Balanced Graph Partition Refinement using the Graph p-Laplacian

D. Pasadakis[1], T.Simpson[1], D. Kourounis[1], K. Fujita[2], T. Yamaguchi[2], T. Ichimura[2], O. Schenk[1]

1] Institute of Computational Science,
Università della Svizzera italiana

2] Earthquake Research Institute &
Department of Civil Engineering,
The University of Tokyo

September 6, 2018

Motivation

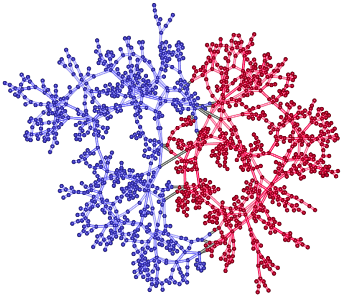

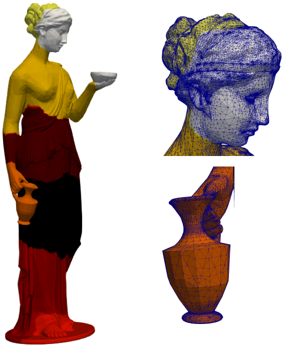

Definitions

For a graph $\mathcal{G}(V, E)$

- Incidence: $\boldsymbol{A} \in \mathbb{R}^{m \times n}$,
- Graph Laplacian: $\boldsymbol{L} \in \mathbb{R}^{n \times n}$.

$$\boldsymbol{A}^\top : \nabla \cdot, \qquad \boldsymbol{A} : \nabla = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 1 & 0 & 1 \\ -1 & 0 & 0 & 1 \end{bmatrix}$$

$$\boldsymbol{L} = \boldsymbol{A}^\top \boldsymbol{A} : \nabla \cdot \nabla = \Delta = \begin{bmatrix} 2 & -1 & 0 & -1 \\ -1 & 3 & -1 & -1 \\ 0 & -1 & 2 & -1 \\ -1 & -1 & -1 & 3 \end{bmatrix}$$

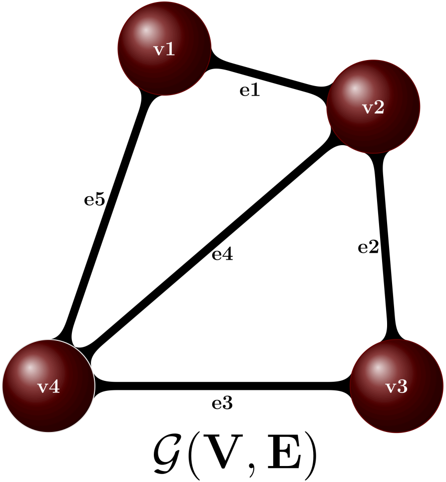

Spectral Graph Partitioning

Calculate an edge separator using the Fiedler eigenvector of $\mathbf{L} \in \mathbb{R}^{n \times n}$.

$$x_i = \begin{cases} 1\,, & i \in V_k \\ 0\,, & i \in \overline{V_k} \end{cases}$$

$$\mathbf{x}^\top \mathbf{x} = \|\mathbf{x}\|_2^2 = |V_k|$$

$\boldsymbol{A}\mathbf{x}$ has nonzero elements only for edges that connect a vertex in V_k with one in $\overline{V}_k$.

$$\text{cut}(V_k, \overline{V}_k) = \|\boldsymbol{A}\mathbf{x}\|_2^2$$

Spectral Graph Partitioning

2-Laplacian partitioning

$$\min_{V_k \subset V} \frac{\|\boldsymbol{A}\hat{\mathbf{x}}\|_2^2}{\|\hat{\mathbf{x}}\|_2^2} \approx \min_{x \in \mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} = \min_{x \in \mathbb{R}^n} \frac{\mathbf{x}^\top \boldsymbol{L}\mathbf{x}}{\mathbf{x}^\top \mathbf{x}} = \lambda_i$$

$(\lambda_1 = 0, \mathbf{v}_1 = c\mathbf{1}) \rightarrow V = V \cup \emptyset$

$(\lambda_2, \boldsymbol{v}_2) \rightarrow$ Fiedler's eigenvector

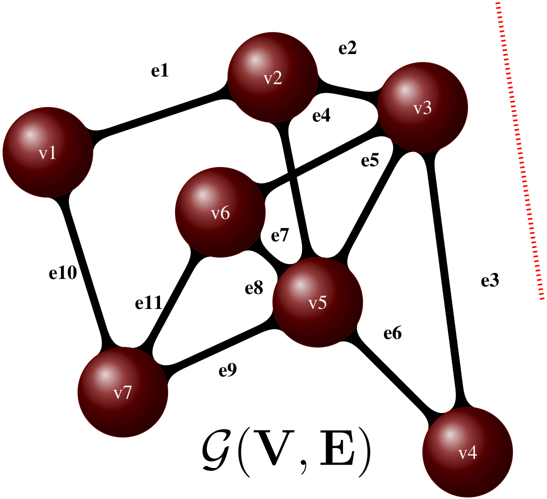

Spectral Graph Partitioning

2-Laplacian partitioning

$$\min_{V_k \subset V} \frac{\|\boldsymbol{A}\hat{\mathbf{x}}\|_2^2}{\|\hat{\mathbf{x}}\|_2^2} \approx \min_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} = \min_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\mathbf{x}^\top \boldsymbol{L}\mathbf{x}}{\mathbf{x}^\top \mathbf{x}}$$

subject to $\mathbf{1}^\top \mathbf{x} = 0$

p-Laplacian partitioning, $p \in (1, 2]$

$$\min_{V_k \subset V} \frac{\|\boldsymbol{A}\hat{\mathbf{x}}\|_p^p}{\|\hat{\mathbf{x}}\|_p^p} \approx \min_{\mathbf{x} \in \mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p} = \min_{\mathbf{x} \in \mathbb{R}^n} \frac{(\boldsymbol{A}\mathbf{x})^\top \phi_p(\boldsymbol{A}\mathbf{x})}{\mathbf{x}^\top \phi_p(\mathbf{x})}$$

subject to $\mathbf{1}^\top \phi_p(\mathbf{x}) = 0$

Spectral Graph Partitioning

2-Laplacian partitioning

$$\min_{V_k \subset V} \frac{\|\mathbf{A}\hat{\mathbf{x}}\|_2^2}{\|\hat{\mathbf{x}}\|_2^2} \approx \min_{\mathbf{x}\in\mathbb{R}^n} \frac{\|\mathbf{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} = \min_{\mathbf{x}\in\mathbb{R}^n} \frac{\mathbf{x}^\top \mathbf{L}\mathbf{x}}{\mathbf{x}^\top \mathbf{x}}$$

$$\text{subject to } \mathbf{1}^\top \mathbf{x} = 0$$

p-Laplacian partitioning, $p \in (1, 2]$

$$\min_{V_k \subset V} \frac{\|\mathbf{A}\hat{\mathbf{x}}\|_p^p}{\|\hat{\mathbf{x}}\|_p^p} \approx \min_{\mathbf{x}\in\mathbb{R}^n} \frac{\|\mathbf{A}\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p} = \min_{\mathbf{x}\in\mathbb{R}^n} \frac{(\mathbf{A}\mathbf{x})^\top \phi_p(\mathbf{A}\mathbf{x})}{\mathbf{x}^\top \phi_p(\mathbf{x})}$$

$$\text{subject to } \mathbf{1}^\top \phi_p(\mathbf{x}) = 0$$

Spectral Graph Partitioning

2-Laplacian partitioning

$$\min_{V_k \subset V} \frac{\|\boldsymbol{A}\hat{\mathbf{x}}\|_2^2}{\|\hat{\mathbf{x}}\|_2^2} \approx \min_{\mathbf{x}\in\mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} = \min_{\mathbf{x}\in\mathbb{R}^n} \frac{\mathbf{x}^\top \boldsymbol{L}\mathbf{x}}{\mathbf{x}^\top\mathbf{x}}$$

p-Laplacian partitioning, $p \in (1, 2]$

$$\min_{V_k \subset V} \frac{\|\boldsymbol{A}\hat{\mathbf{x}}\|_p^p}{\|\hat{\mathbf{x}}\|_p^p} \approx \min_{\mathbf{x}\in\mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p} = \min_{\mathbf{x}\in\mathbb{R}^n} \frac{(\boldsymbol{A}\mathbf{x})^\top \phi_p(\boldsymbol{A}\mathbf{x})}{\mathbf{x}^\top \phi_p(\mathbf{x})}$$

subject to $\mathbf{1}^\top \mathbf{x} = 0$

subject to $\mathbf{1}^\top \phi_p(\mathbf{x}) = 0$

Spectral Graph Partitioning

2-Laplacian partitioning

$$\min_{V_k \subset V} \frac{\|\boldsymbol{A}\hat{\mathbf{x}}\|_2^2}{\|\hat{\mathbf{x}}\|_2^2} \approx \min_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} = \min_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\mathbf{x}^\top \boldsymbol{L} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}}$$

subject to $\mathbf{1}^\top \mathbf{x} = 0$

p-Laplacian partitioning, $p \in (1, 2]$

$$\min_{V_k \subset V} \frac{\|\boldsymbol{A}\hat{\mathbf{x}}\|_p^p}{\|\hat{\mathbf{x}}\|_p^p} \approx \min_{\mathbf{x} \in \mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p} = \min_{\mathbf{x} \in \mathbb{R}^n} \frac{(\boldsymbol{A}\mathbf{x})^\top \phi_p(\boldsymbol{A}\mathbf{x})}{\mathbf{x}^\top \phi_p(\mathbf{x})}$$

subject to $\mathbf{1}^\top \phi_p(\mathbf{x}) = 0$

Promoting Sparse Solutions

2-Laplacian partitioning

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}} \; \frac{\|\mathbf{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2}$$

$$\text{subject to } \mathbf{1}^\top \mathbf{x} = 0$$

p-Laplacian partitioning

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}} \; \frac{\|\mathbf{A}\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p}$$

$$\text{subject to } \mathbf{1}^\top \phi_p(\mathbf{x}) = 0$$

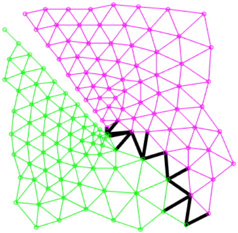

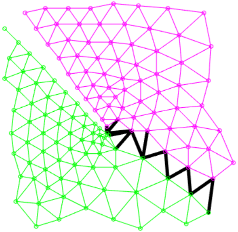

Promoting Sparse Solutions

2-Laplacian partitioning

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}} \frac{\|\mathbf{Ax}\|_2^2}{\|\mathbf{x}\|_2^2}$$

$$\text{subject to } \mathbf{1}^\top \mathbf{x} = 0$$

p-Laplacian partitioning

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}} \frac{\|\mathbf{Ax}\|_p^p}{\|\mathbf{x}\|_p^p}$$

$$\text{subject to } \mathbf{1}^\top \phi_p(\mathbf{x}) = 0$$

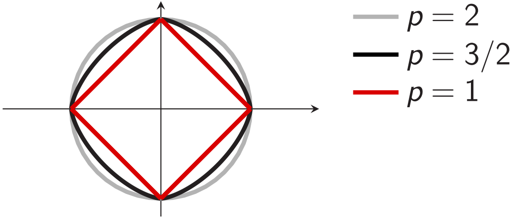

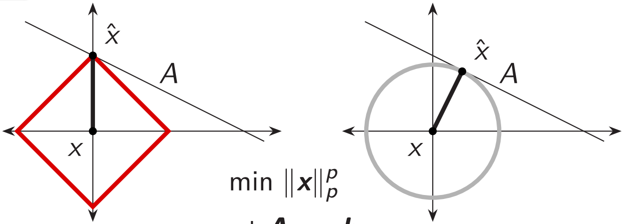

$$\min \|\boldsymbol{x}\|_p^p$$

$$\text{s.t } \boldsymbol{Ax} = \boldsymbol{b}$$

Constrained Optimization Problem

2-norm minimization

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}}\ \frac{\|\mathbf{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2}$$
$$\text{subject to } \mathbf{1}^\top \mathbf{x} = 0$$

p-norm minimization

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

Constrained Optimization Problem

2-norm minimization

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}}\ \frac{\|\mathbf{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2}$$
$$\text{subject to } \mathbf{1}^\top \mathbf{x} = 0$$

p-norm minimization

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

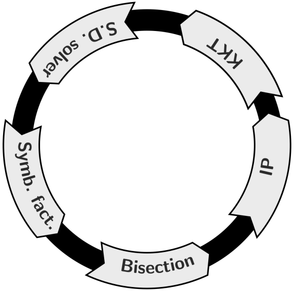

Constrained Optimization Problem

2-norm minimization

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}}\ \frac{\|\mathbf{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2}$$
$$\text{subject to } \mathbf{1}^\top \mathbf{x} = 0$$

p-norm minimization

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

Constrained Optimization Problem

2-norm minimization

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}} \; \frac{\|\mathbf{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2}$$
$$\text{subject to } \mathbf{1}^\top \mathbf{x} = 0$$

p-norm minimization

$$\underset{\mathbf{x}\in\mathbb{R}^n}{\text{minimize}} \; \frac{\|\mathbf{A}\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p}$$
$$\text{subject to } \mathbf{1}^\top \phi_p(\mathbf{x}) = 0$$

$$\widehat{\mathbf{x}} = \mathbf{x} - \frac{\mathbf{1}^\top \mathbf{x}}{n}$$

$$\phi_p(x) = |x|^{p-2} x$$
$$\phi_p^{-1}(x) = |x|^{\frac{1}{p-1}} \operatorname{sign}(x)$$
$$\widehat{\mathbf{x}}_p = \phi_p^{-1}\left(\phi_p(\mathbf{x}) - \frac{\mathbf{1}^\top \phi_p(\mathbf{x})}{n}\right)$$

Decreasing the Value of p

Algorithm 1 *p*-Laplacian Bisection

Input: $\mathbf{x}_0$ ▷ METIS or KaHIP bisection
Output: $\mathbf{x}_p^{\min}$ ▷ *p*-Laplacian bisection

```
function pLaplacian(A, x_0, b^max, β, max_it)
    r_c^min ← RCCut(x_0)
    p = 2,    x = x_0
    for k=0:max_it do
        p_k = 1 + e^(−βk/max_iters)
        x_k^min ← pLaplacianDescent(A, p_k)
    end for
    return x_p^min
end function
```

Ratio Cheeger cut (Cheeger et. al. 1969):

$$\mathrm{RCCut}(V_k, \overline{V}_k) = \frac{\mathrm{cut}(V_k, \overline{V}_k)}{\min\{|V_k|, |\overline{V}_k|\}}.$$

Decreasing the Value of p

Algorithm 2 p-Laplacian Descent

Input: $\mathbf{x}_0$ ▷ approximation of the p-eigenvector
Output: $\mathbf{x}_p^{\min}$ ▷ p-Laplacian bisection

```
function pLaplacianDescent(A, x_0, p)
    while not converged do
        x ← x_0
        x ← x̂_p
        r_c = RCCut(x)
        b_r = ImBal(x)
        if r_c < r_c^min and b_r < b_r^max then
            x_p^min ← x                          ▷ save best solution
            r_c^min ← r_c                        ▷ and minimum cut
        end if
        g ← ∇f(x)
        α ← argmin_α f(φ_p^{-1}(x − αg))
        x ← x − αg
    end while
    return x_p^min
end function
```

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

$$b_r = \frac{\left||V_2| - |\overline{V}_2|\right|}{|V|}$$

$$O(n^3) \to O(m)$$

Triangular Delaunay Meshes

Mesh	Nodes	Edges	METIS/p-Lap	%
sphere3	258	768	58/50	13.7
stufe	1,036	1,868	36/17	**52.7**
airfoil1	4,253	12,289	85/70	17.6
barth4	6,019	17,473	115/86	**25.2**
commanche	7,920	11,880	27/23	14.8
barth5	15,606	45,878	155/140	9.7
crack_dual	20,141	30,043	96/78	**18.7**
brack2	62,631	366,559	787/708	10.0
tandem	94,069	183,212	478/446	6.7
wave	156,317	1,059,331	8,944/8,713	2.5

Triangular Delaunay Meshes

Mesh	Nodes	Edges	METIS/p-Lap	%
sphere3	258	768	58/50	13.7
stufe	1,036	1,868	36/17	**52.7**
airfoil1	4,253	12,289	85/70	17.6
barth4	6,019	17,473	115/86	**25.2**
commanche	7,920	11,880	27/23	14.8
barth5	15,606	45,878	155/140	9.7
crack_dual	20,141	30,043	96/78	**18.7**
brack2	62,631	366,559	787/708	10.0
tandem	94,069	183,212	478/446	6.7
wave	156,317	1,059,331	8,944/8,713	2.5

Triangular Delaunay Meshes

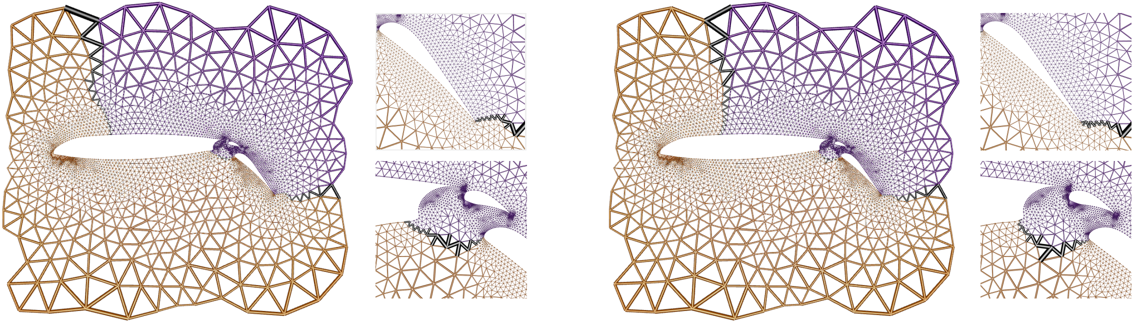

Mesh	Nodes	Edges	METIS/p-Lap	%
airfoil1	4,253	12,289	85/70	17.6

Power Networks

Case	Nodes	Edges	METIS/p-Lap	%	KaHIP/p-Lap	%
1354pegase	1,354	1,991	20/17	15.0	18/16	11.1
1888rte	1,888	2,531	28/25	10.7	19/18	5.3
6470rte	6,470	9,005	43/32	**25.6**	28/25	**10.7**
6495rte	6,495	9,019	35/34	2.9	32/25	**21.9**
6515rte	6,516	9,037	49/32	**34.7**	41/38	**7.3**
9241pegase	9,241	16,049	29/22	**24.1**	16/15	6.3
13659pegase	13,659	20,467	37/32	13.5	21/20	4.7

Power Networks

Case	Nodes	Edges	METIS/p-Lap	%	KaHIP/p-Lap	%
1354pegase	1,354	1,991	20/17	15.0	18/16	11.1
1888rte	1,888	2,531	28/25	10.7	19/18	5.3
6470rte	6,470	9,005	43/32	**25.6**	28/25	**10.7**
6495rte	6,495	9,019	35/34	2.9	32/25	**21.9**
6515rte	6,516	9,037	49/32	**34.7**	41/38	**7.3**
9241pegase	9,241	16,049	29/22	**24.1**	16/15	6.3
13659pegase	13,659	20,467	37/32	13.5	21/20	4.7

Power Networks

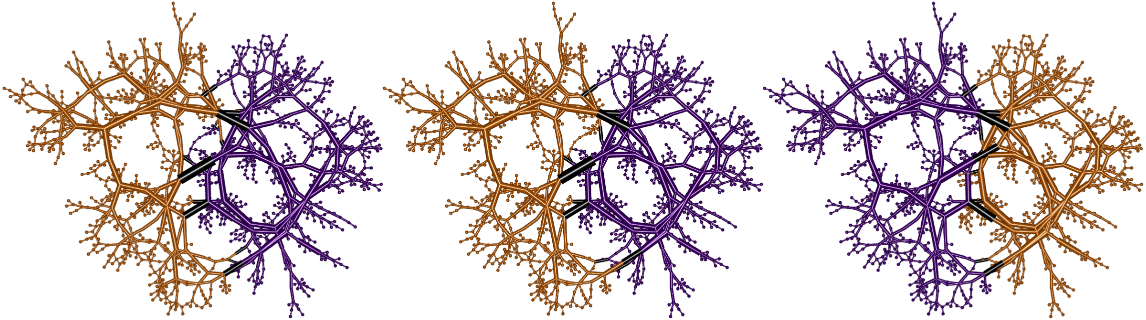

Case	Nodes	Edges	METIS/p-Lap	%	KaHIP/p-Lap	%
1354pegase	1,354	1,991	20/17	15.0	18/16	11.1

Power Networks

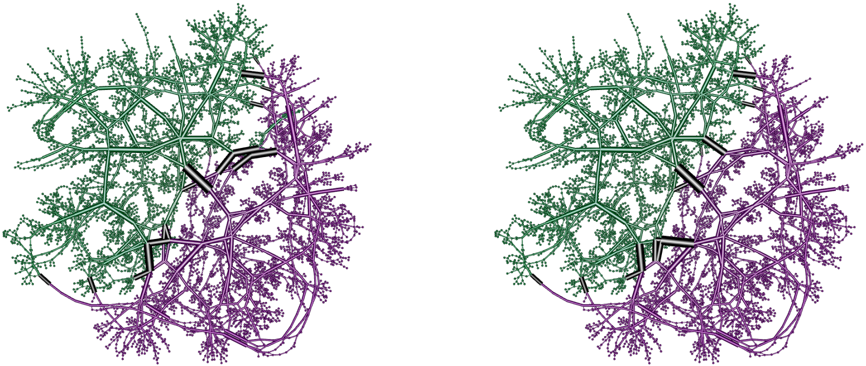

Case	Nodes	Edges	METIS/p-Lap	%	KaHIP/p-Lap	%
6495rte	6,495	9,019	35/34	2.9	32/25	21.9

Social Networks

Case	Nodes	Edges	METIS/p-Lap	%	KaHIP/p-Lap	%
geom	7,343	11,898	61/14	77.1	58/18	68.9
ca-HepPh	9,877	25,973	6659/3974	40.3	5689/4629	18.6
ca-HepTh	12,008	118,489	1578/1324	16.1	2041/1973	3.3
Reuters911	13,332	148,038	26072/18979	27.2	26387/23354	11.5
p2p-Gnutella30	36,682	48,507	2631/2326	11.6	3336/3041	8.9
coAuthorsCiteseer	227,320	814,134	22756/19811	13.0	36388/26761	26.5
amazon0302	262,111	454,587	2837/2437	14.1	2730/2540	7.0
amazon0312	400,727	1,049,624	10665/8876	16.8	14899/7698	48.3
amazon0505	410,236	1,140,454	11680/9123	21.9	14972/8727	41.8
amazon0601	403,394	1,140,070	11572/9127	21.1	11443/9731	15.0
coPapersDBLP	540,486	15,245,729	594631/551400	7.3	818757/492700	39.9

Social Networks

Case	Nodes	Edges	METIS/p-Lap	%	KaHIP/p-Lap	%
geom	7,343	11,898	61/14	77.1	58/18	68.9
ca-HepPh	9,877	25,973	6659/3974	40.3	5689/4629	18.6
ca-HepTh	12,008	118,489	1578/1324	16.1	2041/1973	3.3
Reuters911	13,332	148,038	26072/18979	27.2	26387/23354	11.5
p2p-Gnutella30	36,682	48,507	2631/2326	11.6	3336/3041	8.9
coAuthorsCiteseer	227,320	814,134	22756/19811	13.0	36388/26761	26.5
amazon0302	262,111	454,587	2837/2437	14.1	2730/2540	7.0
amazon0312	400,727	1,049,624	10665/8876	16.8	14899/7698	48.3
amazon0505	410,236	1,140,454	11680/9123	21.9	14972/8727	41.8
amazon0601	403,394	1,140,070	11572/9127	21.1	11443/9731	15.0
coPapersDBLP	540,486	15,245,729	594631/551400	7.3	818757/492700	39.9

Concentrated Results

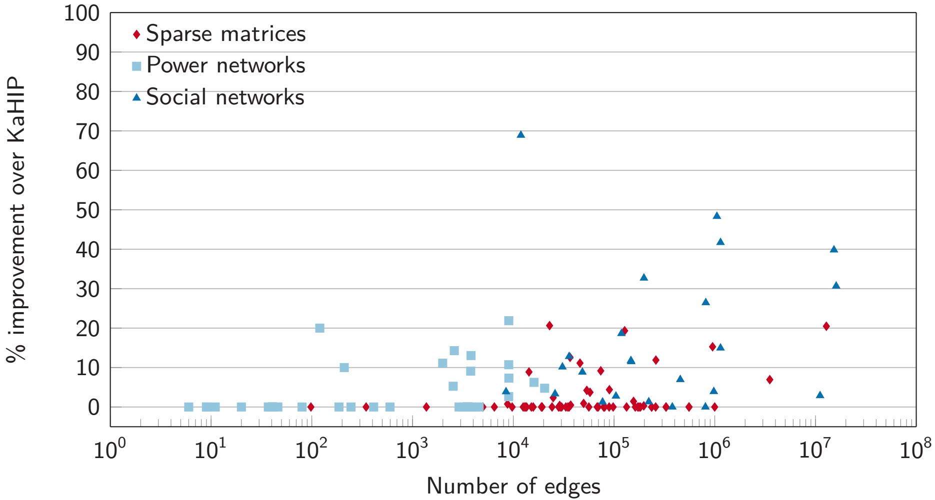

Concentrated Results

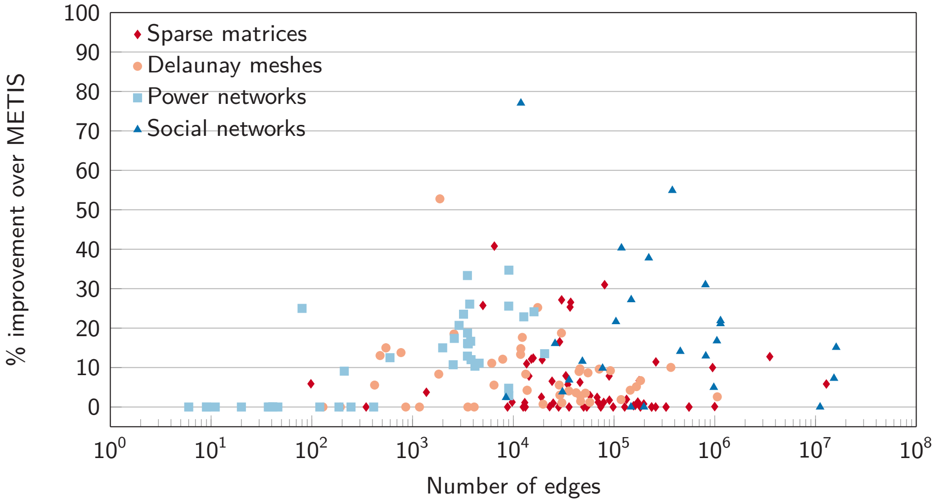

Parallel P-Laplacian Graph Partitioning

15,360,000 elements, 2,618,021 nodes

CPU	Peak FLOPS	(Peak MEM.)	Elapsed time (s)
Intel Xeon	499.2 GFLOPS	68 GB/s	1862.3 (1 core)
E5-2690 v3 CPU	—	—	284.8 (12 cores)
Nvidia P100 GPU	4.7 TFLOPS	732 GB/s	28.2

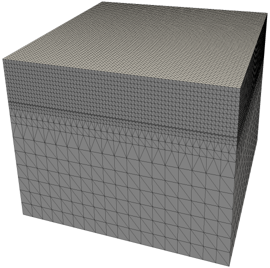

Parallel P-Laplacian Graph Partitioning

15,360,000 elements, 2,618,021 nodes

CPU	Peak FLOPS	(Peak MEM.)	Elapsed time (s)
Intel Xeon	499.2 GFLOPS	68 GB/s	1862.3 (1 core)
E5-2690 v3 CPU	—	—	284.8 (12 cores)
Nvidia P100 GPU	4.7 TFLOPS	732 GB/s	28.2

Parallel P-Laplacian Graph Partitioning

15,360,000 elements, 2,618,021 nodes

Method	edgecut	(METIS%)	$\max_{i\leq 128} \lvert V_i^{128} \rvert$	$b_r^{128}(\%)$
METIS	600,450	—	120,000	0
p-Laplacian	558,492	7.5	120,939	0.78
Hybrid	570,240	5.3	120,660	0.55

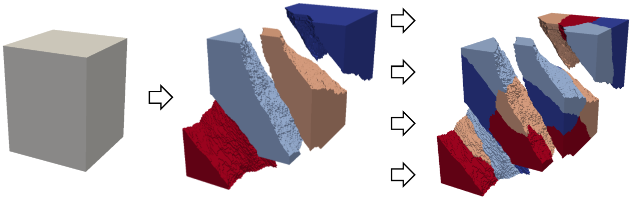

Input mesh

Part 1: Serial partitioning on shared memory computer (METIS)

Part 2: Parallel partitioning on distributed memory computer (p-Laplacian)

Parallel P-Laplacian Graph Partitioning

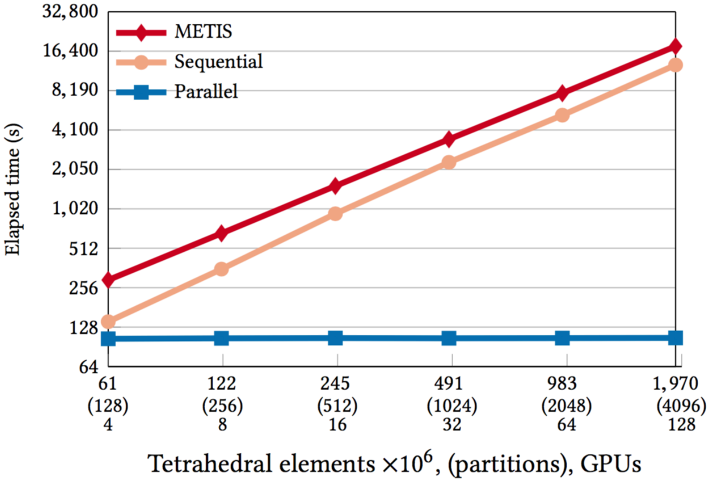

Balanced Graph Partition Refinement using the Graph p-Laplacian

Toby Simpson
Institute of Computational Science, Università della Svizzera italiana
toby.simpson@usi.ch

Dimosthenis Pasadakis
Institute of Computational Science, Università della Svizzera italiana
dimosthenis.pasadakis@usi.ch

Drosos Kourounis
Institute of Computational Science, Università della Svizzera italiana
drosos.kourounis@usi.ch

Kohei Fujita
Earthquake Research Institute & Department of Civil Engineering, The University of Tokyo
Advanced Institute for Computational Science, RIKEN
fujita@eri.u-tokyo.ac.jp

Takuma Yamaguchi
Earthquake Research Institute & Department of Civil Engineering, The University of Tokyo
yamaguchi@eri.u-tokyo.ac.jp

Tsuyoshi Ichimura
Earthquake Research Institute & Department of Civil Engineering, The University of Tokyo
Advanced Institute for Computational Science, RIKEN
ichimura@eri.u-tokyo.ac.jp

Olaf Schenk
Institute of Computational Science, Università della Svizzera italiana
olaf.schenk@usi.ch

ABSTRACT

A continuous formulation of the optimal 2-way graph partitioning based on the p-norm minimization of the graph Laplacian Rayleigh quotient is presented, which provides a sharp approximation to the balanced graph partitioning problem, the optimality of which is known to be NP-hard. The minimization is initialized from a cut provided by a state-of-the-art multilevel recursive bisection algorithm, and then a continuation approach reduces the p-norm from a 2-norm towards a 1-norm, employing for each value of p a feasibility-preserving steepest-descent method that converges on the p-Laplacian eigenvector. A filter favors iterates advancing towards minimum edgecut and partition load imbalance. The complexity of the suggested approach is linear in graph edges. The simplicity of the steepest-descent algorithm renders the overall approach highly scalable and efficient in parallel distributed architectures. Parallel implementation of recursive bisection on multi-core CPUs and GPUs are presented for large-scale graphs with up to 1.9 billion tetrahedra. The suggested approach exhibits improvements of up to 52.8% over METIS for graphs originating from triangular Delaunay meshes, 34.7% over METIS and 21.9% over KaHIP for power network graphs, 40.8% over METIS and 20.6% over KaHIP for sparse matrix graphs, and finally 93.2% over METIS for graphs emerging from social networks.

PASC '18, July 2–4, 2018, Basel, Switzerland

ACM ISBN 978-1-4503-5891-0/18/07…$15.00
https://doi.org/10.1145/3218176.3218232

CCS CONCEPTS

• **Mathematics of computing** → **Graph theory**; **Graph algorithms**; *Combinatorial algorithms*; *Combinatorial optimization*; *Approximation algorithms*; • **Computing methodologies** → *Parallel algorithms*;

KEYWORDS

Combinatorial mathematics; Graph Theory; Spectral Methods; Parallel processing;

ACM Reference Format:
Toby Simpson, Dimosthenis Pasadakis, Drosos Kourounis, Kohei Fujita, Takuma Yamaguchi, Tsuyoshi Ichimura, and Olaf Schenk. 2018. Balanced Graph Partition Refinement using the Graph p-Laplacian. In *PASC '18: Platform for Advanced Scientific Computing Conference, July 2–4, 2018, Basel, Switzerland.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3218176.3218232

1 INTRODUCTION

1.1 Applications using graph partitioning

Graph partitioning is a ubiquitous technique in scientific computing, in which domains are modeled as weighted or unweighted graphs and separated into parts either by the removal of vertices or edges. It has been used in several areas of science, such as the determination of genomic sequences [8, 34], the design of space-efficient circuit placements [14], the organization of databases for efficient retrieval [23], the large scale analysis of finite element meshes [13, 19], and the ordering of sparse matrices for efficient sparse multifrontal factorization [6, 26, 27, 32].

In general, good solutions require that cuts are small and parties have equal size. The problem arises, for instance, when assigning work to a parallel computer. In order to achieve efficiency, the workload (partition size) should be balanced evenly among the

Toby Simpson, Dimosthenis Pasadakis, Drosos Kourounis, Kohei Fujita, Takuma Yamaguchi, Tsuyoshi Ichimura, and Olaf Schenk.
Balanced graph partition refinement using the graph p-laplacian.
In Proceedings of the Platform for Advanced Scientific Computing Conference, PASC '18, pages 8:1–8:11, New York, NY, USA, 2018. ACM.

Conclusions

- Refinement of METIS or KaHIP partition.
- Significant improvements for the "notorious" social graphs.
- Highly scalable parallel implementations in both distributed multi-core CPU platforms and GPU accelerators.
- Hybrid recursive bisection method using multiple GPUs.

Future Pespectives

- Assess relative benefits of runtime performance and convergence.
- Embody the p-Laplacian algorithm in a multilevel hierarchy based framework.
- Classification of graphs according to their topological features.

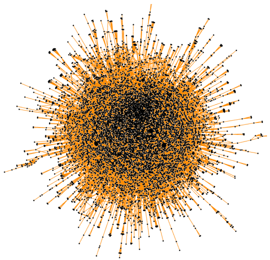

David A. Bader, Henning Meyerhenke, Peter Sanders, Christian Schulz, Andrea Kappes, and Dorothea Wagner.
Benchmarking for Graph Clustering and Partitioning, pages 73–82.
Springer New York, New York, NY, 2014.

Paolo Boldi and Sebastiano Vigna.
The WebGraph framework I: Compression techniques.
In *Proc. of the Thirteenth International World Wide Web Conference (WWW 2004)*, pages 595–601, Manhattan, USA, 2004. ACM Press.

Ralf Diekmann and Robert Preis.
AG-Monien Graph Collection.
`http://www2.cs.uni-paderborn.de/fachbereich/AG/monien/RESEARCH/PART/graphs.html`, January 2018.

Jure Leskovec and Andrej Krevl.
SNAP Datasets: Stanford large network dataset collection.
`http://snap.stanford.edu/data`, June 2014.

Alex Pothen.
Structural engineering matrices from NASA.
`https://sparse.tamu.edu/Pothen`, January 2003.

R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas.
MATPOWER: Steady-State Operations, Planning, and Analysis Tools for Power Systems Research and Education.
IEEE Trans. on Power Syst., 26(1):12–19, Feb 2011.

Additional material

Spectral Graph Partitioning

Calculate an edge separator using the Fiedler eigenvector of $\mathbf{L} \in \mathbb{R}^{n \times n}$.

$$x_i = \begin{cases} 1\,, & i \in V_k \\ 0\,, & i \in \overline{V_k} \end{cases}$$

$$\mathbf{x}^\top \mathbf{x} = \|\mathbf{x}\|_2^2 = |V_k|$$

$\boldsymbol{A}\mathbf{x}$ has nonzero elements only for edges that connect a vertex in V_k with one in $\overline{V}_k$. Thus:

$$\mathrm{cut}(V_k, \overline{V}_k) = \|\boldsymbol{A}\mathbf{x}\|_2^2$$

$$\mathrm{RCut}(V_1, \ldots, V_k) = \frac{1}{2} \sum_{j=1}^{k} \frac{\|\boldsymbol{A}\mathbf{x}_j\|_2^2}{\|\mathbf{x}_j\|_2^2}$$

Allowing $x_i \in \mathbb{R}$:

$$\min_{V_k \subset V} \frac{\|\boldsymbol{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} \approx \min_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} = \min_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\langle \mathbf{x}, \Delta_2 \mathbf{x} \rangle}{\langle \mathbf{x}, \mathbf{x} \rangle}$$

$$\min_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_2^2}{\|\mathbf{x}\|_2^2} = \min_{\boldsymbol{x} \in \mathbb{R}^n} \frac{\mathbf{x}^\top \boldsymbol{L} \mathbf{x}}{\mathbf{x}^\top \mathbf{x}} = \lambda_i$$

$$(\lambda_1 = 0, \mathbf{v}_2 = c\mathbf{1}) \rightarrow V = V \cup \emptyset$$

$$(\lambda_2, \boldsymbol{v}_2) \rightarrow \text{Fiedler's eigenvector}$$

P-Laplacian Graph Partitioning

For $\mathbf{x} \in \mathbb{R}^n, p \in \mathbb{R}$,
$\|\mathbf{x}\|_p^p = \sum_i |x_i|^p, \phi_p(x) = |x|^{p-2}x$:

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

The Rayleigh quotient:

$$\Delta_p(\mathbf{x}) = \boldsymbol{A}^\top \phi_p(\boldsymbol{A}\mathbf{x}) \Rightarrow$$
$$\Delta_p(\mathbf{v}) = \boldsymbol{A}^\top \phi_p(\boldsymbol{A}\mathbf{v}) = \lambda \phi_p(\mathbf{v})$$

Starting from the gradient:

$$\nabla f(\mathbf{x}) \overset{\boldsymbol{x}=\boldsymbol{v}}{\Longrightarrow} \nabla_{\mathbf{x}} \frac{\|\boldsymbol{A}\mathbf{v}\|_p^p}{\|\mathbf{v}\|_p^p} = \frac{p}{\mathbf{v}^\top \phi(\mathbf{v})} \left(\lambda\phi(\mathbf{v}) - \lambda\phi(\mathbf{v})\right) = 0$$

The ϕ_p function

$$\min_{\mathbf{x}\in\mathbb{R}^n} \frac{\|\boldsymbol{A}\mathbf{x}\|_p^p}{\|\mathbf{x}\|_p^p}.$$

This objective is the Rayleigh quotient of the p-Laplacian. For $p \in \mathbb{R}$ the scalar $\phi_p : \mathbb{R} \mapsto \mathbb{R}$ function is defined as:

$$\phi_p(x) = |x|^{p-2}x \;\Rightarrow\; x\phi_p(x) = |x|^p.$$

When applied elementwise to a vector $\mathbf{x} \in \mathbb{R}^n$ the inner product with $\mathbf{x}$ returns the p-norm:

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

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

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

$$\Delta_p(\mathbf{x}) = \boldsymbol{A}^\top \phi_p(\boldsymbol{A}\mathbf{x}).$$

Balanced Graph Cut Criteria (1)

The ratio cut and ratio Cheeger cut for a partition V into $C, \overline{C}$:

$$\text{RCut}(C, \overline{C}) = \frac{\text{cut}(C, \overline{C})}{|C|} + \frac{\text{cut}(C, \overline{C})}{|\overline{C}|},$$

$$\text{RCCut}(C, \overline{C}) = \frac{\text{cut}(C, \overline{C})}{\min\{|C|, |\overline{C}|\}} \quad \Rightarrow$$

$$\text{RCC}(C, \overline{C}) \leq \text{RC}(C, \overline{C}) \leq 2\text{RCC}(C, \overline{C})$$

We define the optimal ratio cut h_{RCC} as:

$$h_{\text{RCC}} = \inf_C \text{RCC}(C, \overline{C})$$

Balanced Graph Cut Criteria (2)

Theorem

Denote by $\lambda_p^{(2)}$ the second eigenvalue of the unnormalised p-Laplacian. For $p > 1$,

$$\left(\frac{2}{\max_i d_i}\right)^{p-1}\left(\frac{h_{\mathrm{RCC}}}{p}\right)^{p} \leq \lambda_p^{(2)} \leq 2^{p-1} h_{\mathrm{RCC}}.$$

Proof in Buehler & Hein, 2009.

Considering the limit $p \to 1$, the bounds on λ_p become tight as $p \to 1$. The second eigenvalue of the unnormalized p-Laplacian approximates the optimal ratio Cheeger cut arbitrarily well in this limit.

Objective: Transform the real-valued second eigenvector of the p-Laplacian into a partitioning of the graph.

Approach: Threshold the second eigenvector $v_p^{(2)}$ to obtain the partitioning. The optimal threshold is determined by minimizing the corresponding Cheeger cut.

$$\arg\min_{C_t = \{i \in V | v_p^{(2)}(i) > t\}} \mathrm{RCC}(C_t, \overline{C}_t)$$

Balanced Graph Cut Criteria (3)

Question: How good the cut values obtained by thresholding the second eigenvector of the p-Laplacian are compared to optimal Cheeger cut values.

Theorem

Denote by h^*_{RCC} the ratio Cheeger cut values obtained by thresholding the second eigenvector $v_p^{(2)}$ of the unnormalized p-Laplacian. Then for $p > 1$,

Proof in Buehler & Hein, 2009.