CFD Parallel Solvers

and Fluid–Structure Interaction

May 8, 2013 | Paolo Crosetto

Table of contents

Motivations

Open source CFD tools

- Widely used and developed in industry/research
- Broad variety of different methods and codes
- Active community distributing and maintaining the source codes
- Often the freely available tools have the same features/performance of the commercial ones!

Motivations

Open source CFD tools

- Widely used and developed in industry/research
- Broad variety of different methods and codes
- Active community distributing and maintaining the source codes
- Often the freely available tools have the same features/performance of the commercial ones!

Drawbacks

- Which tool should I use to solve my problem?
- Often scientists restart the development from scratch, too many codes implement the same thing
- High mortality of CFD codes
- Difficulty to use, lack of documentation.

Need of CFD benchmarks to test the sw and algorithms.

An example:

The Incompressible Navier–Stokes equations:

$$\rho_f \frac{\partial \mathbf{u}}{\partial t} + \rho_f [\mathbf{u} \cdot \nabla_x] \mathbf{u} - \mu_f \Delta_x \mathbf{u} + \nabla_x p = \mathbf{f}_f \qquad \text{in } \Omega$$

$$\nabla_x \cdot \mathbf{u} = 0 \qquad \text{in } \Omega.$$

Saddle point problem:

$$\begin{pmatrix} \mathcal{F} & \mathcal{B}^* \\ \mathcal{B} & O \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ p \end{pmatrix} = \begin{pmatrix} \mathbf{f}_f \\ O \end{pmatrix}$$

$$A\mathbf{x} = \mathbf{b}$$

Equivalent to (Block LU):

$$\begin{pmatrix} \mathcal{F} & O \\ \mathcal{B} & \mathcal{S} \end{pmatrix} \begin{pmatrix} I & \mathcal{F}^{-1}\mathcal{B}^* \\ O & I \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ p \end{pmatrix} = \begin{pmatrix} \mathbf{f}_f \\ O \end{pmatrix}$$

$$LU\mathbf{x} = \mathbf{b}$$

where the operator $\mathcal{S} = \mathcal{B}\mathcal{F}^{-1}\mathcal{B}^*$ corresponds to the Schur complement.

Solution Strategies

Splitting vs. non splitting

Splitting the solutions for **u** and *p*

+ solve s.p.d. system (CG) for the pressure (advantage in memory and time)
- nested iterations to reach convergence
- splitting error

Solving the monolithic system as a whole

- global nonsymmetric undefinite system (GMRES)
+ no nested subiterations
+ via preconditioning one can recover some splitting methods

The separation between the two may not be crystal clear

Discretization

$$\begin{pmatrix} \mathcal{F} & \mathcal{B}^* \\ \mathcal{B} & O \end{pmatrix} \overset{skipping}{\Rightarrow} \begin{pmatrix} F & B^T \\ B & O \end{pmatrix} \tag{1}$$

Skipping the details about:

- time discretization (BDF);
- space discretization (FEM, FVM, FDM);
- nonlinearity (Newton, FP).

At the end of the discretization in general you end up with a block matrix like (1).

Splitting Schemes

Algebraic system:

$$\begin{pmatrix} F & O \\ B & S \end{pmatrix} \begin{pmatrix} I & F^{-1}B^T \\ O & I \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ p \end{pmatrix} = \begin{pmatrix} \mathbf{f}_f \\ O \end{pmatrix}$$

$$LU\mathbf{x} = \mathbf{b}$$

with the Schur complement $S = BF^{-1}B^T$. Preconditioned system

$$P^{-1}LU\mathbf{x} = P^{-1}\mathbf{b}$$

Choice:

$$P = \widetilde{L}\widetilde{U} = \begin{pmatrix} F & O \\ B & \widetilde{S} \end{pmatrix} \begin{pmatrix} I & \widetilde{F}^{-1}B^T \\ O & I \end{pmatrix}$$

Richardson:

$$P\mathbf{x}^{k+1} = P\mathbf{x}^k + \omega(\mathbf{b} - A\bar{\mathbf{x}}^k)$$

$$\begin{pmatrix} F & O \\ B & \widetilde{S} \end{pmatrix} \begin{pmatrix} I & -\widetilde{F}^{-1}B^T \\ O & I \end{pmatrix} \begin{pmatrix} \mathbf{u}^{k+1} - \mathbf{u}^k \\ p^{k+1} - p^k \end{pmatrix} = \begin{pmatrix} \omega_1 \\ \omega_2 \end{pmatrix} \left[\begin{pmatrix} F & B^T \\ B & O \end{pmatrix} \begin{pmatrix} \mathbf{u}^k \\ p^k \end{pmatrix} - \begin{pmatrix} \mathbf{f}_f \\ O \end{pmatrix} \right]$$

Splitting Scheme II

$$\begin{pmatrix} F & O \\ B & \widetilde{S} \end{pmatrix} \begin{pmatrix} I & -\widetilde{F}^{-1}B^T \\ O & I \end{pmatrix} \begin{pmatrix} \delta\mathbf{u}^{k+1} \\ \delta p^{k+1} \end{pmatrix} = \begin{pmatrix} 1 \\ \omega_2 \end{pmatrix} \begin{pmatrix} \mathbf{r}_1^k \\ \mathbf{r}_2^k \end{pmatrix}$$

with the Schur complement $S = BF^{-1}B^T$.

Uzawa	Chorin	Yosida	PCD
	[Gervasio Saleri Veneziani 2006]		[Elman et al. 2004]
$\widetilde{F} = O$	$\widetilde{F} = \frac{\alpha}{\delta t} M$	$\widetilde{F} = F$	$\widetilde{S} = \mathcal{B}\mathcal{B}^* \mathcal{F}_p^{-1}$
$\widetilde{S} = S$	$\widetilde{S} = B\frac{\delta t}{\alpha} M^{-1} B^T$	$\widetilde{S} = B\frac{\delta t}{\alpha} M^{-1} B^T$	$\Rightarrow S^{-1} = F_p (BB^T)^{-1}$
$\omega_2 = 1$	$\omega_2 = 1$	$\omega_2 = 1$	$\omega_2 = 1$

Orthomin	SIMPLE
[Houseaux & al. 2011]	[Klaij, Vuik 2013]
$\widetilde{F} = M^{-1}(\frac{\alpha}{\delta t} + \frac{1}{\tau_1})^{-1}; \widetilde{F} = F$	$\widetilde{F} = \text{diag}(F)$
$\widetilde{S} = BB^T M(\frac{\alpha}{\delta t} + \frac{1}{\tau_1})$	$\widetilde{S} = BB^T \text{ diag}(F)^{-1}$
$\omega_2 = \frac{<\mathbf{r}_2^k, S\widetilde{S}^{-1}\mathbf{r}_2^k>}{\|S\widetilde{S}^{-1}\mathbf{r}_2^k\|}$	$\omega_2 = 1$

Non splitting

Preconditioned Krylov:

$$\widetilde{L}^{-1} A\mathbf{x} = \widetilde{L}^{-1}\mathbf{b}$$

The $\widetilde{L}$ or $\widetilde{L}\widetilde{U}$ shown in the previous slide can be efficiently used for preconditioning a Krylov method (e.g. GMRES) on the monolithic system

Uzawa	Chorin	Yosida	PCD
$\widetilde{F} = O$	$\widetilde{F} = \frac{\alpha}{\delta t} M$	$\widetilde{F} = F$	
$\widetilde{S} = S$	$\widetilde{S} = B\frac{\delta t}{\alpha} M^{-1} B^T$	$\widetilde{S} = B\frac{\delta t}{\alpha} M^{-1} B^T$	$\widetilde{S} = \mathcal{B}\mathcal{B}^T \mathcal{F}_p^{-1}$
max. 2 it. [Murray et al.]	$\omega_2 = 1$	$\omega_2 = 1$	$\Rightarrow S^{-1} = F_p (BB^T)^{-1}$

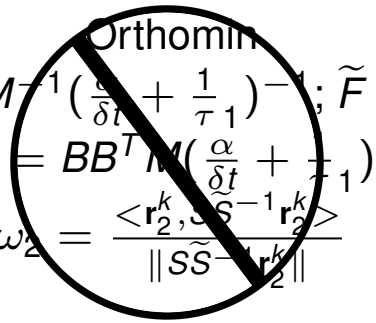

Orthomin

$$\widetilde{F} = M^{-1}(\frac{\alpha}{\delta t} + \frac{1}{\tau_1})^{-1}; \widetilde{F} = F$$
$$\widetilde{S} = BB^T M(\frac{\alpha}{\delta t} + \frac{1}{\tau_1})$$
$$\omega_2 = \frac{<\mathbf{r}_2^k, S\widetilde{S}^{-1}\mathbf{r}_2^k>}{\|S\widetilde{S}^{-1}\mathbf{r}_2^k\|}$$

SIMPLE

$$\widetilde{F} = \text{diag}(F)$$
$$\widetilde{S} = BB^T \text{diag}(F)^{-1}$$
$$\Rightarrow \widetilde{S}^{-1} = \text{diag}(F)(BB^T)^{-1}$$

ILU, AAS, AMG
regardless of the physics

Example Codes Ported on JUQUEEN

Name	Space discr.	Preconditioner	Linear Sol.	Language
Alya	FEM	Chorin, Yosida	Orthomin, Richardson	Fortran
HiFLOW3	FEM	ILU	GMRES	C++
OpenFOAM	FVM	SIMPLE (PISO, PIMPLE)	Richardson	C++
LifeV	FEM	PCD, SIMPLE, Yosida	GMRES	C++
Dolfin	FEM	PCD	Richardson	C++

Benchmark: FDA Nozzle

Simplified medical device. Experimental results available for comparison [Hariharan et al. 2011], [Passerini et al., 2013]

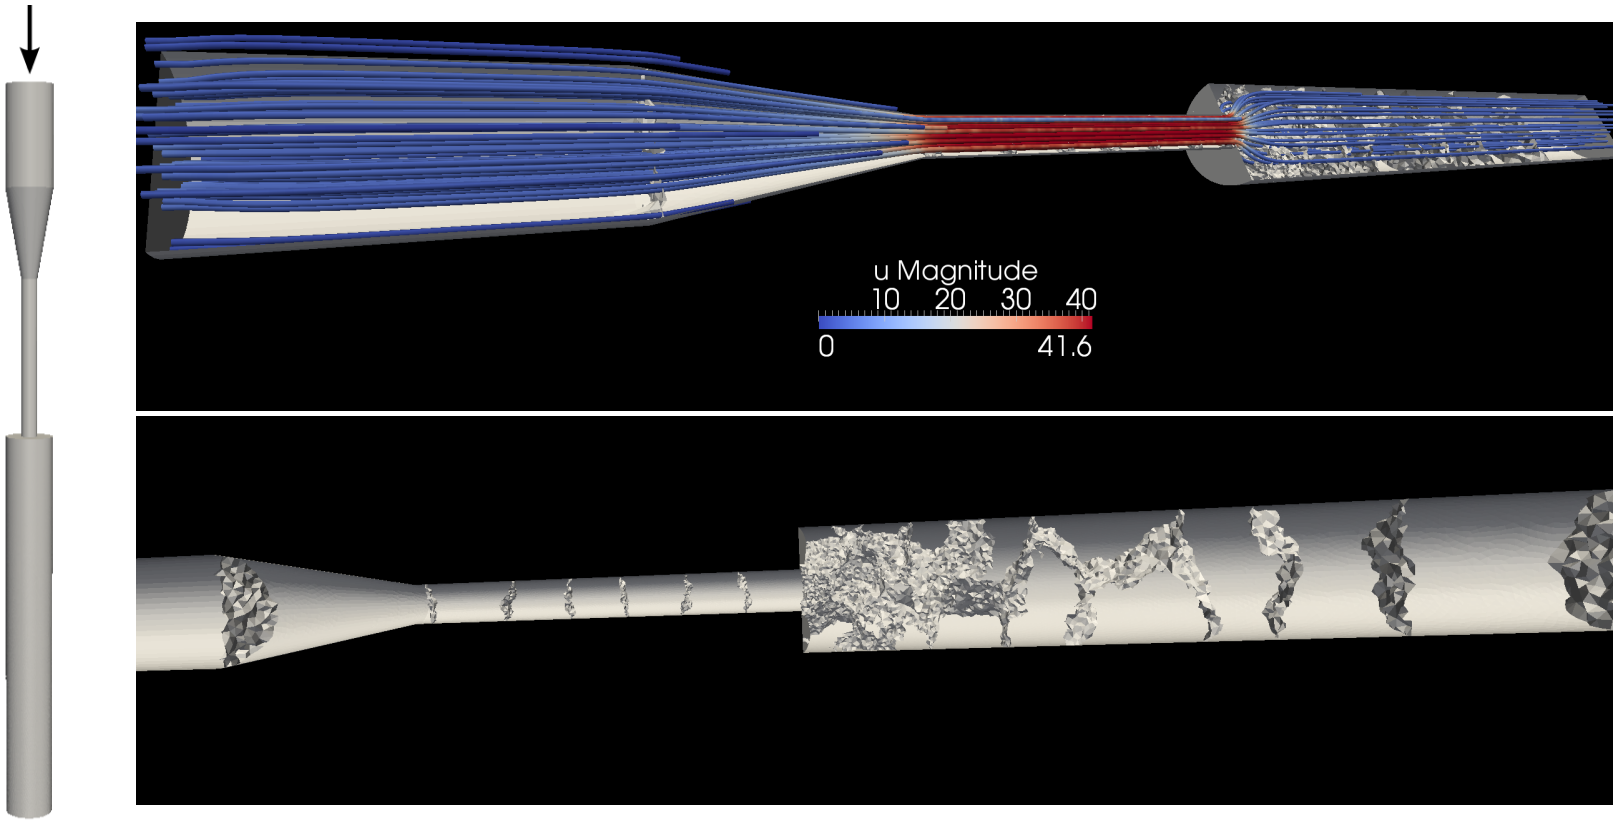

Profiling 2 Solvers (Scalasca)

$\approx$ 916K tetrahedra, on 512 cores (using 16 MPI processes per node).

HiFLOW3 (ILU) with Dr. Staffan Ronnas

	s. (sum)	visits
Preconditioner computation	15.5	4
Preconditioner application	45.5	1896
Dot product in GMRES	83.3	449'220[1]

Tolerances: 10^{-10} for GMRES and 10^{-6} for Newton

- Lot of time wasted in the GMRES dot products!! (in particular in the call to MPI_AllReduce)
- ILU preconditioner for the monolithic system: expensive to build and to apply

Alya (Orthomin) wit Dr. Herbert Owen

	s. (sum)	visits
momentum[2]	0.08	3 (8 iterations)
pressure schur[3]	0.28	3 (11 iterations)
momentum[4]	0.35	3 (70 iterations)

Tolerances: 10^{-9} for the linear solvers, 0.06 for FP

- Monolithic problem split into several small problems (splitting method)
- Simple and fast preconditioners for the sub-problems.

[1] 474 GMRES it. per solve

[2] GMRES, advection

[3] deflated CG

[4] GMRES, computing ω_2

Why?

- GMRES algorithm: Once computed $\widetilde{\mathbf{v}}_{i+1} = A\mathbf{v}_i$ orthogonalize w.r.t the Krylov space:
- projection $\mathbf{v}_{i+1} = \widetilde{\mathbf{v}}_{i+1} - \sum_{j=0}^{j=i} < \widetilde{\mathbf{v}}_{i+1}, \mathbf{v}_j > \mathbf{v}_j$
- if n is the number of iterations, this means $\frac{n^2}{2}$ scalar products
- for CG (or MINRES) the scalar products $< \widetilde{\mathbf{v}}_i, \mathbf{v}_j >$ magically disappear for $j < i - 1$ (three terms recursion formula!!), so that the number of products per iteration is 3, independent of n.
- Monolithic+ILU: $\frac{474^2}{2} = 112330$ “big” scalar products, versus
- Orthomin: we have $\frac{8^2}{2} + \frac{70^2}{2} + 11 \cdot 3 = 2515$ “small” scalar products

Viable solutions?

- Change preconditioner
- Use the “restarted" version
- Improve the scalar product performance

JÜLICH FORSCHUNGSZENTRUM

Table of contents

Motivations

JÜLICH FORSCHUNGSZENTRUM

OpenFOAM: Industrial Gas Turbines

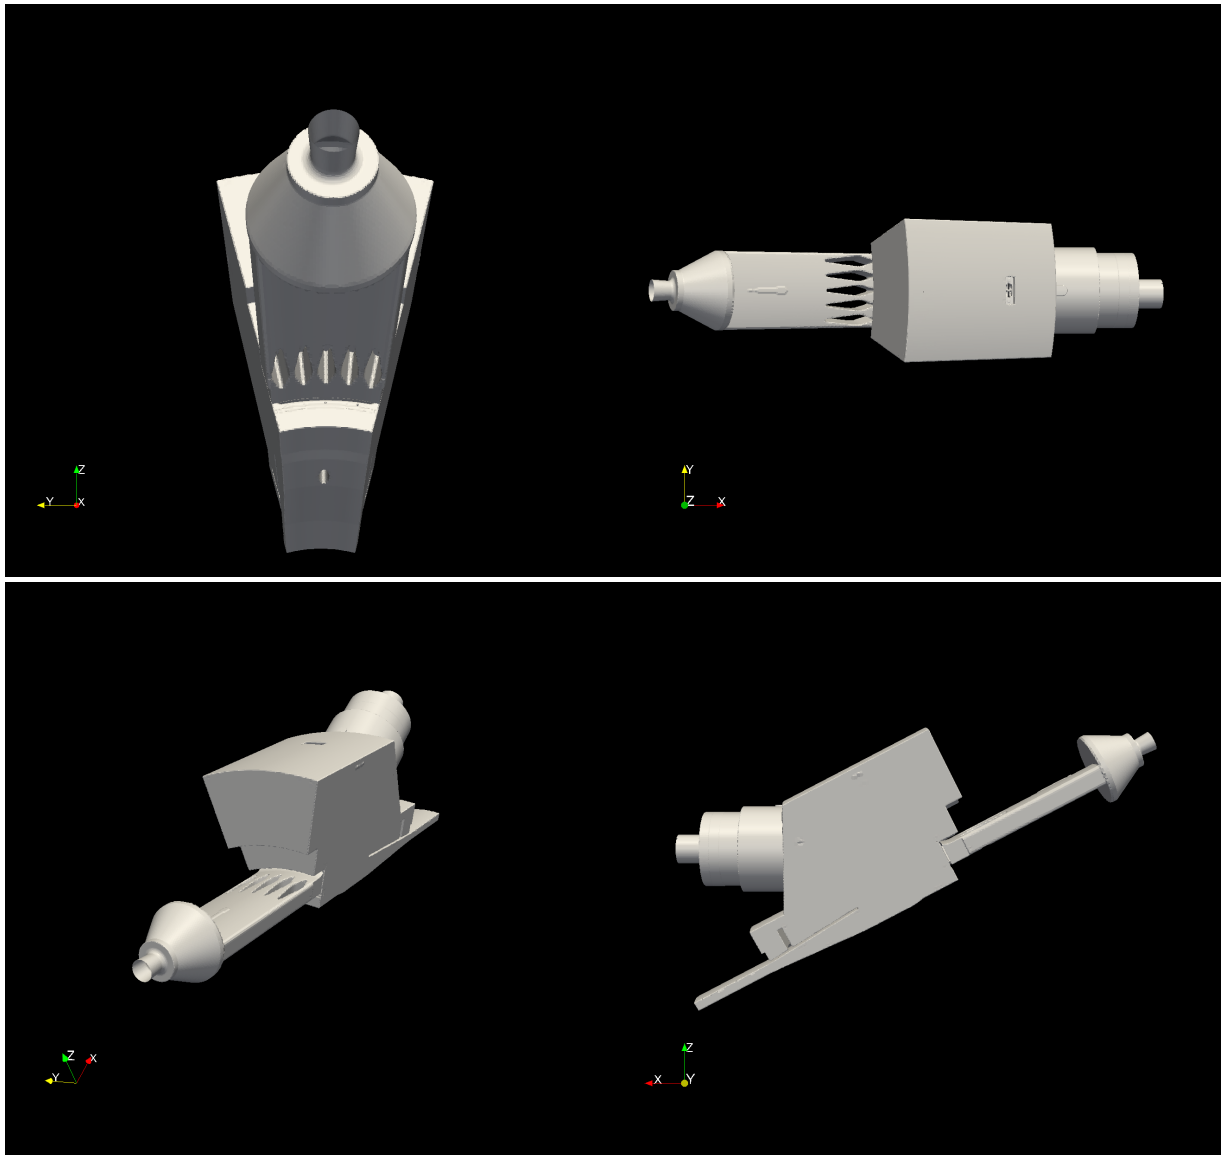

Simulations

- goal: understanding thermo–acoustic instabilities in industrial Siemens gas turbines.
- Model: compressible Navier–Stokes coupled with heat equation and turbulence models.
- Solver: PIMPLE, an implicit variant of the PISO algorithm.
- Mesh: 22’422’878 points:

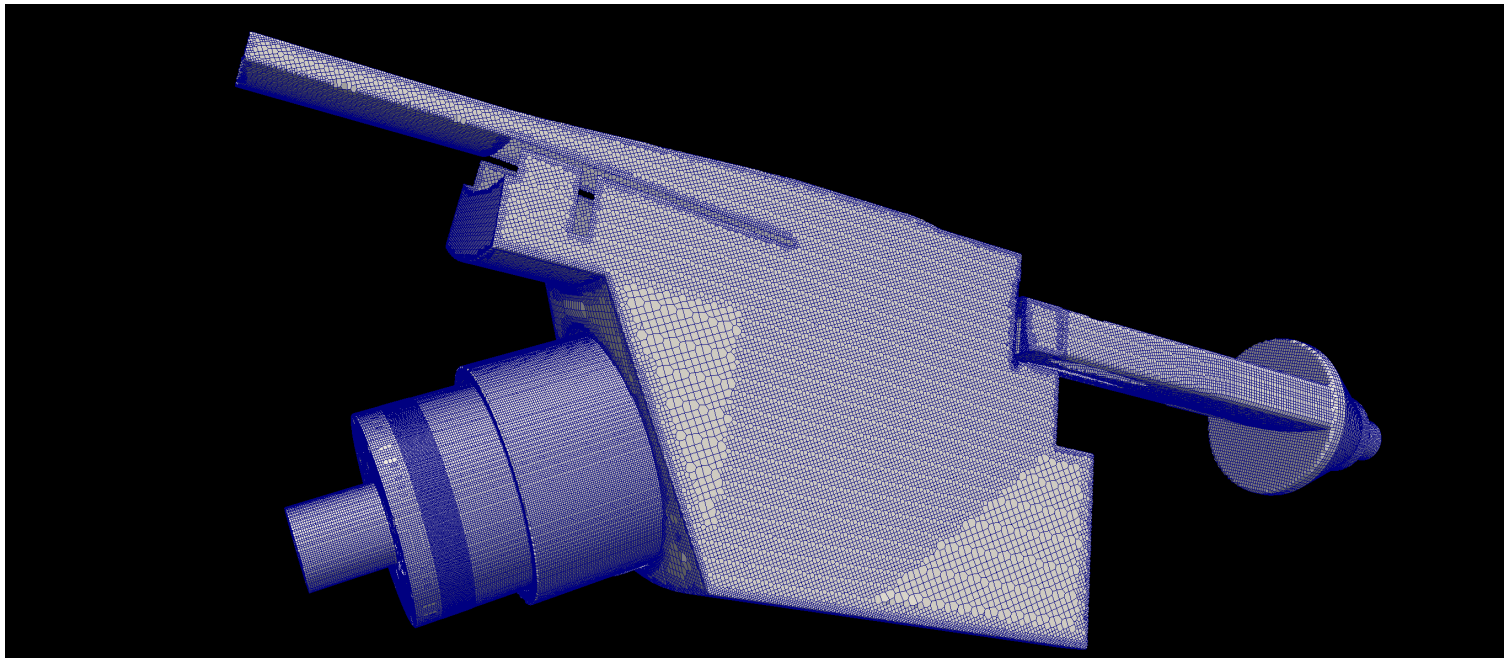

Table of contents

Motivations

With LHTC (P. Reymond), CMCS (A. Quarteroni, S. Deparis)

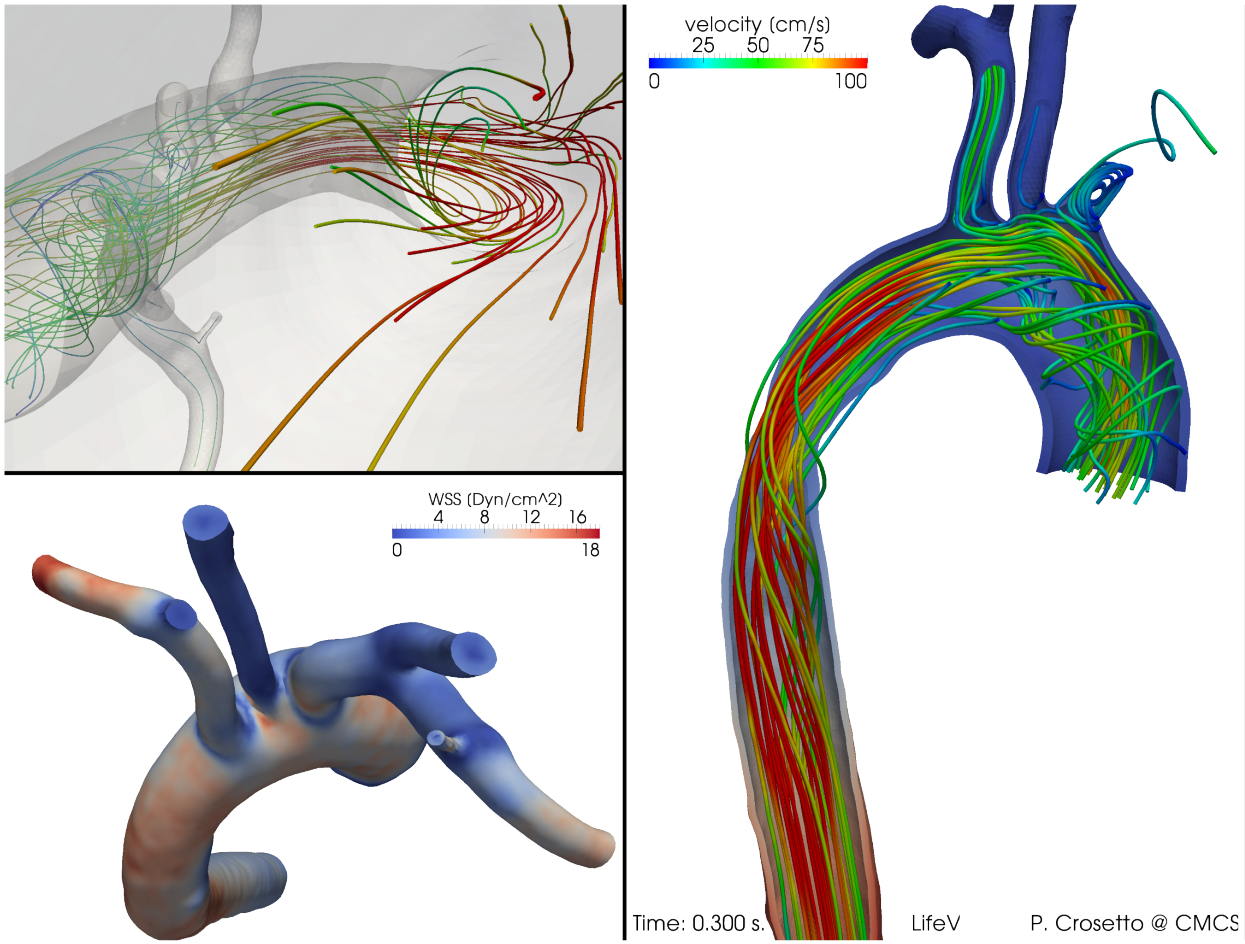

Idea of the ALE formulation

Fluid equations in a moving domain

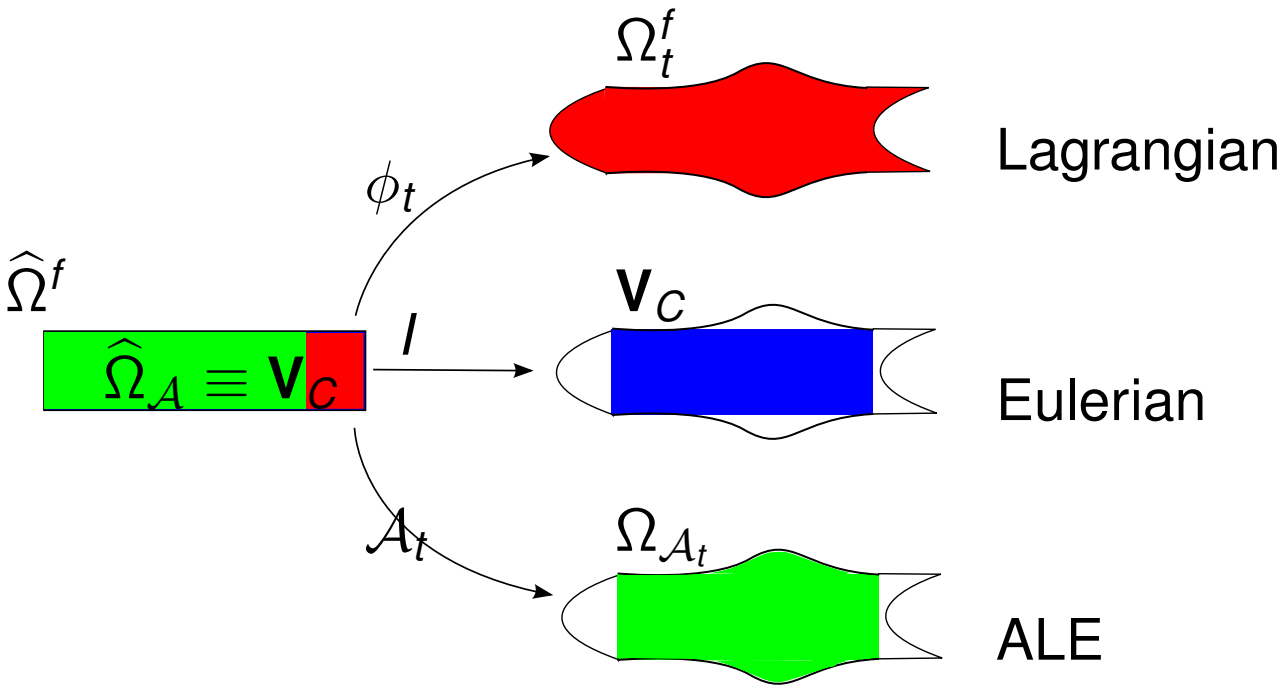

$\phi_t : \widehat{\Omega} \to \Omega_t$ is the deformation map; $\mathcal{A}_t : \widehat{\Omega}_{\mathcal{A}} \to \Omega_{\mathcal{A}_t}$ is the ALE map;

FSI problem: coupling conditions

Continuity of stresses

$$\widehat{\sigma}_f \widehat{\mathbf{n}}^f + \boldsymbol{\Pi} \widehat{\mathbf{n}}^s = 0 \qquad \text{on } \widehat{\Gamma}$$

Continuity of velocities

$$\frac{\partial \widehat{\mathbf{d}}_s}{\partial t} = \mathbf{u} \circ \mathcal{A}_t \qquad \text{on } \widehat{\Gamma}$$

Geometry adherence

$$\frac{\partial \widehat{\mathbf{d}}_s}{\partial t} = \mathbf{w} \circ \mathcal{A}_t \qquad \text{on } \widehat{\Gamma}$$

Three fields formulation

$$\mathrm{F}(\mathbf{u}_f, \widehat{\mathbf{d}}_f, \widehat{\mathbf{d}}_s) = 0 \text{ in } \Omega_{\mathcal{A}_t} \qquad \text{fluid problem, unknowns } \mathbf{u}_f = (\mathbf{u}, p)$$

$$\mathrm{S}(\mathbf{u}_f, \widehat{\mathbf{d}}_s) = 0 \text{ in } \widehat{\Omega}^s \qquad \text{structure problem, unknown } \widehat{\mathbf{d}}_s$$

$$\mathrm{G}(\widehat{\mathbf{d}}_f, \widehat{\mathbf{d}}_s) = 0 \text{ in } \widehat{\Omega}_{\mathcal{A}} \qquad \text{geometry problem, unknown } \widehat{\mathbf{d}}_f$$

Block preconditioners

with S. Deparis

we have a 3 blocks matrix A defining a coupled problem. Idea:

- Gauss-Seidel approximation $P \approx A$;
- Factorization $P = P_F P_S P_G$ (each block containing one of the three problems);
- Precondition each factor using an efficient strategy $\tilde{P} = P_\star(P_F)P_\star(P_S)P_\star(P_G)$.

Examples:

- $(AAS)^3$: algebraic additive Schwarz preconditioners for the three factors;
- $aPCD - (AAS)^2$: approximated PCD ([Elman, JSC, 2005]) for the fluid, AAS for the other factors (PhD work by **Gwenol Grandperrin**);
- $aPCD - (AMG)^2$: approximated PCD for the fluid, algebrainc multigrid for the other two blocks;

Advantages of the factorization :

- modularity;
- memory requirements and computational cost;

note: no nested iterative solutions.

LifeV

- C++ finite element library
- parallel and serial versions
- distributed under LGPL

- CMCS – EPFL
- Math/CS – Emory
- MOX – Polimi
- REO – INRIA

based on MPI, ParMETIS, and Trilinos

- All processors read the same mesh
- ParMetis partitions the mesh
- Finite elements: loop on local mesh
- Linear solver in parallel
- post-processing on hdf5 or separate ensight files

Greenshields-Weller benchmark

(a) From Greenshields & Weller, 2005

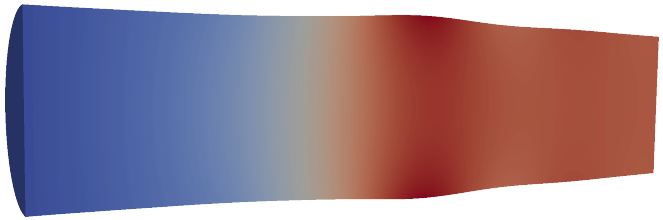

(b) detail of the mesh

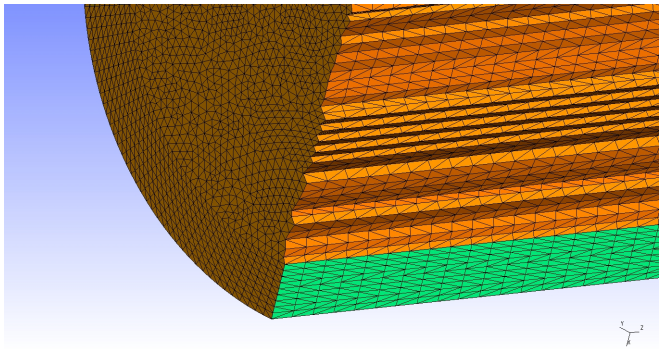

BC: pressure step function at inlet
Comparison between several preconditioners:

- $(AAS)^3$
- $aPCD - (AAS)^2$
- $aPCD - (AMG)^2$

Weak scalability on Cray XE6

Results on 2 meshes, with label $_1$ (finer), and $_2$ (coarser).

h_1/h_2	$procs_1/procs_2$	dof_1/dof_2
0.77	512/256 = 2	2'970'864/1'473'624=2.016

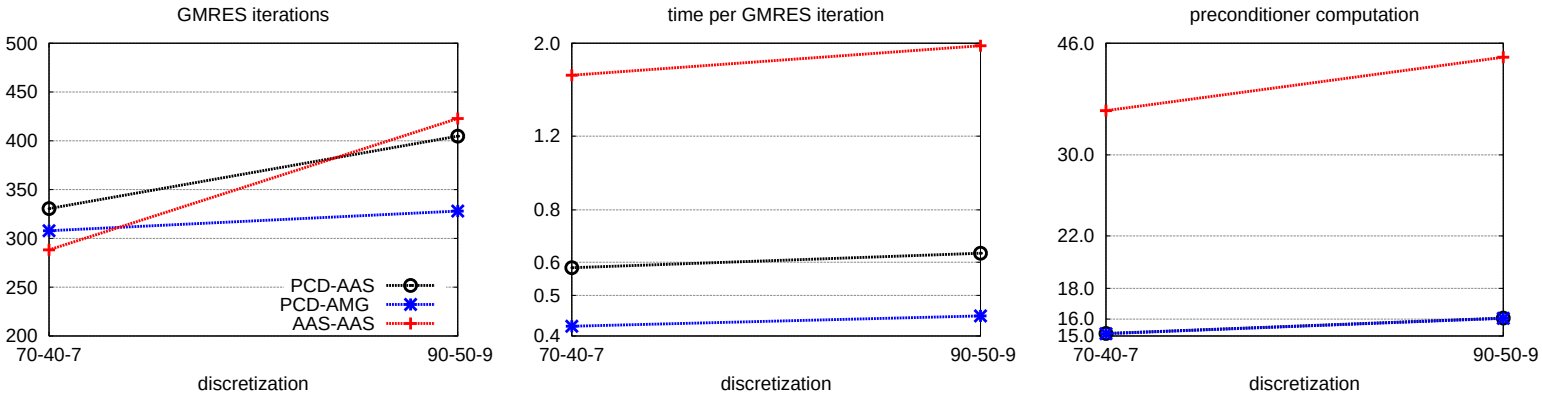

Figure : GCE — Weak scalability results on Cray XE6.

	$(AAS)^3$	$aPCD-(AAS)^2$	$aPCD-(AMG)^2$
$\frac{iter_1}{iter_2}$ (optimal≈1)	1.55	1.27	1.08
$\frac{prec_1}{prec_2}$ (optimal≈1)	1.15	1.10	1.10
$\frac{gmres\_time_1}{gmres\_time_2}$ (optimal≈1)	1.19	1.17	1.15

On Blue Gene/P

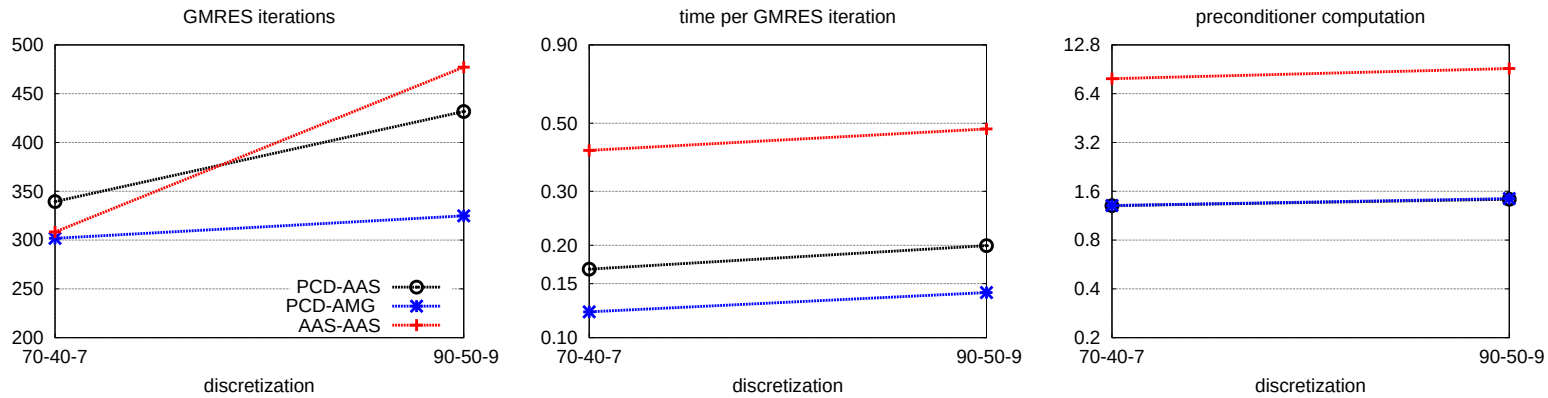

Figure : GCE — Weak scalability results on BG/P.

	$(AAS)^3$	$aPCD-(AAS)^2$	$aPCD-(AMG)^2$
$\frac{iter_1}{iter_2}$ (optimal≈1)	1.47	1.22	1.07
$\frac{prec_1}{prec_2}$ (optimal≈1)	1.22	1.06	1.06
$\frac{gmres\_time_1}{gmres\_time_2}$ (optimal≈1)	1.18	1.08	1.05

\+ **Best weak scalability obtained with the $aPCD-(AMG)^2$;**