



# Nonlinear spectral clustering with C++ GraphBLAS

D. Pasadakis & O. Schenk – V. Vlacic & A.-J. Yzeman

Università della Svizzera Italiana, Institute of Computing, Lugano, Switzerland.  
Computing Systems Lab, Huawei Zurich Research Center, Switzerland.



• D. Pasadakis, C. L. Alappat, O. Schenk, and G. Wellein, “Multiway  $p$ -spectral graph cuts on Grassmann manifolds,” Machine Learning, vol. 111, Feb 2022, DOI: [10.1007/s10994-021-06108-1](https://doi.org/10.1007/s10994-021-06108-1).

• A. N. Yzeman, D. Di Nardo, J. M. Nash, and W. J. Suijlen, “A C++ GraphBLAS: specification, implementation, parallelisation, and evaluation,” 2020, <http://albert-jan.yzeman.net/PDFs/yzeman20.pdf>.

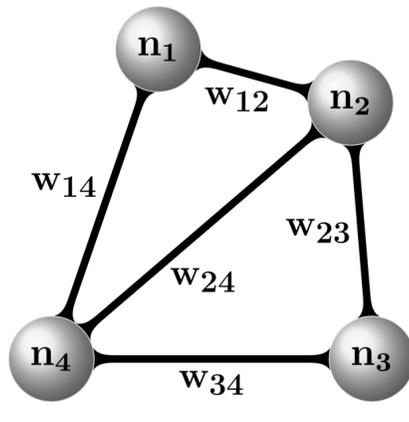
→ We combine their benefits and furnish the first  $p$ -norm spectral clustering algorithm applicable to large-scale data for shared-memory machines.

## Nonlinear spectral clustering

For a graph  $\mathcal{G}(V, E, \mathbf{W})$  with  $n$  vertices and  $m$  edges:

- Scalar function  $\phi_p(x) = |x|^{p-1}\text{sign}(x)$ , and  $p$ -norm  $\|\mathbf{u}\|_p = \sqrt[p]{\sum_{i=1}^n |u_i|^p}$  for  $p \in (1, 2]$ .
- $p$ -Laplacian operator  $(\Delta_p \mathbf{u})_i = \sum_{j \in V} \mathbf{W}_{ij} \phi_p(u_i - u_j)$

$$\mathbf{W} = \begin{bmatrix} 0 & w_{12} & 0 & w_{14} \\ w_{12} & 0 & w_{23} & w_{24} \\ 0 & w_{23} & 0 & w_{34} \\ w_{14} & w_{24} & w_{34} & 0 \end{bmatrix}, \quad d_{ii} = \begin{bmatrix} \sum_j w_{1j} \\ \sum_j w_{2j} \\ \sum_j w_{3j} \\ \sum_j w_{4j} \end{bmatrix}$$



Calculate a multiway partition using  $k$  eigenvectors of  $\Delta_p \in \mathbb{R}^{n \times n}$ .

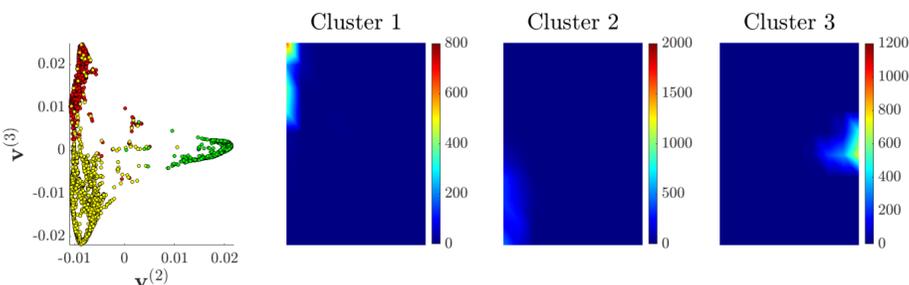
### 2-spectral clustering

$$\min_{\mathbf{U} \in \mathbb{R}^{n \times k}} F_2(\mathbf{U}) = \text{Tr}(\mathbf{U}^T \mathbf{L} \mathbf{U}),$$

s.t.  $\mathbf{U}^T \mathbf{U} = \mathbf{I}$ .

### Mutually orthogonal eigenvectors

$$\begin{array}{c|ccc} & \mathbf{u}_1 & \dots & \mathbf{u}_k \\ \mathbf{U}_1 & u_{11} & \dots & u_{1k} \\ \vdots & \vdots & \dots & \vdots \\ \mathbf{U}_n & u_{n1} & \dots & u_{nk} \end{array}$$




### $p$ -spectral clustering

$$\text{minimize}_{\mathbf{U} \in \mathcal{G}^z(k, n)} F_p(\mathbf{U}) = \sum_l \sum_{i,j=1}^n \frac{w_{ij} |u_i^l - u_j^l|^p}{2 \| \mathbf{u}^l \|_p^p}$$

### Optimization on the manifold

- Software package ROPTLIB. [github.com/whuang08/ROPTLIB](https://github.com/whuang08/ROPTLIB)
- Newton's method on the Grassmann.

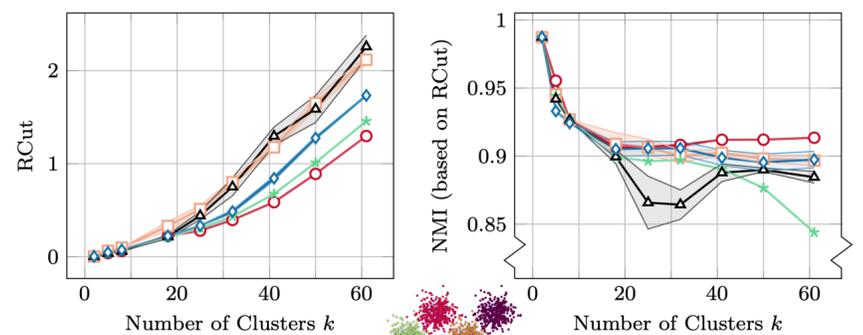


## Numerical results

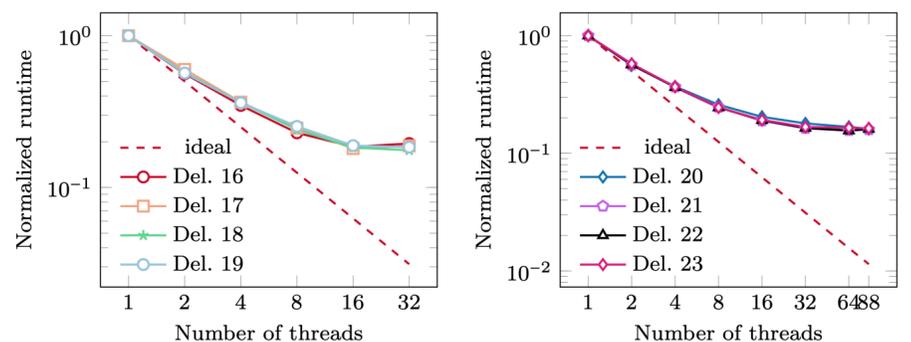
Gaussian datasets with an increasing number of clusters  $k \in \{2, \dots, 61\}$ , nodes  $n = 400 * k$ , and edges  $m \in \{4902, \dots, 160312\}$ .

✓ Superior results over competing spectral methods in terms of balanced graph cut metrics (RCut), and accuracy of classification (NMI).

—○— pGrass —□— Spec —\*— pSpec —▲— kCuts —◇— pMulti —◇— Graclus



Scalability studies in Delaunay graphs after expressing the objective, gradient, Hessian, and the k-means step, in C++ GraphBLAS terms.



→ Up to 32 threads for the mid-scale cases of node size  $n \in [2^{16}, 2^{19}]$ , up to 88 threads for large-scale cases with  $n \in [2^{20}, 2^{23}]$ . Edge distribution is  $m \approx 6 * n$ .

✓ On average, the parallel execution of the algorithm is  $5.5 \times$  faster than its sequential variant for the mid-scale tests, and  $6.4 \times$  faster for the large-scale cases.

✓ Run-time of smallest case (Del. 16) was  $\sim 300$  sec, and of the largest one (Del. 23)  $\sim 20$  hours. Only the GraphBLAS components of the algorithm exhibit excellent weak scalability for the large-scale graphs.

## A C++ GraphBLAS algorithm

### Design

<https://github.com/DmsPas/Multiway-p-spectral-clustering>

- C++11 implementation.
- Algebraic containers for sparse matrices and vectors (grb::Vector).
- Subroutines for I/O from and to the ROPTLIB data structures.
- Algebraic structure of the ring of real numbers to parallelize the SpMV operations (grb::vxm).
- Leveraging the auto-parallelisation shared-memory capabilities.

### Implementation

Compute  $\eta \mapsto \mathcal{H}\eta = (\mathcal{H}^\ell \eta^\ell)_{\ell=1}^k$  for arbitrary  $\eta \in \mathbb{R}^{k \times n}$ .

Algorithm 1 Hessian evaluation

```

Input:  $\eta$ , a  $k \times n$  matrix
        $(D[\ell])_{\ell=1}^k$ , where each  $D[\ell] = \text{diag}(\mathcal{H}^\ell)$ 
        $(H[\ell])_{\ell=1}^k$ , where each  $H[\ell] = \text{diag}(\mathcal{H}^\ell) - \mathcal{H}^\ell$ 
Output:  $\mathbf{r}$ , the result of  $\eta \mapsto \mathcal{H}\eta$ 
1: std::vector<grb::Vector<double>> grb_eta, grb_res
2: grb::Vector<double> v, w
3: ROPTLIBtoGRB( $\eta$ , grb_eta)
4: for  $\ell = 1$  to  $k$  do
5:   grb::set(v, 0)
6:   grb::vxm(v, grb_eta[\ell], H[\ell], reals_ring)
7:   grb::eWiseApply(w, grb_eta[\ell], D[\ell])
   grb::operators::mul<double>()
8:   grb::eWiseApply(grb_res[\ell], w, v,
   grb::operators::subtract<double>())
9: end for
10: GRBtoROPTLIB(grb_res,  $\mathbf{r}$ );
11: return  $\mathbf{r}$ 

```

### Open source library

<https://github.com/Algebraic-Programming/ALP>

- ① Sequential programs.
- ② Nonblocking shared-memory auto-parallelised programs.
- ③ Sequential programs with HyperDAG representations.
- ④ Distributed-memory auto-parallelised implementations.
- ⑤ Hybrid shared- and distributed-memory auto-parallelised programs.