CS 267: Applications of Parallel Computers Lecture 17 - Structured Grids

## James Demmel

www.cs.berkeley.edu/~demmel/cs267_Spr16

## Outline

[^0]

## Solving PDEs

${ }^{\circ}$ Hyperbolic problems (waves):

- Sound wave(position, time)
- Use explicit time-stepping
- Solution at each point depends on neighbors at previous time
${ }^{\circ}$ Elliptic (steady state) problems:
- Electrostatic Potential (position)
- Everything depends on everything else
- This means locality is harder to find than in hyperbolic problems
- Parabolic (time-dependent) problems:
- Temperature(position,time)
- Involves an elliptic solve at each time-step
${ }^{\circ}$ Focus on elliptic problems
- Canonical example is the Poisson equation

$$
\partial^{2} \mathbf{u} / \partial \mathbf{x}^{2}+\partial^{2} \mathbf{u} / \partial \mathbf{y}^{2}+\partial^{2} \mathbf{u} / \partial \mathbf{z}^{2}=\mathbf{f}(\mathbf{x}, \mathbf{y}, \mathbf{z})
$$

03/15/2016
CS267 Lecture 17
4

## Explicit Solution of PDEs

## ${ }^{\circ}$ Often used for hyperbolic PDEs

${ }^{\circ}$ Stability limits size of time-step
${ }^{\circ}$ Computation corresponds to

- Matrix vector multiply
- Combine nearest neighbors on grid
${ }^{\circ}$ Use finite differences with $u[j, i]$ as the solution at
- time $t=i^{*} \delta(i=0,1,2, \ldots)$ and
- position $x=j * h(j=0,1, \ldots, N=1 / h)$
- initial conditions on $u[j, 0]$
- boundary conditions on $\mathrm{u}[0, \mathrm{i}]$ and $\mathrm{u}[\mathrm{N}, \mathrm{i}]$

。 At each timestep $\mathbf{i}=\mathbf{0 , 1 , 2}, \ldots$ For $\mathrm{j}=1$ to $\mathrm{N}-1$
$u[j, i+1]=z^{*} u[j-1, i]+\left(1-2^{*} z\right)^{*} u[j, i]$
$+z^{*} u[j+1, i]$
where $z=C * / h^{2}$


Matrix View of Explicit Method for Heat Equation

- $u[j, i+1]=z^{*} u[j-1, i]+\left(1-2^{*} z\right)^{*} u[j, i]+z^{*} u[j+1, i]$
- $u[:, i+1]=T^{*} u[:, i]$ where $T$ is tridiagonal:


$$
\stackrel{\bullet}{\hookleftarrow}
$$

- L called Laplacian (in 1D)
- For a 2D mesh (5 point stencil) the Laplacian has 5 diagonals
- For a 3D mesh there are 7 diagonals 03/15/2016

CS267 Lecture 17

## 2D Poisson's equation

${ }^{\circ}$ Similar to the 1D case, but the matrix $T$ is now

## ${ }^{\circ}$ 3D is analogous

03/15/2016
CS267 Lecture 17
8

| Algorithms for 2D (3D) Poisson Equation ( $\mathrm{N}=\mathrm{n}^{2}\left(\mathrm{n}^{3}\right)$ vars) |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| Algorithm | Serial | PRAM | Memory | \#Procs |
| ${ }^{\circ}$ Dense LU | $\mathrm{N}^{3}$ | N | $\mathrm{N}^{2}$ | $\mathrm{N}^{2}$ |
| - Band LU | $\mathrm{N}^{2}\left(\mathrm{~N}^{7 / 3}\right)$ | N | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{5 / 3}\right)$ | $\mathrm{N}\left(\mathrm{N}^{4 / 3}\right)$ |
| - Jacobi | $\mathrm{N}^{2}\left(\mathrm{~N}^{5 / 3}\right)$ | $\mathrm{N}\left(\mathrm{N}^{2 / 3}\right)$ | N | N |
| - Explicit Inv. | $\mathrm{N}^{2}$ | $\boldsymbol{\operatorname { l o g }} \mathbf{N}$ | $\mathrm{N}^{2}$ | $\mathrm{N}^{2}$ |
| ${ }^{\circ}$ Conj.Gradient | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{4 / 3}\right)$ | $\mathrm{N}^{1 / 2(1 / 3)} \boldsymbol{\operatorname { l o g }} \mathrm{N}$ | N | N |
| ${ }^{\circ}$ Red/Black SOR | R $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{4 / 3}\right)$ | $\mathrm{N}^{1 / 2}\left(\mathrm{~N}^{1 / 3}\right)$ | N | N |
| ${ }^{\circ}$ Sparse LU | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{2}\right)$ | $\mathrm{N}^{1 / 2}$ | $\mathrm{N}^{*} \log \mathrm{~N}\left(\mathrm{~N}^{4 / 3}\right)$ | N |
| ${ }^{\circ} \mathrm{FFT}$ | $N^{*} \log N$ | $\boldsymbol{\operatorname { l o g }} \mathrm{N}$ | N | N |
| ${ }^{\circ}$ Multigrid | N | $\log ^{2} \mathrm{~N}$ | N | N |
| - Lower bound | N | $\log N$ | N |  |

PRAM is an idealized parallel model with zero cost communication Reference: James Demmel, Applied Numerical Linear Algebra, SIAM, 1997.

$$
03 / 15 / 2016
$$

CS267 Lecture 17


## Convergence of Nearest Neighbor Methods

- Jacobi' s method involves nearest neighbor computation on $n x n$ grid ( $N=\mathbf{n}^{2}$ )
- So it takes $\mathbf{O}(\mathrm{n})=\mathbf{O}(\mathrm{sqrt}(\mathrm{N})$ ) iterations for information to propagate
${ }^{\circ}$ E.g., consider a rhs (b) that is 0 , except the center is 1
- The exact solution looks like:


03/15/2016

CS267 Lecture 17
10

## Jacobi's Method

${ }^{\circ}$ To derive Jacobi's method, write Poisson as: $\mathbf{u}(\mathrm{i}, \mathrm{j})=(\mathbf{u}(\mathrm{i}-1, \mathrm{j})+\mathbf{u}(\mathrm{i}+1, \mathrm{j})+\mathbf{u}(\mathrm{i}, \mathrm{j}-1)+\mathbf{u}(\mathrm{i}, \mathrm{j}+1)+\mathrm{b}(\mathrm{i}, \mathrm{j})) / 4$
${ }^{\circ}$ Let $\mathbf{u}(\mathrm{i}, \mathrm{j}, \mathrm{m})$ be approximation for $\mathbf{u}(\mathrm{i}, \mathrm{j})$ after m steps $u(i, j, m+1)=(u(i-1, j, m)+u(i+1, j, m)+u(i, j-1, m)+$

$$
u(i, j+1, m)+b(i, j)) / 4
$$

${ }^{\circ}$ I.e., $\mathbf{u}(\mathrm{i}, \mathrm{j}, \mathrm{m}+1)$ is a weighted average of neighbors
${ }^{\circ}$ Motivation: $u(i, j, m+1)$ chosen to exactly satisfy equation at (i,j)
${ }^{\circ}$ Steps to converge proportional to problem size, $\mathbf{N}=\mathbf{n}^{2}$
${ }^{\circ}$ Therefore, serial complexity is $\mathbf{O}\left(\mathrm{N}^{2}\right)$

03/15/2016
CS267 Lecture 17
11

Convergence of Nearest Neighbor Methods


True Solution


Takes $O(n)$ steps to propagate information across an nxn grid

## Communication Avoiding Jacobi:

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$


## A3.x. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .

A²:x. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
A.x• . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .

1234
.. 32

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$
- Like Matrix-Powers Kernel, but simpler:
- Don't need to store A explicitly (it's Jacobi)
- Only need to save $A^{k} x$

Parallelizing Jacobi's Method
${ }^{\circ}$ Reduces to sparse-matrix-vector multiply by (nearly) T $U(m+1)=(T / 4-I) * U(m)+B / 4$
${ }^{\circ}$ Each value of $U(m+1)$ may be updated independently - keep 2 copies for iterations $m$ and $m+1$
${ }^{\circ}$ Requires that boundary values be communicated

- each processor owns $n^{2} / p$ elements to update
- amount of data communicated, $n / p^{1 / 2}$ per neighbor, relatively small if $n \gg p$


Want to take s>>1 iterations
All the communication-avoiding techniques for Matrix-powers kernel (i.e. repeated SpMV s)
from Lecture 15 may be used
Reduce communication cost of s iterations to 1 iteration ${ }^{14}$

## Communication Avoiding Jacobi:

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
A3.x. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .

$$
\text { A2 } \mathrm{A}^{-} \text {- }
$$

A•x • - . . . . . . . . . . . . . . . . . . . . . . . . . . . .
x • . . . . . . . . . . . . . . . . . . . . . . . . . . . .

1234 ..
.. 32

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$
- Like Matrix-Powers Kernel, but simpler:
- Don't need to store A explicitly (it's Jacobi)
- Only need to save $A^{k} x$



## Communication Avoiding Jacobi:

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Sequential Algorithm

- Example: A tridiagonal, $n=32, k=3$
- Like Matrix-Powers Kernel, but simpler:
- Don't need to store A explicitly (it's Jacobi)
- Only need to save $A^{k} x$ - move $O(n)$ words instead of $\mathcal{O}(\mathrm{kn})$


## Communication Avoiding Jacobi:

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Parallel Algorithm

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$
- Like Matrix-Powers Kernel, but simpler:
- Don't need to store A explicitly (it's Jacobi)
- Only need to save $A^{k} x$


## Communication Avoiding Jacobi:

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Parallel Algorithm

```
A3.x• . . . . . . . . . . . . . . . P P . . . . . . . . . . 
```




```
x ...............................
1234
```

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$
- Like Matrix-Powers Kernel, but simpler:
- Don't need to store A explicitly (it's Jacobi)
- Only need to save $A^{k} x$


## Communication Avoiding Jacobi:

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Parallel Algorithm

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$
- Like Matrix-Powers Kernel, but simpler:
- Don't need to store A explicitly (it's Jacobi)
- Only need to save $A^{k} x$

Remotely Dependent Entries for [x,Ax, $\left.A^{2} x, A^{3} x\right]$, 2D Laplacian

## References for Optimizing Stencils (1/2)

## ${ }^{\circ}$ Bebop.cs.berkeley.edu

- "Autotuning Stencil Codes for Cache-Based Multicore Platforms", K. Datta, UCB PhD thesis, 2009,
- "Avoiding Communication in Computing Krylov Subspaces,"
J. Demmel, M. Hoemmen, M. Mohiyuddin, K. Yelick, 2007
- "Optimization and Performance Modeling of Stencil Computations on "Optimization and Performance Modeling of Stencil Computations on
Modern Microprocessors", K. Datta, S. Kamil, S. Williams, L. Oliker, Modern Microprocessors", K. Datta, S.
J.Shalf, K. Yelick, SIAM Review, 2008
${ }^{\circ}$ SEJITS - sejits.org (Armando Fox et al @ UCB) "Bringing parallel performance to python with domainspecific selective embedded just-in-time specialization"
${ }^{\circ}$ Pochoir - stencil compiler (Charles Leiserson @ MIT) people.csail.mit.edu/yuantang/
${ }^{\circ}$ Autotuning stencils and multigrid (Mary Hall @ Utah) super-scidac.org/
${ }^{\circ}$ Polyhedral tiling (Michelle Strout @ Colorado)
www.cs.colostate.edu/~mstrout/Papers/pubs-poly.php


## Improvements to Jacobi

${ }^{\circ}$ Similar to Jacobi: $u(i, j, m+1)$ will be computed as a linear combination of neighbors

- Numeric coefficients and update order are different
${ }^{\circ} 2$ improvements
- Use "most recent values" of $u$ that are available, since these are probably more accurate
- Update value of $u(m+1)$ "more aggressively" at each step
${ }^{\circ}$ First, note that while evaluating sequentially - $u(i, j, m+1)=(u(i-1, j, m)+u(i+1, j, m) .$.
some of the values for $\mathbf{m + 1}$ are already available
$\cdot \mathrm{u}(\mathrm{i}, \mathrm{j}, \mathrm{m}+1)=(\mathrm{u}(\mathrm{i}-1, \mathrm{j}$, latest $)+\mathrm{u}(\mathrm{i}+1, \mathrm{j}$, latest $) ..$
where latest is either $m$ or $m+1$

03/15/2016
CS267 Lecture 17
${ }^{\circ}$ Ian Foster et al, on grids (SC2001)

- "Efficient out-of-core algorithms for linear relaxation using blocking covers," C. Leiserson, S. Rao, S. Toledo FOCS, 1993
" Data flow and storage allocation for the PDQ-5 program on the Philco-2000," C. Pfeifer, CACM, 1963

Gauss-Seidel
${ }^{\circ}$ Updating left-to-right row-wise order, we get the Gauss-Seidel algorithm

## for $\mathrm{i}=1$ to n

for $\mathrm{j}=1$ to n
$u(i, j, m+1)=(u(i-1, j, m+1)+u(i+1, j, m)+u(i, j-1, m+1)+u(i, j+1, m)$
$+b(i, j)) / 4$

${ }^{\circ}$ Cannot be parallelized, because of dependencies 28

## Gauss-Seidel

```
Updating left-to-right row-wise order, we get the
Gauss-Seidel algorithm
    for \(\mathrm{i}=1\) to n
    for \(\mathrm{j}=1\) to n
        \(u(i, j, m+1)=(u(i-1, j, m+1)+u(i+1, j, m)+u(i, j-1, m+1)+u(i, j+1, m)\)
                        \(+b(i, j)) / 4\)
```

Cannot be parallelized, because of dependencies, so instead we use a "red-black" order forall black points $\mathbf{u}(\mathrm{i}, \mathrm{j})$ $u(i, j, m+1)=(u(i-1, j, m)+\ldots$ red neighbors forall red points $u(i, j)$ $\mathrm{u}(\mathrm{i}, \mathrm{j}, \mathrm{m}+1)=(\mathrm{u}(\mathrm{i}-1, \mathrm{j}, \mathrm{m}+1)+\ldots$ black neighbors

${ }^{\circ}$ For general graph, use "graph coloring" Can use repeated Maximal Independent Sets to color Can use repeated Maximal Independent Sets to color
Graph $(\mathrm{T})$ is bipartite => 2 colorable (red and black) - Nodes for each color can be updated simultaneously - Same optimizations, using submatrices

Successive Overrelaxation (SOR)
Red-black Gauss-Seidel converges twice as fast as Jacobi, bu there are twice as many paraliel steps, so the same in practice

- To motivate next improvement, write basic step in algorithm as: $\mathbf{u}(\mathrm{i}, \mathrm{j}, \mathrm{m}+1)=\mathbf{u}(\mathrm{i}, \mathrm{j}, \mathrm{m})+\operatorname{correction}(\mathrm{i}, \mathrm{j}, \mathrm{m})$
If "correction" is a good direction to move, then one should move even further in that direction by some factor $w>1$
$u(i, j, m+1)=u(i, j, m)+w * \operatorname{correction}(i, j, m)$

Called successive overrelaxation (SOR)

- Parallelizes like Jacobi

Can prove $w=\mathbf{2} /(\mathbf{1 + \operatorname { s i n } ( \pi / ( n + 1 ) )})$ for best convergence for Poisson - Number of steps to converge $=$ parallel complexity $=\mathrm{O}(\mathrm{n})$, instead of $\mathrm{O}\left(\mathrm{n}^{2}\right)$ for

- Serial complexity $O\left(n^{3}\right)=O\left(N^{3 / 2}\right)$, instead of $O\left(n^{4}\right)=O\left(N^{2}\right)$ for Jacob


## 2D Poisson's equation

${ }^{\circ}$ Similar to the 1D case, but the matrix $T$ is now

${ }^{\circ}$ 3D is analogous
03/15/2016
CS267 Lecture 17
32

| Algorithms for 2D (3D) Poisson Equation ( $\mathrm{N}=\mathrm{n}^{2}\left(\mathrm{n}^{3}\right)$ vars) |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| Algorithm | Serial | PRAM | Memory | \#Procs |
| ${ }^{\circ}$ Dense LU | $\mathrm{N}^{3}$ | N | $\mathrm{N}^{2}$ | $\mathrm{N}^{2}$ |
| - Band LU | $\mathrm{N}^{2}\left(\mathrm{~N}^{7 / 3}\right)$ | N | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{5 / 3}\right)$ | $\mathrm{N}\left(\mathrm{N}^{4 / 3}\right)$ |
| - Jacobi | $\mathrm{N}^{2}\left(\mathrm{~N}^{5 / 3}\right)$ | $\mathrm{N}\left(\mathrm{N}^{2 / 3}\right)$ | N | N |
| - Explicit Inv. | $\mathrm{N}^{2}$ | $\boldsymbol{\operatorname { l o g }} \mathrm{N}$ | $\mathrm{N}^{2}$ | $\mathrm{N}^{2}$ |
| ${ }^{\circ}$ Conj.Gradients | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{4 / 3}\right)$ | $\mathrm{N}^{1 / 2(1 / 3)} \boldsymbol{\operatorname { l o g }} \mathrm{N}$ | N | N |
| ${ }^{\circ}$ Red/Black SOR | R $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{4 / 3}\right)$ | $\mathrm{N}^{1 / 2}\left(\mathrm{~N}^{1 / 3}\right)$ | N | N |
| ${ }^{\circ}$ Sparse LU | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{2}\right)$ | $\mathrm{N}^{1 / 2}$ | $\mathbf{N}^{*} \log \mathrm{~N}\left(\mathrm{~N}^{4 / 3}\right)$ | N |
| - FFT | $N^{*} \log N$ | $\log N$ | N | N |
| - Multigrid | N | $\log ^{2} \mathrm{~N}$ | N | N |
| - Lower bound | N | $\log N$ | N |  |

PRAM is an idealized parallel model with zero cost communication Reference: James Demmel, Applied Numerical Linear Algebra, SIAM, 1997. 03/15/2016 CS267 Lecture 17 ${ }^{33}$

## Multigrid Motivation

${ }^{\circ}$ Recall that Jacobi, SOR, CG, or any other sparse matrix-vector-multiply-based algorithm can only move information one grid cell at a time

- Take at least $\mathbf{n}$ steps to move information across $\mathbf{n} \times \mathrm{n}$ grid
- Therefore, converging in $\mathbf{O}(1)$ steps requires moving information across grid faster than to one neighboring grid cell per step
- One step can't just do sparse-matrix-vector-multiply

| Algorithms for 2D (3D) Poisson Equation ( $\mathrm{N}=\mathrm{n}^{2}\left(\mathrm{n}^{3}\right)$ vars) |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  | Algorithm | Serial | PRAM | Memory | \#Procs |
|  | ${ }^{\circ}$ Dense LU | $\mathrm{N}^{3}$ | N | $\mathrm{N}^{2}$ | $\mathrm{N}^{2}$ |
|  | - Band LU | $\mathrm{N}^{2}\left(\mathrm{~N}^{7 / 3}\right)$ | N | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{5 / 3}\right)$ | $\mathrm{N}\left(\mathrm{N}^{4 / 3}\right)$ |
|  | - Jacobi | $\mathrm{N}^{2}\left(\mathrm{~N}^{5 / 3}\right)$ | $\mathrm{N}\left(\mathrm{N}^{2 / 3}\right)$ | N | N |
|  | - Explicit Inv. | $\mathrm{N}^{2}$ | $\log N$ | $\mathrm{N}^{2}$ | $\mathrm{N}^{2}$ |
|  | - Conj.Gradien | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{4 / 3}\right)$ | $\mathrm{N}^{1 / 2(1 / 3)} \boldsymbol{\operatorname { l o g }} \mathrm{N}$ | N | N |
|  | - Red/Black SOR | $\mathrm{R} \mathrm{N}^{3 / 2}\left(\mathrm{~N}^{4 / 3}\right)$ | $\mathrm{N}^{1 / 2}\left(\mathrm{~N}^{1 / 3}\right)$ | N | N |
|  | - Sparse LU | $\mathrm{N}^{3 / 2}\left(\mathrm{~N}^{2}\right)$ | $\mathrm{N}^{1 / 2}$ | $N^{*} \log N\left(N^{4 / 3}\right)$ | N |
|  | - FFT | $N^{*} \log N$ | $\log N$ | N | N |
|  | - Multigrid | N | $\log ^{2} \mathrm{~N}$ | N | N |
|  | - Lower bound | N | $\boldsymbol{\operatorname { l o g }} N$ | N |  |
| PRAM is an idealized parallel model with zero cost communication |  |  |  |  |  |
| Reference: James Demmel, Applied Numerical Linear Algebra, SIAM, 1997. |  |  |  |  |  |
| 03/15/2016 |  |  | CS267 Lecture 17 |  | 34 |



Big Idea used in multigrid and elsewhere
${ }^{\circ}$ If you are far away, problem looks simpler

- For gravity: approximate earth, distant galaxies, ... by point masses
${ }^{\circ}$ Can solve such a coarse approximation to get an approximate solution, iterating if necessary
- Solve coarse approximation problem by using an even coarser approximation of it, and so on recursively
${ }^{\circ}$ Ex: Graph Partitioning (used to parallelize SpMV)
- Replace graph to be partitioned by a coarser graph
${ }^{\circ}$ Ex: Multigrid for solving PDE in O(n) time
- Use coarser mesh to get approximate solution of Poisson's Eq.
${ }^{\circ}$ Ex: Fast Multipole Method, Barnes-Hut for computing gravitational forces on $n$ particles in $\mathbf{O}(\mathrm{n} \log \mathrm{n})$ time:
- Approximate particles in box by total mass, center of gravity


## Multigrid Overview

${ }^{\circ}$ Basic Algorithm:

- Replace problem on fine grid by an approximation on a coarser grid
- Solve the coarse grid problem approximately, and use the solution as a starting guess for the fine-grid problem, which is then iteratively updated
- Solve the coarse grid problem recursively, i.e. by using a still coarser grid approximation, etc.
${ }^{\circ}$ Success depends on coarse grid solution being a good approximation to the fine grid


03/15/2016
CS267 Lecture 17

Fine and Coarse Approximations


Fine


Coarse


03/15/2016
CS267 Lecture 17

## Multigrid Sketch in 1D

- Consider a $2^{m+1}$ grid in 1D for simplicity
- Let $\mathbf{P}^{(i)}$ be the problem of solving the discrete Poisson equation on a $2^{i+1}$ grid in 1D (2i-1 unknowns plus 2 boundaries)
- Write linear system as $T(i) * x(i)=b(i)$
- $\mathbf{P}^{(\mathrm{m})}, \mathbf{P}^{(\mathrm{m}-1)}, \ldots, \mathbf{P}^{(1)}$ is sequence of problems from finest to coarsest

| $\begin{array}{lllll}2 & 2 & 2 & 2 & 2\end{array}$ | $1 \begin{array}{lll}1 & 1\end{array}$ |  |
| :---: | :---: | :---: |
| $\mathrm{P}^{(3)}: 1 \mathrm{D}$ grid of 9 points | $\mathrm{P}^{(2)}: 1 \mathrm{D}$ grid of 5 points | $\mathrm{P}^{(1)}: 1 \mathrm{D}$ grid of 3 points |
| 7 unknowns | 3 unknowns | 1 unknown |
| Points labeled 2 are part of next coarser grid | Points labeled 1 are part of next coarser grid |  |

## Multigrid Sketch in 2D

- Consider a $2^{\mathrm{m}+1}$ by $2^{\mathrm{m}+1}$ grid in 2 D
- Let $\mathbf{P}^{(i)}$ be the problem of solving the discrete Poisson equation on a $2^{i+1}$ by $2^{i+1}$ grid in 2D
-Write linear system as $T(i)$ * $x(i)=b(i)$
- $P^{(m)}, P^{(m-1)}, \ldots, P^{(1)}$ is sequence of problems from finest to coarsest

$\mathbf{P}^{(3)}: 9$ by 9 grid of points 7 by 7 grid of unkno

Pints labeled 2 are part of next coarser grid

$\mathbf{P}^{(2)}: 5$ by 5 grid of points 3 by 3 grid of unknowns Points labeled 1 are
part of next coarser gri

$\mathbf{P}^{(1)}: 3$ by 3 grid of points 1 by 1 grid of unknowns

03/15/2016
CS267 Lecture 17

## Multigrid Operators

- For problem $P^{(i)}$ at varying coarsening levels ( $\mathbf{i}$, grid size grows with $i$ ):
- b(i) is the Right Hand Side (RHS) and
- $x(i)$ is the current estimated solution
- All the following operators just average values on neighboring grid points (so information moves fast on coarse grids)
- The restriction operator $\mathbf{R}(\mathrm{i})$ maps $\mathbf{P}^{(\mathbf{i})}$ to $\mathbf{P}^{(\mathrm{i}-1)}$
- Restricts problem on fine grid $\mathrm{P}^{(\mathbf{i})}$ to coarse grid $\mathrm{P}^{(\mathrm{i}-1)}$
- Uses sampling or averaging
- $b(i-1)=R(i)(b(i))$
- The interpolation operator $\ln (\mathrm{i}-1)$ maps approx. solution $\mathrm{x}(\mathrm{i}-1)$ to $\mathrm{x}(\mathrm{i})$ Interpolates solution on coarse grid $\mathrm{P}^{(\mathrm{i}-1)}$ to fine grid $\mathrm{P}^{(i)}$
$x(i)=\ln (i-1)(x(i-1))$
- The solution operator $\mathrm{S}(\mathrm{i})$ takes $\mathrm{P}^{(\mathrm{i})}$ and improves solution $\mathrm{x}(\mathrm{i})$ - Uses "weighted" Jacobi or SOR on single level of grid
$\mathrm{x}_{\mathrm{improved}}(\mathrm{i})=\mathrm{S}(\mathrm{i})(\mathrm{b}(\mathrm{i}), \mathrm{x}(\mathrm{i})$ )
- Overall algorithm, then details of operators

03/15/2016
CS267 Lecture 17
43

Multigrid Operators

- For problem $P^{(i)}$ at varying coarsening levels ( $i$, grid size grows with $i$ ):
- $\mathbf{b}(\mathbf{i})$ is the Right Hand Side (RHS) and $\}$ both live on grids of size 2i-1
- All the following operators just average values on neighboring grid points (so information moves fast on coarse grids)
- The restriction operator $\mathrm{R}(\mathrm{i})$ maps $\mathrm{P}^{(\mathrm{i})}$ to $\mathrm{P}^{(\mathrm{i}-1)}$

Restricts problem on fine grid $\mathrm{P}^{(\mathrm{i})}$ to coarse grid $\mathrm{P}^{(\mathrm{i}-1)}$

- Uses sampling or averaging
b(i-1) $=R(i)(b(i))$
- The interpolation operator $\ln (\mathrm{i}-1)$ maps approx. solution $\mathrm{x}(\mathrm{i}-1)$ to $\mathrm{x}(\mathrm{i})$ - Interpolates solution on coarse grid $\mathrm{P}^{(\mathrm{i}-1)}$ to fine grid $\mathrm{P}^{(\mathrm{i})}$

$$
x(i)=\ln (i-1)(x(i-1))
$$

- The solution operator $\mathrm{S}(\mathrm{i})$ takes $\mathrm{P}^{(\mathrm{i})}$ and improves solution $\mathrm{x}(\mathrm{i})$
- Uses "weighted" Jacobi or SOR on single level of grid
- $\mathrm{x}_{\text {improved }}(\mathrm{i})=\mathrm{s}(\mathrm{i})(\mathrm{b}(\mathrm{i}), \mathrm{x}(\mathrm{i})$
- Overall algorithm, then details of operators

03/15/2016
CS267 Lecture 17

## Multigrid Operators

- For problem $\mathbf{P}^{(\mathbf{i})}$ at varying coarsening levels ( $\mathbf{i}$, grid size grows with i ):
- b(i) is the Right Hand Side (RHS) and both live on grids of size 2i-1
$x(i)$ is the current estimated solutio
- All the following operators just average values on neighboring grid points (so information moves fast on coarse grids)
- The restriction operator $\mathrm{R}(\mathrm{i})$ maps $\mathrm{P}^{(\mathrm{i})}$ to $\mathrm{P}^{(\mathrm{i}-1)}$
- Restricts problem on fine grid $\mathrm{P}^{(\mathrm{i})}$ to coarse grid $\mathrm{P}^{(\mathrm{i}-1)}$

Uses sampling or averaging
$b(i-1)=R(i)(b(i))$

- The interpolation operator $\ln (\mathrm{i}-1)$ maps approx. solution $\mathbf{x}(\mathrm{i}-1)$ to $\mathbf{x}(\mathrm{i})$ - Interpolates solution on coarse grid $\mathrm{P}^{(\mathrm{i}-1)}$ to fine grid $\mathrm{P}^{(\mathrm{i})}$ - $x(i)=\ln (i-1) /(x(i-1))$
- The solution operator $\mathrm{S}(\mathrm{i})$ takes $\mathrm{P}^{(\mathrm{i})}$ and improves solution $\mathrm{x}(\mathrm{i})$
- Uses "weighted" Jacobi or SOR on single level of grid
$\mathrm{x}_{\text {improved }}(\mathrm{i})=\mathrm{S}(\mathrm{i})(\mathrm{b}(\mathrm{i}), \mathrm{x}(\mathrm{i})$
- Overall algorithm, then details of operators

03/15/2016
CS267 Lecture 17
44

## Multigrid Operators

- For problem $P^{(i)}$ at varying coarsening levels ( $\mathbf{i}$, grid size grows with $i$ ):
- $b$ (i) is the Right Hand Side (RHS) and
- $x$ (i) is the current estimated solution both live on grids of size $2^{i-1}$
- All the following operators just average values on neighboring grid points (so information moves fast on coarse grids)
- The restriction operator $\mathrm{R}(\mathrm{i})$ maps $\mathrm{P}^{(\mathrm{i})}$ to $\mathrm{P}^{(\mathrm{i}-1)}$

Restricts problem on fine grid $\mathrm{P}^{(\mathrm{i})}$ to coarse grid $\mathrm{P}^{(\mathrm{i}-1)}$
Uses sampling or averaging
$b(i-1)=R(i)(b(i))$

- The interpolation operator $\ln (\mathrm{i}-1)$ maps approx. solution $\mathrm{x}(\mathrm{i}-1)$ to $\mathrm{x}(\mathrm{i})$

Interpolates solution on coarse grid $\mathrm{P}^{(\mathrm{i}-1)}$ to fine grid $\mathrm{P}^{(\mathrm{i})}$
$x(i)=\ln (i-1)(x(i-1))$

- The solution operator $\mathrm{S}(\mathrm{i})$ takes $\mathrm{P}^{(\mathrm{i})}$ and improves solution $\mathbf{x}(\mathbf{i})$
- Uses "weighted" Jacobi or SOR on single level of grid
- $\mathrm{x}_{\text {improved }}$ (i) $=\mathbf{s}(\mathrm{i})(\mathrm{b}(\mathrm{i}), \mathbf{x}(\mathrm{i}))$
- Overall algorithm, then details of operators

03/15/2016
CS267 Lecture 17

## Multigrid V-Cycle Algorithm

## Function MGV ( b(i), x(i) )

... Solve $T(i){ }^{*} x(i)=b(i)$ given $b(i)$ and an initial guess for $x(i)$
... return an improved $\mathbf{x}(\mathrm{i})$
if ( $\mathrm{i}=1$ )
compute exact solution $\mathbf{x}(1)$ of $\mathbf{P}^{(1)} \quad$ only 1 unknown
return $\times(1)$
else
$\mathbf{x}(\mathbf{i})=\mathbf{S}(\mathbf{i})(\mathbf{b}(\mathbf{i}), \mathbf{x}(\mathrm{i}))$
$r(i)=T(i)^{*} x(i)-b(i)$
$d(i)=\ln (i-1)(\operatorname{MGV}(R(i)(r(i)), 0))$
$\mathbf{x}(\mathrm{i})=\mathbf{x}(\mathbf{i})-\mathbf{d}(\mathrm{i})$
$\mathbf{x}(\mathrm{i})=\mathbf{S}(\mathrm{i})(\mathrm{b}(\mathrm{i}), \mathbf{x}(\mathrm{i}))$
return $\mathrm{x}(\mathrm{i})$
solve recursively improve solution by damping
high frequency error compute residual solve $T(i)^{*} d(i)=r(i)$ recursively correct fine grid solution improve solution again

03/15/2016
CS267 Lecture 17
46

## This is called a V-Cycle

## ${ }^{\circ}$ Just a picture of the call graph

${ }^{\circ}$ In time a V-cycle looks like the following

## Mul tigid V -cyele



03/15/2016
CS267 Lecture 17

## Complexity of a V-Cycle

## ${ }^{\circ}$ On a serial machine

- Work at each point in a V-cycle is O(the number of unknowns)
- Cost of Level $i$ is $\left(2^{i}-1\right)^{2}=O\left(4^{i}\right)$ for a 2 D grid
- If finest grid level is $m$, total time is:
$\sum_{i=1}^{m} O\left(4^{i}\right)=O\left(4^{m}\right)$ for a $2 D$ grid
= O(\# unknowns) in general


## ${ }^{\circ}$ On an ideal parallel machine (PRAM)

with one processor per grid point and free communication, each
step in the V-cycle takes constant time

- Total V-cycle time is $\mathbf{O}(\mathrm{m})=\mathbf{O}(\log$ \#unknowns)

03/15/2016
CS267 Lecture 17
48

## Full Multigrid (FMG)

- Intuition:
- improve solution by doing multiple V-cycles
- avoid expensive fine-grid (high frequency) cycles
- analysis of why this works is beyond the scope of this class

Function FMG (b(m), x(m))
... return improved $x(m)$ given initial guess
compute the exact solution $x(1)$ of $P(1)$
for $\mathrm{i}=2$ to m ... from coarse to fine mesh
x(i) = MGV (b(i), In (i-1) (x(i-1) ))
${ }^{\circ}$ In other words:

- Solve the problem with 1 unknow
- Given a solution to the coarser problem, $\mathrm{P}^{(\mathrm{i}-1)}$, map it to starting guess fo P(i)
- Solve the finer problem using the Multigrid V-cycle

Full Multigrid Cost Analysis Ful Mult igidi Cyele


## ${ }^{\circ}$ One V for each call to FMG

- people also use Ws and other compositions
${ }^{\circ}$ Serial time: $\Sigma_{i=1}^{m} O\left(4^{i}\right)=O\left(4^{m}\right)=O$ (\# unknowns)
${ }^{\circ}$ PRAM time: $\Sigma^{\mathrm{m}=1} \mathrm{O}(\mathrm{i})=\mathbf{O}\left(\mathrm{m}^{2}\right)=\mathbf{O}\left(\log ^{2} \#\right.$ unknowns $)$

03/15/2016
CS267 Lecture 17

## Complexity of Solving Poisson's Equation

## Complexity of Solving Poisson's Equation

- Theorem: error after one call to multigrid
- error_after $\leq .5$ *error_before
- independent of \# unknowns
$\rightarrow$ At least 1 bit each time
- Corollary: We can make the error < any fixed tolerance in a fixed number of steps, independent of size of finest grid
- This is the most important convergence property of MG, distinguishing it from all other methods, which converge more slowly for large grids
- Theorem: error after one FMG call
- error_after $\leq .5^{*}$ error_before
- independent of \# unknowns
$\rightarrow$ At least 1 bit each time
- Corollary: We can make the error < any fixed tolerance in a fixed number of steps, independent of size of finest grid
- This is the most important convergence property of MG, distinguishing it from all other methods, which converge more slowly for large grids

The Solution Operator S(i) - Details
${ }^{\circ}$ The solution operator, $\mathrm{S}(\mathrm{i})$, is a weighted Jacobi
${ }^{\circ}$ Consider the 1D problem

${ }^{\circ}$ At level i, pure Jacobi replaces:
$x(j):=1 / 2(x(j-1)+x(j+1)+b(j))$

- Weighted Jacobi uses:

$$
x(j):=1 / 3(x(j-1)+x(j)+x(j+1)+b(j))
$$

${ }^{\circ}$ In 2D, similar average of nearest neighbors
$\square$

Weighted Jacobi chosen to damp high frequency error -, Solitat teso semes "Rough"
Lots of high frequency components Norm $=1.65$

| Error after 1 weighted Jacobi step |
| :---: |
| "Smorn | "Smoother"

Less high frequency component
Norm $=1.06$

Error after 2 weighted Jacobi steps
"Smooth" "Smooth"
Little high frequency component Norm = .92,
won't decrease much more
03/15/2016
CS267 Lecture 17

## The Restriction Operator R(i) - Details

- The restriction operator, $R(i)$, takes
- a problem $P^{(i)}$ with RHS $b(i)$ and
- maps it to a coarser problem $\mathrm{P}^{(\mathrm{i}-1)}$ with RHS $\mathrm{b}(\mathrm{i}-1)$
${ }^{\circ}$ In 1D, average values of neighbors
$\cdot x_{\text {coarse }}(i)=1 / 4 * x_{\text {fine }}(i-1)+1 / 2 * x_{\text {fine }}(i)+1 / 4 * x_{\text {fine }}(i+1)$


${ }^{\circ}$ In 2D, average with all 8 neighbors (N,S,E,W,NE,NW,SE,SW)
03/15/2016 CS267 Lecture 17
56



58


03/15/2016
CS267 Lecture 17
Convergence Picture of Multigrid in 1D


## Convergence Picture of Multigrid in 2D



Right Hand Side


3/15/2016
CS267 Lecture 17

Performance Model of parallel 2D Multigrid (V-cycle)
${ }^{\circ}$ Assume $2^{m+1}$ by $2^{m+1}$ grid of points, $n=2^{m}-1, N=n^{2}$
${ }^{\circ}$ Assume $p=4^{k}$ processors, arranged in $2^{k}$ by $2^{k}$ grid - Processors start with $2^{m-k}$ by $2^{m-k}$ subgrid of unknowns
${ }^{\circ}$ Consider V-cycle starting at level m

- At levels $m$ through $k$ of V -cycle, each processor does some work
- At levels $k-1$ through 1 , some processors are idle, because a $2^{k-1}$ by $2^{k-1}$ grid of unknowns cannot occupy each processor
- Cost of one level in V-cycle
- If level $\mathrm{j}>\mathrm{k}$, then cost $=$
$\mathrm{O}\left(4^{-\mathrm{k}}\right) \quad$.... Flops, proportional to the number of grid points/processor $+\mathrm{O}(1) \alpha \quad$.... Send a constant \# messages to neighbors $+O\left(2^{j-k}\right) \beta \quad$.... Number of words sent
- If level $\mathrm{j}<\mathrm{k}$, then cost =
$\mathrm{O}(1) \quad$.... Flops, proportional to the number of grid points/processor
$+\mathrm{O}(1) \quad \alpha \quad \ldots$. Send a constant \# messages to neighbors
$+\mathrm{O}(1) \beta \quad$.... Number of words sent
${ }^{\circ}$ Sum over all levels in all V-cycles to get complexity 03/15/2016

CS267 Lecture 17

Comparison of Methods (in O(.) sense)

|  | \# Flops | \# Messages | \# Words sent |
| :--- | :--- | :--- | :--- |
| MG | $N / p+$ <br> $\log p^{*} \log N$ | $(\log N)^{2}$ | $(N / p)^{1 / 2}+$ <br> $\log p^{*} \log N$ |
| FFT | $N \log N / p$ | $p^{1 / 2}$ | $N / p$ |
| SOR | $N^{3 / 2} / p$ | $N^{1 / 2}$ | $N / p$ |

${ }^{\circ}$ SOR is slower than others on all counts
${ }^{\circ}$ Flops for MG depends on accuracy of MG
${ }^{\circ}$ MG communicates less total data (bandwidth)

- Total messages (latency) depends ...
- This coarse analysis can't say whether MG or FFT is better when $\underset{03 / 15 / 2016}{\alpha \gg} \beta$

CS267 Lecture 17

## Multigrid on an Adaptive Mesh

${ }^{\circ}$ In practice, we don't go all the way to $\mathrm{P}^{(1)}$
${ }^{\circ}$ In sequential code, the coarsest grids are negligibly cheap, but on a parallel machine they are not.

- Consider 1000 points per processor, so flops $=\mathbf{O}(1000)$
- In 2D, the surface to communicate is $4 \times 1000^{1 / 2} \sim=128$, or $13 \%$
- In 3D, the surface is $1000-8^{3} \sim 500$, or $50 \%$
- See Tuminaro and Womble, SIAM J. Sci. Comp.,
v14, n5, 1993 for analysis of MG on 1024 nCUBÉ2
- on $64 \times 64$ grid of unknowns, only 4 per processor
efficiency of 1 V -cycle was .02 , and on FMG . 008
- on $1024 \times 1024$ grid
efficiencies were .7 (MG V-cycle) and . 42 (FMG) although worse parallel efficiency, FMG is 2.6 times faster than V-cycles alone
- nCUBE had fast communication, slow processors
- Today: Same problem in Chombo @ LBL
- Communication of coarsest meshes
${ }^{\circ}$ For problems with very large dynamic range, another level of refinement is needed
${ }^{\circ}$ Build adaptive mesh and solve multigrid (typically) at each level

${ }^{\circ}$ Can't afford to use finest mesh everywhere

03/15/2016
CS267 Lecture 17
64

Multiblock Applications
${ }^{\circ}$ Solve system of equations on a union of rectangles - subproblem of AMR
${ }^{\circ}$ E.g.,


## Support for AMR

${ }^{\circ}$ Domains in Titanium designed for this problem
${ }^{\circ}$ Kelp, Boxlib, and AMR++ are libraries for this
${ }^{\circ}$ Primitives to help with boundary value updates, etc.


03/15/2016
CS267 Lecture 17
67

Multigrid on an Unstructured Mesh
Adaptive Mesh Refinement
${ }^{\circ}$ Data structures in AMR
${ }^{\circ}$ Usual parallelism is to assign grids on each level to processors
${ }^{\circ}$ Load balancing is a problem


03/15/2016
66

Another approach to variable activity is to use an unstructured mesh that is more refined in areas of interest
${ }^{\circ}$ Adaptive rectangular or unstructured?

- Numerics easier on rectangular
- Supposedly easier to implement (arrays without indirection) but boundary cases tend to dominate
code


Multigrid on an Unstructured Mesh
${ }^{\circ}$ Need to partition graph for parallelism
${ }^{\circ}$ What does it mean to do Multigrid anyway?
${ }^{\circ}$ Need to be able to coarsen grid (hard problem)

- Can't just pick "every other grid point" anymore
- Use "maximal independent sets" again
- How to make coarse graph approximate fine one
${ }^{\circ}$ Need to define $\mathbf{R}()$ and $\operatorname{In}()$
- How do we convert from coarse to fine mesh and back?


## ${ }^{\circ}$ Need to define $\mathbf{S}()$

- How do we define coarse matrix (no longer formula, like Poisson)


## - Dealing with coarse meshes efficiently

- Should we switch to using fewer processors on coarse meshes?
- Should we switch to another solver on coarse meshes? 03/15/2016 CS267 Lecture 17

Example of Prometheus meshes


Study failure modes of trabecular bone under stress


[^1]03/15/2016
CS267 Lecture 17



[^0]:    ${ }^{\circ}$ Review of Poisson Equation

    - Jacobi's method
    ${ }^{\circ}$ Red-Black SOR method
    ${ }^{\circ}$ Conjugate Gradient (topic of Lecture 15)
    ${ }^{\circ}$ Multigrid
    ${ }^{\circ}$ (Later lecture: FFTs)

[^1]:    Source: M. Adams, H. Bayraktar, T. Keaveny, P. Papadopoulos, A. Gupta

