Spring Embeddings of Planar Graphs

John Urschel

Department of Mathematics, Massachusetts Institute of Technology

joint with Ludmil Zikatanov

July 10, 2020

Tutte's Spring Embedding Theorem

Given a three-connected planar graph G, if a face of G is embedded in $\mathbb{R}^2$ as a convex polygon and every other vertex is placed at the mass center of its neighbors, then this embedding is planar and uniquely determined by the embedding of the face.

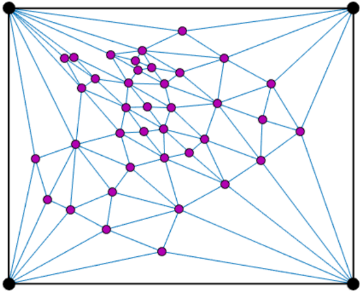

Energy Minimization in Spring Embeddings

Let Γ, $n_\Gamma := |\Gamma|$, be the vertices of the "outer" face of G. Given an embedding $X_\Gamma \in \mathbb{R}^{n_\Gamma \times 2}$ of the outer face, the Tutte condition

$$X_{i,\cdot} = \frac{1}{d(i)} \sum_{j \in N(i)} X_{j,\cdot} \qquad i \in V(G) \setminus \Gamma$$

also minimizes the sum of squared edge lengths, conditional on the embdding X_Γ of Γ.

Graph Laplacian

The graph Laplacian $L_G \in \mathbb{R}^{n \times n}$ of G is the symmetric matrix defined by

$$\langle L_G x, x \rangle = \sum_{\{i,j\} \in E} (x_i - x_j)^2,$$

and, in general, a matrix is the graph Laplacian of some weighted graph if it is symmetric diagonally dominant, has non-positive off-diagonal entries, and the vector $\mathbf{1} := (1, ..., 1)^T$ lies in its nullspace.

Block Notation

Let $V = \{1, ..., n\}$ and $\Gamma = \{n - n_\Gamma + 1, ..., n\}$. We can write both the Laplacian and embedding of G in block-notation, differentiating between interior and boundary vertices as follows:

$$L_G = \begin{pmatrix} L_o + D_o & -A_{o,\Gamma} \\ -A_{o,\Gamma}^T & L_\Gamma + D_\Gamma \end{pmatrix} \in \mathbb{R}^{n \times n}, \quad X = \begin{pmatrix} X_o \\ X_\Gamma \end{pmatrix} \in \mathbb{R}^{n \times 2},$$

where L_o and L_Γ are the Laplacians of $G[V \backslash \Gamma]$ and $G[\Gamma]$, respectively.

Energy Minimization in Spring Embeddings

Using block notation, the system of equations for the Tutte spring embedding of some convex embedding X_Γ is given by

$$X_o = (D_o + D[L_o])^{-1}[(D[L_o] - L_o)X_o + A_{o,\Gamma}X_\Gamma],$$

where $D[A]$ is the diagonal matrix with diagonal entries given by the diagonal of A. The unique solution to this system is

$$X_o = (L_o + D_o)^{-1}A_{o,\Gamma}X_\Gamma.$$

X_o not only guarantees a planar embedding of G, but also minimizes Hall's energy, namely,

$$\arg\min_{X_o} h(X) = (L_o + D_o)^{-1}A_{o,\Gamma}X_\Gamma,$$

where $h(X) := \mathrm{Tr}(X^T L X)$.

Spring Embeddings vs Spectral Layouts

Embedding a graph using the two minimal non-trivial eigenvectors of the graph Laplacian minimizes the sum of squared edge lengths ($\mathrm{Tr}(X^T L X)$) subject to $X^T X = I$ and $X^T 1 = 0$.

The problem is that, for planar graphs, the resulting embedding can often be far from planar.

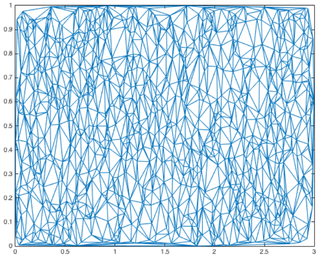

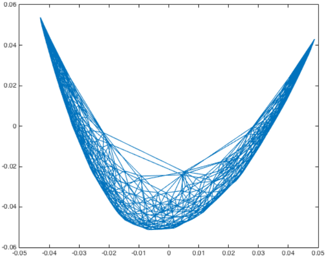

Spring Embeddings vs Spectral Layouts

Spring embeddings give the best of both worlds. It produces a planar embedding while minimizing the sum of squared edges, conditional on the boundary embedding.

But how to embed the boundary?

A Natural Choice

A natural choice is to embed the outer face using the restriction of two minimal non-trivial eigenvectors of the graph Laplacian.

Unfortunately, the restriction to the boundary is also often non-convex.

Another choice is the two minimal non-trivial eigenvectors of the graph Laplacian of the boundary L_Γ.

Unfortunately, this results in a regular n_Γ-gon, and fails to take into account the dynamics of the interior.

A Minimization Problem

Let $\mathcal{X}$ be the set of all convex, planar embeddings X_Γ that satisfy $X_\Gamma^T X_\Gamma = I$ and $X_\Gamma^T \mathbf{1} = 0$. Consider the optimization problem

$$\min \ \mathrm{Tr}(X^T L X) \quad s.t. \quad X_\Gamma \in \mathrm{cl}(\mathcal{X}), \tag{0.1}$$

where $\mathrm{cl}(\cdot)$ is the closure of a set.

The normalizations $X_\Gamma^T \mathbf{1} = 0$ and $X_\Gamma^T X_\Gamma = I$ ensure that the solution does not degenerate into a single point or line.

A Minimization Problem

Given some choice of X_Γ, by Tutte's theorem the minimum value of $\mathrm{Tr}(X^T L X)$ is attained when $X_o = (L_o + D_o)^{-1} A_{o,\Gamma} X_\Gamma$, and given by

$$\begin{aligned}
\mathrm{Tr}(X^T L X) &= \mathrm{Tr}\Bigg[\left([(L_o + D_o)^{-1} A_{o,\Gamma} X_\Gamma]^T \quad X_\Gamma^T \right) \begin{pmatrix} L_o + D_o & -A_{o,\Gamma} \\ -A_{o,\Gamma}^T & L_\Gamma + D_\Gamma \end{pmatrix} \\
&\qquad \begin{pmatrix} (L_o + D_o)^{-1} A_{o,\Gamma} X_\Gamma \\ X_\Gamma \end{pmatrix} \Bigg] \\
&= \mathrm{Tr}\left(X_\Gamma^T \left[L_\Gamma + D_\Gamma - A_{o,\Gamma}^T (L_o + D_o)^{-1} A_{o,\Gamma} \right] X_\Gamma \right) \\
&= \mathrm{Tr}\left(X_\Gamma^T S_\Gamma X_\Gamma \right),
\end{aligned}$$

where S_Γ is the Schur complement of L_G with respect to $V \backslash \Gamma$,

$$S_\Gamma = L_\Gamma + D_\Gamma - A_{o,\Gamma}^T (L_o + D_o)^{-1} A_{o,\Gamma}.$$

A Minimization Problem, Reparameterized

Given that

$$\min_{X_o} \mathrm{Tr}(X^T L X) = \mathrm{Tr}\left(X_\Gamma^T S_\Gamma X_\Gamma\right),$$

we can treat X_o as a function of X_Γ and instead consider the optimization problem

$$\min \ \mathrm{Tr}\big(X_\Gamma^T S_\Gamma X_\Gamma\big) \quad s.t. \quad X_\Gamma \in \mathrm{cl}(\mathcal{X}). \tag{0.2}$$

This immediately implies that, if the minimal two non-trivial eigenvectors of S_Γ produce a convex embedding, then this is the exact solution of (0.2).

Do we have any reason to think that this embedding would be planar or convex?

The Schur Complement

Spectral Layout

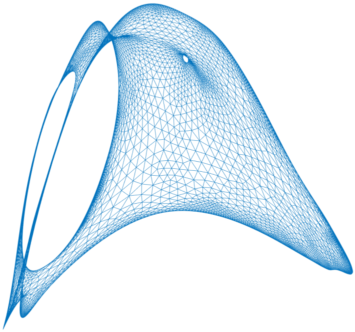

Schur Complement Spring Embedding

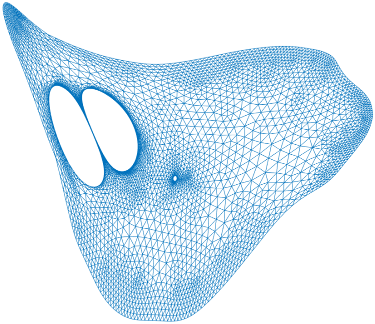

The Schur Complement

Proposition

Let $G = (V, E)$, $n = |V|$, be a graph and $L_G \in \mathbb{R}^{n \times n}$ the associated graph Laplacian. Let L_G and vectors $v \in \mathbb{R}^n$ be written in block form

$$L(G) = \begin{pmatrix} L_{11} & L_{12} \\ L_{21} & L_{22} \end{pmatrix}, \quad v = \begin{pmatrix} v_1 \\ v_2 \end{pmatrix},$$

where $L_{22} \in \mathbb{R}^{m \times m}$, $v_2 \in \mathbb{R}^m$, and $L_{12} \neq 0$. Then

(1) $S = L_{22} - L_{21} L_{11}^{-1} L_{12}$ *is a graph Laplacian,*

(2) $\sum_{i=1}^{m} (e_i^T L_{22} \mathbf{1}_m) e_i e_i^T - L_{21} L_{11}^{-1} L_{12}$ *is a graph Laplacian,*

(3) $\langle Sw, w \rangle = \inf\{\langle Lv, v \rangle | v_2 = w\}$.

Schur Complement

Based on the previous proposition, we see the semi-norm for the Schur complement corresponds to the energy semi-norm of the harmonic extension of the boundary vertices.

The Schur complement is a Laplacian and is the sum of the boundary Laplacian L_Γ and another Laplacian $D_\Gamma - A_{o,\Gamma}^T (L_o + D_o)^{-1} A_{o,\Gamma}$ representing the dynamics of the interior.

For graphs with a certain amount of structure (i.e., triangulations of convex regions, etc.), we can mathematically quantify the behavior of S_Γ, and show that S_Γ is spectrally equivalent to $L_\Gamma^{1/2}$.

Trace Theorems for Lipschitz Domains from PDE Theory

Let Ω be a domain in $\mathbb{R}^d$ with boundary which is locally a graph of a Lipschitz function. We denote by $H^1(\Omega)$ the Sobolev space of square integrable functions with square integrable gradient. The norm on this space is defined as

$$\|u\|_{1,\Omega}^2 = \|\nabla u\|_{L^2(\Omega)}^2 + \|u\|_{L^2(\Omega)}^2, \quad \text{where} \quad \|u\|_{L_2(\Omega)}^2 = \int_\Omega u^2 dx,$$

We introduce the following norm for functions defined on Γ:

$$\|\varphi\|_{1/2,\Gamma}^2 = \|\varphi\|_{L_2(\Gamma)}^2 + \iint_{\Gamma\times\Gamma} \frac{(\varphi(x)-\varphi(y))^2}{|x-y|^d} dxdy.$$

and we denote by $H^{1/2}(\Gamma)$, $\Gamma = \partial\Omega$ the Sobolev space of functions defined on the boundary Γ for which the $\|\cdot\|_{1/2,\Gamma}$ is finite.

Trace Theorems for Lipschitz Domains from PDE Theory

Theorem (Trace Theorem)

Let $\Omega \subset \mathbb{R}^d$ be a Lipschitz domain and $u \in H^1(\Omega)$ be a given function. Then

(i) *The trace operator $\gamma : H^1(\Omega) \mapsto H^{1/2}(\Gamma)$, $\gamma u = u\big|_\Gamma$ is bounded, namely there exists a constant c such that*

$$\|\gamma u\|_{1/2,\Gamma} \leq c\|u\|_{1,\Omega}$$

(ii) *The trace operator has a continuous right inverse, namely, for any $\varphi \in H^{1/2}(\Gamma)$ there exists $u \in H^1(\Omega)$ such that*

$$\gamma u = \varphi, \quad \text{and} \quad \|u\|_{1,\Omega} \leq c_2\|\varphi\|_{1/2,\Omega}.$$

(iii) *The following norm equivalence holds:*

$$\|\varphi\|_{1/2,\Omega} \eqsim \inf\{\|u\|_{1,\Omega} \mid \gamma u = \varphi\}.$$

Spectral Equivalence

If we can prove a discrete version of a trace theorem for energy semi-norms, i.e. show that, for any $\varphi \in \mathbb{R}^{n_\Gamma}$,

$$\frac{1}{C_1} \langle \widetilde{L}_\Gamma \varphi, \varphi \rangle \leq \langle S_\Gamma \varphi, \varphi \rangle \leq C_2 \langle \widetilde{L}_\Gamma \varphi, \varphi \rangle,$$

where

$$\langle \widetilde{L}_\Gamma \varphi, \varphi \rangle = \sum_{\substack{p,q \in \Gamma, \\ p<q}} \frac{(\varphi(p) - \varphi(q))^2}{d_\Gamma^2(p,q)}$$

and C_1 and C_2 are constants that do not depend on n, then to show our desired result, we just have to show $L_\Gamma^{1/2}$ is spectrally equivalent to $\widetilde{L}_\Gamma$.

Spectral Equivalence

Theorem

If there exists a planar spring embedding X of $(G,\Gamma) \in \mathcal{G}_n^{f \leq c_1}$ for which

(1) $K = conv\left(\{[X_\Gamma]_{i,\cdot}\}_{i=1}^{n_\Gamma}\right)$ *satisfies*

$$\sup_{u \in K} \inf_{v \in \partial K} \sup_{w \in \partial K} \frac{\|u - v\|}{\|u - w\|} \geq c_2 > 0,$$

(2) X *satisfies*

$$\max_{\substack{\{i_1,i_2\} \in E \\ \{j_1,j_2\} \in E}} \frac{\|X_{i_1,\cdot} - X_{i_2,\cdot}\|}{\|X_{j_1,\cdot} - X_{j_2,\cdot}\|} \leq c_3 \quad \text{and} \quad \min_{\substack{i \in V \\ j_1,j_2 \in N(i)}} \angle X_{j_1,\cdot}\, X_{i,\cdot}\, X_{j_2,\cdot} \geq c_4 > 0,$$

then, for any $\varphi \in \mathbb{R}^{n_\Gamma}$,

$$\frac{1}{C_1} \langle L_\Gamma^{1/2} \varphi, \varphi \rangle \leq \langle S_\Gamma \varphi, \varphi \rangle \leq C_2 \langle L_\Gamma^{1/2} \varphi, \varphi \rangle,$$

where C_1 *and* C_2 *are constants that depend on* c_1, c_2, c_3, *and* c_4.

(Part of) the Proof

Let us define

$$\ell(i,j) = \min\{i - j \mod n,\ \ j - i \mod n\}.$$

By definition, Γ induces a cycle, so $\tilde{L}(i,j) = -\ell(i,j)^{-2}$ for $i \neq j$.

Because L_Γ is the cycle graph, its spectral decomposition is well known,

$$L_\Gamma = \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor} \lambda_k(L_\Gamma)\left(\frac{x_k x_k^T}{\|x_k\|^2} + \frac{y_k y_k^T}{\|y_k\|^2}\right),$$

where $\lambda_k = 2 - 2\cos\left(\frac{2\pi k}{n}\right)$ and $x_k = \sin\left(\frac{2\pi ki}{n}\right)$, $y_k = \cos\left(\frac{2\pi ki}{n}\right)$, $i = 1, \ldots, n$. We can use the spectral decomposition to write $L_\Gamma^{1/2}$ explicitly.

(Part of) the Proof

Namely, (suppose wlog that n is odd)

$$L_\Gamma = \sum_{k=1}^{\frac{n-1}{2}} \lambda_k(L_\Gamma)\left(\frac{x_k x_k^T}{\|x_k\|^2} + \frac{y_k y_k^T}{\|y_k\|^2}\right)$$

where $\lambda_k = 2 - 2\cos\left(\frac{2\pi k}{n}\right)$ and $x_k = \sin\left(\frac{2\pi k i}{n}\right)$, $y_k = \cos\left(\frac{2\pi k i}{n}\right)$, $i = 1, ..., n$. We first compute the value of $\|x_k\|^2$ and $\|y_k\|^2$. We have

$$\begin{aligned}\|x_k\|^2 &= \sum_{i=1}^{n} \sin^2(\frac{2\pi k}{n} i) = \frac{n}{2} - \frac{1}{2}\sum_{i=1}^{n} \cos(\frac{4\pi k}{n} i) \\ &= \frac{n}{2} - \frac{1}{4}(\frac{\sin(2\pi k(2+\frac{1}{n}))}{\sin(\frac{2\pi k}{n})} - 1) = \frac{n}{2},\end{aligned}$$

which, from $\|x_k\|^2 + \|y_k\|^2 = n$, implies that $\|y_k\|^2 = \frac{n}{2}$ as well.

(Part of) the Proof

This gives

$$\begin{aligned} L_\Gamma^{1/2}(i,j) &= \frac{2\sqrt{2}}{n}\sum_{k=1}^{\lfloor \frac{n}{2}\rfloor}\left[1-\cos\left(\frac{2k\pi}{n}\right)\right]^{1/2} \\ &\times \left[\sin\left(\frac{2\pi ki}{n}\right)\sin\left(\frac{2\pi kj}{n}\right)-\cos\left(\frac{2\pi ki}{n}\right)\cos\left(\frac{2\pi kj}{n}\right)\right] \\ &= \frac{2\sqrt{2}}{n}\sum_{k=1}^{\lfloor \frac{n}{2}\rfloor}\left[1-\cos\left(\pi\frac{2k}{n}\right)\right]^{1/2}\cos\left(\ell(i,j)\pi\frac{2k}{n}\right). \end{aligned}$$

By the properties of a Riemann sum, we have

$$\begin{aligned} \left|\frac{L_\Gamma^{1/2}(i,j)}{\sqrt{2}}-\int_0^1(1-\cos(\pi x))^{1/2}\cos(\ell\pi x)\right| &= \left|\frac{L_\Gamma^{1/2}(i,j)}{\sqrt{2}}-\frac{-2\sqrt{2}}{\pi(4\ell^2-1)}\right| \\ &\leq \frac{C}{n} \end{aligned}$$

for some C. This immediately implies that $L_\Gamma^{1/2}$ is spectrally equivalent to the Laplacian $\hat{L}$ with $\hat{L}(i,j)=(4\ell(i,j)^2-1)^{-1}, i\neq j$.

From Analysis to Algorithm

The above analysis inspires a simple meta-algorithm.

Compute the minimal two eigenpairs of S_Γ. If the resulting embedding of Γ is planar and convex, we are done.

Otherwise, use an initial guess (a convex version of the above embedding, or a regular polygon, etc.) and iteratively smooth the approximation using S_Γ. The spectral equivalence implies an initial guess with objective function within a factor of $C_1 C_2$ of the optimum.

An Obligatory Message about S_Γ and S_Γ^{-1}

The Schur complement S_Γ is a dense matrix and requires the inversion of a large matrix, but can be represented as the composition of functions of sparse matrices.

In practice, S_Γ should NEVER be formed explicitly.

The operation of applying S_Γ to a vector x should occur in two steps:
1) Solve $(L_o + D_o)y = A_{o,\Gamma}x$ should be solved for y.
2) Compute $S_\Gamma x = (L_\Gamma + D_\Gamma)x - A_{o,\Gamma}^T y$.

Each application of S_Γ is therefore an $O(n \log n)$ procedure (using an $O(n \log n)$ Laplacian solver).

An Obligatory Message about S_Γ and S_Γ^{-1}

The application of the inverse S_Γ^{-1} defined on the subspace $\{x \,|\, \langle x, \mathbf{1} \rangle = 0\}$ also requires the solution of a Laplacian system.

The action of S_Γ^{-1} on a vector $x \in \{x \,|\, \langle x, \mathbf{1} \rangle = 0\}$ is given by

$$S_\Gamma^{-1} x = \begin{pmatrix} 0 & I \end{pmatrix} \begin{pmatrix} L_o + D_o & -A_{o,\Gamma} \\ -A_{o,\Gamma}^T & L_\Gamma + D_\Gamma \end{pmatrix}^{-1} \begin{pmatrix} 0 \\ x \end{pmatrix}.$$

Given that the application of S_Γ^{-1} has the same complexity as an application S_Γ, the inverse power method is naturally preferred over the shifted power method for smoothing.

Examples

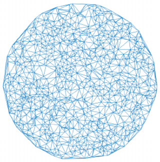

(A) Circle

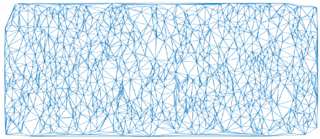

(B) 3-by-1 Rectangle

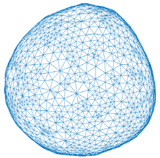

(C) Spectral Layout

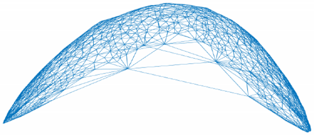

(D) Spectral Layout

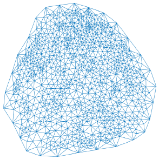

(E) Schur Complement Layout

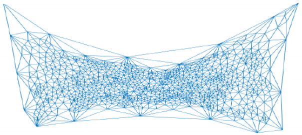

(F) Schur Complement Layout

Thank You!

Thank you to MGGG Redistricting Lab and Moon Duchin for having me.

The associated preprint is

John C. Urschel, Ludmil T. Zikatanov. **Discrete Trace Theorems and Energy Minimizing Spring Embeddings of Planar Graphs**.

and can be found on arXiv.

If you have any further questions, feel free to send me an email at urschel 'at' mit 'dot' edu.