Analysis of the multi-level, shifted Laplace preconditioned method for the Helmholtz equation

ESF OPTPDE Workshop: Fast Solvers for Simulation, Inversion, and Control of Wave Propagation Problems

C. Vuik, A.H. Sheikh, D. Lahaye, H. Knibbe, and C.W. Oosterlee

27 September 2011

Contents

The Helmholtz equation

The Helmholtz equation without damping

$$-\Delta \mathbf{u}(x,y) - k^2(x,y)\mathbf{u}(x,y) = \mathbf{g}(x,y) \text{ in } \Omega$$

$\mathbf{u}(x,y)$ is the pressure field,
$\mathbf{k}(x,y)$ is the wave number,
$\mathbf{g}(x,y)$ is the point source function and
Ω is domain bounded by Absorbing boundary conditions

$$\frac{\partial \mathbf{u}}{\partial n} - \iota \mathbf{u} = 0$$

n is normal direction to respective boundary.

Problem description

- Second order Finite difference stencil:

$$\begin{bmatrix} & -1 & \\ -1 & 4 - k^2h^2 & -1 \\ & -1 & \end{bmatrix}$$

- Linear system $Au = g$: properties

  Sparse & complex valued

  Symmetric & Indefinite for large k

- Is traditionally solved by Krylov methods, which exploit sparsity.

Preconditioning

- ILU and variants
- From Laplace to complex shifted Laplace preconditioner (2005)
- Shifted Laplace preconditioner (SLP)

$$M := -\Delta \mathbf{u} - (\beta_1 - \iota\beta_2)k^2\mathbf{u}$$

- Results shows $(\beta_1, \beta_2) = (1, 0.5)$ is the optimal shift
- What does SLP do?

Shifted Laplace Preconditioner

- Introduces damping, Multigrid approximation
- Norm of spectrum of preconditioned operator bounded above by 1
- Spectrum goes to zero, as k increases.

Spectrum of $M^{-1}(1, 0.5)A$ for

$k = 30$ and $k = 120$

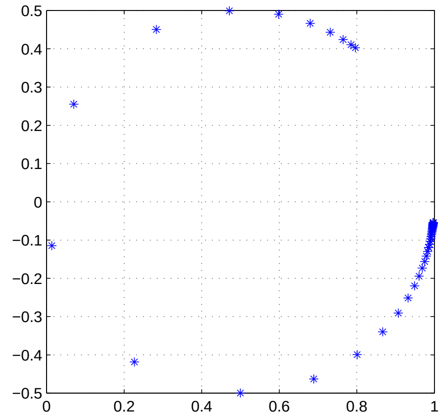

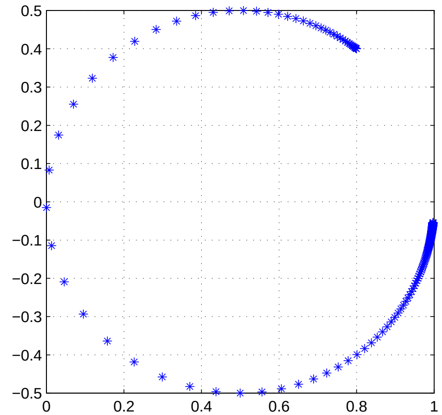

Some Restuls at a Glance

Number of GMRES iterations. Shifts in preconditioner are $(1, 0.5)$

Grid	$k = 10$	$k = 20$	$k = 30$	$k = 40$	$k = 50$	$k = 100$
$n = 32$	**10**	17	28	44	70	13/14
$n = 64$	10	**17**	28	36	45	173/163
$n = 96$	10	17	**27**	35	43	36/97
$n = 128$	10	17	27	**35**	43	36/85
$n = 160$	10	17	27	35	**43**	25/82
$n = 320$	10	17	27	35	42	**80**

Deflation improves!!!

Number of GMRES iterations. Shifts in preconditioner are $(1, 0.5)$

Grid	$k = 10$	$k = 20$	$k = 30$	$k = 40$	$k = 50$	$k = 100$
$n = 32$	**5/10**	8/17	14/28	26/44	42/70	13/14
$n = 64$	4/10	**6/17**	8/28	12/36	18/45	173/163
$n = 96$	3/10	5/17	**7/27**	9/35	12/43	36/97
$n = 128$	3/10	4/17	6/27	**7/35**	9/43	36/85
$n = 160$	3/10	4/17	5/27	6/35	**8/43**	25/82
$n = 320$	3/10	4/17	4/27	5/35	5/42	**10/80**

Erlangga and Nabben, 2008

with / without deflation.

Deflation: Definiton

For any deflation deflation subspace matrix

$$Z \in R^{n \times r}, \quad \text{with deflation vectors} \quad Z = [z_1, ..., z_r], \quad rank Z = r$$

$$P = I - AQ, \quad \text{with} \quad Q = ZE^{-1}Z^T \ \text{and} E = Z^T AZ$$

Solve $PAu = Pg$ preconditioned by M^{-1} or $M^{-1}PA = M^{-1}Pg$

For e.g. say,

$$\textbf{spec}(A) = \{\lambda_1, \lambda_2, \lambda_3, ..., \lambda_n\}$$

and if Z is matrix with columns as the r eigenvectors then

$$\textbf{spec}(PA) = \{0, ..., 0, \lambda_{r+1}, ...\lambda_n\}$$

We use multigrid inter-grid transfer operator (Prolongation) as deflation matrix.

Deflation

Setting $Z = I_h^{2h}$ and $Z^T = I_{2h}^h$ then

$$P = I - AQ, \quad \text{with} \quad Q = I_h^{2h} E^{-1} I_{2h}^h \ \text{and} E = I_{2h}^h A I_h^{2h}$$

where

P can be read as coarse grid correction and

Q the coarse grid operator

Fourier Analysis

Dirichlet boundary conditions for analysis.
With above deflation,

$$\textbf{spec}(PM^{-1}A) = f(\beta_1, \beta_2, k, h)$$

is a complex valued function.
Setting $kh = 0.625$,

- The non-zero eigenvalues of $PM^{-1}A$ near zero are wrapped and clustered around 1 with a few outliers.
- Spectrum remains almost the same, when imaginary shift is varied from 0.5 to 1.

Fourier Analysis

Analysis shows spectrum clustered around 1 with few outliers.

$k = 30$

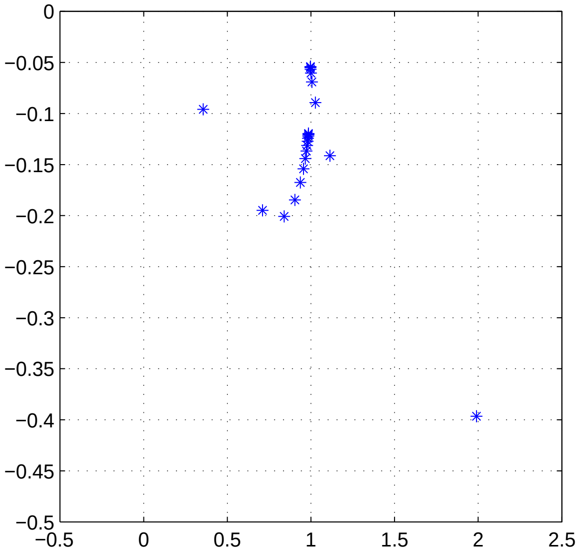

$k = 120$

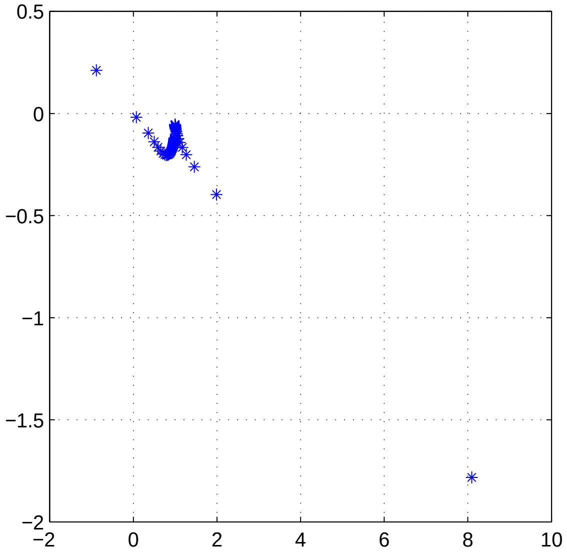

Fourier Analysis

Analysis tells increase in imaginary shift does not change spectrum.

$(\beta_1, \beta_2) = (1, 0.5)$

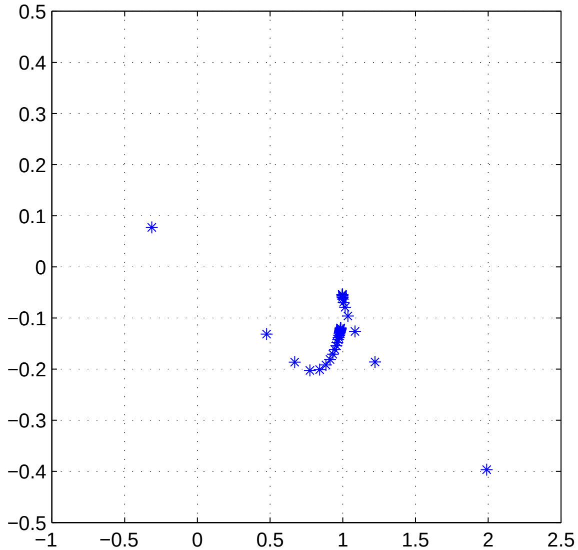

$(\beta_1, \beta_2) = (1, 1)$

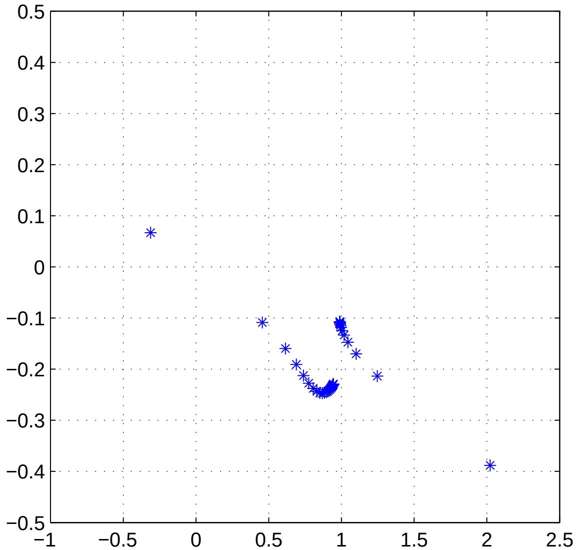

Numerical results

Sommerfeld boundary conditions are used for test problem.

Increase in imaginary shift in SLP?

Constant wavenumber problem

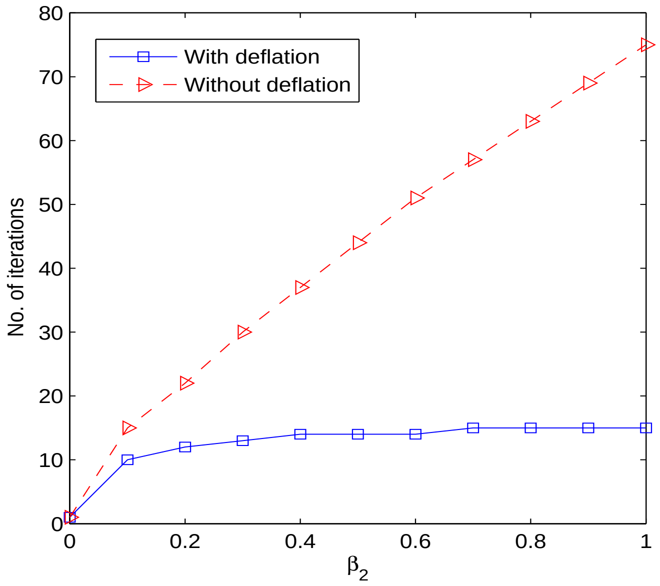

Wedge problem

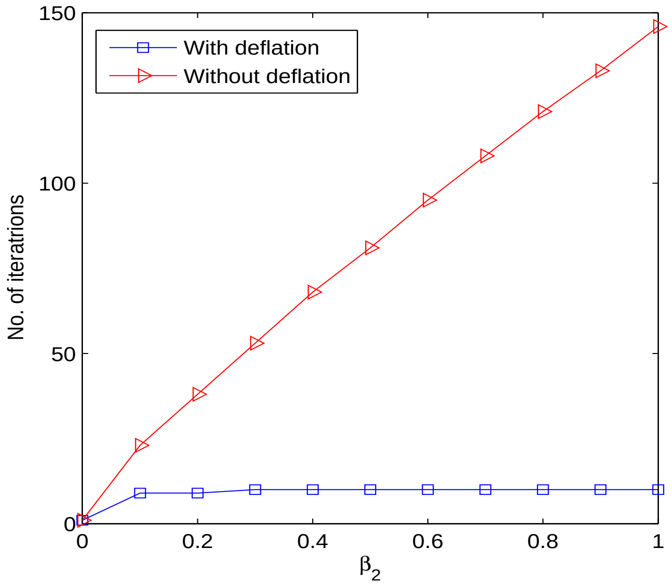

Numerical results

Number of GMRES iterations with/without deflation. Shift in preconditioner is $(1, 0.5)$

Grid	$k = 10$	$k = 20$	$k = 30$	$k = 40$	$k = 50$	$k = 100$
$n = 32$	**5/10**	8/17	14/28	26/44	42/70	13/14
$n = 64$	4/10	**6/17**	8/28	12/36	18/45	173/163
$n = 96$	3/10	5/17	**7/27**	9/35	12/43	36/97
$n = 128$	3/10	4/17	6/27	**7/35**	9/43	36/85
$n = 160$	3/10	4/17	5/27	6/35	**8/43**	25/82
$n = 320$	3/10	4/17	4/27	5/35	5/42	**10/80**

Numerical results

Number of GMRES iterations with/without deflation to solve a Wedge problem. Shift in preconditioner is $(1, 0.5)$

Grid	$freq = 10$	$freq = 20$	$freq = 30$	$freq = 40$	$freq = 50$
74×124	**7/33**	20/60	79/95	267/156	490/292
148×248	5/33	**9/57**	17/83	42/112	105/144
232×386	5/33	7/57	**10/81**	25/108	18/129
300×500	4/33	6/57	8/81	**12/105**	18/129
374×624	4/33	5/57	7/80	9/104	**13/128**

Numerical results

Number of GMRES outer-iterations in multilevel algorithm.
$(\beta_1, \beta_2) = (1, 0.5)$ $kh = .3125$ or 20 gp/wl and MG Vcycle(1,1) for SLP

Grid	$k = 10$	$k = 20$	$k = 40$	$k = 80$	$k = 160$
MLMGV(4,2,1)	9	11	16	27	100+
MLMGV(6,2,1)	9	10	14	21	47
MLMGV(8,2,1)	9	10	13	20	38
MLMGV(8,3,2)	9	10	13	19	37

Fourier Analysis

Spectrum of A, $M^{-1}A$ and $PM^{-1}A$ (from left to right) in bar-graph.

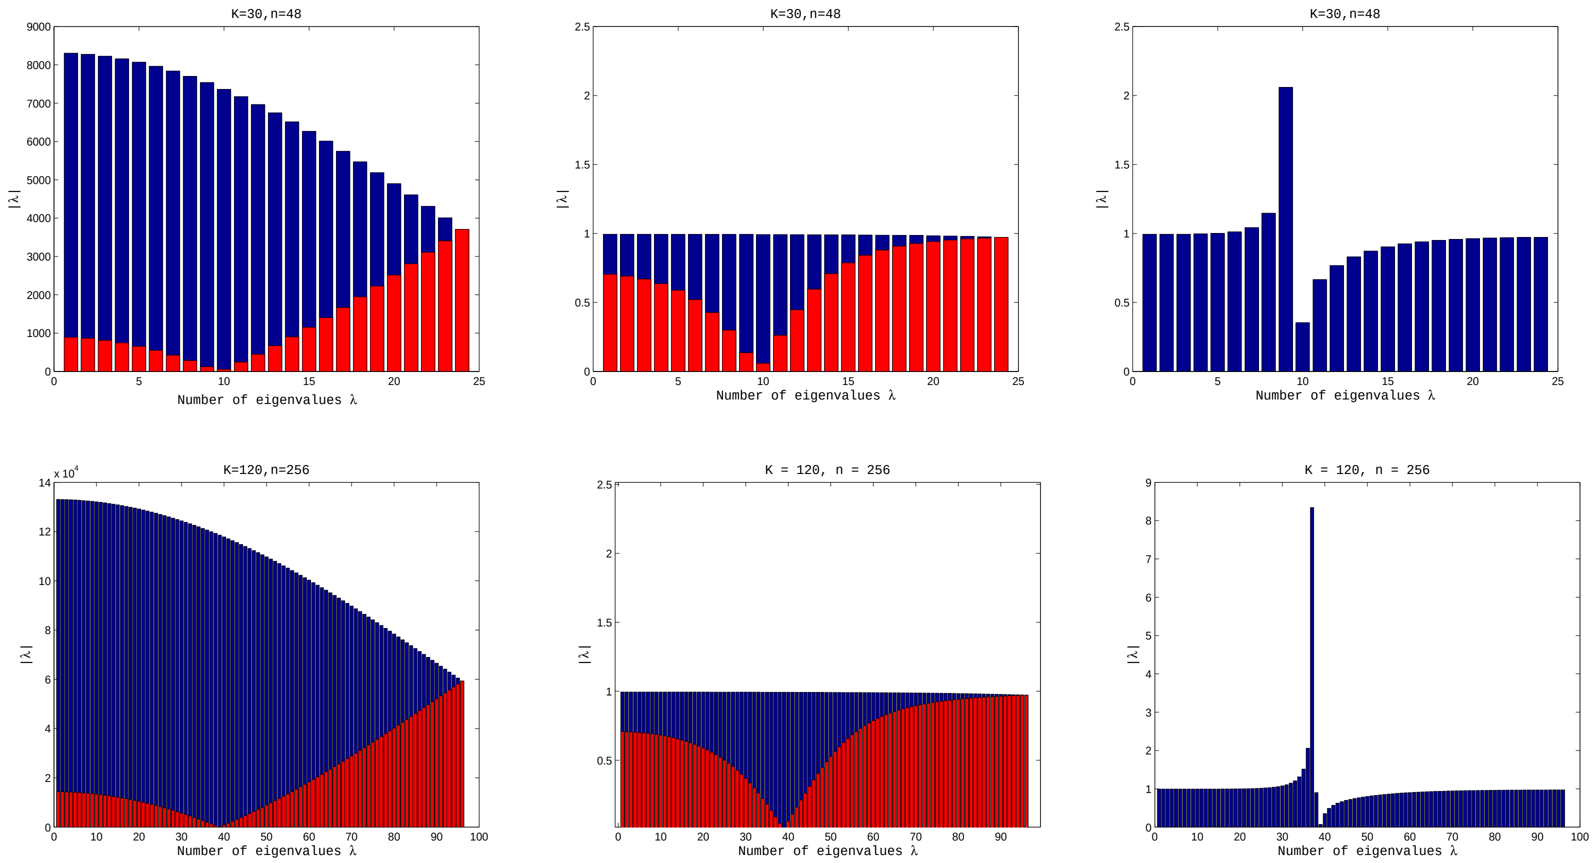

Delft Institute of Applied Mathematics

TU Delft

Implementation on Multiple GPU's

Bi-CGSTAB preconditioned by shifted Laplace multigrid method.

Equation solved in preconditioner is

$$-\frac{\partial^2\phi}{\partial x^2} - \frac{\partial^2\phi}{\partial y^2} - \frac{\partial^2\phi}{\partial z^2} - (\beta_1 - \beta_2 i)k^2\phi = g, \qquad (1)$$

$\beta_1, \beta_2 \in \mathbb{R}$, with the same boundary conditions as the original problem.

Multigrid components:

- Matrix-dependent prolongation (2D: de Zeeuw, 1990, 3D: Zhebel, 2006)
- Standard restriction
- Multi-coloured Gauss-Seidel as a smoother

Preconditioner is computed in single precision.

Little-Green Machine

Little-Green Machine

20 general computing nodes

- 2 Intel quadcore E5620
- 24 GB RAM
- 2 TB disk
- 2 NVIDIA GTX480

Funded by

- University of Leiden
- NWO project number 612.071.305
- TU Delft (DCSE www.cse.tudelft.nl)
- KNMI

NVidia Computer

8 GPUs, each GPU has 448 cores, 3 GB RAM

12 cores (2 Westmeres), 48 GB RAM

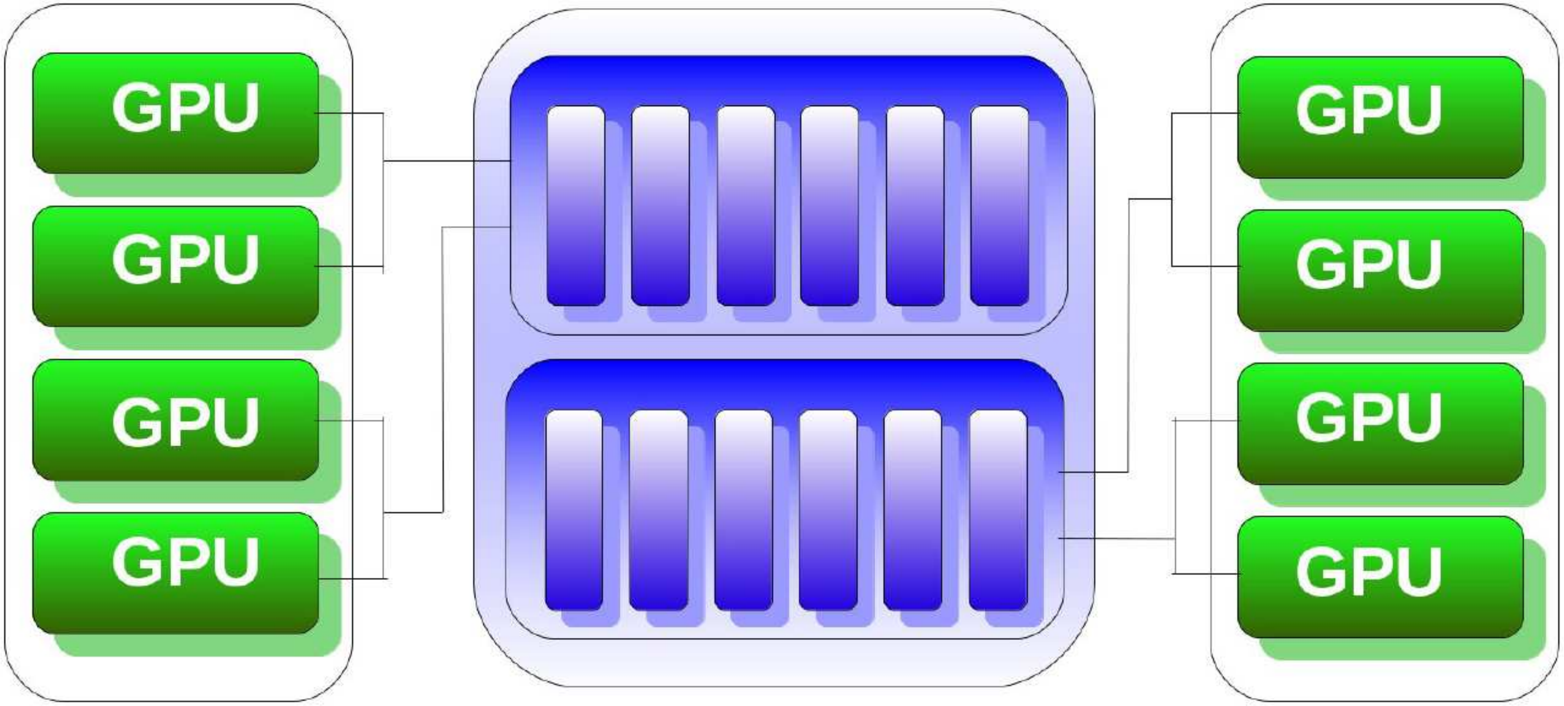

Gauss-Seidel Smoother

Four color Gauss-Seidel

Size	Time 8-cores (ms)	Time GPU (ms)	CPU/GPU
10,000	4	0.6	**7**
100,000	23.4	0.8	**29**
1,000,000	164.5	2.5	**66**
5,000,000	625.9	10.3	**61**
20,000,000	3733.9	39.4	**95**

Multi-GPU Approach

1. Data-parallel approach (e.g. vector operations on multi-GPU)
   (a) Relatively easy to implement
   (b) CPU→GPU→CPU data transfer
2. Split of the algorithm (e.g. solver on one GPU, preconditioner on the another one)
   (a) No or little data transfers
   (b) Find the best way to split the algorithm
3. Domain-Decomposition approach (e.g. each domain on a different GPU)
   (a) Exchange of halos (still data transfer)
   (b) Can affect convergence of the preconditioned method

Multi-GPU Issues

- Limited GPU memory size so need multiple GPUs for large problems.
- Efficient memory reusage to avoid allocation/deallocation, e.g. pool of GPU-vectors.
- Limit communications CPU$\rightarrow$GPU and GPU$\rightarrow$CPU.
- Each GPU need separate texture reference.
- Cublas vectors limited to 512 MB.

Bi-CGSTAB

Timings for Bi-CGSTAB, single precision

n	12-cores	1 GPU	**Speedup**	8-GPU	**Speedup**
5,000,000	24 s	0.8 s	**29.8**	2.3 s	**10.5**
15,000,000	82 s	2 s	**38.1**	5.8 s	**14.2**
100,000,000	395 s	-	-	28.6 s	**13.8**

Bi-CGSTAB with SLP and multigrid

Wedge problem, size $350 \times 350 \times 350 \approx 43,000,000$ unknowns

	Bi-CGSTAB (DP)	Preconditioner (SP)	Total
12-cores	94 s	690 s	784 s
1 GPU	13 s	47 s	60 s
Speedup CPU/GPU	**7.2**	**14.7**	**13.1**
8 GPUs	83 s	86 s	169 s
Speedup CPU/GPUs	**1.1**	**7.9**	**4.6**
2 GPUs+split	12 s	38 s	50 s
Speedup CPU/GPU	**7.8**	**18.2**	**15.5**

Conclusions

- The two grid method is parameter independent
- Numerical results confirm analysis
- Flexibility to increase imaginary shift, when deflation is combined with SLP.
- Further research Multilevel scheme, applying similar algorithm for coarse problem in deflation, does not behave similarly, so not scalable in amount of work
- Multiple GPU’s can be used to solve realistic problems with a decrease of wall clock time of a factor 10-20

References

- Y.A. Erlangga and R. Nabben. On a multilevel Krylov method for the Helmholtz equation preconditioned by shifted Laplacian. ETNA, 2008.
- M.B. van Gijzen, Y.a. Erlangga and C. Vuik. Spectral analysis of the discrete Helmholtz operator preconditioned withe a shifted Laplacian. SIAM J.of Sc. Comp. 2007.
- J.M. Tang. Two level preconditioned Congugate Gradient methods with applications to bubbly flow problems. PhD Thesis, DIAM TU Delft 2008.
- A.H. Sheikh, D. Lahaye and C. Vuik. A scalable Helmholtz solver combining the shifted Laplace preconditioner with multigrid deflation. submitted
- U. Trottenberg, C.W. Oosterlee and A. Schuller. Multigrid. Academic Press London 2000.
- H. Knibbe and C.W. Oosterlee and C. Vuik GPU implementation of a Helmholtz Krylov solver preconditioned by a shifted Laplace multigrid method. Journal of Computational and Applied Mathematics, 236, pp. 281-293, 2011

Thank You for Your Attention