CFD in Manufacturing and Medical Applications

Petros Koumoutsakos

Chair of Computational Science

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

www.cse-lab.ethz.ch

Modeling and Technology

- No aircraft is flown without having been designed with complex, mechanistic simulations

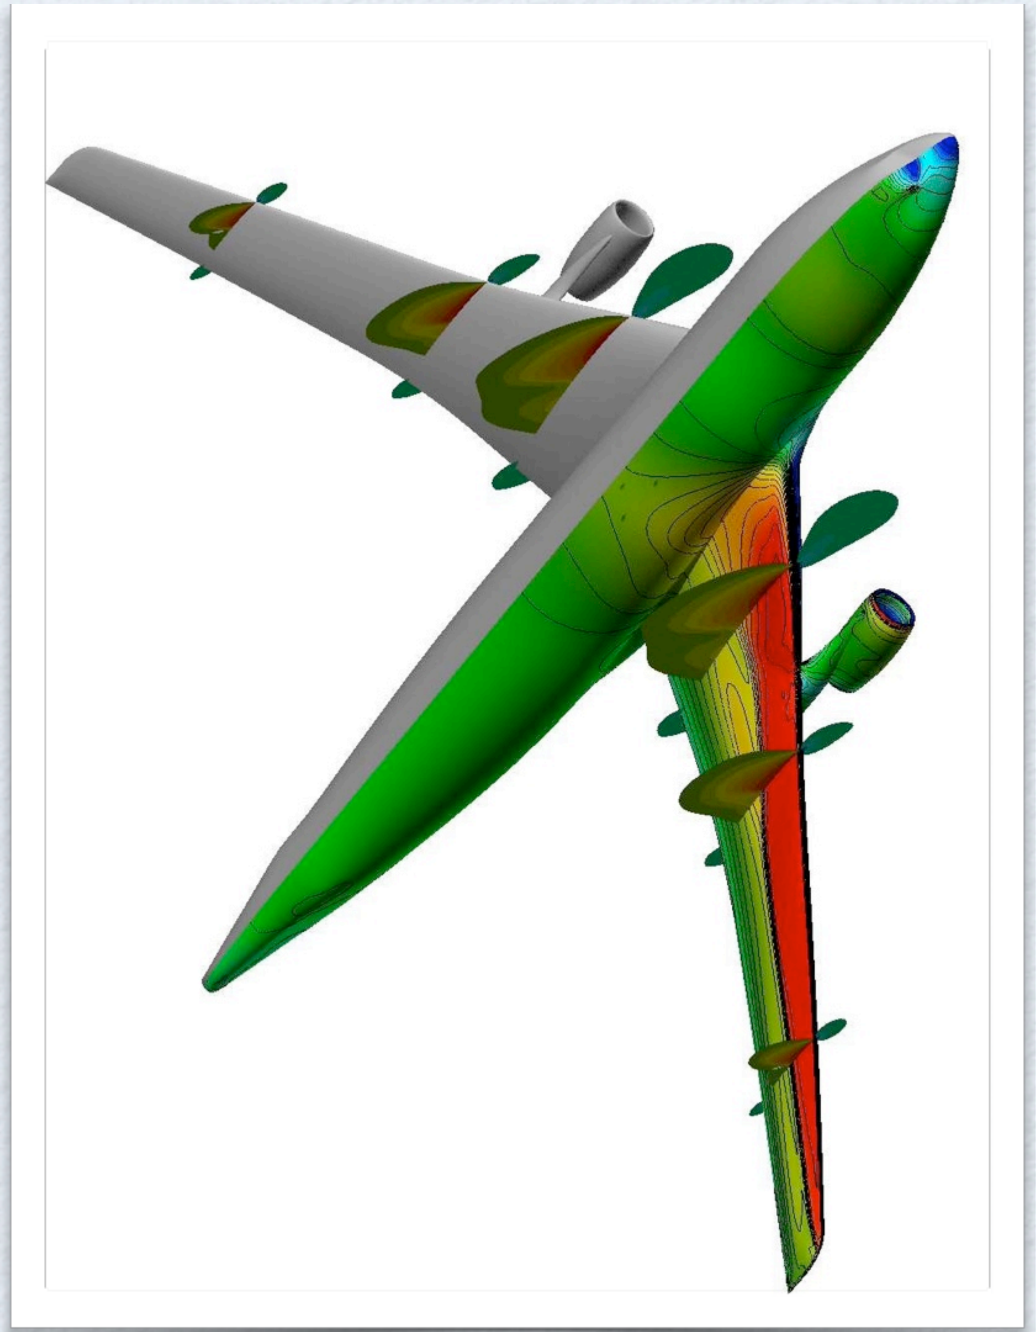

Modeling and Medicine

- Heuristics and Data
- Models ?

Dreamstime.com

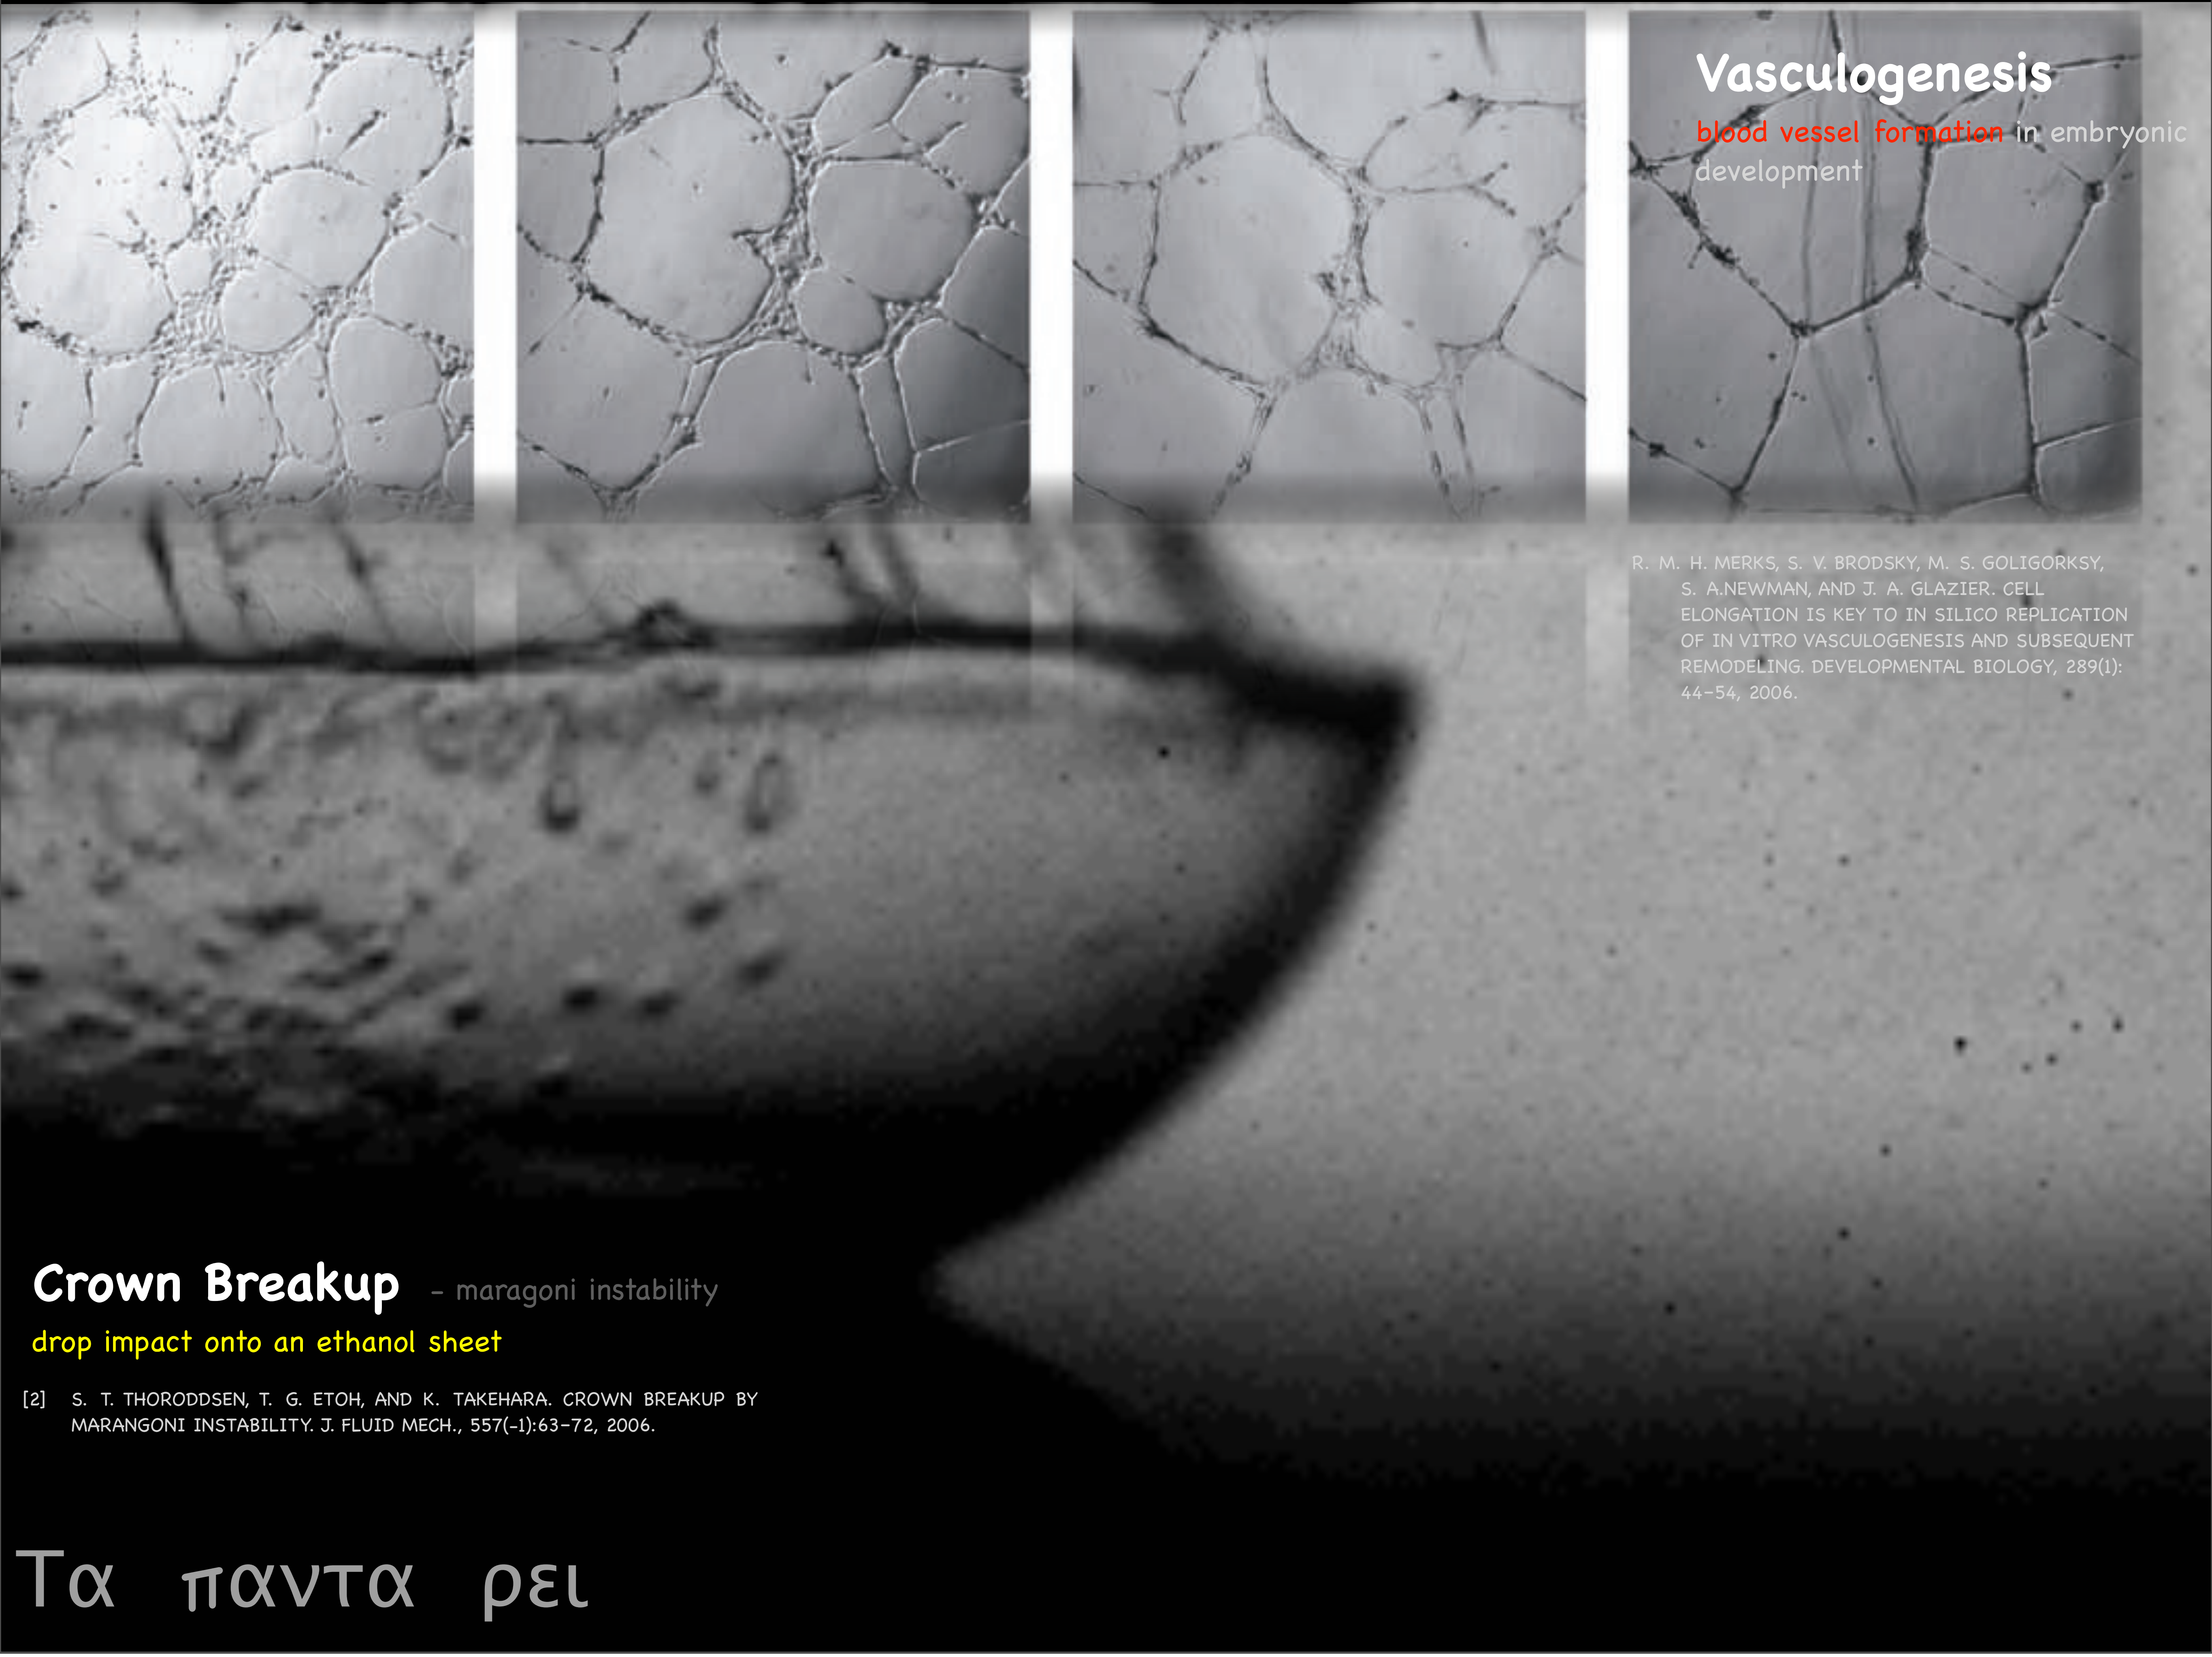
Vasculogenesis
blood vessel formation in embryonic development
R. M. H. MERKS, S. V. BRODSKY, M. S. GOLIGORKSY, S. A.NEWMAN, AND J. A. GLAZIER. CELL ELONGATION IS KEY TO IN SILICO REPLICATION OF IN VITRO VASCULOGENESIS AND SUBSEQUENT REMODELING. DEVELOPMENTAL BIOLOGY, 289(1): 44–54, 2006.
Crown Breakup – maragoni instability
drop impact onto an ethanol sheet
[2] S. T. THORODDSEN, T. G. ETOH, AND K. TAKEHARA. CROWN BREAKUP BY MARANGONI INSTABILITY. J. FLUID MECH., 557(-1):63–72, 2006.
Τα παντα ρει

16384 Cores - 10 Billion Particles - 60% efficiency

Runs at IBM Watson Center - BLue Gene/L

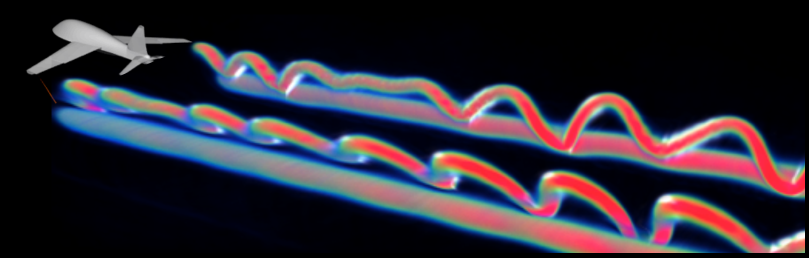

Chatelain P., Curioni A., Bergdorf M., Rossinelli D., Andreoni W., Koumoutsakos P., Billion Vortex Particle Direct Numerical Simulations of Aircraft Wakes, Computer Methods in Applied Mech. and Eng. 197/13-16, 1296-1304, 2008

Tumor Induced Angiogenesis

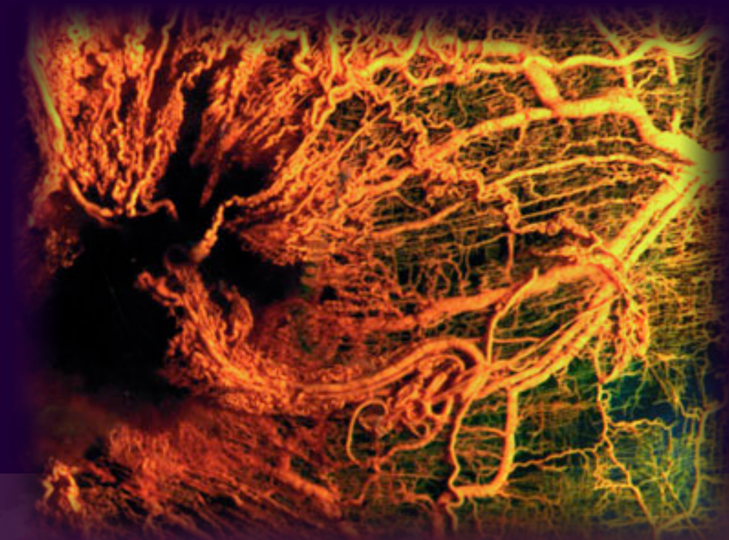

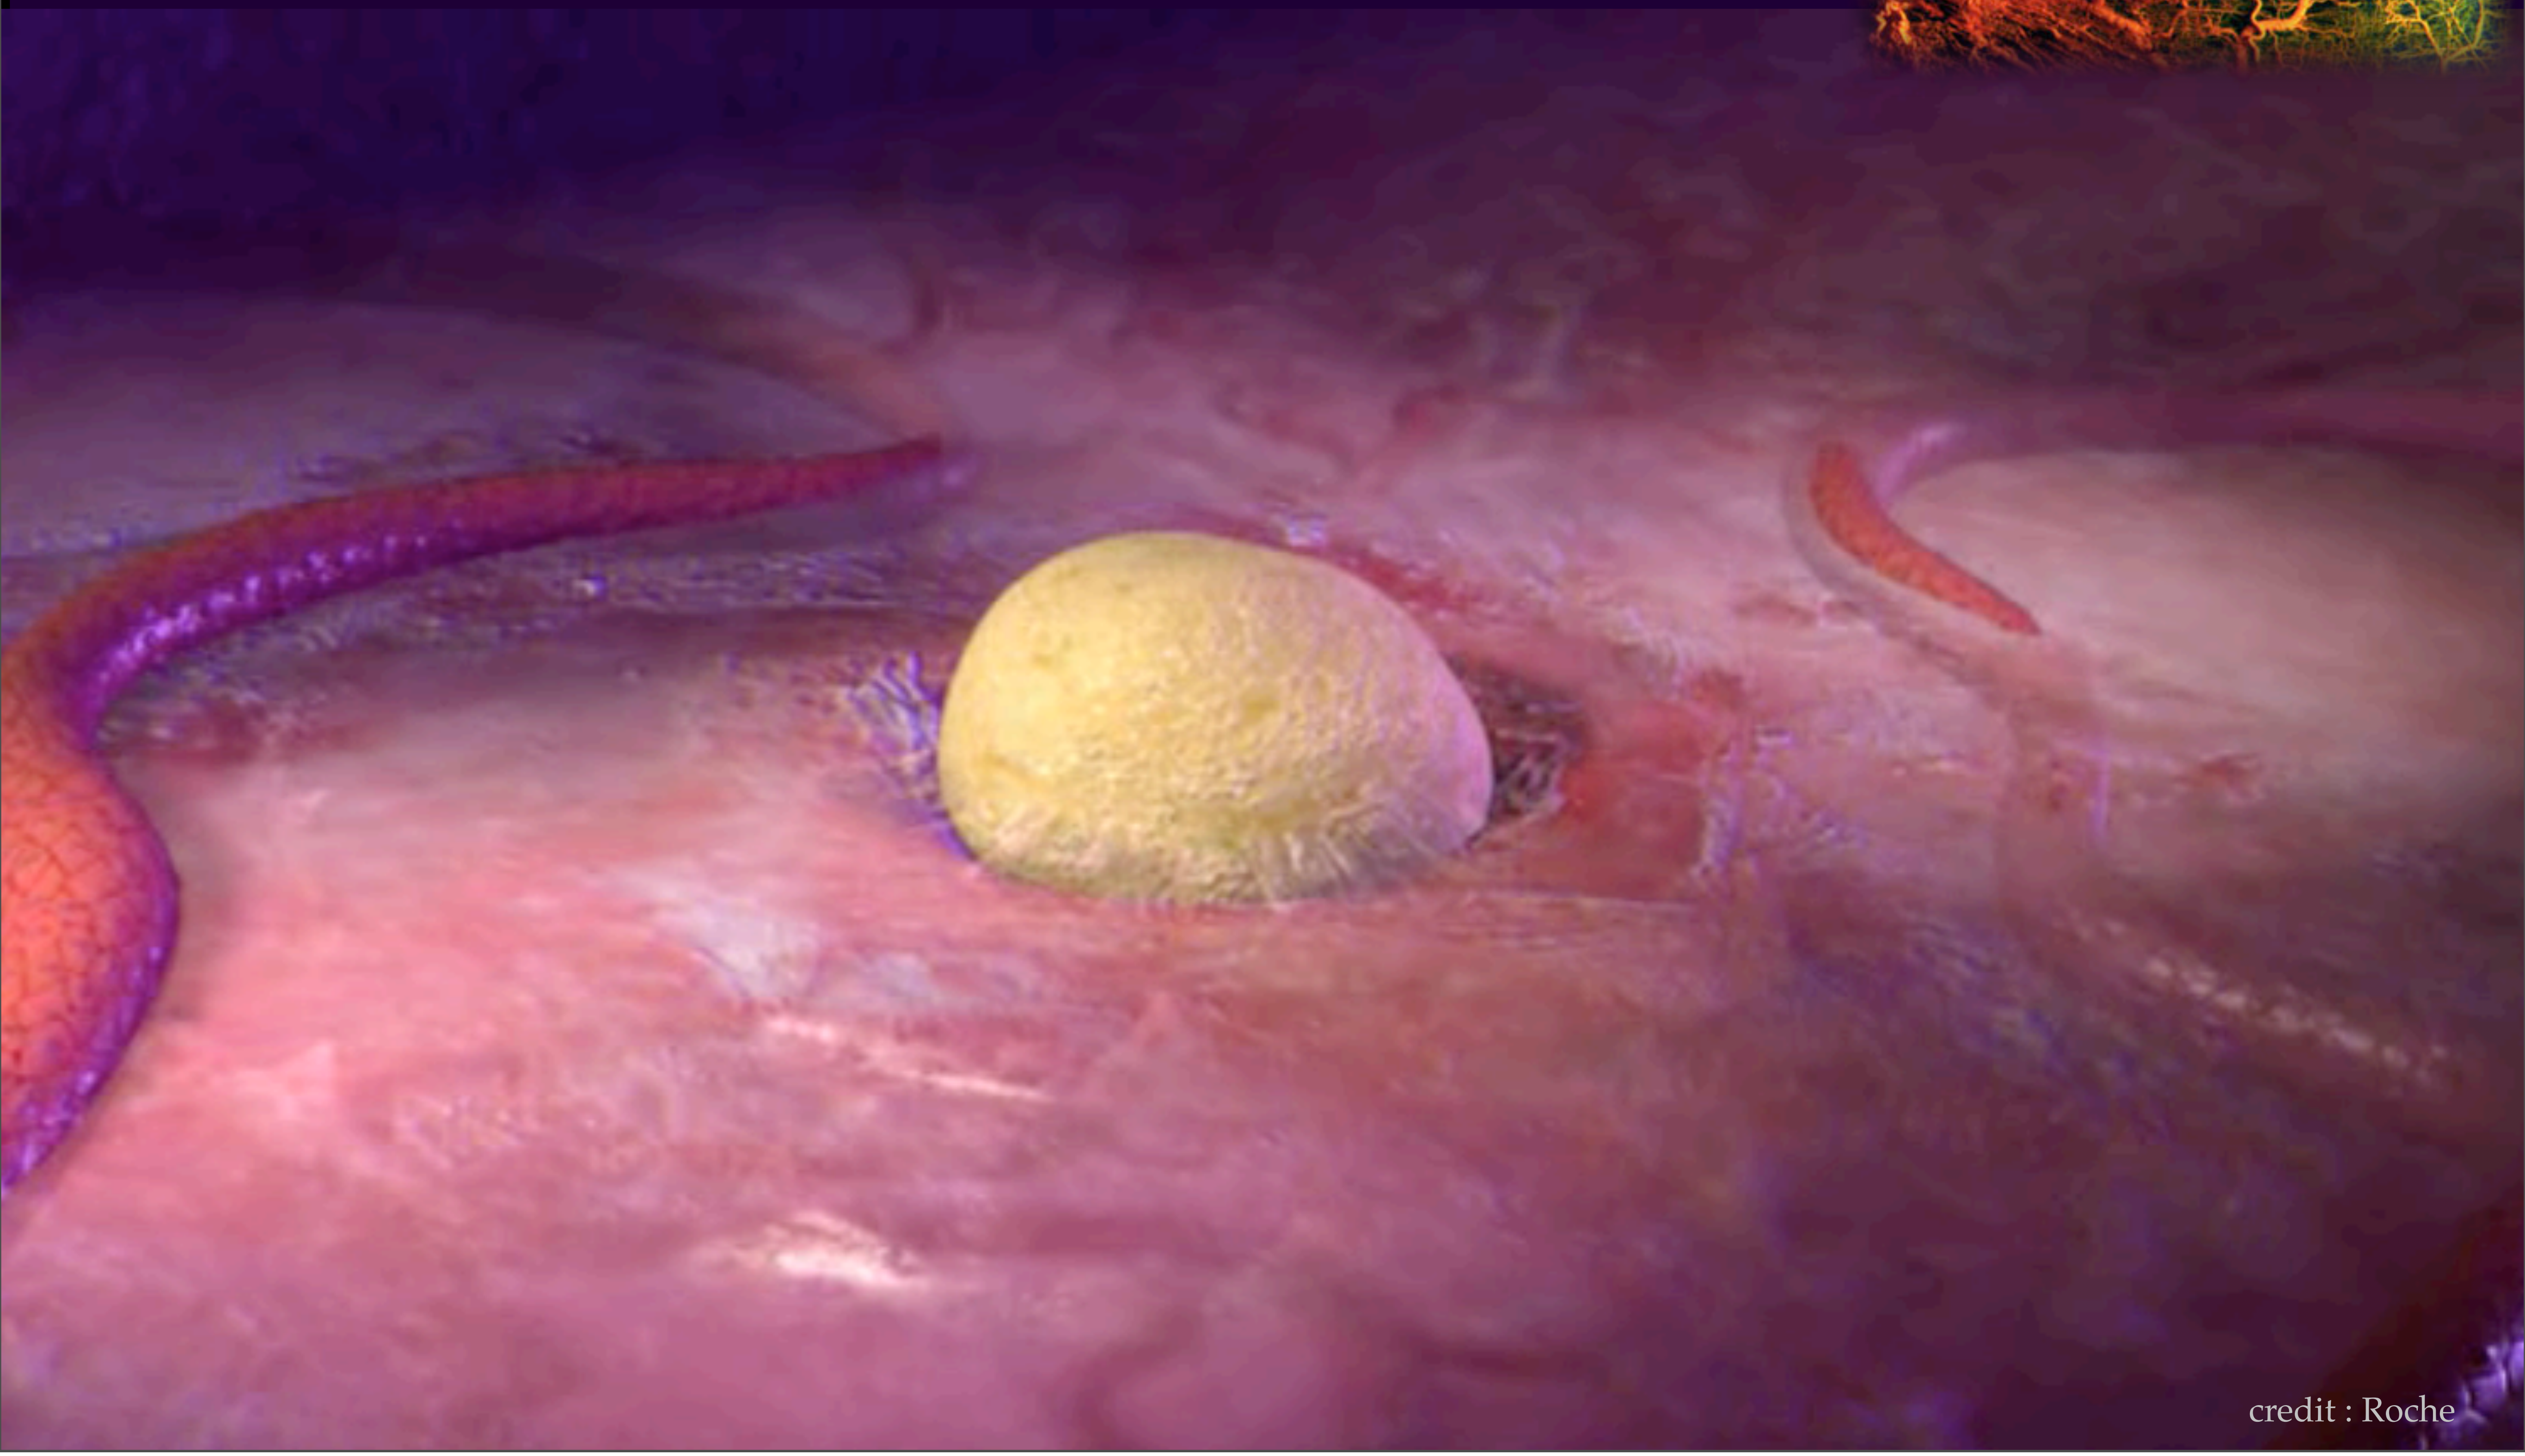

CFD: Then and Now

Re = 9500 ~ 10^6 particles

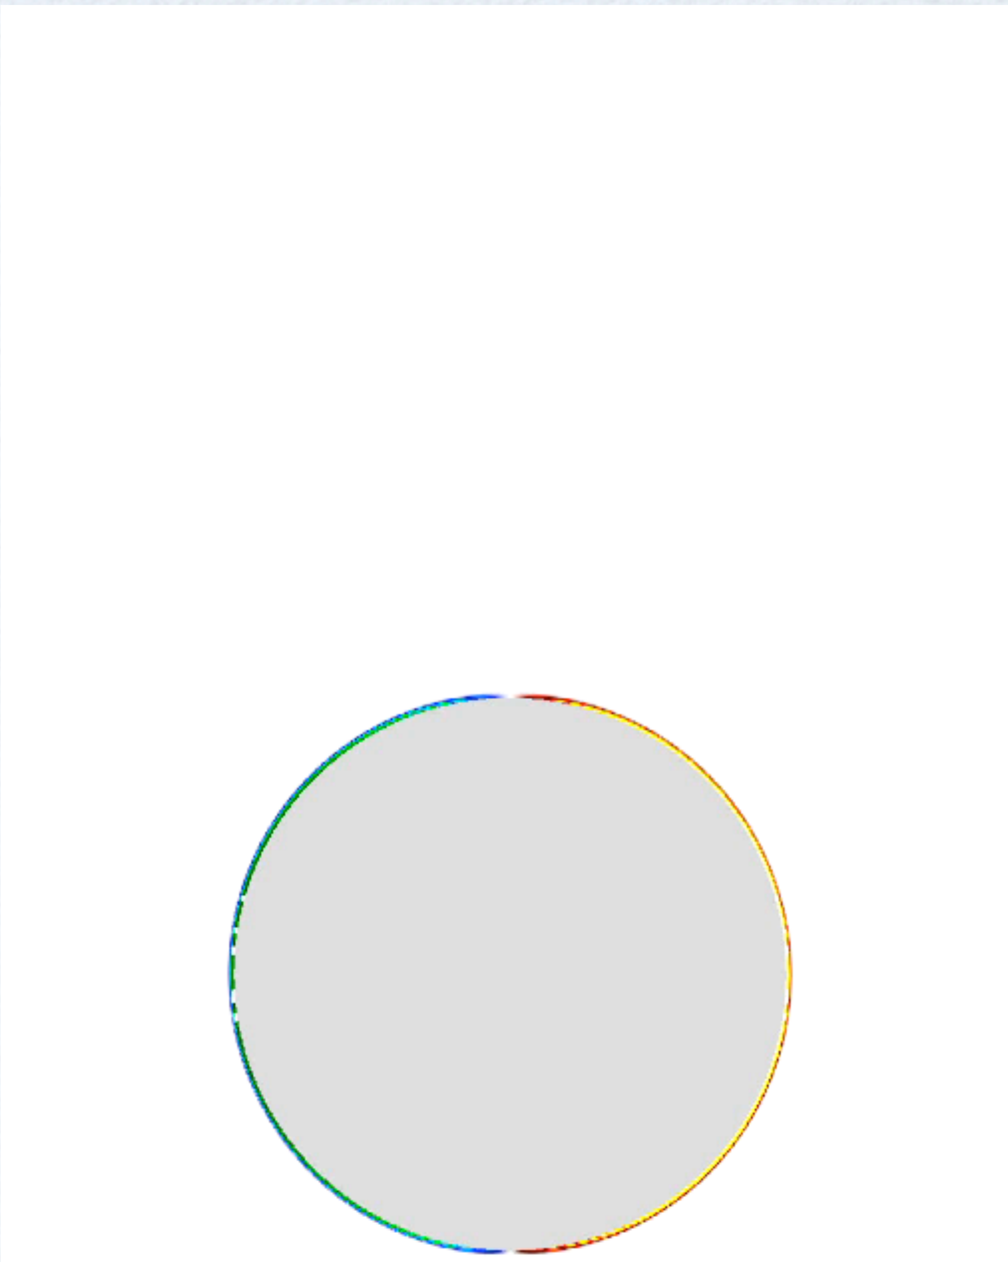

1995 **20 Days** on CRAY YMP

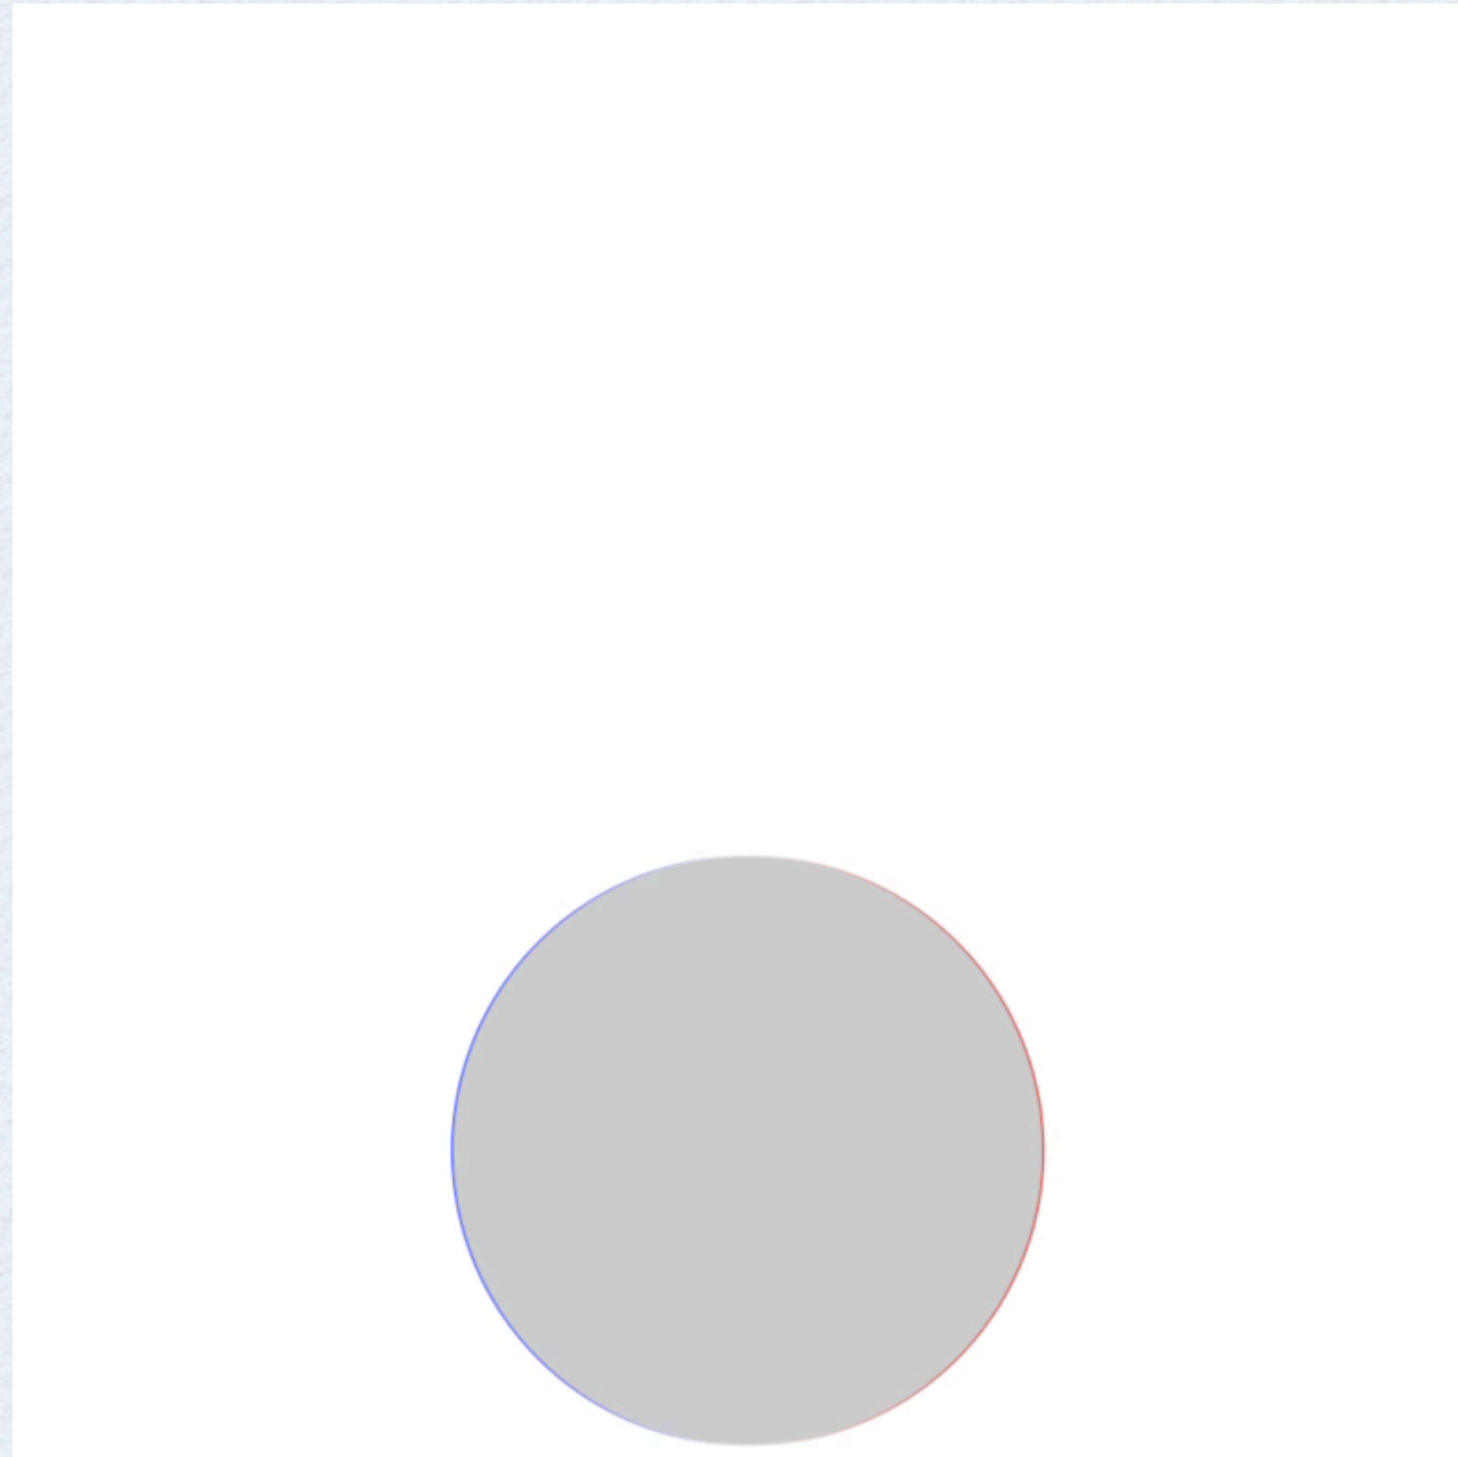

2009 **150sec** on GPU

Rossinelli D., et.al., GPU accelerated simulations of bluff body flows using vortex particle methods, **Journal of Computational Physics**, 229, 9, 3316-3333, **2010**

Outline

- COMPLEX DEFORMING GEOMETRIES
  - Meshing or Meshless ?
- FAST AND ACCURATE SIMULATIONS
  - Multiresolution and GPUs
- APPLICATIONS
  - Fish Hydrodynamics
  - Tumor induced Angiogenesis

PARTICLES : Lagrangian, Conservation and Other Laws

SPH, Vortex Methods

$$\rho_p \frac{D\mathbf{u_p}}{Dt} = (\nabla \cdot \sigma)_p$$

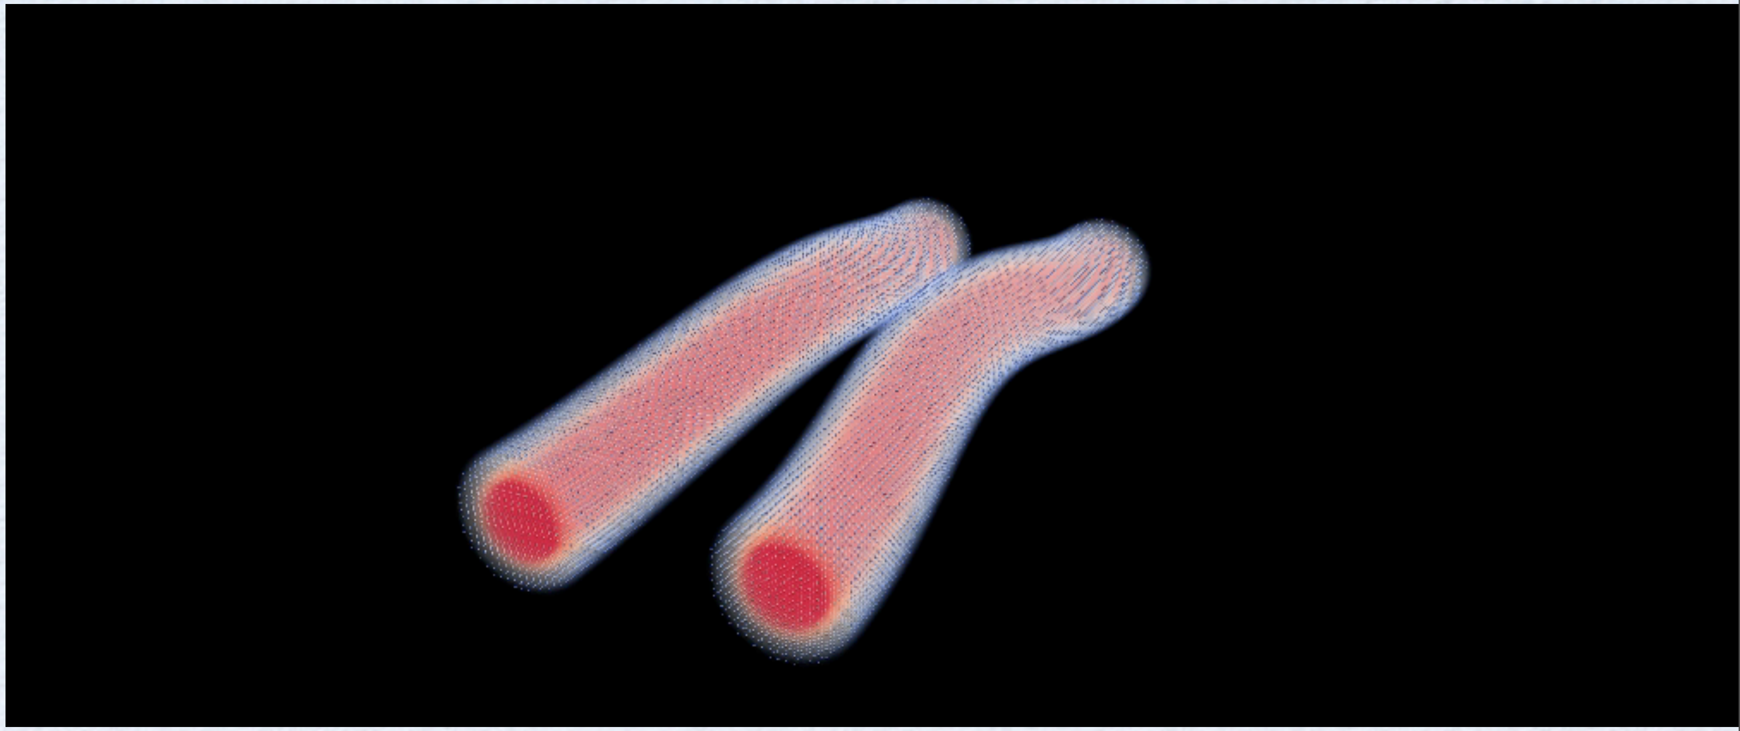

$$\frac{d\mathbf{x_p}}{dt} = \mathbf{u}_p$$

$$m \frac{d\mathbf{u_p}}{dt} = F_p$$

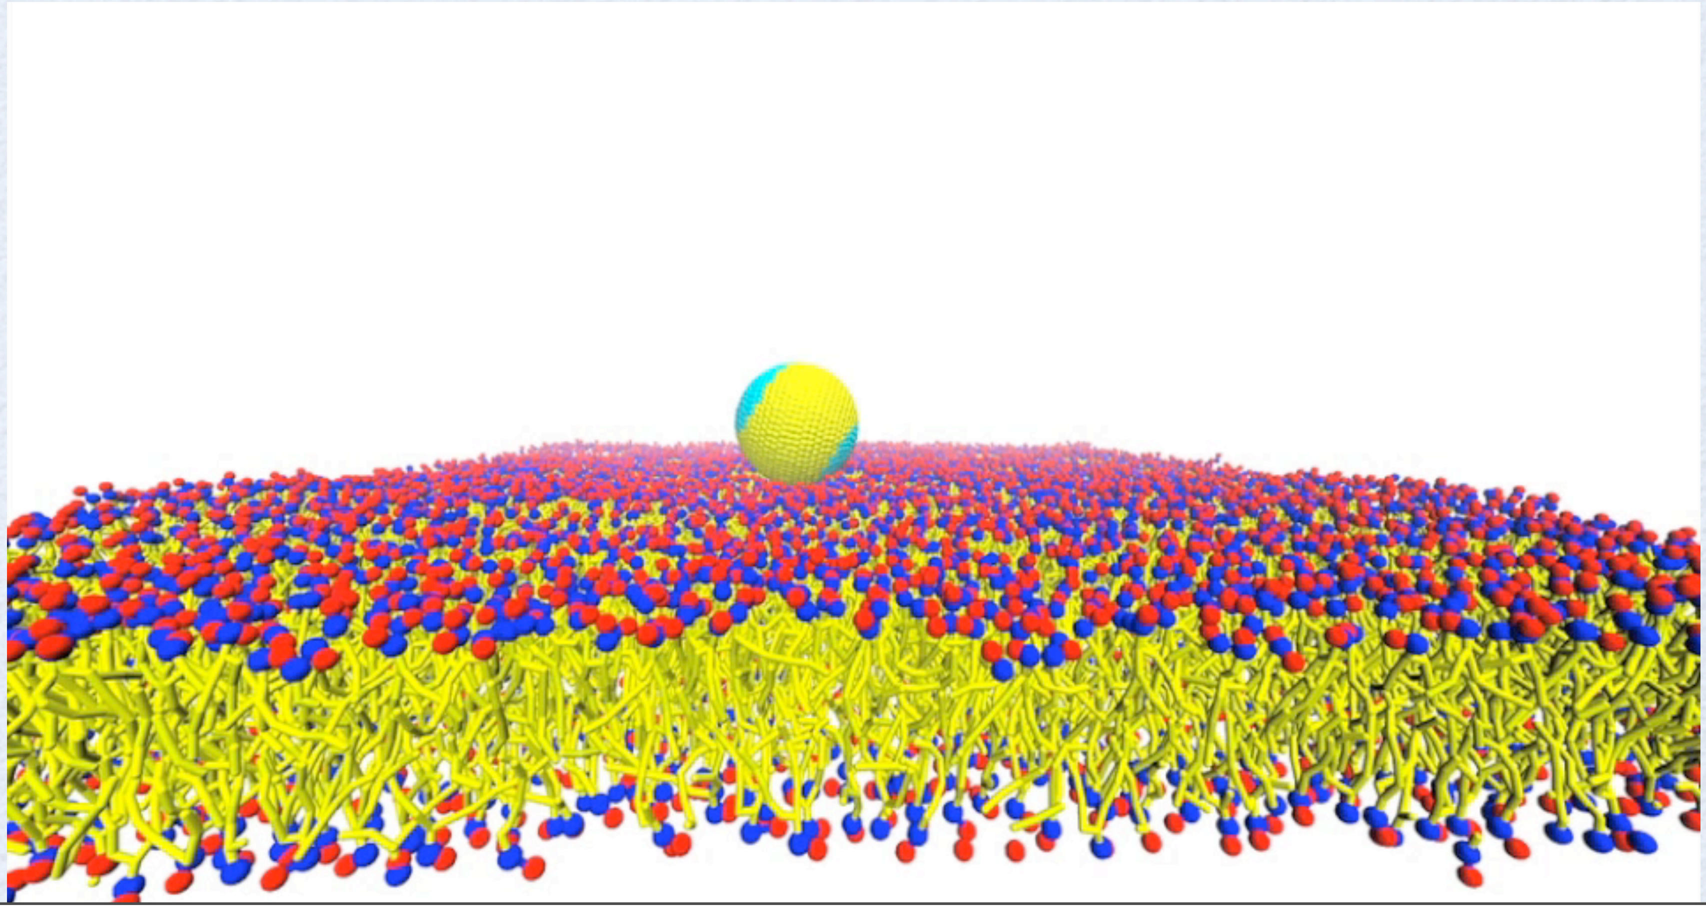

MD, DPD, CGMD

PARTICLE APPROXIMATIONS

Function Mollification

$$\Phi_\epsilon(x) = \int \Phi(y)\, \zeta_\epsilon(x-y)\, dy$$

Smooth Particle Quadrature

$$\Phi^h_\epsilon(x,t) = \sum_{p=1}^{N_p} h_p^d\, \Phi_p(t)\, \zeta_\epsilon(x - x_p(t))$$

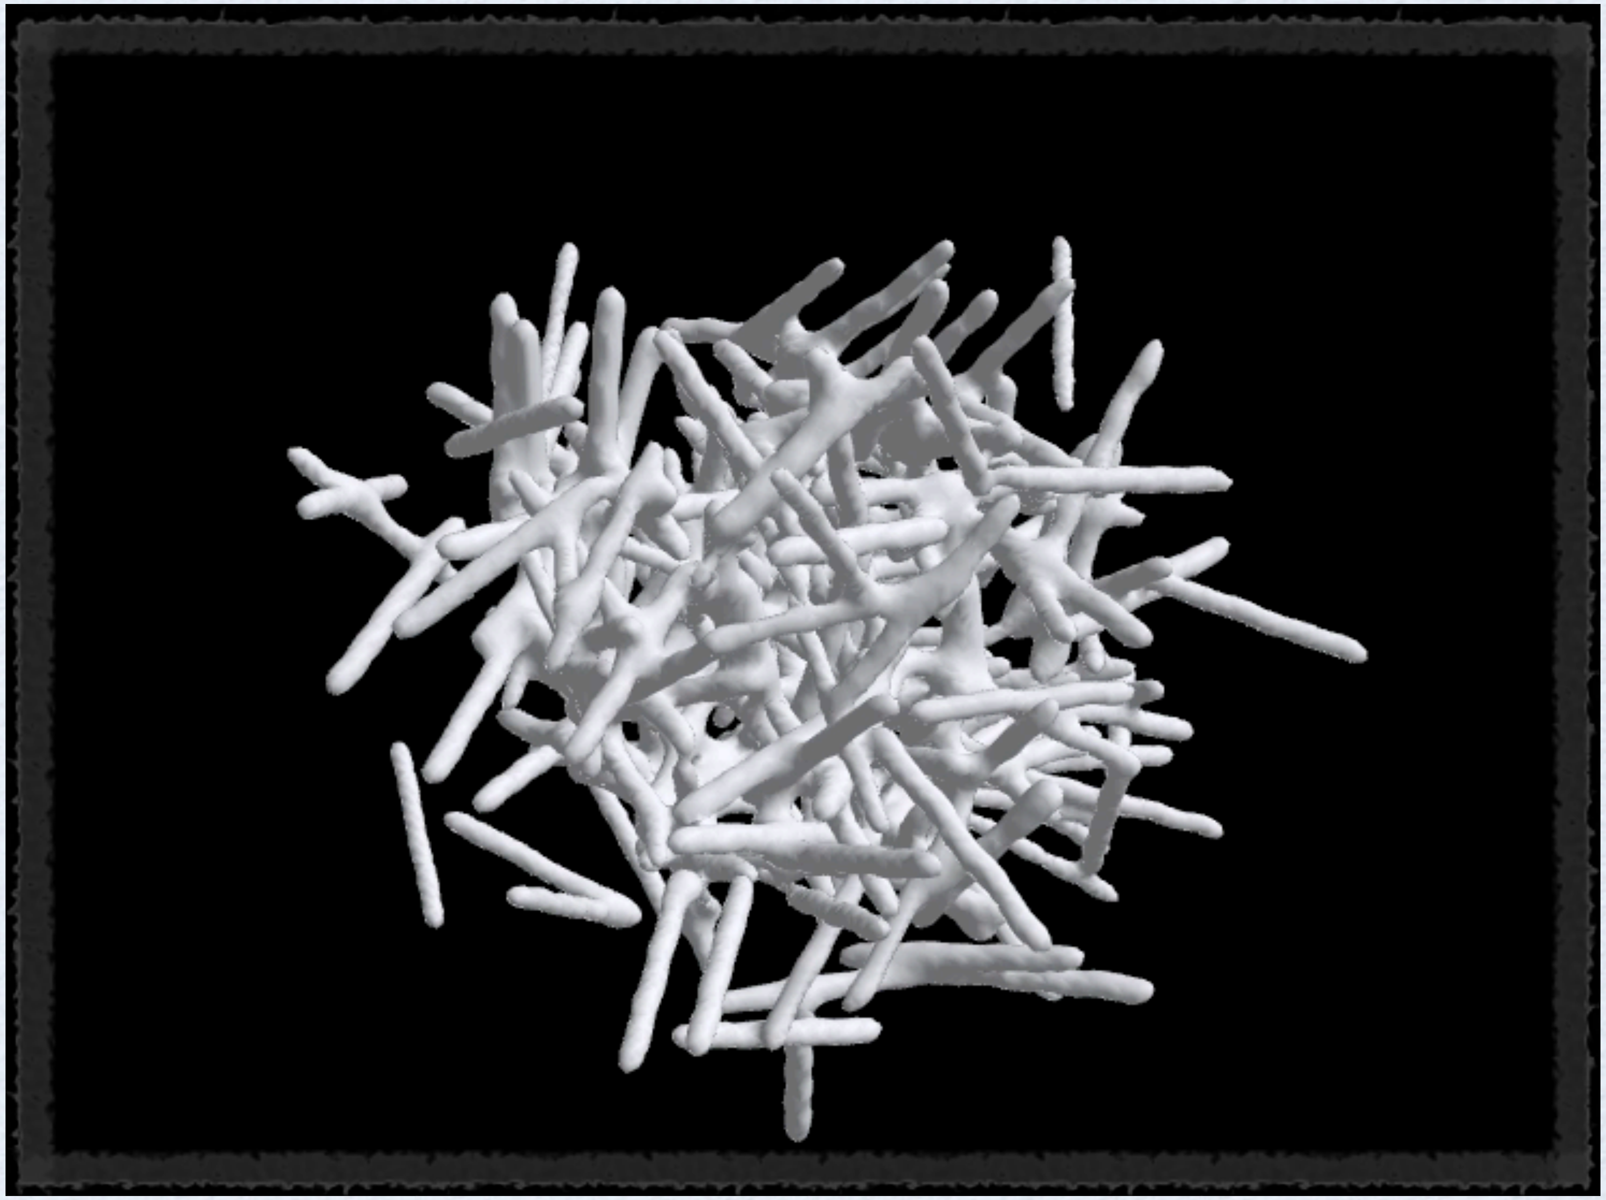

are Particles MESH Free ?

SURFACES -> LEVEL SETS

$$\Gamma(t) = \{\mathbf{x} \in \Omega \mid \phi(\mathbf{x}, t) = 0\}$$

$$|\nabla \phi| = 1$$

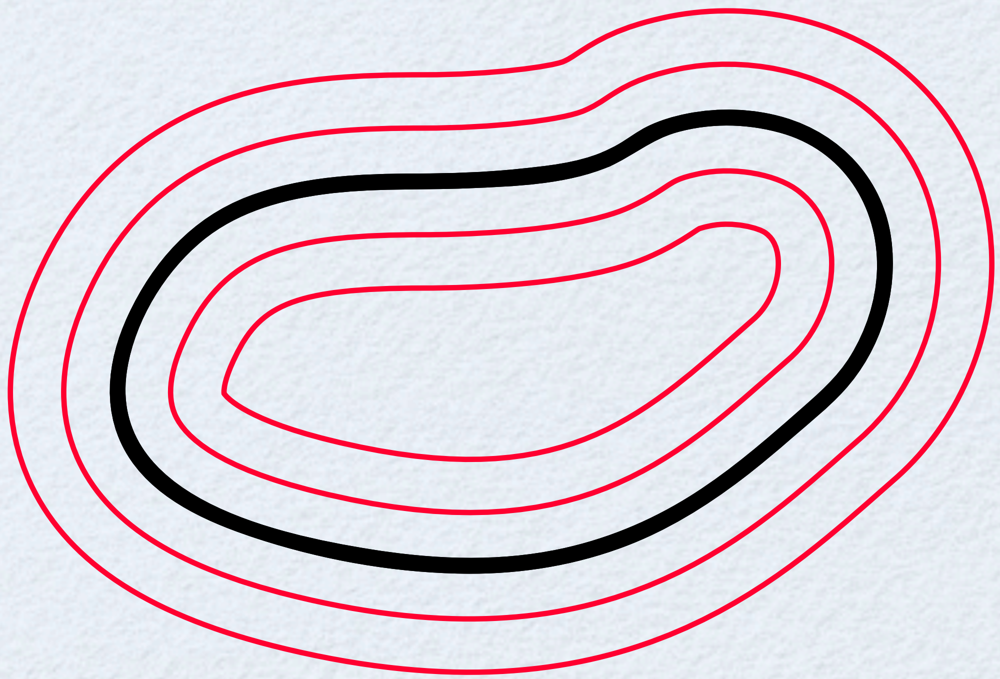

EVOLVING LEVEL SETS

$$\frac{\partial \Phi}{\partial t} + u \cdot \nabla \Phi = 0$$

PARTICLES

$$\Phi_\epsilon^h(x,t) = \sum_{p=1}^{N_p} h_p^d \, \Phi_p(t) \, \zeta_\epsilon(x - x_p(t))$$

Lagrangian Surface Transport

$$\frac{dx_p}{dt} = \mathbf{u_p} \qquad \frac{D\Phi_p}{Dt} = 0$$

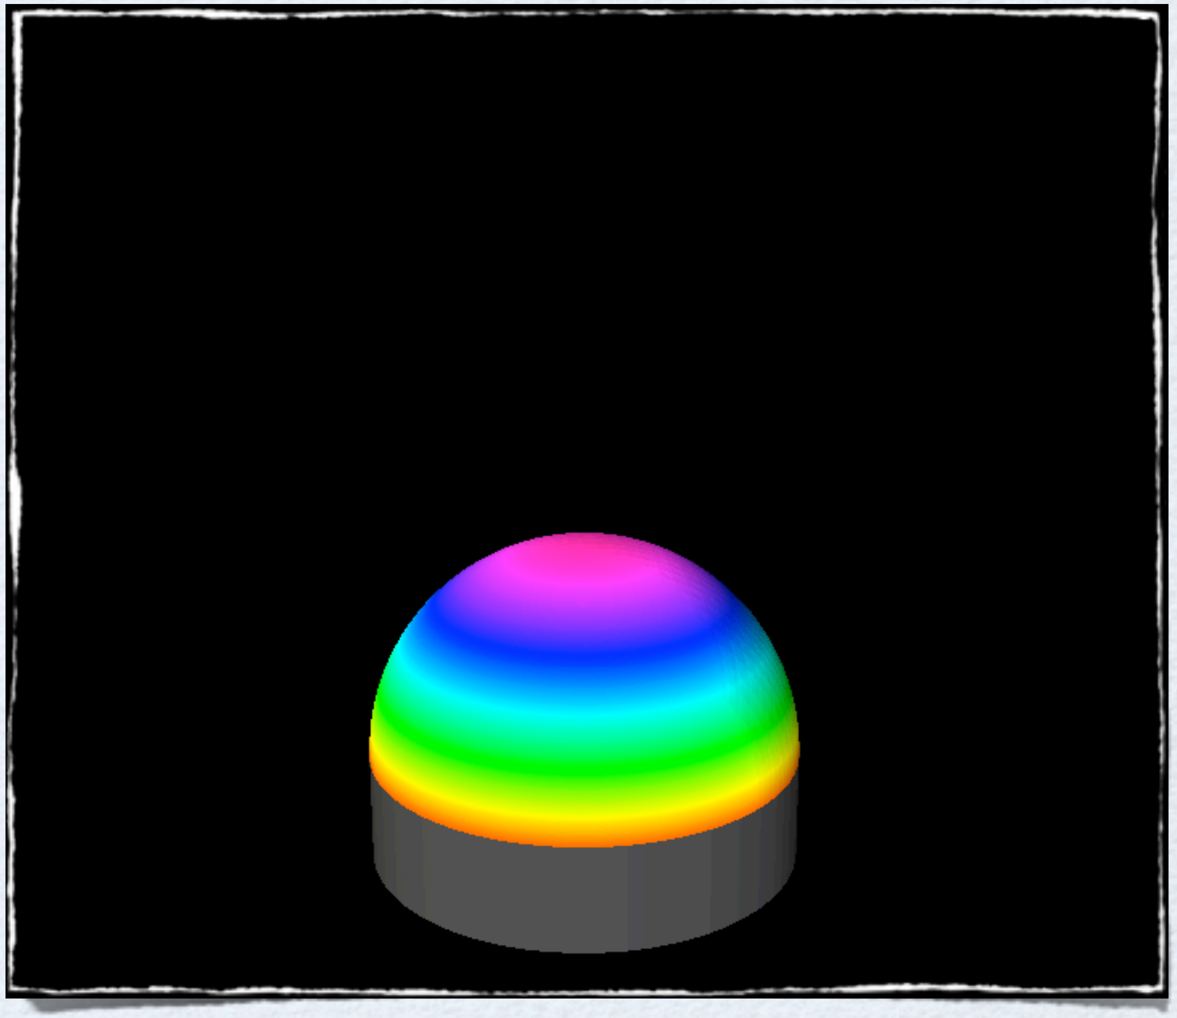

S. E. Hieber and P. Koumoutsakos. A Lagrangian particle level set method. **J. Computational Physics**, 210:342-367, 2005

Lagrangian vs Eulerian Descriptions

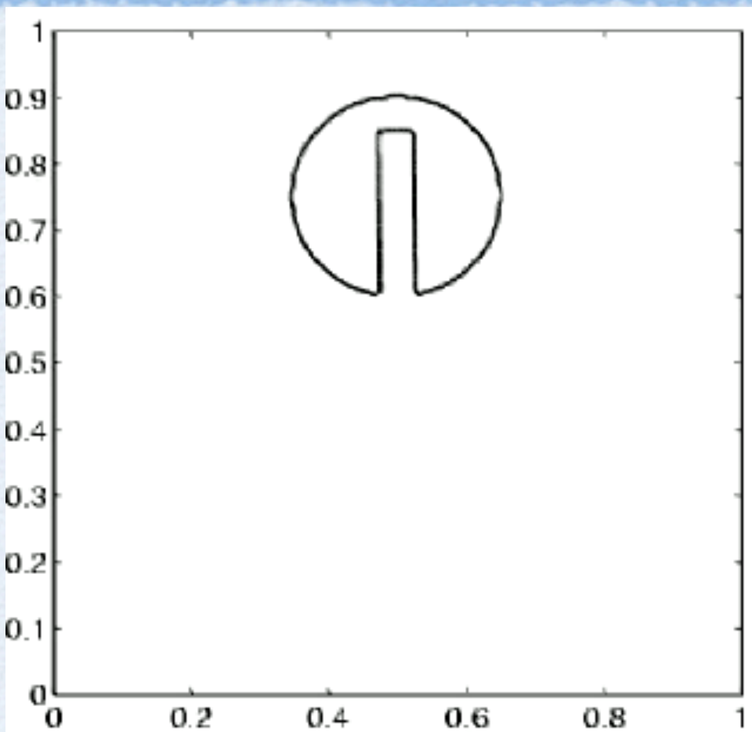

- **PARTICLE LEVEL SETS** exact for rigid body motion

$$\Phi(\mathbf{x}, t) = \Phi_0(\mathbf{x} - \mathbf{u}t)$$

Particle methods PERFECT for **linear** advection

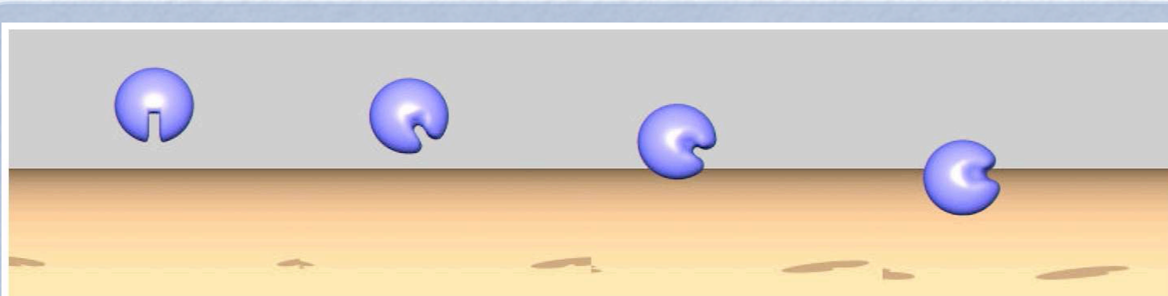

Hubrid Particle-Grid Level Sets
(Enright and Fedkiw, 2002)

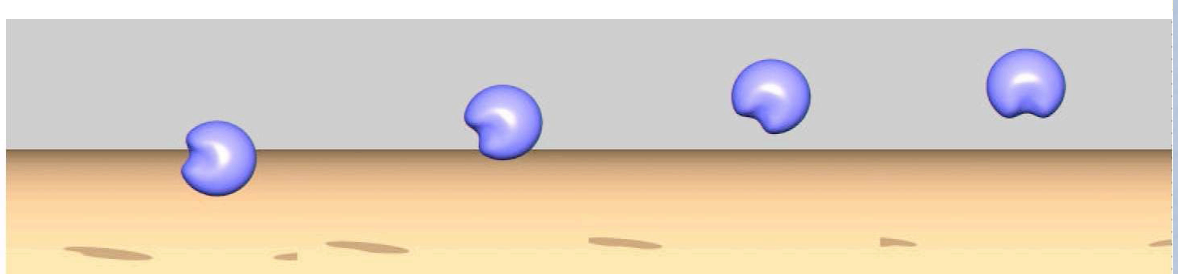

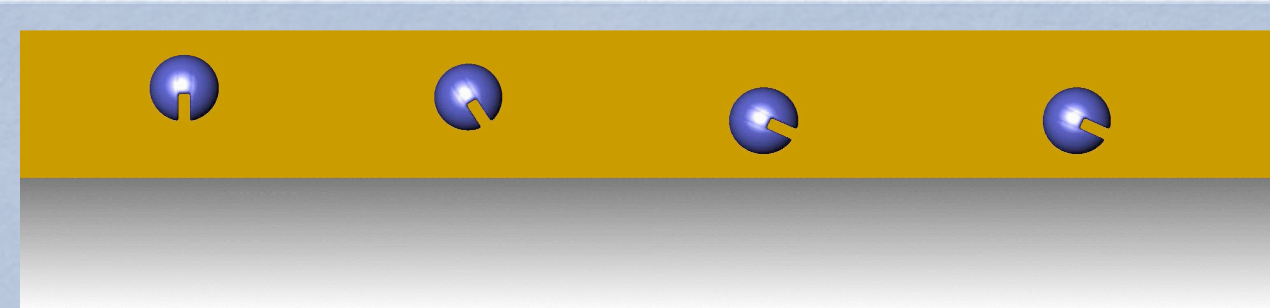

Lagrangian Particle Level Sets
(Hieber and Koumoutsakos, 2005)

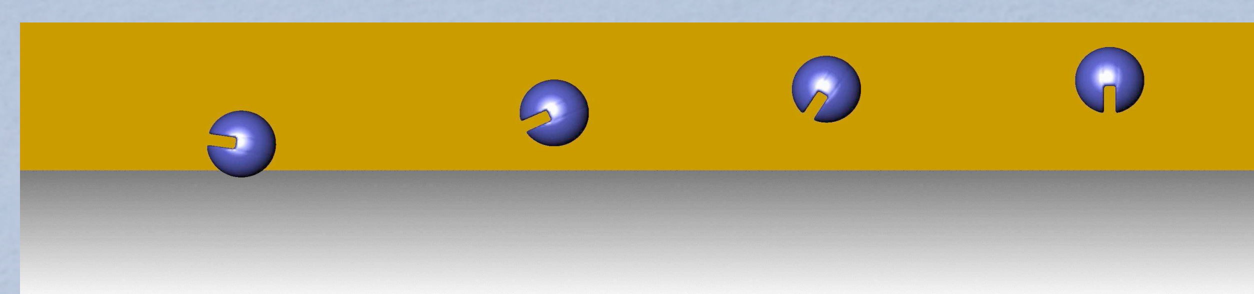

LAGRANGIAN DISTORTION

• loss of overlap -> loss of convergence

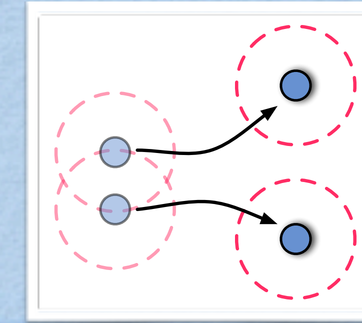

Particles follow flow trajectories - Location distortion

EXAMPLE :

Incompressible 2D Euler Equations

$$\omega = \nabla \times \mathbf{u} \quad \nabla \cdot \mathbf{u} = 0$$

$$\frac{D\omega}{Dt} = 0$$

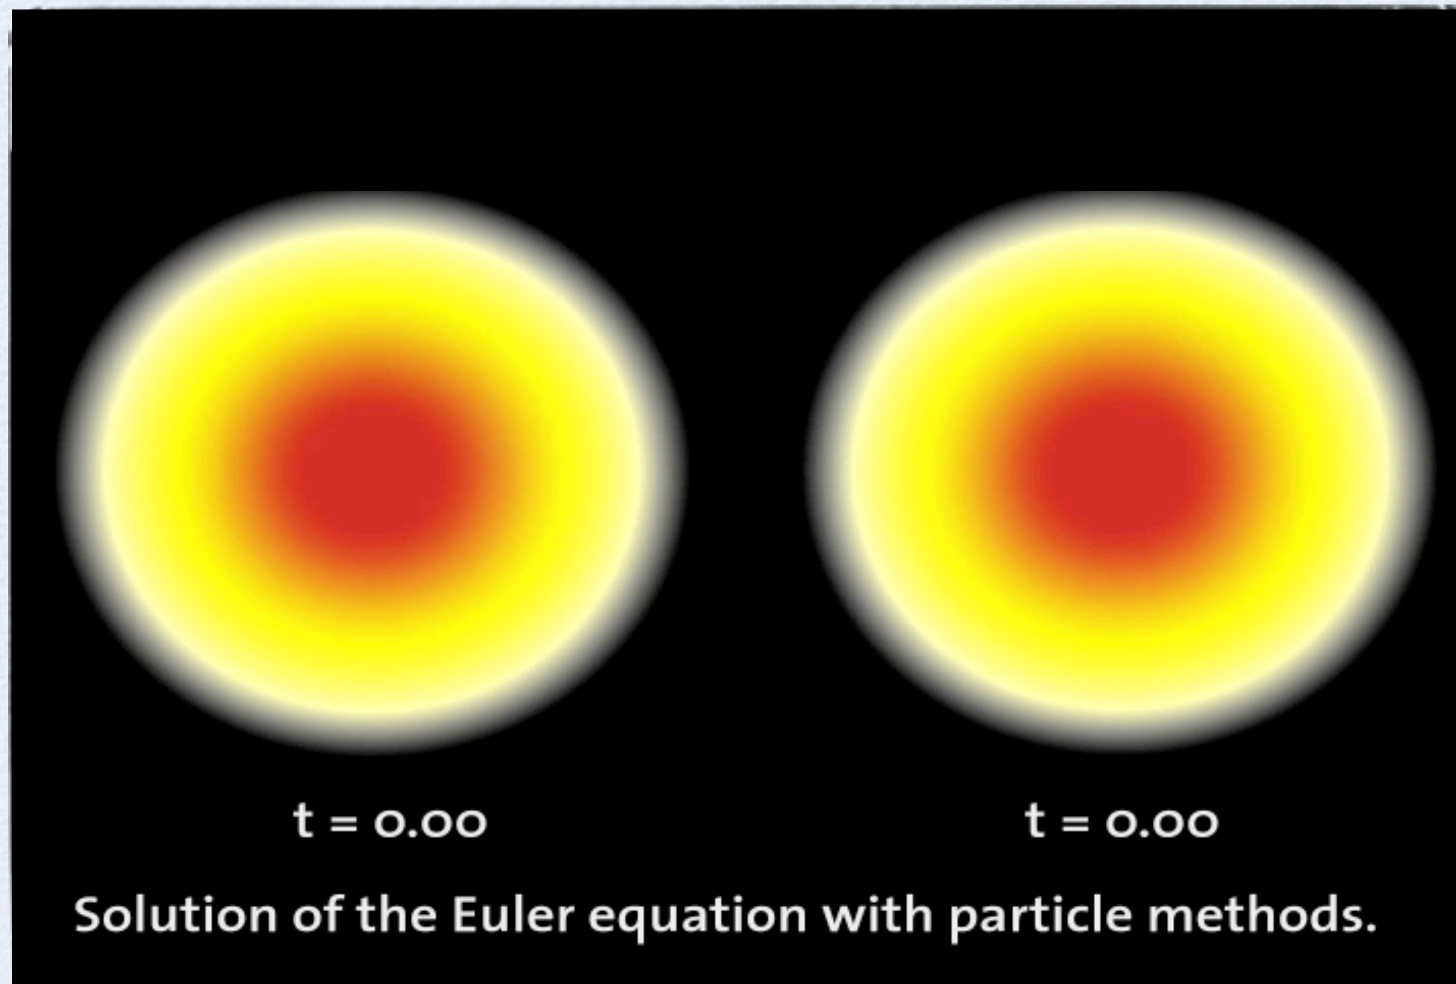

There is an exact axisymmetric solution

Are Particle Methods Grid Free ?

How to fix it ?

- Modify the smoothing kernels (SPH - Monaghan) — DOES NOT WORK
- Re-distribute particles with Voronoi Meshes (ALE - Russo) — EXPENSIVE - UNSTABLE
- Re-initialise particle strengths (WRKPM - Liu, Belytchko) — EXPENSIVE

REMESHING : Re-project particles on a mesh

- NO MESH-FREE particle methods
- Can use all the “tricks” of mesh based methods
- High CFL
- Multiresolution & Multiscaling
- ....

Particle Remeshing

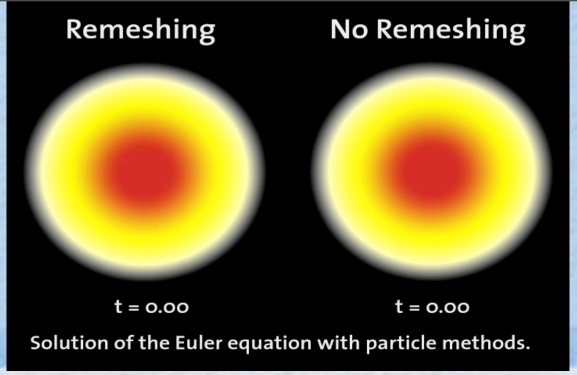

Koumoutsakos, J. Comp. Phys., 1997

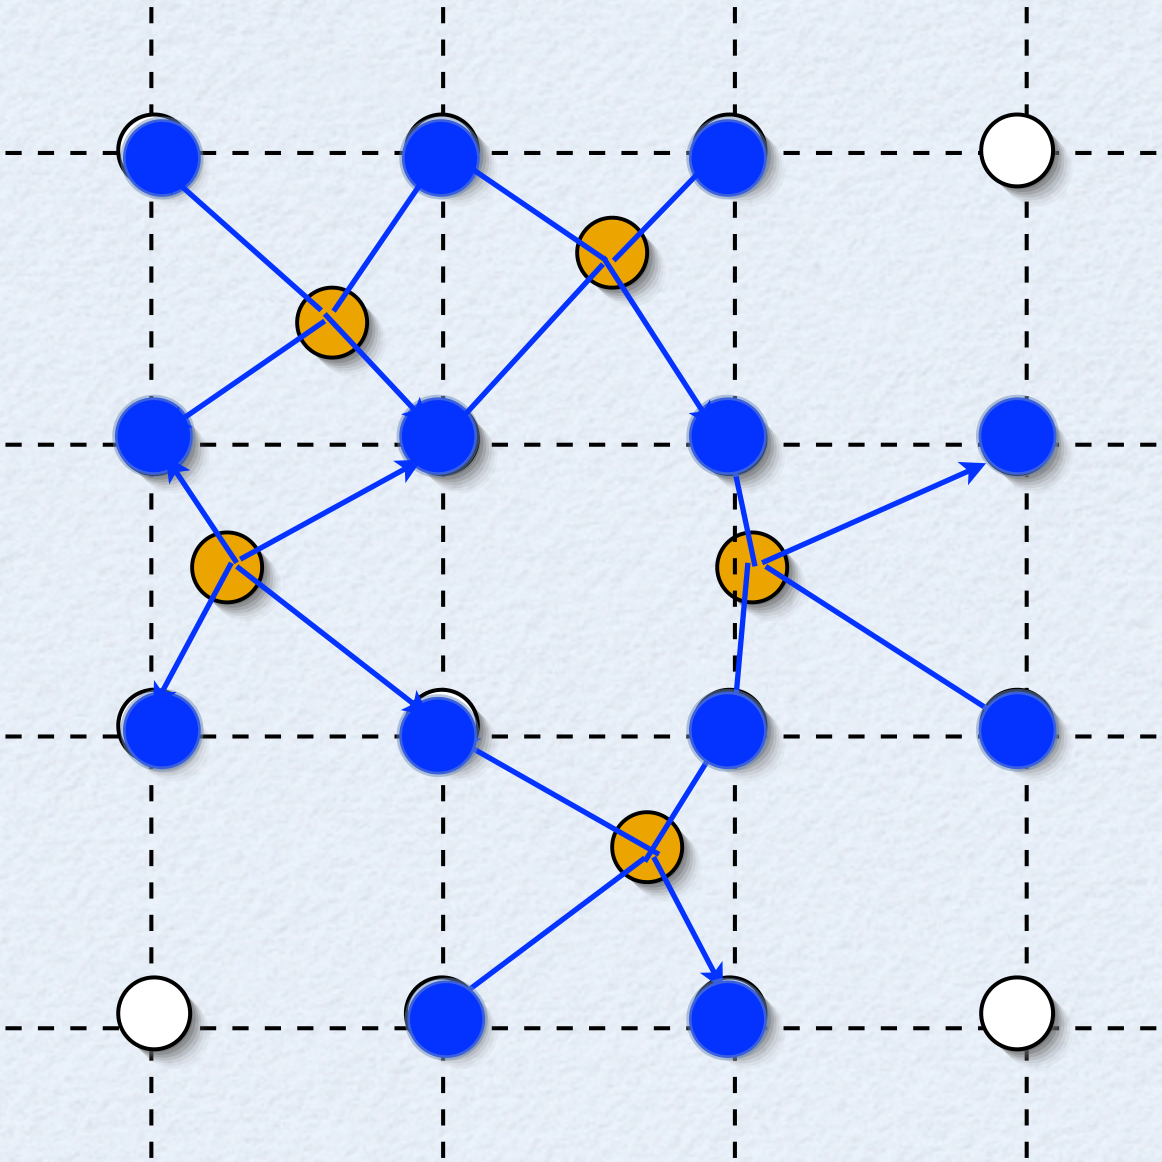

Moment Conserving Interpolation : $Q_p^{\text{new}} = \sum_{p'} Q_{p'} M(j\,h - x_{p'})$

remeshed PARTICLE METHODS (rPM)

1. ADVECT : Particles ->Large CFL

2. REMESH : Particles to Mesh -> Gather/Scatter

3. SOLVE: Poisson/Derivatives on Mesh ->FFTw/Ghosts

4: RESAMPLE: Mesh Nodes BECOME Particles

VORTEX RING COLLISION, Re = 1800

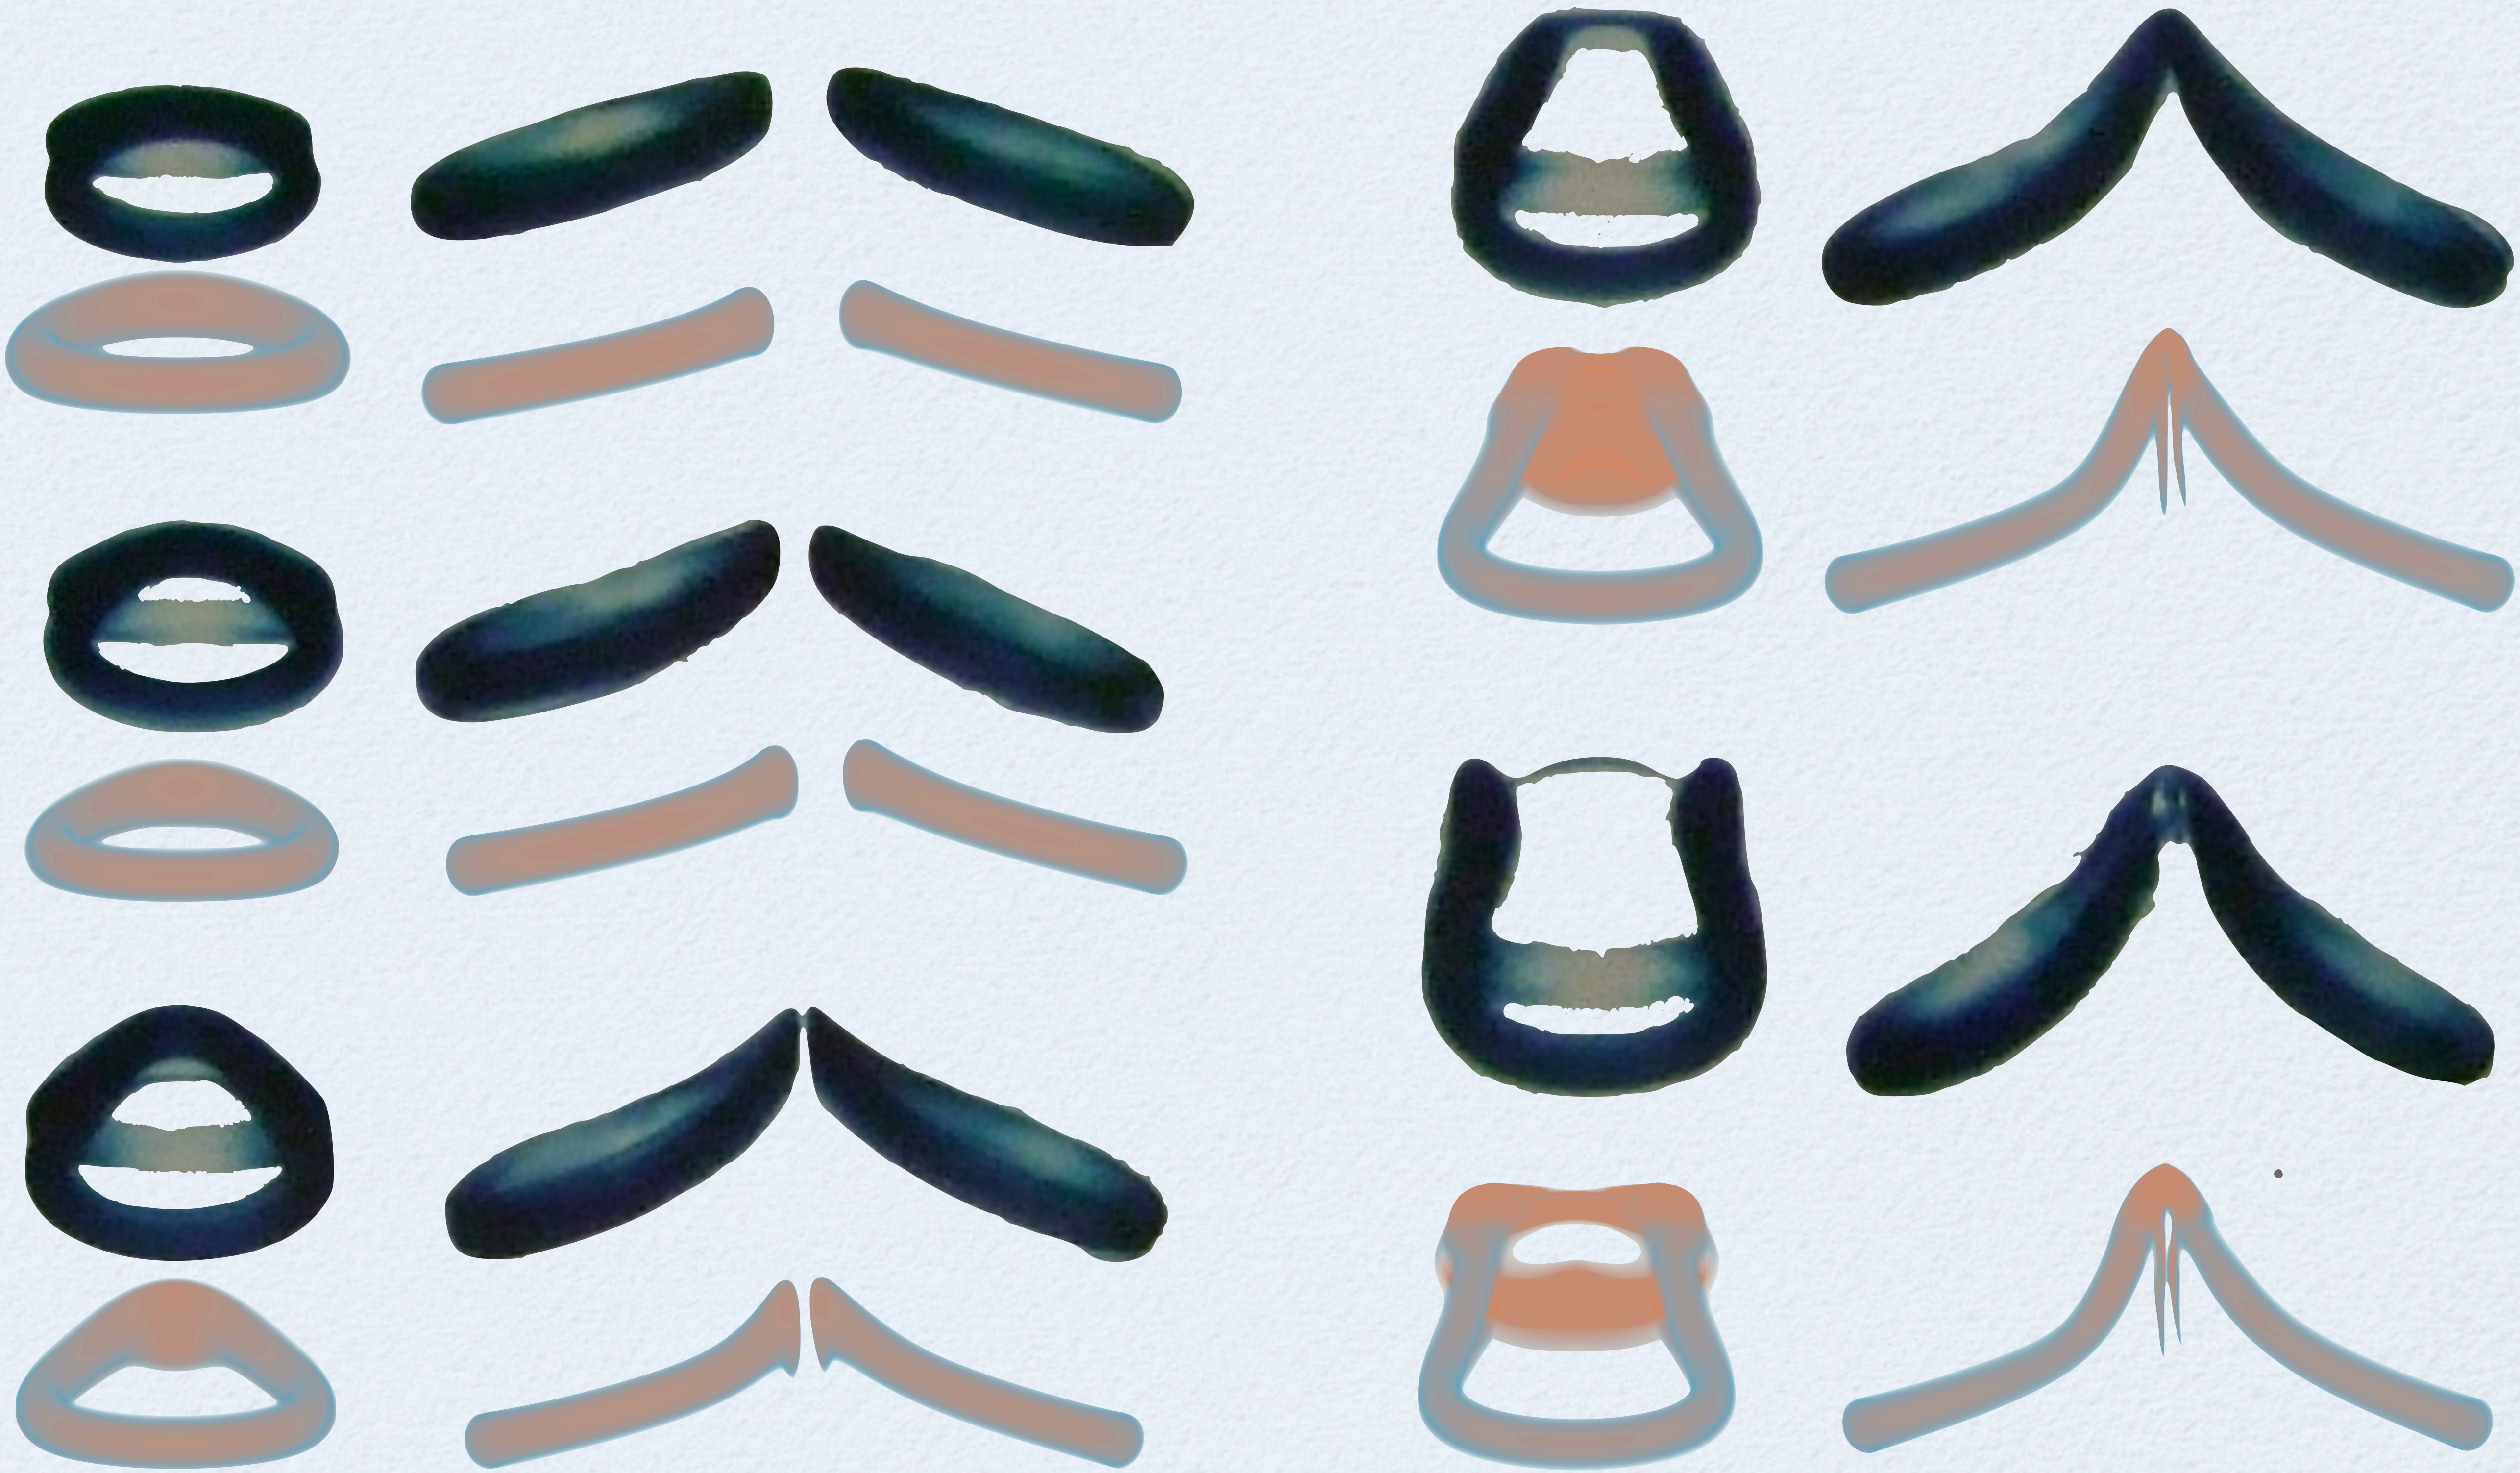

Experiments : P. Schatzle & D. Coles (1986)

VORTEX DYNAMICS at High Re

van Rees W.M., Leonard A., Pullin D.I., Koumoutsakos P., A comparison of vortex and pseudo-spectral methods for the simulation of periodic vortical flows at high Reynolds numbers, **J. of Comp.Physics**, 2011

rPM : ADAPTIVE

yet inefficient !

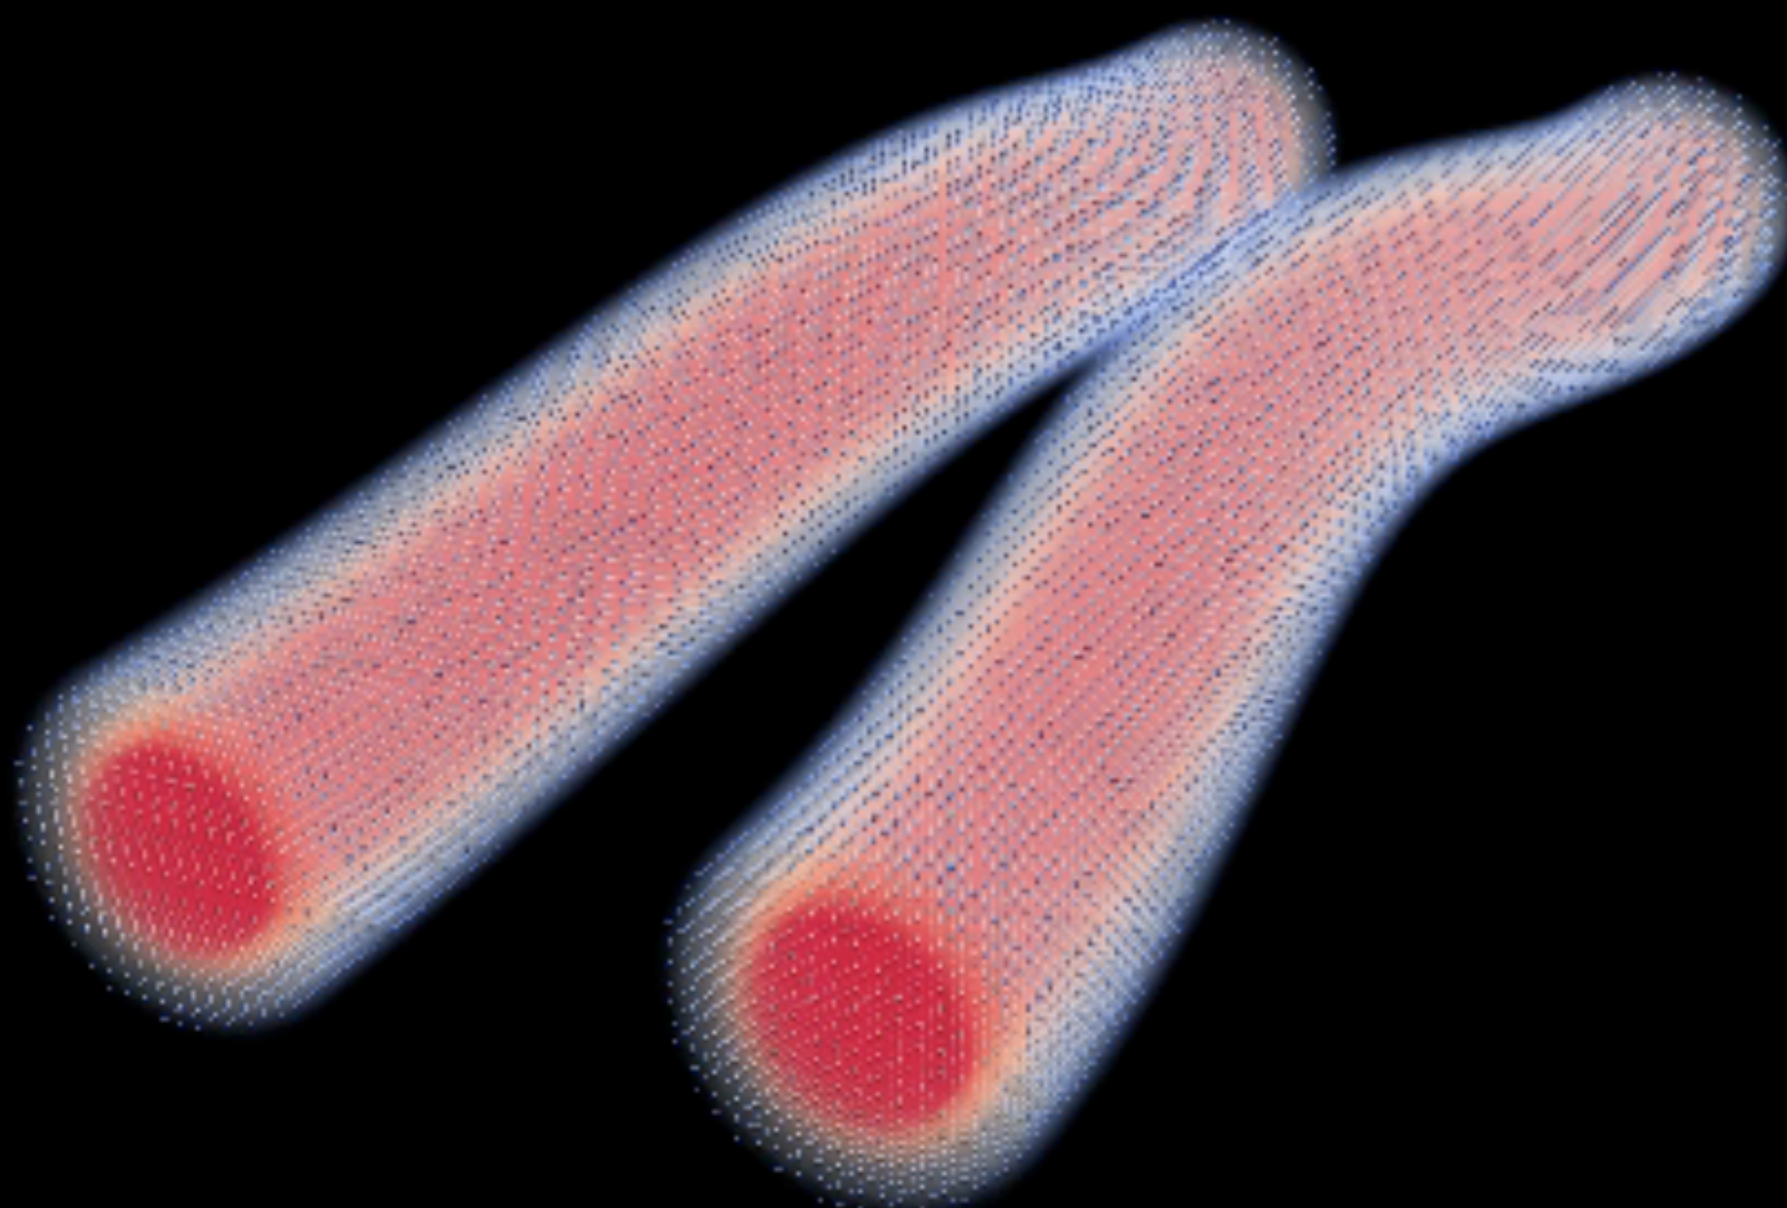

Adaptive Mesh Refinement

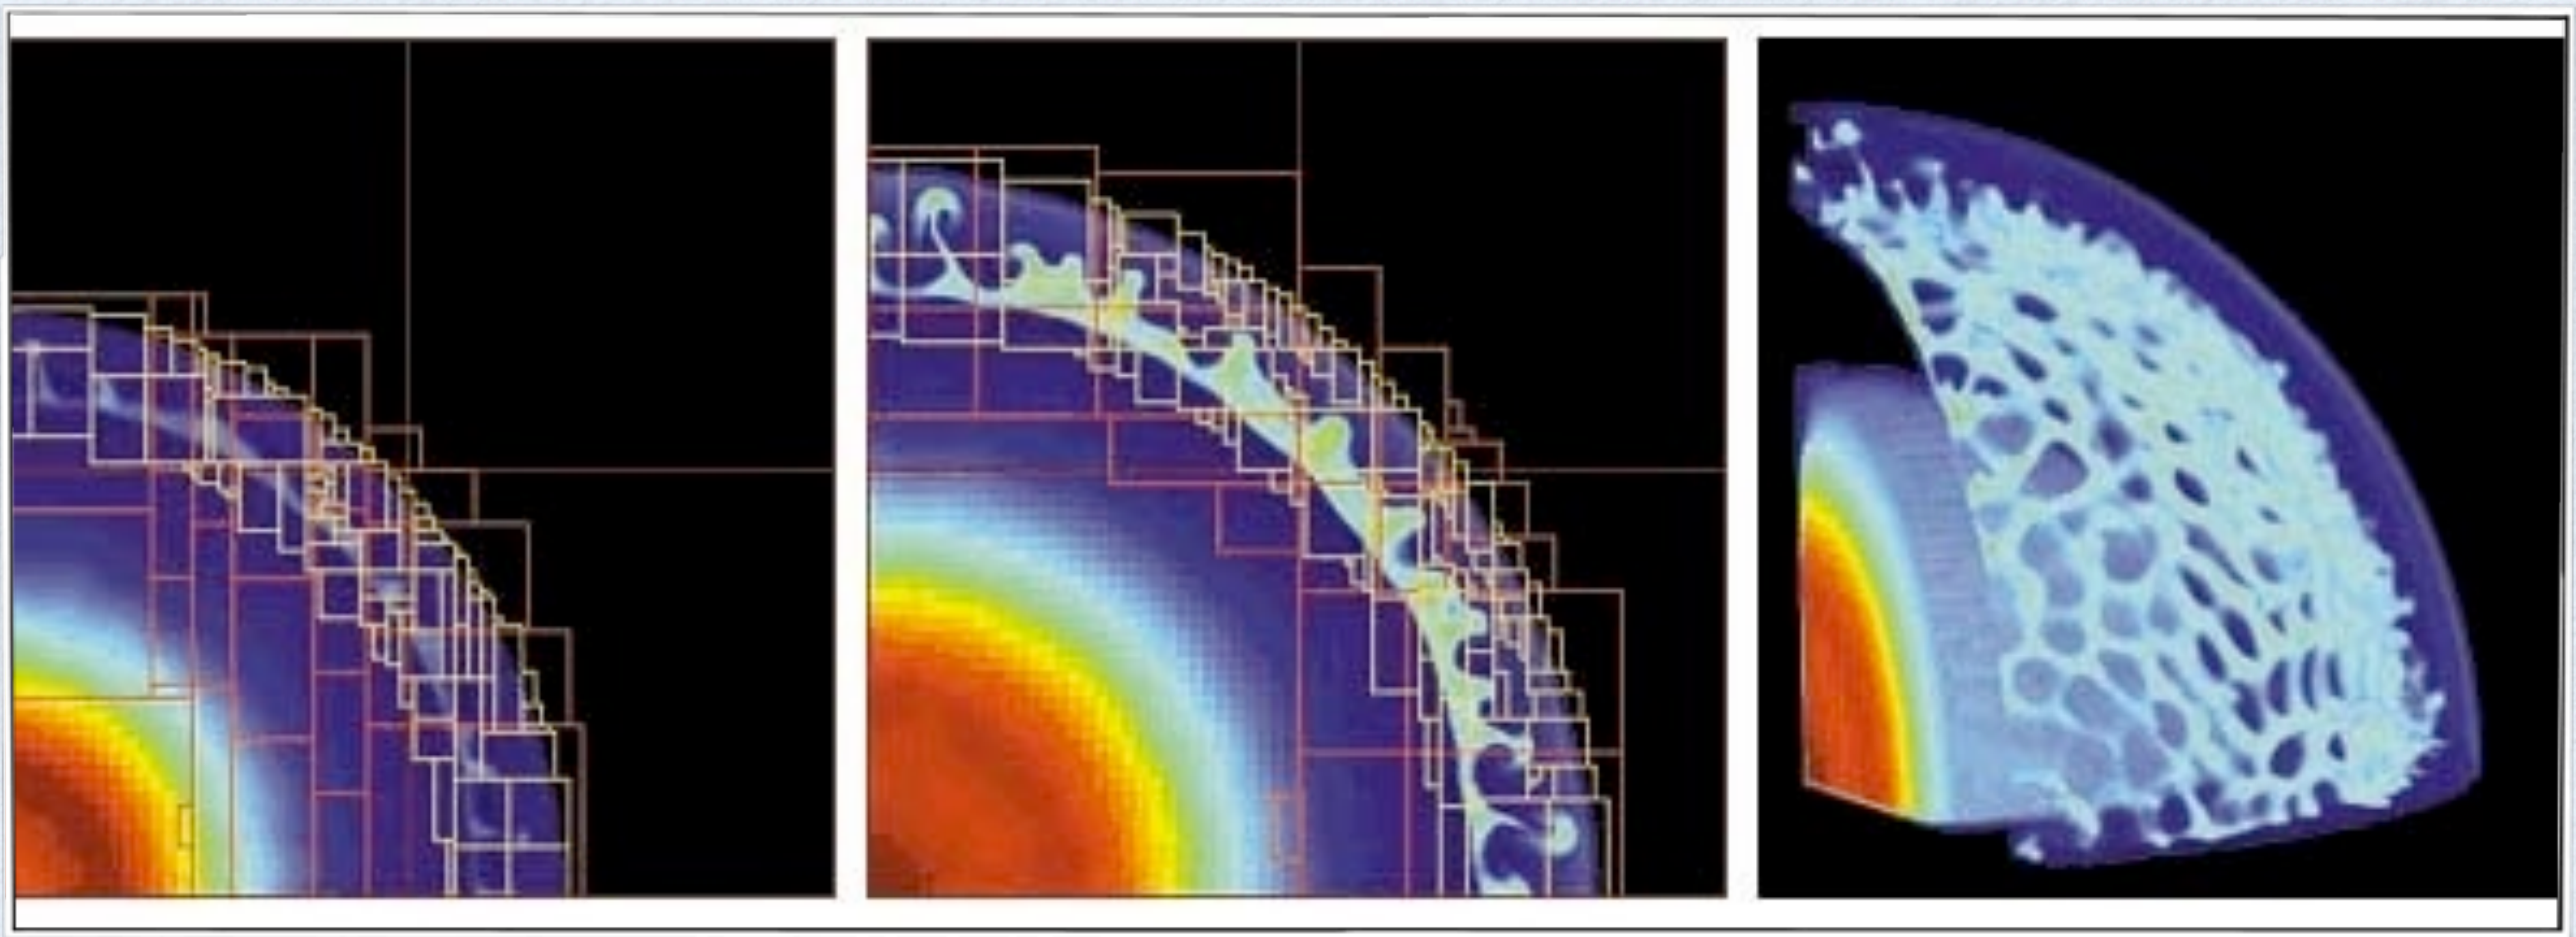

- Support of unstructured grids
- Different mesh orientations

- Low compression rate (Gradient, curvature)
- No explicit control on the compression error

Berger, Colella,J. Comp. Phys., 1989

Wavelet Compression

50:1

WAVELET PARTICLE METHOD

While particles are on grid locations

mollification kernel ⟷ basis/scaling function

Multiresolution analysis (MRA) $\{\mathcal{V}^l\}_{l=0}^{L}$ of particle quantities

Refineable kernels
as basis functions of $\mathcal{V}^l$

$$\zeta_k^l = \sum_j h_{j,k}^l \, \zeta_j^{l+1}$$

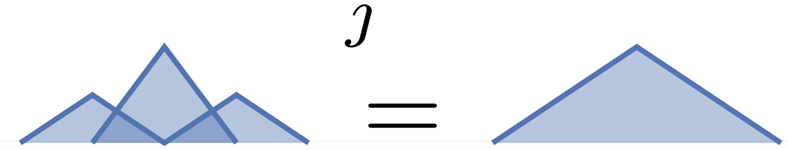

Wavelets as basis functions of the complements $\mathcal{W}^l$

$$\zeta_k^{l+1} = \sum_j \tilde{h}_{j,k}^l \zeta_j^l + \sum_j \tilde{g}_{j,k}^l \psi_j^l$$

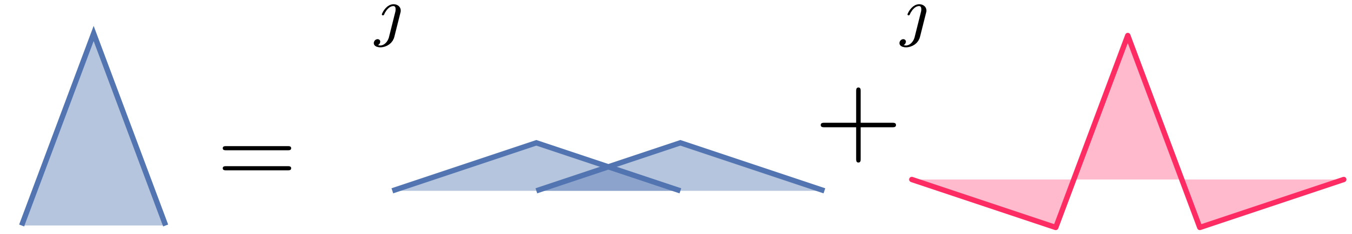

PARTICLETS : REMESHED PARTICLES + WAVELETS

M. Bergdorf, P. Koumoutsakos. A Lagrangian Particle-Wavelet Method, Multiscale Modeling and Simulation: A SIAM Interdisciplinary Journal, 5(3), 980-995, 2006

$$q^L = \sum_k c_k^0 \, \zeta_k^0 + \sum_{l<L} \sum_k d_k^l \, \psi_k^l$$

"ground" level (c_k^0)

detail coefficients (d_k^l)

wavelets (ψ_k^l)

Wavelet Active Points

=

Active Grid Points

1. Remesh
2. Wavelets- Compress/Adapt
3. Convect
4. Wavelets Reconstruct
5. GOTO 1

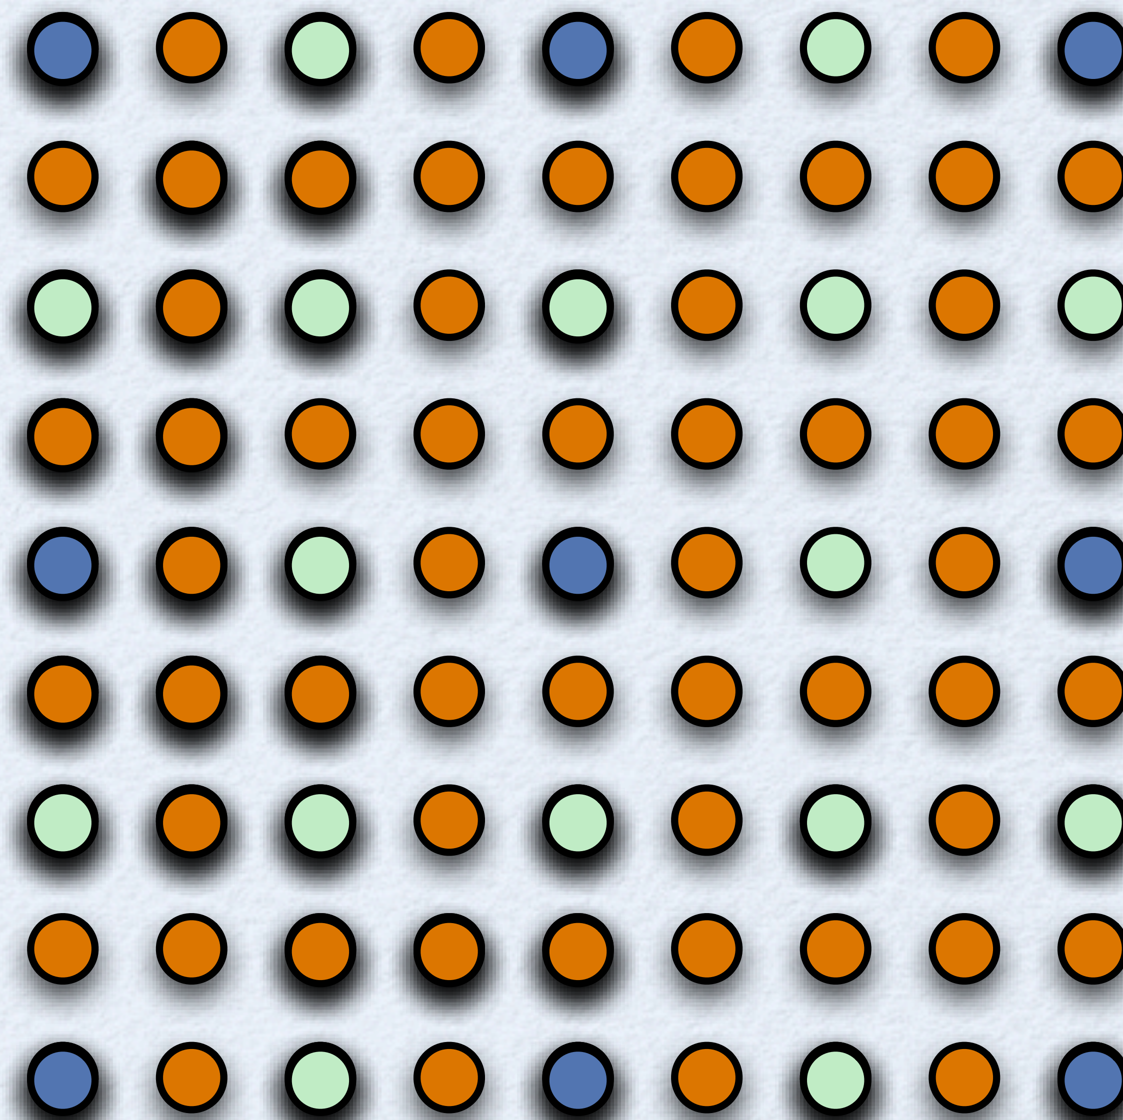

Friday, October 7, 11

Wavelet-adapted grids

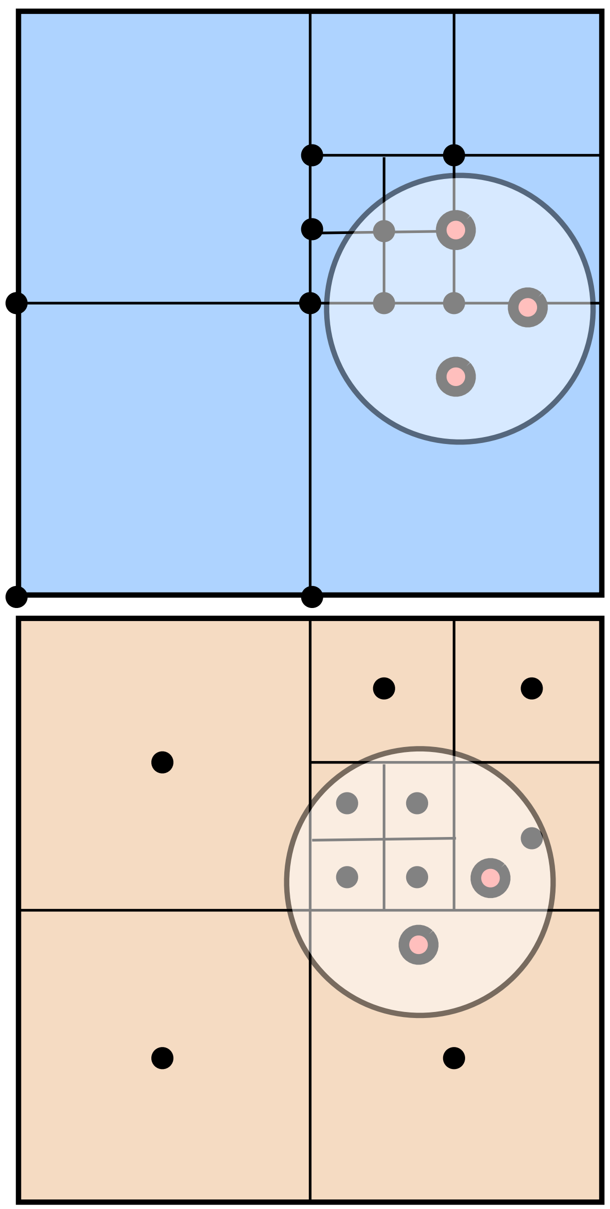

PDE:

$$\frac{\partial \phi}{\partial t} + \boldsymbol{u} \cdot \nabla \phi = 0$$

Spatial Differences = filtering operations:

$$F(c_k^l) = \sum_{j=s_f}^{e_f - 1} c_{k+j}^l \beta_j^l, \quad \beta_j^l \text{ function of } \{c_m^l\}$$

GHOSTS : **easy to compute** – (locally) uniform filtering of the grid

MULTIRESOLUTION LEVEL SETS

M. Bergdorf, P. Koumoutsakos. A Lagrangian Particle-Wavelet Method, Multiscale Modeling and Simulation: A SIAM Interdisciplinary Journal, 5(3), 980-995, 2006

Enright, Fedkiw et al, 2002
dof = # grid points + aux. particles at t=0.0

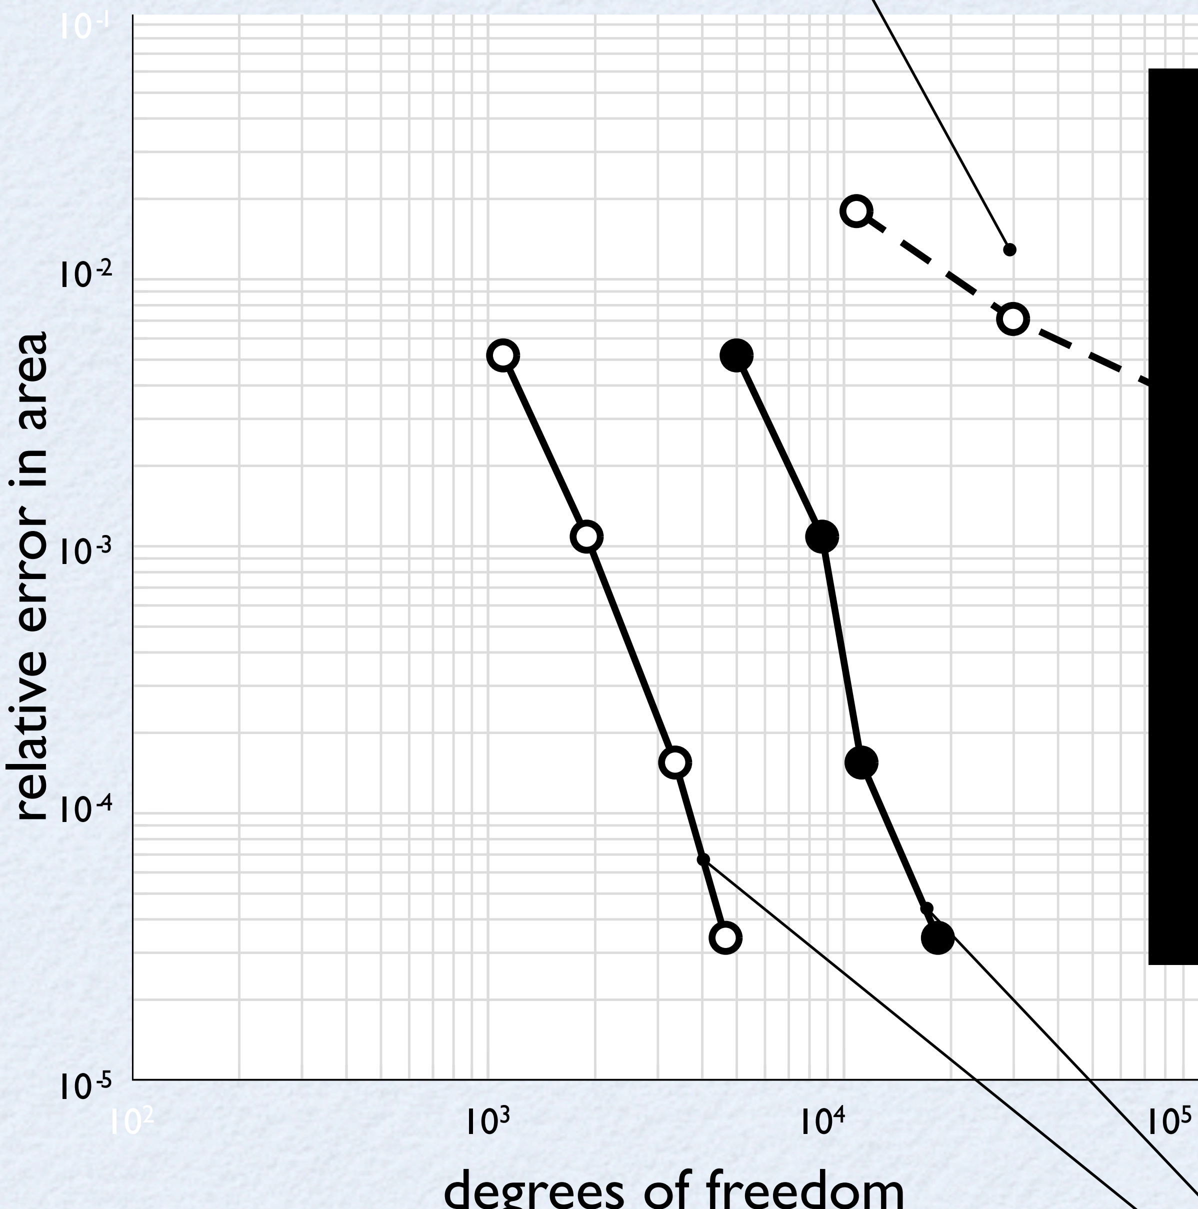

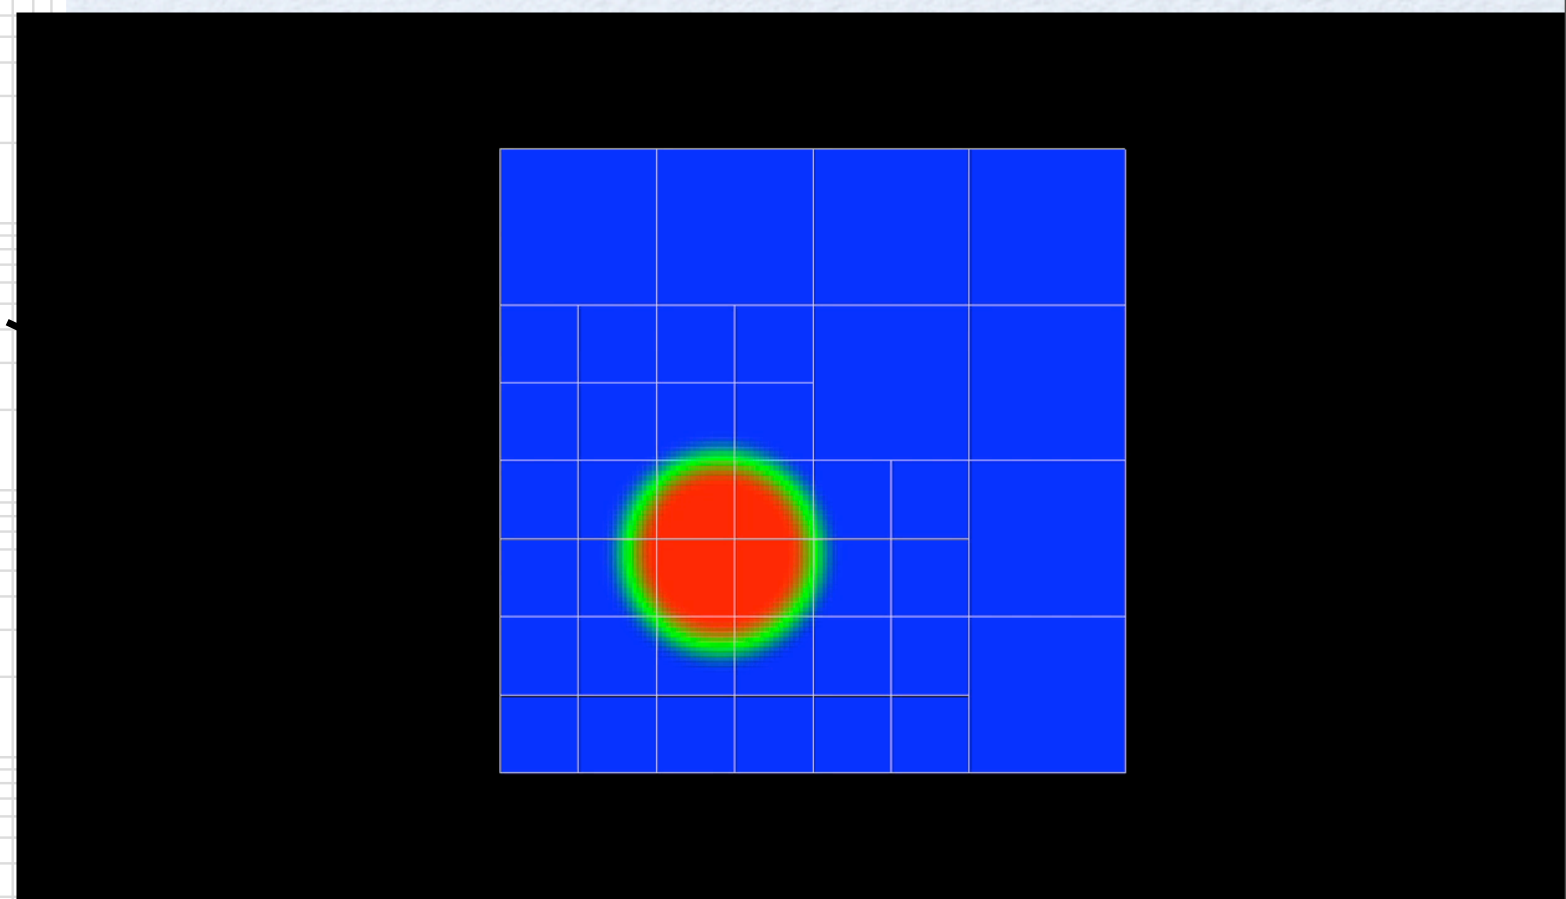

Present Method
dof = # active gp/particles at t=0.0
dof = # active gp/particles at final time

$\mathrm{CFL}_{max} \sim 40$

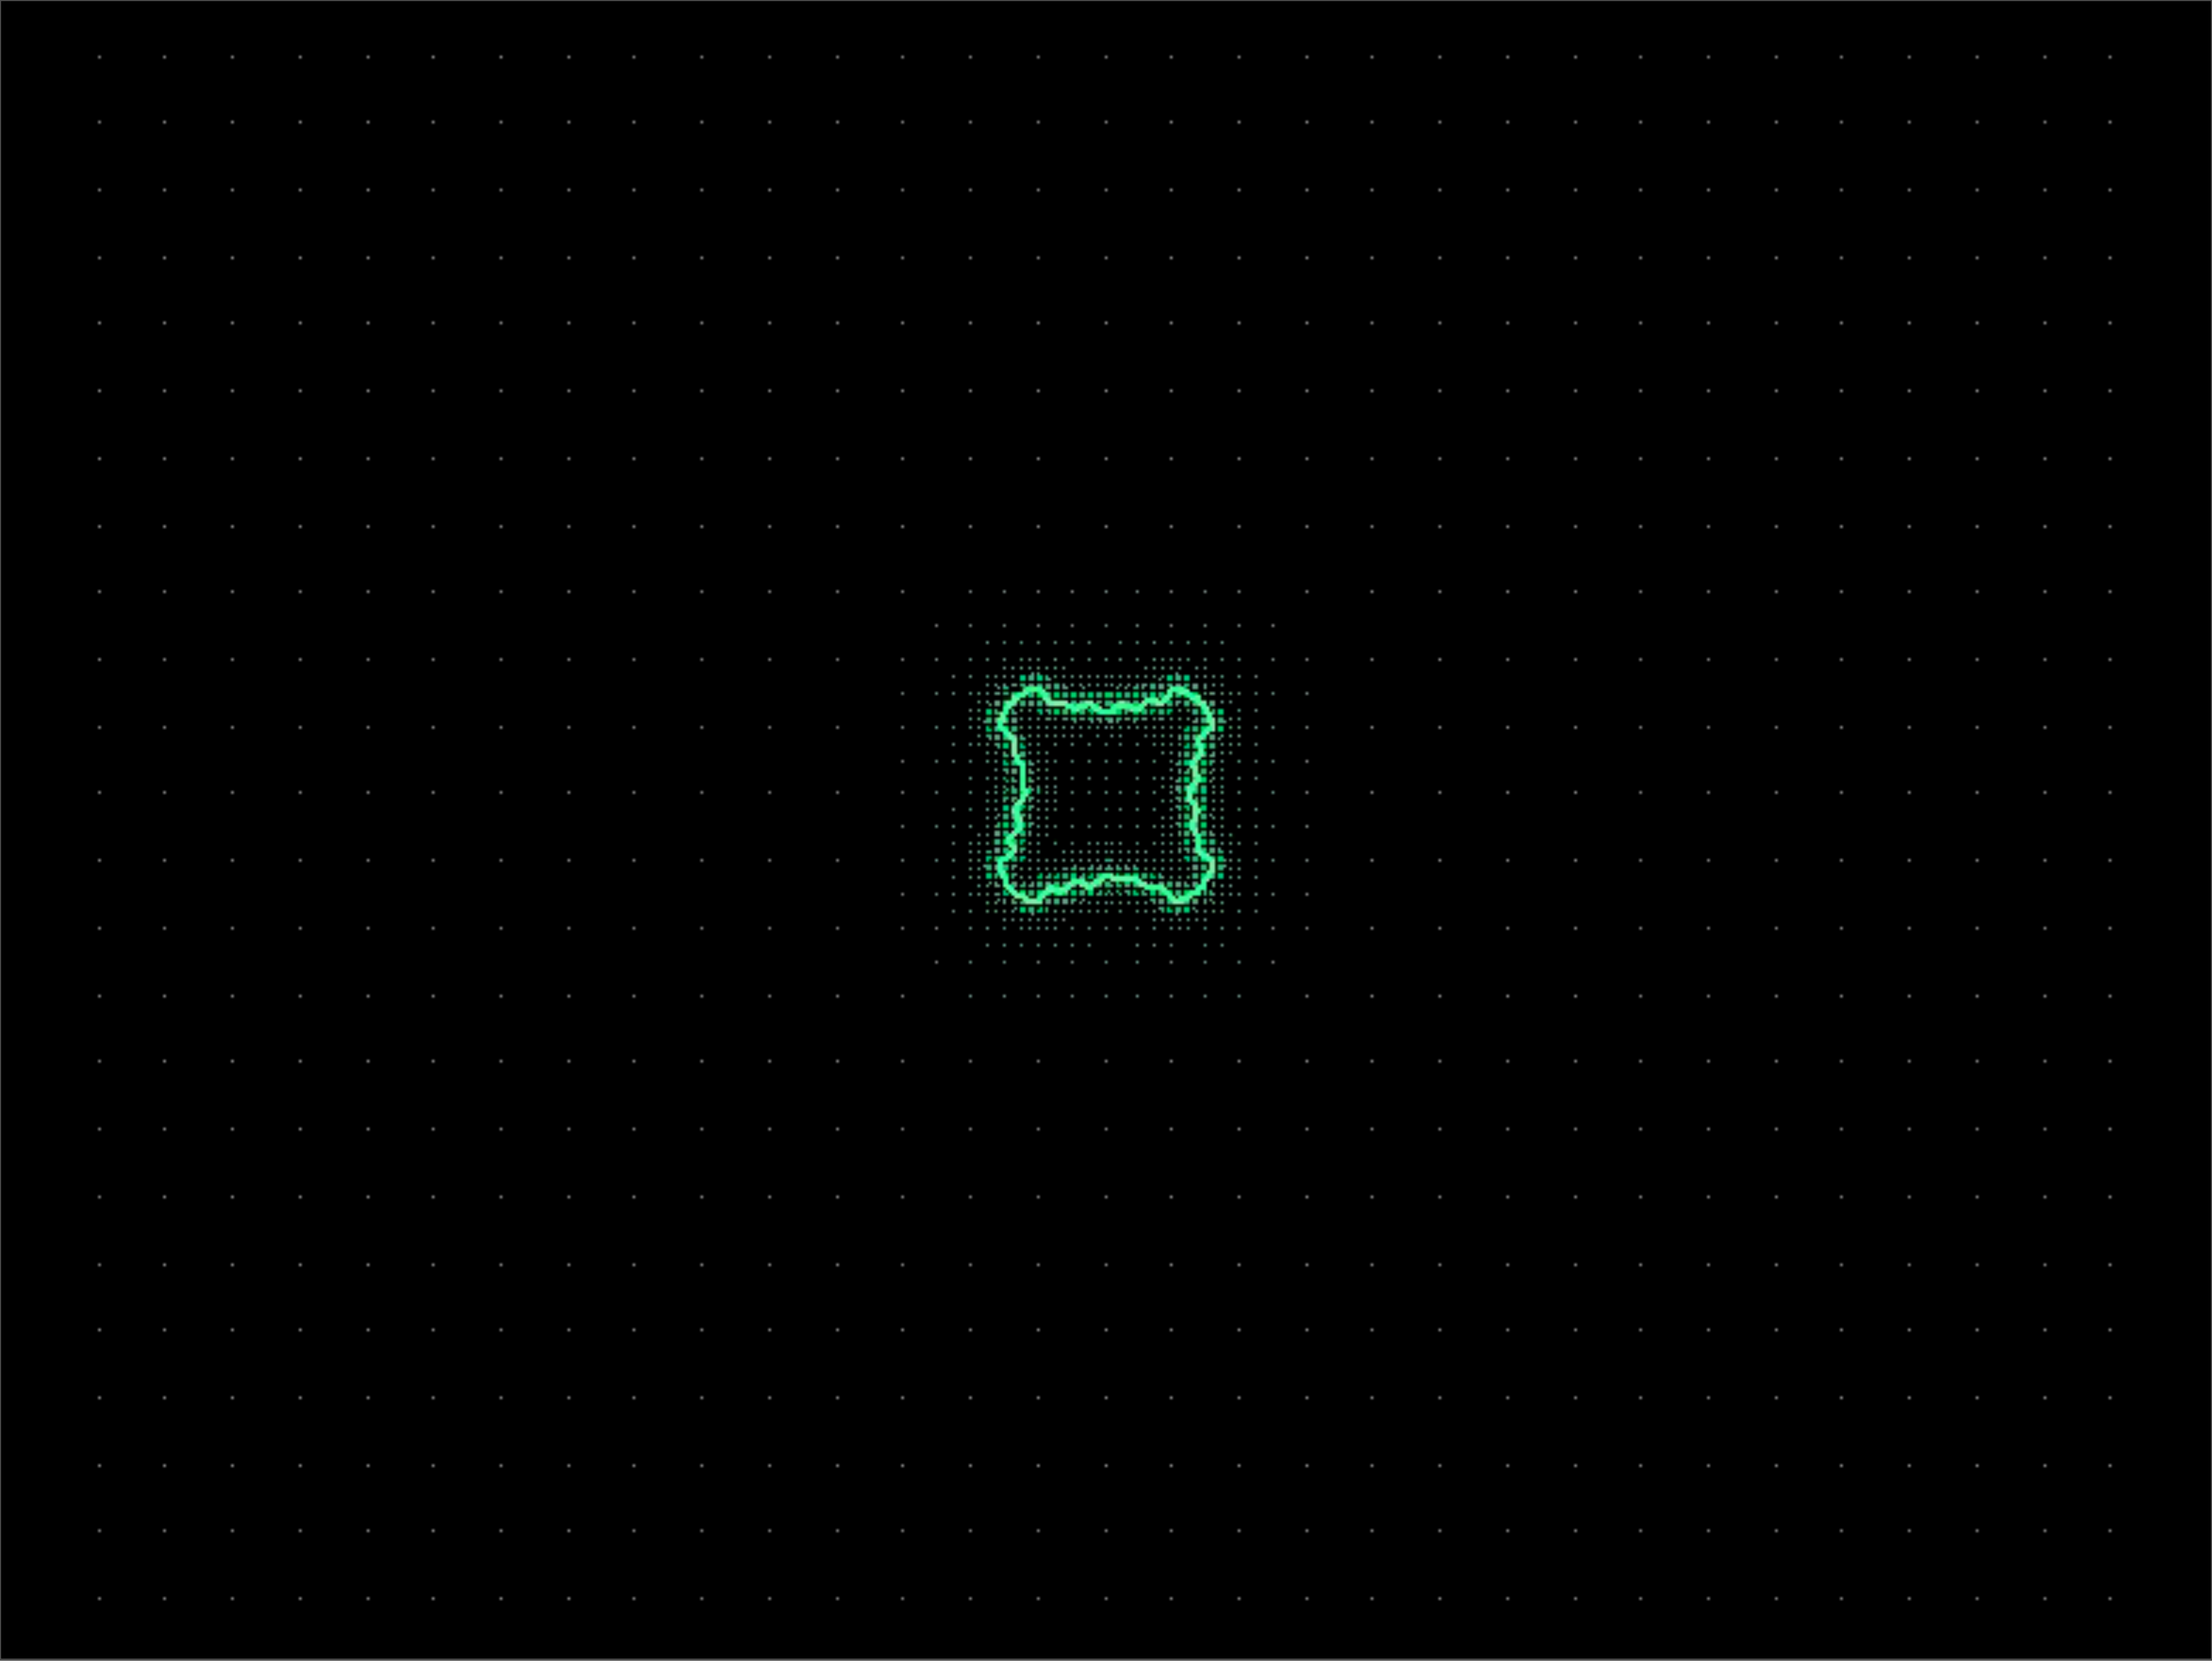

Shock Bubble Interaction

(M=3, At=0.8)

FINEST RESOLUTION EQUIVALENT
8000 x 8000 uniform grid
~40 times smaller adaptive

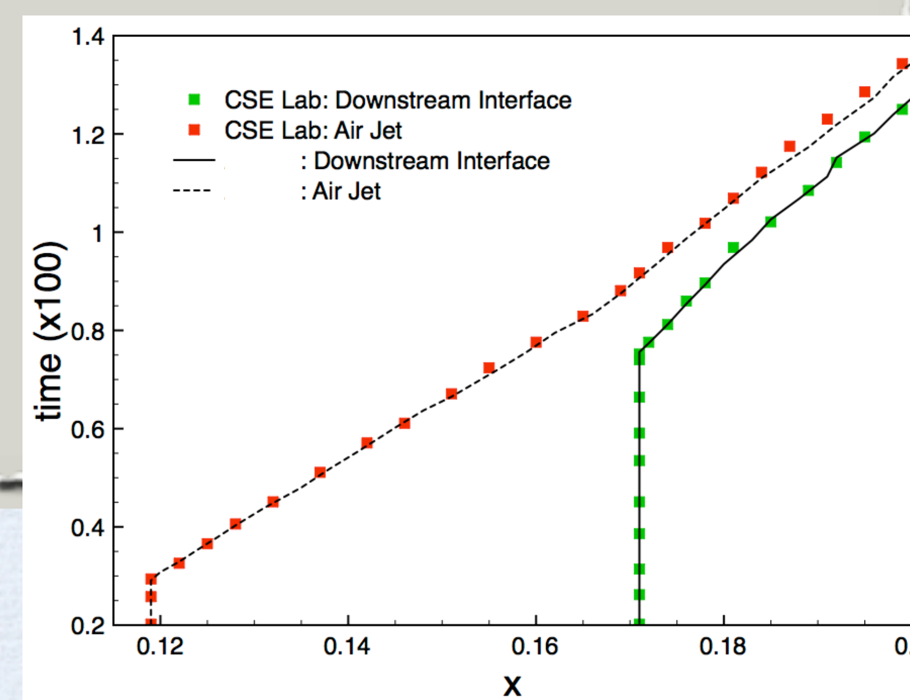

Hejazialhosseini B., Rossinelli D., Bergdorf M.,Koumoutsakos P., High order Finite Volume methods on Wavelet-adapted Grids with Local Time-Stepping on Multicore Architectures for the Simulation of Shock-Bubble Interactions, **J. of Comp. Physics**, 2010

Block Grid for Multi/Many-core:

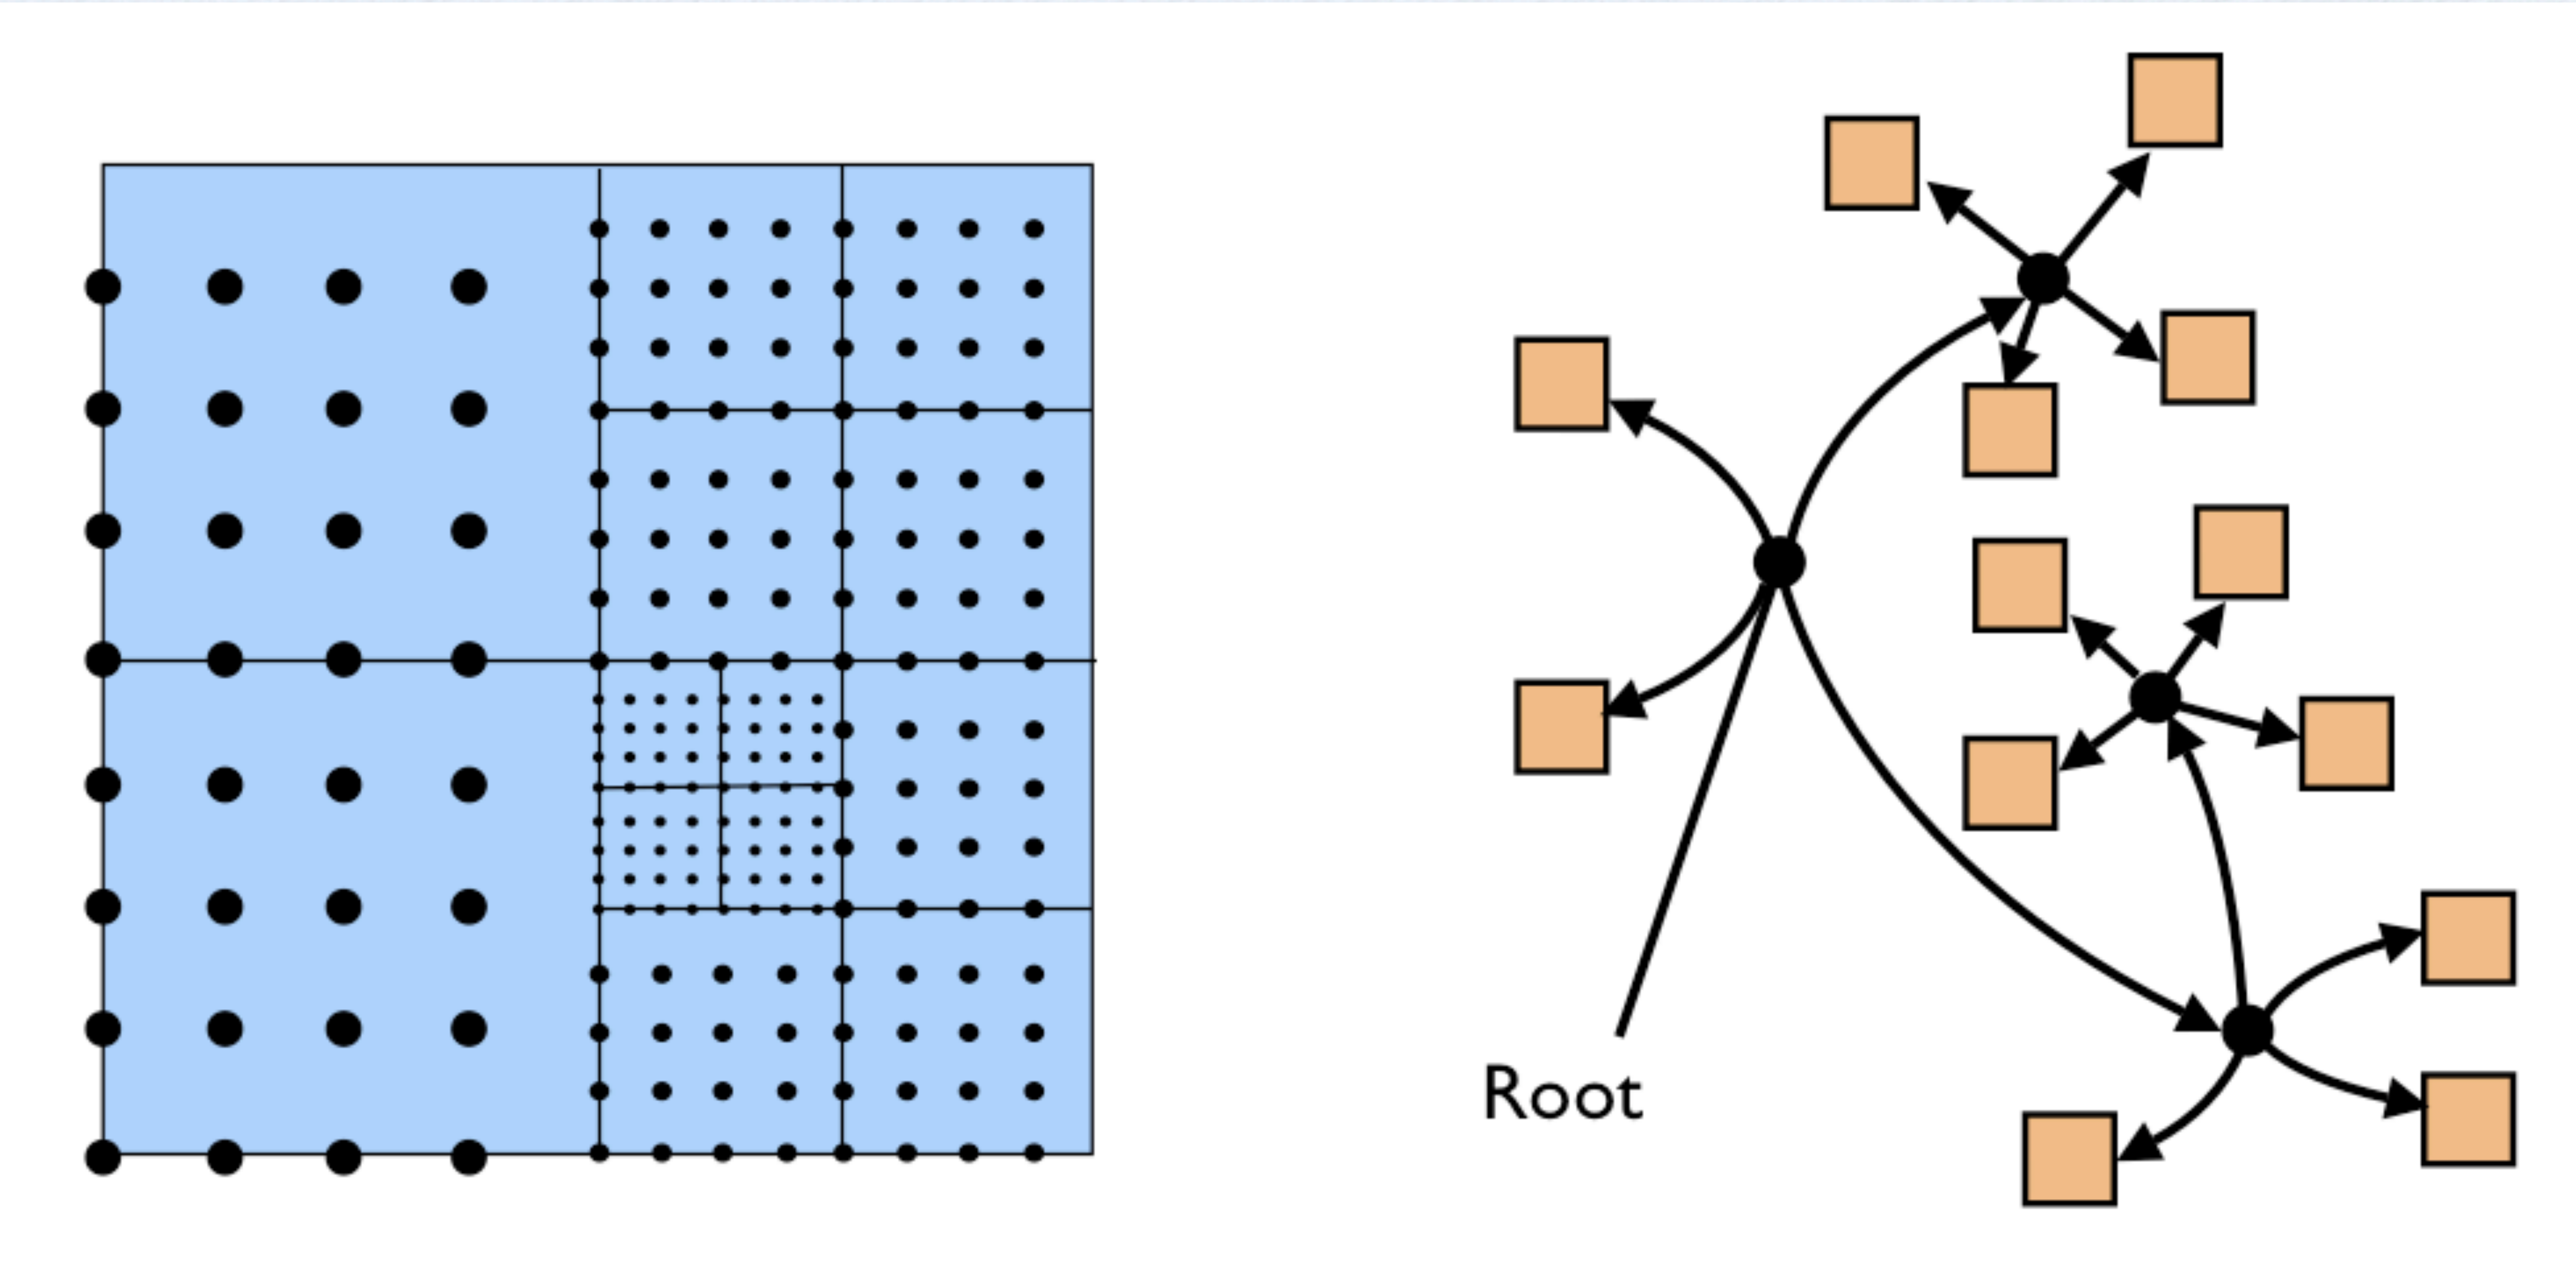

- Neighbors look-up: less memory indirections
- Less #ghosts
- Within a block: random access

Multiresolution + MultiCore + GPU

- MULTIPLE TASKS

  1. task parallel,ghost computing -> **multi-core**
  2. fine-grained data parallelism for RHS -> **GPUs**
  3. Integration step -> **multi-core**

$$\mathbf{q}^{\mathbf{new}} = \mathbf{q}^{\mathbf{old}} + \delta t \mathbf{F}\left(\mathbf{q}^{\mathbf{old}}, \nabla \mathbf{q}^{\mathbf{old}}\right)$$

CUDA/OpenCL

- **How much** ***faster*** than CPU-only execution?
- **How much** ***different*** are CPU/GPU and CPU-only solutions?

Wavelet Blocks on GPUs

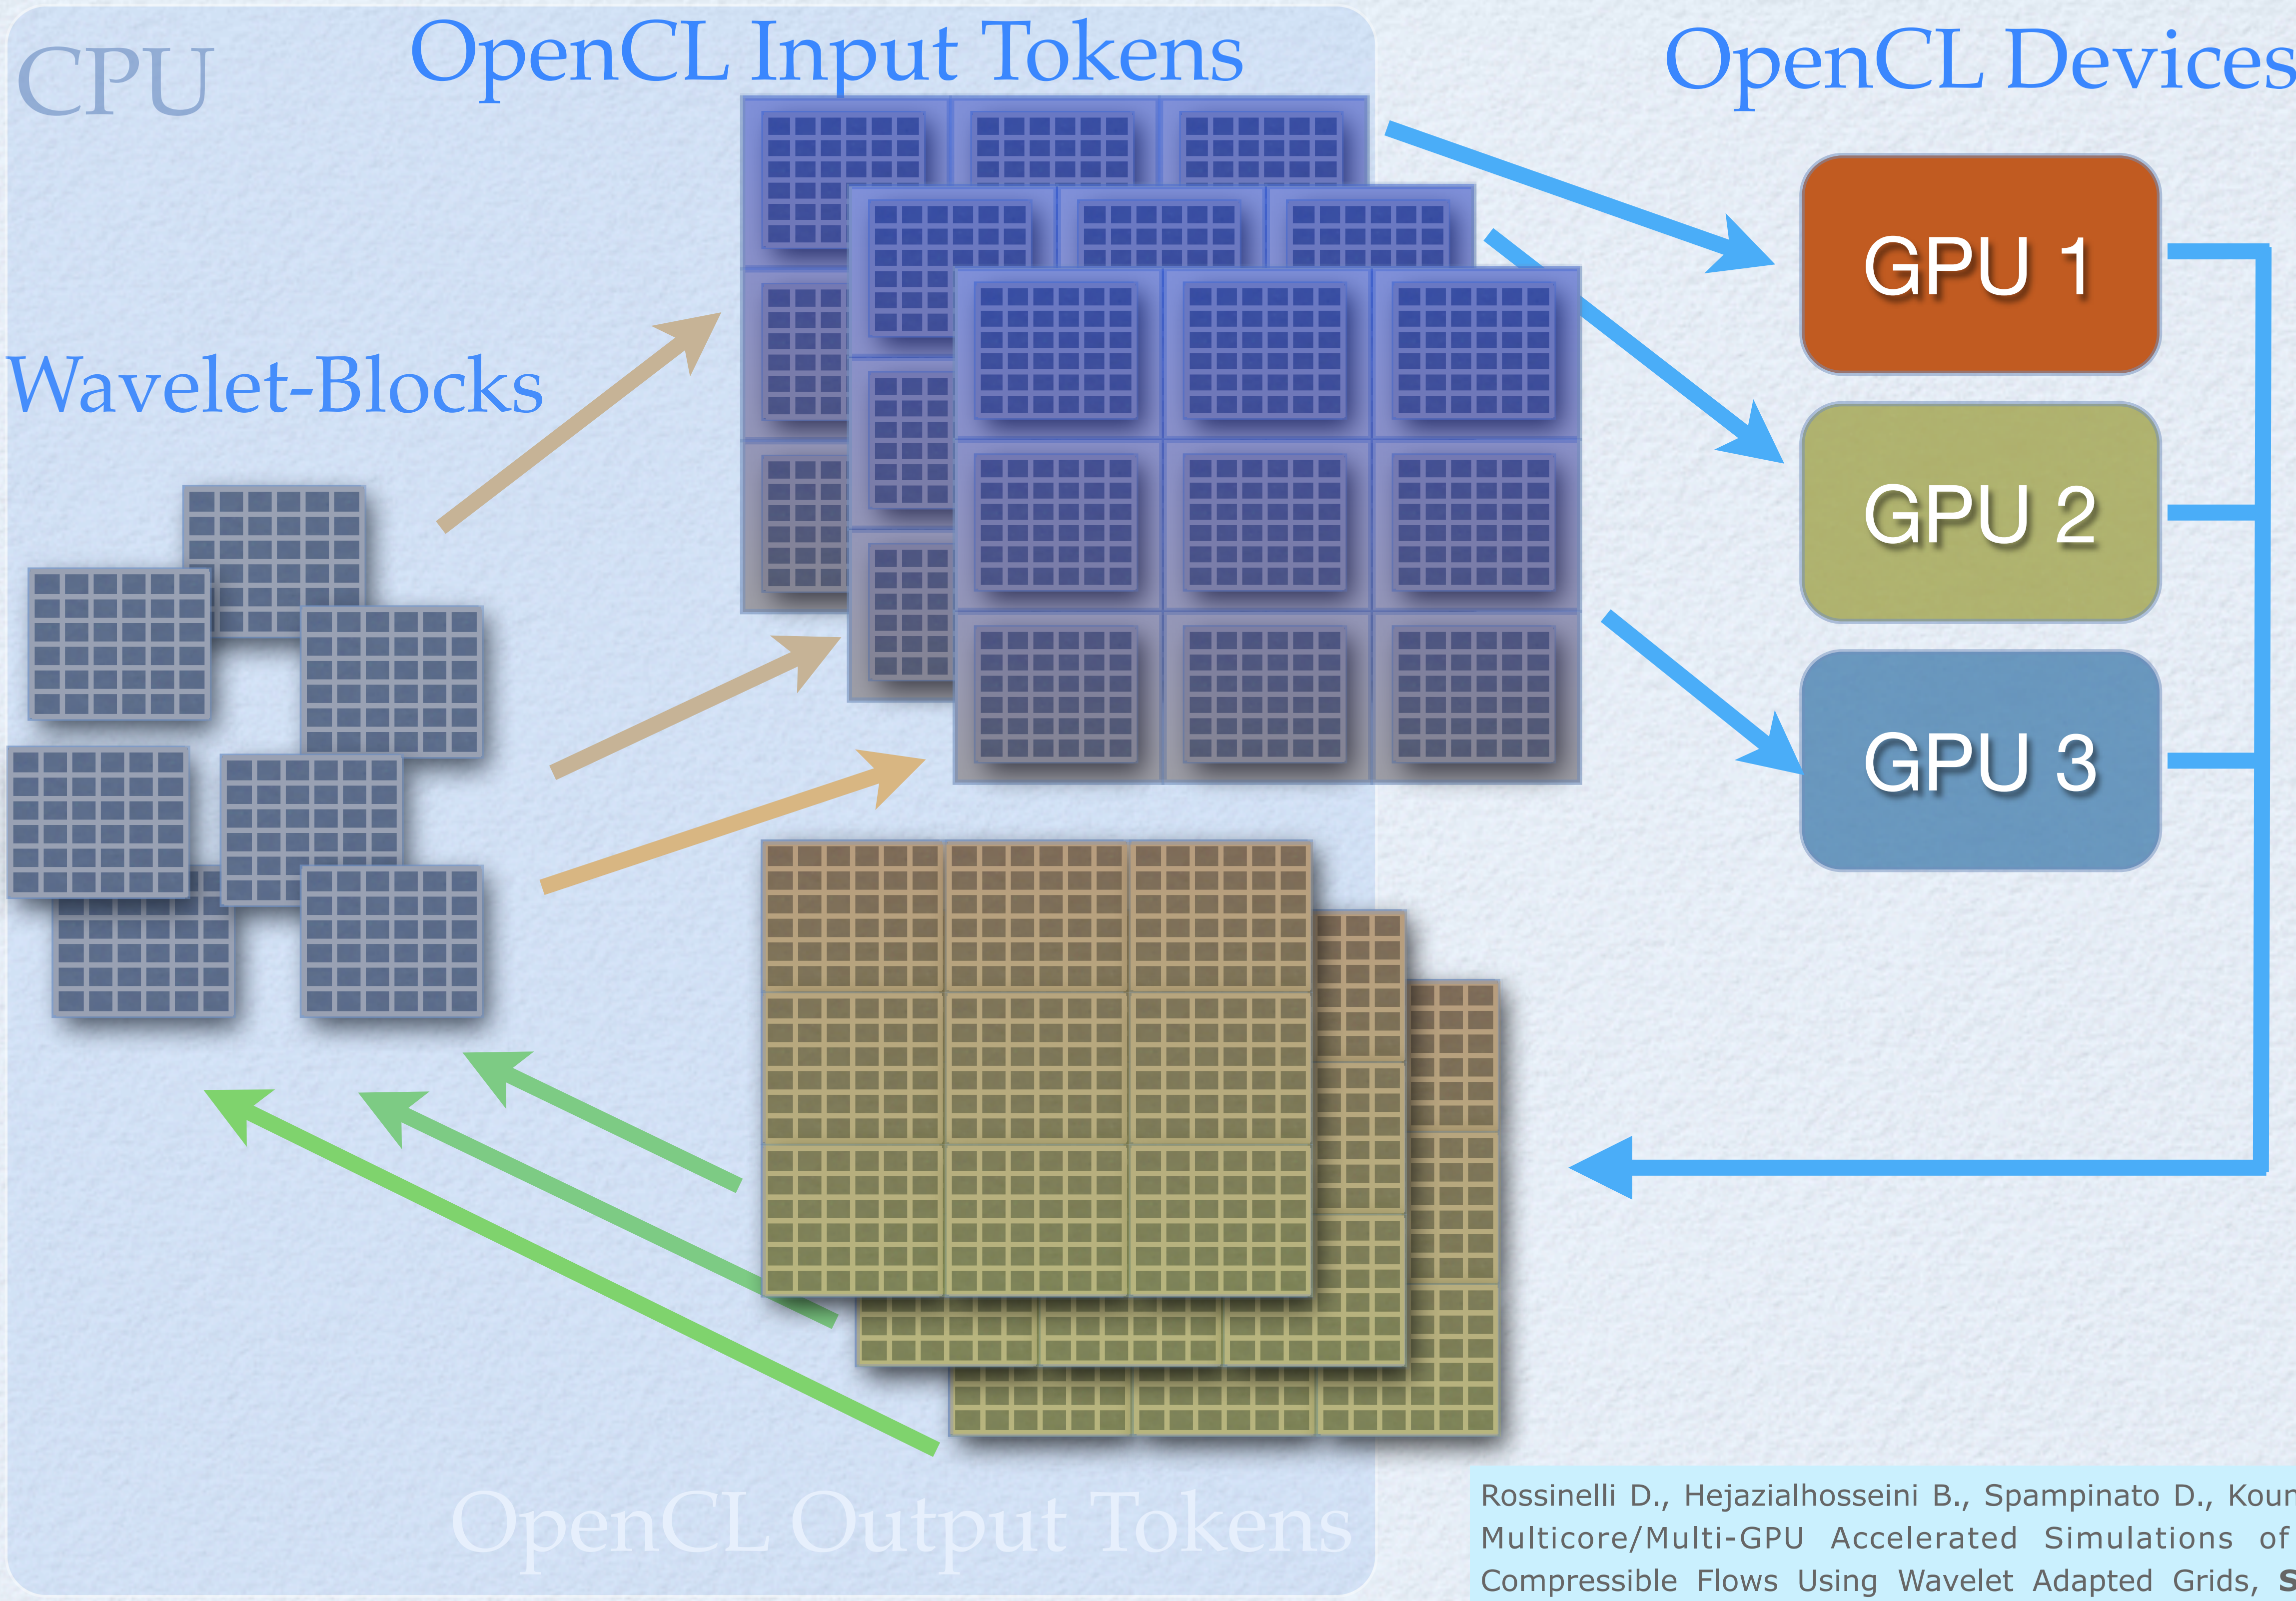

Rossinelli D., Hejazialhosseini B., Spampinato D., Koumoutsakos P., Multicore/Multi-GPU Accelerated Simulations of Multiphase Compressible Flows Using Wavelet Adapted Grids, **SIAM J. Sci. Comput.,** 33, pp. 512-540, 2011

Multiple kernels for the GPU

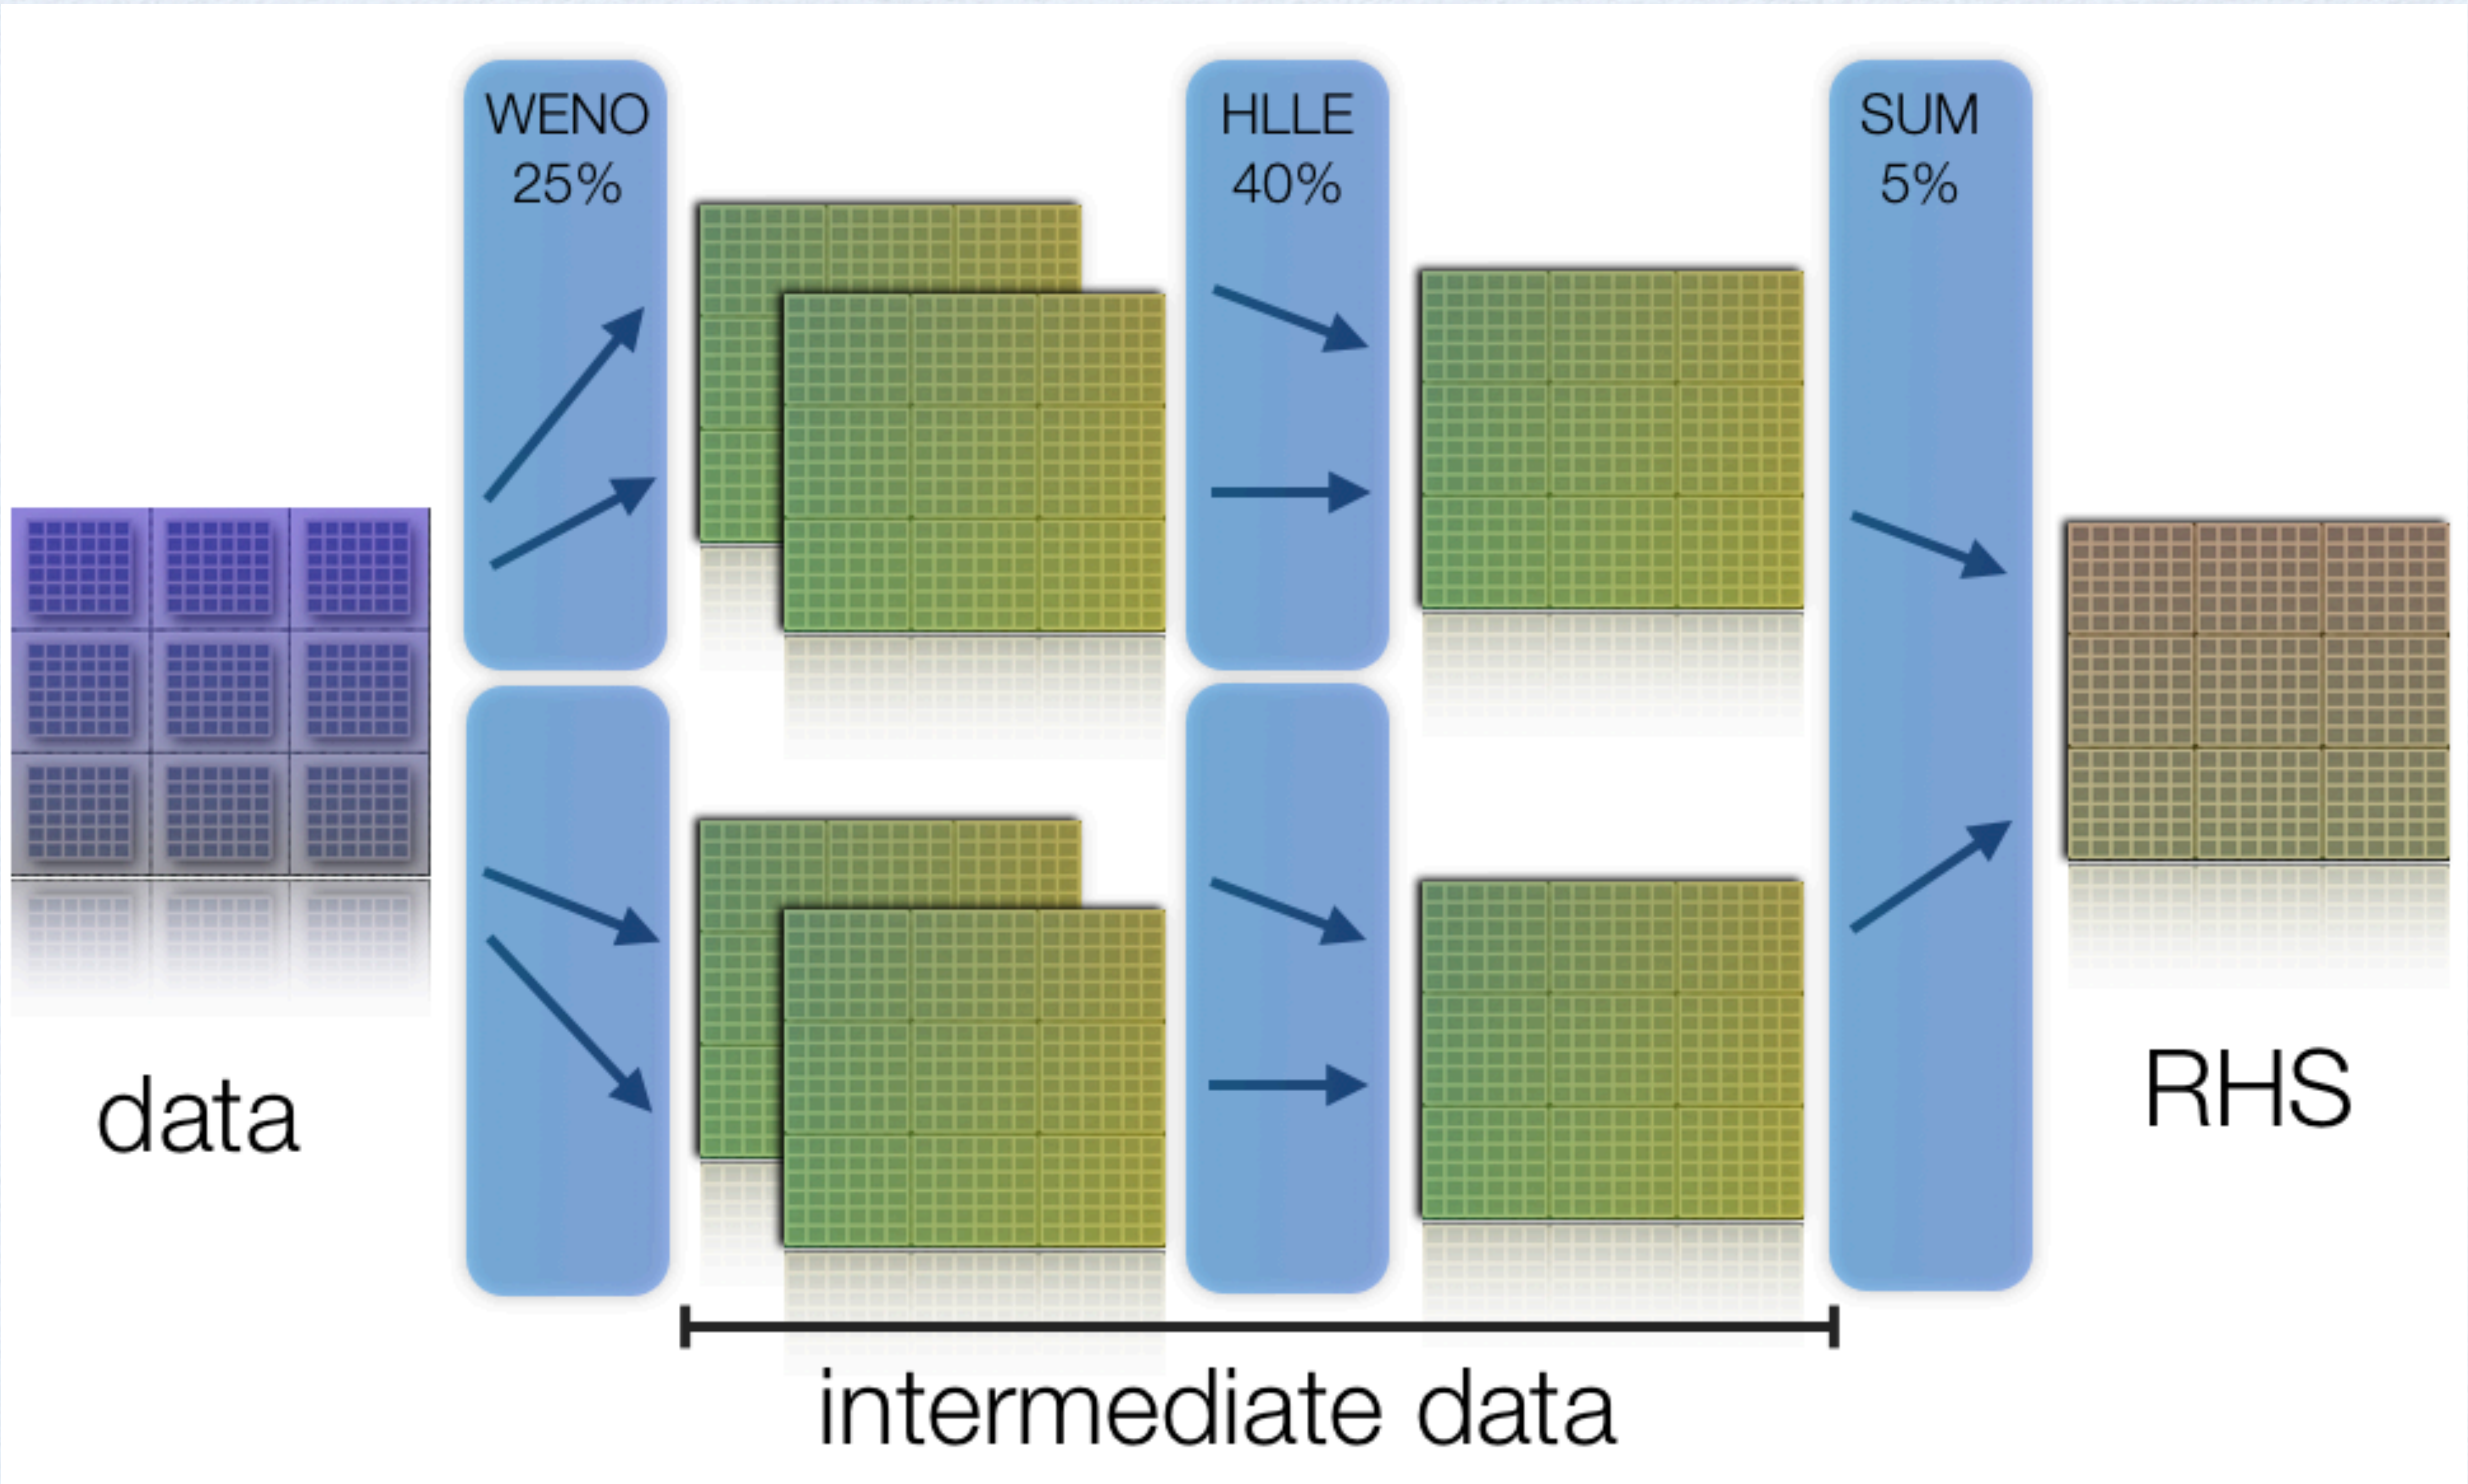

Performance I : Strong Scaling

Strong scaling (effective 8000^2 - actual 40x less) vs. #GPUs, #CPU cores

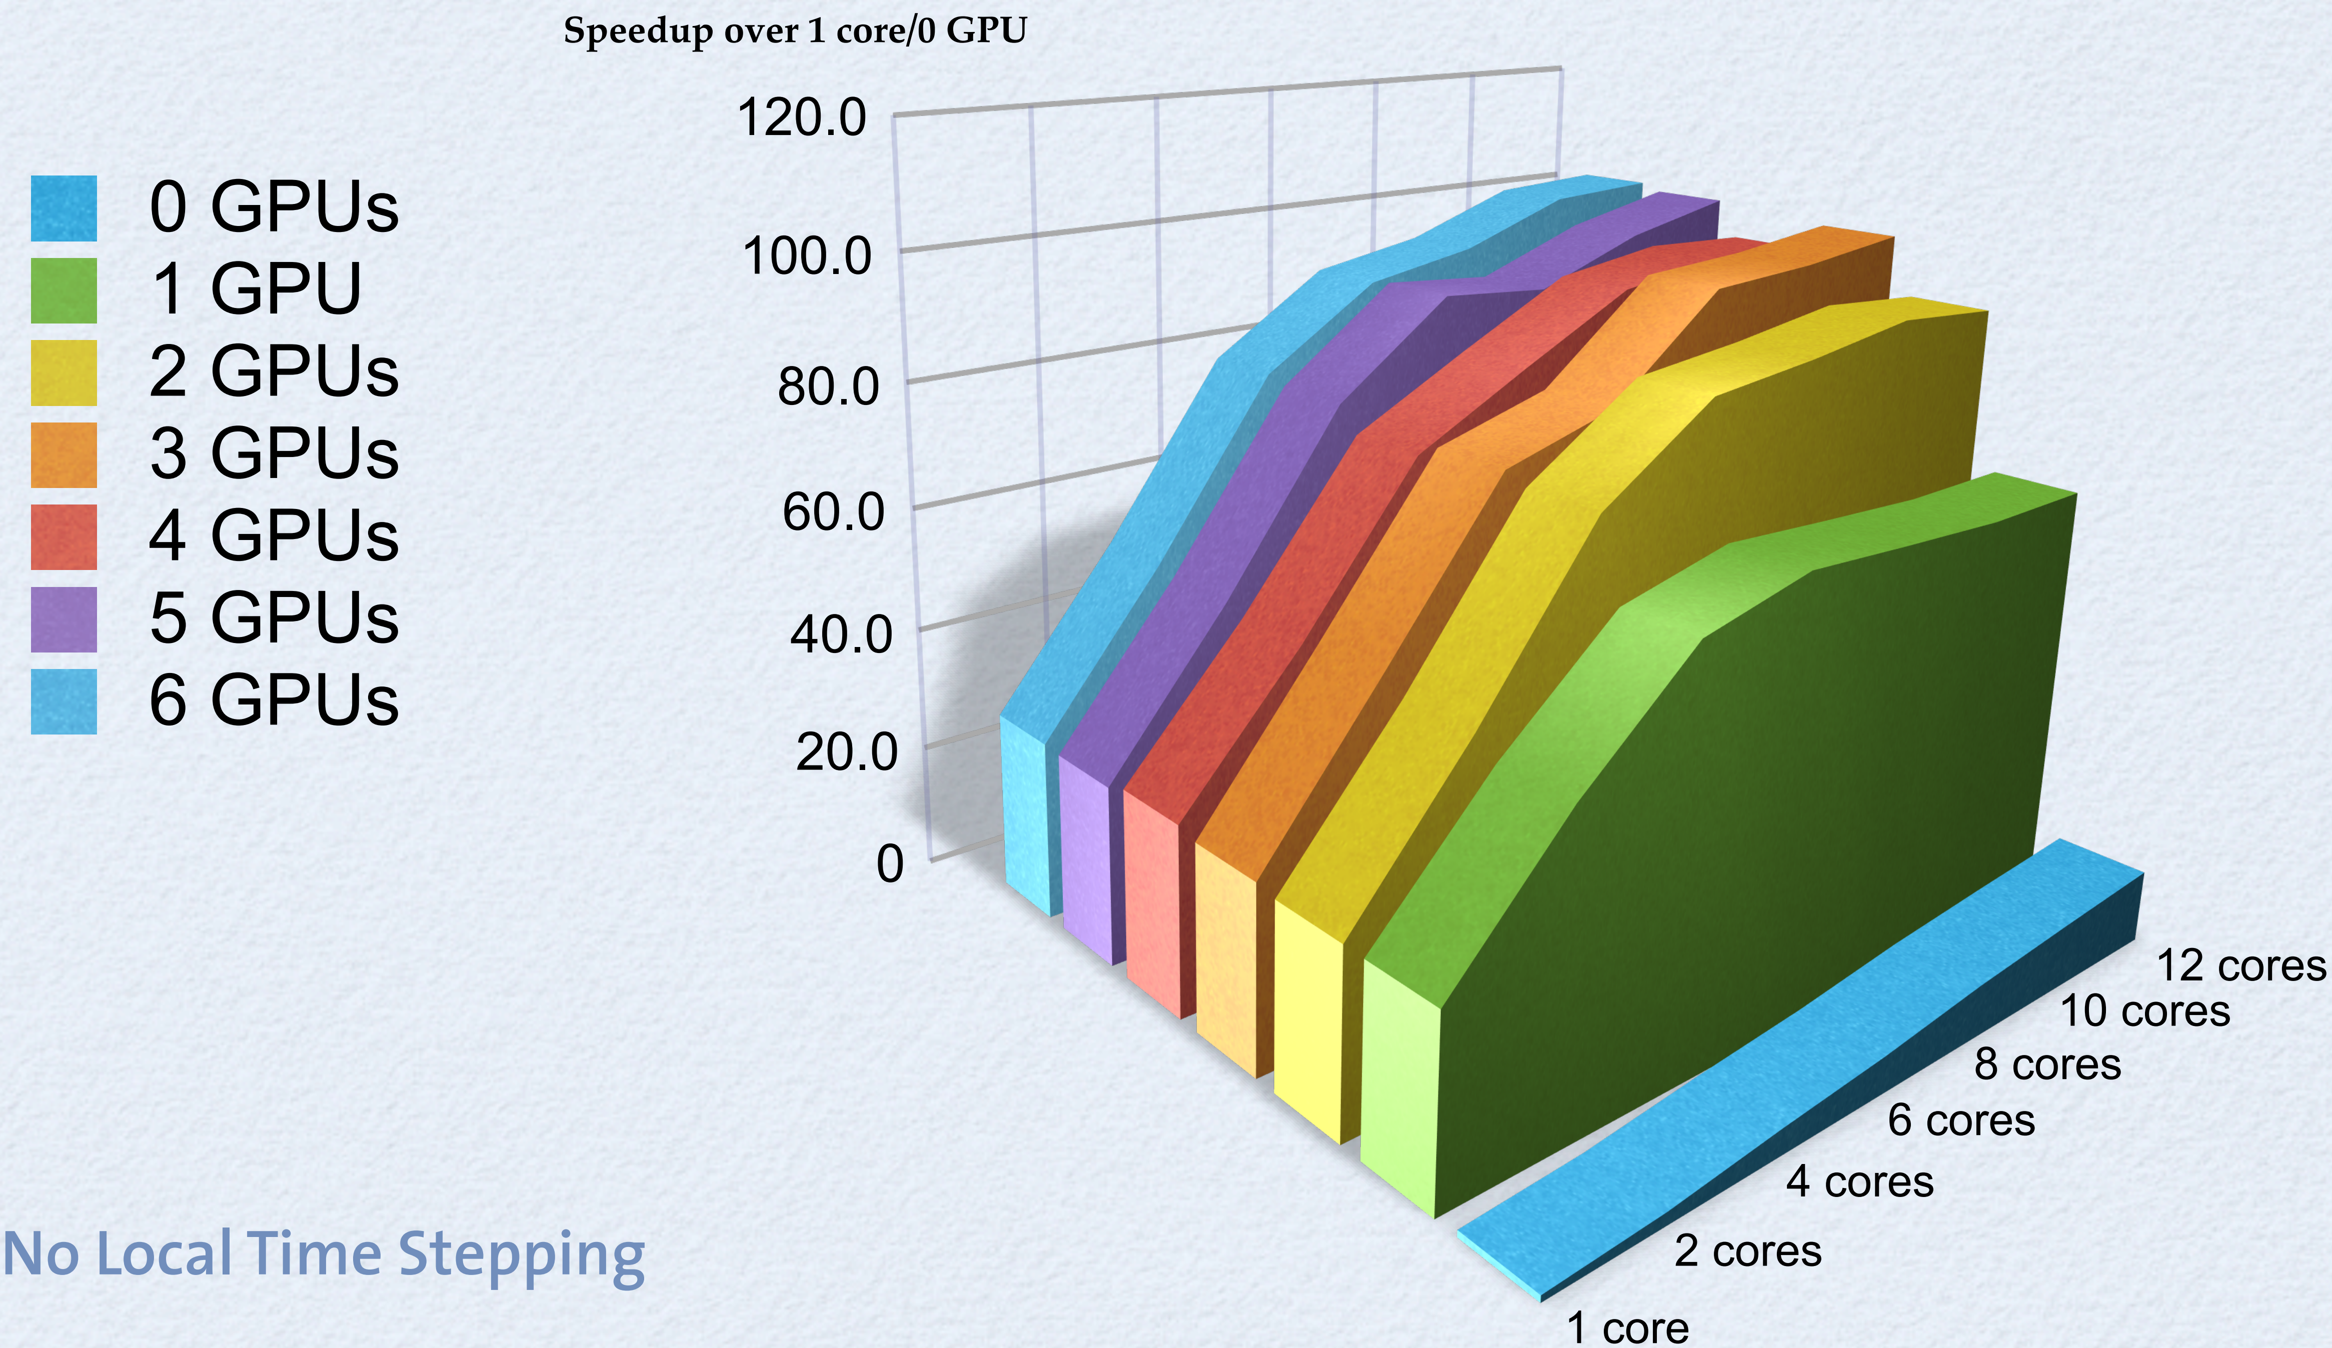

No Local Time Stepping

Performance II : Time to Solution

Compared to a space adaptive, single-threaded solver:

- Algorithms : Local Time Stepping: **24X**
- Ghost Reconstruction : CPU optimization (vectorization): **1.8X**
- Ghost Reconstruction : Task-based parallelism (via TBB): **8X** (over 12)
- GPUs as accelerators: **3X**

▶ Overall Reduction in time to solution: **~ 1000**

A comparison of CHOMBO vs MRAG

shock-bubble interaction

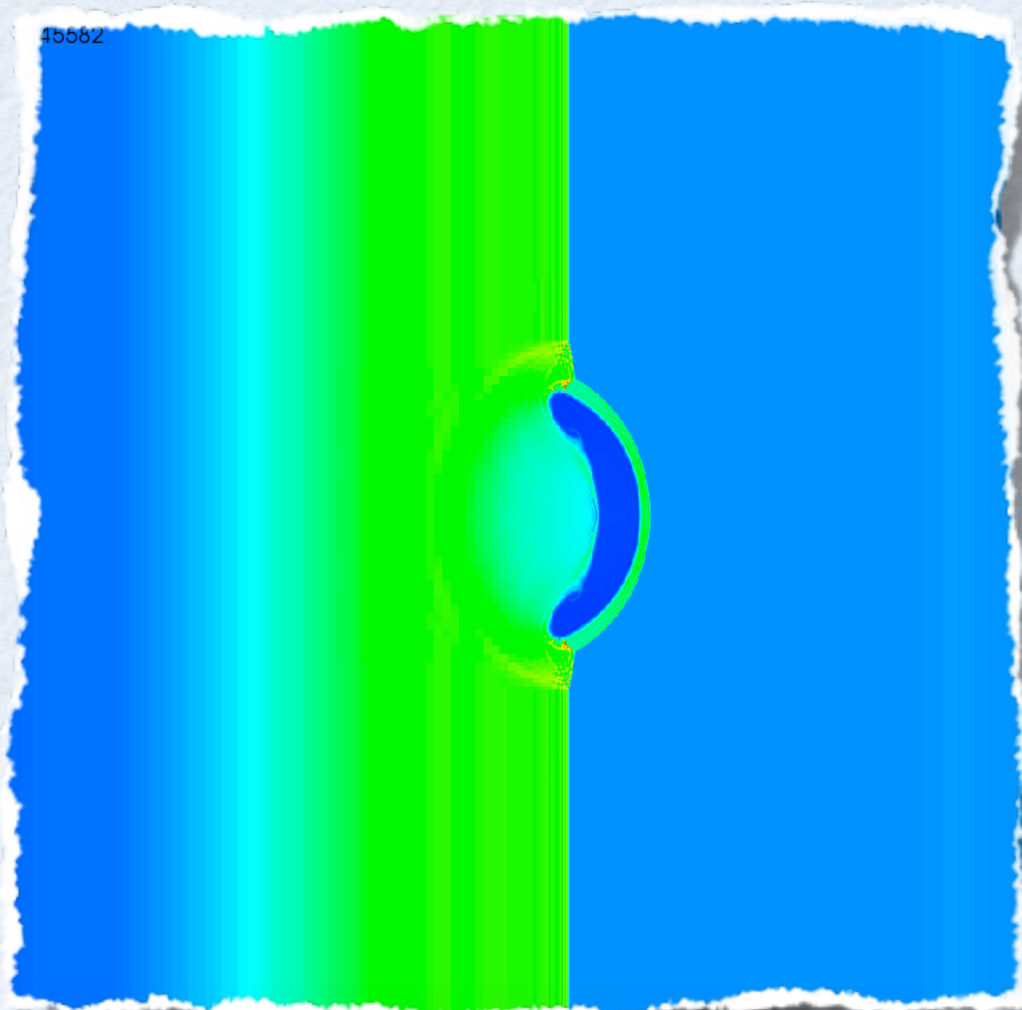

density field

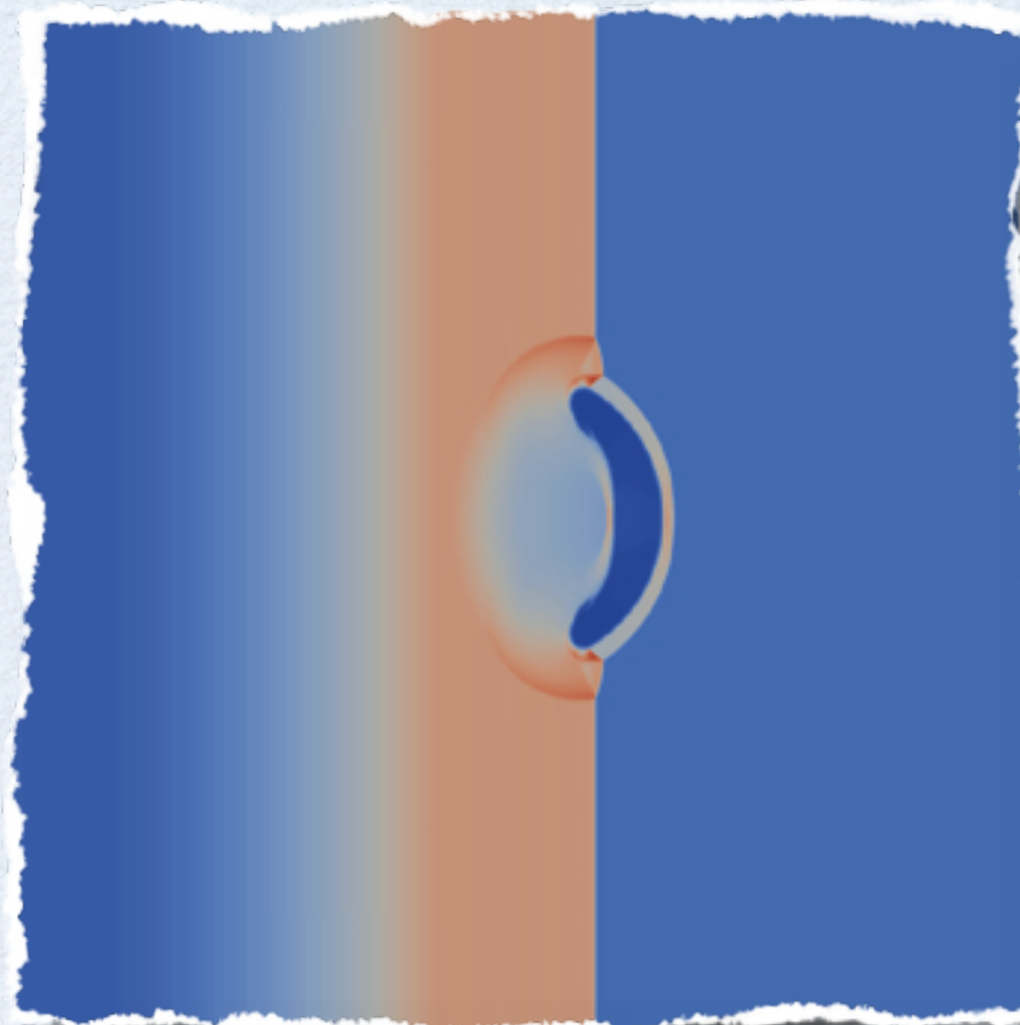

Chombo:
91 min, 230 MB

single-phase
2nd order PPM

MRAG (home grown, swiss quality stuff)**:**
56 min, 244 MB (+ 1 GPU: 7 min)

multi-phase
5th order WENO scheme

Rossinelli D., Hejazialhosseini B., Spampinato D., Koumoutsakos P., Multicore/Multi-GPU Accelerated Simulations of Multiphase Compressible Flows Using Wavelet Adapted Grids, **SIAM J. Sci. Comput.**, 33, pp. 512-540, 2011


BOUNDARIES + ALGORITHMS

Boundary Conditions = Coupling

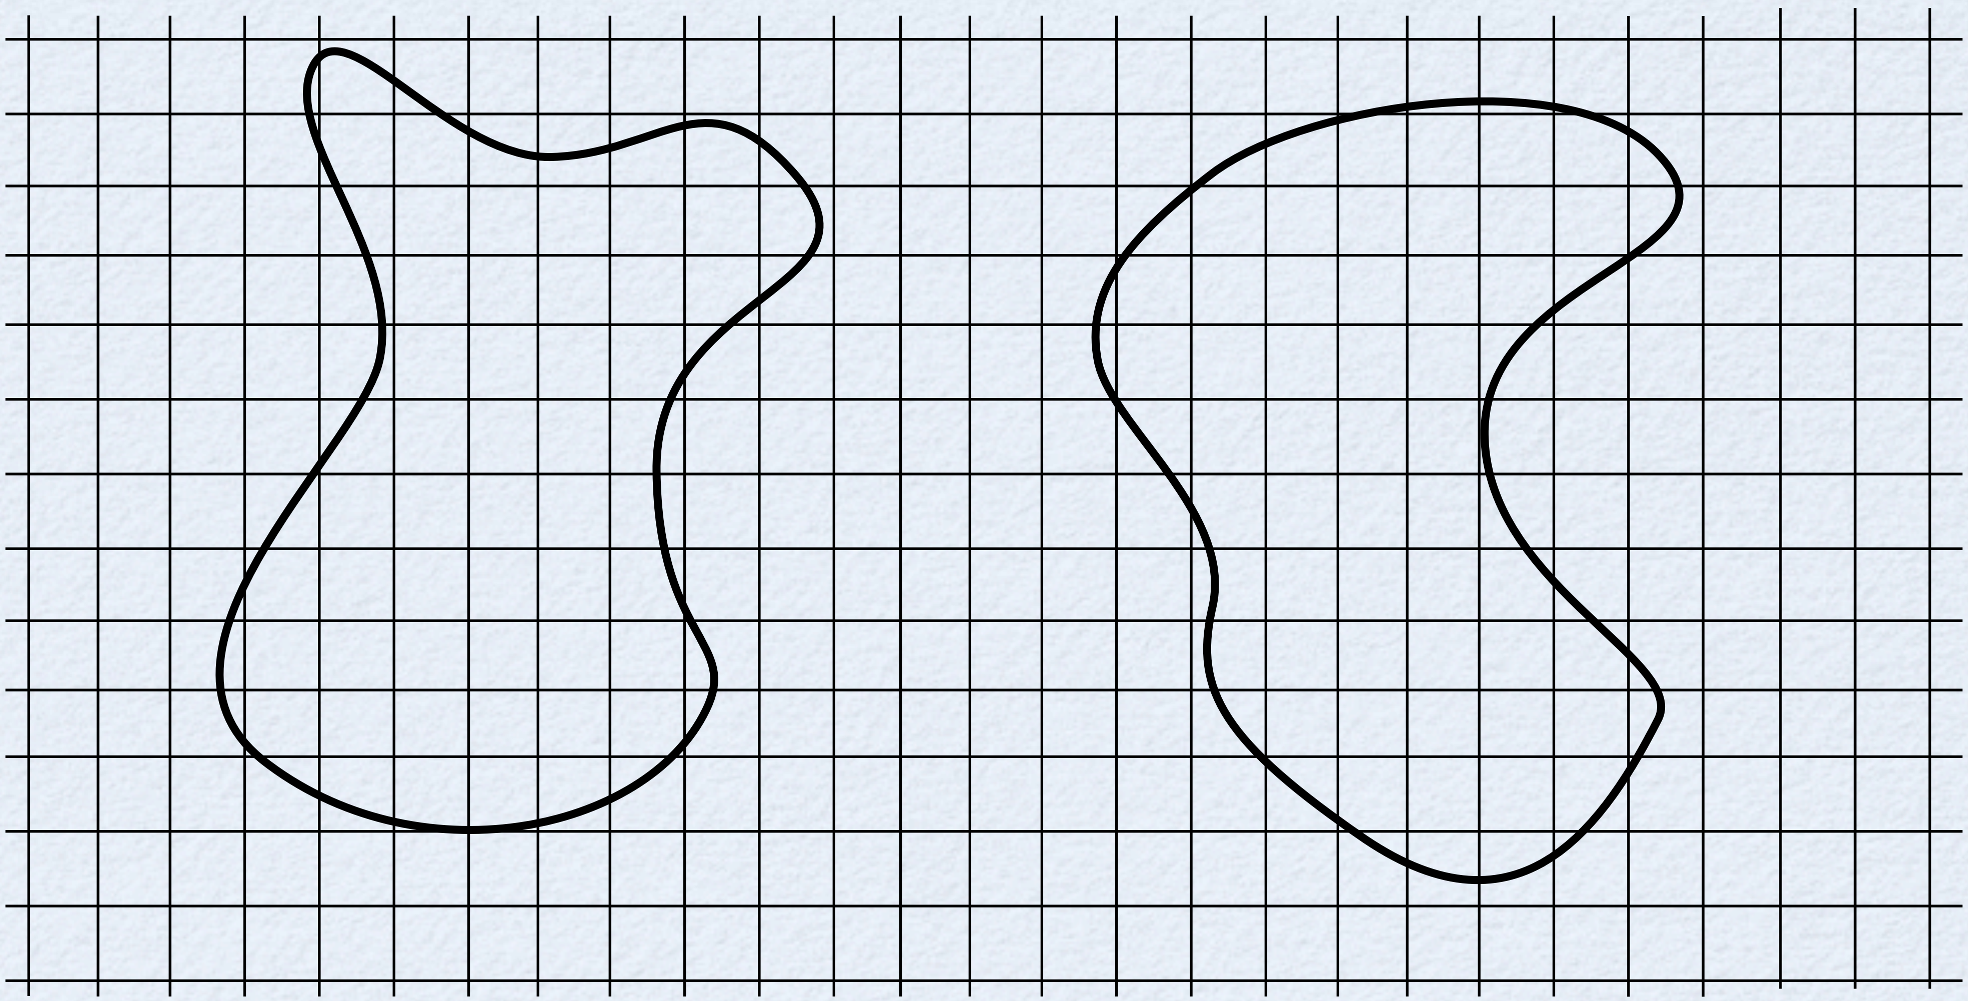

$$\rho \frac{D\mathbf{u}}{Dt} = \nabla \cdot \boldsymbol{\sigma} \; + f(\text{enforces b.c.})$$

Boundary Conditions = Coupling

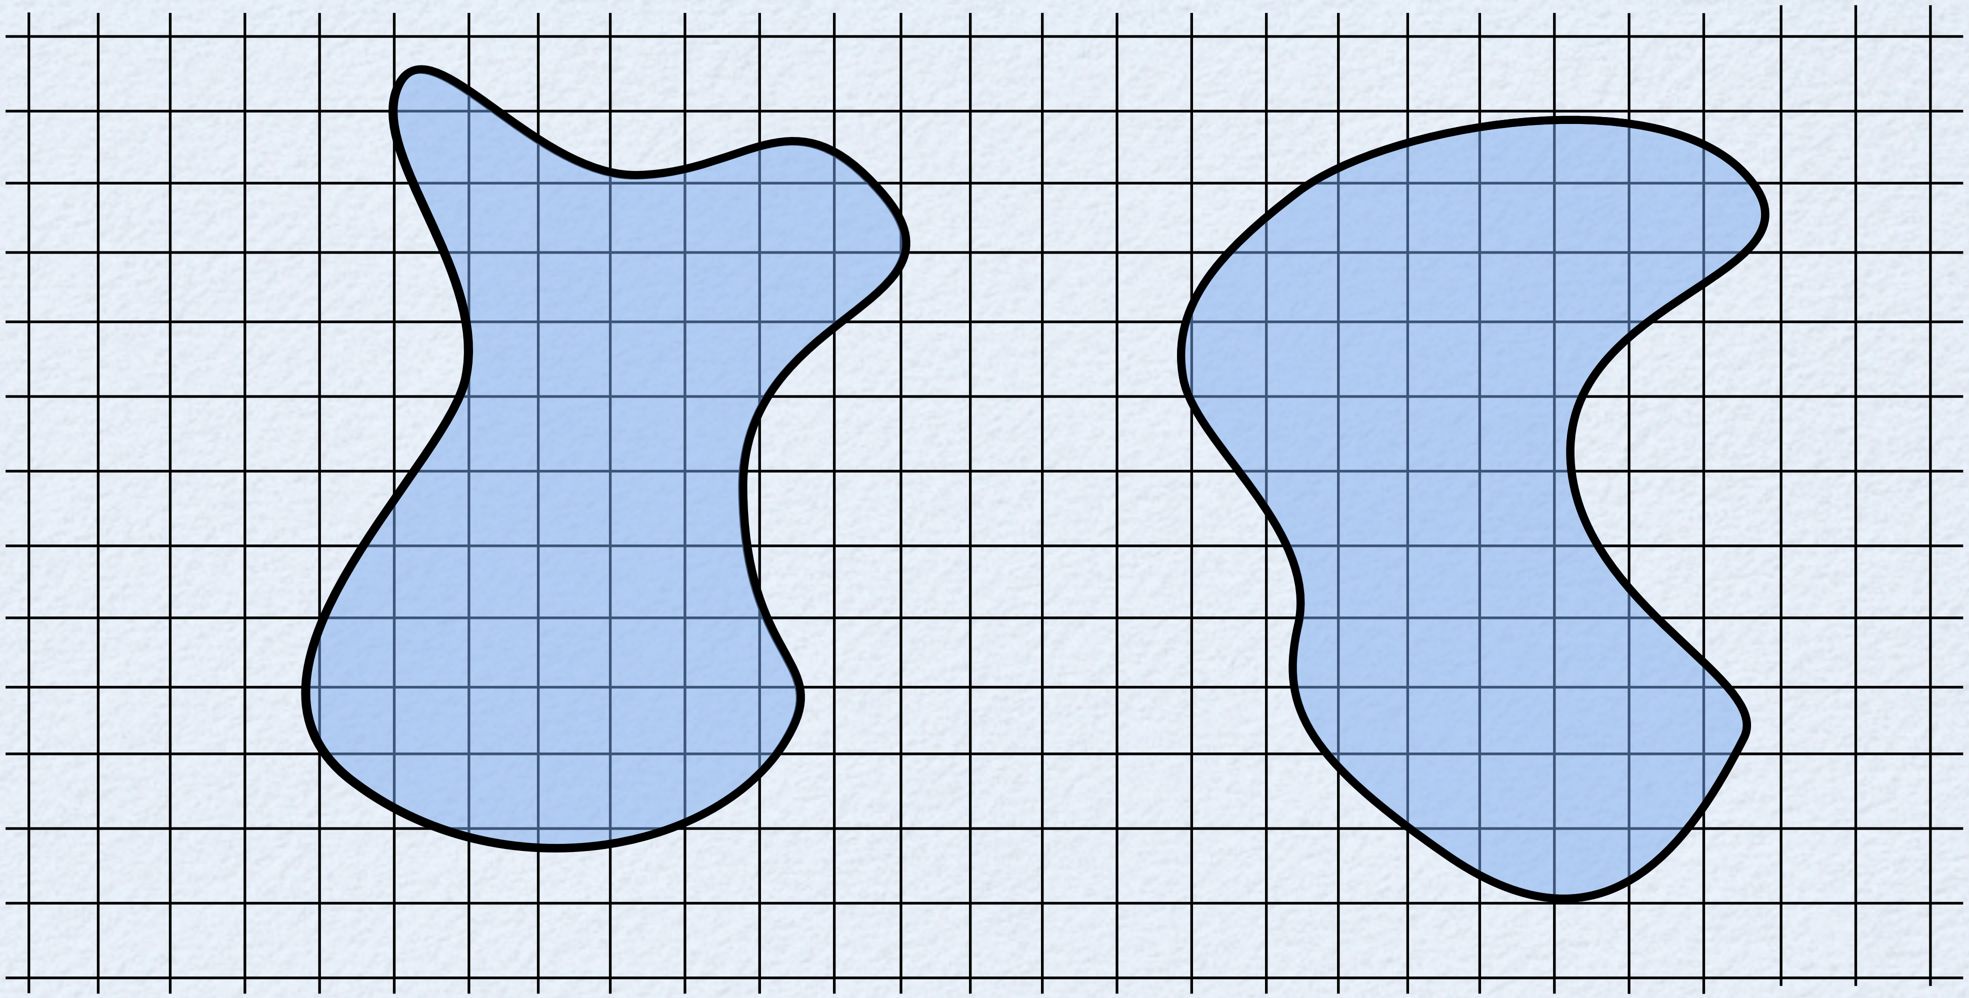

$$\rho \frac{D\mathbf{u}}{Dt} = \nabla \cdot \boldsymbol{\sigma} + f(\text{enforces b.c.})$$

Penalization Method: $f(\mathbf{x}) = \lambda \chi_S (\mathbf{u}_S - \mathbf{u})$

Boundary Conditions = Coupling

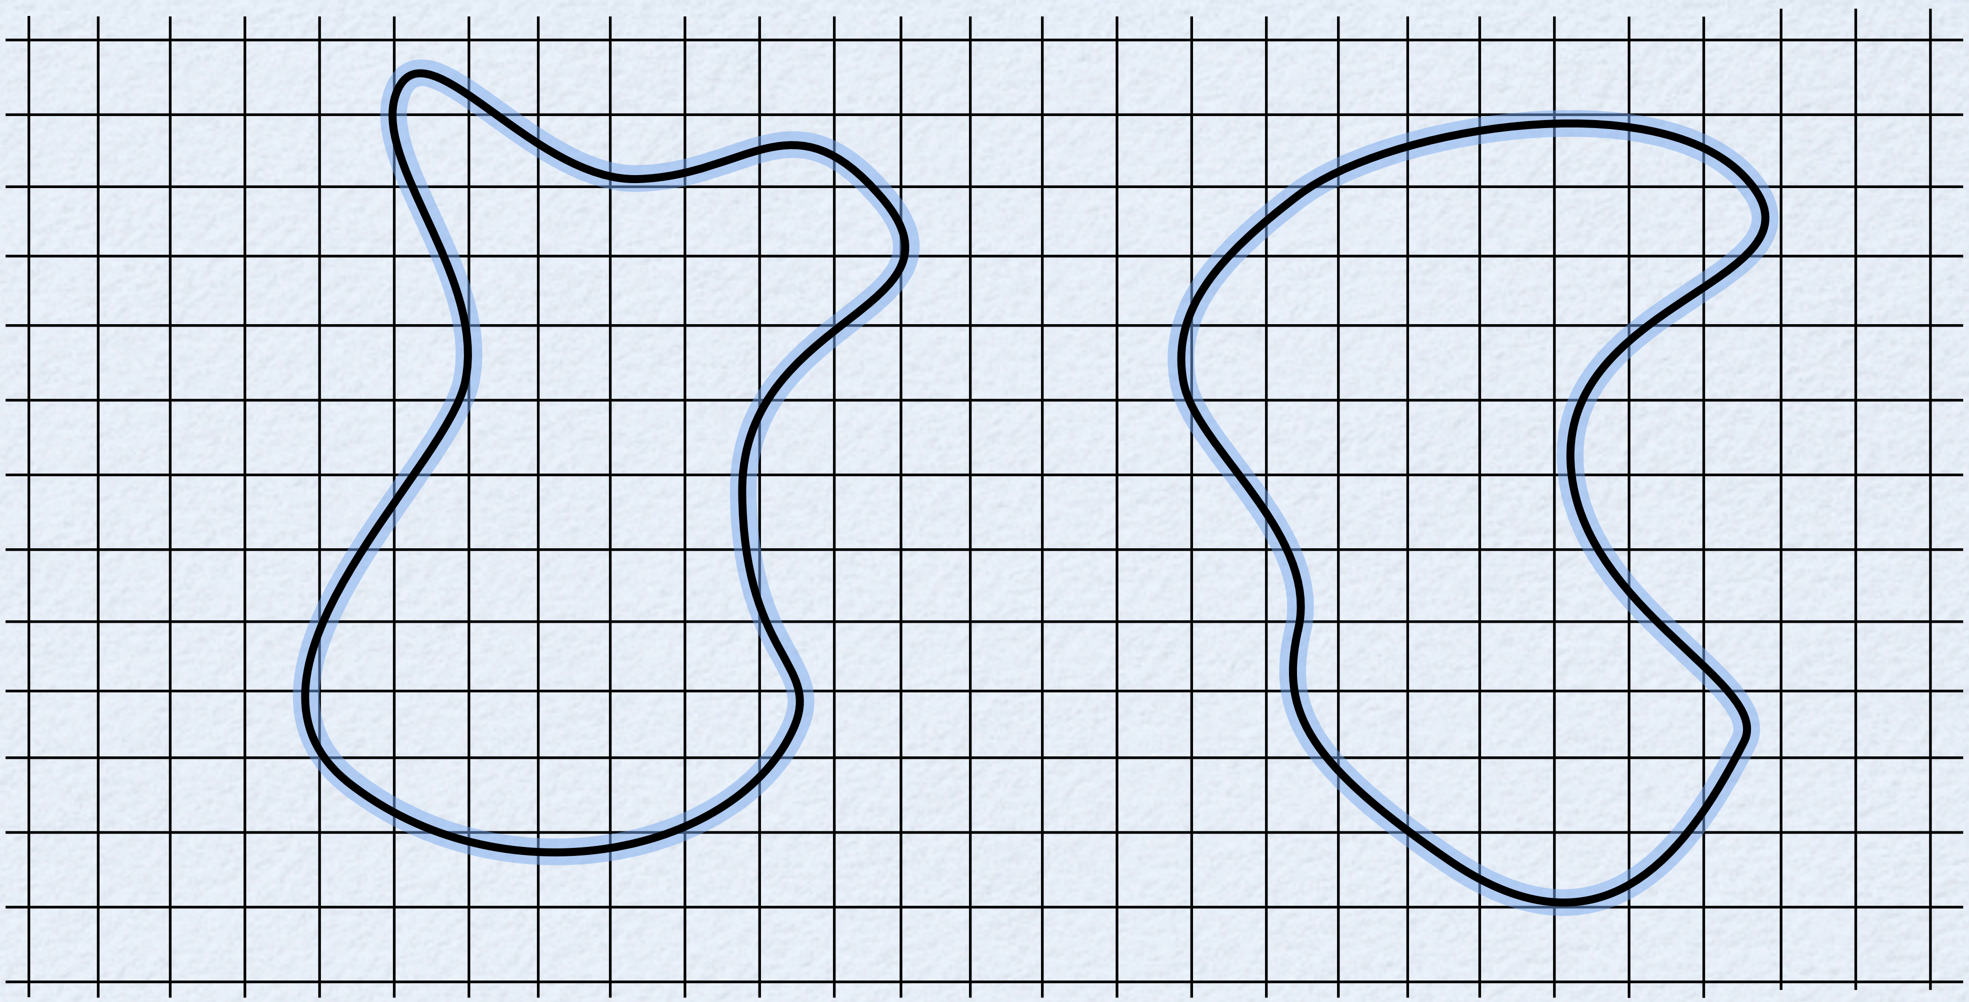

$$\rho \frac{D\mathbf{u}}{Dt} = \nabla \cdot \boldsymbol{\sigma} + f(\text{enforces b.c.})$$

Penalization Method: $f(\mathbf{x}) = \lambda \chi_S (\mathbf{u}_S - \mathbf{u})$

Immersed Boundary Method: $f(\mathbf{x}) = \kappa \delta_S (\mathbf{x}_S - \mathbf{x})$

Boundary Conditions = Coupling

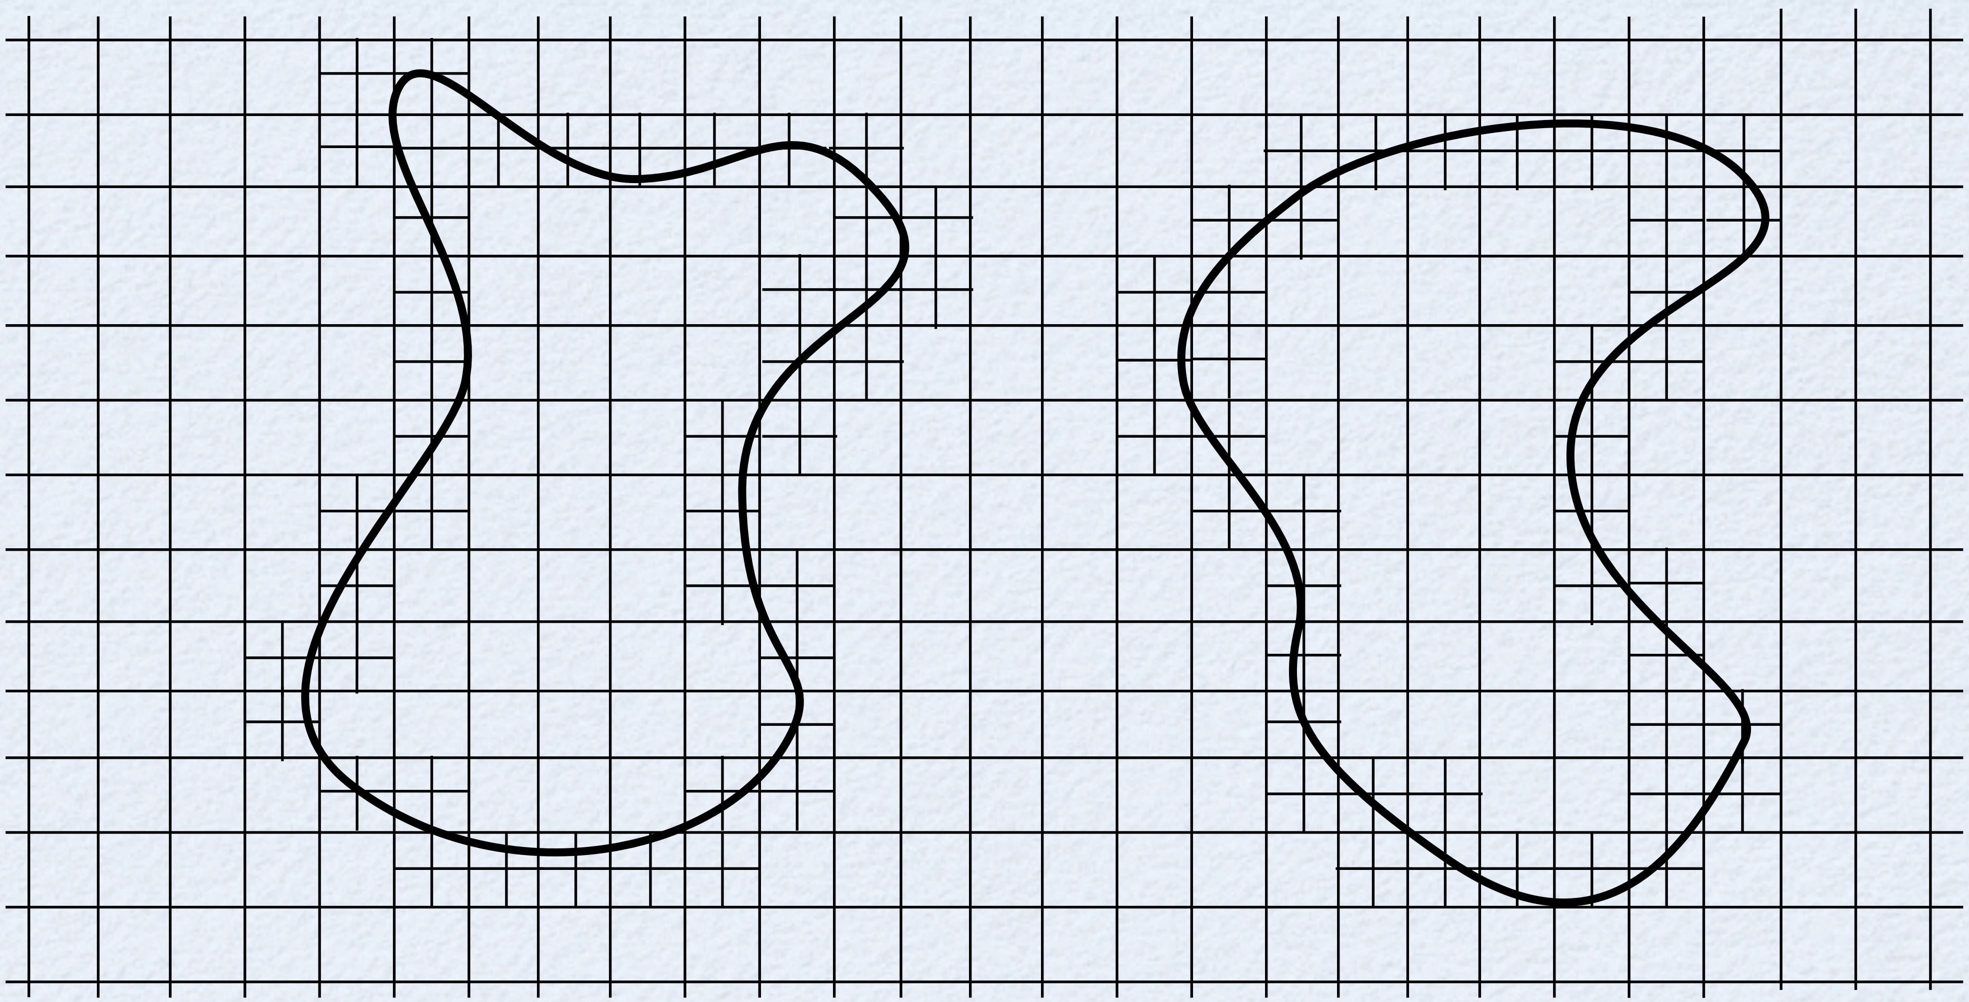

$$\rho \frac{D\mathbf{u}}{Dt} = \nabla \cdot \boldsymbol{\sigma} \; + f(\text{enforces b.c.})$$

Penalization Method: $f(\mathbf{x}) = \lambda \chi_S(\mathbf{u}_S - \mathbf{u})$

Immersed Boundary Method: $f(\mathbf{x}) = \kappa \delta_S(\mathbf{x}_S - \mathbf{x})$

Boundary Conditions = Coupling

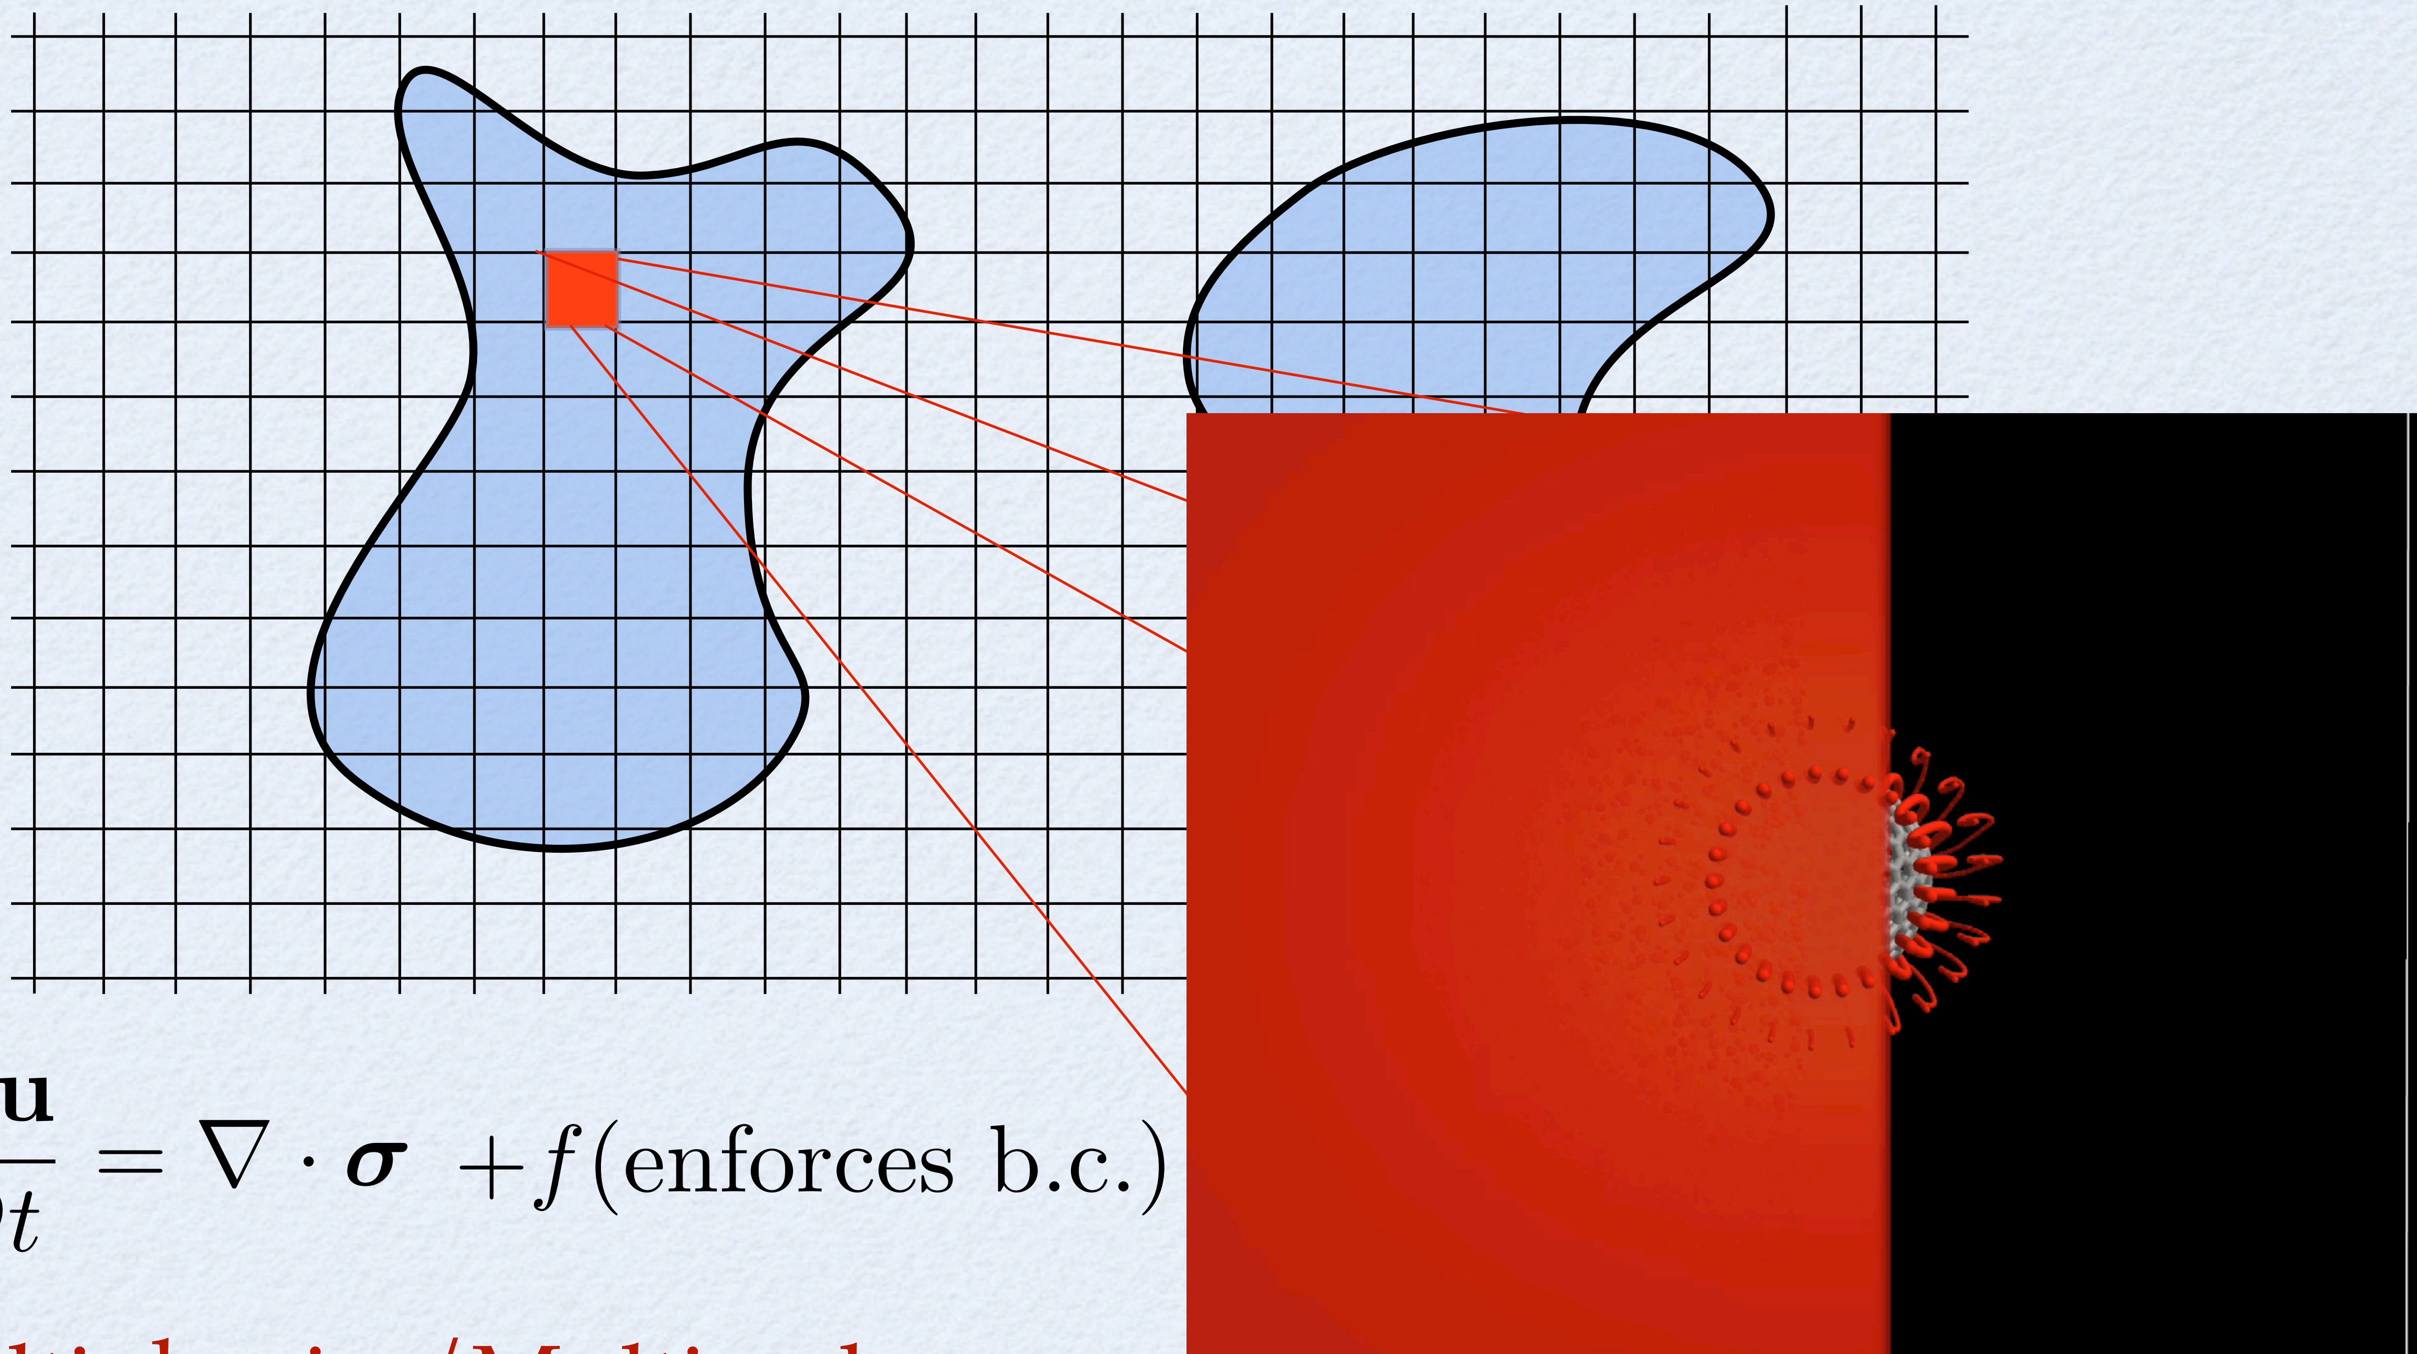

$$\rho \frac{D\mathbf{u}}{Dt} = \nabla \cdot \boldsymbol{\sigma} + f(\text{enforces b.c.})$$

Multiphysics/Multiscale

$$f(\mathbf{x}) = (\text{result from Molecular Simulations})$$

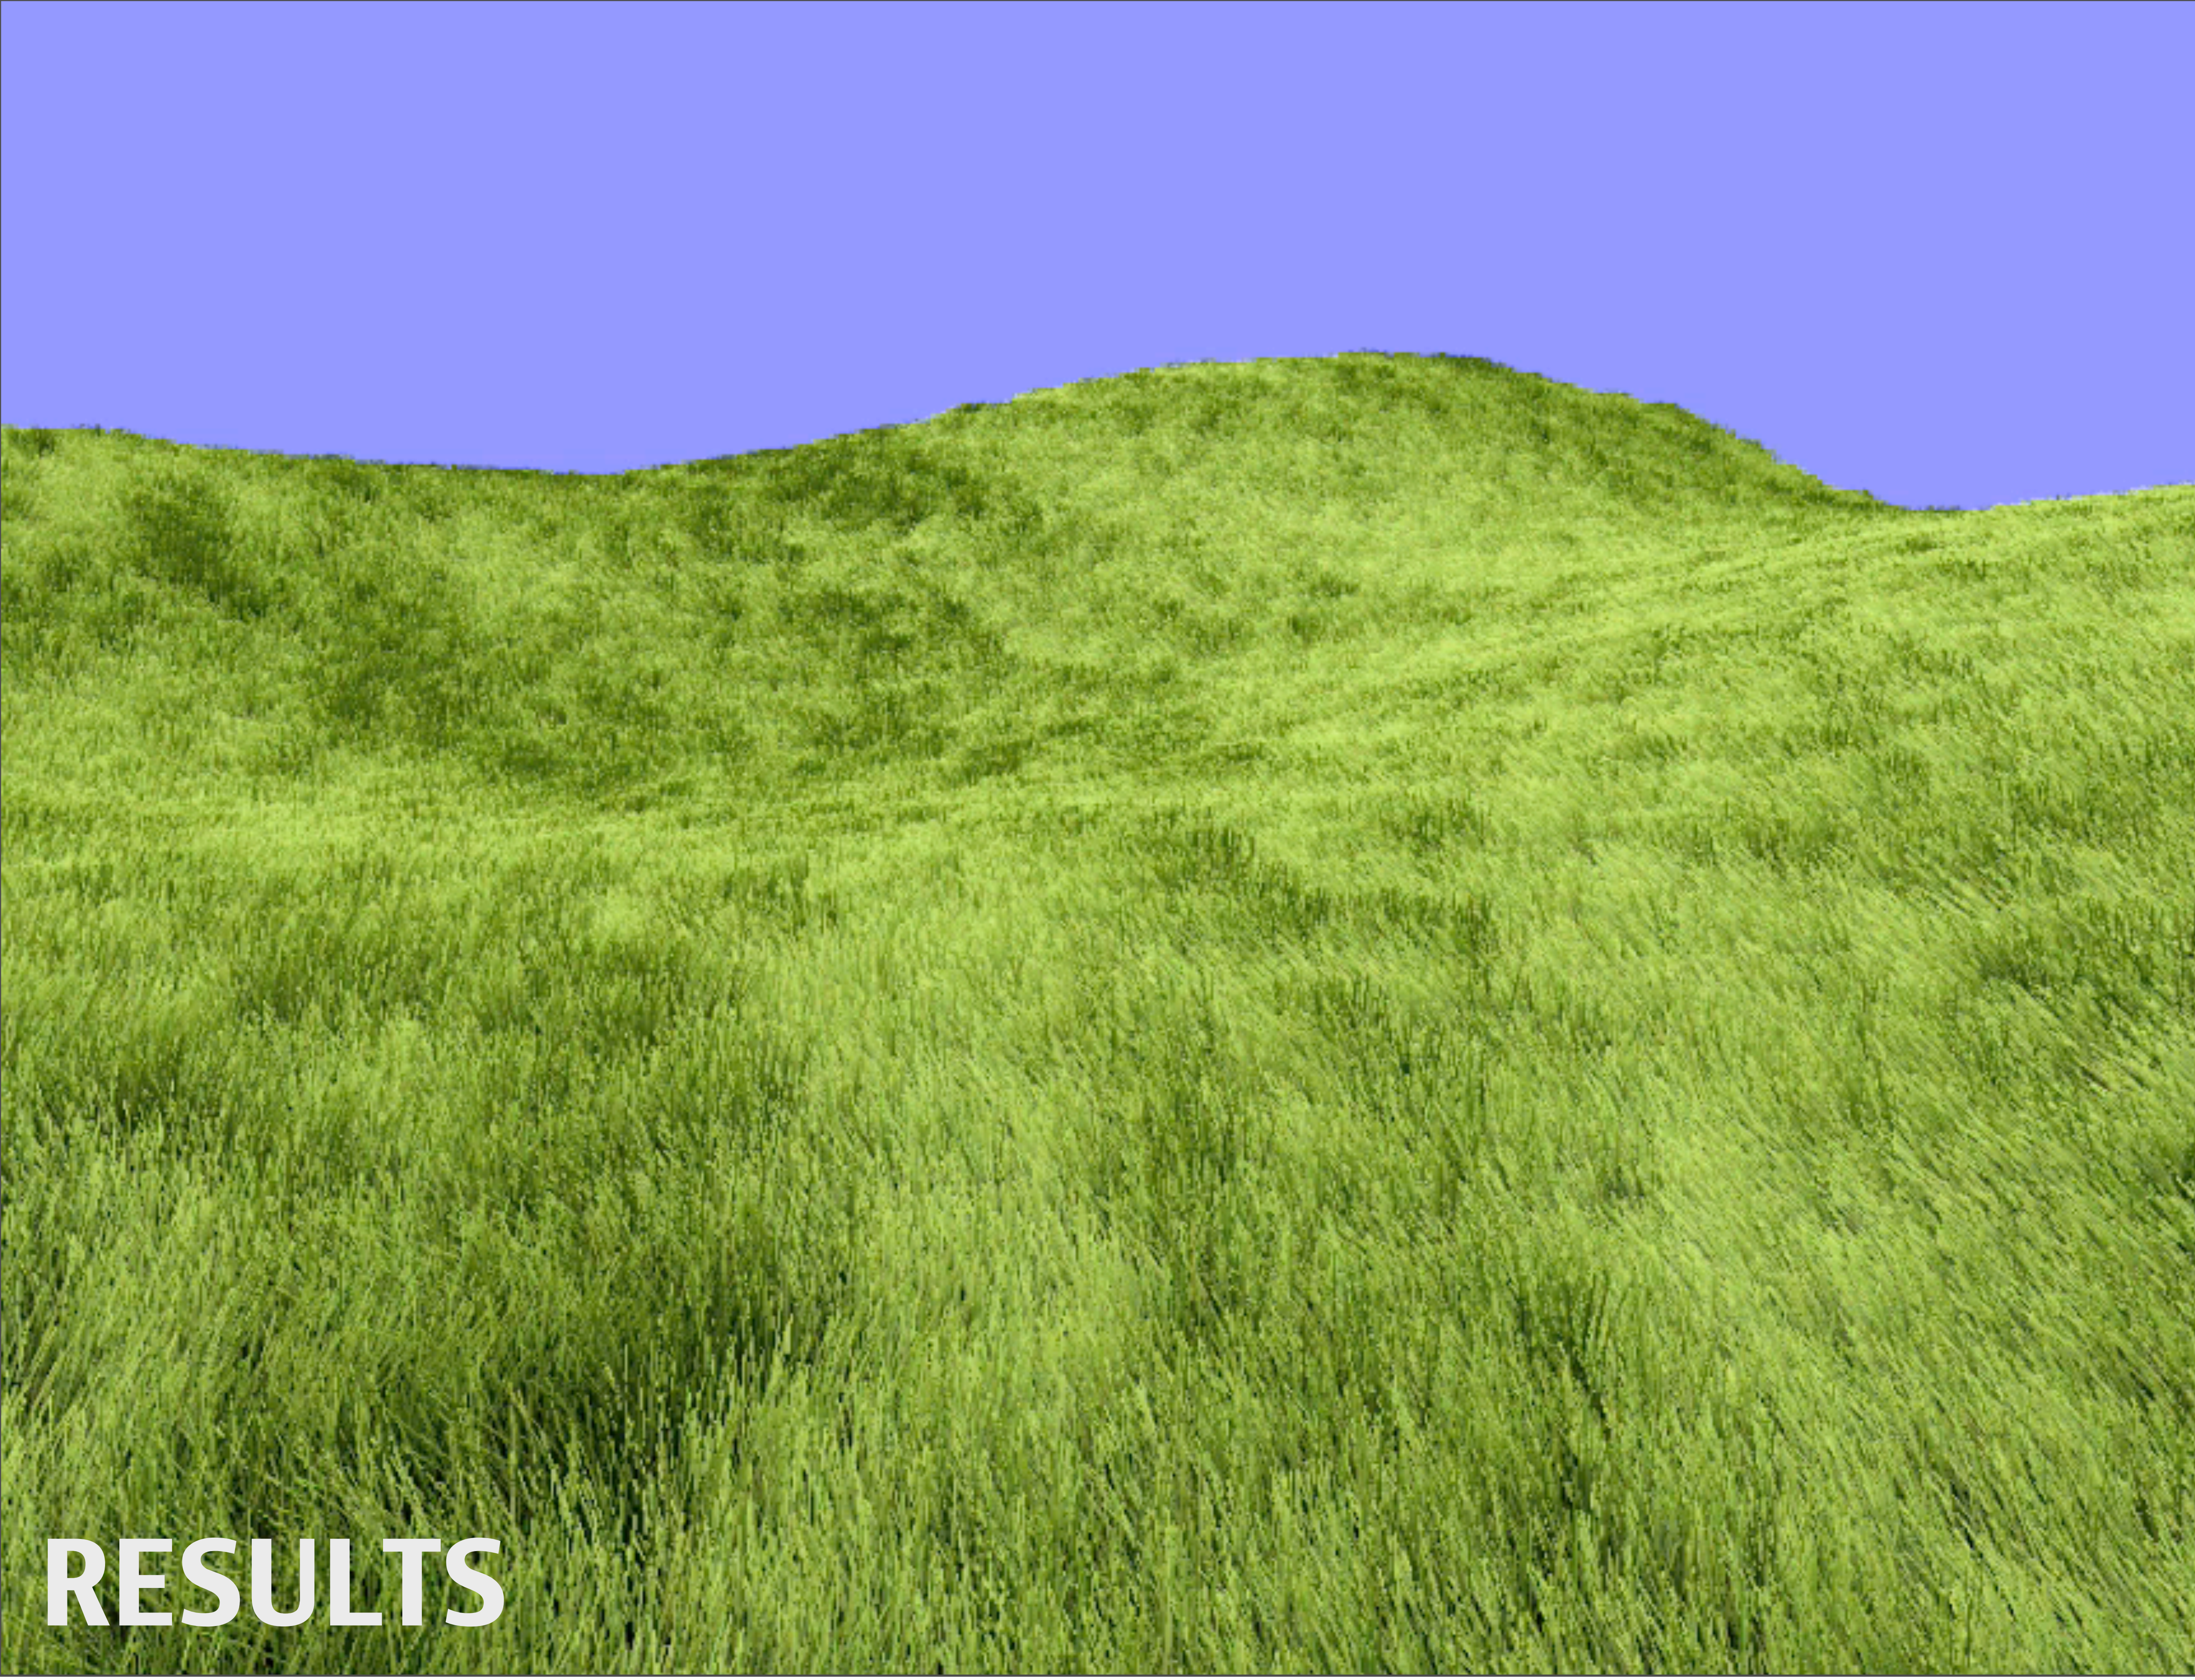

RESULTS

SPHERE @ Re = 1000 with Effective Resolution 1024^3

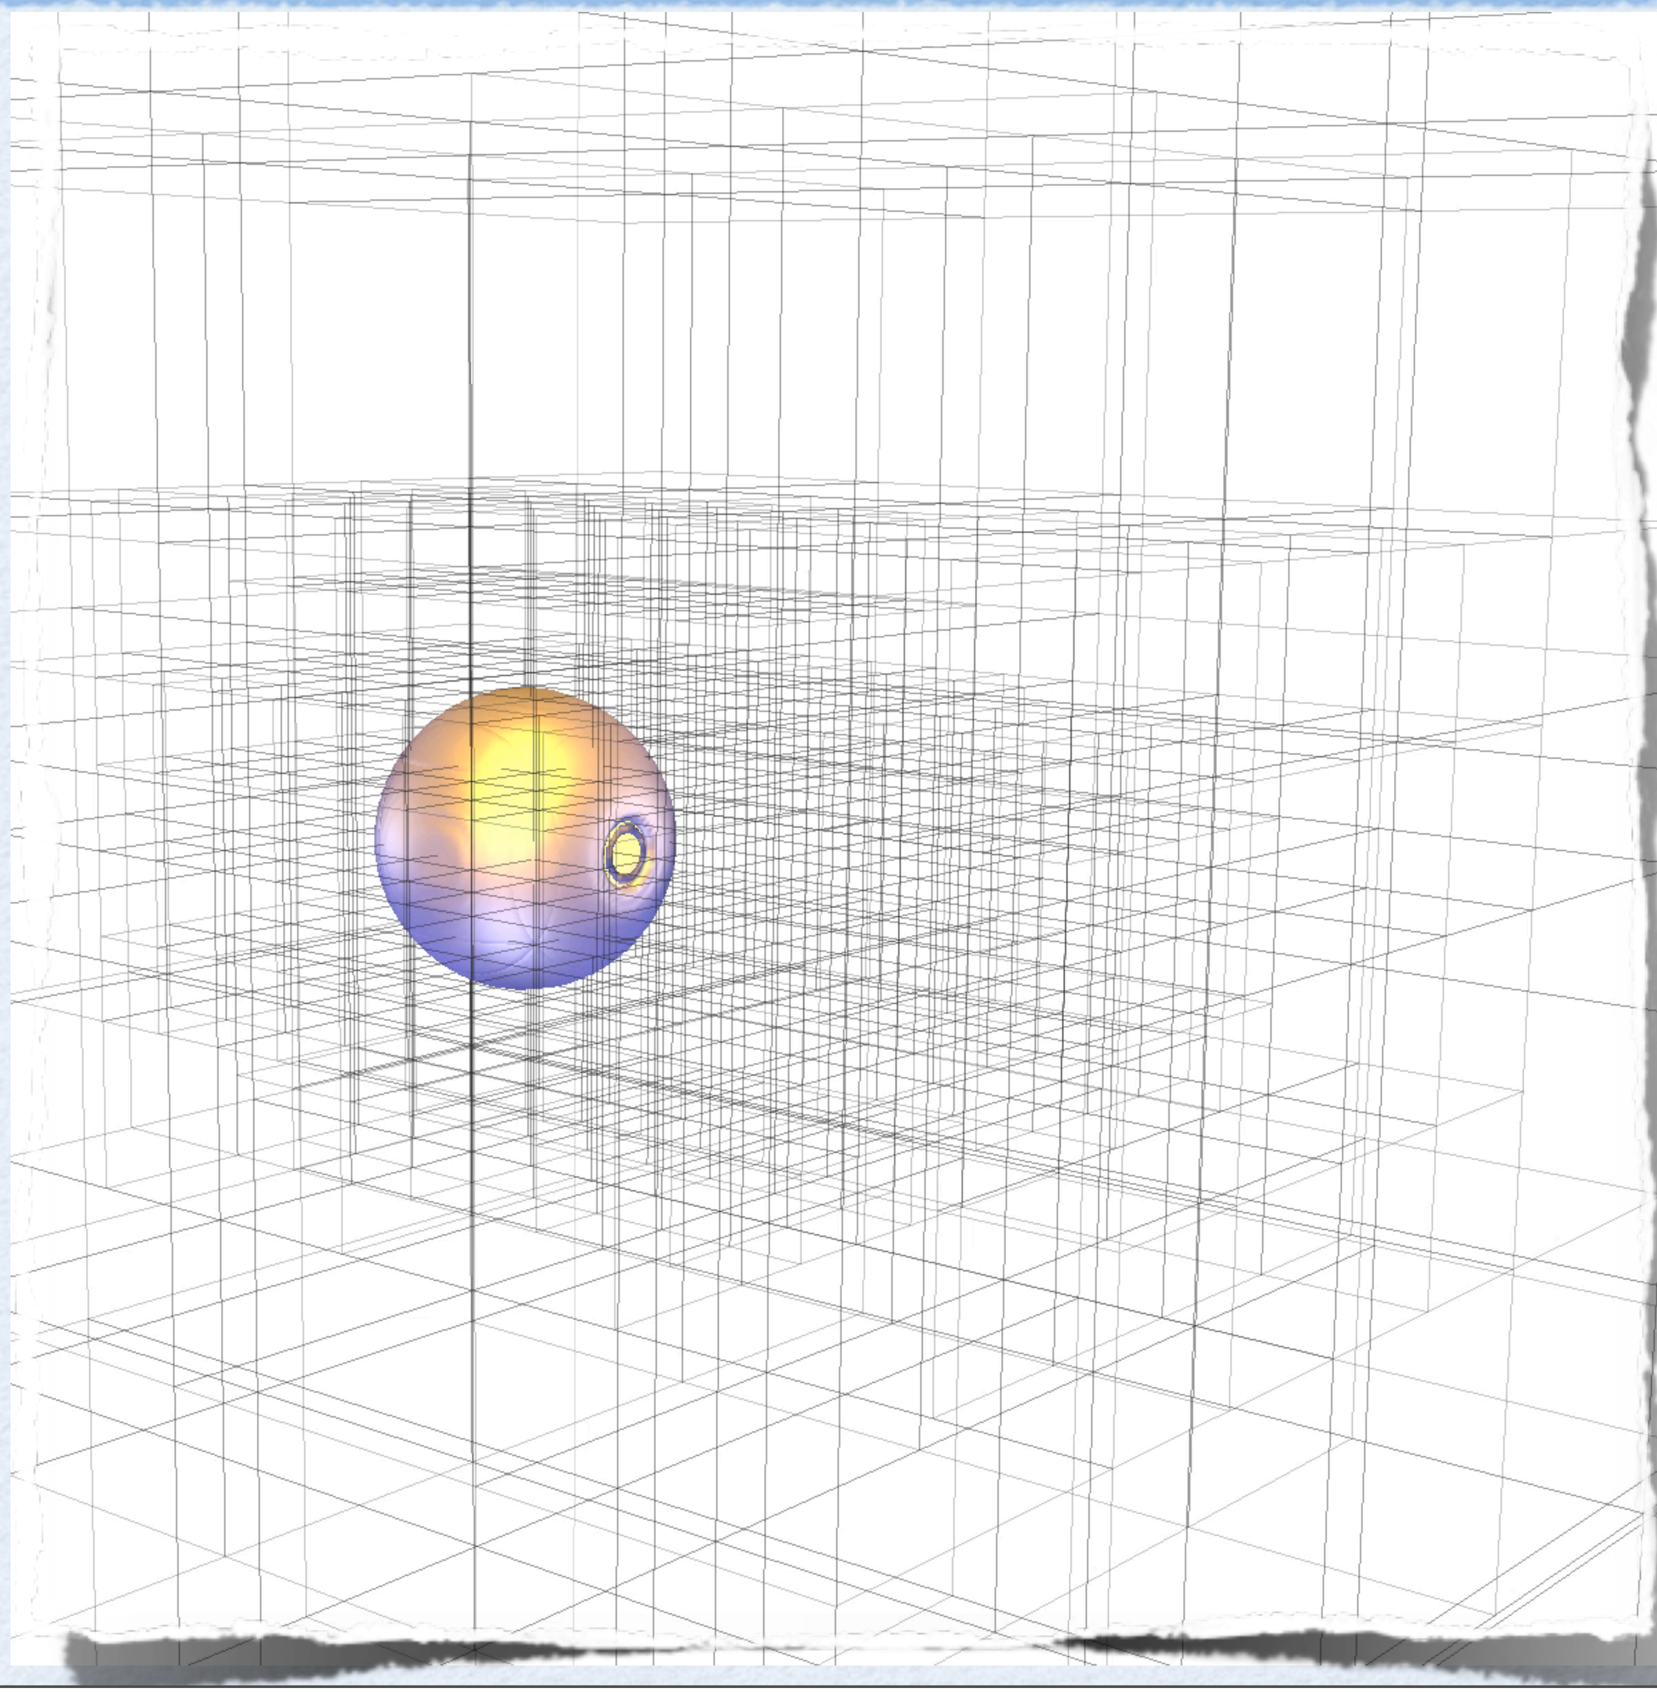

TIMINGS : 4 days on 3
cores, 2.4 GHz - OpenM
and MPI and TBB

Multi-body Simulations

TIMINGS : 3hours on 16 cores - TBB + SSE - Solver reaches 70% of the peak performence 180 GFlops over 210

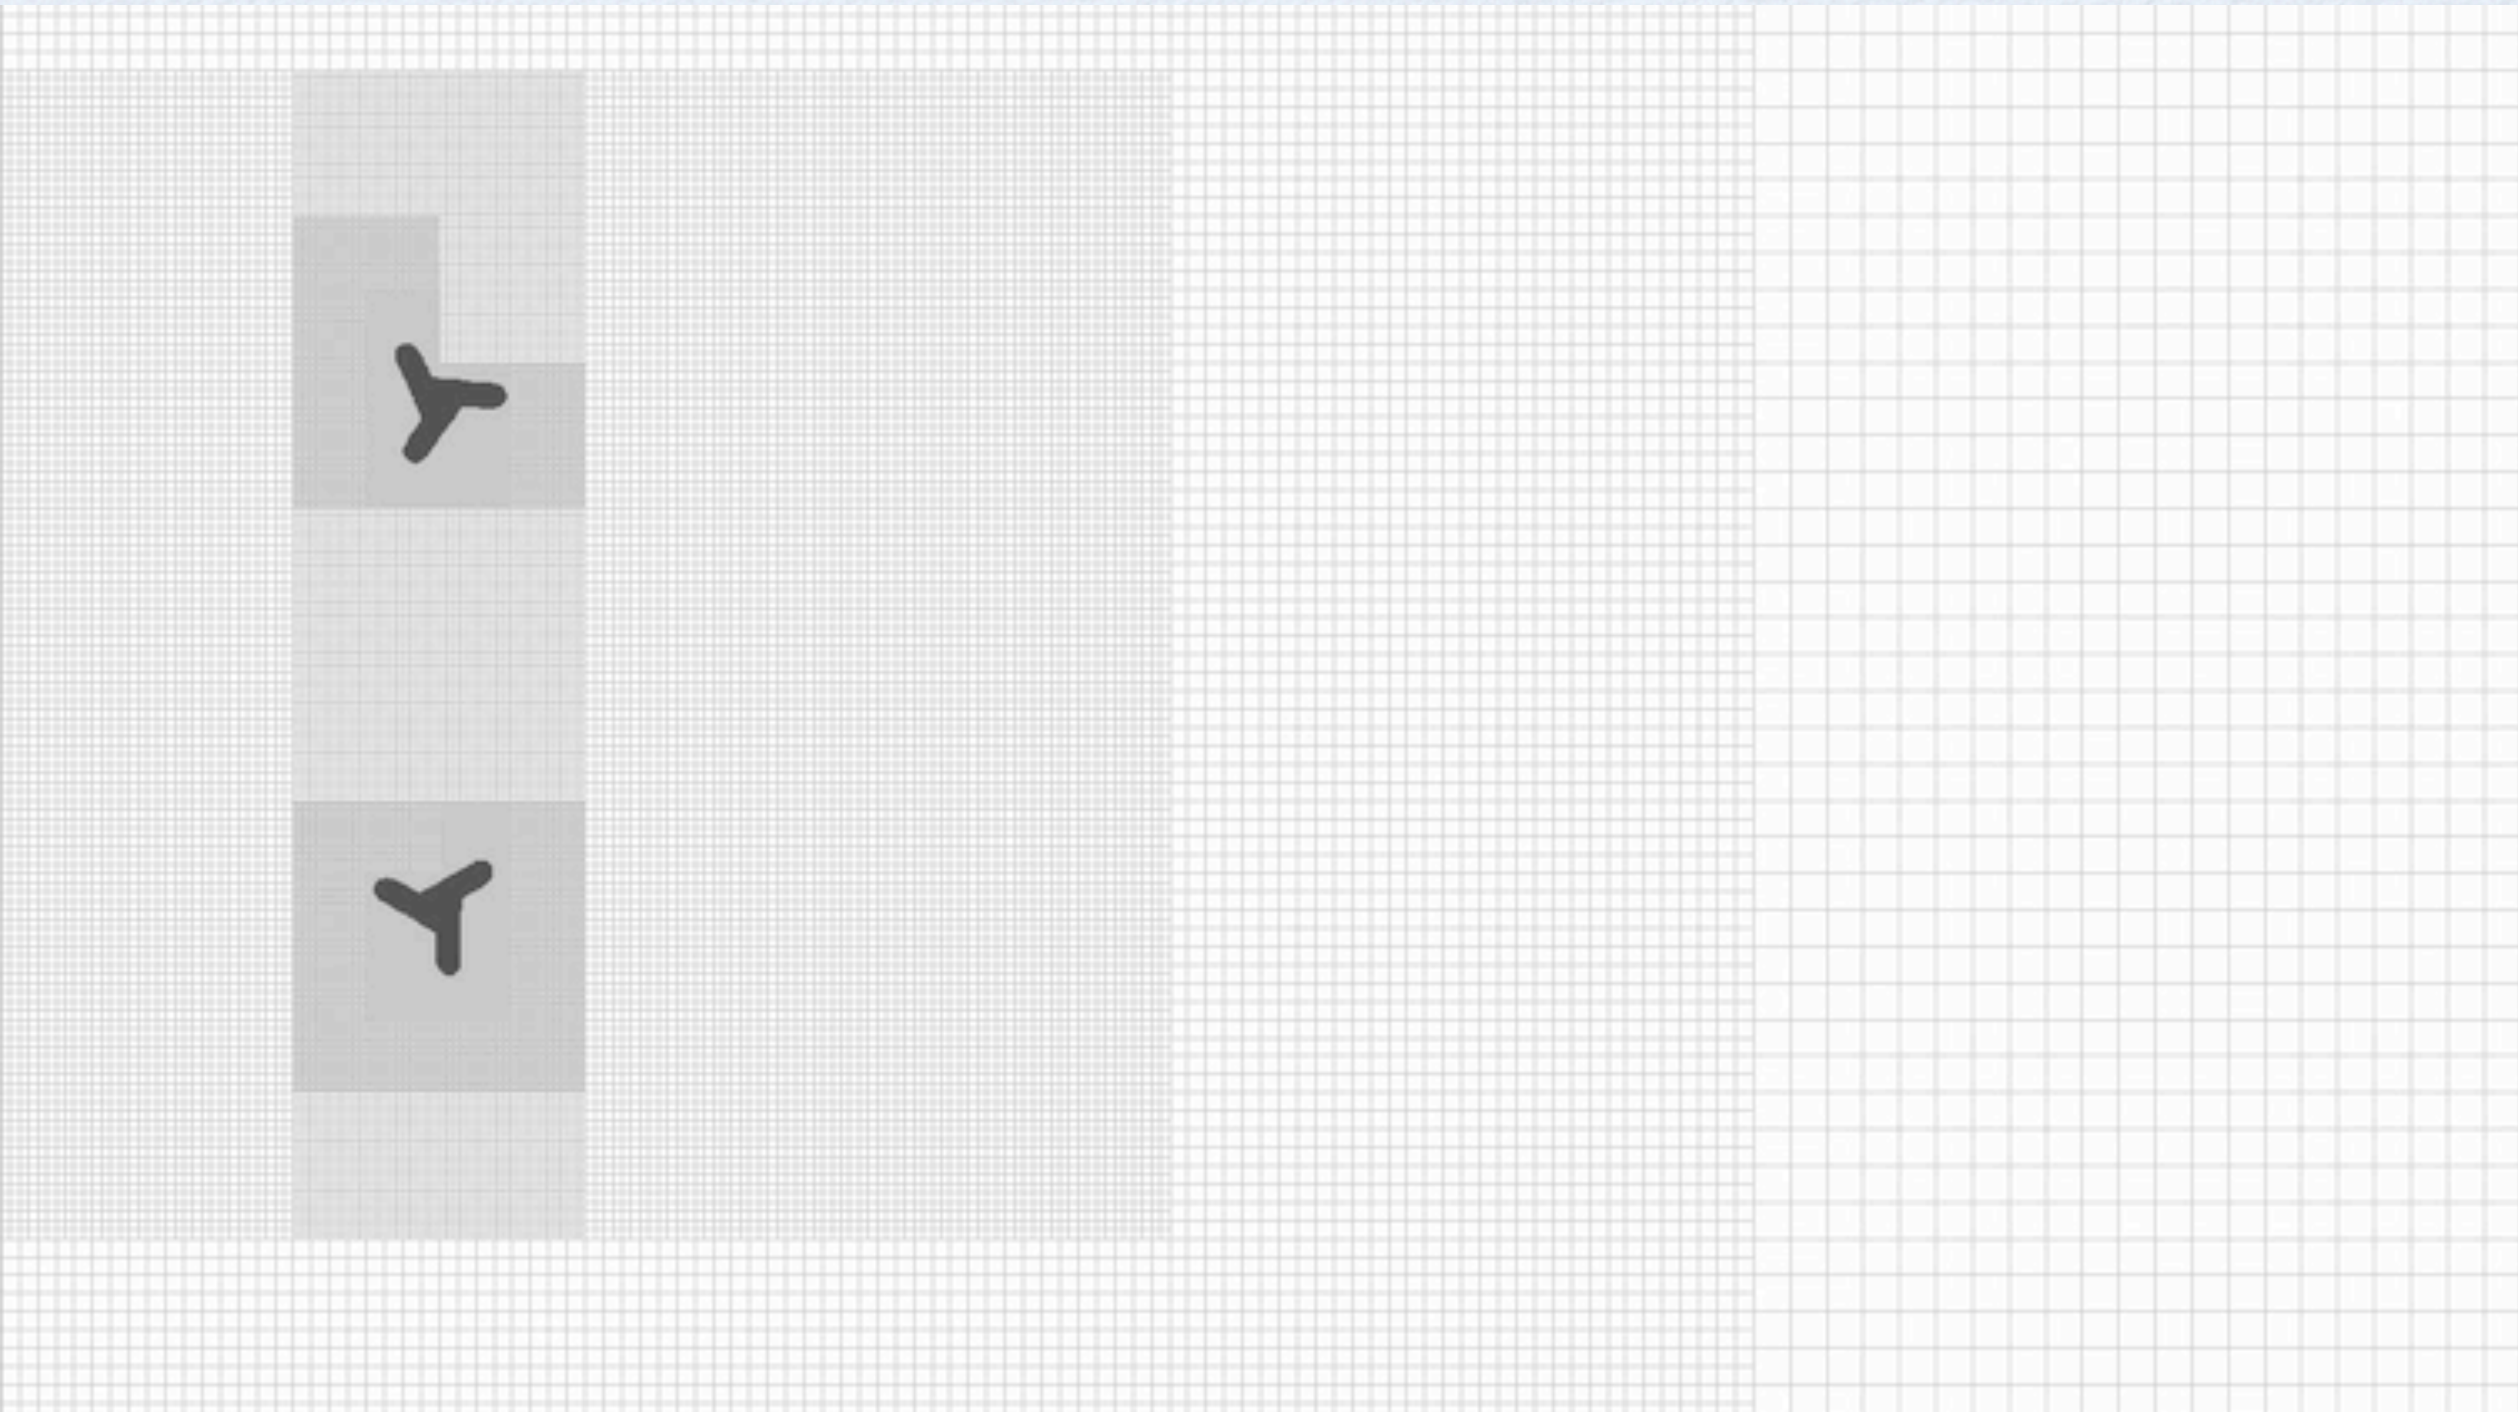

Fish Schooling

Gazzola M., Chatelain P., van Rees W.M., Koumoutsakos P., Simulations of single and multiple swimmers with non-divergence free deforming geometries, **J. of Comput. Physics**, 2011

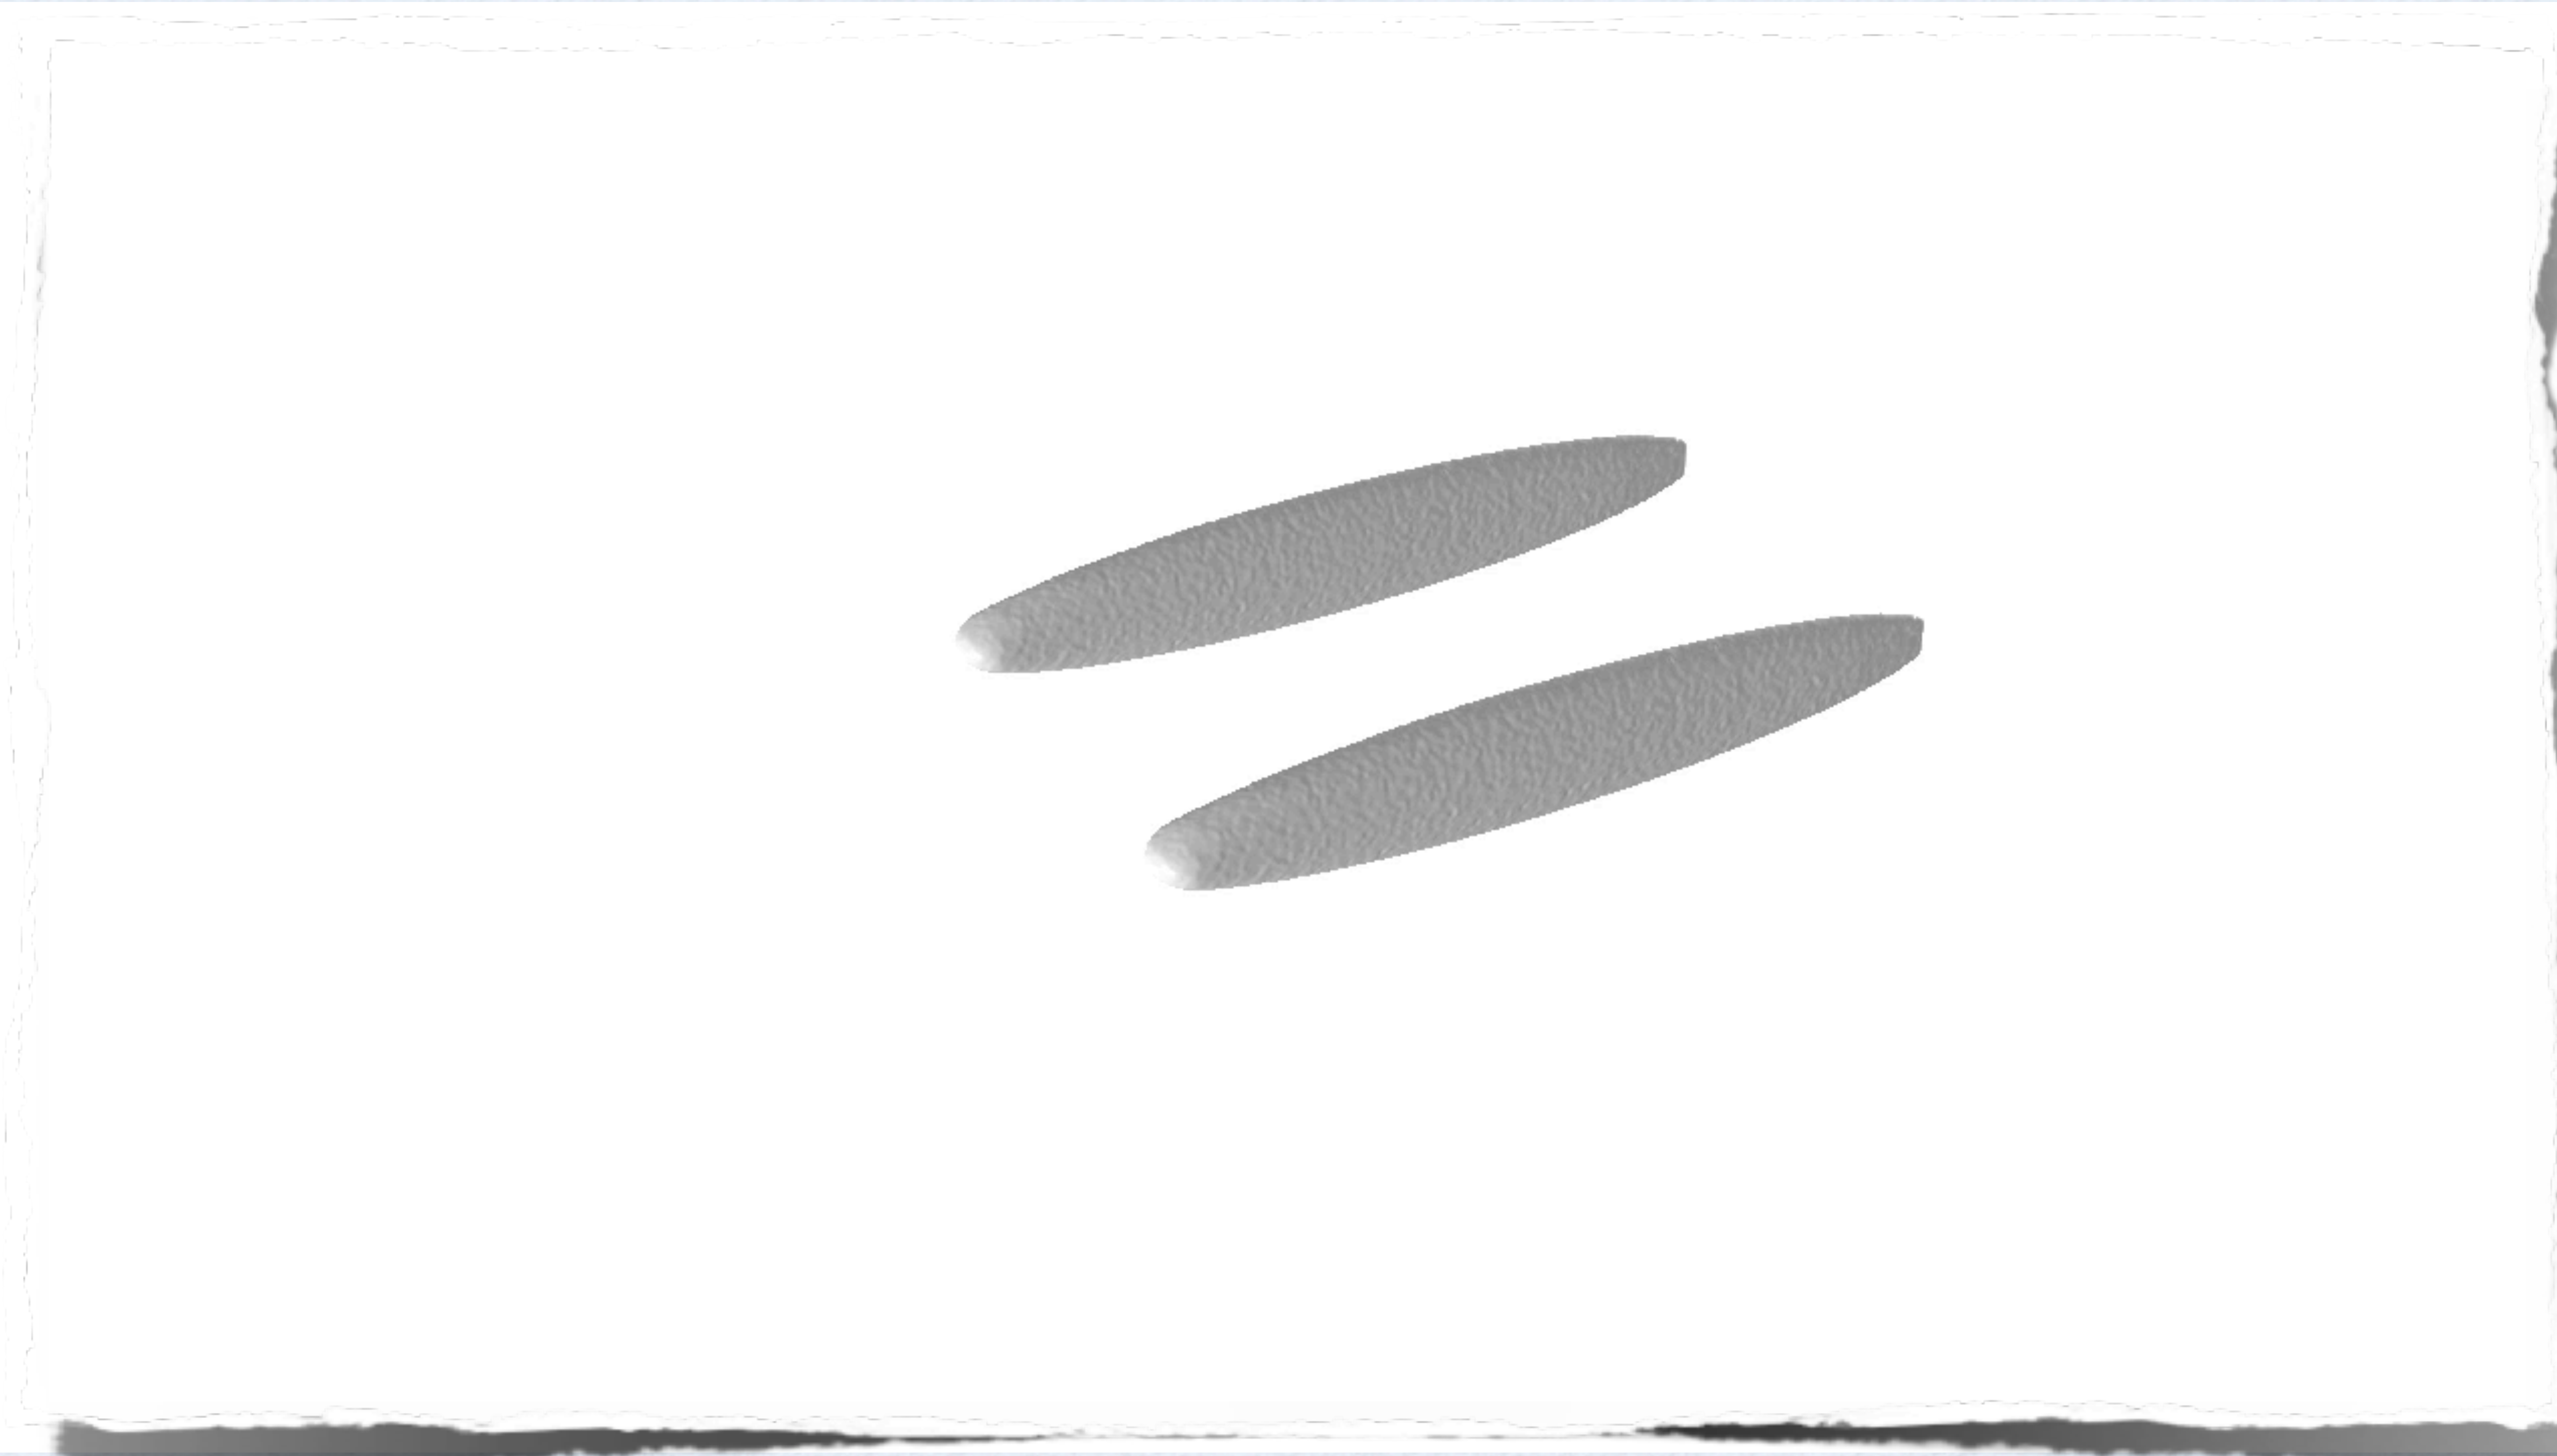

2 FISH (OBVIOUSLY)

Fast Swimmers
Shape Optimization

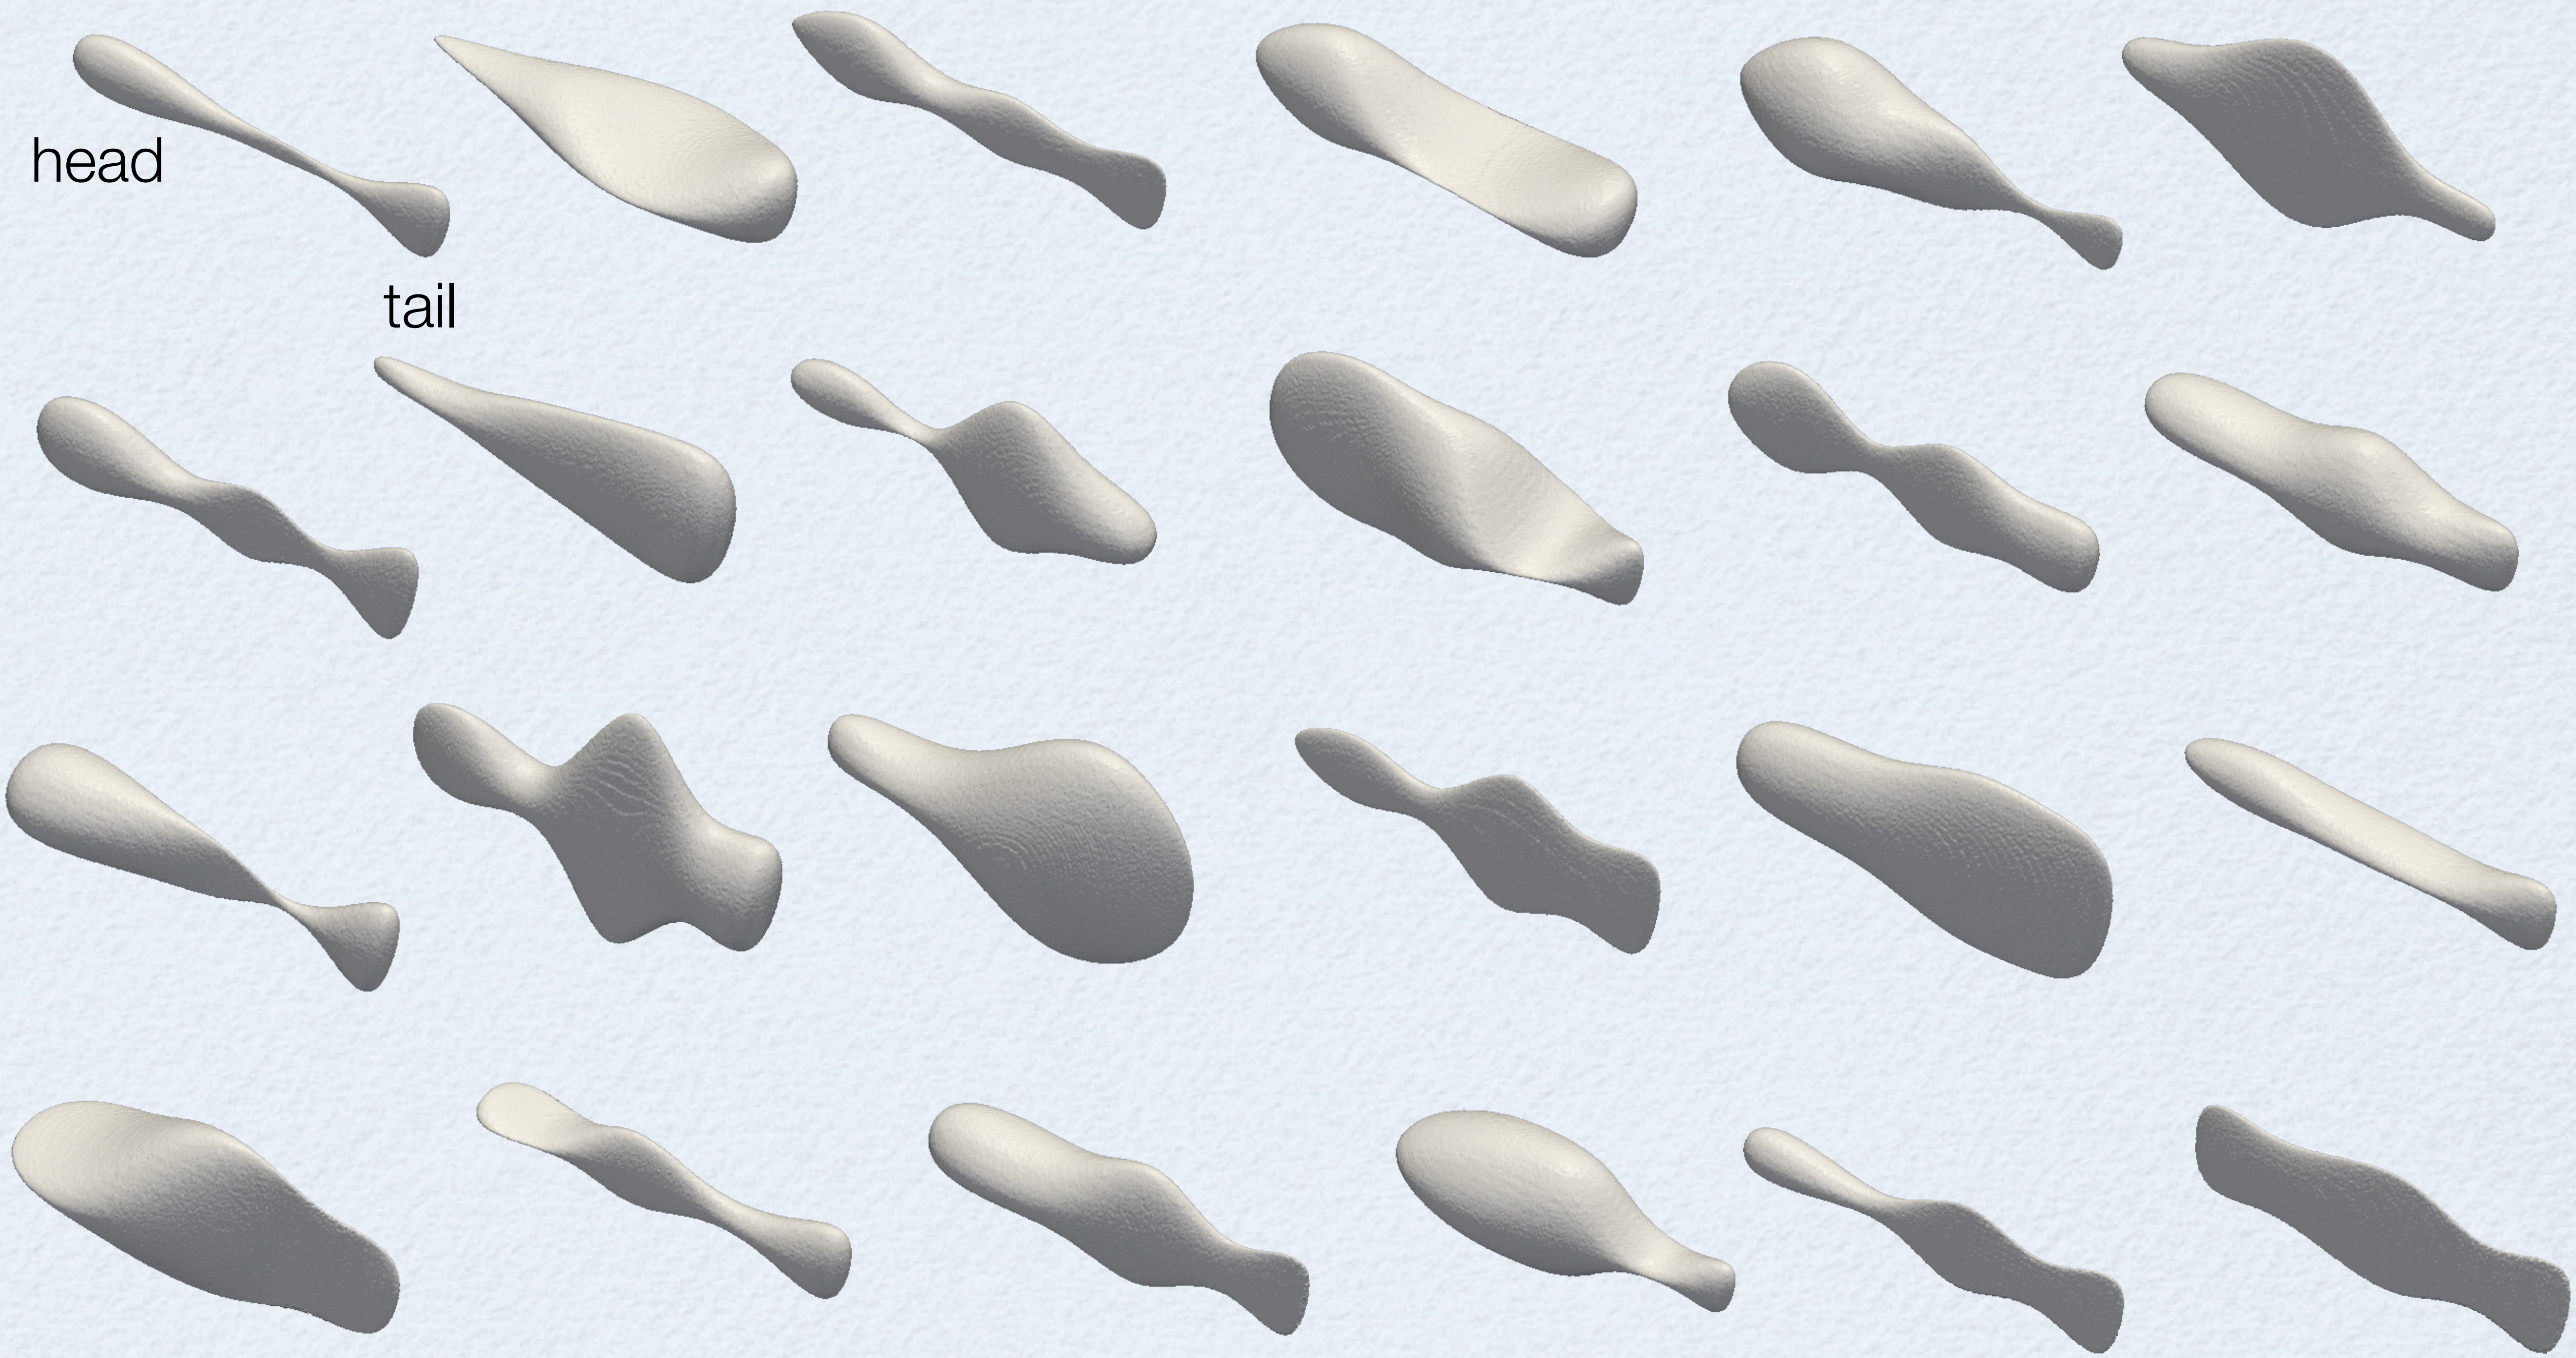

Mean Shape During Evolution

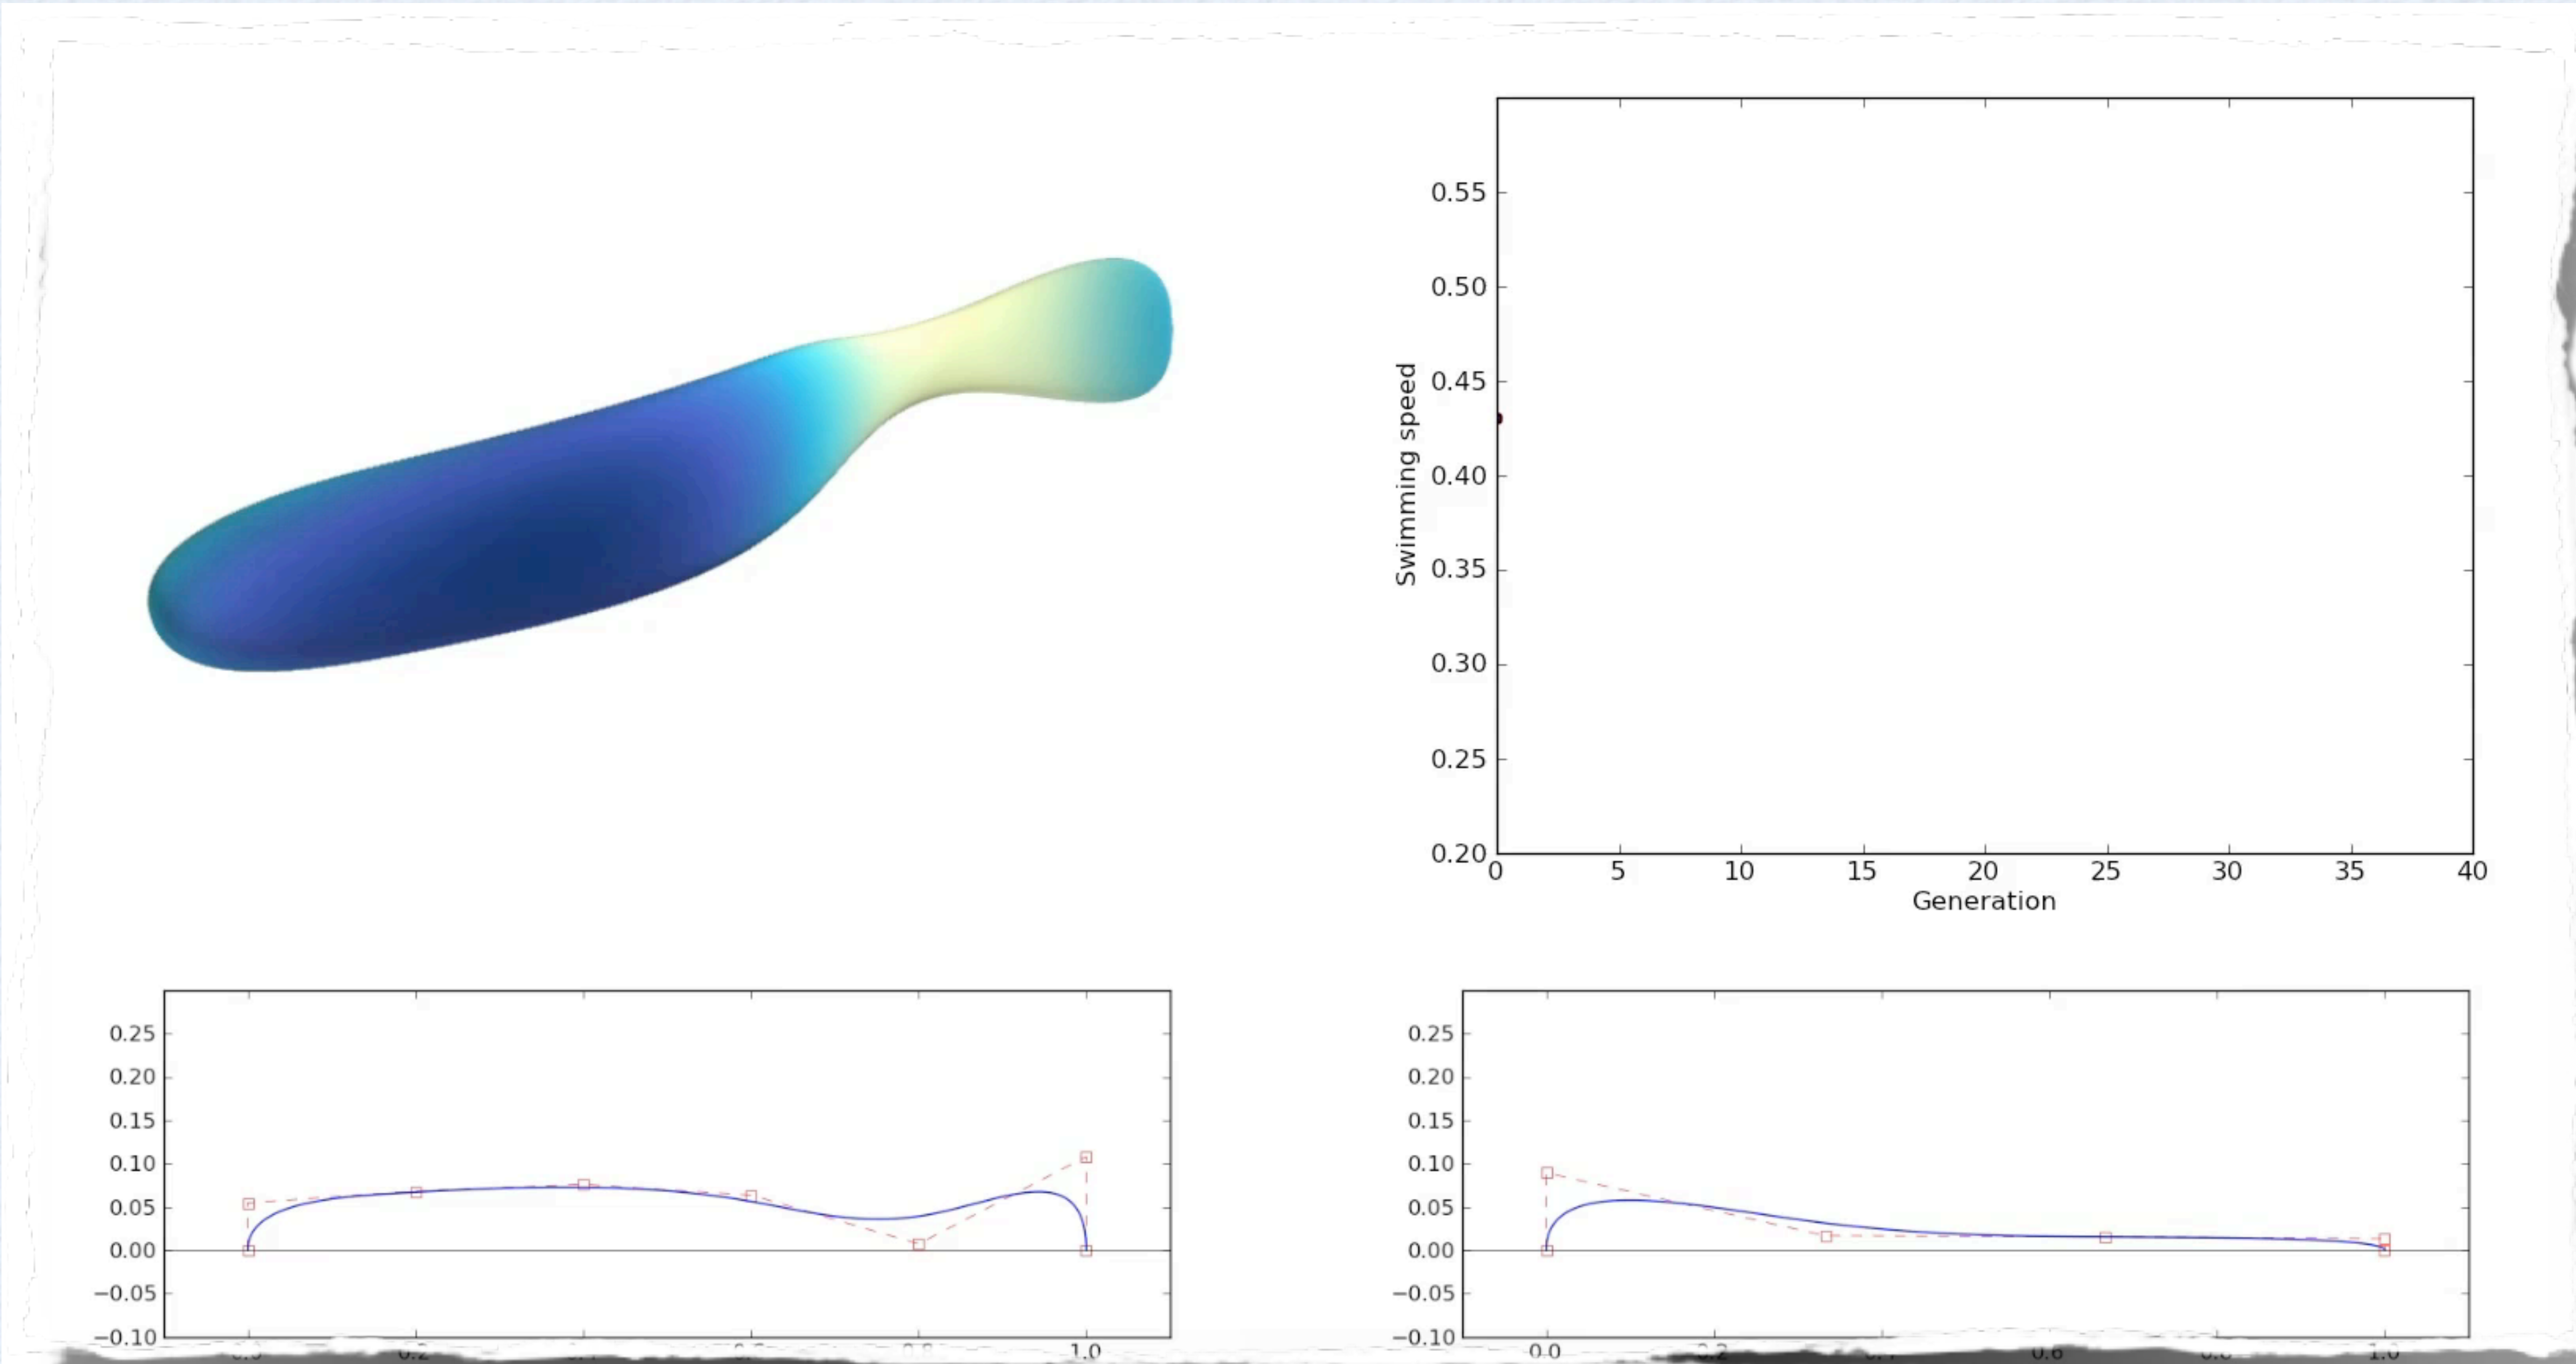

How to escape fast ?

Best Result of an Optimization for escape speed

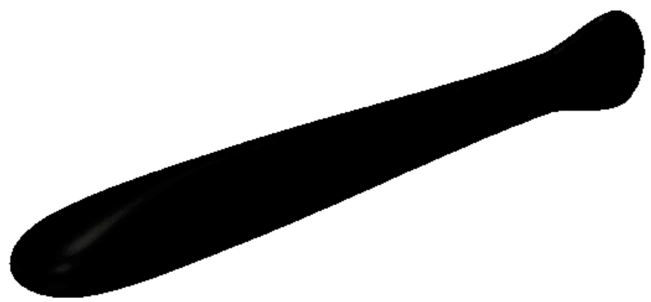

COMPRESSIBLE FLOWS

Moving Boundaries

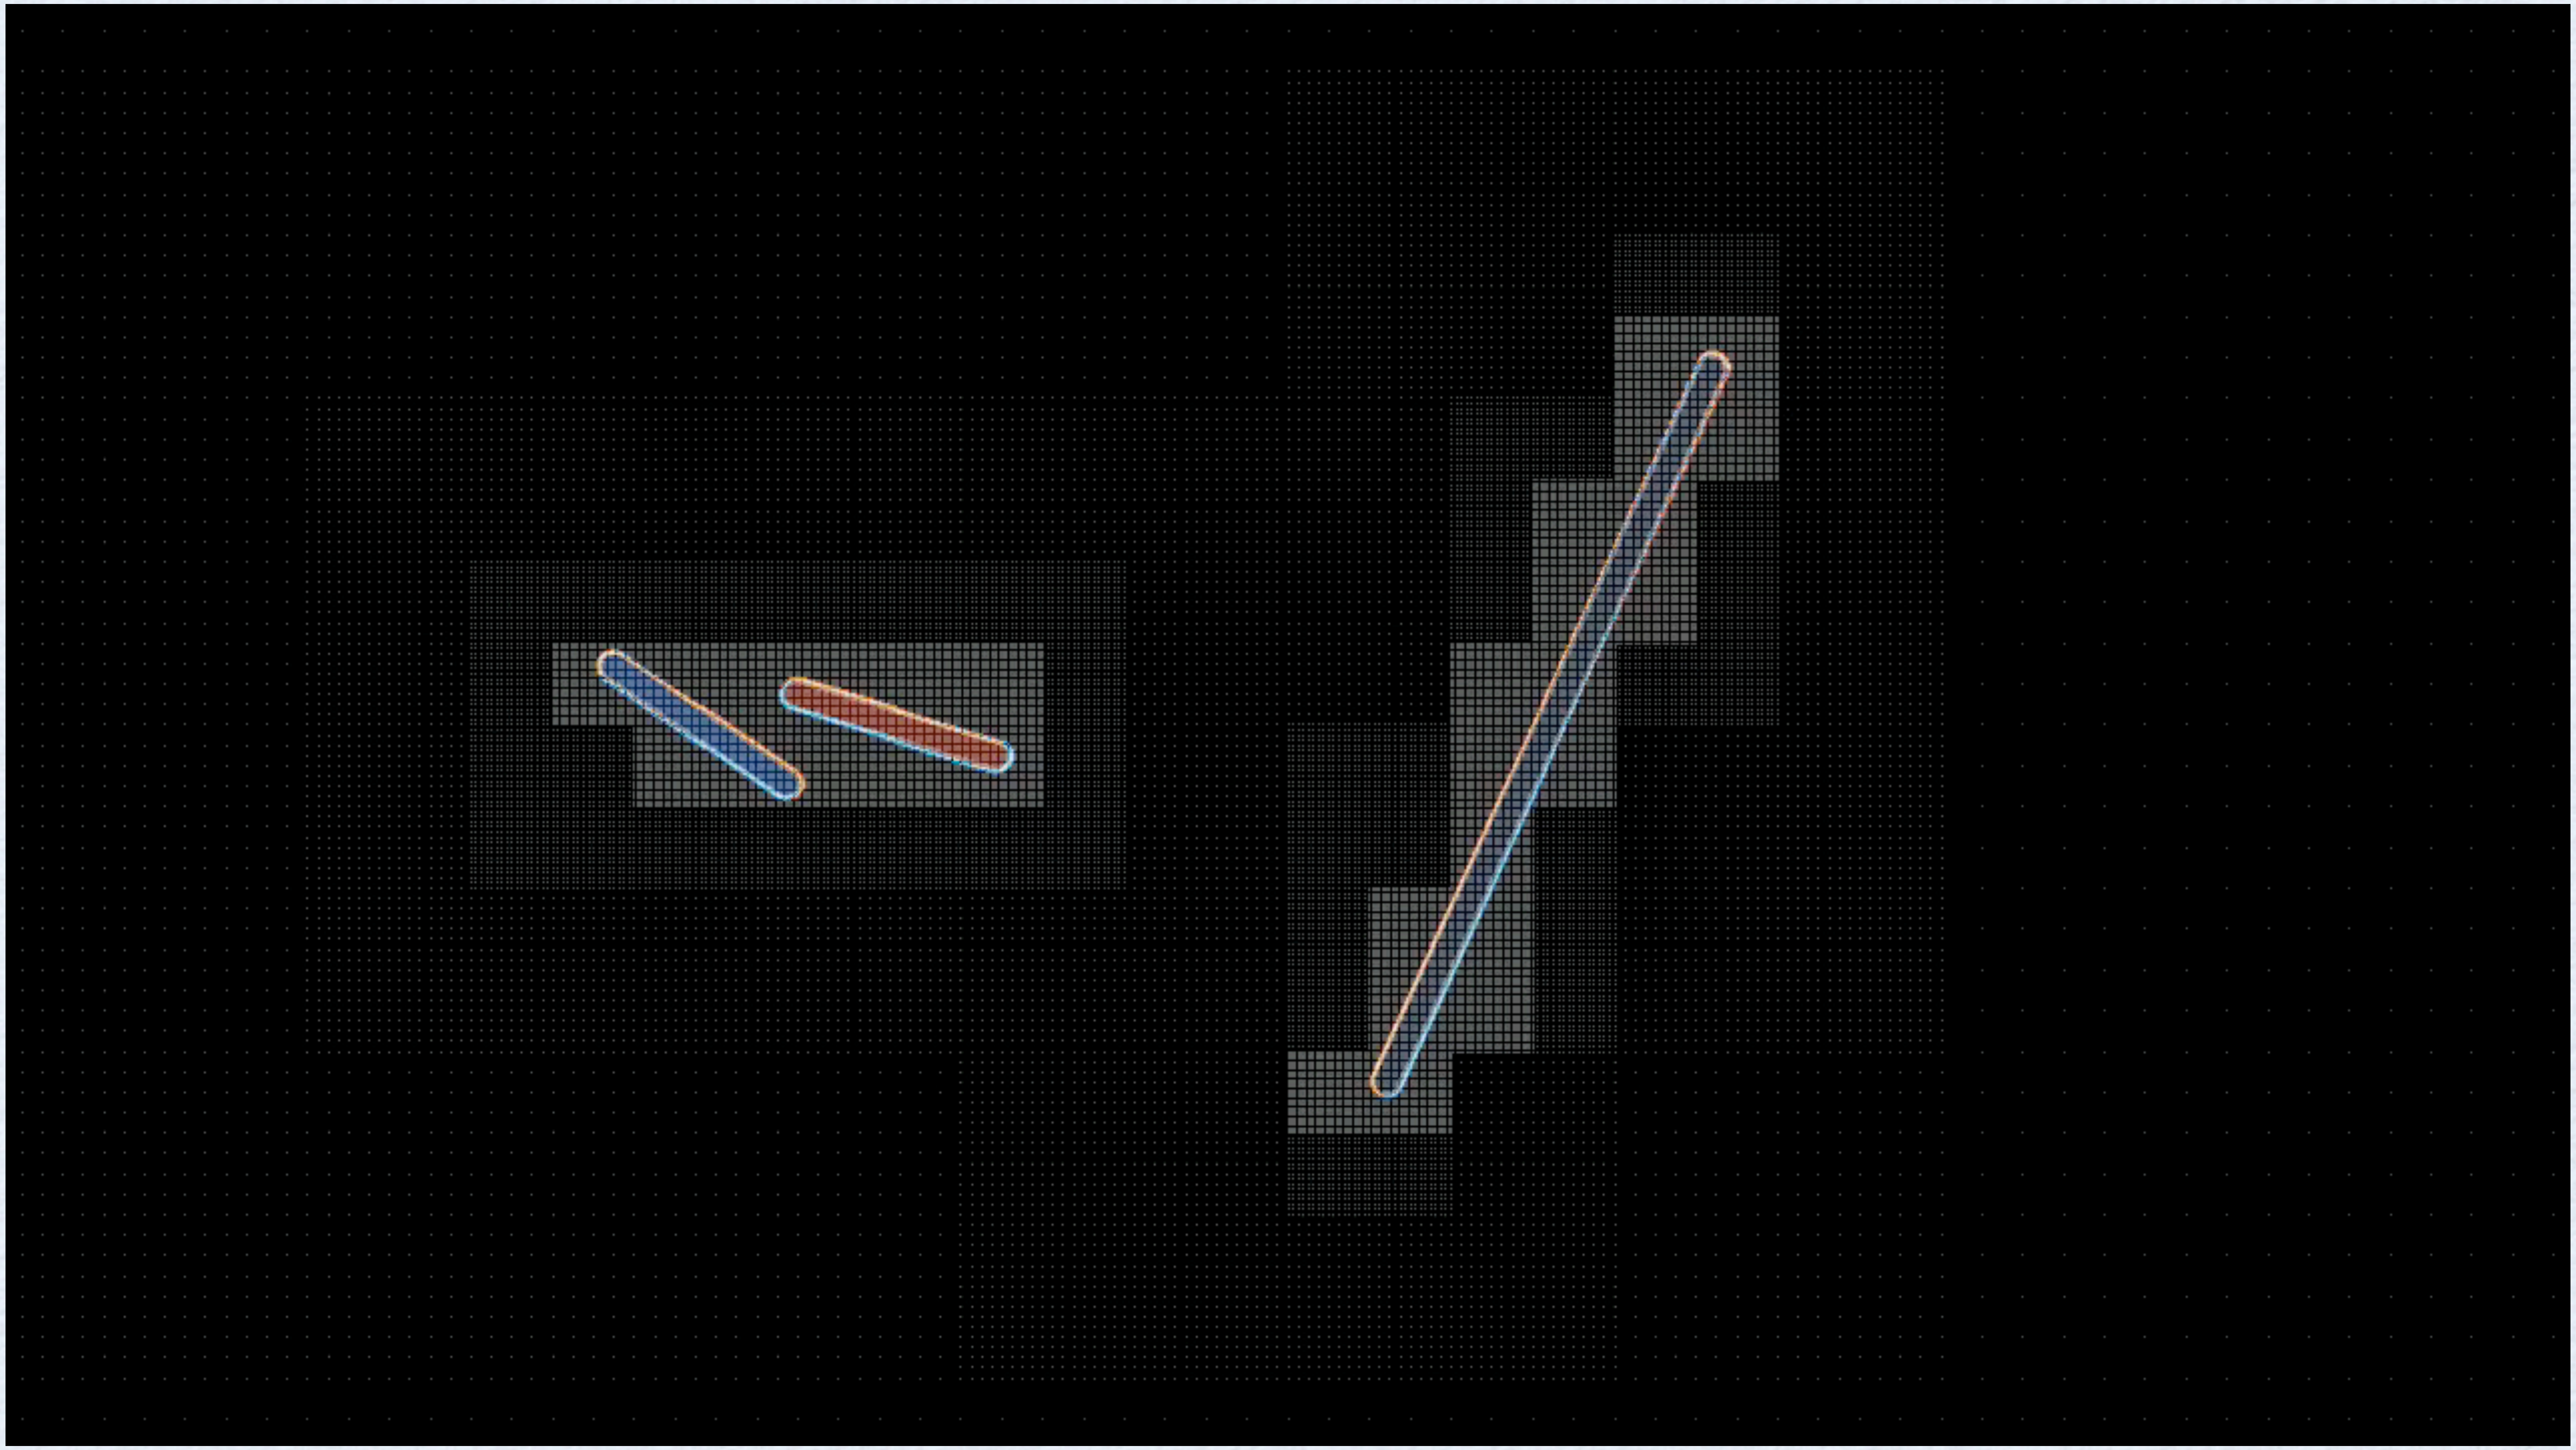

Shock – Ballut Interactions

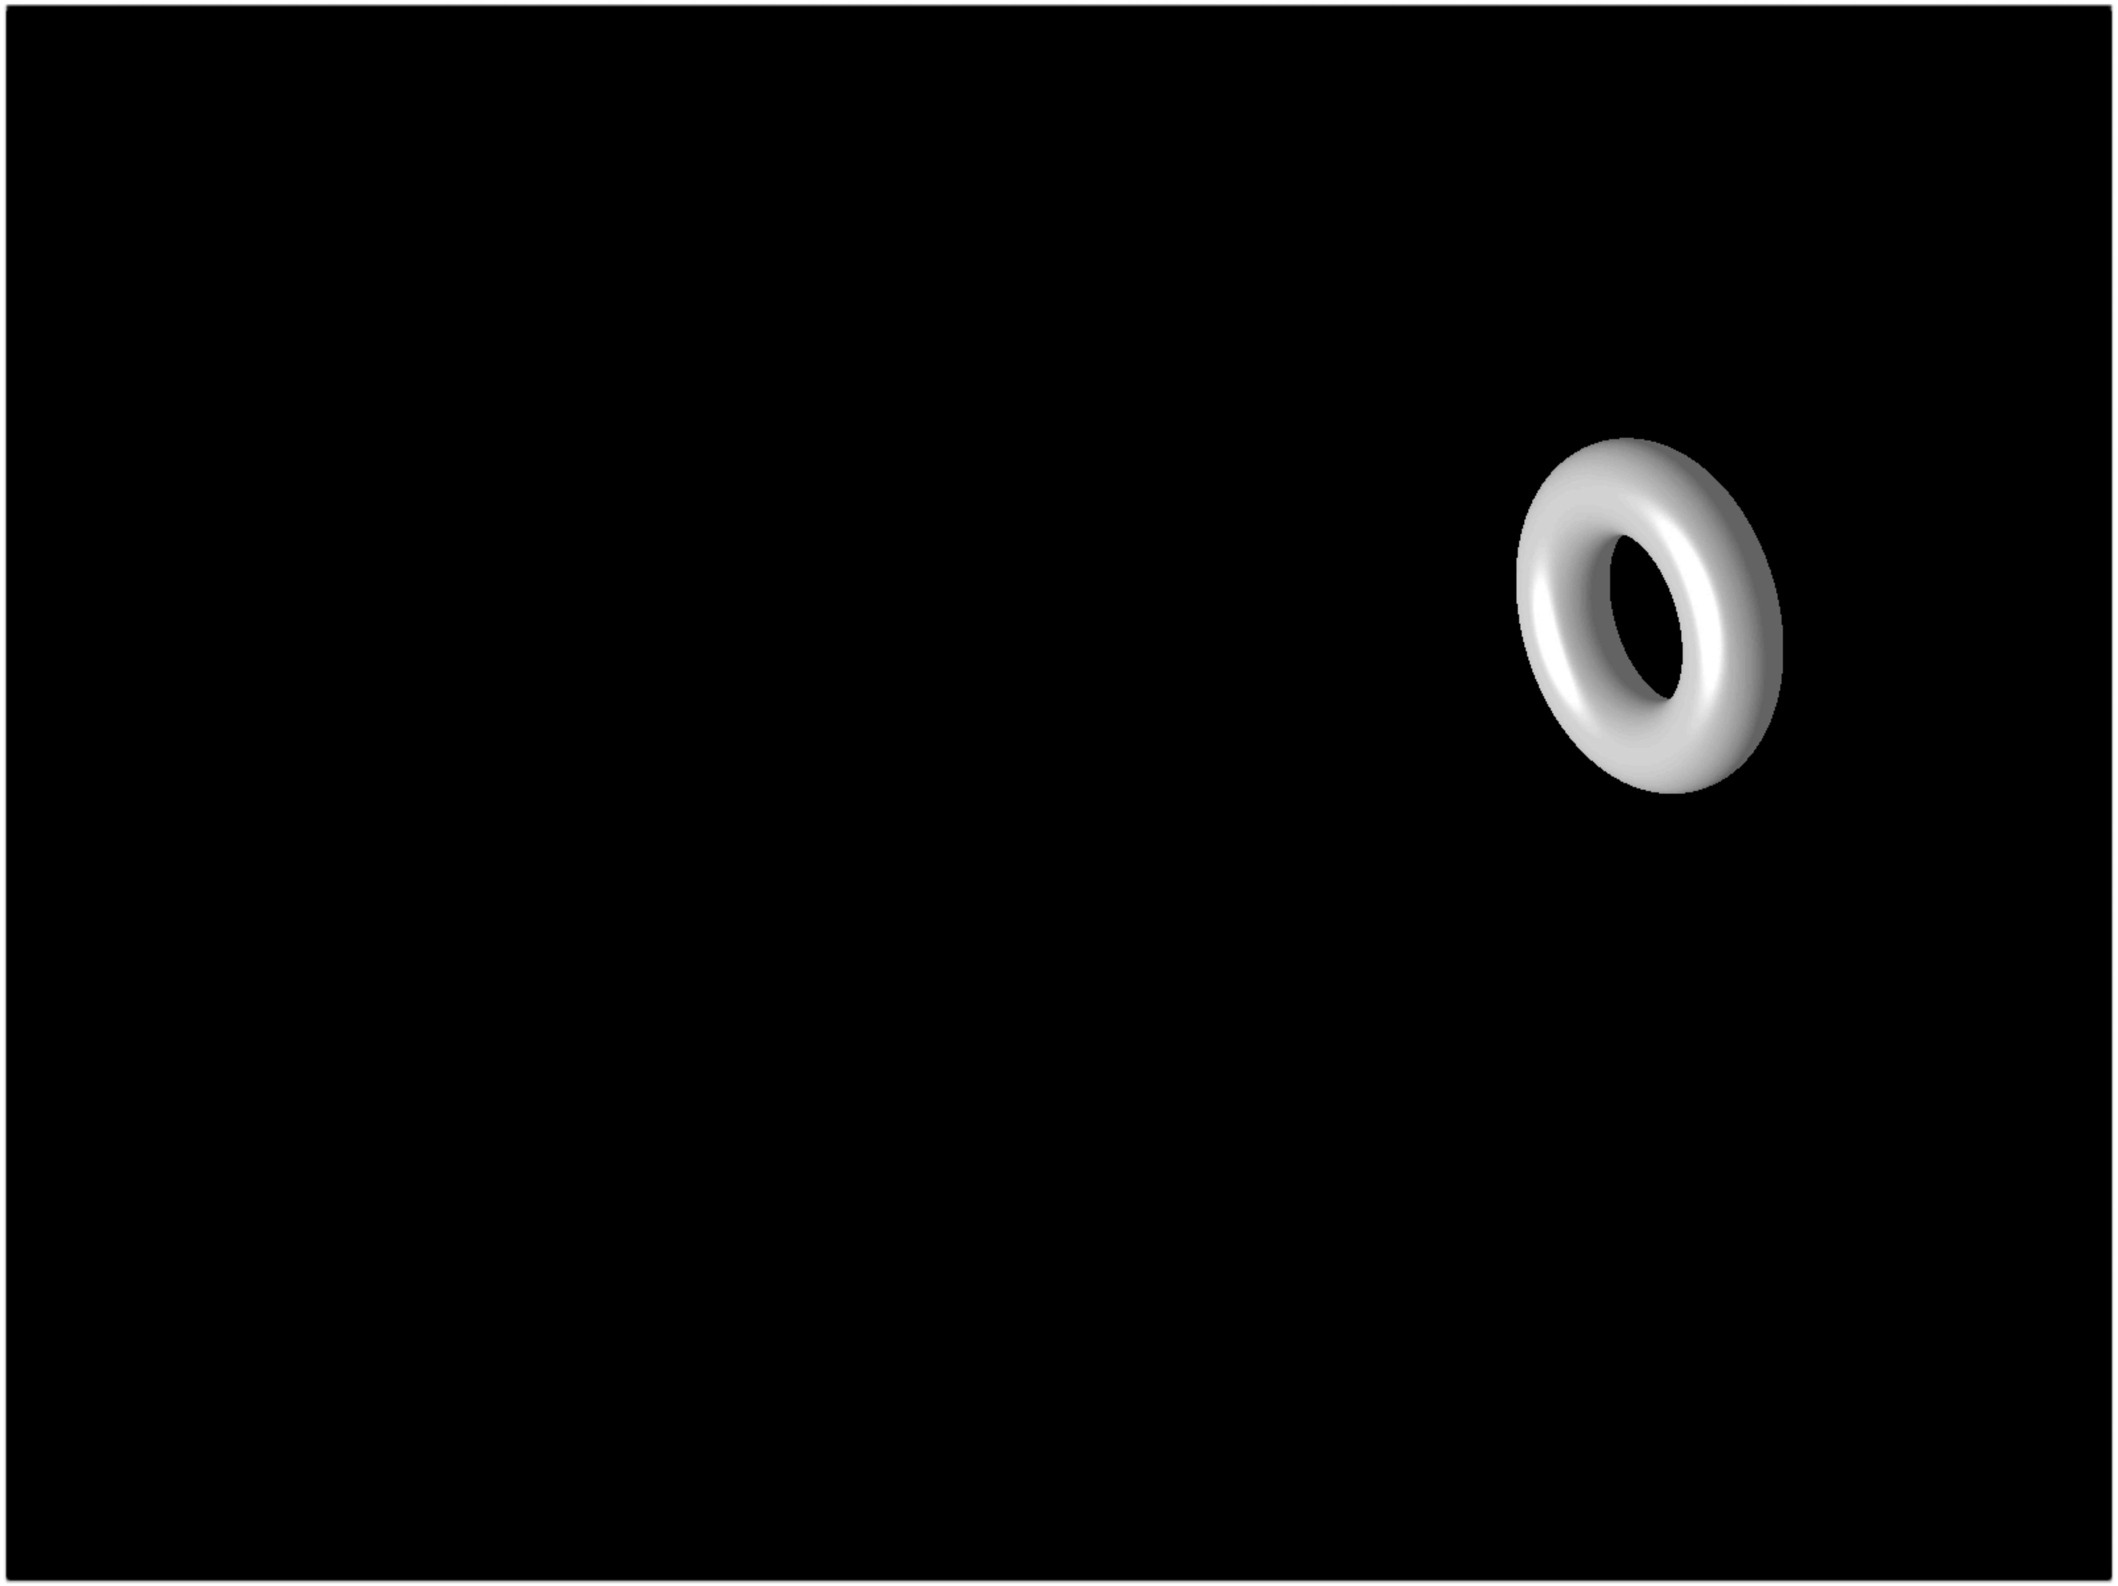

Biological and medical simulations

A Model of Sprouting Angiogenesis

Mechanism:

endothelial cells migrate towards source of growth factors

- form cords
- proliferate
- branch / fuse

Growth factor: VEGF

exists in two forms:

- soluble
- bound to the matrix (bVEGF)

Release of bVEGF

endothelial cells secrete proteinases
proteinases cleave bVEGF → soluble

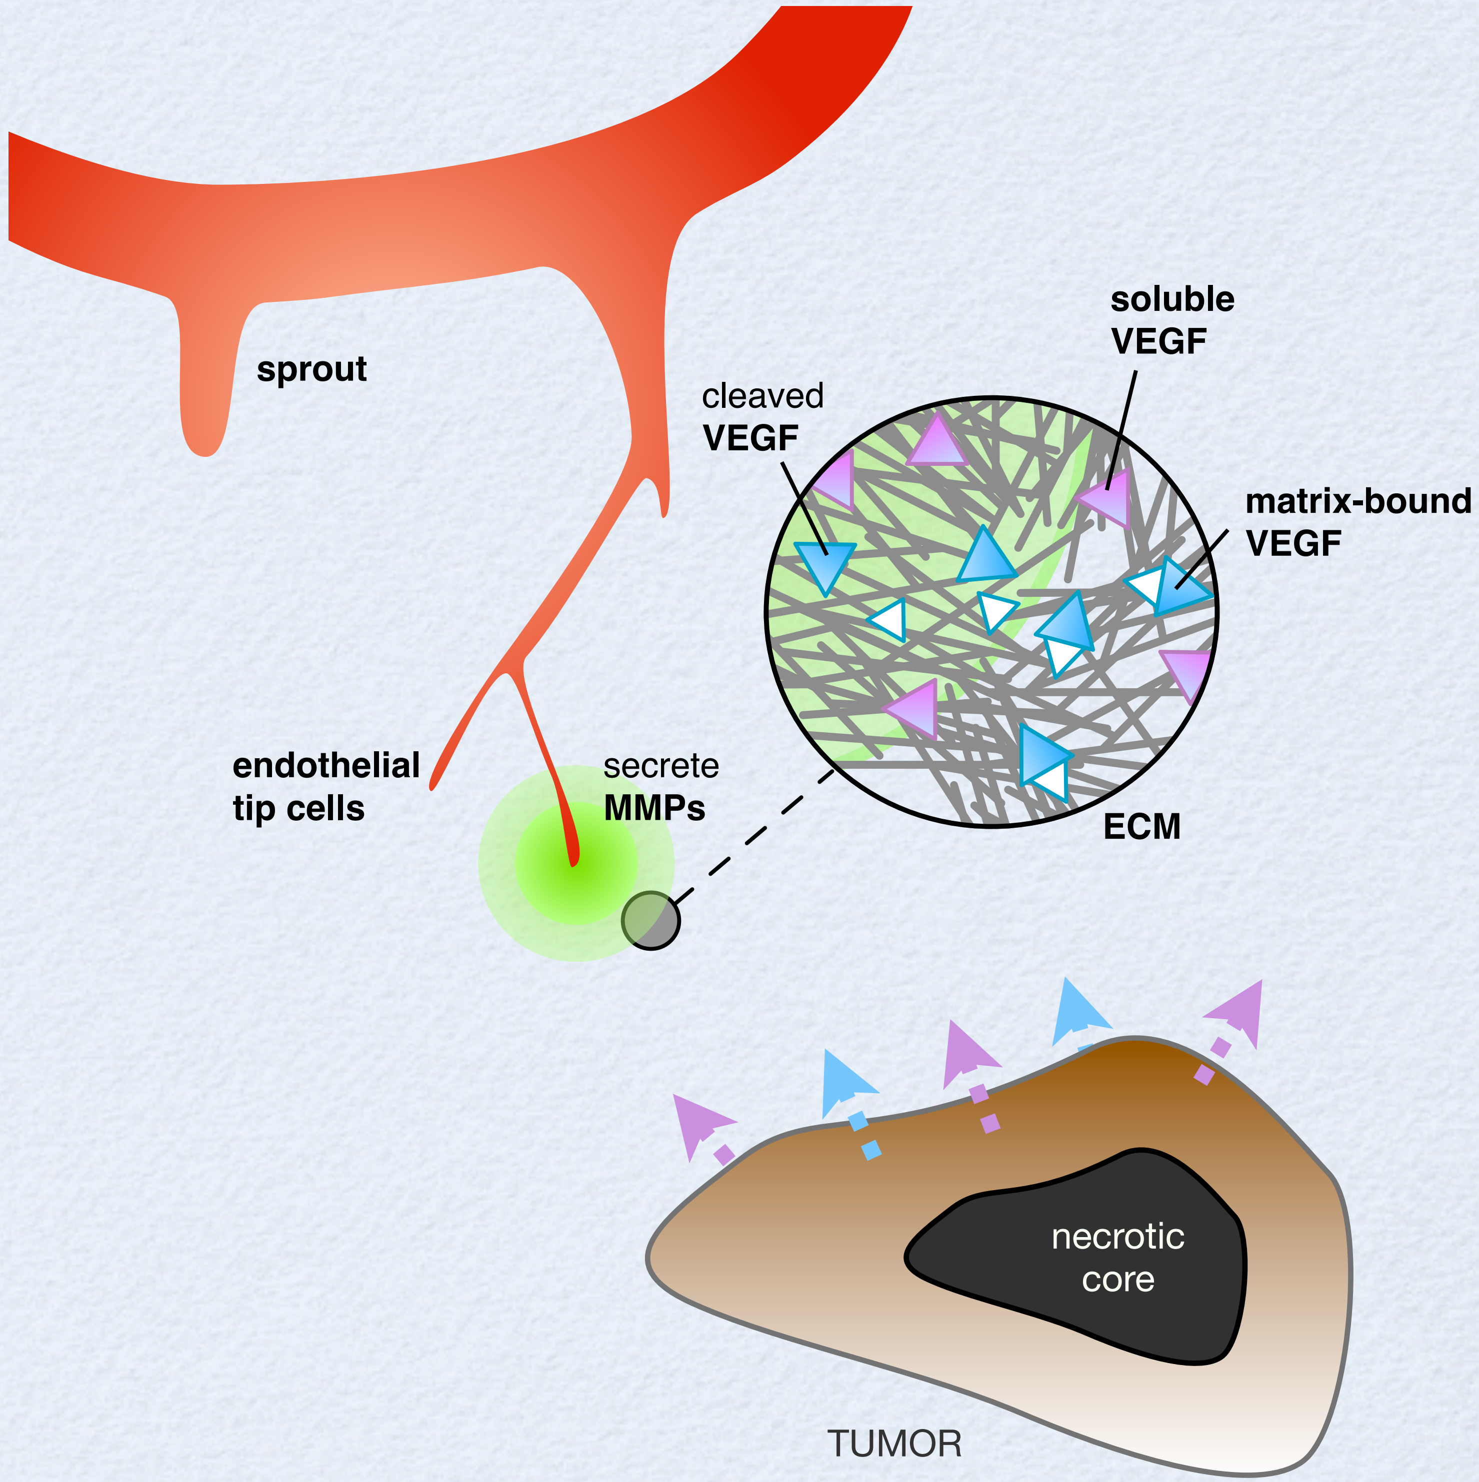

Multi-scale Modeling of Angiogenesis

Milde F., Bergdorf M., Koumoutsakos P., A Hybrid Model for 3D Simulations of Sprouting Angiogenesis, **Biophysical J.**,2008

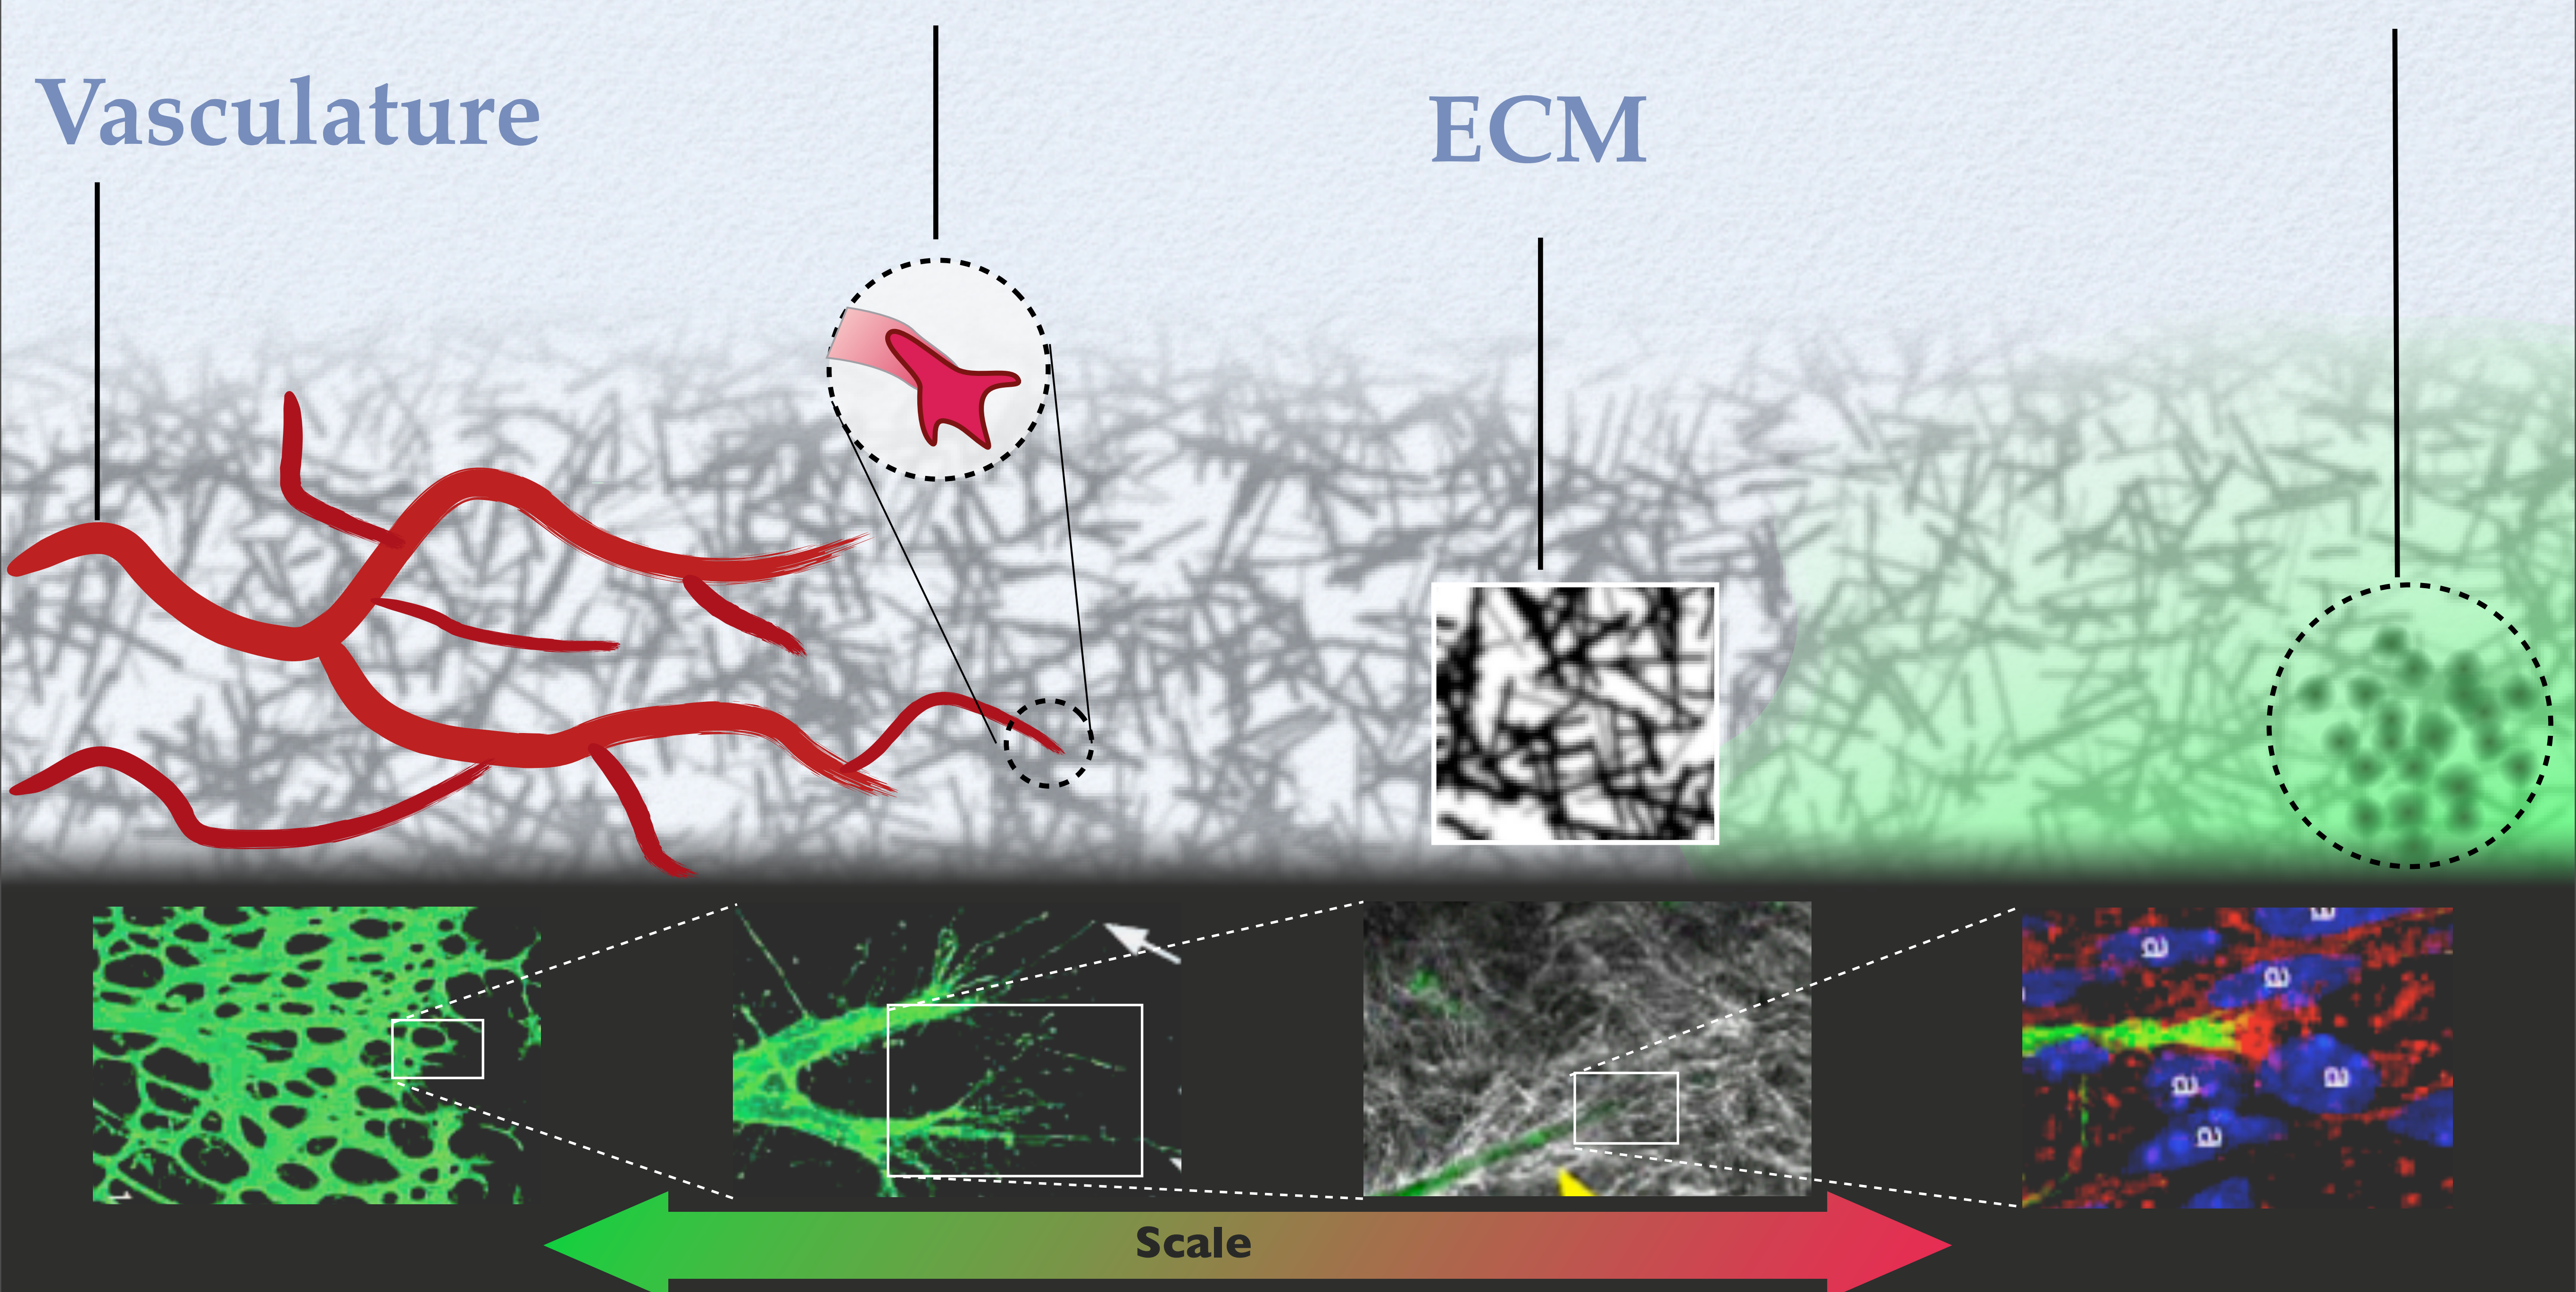

[1] H. GERHARDT, M. GOLDING, M.FRUTTIGER, C. RUHRBERG, A. LUNDKVIST A. ABRAMSSON, M. JELTSCH C. MICHELL, .ALITALO, D. SHIMA AND C. BETSHOLTZ, VEGF GUIDES ANGIOGENIC SPROUTING UTILIZING ENDOTHELIAL TIP CELL FILOPODIA, J. CELL. BIOL., 2003

Modeling the Matrix

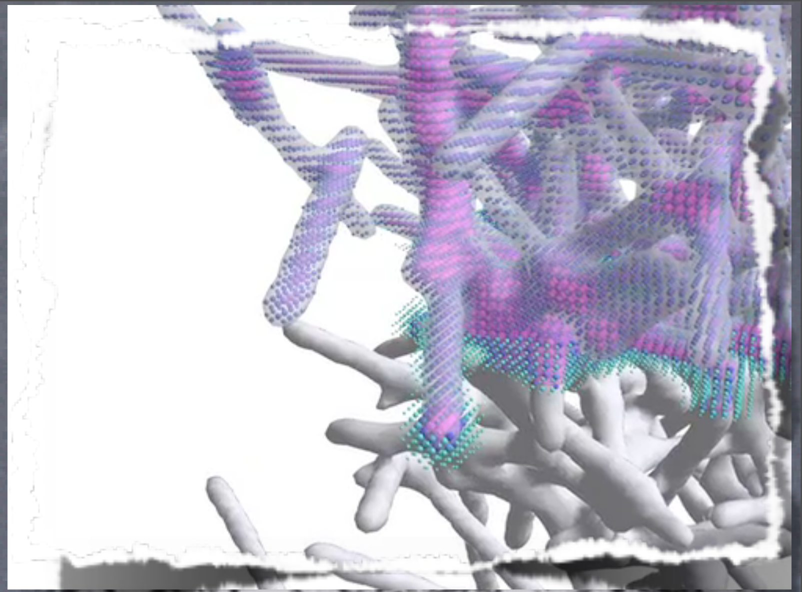

Fibers:

- straight
- random direction
- distribution of lengths

$$l = l_0 \, 2^{m\,z}$$
$$\alpha \in \mathcal{U}([0,\pi])$$
$$z \in \mathcal{N}(0,1)$$

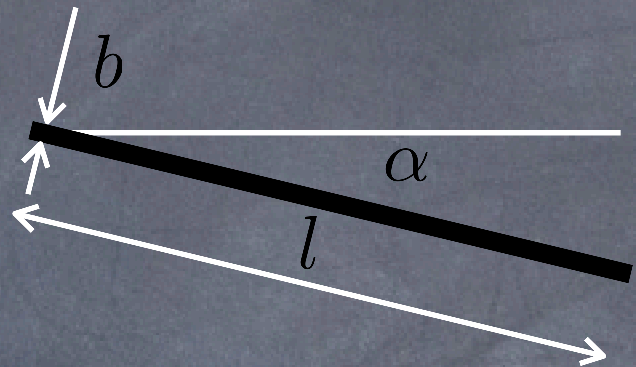

Indicator field $: e$

- unity where fibers present
- smoothed (implicit filopodia)

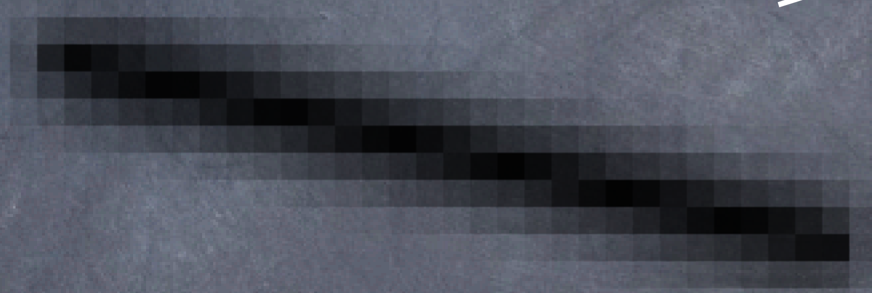

Randomly oriented collagen fibrils in cartilage ECM imaged by TEM.
http://www.shcc.org

Angiogenesis : *in silico*

Effect of Matrix structure on branching

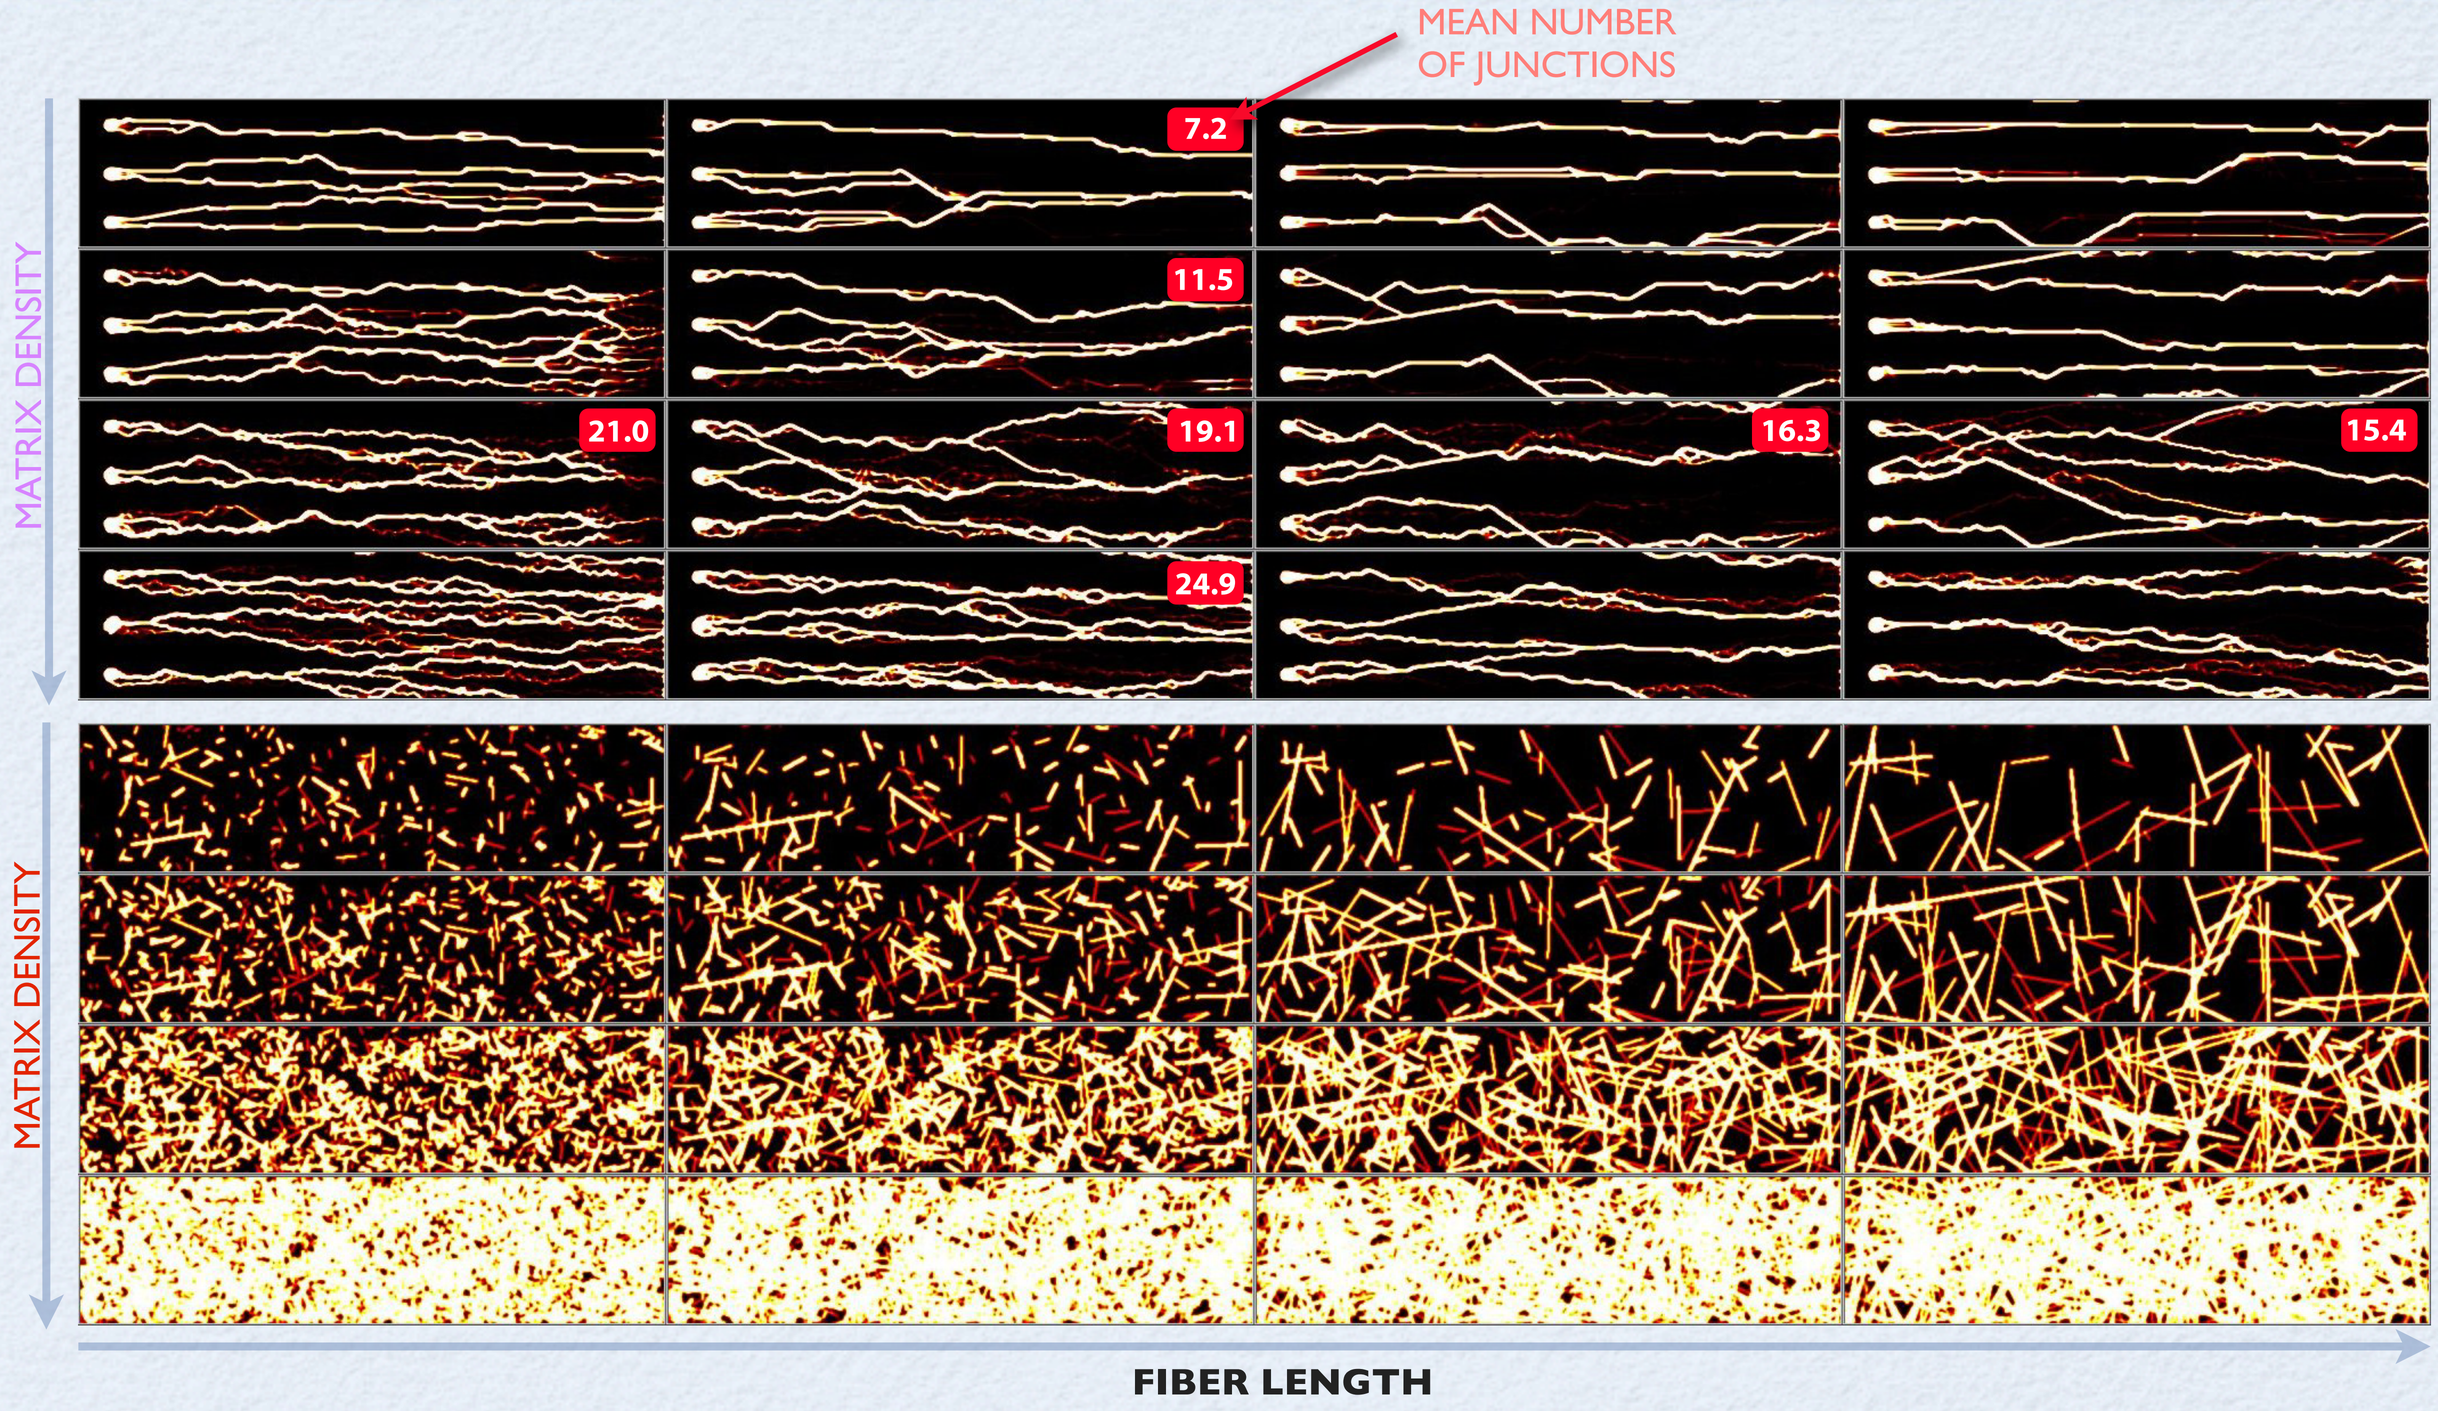

statistics over n = 50 different matrices
junctions identified with AngioQuant

Spatially Adaptive Stochastic Simulations of Gliomas

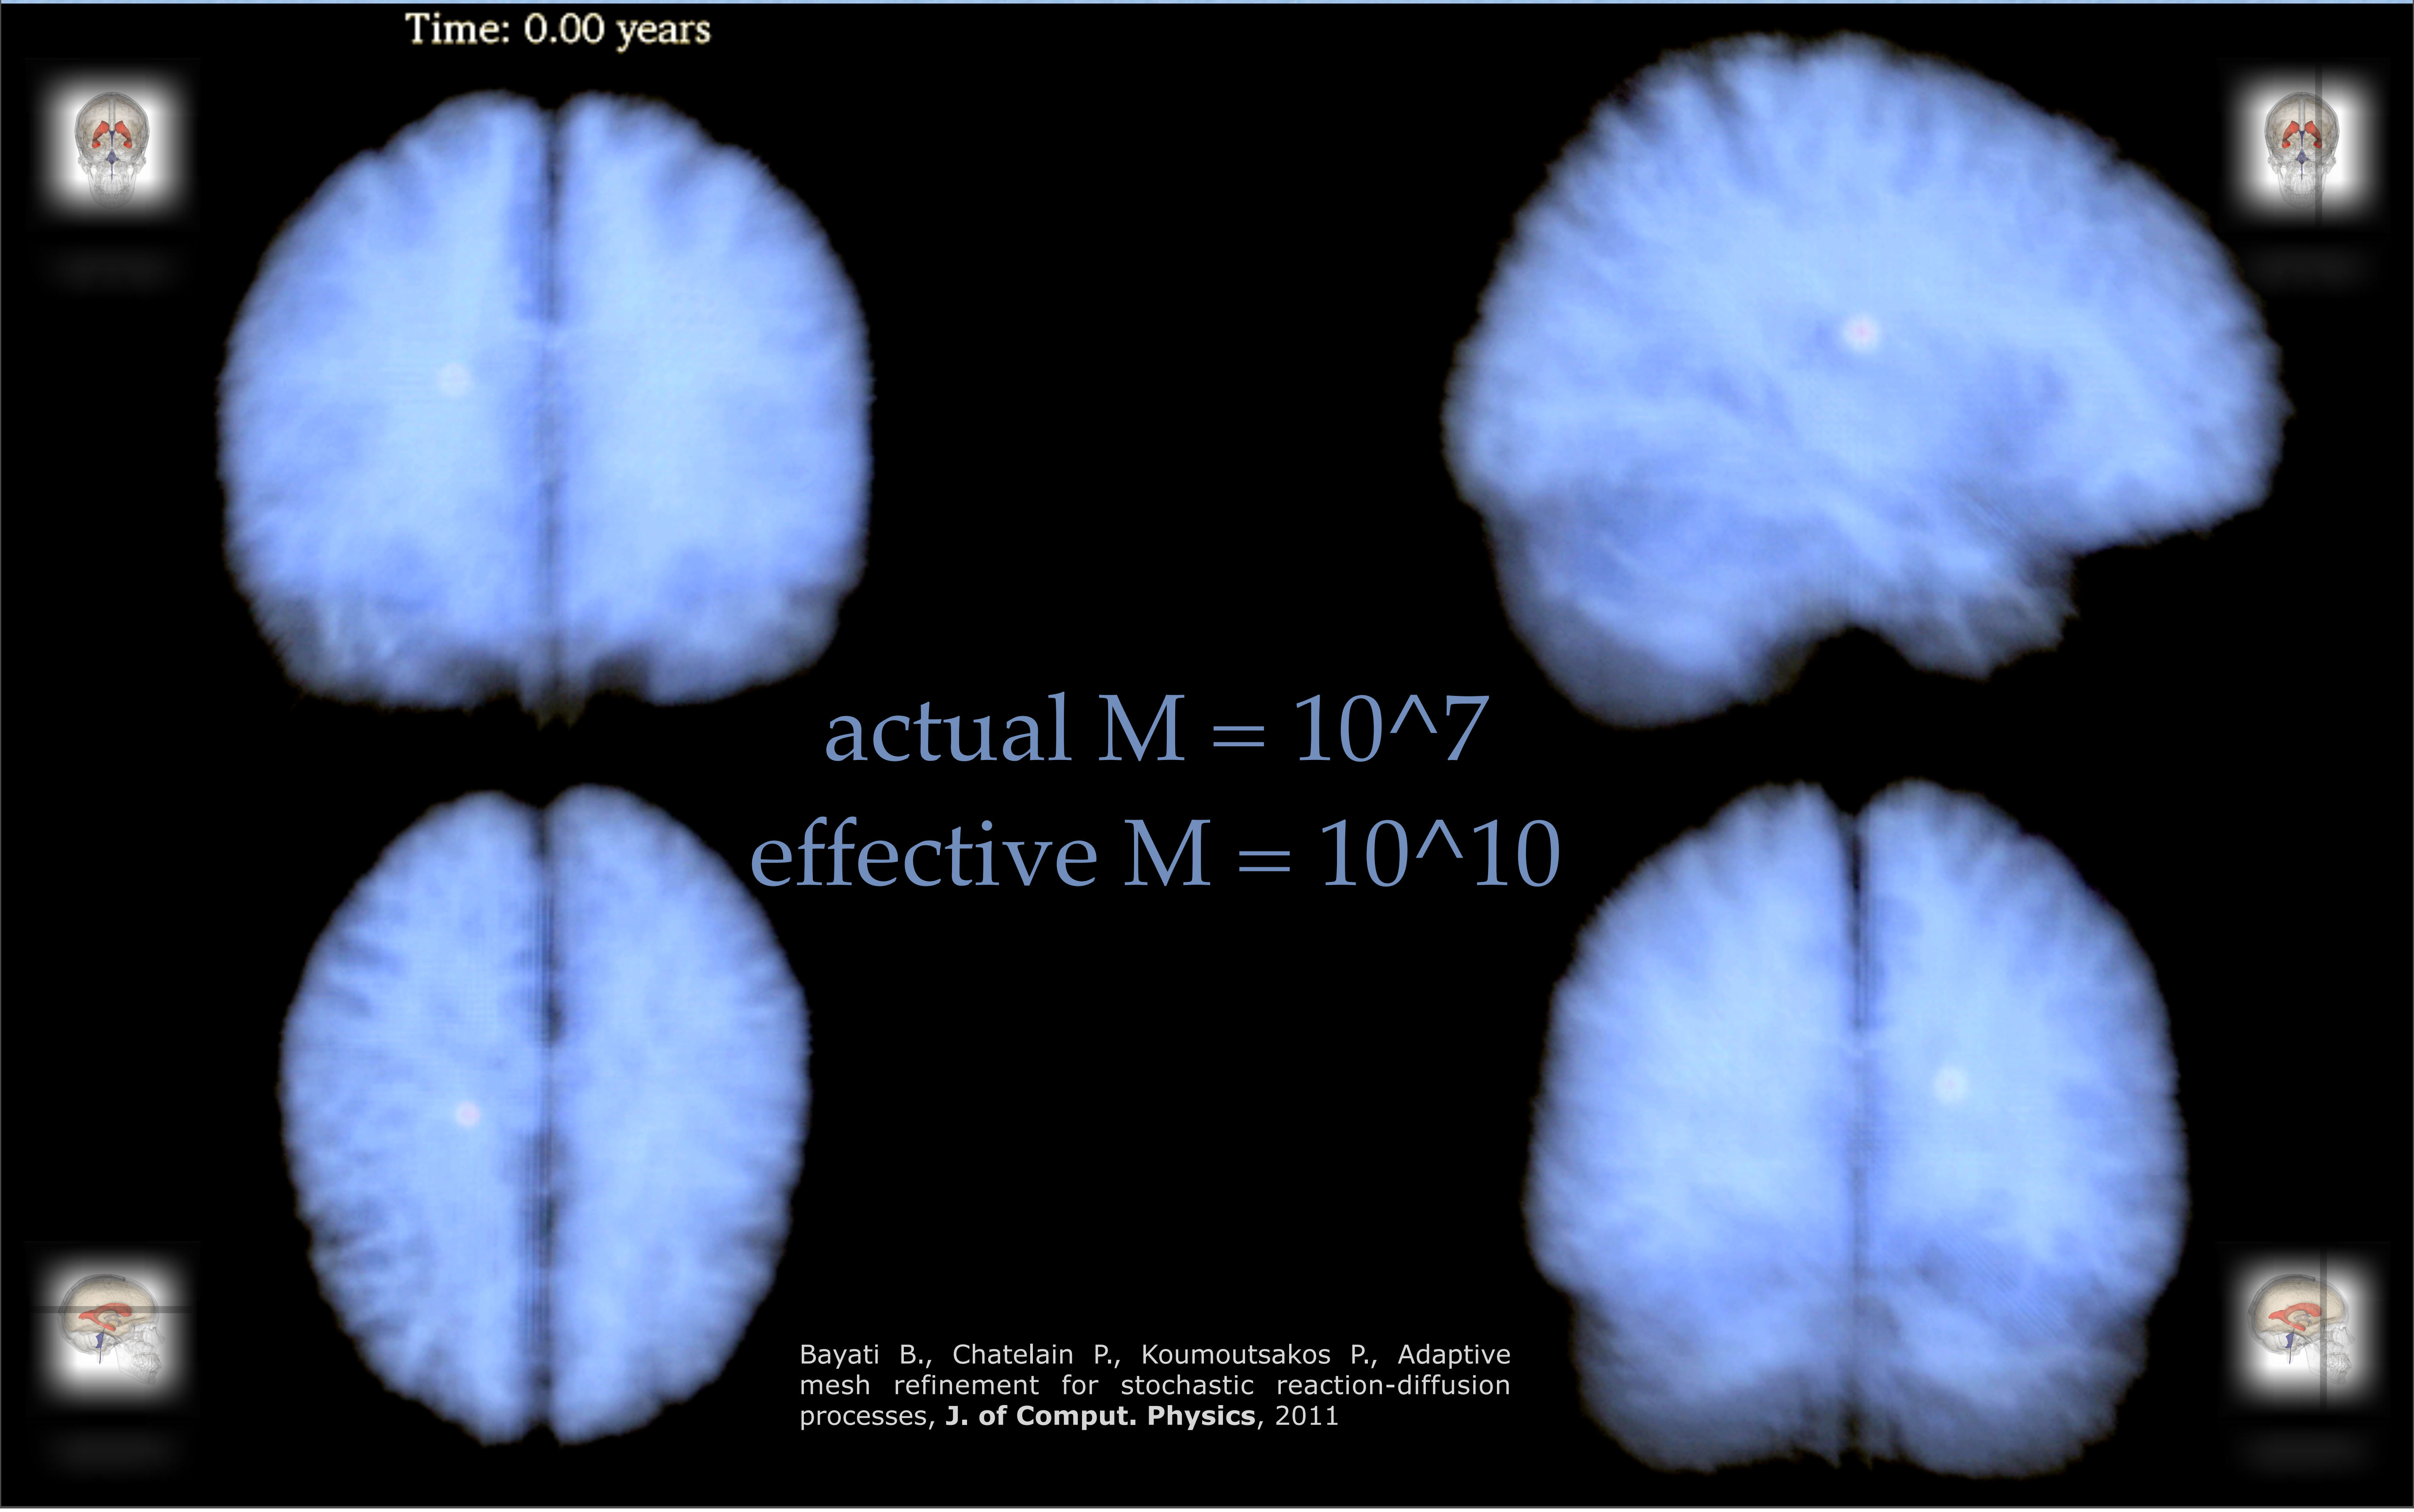

actual M = 10^7

effective M = 10^10

Bayati B., Chatelain P., Koumoutsakos P., Adaptive mesh refinement for stochastic reaction-diffusion processes, **J. of Comput. Physics**, 2011

Last Words

CHALLENGES

- Fast and/or Green Multi-scale Algorithms
- SIMULATIONS ARE DATA : UQ+P

APPLICATIONS

- Biology, Nanotechnology and Fluids : Bridge Gaps and Disciplines

THANKS

- ETHZ + CSCS
- Swiss National Science Foundation
- EU
- NVIDIA (ETHZ a CUDA Research Center)


Walther
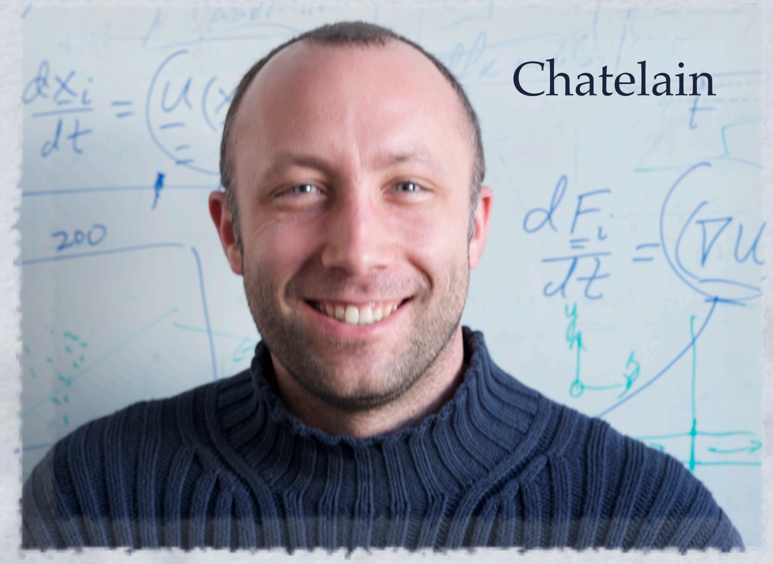
Chatelain

Bayati

Bergdorf

Rossinelli

Gazzola

Hedjazialhosseini

van Rees

Milde