Two-level Algebraic Multigrid for the Helmholtz Problem

Petr Vaněk, Jan Mandel, and Marian Brezina

1. Introduction

An algebraic multigrid method with two levels is applied to the solution of the Helmholtz equation in a first order least squares formulation, discretized by Q1 finite elements with reduced integration. Smoothed plane waves in a fixed set of directions are used as coarse level basis functions. The method is used to investigate numerically the sensitivity of the scattering problem to a change of the shape of the scatterer.

Multigrid methods for the solution of the Helmholtz equation of scattering are known in the literature. A common disadvantage of multigrid methods is that the coarsest level must be fine enough to capture the wave character of the problem, or the iterations diverge [6, 7, 10]. One way to overcome this limitation is to use coarse basis functions derived from plane waves [8]. However, manipulating such functions becomes expensive since the cost does not decrease with the number of variables on the coarse levels. It should be noted that functions derived from plane waves can also be used as basis functions for the discretization of the Helmholtz equation itself; such methods are known under the names of the Microlocal Discretization [5], Partition of Unity Finite Element Method [1], or the Finite Element Ray Method [11].

We propose a two-level method with coarse space basis functions defined as a plane waves within an aggregate of nodes, zero outside the aggregate, and then smoothed by a Chebyshev type iteration using the original fine level matrix. This results in a method with good computational complexity and scalability. This method falls under the abstract framework of black-box two-level iterative methods based on the concept of smoothed aggregations [16]. The objective of the smoothing of the coarse basis functions is to reduce their energy [9, 15, 16]. For a related theoretical analysis of such two-level methods with high order polynomial

smoothing of coarse basis functions, see [3], where a convergence result uniform both with respect to coarse and fine level meshsize was proved for second order eliptic problems discretized on unstructured meshes.

2. Problem Formulation and Discretization

In this section, we describe the first order least squares formulation and the discretization of the two-dimensional scattering problem on a bounded domain.

Denote by H^{div} the space of all vector functions with divergence in L^2. Let $\mathcal{O} \subset \Omega$ be an obstacle in a domain $\Omega \subset \mathbb{R}^2$ that is sufficiently large with respect to the size of $\mathcal{O}$. To be more precise, Ω is assumed large enough to allow the scatteded field to be almost radial near $\partial\Omega$.

We seek a complex pressure $p \in H^1(\Omega \setminus \mathcal{O})$ and its (complex) gradient $u \in H^{div}(\Omega \setminus \mathcal{O})$ minimizing the convex functional

(1)

$$F(p, \mathbf{u}) = w_1 \int_{\Omega \setminus \mathcal{O}} \|\nabla p - \mathbf{u}\|^2 + w_2 \int_{\Omega \setminus \mathcal{O}} |\text{div}(\mathbf{u}) + k^2 p|^2 + w_3 \int_{\partial\Omega} |\mathbf{u} \cdot \mathbf{r} - ikp|^2,$$

where w_1, w_2, w_3 are positive constants and $\mathbf{r} \in \mathbb{R}^2$ is the normalized radiusvector of the point $\mathbf{x} \in \partial\Omega$, i.e. $\mathbf{r} = \mathbf{x}/\|\mathbf{x}\|$. The pressure p is subject to the Dirichlet boundary condition on the boundary of the obstacle,

$$p(\mathbf{x}) = -\exp(ik\mathbf{d} \cdot \mathbf{x}/\|\mathbf{d}\|), \qquad \mathbf{x} \in \partial\mathcal{O},$$

where $\mathbf{d} \in \mathbb{R}^2$ is the direction of the incident wave. The first two integrals are a first order least square formulation corresponding to the Helmholtz equation

$$\Delta p + k^2 p = 0,$$

and the boundary integral enforces the radiation boundary condition

$$\frac{\partial p}{\partial \mathbf{r}} = ikp \quad \text{on } \partial\Omega.$$

Hence, all the integrals in (1) vanish for the minimizer of F and the solution of the minimization problem (1) is independent of the weights w_1, w_2, w_3. We choose the uniform Q1 finite elements for the discretization of the continuous minimization problem, and minimize the functional over the finite element space. As the derivatives of the gradient of a Q1-function are not Q1-functions themselves, the discretization of the integrals in (1) creates an undesirable "artificial viscosity" resulting in a damping of the numerical solution. In fact, it is easy to see on a one-dimensional example, that the restriction of a plane wave on the mesh is not a solution of the discrete problem. To avoid this, the volume integrals have been discretized by the one point quadrature formula with the node in the middle of each element. This restores the property that the discrete system allows plane waves as solutions. Such *inexact integration* is frequently used to eliminate locking caused by similar integrals in plate and shell finite elements. Since $u = \nabla p$ satisfies $\nabla \times u = 0$, adding the term

(2)

$$\int_{\Omega \setminus \mathcal{O}} |\nabla \times u|^2$$

to the functional F in (1), as suggested in [4, 8], does not change the solution and makes the functional F coercive in certain cases.

In the discrete case, adding the integral (2) to F changes the discrete solution and may make it more accurate. We have not found an advantage to adding the integral (2) in our experiments. The results with reduced integration alone were satisfactory. Also, in the discrete case, the solution is no longer independent of the weights w_1, w_2, w_3. Based on our computations, we found no significant advantage to other than unit weights, which are used in all computations reported here.

The discretization described above results in a Hermitian matrix with 3 complex degrees of freedom per node. Because we wanted to take advantage of an existing real code, our implementation treated the real and imaginary parts of the solution as real unknowns, resulting in a real symmetric problem with 6 real degrees of freedom per node. Further efficiency could be gained by an implementation in the complex arithmetics.

3. Algebraic Multigrid

In this section, we describe the algebraic method used for solving the discretized scattering problem. It is a variant of the method introduced in [3, 16]. We will present the method as a variant of the multiplicative Schwarz method [2, 12], and write it in terms of matrices.

3.1. The Multiplicative Schwarz Method.
Let A be a symmetric, positive definite $n \times n$ matrix, and N_j, $j = 0, \ldots m$, be full rank matrices with n rows. Consider the following iterative method for the solution of the linear algebraic system $A\mathbf{x} = \mathbf{f}$:

$$(3) \qquad \mathbf{z} \quad \leftarrow \quad \mathbf{x}^i$$

$$(4) \qquad \mathbf{z} \quad \leftarrow \quad \mathbf{z} + N_i(N_i^T A N_i)^{-1} N_i^T(\mathbf{f} - A\mathbf{z}), \qquad i = 1, \ldots, m$$

$$(5) \qquad \mathbf{z} \quad \leftarrow \quad \mathbf{z} + N_0(N_0^T A N_0)^{-1} N_0^T(\mathbf{f} - A\mathbf{z}),$$

$$(6) \qquad \mathbf{z} \quad \leftarrow \quad \mathbf{z} + N_i(N_i^T A N_i)^{-1} N_i^T(\mathbf{f} - A\mathbf{z}), \qquad i = m, \ldots, 1$$

$$(7) \qquad \mathbf{x}^{i+1} \quad \leftarrow \quad \mathbf{z}$$

Since (3)-(7) is a consistent stationary iterative method for the system $A\mathbf{x} = \mathbf{f}$, it can be written as $\mathbf{x}^{i+1} = \mathbf{x}^i + N(\mathbf{f} - A\mathbf{x}^i)$. We use the operator N, that is, the output of (3)-(7) with $\mathbf{x}^i = 0$, as a preconditioner in the method of conjugate gradients. The operator N is symmetric and positive definite, since [2, 12]

$$(8) \qquad NA = (I - \Pi_1) \cdots (I - \Pi_m)(I - \Pi_0)(I - \Pi_m) \cdots (I - \Pi_1),$$

where Π_i is the A-orthogonal projection onto the range of N_i.

3.2. Two-level Algebraic Multigrid by Smoothed Aggregation.
It remains to specify the matrices N_i. We first construct the matrix N_0. This matrix is denoted $N_0 = P$ and called the *prolongator*. The prolongator is constructed in two steps. In the first step, we construct a *tentative prolongator* capturing precisely a selected set of functions (in the same sense as, for example, P1 finite elements capture linear functions). In the second step, we suppress high-energy components in the range of the tentative prolongator by smoothing it using a proper *prolongator smoother*. The matrices N_i, $i \neq 0$, are injections that correspond to overlapping blocks of unknowns. The blocks are derived from the nonzero structure of the smoothed prolongator. As a prolongator smoother we use a properly chosen polynomial in the stiffness matrix A. The degree of the prolongator smoother

determines the smoothness of the coarse space as well as the amount of overlaps of the blocks in the overlapping Schwarz method.

To construct the prolongator P, we first decompose the set of all nodes, where an essential boundary condition is not imposed, into a disjoint covering

$$(9) \qquad \{1, \ldots, n\} = \bigcup_{i=1}^{m} \mathcal{A}_i, \quad \mathcal{A}_i \cap \mathcal{A}_j = \emptyset \text{ for } i \neq j.$$

We use a simple greedy algorithm that chooses aggregates of nodes that are connected via nonzero terms of A, whenever possible, cf. [13, 16].

Let us consider the set of functions $\{f^i\}_{i=1}^{n_k}$ we want to be captured by the coarse space functions. These functions should approximate the kernel of the constrained problem well (e.g., constants in the case of Poisson equation, 6 rigid body modes in the case of 3D elasticity). For solving the scattering problem here, our choice of the set of functions $\{f^i\}_{i=1}^{n_k}$ are the plane waves

$$f_{\mathbf{d}}(\mathbf{x}) = e^{-ik\mathbf{x}\cdot\mathbf{d}/\|\mathbf{d}\|},$$

where $\mathbf{x} \in \mathbb{R}^2$ is the position and $\mathbf{d} \in \mathbb{R}^2$ is the direction of the plane wave. We choose a finite subset of plane waves; the numerical experiments in this paper are performed with a coarse space built from plane waves in the $n_k = 8$ directions $\mathbf{d} = (1,0)$, $(-1,0)$, $(0,1)$, $(0,-1)$, $(1,1)$, $(-1,-1)$, $(-1,1)$, $(1,-1)$. Since plane waves are not contained in the fine level space exactly, the vectors $\{\hat{\mathbf{f}}^i\}_{i=1}^{n_k}$ are constructed as grid interpolations of the functions $f_{\mathbf{d}}$ and their gradients.

For each function f^i, let $\hat{\mathbf{f}}^i$ be its discrete representation in the finite element basis. Each node has 6 degrees of freedom, namely, the real and imaginary parts of p, $\partial_x p$, and $\partial_y p$. The decomposition (9) induces a decomposition of the degrees of freedom into disjoint sets

$$\{1, \ldots, n_d\} = \bigcup_{i=1}^{m} \mathcal{D}_i, \qquad \mathcal{D}_i \cap \mathcal{D}_j = \emptyset \text{ for } i \neq j,$$

where $\mathcal{D}_i$ is the set of all degrees of freedom associated with the nodes of $\mathcal{A}_i$. We then construct a *tentative prolongator* as the block matrix

$$(10) \qquad \hat{P} = [\hat{\mathbf{p}}_1, \ldots, \hat{\mathbf{p}}_m]$$

where the j-th row of the block column $\hat{\mathbf{p}}_i$ equals to the j-th row of $\hat{\mathbf{f}}^i$ if $j \in \mathcal{D}_i$, and is zero otherwise. The prolongator P is then defined by

$$P = s(A)\hat{P},$$

where s is the polynomial of given degree d such that $s(0) = 1$ and $\max_{0 \leq \lambda \leq \hat{\rho}} |s(\lambda)^2 \lambda|$ is minimal, with $\hat{\rho}$ an easily computable upper bound on the spectral radius of A. It is easy to show [14] that p is a shifted and scaled Chebyshev polynomial, equal to

$$s(\lambda) = \left(1 - \frac{\lambda}{r_1}\right) \cdots \left(1 - \frac{\lambda}{r_d}\right), \qquad r_k = \frac{\hat{\rho}}{2}\left(1 - \cos\frac{2k\pi}{2n+1}\right).$$

This construction of P attempts to minimize the energy of the columns of P. Indeed, any column of the prolongator P is $s(A)\hat{\mathbf{p}}$, where $\hat{\mathbf{p}}$ is a column of the tentative prolongator $\hat{P}$, and its squared energy norm is

$$\|s(A)\hat{\mathbf{p}}\|_A^2 = (s(A)\hat{\mathbf{p}})^T A(s(A)\hat{\mathbf{p}}) = \hat{\mathbf{p}}^T s^2(A)A\hat{\mathbf{p}}.$$

TABLE 1. Model problem for performance experiments

Computational domain:	$[-2, 2] \times [-2, 2]$
Obstacle:	$[-0.3, 0.3] \times [-0.3, -0.3]$
Dir. of the inc. wave:	$\mathbf{d} = (1, 1)$
k	varying
Mesh:	regular 400×400 square mesh
Num. of dofs:	964,806

If the columns of $\hat{P}$ have bounded euclidean norm (which they do here), the construction above gives an optimal bound on the energy of the columns of P, uniform in the choice of $\hat{P}$. See [**3, 14**] for more details, and [**9**] for a more direct approach to minimizing the energy of the columns of P.

The iteration (5) has now the interpretation of a *coarse grid correction*, in terms of multigrid methods. It remains to choose the matrices N_i in the *pre-smoothing* (4) and *post-smoothing* (6). Our choice is equivalent to a multiplicative overlapping Schwarz method. The overlapping blocks are derived from the nonzero structure of the prolongator P by choosing N_i to consist of those columns j of the n_d by n_d identity matrix, for which the j-the row of the block column $\mathbf{p}_i$ from (10) is not zero. The iterations (4) and (6) are now simply block relaxation with overlapping blocks, with block i consisting of all variables that may become nonzero in the range of the prolongator block column $\mathbf{p}_i$.

3.3. Parallel Implementation. The smoothing iterations (4) and (6) are parallelized using a generalization of the well known red-black ordering for Gauss-Seidel iteration. First, observe that if $N_i^T A N_j = 0$, then the projections Π_i and Π_j from (8) commute, and so the iteration steps i and j in (4) or (6) are independent and can be performed concurrently.

The method proceeds as follows. In the setup phase, a coloring of the adjacency graph of the matrix $(N_i^T A N_j)_{i,j}$ is found by a simple greedy algorithm. (Of course, the graph is constructed by performing symbolic matrix multiplication only.) Then $N_i^T A N_j = 0$ for all indices i and j assigned the same color. The smothing iteration (4) is then reordered so that all iteration steps with the same color are done concurrently. Then the pre-smoothing (4) becomes

$$\mathbf{z} \leftarrow \mathbf{z} + \left(\sum_{i \in \mathcal{C}_k} N_i (N_i^T A N_i)^{-1} N_i^T \right)(\mathbf{f} - A\mathbf{z}), \qquad k = 1, \dots, n_c,$$

where n_c the number of colors and $\mathcal{C}_k$ is the set of all indices of color k. The post-smoothing (6) is same except that the colors are processed in the reverse order.

4. Performance Results

All experiments reported in this section have been carried out on a SGI ORIGIN 2000 with 64 R10000 processors and 4GB of memory. As the stopping condition, we have used

$$\langle A P \mathbf{r}^i, A \mathbf{r}^i \rangle^{1/2} \leq \frac{10^{-5}}{\text{cond}} \langle A P \mathbf{r}^0, A \mathbf{r}^0 \rangle^{1/2},$$

where P is the preconditioner, cond is a condition number estimate computed from the conjugate gradient coefficients, and $\mathbf{r}^i$ is the residual after the i-th iteration. In order to illustrate the effect of this stopping condition, we provide achieved relative

TABLE 2. Performance of the method with varying k. H denotes the side of the aggregates (squares). Mesh 400×400 elements. Number of degrees of freedom 964,806. Number of processors 16. $H/h = 30$.

k	els./wave length	H/wave length	cond.	achieved rel. residual	setup time [s] CPU/WALL	iter. time [s] CPU/WALL
251.327	5	6.0	6.565	2.447×10^{-6}	49/64	115/131
125.663	10	3,0	15.63	1.670×10^{-6}	48/63	117/118
83.7758	15	2.0	31.67	1.729×10^{-6}	52/71	169/171
62.8318	20	1.5	60.70	1.112×10^{-6}	53/73	243/246
50.2654	25	1.2	103.9	8.939×10^{-7}	48/64	330/333
41.8879	30	1.0	170.9	7.271×10^{-7}	50/65	446/451

TABLE 3. Performance of the method depending on the size H of the aggregates. Fixed $k = 125.6637$. Mesh 400×400 elements. Number of degrees of freedom 964,806. Number of processors 8.

H/h (els)	H/wave length	cond.	memory [MB]	setup [s] CPU/WALL	iter [s] CPU/WALL
10	1.0	36.63	1,245	232/248	332/335
15	1.5	26.80	1,088	91/108	210/213
20	2.0	20.98	1,077	74/87	169/171
25	2.5	17.18	1,055	75/88	155/155
30	3.0	15.63	1,106	71/86	155/157
35	3.5	13.88	1,140	77/92	144/145
40	4.0	12.48	1,204	93/108	155/154
45	4.5	11.06	1,237	90/106	136/138
50	5.0	10.32	1,298	104/119	140/142

residuals (mesured in Euclidean norm) in Table 2. In all experiments presented here we used a prolongator smoother of degree 4.

Three different types of experiments were done: testing the method in the case of varying k, varying the coarse space size, i.e. the size of the aggregates $\mathcal{A}_i$, and parallel scalability. In all experiments, we have used the scattering model problem described in Table 1.

Due to the regular geometry of our testing problem, we were able to use a system of square aggregates. The size of the side of those squares is denoted by H.

Our results are summarized in Tables 2 and 3, and graphs in Figures 1 and 2. In all the experiments, the actual iteration history was well characterized by the condition number of the preconditioned matrix (we observed no spectral clustering). In other words, the error was being reduced by the factor of about $(\sqrt{\text{cond}} - 1)/(\sqrt{\text{cond}} + 1)$ per iteration. The runtime condition number estimates are reported.

We observe that the condition number of the preconditioned problem depends on H/λ_k, the ratio of the subdomain size to the wavelength $\lambda_k = 2\pi/k$. The condition number decreases with this ratio increasing (due either to increasing the frequency or to enlarging the subdomain size). We offer the following heuristic explanation: In the case of Helmholtz equation, the low-energy functions are waves
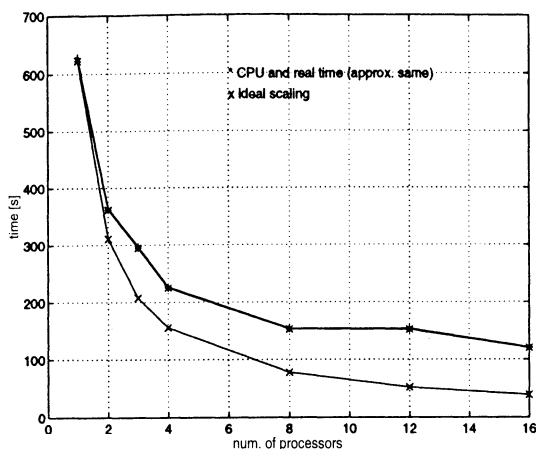
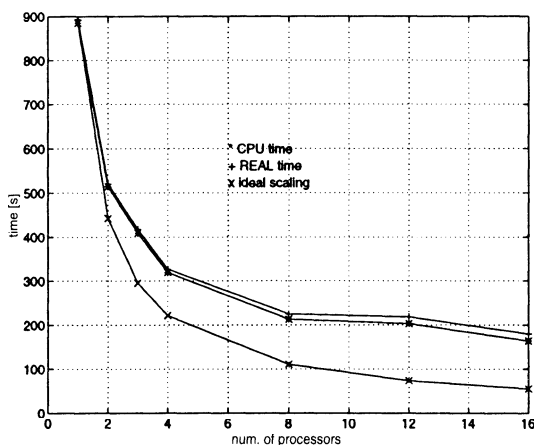FIGURE 1. Parallel scalability of the iteration phase. $k = 125.6637$, $H = 30h$.

FIGURE 2. Parallel scalability of the method (setup+iter). $k = 125.6637$, $H = 30h$.

with wavelength close to the parameter λ_k. For the problem with λ_k small compared to the subdomain size, the direct subdomain solvers (playing the role of a smoother) are capable of approximating such functions well.

The performance tests have shown that using 8 processors, we gain parallel speedup of about 4 times, and the effect of adding more processors seems to be decreasing. We believe that this scalability issue reflects the early stage of parallelization of the code at the time of writing.

References

1. I. Babuška and J. M. Melenk, *The partition of unity finite element method*, Int. J. Numer. Meths. Engrg. **40** (1997), 727–758.
2. Petter E. Bjørstad and Jan Mandel, *Spectra of sums of orthogonal projections and applications to parallel computing*, BIT **31** (1991), 76–88.
3. Marian Brezina and Petr Vaněk, *One black-box iterative solver*, UCD/CCM Report 106, Center for Computational Mathematics, University of Colorado at Denver, 1997, http://www-math.cudenver.edu/ccmreports/rep106.ps.gz.

4. Zhiqiang Cai, Thomas A. Manteuffel, and Stephen F McCormick, *First-order system least squares for second-order partial differential equations: Part ii.*, SIAM J. Numer. Anal. **34** (1997), 425–454.

5. Armel de La Bourdonnaye, *Une Méthode de discrétisation microlocale et son application à un problème de diffraction*, C. R. Acad. Sci. Paris, Serie I **318** (1994), 385–388.

6. Charles I. Goldstein, *Multigrid preconditioners applied to the iterative solution of singularly perturbed elliptic boundary value problems and scattering problems*, Innovative numerical methods in engineering, Proc. 4th Int. Symp., Atlanta/Ga., 1986 (Berlin) (R.P. Shaw, J. Periaux, A. Chaudouet, J. Wu, C. Marino, and C.A. Brebbia, eds.), Springer-Verlag, 1986, pp. 97–102.

7. W. Hackbusch, *A fast iterative method for solving Helmholtz's equation in a general region*, Fast Elliptic Solvers (U. Schumann, ed.), Advance Publications, London, 1978, pp. 112–124.

8. B. Lee, T. Manteuffel, S. McCormick, and J. Ruge, *Multilevel first-order system least squares (FOSLS) for Helmholtz equation*, Procs. 2nd International Conf. on Approx. and Num. Meths. for the Solution of the Maxwell Equations, Washington, D.C, John Wiley and Sons, 1993.

9. Jan Mandel, Marian Brezina, and Petr Vaněk, *Energy optimization of algebraic multigrid bases*, UCD/CCM Report 125, Center for Computational Mathematics, University of Colorado at Denver, February 1998, http://www-math.cudenver.edu/ccmreports/rep125.ps.gz.

10. Jan Mandel and Mirela Popa, *A multigrid method for elastic scattering*, UCD/CCM Report 109, Center for Computational Mathematics, University of Colorado at Denver, September 1997, http://www-math.cudenver.edu/ccmreports/rep109.ps.gz.

11. Petr Mayer and Jan Mandel, *The finite element ray method for the Helmholtz equation of scattering: First numerical experiments*, UCD/CCM Report 111, Center for Computational Mathematics, University of Colorado at Denver, October 1997.

12. B. F. Smith, P. E. Bjørstad, and W. D. Gropp, *Domain decomposition: Parallel multilevel methods for elliptic partial differential equations*, Cambridge University Press, New York, 1996.

13. P. Vaněk, *Acceleration of convergence of a two level algorithm by smoothing transfer operators*, Appl. Math. **37** (1992), 265–274.

14. Petr Vaněk, Marian Brezina, and Jan Mandel, *Algebraic multigrid for problems with jumps in coefficients*, In preparation.

15. Petr Vaněk, Jan Mandel, and Marian Brezina, *Algebraic multigrid on unstructured meshes*, UCD/CCM Report 34, Center for Computational Mathematics, University of Colorado at Denver, December 1994, http://www-math.cudenver.edu/ccmreports/rep34.ps.gz.

16. _____, *Algebraic multigrid based on smoothed aggregation for second and fourth order problems*, Computing **56** (1996), 179–196.

17. Petr Vaněk, Jan Mandel, and Marian Brezina, *Solving a two-dimensional Helmholtz problem by algebraic multigrid*, UCD/CCM Report 110, Center for Computational Mathematics, University of Colorado at Denver, October 1997, http://www-math.cudenver.edu/ccmreports/rep110.ps.gz.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF COLORADO AT DENVER, DENVER, CO 80217-3364, AND UNIVERSITY OF WEST BOHEMIA, PLZEŇ, CZECH REPUBLIC
E-mail address: pvanek@math.cudenver.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF COLORADO AT DENVER, DENVER, CO 80217-3364, AND CENTER FOR SPACE STRUCTURES, UNIVERSITY OF COLORADO AT BOUNDER, BOULDER CO 80309-0429
E-mail address: jmandel@math.cudenver.edu

DEPARTMENT OF APPLIED MATHEMATICS, UNIVERSITY OF COLORADO AT BOULDER, BOULDER CO 80309-0526
E-mail address: mbrezina@math.cudenver.edu