


C2S@Exa (1st year meeting)  
December, Sophia Antipolis



# Parallel Hierarchical Hybrid Algebraic Linear Solvers : current and future

HiePACS Inria Project  
Joint Inria-CERFACS lab  
INRIA Bordeaux Sud-Ouest

# 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


# 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



# Two Hybrid Linear Solvers

## Algebraic non-overlapping domain decomposition

- ▶ Partitioning of the adjacency graph of the sparse matrix
- ▶ Perform a partial Gaussian elimination (sparse direct solution on the internal variables)
- ▶ Solve the Schur complement system using a preconditioned Krylov subspace method
- ▶ Backsolve for the internal variables

## Two parallel implementations

- ▶ HIPS (Hierarchical Iterative Parallel Solver)  
Incomplete LU factorization of the Schur complement based on a hierarchical interface decomposition ordering
- ▶ MaPHyS (Massively Parallel Hybrid Solver)  
Algebraic additive Schwarz preconditioner for the Schur complement

# MaPHYs: Algebraic Additive Schwarz preconditioner [

L.Carvalho, L.G., G.Meurant- NLAA - 01]

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

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

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

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


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]

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


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

# MaPHyS: main future features

## Convergence history



## Time history




## Ongoing/future software development


- ▶ Improved software flexibility (rely on MUMPS, PastiX, Scotch, ...)
- ▶ Hybrid MPI-Thread implementation (PhD thesis - DIP Inria/Total) on top of PaStiX
- ▶ Implementation on top of runtime systems - cf Emmanuel's talk

# MaPHyS: main MPI-Threads features

## Time consumption

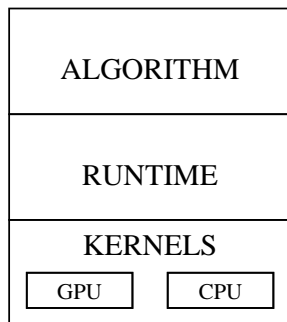


## Memory consumption



Audi test case: 32 cores - 4 to 32 domains/MPI processes

# Multiple layer approach



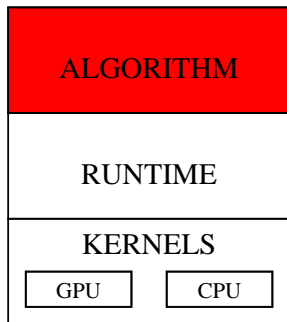
**Governing ideas:** Enable advanced numerical algorithms to be executed on a scalable unified runtime system for exploiting the full potential of future exascale machines.

**Basics:**

- ▶ Graph of tasks
- ▶ Out-of-order scheduling
- ▶ Fine granularity



# Algorithms



**Governing ideas:** Design high-level algorithms

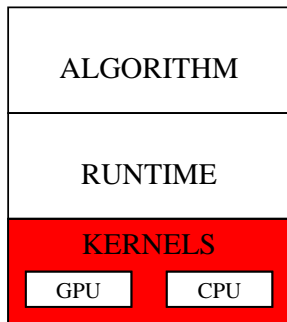
**Main challenges:**

- ▶ Increase concurrency
- ▶ Control granularity of tasks
- ▶ Trade off numerical accuracy and stability with performance

**Fundings and collaborations:**

- ▶ National: Total, ANR-SFGPU
- ▶ International: AT-FastLA (Berkeley, Stanford), AT-MORSE (UTK, KAUST, UC Denver)

# Kernels



**Governing ideas:** Use optimized low-level kernels

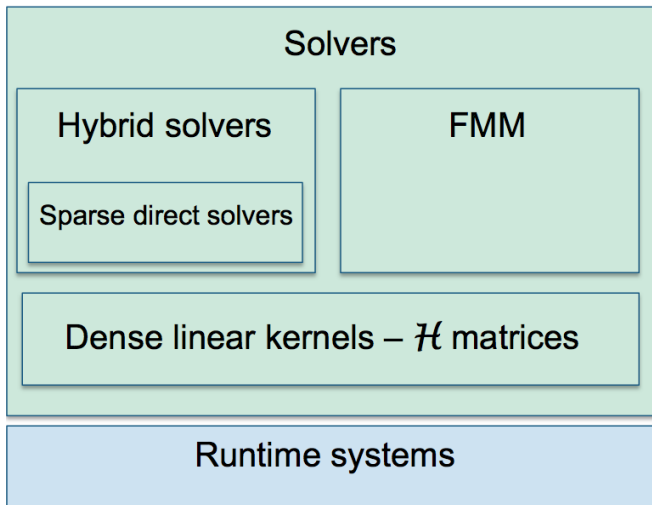
**Main challenges:**

- ▶ Possibly use existing kernels
- ▶ Otherwise design new kernels for complex hardware
- ▶ Automatic generation


**Projects and collaborations:**

- ▶ INRIA Bordeaux: MANAO
- ▶ International: AT-FastLA (Berkeley, Stanford), AT-MORSE (UTK, KAUST, UC Denver)


# Targeted solver stack



# A first example in dense linear algebra




# A first example in dense linear algebra



+ 200 GFlop/s but 12 cores = 150 GFlop/s

# A second example in sparse linear algebra



$LL^T$  factorization on Audi test case

# What's next

## Ongoing/future numerical development

- ▶ Improved numerical robustness - hierarchical toward global preconditioning (ADT Maphys@Exa via C2S@Exa)
- ▶ Coarse space mechanisms through augmentation and/or deflation (FP7 Exa2CT via C2S@Exa)
- ▶ New Krylov subspace methods (block variants, hidden/avoiding communications, ...) - shared with sparse direct
- ▶ H-matrix arithmetic (Stanford/Berkeley collaboration via AT FASTLA)