Evaluation of Overlapping Restricted Additive Schwarz Preconditioning for Parallel Solution of Very Large Power Flow Problems

Shrirang Abhyankar, Barry Smith, Emil Constantinescu

Mathematics and Compute Science Division
Argonne National Laboratory

November 22, 2013

Overview

Overview

1. Power flow background
2. Parallel solution of power flow equations using an overlapping Schwarz preconditioned GMRES
3. Simulation results
4. Summary

Overview

1. Power flow background
2. Parallel solution of power flow equations using an overlapping Schwarz preconditioned GMRES
3. Simulation results
4. Summary

Overview

1. Power flow background
2. Parallel solution of power flow equations using an overlapping Schwarz preconditioned GMRES
3. Simulation results
4. Summary

Power flow equations are the backbone of power system analysis

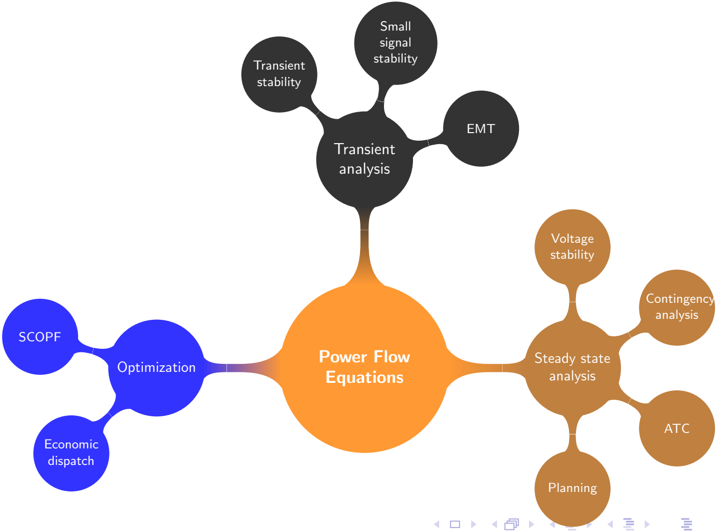

Power flow equations are the backbone of power system analysis

All applications need fast PF solution!!

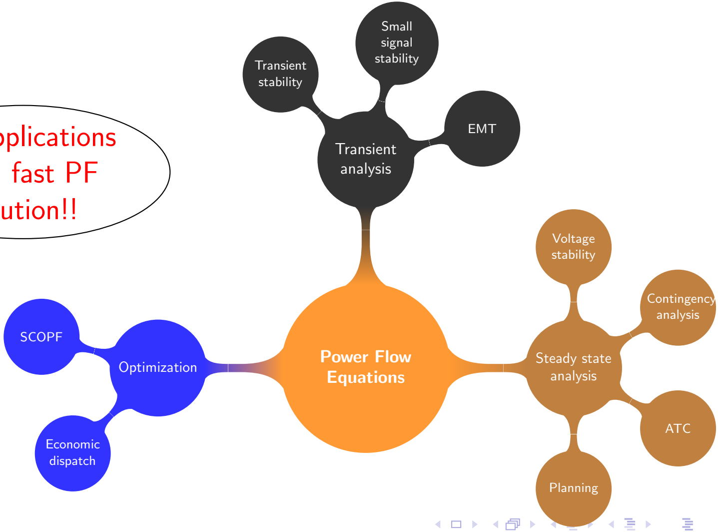

Power flow equations and solution

Equations in power balance form

$$P_i^{inj} - \sum_{k=1}^{n} V_i V_k (G_{ik} cos(\theta_{ik}) + B_{ik} sin(\theta_{ik})) = \Delta P_i,$$

$$Q_i^{inj} - \sum_{k=1}^{n} V_i V_k (G_{ik} sin(\theta_{ik}) - B_{ik} cos(\theta_{ik})) = \Delta Q_i,$$

$$i = 1, \ldots, \text{nbus}$$

Solution via Newton's method

$$\Delta P(\theta, V) = 0$$

$$\Delta Q(\theta, V) = 0$$

Linear system solution: Main computational bottleneck

$$\begin{bmatrix} \frac{\partial \Delta P}{\partial \theta} & \frac{\partial \Delta P}{\partial V} \\ \frac{\partial \Delta Q}{\partial \theta} & \frac{\partial \Delta Q}{\partial V} \end{bmatrix} \begin{bmatrix} \Delta \theta \\ \Delta V \end{bmatrix} = - \begin{bmatrix} \Delta P \\ \Delta Q \end{bmatrix}$$

Parallel solution of power flow equations

- Several previously proposed solution schemes
  - Parallel direct solution via rearrangement of forward-backward substitution [Wu and Bose:1995]
  - Partition to BBDF form [Chen and Chen:2000]
  - Parallel direct solution [Tu and Flueck:2002]
  - GMRES with fast-decoupled matrix as preconditioner [Tu and Flueck:2002]
  - GPU [Gopal et. al.:2007], [Kamiabad:2011]
  - Using OpenMP [Dag and Soykan:2011]
- Proposed approach
  - Domain decomposition of power flow equations
  - Iterative linear solution via GMRES
  - Overlapping additive Schwarz preconditioning

Parallel solution of power flow equations

- Several previously proposed solution schemes
  - Parallel direct solution via rearrangement of forward-backward substitution [Wu and Bose:1995]
  - Partition to BBDF form [Chen and Chen:2000]
  - Parallel direct solution [Tu and Flueck:2002]
  - GMRES with fast-decoupled matrix as preconditioner [Tu and Flueck:2002]
  - GPU [Gopal et. al.:2007], [Kamiabad:2011]
  - Using OpenMP [Dag and Soykan:2011]
- Proposed approach
  - Domain decomposition of power flow equations
  - Iterative linear solution via GMRES
  - Overlapping additive Schwarz preconditioning

Domain decomposition methods

- Divide the domain W

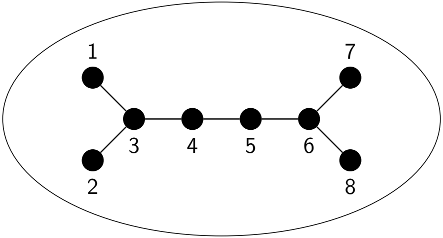

- Typically each subdomain W_i^0 is the computational assignment for a processor (power flow equations for a subnetwork)
- Posed as a graph partitioning problem
  - Several algorithms: Combinatorial, Spectral, Geometric, Multilevel
  - Several libraries: Metis, ParMetis, Chaco
- We use Chaco (Multilevel algorithm) in this work

Domain decomposition methods

- Divide the domain W in N non-overlapping subdomains W_i^0.

$$W = \bigcup_{i=1}^{N} W_i^0$$

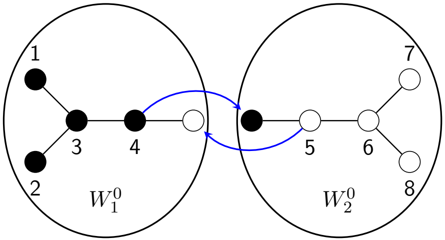

- Typically each subdomain W_i^0 is the computational assignment for a processor (power flow equations for a subnetwork)
- Posed as a graph partitioning problem
  - Several algorithms: Combinatorial, Spectral, Geometric, Multilevel
  - Several libraries: Metis, ParMetis, Chaco
- We use Chaco (Multilevel algorithm) in this work

Domain decomposition methods

- Divide the domain W in N non-overlapping subdomains W_i^0.

$$W = \bigcup_{i=1}^{N} W_i^0$$

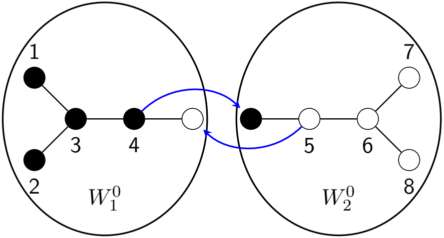

- Typically each subdomain W_i^0 is the computational assignment for a processor (power flow equations for a subnetwork)
- Posed as a graph partitioning problem
  - Several algorithms: Combinatorial, Spectral, Geometric, Multilevel
  - Several libraries: Metis, ParMetis, Chaco
- We use Chaco (Multilevel algorithm) in this work

Generalized Minimal Residual Algorithm (GMRES)

- Krylov subspace based linear solver for square, non symmetric systems
- Approximates solution by minimizing the residual over the Krylov subspace

$$K_n = K(A, b) = [b, Ab, A^2b, \dots, A^{m-1}b]$$

- Efficient in parallel as compared to a direct solver
  - Consist of matrix-vector products and reductions
  - No factorization involved

Preconditioning

- Krylov methods need a "preconditioner" to accelerate convergence since it depends on spectral properties of the matrix.
- A *preconditioner* M^{-1} is matrix such that $M^{-1}A$ or AM^{-1} has a better condition number

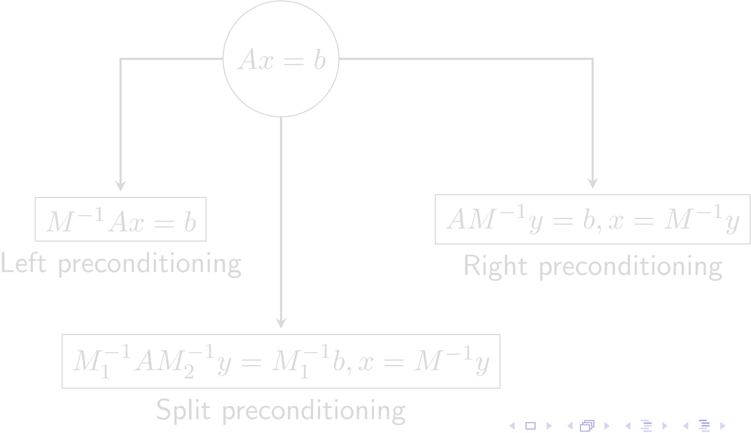

Preconditioning

- Krylov methods need a "preconditioner" to accelerate convergence since it depends on spectral properties of the matrix.
- A *preconditioner* M^{-1} is matrix such that $M^{-1}A$ or AM^{-1} has a better condition number

$$Ax = b$$

$$\boxed{M^{-1}Ax = b}$$

Left preconditioning

$$\boxed{AM^{-1}y = b, x = M^{-1}y}$$

Right preconditioning

$$\boxed{M_1^{-1}AM_2^{-1}y = M_1^{-1}b, x = M^{-1}y}$$

Split preconditioning

Overlapping additive Schwarz preconditioning (ASM)

- Extend the overlap δ to include neighboring nodes from other processors.

$$W = \bigcup_{i=1}^{N} W_i^\delta$$

- Subdomain operator

$$A_i = R_i^\delta A R_i^\delta$$

- R_i^δ is the "restriction" operator (mapping to subdomain)
- Overlapping additive Schwarz preconditioner

$$\boxed{M_{AS}^{-1} = \sum R_i^\delta A_i^{-1} R_i^\delta}$$

- A_i is not invertible but its restriction to the subspace is invertible

Overlapping Additive Schwarz preconditioning (ASM)

Partitioned domain on 2 procesors

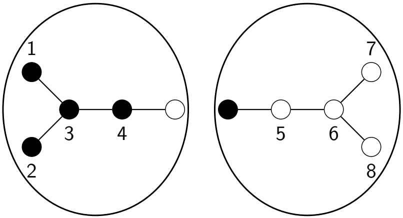

$$A \equiv \begin{bmatrix} \bullet & & \bullet & & & & & \\ & \bullet & \bullet & & & & & \\ \bullet & \bullet & \bullet & \bullet & & & & \\ & & \bullet & \bullet & \bullet & & & \\ & & & \circ & \circ & \circ & & \\ & & & & \circ & \circ & \circ & \circ \\ & & & & & \circ & \circ & \\ & & & & & \circ & & \circ \end{bmatrix}$$

Overlapping Additive Schwarz preconditioning (ASM)

Overlap δ=0 (Block-Jacobi preconditioner)

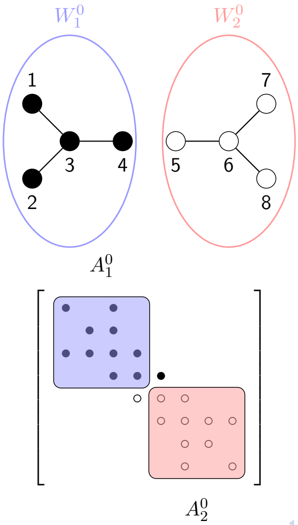

Overlapping Additive Schwarz preconditioning (ASM)

Overlap δ=1

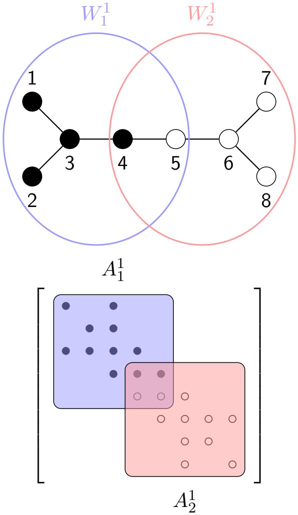

Overlapping Additive Schwarz preconditioning (ASM)

Overlap δ=2

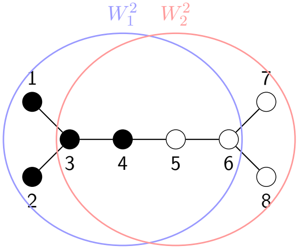

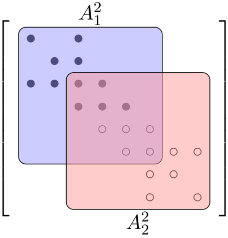

Restricted additive Schwarz preconditioning (RASM)

- Proposed by Cai and Sarkis
- Simple modification to ASM

$$\boxed{M_{RAS}^{-1} = \sum R_i^0 A_i^{-1} R_i^\delta}$$

- Saves half the communication
- Also has a reduction in iteration count compared to ASM

Machine and software details

- 64 core shared memory machine
- Four 2.2 GHz AMD Opteron 6274 processors with 16 cores each
- Built using PETSc (Portable Extensible Toolkit for Scientific Computing) library

Table: Inventory of test cases

casename	Buses	Gens	Lines
case22996	22,996	2,416	27,408
case51741	51,741	5,436	61,686
case91984	91,984	9,664	109,680

- Cases created using MATPOWER's 2000+ test cases in a structured grid formation

case22996: Overlapping RASM schemes more scalable than block-Jacobi and parallel direct solution

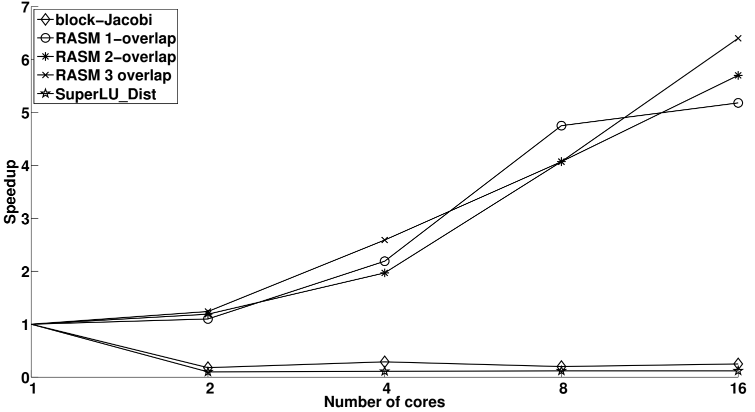

Figure: Speedup with different linear solution schemes

case22996: Overlapping RASM schemes more scalable than block-Jacobi and parallel direct solution

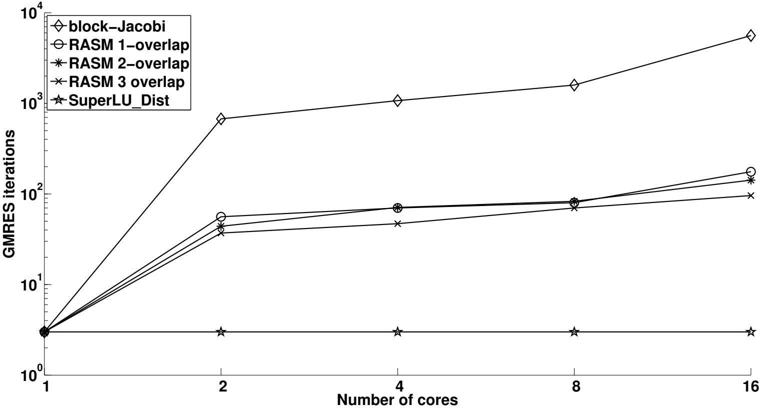

Figure: Linear solver iterations

case22996: Overlapping RASM schemes more scalable than block-Jacobi and parallel direct solution

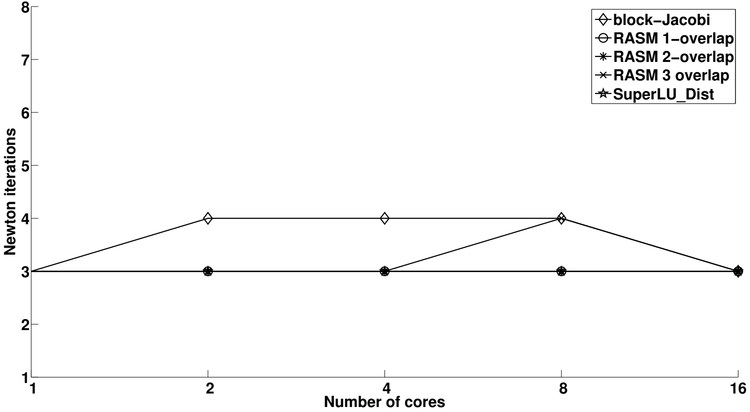

Figure: Newton iterations

case51741 scalability results

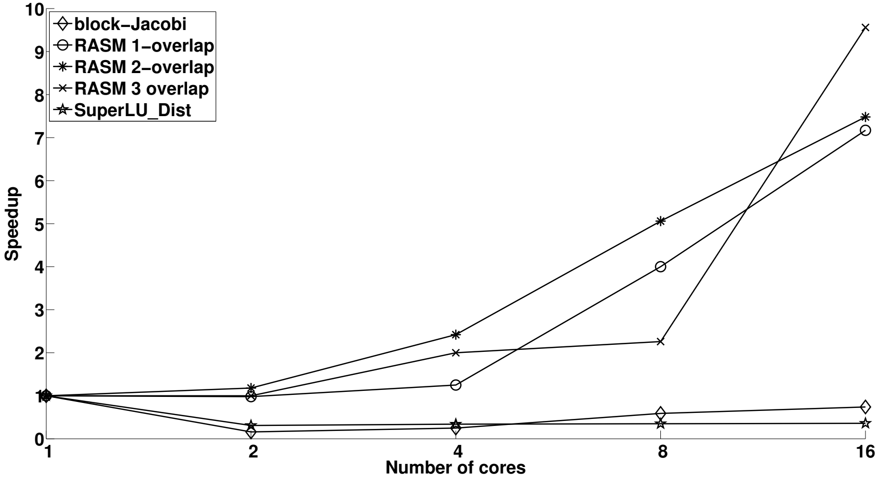

Figure: case51741:Speedup

case91984 scalability results

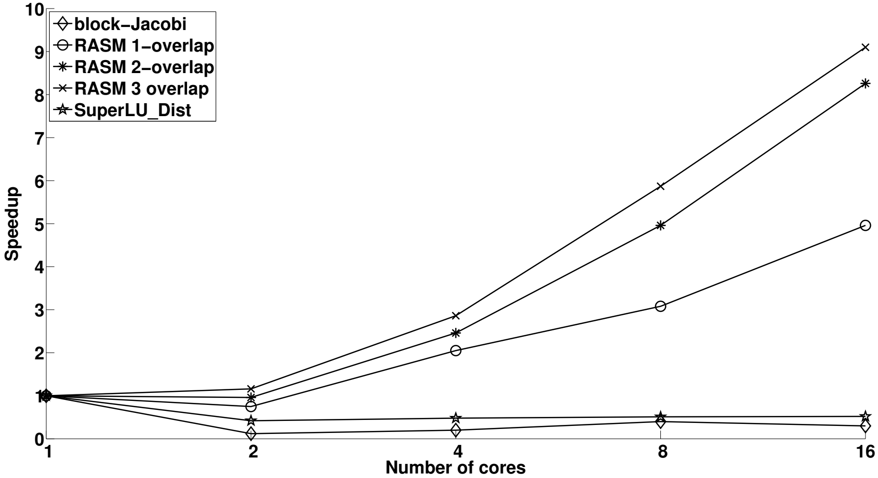

Figure: case91984:Speedup

Summary

- Presented a parallel solution of nonlinear AC power flow equations
- Accelerated the linear solve involved in the power flow solution using a restricted overlapping additive Schwarz (RASM) preconditioned GMRES scheme
- Presented scalability results for three large test cases demonstrating the scalability of RASM preconditioned GMRES