


High-resolution fluid effects




*Rasmussen et al, Smoke Simulation for Large Scale Phenomena
(SIGGRAPH 2003)*

High-resolution fluid effects



*Nielsen et al, Out-Of-Core and Compressed Level Set Methods
(TOG 2007)*

High-resolution fluid effects



Horvath & Geiger, Directable, high-resolution simulation of fire on the GPU (SIGGRAPH 2009)

Introduction

So, what's holding us back?

- Curse of dimensionality : 8x cost to double the linear resolution (typically more, if the time step needs to be shrunk, too)

Introduction

So, what's holding us back?

- Curse of dimensionality : 8x cost to double the linear resolution (typically more, if the time step needs to be shrunk, too)
- Footprint : Simulations can't fit on a single system's RAM (out-of-core processing, cluster computing are possible remedies)

Introduction

So, what's holding us back?

- Curse of dimensionality : 8x cost to double the linear resolution (typically more, if the time step needs to be shrunk, too)
- Footprint : Simulations can't fit on a single system's RAM (out-of-core processing, cluster computing are possible remedies)
- Poor scalability : Kernels with worse-than-linear complexity dominate

Introduction

Cost and scalability of fluids simulation components :

- Most components : linear complexity & admit efficient parallelization
e.g. velocity/particle update, levelset reinitialization, extrapolation, etc.

Introduction

Cost and scalability of fluids simulation components :

- Most components : linear complexity & admit efficient parallelization
e.g. velocity/particle update, levelset reinitialization, extrapolation, etc.
- Poisson pressure projection becomes the bottleneck at high resolution
 - Objective : Solve a discrete Poisson equation $\Delta \mathbf{x} = \mathbf{f}$
 - Common solver in graphics : Preconditioned conjugate gradient

Introduction

Cost and scalability of fluids simulation components :

- Most components : linear complexity & admit efficient parallelization
e.g. velocity/particle update, levelset reinitialization, extrapolation, etc.
- Poisson pressure projection becomes the bottleneck at high resolution
- Objective : Solve a discrete Poisson equation $\Delta \mathbf{x} = \mathbf{f}$
- Common solver in graphics : Preconditioned conjugate gradient
- Popular Poisson preconditioner : Incomplete Cholesky Factorization
 - Introduced in graphics : *Foster & Fedkiw, "Practical animation of liquids", 2001*
 - Difficult to parallelize without compromising preconditioning efficiency
 - Still too many iterations required for high resolution simulations

Introduction

Poisson solvers with even more favorable complexity?

- “Fast Poisson Solvers” : *Cyclic reduction, FFT, FACR, etc.*
- Extremely fast, when applicable
- Not appropriate for irregularly shaped domains

Introduction

Poisson solvers with even more favorable complexity?

- “Fast Poisson Solvers” : *Cyclic reduction, FFT, FACR, etc.*
 - Extremely fast, when applicable
 - Not appropriate for irregularly shaped domains
- Multigrid methods
 - Potentially $O(N)$ asymptotic complexity
 - Good parallel potential

Introduction

Poisson solvers with even more favorable complexity?

- “Fast Poisson Solvers” : *Cyclic reduction, FFT, FACR, etc.*
 - Extremely fast, when applicable
 - Not appropriate for irregularly shaped domains
- Multigrid methods
 - Potentially $O(N)$ asymptotic complexity
 - Good parallel potential
 - Complications can compromise convergence performance
 - Irregular domain shapes, highly variable boundary conditions
 - Elaborate topological features (bubbles, fingers, slits, etc)


Problem description

*The Pressure Poisson equation
(for projecting a velocity field
to its divergence-free component)*

$$\Delta p = f \quad \text{in } \Omega$$

$$p(\mathbf{x}) = \alpha(\mathbf{x}) \quad \text{on } \Gamma_D$$

$$p_n(\mathbf{x}) = \beta(\mathbf{x}) \quad \text{on } \Gamma_N$$




— Dirichlet Boundary (Γ_D)


— Neumann Boundary (Γ_N)

□ Computational Domain (Ω)

Problem description



- Dirichlet Boundary (Γ_D)
- Neumann Boundary (Γ_N)
- Computational Domain (Ω)



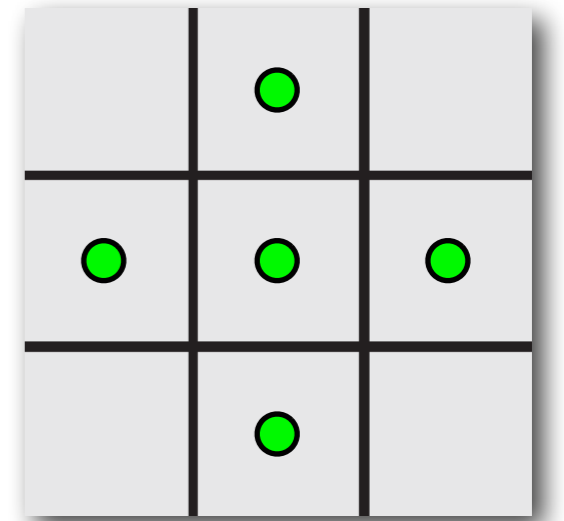
- Interior Cells
- Dirichlet Cells
- Neumann Cells

“Voxelized” Poisson problem

Problem description

Interior point discretization

$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$




Problem description

Interior point discretization

$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$

Discretization near a Dirichlet boundary

$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$




Problem description

Interior point discretization

$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$

Discretization near a Dirichlet boundary

$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$




Problem description

Interior point discretization

$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$

Discretization near a Dirichlet boundary

$$\frac{-4u_{ij} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij} - \frac{u_{i+1,j}}{h^2}$$




Problem description

Interior point discretization

$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$

Discretization near a Dirichlet boundary

$$\frac{-4u_{ij} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij} - \frac{\alpha_{i+1,j}}{h^2}$$



Problem description

Interior point discretization




$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$

Discretization near a Dirichlet boundary

$$\frac{-4u_{ij} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij} - \frac{\alpha_{i+1,j}}{h^2}$$

Discretization near a Neumann boundary

$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$



Problem description

Interior point discretization




$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$

Discretization near a Dirichlet boundary

$$\frac{-4u_{ij} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij} - \frac{\alpha_{i+1,j}}{h^2}$$

Discretization near a Neumann boundary

$$\frac{-3u_{ij} - u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$



Problem description

Interior point discretization




$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$

Discretization near a Dirichlet boundary

$$\frac{-4u_{ij} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij} - \frac{\alpha_{i+1,j}}{h^2}$$

Discretization near a Neumann boundary

$$\frac{-3u_{ij} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij} - \frac{u_{i+1,j} - u_{ij}}{h^2}$$



Problem description

Interior point discretization




$$\frac{-4u_{ij} + u_{i+1,j} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij}$$

Discretization near a Dirichlet boundary

$$\frac{-4u_{ij} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij} - \frac{\alpha_{i+1,j}}{h^2}$$

Discretization near a Neumann boundary


$$\frac{-3u_{ij} + u_{i-1,j} + u_{i,j+1} + u_{i,j-1}}{h^2} = f_{ij} - \frac{\beta_{i+\frac{1}{2},j}}{h}$$




Problem description

Properties of the resulting discretization

- Symmetric, sparse (banded) system
- Negative semi-definite (strictly definite with any Dirichlet boundary)
- Symmetric Krylov solvers are applicable, i.e. preconditioned CG




Sample 2D domain (512x512 resolution)




Exact solution

End frame [last valid frame]:



Conjugate Gradients (w/o preconditioning)

End frame [last valid frame]:



Conjugate Gradients (with a stock preconditioner)

Problem description

Properties of the resulting discretization


- Symmetric, sparse (banded) system
- Negative semi-definite (strictly definite with any Dirichlet boundary)
- Symmetric Krylov solvers are applicable, i.e. preconditioned CG
- *Convergence deterioration is even more pronounced in 3 dimensions, at higher resolutions and with more elaborate domain geometries*

Multigrid

What about multigrid solvers?


- Well suited for Poisson-type problems
- Should be able to provide resolution-independent convergence

frame 0




Multigrid iterations towards convergence (1 V-Cycle/frame)

frame 0



Multigrid iterations towards convergence (random initial guess)

frame 0



Conjugate Gradients (with parallel multigrid preconditioner)

Multigrid


What about multigrid solvers?

- Well suited for Poisson-type problems
- Should be able to provide resolution-independent convergence

Multigrid: relaxation

$$\Delta u = f$$

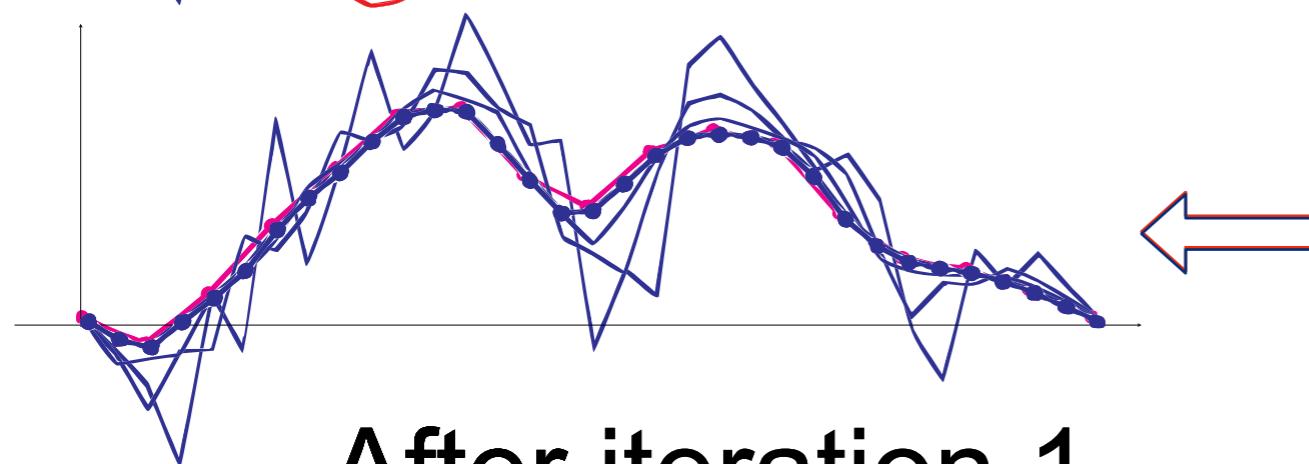
$$u(0) = a, u(1) = b$$



Relaxation:

a simple iterative
Initial guess
solver

Exact solution
Gauss-Seidel



After iteration 1
iteration 3


Error equation: $\Delta u_{\text{current}} - u$

easier to solve

Error equation:


$$\Delta e = r \quad (= \Delta u_{\text{current}} - f)$$

Multigrid cycle



- Relaxation
- Restriction
- Coarse grid solver
- Prolongation

Multigrid V-cycle



- Relaxation
- Restriction
- Coarse grid solver
- Prolongation
- Recursive coarsening
- Multigrid V-Cycle - $O(N)$
- Resolution independent numerical efficiency

Multigrid

What about multigrid solvers?

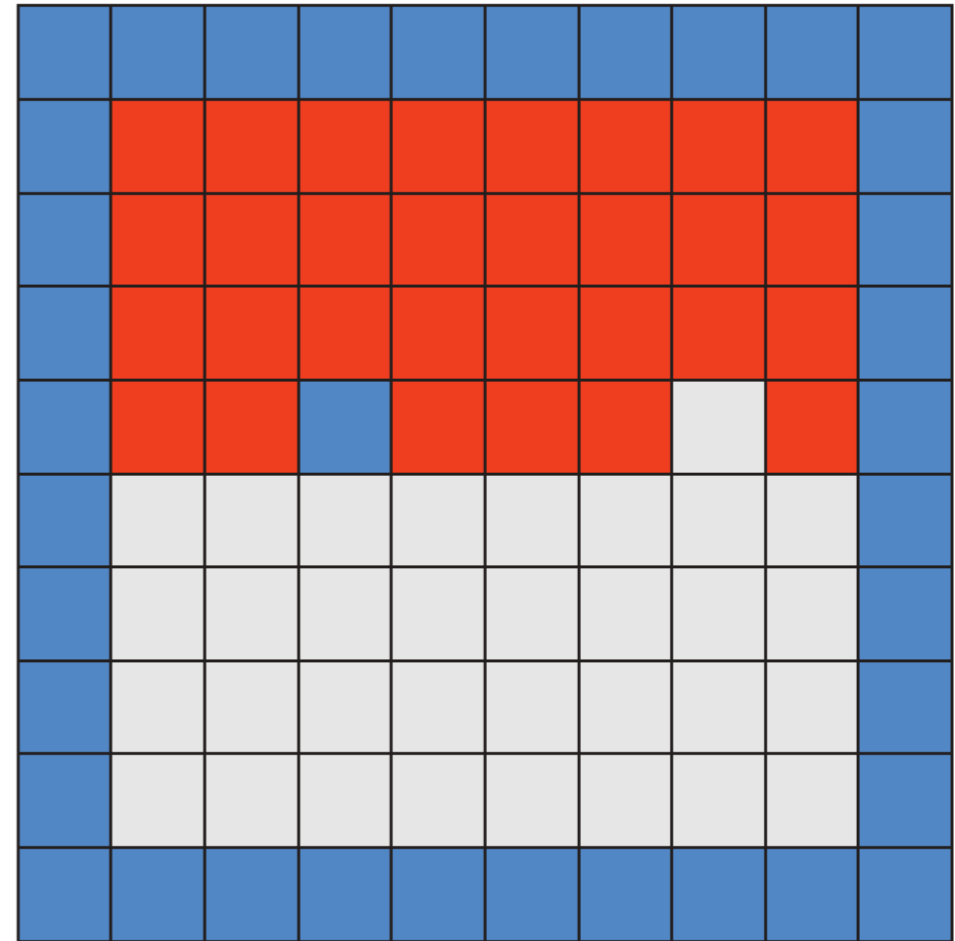


- Well suited for Poisson-type problems
- Should be able to provide resolution-independent convergence

For a multigrid solver, we need 3 algorithmic components

- A hierarchy of discretizations
- Transfer operators (restriction - prolongation)
- A “smoothing” routine




Multigrid

Creating a hierarchy of discretizations ...



Multigrid




Creating a hierarchy of discretizations ...



Every group of 4 cells is coarsened into 1 cell of at the immediately coarser level

Multigrid

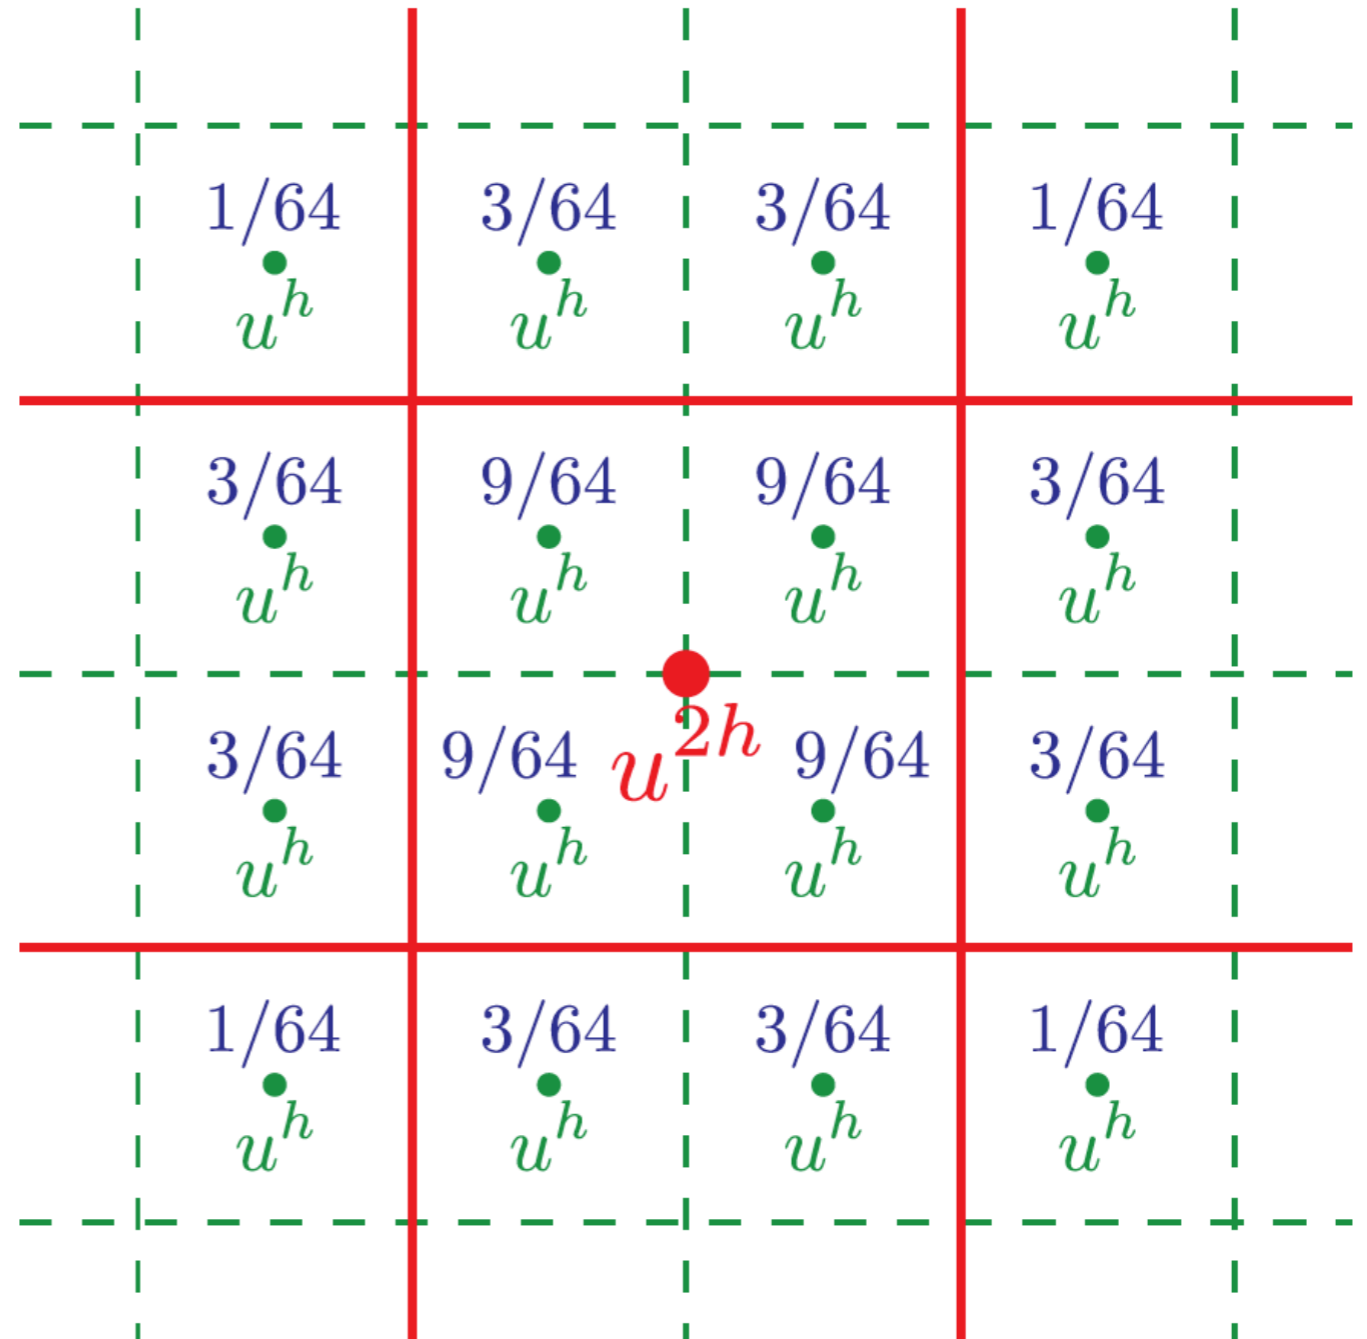
Creating a hierarchy of discretizations ...



*When fine grid cells have more than one type, the coarse cell becomes (by priority)
Dirichlet -- Interior (if no Dirichlet) -- Neumann (if no Interior or Dirichlet)*

Multigrid


*Inter-grid transfer operators
(Prolongation - Restriction)*




Multigrid

Smoothing procedure :

- *Perform N sweeps of relaxation (e.g. Gauss-Seidel method) on a band around the boundary*
- *Perform 1 sweep in the interior*
- *Perform N more sweeps on the boundary band*




frame 0




Stable Multigrid V-cycle (30 boundary smoothing sweeps)

frame 0




MG-PCG (same boundary smoothing effort as stable V-cycle)

frame 0



MG-PCG (1/3 of the smoothing effort needed for stable V-cycle)

frame 0



MG-PCG (1/30 of the smoothing effort needed for stable V-cycle)

Multigrid

No free lunch ...

- Convergence & Stability may necessitate an impractically intensive boundary smoothing effort

N = 30 for stability !

Smoothing procedure :

- Perform **N sweeps** of relaxation (e.g. Gauss-Seidel method) on a band around the boundary
- Perform 1 sweep in the interior
- Perform **N more sweeps** on the boundary band

Iteration 0

N



*V-Cycle iteration, 10 boundary iterations per cycle
(1/3 of the effort needed for convergent method)*

Multigrid

No free lunch ...

- Convergence & Stability may necessitate an impractically intensive boundary smoothing effort
- For moderate resolutions (e.g. 128^3) one V-cycle iteration may cost as much as 20-30 iterations of (unpreconditioned) CG

Multigrid

No free lunch ...

- Convergence & Stability may necessitate an impractically intensive boundary smoothing effort
- For moderate resolutions (e.g. 128^3) one V-cycle iteration may cost as much as 20-30 iterations of (unpreconditioned) CG
- Boundary cost goes away *asymptotically*, but only at impractically high resolutions




Multigrid

No free lunch ...

- Convergence & Stability may necessitate an impractically intensive boundary smoothing effort
- For moderate resolutions (e.g. 128^3) one V-cycle iteration may cost as much as 20-30 iterations of (unpreconditioned) CG
- Boundary cost goes away *asymptotically*, but only at impractically high resolutions
- Even worse, when considering parallelism

Multigrid

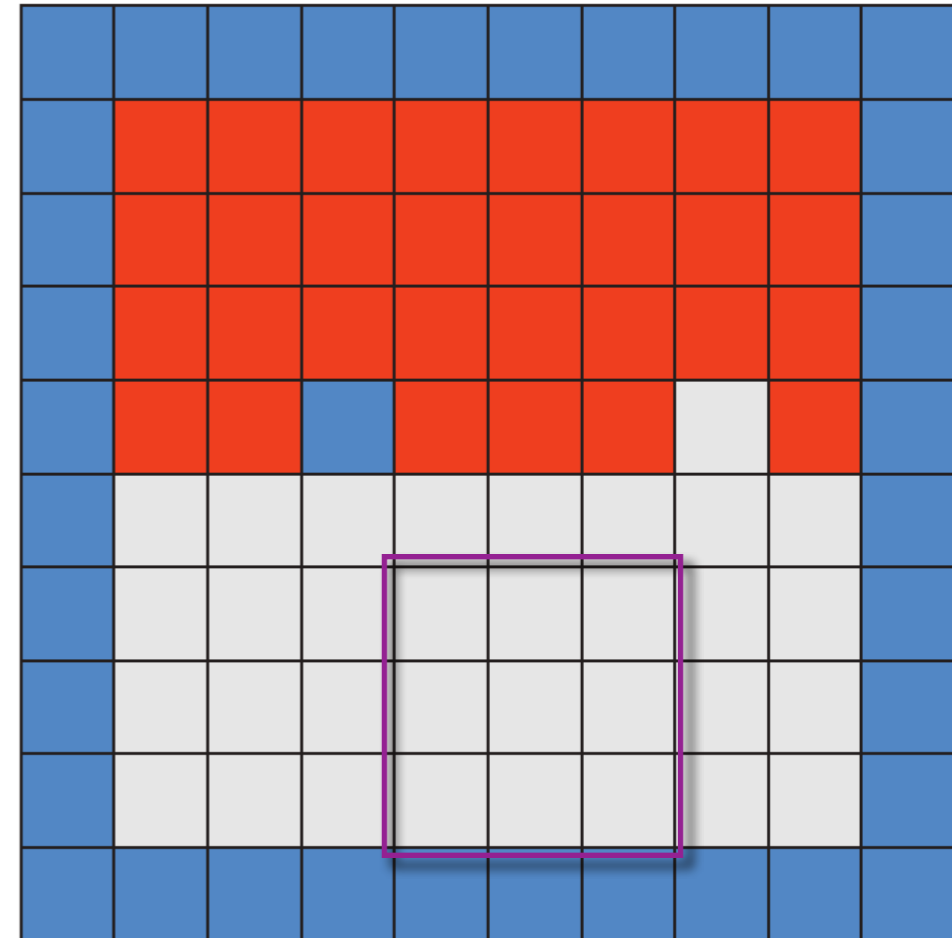


Problem #1 : Geometric discrepancies lead to instability



*Remedies : Intensive boundary smoothing, algebraic coarsening, specialized transfer operators ...
... or using MG as a preconditioner for a Krylov method*

Multigrid

Problem #2 : Topological discrepancies lead to stagnation



*Remedies : Algebraic coarsening, recombined iterants,
using a fully Algebraic Multigrid method ...
... or using MG as a preconditioner for a Krylov method*

Multigrid preconditioning

Basic concepts of Preconditioned Conjugate Gradients :

- Basic problem : $\mathbf{Ax} = \mathbf{b}$

Multigrid preconditioning

Basic concepts of Preconditioned Conjugate Gradients :

- Basic problem : $\mathbf{Ax} = \mathbf{b}$
- Consider a matrix \mathbf{M} such that $\mathbf{MA} \approx \mathbf{I}$

Multigrid preconditioning

Basic concepts of Preconditioned Conjugate Gradients :

- Basic problem : $\mathbf{Ax} = \mathbf{b}$
- Consider a matrix \mathbf{M} such that $\mathbf{MA} \approx \mathbf{I}$
- This matrix needs to be symmetric and positive definite.
In this case we can write : $\mathbf{M} = \mathbf{U}^2$ where \mathbf{U} is symmetric


Multigrid preconditioning

Basic concepts of Preconditioned Conjugate Gradients :

- Basic problem : $\mathbf{Ax} = \mathbf{b}$
- Consider a matrix \mathbf{M} such that $\mathbf{MA} \approx \mathbf{I}$
- This matrix needs to be symmetric and positive definite.
In this case we can write : $\mathbf{M} = \mathbf{U}^2$ where \mathbf{U} is symmetric
- The PCG method is algebraically equivalent to applying the Conjugate Gradients method on the modified system :


$$(\mathbf{UAU})(\mathbf{U}^{-1}\mathbf{x}) = \mathbf{Ub}$$

End frame [last valid frame]:



Conjugate Gradients (w/o preconditioning)

End frame [last valid frame]:



Conjugate Gradients (with a stock preconditioner)

Preconditioned Conjugate Gradients

```
1: procedure MGPCG( $\mathbf{r}, \mathbf{x}$ )
2:    $\mathbf{r} \leftarrow \mathbf{r} - \mathcal{L}\mathbf{x}$ ,  $\boldsymbol{\mu} \leftarrow \bar{\mathbf{r}}$ ,  $v \leftarrow \|\mathbf{r} - \boldsymbol{\mu}\|_\infty$ 
3:   if ( $v < v_{\max}$ ) then return
4:    $\mathbf{r} \leftarrow \mathbf{r} - \boldsymbol{\mu}$ ,  $\mathbf{p} \leftarrow \mathcal{M}^{-1}\mathbf{r}^{(\dagger)}$ ,  $\rho \leftarrow \mathbf{p}^T \mathbf{r}$ 
5:   for  $k = 0$  to  $k_{\max}$  do
6:      $\mathbf{z} \leftarrow \mathcal{L}\mathbf{p}$ ,  $\sigma \leftarrow \mathbf{p}^T \mathbf{z}$ 
7:      $\alpha \leftarrow \rho / \sigma$ 
8:      $\mathbf{r} \leftarrow \mathbf{r} - \alpha\mathbf{z}$ ,  $\boldsymbol{\mu} \leftarrow \bar{\mathbf{r}}$ ,  $v \leftarrow \|\mathbf{r} - \boldsymbol{\mu}\|_\infty$ 
9:     if ( $v < v_{\max}$  or  $k = k_{\max}$ ) then
10:        $\mathbf{x} \leftarrow \mathbf{x} + \alpha\mathbf{p}$ 
11:       return
12:     end if
13:      $\mathbf{r} \leftarrow \mathbf{r} - \boldsymbol{\mu}$ ,  $\mathbf{z} \leftarrow \mathcal{M}^{-1}\mathbf{r}^{(\dagger)}$ ,  $\rho^{\text{new}} \leftarrow \mathbf{z}^T \mathbf{r}$ 
14:      $\beta \leftarrow \rho^{\text{new}} / \rho$ 
15:      $\rho \leftarrow \rho^{\text{new}}$ 
16:      $\mathbf{x} \leftarrow \mathbf{x} + \alpha\mathbf{p}$ ,  $\mathbf{p} \leftarrow \mathbf{z} + \beta\mathbf{p}$ 
17:   end for
18: end procedure
```

Multigrid preconditioning

Basic concepts of Preconditioned Conjugate Gradients :

- Basic problem : $\mathbf{Ax} = \mathbf{b}$
- Consider a matrix \mathbf{M} such that $\mathbf{MA} \approx \mathbf{I}$
- This matrix needs to be symmetric and positive definite.
In this case we can write : $\mathbf{M} = \mathbf{U}^2$ where \mathbf{U} is symmetric
- The PCG method is algebraically equivalent to applying the Conjugate Gradients method on the modified system :

$$(\mathbf{UAU})(\mathbf{U}^{-1}\mathbf{x}) = \mathbf{Ub}$$

- The final algorithm requires only a routine for computing \mathbf{Mv} (for any input vector) to be specified

Multigrid preconditioning

Using a multigrid V-cycle to define a preconditioner :

Multigrid preconditioning

Using a multigrid V-cycle to define a preconditioner :

- Define a hierarchy of discrete Poisson problems *with zero Dirichlet and Neumann Boundary conditions at all levels.*

Multigrid preconditioning

Using a multigrid V-cycle to define a preconditioner :

- Define a hierarchy of discrete Poisson problems *with zero Dirichlet and Neumann Boundary conditions at all levels.*
- Use the input vector \mathbf{v} as the right hand side at the top level. *Use a zero initial guess at every level of the V-cycle.*

Multigrid preconditioning

Using a multigrid V-cycle to define a preconditioner :

- Define a hierarchy of discrete Poisson problems *with zero Dirichlet and Neumann Boundary conditions at all levels.*
- Use the input vector \mathbf{v} as the right hand side at the top level. *Use a zero initial guess at every level of the V-cycle.*
- Perform l V-cycle. The result is the desired product $\mathbf{M}\mathbf{v}$.

Multigrid preconditioning

Using a multigrid V-cycle to define a preconditioner :

- Define a hierarchy of discrete Poisson problems *with zero Dirichlet and Neumann Boundary conditions at all levels.*
- Use the input vector \mathbf{v} as the right hand side at the top level. *Use a zero initial guess at every level of the V-cycle.*
- Perform l V-cycle. The result is the desired product $M\mathbf{v}$.

A valid CG preconditioner (symmetric, definite) is obtained if :


- Restriction - prolongation are defined as adjoint operators.
- Jacobi (or damped Jacobi) is used instead of Gauss-Seidel to relax the interior of the domain.
- Boundary band is traversed by the smoother in opposite orders during the downstroke and upstroke of the V-cycle.

Multigrid preconditioning

MG-preconditioned CG benefits :


- Remains stable, even if an *unstable* V-cycle is used for preconditioning

frame 0




Stable Multigrid V-cycle (30 boundary smoothing sweeps)

frame 0




MG-PCG (same boundary smoothing effort as stable V-cycle)

frame 0



MG-PCG (1/3 of the smoothing effort needed for stable V-cycle)

frame 0



MG-PCG (1/30 of the smoothing effort needed for stable V-cycle)

Multigrid preconditioning

MG-preconditioned CG benefits :

- Remains stable, even if an *unstable* V-cycle is used for preconditioning
- Unstable V-cycles simply require a few extra PCG iterations

Multigrid preconditioning

MG-preconditioned CG benefits :

- Remains stable, even if an *unstable* V-cycle is used for preconditioning
- Unstable V-cycles simply require a few extra PCG iterations
- Intensity of the boundary treatment can be tuned to *moderate-difficulty scenarios*, and remain stable even in highly complicated cases

Multigrid preconditioning

MG-preconditioned CG benefits :

- Remains stable, even if an *unstable* V-cycle is used for preconditioning
- Unstable V-cycles simply require a few extra PCG iterations
- Intensity of the boundary treatment can be tuned to *moderate-difficulty scenarios*, and remain stable even in highly complicated cases
- With a well-designed (albeit unstable) V-cycle, PCG converges as quickly as the best-case scenario multigrid cycle, in practice.

Results and performance

