The VPH Tianhe-2 Challenge

The challenge

The VPH Institute has agreed with the National Supercomputing Center in Guangzhou (NSCC-Guangzhou) - China, to launch a grand challenge to the VPH community worldwide.

NSCC-Guangzhou is the home of Tianhe-2, a 33.86 petaflops supercomputer, which was developed by a team of 1300 scientists and engineers, and according to the TOP500 list is currently the world fastest supercomputer.

For more information visit: http://en.wikipedia.org/wiki/Tianhe-2

The scope of this challenge is to identify a limited set of codes that if ported to the Tianhe-2 HPC platform would in principle allow breakthroughs in computational biomedicine.

The codes we aim to identify need to have some potential:

- from one side, they should be of interest for a significant number of VPH researchers worldwide
- from the other side, they should enable the solution of problems that otherwise would be difficult (or impossible) to solve with more conventional computers because of their excessive solution time.

If we will be able to find codes that respond to these requirements, the Guangzhou team will consider investing some of their experts' time to port such codes to the Tianhe-2 HPC platform, in collaboration with the code developers. If accomplished, this initiative will bring a great value not only to our research community, but also to computational biomedicine at large. The VPH Institute is proud to have been selected as facilitators for this exciting process of exploring what are the codes that could be suitable for this grand challenge.

As a first step we asked the members of the VPH Institute to compile the list of problems, methods, and codes that they think should be ported to the Tianhe-2 HPC platform. Here below are the results. The questions we posed were:

Problems: What are the computational biomedicine problems that you experience in your daily research?

Methods: Could these problems be solved with numerical methods?

Codes: Are there open source codes or commercial codes that can solve them efficiently?

Problems

What are the computational biomedicine problems that could be solved only if an appropriate solver was ported to an extremely large supercomputing platform?

1. **Biomechanics, solids, mono-scale**: whole body finite element models. The USA Visible Human female has datasets of 57 billion voxels with a spatial resolution of 330 microns. If we could solve a Cartesian finite model at this resolution we could explore how the various tissues types and structures contribute to the passive biomechanics of the whole body. This would have huge implications for biomedical,

defence, and sport [*Marco Viceconti, Sheffield*].

2. **Mechanobiology, solids, multiscale**: if we could solve in parallel 10 million ODE models coupled each to a cell of a large scale finite element model we could explore a number of mechanobiology problems for bones and muscles where the tissue biomechanics PDE model is coupled with cell-tissue ODE model [*Marco Viceconti, Sheffield*].

3. **Estimating ionic conductance and molecular permeability of gap junction channels from numerical simulations that combine quantum mechanics, molecular dynamics and coarse-grained models**. Gap junction channels provide direct electrical and metabolic communication between adjacent cells. Alterations of channel permeability due to point mutations or deletions in connexin proteins are thought to underlie several other human diseases. A theoretical rationalization of the experimental results is of fundamental importance in order to understand the molecular mechanisms that regulate the permeation in gap junction channels and their pathological alterations [*Fabio Mammano, BMS, Sheffield*].

4. **Multiscales and multilevels simulation of human bone from macroscale (whole organ) to nanoscopic scale (collagen, mineral, cross links, ...) considering the bone remodeling and the bone cells activities.** This needs to model the problem with different metamodel containing different blocs and to link them into one multiscale integrated remodeling-to-fracture metamodel to describe comprehensive and enhanced human bone behaviour. This would be very beneficial to develop computer models in order to investigate among the others the effects of treatments strategies to prevent osteoporotic fracture. Abaqus code can be used for these developments [*Prof. Ridha Hambli, Polytechnic Institute, Orleans, France*].

5. **Longer-term blood flow predictions for patient-specific vascular networks of the Circle of Willis**. Here we require a 250 million lattice site model, to be run under a diverse set of flow conditions to reflect fluctuations in blood flow for a patient leading a regular life. We use the results of this extensive analysis to identify the activity regimes where patients may experience unusual levels of wall stress or flow velocities in one or more cerebral arteries. This research requires exascale resources at its largest scale and has the potential to help mitigate stroke risks for patients with aneurysms. [*Peter Coveney and Derek Groen, UCL, London, UK*]

6. **Understanding how cardiac morphogenesis is driven by cellular signalling in response to signals from tissue stress fields.** Using experimental measurements of shape change and tissue structure coupled with continuum mechanics models of the growing heart to predict cell proliferation, growth, apoptosis and migration and investigate how these influence the looping phenomenon in the growing heart. Continuum mechanics models will provide a mathematical framework for solving the coupled physical laws that govern electrical activation processes, large deformation mechanics, 3D fluid flow and fluid-structure interaction. These multi-scale, multi-physics models of the growing heart generate very large computational problems, which require the use of high performance computing resources [Peter Hunter, University of Auckland, Auckland, NZ].

7. **Accurate simulations of subcellular calcium dynamics for both healthy and pathological situations.** It can be estimated that a single simulation of one sarcomere with a uniform 1 nanometer spatial resolution and for 1 ms in time requires more than 10^19 floating-point operations. Using Tianhe-2, we have already run proof-of-concept 3D simulations of one sarcomere using a spatial resolution of 3

nanometer, see Ref[4]. To reach the ultimate 1-nanometer resolution, it requires a large amount of computing time on Tianhe-2, together with further improvements of the existing in-house software code that was jointly developed by Simula Research Lab (Norway) and NUDT (China) [*Xing Cai, Simula Research Lab, Norway*].

8. **Multiphysical modeling of intestine chemo-electro-viscoelasticity to explore the clinical problem of paralytic ileus following main abdominal surgeries.** Using experimental measurements of thermo-electrical and electro-mechanical patterns during intestine dysrhythmia, we could translate most of the studies applied to the heart (see point 6) to the much less explored area of gastroenterology from a mathematical point of view. The main result of 3D patient-specific HPC intestine computational analysis will consist in identifying the best clinical practice for reducing the hospital stay thus reducing the overall healthcare costs. Moreover, a precise indication to where perform surgeries will be fundamental for colon cancer. Collateral applications will be a fine estimation of gastrointestinal energetics and dissipation [*Simonetta Filippi, University Campus Bio-Medico of Rome, Rome, Italy*].
9. **Patient-specific multi-scale models of cardiac electromechanics that include cellular mechanisms of excitation-contraction coupling, 3D nonlinear anisotropic microstructural models of tissue mechanics and conductivities, 4-chamber cardiac geometry and fiber architecture and whole body circulatory dynamics models**. These models with 1-10 million degrees of freedom must be optimized to multimodal patient data sets including cardiac CT or MRI imaging CT, electroanatomic maps, cardiac catheterization, SPECT perfusion scans, cardiac echo and electrocardiograms. The aim of this project is to understand the therapeutic mechanisms of cardiac resynchronization therapy in patients with dyssynchronous heart failure. Patient data sets have been acquired with informed consent. [*Andrew mcCulloch - University of California San Diego*]
10. **Patient-specific modeling of atrial fibrillation using finite element meshes with 100 million degrees of freedom with parameters optimized to electroanatomic maps are being used to understand mechanisms of a new strategy for AF ablation.** Large patient data sets have been acquired with informed consent.[*Andrew mcCulloch - University of California San Diego*]
11. **Long time scale multi-scale models of hypertrophic growth and remodeling in the rodent heart for 4-8 weeks following aortic banding**. Data sets are being acquired now. [*Andrew mcCulloch - University of California San Diego*]
12. **Multi-scale models of cardiac gene mutations in mice that span from MD simulations to molecular electrostatics to Brownian dynamics to Markov models of cardiac contractile regulatory proteins especially troponins and PKA that link to 3D multi-scale models of mouse heart mechanics**. [*Andrew mcCulloch - University of California San Diego*]
13. **High resolution models of intracellular second messenger signaling in 3D cardiac myocyte models reconstructed using electron tomography.** Two papers have been published and additional high quality data sets are being obtained. [*Andrew mcCulloch - University of California San Diego*]
14. **Multi-scale modeling of function, metabolism, degeneration, and regeneration of livers to improve disease diagnosis, prognosis and treatment, and evaluation of drug toxicity.** Relevance: The liver is the main detoxifying organ removing food toxins and drugs from the blood. It is able to regenerate up to about 2/3 of its mass and complex functional architecture. More than 100 liver diseases exist, for example alcohol, virus infections, obesity or drugs can lead to liver cirrhosis representing the 11th

frequent cause of death in the US. Anatomy: The human liver consists two anatomical lobes, each organized into hundreds of thousands of lobules, small functional subunits with a complex architecture for maximal exchange of metabolites between hepatocytes (liver parenchymal cells) and the blood. Many liver enzymes are known to be spatially patterned (zonated). State in modeling: An agent-based biophysical model representing each individual hepatocyte exists, able to mimic the regeneration process in one lobule on the tissue scale. Signaling pathway models, metabolic models, models of blood flow in the larger liver veins and arteries and the smallest capillaries exist as well as a models of the extrahepatic compartment. Aim: The different models should be integrated, later refined, into one multiscale model spanning intracellular molecular, spatial temporal tissue, organ and body scale taking into account spatial and zonal inhomogeneities to improve clinical decision making in disease management, e.g., of non-alcoholic fatty liver disease, or/and improve regeneration after drug-induced damage or partial organ removal. [*AG Preusser, Fraunhofer MEVIS Bremen; AG Drasdo, Inria Paris/IZBI Leipzig*]

Methods

Which numerical methods are necessary to solve the problems above?

1. Finite element method to solve 3D elasticity problems [*Marco Viceconti, Sheffield*].
2. Finite element method to solve 3D elasto-plasticity problems [*Marco Viceconti, Sheffield*].
3. Coupled PDE-ODE problems, where each PDE cell is couple to a different initial value for the same ODE model [*Marco Viceconti, Sheffield*].
4. Numerical simulations combining quantum mechanics, molecular dynamics and coarse grained models can be used to predict conductance and permeability of gap junction channels (and other membrane channels too) [*Fabio Mammano, BMS, Sheffield*].
5. We need multiscale software for many of the problems i.e. coupling continuum solvers (typically FEM/FVM) to atomistic methods (MD/DPD). This means overcoming disparities in both temporal and spatial scales and techniques remain in their infancy [*David Emerson, SCD, Daresbury*]
6. To efficiently use Tianhe-2 will require software that can exploit the Xeon Phi co-processor. Very few codes can currently do this. In addition, the code will have to be able to scale very well. I cannot see how commercial codes will be able to exploit the hardware and the license costs could be prohibitive. [*David Emerson, SCD, Daresbury*]
7. We will need Abaqus code with FORTRAN to develop bone-specific material behaviour related to each scale [*Prof. Ridha Hambli, Polytechnic Institute, Orleans, France*]
8. Fluid-Structure interaction problems [Eduardo Soudah, CIMNE]
9. Sophisticated flow predictions in cerebrovascular networks can be made with a multi-model computational approach centred on using a highly parallel lattice-Boltzmann simulation. In addition, it requires an analysis environment that extracts key findings from the simulation results. [*Peter Coveney and Derek Groen, UCL, London, UK*]
10. Finite element methods for coupled equations of electrical propagation, large deformation elasticity

with active contraction, fluid dynamics and fluid-structure interaction. Further coupling with DAEs associated with growth, cellular signalling, cellular electro-mechanics and constitutive laws. On Tianhe-2 the continuum PDE equations could be solved in parallel using domain decomposition with MPI or hybrid MPI/OpenMP and the Phi co-processors could be used to accelerate the solution of DAEs [*Peter Hunter, University of Auckland, Auckland, NZ*].

11. 3D finite difference based computations of five coupled reaction-diffusion equations, in addition to two ordinary differential equations. Moreover, each Ryanodine receptor is subject to a stochastic Markov-chain computation per time step, deciding whether calcium release happens or not. Adaptive mesh refinement will be attempted as one improvement of the current numerical strategy to facilitate 1-nanometer-resolution 3D simulations on Tianhe-2 [*Xing Cai, Simula Research Lab, Norway*].
12. Finite element methods as in point 10 enriched with time dependent processes characterized by viscoelasticity time constants for 3D electro-mechanical simulations in presence of multiple diffusive fields, i.e. temperature. Parallel segregated solvers could be used based on the different time scales of the different physics involved [*Simonetta Filippi, University Campus Bio-Medico of Rome, Rome, Italy*].
13. Problem 9/10/11: use coupled multi-scale finite element methods that currently exploit GPU acceleration [*Andrew mcCulloch - University of California San Diego*]
14. Problem 12: use Molecular Dynamics, Molecular Electrostatics, Brownian Dynamics, Monte carlo Methods and Finite Element Methods [*Andrew mcCulloch - University of California San Diego*]
15. Problem 13: Uses new FE methods a for reaction-diffusion problems. [*Andrew mcCulloch - University of California San Diego*]
16. Problem 14: The system of equations to be solved consists of several advection, diffusion, reaction equations that are coupled on multiple scales [5]. Thus, for the numerical solution of the problem a combination of FEM, ODE and stochastic ODE integrators and agent-based methods [6, 12] will be utilized. [*AG Preusser, Fraunhofer MEVIS Bremen; AG Drasdo, Inria Paris/IZBI Leipzig*]

Codes

Which open source or commercial codes do solve efficiently the numerical problems above, and thus would be worth to be ported to an extremely large HPC platform?

1. Ansys Multiphysics: http://www.ansys.com/ [*Marco Viceconti, Sheffield*].
2. ParFE: http://parfe.sourceforge.net/ [*Marco Viceconti, Sheffield*].
3. OpenCMISS: http://physiomeproject.org/software/opencmiss/ [*Marco Viceconti, Sheffield*] [*Peter Hunter, University of Auckland, Auckland, NZ*].
4. Chaste: http://www.cs.ox.ac.uk/chaste/ [*Marco Viceconti, Sheffield*].
5. FEBio: http://www.febio.org [*Marco Viceconti, Sheffield*].
6. LifeV: http://www.lifev.org and https://cmcsforge.epfl.ch/projects/lifev. It is a C++ parallel (MPI+OpenMP) finite element library able to solve fluid-structure interaction, multiscale, Darcy flows in porous media, heart electro-mechanical coupling problems [7,8,9,10,11]. It has been used in the Eu

(http://www.mathcard.eu/). It's currently used on BG/Q "Fermi" supercomputer at CINECA (Italy) [*Rocco Michele Lancellotti, MOX laboratory, Politecnico di Milano*].

7. Gromacs: http://www.gromacs.org/ [*Fabio Mammano, BMS, Sheffield]*.
8. Gaussian: http://www.gaussian.com/ [*Fabio Mammano, BMS, Sheffield*].
9. OpenFOAM: http://www.openfoam.com/ [*David Emerson, SCD, Daresbury*]
10. Code_Saturne: http://code-saturne.org/ [*David Emerson, SCD, Daresbury*]
11. Abaqus code and Fortran compiler to run abaqus routines and develop material models [*Prof. Ridha Hambli, Polytechnic Institute, Orleans, France*].
12. Code_Aster: http://www.code-aster.org/V2/spip.php?rubrique4 Structural analysis, Open Source [*Alessandro Chiarini, SCS, Bologna*]
13. KRATOS Multi-physics: www.cimne.com/kratos [*Eduardo Soudah, CIMNE*]
14. The HemeLB lattice-Boltzmann simulation environment (http://ccs.chem.ucl.ac.uk/hemelb) [1,2] is optimized for vascular networks, and can make effective use of peta- and exascale resources when applied to model a range of velocity profiles. HemeLB also comes with an extensive environment for data analysis, supports coupling to other models (e.g., 1D flow [3] and vascular remodelling models) for multiscale simulation and provides capabilities for monitoring and interactively steering simulations. To incorporate a range of flow conditions we will require a 10 petaflop+ system such as Tianhe-2 to accurately model bloodflow in a larger network of arteries. [*Peter Coveney and Derek Groen, UCL, London, UK*]
15. An in-house code that is jointly maintained by Simula Research Lab (Norway) and NUDT (China). The code has already been tuned for running on Tianhe-2, see Ref [4]. The so-far largest test run on Tianhe-2 used 4096 nodes, i.e., 12288 Xeon Phi coprocessors, achieving 1.27 Pflops in double precision [*Xing Cai, Simula Research Lab, Norway*].
16. Problem 9/10/11: use Continuity 6.3 (http://www.continuity.ucsd.edu) which is public domain, primarily Python and C with some legacy f77 and run-time generated C and Cuda code with JIT compilation. Supported by the NIH for 20 years. [*Andrew mcCulloch - University of California San Diego*]
17. Problem 12: use aMD, APBS, BrownDye and Continuity that are optimized for distributed memory parallelization primarily using MPI but also include some GPU acceleration (nVidia). All of these codes are freely available from the National Biomedical Computation Resource http://*www.nbcr.net [Andrew mcCulloch - University of California San Diego*]
18. Problem 13: Uses SubCell https://launchpad.net/subcell which is Open Source and Python and uses the FEniCS open source library. [*Andrew mcCulloch - University of California San Diego*]
19. Problem 14: We utilise in-house developed codes on C/C++ basis. A modular, open source software code (based on former CellSys, [13]) for spatial temporal tissue units integrating intracellular signaling (ODE) and flows is short before delivery. The code includes an image processing/analysis tool that can directly feed the simulation code with histological images and is to be integrated with the code for organ-scale simulations [5]. [*AG Preusser, Fraunhofer MEVIS Bremen; AG Drasdo, Inria Paris/IZBI Leipzig*]

References

Groen et al.: "Analysing and modelling the performance of the HemeLB lattice-Boltzmann simulation environment", JoCS, 2013, DOI: 10.1016/j.jocs.2013.03.002.

Nash et al.: "Choice of boundary condition for lattice-Boltzmann simulation of moderate Reynolds number flow in complex domains", PRE, 2014. DOI: 10.1103/PhysRevLett.112.068103.

Groen et al. "Flexible composition and execution of high performance, high fidelity multiscale biomedical simulations", Interface Focus, 2014. DOI: 10.1098/rsfs.2012.0087.

Chai et al.: "Towards simulation of subcellular calcium dynamics at nanometre resolution", The international journal of high performance computing applications, 2013, DOI: 10.1177/1094342013514465.

Schwen et al., Spatio-Temporal Simulation of First Pass Drug Perfusion in the Liver. PLOS Computational Biology, 10(3), e1003499, 2014. DOI: 10.1371/journal.pcbi.1003499

Schliess et al., Integrated metabolic spatial-temporal model for the prediction of ammonia detoxification during liver damage and regeneration, Hepatology (2014), in press. DOI: 10.1002/hep.27136

Deparis et al., "Parallel preconditioners for the unsteady Navier-Stokes equations and applications to hemodynamic simulations", Computers & Fluids, 01/2014, http://dx.doi.org/10.1016/j.compfluid.2013.10.034.

Nobile et al., "Time accurate partitioned algorithms for the solution of fluid-structure interaction problems in haemodynamics", Computer & Fluid, 11/2013, http://dx.doi.org/10.1016/j.compfluid.2013.07.031.

Rossi et al., "Thermodynamically consistent orthotripic activation model capturing ventricular systolic wall thickening in cardiac electromechanics", European Journal of Mechanics - A/Solids, in press, http://dx.doi.org/10.1016/j.euromechsol.2013.10.009.

Faggiano et al., "Helical flows and asymmetry of blood jet in dilated ascending aorta with normally functioning bicuspid valve", Biomechanics and Modeling in Mechanobiology, 08/2013, DOI: 10.1007/s10237-012-0444-1.

Palamara et al., "An effective algorithm for the generation of patient-specific Purkinje networks in computational electrocardiology", submitted.

Hoehme, S., et. al. (2010). Prediction and validation of cell alignment along microvessels as order principle to restore tissue architecture in liver regeneration. Proc. Natl. Acad. Sci. (USA), 107(23), 10371-10376.

Hoehme S and Drasdo D. (2010). A cell-based simulation software for multicellular systems. Bioinformatics, 2010 Oct 15;26(20):2641-2