

PFEM performance meeting

Last updated on Monday 3rd May, 2021

Simon FÉVRIER

Table of contents

1. Available solvers
2. Performance with number of nodes
3. OpenMP performance

Available solvers

Monolithic PSPG (implicit)

Add a term proportional to the momentum equation to the continuity equation:

$$\begin{bmatrix} \frac{1}{\Delta t} \tilde{\mathbf{M}}_{\rho_0}^{n+1} + \mathbf{K}_{\mu}^{n+1} & -\mathbf{D}^{\top n+1} \\ \frac{1}{\Delta t} \mathbf{C}_{\tau}^{n+1} + \mathbf{D}^{n+1} & \mathbf{L}_{\tau/\rho}^{n+1} \end{bmatrix} \begin{pmatrix} \mathbf{v}^{n+1} \\ \mathbf{p}^{n+1} \end{pmatrix} = \begin{bmatrix} \mathbf{f}_{\text{ext}}^{n+1} + \frac{1}{\Delta t} \tilde{\mathbf{M}}_{\rho_0}^{n+1} \mathbf{v}^n \\ \mathbf{h}_{\tau}^{n+1} + \frac{1}{\Delta t} \mathbf{C}_{\tau}^{n+1} \mathbf{v}^n \end{bmatrix}$$

$$\mathbf{x}^{n+1} = \mathbf{x}^n + \mathbf{v}^{n+1} \Delta t$$

- uses Eigen::SparseLU linear solver \rightarrow no parallel computing,
- matrix not symmetric \rightarrow cannot use iterative solver from Eigen,
- matrix of size $(\text{mesh dim} + 1) \times N_{\text{nodes}}$.

Fractional Step (implicit)

Split the computation in two step: compute an approximation of the velocity, then correct the pressure to obtain divergence free velocity:

$$\tilde{\mathbf{M}}_{\rho_0}^{n+1} \frac{\tilde{\mathbf{v}}^{n+1} - \mathbf{v}^n}{\Delta t} + \mathbf{K}_{\mu}^{n+1} \mathbf{v}^{n+1} - \gamma_{FS} \mathbf{D}^{\top, n+1} \mathbf{p}^n = \mathbf{f}_{\text{ext}}^{n+1}$$

$$\frac{\Delta t}{\rho_0} \mathbf{L} \delta \mathbf{p} = -\mathbf{D} \tilde{\mathbf{v}}^{n+1}$$

$$\tilde{\mathbf{M}}_{\rho_0}^{n+1} \frac{\mathbf{v}^{n+1} - \tilde{\mathbf{v}}^{n+1}}{\Delta t} = \mathbf{D}^{\top, n+1} \delta \mathbf{p}$$

$$\mathbf{x}^{n+1} = \mathbf{x}^n + \mathbf{v}^{n+1} \Delta t \quad \mathbf{p}^{n+1} = \gamma_{FS} \mathbf{p}^n + \delta \mathbf{p}$$

- all matrices are symmetric,
- iterative solver can thus be used (e.g. Eigen::ConjugateGradient,
- smaller matrices.

Weakly compressible (explicit)

Express the density as a function of the pressure and use an explicit compressible solver:

$$\bar{\mathbf{v}}^{n+1} = \mathbf{v}^n + 0.5\Delta t \mathbf{a}_n,$$

$$\bar{\mathbf{x}}^{n+1} = \mathbf{x}^n + \Delta t \mathbf{v}_n + 0.5\Delta t^2 \mathbf{a}_n = \mathbf{x}_n + \Delta t \bar{\mathbf{v}}^{n+1},$$

$$\mathbf{M}\rho^{n+1} = \mathbf{M}\rho^n - \Delta t \mathbf{D}\rho \bar{\mathbf{v}}^{n+1},$$

$$\left[\mathbf{p}^{n+1} \right]_i = \rho^{-1} \left(\left[\rho^{n+1} \right]_i \right),$$

$$\tilde{\mathbf{M}}_\rho \mathbf{a}^{n+1} = \mathbf{f}_{\text{ext}}^{n+1} - \mathbf{K}_\mu \bar{\mathbf{v}}^{n+1} + \mathbf{D}^T \mathbf{p}^{n+1}$$

$$\mathbf{v}^{n+1} = \bar{\mathbf{v}}^{n+1} + 0.5\Delta t \mathbf{a}_{n+1}$$

$$\mathbf{x}^{n+1} = \bar{\mathbf{x}}^{n+1}$$

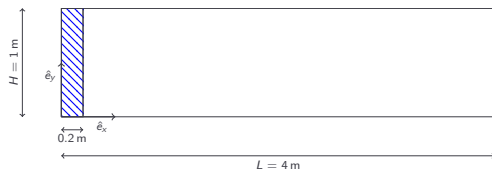
$$\Delta t = C \min_e \left(\frac{h_e}{\max(\|\vec{\mathbf{v}}\|, \sqrt{K/\rho})} \right)$$

- mass matrices are lumped \rightarrow trivial inversion,
- good parallel properties.

Performance with number of nodes

2. Performance with number of nodes

Solving time - Problem used



Initial geometry of the flow between two plates problem.

ρ	1000 kg m^{-3}
μ	0.001 Pa s

Incompressible fluid parameters

ρ^*	1000 kg m^{-3}
μ	0.001 Pa s
K_0	$2\,200\,000 \text{ Pa}$
K'_0	7.6

Weakly compressible fluid parameters

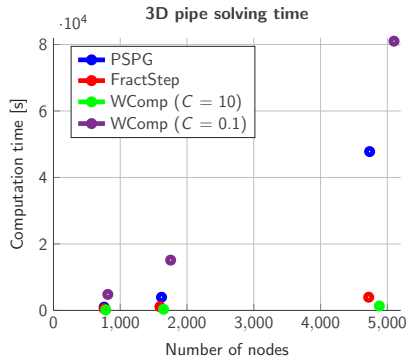
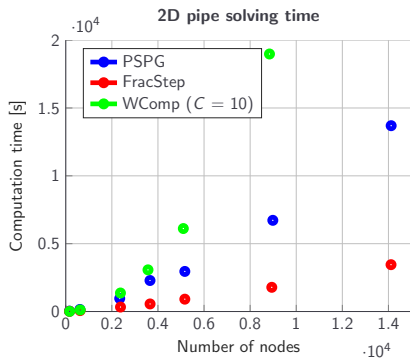
$h_{2D}(m)$	0.2	0.1	0.05	0.04	0.033	0.025	0.02
$h_{3D}(m)$	0.2	0.15	0.1				

Mesh characteristic size

N.B.: computation time computed on the whole run.

2. Performance with number of nodes

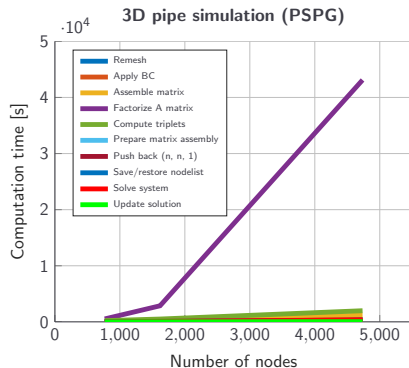
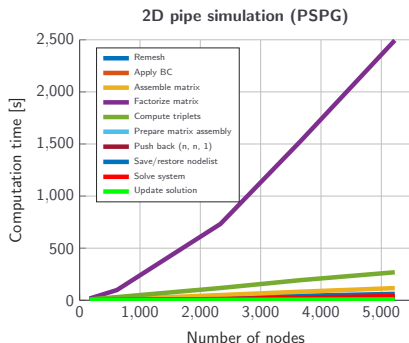
2D and 3D comparison of whole solvers



- In 2D, fractional step implicit solver is the best,
- In 3D, there is a way to find a combination of C and K_0 parameter such that the weakly compressible explicit code compete with the fractional step implicit one,
- PSPG code loose in each case.

2. Performance with number of nodes

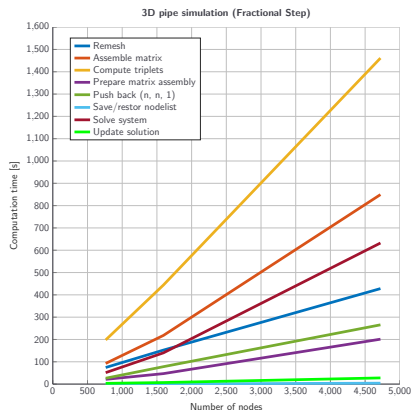
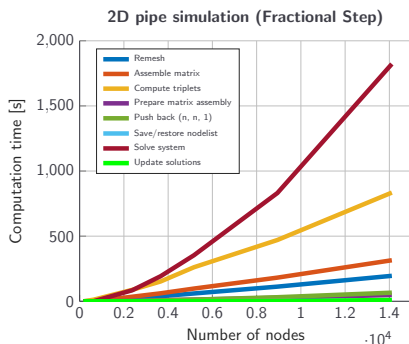
Repartition of computation time in PSPG



- the factorize matrix step takes most of the computation time → no need to check parallel efficiency of PSPG code,
- without finding a better solver than `Eigen::SparseLU`, nothing can be done.

2. Performance with number of nodes

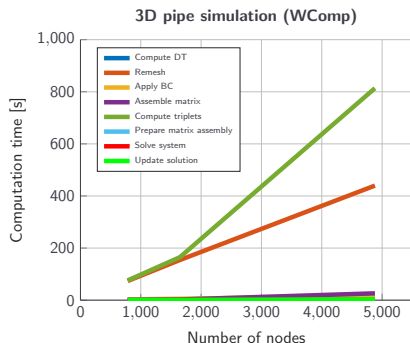
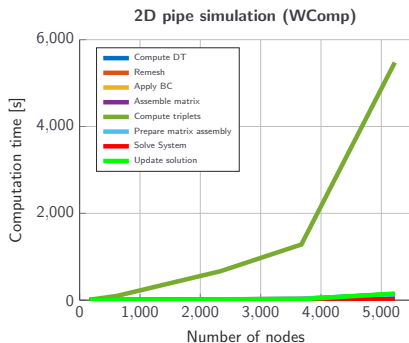
Repartition of computation time in Fractional Step



- no factorization step,
- "Solve system" and "Compute triplets" are parallel, "Assemble matrices" and "Remesh" are not.

2. Performance with number of nodes

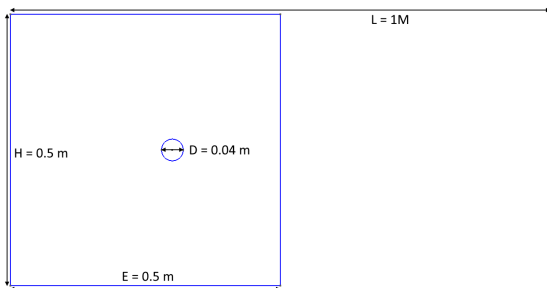
Repartition of computation time in Weakly Compressible



- "Compute triplets" takes a lot of time,
- but it is parallel.

2. Performance with number of nodes

Remeshing time - Problem used 2



Initial geometry of the problem (initial velocity at the left boundary, 0.04 m s^{-1}).

ρ	1000 kg m^{-3}
μ	0.015 Pa s

Incompressible fluid parameters

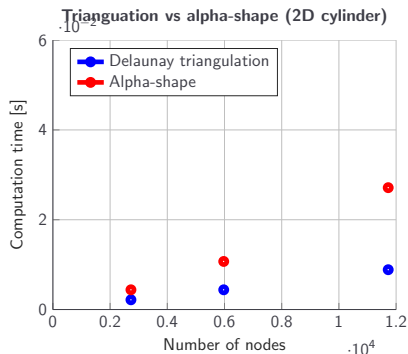
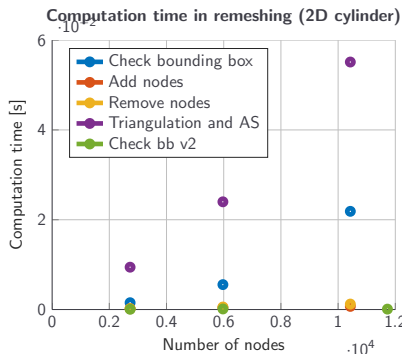
$h_{2D}(m)$	0.15	0.1	0.75
-------------	------	-----	------

Mesh characteristic size

N.B.: computation time computed on one function execution.

2. Performance with number of nodes

Repartition of computation time in remeshing



- the `Mesh::triangulateAlphaShape2D` function takes a little bit more time than what CGAL do (expected but can it be decreased ?),
- the `Mesh::checkBoundingBox` takes a lot of time but it can be fixed by changing the order of the functions in `Mesh::remesh`.

2. Performance with number of nodes

Further work

Welcome x r003hs x

Hotspots by CPU Utilization

Analysis Configuration Collection Log Summary Bottom-up Caller/Callee Top-down Tree Platform

Elapsed Time: 138.029s

CPU Time: 135.757s

Total Thread Count: 1

Paused Time: 0s

Top Hotspots

This section lists the most active functions in your application. Optimizing these hotspot functions typically results in improving overall application performance.

Function	Module	CPU Time
malloc	msvcr.dll	47.180s
free	msvcr.dll	33.106s
Eigen::internal::general_matrix_vector_product<long long, double, Eigen::internal::const_blas_data_mapper<double, long long, (int)0>, (bool)0, double, Eigen::internal::const_bias_data_mapper<double, long long, (int)1>, (bool)0, (int)0>::run	libpfemSimulation.dll	2.227s
Eigen::internal::general_matrix_vector_product<long long, double, Eigen::internal::const_blas_data_mapper<double, long long, (int)1>, (int)1, (bool)0, double, Eigen::internal::const_bias_data_mapper<double, long long, (int)0>, (bool)0, (int)0>::run	libpfemSimulation.dll	2.203s
Eigen::internal::general_matrix_vector_product<long long, double, Eigen::internal::const_blas_data_mapper<double, long long, (int)0>, (bool)0, double, Eigen::internal::const_bias_data_mapper<double, long long, (int)1>, (bool)0, (int)0>::run	libpfemSimulation.dll	1.940s
[Others]	N/A*	49.098s

Intel VTune Profiler hotspot analysis of weakly compressible explicit code.

- comes from the use of `Eigen::MatrixXd` which are dynamic and thus needs to allocate memory,
- can be fixed by using fixed size matrices but refactoring of code needed to handle correctly 2D and 3D cases (test: 1712 \rightarrow 1000 s on a pipe problem with $h = 0.04$),
- implicit codes have less this problem (6s/26s).

OpenMP performance

Speed-up ratio and parallel efficiency definition

Let τ be a vector containing the execution time of the program for a certain number of thread such that τ_i is the execution time using i threads.

The speed-up ratio is defined as:

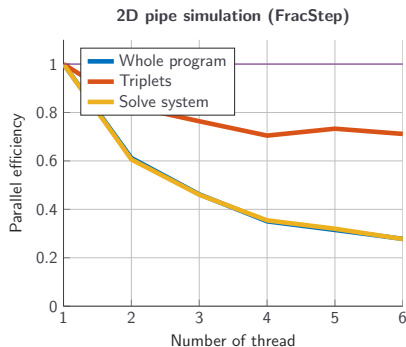
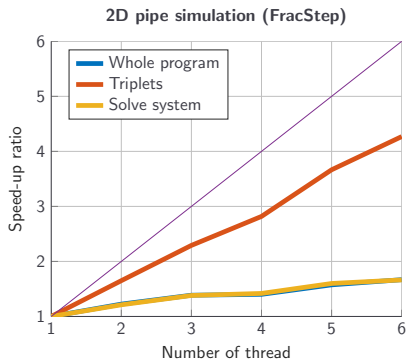
$$\gamma = \frac{\tau_1}{\tau_i}$$

The parallel efficiency is defined as:

$$\eta = \frac{\tau_1}{i\tau_i}$$

3. OpenMP performance

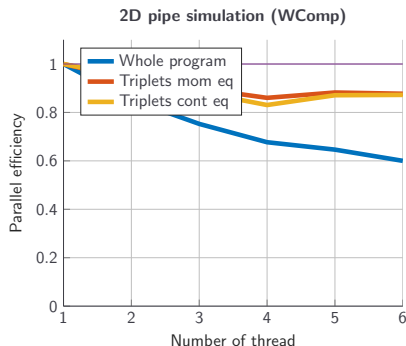
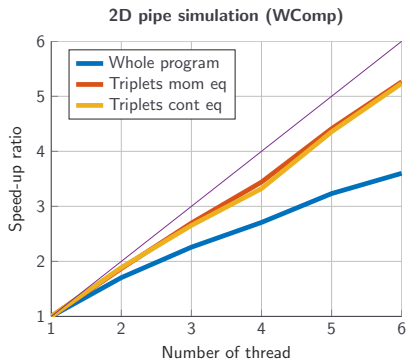
Speed-up ratio and parallel efficiency of Fractional Step solver



- efficiency is killed by `Eigen::ConjugateGradient` poor performance,
- look to improve triplets computation efficiency.

3. OpenMP performance

Speed-up ratio and parallel efficiency of Weakly compressible solver



- global efficiency is decreased due to serial part of the code.