Parallel Numerical Algorithms
Chapter 7 – Differential Equations
Section 7.2 – Partial Differential Equations

Michael T. Heath

Department of Computer Science
University of Illinois at Urbana-Champaign

CS 554 / CSE 512

Outline

Outline

Types of Partial Differential Equations

- PDEs are typically classified in three major types:
    - Elliptic — $\nabla^2 u(\vec{x}) = f(\vec{x})$
    - Parabolic — $\partial u(\vec{x}, t)/\partial t = k\nabla^2 u(\vec{x}, t)$
    - Hyperbolic — $\partial u(\vec{x}, t)/\partial t = \partial f(u(\vec{x}, t))/\partial \vec{x}$
- PDEs in practice are often mixture of these
- Numerical methods for time-dependent equations fall into two broad categories
    - Explicit methods define solution at next time step solely in terms of previous time steps
    - Implicit methods by contrast include values of current time step and hence involve solving system of linear or nonlinear equations at each time step
- We'll begin by looking at Elliptic problems. All stable methods for elliptic problems are implicit (there is no timelike direction)

Numerical Methods for Partial Differential Equations

- Partial differential equations are typically solved numerically by finite difference, finite element, finite volume, or spectral discretization

- Such discretization yields system of linear or nonlinear algebraic equations whose solution gives approximate solution to PDE

- Solving linear or nonlinear algebraic system is one major source of parallelism in solving PDEs numerically

- We will consider *domain decomposition* methods that exploit natural parallelism in PDE and its discretization

Alternating Schwarz Method

- Consider elliptic PDE

$$L\,u = f$$

on domain $\Omega = \Omega_1 \cup \Omega_2$, with boundary condition

$$u = g$$

on $\partial\,\Omega$

Alternating Schwarz Method

Given initial guess $u_2^{(0)}$ on Ω_2, for $k = 0, 1, \ldots$

- On Ω_1, solve

$$L\,u_1^{(k+1)} = f$$

with boundary conditions

$$
\begin{aligned}
u_1^{(k+1)} &= g \text{ on } \partial\,\Omega_1 \setminus \Gamma_1 \\
u_1^{(k+1)} &= u_2^{(k)} \text{ on } \Gamma_1
\end{aligned}
$$

- On Ω_2, solve

$$L\,u_2^{(k+1)} = f$$

with boundary conditions

$$
\begin{aligned}
u_2^{(k+1)} &= g \text{ on } \partial\,\Omega_2 \setminus \Gamma_2 \\
u_2^{(k+1)} &= u_1^{(k+1)} \text{ on } \Gamma_2
\end{aligned}
$$

Alternating Schwarz Method

- Alternating iterations continue until convergence to solution u on entire domain Ω

- Schwarz proposed this method in 1870 to deal with regions for which analytical solutions are not known

- Today it is of interest, in discretized form, for suggesting one of two major paradigms for solving PDEs numerically by *domain decomposition*

  - Overlapping subdomains (Schwarz)
  - Non-overlapping subdomains (Schur)

Discretized Schwarz Method

- Discretization yields $n \times n$ symmetric positive definite linear algebraic system

$$Ax = b$$

- For $i = 1, 2$, let S_i be set of n_i indices of grid points in interior of Ω_i, where $n_i = |S_i|$

- Because subdomains overlap, $S_1 \cap S_2 \neq \varnothing$ and $n_1 + n_2 > n$

- For $i = 1, 2$, let R_i be $n_i \times n$ Boolean *restriction* matrix such that for any n-vector v, $v_i = R_i v$ contains precisely those components of v corresponding to indices in S_i (i.e., nodes in Ω_i)

Discretized Schwarz Method

- Conversely, $n \times n_i$ *extension* matrix $\boldsymbol{R}_i^T$ expands n_i-vector $\boldsymbol{v}_i$ to n-vector $\boldsymbol{v}$ whose components corresponding to indices in S_i are same as those of $\boldsymbol{v}_i$, and whose remaining components are all zero

- Principal submatrices of $\boldsymbol{A}$ of order n_1 and n_2 corresponding to two subdomains are given by

$$
\begin{aligned}
\boldsymbol{A}_1 &= \boldsymbol{R}_1 \boldsymbol{A} \boldsymbol{R}_1^T \\
\boldsymbol{A}_2 &= \boldsymbol{R}_2 \boldsymbol{A} \boldsymbol{R}_2^T
\end{aligned}
$$

Discretized Schwarz Method

- For discretized problem, *alternating Schwarz iteration* takes form

$$
\begin{aligned}
\boldsymbol{x}_{k+\frac{1}{2}} &= \boldsymbol{x}_k + \boldsymbol{R}_1^T \boldsymbol{A}_1^{-1} \boldsymbol{R}_1 (\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}_k) \\
\boldsymbol{x}_{k+1} &= \boldsymbol{x}_{k+\frac{1}{2}} + \boldsymbol{R}_2^T \boldsymbol{A}_2^{-1} \boldsymbol{R}_2 (\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}_{k+\frac{1}{2}})
\end{aligned}
$$

- This method is analogous to block Gauss-Seidel, but with overlapping blocks

- Overall error $\boldsymbol{e}_k = \boldsymbol{x} - \boldsymbol{x}_k$ updated as $\boldsymbol{e}_{k+1} = \boldsymbol{G}\boldsymbol{e}_k$ where

$$
\boldsymbol{G} = (\boldsymbol{I} - \boldsymbol{R}_2^T \boldsymbol{A}_2^{-1} \boldsymbol{R}_2 \boldsymbol{A})(\boldsymbol{I} - \boldsymbol{R}_1^T \boldsymbol{A}_1^{-1} \boldsymbol{R}_1 \boldsymbol{A})
$$

so this is known as *multiplicative Schwarz method*

Discretized Schwarz Method

- We have as yet achieved no parallelism, since two subproblems must be solved sequentially for each iteration, but instead of Gauss-Seidel, we can use the block Jacobi approach

$$
\begin{aligned}
\boldsymbol{x}_{k+\frac{1}{2}} &= \boldsymbol{x}_k + \boldsymbol{R}_1^T \boldsymbol{A}_1^{-1} \boldsymbol{R}_1 (\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}_k) \\
\boldsymbol{x}_{k+1} &= \boldsymbol{x}_{k+\frac{1}{2}} + \boldsymbol{R}_2^T \boldsymbol{A}_2^{-1} \boldsymbol{R}_2 (\boldsymbol{b} - \boldsymbol{A}\boldsymbol{x}_k)
\end{aligned}
$$

whose subproblems can be solved simultaneously

- With either Gauss-Seidel or Jacobi version, it can be shown that iteration converges at rate independent of mesh size, provided overlap area between subdomains is sufficiently large (and mesh is refined uniformly)

Discretized Schwarz Method

- Eliminating $x_{k+\frac{1}{2}}$ in Jacobi version, we obtain

$$x_{k+1} = x_k + (R_1^T A_1^{-1} R_1 + R_2^T A_2^{-1} R_2)(b - A x_k)$$

which is just Richardson iteration with *additive Schwarz preconditioner*

$$M^{-1} = R_1^T A_1^{-1} R_1 + R_2^T A_2^{-1} R_2$$

- Symmetry of preconditioned system means it can be used in conjunction with conjugate gradient method to accelerate convergence

Discretized Schwarz Method

- Multiplicative Schwarz iteration matrix is not symmetric, but can be made symmetric by additional step with A_1^{-1} each iteration

$$
\begin{aligned}
\boldsymbol{x}_{k+\frac{1}{3}} &= \boldsymbol{x}_k + \boldsymbol{R}_1^T \boldsymbol{A}_1^{-1} \boldsymbol{R}_1 (\boldsymbol{b} - \boldsymbol{A} \boldsymbol{x}_k) \\
\boldsymbol{x}_{k+\frac{2}{3}} &= \boldsymbol{x}_{k+\frac{1}{3}} + \boldsymbol{R}_2^T \boldsymbol{A}_2^{-1} \boldsymbol{R}_2 (\boldsymbol{b} - \boldsymbol{A} \boldsymbol{x}_{k+\frac{1}{3}}) \\
\boldsymbol{x}_{k+1} &= \boldsymbol{x}_{k+\frac{2}{3}} + \boldsymbol{R}_1^T \boldsymbol{A}_1^{-1} \boldsymbol{R}_1 (\boldsymbol{b} - \boldsymbol{A} \boldsymbol{x}_{k+\frac{2}{3}})
\end{aligned}
$$

which yields symmetric preconditioner that can be used in conjunction with conjugate gradient method to accelerate convergence

Many Overlapping Subdomains

- To achieve higher degree of parallelism with Schwarz method, we can apply two-domain algorithm recursively or use many subdomains

- If there are p overlapping subdomains, then define matrices $\boldsymbol{R}_i$ and $\boldsymbol{A}_i$ as before, $i = 1, \ldots, p$

- Additive Schwarz preconditioner then takes form

$$\sum_{i=1}^{p} \boldsymbol{R}_i^T \boldsymbol{A}_i^{-1} \boldsymbol{R}_i$$

Many Overlapping Subdomains

- Resulting generalization of block-Jacobi iteration is highly parallel, but not algorithmically scalable because convergence rate degrades as p grows

- Convergence rate can be restored by using coarse grid correction to provide global coupling

- If R_0 and R_0^T are restriction and interpolation matrices between coarse and fine grids, and $A_0 = R_0 A R_0^T$, then additive Schwarz preconditioner becomes

$$\sum_{i=0}^{p} R_i^T A_i^{-1} R_i$$

Many Overlapping Subdomains

- Multiplicative Schwarz iteration for p domains is defined analogously

- As with classical Gauss-Seidel vs. Jacobi, multiplicative Schwarz has faster convergence rate than corresponding additive Schwarz (though it still requires coarse grid correction to remain scalable)

- But unfortunately, multiplicative Schwarz appears to provide no parallelism, as p subproblems per iteration must be solved sequentially

- As with classical Gauss-Seidel, parallelism can be introduced by coloring subdomains to identify independent subproblems that can be solved simultaneously

Many Overlapping Subdomains

Non-Overlapping Subdomains

- We now consider adjacent subdomains whose only points in common are along their mutual boundary

- We partition indices of unknowns in corresponding discrete linear system into three sets, S_1 and S_2 corresponding to interior nodes in Ω_1 and Ω_2, respectively, and S_3 corresponding to interface nodes in Γ

Non-Overlapping Subdomains

- Partitioning matrix and right-hand-side vector accordingly, we obtain symmetric block linear system

$$\begin{bmatrix} A_{11} & 0 & A_{13} \\ 0 & A_{22} & A_{23} \\ A_{13}^T & A_{23}^T & A_{33} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$$

- Zero blocks result from assumption that nodes in Ω_1 are not directly connected to nodes in Ω_2, but only through interface nodes in Γ

Schur Complement

- Block LU factorization of matrix A yields

$$A = \begin{bmatrix} I & 0 & 0 \\ 0 & I & 0 \\ A_{13}^T A_{11}^{-1} & A_{23}^T A_{22}^{-1} & I \end{bmatrix} \begin{bmatrix} A_{11} & 0 & A_{13} \\ 0 & A_{22} & A_{23} \\ 0 & 0 & S \end{bmatrix}$$

where *Schur complement* matrix S is given by

$$S = A_{33} - A_{13}^T A_{11}^{-1} A_{13} - A_{23}^T A_{22}^{-1} A_{23}$$

Schur Complement

- We can now determine interface unknowns x_3 by solving system

$$Sx_3 = \hat{b}_3$$

where

$$\hat{b}_3 = b_3 - A_{13}^T A_{11}^{-1} b_1 - A_{23}^T A_{22}^{-1} b_2$$

- Remaining unknowns are then given by

$$
\begin{aligned}
x_1 &= A_{11}^{-1}(b_1 - A_{13}x_3) \\
x_2 &= A_{22}^{-1}(b_2 - A_{23}x_3)
\end{aligned}
$$

which can be computed simultaneously

Schur Complement

- Schur complement matrix S is expensive to compute and is generally dense even if A is sparse

- But if Schur complement system $S x_3 = \hat{b}_3$ is solved iteratively, then S need not be formed explicitly

- Matrix-vector multiplication by S requires solution in each subdomain, implicitly involving A_{11}^{-1} and A_{22}^{-1}, which can be done independently in parallel

- Conditioning of S is generally better than that of A, typically $\mathcal{O}(h^{-1})$ instead of $\mathcal{O}(h^{-2})$ for mesh size h, but *interface preconditioner* is still needed to accelerate convergence

Many Non-Overlapping Subdomains

- To achieve higher degree of parallelism with Schur method, we can apply two-domain algorithm recursively or use many subdomains

- If there are p non-overlapping subdomains, let I be set of indices of interior nodes of subdomains and B be set of indices of interface nodes separating subdomains

- Then discrete linear system has block form

$$\begin{bmatrix} \boldsymbol{A}_{II} & \boldsymbol{A}_{IB} \\ \boldsymbol{A}_{IB}^T & \boldsymbol{A}_{BB} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_I \\ \boldsymbol{x}_B \end{bmatrix} = \begin{bmatrix} \boldsymbol{b}_I \\ \boldsymbol{b}_B \end{bmatrix}$$

where $\boldsymbol{A}_{II} = \mathrm{diag}(\boldsymbol{A}_{11}, \ldots, \boldsymbol{A}_{pp})$ is block diagonal

Many Non-Overlapping Subdomains

- Block LU factorization of matrix A yields system

$$S x_B = \hat{b}_B$$

  where Schur complement matrix S is given by

$$S = A_{BB} - A_{IB}^T A_{II}^{-1} A_{IB}$$

  and $\hat{b}_B = b_B - A_{IB}^T A_{II}^{-1} b_I$

- As before, this system can be solved iteratively without forming S explicitly, and interface preconditioner is used to accelerate convergence

Many Non-Overlapping Subdomains

- Interior unknowns are then given by

$$\boldsymbol{x}_I = \boldsymbol{A}_{II}^{-1}(\boldsymbol{b}_I - \boldsymbol{A}_{IB}\boldsymbol{x}_B)$$

- All solves involving $\boldsymbol{A}_{II}^{-1}$, both in iterative phase for computing interface unknowns and subsequent computation of interior unknowns, can be performed on all subdomains in parallel because $\boldsymbol{A}_{II}$ is block diagonal

Parallel Computation with Grids

- Two basic approaches to parallel numerical solution of PDEs

  - Domain decomposition based on original PDE, which yields multiple problems to be solved in parallel, each on different subdomain

  - Parallel implementation of serial algorithm for solving discretized version of original problem

- Either approach ultimately leads to distribution of discrete mesh or grid across processors

- Communication between processors required to provide interface between subdomains or for parallel solution of discrete problem (e.g., matrix-vector multiplication for iterative methods)

Parallel Iterative Methods

- For some iterative solvers, such as Jacobi and red-black Gauss-Seidel, serial and parallel versions are equivalent and produce same results

- In parallel setting, with grid partitioned across processors, additional options arise that are not relevant in serial context, such as using standard Gauss-Seidel within each processor and Jacobi between processors

- Although iterative methods such as Jacobi and red-black Gauss-Seidel often yield good parallel efficiency, their relatively slow asymptotic convergence rate limits their usefulness in practice

Smoothers

- Stationary iterative methods, such as Jacobi and Gauss-Seidel, usually make fairly rapid initial progress in reducing error before settling into slow asymptotic phase

- In particular, they reduce high-frequency (i.e., oscillatory) components of error rapidly, but reduce low-frequency (i.e., smooth) components of error much more slowly

- For this reason, such methods are sometimes called *smoothers*

Smoothers

- Smooth or oscillatory components of error are relative to grid on which solution is defined

- Component that appears smooth on fine grid may appear oscillatory when sampled on coarser grid

- If we apply smoother on coarser grid, then we may make rapid progress in reducing this (now oscillatory) component of error

- After few iterations of smoother, results can then be interpolated back to fine grid to produce solution that has both higher-frequency and lower-frequency components of error reduced

Multigrid

- This idea can be extended to multiple levels of grids, so that error components of various frequencies can be reduced rapidly, each at appropriate level

- Transition from finer grid to coarser grid involves *restriction* or *injection*

- Transition from coarser grid to finer grid involves *interpolation* or *prolongation*

Residual Equation

- If $\hat{x}$ is approximate solution to $Ax = b$, with residual $r = b - A\hat{x}$, then error $e = x - \hat{x}$ satisfies *residual equation*

$$Ae = r$$

- Thus, in improving approximate solution we can work with just this *residual equation* involving error and residual, rather than solution and original right-hand side

- One advantage of residual equation is that zero is reasonable starting guess for its solution

Two-Grid Algorithm

1. On fine grid, use few iterations of *smoother* to compute approximate solution $\hat{x}$ for system $Ax = b$

2. Compute residual $r = b - A\hat{x}$

3. *Restrict* residual to coarse grid

4. On coarse grid, use few iterations of *smoother* on residual equation to obtain coarse-grid approximation to error

5. *Interpolate* coarse-grid correction to fine grid to obtain improved approximate solution on fine grid

6. Apply few iterations of *smoother* to corrected solution on fine grid

Multigrid Algorithm

- Multigrid method results from recursion in Step 4: coarse grid correction is itself improved by using still coarser grid, and so on down to some bottom level

- Computations become progressively cheaper on coarser and coarser grids because systems become successively smaller

- In particular, direct method may be feasible on coarsest grid if system is small enough

- There are many possible strategies for cycling through various grid levels

Multigrid Cycles

- *V-cycle* goes from finest grid goes down through successive levels to coarsest grid and then back up to finest grid

- *W-cycle* zig-zags among lower level (and cheaper) grids before going back up to finest grid

- *Full multigrid* bootstraps coarse solution up through grid levels, ultimately reaching finest grid

Multigrid

- By exploiting strengths of underlying iterative smoothers and avoiding their weaknesses, multigrid methods are capable of extraordinarily good performance

- At each level, smoother reduces oscillatory component of error rapidly, at rate independent of mesh size h, since few iterations of smoother, often only one, are performed at each level

- Since all components of error appear oscillatory at some level, convergence rate of entire multigrid scheme should be rapid and independent of mesh size, in contrast to most other iterative methods

Multigrid

- Moreover, cost of entire cycle of multigrid is only modest multiple of cost of single sweep on finest grid

- As result, multigrid methods are among most powerful methods available for solving sparse linear systems arising from PDEs

- In many cases, they are capable of converging to within truncation error of discretization at cost proportional to number of grid points, which is much faster than most other methods

Parallel Multigrid

- Many aspects of multigrid algorithms are readily implemented in parallel
    - Point smoothers, such as Jacobi and multi-color Gauss-Seidel
    - Residual computation
    - Restriction (fine-to-coarse)
    - Interpolation (coarse-to-fine)
- Other aspects are more problematic
    - Sequential cycling through grids
    - Parallel efficiency for coarse grids

Parallel Multigrid

- Compared with fine grids, coarse grids
  - Inherently allow less parallelism
  - Incur higher communication cost relative to computation
  - Risk poor load balance (some processors may even be idle for coarsest grids)
  - Do not necessarily grow as overall problem grows
- For these reasons, parallel implementations of multigrid try to minimize time spent on coarse grids

Parallel Multigrid

For model problem with n grid points on finest grid, depth of parallel multigrid is

- $\Theta(\log n)$ for V-cycle

- $\Theta(\log^2 n)$ for FMG

- $\Theta(\sqrt{n})$ for W-cycle

Parallel Multigrid

- For coarse grids, communication/computation ratio could be improved by using *fewer* processors, and load balance could be improved by redistributing work across processors

- However, such measures affect other aspects of algorithm negatively; for example, restriction and interpolation would no longer be local operations within individual processors

Parallel Multigrid

- Because of parallel inefficiencies associated with coarse grids, alternatives have been proposed to enhance parallelism in multigrid

- *Additive multigrid* performs smoothing on all grid levels simultaneously, but convergence is not guaranteed, so it is used as preconditioner

- *Parallel superconvergent multigrid* performs smoothing on multiple grids at each level simultaneously, thereby (hopefully) accelerating convergence

Parallel Multigrid

- Such variants of multigrid motivated by parallelism are not equivalent to serial multigrid and sacrifice some of its serial efficiency to gain greater parallelism

- Whether such strategies actually reduce overall time to solution depends on specific problem and parallel system

- Even with "classical" multigrid, serial superiority of FMG over V-cycle may outweigh parallel superiority of V-cycle over FMG

Multigrid and Domain Decomposition

- Domain decomposition with coarse problem can be viewed as two-level multigrid

- As with domain decomposition, restriction operator should often be transpose of interpolation operator, whose choice is critical for success of both methods

- Parallel solution of coarse problem is same

- Domain decomposition involves less communication per iteration than parallel multigrid but may require (a constant factor) more iterations

Scalability in Solving PDEs

- How scalable are numerical methods for solving PDEs?

- Consider 3-D Poisson equation, discretized with finite differences or low-order finite elements on $n \times n \times n$ grid

- Decompose in three dimensions, yielding cubes of size $n/p^{1/3} \times n/p^{1/3} \times n/p^{1/3}$

- Method involves one ghost cell exchange, one dot product (e.g., for CG), and local evaluation of matrix-vector product

Cost Model

- Per iteration cost is

$$T = \Theta\left(\alpha \log p + \beta \frac{n^2}{p^{2/3}} + \gamma \frac{n^3}{p} \right)$$

- If used as preconditioner or as part of parallel multigrid or domain decomposition, cost is similar and number of iterations can be independent of p

Summary and Suggestions

- Use redundant computation to reduce communication, increasing parallel efficiency
    - Limits *maximum* efficiency

- Consider alternate approximation or solution methods to improve concurency
    - E.g., domain decomposition; some block decompositions

- Parallelize best methods
    - E.g., multigrid, Krylov methods with high-quality preconditioners, not block Jacobi

References – Domain Decomposition

- T. F. Chan and T. P. Mathew, Domain decomposition algorithms, *Acta Numerica* 3:61-143, 1994

- A. Quarteroni and A. Valli, *Domain Decomposition Methods for Partial Differential Equations*, Oxford Univ. Press, 1999

- B. F. Smith, Domain decomposition methods for partial differential equations, D. E. Keyes, A. Sameh, and V. Venkatakrishnan, eds., *Parallel Numerical Algorithms*, pp. 225-243, Kluwer, 1997

- B. F. Smith, P. E. Bjørstad, and W. D. Gropp, *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*, Cambridge Univ. Press, 1996

- A. Toselli and O. Widlund, *Domain Decomposition Methods: Algorithms and Theory*, Springer, 2005

References – Multigrid

- P. Bastian, W. Hackbusch, and G. Wittum, Additive and multiplicative multigrid — a comparison, *Computing* 60:345-364, 1998

- J. H. Bramble, *Multigrid Methods*, Pitman, 1993

- W. L. Briggs, V. E. Henson and S. F. McCormick, *A Multigrid Tutorial*, 2nd ed., SIAM, 2000

- W. Hackbusch, *Multigrid Methods and Applications*, Springer-Verlag, 1985

- U. Trottenberg, C. Oosterlee, and A. Schüller, *Multigrid*, Academic Press, 2001

- P. Wesseling, *An Introduction to Multigrid Methods*, John Wiley & Sons, 1992

References – Parallel Multigrid

- A. Brandt, Multigrid solvers on parallel computers, M. H. Schultz, ed., *Elliptic Problem Solvers*, pp. 39-83, Academic Press, 1981

- T. F. Chan and Y. Saad, Multigrid algorithms on the hypercube multiprocessor, *IEEE Trans. Comput.*, 35:969-977, 1986

- T. F. Chan and R. Schreiber, Parallel networks for multigrid algorithms: architecture and complexity, *SIAM J. Sci. Stat. Comput.* 6:698-711, 1985

- A. Greenbaum, A multigrid method for multiprocessors, *Appl. Math. Comput.* 19:75-88, 1986

- H. Hoppe and H. Muhlenbein, Parallel adaptive full-multigrid methods on message-based multiprocessors, *Parallel Computing* 3:269-287, 1986

- J. E. Jones and S. F. McCormick, Parallel multigrid methods, D. E. Keyes, A. Sameh, and V. Venkatakrishnan, eds., *Parallel Numerical Algorithms*, pp. 203-224, Kluwer, 1997