Università della Svizzera italiana

Faculty of Informatics

Institute of Computational Science ICS

Load-Balanced Partition Refinement with the Graph p-Laplacian

Toby Simpson & Olaf Schenk,
Institute of Computational Science, Università della Svizzera italiana, Lugano, Switzerland

Spectral Graph Partitioning

For an unweighted graph $G(V, E)$ with n vertices and m edges, the degree D, adjacency W and incidence A matrices lead to the graph Laplacian operator L.

$$(D - W)\mathbf{x} = A^T A\mathbf{x} = L\mathbf{x}$$

An edge separator for a vertex subset $V_k \cup \overline{V_k} = V$ has size:

$$\text{cut}(V_k, \overline{V_k}) = \sum_{i \in V_k, j \in \overline{V_k}} w_{ij}$$

The Ratio Cut balances the partition by cardinality. For an indicator vector $\mathbf{x} \in \mathbb{R}^n$ of V_k, it is approximated by the normalised edge gradients $A\mathbf{x} \in \mathbb{R}^m$.

$$\text{RCut}(V_k, \overline{V_k}) = \frac{\text{cut}(V_k, \overline{V_k})}{|V_k||\overline{V_k}|} \approx \frac{1}{2}\frac{\|A\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2}$$

The minimisation problem is the Rayleigh quotient of the symmetric graph Laplacian matrix and local solutions are eigenpairs $(\lambda_i, \mathbf{v}_i)$.

$$\arg\min_{\mathbf{x} \in \mathbb{R}^n} \frac{\mathbf{x}^T L \mathbf{x}}{\mathbf{x}^T \mathbf{x}} = \lambda_i$$

The first eigenvector is constant $\mathbf{v}_1 = c\mathbf{1}$ and represents the trivial solution $V = V \cup \emptyset$. We exclude it by orthogonality and define the Spectral Graph Partitioning [1] problem as follows.

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} \quad \text{subject to} \quad \mathbf{1}^T\mathbf{x} = 0$$

The global solution $\mathbf{v}_2$ is Fiedler's famous eigenvector. The sign change indicates the narrowest region in the eigenspace, the shortest cut.

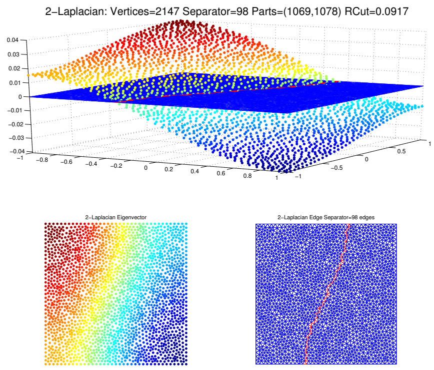

The Graph p-Laplacian

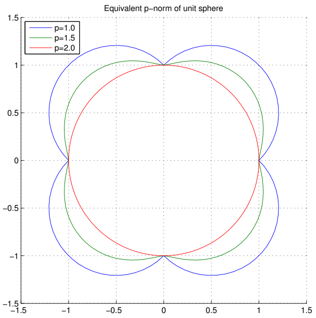

Minimising a vector p-norm $\|\mathbf{x}\|_p^p = \sum_i |x_i|^p$, $p \in (1, 2]$ pushes elements of the solution towards discrete values.

We define the scalar function $\phi_p(x) = |x|^{p-1}\,\text{sign}(x)$. When applied element-wise to a vector $\mathbf{x} \in \mathbb{R}^n$ the inner product returns the p-norm.

$$\mathbf{x}^T\phi_p(\mathbf{x}) = \sum_i |x_i|^p = \|\mathbf{x}\|_p^p$$

The ϕ_p function allows us to re-define the minimisation problem in the p-norm and recover the discreteness lost in the Ratio cut approximation.

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p} = \min_{\mathbf{x} \in \mathbb{R}^n} \frac{(A\mathbf{x})^T\phi_p(A\mathbf{x})}{\mathbf{x}^T\phi_p(\mathbf{x})}$$

This is the Rayleigh quotient for the nonlinear eigenvalue problem that defines the p-Laplacian operator.

$$A^T\phi_p(A\mathbf{x}) = \lambda\phi_p(\mathbf{x})$$

We exclude the trivial solution with a nonlinear constraint and define the p-Laplacian Graph Partitioning [2] problem as follows.

$$\min_{\mathbf{x} \in \mathbb{R}^n} \frac{\|A\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p} \quad \text{subject to} \quad \mathbf{1}^T\phi_p(\mathbf{x}) = 0$$

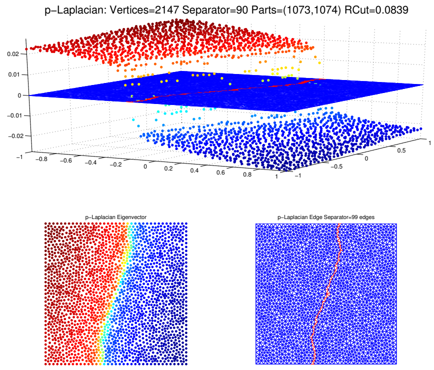

Solving the Eigenvalue Problem

The algorithm was adapted to reduce complexity from $\mathcal{O}(n^2)$ vertices to $\mathcal{O}(m)$ edges for large sparse matrices.

- The nonlinear constraint is handled by feasible projection, using the invertible ϕ_p function.

$$\phi_p^{-1}(x) = |x|^{\frac{1}{p-1}}\,\text{sign}(x)$$

$$\hat{\mathbf{x}} = \phi_p^{-1}\left(\phi_p(\mathbf{x}) - \frac{\mathbf{1}^T\phi_p(\mathbf{x})}{n}\right)$$

- Initialising from a 2-Laplacian eigenvector is too costly to compute. Instead a vector is synthesised using METIS [3] software.
- The algorithm is enclosed in an outer loop that decreases $p \to 1$ improving discreteness and balance during the solve.

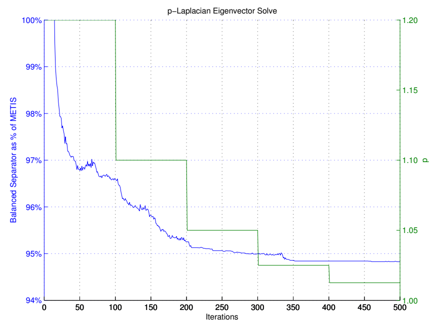

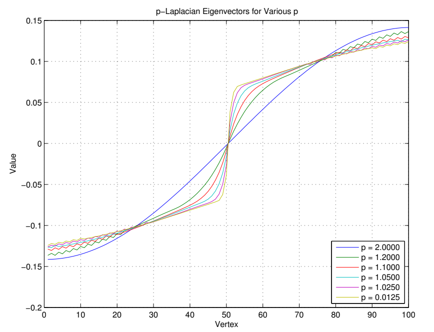

Numerical Experiments

The p-Laplacian refinement was tested on Delaunay triangulations of 3-dimensional random point clouds.

- The algorithm ran over five sets of six graphs, each with 1 to 6 million vertices (15 to 100 million edges) and calculated the balanced cut metric before and after refinement.
- Results show a consistent improvement of around 4% over the METIS cut, increasing with scale. A similar experiment with 2-dimensional meshes yields an improvement of around 8%.
- Times for the separator refinement are linear in edges $\mathcal{O}(m)$ as expected.

When tested on benchmark sparse matrices from the UFL [4] collection, there was a similar improvement of around 4% with more variation due to the different structures encountered.

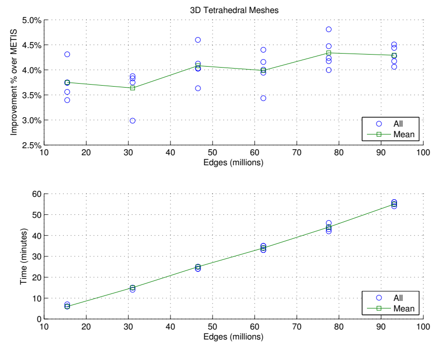

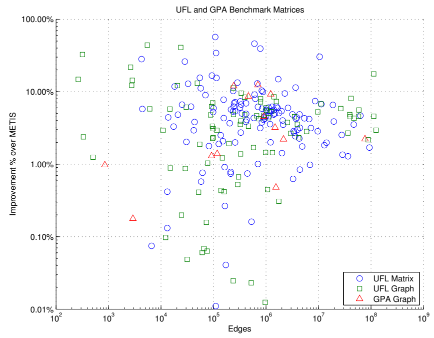

References

[1] Alex Pothen, Horst D. Simon, and Kan-Pu Liou. Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Matrix Anal. Appl.*, 11(3):430–452, May 1990.

[2] Thomas Bühler and Matthias Hein. Spectral clustering based on the graph p-laplacian. In *Proceedings of the 26th Annual International Conference on Machine Learning*, ICML '09, New York, 2009. ACM.

[3] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput.*, 20(1):359–392, December 1998.

[4] Timothy A. Davis and Yifan Hu. The university of florida sparse matrix collection. *ACM Trans. Math. Softw.*, 38(1):1:1–1:25, December 2011.