

Solveurs linéaires parallèles par décomposition de domaine algébrique



Luc Giraud


HiePACS project - INRIA Bordeaux Sud-Ouest
joint INRIA-CERFACS lab. on High Performance Computing

Journée GNR MOMAS / GDR Calcul

Paris, 5 Mai 2010

- 1 Background
- 2 A parallel algebraic domain decomposition solver
- 3 Parallel and numerical scalability on 3D academic problems
- 4 Two-level parallel implementation features and performances
- 5 Prospectives

Motivations



The “spectrum” of linear algebra solvers

Direct

- Robust/accurate for general problems
- BLAS-3 based implementations
- Memory/CPU prohibitive for large 3D problems
- Limited parallel scalability

Iterative

- Problem dependent efficiency/controlled accuracy
- Only mat-vect required, fine grain computation
- Less memory computation, possible trade-off with CPU
- Attractive “build-in” parallel features

An effort for combining advantages of those solvers is needed


Goal: design Hybrid Linear Solvers

Develop robust scalable parallel hybrid direct/iterative linear solvers

- Exploit the efficiency and robustness of the sparse direct solvers
- Develop robust parallel preconditioners for iterative solvers
- Take advantage of the natural scalable parallel implementation of iterative solvers


Domain Decomposition (DD)

- Natural approach for PDE's
- Extend to general sparse matrices
- Partition the problem into subdomains, subgraphs
- Use a direct solver on the subdomains
- Robust preconditioned iterative solver



Overlapping Domain Decomposition

Classical Additive Schwarz preconditioners



- Goal: solve linear system $\mathcal{A}x = b$
- Use iterative method
- Apply the preconditioner at each step
- The convergence rate deteriorates as the number of subdomains increases


$$\mathcal{A} = \begin{pmatrix} \mathcal{A}_{1,1} & \mathcal{A}_{1,\delta} & \mathcal{A}_{\delta,2} \\ \mathcal{A}_{\delta,1} & \mathcal{A}_{\delta,\delta} & \mathcal{A}_{\delta,2} \\ \mathcal{A}_{\delta,2} & \mathcal{A}_{\delta,2} & \mathcal{A}_{2,2} \end{pmatrix} \Rightarrow \mathcal{M}_{AS}^{\delta} = \begin{pmatrix} \boxed{\mathcal{A}_{1,1}} & \boxed{\mathcal{A}_{1,\delta}} & -1 & \\ \mathcal{A}_{\delta,1} & \boxed{\mathcal{A}_{\delta,\delta}} & \mathcal{A}_{\delta,2} & -1 \\ & \boxed{\mathcal{A}_{\delta,2}} & \boxed{\mathcal{A}_{2,2}} & \end{pmatrix}$$

Classical Additive Schwarz preconditioners N subdomains case

$$\mathcal{M}_{AS}^{\delta} = \sum_{i=1}^N (\mathcal{R}_i^{\delta})^T (\mathcal{A}_i^{\delta})^{-1} \mathcal{R}_i^{\delta}$$

Non-overlapping Domain Decomposition

Schur complement reduced system



- Goal: solve linear system $\mathcal{A}x = b$
- Apply partially Gaussian elimination
- Solve the reduced system $Sx_\Gamma = f$
- Then solve $\mathcal{A}_i x_i = b_i - \mathcal{A}_{i,\Gamma} x_\Gamma$

$$\begin{pmatrix} \mathcal{A}_{1,1} & 0 & \mathcal{A}_{1,\Gamma} \\ 0 & \mathcal{A}_{2,2} & \mathcal{A}_{2,\Gamma} \\ 0 & 0 & S \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_\Gamma \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_\Gamma - \sum_{i=1}^2 \mathcal{A}_{\Gamma,i} \mathcal{A}_{i,i}^{-1} b_i \end{pmatrix}$$


Solve $\mathcal{A}x = b \implies$ solve the reduced system $Sx_\Gamma = f \implies$ then solve $\mathcal{A}_i x_i = b_i - \mathcal{A}_{i,\Gamma} x_\Gamma$

where $S = \mathcal{A}_{\Gamma,\Gamma} - \sum_{i=1}^2 \mathcal{A}_{\Gamma,i} \mathcal{A}_{i,i}^{-1} \mathcal{A}_{i,\Gamma}$,

and $f = b_\Gamma - \sum_{i=1}^2 \mathcal{A}_{\Gamma,i} \mathcal{A}_{i,i}^{-1} b_i$.

Nonoverlapping Domain Decomposition

Schur complement reduced system



$$\Gamma = k U l \cup m U n$$

Distributed Schur complement

$$\overbrace{\begin{pmatrix} S_{kk}^{(\iota)} & S_{k\ell} \\ S_{\ell k} & S_{\ell\ell}^{(\iota)} \end{pmatrix}}^{\Omega_l} \quad \overbrace{\begin{pmatrix} S_{\ell\ell}^{(\iota+1)} & S_{\ell m} \\ S_{m\ell} & S_{mm}^{(\iota+1)} \end{pmatrix}}^{\Omega_{l+1}} \quad \overbrace{\begin{pmatrix} S_{mm}^{(\iota+2)} & S_{mn} \\ S_{nm} & S_{nn}^{(\iota+2)} \end{pmatrix}}^{\Omega_{l+2}}$$

In an assembled form: $S_{\ell\ell} = S_{\ell\ell}^{(\iota)} + S_{\ell\ell}^{(\iota+1)} \Rightarrow S_{\ell\ell} = \sum_{\iota \in \text{adj}} S_{\ell\ell}^{(\iota)}$

Non-overlapping Domain Decomposition

Algebraic Additive Schwarz preconditioner [L.Carvalho, L.G., G.Meurant - 01]

$$S = \sum_{i=1}^N \mathcal{R}_{\Gamma_i}^T S^{(i)} \mathcal{R}_{\Gamma_i}$$

$$S = \begin{pmatrix} \ddots & & & & & & \\ & S_{kk} & S_{kl} & & & & \\ & S_{lk} & S_{ll} & S_{lm} & & & \\ & & S_{ml} & S_{mm} & S_{mn} & & \\ & & & S_{nm} & S_{nn} & & \end{pmatrix} \Rightarrow \mathcal{M} = \begin{pmatrix} \ddots & & & & & & \\ & \boxed{S_{kk} \quad S_{kl}} & & -1 & & & \\ & \boxed{S_{lk}} & \boxed{S_{ll}} & \boxed{S_{lm}} & & & -1 \\ & & \boxed{S_{ml}} & \boxed{S_{mm}} & \boxed{S_{mn}} & & \\ & & & \boxed{S_{nm}} & \boxed{S_{nn}} & & \end{pmatrix}$$

$$\mathcal{M} = \sum_{i=1}^N \mathcal{R}_{\Gamma_i}^T (\bar{S}^{(i)})^{-1} \mathcal{R}_{\Gamma_i}$$

where $\bar{S}^{(i)}$ is obtained from $S^{(i)}$

$$\underbrace{S^{(i)} = \begin{pmatrix} S_{kk}^{(\iota)} & S_{kl} \\ S_{lk} & S_{ll}^{(\iota)} \end{pmatrix}}_{\text{local Schur}} \Rightarrow \bar{S}^{(i)} = \underbrace{\begin{pmatrix} S_{kk} & S_{kl} \\ S_{lk} & S_{ll} \end{pmatrix}}_{\text{local assembled Schur}}$$


$$\sum_{\iota \in \text{adj}} S_{ll}^{(\iota)}$$

Similarity with Neumann-Neumann preconditioner [J.F. Bourgat, R. Glowinski, P. Le Tallec and M. Vidrascu - 89] [Y.H. de Roek, P. Le Tallec and M. Vidrascu - 91]

Parallel preconditioning features

$$S^{(i)} = A_{\Gamma_i \Gamma_i}^{(i)} - A_{\Gamma_i \Gamma_j} A_{\Gamma_j \Gamma_j}^{-1} A_{\Gamma_j \Gamma_i}$$

$$M_{AS} = \sum_{i=1}^{\# \text{domains}} R_i^T (\bar{S}^{(i)})^{-1} R_i$$



$$\bar{S}^{(i)} = \begin{pmatrix} S_{mm} & S_{mg} & S_{mk} & S_{ml} \\ S_{gm} & S_{gg} & S_{gk} & S_{gl} \\ S_{km} & S_{kg} & S_{kk} & S_{kl} \\ S_{lm} & S_{lg} & S_{lk} & S_{ll} \end{pmatrix}$$

Assembled local Schur complement

$$S^{(i)} = \begin{pmatrix} S_{mm}^{(i)} & S_{mg} & S_{mk} & S_{ml} \\ S_{gm} & S_{gg}^{(i)} & S_{gk} & S_{gl} \\ S_{km} & S_{kg} & S_{kk}^{(i)} & S_{kl} \\ S_{lm} & S_{lg} & S_{lk} & S_{ll}^{(i)} \end{pmatrix}$$

local Schur complement

$$S_{mm} = \sum_{j \in \text{adj}(m)} S_{mm}^{(j)}$$

Parallel implementation

- Each *subdomain* $\mathcal{A}^{(i)}$ is handled by one *processor*

$$\mathcal{A}^{(i)} \equiv \begin{pmatrix} \mathcal{A}_{\mathcal{I}_i \mathcal{I}_i} & \mathcal{A}_{\mathcal{I}_i \Gamma_i} \\ \mathcal{A}_{\mathcal{I}_i \Gamma_i} & \mathcal{A}_{\Gamma_i \Gamma_i}^{(i)} \end{pmatrix}$$

- Concurrent partial factorizations are performed on each processor to form the so called “local Schur complement”

$$\mathcal{S}^{(i)} = \mathcal{A}_{\Gamma_i \Gamma_i}^{(i)} - \mathcal{A}_{\Gamma_i \mathcal{I}_i} \mathcal{A}_{\mathcal{I}_i \mathcal{I}_i}^{-1} \mathcal{A}_{\mathcal{I}_i \Gamma_i}$$

- The reduced system $\mathcal{S}x_{\Gamma} = f$ is solved using a distributed Krylov solver
 - One matrix vector product per iteration each processor computes $\mathcal{S}^{(i)}(x_{\Gamma}^{(i)})^k = (y^{(i)})^k$
 - One local preconditioner apply $(\mathcal{M}^{(i)})(z^{(i)})^k = (r^{(i)})^k$
 - Local neighbor-neighbor communication per iteration
 - Global reduction (dot products)
- Compute simultaneously the solution for the interior unknowns

$$\mathcal{A}_{\mathcal{I}_i \mathcal{I}_i} x_{\mathcal{I}_i} = b_{\mathcal{I}_i} - \mathcal{A}_{\mathcal{I}_i \Gamma_i} x_{\Gamma_i}$$

Algebraic Additive Schwarz preconditioner

Main characteristics in $2D$

- The ratio interface/interior is small
- Does not require large amount of memory to store the preconditioner
- Computation/application of the preconditioner are fast
- They consist in a call to LAPACK/BLAS-2 kernels

Main characteristics in $3D$

- The ratio interface/interior is large
- The storage of the preconditioner might not be affordable
- The construction of the preconditioner can be computationally expensive
- Need **cheaper** Algebraic Additive Schwarz form of the preconditioner

What tricks exist to construct cheaper preconditioners

Sparsification strategy through dropping

$$\widehat{\mathbf{S}}_{k\ell} = \begin{cases} \bar{\mathbf{s}}_{k\ell} & \text{if } \bar{\mathbf{s}}_{k\ell} \geq \xi(|\bar{\mathbf{s}}_{kk}| + |\bar{\mathbf{s}}_{\ell\ell}|) \\ 0 & \text{else} \end{cases}$$

Approximation through ILU - [INRIA PhyLeas - A. Haidar, L.G., Y.Saad - 10]



$$pILU(A^{(i)}) \equiv pILU \begin{pmatrix} A_{ii} & A_{i\Gamma_i} \\ A_{\Gamma_i i} & A_{\Gamma_i \Gamma_i}^{(i)} \end{pmatrix} \equiv \begin{pmatrix} \tilde{L}_i & 0 \\ A_{\Gamma_i} \tilde{U}_i^{-1} & I \end{pmatrix} \begin{pmatrix} \tilde{U}_i & \tilde{L}_i^{-1} A_{i\Gamma} \\ 0 & \tilde{\mathbf{S}}^{(i)} \end{pmatrix}$$

Mixed arithmetic strategy

- Compute and store the preconditioner in 32-bit precision arithmetic **Is accurate enough?**
- Limitation when the conditioning exceeds the accuracy of the 32-bit computations **Fix it!**
- **Idea:** Exploit 32-bit operation whenever possible and resort to 64-bit at critical stages
- **Remarks:** the backward stability result of GMRES indicates that it is hopeless to expect convergence at a backward error level smaller than the 32-bit accuracy [C.Paige, M.Rozložník, Z.Strakoš - 06]
- **Idea:** To overcome this limitation we use FGMRES [Y.Saad - 93]

Academic model problems

Problem patterns




Diffusion equation ($\epsilon = 1$ and $v = 0$) and convection-diffusion equation

$$\begin{cases} -\epsilon \operatorname{div}(K \cdot \nabla u) + v \cdot \nabla u & = f & \text{in } \Omega, \\ u & = 0 & \text{on } \partial\Omega. \end{cases}$$


- Heterogeneous problems
- Anisotropic-heterogeneous problems
- Convection dominated term

Numerical behaviour of sparse preconditioners

Convergence history of PCG




Time history of PCG




- 3D heterogeneous diffusion problem with 43 M dof mapped on 1000 processors
- For ($\xi \ll \ll$) the convergence is marginally affected while the memory saving is significant 15%
- For ($\xi \gg \gg$) a lot of resources are saved but the convergence becomes very poor 1%
- Even though they require more iterations, the sparsified variants converge faster as the time per iteration is smaller and the setup of the preconditioner is cheaper.

Numerical behaviour of mixed preconditioners

Convergence history of PCG




Time history of PCG




- 3D heterogeneous diffusion problem with 43 M dof mapped on 1000 processors
- 64-bit and mixed computation both attained an accuracy at the level of 64-bit machine precision
- The number of iterations slightly increases
- The mixed approach is the fastest, down to an accuracy that is problem dependent

Scaled scalability on massively parallel platforms

Numerical scalability



Parallel performance



- The solved problem size varies from 2.7 up to 74 M dof
- Control the grow in the # of iterations by introducing a coarse space correction
- The computing time increases slightly when increasing # sub-domains
- Although the preconditioners do not scale perfectly, the parallel time scalability is acceptable
- The trend is similar for all variants of the preconditioners using CG Krylov solver

Summary on the model problems

[L.Giraud, A.Haidar, L.T.Watson - 08] [L.Giraud, A.Haidar, Y.Saad - 10]

Sparse preconditioner

- For reasonable choice of the dropping parameter ξ the convergence is marginally affected
- The sparse preconditioner outperforms the dense one in time and memory

Mixed preconditioner

- Mixed arithmetic and 64-bit both attained an accuracy at the level of 64-bit machine precision
- Mixed preconditioner does not delay too much the convergence

Approximate preconditioner

- The convergence is marginally affected while the memory saving is significant
- The approximate variant converge faster as the time per iteration is smaller and the setup of the preconditioner is cheaper.
- This preconditioner require some tuning for very hard problem (structural mechanics...)


On the weak scalability

- Although these preconditioners are local, possibly not numerically scalable, they exhibit a fairly good parallel time scalability (possible fix for elliptic problems)
- The trends that have been observed on this choice of model problem have been observed on many other problems

Application areas

- Structural mechanics : real SPD and symmetric indefinite linear systems.
- Electromagnetism : complex symmetric non-Hermitian.
- Seismic : complex symmetric non-Hermitian.


Fuselage of 6.5 M dof




- Linear elasticity
- Composed of its skin, stringers and frames
- Midlinn shell elements are used
- Each node has 6 unknowns
- One extremity is fixed
- On the other extremity a rigid body element is added
- A force perpendicular to the axis is applied

Numerical behaviour of sparse preconditioners

Convergence history



Time history



- Fuselage problem of 6.5 M dof mapped on 16 processors
- The sparse preconditioner setup is 4 times faster than the dense one (19.5 v.s. 89 seconds)
- In term of global computing time, the sparse algorithm is about twice faster
- The attainable accuracy of the hybrid solver is comparable to the one computed with the direct solver

Exploiting 2-levels of parallelism - motivations

“The numerical improvement”

- Classical parallel implementations (*1-level*) of DD assign one subdomain per processor
- Parallelizing means increasing the number of subdomains
- Increasing the number of subdomains often leads to increasing the number of iterations
- To avoid this, one can instead of increasing the number of subdomains, keeping it small while handling each subdomain by more than one processor introducing *2-levels* of parallelism

“The parallel performance improvement”

- Large 3D systems often require a huge amount of data storage
- On SMP node: classical *1-level parallel* can only use a subset of the available processors
- Thus some processors are “wasted”, as they are “idle” during the computation
- The “idle” processors might contribute to the computation and the simulation runs closer to the peak of per-node performance by using *2-levels* of parallelism

Numerical improvement benefits

Fuselage of 6.5Mdof

# total processors	Algo	# subdomains	# processors/ subdomain	# iter	iterative loop time
16 processors	<i>1-level parallel</i>	16	1	147	77.9
	<i>2-level parallel</i>	8	2	98	51.4
32 processors	<i>1-level parallel</i>	32	1	176	58.1
	<i>2-level parallel</i>	16	2	147	44.8
	<i>2-level parallel</i>	8	4	98	32.5
64 processors	<i>1-level parallel</i>	64	1	226	54.2
	<i>2-level parallel</i>	32	2	176	40.1
	<i>2-level parallel</i>	16	4	147	31.3
	<i>2-level parallel</i>	8	8	98	27.4

- Reduce the number of subdomains \implies reduce the number of iterations
- Though the subdomain size increases, the time of the iterative loop decreases as:
 - The number of iterations decreases
 - Each subdomain is handled in parallel
 - All the iterative kernels are efficiently computed in parallel
- The speedup factors of the iterative loop vary from 1.3 to 1.8
- Very attractive especially when the convergence rate depends on the # of subdomains
- Might be of great interest when embeded into nonlinear solver

Parallel performance benefits

Fuselage of 6.5Mdof

# subdomains or SMP node	Algo	proc/subdom or "working"	Precond setup time	# iter	iterative loop time	Total time
8	<i>1-level</i>	1	208.0	98	94.1	525.1
	<i>2-level</i>	2	124.6		51.5	399.1
		4	70.8		32.5	326.4
16	<i>1-level</i>	1	89.0	147	77.9	217.2
	<i>2-level</i>	2	52.7		44.8	147.8
		4	30.4		31.3	112.0
32	<i>1-level</i>	1	30.0	176	58.1	124.1
	<i>2-level</i>	2	20.4		40.8	97.2
		4	13.0		22.7	71.7

- When running large simulations that need all the memory available on the nodes
- The *1-level parallel* algo "wastes" resource performance (it lose 48 "idle" processors on 16 SMP)
- The *2-level parallel* algo exploits the computing facilities of the remaining "idle" processors
- The *2-level parallel* algo runs closer to the peak of per-node performance
- The preconditioner setup time benefits vary from 1.5 to 3
- The speedup factors of the iterative loop vary from 1.8 to 2.7

Toward a “black-box” parallel solver

MAPHYS package- ongoing work

- Apply ideas to adjacency graph of sparse matrices no longer to meshes (ANR-CIS Solstice).
- Replace full-MPI two-level parallelism by mixed multi-threaded MPI parallel implementation to better comply with NUMA cluster features (PasTiX, Super_LU).
- Improve the solver capability for symmetric indefinite et fully unsymmetric (France-Berkeley Fund pending proposal).
- Perform a complexity analysis to study the computational scalability

<http://www.inria.fr/recherche/equipes/hiepac.fr.html>

Credit to co-workers

- Numerical methodologies:
E. Agullo (INRIA), A. Guermouche (INRIA), A. Haidar (ICL, Univ. Tennessee), Y. Lee (INRIA), J. Roman (INRIA), Y. Saad (Univ. Minnesota). MUMPS & PaStiX developers.
- Applications:
H. Benhadjali (SEISCOPE), D. Goudin (CEA-CESTA), S. Operto (Géosciences Azur), S. Pralet (SAMTECH), J. Virieux (LGIT).

Merci pour votre attention

Questions ?