Laplacian eigenvalues and optimality: II. The Laplacian of a graph

R. A. Bailey and Peter J. Cameron

Groups and Graphs, Designs and Dynamics
Yichang, China, August 2019

The Laplacian of a graph

This lecture will be about the Laplacian matrix of a graph and its eigenvalues, and their relation to some graph parameters.

The Laplacian of a graph

This lecture will be about the Laplacian matrix of a graph and its eigenvalues, and their relation to some graph parameters. This is not a complete account of the theory, but concentrates mainly on the things that are most relevant for experimental design.

The Laplacian of a graph

This lecture will be about the Laplacian matrix of a graph and its eigenvalues, and their relation to some graph parameters. This is not a complete account of the theory, but concentrates mainly on the things that are most relevant for experimental design.

For further reading, we recommend

- B. Bollobás, *Modern Graph Theory*, Springer, (especially chapters II and IX)

The Laplacian of a graph

This lecture will be about the Laplacian matrix of a graph and its eigenvalues, and their relation to some graph parameters. This is not a complete account of the theory, but concentrates mainly on the things that are most relevant for experimental design.

For further reading, we recommend

- B. Bollobás, *Modern Graph Theory*, Springer, (especially chapters II and IX)
- B. Mohar, Some applications of Laplace eigenvalues of graphs, pp.227–275 in *Graph Symmetry: Algebraic Methods and Applications* (ed. G. Hahn and G. Sabidussi), Kluwer, 1997.

Which graph is best?

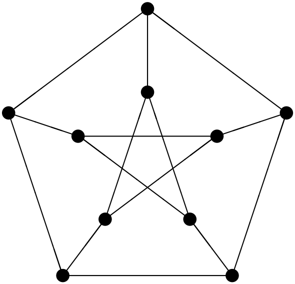

Which graph is best?

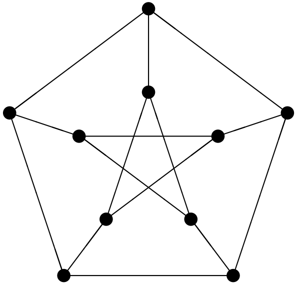

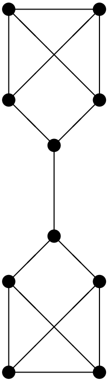

Which graph is best?

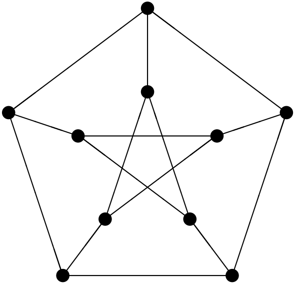

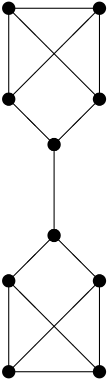

Of course the question is not well defined. But which would you choose for a network, if you were concerned about connectivity, reliability, etc.?

What makes a good network?

- No two vertices should be too far apart.

What makes a good network?

- No two vertices should be too far apart.
- There should be several alternative routes between two vertices.

What makes a good network?

- No two vertices should be too far apart.
- There should be several alternative routes between two vertices. (But should these routes be disjoint?)

What makes a good network?

- No two vertices should be too far apart.
- There should be several alternative routes between two vertices. (But should these routes be disjoint?)
- There should be no “bottlenecks”.

What makes a good network?

- No two vertices should be too far apart.
- There should be several alternative routes between two vertices. (But should these routes be disjoint?)
- There should be no "bottlenecks".
- Loss of a small part of the network should not result in disconnection.

What makes a good network?

- No two vertices should be too far apart.
- There should be several alternative routes between two vertices. (But should these routes be disjoint?)
- There should be no "bottlenecks".
- Loss of a small part of the network should not result in disconnection.

Of course, we are resource-limited, else we would just put multiple edges between any two nodes.

Trees and forests

A forest is a graph without cycles. A tree is a connected forest.

Trees and forests

A forest is a graph without cycles. A tree is a connected forest. A tree on n vertices has $n-1$ edges. More generally, a forest on n vertices which has k connected components has $n-k$ edges.

Trees and forests

A forest is a graph without cycles. A tree is a connected forest. A tree on n vertices has $n-1$ edges. More generally, a forest on n vertices which has k connected components has $n-k$ edges. Given a graph G, a spanning subgraph of G is a graph consisting of all the vertices and some of the edges of G. If it is a tree (or forest), we call it a spanning tree (or spanning forest) of G.

Trees and forests

A forest is a graph without cycles. A tree is a connected forest. A tree on n vertices has $n-1$ edges. More generally, a forest on n vertices which has k connected components has $n-k$ edges. Given a graph G, a spanning subgraph of G is a graph consisting of all the vertices and some of the edges of G. If it is a tree (or forest), we call it a spanning tree (or spanning forest) of G.

Every connected graph has a spanning tree. *Cayley's Theorem* says that the complete graph K_n has n^{n-2} spanning trees.

Which graph is best connected?

Here are some ways of measuring the “connectivity” of a graph.

Which graph is best connected?

Here are some ways of measuring the "connectivity" of a graph.

- How many spanning trees does it have? The more spanning trees, the better connected. The first graph has 2000 spanning trees, the second has 576.

Which graph is best connected?

Here are some ways of measuring the "connectivity" of a graph.

- How many spanning trees does it have? The more spanning trees, the better connected. The first graph has 2000 spanning trees, the second has 576.
- Electrical resistance. Imagine that the graph is an electrical network with each edge being a 1-ohm resistor. Now calculate the resistance between each pair of terminals, and sum over all pairs; the lower the total, the better connected. In the first graph, the sum is 33; in the second, it is $206/3$, more than twice as large.

Which graph is best connected?

- Isoperimetric number. This is defined to be

$$\iota(G) = \min\left\{\frac{|\partial S|}{|S|} : S \subseteq V(G), 0 < |S| \leq n/2\right\},$$

where $n = |V(G)|$ and for a set S of vertices, ∂S is the set of edges from S to its complement. Large isoperimetric number means that there are many edges out of any set of vertices. The isoperimetric number for the first graph is 1 (there are just five edges between the inner and outer pentagons), that of the second graph is $1/5$ (there is just one edge between the top and bottom pieces).

The Laplacian of a graph

Let G be a graph on n vertices. (Multiple edges are allowed but loops are not.)

The Laplacian of a graph

Let G be a graph on n vertices. (Multiple edges are allowed but loops are not.)
The Laplacian matrix of G is the $n \times n$ matrix $L = L(G)$ whose (i, i) entry is the number of edges containing vertex i, while for $i \neq j$ the (i, j) entry is the negative of the number of edges joining vertices i and j.

The Laplacian of a graph

Let G be a graph on n vertices. (Multiple edges are allowed but loops are not.)

The Laplacian matrix of G is the $n \times n$ matrix $L = L(G)$ whose (i, i) entry is the number of edges containing vertex i, while for $i \neq j$ the (i, j) entry is the negative of the number of edges joining vertices i and j.

This is a real symmetric matrix; its eigenvalues are the Laplacian eigenvalues of G. Note that its row sums are zero.

The weighted Laplacian

Suppose that we have positive weights $w(e)$ on the edges of G. Then the weighted Laplacian has the (i, i) entry the sum of weights of edges containing i, and whose (i, j) entry for $i \neq j$ is minus the sum of the weights of edges joining i to j.

The weighted Laplacian

Suppose that we have positive weights $w(e)$ on the edges of G. Then the weighted Laplacian has the (i, i) entry the sum of weights of edges containing i, and whose (i, j) entry for $i \neq j$ is minus the sum of the weights of edges joining i to j.
If the weights are rational, then we may multiply them by the least common multiple of the denominators to make them integers. Then replace an edge of weight w by w edges, to obtain a multigraph with the same Laplacian.

The weighted Laplacian

Suppose that we have positive weights $w(e)$ on the edges of G. Then the weighted Laplacian has the (i, i) entry the sum of weights of edges containing i, and whose (i, j) entry for $i \neq j$ is minus the sum of the weights of edges joining i to j.
If the weights are rational, then we may multiply them by the least common multiple of the denominators to make them integers. Then replace an edge of weight w by w edges, to obtain a multigraph with the same Laplacian.
For general real weights, replace them first by rational approximations.

The weighted Laplacian

Suppose that we have positive weights $w(e)$ on the edges of G. Then the weighted Laplacian has the (i, i) entry the sum of weights of edges containing i, and whose (i, j) entry for $i \neq j$ is minus the sum of the weights of edges joining i to j.
If the weights are rational, then we may multiply them by the least common multiple of the denominators to make them integers. Then replace an edge of weight w by w edges, to obtain a multigraph with the same Laplacian.
For general real weights, replace them first by rational approximations.
We will not consider weighted Laplacians.

Relation to classical Laplacian

The classical Laplacian is a second-order differential operator defined on functions on a manifold, closely related to potential theory, the wave equation, etc.

Relation to classical Laplacian

The classical Laplacian is a second-order differential operator defined on functions on a manifold, closely related to potential theory, the wave equation, etc.
A manifold can be triangulated, that is, approximated by a graph drawn in it. If we take the weight of an edge to be inversely proportional the square of its length, then the weighted Laplacian of the graph is an approximation to the Laplacian of the manifold.

Relation to classical Laplacian

The classical Laplacian is a second-order differential operator defined on functions on a manifold, closely related to potential theory, the wave equation, etc.

A manifold can be triangulated, that is, approximated by a graph drawn in it. If we take the weight of an edge to be inversely proportional the square of its length, then the weighted Laplacian of the graph is an approximation to the Laplacian of the manifold.

In the other direction, given a graph, we can build a manifold reflecting its structure. Given a d-valent vertex, take a sphere with d holes; glue spheres corresponding to the vertices of an edge together along the corresponding holes.

Relation to classical Laplacian

The classical Laplacian is a second-order differential operator defined on functions on a manifold, closely related to potential theory, the wave equation, etc.
A manifold can be triangulated, that is, approximated by a graph drawn in it. If we take the weight of an edge to be inversely proportional the square of its length, then the weighted Laplacian of the graph is an approximation to the Laplacian of the manifold.
In the other direction, given a graph, we can build a manifold reflecting its structure. Given a d-valent vertex, take a sphere with d holes; glue spheres corresponding to the vertices of an edge together along the corresponding holes.
We won't pursue this any further.

Adjacency matrix and Laplacian

The usual adjacency matrix $A(G)$ of a graph G on n vertices has rows and columns indexed by vertices; the (i, j) entry is the number of edges connecting i to j.

Adjacency matrix and Laplacian

The usual adjacency matrix $A(G)$ of a graph G on n vertices has rows and columns indexed by vertices; the (i, j) entry is the number of edges connecting i to j.
Note that we can allow loops here (though it is not clear whether a loop contributes 1 or 2 to the corresponding diagonal entry!) For the Laplacian, we forbid loops.

Adjacency matrix and Laplacian

The usual adjacency matrix $A(G)$ of a graph G on n vertices has rows and columns indexed by vertices; the (i, j) entry is the number of edges connecting i to j.
Note that we can allow loops here (though it is not clear whether a loop contributes 1 or 2 to the corresponding diagonal entry!) For the Laplacian, we forbid loops.
If G is a regular graph with valency d, then $L(G) = dI - A(G)$; so the Laplacian eigenvalues are obtained by subtracting the adjacency matrix eigenvalues from d.

Adjacency matrix and Laplacian

The usual adjacency matrix $A(G)$ of a graph G on n vertices has rows and columns indexed by vertices; the (i, j) entry is the number of edges connecting i to j.

Note that we can allow loops here (though it is not clear whether a loop contributes 1 or 2 to the corresponding diagonal entry!) For the Laplacian, we forbid loops.

If G is a regular graph with valency d, then $L(G) = dI - A(G)$; so the Laplacian eigenvalues are obtained by subtracting the adjacency matrix eigenvalues from d.

If G is not regular, there is no such simple relationship between the eigenvalues of the two matrices.

Positive semi-definiteness

A square matrix A is positive semidefinite (aka non-negative definite) if $vAv^\top \geq 0$ for all row vectors v. If A is symmetric, this is equivalent to saying that all its eigenvalues are non-negative.

Positive semi-definiteness

A square matrix A is positive semidefinite (aka non-negative definite) if $vAv^\top \geq 0$ for all row vectors v. If A is symmetric, this is equivalent to saying that all its eigenvalues are non-negative.

Theorem

The Laplacian of a graph is positive semidefinite.

Positive semi-definiteness

A square matrix A is positive semidefinite (aka non-negative definite) if $vAv^\top \geq 0$ for all row vectors v. If A is symmetric, this is equivalent to saying that all its eigenvalues are non-negative.

Theorem

The Laplacian of a graph is positive semidefinite.

For L is the sum of submatrices $\begin{pmatrix} +1 & -1 \\ -1 & +1 \end{pmatrix}$, one for each edge (this 2×2 matrix in the positions indexed by the two vertices of the edge, with zeros elsewhere). This matrix is positive semidefinite (its eigenvalues are 2 and 0.)

Positive semi-definiteness

A square matrix A is positive semidefinite (aka non-negative definite) if $vAv^\top \geq 0$ for all row vectors v. If A is symmetric, this is equivalent to saying that all its eigenvalues are non-negative.

Theorem

The Laplacian of a graph is positive semidefinite.

For L is the sum of submatrices $\begin{pmatrix} +1 & -1 \\ -1 & +1 \end{pmatrix}$, one for each edge (this 2×2 matrix in the positions indexed by the two vertices of the edge, with zeros elsewhere). This matrix is positive semidefinite (its eigenvalues are 2 and 0.)
We'll see another argument for this later.

Positive semi-definiteness

A square matrix A is positive semidefinite (aka non-negative definite) if $vAv^\top \geq 0$ for all row vectors v. If A is symmetric, this is equivalent to saying that all its eigenvalues are non-negative.

Theorem

The Laplacian of a graph is positive semidefinite.

For L is the sum of submatrices $\begin{pmatrix} +1 & -1 \\ -1 & +1 \end{pmatrix}$, one for each edge (this 2×2 matrix in the positions indexed by the two vertices of the edge, with zeros elsewhere). This matrix is positive semidefinite (its eigenvalues are 2 and 0.)
We'll see another argument for this later.
It follows that the eigenvalues of L are all non-negative.

The multiplicity of zero

Theorem

The multiplicity of 0 as an eigenvalue of L is equal to the number of connected components of the graph.

The multiplicity of zero

Theorem

The multiplicity of 0 *as an eigenvalue of* L *is equal to the number of connected components of the graph.*

An eigenvector with zero eigenvalue is a function on the vertices whose value at i is the weighted average of its values on the neighbours of i, each neighbour weighted by the number of edges joining it to i. (If you know about harmonic functions, you will recognise this!) Considering a vertex where the maximum modulus is achieved, we see that the same value occurs on all neighbours, so the function is constant on connected components.

The multiplicity of zero

Theorem

The multiplicity of 0 *as an eigenvalue of* L *is equal to the number of connected components of the graph.*

An eigenvector with zero eigenvalue is a function on the vertices whose value at i is the weighted average of its values on the neighbours of i, each neighbour weighted by the number of edges joining it to i. (If you know about harmonic functions, you will recognise this!) Considering a vertex where the maximum modulus is achieved, we see that the same value occurs on all neighbours, so the function is constant on connected components.

In particular, if the graph is connected (as we always assume), the zero eigenvalue (called “trivial”) has multiplicity 1; the other eigenvalues are nontrivial. The eigenvectors for the trivial eigenvalue are the constant vectors.

On average

Note that the sum of the eigenvalues is the trace of L, which is the sum of the vertex valencies, or twice the number of edges. So the average of the non-trivial eigenvalues is $2|E(G)|/(|V(G)|-1)$; it depends just on the numbers of vertices and edges, and the detailed structure of the graph has no effect.

On average

Note that the sum of the eigenvalues is the trace of L, which is the sum of the vertex valencies, or twice the number of edges. So the average of the non-trivial eigenvalues is $2|E(G)|/(|V(G)|-1)$; it depends just on the numbers of vertices and edges, and the detailed structure of the graph has no effect.

We'll see that other means, in particular the geometric and harmonic means, of the non-trivial eigenvalues, give us important information!

Examples

The Petersen graph is strongly regular; its adjacency matrix A satisfies $A^2 + A - 2I = J$, where J is the all-1 matrix; its eigenvalues are 3, 1 and -2, and so the Laplacian eigenvalues are 0, 2 and 5, with multiplicities 1, 5 and 4 respectively.

Examples

The Petersen graph is strongly regular; its adjacency matrix A satisfies $A^2 + A - 2I = J$, where J is the all-1 matrix; its eigenvalues are 3, 1 and -2, and so the Laplacian eigenvalues are 0, 2 and 5, with multiplicities 1, 5 and 4 respectively.

For the other graph in our introductory example, the Laplacian eigenvalues are 0, 2, 3 (multiplicity 2), 4 (multiplicity 2), 5, and the roots of $x^3 - 9x^2 + 20x - 4$ (which are approximately 0.2215, 3.2892, and 5.4893).

The Rayleigh principle

Recall that eigenvectors corresponding to distinct eigenvalues of a symmetric matrix are orthogonal.

The Rayleigh principle

Recall that eigenvectors corresponding to distinct eigenvalues of a symmetric matrix are orthogonal.
Note that I am using *row vectors* here!

The Rayleigh principle

Recall that eigenvectors corresponding to distinct eigenvalues of a symmetric matrix are orthogonal.
Note that I am using *row vectors* here!

Theorem

Let λ_1, λ_2 be the smallest and second-smallest eigenvalues of the symmetric matrix A, and suppose that λ_1 is a simple eigenvalue with eigenvector u. Let v be any non-zero vector orthogonal to u. Then

$$\frac{vAv^\top}{vv^\top} \geq \lambda_2,$$

with equality if and only if v is an eigenvector with eigenvalue λ_2.

The Rayleigh principle

Recall that eigenvectors corresponding to distinct eigenvalues of a symmetric matrix are orthogonal.
Note that I am using *row vectors* here!

Theorem

Let λ_1, λ_2 be the smallest and second-smallest eigenvalues of the symmetric matrix A, and suppose that λ_1 is a simple eigenvalue with eigenvector u. Let v be any non-zero vector orthogonal to u. Then

$$\frac{vAv^\top}{vv^\top} \geq \lambda_2,$$

with equality if and only if v is an eigenvector with eigenvalue λ_2.

This is obvious when v is expressed as a linear combination of eigenvectors of A.

The Rayleigh principle

Recall that eigenvectors corresponding to distinct eigenvalues of a symmetric matrix are orthogonal.
Note that I am using *row vectors* here!

Theorem

Let λ_1, λ_2 be the smallest and second-smallest eigenvalues of the symmetric matrix A, and suppose that λ_1 is a simple eigenvalue with eigenvector u. Let v be any non-zero vector orthogonal to u. Then

$$\frac{vAv^\top}{vv^\top} \geq \lambda_2,$$

with equality if and only if v is an eigenvector with eigenvalue λ_2.

This is obvious when v is expressed as a linear combination of eigenvectors of A.
There is an extension to any eigenvalue.

The cutset lemma

Theorem

Let G be a connected graph on n vertices, and E a set of m edges whose removal disconnects G into vertex sets of sizes n_1 and n_2, with $n_1 + n_2 = n$. Let μ be the smallest non-trivial eigenvalue of L. Then $m \geq \mu n_1 n_2 / (n_1 + n_2)$.

The cutset lemma

Theorem

Let G be a connected graph on n vertices, and E a set of m edges whose removal disconnects G into vertex sets of sizes n_1 and n_2, with $n_1 + n_2 = n$. Let μ be the smallest non-trivial eigenvalue of L. Then $m \geq \mu n_1 n_2 / (n_1 + n_2)$.

For let V_1 and V_2 be the vertex sets in the theorem, and let v be the vector with value n_2 on vertices in V_1, and $-n_1$ in vertices in V_2. These values are chosen so that v is orthogonal to the all-1 vector (the trivial eigenvector). Clearly, $vv^\top = n_1 n_2^2 + n_2 n_1^2 = (n_1 + n_2) n_1 n_2$.

The cutset lemma

Theorem

Let G be a connected graph on n vertices, and E a set of m edges whose removal disconnects G into vertex sets of sizes n_1 and n_2, with $n_1 + n_2 = n$. Let μ be the smallest non-trivial eigenvalue of L. Then $m \geq \mu n_1 n_2 / (n_1 + n_2)$.

For let V_1 and V_2 be the vertex sets in the theorem, and let v be the vector with value n_2 on vertices in V_1, and $-n_1$ in vertices in V_2. These values are chosen so that v is orthogonal to the all-1 vector (the trivial eigenvector). Clearly,
$vv^\top = n_1 n_2^2 + n_2 n_1^2 = (n_1 + n_2) n_1 n_2$.
I claim that $v^\top L v = m(n_1 + n_2)^2$. Recall that L is the sum of submatrices corresponding to edges; we have to add the contributions of these. An edge within one of the parts contributes 0; one between the parts contributes $(n_1 + n_2)^2$. The claim follows.

The cutset lemma

Theorem

Let G be a connected graph on n vertices, and E a set of m edges whose removal disconnects G into vertex sets of sizes n_1 and n_2, with $n_1 + n_2 = n$. Let μ be the smallest non-trivial eigenvalue of L. Then $m \geq \mu n_1 n_2 / (n_1 + n_2)$.

For let V_1 and V_2 be the vertex sets in the theorem, and let v be the vector with value n_2 on vertices in V_1, and $-n_1$ in vertices in V_2. These values are chosen so that v is orthogonal to the all-1 vector (the trivial eigenvector). Clearly, $vv^\top = n_1 n_2^2 + n_2 n_1^2 = (n_1 + n_2) n_1 n_2$.
I claim that $v^\top L v = m(n_1 + n_2)^2$. Recall that L is the sum of submatrices corresponding to edges; we have to add the contributions of these. An edge within one of the parts contributes 0; one between the parts contributes $(n_1 + n_2)^2$. The claim follows.
The theorem follows from the Rayleigh principle.

Isoperimetric number

Theorem

Let G be a connected graph whose smallest nontrivial Laplacian eigenvalue is μ. Then the isoperimetric number $\iota(G)$ is at least $\mu/2$.

Isoperimetric number

Theorem

Let G be a connected graph whose smallest nontrivial Laplacian eigenvalue is μ. Then the isoperimetric number $\iota(G)$ is at least $\mu/2$.

For let S be a set of at most half the vertices, and let $|S| = n_1$, $|V \setminus S| = n_2$, and $|\partial(S)| = m$. Then by the cutset lemma,

$$\frac{|\partial(S)|}{|S|} = \frac{m}{n_1} \geq \frac{\mu n_2}{n_1 + n_2} \geq \frac{\mu}{2}.$$

Isoperimetric number

Theorem

Let G be a connected graph whose smallest nontrivial Laplacian eigenvalue is μ. Then the isoperimetric number $\iota(G)$ is at least $\mu/2$.

For let S be a set of at most half the vertices, and let $|S| = n_1$, $|V \setminus S| = n_2$, and $|\partial(S)| = m$. Then by the cutset lemma,

$$\frac{|\partial(S)|}{|S|} = \frac{m}{n_1} \geq \frac{\mu n_2}{n_1 + n_2} \geq \frac{\mu}{2}.$$

So, on one of our criteria, a good network is one whose smallest nontrivial Laplacian eigenvalue is as large as possible.

Examples

In our two examples, the smallest nontrivial Laplacian eigenvalues are 2 (for the Petersen graph) and 0.2215 (for the other graph).

Examples

In our two examples, the smallest nontrivial Laplacian eigenvalues are 2 (for the Petersen graph) and 0.2215 (for the other graph).

Note that the Petersen graph has isoperimetric number 1, meeting the bound of half the least non-trivial eigenvalue. So the vector which is $+1$ on the outer pentagon and -1 on the inner pentagram is an eigenvector.

Examples

In our two examples, the smallest nontrivial Laplacian eigenvalues are 2 (for the Petersen graph) and 0.2215 (for the other graph).
Note that the Petersen graph has isoperimetric number 1, meeting the bound of half the least non-trivial eigenvalue. So the vector which is $+1$ on the outer pentagon and -1 on the inner pentagram is an eigenvector.
In the other graph, the true value is a bit more than half the smallest eigenvalue.

Expanders

Loosely, an expander is a regular connected graph whose smallest non-trivial Laplacian eigenvalue is large. The above result shows that expanders have large isoperimetric numbers.

Expanders

Loosely, an expander is a regular connected graph whose smallest non-trivial Laplacian eigenvalue is large. The above result shows that expanders have large isoperimetric numbers. More precisely, a sequence $(G_n : n \in \mathbb{N})$ is a sequence of expanders if there is a constant $c > 0$ such that the smallest non-trivial Laplacian eigenvalue of every graph G_n is at least c.

Expanders

Loosely, an expander is a regular connected graph whose smallest non-trivial Laplacian eigenvalue is large. The above result shows that expanders have large isoperimetric numbers. More precisely, a sequence $(G_n : n \in \mathbb{N})$ is a sequence of expanders if there is a constant $c > 0$ such that the smallest non-trivial Laplacian eigenvalue of every graph G_n is at least c. It is known that a random regular graph is an expander with high probability; but explicit constructions are more difficult. The first constructions were given by Margulis and by Lubotzky, Phillips and Sarnak, and depend on substantial number-theoretic and group-theoretic background.

Cheeger's inequality

Cheeger's inequality is a result about Laplacians of manifolds; it has a discrete analogue. It gives a bound in the other direction between the isoperimetric number and the smallest nontrivial Laplacian eigenvalue.

Cheeger's inequality

Cheeger's inequality is a result about Laplacians of manifolds; it has a discrete analogue. It gives a bound in the other direction between the isoperimetric number and the smallest nontrivial Laplacian eigenvalue.

Theorem

Let G be a connected graph; let Δ be the maximum valency of G, and μ the smallest nontrivial Laplacian eigenvalue. Then

$$\iota(G) \leq \sqrt{2\Delta\mu}.$$

Cheeger's inequality

Cheeger's inequality is a result about Laplacians of manifolds; it has a discrete analogue. It gives a bound in the other direction between the isoperimetric number and the smallest nontrivial Laplacian eigenvalue.

Theorem

Let G be a connected graph; let Δ be the maximum valency of G, and μ the smallest nontrivial Laplacian eigenvalue. Then

$$\iota(G) \leq \sqrt{2\Delta\mu}.$$

Mohar improved the upper bound to $\sqrt{(2\Delta - \mu)\mu}$ if the graph is connected but not complete.

Incidence matrix

Choose a fixed but arbitrary orientation of the edges of the graph G. Define the vertex-edge incidence matrix Q to have rows indexed by vertices, columns by edges, and (v, e) entry $+1$ if v is the head of the edge e, -1 if v is the tail of e, and 0 otherwise.

Incidence matrix

Choose a fixed but arbitrary orientation of the edges of the graph G. Define the vertex-edge incidence matrix Q to have rows indexed by vertices, columns by edges, and (v, e) entry $+1$ if v is the head of the edge e, -1 if v is the tail of e, and 0 otherwise.

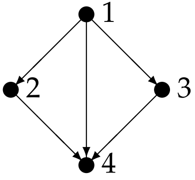

$$\begin{pmatrix} -1 & -1 & -1 & 0 & 0 \\ +1 & 0 & 0 & -1 & 0 \\ 0 & 0 & +1 & 0 & -1 \\ 0 & +1 & 0 & +1 & +1 \end{pmatrix}$$

Incidence matrix and Laplacian

Theorem

Let G have incidence matrix Q and Laplacian L. Then $QQ^\top = L$.

Incidence matrix and Laplacian

Theorem

Let G have incidence matrix Q and Laplacian L. Then $QQ^\top = L$.

For the (v, v) entry of $QQ^\top$ is the number of edges with either head or tail at v; and the (v, w) entry is the sum of -1 for all edges with head at v and tail at w or *vice versa*.

Incidence matrix and Laplacian

Theorem

Let G have incidence matrix Q and Laplacian L. Then $QQ^\top = L$.

For the (v, v) entry of $QQ^\top$ is the number of edges with either head or tail at v; and the (v, w) entry is the sum of -1 for all edges with head at v and tail at w or *vice versa*.

This shows, again, that L is positive semidefinite. And note that the orientation doesn't matter.

The Moore–Penrose inverse

Let A be a real symmetric matrix. Then we have a spectral decomposition of A:

$$A = \sum_{\lambda \in \Lambda} \lambda P_\lambda,$$

where Λ is the set of eigenvalues of A, and P_λ is the orthogonal projector onto the space of eigenvectors with eigenvalue λ.

The Moore–Penrose inverse

Let A be a real symmetric matrix. Then we have a spectral decomposition of A:

$$A = \sum_{\lambda \in \Lambda} \lambda P_\lambda,$$

where Λ is the set of eigenvalues of A, and P_λ is the orthogonal projector onto the space of eigenvectors with eigenvalue λ.
We define the Moore–Penrose inverse of A by

$$A^- = \sum_{\lambda \neq 0} \lambda^{-1} P_\lambda.$$

In other words, we invert A where we can.

The Moore–Penrose inverse

Let A be a real symmetric matrix. Then we have a spectral decomposition of A:

$$A = \sum_{\lambda \in \Lambda} \lambda P_\lambda,$$

where Λ is the set of eigenvalues of A, and P_λ is the orthogonal projector onto the space of eigenvectors with eigenvalue λ.
We define the Moore–Penrose inverse of A by

$$A^- = \sum_{\lambda \neq 0} \lambda^{-1} P_\lambda.$$

In other words, we invert A where we can.
The Moore–Penrose inverse is a quasi-inverse of A in the sense of ring theory: that is,

$$A^- A A^- = A^-, \qquad A A^- A = A.$$

The Moore–Penrose inverse of the Laplacian

We will see a lot of the matrix L^{-}, where L is the Laplacian of a graph G on n vertices.

The Moore–Penrose inverse of the Laplacian

We will see a lot of the matrix L^{-}, where L is the Laplacian of a graph G on n vertices.
If G is connected, then the projector onto the trivial eigenspace is J/n, where J is the all-1 matrix. So adding this to L changes the trivial eigenvalue to 1, and subtracting it takes the 1 off again.

The Moore–Penrose inverse of the Laplacian

We will see a lot of the matrix L^-, where L is the Laplacian of a graph G on n vertices.
If G is connected, then the projector onto the trivial eigenspace is J/n, where J is the all-1 matrix. So adding this to L changes the trivial eigenvalue to 1, and subtracting it takes the 1 off again.
In other words,

$$L^- = (L + J/n)^{-1} - J/n.$$

Electrical networks

As mentioned earlier, we regard the graph G as an electrical network, where we regard each edge as a one-ohm resistor. Given any two vertices i and j, the effective resistance $R(i, j)$ between i and j is the voltage of a battery which, when connected to the two vertices, causes a current of 1 ampere to flow.

Electrical networks

As mentioned earlier, we regard the graph G as an electrical network, where we regard each edge as a one-ohm resistor. Given any two vertices i and j, the effective resistance $R(i, j)$ between i and j is the voltage of a battery which, when connected to the two vertices, causes a current of 1 ampere to flow.

In order to calculate $R(i, j)$, we need to calculate the currents and potentials in the network. This requires some basic results from 19th century physics: Ohm's and Kirchhoff's laws.

Ohm's Law and Kirchhoff's Laws

- Ohm's Law: the potential drop in each edge is the product of the current and the resistance (and so is equal to the current since we have set all resistances to 1).

Ohm's Law and Kirchhoff's Laws

- Ohm's Law: the potential drop in each edge is the product of the current and the resistance (and so is equal to the current since we have set all resistances to 1).
- Kirchhoff's Voltage Law: the sum of the potential drops on any path between vertices i and j is independent of the choice of path.

Ohm's Law and Kirchhoff's Laws

- Ohm's Law: the potential drop in each edge is the product of the current and the resistance (and so is equal to the current since we have set all resistances to 1).
- Kirchhoff's Voltage Law: the sum of the potential drops on any path between vertices i and j is independent of the choice of path.
- Kirchhoff's Current Law: if vertex i is not connected to the battery, then the sum of the currents flowing into i is equal to the sum of the currents flowing out.

Resistance distance

I begin by showing that effective resistance defines a metric on the vertex set of the graph.

Resistance distance

I begin by showing that effective resistance defines a metric on the vertex set of the graph.

Theorem

R is a metric on the vertex set of the graph.

Resistance distance

I begin by showing that effective resistance defines a metric on the vertex set of the graph.

Theorem

R is a metric on the vertex set of the graph.

The proof depends on a lemma:

Lemma

If we connect the terminals of a battery to two vertices i and j, then the potential of any other vertex lies between the potentials of i and j.

Proof of the lemma

First observe that, if h is any vertex other than a terminal, then the net current into h is zero; by Ohm's law, this implies that the potential on h is the average of the potentials of its neighbours. So potential is a harmonic function (thinking of the graph as a "manifold" whose "boundary" is the pair of terminals).

Proof of the lemma

First observe that, if h is any vertex other than a terminal, then the net current into h is zero; by Ohm's law, this implies that the potential on h is the average of the potentials of its neighbours. So potential is a harmonic function (thinking of the graph as a "manifold" whose "boundary" is the pair of terminals).

Let h be a vertex with smallest potential, and suppose that it is a non-terminal vertex. Then all neighbours of v must have the same potential as h.

Proof of the lemma

First observe that, if h is any vertex other than a terminal, then the net current into h is zero; by Ohm's law, this implies that the potential on h is the average of the potentials of its neighbours. So potential is a harmonic function (thinking of the graph as a "manifold" whose "boundary" is the pair of terminals).

Let h be a vertex with smallest potential, and suppose that it is a non-terminal vertex. Then all neighbours of v must have the same potential as h.

By induction and connectedness, all vertices have the same potential, a contradiction.

Proof of the lemma

First observe that, if h is any vertex other than a terminal, then the net current into h is zero; by Ohm's law, this implies that the potential on h is the average of the potentials of its neighbours. So potential is a harmonic function (thinking of the graph as a "manifold" whose "boundary" is the pair of terminals).
Let h be a vertex with smallest potential, and suppose that it is a non-terminal vertex. Then all neighbours of v must have the same potential as h.
By induction and connectedness, all vertices have the same potential, a contradiction.
So the smallest and greatest potential are realised at terminals, as required.

Proof of the theorem

Clearly $R(i,j)$ is non-negative, and zero only if $i = j$; also it is symmetric in i and j. We have to prove the triangle inequality: $R(i,j) + R(j,k) \geq R(i,k)$.

Proof of the theorem

Clearly $R(i,j)$ is non-negative, and zero only if $i = j$; also it is symmetric in i and j. We have to prove the triangle inequality: $R(i,j) + R(j,k) \geq R(i,k)$.
Let $r_1 = R(i,j)$, and $r_2 = R(j,k)$. We define two current flows in the graph as follows. For the first one, we connect a battery with voltage r_1 to i and j (with i at the higher potential). This causes unit current to flow out of i and into j. If P_1 denotes the corresponding potential, then $P_1(i) - P_1(j) = r_1$, and $P_1(j) - P_1(k) \leq 0$, by the lemma. So $P_1(i) - P_1(k) \leq r_1$.

Similarly, a battery of voltage r_2 connected to j and k causes a unit current to flow out of j and into k; the potential satisfies $P_2(j) - P_2(k) = r_2$, and $P_2(i) - P_2(j) \leq 0$. So $P_2(i) - P_2(k) \leq r_2$.

Similarly, a battery of voltage r_2 connected to j and k causes a unit current to flow out of j and into k; the potential satisfies $P_2(j) - P_2(k) = r_2$, and $P_2(i) - P_2(j) \leq 0$. So $P_2(i) - P_2(k) \leq r_2$. Since Ohm's and Kirchhoff's Laws are linear, we can add these two solutions. The resulting solution has a unit current flowing out of i and into k. With $P = P_1 + P_2$, we thus have $P(k) - P(i) = R(i,k)$, and so

$$R(i,k) = (P_1(k) - P_1(i)) + (P_2(k) - P_2(i)) \leq r_1 + r_2.$$

Resistance distance

In view of this theorem, we refer to the metric $R(i, j)$ as the resistance distance.

Resistance distance

In view of this theorem, we refer to the metric $R(i, j)$ as the resistance distance.
This is an important metric, with many applications. Michael Kagan's conference talk gave much further information from both the points of view of mathematics and physics, including other methods of computing the metric, and applications to a process called resistance distance transform of a graph, related to Weisfeiler–Lehman stabilisation.

Resistance distance and the Laplacian

We now show how the resistance distance can be calculated from the Laplacian matrix of the graph.

Resistance distance and the Laplacian

We now show how the resistance distance can be calculated from the Laplacian matrix of the graph.

Theorem

Let G be a connected graph with Laplacian L. Then the effective resistance between i and j is

$$L^-_{ii} + L^-_{jj} - L^-_{ij} - L^-_{ji},$$

where L^- is the Moore–Penrose inverse of L.

Proof of the theorem

Kirchhoff's voltage law and Ohm's Law are taken care of if we take a vector p of potentials with components indexed by vertices, and require that the current on the edge e is equal to the potential difference between its ends. (As before, we take a fixed ordering of each edge, and take the current to be negative if it flows from head to tail of the edge.) Note that p is defined up to adding a constant vector.

Proof of the theorem

Kirchhoff's voltage law and Ohm's Law are taken care of if we take a vector p of potentials with components indexed by vertices, and require that the current on the edge e is equal to the potential difference between its ends. (As before, we take a fixed ordering of each edge, and take the current to be negative if it flows from head to tail of the edge.) Note that p is defined up to adding a constant vector.

This is expressed by the requirement that pQ is the vector of currents in the edges, where Q is the vertex-edge incidence matrix.

Then $pQQ^\top = pL$ is a vector whose ith entry is the sum of the signed currents into the vertex i. So Kirchhoff's current law says that $pQQ^\top$ has all entries zero except at the two vertices connected to the battery. If the current is 1 ampere, the entries in $pL = pQQ^\top$ are $+1$ and -1 on these two vertices. Let us write $pL = f_i - f_j$, where f_i is the unit basis vector corresponding to vertex i.

Then $pQQ^\top = pL$ is a vector whose ith entry is the sum of the signed currents into the vertex i. So Kirchhoff's current law says that $pQQ^\top$ has all entries zero except at the two vertices connected to the battery. If the current is 1 ampere, the entries in $pL = pQQ^\top$ are $+1$ and -1 on these two vertices. Let us write $pL = f_i - f_j$, where f_i is the unit basis vector corresponding to vertex i.

Now $f_i - f_j$ is orthogonal to the all-1 vector, so $(f_i - f_j)L^- = p$. This gives the vector of potentials. The potential difference between i and j is the dot product of this vector with $f_i - f_j$, that is, $xL^-x^\top$, where $x = f_i - f_j$. This is the potential difference required to make a current of 1 ampere flow; hence it is the effective resistance between i and j. This can be written

$$R(i,j) = L^-_{ii} + L^-_{jj} - L^-_{ij} - L^-_{ji},$$

as required.

The average pairwise resistance

One of our criteria for a good network is that the average pairwise resistance between two vertices should be small. The next theorem shows that this is equivalent to maximizing the harmonic mean of the nontrivial Laplacian eigenvalues.

The average pairwise resistance

One of our criteria for a good network is that the average pairwise resistance between two vertices should be small. The next theorem shows that this is equivalent to maximizing the harmonic mean of the nontrivial Laplacian eigenvalues.

Theorem

The average pairwise resistance is equal to 2 *divided by the harmonic mean of the nontrivial Laplacian eigenvalues.*

Proof of the resistance theorem

The sum of the resistances between all ordered pairs of vertices is

$$\sum_{i \neq j} R(i,j) = 2(n-1)\,\mathrm{Trace}(L^{-}) - 2\sum_{i \neq j} L^{-}_{ij} = 2n\,\mathrm{Trace}(L^{-}),$$

since the sum of all elements of L^{-} is zero (as the all-1 vector is an eigenvector with eigenvalue 0).

Proof of the resistance theorem

The sum of the resistances between all ordered pairs of vertices is

$$\sum_{i\neq j} R(i,j) = 2(n-1)\,\mathrm{Trace}(L^-) - 2\sum_{i\neq j} L^-_{ij} = 2n\,\mathrm{Trace}(L^-),$$

since the sum of all elements of L^- is zero (as the all-1 vector is an eigenvector with eigenvalue 0).
So the average pairwise resistance is $2\,\mathrm{Trace}(L^-)/(n-1)$.

Proof of the resistance theorem

The sum of the resistances between all ordered pairs of vertices is

$$\sum_{i \neq j} R(i,j) = 2(n-1)\,\mathrm{Trace}(L^-) - 2\sum_{i \neq j} L^-_{ij} = 2n\,\mathrm{Trace}(L^-),$$

since the sum of all elements of L^- is zero (as the all-1 vector is an eigenvector with eigenvalue 0).
So the average pairwise resistance is $2\,\mathrm{Trace}(L^-)/(n-1)$.
Now the trace of L^- is the sum of the reciprocals of the non-zero eigenvalues of L, and so we are done.

Examples

For the Petersen graph, the harmonic mean of the non-trivial eigenvalues is

$$(((5 \cdot 1/2) + (4 \cdot 1/5))/9)^{-1} = 30/11,$$

so the average resistance is $11/15$.
For the other graph, a similar calculation gives $135/103$, so the average resistance is $206/135$.

Example: Petersen graph

For the Petersen graph, we can exploit symmetry to calculate the resistance between two terminals. Two vertices equivalent under a symmetry fixing the terminals must be at the same potential, and so edges between them can be neglected.

Example: Petersen graph

For the Petersen graph, we can exploit symmetry to calculate the resistance between two terminals. Two vertices equivalent under a symmetry fixing the terminals must be at the same potential, and so edges between them can be neglected.
If the terminals i and j are joined, the graph reduces to a pentagon $i = i_0, i_1, i_2, i_3, i_4 = j$, with one edge from i to j, two from i_0 to i_1 and from i_3 to i_4, and four from i_1 to i_2 and i_2 to i_3.
So the resistance of the path $(i_0, i_1, i_2, i_3, i_4)$ is $1/2 + 1/4 + 1/4 + 1/2 = 3/2$. This is in parallel with a single edge, so the overall resistance is $1/(1 + 2/3) = 3/5$.

Example: Petersen graph

For the Petersen graph, we can exploit symmetry to calculate the resistance between two terminals. Two vertices equivalent under a symmetry fixing the terminals must be at the same potential, and so edges between them can be neglected.
If the terminals i and j are joined, the graph reduces to a pentagon $i = i_0, i_1, i_2, i_3, i_4 = j$, with one edge from i to j, two from i_0 to i_1 and from i_3 to i_4, and four from i_1 to i_2 and i_2 to i_3.
So the resistance of the path $(i_0, i_1, i_2, i_3, i_4)$ is $1/2 + 1/4 + 1/4 + 1/2 = 3/2$. This is in parallel with a single edge, so the overall resistance is $1/(1 + 2/3) = 3/5$.
Similar but slightly more complicated arguments give the resistance between non-adjacent terminals as $4/5$. (Exercise: prove this using symmetry.)

Example: Petersen graph

For the Petersen graph, we can exploit symmetry to calculate the resistance between two terminals. Two vertices equivalent under a symmetry fixing the terminals must be at the same potential, and so edges between them can be neglected.
If the terminals i and j are joined, the graph reduces to a pentagon $i = i_0, i_1, i_2, i_3, i_4 = j$, with one edge from i to j, two from i_0 to i_1 and from i_3 to i_4, and four from i_1 to i_2 and i_2 to i_3. So the resistance of the path $(i_0, i_1, i_2, i_3, i_4)$ is $1/2 + 1/4 + 1/4 + 1/2 = 3/2$. This is in parallel with a single edge, so the overall resistance is $1/(1 + 2/3) = 3/5$.
Similar but slightly more complicated arguments give the resistance between non-adjacent terminals as $4/5$. (Exercise: prove this using symmetry.)
So the total is $15 \cdot 3/5 + 30 \cdot 4/5 = 33$, and the average is $33/45 = 11/15$, agreeing with the eigenvalue calculation.

The Matrix-Tree Theorem

Theorem

Let G be a connected graph on n vertices. Then the following three quantities are equal:

1. *the number of spanning trees of G;*
2. $(\lambda_2 \cdots \lambda_n)/n$, *where* $\lambda_2, \ldots, \lambda_n$ *are the nontrivial Laplacian eigenvalues of G;*
3. *any cofactor of $L(G)$ (that is, the determinant of the matrix obtained by deleting row i and column j, multiplied by $(-1)^{i+j}$).*

The Matrix-Tree Theorem

Theorem

Let G be a connected graph on n vertices. Then the following three quantities are equal:

1. *the number of spanning trees of G;*
2. $(\lambda_2 \cdots \lambda_n)/n$*, where $\lambda_2, \ldots, \lambda_n$ are the nontrivial Laplacian eigenvalues of G;*
3. *any cofactor of $L(G)$ (that is, the determinant of the matrix obtained by deleting row i and column j, multiplied by $(-1)^{i+j}$).*

Since one of our criteria for a good network is a large number of spanning trees, this is equivalent to maximizing the geometric mean of the non-trivial Laplacian eigenvalues.

The Cauchy–Binet formula

The proof depends on the Cauchy–Binet formula , which says the following:

Theorem

Let A be an $m \times n$ matrix, and B an $n \times m$ matrix, where $m < n$. Then

$$\det(AB) = \sum_{X} \det(A(X)) \det(B(X)),$$

where X ranges over all m-element subsets of $\{1, \ldots, n\}$. Here $A(X)$ is the $m \times m$ matrix whose columns are the columns of A with index in X, and $B(X)$ is the $m \times m$ matrix whose rows are the rows of B with index in X.

The Cauchy–Binet formula

The proof depends on the Cauchy–Binet formula , which says the following:

Theorem

Let A be an $m \times n$ matrix, and B an $n \times m$ matrix, where $m < n$. Then

$$\det(AB) = \sum_X \det(A(X)) \det(B(X)),$$

where X ranges over all m-element subsets of $\{1, \ldots, n\}$. Here $A(X)$ is the $m \times m$ matrix whose columns are the columns of A with index in X, and $B(X)$ is the $m \times m$ matrix whose rows are the rows of B with index in X.

The proof is an exercise.

Proof of the Matrix-Tree Theorem

Let Q be the incidence matrix of G, so that $QQ^\top = L$. Let i be any vertex of G, and let $N = Q_i$ be the matrix obtained by deleting the row of Q indexed by i.

Proof of the Matrix-Tree Theorem

Let Q be the incidence matrix of G, so that $QQ^\top = L$. Let i be any vertex of G, and let $N = Q_i$ be the matrix obtained by deleting the row of Q indexed by i.

It can be shown that, if X is a set of $n-1$ edges, then $\det(N(X))$ is ± 1 if X is the edge set of a spanning tree, and is 0 otherwise. (The proof of this is also an exercise.)

Proof of the Matrix-Tree Theorem

Let Q be the incidence matrix of G, so that $QQ^\top = L$. Let i be any vertex of G, and let $N = Q_i$ be the matrix obtained by deleting the row of Q indexed by i.

It can be shown that, if X is a set of $n-1$ edges, then $\det(N(X))$ is ± 1 if X is the edge set of a spanning tree, and is 0 otherwise. (The proof of this is also an exercise.)

By the Cauchy–Binet formula, $\det(NN^\top)$ is equal to the number of spanning trees. But $NN^\top$ is the principal cofactor of $L(G)$ obtained by deleting the row and column indexed by i.

Matrices with row and column sum zero

To finish the proof, let A be any matrix with row and column sums zero, and let $B = A + J$, where J is the all-1 matrix. We evaluate $\det(B)$.

Matrices with row and column sum zero

To finish the proof, let A be any matrix with row and column sums zero, and let $B = A + J$, where J is the all-1 matrix. We evaluate $\det(B)$.

- Replace the first row by the sum of all the rows; this makes the entries in the first row n and doesn't change the other entries; the determinant is unchanged.

Matrices with row and column sum zero

To finish the proof, let A be any matrix with row and column sums zero, and let $B = A + J$, where J is the all-1 matrix. We evaluate $\det(B)$.

- Replace the first row by the sum of all the rows; this makes the entries in the first row n and doesn't change the other entries; the determinant is unchanged.
- Replace the first column by the sum of all the columns. This makes the first entry n^2, and the other entries in this column n, and doesn't change the other entries of the matrix; the determinant is unchanged.

Matrices with row and column sum zero

To finish the proof, let A be any matrix with row and column sums zero, and let $B = A + J$, where J is the all-1 matrix. We evaluate $\det(B)$.

- Replace the first row by the sum of all the rows; this makes the entries in the first row n and doesn't change the other entries; the determinant is unchanged.
- Replace the first column by the sum of all the columns. This makes the first entry n^2, and the other entries in this column n, and doesn't change the other entries of the matrix; the determinant is unchanged.
- Subtract $1/n$ times the first row from each other row. The elements of the first column, other than the first, become 0; we subtract 1 from all elements not in the first row or column of B, leaving the entries of A; and the determinant is unchanged.

We conclude $\det(B)$ is n^2 times the $(1,1)$ cofactor of A.

We conclude $\det(B)$ is n^2 times the $(1,1)$ cofactor of A.
It is easily checked that the argument works for any cofactor of A. So all cofactors of A are equal.

We conclude $\det(B)$ is n^2 times the $(1,1)$ cofactor of A.
It is easily checked that the argument works for any cofactor of A. So all cofactors of A are equal.
Finally, the all-1 vector is an eigenvector of B with eigenvalue n, while its other eigenvalues are the same as those of A. Thus $\det(B)$ is n times the product of the nontrivial eigenvalues of A.

Cayley's formula

The Matrix-Tree theorem gives us a simple proof of the famous formula of Cayley:

Theorem

The number of spanning trees in the complete graph on n vertices is equal to n^{n-2}.

For the Laplacian of the complete graph is $nI - J$, where J is the all-1 matrix; its non-trivial Laplacian eigenvalues are all equal to n, and so the number of spanning trees is $n^{n-1}/n = n^{n-2}$.

Cayley's formula

The Matrix-Tree theorem gives us a simple proof of the famous formula of Cayley:

Theorem

The number of spanning trees in the complete graph on n vertices is equal to n^{n-2}.

For the Laplacian of the complete graph is $nI - J$, where J is the all-1 matrix; its non-trivial Laplacian eigenvalues are all equal to n, and so the number of spanning trees is $n^{n-1}/n = n^{n-2}$.

In our two examples, the number of spanning trees are 2000 and 576 respectively. (Exercise: prove this using our earlier calculations.)

The Jacobian group

The Jacobian group $\mathrm{Jac}(G)$ of a graph G on n vertices (aka Picard group, critical group, or sandpile group) is defined to be the quotient $\mathbb{Z}^n\,/\,\mathrm{rowspace}(L(G))$ of $\mathbb{Z}^n$ by $\mathrm{rowspace}(L(G))$, the space spanned by the rows of G.

The Jacobian group

The Jacobian group $\mathrm{Jac}(G)$ of a graph G on n vertices (aka Picard group, critical group, or sandpile group) is defined to be the quotient $\mathbb{Z}^n / \mathrm{rowspace}(L(G))$ of $\mathbb{Z}^n$ by $\mathrm{rowspace}(L(G))$, the space spanned by the rows of G.

$\mathrm{Jac}(G)$ is a finitely generated abelian group. If G is connected, it has one factor which is the infinite cyclic group $\mathbb{Z}$; so it has the form $\mathbb{Z} \oplus A$, where A is a finite abelian group whose order is $T(G)$, the number of spanning trees of G.

The Jacobian group

The Jacobian group $\mathrm{Jac}(G)$ of a graph G on n vertices (aka Picard group, critical group, or sandpile group) is defined to be the quotient $\mathbb{Z}^n / \operatorname{rowspace}(L(G))$ of $\mathbb{Z}^n$ by $\operatorname{rowspace}(L(G))$, the space spanned by the rows of G.

$\mathrm{Jac}(G)$ is a finitely generated abelian group. If G is connected, it has one factor which is the infinite cyclic group $\mathbb{Z}$; so it has the form $\mathbb{Z} \oplus A$, where A is a finite abelian group whose order is $T(G)$, the number of spanning trees of G.

For example, $\mathrm{Jac}(K_n)$ is the direct sum of $n-2$ copies of the cyclic group $\mathbb{Z}/(n)$ of order n, refining Cayley's theorem.

Thus $\mathrm{Jac}(G)$ is a more refined invariant than $T(G)$, since different graphs may have the same number of spanning trees but different structures of the group A (as direct sum of cyclic groups).

Thus Jac(G) is a more refined invariant than $T(G)$, since different graphs may have the same number of spanning trees but different structures of the group A (as direct sum of cyclic groups).
In his invited lecture, Alexander Mednykh explained this material, and gave details of the calculation of the structure of Jac(G) for a number of graphs G, including circulants and cyclic covers with finite valency.

Thus $\mathrm{Jac}(G)$ is a more refined invariant than $T(G)$, since different graphs may have the same number of spanning trees but different structures of the group A (as direct sum of cyclic groups).
In his invited lecture, Alexander Mednykh explained this material, and gave details of the calculation of the structure of $\mathrm{Jac}(G)$ for a number of graphs G, including circulants and cyclic covers with finite valency.
However, we do not know an application of the Jacobian group in the theory of optimal design. Perhaps one of our readers can find one …

Markov chains

A Markov chain on a finite state space S is a sequence of random variables with values in S which has no memory: the state at time $n + 1$ depends only on the state at time n.

Markov chains

A Markov chain on a finite state space S is a sequence of random variables with values in S which has no memory: the state at time $n+1$ depends only on the state at time n.
A Markov chain is defined by a transition matrix P, with rows and columns indexed by S, where p_{ij} is the probability of moving from state i to state j in one time step.

Markov chains

A Markov chain on a finite state space S is a sequence of random variables with values in S which has no memory: the state at time $n+1$ depends only on the state at time n.
A Markov chain is defined by a transition matrix P, with rows and columns indexed by S, where p_{ij} is the probability of moving from state i to state j in one time step.
As usual, the entries of P are non-negative and the row sums are equal to 1.

Random walks

An important example of a Markov chain is the random walk on a graph G. The state space is the vertex set $V(G)$. At time n, if the process is at vertex i, it chooses at random (with equal probabilities) an edge containing i, and at stage $n+1$ moves to the other end of this edge.

Random walks

An important example of a Markov chain is the random walk on a graph G. The state space is the vertex set $V(G)$. At time n, if the process is at vertex i, it chooses at random (with equal probabilities) an edge containing i, and at stage $n+1$ moves to the other end of this edge.

If the graph has no loops, then the probability of moving from i to j is $-L_{ij}/L_{ii}$, where L is the Laplacian. In particular, if the graph is regular with degree d, then $P = I - L/d$.

Random walks

An important example of a Markov chain is the random walk on a graph G. The state space is the vertex set $V(G)$. At time n, if the process is at vertex i, it chooses at random (with equal probabilities) an edge containing i, and at stage $n+1$ moves to the other end of this edge.

If the graph has no loops, then the probability of moving from i to j is $-L_{ij}/L_{ii}$, where L is the Laplacian. In particular, if the graph is regular with degree d, then $P = I - L/d$.

More generally, $P = I - D^{-1}L$, where D is the diagonal matrix whose (i, i) entry is the number of edges incident with i.

Theory of Markov chains

If a Markov chain has transition matrix P, then the (i, j) entry of P^m is the probability of moving from i to j in m steps.

Theory of Markov chains

If a Markov chain has transition matrix P, then the (i, j) entry of P^m is the probability of moving from i to j in m steps.
The Markov chain is connected if, for any i and j, there exists m such that $(P^m)_{ij} \neq 0$; it is aperiodic if the the greatest common divisors of the values of m for which $(P^m)_{ii} \neq 0$ for some i is 1.

Theory of Markov chains

If a Markov chain has transition matrix P, then the (i, j) entry of P^m is the probability of moving from i to j in m steps.
The Markov chain is connected if, for any i and j, there exists m such that $(P^m)_{ij} \neq 0$; it is aperiodic if the the greatest common divisors of the values of m for which $(P^m)_{ii} \neq 0$ for some i is 1.
A random walk on a graph G is connected if and only if G is connected, and is aperiodic if and only if G is not bipartite.

Theorem

A connected aperiodic Markov chain has a unique limiting distribution, to which it converges from any starting distribution.

Theorem

A connected aperiodic Markov chain has a unique limiting distribution, to which it converges from any starting distribution.

Since the row sums of P are all 1, we see that $Pp^\top = p^\top$, where p is the all-1 vector; our assumptions imply that the multiplicity of 1 as eigenvalue is 1. Now left and right eigenvalues are equal, so there is a vector $q \neq 0$ such that $qP = q$. It can be shown that the entries of q are non-negative; we can normalise it so that their sum is 1. Then q is a probability distribution which is fixed by P, so it is the unique stationary distribution.

Convergence

Suppose that P is symmetric. Then we can write $P = \sum \lambda P_\lambda$ where λ runs over the eigenvalues, and P_λ is the projection onto the λ eigenspace. Then $P^m = \sum \lambda^m P_\lambda$.

Convergence

Suppose that P is symmetric. Then we can write $P = \sum \lambda P_\lambda$ where λ runs over the eigenvalues, and P_λ is the projection onto the λ eigenspace. Then $P^m = \sum \lambda^m P_\lambda$.
It is also true, by the Perron–Frobenius theorem, that every eigenvalue λ satisfies $|\lambda| \leq 1$. If the Markov chain is irreducible and aperiodic, then 1 is a simple eigenvalue, and all other eigenvalues have modulus strictly less than 1.

Convergence

Suppose that P is symmetric. Then we can write $P = \sum \lambda P_\lambda$ where λ runs over the eigenvalues, and P_λ is the projection onto the λ eigenspace. Then $P^m = \sum \lambda^m P_\lambda$.

It is also true, by the Perron–Frobenius theorem, that every eigenvalue λ satisfies $|\lambda| \leq 1$. If the Markov chain is irreducible and aperiodic, then 1 is a simple eigenvalue, and all other eigenvalues have modulus strictly less than 1.

Now let x be any non-negative vector whose coordinates sum to 1. We can regard x as the initial probability distribution. Then we have

$$xP^m = \sum \lambda^m x P_\lambda \to xP_1$$

as $m \to \infty$. So $xP_1 = q$ is the limiting distribution, and the convergence to q is like μ^m where μ is the second-largest modulus of an eigenvalue. So the convergence is exponential if μ is not close to 1.

Random walks revisited

For a random walk, we have $P = I - D^{-1}L$. Then

$$D^{1/2}PD^{-1/2} = I - D^{-1/2}LD^{-1/2}.$$

This matrix is symmetric, and is similar to P; so P is indeed diagonalizable. However, the analysis is a bit more complicated, and not given here.

Random walks revisited

For a random walk, we have $P = I - D^{-1}L$. Then

$$D^{1/2}PD^{-1/2} = I - D^{-1/2}LD^{-1/2}.$$

This matrix is symmetric, and is similar to P; so P is indeed diagonalizable. However, the analysis is a bit more complicated, and not given here.
Its eigenvalues are $1 - \lambda$, where λ is an eigenvalue of the positive semidefinite matrix $D^{-1/2}LD^{-1/2}$, so for rapid convergence we require that the smallest positive eigenvalue of this matrix should be as large as possible.

Random walks revisited

For a random walk, we have $P = I - D^{-1}L$. Then

$$D^{1/2}PD^{-1/2} = I - D^{-1/2}LD^{-1/2}.$$

This matrix is symmetric, and is similar to P; so P is indeed diagonalizable. However, the analysis is a bit more complicated, and not given here.
Its eigenvalues are $1 - \lambda$, where λ is an eigenvalue of the positive semidefinite matrix $D^{-1/2}LD^{-1/2}$, so for rapid convergence we require that the smallest positive eigenvalue of this matrix should be as large as possible.
Thus the problem is a twisted version of the usual problem about the smallest non-trivial Laplacian eigenvalue. If the graph is regular, so that $D = dI$, then it reduces exactly to the former problem.

Other results

The smallest nontrivial Laplacian eigenvalue μ of a graph G is an important parameter which occurs in many other situations.

Other results

The smallest nontrivial Laplacian eigenvalue μ of a graph G is an important parameter which occurs in many other situations. For example, a recent result of Krivelevich and Sudakov asserts that, in a regular graph of valency d on n vertices, if μ is sufficiently large in terms of n and d, then the graph is Hamiltonian.

Summing up

We saw that three important parameters of a connected graph, which are determined by its Laplacian spectrum, are:

- the harmonic mean of the non-trivial Laplacian eigenvalues, which tells us about the average resistance between pairs of vertices;

Summing up

We saw that three important parameters of a connected graph, which are determined by its Laplacian spectrum, are:

- the harmonic mean of the non-trivial Laplacian eigenvalues, which tells us about the average resistance between pairs of vertices;
- the geometric mean of the non-trivial Laplacian eigenvalues, which tells us about the number of spanning trees;

Summing up

We saw that three important parameters of a connected graph, which are determined by its Laplacian spectrum, are:

- the harmonic mean of the non-trivial Laplacian eigenvalues, which tells us about the average resistance between pairs of vertices;
- the geometric mean of the non-trivial Laplacian eigenvalues, which tells us about the number of spanning trees;
- the smallest non-trivial Laplacian eigenvalue, which is related to the isoperimetric number and the rate of convergence of the random walk on the graph.

Summing up

We saw that three important parameters of a connected graph, which are determined by its Laplacian spectrum, are:

- the harmonic mean of the non-trivial Laplacian eigenvalues, which tells us about the average resistance between pairs of vertices;
- the geometric mean of the non-trivial Laplacian eigenvalues, which tells us about the number of spanning trees;
- the smallest non-trivial Laplacian eigenvalue, which is related to the isoperimetric number and the rate of convergence of the random walk on the graph.

In the next lecture, we will see that these are also important parameters in experimental design!