GPU Computing in the Cloud. A future of HPC?

Lukasz Miroslaw, Ph.D.
CEO
Vratis LTD

3rd PRACE Meeting, Stockholm, 27-28.03.2011

Contents

- A few words about Vratis and SpeedIT technology.
- GPU Computing in the cloud from SME perspective.

Vratis

- Founded in 2006 as spin-off of Wroclaw University Of Technology.
- Main activity : image processing, data mining and pattern recognition in biology, medicine and industry, **hardware acceleration in science.**
- Collaboration
  - SlidePath Ltd. (Digital Pathology)
  - Visimatik Ltd., ETH Zurich (Image Processing)
  - Engys, IT'IS Foundation, Accelereyes (**SpeedIT**)

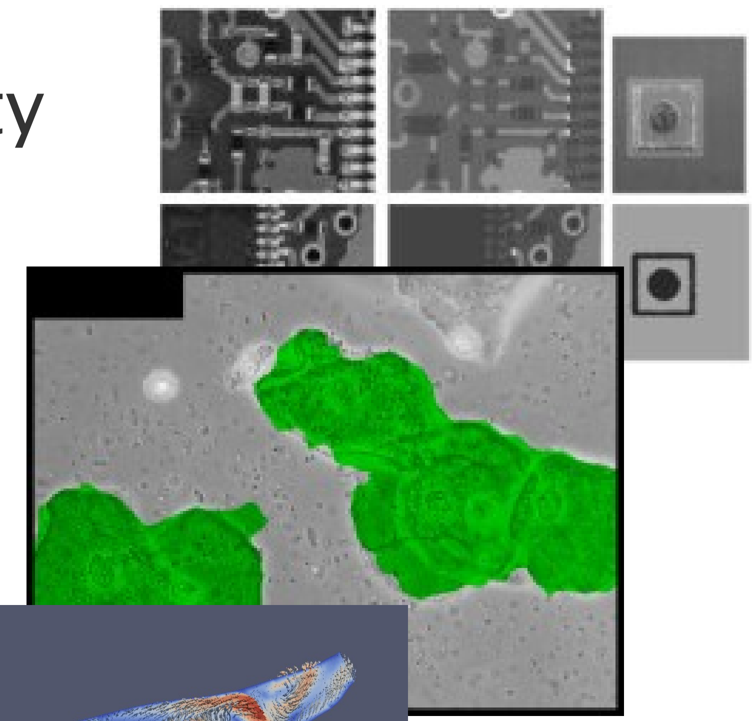

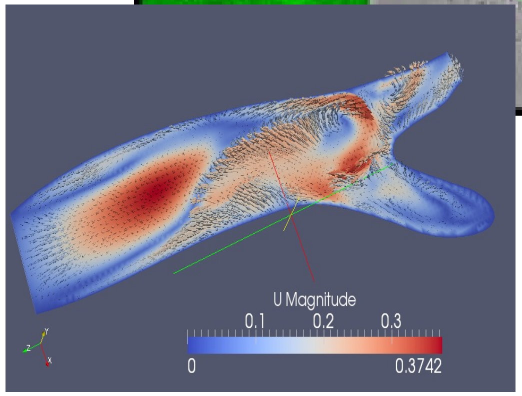

- Product released in 2009.
- Objective: hardware acceleration of Computational Fluid Dynamics (CFD) simulations.
- Motivation: time, cost and size constraints of HPC systems.
  - Typical simulations need > 16 GB space
  - Hospitals have no infrastructure for expensive clusters and supercomputers.
  - Typical duration of patient-specific simulations: several days or weeks.

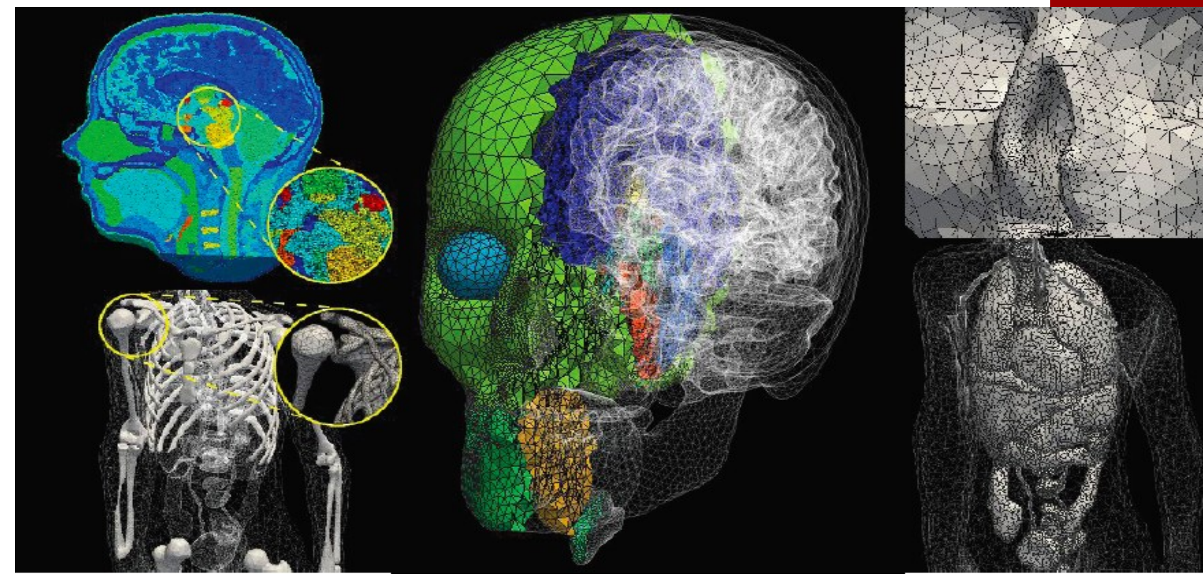

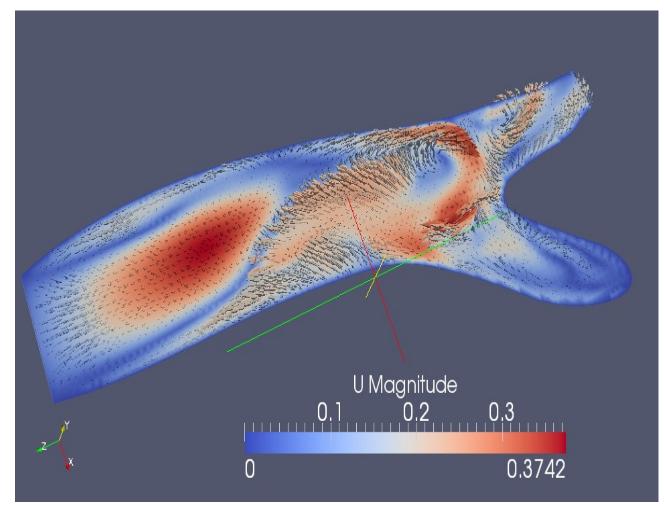

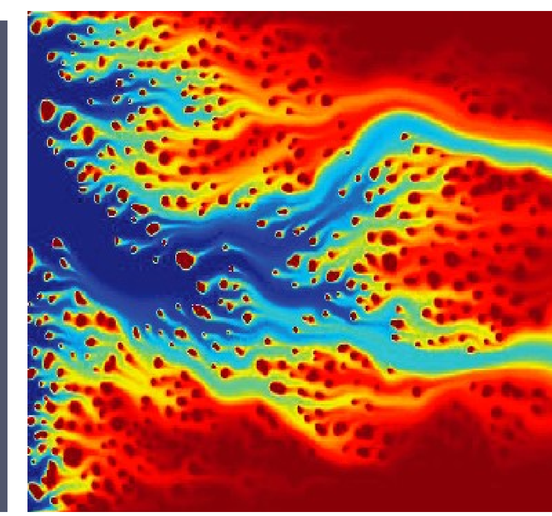

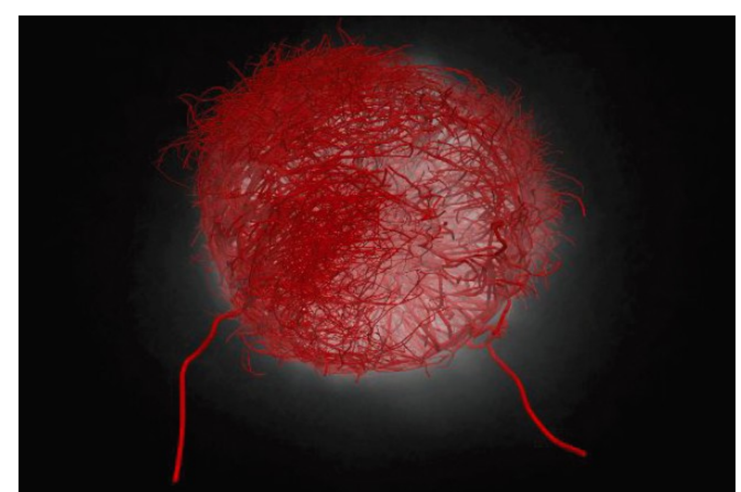

SpeedIT Cloud

Proposed solution: hardware accelerator that can be placed close to the patient.

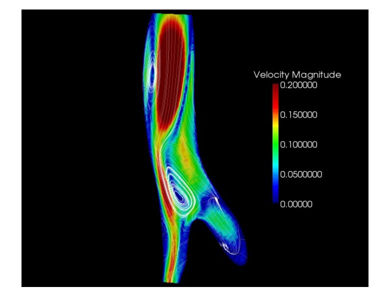

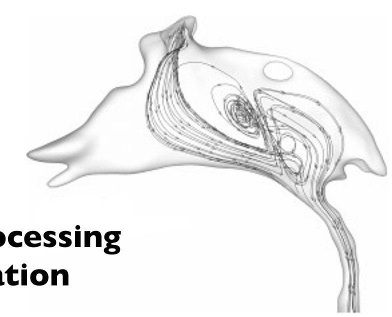

Data acquisition (MRI, CT)
Determination of the geometry

Model building
Solving of PDE Equations

Post-processing
Visualisation

Partners:

Vratis

IT IV FOUNDATION

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

UNIVERSITAS WRATISLAVIENSIS

SpeedIT Toolkit

- GPU-based library of iterative solvers commonly used in CFD simulations:
  - BiCGSTAB for non-symmetric matrices.
  - CG for symmetric matrices.
  - Preconditioners: diagonal, complex operations.
  - OpenFOAM plugin that supports multi-gpu systems.
  - Classic and eXtreme version.

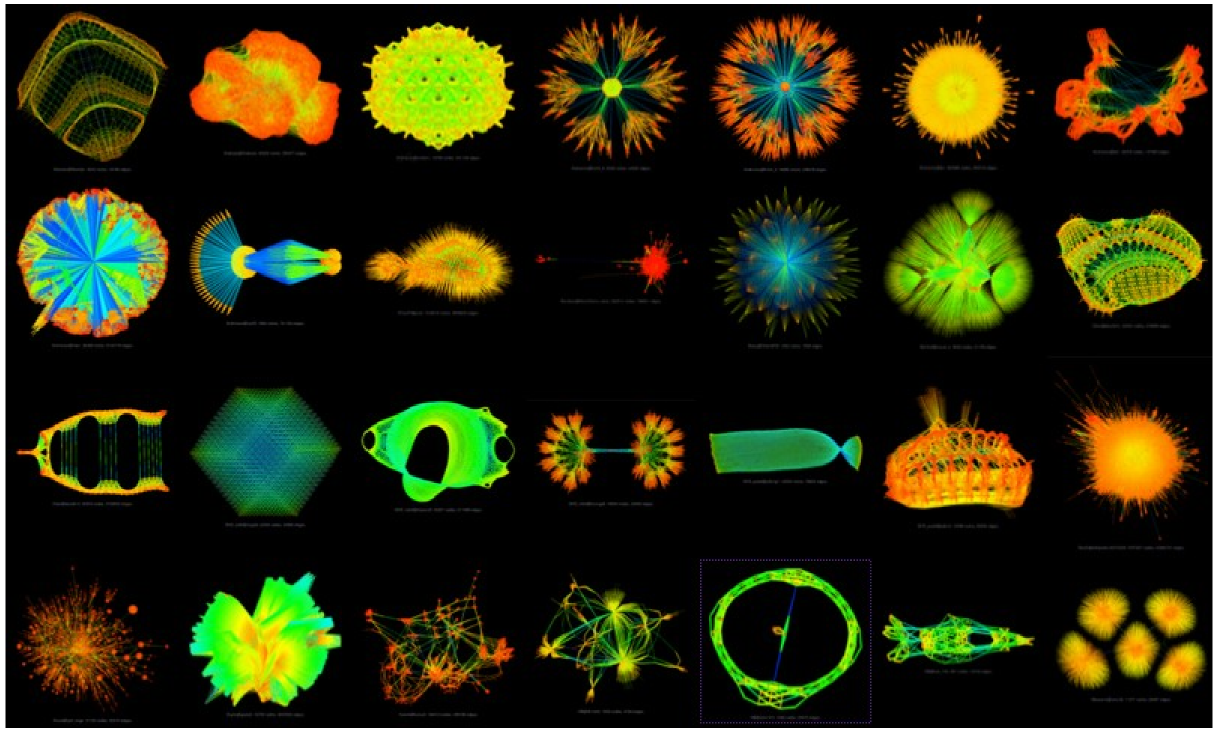

- Testing environment:
  - The University of Florida Sparse Matrix Collection
  - and matrices from our CFD problems
  - single & double precision, CSR matrix format,
  - performance with & without transfer times
  - GPU GTX 275, CPU Intel Core 2 Quad 2.66 GHz

Test matrices

Matrix name	Number of nonzero elements	Number of rows	Memory double precision [MB]	Memory single precision [MB]
af_shell10	52672325	1508065	608,54	407,61
apache1	542184	80800	6,51	4,44
audikw_1	77651847	943695	892,25	596,04
bmw3_2	11288630	227362	130,06	86,99
bmwcra_1	10644002	148770	122,38	81,77
bone010	71666325	986703	823,92	550,53
bone010_M	23888775	986703	277,15	186,02
boneS10	55468422	914898	638,28	426,68
boneS10_M	18489474	914898	215,09	144,55
cage14	27130349	1505785	316,23	212,73
crankseg_1	10614210	52804	121,67	81,18
crankseg_2	14148858	63838	162,16	108,19

Matrix name	Number of nonzero elements	Number of rows	Memory double precision [MB]	Memory single precision [MB]
F1	26837113	343791	308,44	206,06
Freescale1	18920347	3428755	229,61	157,43
hood	10768436	220542	124,08	83,00
inline_1	36816342	503712	423,25	282,81
ldoor	46522475	952203	536,04	358,57
msdoor	20240935	415863	233,23	156,01
nd12k	14220946	36000	162,88	108,63
nd24k	28715634	72000	328,90	219,36
p029_66k_run9	10646210	323029	123,07	82,46
run20	3928604	165911	45,59	30,61
test_nliter	24271103	713969	280,48	187,90

Test matrices

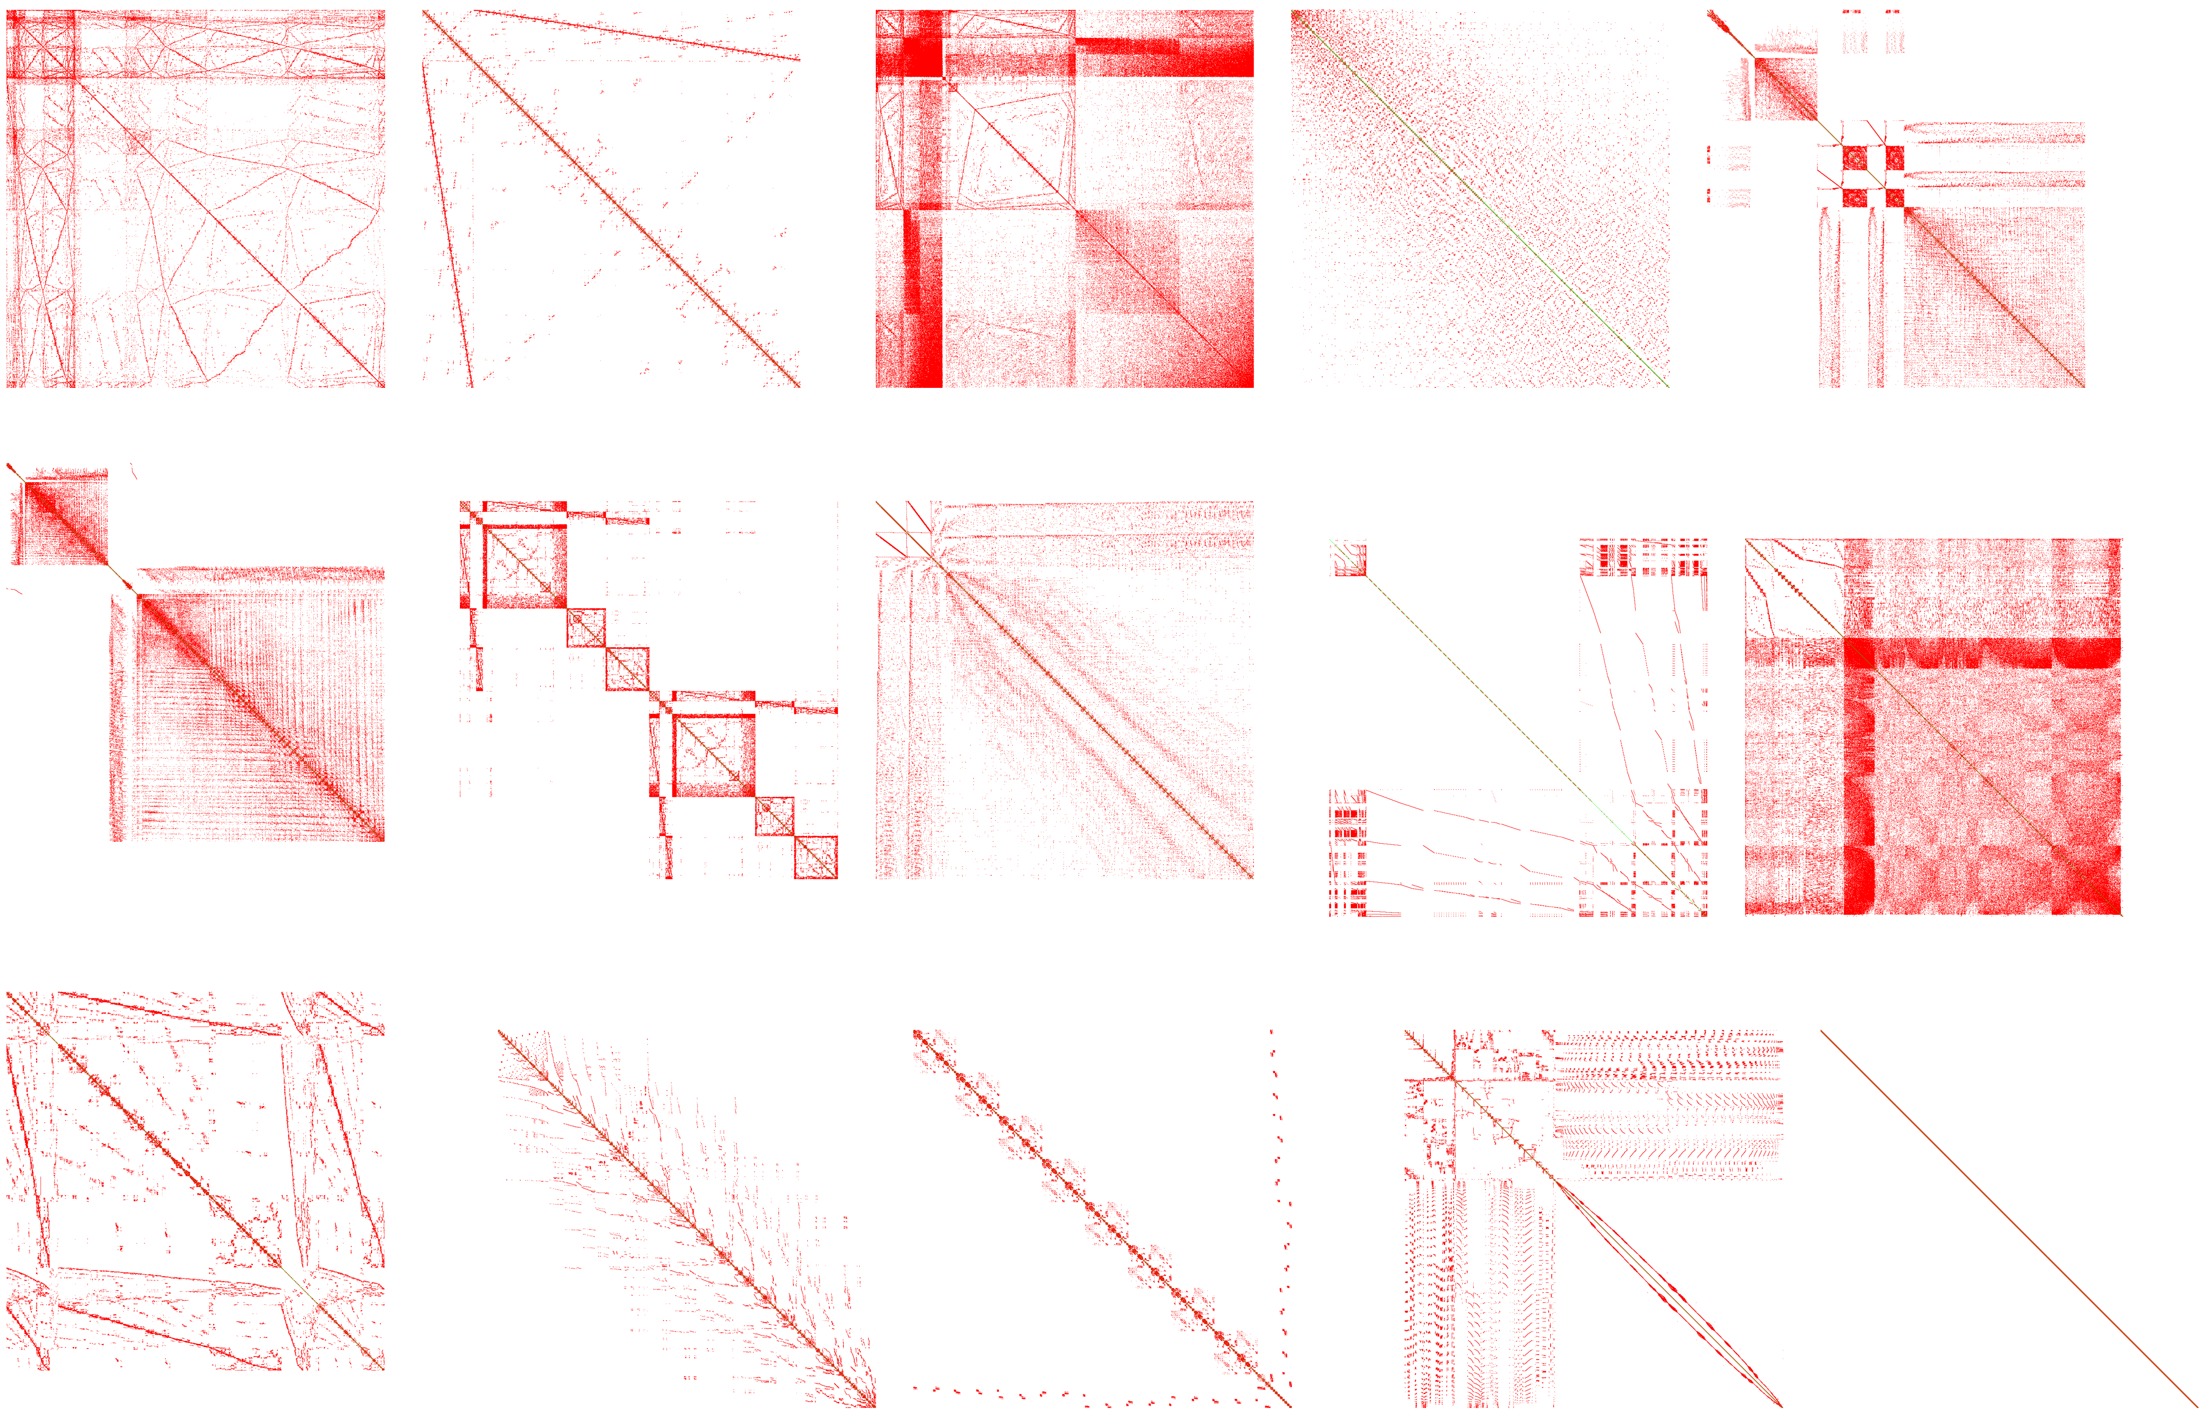

Iterative solvers

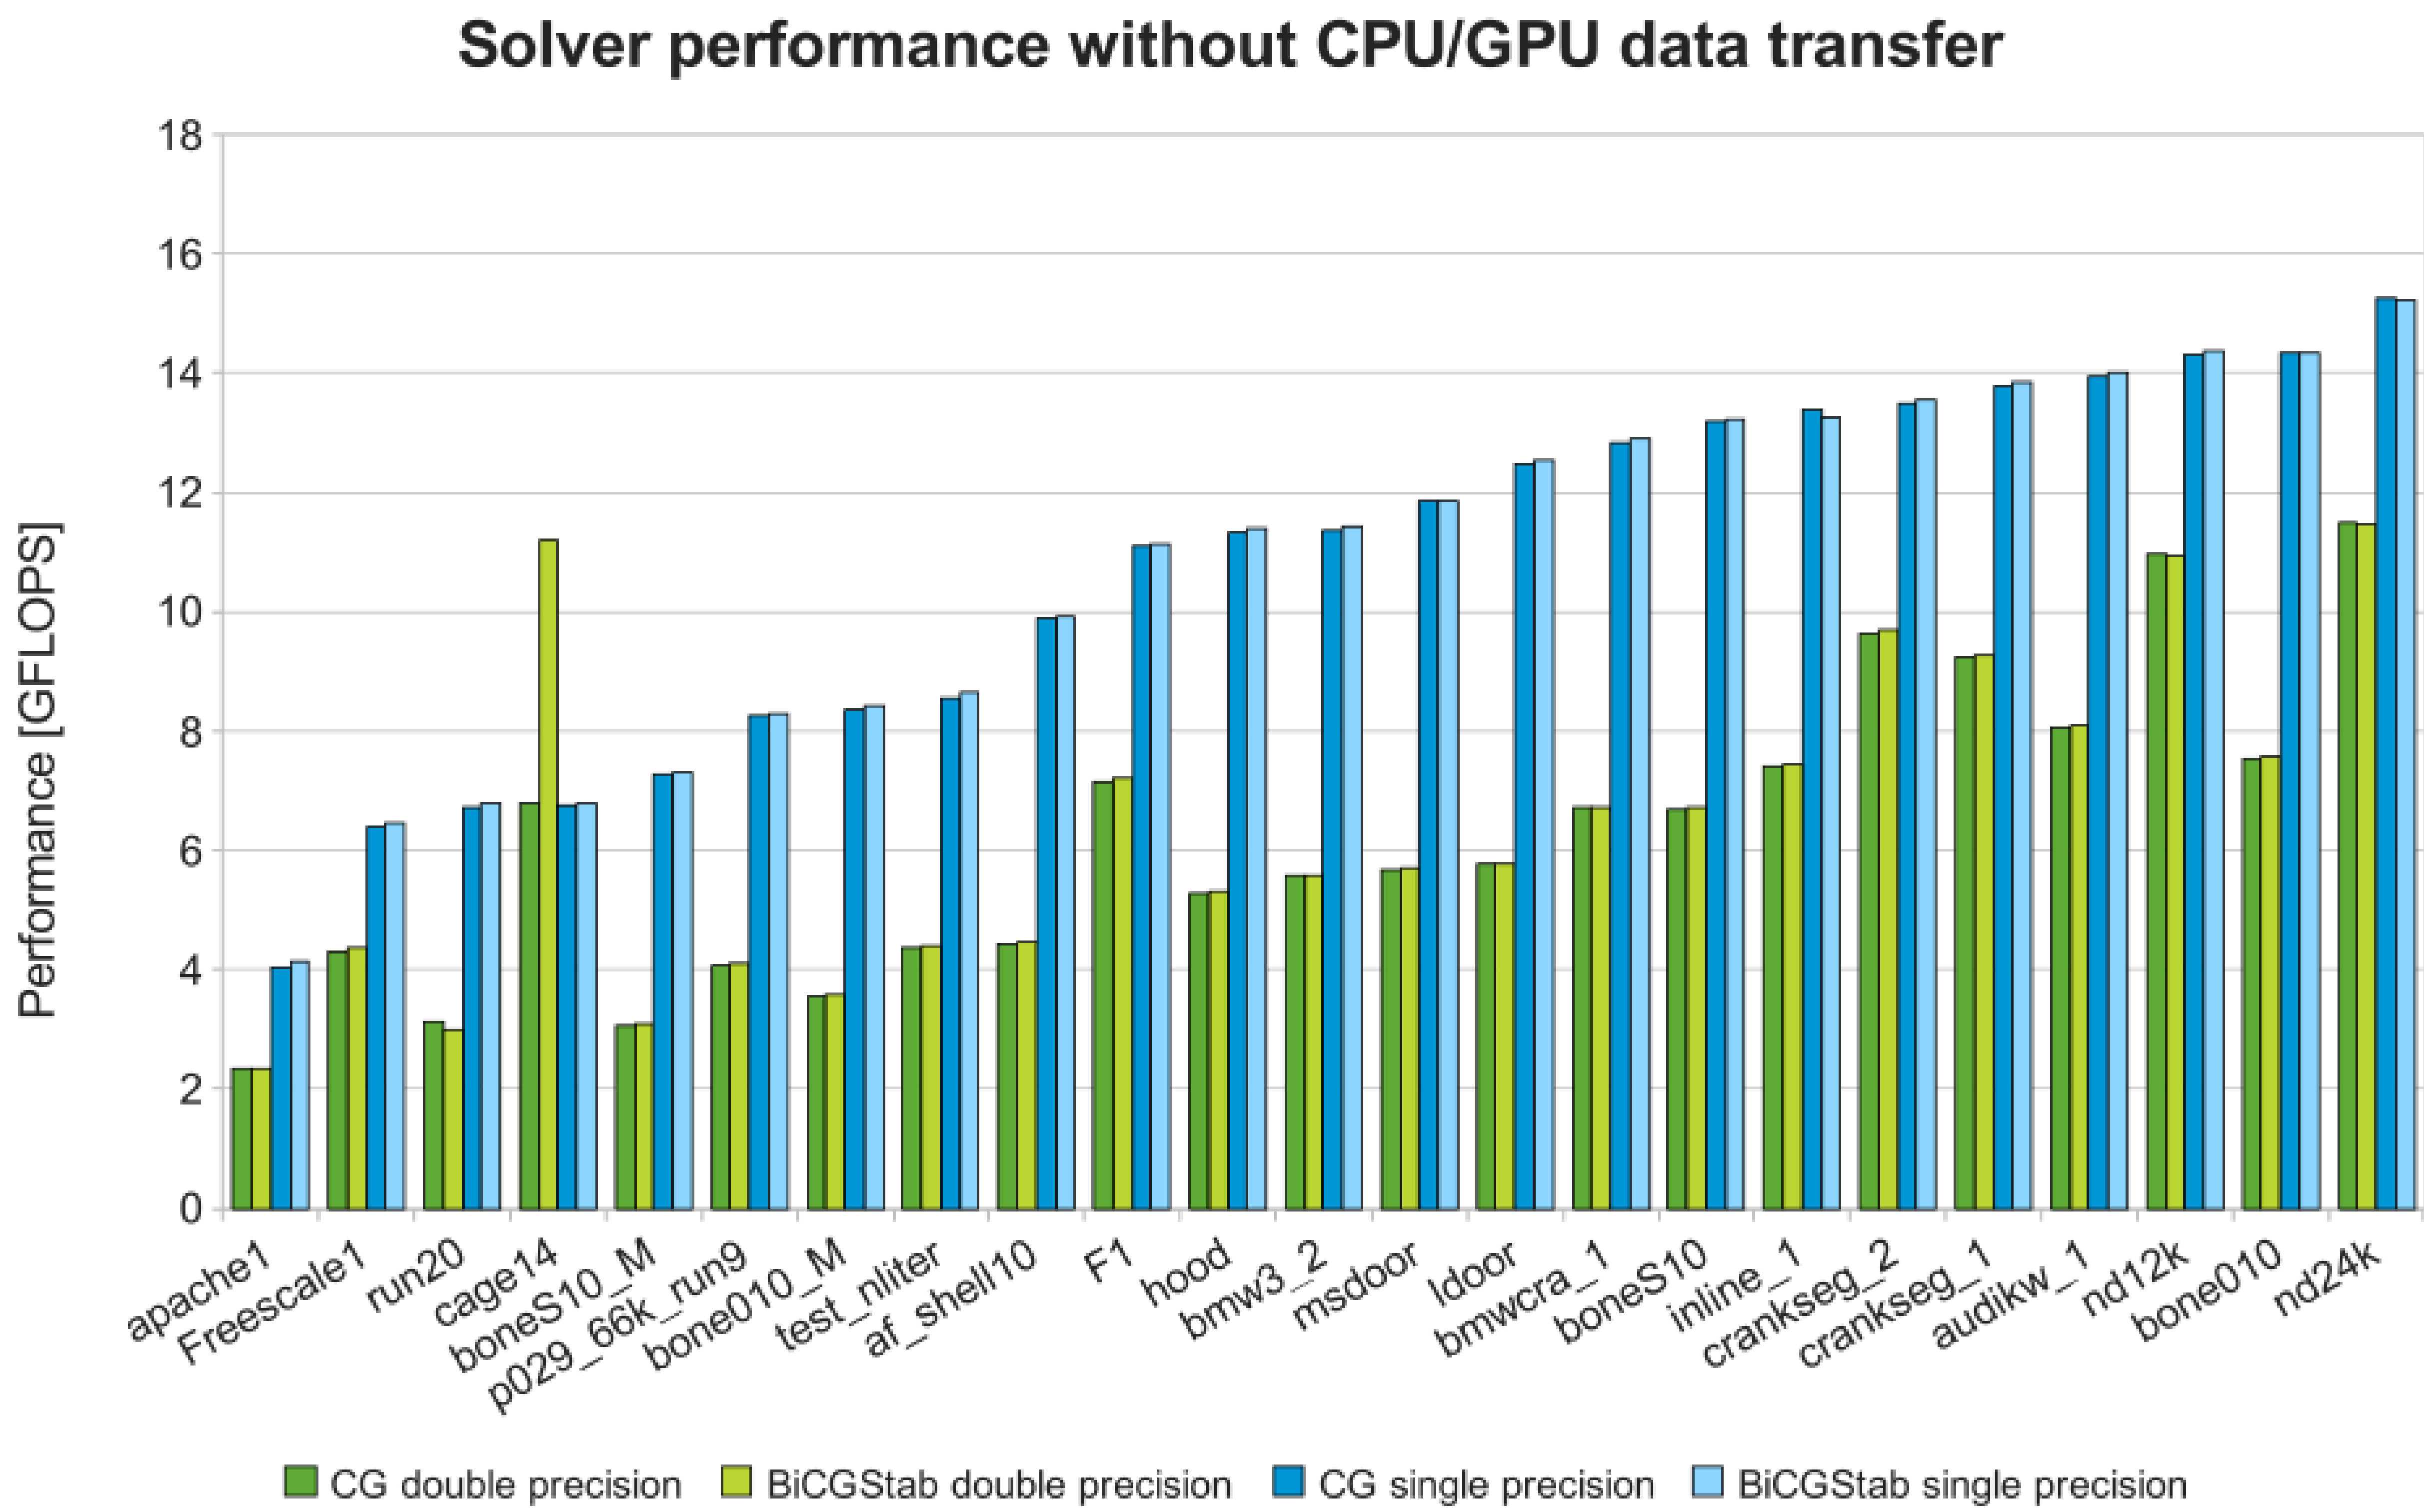

Iterative solvers

Matrix name	CPU implementation: Number of iterations	CPU implementation: Final tolerance	CPU implementation: Time [s]	**SpeedIT Extreme**: si_cdcsrbicgstab: Number of iterations	**SpeedIT Extreme**: si_cdcsrbicgstab: Final tolerance	**SpeedIT Extreme**: si_cdcsrbicgstab: Time [s]	Speedup
apache1	2014	2,88E-11	18,37	2157	5,17E-11	3,51	5,24
bone010_M	33	9,72E-11	9,09	33	9,72E-11	1,45	6,28
boneS10_M	32	7,13E-11	7,56	32	7,13E-11	1,32	5,73
cage14	12	1,10E-11	4,83	10	4,41E-11	0,89	5,41
crankseg_2	545	9,48E-11	67,39	1332	6,83E-11	8,11	8,31
F1	4396	2,88E-11	1339,59	5516	9,97E-11	87,4	15,33
hood	6165	7,67E-11	686,86	9811	9,72E-11	86,99	7,9
run20	615	9,69E-11	30,76	635	7,41E-11	4,23	7,27

SpeedIT 1.2 vs. CUDA 3.2

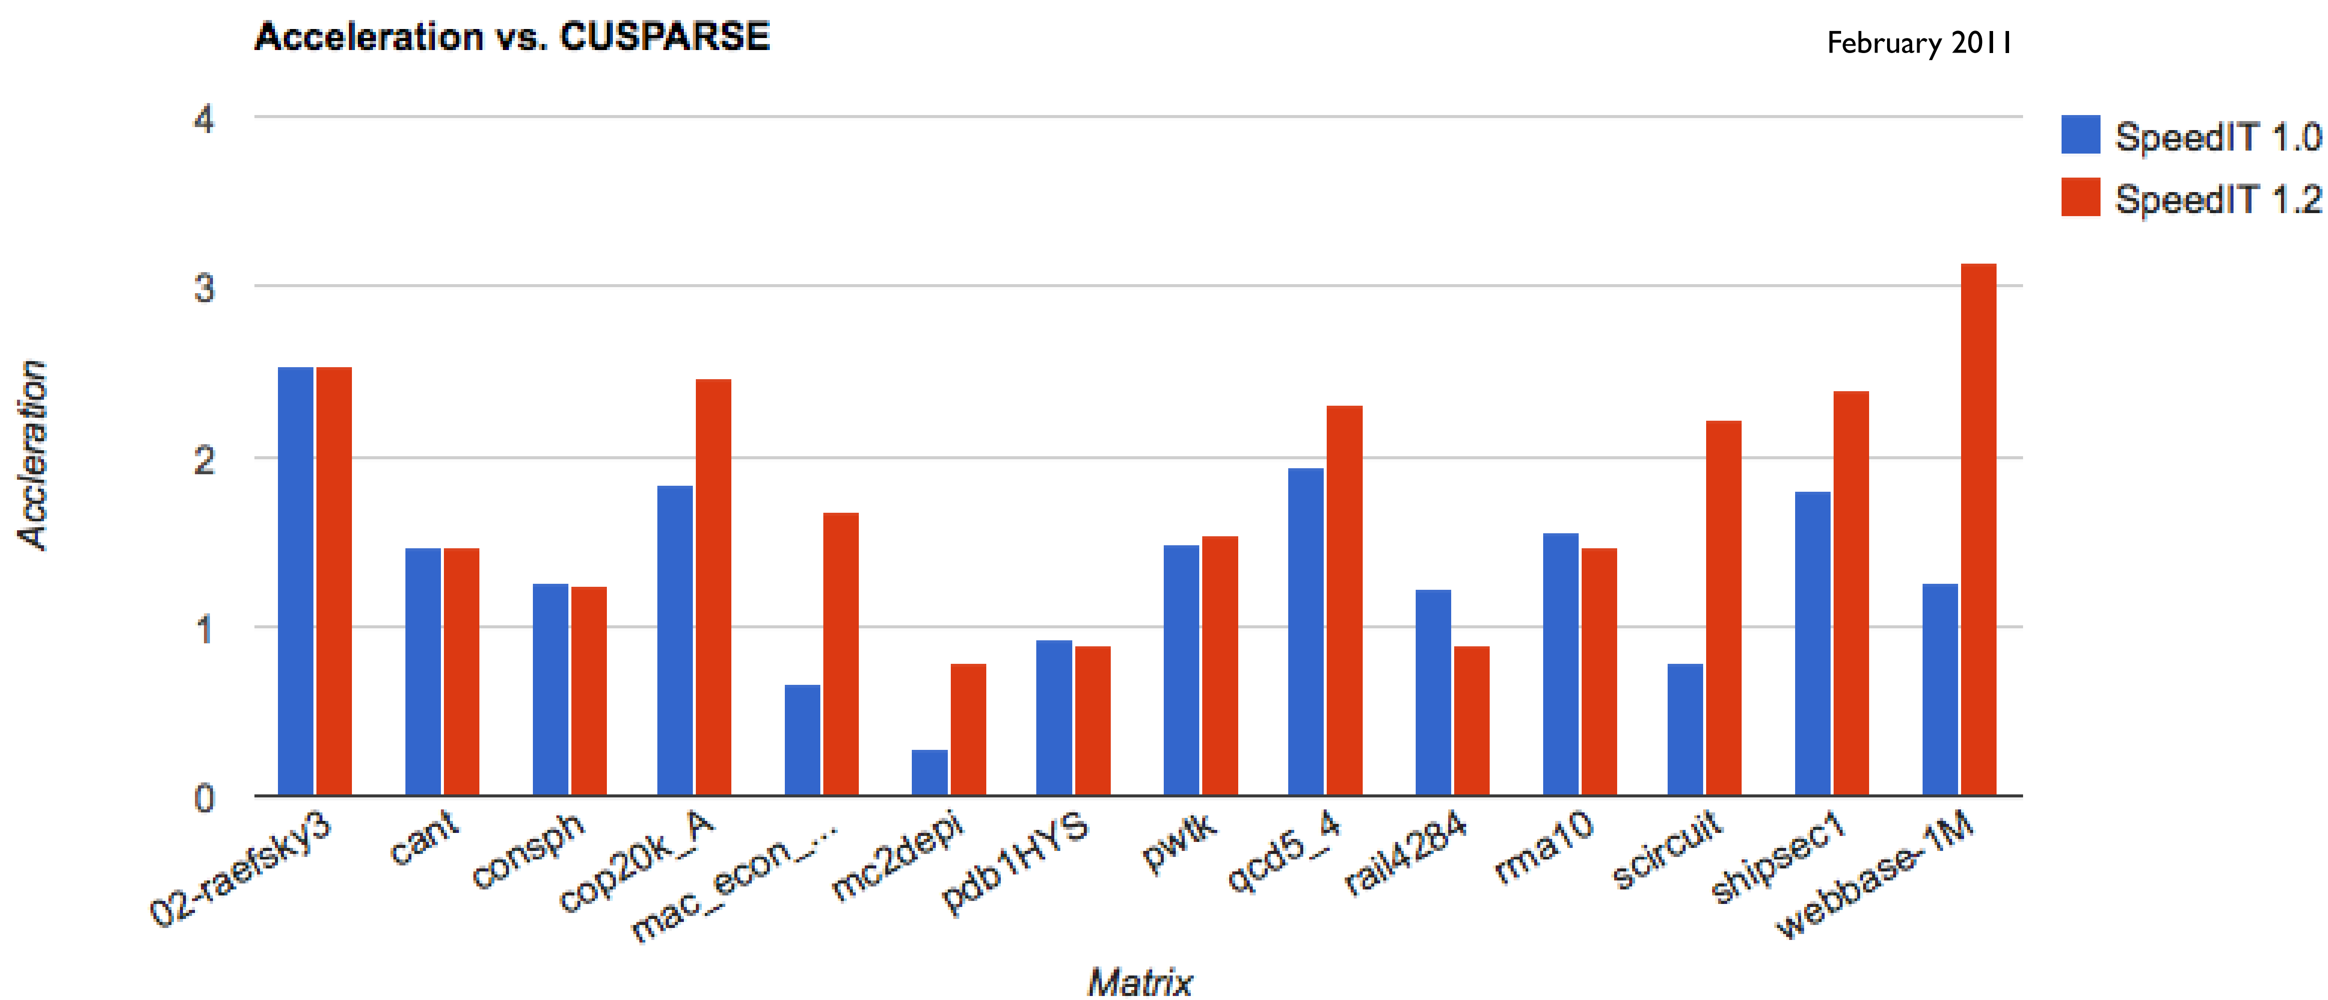

CUSPARSE developed by NVIDIA CUDA

Integration with CFD code

OpenFOAM

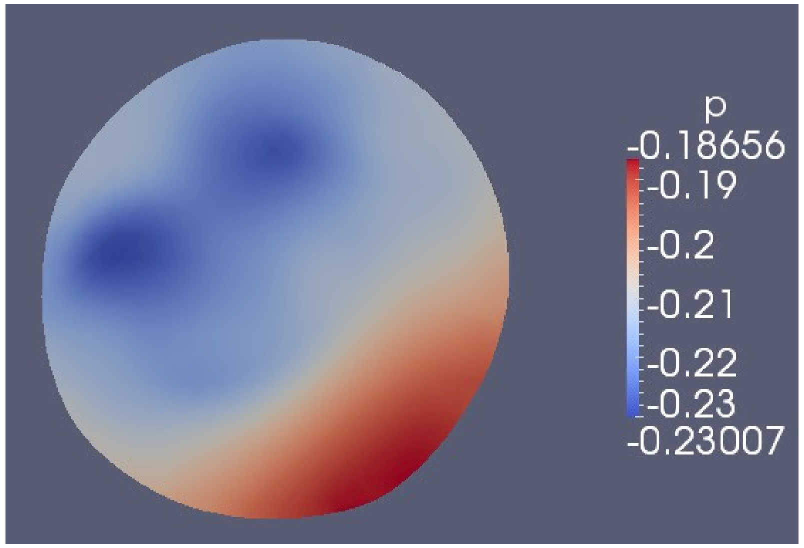

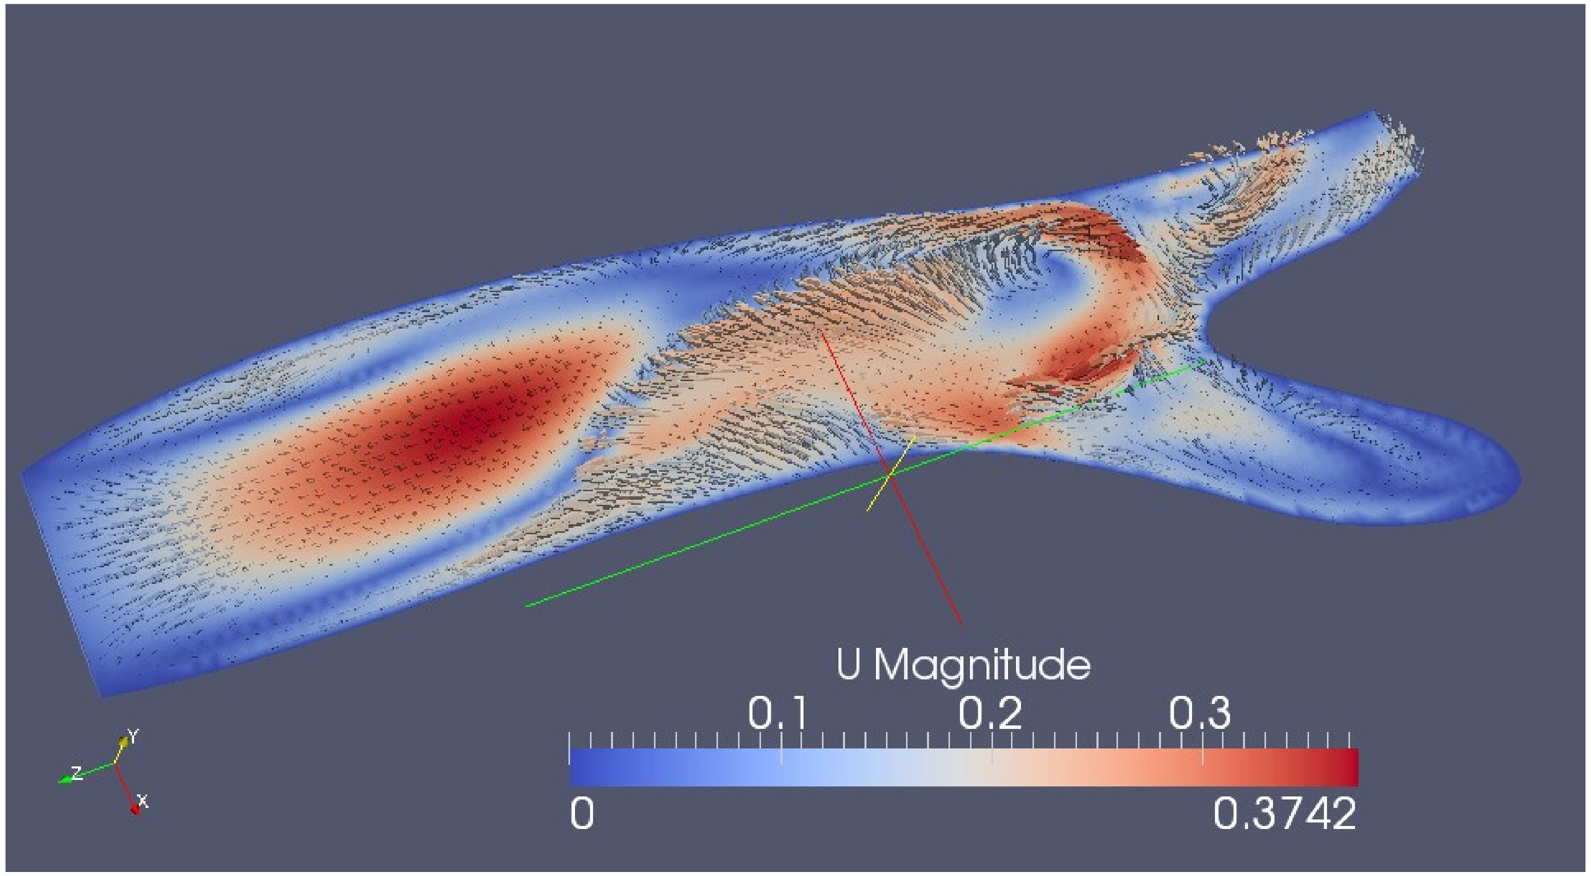

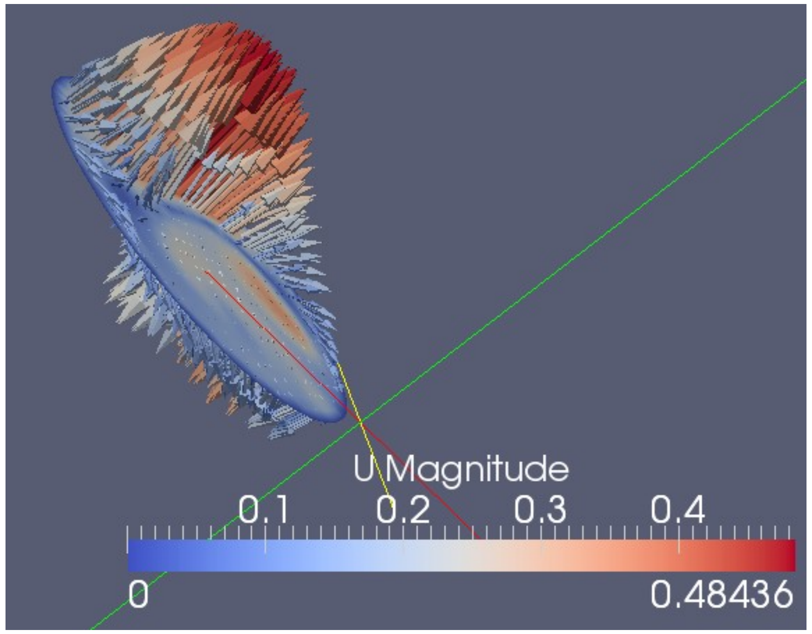

$$\frac{\partial \boldsymbol{u}}{\partial t} + \nabla \cdot (\boldsymbol{u}\boldsymbol{u}) + \nabla \cdot (\nu \nabla \boldsymbol{u}) = -\nabla p$$

OpenFOAM speed-up

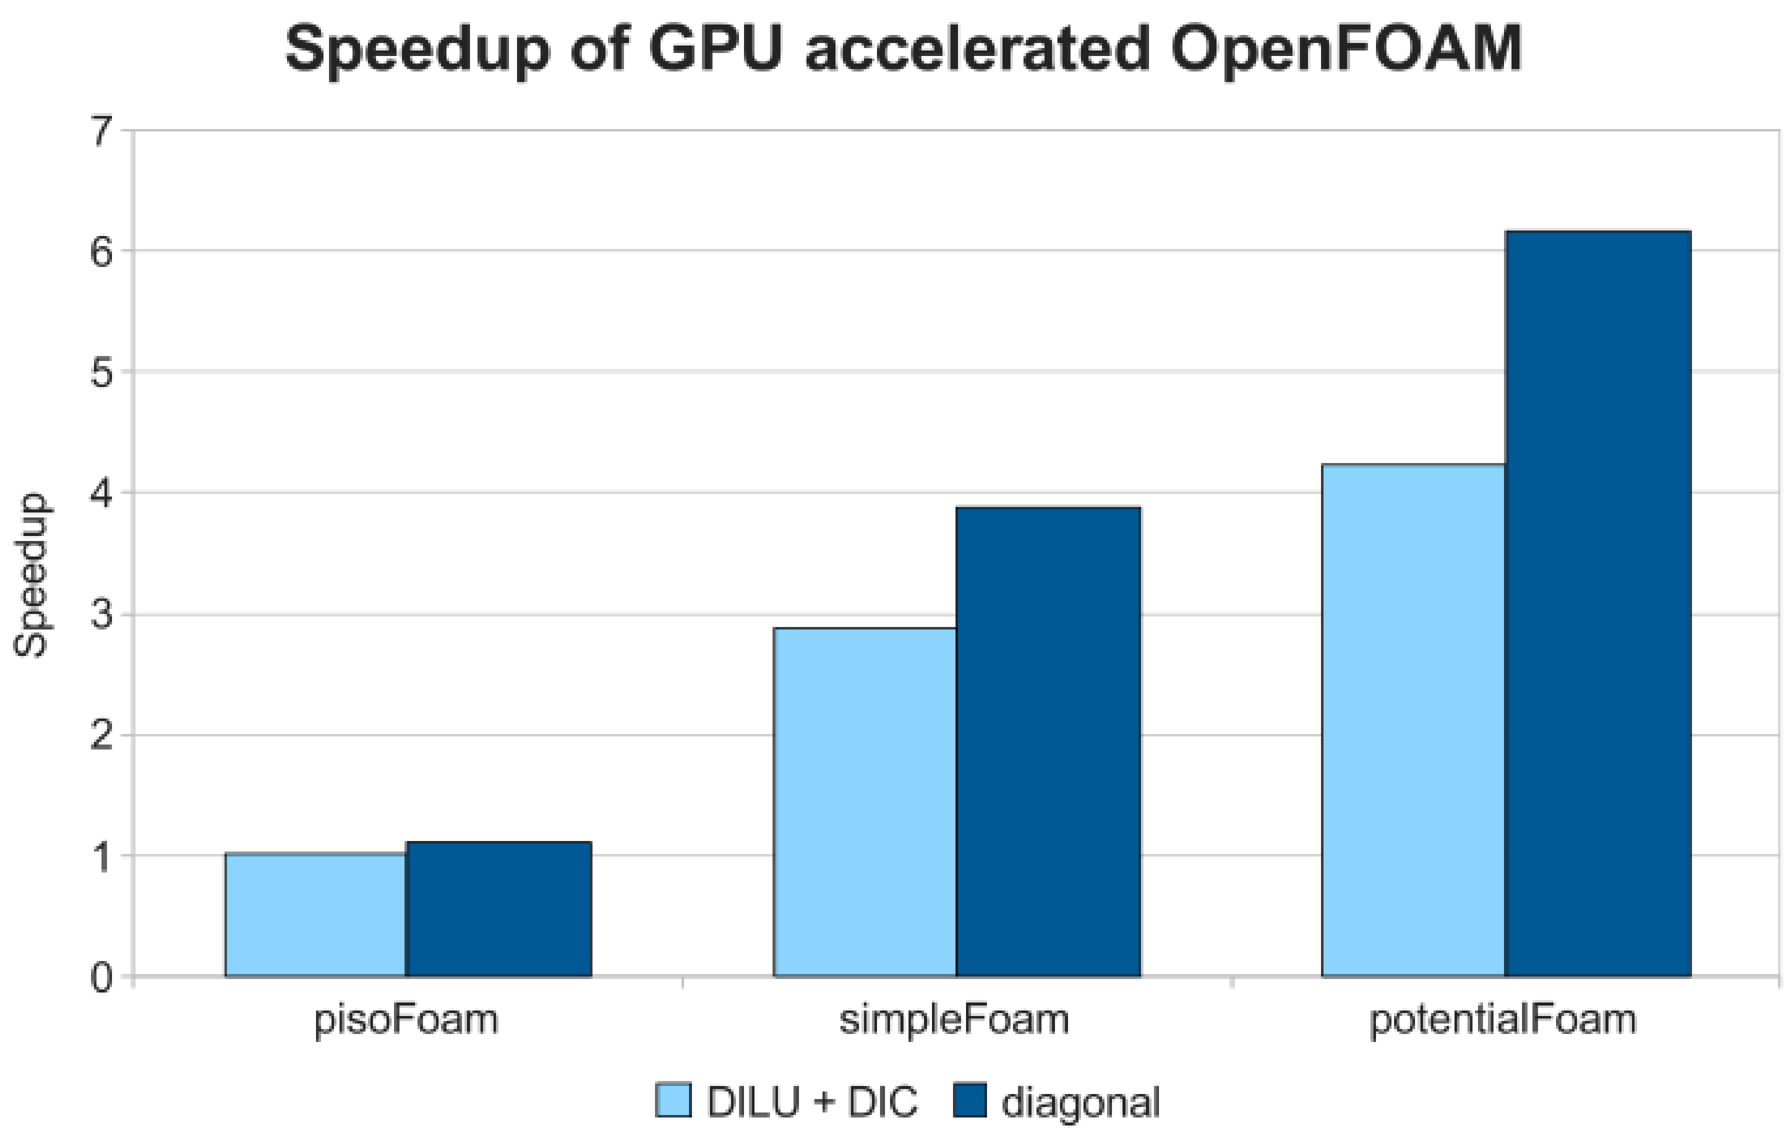

Z. Malecha, L. Miroslaw, T. Tomczak, Z. Koza, M. Matyka, W. Tarnawski, D. Szczerba. GPU-Based Simulation of 3D Blood Flow in Abdominal Aorta Using OpenFoam. Archive of Mechanics 2011 (accepted).

Multi-GPU in OpenFOAM

AhmedBody3D case: Mesh 2.2M cells, 4 domains. 0.55M cells per processor which is near optimal for simpleFOAM runs.

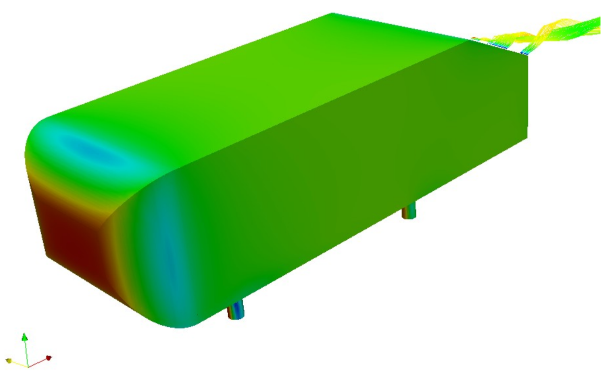

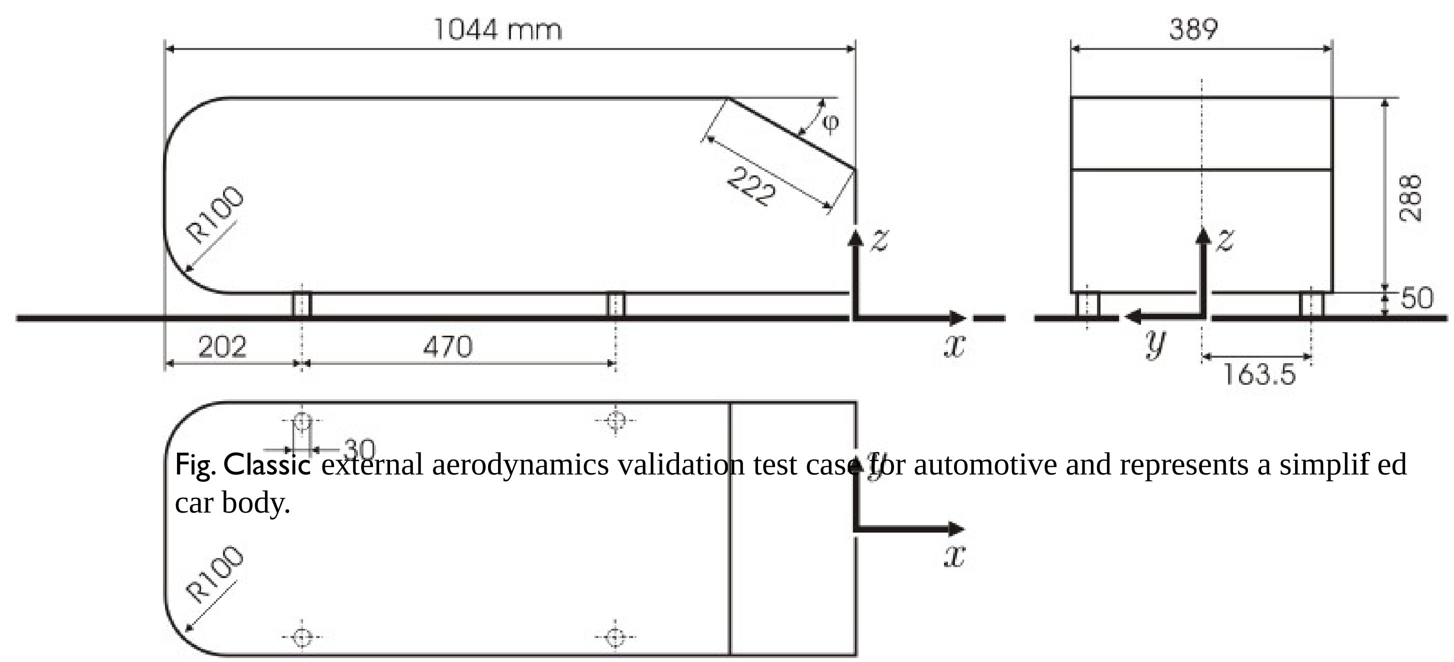

Fig. Classic external aerodynamics validation test case for automotive and represents a simplif ed car body.

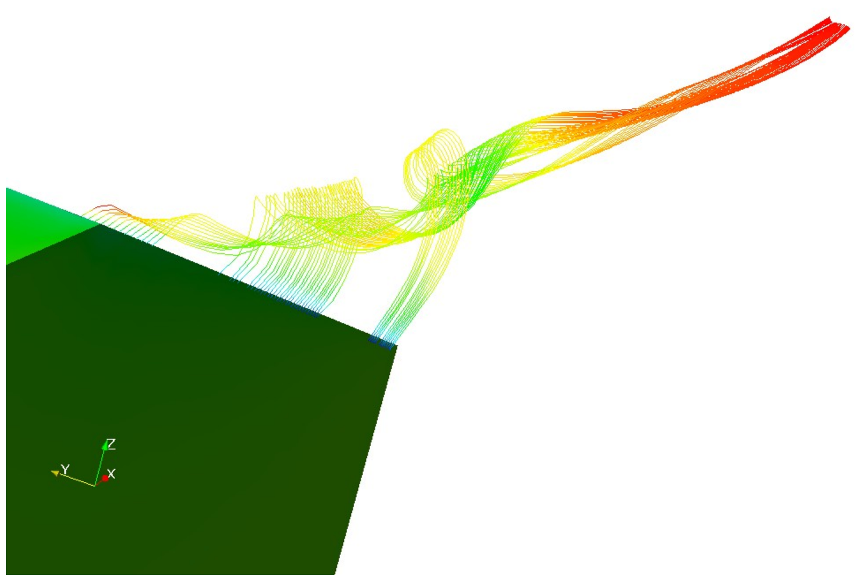

With collaboration with

engys

Multi-GPU in OpenFOAM

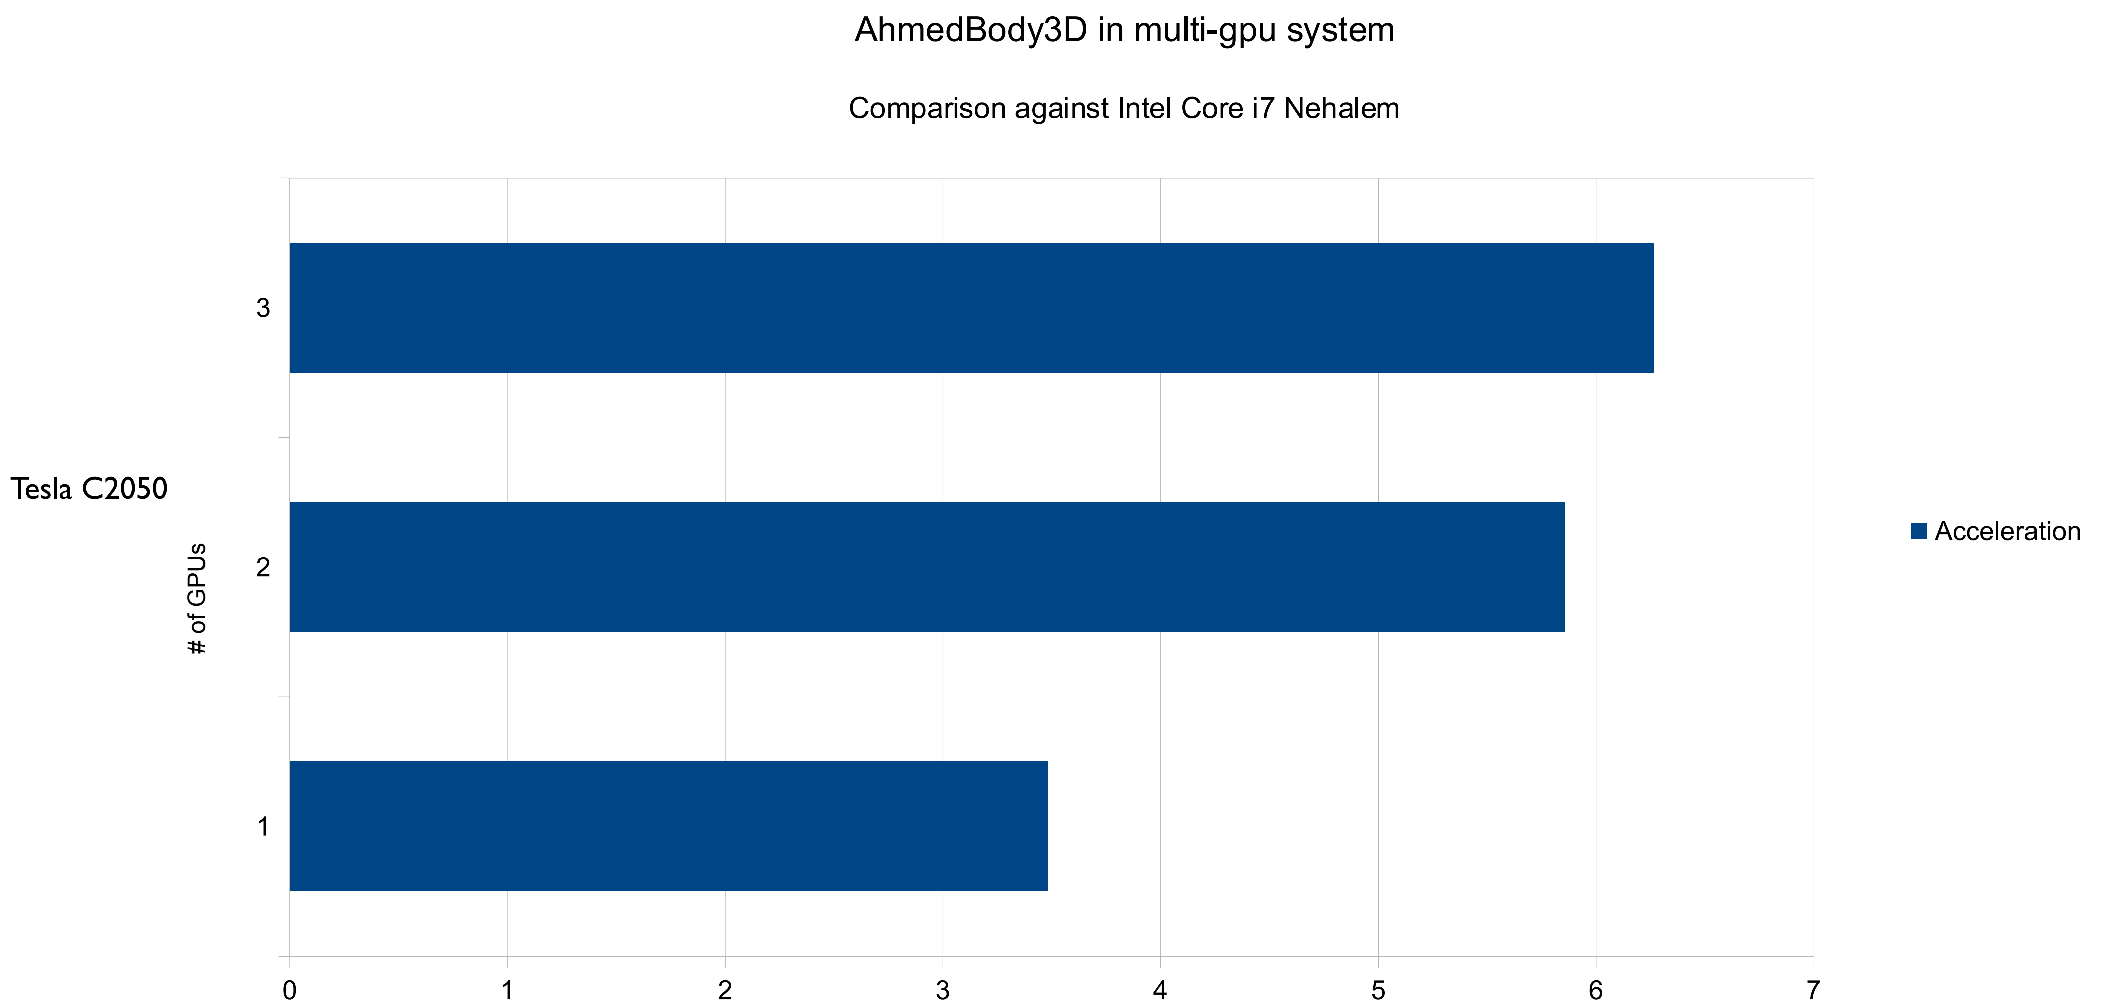

Fig. AhmedBoy3D simulation. The simulations were generated on multi-core systems.

With collaboration with engys FluiDyna INNOVATIVE ENGINEERING AND BEYOND

Evaluators, customers

ThyssenKrupp

STANFORD UNIVERSITY

UNIVERSITY OF CAMBRIDGE

SIEMENS

ACCELEREYES

engys

IT'IS FOUNDATION

Lucerne University of Applied Sciences and Arts
HOCHSCHULE LUZERN

Schlumberger

icon
Technology • Process • Consulting
www.iconCFD.com

Penn
University of Pennsylvania

PENNSTATE
1855

UNI BASEL

上海大学
SHANGHAI UNIVERSITY

IOWA STATE UNIVERSITY

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

CAPE HORN ENGINEERING

Market needs more!

- Realistic models in biomedicine take >> 24 GB.
  - Multi-GPU systems are a must.
- CPU-GPU memcopy limits acceleration for typical CFD problems.
  - GPU Direct was not (?) so successful.
    - Vide NVIDIA GPU Direct Technology Brief,
    - HPC Advisory Council Switzerland Workshop 2011(Towards Transparent and Eff cient GPU Communication on Inf niBand Clusters, CSCS)
  - CUDA 4.0, GPU Direct 2.0 for MPI is a next milestone.
    - Some custom MPI conf gurations do not work (RDMA is needed)
  - CFD Codes need to be rewritten to f t to new architectures.
    - CUDA API were not always backward compatible.

Cloud Computing. SME perspective.

- Privacy is still a concern.
- Porting to CUDA is not trivial.
- Microtransaction policy is an issue.

" (...) Cloud computing is the next stage in the evolution of the Internet that dynamically delivers information to end users where they want it, when they want it. And, in doing that there needs to be a sophisticated balance of computing hardware and software to power that kind of abundant and flexible data."

Tracey Blanton, a product marketing manager

AMD

- Clouds Considered at Upcoming GPU Technology Conference

To cloud or not to cloud? That is a question.

HPC	HPC in the cloud
▪ Expensive solutions. High risk of investment. ▪ Instant access to data.	▪ Micropayments, more flexibility. ▪ Problem with transfers of GB/TB of data.
▪ Security is well established.	▪ How to promise the customer the data are secure?

Acknowledgments

- Maciej Matyka
  PhD, University Wroclaw,
  Max Planck Institute for Marine Microbiology
  GPU, Theoretical Physics and Visualization
- Zbigniew Koza
  Prof. at University Wroclaw Theoretical Physics
- Dominik Szczerba
  PhD, IT'IS Foundation
  Biomedical simulations
- Vartan Kurtcuoglu
  PhD, Group leader at LTNT Lab
  ETH Zurich

UNIVERSITAS WRATISLAVIENSIS

IT'IS FOUNDATION

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

engys

FluiDyna INNOVATIVE ENGINEERING AND BEYOND

Universität Stuttgart

Don’t wait. Just SpeedIT.

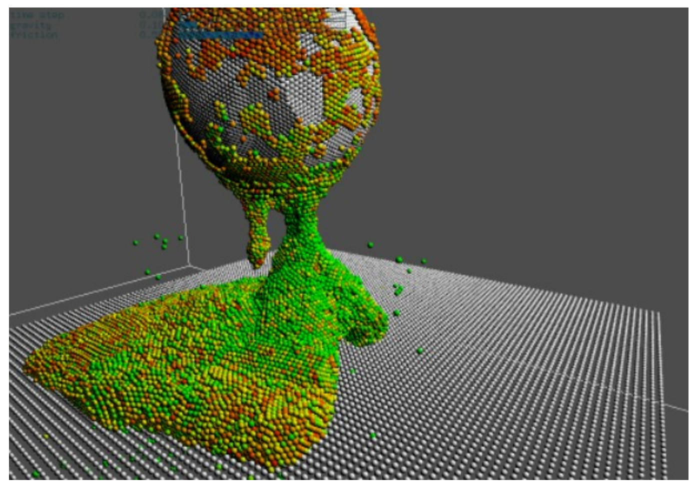

Vratis Ltd.
Wroclaw Technology Park
Muchoborska 18
Wroclaw, Poland

info@vratis.com
www.vratis.com