

Glimberg, Stefan Lemvig; Engsig-Karup, Allan Peter

Publication date: 2011

Document Version Publisher's PDF, also known as Version of record

Link back to DTU Orbit

Citation (APA): Glimberg, S. L. (Author), & Engsig-Karup, A. P. (Author). (2011). A Fast Mixed-Precision Strategy for Iterative Gpu-Based Solution of the Laplace Equation. Sound/Visual production (digital) http://www.siam.org/meetings/pd11/

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.

- · You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Stefan L. Glimberg

Section of Scientific Computing Department of Informatics and Mathematical Modelling Technical University of Denmark

SIAM Analysis of Partial Differential Equations November 14th, 2011



A Fast Mixed-precision Strategy for Iterative GPU-based Solution of the Laplace Equation

Technical University of Denmark

Potential Flow Model		
●000		



A Fast Mixed-precision Strategy for Iterative GPU-based Solution of the Laplace Equation

Potential Flow Model		
0000		

Fully Nonlinear Free Surface Water Waves

The potential flow equations describe fully nonlinear water waves under the assumption of inviscid and irrotational flow.

2D Potential Flow Equations



Wave parameters

- η surface elevation
- ϕ potential ($u = \nabla \phi$)
- h still water depth
- $k = 2\pi/L$ wave number
- kh dispersion
- H/L nonlinearity

Potential Flow Model		
0000		

Fully Nonlinear Free Surface Water Waves

The potential flow equations describe fully nonlinear water waves under the assumption of inviscid and irrotational flow.

2D Potential Flow Equations



$$\begin{split} &\partial_t \eta = -\partial_x \eta \, \partial_x \tilde{\phi} + \tilde{\omega} (1 + (\partial_x \eta)^2) \\ &\partial_t \tilde{\phi} = -g\eta - \frac{1}{2} ((\partial_x \tilde{\phi})^2 - \tilde{\omega}^2 (1 + (\partial_x \eta)^2)) \\ &\tilde{\omega} = \partial_z \tilde{\phi}, \quad \tilde{\phi} = \phi|_{z=\eta} \end{split}$$

For $\tilde{\omega}$ to be computed, we need to know the potential in the entire domain.

$$\begin{split} \phi &= \tilde{\phi}, \quad z = \eta \\ \partial_{xx} \phi &+ \partial_{zz} \phi = 0, \quad -h \leq z < \eta \\ \partial_z \phi &+ \partial_x h \, \partial_x \phi = 0, \quad z = -h \end{split}$$

A Fast Mixed-precision Strategy for Iterative GPU-based Solution of the Laplace Equation

Potential Flow Model		
0000		

σ -Transformed Laplace Equation

$$\sigma(x, z, t) = \frac{z + h(x)}{\eta(x, t) + h(x)}$$



 $\Phi = \tilde{\phi}, \quad \sigma = 1$ $\partial_{xx}\Phi + \partial_{xx}\sigma(\partial_{\sigma}\Phi) + 2\partial_{x}\sigma(\partial_{x\sigma}\Phi) + ((\partial_{x}\sigma)^{2} + (\partial_{z}\sigma)^{2})\partial_{\sigma\sigma}\Phi = 0, \quad 0 \leq \sigma$ $(\partial_{z}\sigma + \partial_{x}h\partial_{x}\sigma)\partial_{\sigma}\Phi + \partial_{x}h\partial_{x}\Phi = 0, \quad \sigma \neq 0$

Potential Flow Model		
0000		

Linear Free Surface Water Waves

If wave amplitudes are small $\eta < \epsilon$, then the total water depth is almost the same as the still water depth $(\eta + h \approx h)$. If also the derivatives in η and h are assumed to be zero, the free surface equations take linear form.

Linearized Laplace Equation

$$egin{aligned} \Phi &= ilde{\phi}, \quad \sigma = 1 \ \partial_{xx} \Phi + (\partial_z \sigma)^2 \partial_{\sigma\sigma} \Phi &= 0, \quad 0 \leq \sigma < 1 \ \partial_z \sigma \, \partial_\sigma \Phi &= 0, \quad \sigma = 0 \end{aligned}$$

These equations might serve as an approximation for the fully nonlinear equations and can thus be used for preconditioning.

GPU Computing	
0000	



A Fast Mixed-precision Strategy for Iterative GPU-based Solution of the Laplace Equation

Technical University of Denmark

GPU Computing	
0000	

Motivation for GPU computing

There are several good reasons to consider Graphical Processing Units for high-performance computing

- Massively parallel architecture, \sim 500 cores.
- Teraflops of floating point performance
- Moderate prices \$100 - \$2,000. A personal super computer
- Fairly easy to get started (CUDA, OpenCL)
- Number 2 and 4 on top500 are based on GPUs



< A >



	GPU Computing		
0000	0000	0000	000000

A GPU-based Framework for PDE Solvers

We have build a highly generic heterogenous CPU-GPU framework for fast PDE solver prototyping.

Framework Objectives

4

5

- Remove all GPU-specific code for the non-expert GPU programmer
- While maintaining the possibility to customize code at kernel level

```
gpulab::vector<float,host_memory>
                                         x_h(100,3.f); // Create host vector x, size 100, value 3
2
    gpulab::vector<float,device_memory>
                                                      // Create device vector x, transfer host data
                                         x_d(x_h);
3
    gpulab::vector<float,device_memory>
                                         y_d(x_d);
                                                       // Create device vector y, copy device data
    y_d.axpy(4.f,x_d);
                                                       // Do y = a*x+y on the device
    y_d.nrm2();
                                                       // Calculate the 2-norm on the device
```

Ideas are based on the C++ standard library, Thrust, and CUSP that exist for GPUs

GPU Computing	
0000	

Key library components

• Regular grid objects, 1D, 2D, 3D.

1	grid_dim <int> dim(100,100);</int>	// 100x100 grid
2	<pre>grid_dim<double> phys0(0.,0.);</double></pre>	// Domain starts in x=0, y=0
3	<pre>grid_dim<double> phys1(1.,1.);</double></pre>	// Domain end in x=1, y=1
4	grid_properties <int, double=""> grid_props(dim,</int,>	phys0, phys1);
5	<pre>grid<double,device_memory> u(grid_props);</double,device_memory></pre>	// Create u
6	<pre>grid<double,device_memory> f(grid_props);</double,device_memory></pre>	// Create f



< □ > < 同 >

GPU Computing	
0000	

Key library components

- Regular grid objects, 1D, 2D, 3D.
- Compact stencil-based flexible order FD operators

```
grid_dim <int > dim (100,100);
                                                  // 100x100 grid
2
    grid_dim<double> phys0(0.,0.);
                                                  // Domain starts in x=0, y=0
3
    grid_dim<double> phys1(1.,1.);
                                                  // Domain end in x=1, y=1
4
    grid_properties <int, double > grid_props(dim, phys0, phys1);
5
    grid<double,device_memory> u(grid_props);
                                                 // Create u
6
    grid<double,device_memory> f(grid_props);
                                                 // Create f
7
8
    FD::stencil_2d<double> A(2,4);
                                                  // Second order derivative, fourth order accuracy
q
    A.matvec(u.f);
                                                  // Calculate f = du/dxx + du/dyy
```



Image: Image:

GPU Computing	
0000	

Key library components

- Regular grid objects, 1D, 2D, 3D.
- Compact stencil-based flexible order FD operators
- Iterative methods for solving large systems of eqs.

```
grid_dim <int > dim (100,100);
                                                   // 100x100 grid
2
     grid_dim<double> phys0(0.,0.);
                                                   // Domain starts in x=0, y=0
3
     grid_dim<double> phys1(1.,1.);
                                                   // Domain end in x=1, y=1
4
     grid_properties <int, double > grid_props(dim, phys0, phys1);
5
     grid<double,device_memory> u(grid_props);
                                                   // Create u
6
     grid<double,device_memory> f(grid_props);
                                                  // Create f
7
8
    FD::stencil_2d<double> A(2,4);
                                                   // Second order derivative, fourth order accuracy
q
    A.matvec(u.f);
                                                   // Calculate f = du/dxx + du/dyy
10
11
    monitor m(iter.rtol.atol):
                                                   // Stopping criteria
     solvers::cg cg solver(A.m):
                                                   // Create a CG solver from A
13
     cg solver.solve(u.f);
                                                   // Solve Au = f
```



GPU Computing	
0000	

Key library components

- Regular grid objects, 1D, 2D, 3D.
- Compact stencil-based flexible order FD operators
- Iterative methods for solving large systems of eqs.
- Effective preconditioning strategies

```
1
    grid_dim <int > dim (100,100);
                                                   // 100x100 grid
2
     grid_dim<double> phys0(0.,0.);
                                                   // Domain starts in x=0, y=0
3
     grid_dim<double> phys1(1.,1.);
                                                  // Domain end in x=1, y=1
4
     grid_properties <int, double > grid_props(dim, phys0, phys1);
5
     grid<double,device_memory> u(grid_props);
                                                  // Create u
6
     grid<double,device_memory> f(grid_props);
                                                  // Create f
7
8
    FD::stencil_2d<double> A(2,4);
                                                   // Second order derivative, fourth order accuracy
q
    A.matvec(u.f);
                                                   // Calculate f = du/dxx + du/dyy
10
    monitor m(iter.rtol.atol):
11
                                                   // Stopping criteria
12
     solvers::cg cg solver(A.m):
                                                   // Create a CG solver from A
13
     cg solver.solve(u.f):
                                                   // Solve Au = f
14
15
     FD::stencil 2d<double> P(2.2):
                                                   // Second order derivative. second order accuracy
16
     cg solver.set preconditioner(P):
                                                   // Add the preconditioner
     cg solver.solve(u,f):
17
                                                   // Solve PAu = Pf
```

イロト イポト イヨト イヨト

	Iterative Solver Analysis	
	0000	



A Fast Mixed-precision Strategy for Iterative GPU-based Solution of the Laplace Equation

Technical University of Denmark

	Iterative Solver Analysis	
	0000	

Defect Correction Method

We found that the Defect Correction method works well for our Laplace problem

- High-order approximations (accuracy)
- Minimal storage overhead (problem size)
- Minimal global synchronization and reduction steps (parallelizable)
- Effective as GMRES in practice (effective)

Textbook Recipe

Algorithm: DC Method for approximate solution of Ax = bChoose $x^{[0]}$ 1 /* initial guess */ 2 k = 03 Repeat $r^{[k]} = b - A x^{[k]}$ 4 /* high order defect */ 5 Solve $M\delta^{[k]} = r^{[k]}$ /* preconditioner */ $x^{[k+1]} = x^{[k]} + \delta^{[k]}$ 6 /* defect correction */ 7 k = k + 18 **Until** convergence or $k > k_{max}$

lah

	Iterative Solver Analysis	
	0000	

Defect Correction Method

We found that the Defect Correction method works well for our Laplace problem

- High-order approximations (accuracy)
- Minimal storage overhead (problem size)
- Minimal global synchronization and reduction steps (parallelizable)
- Effective as GMRES in practice (effective)

Textbook Recipe

Algorithm: DC Method for approximate solution of Ax = bChoose $x^{[0]}$ 1 /* initial guess */ 2 k = 03 Repeat $r^{[k]} = h - A x^{[k]}$ 4 /* high order defect */ 5 Solve $M\delta^{[k]} = r^{[k]}$ /* preconditioner */ $x^{[k+1]} = x^{[k]} + \delta^{[k]}$ 6 /* defect correction */ 7 k = k + 18 **Until** convergence or $k > k_{max}$

lah

		Iterative Solver Analysis	
0000	0000	0000	000000

Analysis of Defect Correction Convergence

Rewriting DC into the form of a stationary iterative method

$$x^{[k+1]} = x^{[k]} + \mathcal{M}^{-1}(b - \mathcal{A}x^{[k]})$$
(1)

$$= (1 - \mathcal{M}^{-1}\mathcal{A})x^{[k]} + \mathcal{M}^{-1}b$$
(2)

$$=\mathcal{G}x^{[k]}+c, \quad k=0,1,...$$
 (3)

where G is called the iteration matrix. From stationary iterative theory we know that to ensure convergence towards the exact solution we must have

$$\rho(\mathcal{G}) < 1,$$

where $\rho(\mathcal{G})$ is the spectral radius of \mathcal{G} , i.e. the maximum absolute eigenvalue of \mathcal{G} . Closer to 0 means better convergence.

	Iterative Solver Analysis	
	0000	

Analysis of Defect Correction Convergence

We can now predict attainable convergence rates for various free surface setups using linear flexible-order preconditioners.



Dispersion (kh) expresses ratio between water depth and wave length and relates to the condition number of the Laplacian matrix.



	Mixed Precision
	00000



A Fast Mixed-precision Strategy for Iterative GPU-based Solution of the Laplace Equation

	Mixed Precision

Mixed Precision

Definition

• An algorithm that mixes different machine precision numbers in its calculations – while maintaining a high precision solution.

Advantages

Bandwith bound

- 1 double = 2 floats = 64 bits
- Less storage at all levels
- Less bandwith required

Compute bound

- 1 double multiplier \approx 4 float multipliers
- 1 double adder pprox 2 float adder
- On many GPUs 1:8



	Mixed Precision

Mixed Precision

Definition

• An algorithm that mixes different machine precision numbers in its calculations – while maintaining a high precision solution.

Example

• Single precision roundoff error:

 $c = 0.5 + 0.5 + 0.00000004 - 0.00000003 = 1.000000001 = 1_{fl}$

• Mixed precision fix:

 $a = 0.5 + 0.5 = 1_{fl}$ $b = 0.000000004 - 0.00000003 = 0.00000001_{fl}$ $c = a + b = 1.000000001_{dl}$



	Mixed Precision
	000000

Mixed Precision Defect Correction

The same principle holds for the defect correction update - and all other refinement processes in general.

/* Double Precision */

/* Single Precision */

/* Double Precision */

Mixed Precision DC

- Choose $x^{[0]}$ 1
- 2 k = 0
- 3 Repeat
- $r^{[k]} = b A x^{[k]}$ 4
- 5 Solve $M\delta^{[k]} = r^{[k]}$ $x^{[k+1]} = x^{[k]} + \delta^{[k]}$ 6
- 7 k = k + 1
- **Until** convergence or $k > k_{max}$ 8

Remember, much work lies within the preconditioner!

	Mixed Precision
	000000

Mixed Precision GPU-based Performance Results

Timings per Defect Correction iteration. Using 6^{th} order accurate stencil, preconditioned with a linear 2^{nd} order accurate multigrid approach, DC+MG-RB-GS-1V(2,2).





			Mixed Precision
0000	0000	0000	000000

Mixed Precision Convergence

The residual error at every iteration confirms that the mixed precision algorithm in fact maintains high accuracy.





	Mixed Precision
	000000



Allan P. Engsig-Karup.

Efficient low-storage solution of unsteady fully nonlinear water waves using a defect correction method.

Submitted to: Journal of Scientific Computing, 2011.



Allan Peter Engsig-Karup, Morten Gorm Madsen, and Stefan Lemvig Glimberg. A massively parallel gpu-accelerated model for analysis of fully nonlinear free surface waves. International Journal for Numerical Methods in Fluids, 2011.

