology.

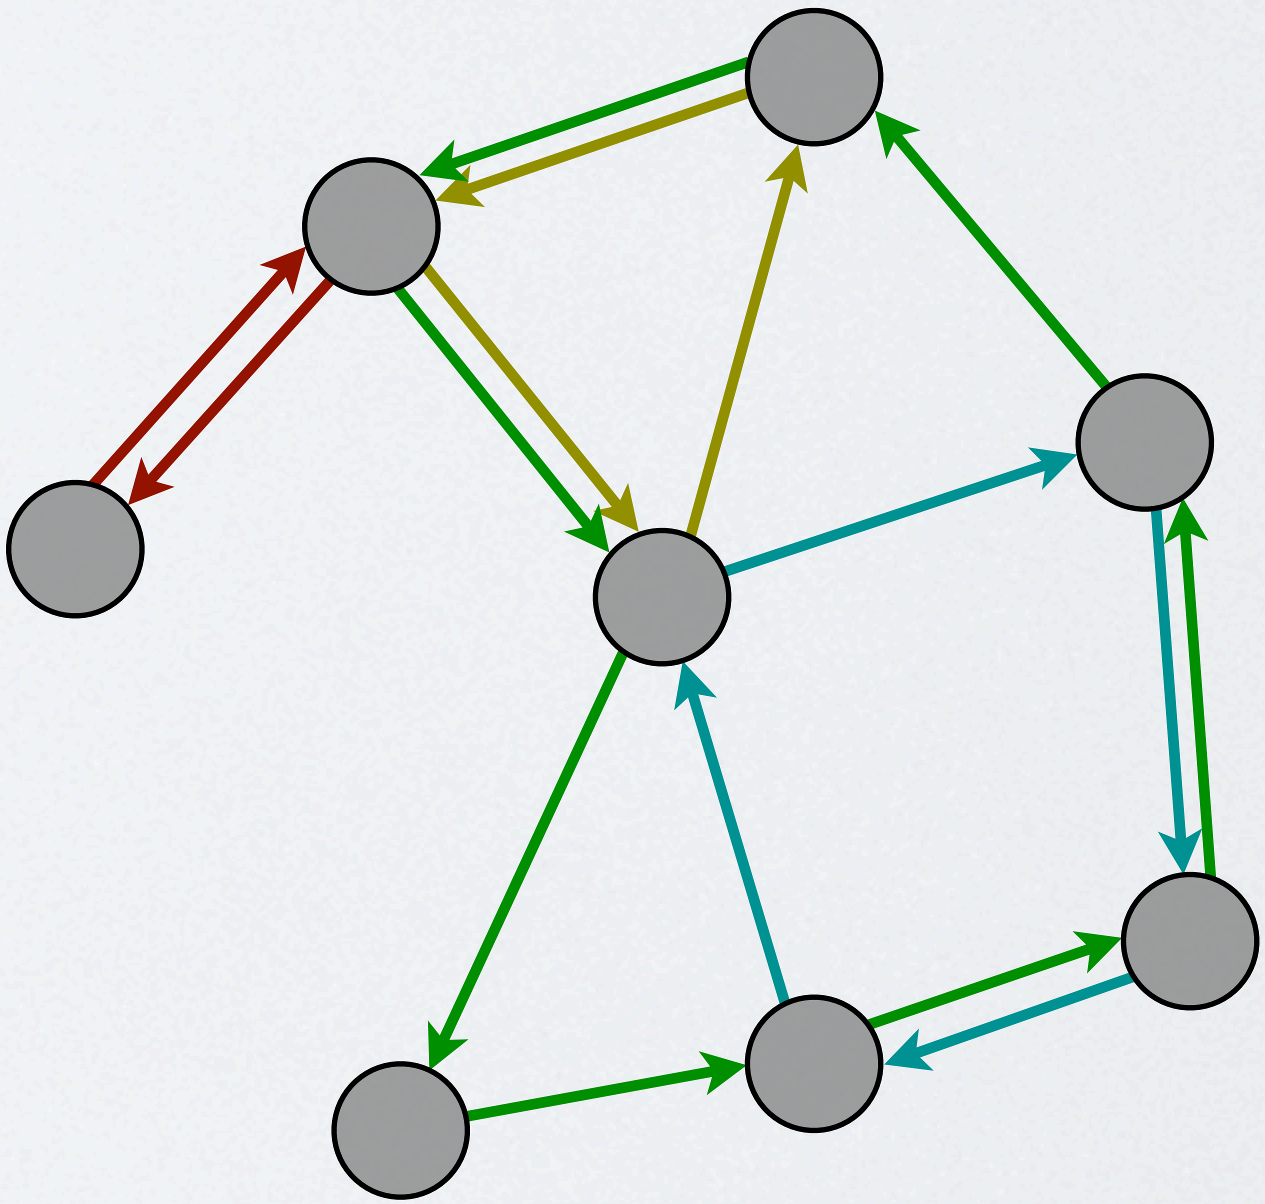

The Geometry of Indirect Reciprocity

- Reciprocity creates loops in the contribution topology.
- However, altruism requires non-cyclic contribution flows

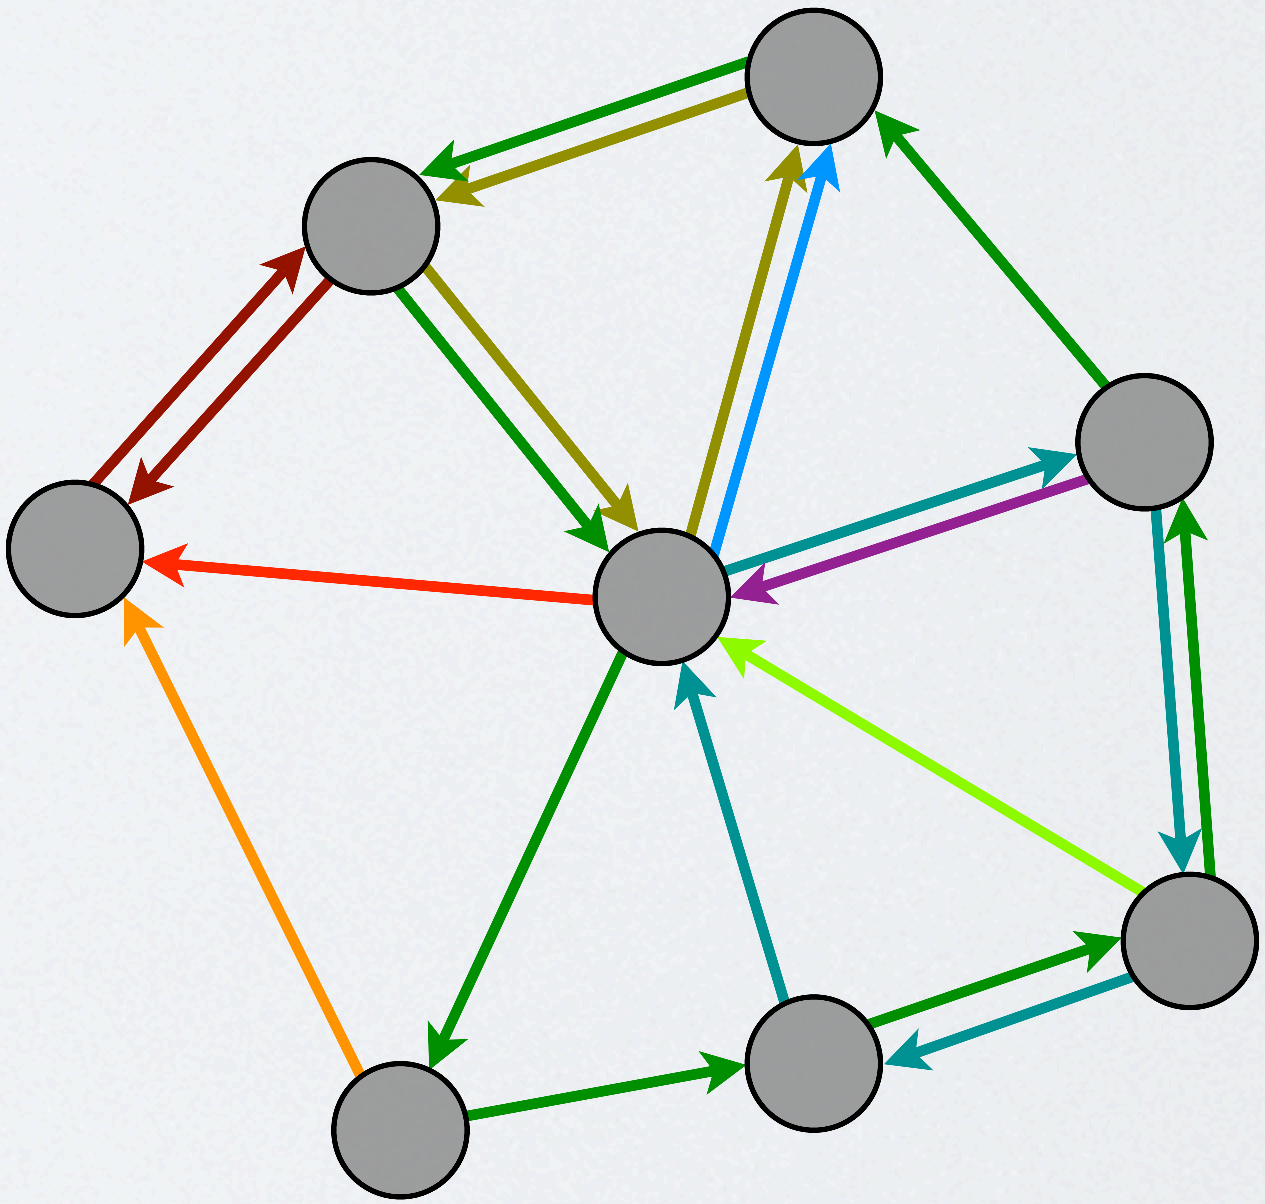

The Geometry of Indirect Reciprocity

- Reciprocity creates loops in the contribution topology.
- However, altruism requires non-cyclic contribution flows

- How can we model these contribution flows?

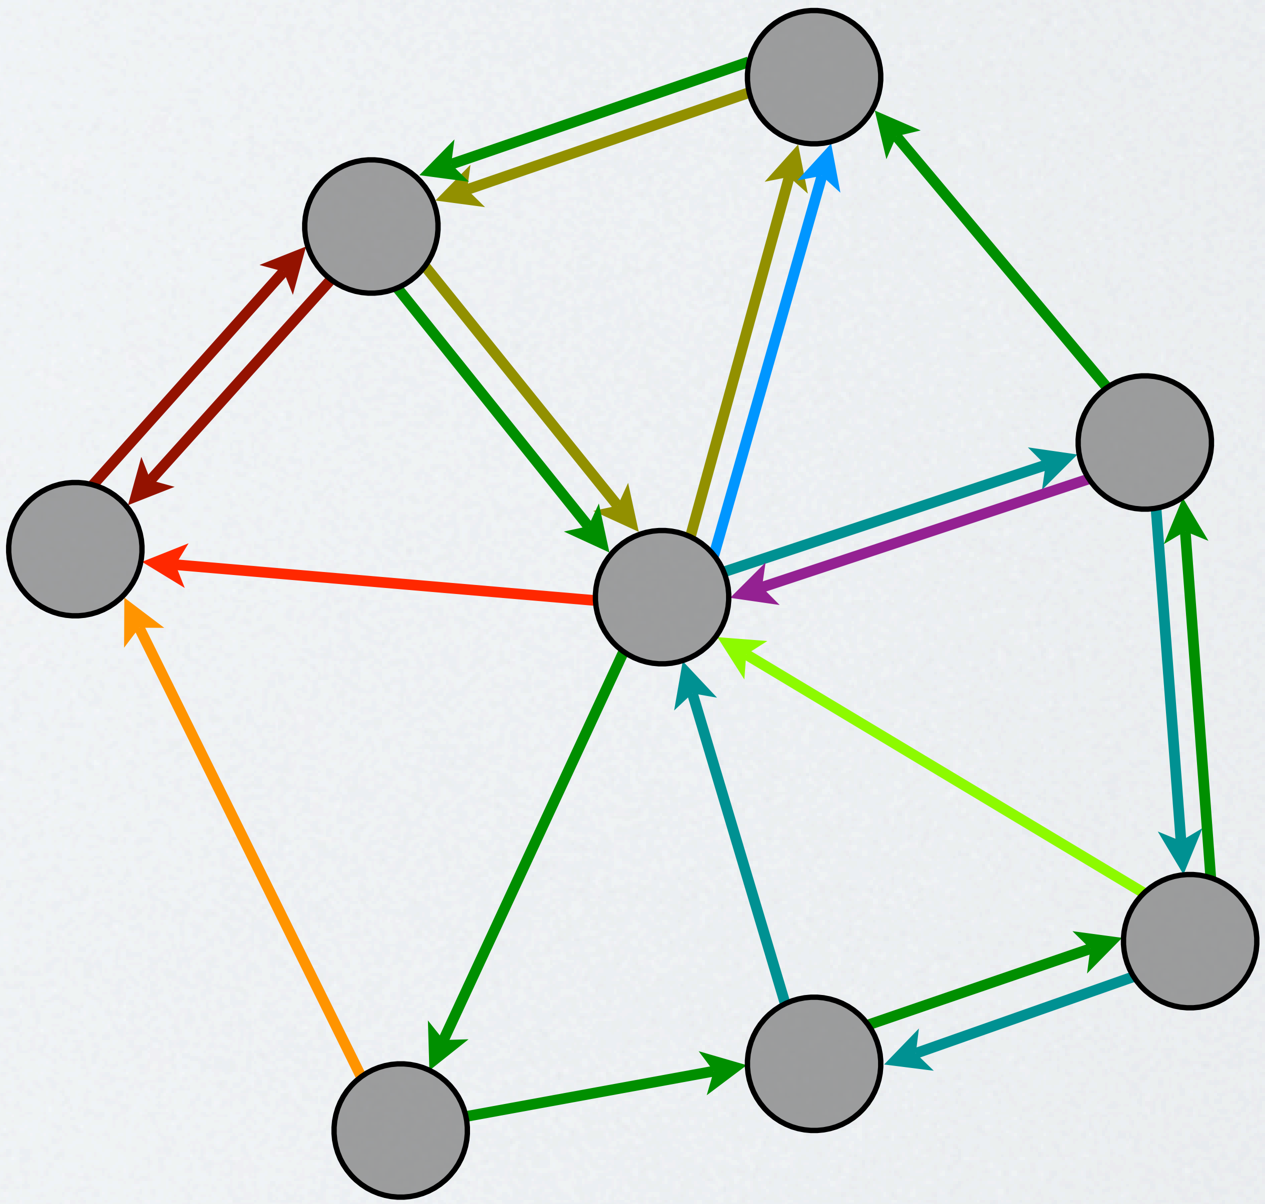

Functions in Graphs

- Domain:
  - Nodes (N)
  - Links (L)
  - Cycles (C)
- Range:
  - Reals ($\mathbb{R}$)

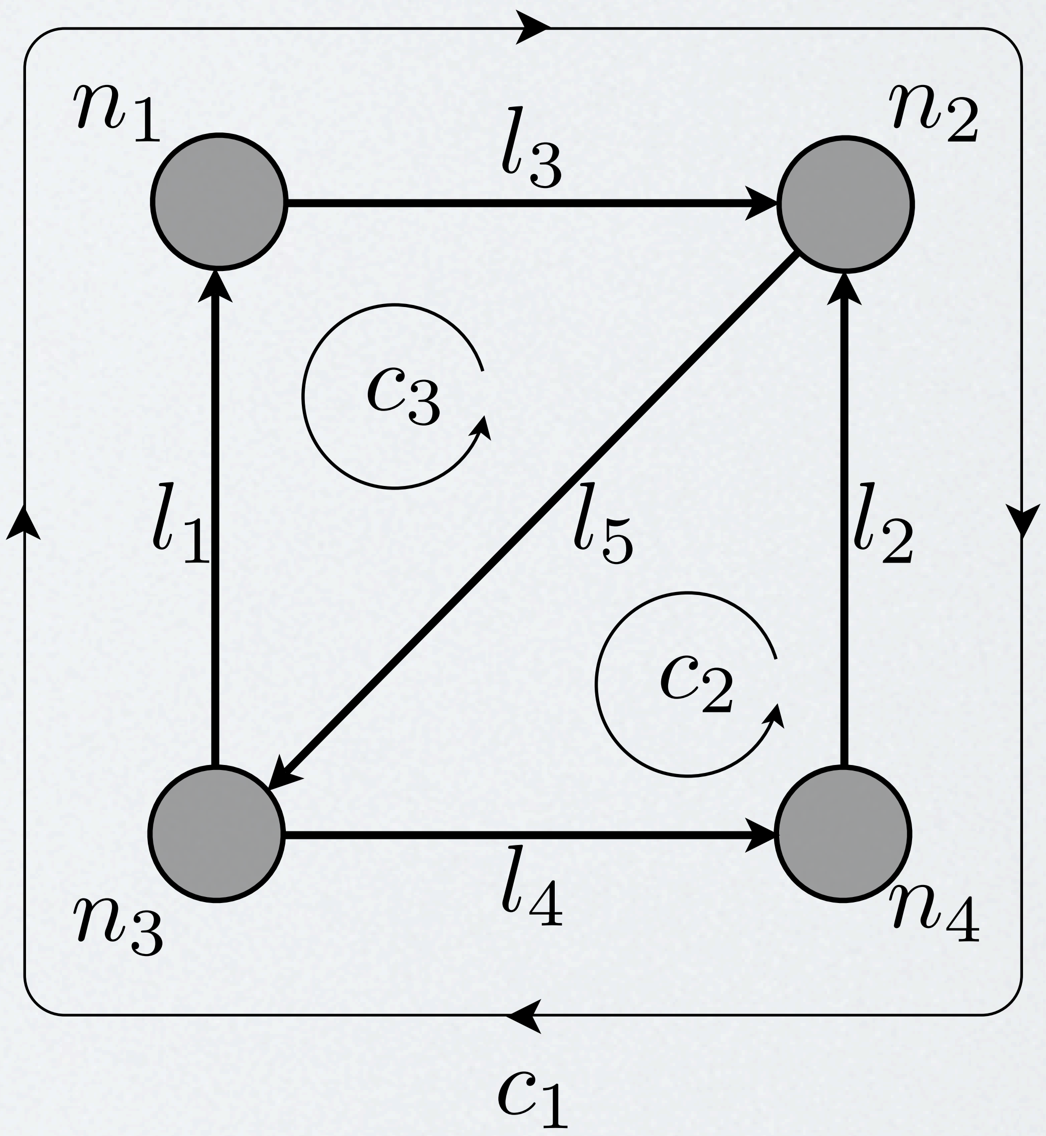

Differential Operators in Graphs

- They operate over node, link and cycle functions

- Equivalent to the well known vector operators:
  - Divergence
  - Gradient
  - Curl
  - Laplacian

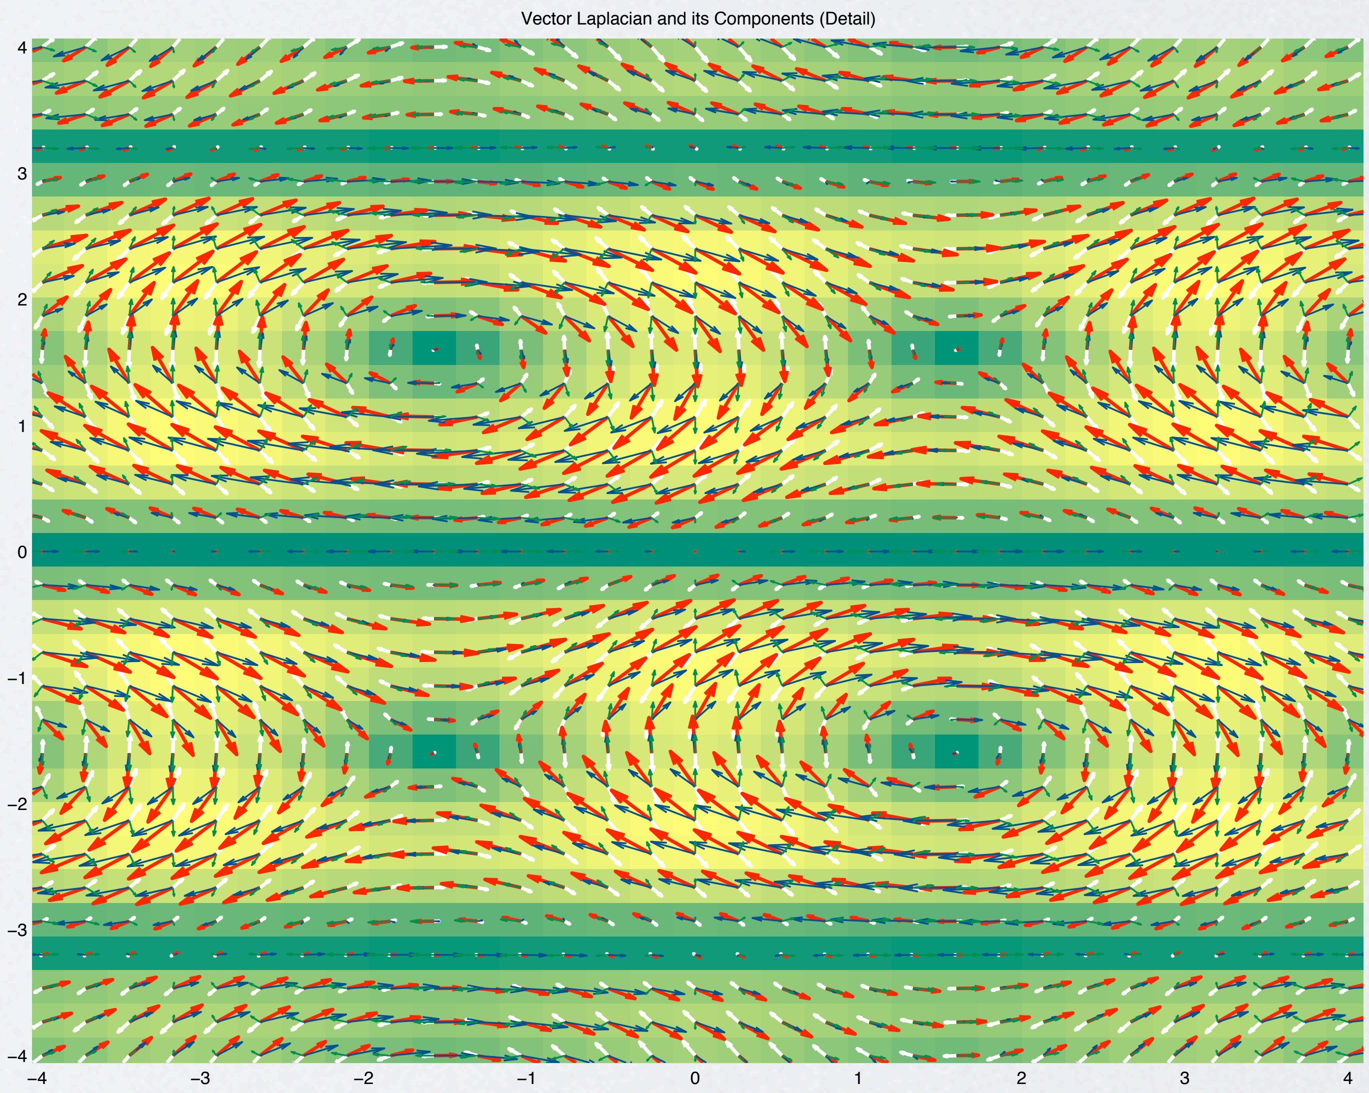

The Divergence

$$D(n_i, l_j) = \begin{cases} 1 & \text{if link } l_j \text{ is outgoing from node } n_i \\ -1 & \text{if link } l_j \text{ is incoming to node } n_i \end{cases}$$

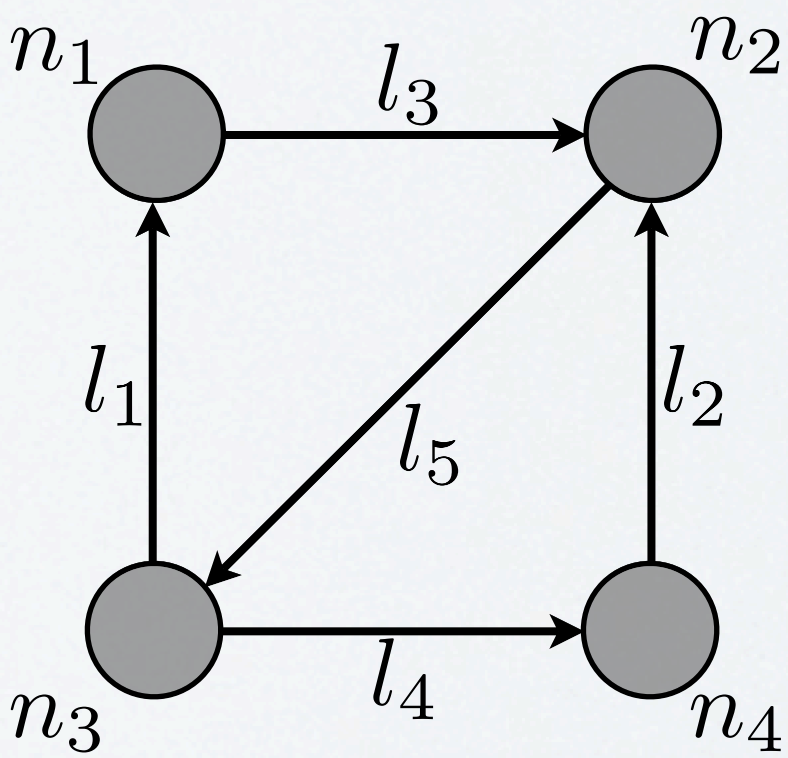

	l_1	l_2	l_3	l_4	l_5
n_1	1	0	-1	0	0
n_2	0	1	1	0	-1
n_3	-1	0	0	-1	1
n_4	0	-1	0	1	0

Calculating the Divergence

- If we have a link function f, we calculate its divergence simply by:

$$Df = \begin{pmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \end{pmatrix} = \begin{pmatrix} 1 & 0 & -1 & 0 & 0 \\ 0 & 1 & 1 & 0 & -1 \\ -1 & 0 & 0 & -1 & 1 \\ 0 & -1 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \\ f_5 \\ f_6 \end{pmatrix}$$

Calculating the Divergence

- If we have a link function f, we calculate its divergence simply by:

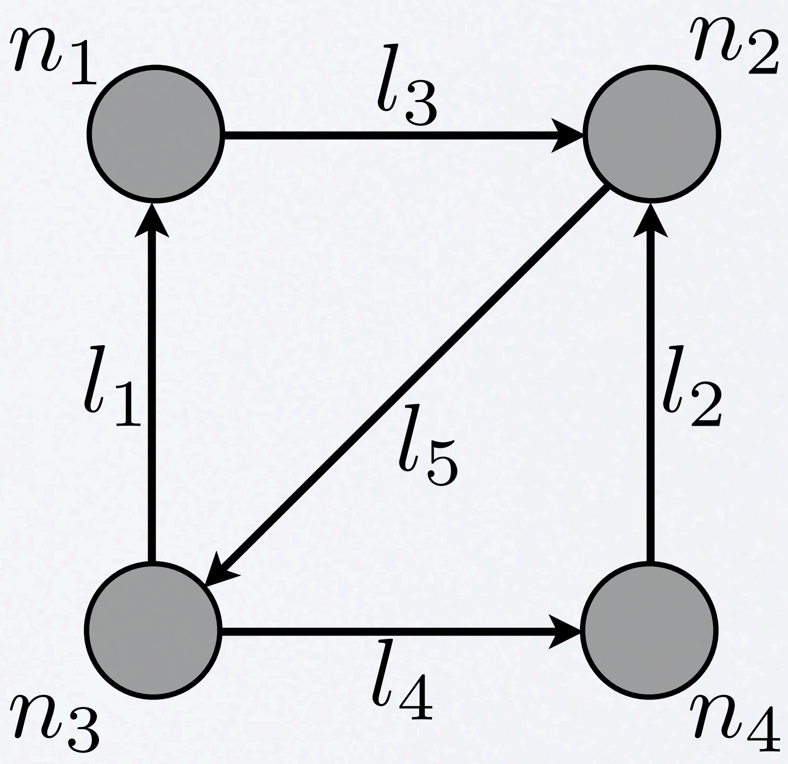

$$\begin{pmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \end{pmatrix} = \begin{pmatrix} f_1 - f_3 \\ f_2 + f_3 - f_6 \\ f_6 - f_1 - f_5 \\ f_5 - f_2 \end{pmatrix}$$

The Gradient

- It is just the transpose of the divergence

$$G = D^T$$

- If we have a node function F, we calculate its gradient simply by:

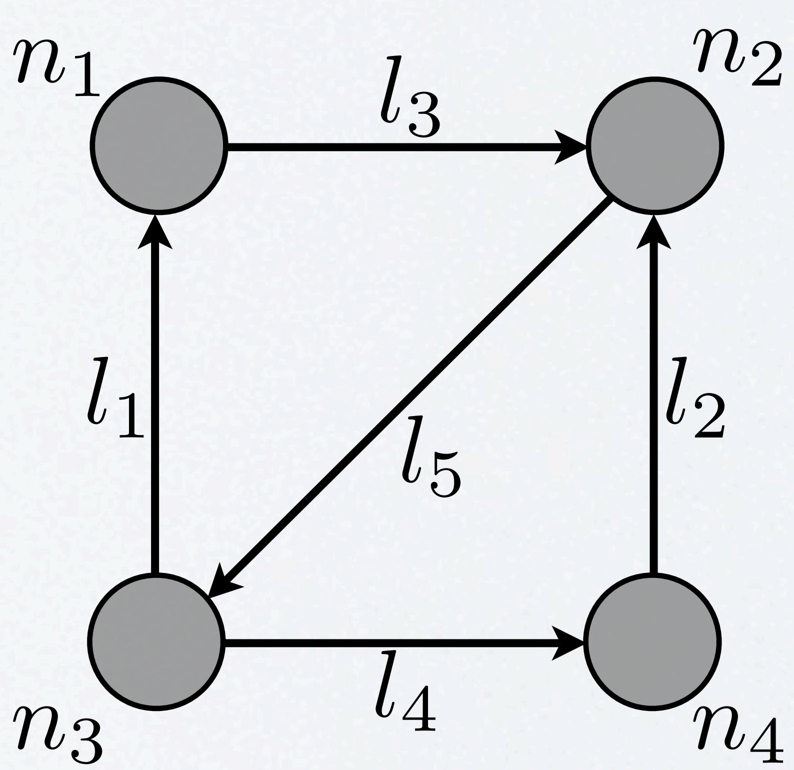

$$\begin{pmatrix} g_1 \\ g_2 \\ g_3 \\ g_4 \\ g_5 \end{pmatrix} = \begin{pmatrix} F_1 - F_3 \\ F_2 - F_4 \\ F_2 - F_1 \\ F_4 - F_3 \\ F_3 - F_2 \end{pmatrix}$$

The Rotational Operators

- They require knowledge of the *cycle structure* of the graph $\mathcal{G}$:
  - Generate $\mathcal{G}'$, an undirected version of $\mathcal{G}$
  - Embed $\mathcal{G}'$ in a surface with minimum *genus*
  - Recover a *cellular cycle basis* from the embedding
  - Define an *orientation* for the cycle basis
  - Use this oriented cycle basis to define the curl

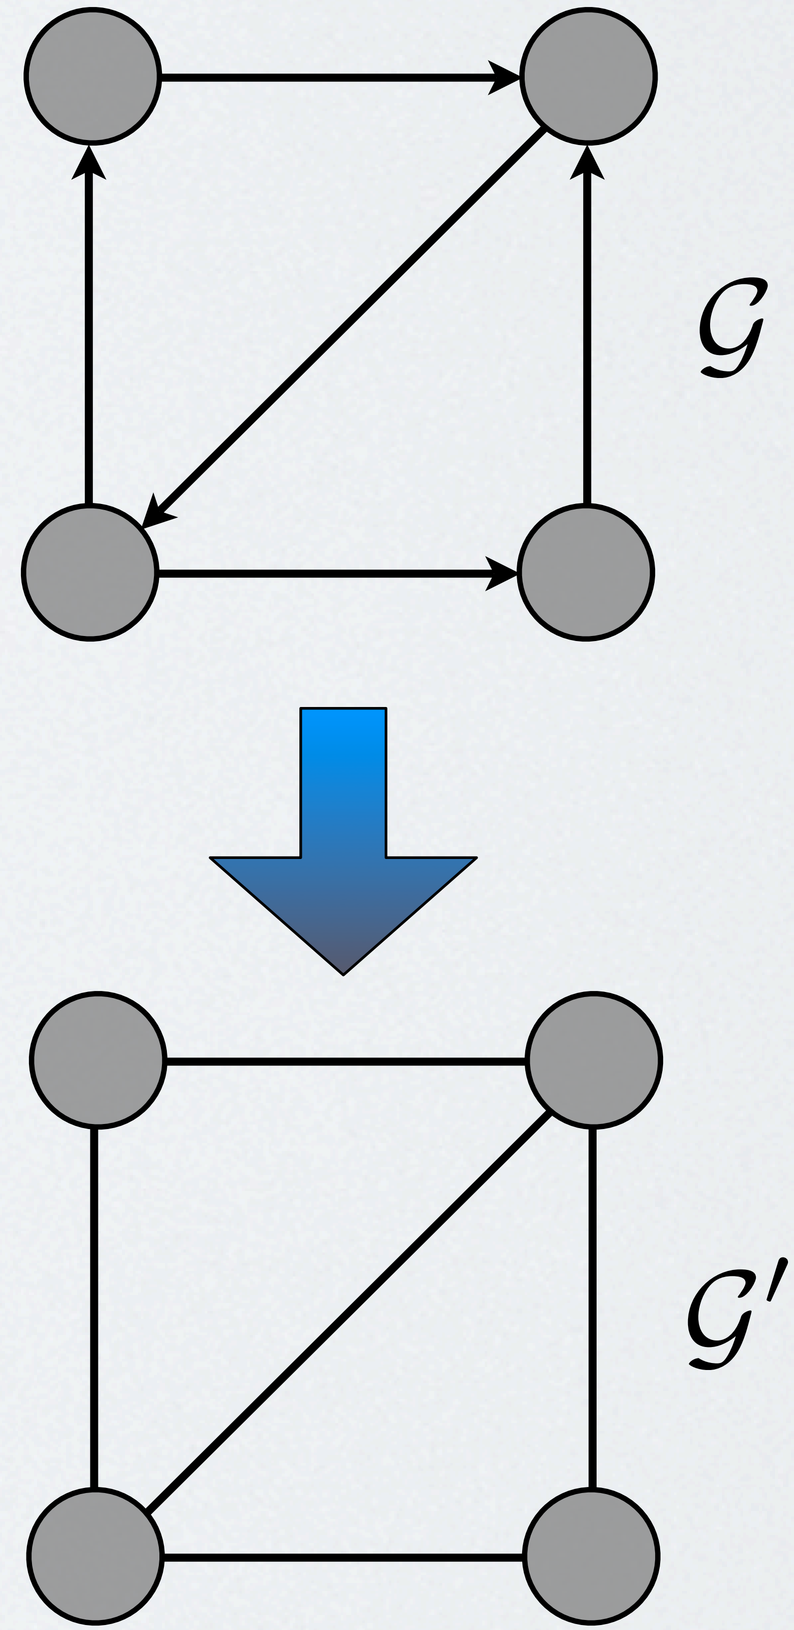

Graph Surface Embedding

- An embedding of $\mathcal{G}'$ on a surface $\mathcal{S}$ is a way of drawing $\mathcal{G}'$ on $\mathcal{S}$ so that there are no edge crossings.

- Links become ***lines*** in $\mathcal{S}$
- Nodes become ***points*** in $\mathcal{S}$

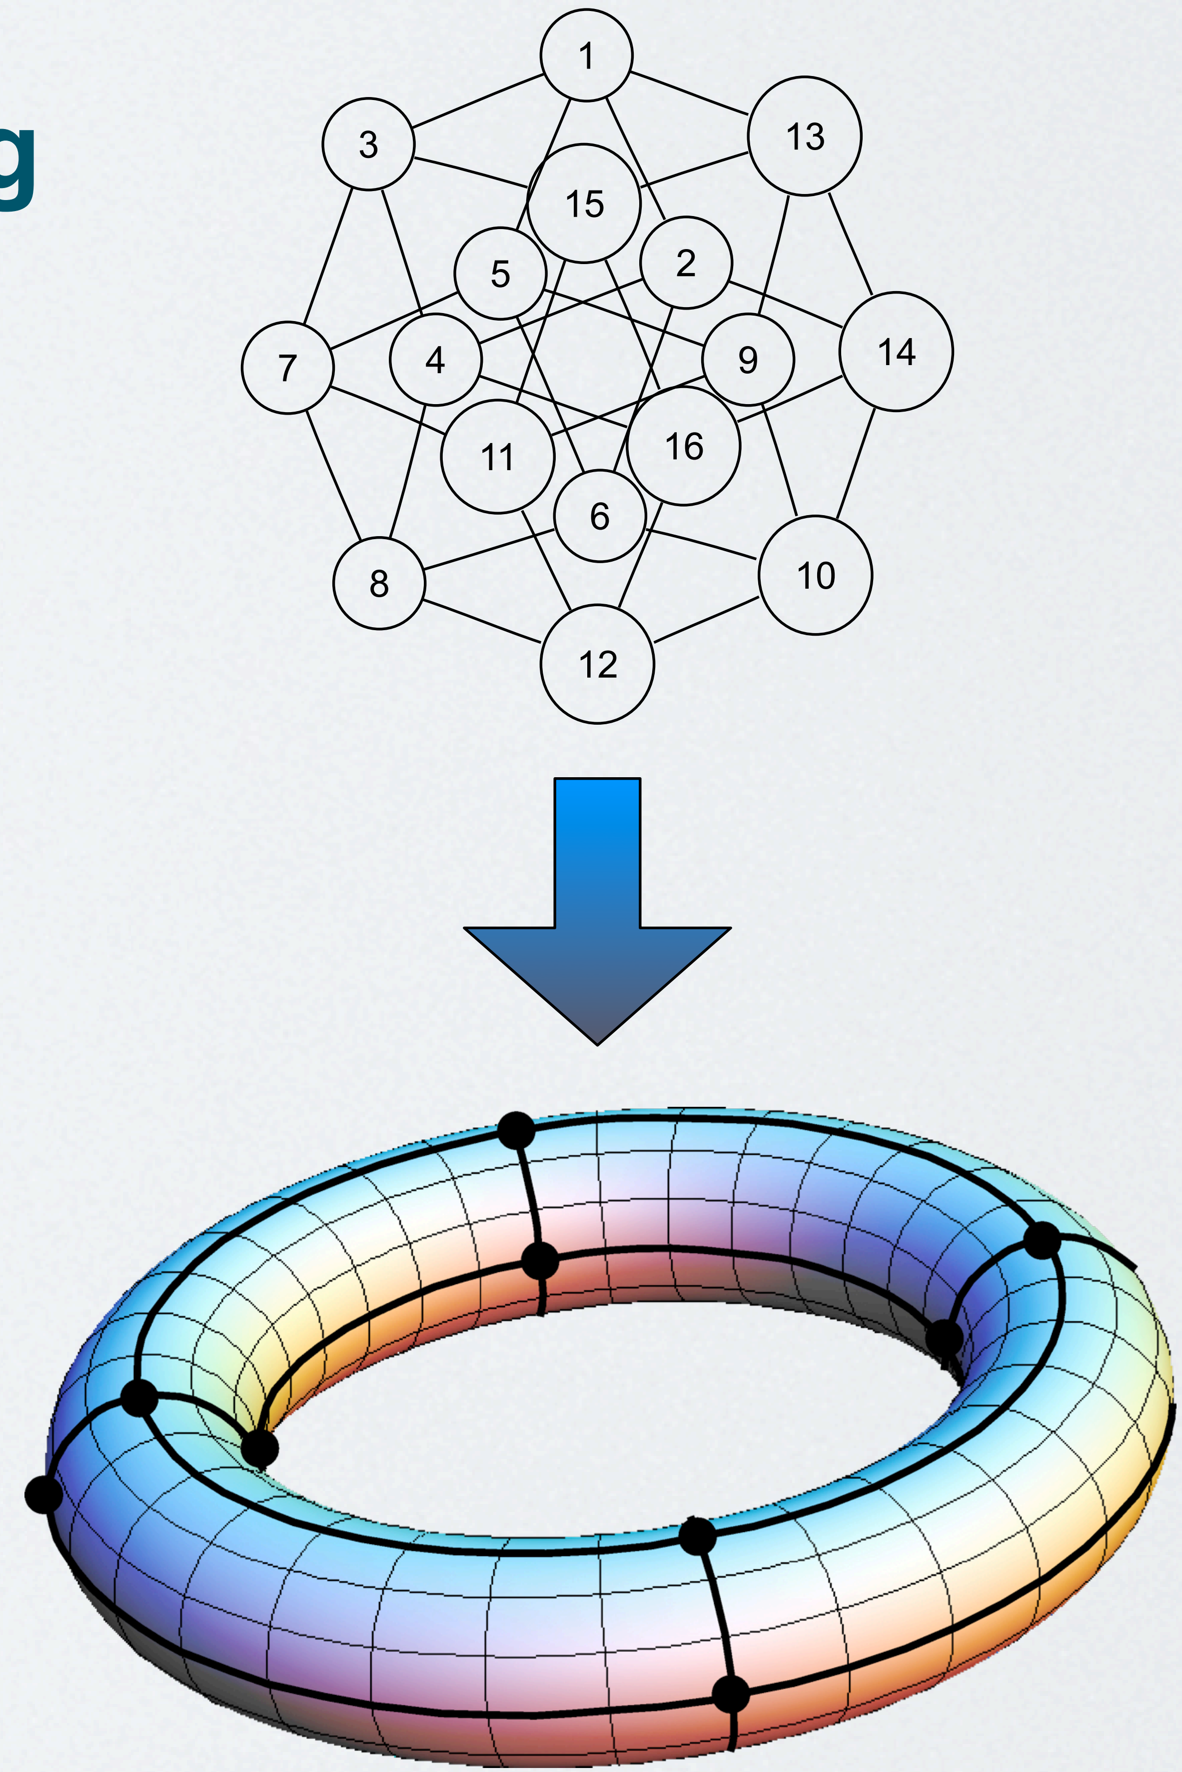

Minimum Genus Embedding

- A surface embedding on which $\mathcal{S}$ has the minimum number of holes possible
- We focus on ***orientable, closed*** surfaces, although the embedding can be done on non-orientable surfaces as well

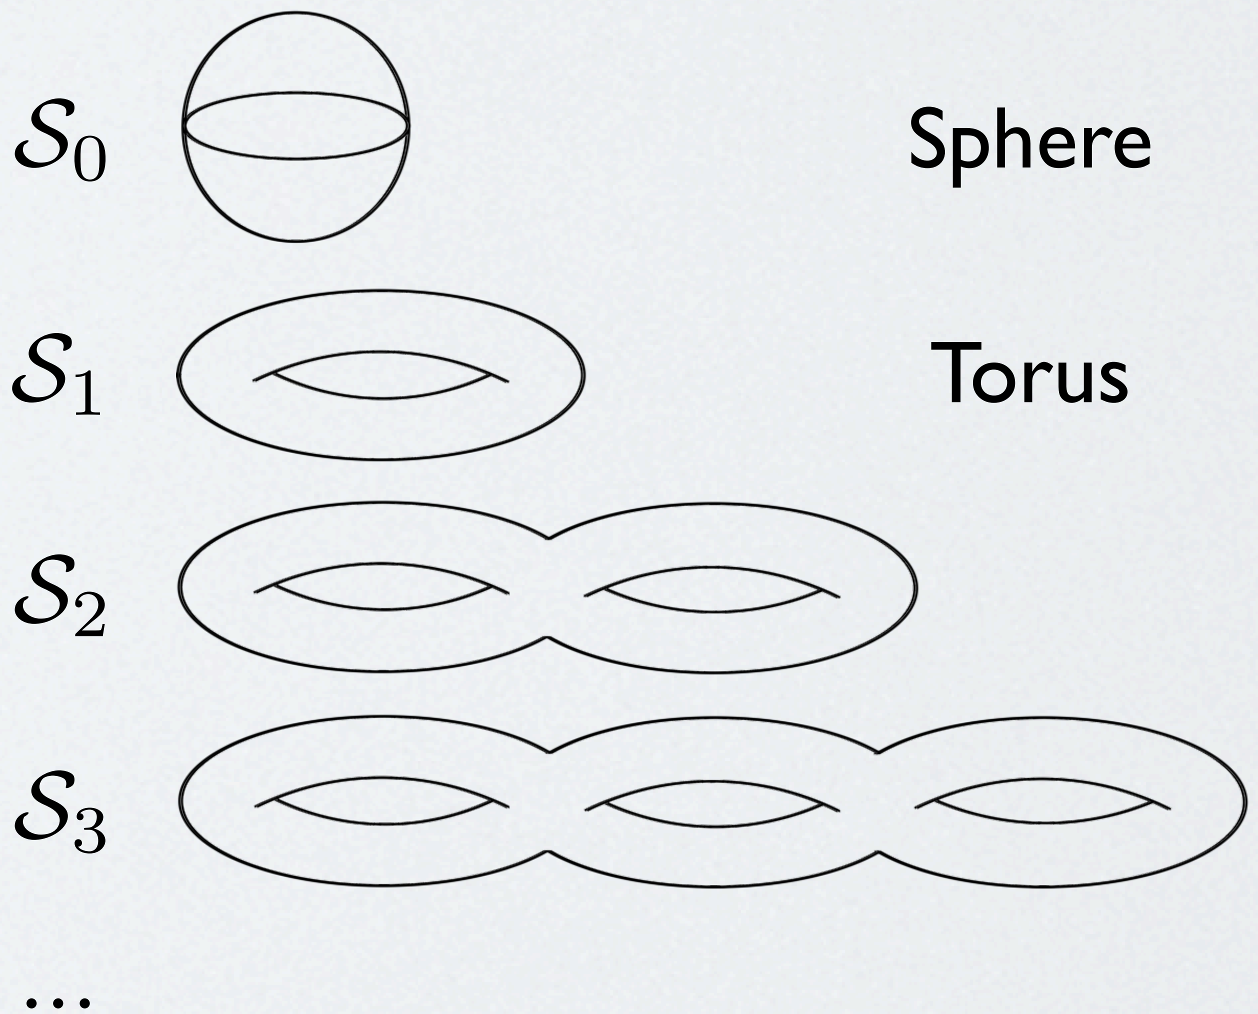

Cellular Cycle Basis

- A minimum genus embedding provides a ***cellular cycle system***, where:
  - Every link belongs to exactly two cycles, a *left* cycle and a *right* cycle
  - Areas bordered by links become polygonal ***faces***
    - In a planar graph, each face defines a cellular cycle
  - The network becomes a ***polyhedron***

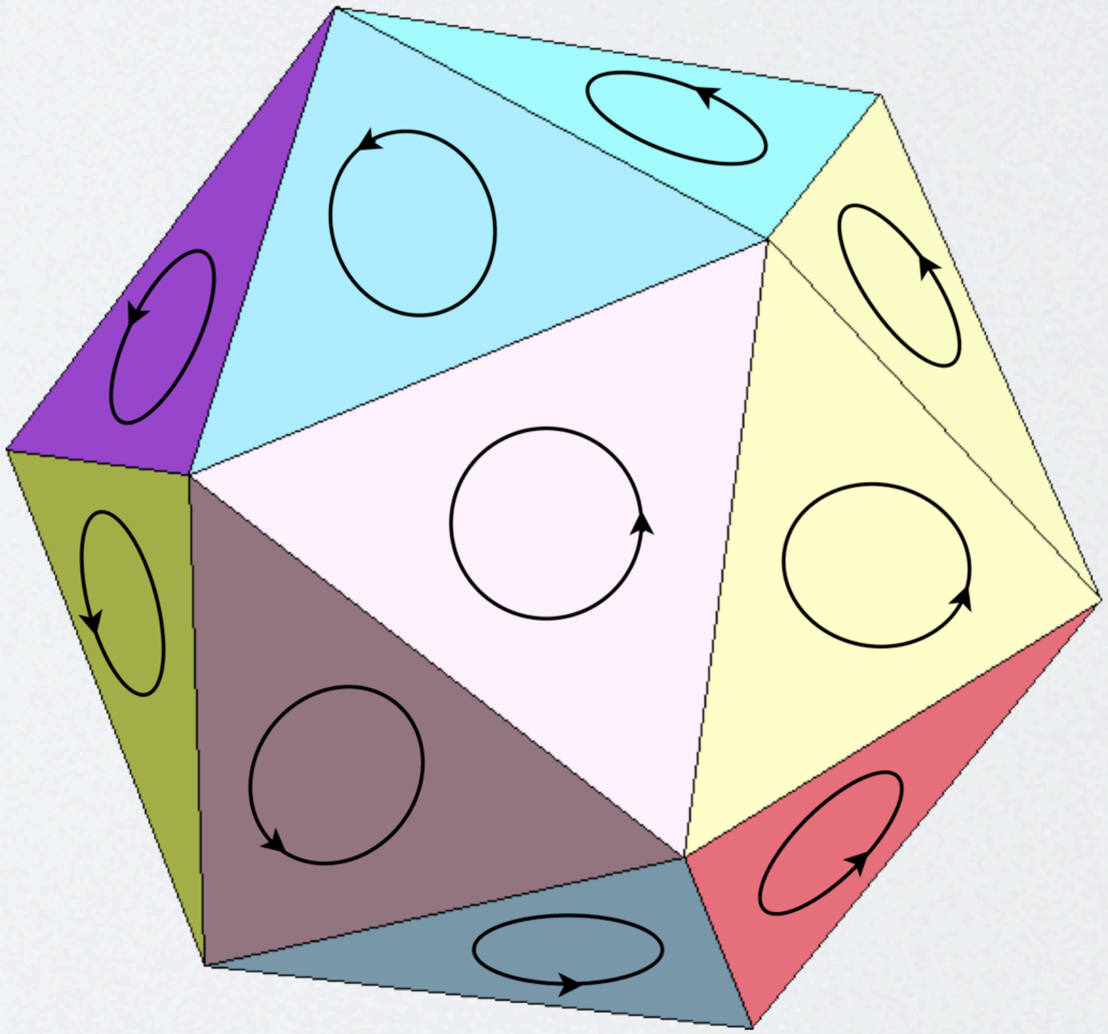

The Curl

$$C(c_i, l_j) = \begin{cases} 1 & \text{if link } l_j \text{ is positively oriented along cycle } c_i \\ -1 & \text{if link } l_j \text{ is negatively oriented along cycle } c_i \end{cases}$$

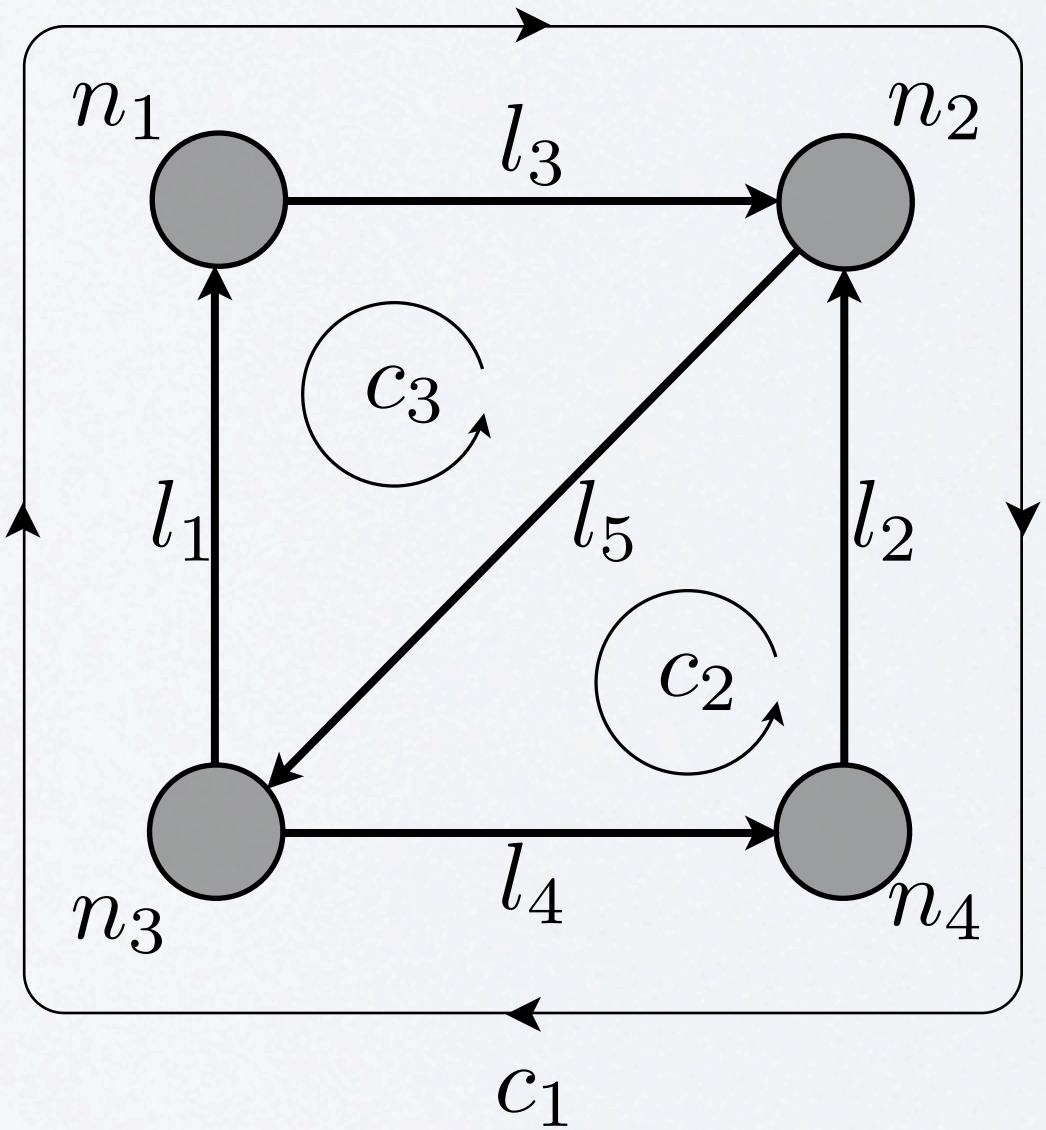

	l_1	l_2	l_3	l_4	l_5
c_1	1	-1	1	-1	0
c_2	0	1	0	1	1
c_3	-1	0	-1	0	-1

Calculating the Curl

- For a given link function f, we have that Cf can be calculated as:

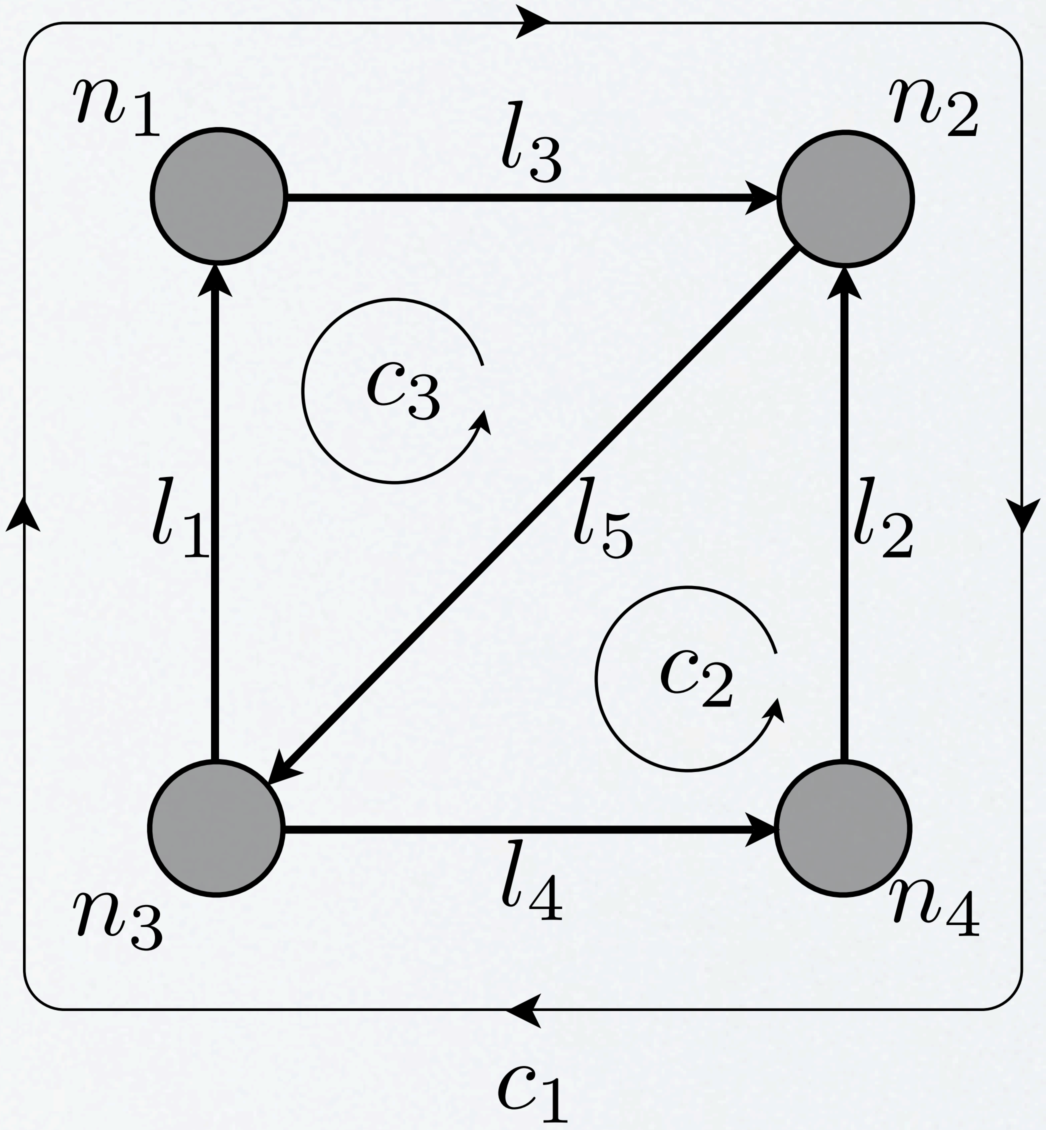

$$\begin{pmatrix} c_1 \\ c_2 \\ c_3 \end{pmatrix} = \begin{pmatrix} f_1 + f_3 - f_2 - f_4 \\ f_2 + f_5 + f_4 \\ -f_1 - f_5 - f_3 \end{pmatrix}$$

The Adjoint Curl

- It is just the transpose of the curl

$$S = C^T$$

- If we have a cycle function F, we have for SF:

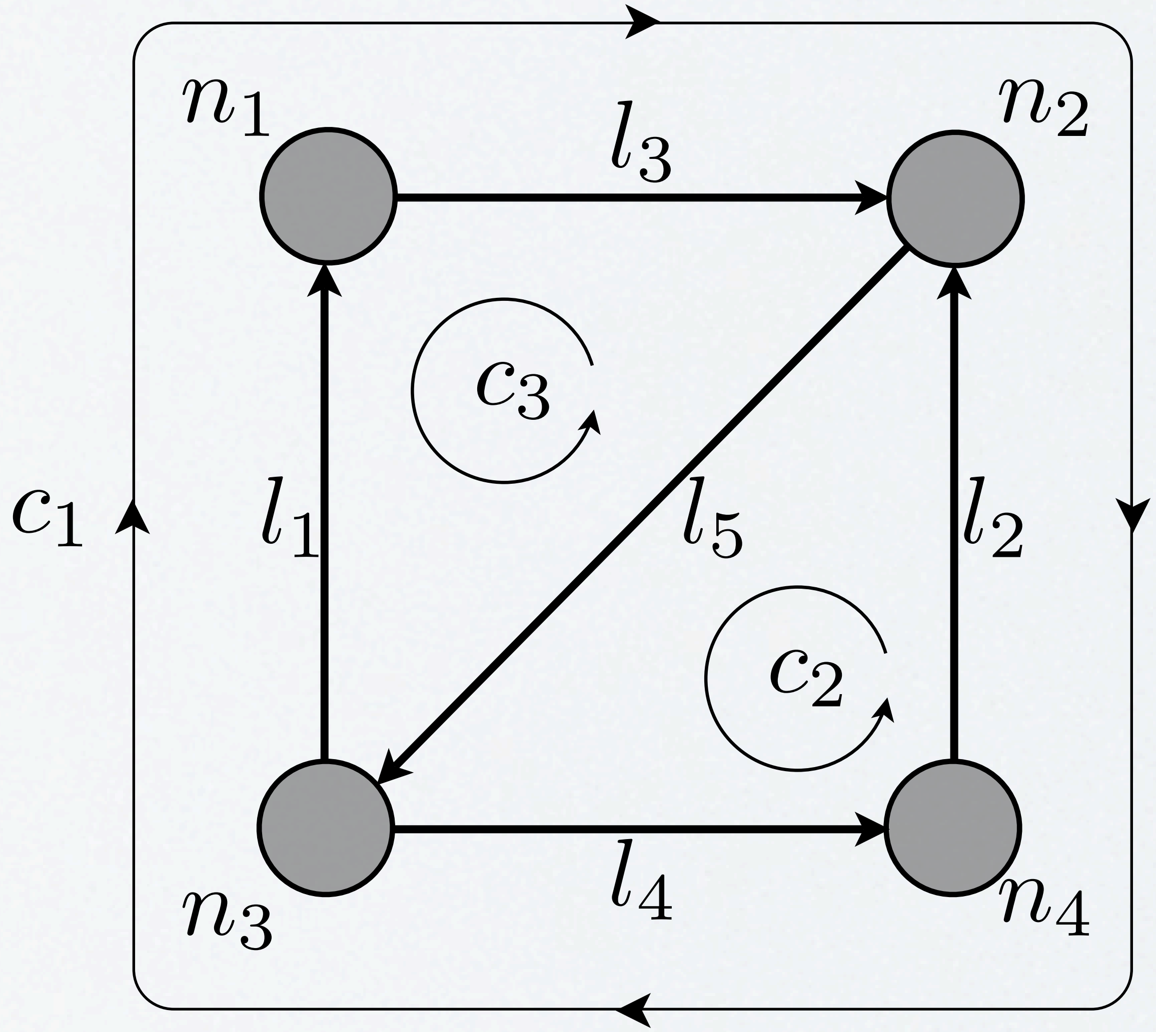

$$\begin{pmatrix} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \end{pmatrix} = \begin{pmatrix} F_1 - F_3 \\ F_2 - F_1 \\ F_1 - F_3 \\ F_2 - F_1 \\ F_2 - F_3 \end{pmatrix}$$

Gradients are Irrotational

- For any node function F we have that:

$$CGF = 0$$

- This is because every row of C (a cycle c_i) is orthogonal to every column in G (a node n_j)
  - Two cases:
    - If n_j does not belong to c_i, they will have no common nonzero entries and $c_i \cdot n_j = 0$.
    - If n_j does belong to c_i, we know that they have exactly two common nonzero entries, corresponding to the two links in c_i, incident on n_j.

Gradients are Irrotational

- In this case, we have two sub-cases:
  - *Sub-case 1*: Both links incident to n_j have the same orientation with respect to c_i
    - Their entries in will c_i have the same sign, but their entries in n_j will have opposite signs (one outgoing, one incoming)
  - *Sub-case 2*: The 2 links incident to n_j have opposite orientations with respect to c_i
    - Their entries in n_j will have the same sign, but their entries in c_i will have opposite signs (one with c_i , one against it)

Sub-case ***1***

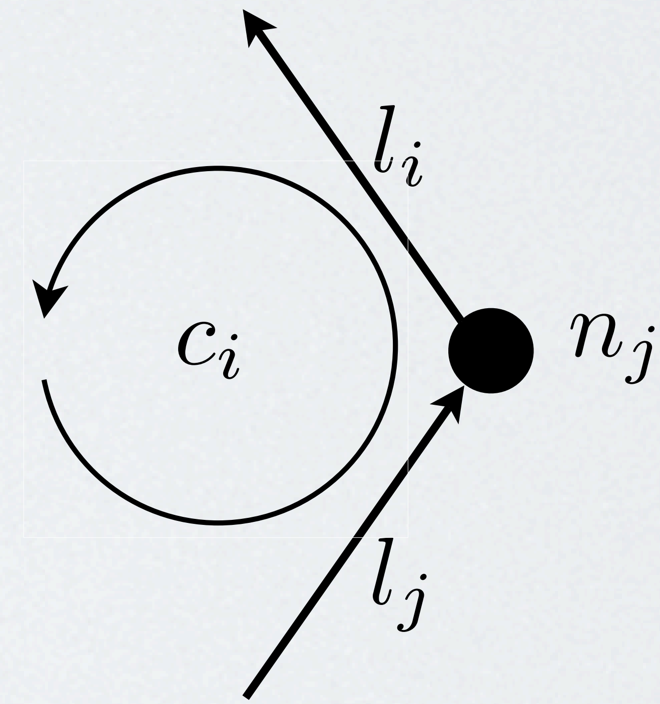

Sub-case **2**

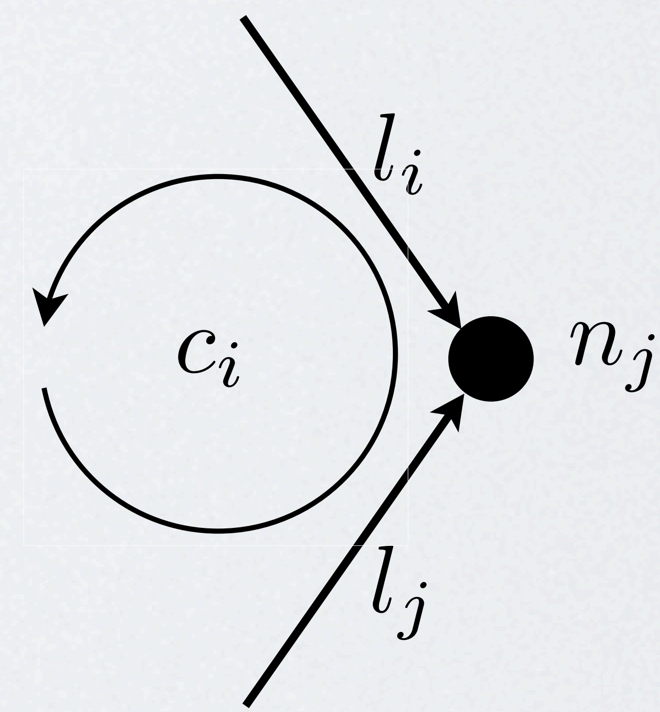

Gradients are Irrotational

- Thus, every link function f that has zero curl can be represented as the gradient of a node potential ϕ:

$$Cf = 0 \;\Rightarrow\; f = G\phi$$

Adjoint Curl functions are Incompressible

- For any cycle function F, we have that:

$$DSF = 0$$

- Proof:
  - We have that:

$$(CG)^T = G^T C^T = DS$$

  - And thus:

$$CG = 0 \iff DS = 0$$

Adjoint Curl functions are Incompressible

- Thus, every link function f that has zero divergence can be represented as the curl of a cycle potential ψ:

$$Df = 0 \;\Rightarrow\; f = C\psi$$

Second-Order Differential Operators

- By combining D, G, C and S we obtain second-order operators.

- The eigenvectors of these operators provide basis for node, cycle and link functions

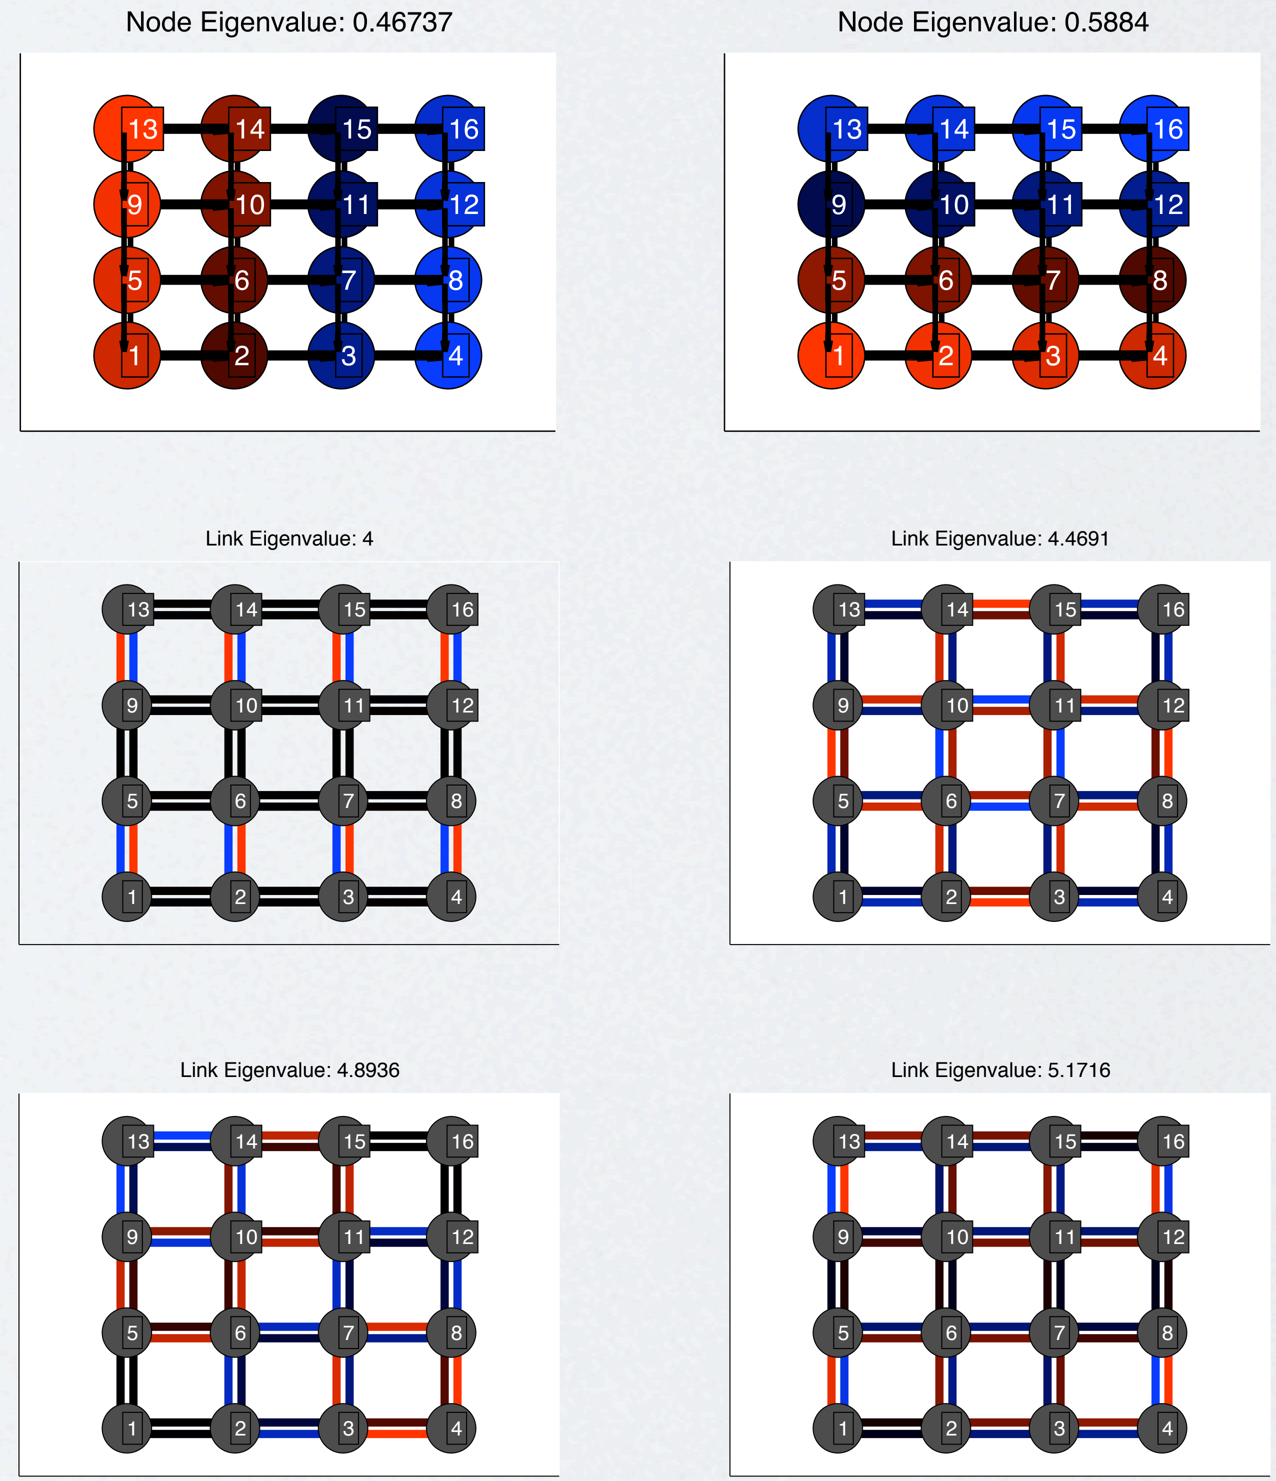

The Node Laplacian

- The divergence of the gradient
  - Maps node functions to node functions

$$\mathcal{L}_N = DG = DD^T$$

  - Measures the difference between the value of a node function in a node and its average value in the neighborhood of the node
  - Its eigenvectors provide a basis for node functions: a ***node basis***

$$\begin{pmatrix} L_1 \\ L_2 \\ L_3 \\ L_4 \end{pmatrix} = \begin{pmatrix} 2F_1 - F_2 - F_3 \\ 3F_2 - F_1 - F_3 - F_4 \\ 3F_3 - F_1 - F_2 - F_4 \\ 2F_4 - F_2 - F_3 \end{pmatrix}$$

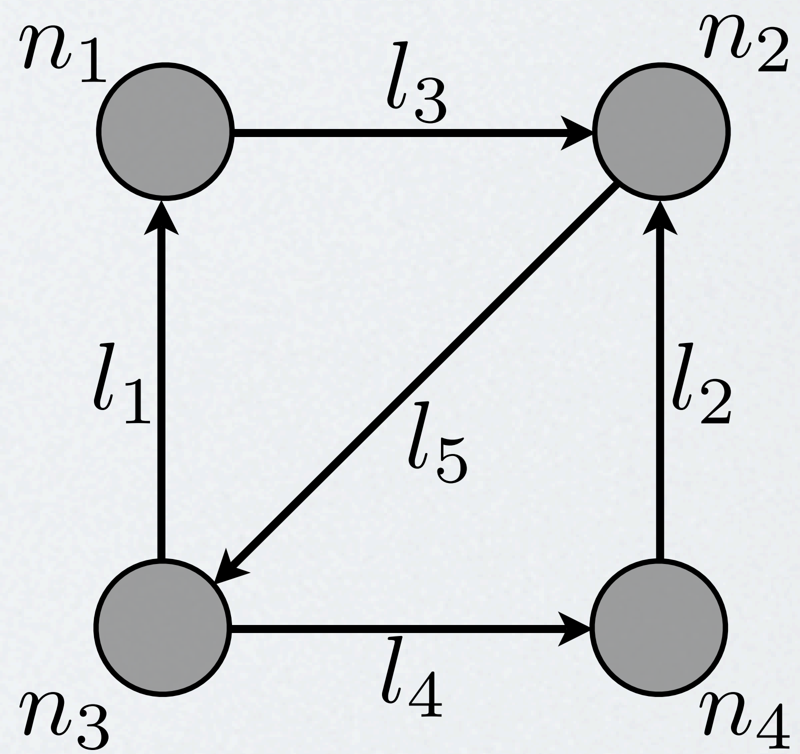

Node Laplacian Eigenfunctions

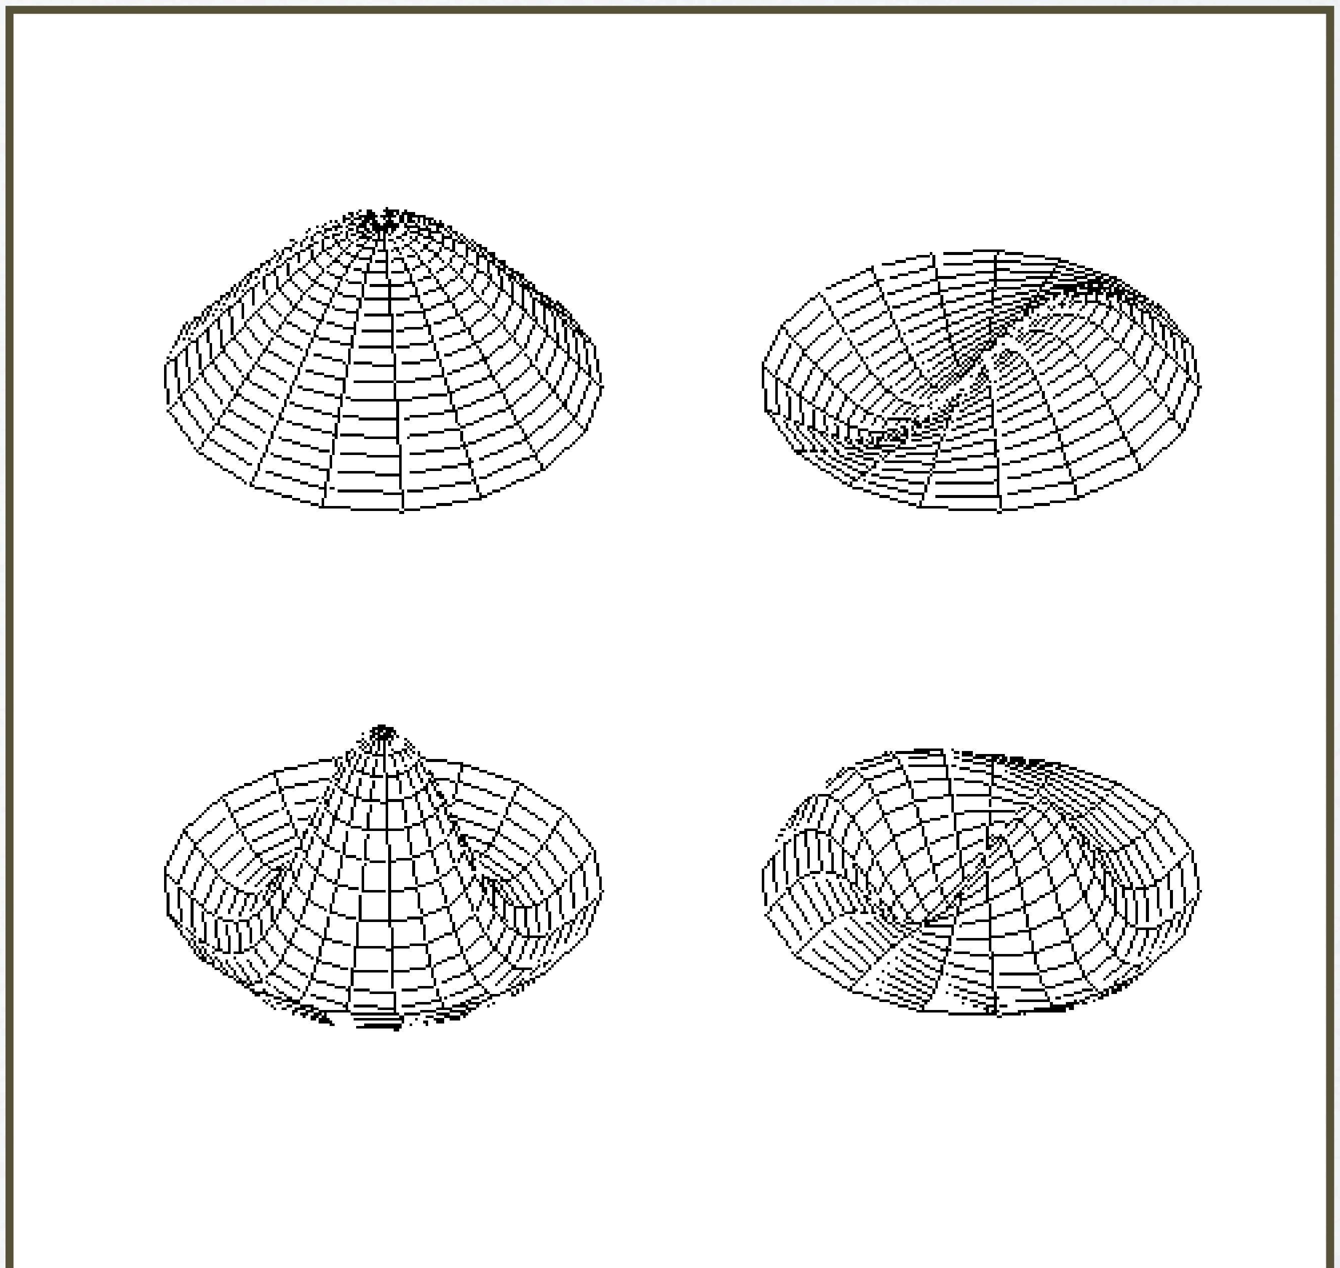

http://www.kettering.edu/~drussell/Demos/MembraneCircle/Circle.html

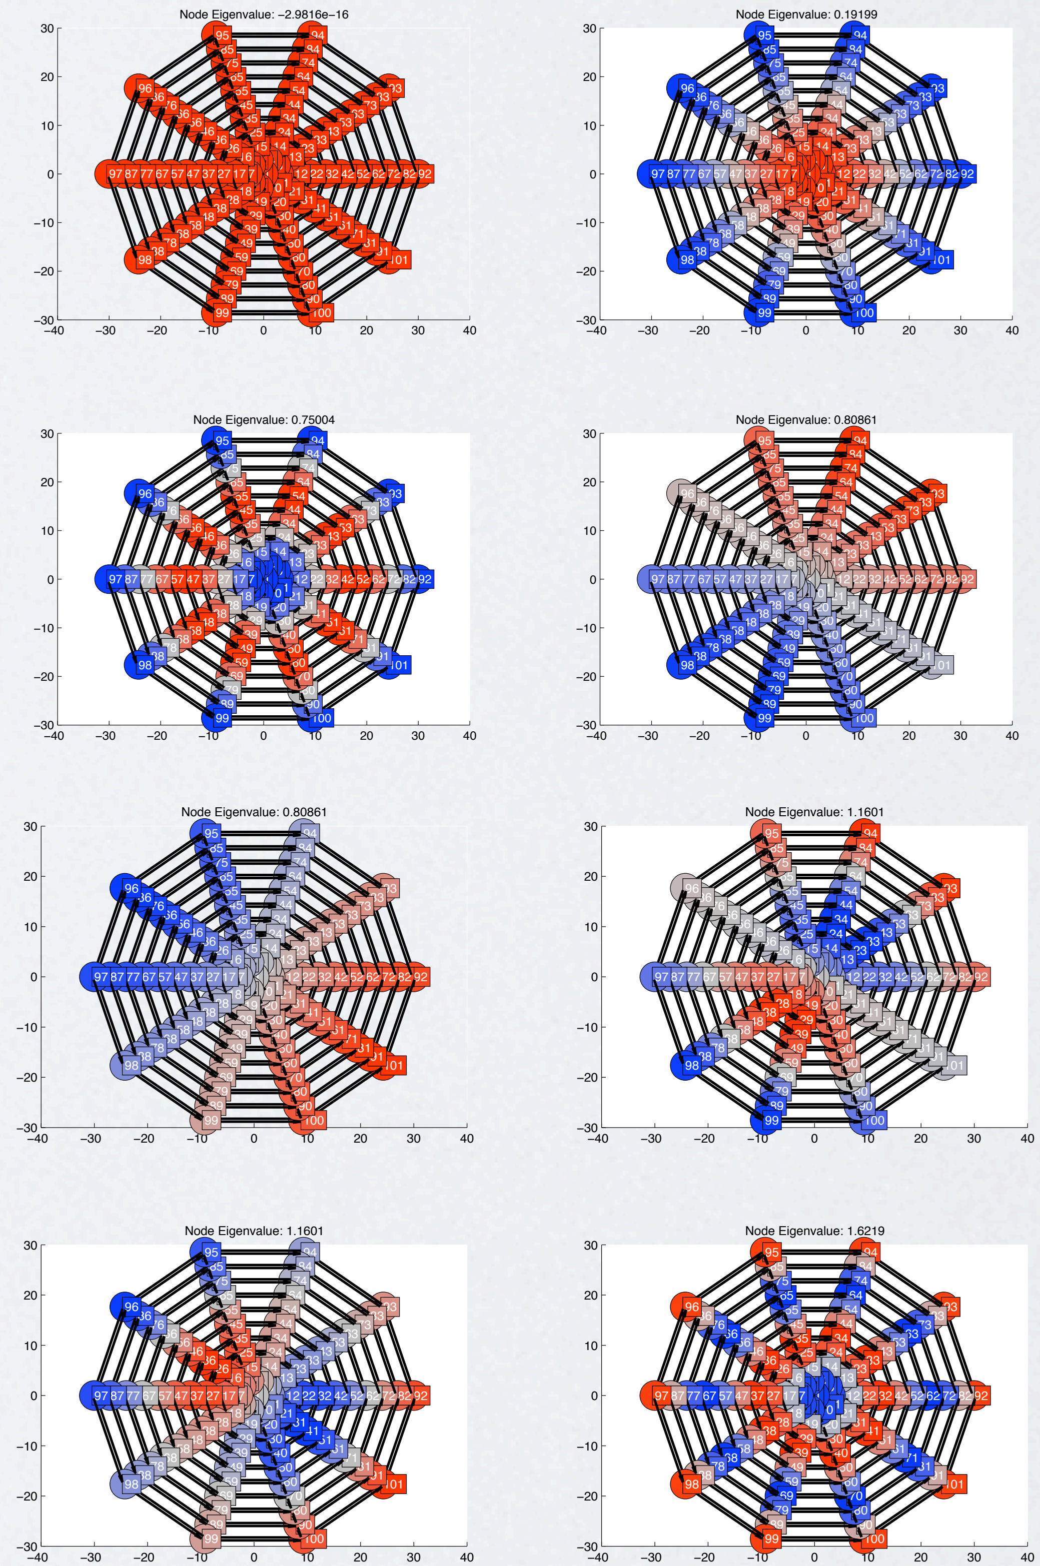

The Irrotational Laplacian

- The divergence of the gradient
  - Maps link functions to link functions

$$\mathcal{L}_I = GD = D^T D$$

  - Its eigenvectors span the ***cut-set subspace***
    - They provide a basis for link functions defined over cut-sets (a ***cut-set basis***)

$$\begin{pmatrix} \ell_1 \\ \ell_2 \\ \ell_3 \\ \ell_4 \\ \ell_5 \end{pmatrix} = \begin{pmatrix} 2f_1 + f_4 - f_3 - f_5 \\ 2f_2 + f_3 - f_4 - f_5 \\ 2f_3 + f_2 - f_1 - f_5 \\ 2f_4 + f_1 - f_2 - f_5 \\ 2f_5 - f_1 - f_2 - f_3 - f_4 \end{pmatrix}$$

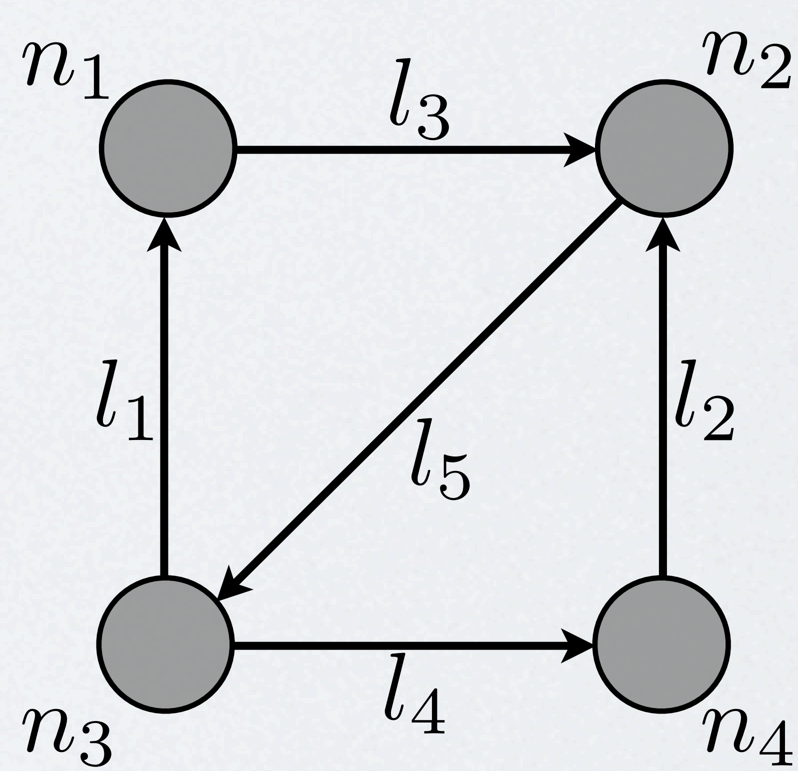

The Irrotational Laplacian

- The divergence of the gradient
  - Maps link functions to link functions

$$\mathcal{L}_I = GD = D^T D$$

  - Its eigenvectors span the ***cut-set subspace***
    - They provide a basis for link functions defined over cut-sets (a ***cut-set basis***)

Link Eigenvalue: 1.1716

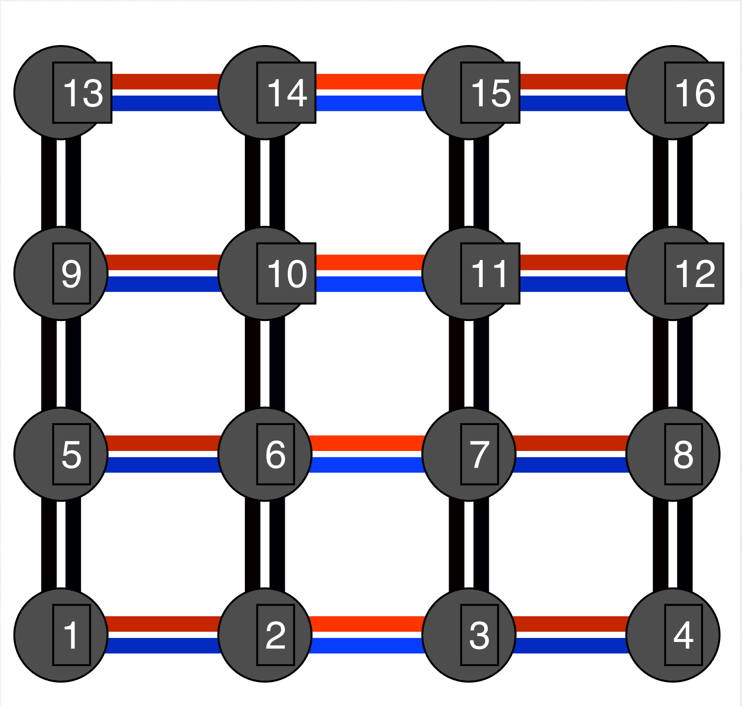

Link Eigenvalue: 1.1716

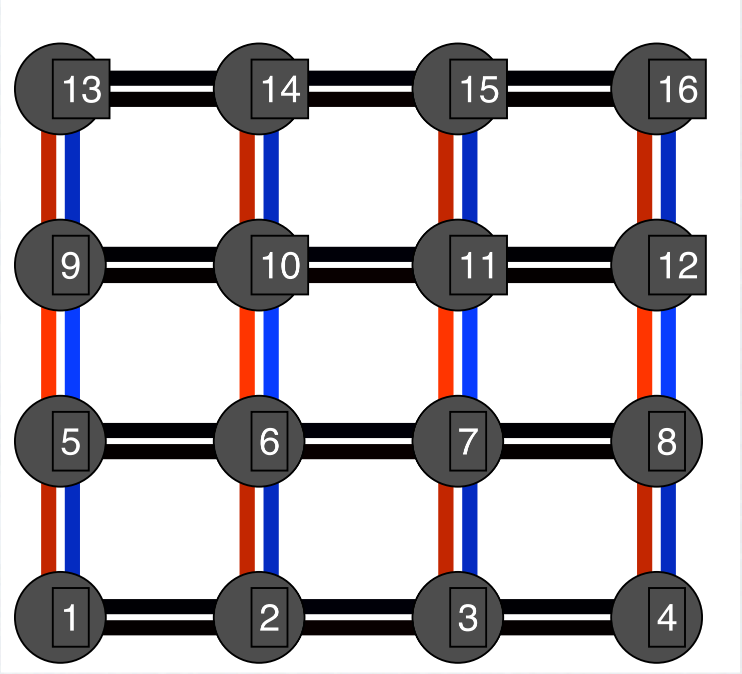

The Solenoidal Laplacian

- The adjoint curl of the curl
  - Maps link functions to link functions

$$\mathcal{L}_S = SC = C^T C$$

  - Its eigenvectors span the ***cycle subspace***:
    - They provide a basis for link functions defined over cycles (a ***cycle basis***)

$$\begin{pmatrix} \ell_1 \\ \ell_2 \\ \ell_3 \\ \ell_4 \\ \ell_5 \end{pmatrix} = \begin{pmatrix} 2f_1 + 2f_3 + f_5 - f_2 - f_4 \\ 2f_2 + 2f_4 + f_5 - f_1 - f_3 \\ 2f_3 + 2f_1 + f_5 - f_2 - f_4 \\ 2f_4 + 2f_2 + f_5 - f_1 - f_3 \\ 2f_5 + f_1 + f_2 + f_3 + f_4 \end{pmatrix}$$

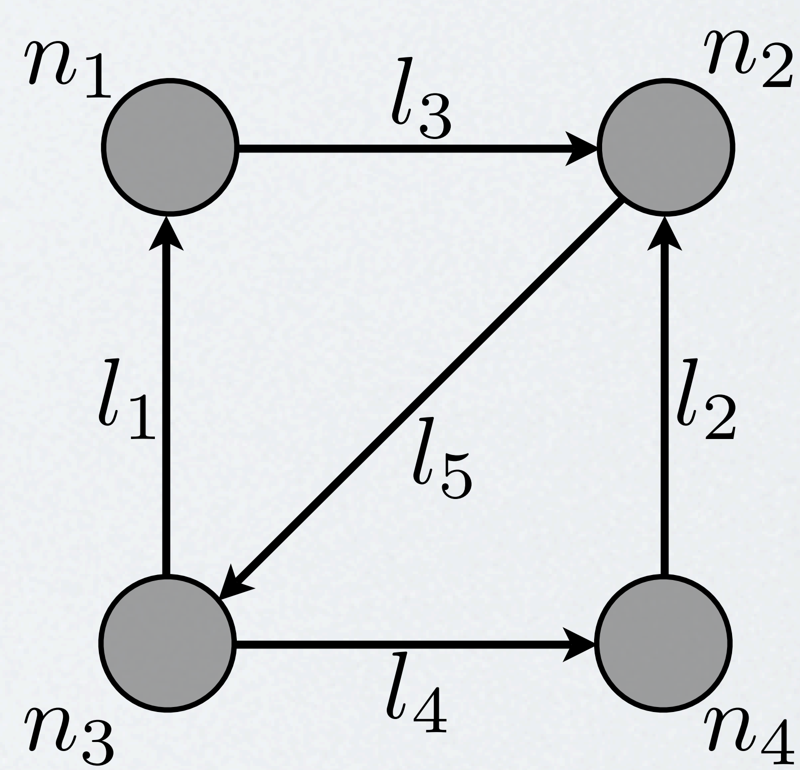

The Solenoidal Laplacian

- The adjoint curl of the curl
  - Maps link functions to link functions

$$\mathcal{L}_S = SC = C^T C$$

  - Its eigenvectors span the ***cycle subspace***:
    - They provide a basis for link functions defined over cycles (a ***cycle basis***)

Link Eigenvalue: 0.46737

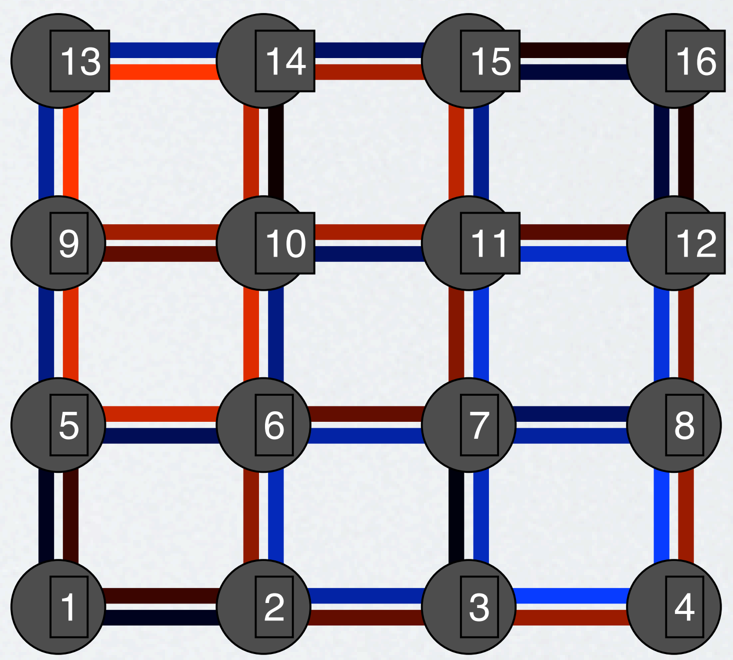

Link Eigenvalue: 0.46737

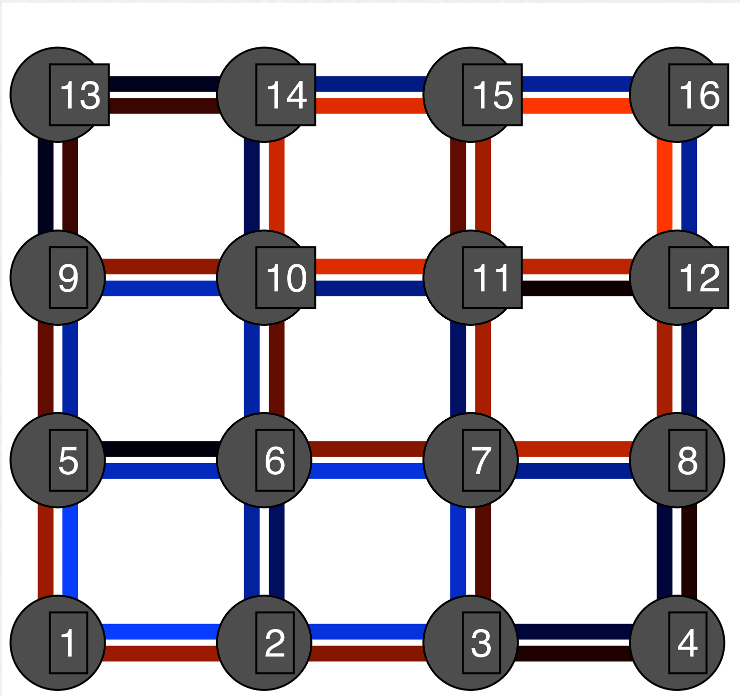

Link Laplacian Eigenfunctions

- It is easy to prove that the cycle and the cut-set subspaces are orthogonal.
  - We begin with the eigen-decompositions:

$$\mathcal{L}_S = U_S \Lambda_S U_S^T \qquad \mathcal{L}_I = U_I \Lambda_I U_I^T$$

- Given that $U_S^T U_S = I$ and $U_I^T U_I = I$, we have that:

$$U_S^T \mathcal{L}_S = \Lambda_S U_S^T \qquad \mathcal{L}_I U_I = U_I \Lambda_I$$

$$U_S^T \mathcal{L}_S \mathcal{L}_I U_I = \Lambda_S U_S^T U_I \Lambda_I$$

- But $\mathcal{L}_S \mathcal{L}_I = SCGD = 0$, because $CG = 0$.

Link Laplacian Eigenfunctions

- Thus, we have that:

$$\Lambda_S U_S^T U_I \Lambda_I = 0$$

- Thus, for all eigenvalues, the eigenvectors of the solenoidal Laplacian (the columns of U_S) are orthogonal to the eigenvectors of the irrotational Laplacian (the columns of U_I).

- The cycle subspace and the cut-set subspace are orthogonal.

The Link Laplacian

- We define the link Laplacian following the usual vector Laplacian from calculus:

$$\nabla^2 \mathbf{F} = \nabla(\nabla \cdot \mathbf{F}) - \nabla \times (\nabla \times \mathbf{F})$$

- This is equivalent to:

$$\mathcal{L}_L = \mathcal{L}_I - \mathcal{L}_S = GD - SC$$

- The link Laplacian maps link functions to link functions

Link Laplacian Eigenfunctions

Link Eigenvalue: 0.46737

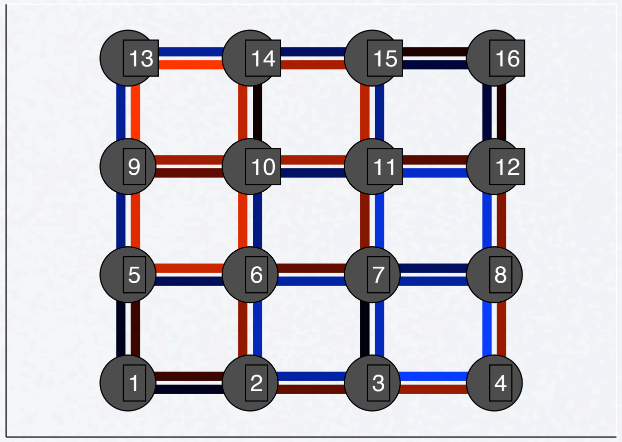

Link Eigenvalue: 0.46737

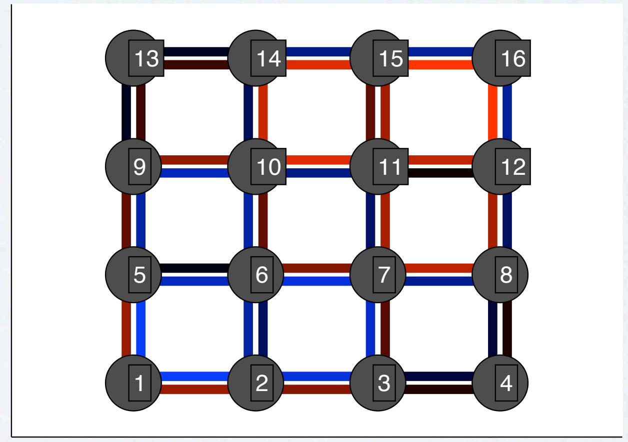

Link Eigenvalue: 0.5884

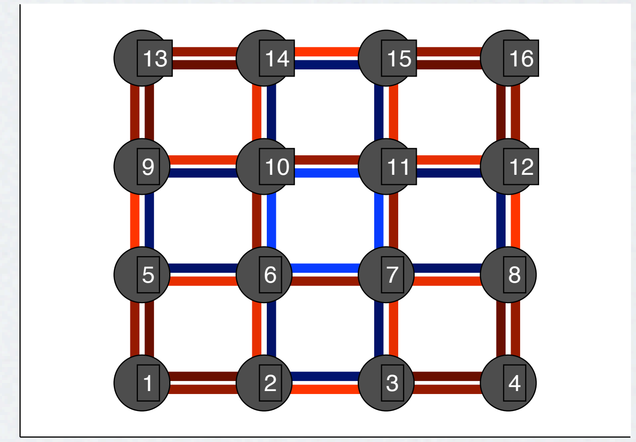

Link Eigenvalue: 0.76393

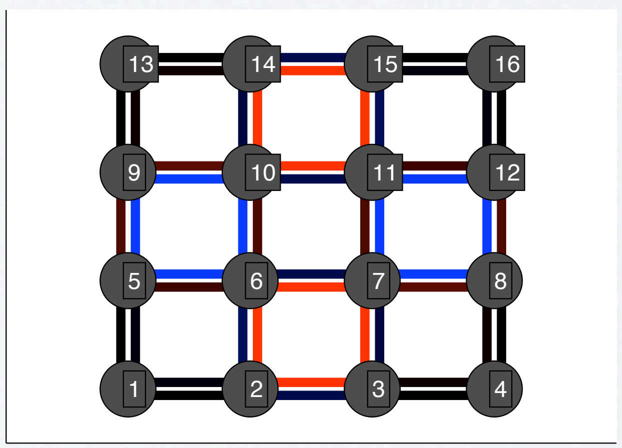

Link Eigenvalue: 0.76393

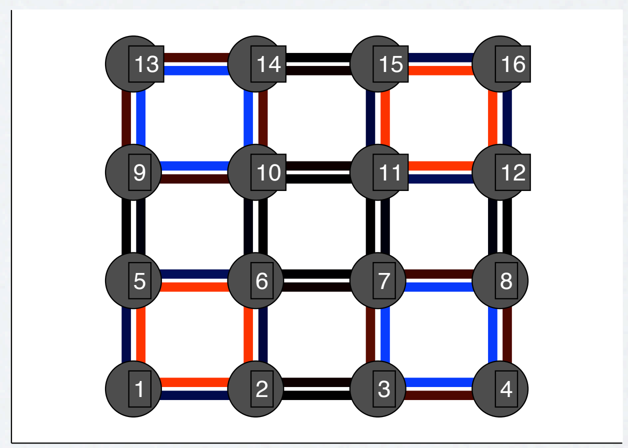

Link Eigenvalue: 1.1064

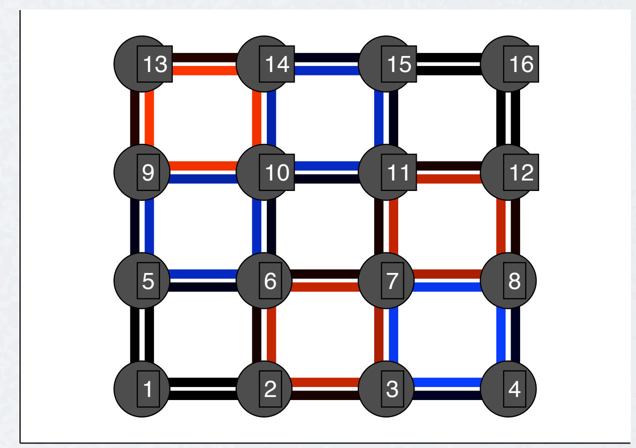

Link Eigenvalue: 1.1064

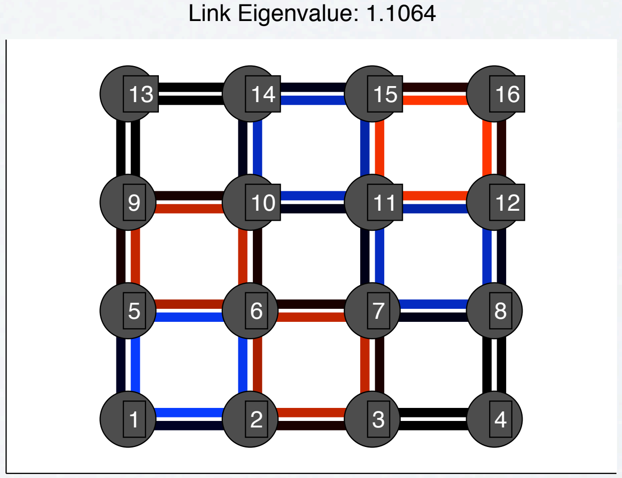

Link Eigenvalue: 1.1716

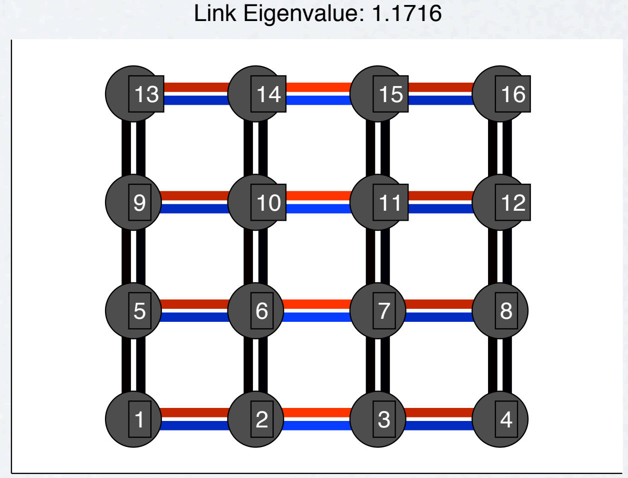

Link Eigenvalue: 1.1716

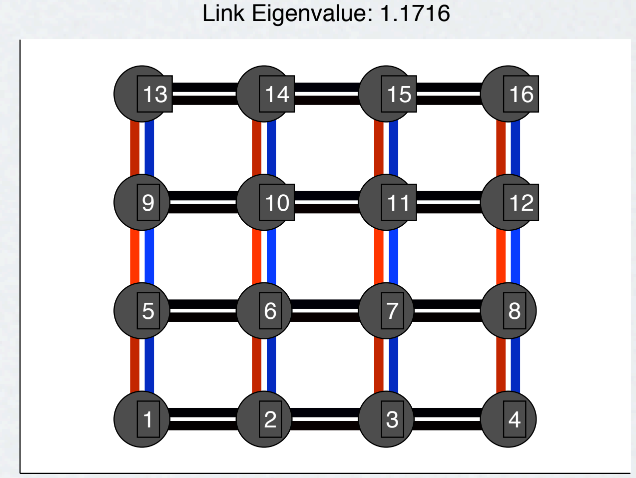

The Rank of $\mathcal{L}_S$, $\mathcal{L}_I$ and $\mathcal{L}_L$

- G and $\mathcal{L}_I$ have rank $|\mathsf{N}| - 1$
- C and $\mathcal{L}_S$ have rank $|\mathsf{C}| - 1$
- Thus, $\mathcal{L}_L$ has rank $|\mathsf{N}| + |\mathsf{C}| - 2$

- For planar graphs, the rank of $\mathcal{L}_L$ equals $|\mathsf{L}|$, due to Euler's Formula:

$$V - E + F = 2$$

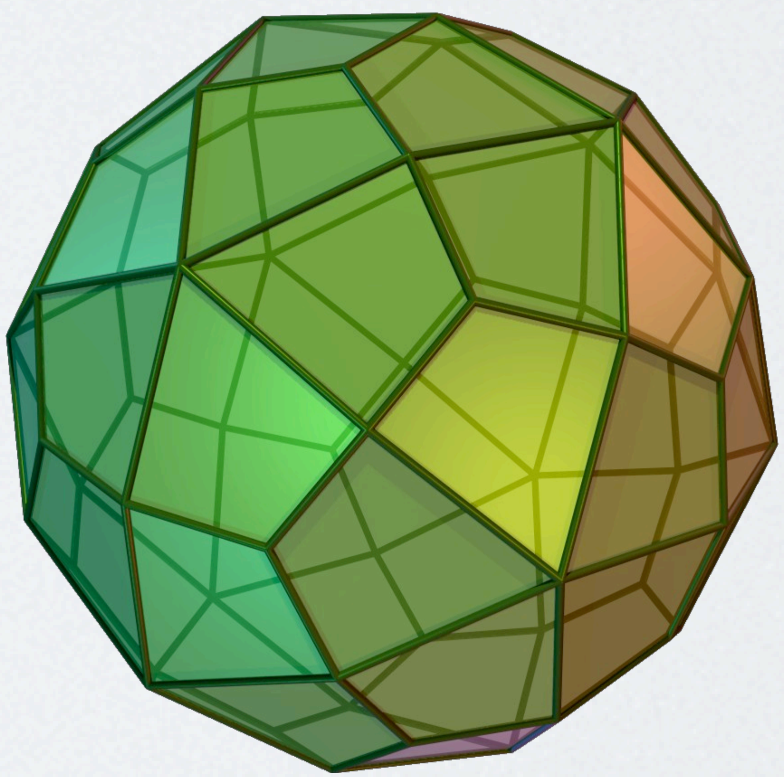

Modeling Indirect Reciprocity

- Any contribution field f can be expressed as the sum of two orthogonal components:
  - f_ψ, a superposition of flows along cycles
    - Incompressible (zero divergence)
    - Modeled through a cycle potential ψ.
  - f_ϕ, a superposition of flows through cut-sets
    - Irrotational (zero curl)
    - Modeled through a node potential ϕ.

Modeling Indirect Reciprocity

- To obtain f_ψ from f, we use the cycle projector P_ψ:

$$P_\psi = \hat{U}_S \hat{U}_S^T$$

  - Thus: $$f_\psi = P_\psi f$$

- To obtain f_ϕ from f, we use the cut-set projector P_ϕ:

$$P_\phi = \hat{U}_I \hat{U}_I^T$$

  - Thus: $$f_\phi = P_\phi f$$

- We obtain $\hat{U}_S$ and $\hat{U}_I$ by selecting from U_S or U_I the eigenvectors corresponding to nonzero eigenvalues

Calculating Potentials

- For the cut-set potential ϕ we have that:

$$P_\phi f = G\phi$$

- Since we assume that we are dealing with a connected graph, the rank of G is $|\mathrm{N}| - 1$.
  - We perform an SVD on G and discard the singular vectors related to the zero eigenvalues. We have:

$$G = \hat{U}_I \hat{\Lambda}_I^{\frac{1}{2}} \hat{V}_I^T$$

$$\hat{U}_I^T f = \hat{\Lambda}_I^{\frac{1}{2}} \hat{V}_I^T \phi$$

Calculating Potentials

- As $\hat{\Lambda}$ has full rank, we can solve for ϕ:

$$\boxed{\phi = \hat{V}_I \hat{\Lambda}_I^{-\frac{1}{2}} \hat{U}_I^T f}$$

- In the same way, if we perform SVD on S and discard zero eigenvalues:

$$S = \hat{U}_S \hat{\Lambda}_S^{\frac{1}{2}} \hat{V}_S^T$$

- Following an identical procedure, we find that:

$$\boxed{\psi = \hat{V}_S \hat{\Lambda}_S^{-\frac{1}{2}} \hat{U}_S^T f}$$

Conclusions

- Indirect Reciprocity
  - Is important for the practical deployment of overlay networks
  - Implies contribution flows built through the superposition of cycles
- Differential Operators
  - Provide basis for the cut-set and cycle spaces
  - Allow contribution fields to be decomposed in these components
- Applications?

Thank You!

Planarity and Embedding on the Sphere