ICM: current developments in HPC area

Marek Niezgódka
marekn@icm.edu.pl

Kobe, HPC User Forum, July 16, 2014

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

HPC in Poland: 2014 update

- Funding 2008-2014: $200M+
- Main components:
  - Capacity infrastructure:
    - Cracow (Cyfronet)
    - Poznan (PSNC)
    - Gdansk (TASK)
    - Wroclaw (WNSC)
    - Governmental site (energy)
  - Capability facility:
    - Warsaw (ICM)

Coordinated concept: competence centers

- High Performance Networking: PSNC
- Capacity services: Cyfronet
- Capability and data-driven services: ICM

- Concerted POWIEW and PL-GRID R&D programs:
  - Complementary computing infrastructures
  - Joint job allocation concept
  - National services
- Aggregated computing capacity, 7/2014:
  - Cyfronet: .5 PF (2+ PF coming)
  - PSNC: 150 TF (subject expansion)
  - ICM: .5 PF (2-3 PF coming)
  - In addition, TASK & WNSC: each appr. 100 TF

ICM: compu-infrastructure

- As of July, 2014:
  - x86-based compu-clusters: appr. 10K cores
  - Blue Gene Q: 16K cores
  - Power 775: 2.5K cores
  - Blue Gene P: 4K cores (development & education)
- 2014 (pending):
  - Capacity compu-server: .2-.3 PF
- 2015 (pending, OCEAN facility):
  - HighPerformance data analytics system
  - Capability computing system
  - HighPerformance data storage

ICM, a few facts:

- founded in 1993, as a centre for:
  - HPC infrastructure: operations and development programs on national scale
  - National research information infrastructure
  - research in computational and information sciences

- one of key contributing institutions to the national research ICT infrastructure development in Poland
- ICM has set up:
  - national academic applications software system (1996)
  - national w3caching program (1996)
  - national research database accessibility (1997)
  - open-access multi-scale numerical weather prediction system (1997)
  - a group of leading-edge experimental labs (2005), extended into a university centre for new technology (2012)

  and also

  - annual festivals of science (from 1997)

Some of ICM's contributions to large research infrastructures and their development

- National virtual library of science (1997): content, software system
- National academic information infrastructure: INFONA system
- Polish Research Bibliography and Polish Citation Index
- EU: DRIVER and OpenAIRE open repository infrastructures
- EU: EuDML (European Digital Mathematics Library)
- EU: UNICORE – grid infrastructure (security functionalities)

ICM: evolution

- Centre for computational sciences (1993+):
  - Mathematical modelling
  - Foundations: physics, chemistry, biology, …
- Promotion and implementation of open publishing models (2004+) – Poland and Europe:
  - software
  - Publishing and scholarly communication
  - Research data
  - e-infrastructures
- OCEAN: Data Science and data-driven sciences centre (2013+)

Principle: INTERDISCIPLINARITY

<u>OCEAN:</u> Centre for Data, their Analysis and Computational Modelling

Infrastructure:

- Delocalized data centre (2 remote locations)
- Data storage systems
- Critical processing and big data analytics systems (*time-critical processing*)
- Extreme computing (*capability computing*)

Concept:

- Centre for data science and data-driven services
- Dual set-up:
  - research-oriented profile
  - governmental service infrastructure

ICM: integrated OCEAN concept

- Distributed location:
  - Current site – future back-up facility (2015)
  - New site:
    - Under construction
    - Operations: Q3, 2015
    - Total floor: 6000+ sq.m
    - Power supply: 4MW (10MW pending)
    - Energy-efficient facility
  - Concept:
    - Data storage services: initial ~ 40 PB fast storage
    - HP Data Analytics: ~ 850 GTeps
    - Capability computing: ~ 3 PF

OCEAN: program

- National data infrastructures:
  - Research: INFONA integrated knowledge system
  - Government: CRIP public information repository system
- HighPerformance data analytics:
  - large-scale data-sets
  - time-critical dependable services
  - development of specialized solutions and infrastructure for big-data analytics
- Capability computing:
  - time-critical applications
  - data-driven computing
  - development of specialized solutions for big-data driven computing

OCEAN: core development and service areas

- Large operational neworks (interactive structured populations)
- Sectors:
  - Natural environment (weather, pollutions, hazards)
  - Energy (intelligent energy networks - *smart grids*)
  - Transportation (air and surface)
  - Logistics (delivery services)
  - Healthcare and personalized medicine (also biobanks)
- Public administration:
- Central Repository of Public Information CRIP (for government and public administration -created data and information)
- Intelligent city (Warsaw)
- Social processes related to new technologies and digital society transformation

OCEAN: selected reference areas

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Numerical weather forecasting and dependent services (meteo.pl)

- Multi-model multi-grid processing
- Top horizontal resolution in development 1km
- Vertical resolution: 70+ layers

- Services:
  - Severe weather warning systems
  - Energy sector
  - Transportation and logistics
  - Agriculture
- R&D:
  - Energy smart grids
  - Airspace management
  - Natural hazards
  - Precise agriculture

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Large-scale visual data processing and analysis

- **Visualization**
  - Translation from data to image
  - Natural and parallel information interpreted by human
  - Information presentation
  - More comprehensive overview of data
  - Key role in data interpretation
- **Visual Analysis**
  - Enables key information extraction
  - Reveals the unexpected information
  - Supports model building
  - Allows interactive analysis

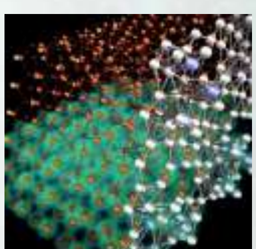

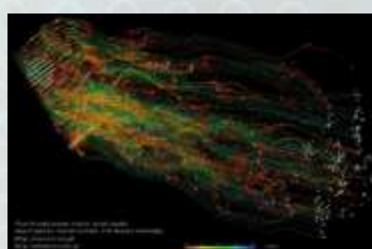

Visual Computing and Data Analysis

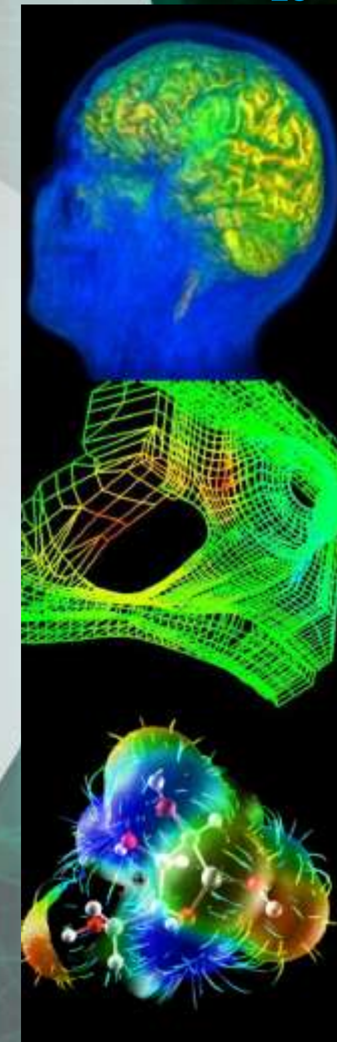

New solutions: VisNow Visual Analysis Software

- Distributed Visualization Engine (MPP, SMP, GPGPU)
- In situ Visualization
- **Visualization of simulation results**
    - Cosmology
    - Turbulent flows
    - Neurobiology
    - Biomedicine
- visnow.icm.edu.pl (Open Source license)

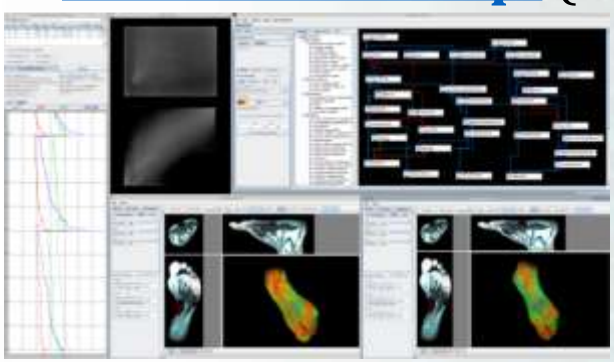

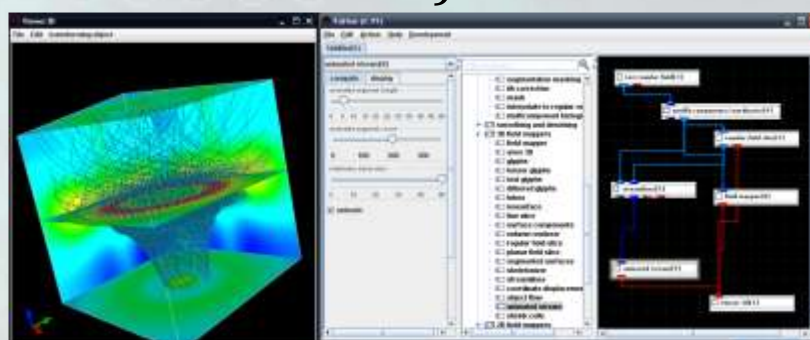

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Software and services

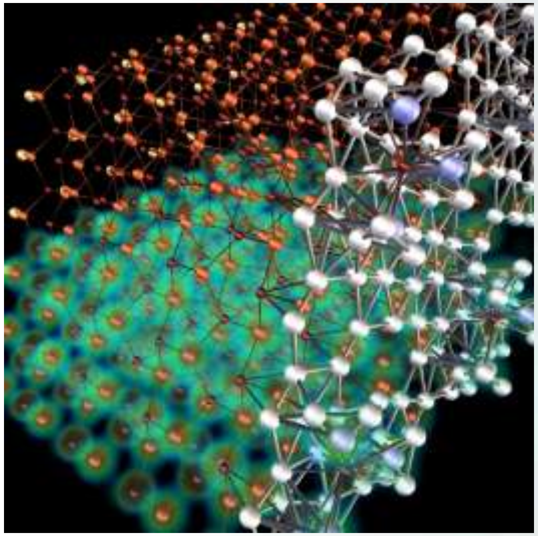

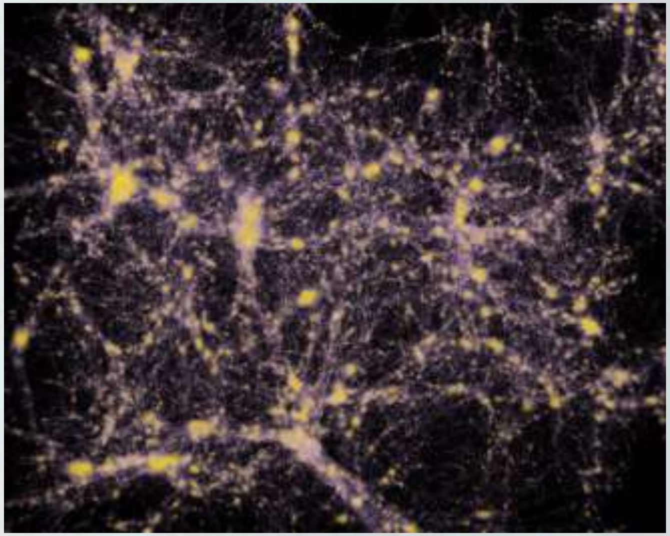

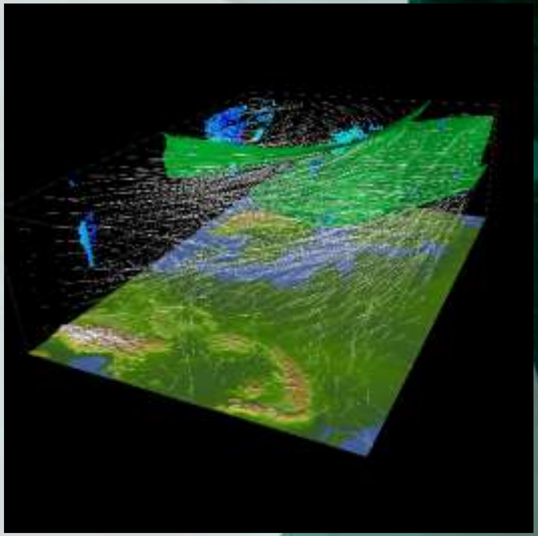

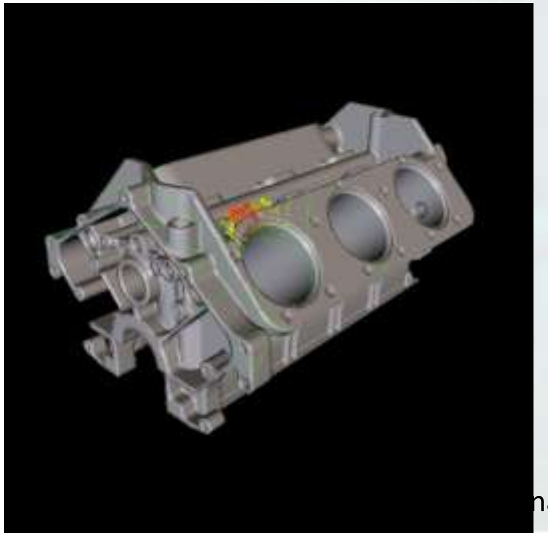

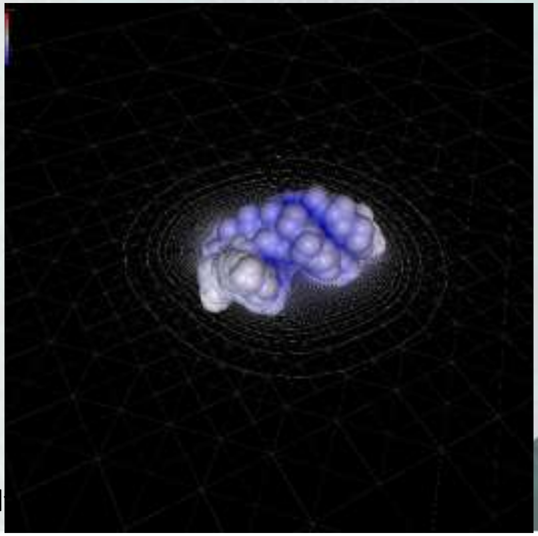

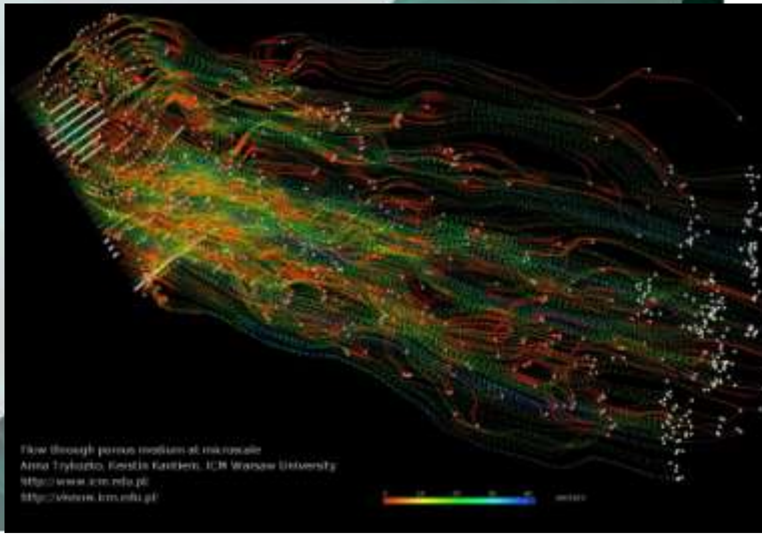

VisNow: software and services

- **MolDyAna**
  - Interactive tool for **molecular dynamics** data analysis and visualization
  - Dynamics animations
  - Geometry measurements
  - Density charts
  - Trajectory analysis
  - Spectral analysis

http://moldyana.icm.edu.pl

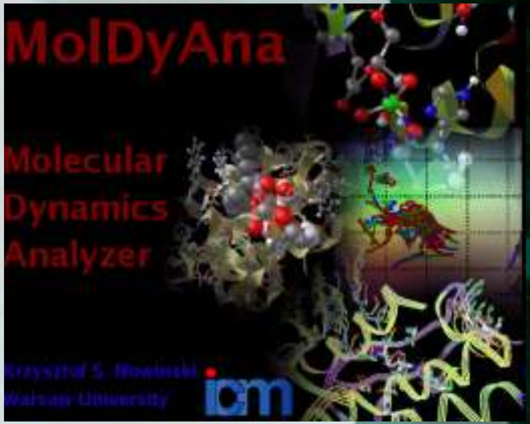

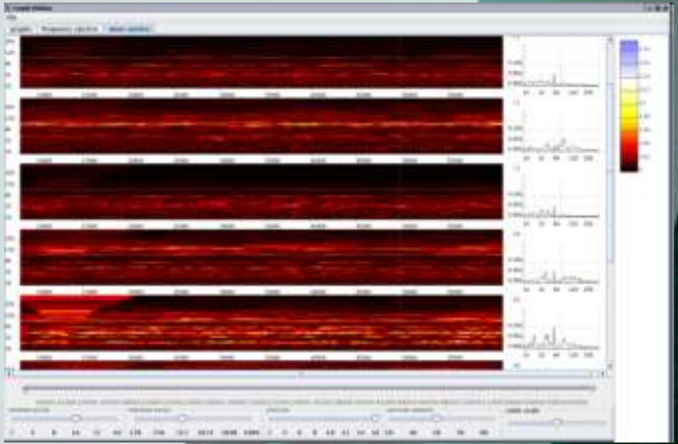

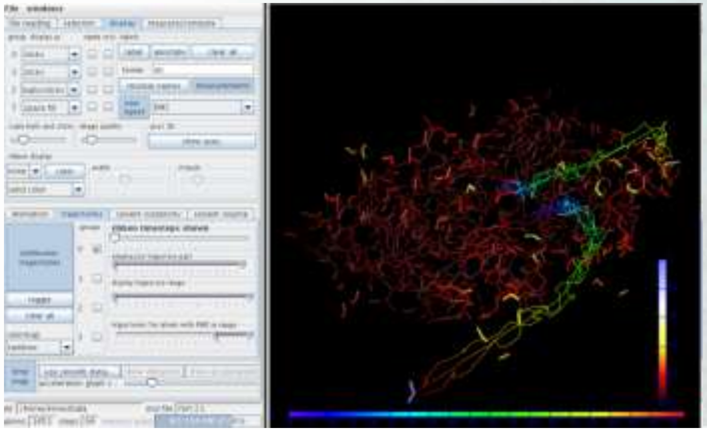

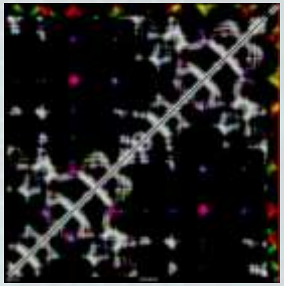

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Case studies: computing

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Case study: Simulations of biological cellular processes at tissue scale

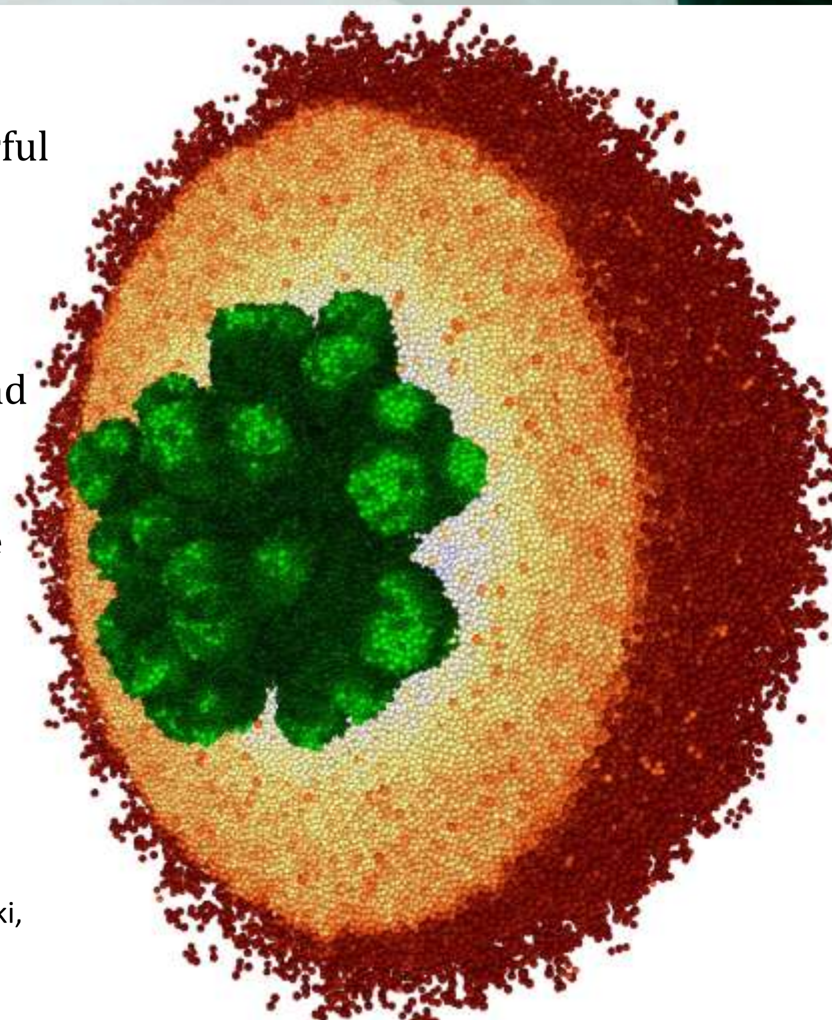

- **Main objective**: set up a methodology and powerful computational algorithms to enable large scale modelling of cellular biosystems
- **Multiscale approach**: simulation of cellular biosystems dynamics, individual cell processes and environment
- **Hybrid computational model**: solving off-lattice many body system, PDE's describing cellular environment and their interactions
- **Challenge**: enabling simulations on clinically detectable scales (10^9 cells = $1cm^3$ tissue)

Large Scale Parallel Simulations of 3-D Cell Colony Dynamics, M.Cytowski, Z.Szymańska, IEEE Computing in Science & Engineering, 2014

Case study: Copernicus Complexio

- **Zoom-in type simulation:**
  - **Ultra high-resolution**: over 13 bilion particles inside computational domain (green)
  - **Medium resolution**: transition zone (blue)
  - **Lowest resolution**: sufficiently large segment of the Universe (red)
- **Technical details:** 70 TF & 10 TB computational partition, approx. 4 weeks wall clocktime
- **Analysis**: results of the CoCo simulation are being analysed by international cosmological consortium
- **Visualization**: http://vimeo.com/76812335

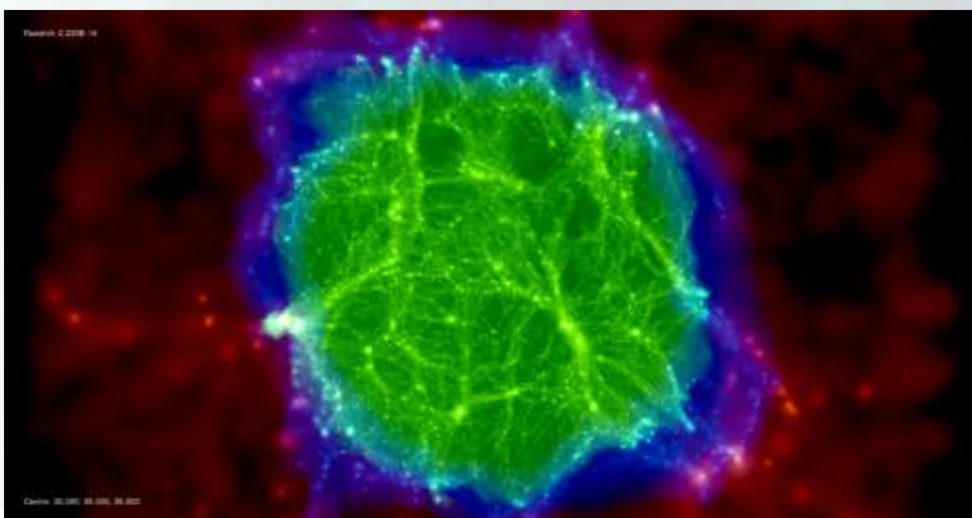

Modelling of the Structure of the Universe

- Large scale simulations with N-body codes (Gadget3, GotPM)
- **Warsaw Universe Simulation**, 2048^3 particles
- MPP-based results analysis:
  - Statistical methods
  - Topological classification
  - Geometrical classification
  - Delaunay Tesselation, Alpha-shapes and Betti numbers

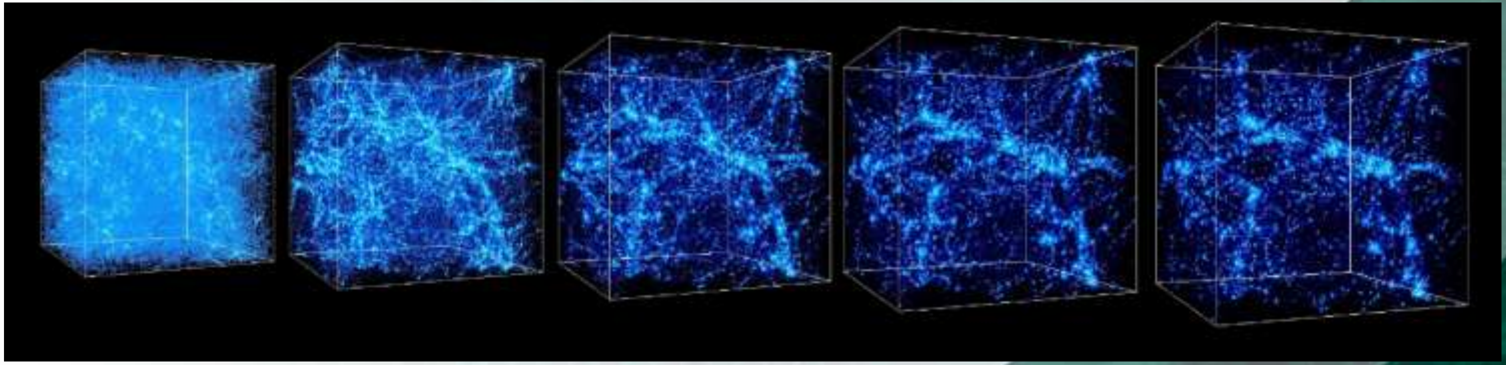

Case study: Numerical Weather Forecasting

- **Operational numerical weather forecasting for Central Europe**

  http://meteo.pl

- **Unified Model 60h forecasting,** current resolution approx. 4km, planned resolution approx. 1km
- **IBM Power 775** – 5x speedup over x86 cluster
- **Technical details:**
  - 4 runs per day
  - 24 nodes, 768 Power7 cores
  - approx. 20’ walltime per run

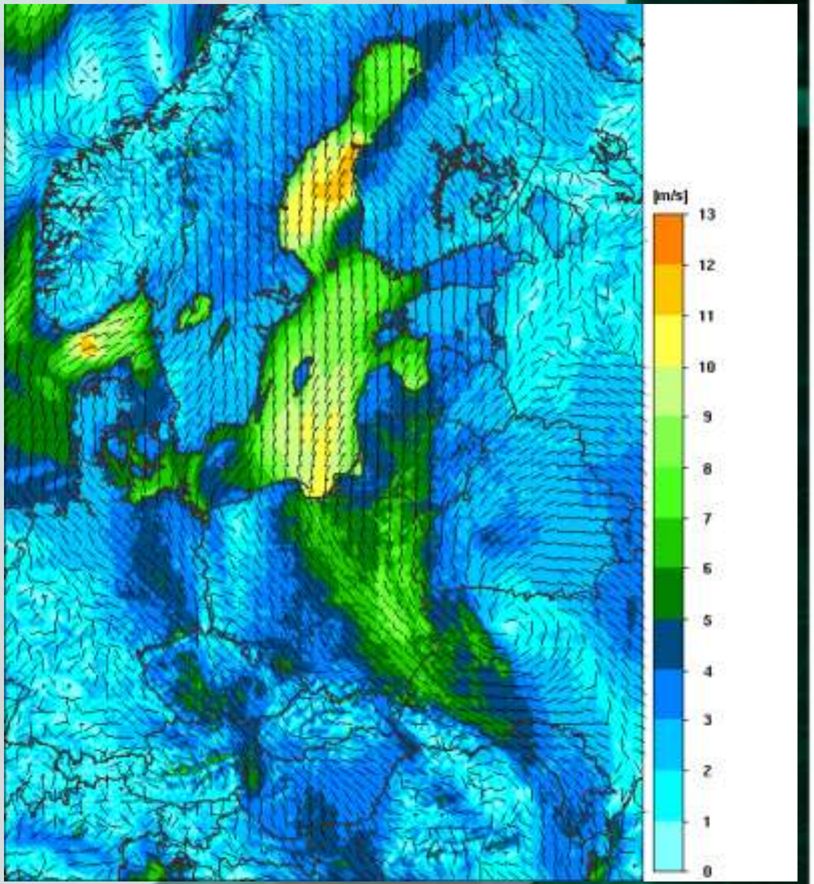

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

High-throughput modelling of functionally and therapeutically relevant spatial RNA structures

- Modelling of spatial RNA structures
- Design of new therapeutic targets aimed at the RNA
- Validation of ICM's RNAComposer software for fully automated spatial RNA models generation from structural fragments

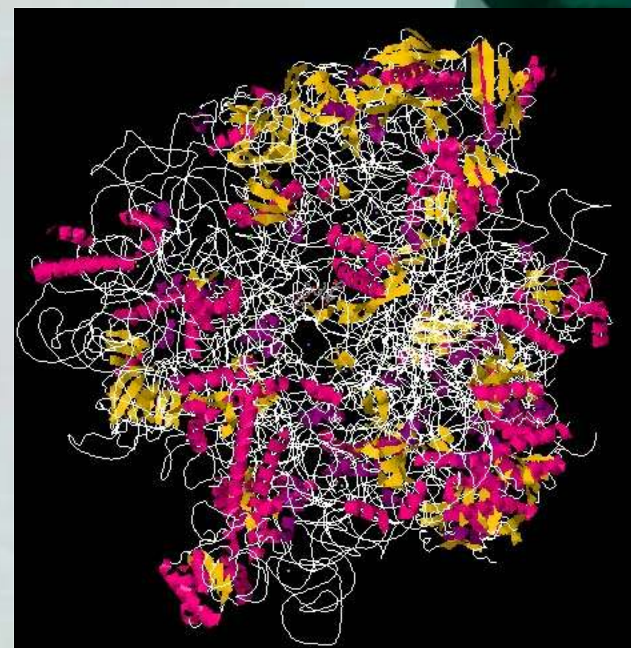

Infectious disease spread over Poland: computational Agent Based Model and its predictions

- Simulations based on detailed, statistically relevant artificial social structure and mathematical model of the infection
- Model predicts space-distributed epidemic outbreak:

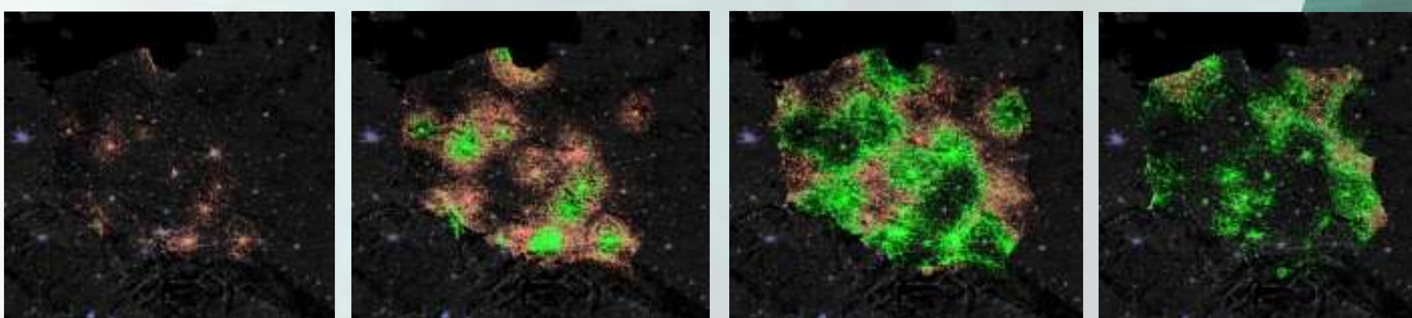

- and location of the epidemic peak in time, dependent on different socio-biological circumstances:

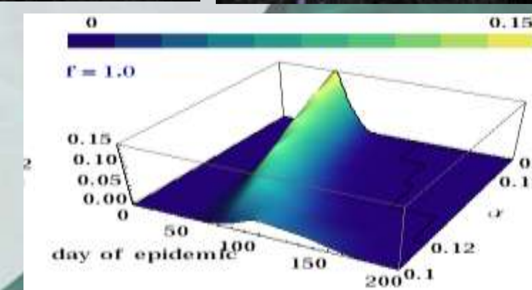

POWIEW program

http://wielkiewyzwania.pl

Large scale scientific computing projects:

- Numerical Weather Prediction
- Semiconductor modeling
- Modelling and visualization of RNA structures
- Neuroinformatics simulations
- Modelling of the structure of the Universe
- Molecular modeling
- Reservoir modeling
- Astrophysics

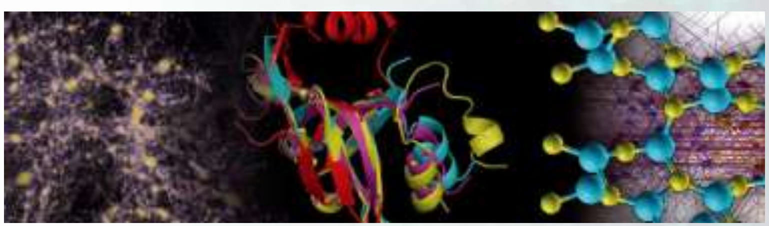

Main computing architectures:

- MPP systems – IBM BlueGene/P & BlueGene/Q
- Fat Node systems – IBM POWER7 IH
- SMP systems – HP Blade Center Versatile SMP (vSMP)
- GPU-based hybrid systems – HP SL390s nVidia Fermi

Marek Niezgódka

marekn@icm.edu.pl

www.icm.edu.pl

ICM: driving research areas

- **Human physiology:**
  - Towards personalised medicine: cardiac system, circulation
  - Multiscale developments: from cellular level upwards
- **Materials science and engineering**:
  - Design of new functional materials and biomaterials down to nanoscale: characterization, process modelling
- **Methodology:**
  - New concepts of computational algorithms and data structures for future architectures (towards exascale)
  - Nonlinear process dynamics in systems of high complexity:
    - spatial structure formation in systems over complex, possibly variable geometry/topology, applications to population dynamics
    - irreversible and nonlocal phenomena
    - decision making and control over multiple time-scales
    - stochastic networks

Modeling flows through porous media

- Main interest:
  - Modeling flows at pore scale   (large range of velocities)
  - Upscaling to core scale: effective parametres and model fitting
  - Studying inertia effects at pore scale (non-Darcy core-scale models)
  - Modeling pore-clogging processes and their impact on core-scale parameters
  - Computations based on realistic geometries (microimaging with X-ray computed microtomography)
- Porous media virtual laboratory:
  - Pore scale
  - Upscaling and dual-scale model fitting

Challenges in 3D pore-scale simulations: SIZE!!!

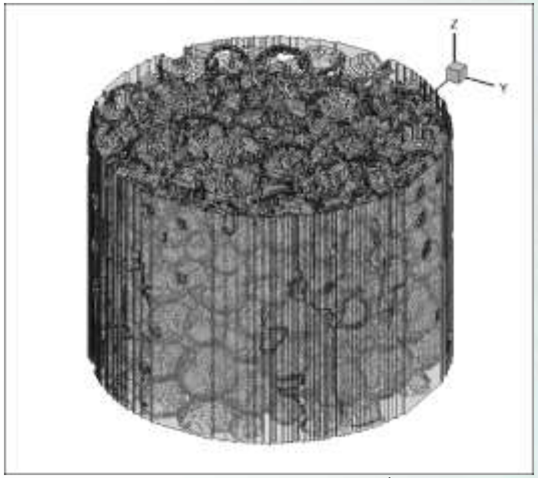

Glass beads[1]

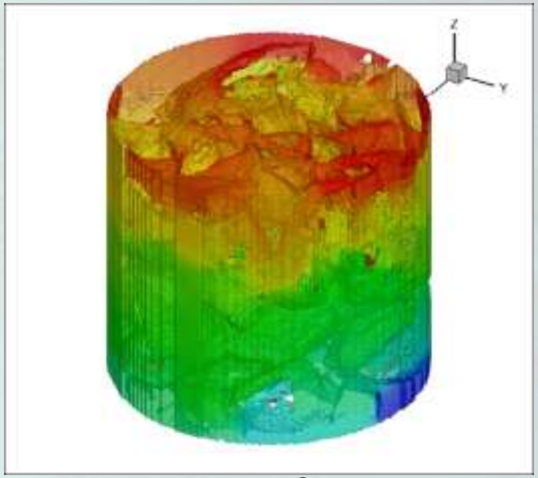

Tuff[2]

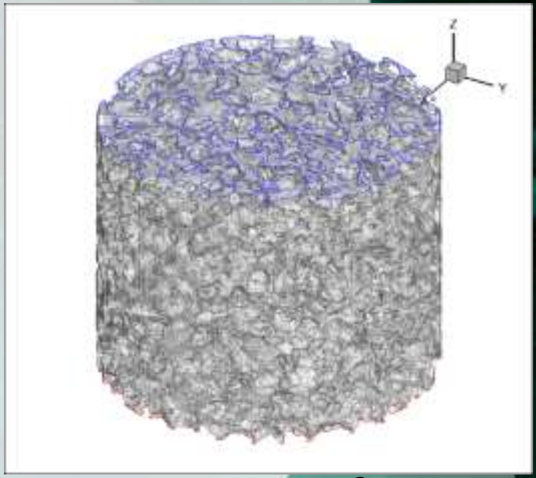

Sandstone[3]

- Complexity ↑30M cells, 10Gb files, challenging visualization and post-processing
- Grid reduction (data coarsening) necessary
- Sampling vs. *Reference Elementary Volume* size
- Mesh refinement

1[DWildenschild], 414x414x300, largest problem: 30Mcells, 9Gb, voxel 34μ.

2[DWildenschild], 431x436x380, 23Mcells, too small to be *REV*

3[BLingquist], 731x731x600, largest problem 8Gb, 27Mcells, voxel=16 μ.

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Upscaling from pore to core scale

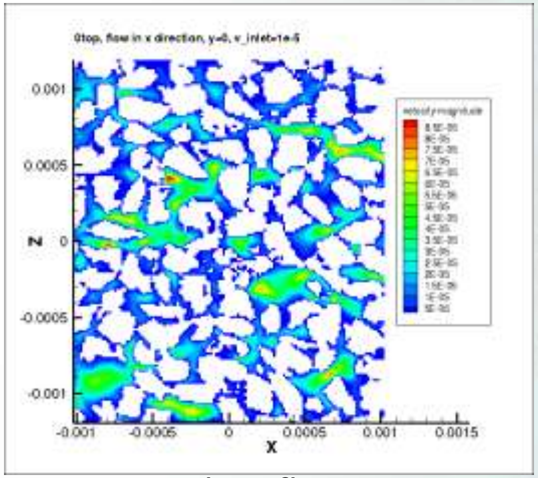

slow flow

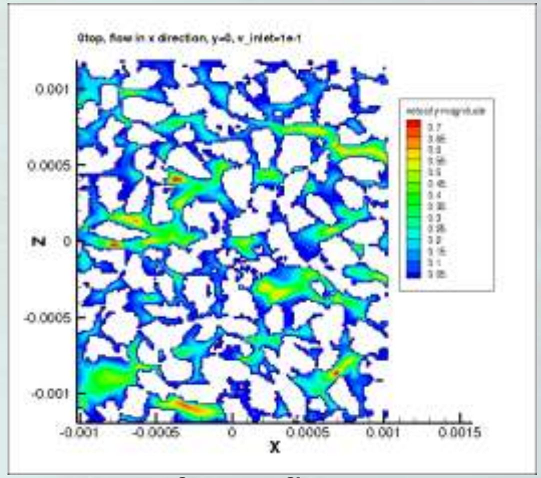

faster flow

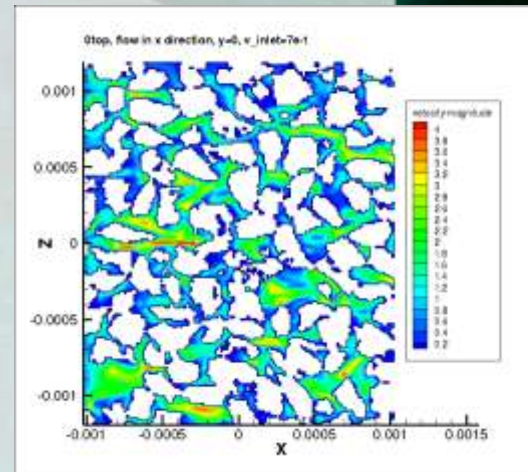

fast flow

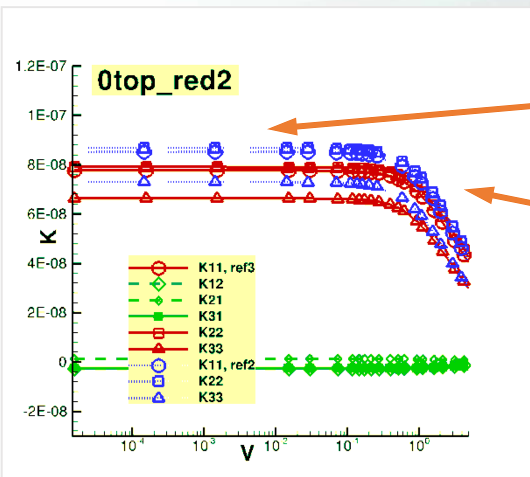

SUMMARY in terms of upscaled values

- Darcy's regime: constant ratio of P (upscaled pressure) and V (upscaled velocity)
- Inertia effects – non-Darcy models
- Accounting for anisotropy

M. Peszyńska, A. Trykozko: *Pore-to-core simulations of flow with large velocities using continuum models and imaging data*, Computational Geosciences, Vol. 17, nr 4, (2013), 623-645

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Impact of pore space clogging on core parameters

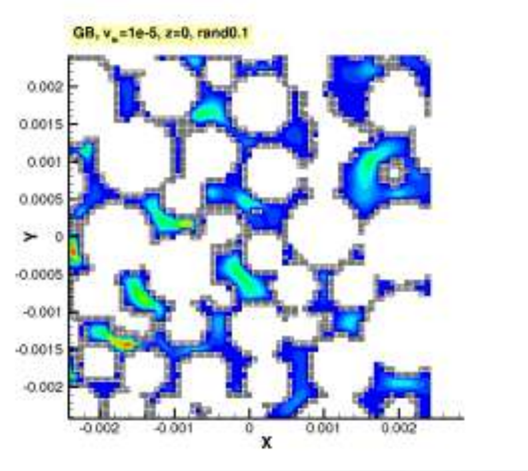

- Changes in pore geometries may result for a variety of processes: biofilm growth, sedimentation, reactive flows, antrophogenic actitivite.
- Significant changes in permeability due to cloggin.
- Figures: random model of clogging (irregular pore lining)

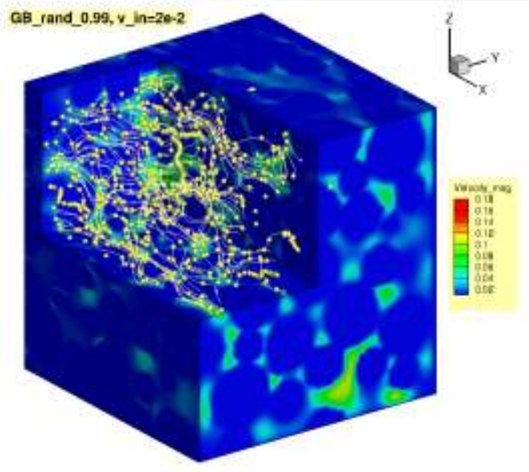

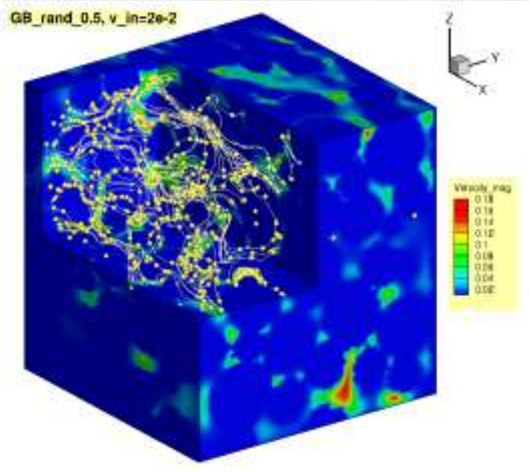

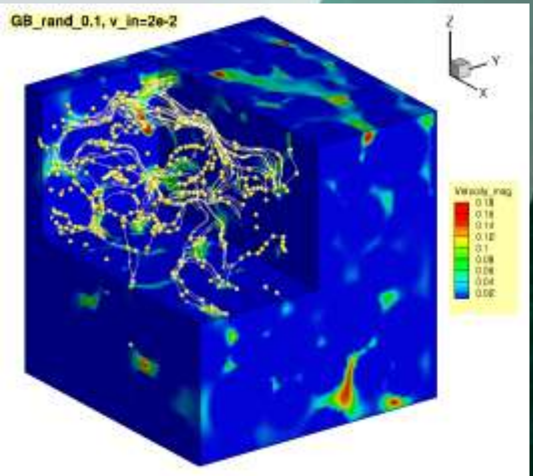

A. Trykozko, M. Peszyńska: *Pore-scale simulations of pore clogging and upscaling with large velocities*. GAKUTO International Series, Mathematical Sciences and Applications, Vol. 36 (2013), 277-300.

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Large scale visualization in HPC

- **Big calculations, Big data, Big challenges**
  - Technology growth rate – computing power and storage size outrank display resolution
  - Rapidly growing size of computational problems in HPC
  - Data amounts exceeding full analysis possibilities
  - Duality of Big Data
    - **Small number of large data sets**
    - **Large number of small data sets**

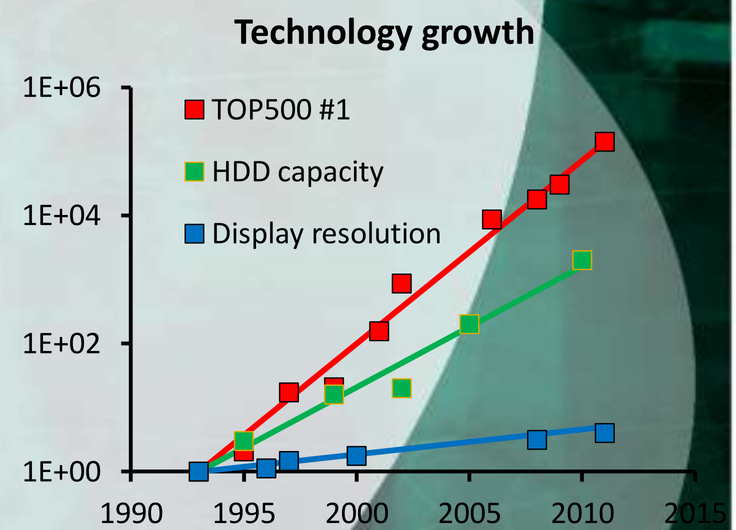

- **Solution – Large scale visualization**
  - **HPC** infrastructure based
  - **Distributed** and **parallel** processing
  - Adequate **new programming paradigms**
  - **Information extraction** and data simplification methods
  - Suitable **software tools** and services

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Marek Niezgódka

marekn@icm.edu.pl

www.icm.edu.pl

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Advances

New feasibility ranges to scientific discovery due to:

- recent spectacular advances in computing technology
- emerging new computer architectures and paradigms

Problems featuring complexity yesterday far beyond any quantitative scientific treatment, soon will get treatable by computational modelling and simulation:

- a breakthrough necessary in:
  - concepts underlying algorithmic structures
  - a new generation of computational models

Core of emerging research foundations:

- concepts:
  - breaking traditional linear (or linearised) paradigms
  - setting up novel adaptive constructions
  - built on top of mixed, hybrid computing architectures
  - data-driven discovery process

Next-future computational sciences research scenarios

- **Several scales of resolution** within a single common framework:
  - phenomenological models (up to macroscale) integrated with those based solely on the first principles
- Admitted **high structural complexity** of the models:
  - apart from rather homogeneous large systems, strongly heterogeneous systems (up to multi-scale complex operational networks) can also be treated
- **Virtual experimentation** will be facilitated:
  - also covering extremely complex interactive systems
  - that may reduce a necessity of high-risk (possibly harmful) real-world experiments
  - leading to data-driven discovery

Next-future computational sciences research scenarios, cont.

- Realistic:
  - integration of large-scale datasets, in particular those acquired in real-time modes including imaging and signalling, into online computing

- Novel architectures will enable replacing computational models built on sequential von Neumann paradigms by:
  - hybrid, adaptive approaches
  - optimized for massively parallel heterogeneous architectures

Validaction of interscale model consistency

- Cell populations dynamics
- Degenerate tissue growth
- Cell populations dynamics, pathological cell populations growth spreading mechanisms
  - Human agents and their behaviour: transport-related reference patterns
  - Process dynamics in large-scale networked environments: local and global passenger-centric model
  - Information transfer trees: hybrid and continuous models in information spaces

Modelling of growth and structure transformations in heterogeneous systems

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Mathematics within multiscale modelling framework

Nonlinear dynamics of spatial developments:

- Phase transformations
- Chemical reactions/ process thermodynamics
- Rheology / memory / nonlocal interaction mechanisms
- Structured populations
- Coupled multiscale systems
- Geometric / topological evolution
- Uncertainty components / **sensitivity** aspects

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical and Computational Modelling
www.icm.edu.pl

Specific applied modeling problems

Phase separation and related phenomena:

- nonlinear variational problems
- process dynamics
- structure formation

Phase transformations, crystal growth, physiology:

- biomedical processes: blood circulation, tumor growth
- process design & control

Population dynamics:

- structured populations
- operational (stochastic) networks

Research: opportunity areas

- Supply/exchange of **image data**, accelerating solutions
- Processing: **multimodal solutions** (integrating diverse categories of image/signalling data)
- **Virtual prototyping**/design in diagnostics and therapy planning (computer simulation)
- **Adaptive** (self-learning) **supporting systems for diagnostics and therapy**:
  - virtual cell
  - virtual patient
  - virtual organism

Applications in biomedical process modelling and physiology

Visual neurocomputing:

3D reconstruction of multimodal data
for therapy/surgery planning:

effective vs. real medium computational model

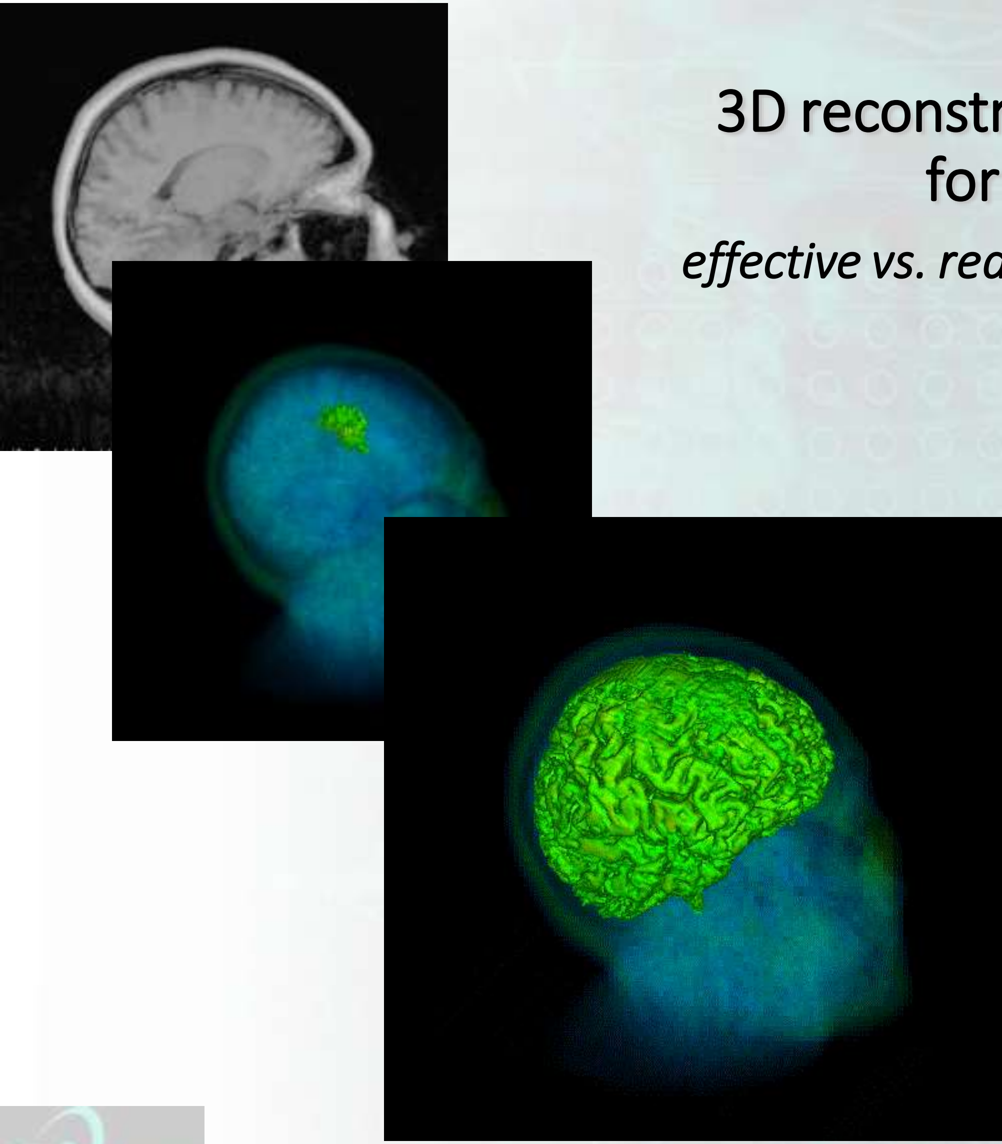

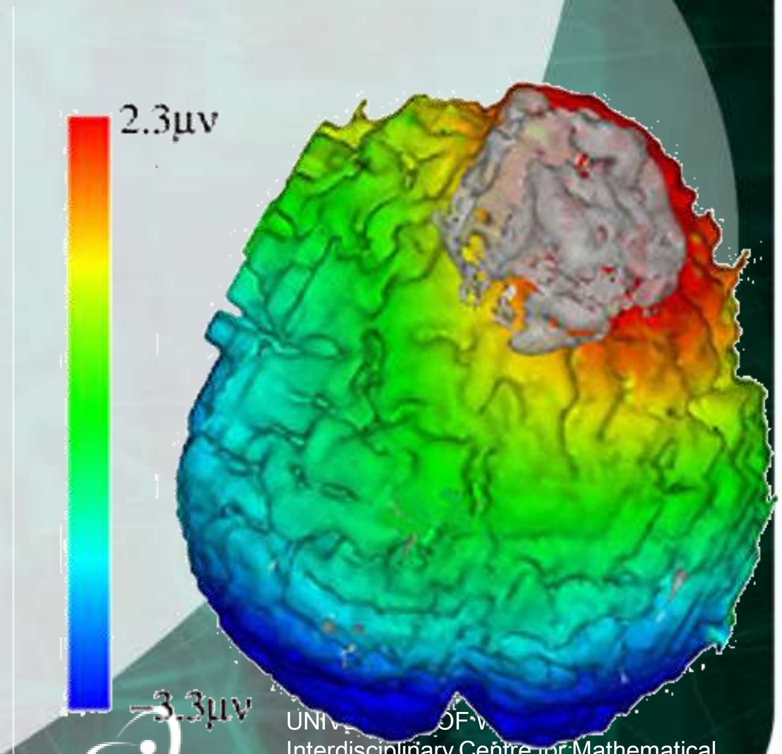

Thrombosis: an outline of temporal sequencing and effects

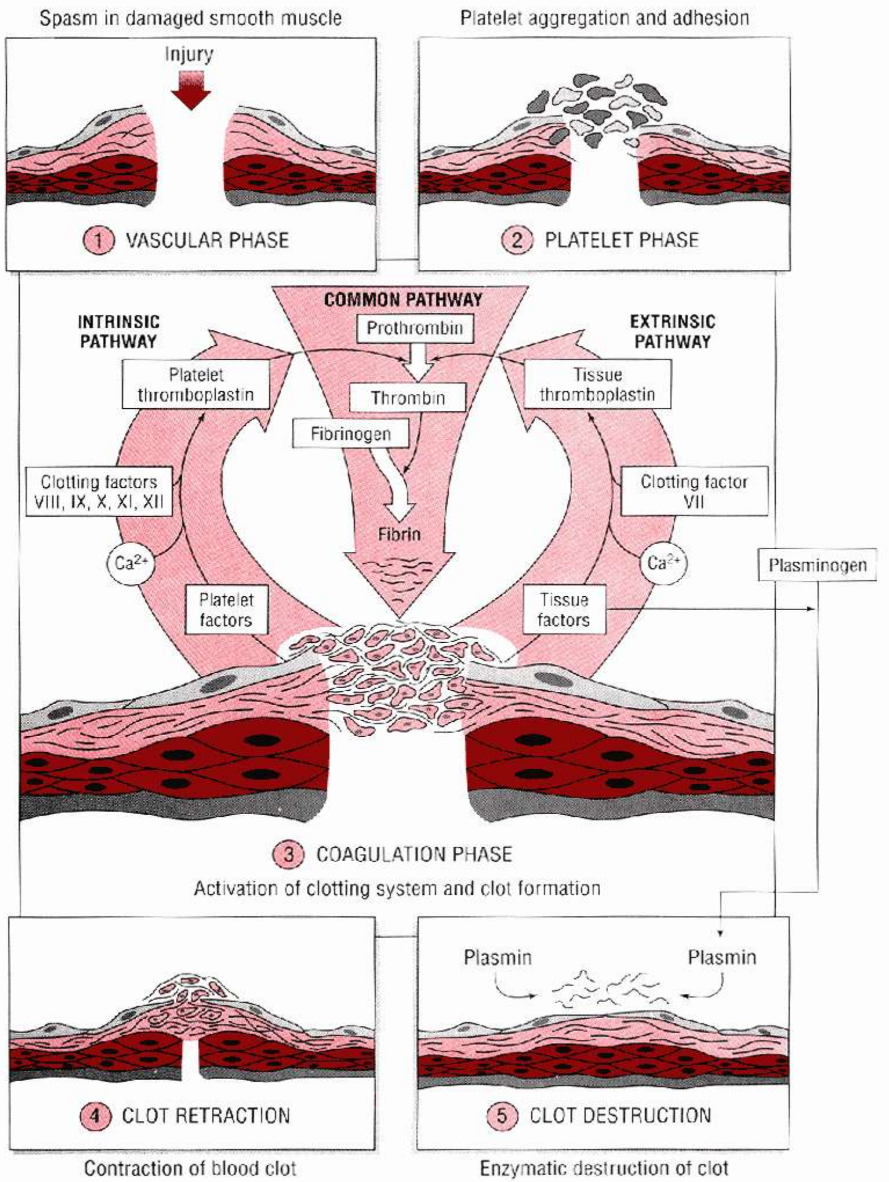

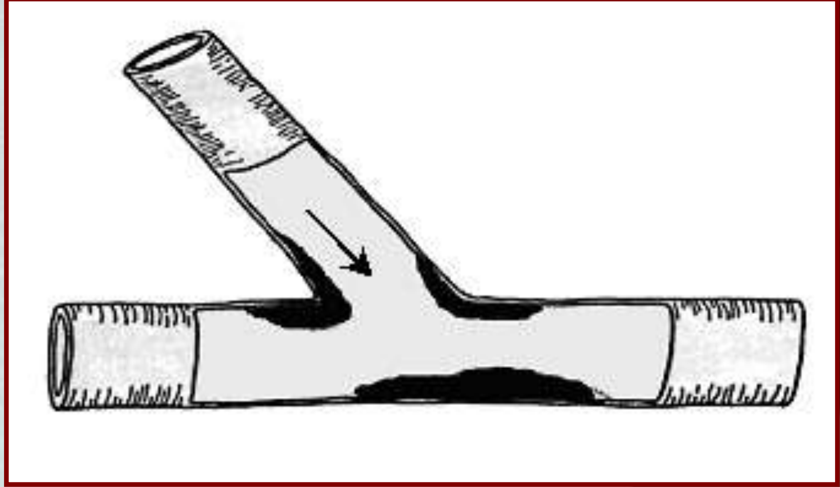

Process dynamics in blood circulation

- variable geometry
- free boundaries
- complex flow nature
- chemical reactivity
- multiscale
- granularity
- role of mechanics
- rheological features
- thermal sensitivity **(→HSPs)**

Case studies: references

- **Prion spread in brain tissue:**
  - switching between different transformation mechanisms (***simple conversion vs. nucleated polymerisation***)
  - role of scrapie prions
  - multiscale irreversible nature
  - nonlocal interactions
  - hybrid geometry
  - free boundary of disease spread

- Role of **heat shocks** and **Heat Shock Proteins** (HSPs) in degenerative developments of biosystems:
  - irreveversible effects
  - process stabilisation

Few open problems for HSP modelling

- Phenomenological model generator for HSP interventions (reactivity hypotheses)
- Nonlocal-in-space-and-time interaction of HSP with other agents (nonlinear diffusion with nonmonotonicities, memory terms, nonlocal contributions)
- Formation of critical HSP concentration profiles (free boundaries)

Special prospective application area: hyperthermia treatments in medicine

UNIVERSITY OF WARSAW
Interdisciplinary Centre for Mathematical
and Computational Modelling
www.icm.edu.pl

Personalized medicine: hyperthermia, an outline

- Thermal treatment based therapeutic procedures:
  - hyperthermia & hypothermia
  - ablation
- Multimodal procedures:
  - immune response stimulation (involvement of HSPs)
- Research framework: some references

Hyperthermia: objectives

- Regional/local hyperthermia of deep-seated tumors with the aid of a phased array of antennas surrounding the patient
- Distribution of the power absorbed controlled by means of amplitude, voltage and temporal profile of antennas' driving voltages
- Space between antennas and body filled with water bolus to prevent excessive skin heating
- Modelling and optimization challenges:
  - modelling and calculation of the EM field
  - and the forcing temperature,
  - optimization of the channel adjustments to reach favorable inference

HSP: more on their activating role

- Expression of HSPs, in particular that of HSP70, protects cells from heat-induced apoptosis

  → until quite recently HSP expression has been viewed a complicating factor in hyperthermia

- Recent results indicate substantial contribution of larger HSPs (HSP70 and higher) to immune reactions
- *Involving HSP expression induced by hyperthermia into tumor immunity, novel anti-cancer immunotherapy based on this concept can be developed*
- Underlying idea:
  - Tumor-specific hyperthermia system that can heat the local tumor region to the desired temperature without damaging normal tissue would be highly advantageous.
  - Posssible heating technology via intracellular hyperthermia: based on using magnetic nanoparticles (approach that can induce necrotic cell death via HSP expression which activates anti-tumor immunity)

Hyperthermia concept as induced by HSPs expression

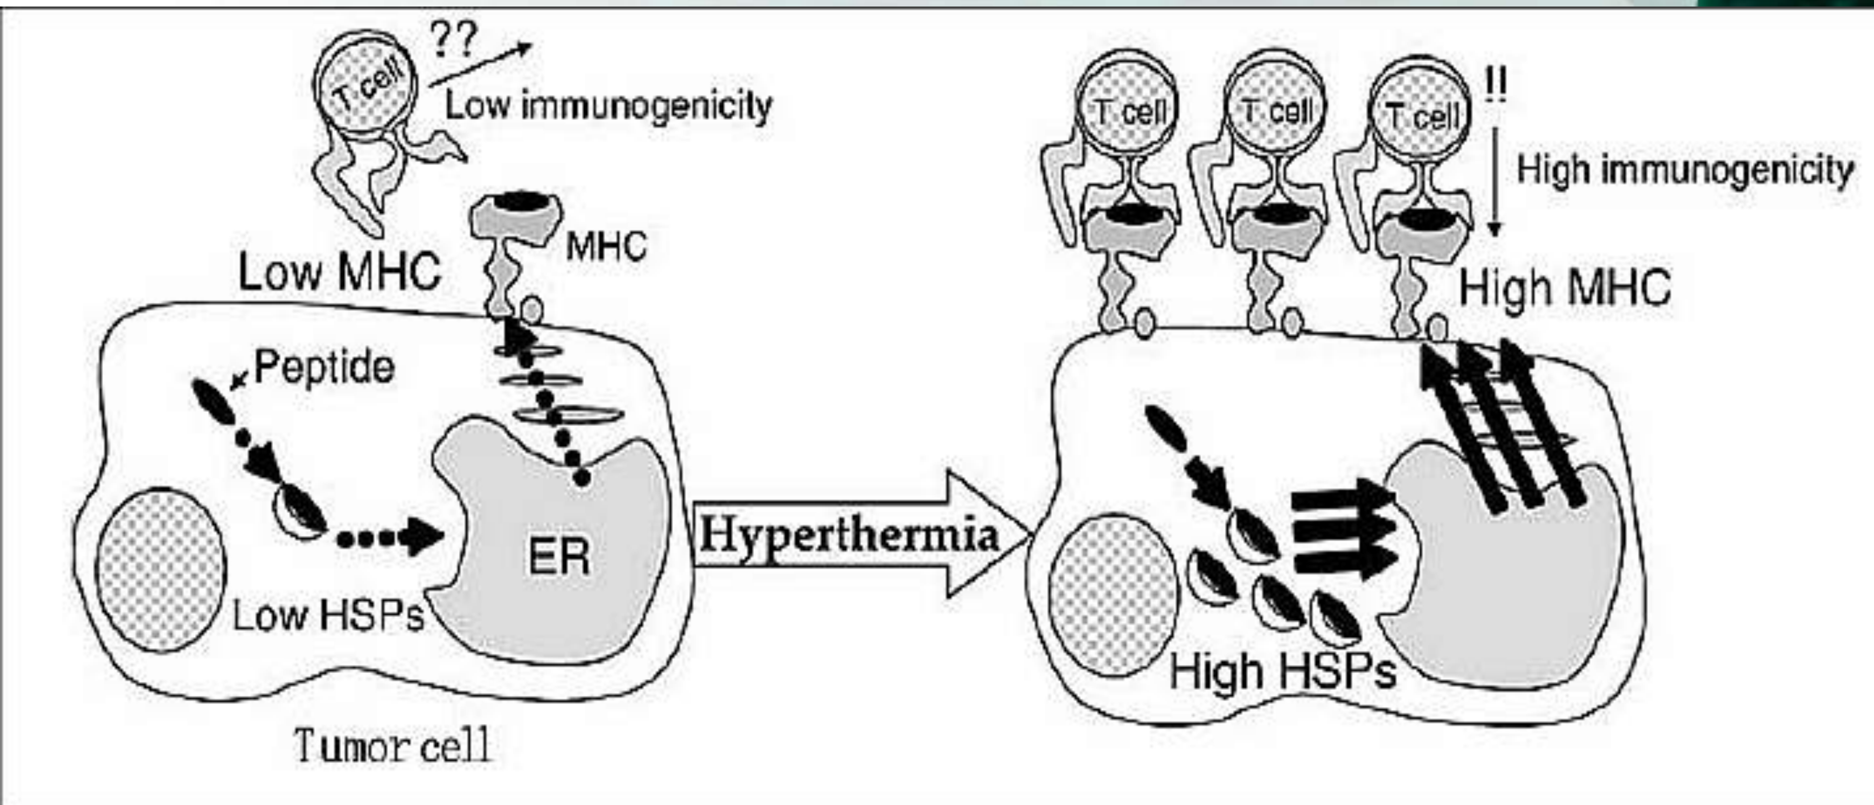

Problems

- Construction of a phenomenological model generator for HSP interventions:
  - iterative formulation and validation of reactivity hypotheses
  - experiment design support
- Nonlocal-in-space interaction of HSP with other agents
- Formation of critical HSP concentration profiles: free boundaries

Protein aggregation

- Without HSP mediation:

  a competition between thermodynamic driving forces and spontaneous protein refolding

- Hyperthemia:

  heating tissues to 42-43 ºC,

  a method of cancer therapy

Parallel computing – challenges

- Parallel programming is still difficult especialy while traditional programming paradigms are used
- There is need for new programing paradigms such as Partitioned Global Address Space (PGAS)
- There is need for new programming languages which can easily express parallelism (Fortress, X10, Click…)
- HPC marked has to open for new languages widely used for data analysis such as Java
- New parallel algorithms for all application areas are necessary

PCJ - Parallel Computations in Java

Java library developed at ICM

- pcj.icm.edu.pl

Programming paradigm:

- partitioned global address space (PGAS)

Features

- Does not require modification of JVM
- Work on *almost all operating system that have JVM*
    - *eg. IBM Java 1.7 on Power7 architecture*
- Uses newest Java SE 7 (NIO, SDP, . . .)
- Works with Java SE 8
- Does not require other libraries!

PCJ – Parallel Computations in Java

Basic functionality:

- Synchronization of tasks
- getting values
- putting values

Advanced functionality:

- broadcasting values
- monitoring variables
- parallel I/O
- creating groups of nodes
- working with groups.

PCJ for HPC and BigData

- PCJ performance is competitive compare to standard solutions based on MPI
- For single node PCJ performance is competitive compared to Java 8 parallel streams
- PCJ runs on multiple nodes (Java 8 is limitted to single JVM)
- PCJ has very good scalability and has been run on 10k cores
- PCJ allows for very efficient I/O operations on distributed and network file systems.
- PCJ can be used to parallelize data analysis codes in Java