Multiscale biofluidics with MUPHY: development strategies for heterogeneous architectures

M. Bernaschi, M. Bisson, S. Melchionna

Consiglio Nazionale delle Ricerche, Italy

Biofluids

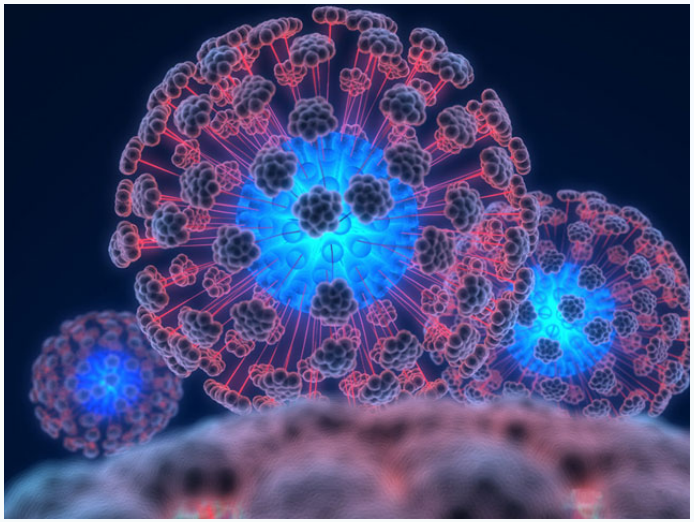

- Physiological carriers (swimmers,nutrients,wastes,etc)

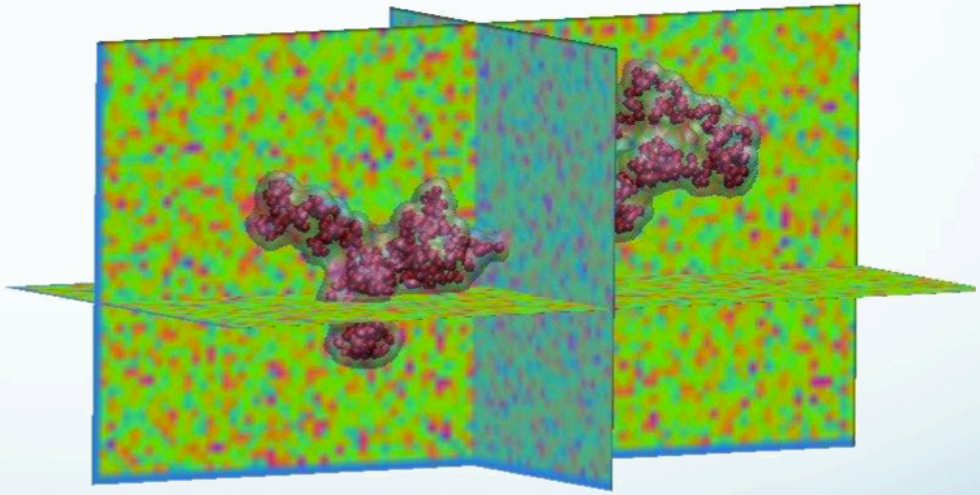

- Cellular environments

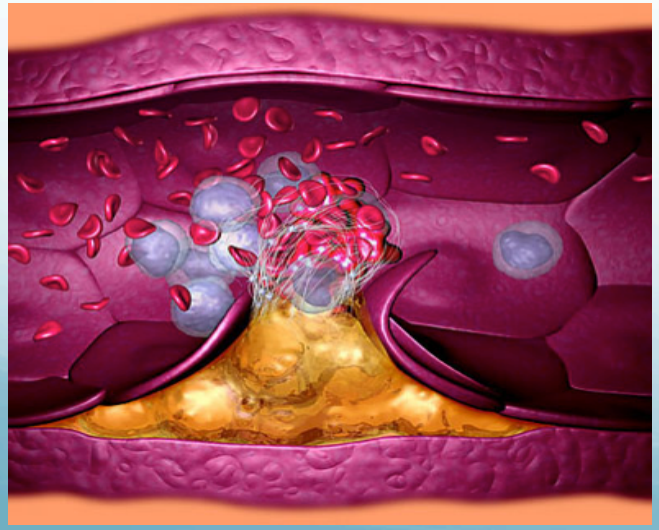

- Biofluid-mechanics

Dissecting Life Sciences

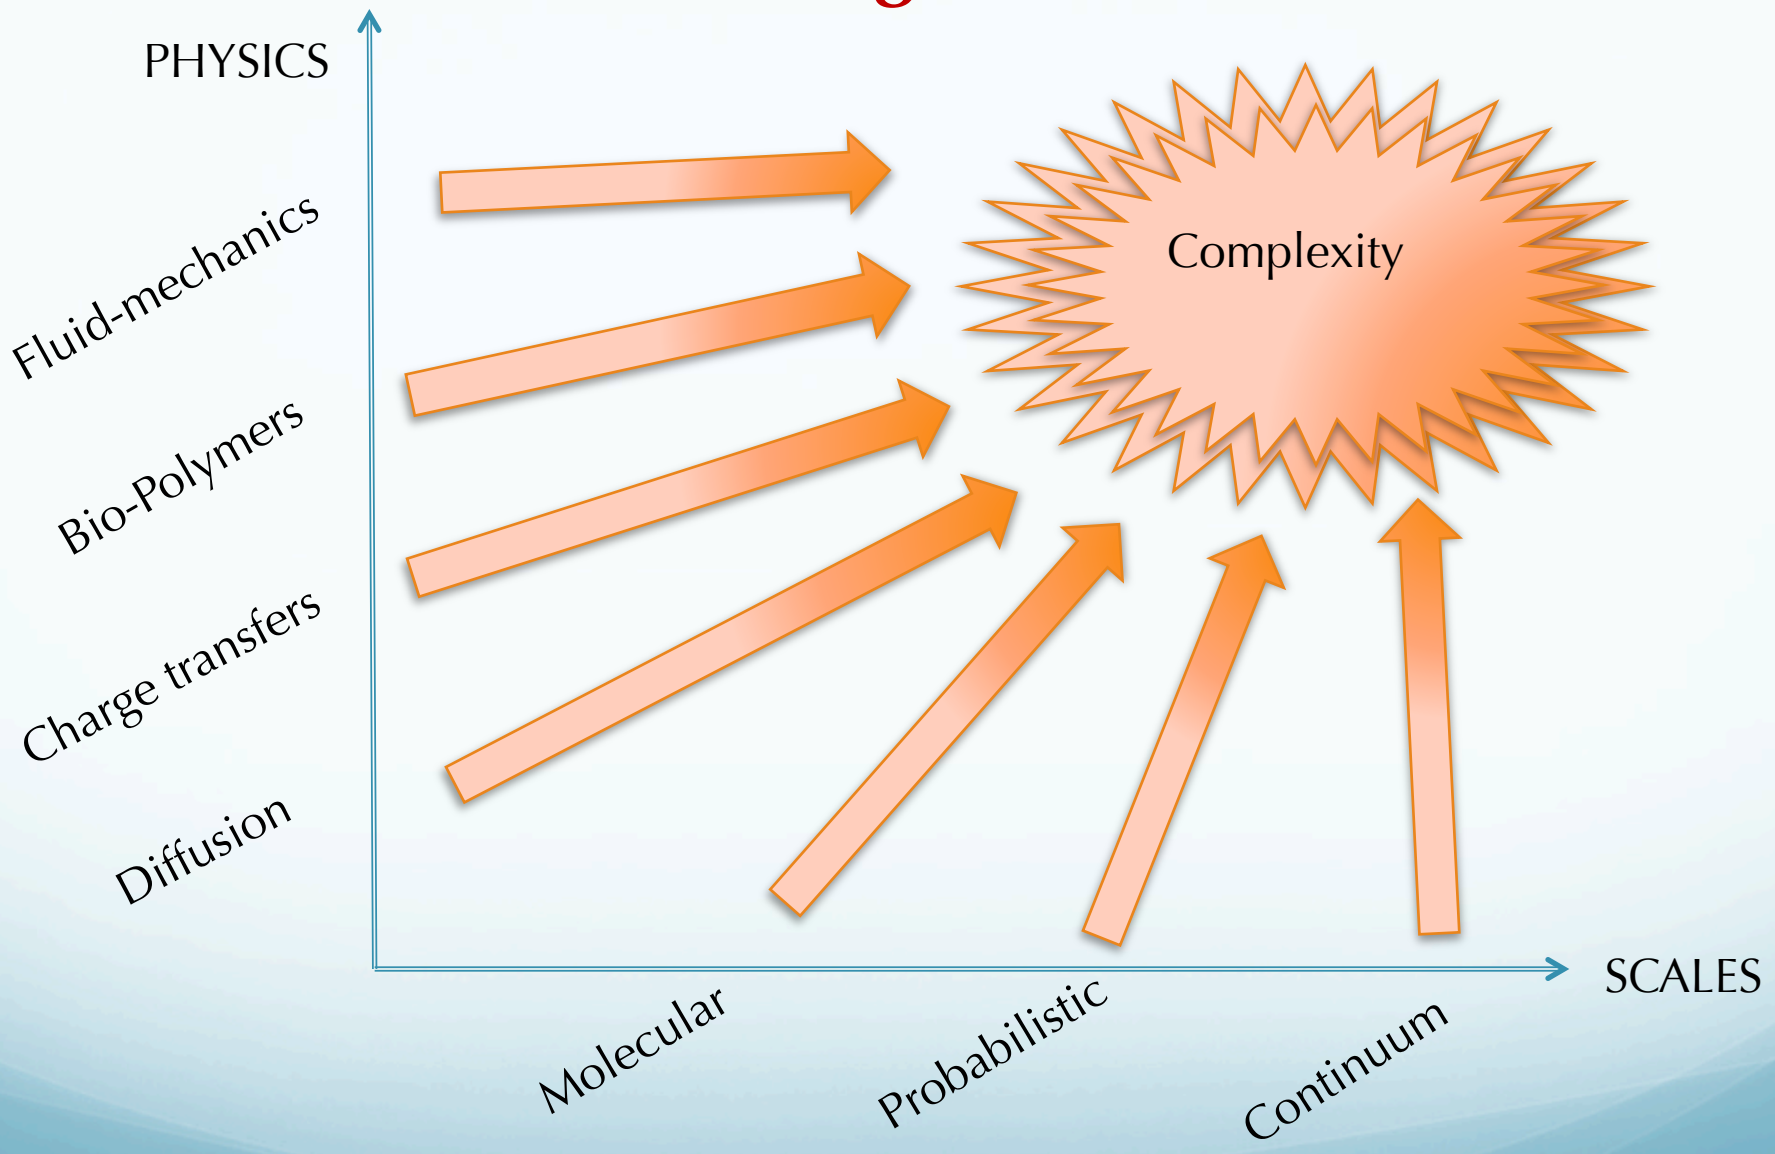

Fact: Complexity requires interdisciplinary modeling

Computer Science

Key to turn simulation techniques into efficient algorithms
on leading-edge computational platforms

Result: unprecedented transformative advances

CARDIOVASCULAR HEMODYNAMICS

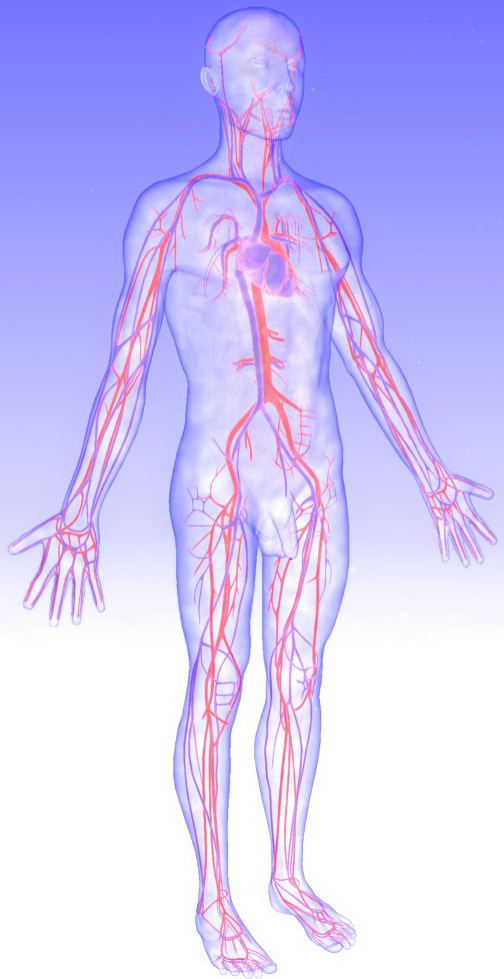

Plaque rupture leads to flow interruption, heart attack and stroke

It is essential to forecast where and when plaques form

Why Multi-Scale ?

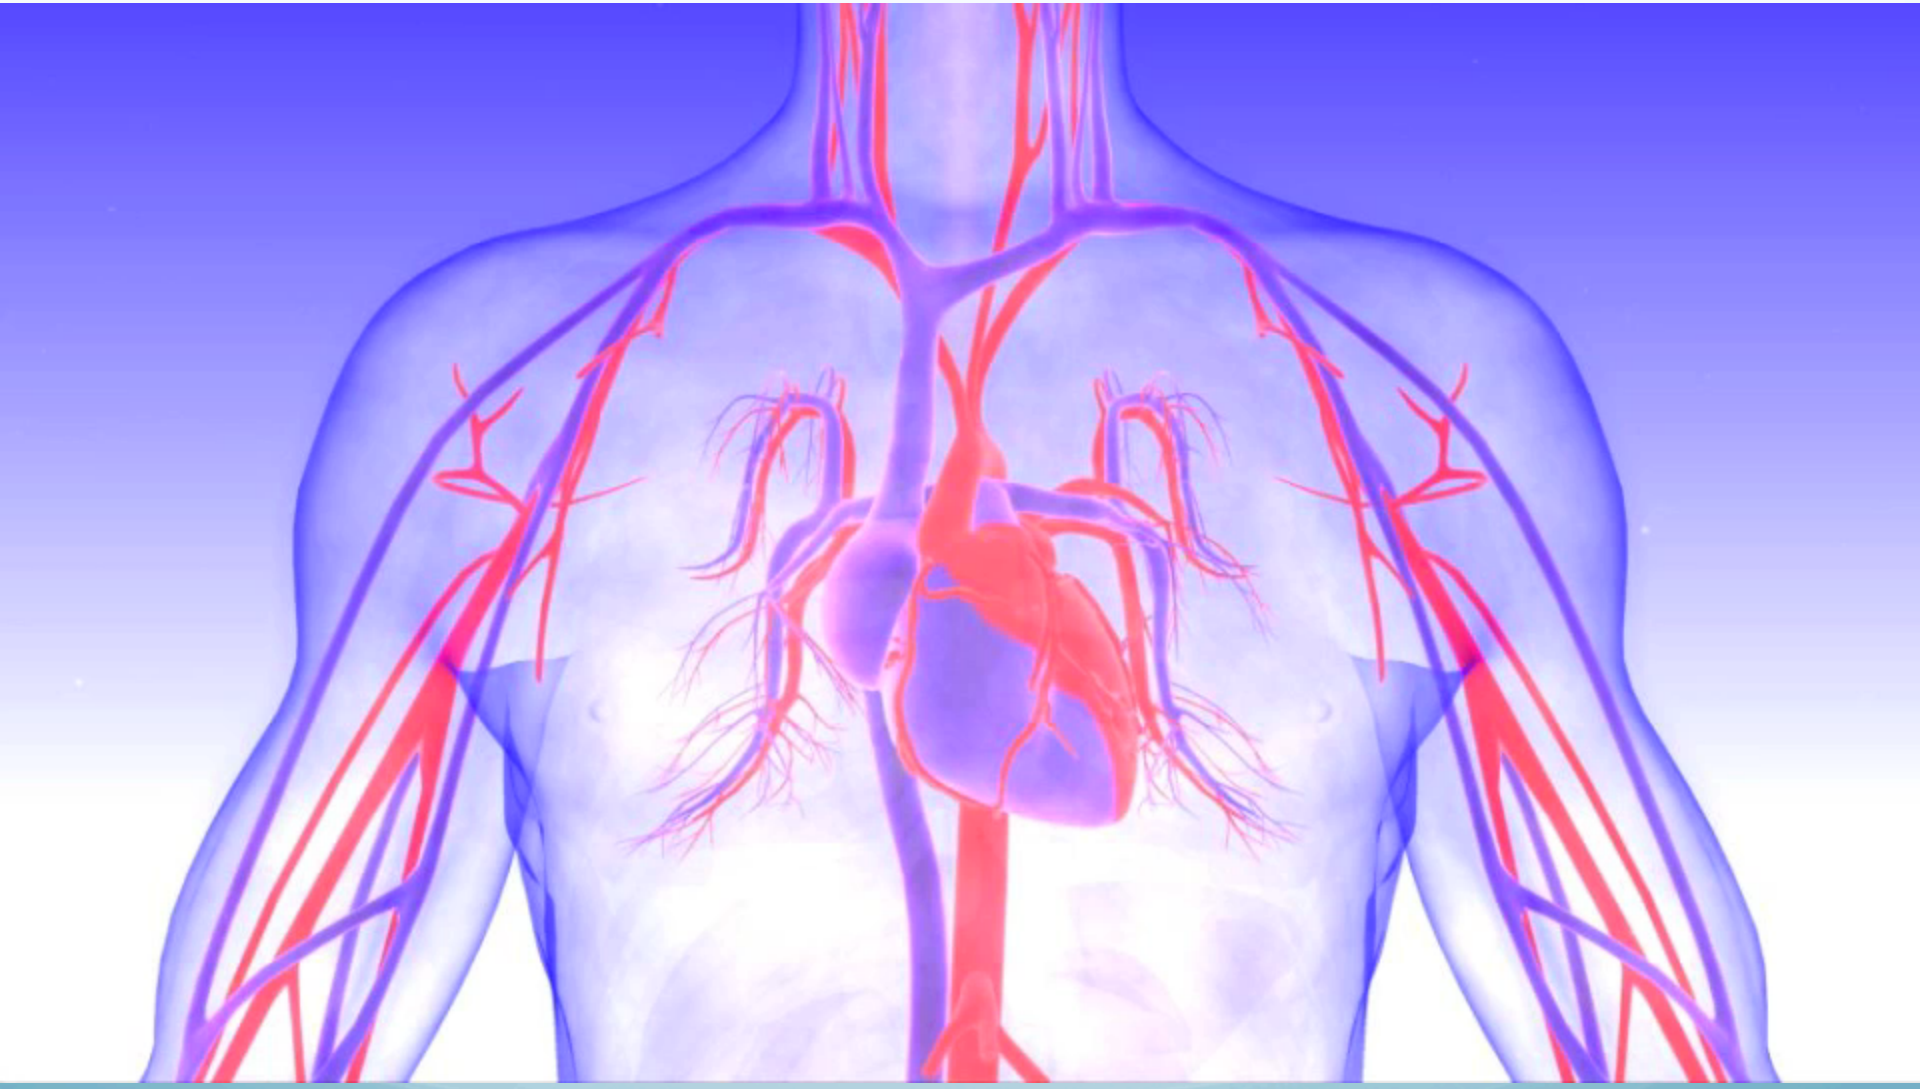

Getting Started

SW engineering design model (not layered stacks!)

Co-Modeling: enable concurrency of physics

Abstraction from morphology/models

Heterogeneous HPComputing

Wide Range of scales

Complex Geometries

Computational Flexibility

Time to solution is critical

The Concurrent Multi-Scale method

Lattice Boltzmann + Molecular Dynamics

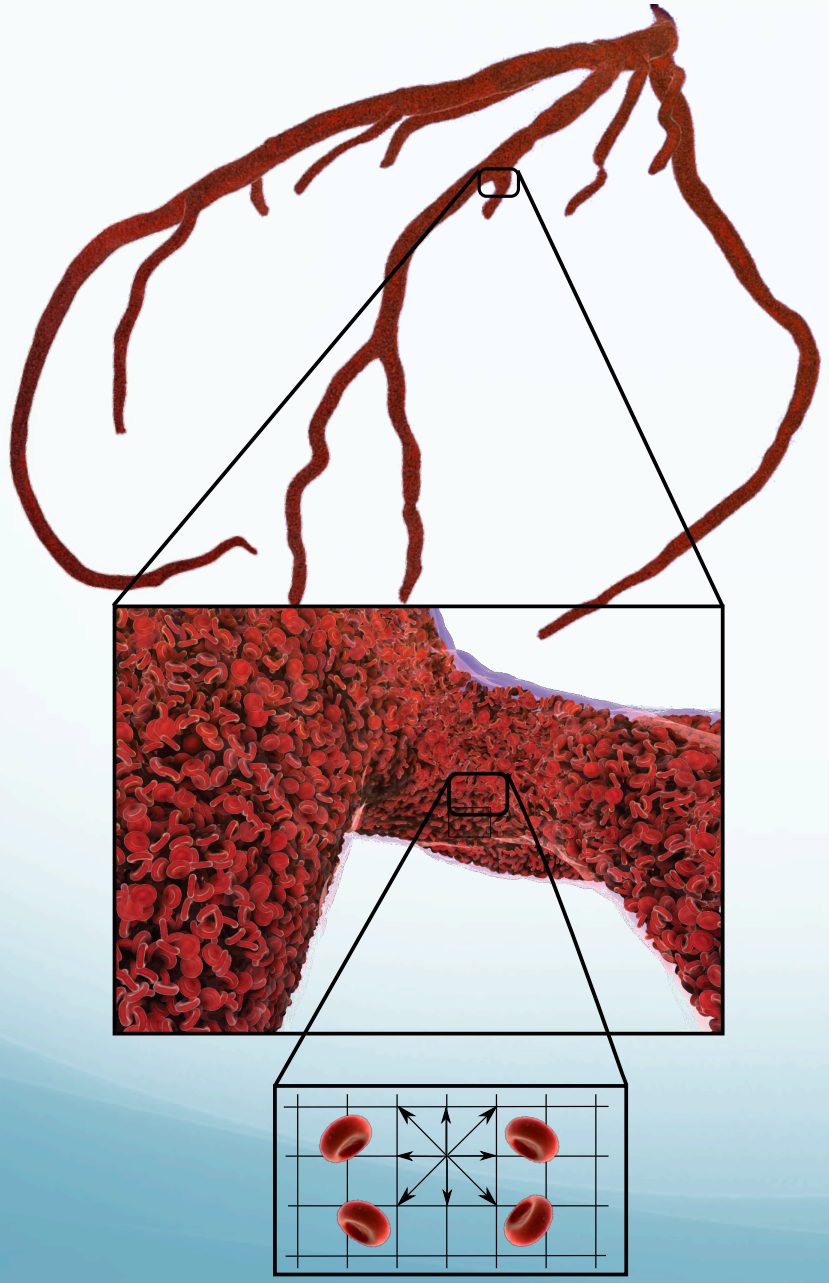

LB Populations : Collide & Stream

+

MD Particles : Interact & Move

+

Coupling : Drag / Vorticity exchange

=

Multi-scale Computing

M. Fyta et al, Comput. Sci. & Eng. 2008
M. Bernaschi et al, Comput. Phys. Comm. 2009
M. Bernaschi et al, Concurr. & Comput 2009
S. M. et al, Comput. Phys. Comm. 2010
S. M., Macromol. Theor. & Sim. 2011
M. Bisson et al, Commun. Comput. Phys. 2011

LB $$(\partial_t + \vec{v}\partial_x) f(\vec{x},\vec{v},t) = -\omega(f - f^{eq})(\vec{x},\vec{v},t)$$

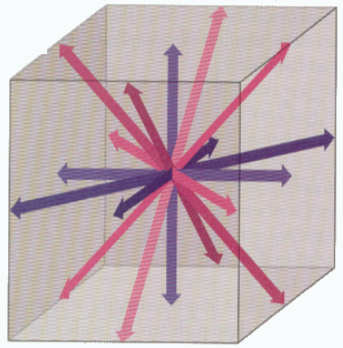

MD $$\frac{d^2}{dt^2}\begin{pmatrix} MR_i \\ \mathrm{I}\Omega_i \end{pmatrix} = \begin{pmatrix} F_i \\ T_i \end{pmatrix}$$

LB-MD

$$F_i^H = \int dx \tilde{\delta}(x - R_i)\varphi(x)$$
$$T_i^H = \int dx \tilde{\delta}(x - R_i)\psi(x)$$

$$\partial_t f = \sum_i \left[\varphi(x) + (x - R_i) \wedge \psi(x)\right]\tilde{\delta}(x - R_i)$$

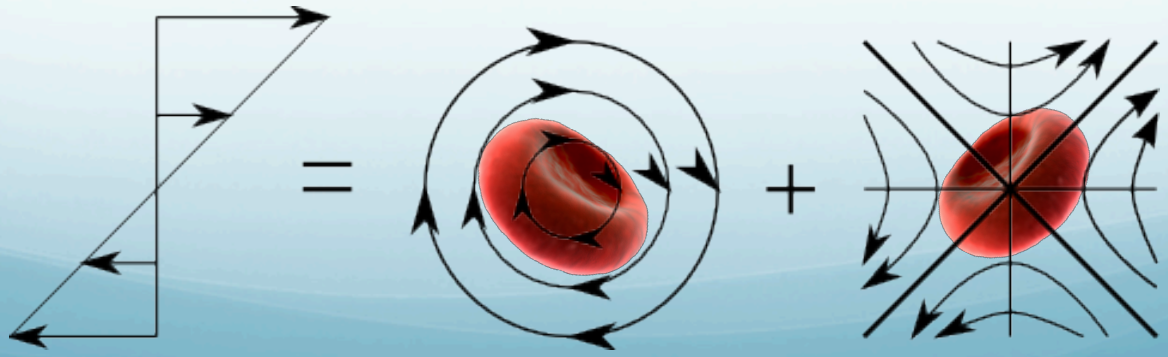

MUPHY : Multi-Physics Simulator

- Library of particle and molecular representations & fluid types
- Broad spectrum of ***fluid-particle*** coupling mechanisms
- Multi-Platform
- Independence from geometrical complexity and parity
- Checkpointing & parallel I/O

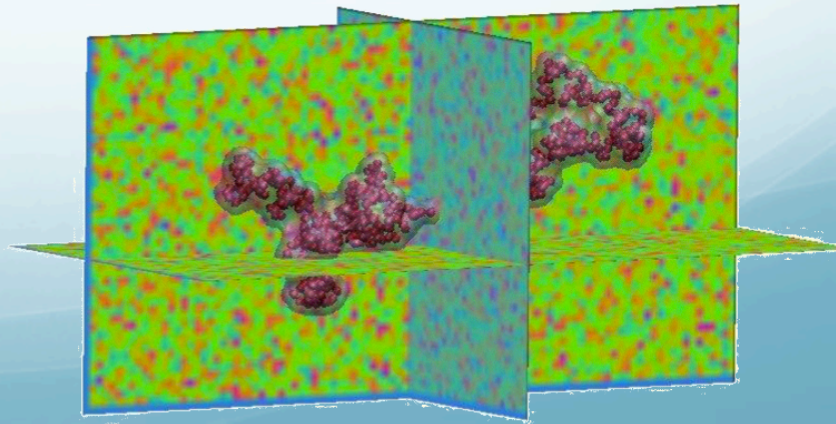

M. Bernaschi et al, Comput. Phys. Comm. 2009
S. M. et al, Comput. Phys. Comm. 2010

MUPHY Computing Architecture

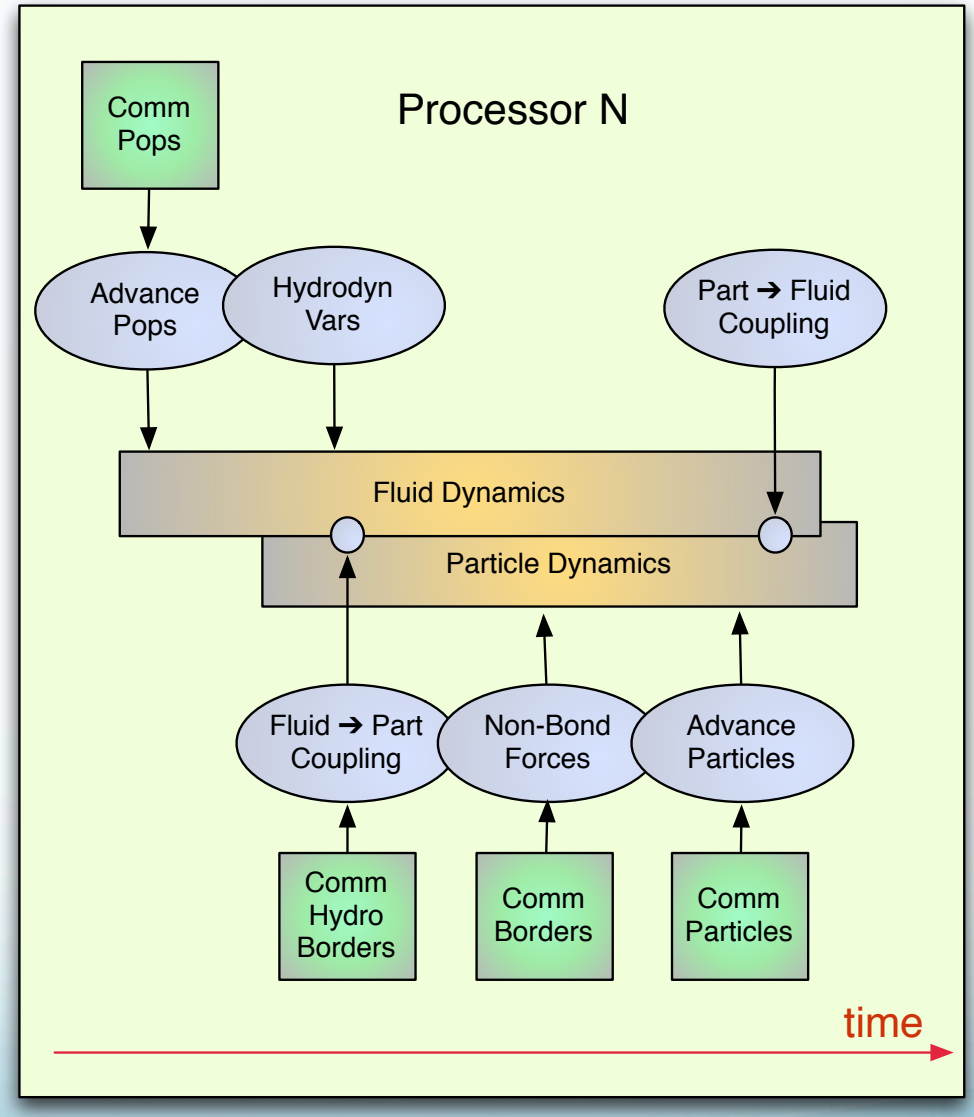

At each step fluid and particles exchange information

Major Algorithmic Challenges

- Domain Partitioning of Huge Irregular Geometries
- Parallel Molecular Dynamics in Irregular Domain
- Efficient, Scalable and Portable Communication Pattern

Domain Partitioning

- ✓ Same domain decomposition for LB and MD (preserve data locality for fluid-particle interaction)
- ✓ Irregular geometry prevents simple cartesian decomposition
- ✓ Standard Graph-based partitioning tools (Metis, Scotch) generate suboptimal solutions:
  - good load balancing (even distribution of points)
  - poor scalability (too many domain neighbors)

Our solution to domain partitioning

Two-step strategy:

1. Graph-based partitioning up to 256 sub-domains
2. Further flooding-based partitioning of each sub-domain

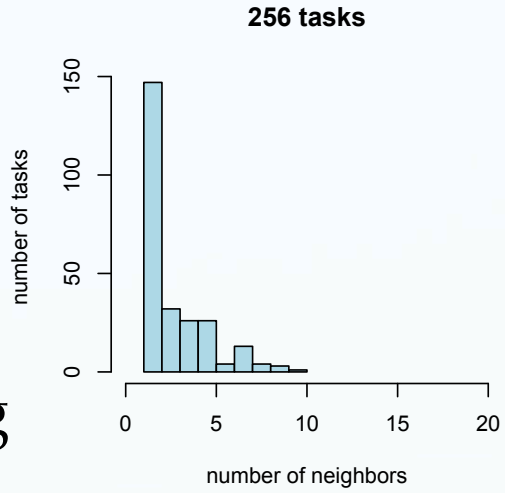

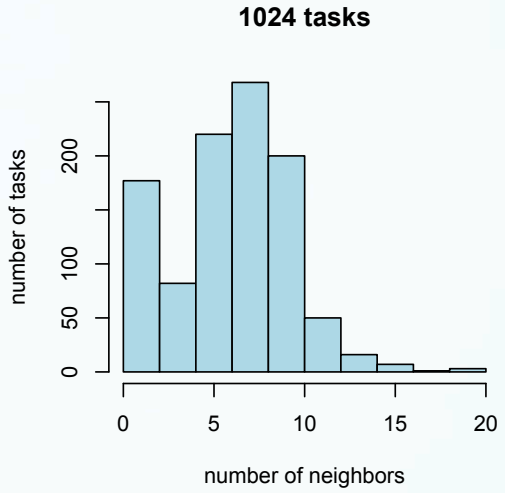

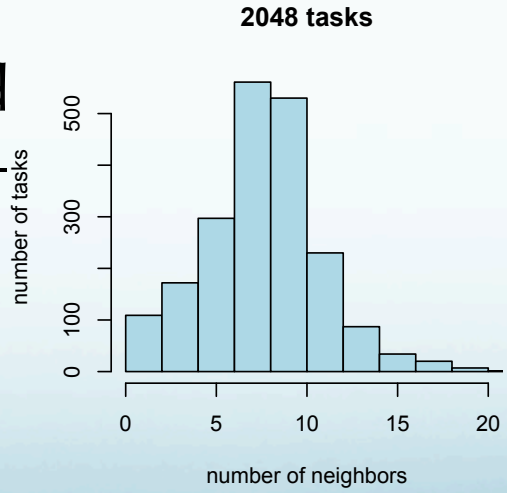

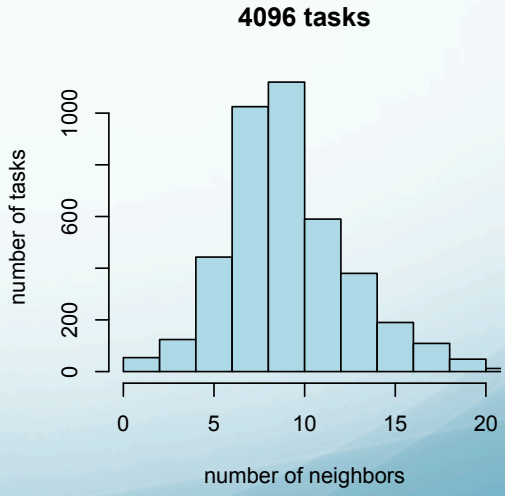

Guaranteed good load balancing and a limited number of domain neighbors at the same time.

Particle Dynamics on Irregular Domains

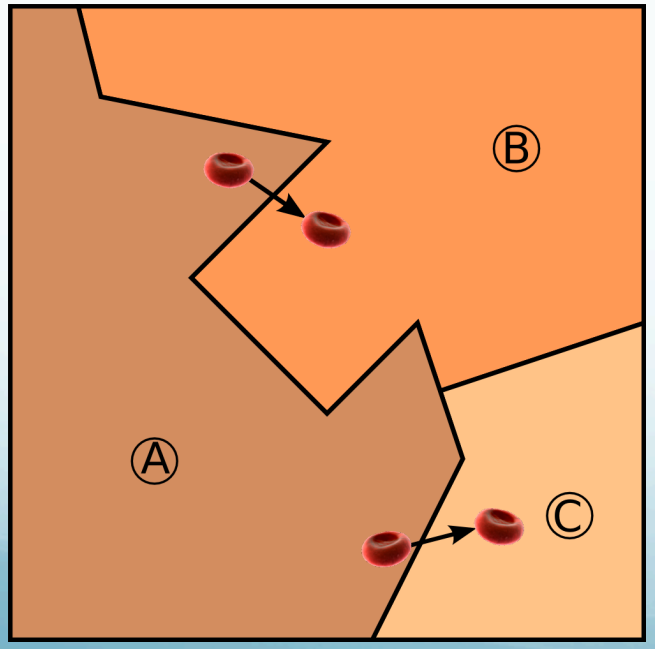

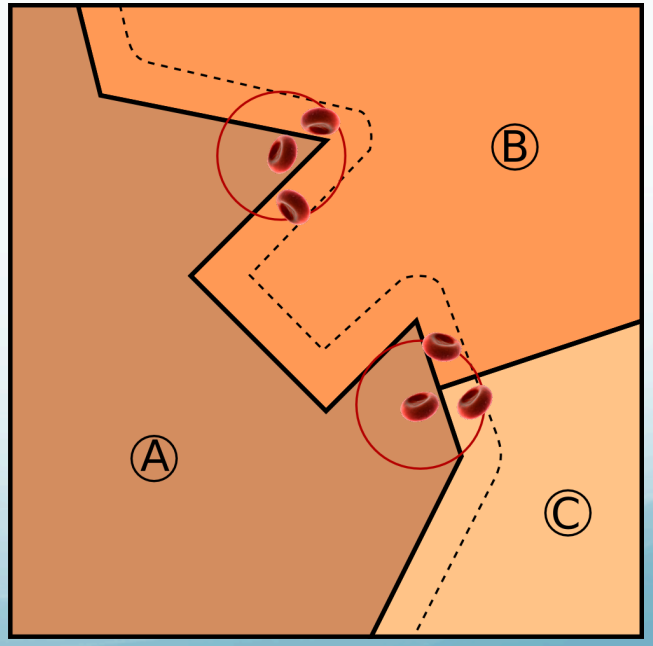

Particle migration and force computation on irregular domains

Digital Lagrangian Tracking

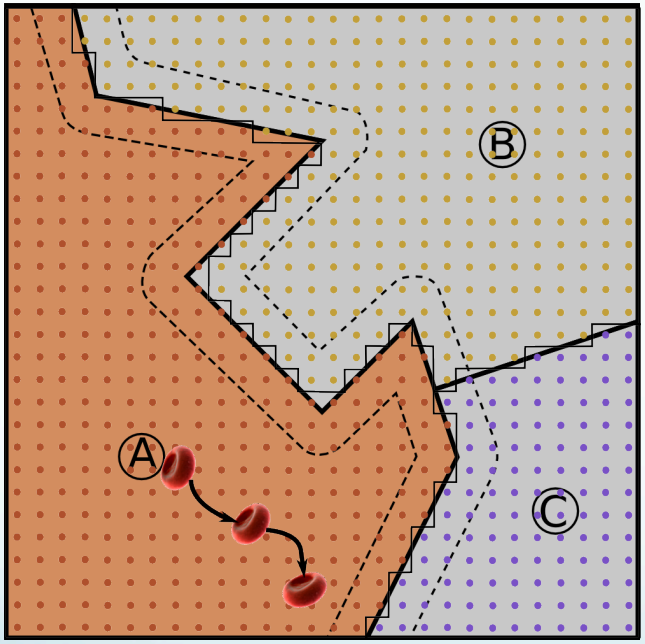

Tiling

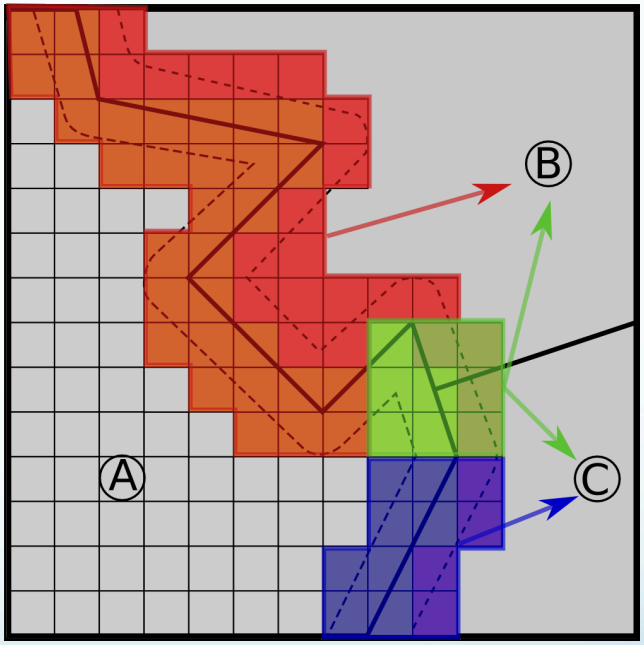

$O(N_{RBC})$ effort

M. Bisson et al, Commun. Comput. Phys. 2011
ibidem, Commun. Comput. Phys. 2011

Target Clinical Task

Mesh resolution = 10 μm → 2×10^{10} lattice degrees of freedom

Time resolution = 1 μs → $10^4 \div 10^6$ timesteps (1 heart beat)

RBC dynamics at 50% hematocrit → 4.5×10^8 RBC's

Bounding box: 6000 x 8000 x 6000 ~ 10^{11} nodes

1% active nodes

Preprocessing (GMUPHY)

- DICOM sculpting:
  - window width (contrast)
  - level (tissue selection)
- Image manipulation :
  - image crop
  - free-hand lasso tool
- Isosurface w Marching Cubes
- Remove isosurface artifacts
- Shrinkage-free smoothing (λ/μ algos)

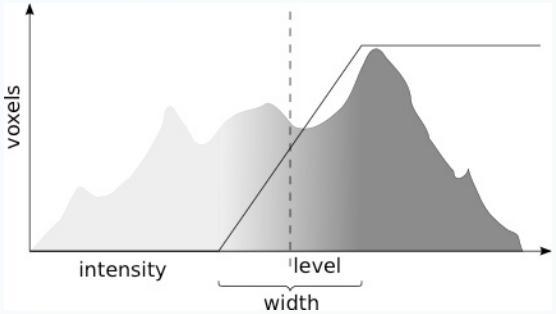

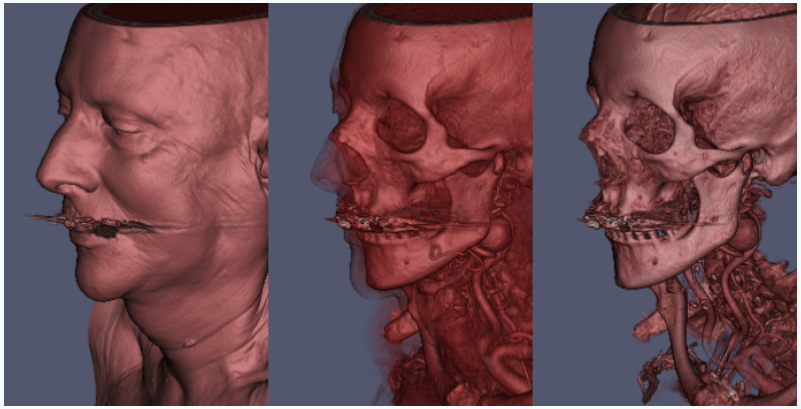

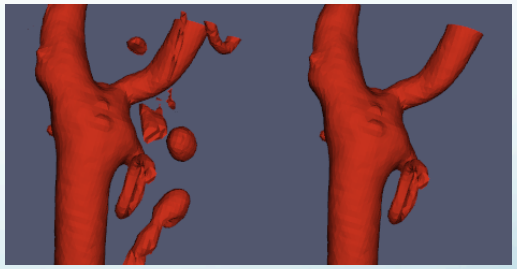

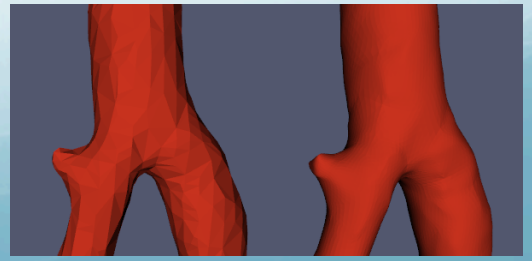

GMUPHY(2)

At given resolution, cartesian mesh generated with

- O(V) space and O(V log V) time

Validation by visual inspection:

Remove imperfections
Validate mesh

Multiscale Starts. Visualise:

Stresses
Fluid patterns
Streamlines
...

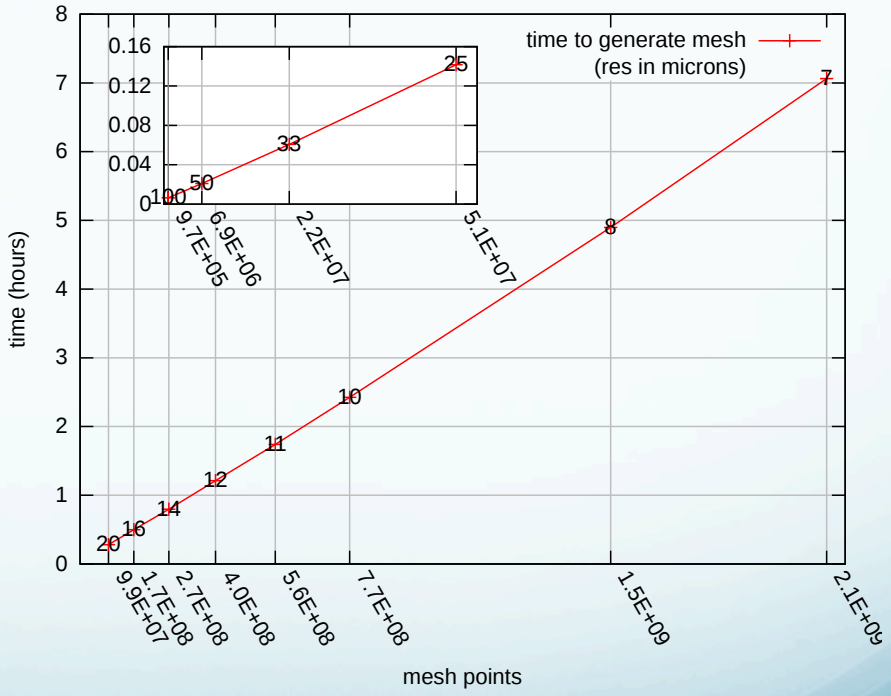

Comparing Architectures

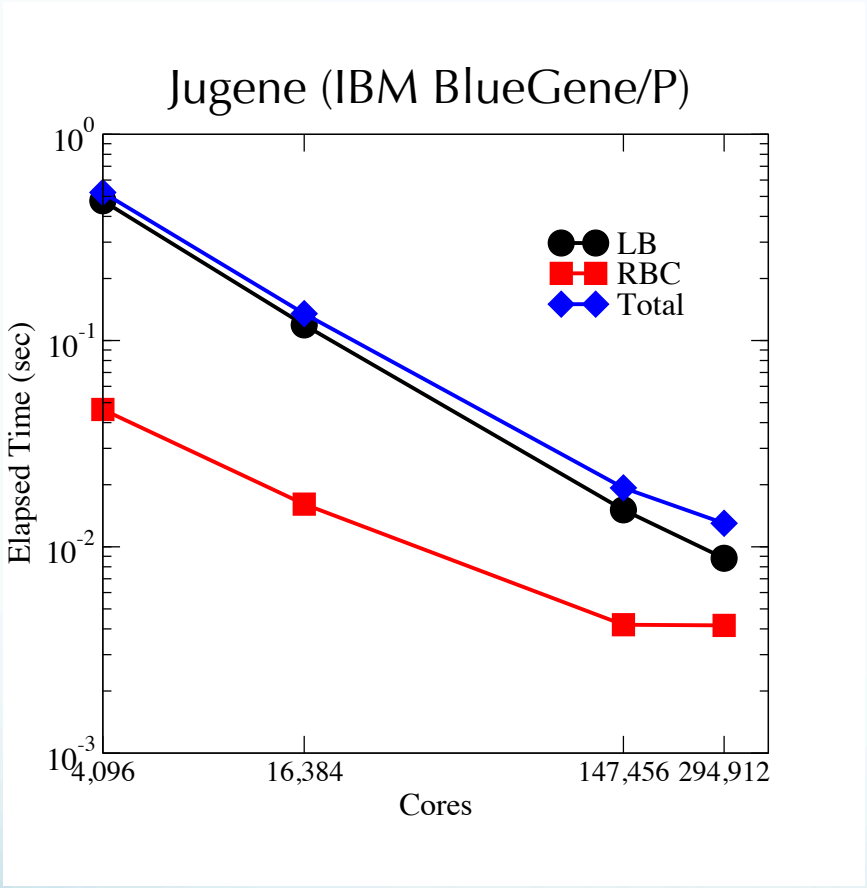

Gordon Bell Final 2010

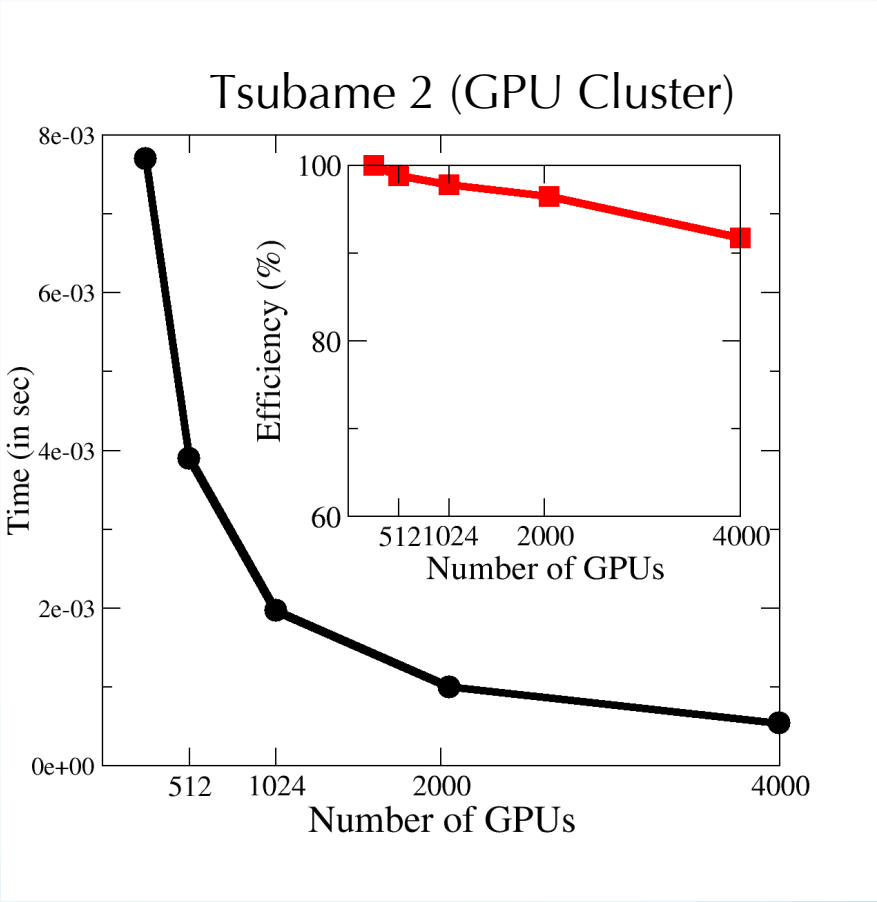

Gordon Bell Final 2011
Honorable Mention

Time (4000 GPUs) / Time (BG/P 300,000 cores) = 1 / 7

Communication scheme

CUDA supports independent *streams* of execution
So it is possible to *overlap* communication of boundaries and computation (in the *bulk*).

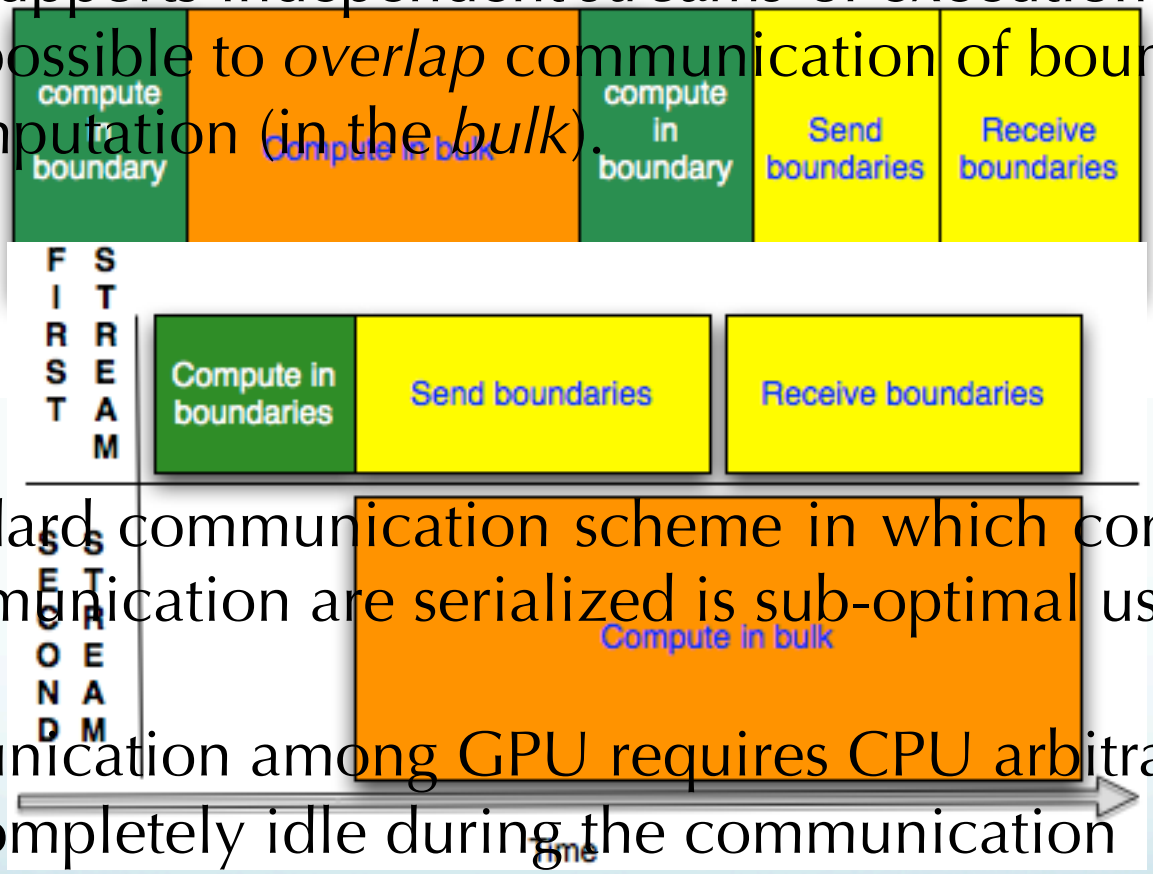

The standard communication scheme in which computation and communication are serialized is sub-optimal using GPU

✓Communication among GPU requires CPU arbitration
✓GPU completely idle during the communication

CPU acts as a communication co-processor for the GPU!

Target Clinical Task (2)

Sustained **600** **Teraflops**

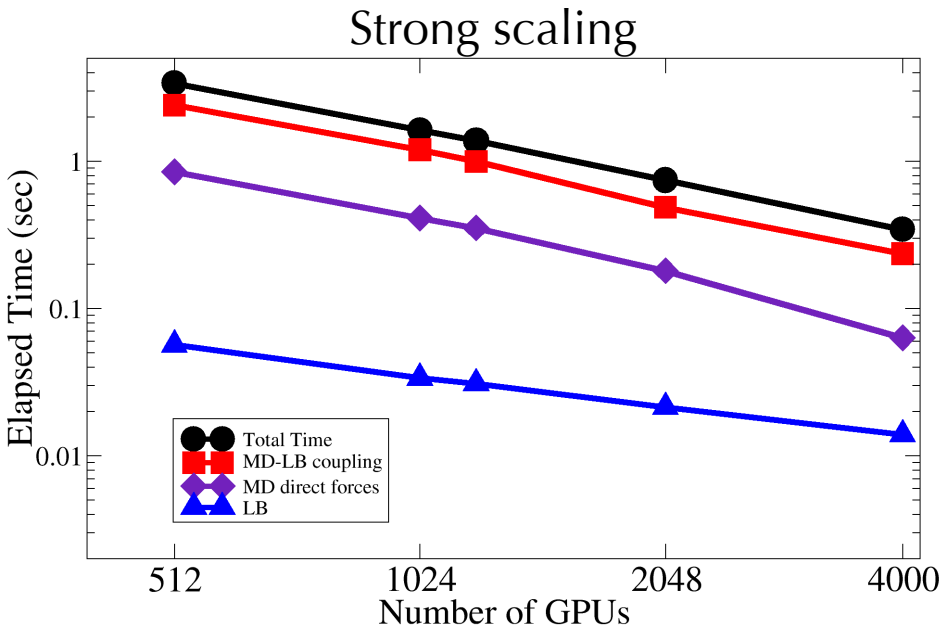

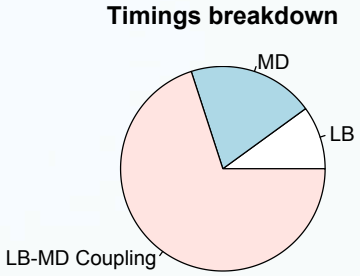

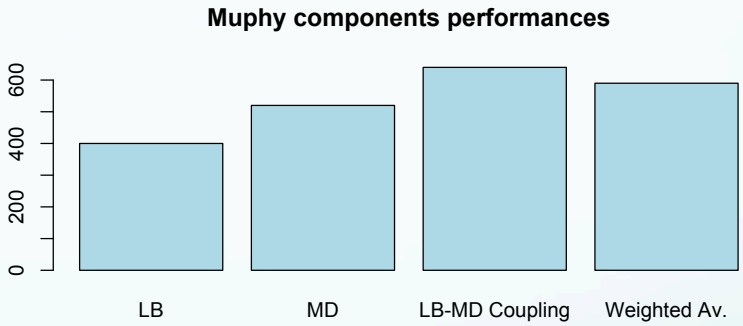

GPUs	Total Time (sec)	MD-LB coupling	MD	LB
256	3.3902	69.73 %	25.48 %	1.68 %
1024	1.6268	72.09 %	25.00 %	2.08 %
1200	1.3812	70.58 %	25.20 %	2.24 %
2048	0.7446	63.04 %	21.56 %	2.88 %
4000	0.3440	62.29 %	15.10 %	4.04 %

Parallel Efficiency & Weak Scaling

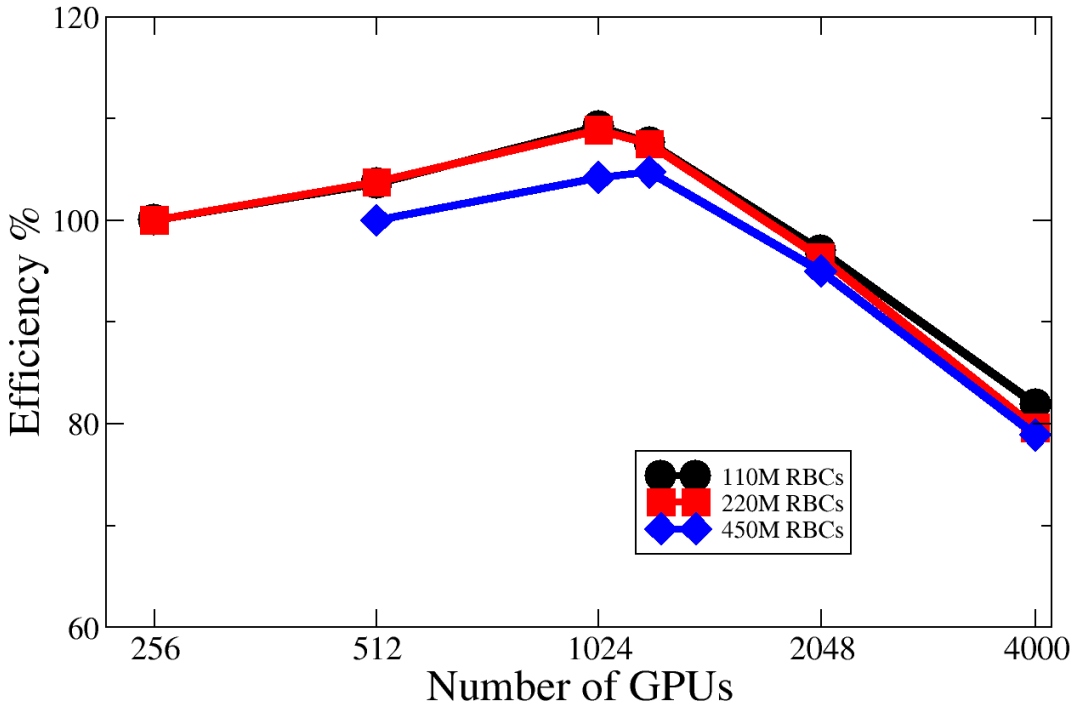

Routine	Calls	Avg Bytes	Avg Time	Bandwidth
Send	38012346	43600	76 μsec	570 MB/sec
Irecv	38467062	45000	67 μsec	670 MB/sec
Waitall	7635000	N/A	750 μsec	N/A
Allreduce	4000	16	675 μsec	N/A

MUPHY at work for Clinical Guidance

Plaques growth evolutions: Single and Multiple Stenoses (10 cases)

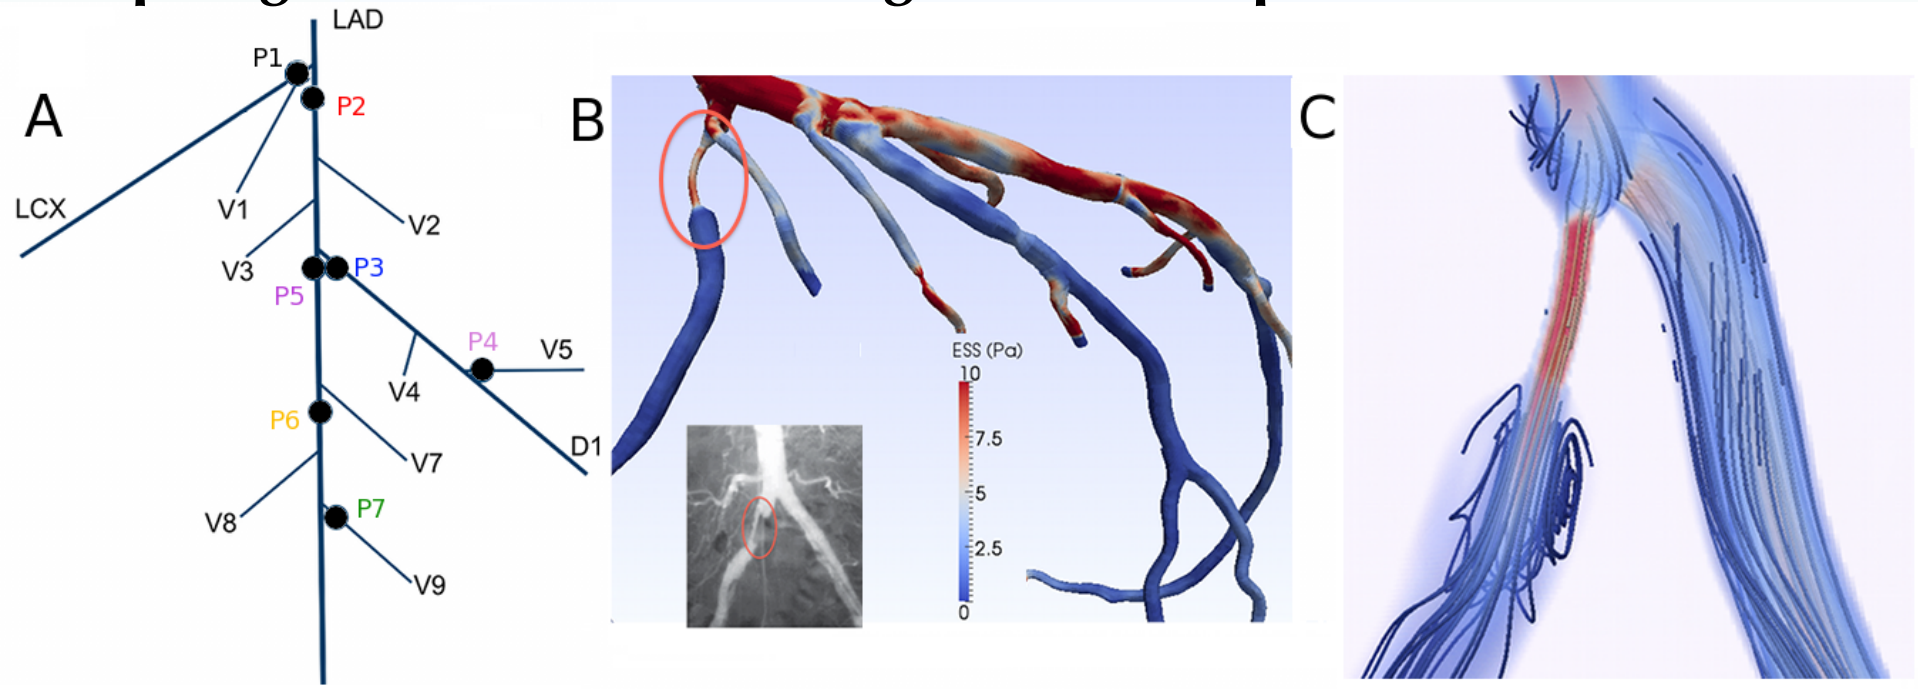

- 1200 GPUs at 45% hematocrit : 30 minutes wall time
- Compute pressure unbalance (O_2 deficit in myocardium)
- Assess patient-specific network robustness !

Conclusions & Outlook

- Scalable Multi-Scale algorithmics on irregular morphologies
- Multi-Platform sw for massively parallel and commodity platforms
- Computational personalized medicine in leading hospitals
- Exascale horizon:

Quantitative match by including the full vascular periphery

From red blood cell to full-heart bio-mechanics

Acknowledgments

E. Kaxiras, C. Feldman, F. Rybicky, P. Stone, J. Lätt,
J. Sircar, A. Peters, M. Borkin, D. Chiappino, F. Recchia,
P. Pontrelli, G. Amati, R. Verzicco. M. Fatica,
T. Endo, S. Matsuoka, S. Succi

Harvard Medical School

Fondazione Monasterio - Ospedale del Cuore

École Polytechnique Fédérale de Lausanne

Jülich Supercomputing Center

CASPUR Supercomputing Center

ScienceVisuals