Multiscale Spectral Graph Partitioning and Image Segmentation

Andrew V. Knyazev

University of Colorado at Denver and Health Sciences Center

Workshop on Algorithms for Modern Massive Data Sets

Stanford University and Yahoo! Research June 21–24, 2006

Supported by the NSF and the Intelligence Technology Innovation Center through the joint "Approaches to Combat Terrorism" Program Solicitation NSF 03-569.

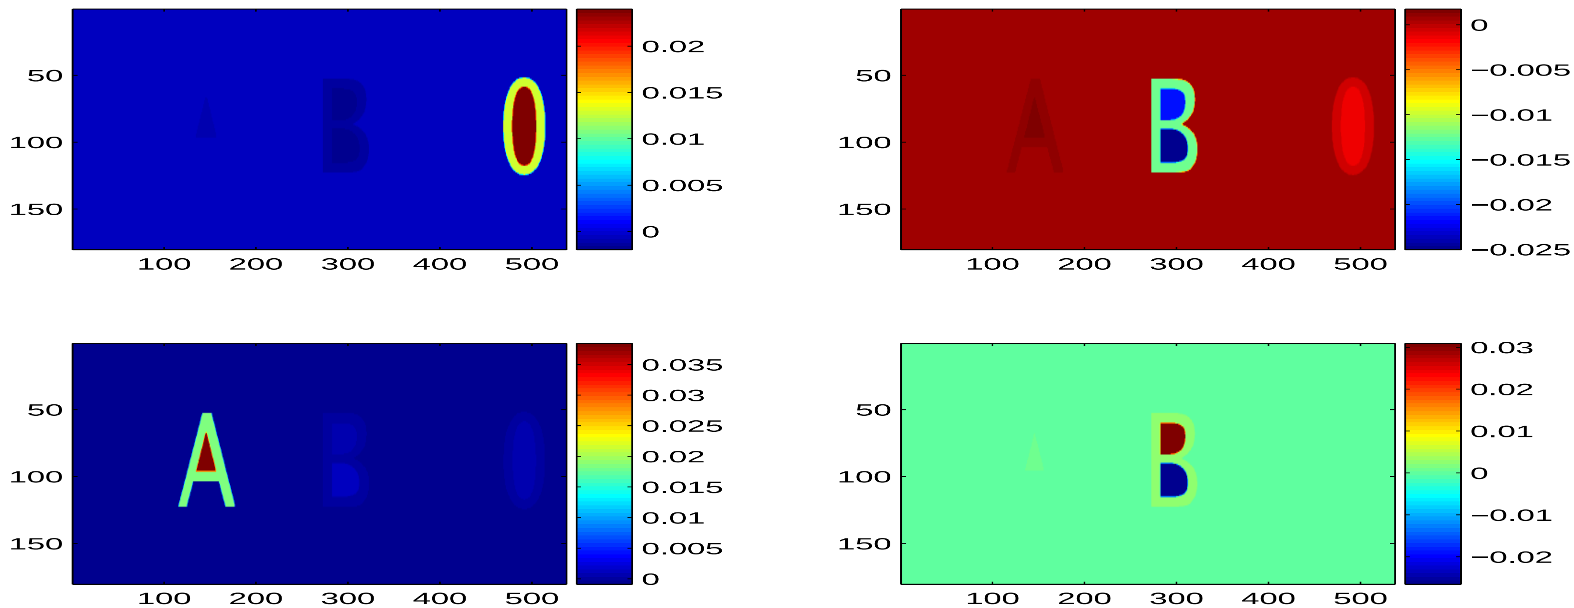

Spectral Fiedler vectors 2–4 for $T(A - \lambda B)x = 0$ synthetic picture.

Abstract. Spectral methods for graph partitioning, based on numerical solution of eigenvalue problems with the graph Laplacian, are well known to produce high quality partitioning, but are also considered to be expensive. We discuss modern preconditioned eigensolvers for computing the Fiedler vector of large scale eigenvalue problems. The ultimate goal is to find a method with a linear complexity, i.e. a method with computational costs that scale linearly with the problem size. We advocate the locally optimal block preconditioned conjugate gradient method (LOBPCG), suggested by the presenter, as a promising candidate, if matched with a high quality preconditioner. We provide preliminary numerical results, e.g., we show that a Fiedler vector for a 24 megapixel image can be computed in seconds on IBM's BlueGene/L using LOBPCG in our BLOPEX software with Hypre algebraic multigrid preconditioning.

Image Segmentation Reduced to Graph Partitioning

The set of points in an arbitrary feature space is presented as a weighted undirected graph $G = (V, E)$. Nodes of the graph are the points in the feature space. An edge is formed between every pair of nodes and the weight on each edge $W(i, j)$ is a function of the similarity between nodes i and j. A graph $G = (V, E)$ is partitioned into two disjoint complementary sets A and $B = V - A$, by removing the edges connecting the two parts.

For a given image, we construct the graph $G = (V, E)$ by taking each pixel as a node and defining the edge weight using a special function of similarity of different pixels. This weighted graph serves as a global image feature descriptor. Image segmentation is thus reduced to graph partitioning.

It is too expensive to compare each pixel against each other pixel on the picture. Usually, we only compare neighboring pixels.

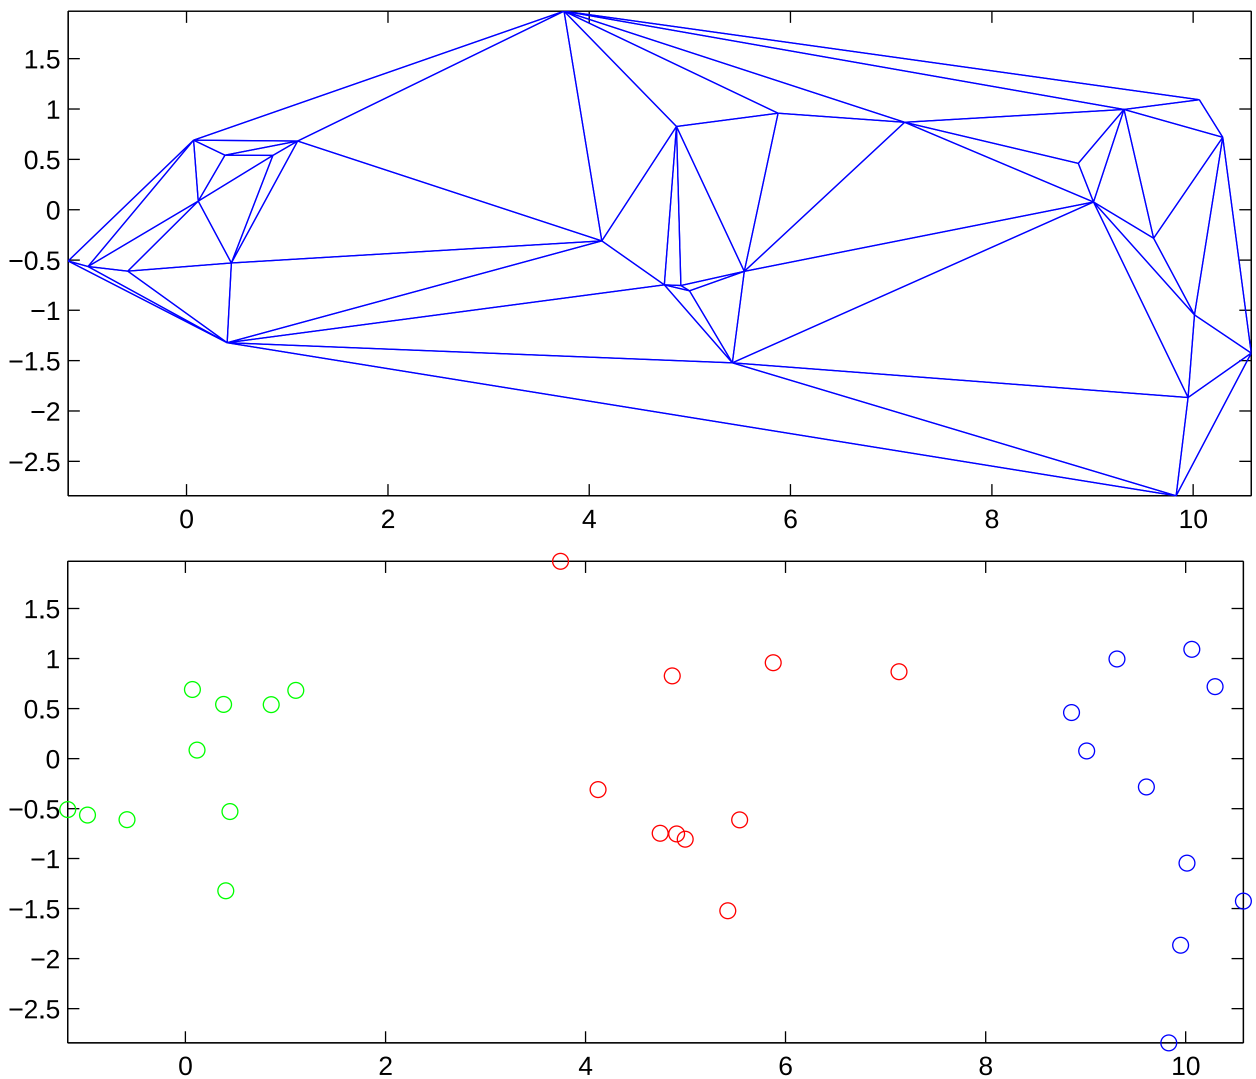
1.5
1
0.5
0
−0.5
−1
−1.5
−2
−2.5
0
2
4
6
8
10

Spectral Graph Partitioning.

$L = D - W$ is called the graph Laplacian or the stiffness matrix, where W is the symmetric graph association matrix with the zero diagonal and D is a diagonal matrix defined as a row sum of W. We solve numerically one of the following two eigenvalue problems

$$Ly = \lambda y \text{ (Spectral), or } Ly = \lambda Dy \text{ (NCuts) }, \qquad (1)$$

The eigenvector, corresponding to the second smallest eigenvalue, is called the Fiedler vector. The signs of its components determine the graph bipartition. The Fiedler eigenvector has a simple and intuitive mechanical interpretation — it is the vibration mode corresponding to the smallest nontrivial frequency of the system of masses connected by springs. The masses that are tightly connected have the tendency to oscillate together in the low frequency modes and thus form clusters.

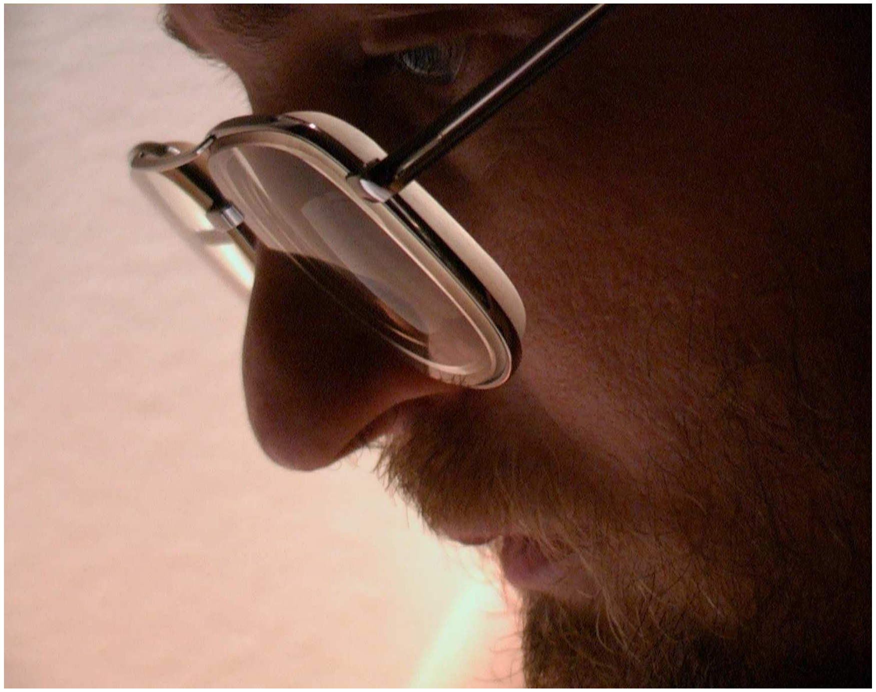

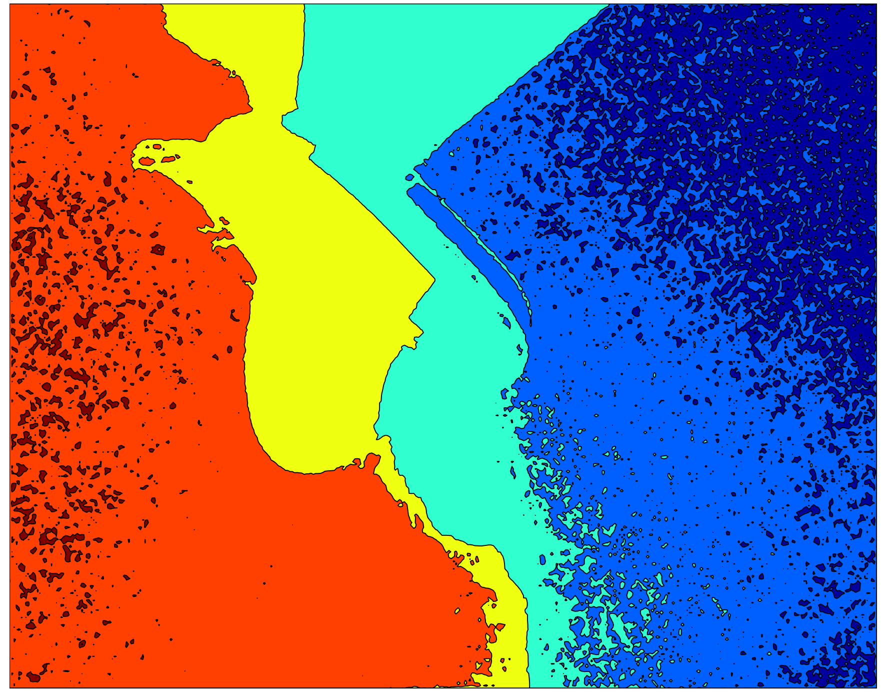

The original 1128×1488 image and its resized 282×372 Fiedler vectors, see also the next slide.

We observe that the Fiedler vectors capture the main features of the original image and thus are suitable for k–way partitioning.

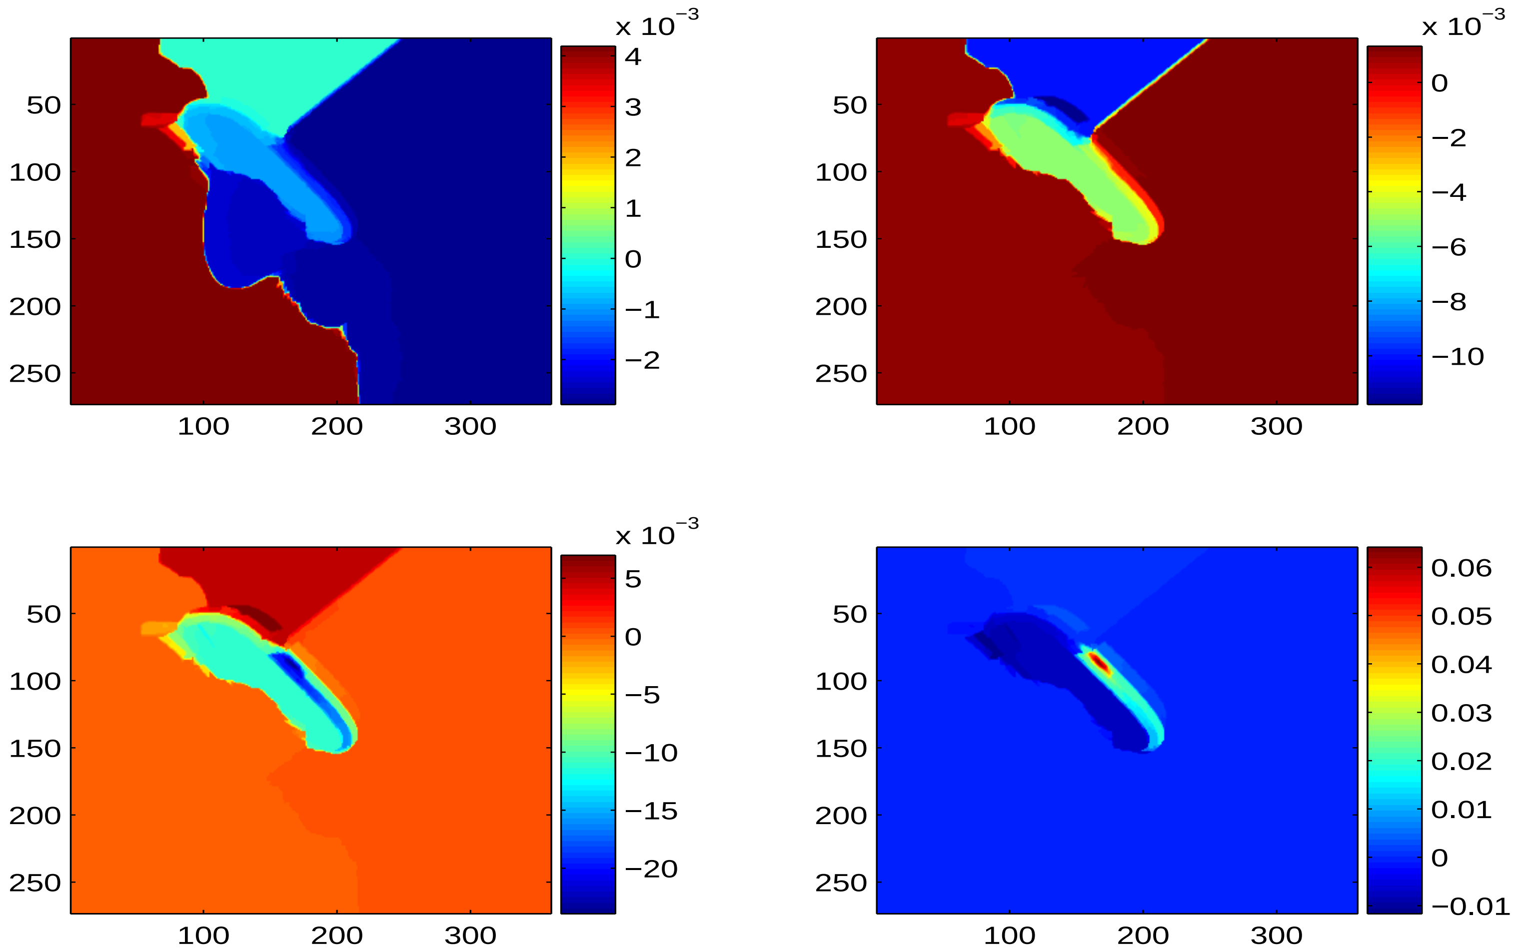
x 10^-3
4
3
2
1
0
−1
−2
50
100
150
200
250
100
200
300
x 10^-3
0
−2
−4
−6
−8
−10
x 10^-3
5
0
−5
−10
−15
−20
0.06
0.05
0.04
0.03
0.02
0.01
0
−0.01

LOBPCG

Locally Optimal Block Preconditioned Conjugate Gradient Method

The algorithm is described in:

- A. V. Knyazev, Toward the Optimal Preconditioned Eigensolver: Locally Optimal Block Preconditioned Conjugate Gradient Method. SIAM Journal on Scientific Computing 23 (2001), no. 2, pp. 517-541.
- I. Lashuk, M. E. Argentati, E. Ovchinnikov and A. V. Knyazev, Preconditioned Eigensolver LOBPCG in hypre and PETSc. Proceedings of the 16th International Conference on Domain Decomposition Methods, (2005). To appear in Lecture Notes in Computational Science and Engineering, Springer.

BLOPEX–LOBPCG Software

The LOBPCG is publicly available in MATLAB, see
`http://math.cudenver.edu/~aknyazev/software/CG/`
The LOBPCG C version using MPI libraries for massively parallel computers is publicly available within our Block Locally Optimal Preconditioned Eigenvalue Xolvers (BLOPEX) package. BLOPEX is built-in in Hypre, see http://www.llnl.gov/CASC/hypre/, and is included as an external package in PETSc, see http://www-unix.mcs.anl.gov/petsc/.

Some currently available LOBPCG software by others

- Earth Simulator CDIR/MPI (Yamada et al., Fermion-Habbard Model)
- SLEPc interface to Hypre LOBPCG (Jose Roman, SLEPc)
- C++ (Rich Lehoucq and Ulrich Hetmaniuk, Anasazi Trilinos)
- C (A.Stathopoulos/O.Marques, PRIMME, real/complex Hermitian)
- Fortran 77 (Randolph Bank, PLTMG)
- Python (Peter Arbenz and Roman Geus, PYFEMax)
- C++ (Sabine Zaglmayr and Joachim Schberl, NGSolve)
- Fortran 90 (Gilles Zèrah, ABINIT, complex Hermitian)
- Fortran 90 (S. Tomov and J. Langou, PESCAN, complex Hermitian)
- (A. Borzì and G. Borzì, AMG)

Eigenvalue Solvers with Multiscale Preconditioning

With an appropriate selection of a preconditioning technique, preconditioned eigensolvers, such as LOBPCG, can converge almost as as fast as the shift–and–invert Lanczos method at the computational costs per iteration proportional to the number of unknowns. The overall complexity is linear, in other words, the computational costs scale linearly with the problem size

Multiscale preconditioning (geometric and algebraic) is known to be efficient for PDEs. Would it work for graph, in particular image, Laplacian?

Numerical Results with Hypre AMG preconditioning

The MATLAB Graph Analysis Toolbox by Leo Grady is available for download at http://www.mathworks.com and provides the image-matrix interface for the spectral partitioning.

Direct solvers in the shift-and-invert approach turn out to be practically efficient for 2D images that are not too big. The main limitation is the memory for the factors. Here, we investigate Hypre algebraic and geometric (structured) multigrid preconditioning, and run tests on massively parallel computers, using our LOBPCG C code, which is a part of the software library Block Locally Optimal Eigenvalue Xolvers (BLOPEX). On BlueGene/L 1024 CPU we can compute the Fiedler vector of a 24 megapixel image in seconds (including the algebraic multigrid setup).

Scalability of BLOPEX-AMG on IBM BlueGene/L

N Proc	N Iter.	Prec. setup (sec)	Apply Prec. (sec)	Lin. Alg. (sec)
32	12	66	30	6
64	14	32	18	3
128	12	18	8	1
256	12	10	4	0.5
512	21	5.4	4	0.5
1024	13	4	2	0.2

TABLE 1: **Scalability Data for 24 megapixel image segmentation**

BLOPEX in PETSc using Hypre AMG, Block size: 1, LOBPCG tolerance: 10^{-6}.

NCAR's single-rack Blue Gene/L with 1024 compute nodes, organized in 32 I/O nodes with 32 compute nodes each. One node is a dual-core chip, containing two 700MHz PowerPC-440 CPUs and 512MB of memory. We run 1 CPU per node.

Scalability of BLOPEX-SMG on IBM BlueGene/L

N Proc	Matrix Size	N Iter.	Prec. Setup (sec)	Solve (sec)
8	4.096 M	10	7	74
64	32.768 M	8	11	67
512	0.262144 B	7	19	61

TABLE 2: *Scalability for 3D Laplacian* $80 \times 80 \times 80 = 512,000$ *mesh per CPU*

BLOPEX in Hypre struct with SMG, Block size: 1, LOBPCG tolerance: 10^{-8}. Uniform cube partitioning. 1 CPU per node.

SMG is not yet adapted for image Laplacian.

Conclusion

1. LOBPCG may be a valuable alternative to the classical Lanczos method for spectral graph partitioning and image segmentation.
2. LOBPCG in BLOPEX is the only currently available package that solves eigenvalue problems using Hypre and PETSc preconditioner.
3. Multiscale preconditioning of Hypre seems to work well for image segmentation.
4. Eigenvalue problems of record sizes can be solved with BLOPEX using Hypre multiscale preconditioning.