

HIGH-PERFORMANCE COMPUTING

Pierre Jolivet, IRIT-CNRS
`pierre.jolivet@enseeiht.fr`

CSMA Juniors
May 15

INTRODUCTION

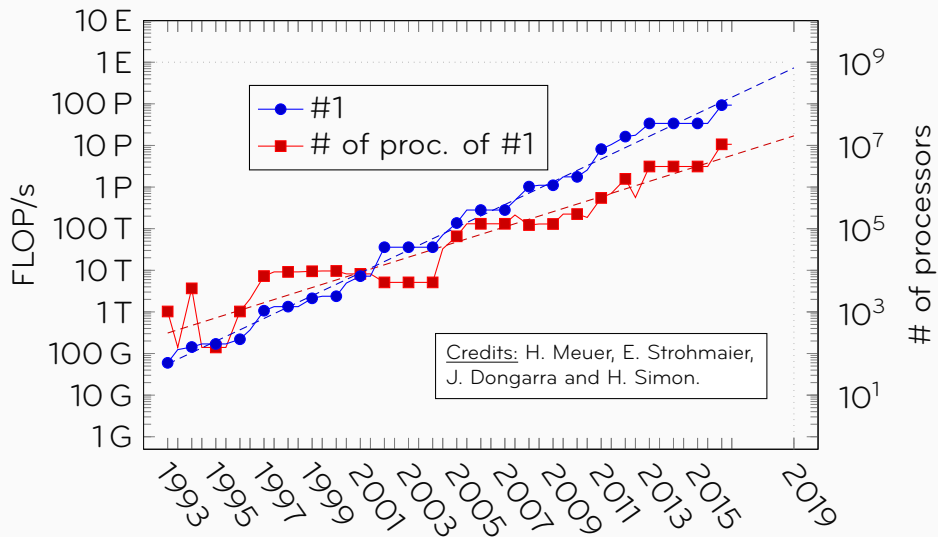
- platforms are ranked according to various metrics
- Floating Point Operations per second is critical
- many other important characteristics
- frequency-scaling has ended since circa 2005

- platforms are ranked according to various metrics
- Floating Point Operations per second is critical
- many other important characteristics
- frequency-scaling has ended since circa 2005

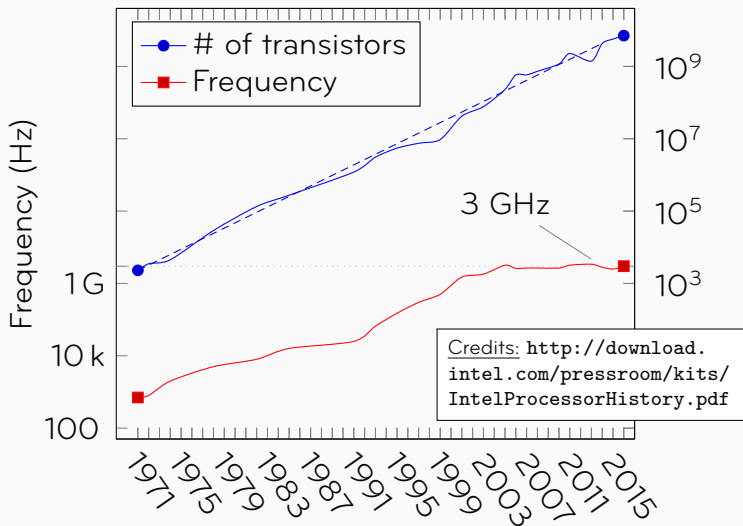
Most used paradigms/interfaces

- Message Passing Interface (<http://mpi-forum.org>)
- OpenMP (<http://www.openmp.org>)

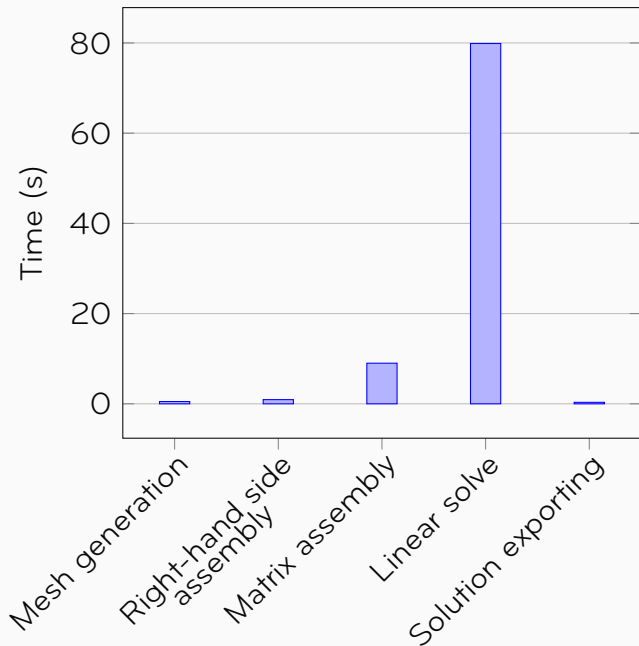
TOP 500



f-SCALING



IMPLICIT CODE



METHODOLOGY AND AVAILABLE LIBRARIES

FOCUS OF THIS MINI-COURSE

- efficient strategies for solving linear systems $Ax = b$
- appropriate use of modern architectures
- introduction to most commonly used libraries for HPC

STRATEGIES FOR SOLVING LINEAR SYSTEMS

Projection methods
Relaxation-based methods

Iterative methods

Multifrontal factorizations
Supernodal factorizations

Direct methods

Sparse MV & dot products

Low





Level of parallelism

Memory consumption

Robustness w.r.t. $\kappa(A)$

Intervention of the user

Dense MM products

High





STRATEGIES FOR SOLVING LINEAR SYSTEMS

Domain decomposition methods
Multigrid methods

Projection methods
Relaxation-based methods

"Hybrid" methods

Multifrontal factorizations
Supernodal factorizations

Iterative methods

Direct methods

Sparse MV & dot products

Low

~~X~~

✓

Level of parallelism

Memory consumption

Robustness w.r.t. $\kappa(A)$

Intervention of the user

Dense MM products

High

✓

~~X~~

Left preconditioning

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

Right preconditioning

$$AM^{-1}u = b$$

with $x = M^{-1}u$

Left preconditioning

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

Right preconditioning


$$AM^{-1}u = b$$

$$\text{with } x = M^{-1}u$$

$\implies M^{-1}$ must be cheap to apply!


A SIMPLE EXAMPLE

Instead of using the inverse of the distributed matrix




A SIMPLE EXAMPLE

“Invert” only the diagonal blocks to define M^{-1}



A SIMPLE EXAMPLE

“Invert” only the diagonal blocks to define M^{-1}



Known as the block Jacobi method (the most basic DDM)

GENERAL GUIDELINES

- if you want performance
- if you don't want to waste time
 - ⇒ do not reinvent the wheel

GENERAL GUIDELINES

- if you want performance
- if you don't want to waste time

⇒ do not reinvent the wheel

Two kinds of libraries in the HPC ecosystem:

- black-box solvers (direct methods, algebraic multigrid methods, iterative methods)
- general-purpose backends (collections of linear/nonlinear solvers, time-stepping methods)

Coarse-grained parallelism

- less synchronization
 - low amount of communication
- ⇒ domain decomposition over MPI

Coarse-grained parallelism

- less synchronization
- low amount of communication

⇒ domain decomposition over MPI

- METIS (<http://glaros.dtc.umn.edu/gkhome/metis/metis/overview>)
- SCOTCH (<http://www.labri.fr/perso/pelegrin/scotch>)

Coarse-grained parallelism

- less synchronization
 - low amount of communication
- ⇒ domain decomposition over MPI

- METIS (<http://glaros.dtc.umn.edu/gkhome/metis/metis/overview>)
- SCOTCH (<http://www.labri.fr/perso/pelegrin/scotch>)

Fined-grained parallelism

- more synchronization
 - better load-balance
- ⇒ graph coloring over threads

Discretization kernel

- after domain decomposition/coloring
- depends on the physics
- yields a linear system

Discretization kernel

- after domain decomposition/coloring
- depends on the physics
- yields a linear system

- MUMPS (<http://mumps.enseeiht.fr>)
- PaStiX (<http://pastix.gforge.inria.fr>)
- UMFPACK (<http://faculty.cse.tamu.edu/davis/suitesparse.html>)
- CHOLMOD
- MKL PARDISO (<https://software.intel.com>)
- PARDISO (<http://www.pardiso-project.org>)
- SuperLU (<http://crd-legacy.lbl.gov/~xiaoye/SuperLU>)

Discretization kernel

- after domain decomposition/coloring
- depends on the physics
- yields a linear system

- MUMPS (<http://mumps.enseeiht.fr>)
- PaStiX (<http://pastix.gforge.inria.fr>)
- UMFPACK (<http://faculty.cse.tamu.edu/davis/suitesparse.html>)
- CHOLMOD
- MKL PARDISO (<https://software.intel.com>)
- PARDISO (<http://www.pardiso-project.org>)
- SuperLU (<http://crd-legacy.lbl.gov/~xiaoye/SuperLU>)

- PETSc (<http://www.mcs.anl.gov/petsc>)
- Trilinos (<https://trilinos.org>)

- *hypre* (<http://www.llnl.gov/casc/hypre>)
- DUNE (<https://www.dune-project.org>)
- PARALUTION (<http://www.paralution.com>)
- HPDDM (<https://github.com/hpddm/hpddm>)

- PETSc (<http://www.mcs.anl.gov/petsc>)
- Trilinos (<https://trilinos.org>)

- *hypre* (<http://www.llnl.gov/casc/hypre>)
- DUNE (<https://www.dune-project.org>)
- PARALUTION (<http://www.paralution.com>)
- HPDDM (<https://github.com/hpddm/hpddm>)

COMPLETE DISTRIBUTED FINITE ELEMENT SOLVER

GOAL OF THE EXERCISE

- see the effect of mesh partitioning
- understand the consequences for a finite element kernel
- use a distributed linear algebra backend
- play with different solvers/preconditioners

GOAL OF THE EXERCISE

- see the effect of mesh partitioning
- understand the consequences for a finite element kernel
- use a distributed linear algebra backend
- play with different solvers/preconditioners

Dependencies

- Gmsh
- MPI
- FreeFem++
- PETSc

`http://jolivet.perso.enseeiht.fr/CSMA17_hpc.tar.gz`

1. mesh generation

STEPS RELATED TO THE MESH

1. mesh generation
2. mesh decomposition

STEPS RELATED TO THE MESH

1. mesh generation
2. mesh decomposition
3. local mesh coloring

1. ghost elements

1. ghost elements
2. local finite element numbering

1. ghost elements
2. local finite element numbering
3. global numbering

- $-\operatorname{div} \sigma = f$ in Ω , + suitable BCs
- $\varepsilon = \frac{1}{2}(\nabla u + \nabla u^T)$
- $\sigma = C : \varepsilon$

$$\int_{\Omega} \lambda \operatorname{div} u \cdot \operatorname{div} v + \int_{\Omega} 2\mu \varepsilon(u) \cdot \varepsilon(v) - \int_{\Omega} f v = 0$$

1. global matrix assembly

STEPS RELATED TO THE SOLUTION PHASE

1. global matrix assembly
2. preconditioner definition

STEPS RELATED TO THE SOLUTION PHASE

1. global matrix assembly
2. preconditioner definition
3. solution of the linear system

STEPS RELATED TO THE SOLUTION PHASE

1. global matrix assembly
2. preconditioner definition
3. solution of the linear system
4. saving the solution to disk

- parallelism is necessary for large-scale computations
- can be tricky on legacy codes, but there are tools
- solutions are tailored for each application

- parallelism is necessary for large-scale computations
- can be tricky on legacy codes, but there are tools
- solutions are tailored for each application

Thank you!