

UCLA
COMPUTATIONAL AND APPLIED MATHEMATICS

Multilevel Spectral Partitioning of Unstructured Grids

Tony F. Chan
Susie Go
Jun Zou

December 1998
CAM Report 98-52

Department of Mathematics
University of California, Los Angeles
Los Angeles, CA. 90095-1555

Multilevel Spectral Partitioning of Unstructured Grids

Tony F. Chan * Susie Go † Jun Zou ‡

October 30, 1998

1 Introduction

The graph partitioning problem is an important component of parallel computing (e.g. for constructing subdomains in domain decomposition methods) and as a result, many partitioning methods and associated sophisticated software packages have been developed recently including algorithms based on greedy, coordinate, inertial, multilevel spectral and graph bisection. However, there is usually an unavoidable tradeoff between quality and speed. Of these, spectral partitioning has the reputation of being the most time consuming, but it also consistently produces high-quality partitions which for some applications can offset its costliness with better convergence rates of iterative solvers. We will show that when the graph is a standard finite element mesh, the graph Laplacian is spectrally equivalent (up to a diagonal scaling) to the mesh Laplacian. This equivalence can be exploited to adapt recently developed multilevel elliptic algorithms for unstructured grids to solving the graph partitioning problem with a true multigrid (MG) convergence rate. We also give some preliminary results of such an approach.

In Section 2, the spectral partitioning algorithm is briefly discussed. The main cost of the spectral method is in solving for the Fiedler vector (the eigenvector corresponding to the lowest non-trivial eigenvalue of the associated graph Laplacian). Attempts at accelerating the computation of the Fiedler vector by multilevel methods encounters the difficulty of adapting standard elliptic multilevel algorithms to solving the discrete graph Laplacian eigenvalue problem.

*Department of Mathematics, University of California at Los Angeles, Los Angeles, CA 90095-1555, USA. E-mail: chan@math.ucla.edu. The work of this author was partially supported by grants from NSF under grant number ASC-9720257 and from Sandia National Laboratory under LG 4440.

†Research Institute of Advanced Computer Science (RIACS), Mail Stop 19-39, NASA Ames Research Center, Moffett Field, CA 94035-1000, USA. E-mail: susie@riacs.edu. Support from RIACS under contract number NCC2-1006.

‡Department of Mathematics, The Chinese University of Hong Kong, Shatin, Hong Kong. E-mail: zou@math.cuhk.edu.hk. The work of this author was partially supported by Hong Kong RGC Grant No. CUHK 338/96E and CUHK 4004/98P.

Many approaches rely on piecewise constant interpolations to develop multilevel algorithms. Thus, it is difficult to construct algorithms which are optimal for calculating solutions involving the Laplacian operator.

If the graph is a grid, we can use the spectral equivalence of the mesh and graph Laplacian to develop efficient ways to find the Fiedler vector. There exist techniques for dealing with the mesh Laplacian on unstructured grids; we can now apply them to the case of the graph Laplacian. One possibility is described in Section 3. A multigrid approach for solving elliptic eigenvalue problems is adapted for use on unstructured grids and on discrete problems. We give some numerical results which show that the algorithm, when appropriately defined, yields optimal convergence rates for finding the Fiedler vector.

2 Spectral equivalence of the graph and mesh Laplacian

We briefly define the spectral partitioning method for bisecting a graph. For a more complete description of partitioning algorithms, we refer the reader to Pothén's survey [9]. Let $G = (V, E)$ be an undirected graph, where $V = \{v_i\}_{i=1}^m$ is the set of vertices of the graph, and $E = \{(v_i, v_j)\}$ is the set of edges. An associated matrix called the graph Laplacian, L , is defined as:

$$L_{ij} = \begin{cases} \deg(v_i) & \text{if } i = j \\ -1 & \text{if } (v_i, v_j) \in E \\ 0 & \text{otherwise,} \end{cases} \quad (1)$$

where $\deg(v_i)$ is the number of edges in E containing vertex v_i .

The spectral bisection method works by solving the discrete Laplacian eigenvalue problem, $Lu = \lambda u$, for the eigenvector corresponding to the second smallest eigenvalue, known as the Fiedler vector. One way of defining a bisection of the graph is found by sorting the components of the resulting Fiedler vector and defining the median value of the components as the so-called splitting value. This median cut partitioning was shown to be optimal in [3]. A bisection of the graph is found by assigning those vertices whose value is less than the splitting value to one subgraph and assigning the rest to the other subgraph.

We now show that for standard finite element meshes, the graph Laplacian is spectrally equivalent, up to a diagonal scaling, to the mesh Laplacian. This equivalence can then be used to show that a multilevel elliptic eigenvalue solver can be used to solve the graph partitioning problem with true multigrid convergence rates.

Assume \mathcal{T}^h is a shape regular finite element triangulation of a domain Ω in R^n ($n = 2$ or 3) and $G = (V, E)$ is the graph induced from \mathcal{T}^h . Let L be the Laplacian of the graph G and A be the stiffness matrix arising from the finite element discretization of the continuous Laplacian operator Δ on the triangulation \mathcal{T}^h .

Theorem 2.1 *The graph Laplacian L is spectrally equivalent to the stiffness matrix defined by the quadratic form:*

$$\sum_{K \in \mathcal{T}^h} h_K^{2-n} \int_K |\nabla u|^2 dx. \quad (2)$$

Proof. Let K be any triangle ($n = 2$) or tetrahedron ($n = 3$) in \mathcal{T}^h , h_K be the diameter of K , and \hat{K} be the reference element with $n + 1$ vertices \hat{x}_i , for $i = 1, \dots, n + 1$.

Let $F : \hat{K} \rightarrow K$ be the affine mapping defined by $F(\hat{x}) = B\hat{x} + b$ such that the vertices $\hat{x}_i, i = 1, \dots, n + 1$ of the element \hat{K} are mapped to the vertices $x_i, i = 1, \dots, n + 1$ of the element K . Let u be any linear function on K , with nodal values u_i at x_i . Clearly, \hat{u} has the same values at its vertices as u . Moreover, we have (cf. [6])

$$\|\nabla u\|_{L^2(K)}^2 \leq C |\det(B)| \|B^{-1}\|^2 \|\nabla \hat{u}\|_{L^2(\hat{K})}^2 \quad (3)$$

where $\|\cdot\|$ is the spectral norm of a matrix.

By the definition of the matrix norm and the shape regularity assumption, we can derive the following bounds [6]

$$\|B^{-1}\| \leq h_K^{-1}, \quad \|B\| \leq Ch_K, \quad (4)$$

and

$$|\det(B)| \leq \text{meas}(K) \leq Ch_K^n, \quad |\det(B^{-1})| \leq Ch_K^{-n}. \quad (5)$$

Then from (3), it follows that

$$\|\nabla u\|_{L^2(K)}^2 \leq Ch_K^{n-2} \|\nabla \hat{u}\|_{L^2(\hat{K})}^2. \quad (6)$$

Now consider the finite dimensional quotient space $P_1(\hat{K})/R$. It is easy to check that

$$\|\hat{u}\| \stackrel{\text{def}}{=} \left(\sum_{\hat{x}_i, \hat{x}_j \in \hat{K}} (\hat{u}(\hat{x}_i) - \hat{u}(\hat{x}_j))^2 \right)^{1/2}$$

is a norm on $P_1(\hat{K})/R$. We also know that $\|\nabla \hat{u}\|_{L^2(\hat{K})}^2$ is a norm in $P_1(\hat{K})/R$. Therefore, they are equivalent and from (6), this implies

$$h_K^{2-n} \|\nabla u\|_{L^2(K)}^2 \leq C \|\nabla \hat{u}\|_{L^2(\hat{K})}^2 \leq C \sum_{x_i, x_j \in K} (u(x_i) - u(x_j))^2 \quad (7)$$

since $\hat{u}(\hat{x}_i) = u(x_i)$.

Similarly as (3), we have

$$\|\nabla \hat{u}\|_{L^2(\hat{K})}^2 \leq C |\det(B^{-1})| \|B\|^2 \|\nabla u\|_{L^2(K)}^2 \leq Ch_K^{2-n} \|\nabla u\|_{L^2(K)}^2, \quad (8)$$

combining this with (7) gives

$$h_K^{2-n} \|\nabla u\|_{L^2(K)}^2 \leq \sum_{x_i, x_j \in K} (u(x_i) - u(x_j))^2 \leq h_K^{2-n} \|\nabla u\|_{L^2(K)}^2. \quad (9)$$

From (9), we obtain

$$\sum_{K \in \mathcal{T}^h} h_K^{2-n} \int_K |\nabla u|^2 dx \simeq \sum_{i,j \in E} (u_i - u_j)^2 \quad (10)$$

where each edge is just counted once.

But note that $u^T L u = \sum_{(i,j) \in E} (u_i - u_j)^2$, hence we have proved the equivalence.

3 Full approximation schedule multigrid algorithm for elliptic eigenvalue problems

The spectral equivalence of the graph Laplacian and the mesh Laplacian motivates us to use the same tools for solving the discrete eigenvalue problem of spectral partitioning as we use for a continuous elliptic eigenproblem. Two possible alternatives for using the mesh Laplacian to solve for Fiedler vector would be to:

1. Use the mesh Laplacian A to precondition the graph Laplacian L in some iterative method (e.g. inverse iteration),
2. Use a multigrid method directly by applying the same interpolants of the mesh problem to define a hierarchy of coarse problems for the associated discrete problem.

We adapt an eigensolver which uses the full approximation scheme (FAS) multigrid method [2] to develop a multigrid partitioner. The multigrid eigensolver treats the eigenvalue problem as a non-linear problem, thus FAS is used. The two-level algorithm for solving for the smallest eigenvalue and corresponding eigenvector of $Au_h - \lambda u_h = 0, \|u_h\| = 1$ is given by:

Algorithm 3.1 *Two-level FAS multigrid:*

1. Given an initial guess, λ^0 and u^0 ,
2. Do one V-cycle FAS multigrid step:
 - (a) Presmooths.
 - (b) Restrict solution and calculate f_H .
 - (c) Solve coarse problem $A_H u_H - \lambda u_H = f_H$.
 - (d) Update λ using the Rayleigh-quotient modified for the extended eigenvalue problem and normalize vector.

- (e) FAS interpolate solution to fine level.
- (f) Postsmooths.

The method can be extended recursively to yield a multigrid algorithm. To solve for multiple eigenvalues/eigenvectors, Ritz acceleration can be added to the method after all the values and vectors desired have been found. The algorithm for finding q vector/value pairs:

Algorithm 3.2 *Multiple eigenvector/value solver:*

1. Given initial guesses, $(\vec{u}_j, \lambda_j), j = 1, \dots, q,$
2. While not converged,
 - (a) For $i = 1,$ to q
 - i. Call FAS multigrid to solve for $(\vec{u}_i, \lambda_i).$
 - (b) Ritz orthogonalize the set of q vector/value pairs.

We can use all the same machinery as the continuous elliptic eigenvalue solver: the same FAS interpolation and the same coarse grid solution process. Since spectral partitioning requires computing the Fiedler vector, the FAS MG method is used to solve the discrete problem, but modified so as to seek only the second eigenvector, since the first eigenvector/value pair is known to be $(\vec{1}, 0).$

We use a grid hierarchy which is generated via maximal independent sets of the fine grid and then retriangulated to create coarser grids. Interpolation operators are defined to be piecewise linear in the regions where both fine and coarse grids overlap. In the case of non-matching boundaries, they can be defined to extend by zero where the coarse grid exists but the fine grid does not. Where the coarse grid does not cover the fine grid, they must provide a non-zero extension, for example, by linear interpolation with respect to a nearby edge or element, or by modifying the coarse grid so that it covers the fine grid (we refer the reader to [4] for more details).

Figure 1 shows the FAS MG convergence histories for solving the discrete Laplacian eigenvalue problem $Lu = \lambda u$ on unstructured grids of varying mesh sizes. It can be observed that the method is independent of mesh size.

Next, we show a comparison for solving the spectral bisection partitioning problem using the FAS MG scheme with various interpolants. Since the spectral bisection problem has Neumann boundary conditions, care must be taken in the definition of interpolation operators at non-matching boundaries (see [4, 5] for details). We observe that using a standard linear interpolant with zero extension cannot be accurate enough to define Neumann boundary conditions when non-matching boundaries exist. In addition, note that piecewise constant interpolants also cannot be used.

The edges cut and times required for bisecting the airfoil using various popular partitioners [1, 7, 8] are compared in Table 1. Options for the partitioners were set to generate a spectral bisection to an eigentolerance of $10^{-6},$ with local refinement turned off. The bisection found by Metis has a different number of

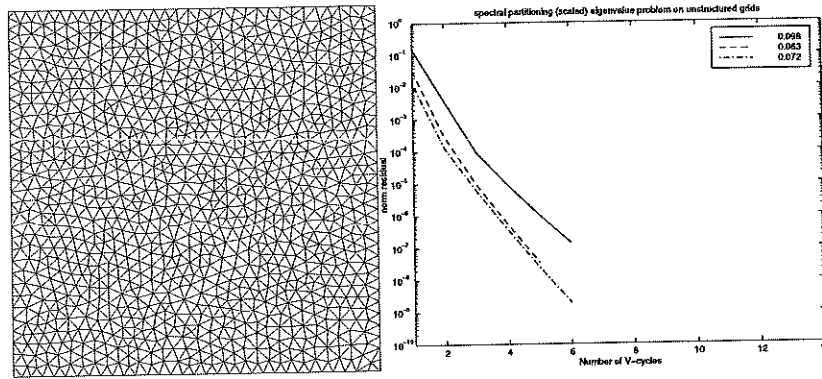


Figure 1: FAS MG for the spectral bisection problem on the unit square with unstructured grids of varying mesh sizes. Solid line (310 nodes), dashed line (1064 nodes), and dash-dotted line (4164 nodes). The grid with 1064 nodes is shown on the left.

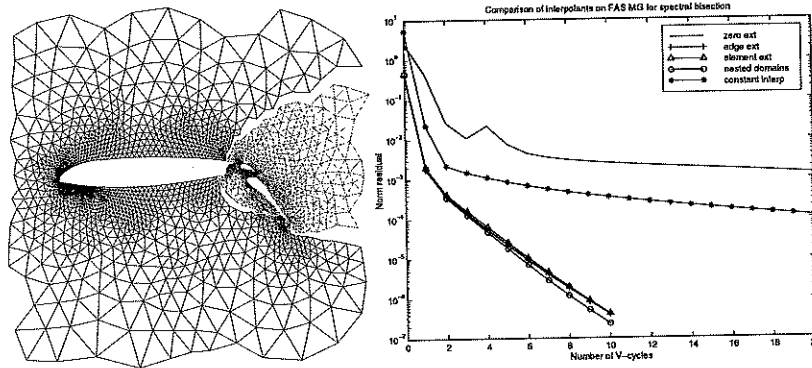


Figure 2: 4-level FAS MG convergence histories for the spectral bisection problem on an unstructured airfoil grid (left) using different interpolants.

cut edges because spectral bisection is used to generate a coarse bisection which is then “uncoarsened” to yield a bisection of the original graph. Depending on the type of maximal matching chosen for the coarsening step, the number of cut edges ranged from 114 to 153 when the original airfoil grid was coarsened down to 100 nodes.

We should comment that the current FAS MG implementation was not optimized in any way and that the main objective was in achieving provable optimal multigrid performance, that is, obtaining a solution process which was grid-size independent. It was also assumed that the interpolation operators are immediately available and the timings do not include the time to generate the Galerkin coarse operators.

Finally, Table 2 shows the relationship between the tolerance and the number of V-cycles used, as well as the number of edges cut. It should be noted that there is little or no advantage to solving the partitioning problem to extremely low tolerance. This was also observed previously, by the authors of Chaco [7], where they recommend an eigentolerance of between 10^{-3} and 10^{-6} in most cases, for generating sufficiently accurate approximations for partitions using spectral methods.

Table 1: Performance of various partitioners implementing spectral bisection on an airfoil grid (4253 nodes).

Partitioner	# of Edges cut (%)	Time (sec)
Chaco	132 (1.07%)	12.34
FAS MG	132 (1.07%)	7.98
Inverse iteration	132 (1.07%)	5.60
Metis	122 (0.99%)	0.59

Table 2: Spectral bisection using 4-level FAS MG on an airfoil grid.

Reduct. residual	# V-cycles	# edges cut
10^{-0}	0	191
10^{-1}	1	106
10^{-2}	1	106
10^{-3}	2	138
10^{-4}	5	136
10^{-5}	7	132
10^{-6}	10	132

References

- [1] S.T. Barnard and H.D. Simon. A fast multilevel implementation of recursive spectral bisection. *Proc. 6th SIAM Conf. Parallel Proc. for Sci. Comp.*, pages 711–718, 1993.
- [2] A. Brandt, S. McCormick, and J. Ruge. Multigrid methods for differential eigenproblems. *SIAM J. Sci. Stat. Comput.*, 4(2):244–260, 1983.
- [3] T. F. Chan, P. Ciarlet, Jr., and W. K. Szeto. On the optimality of the median cut spectral bisection graph partitioning method. *SIAM J. Sci. Comput.*, 18(3):943–948, 1997.
- [4] T. F. Chan, S. Go, and J. Zou. Boundary treatments of multilevel methods on unstructured meshes. Technical Report CAM 96–30, Department of Mathematics, University of California at Los Angeles, September 1996. Accepted for publication in *SIAM J. Sci. Comput.*

- [5] T. F. Chan, B. Smith, and J. Zou. Overlapping Schwarz methods on unstructured meshes using non-matching coarse grids. *Numer. Math.*, 73(2):149–167, April 1996.
- [6] P. Ciarlet. *The Finite Element Method for Elliptic Problems*. North-Holland, 1978.
- [7] B. Hendrickson and R. Leland. The Chaco User’s Guide Version 2.0. Technical Report SAND94-2692, Sandia Natl. Lab., Albuquerque, N.M., July. 1995.
- [8] G. Karypis and V. Kumar. Metis: A software package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices, 1997. available at <http://www-users.cs.umn.edu/~karypis/metis/metis.html>.
- [9] A. Pothen. Graph partitioning algorithms with applications to scientific computing. In D. E. Keyes, A. H. Sameh, and V. Venkatakrisnan, editors, *In Parallel Numerical Algorithms*. Kluwer Academic Press, 1996.