D6.4 - Exemplar scientific simulations

WP6: Co-design via applications

Project Acronym	CRESTA
Project Title	Collaborative Research Into Exascale Systemware, Tools and Applications
Project Number	287703
Instrument	Collaborative project
Thematic Priority	ICT-2011.9.13 Exascale computing, software and simulation

Due date (if applicable):	M30
Submission date:	31/03/2014
Project start date:	01/10/2011
Project duration:	39 months
Lead organization	CSC
Version:	1.0
Status	Final
Author(s):	Mark Abraham (KTH), Mikko Byckling (CSC), Willem Deconinck (ECMWF), Jing Gong (KTH), Derek Groen (UCL), Mats Hamrud (ECMWF), Berk Hess (KTH), Petter Johansson (KTH), Timo Kiviniemi (AALTO), Timo Krappel (USTUTT), Erik Lindahl (KTH), Magnus Lundborg (KTH), George Mozdzynski (ECMWF), Adam Peplinski (KTH), Iman Pouya (KTH), Stephen Sachs (CRAY UK), Sebastian Schmieschek (UCL), Jan Westerholm (ABO), Jan Åström (CSC)
Reviewer(s):	Alan Gray (UEDIN), David Henty (UEDIN)

Software dissemination level	
PU	*PU - Public*

Copyright © CRESTA Consortium Partners 2014

Version History

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	26/02/2014	IFS contribution	George Mozdzynski (ECMWF), Mats Hamrud (ECMWF), Willem Deconinck (ECMWF)
0.2	28/02/2014	ELMFIRE contribution	Timo Kiviniemi (AALTO), Jan Westerholm (ABO)
0.3	28/02/2014	Nek5000 contribution	Adam Peplinski (KTH), Jing Gong (KTH)
0.4	28/02/2014	OpenFOAM contribution	Timo Krappel (USTUTT), Stephen Sachs (CRAY UK)
0.5	28/02/2014	GROMACS contribution	Mark Abraham (KTH), Berk Hess (KTH), Petter Johansson (KTH), Erik Lindahl (KTH), Magnus Lundborg (KTH), Iman Pouya (KTH)
0.6	03/03/2014	Corrections	Mikko Byckling (CSC)
0.7	04/03/2014	HemeLB contribution	Sebastian Schmieschek (UCL), Derek Groen (UCL)
0.8	04/03/2014	Corrections	Mikko Byckling (CSC)
0.9	04/03/2014	Submitted for internal review	Mikko Byckling (CSC)
0.9.1	19/03/2014	Internal review	Alan Gray (EPCC), David Henty (UEDIN)
0.9.2	20/03/2014	Addressed reviewer comments	Mikko Byckling (CSC)
0.9.3	20/03/2014	Minor corrections	Mikko Byckling (CSC)
1.0	21/03/2014	Final version for EC review	Mikko Byckling (CSC)

Copyright © CRESTA Consortium Partners 2014

Table of Contents

Copyright © CRESTA Consortium Partners 2014

Index of Figures

Copyright © CRESTA Consortium Partners 2014

Index of Tables

Copyright © CRESTA Consortium Partners 2014

Copyright © CRESTA Consortium Partners 2014

1 Executive Summary

This document describes exemplar scientific simulations for each CRESTA application. In the final two deliverables of work package 6, i.e., “D6.1.2 Roadmap to exascale (Update 2)” and “D6.5 Peta to exascale enabled applications”, these simulations will be used to show the success of the software development carried out during the project.

In the following, we give a short overview of each simulation and their scientific impact for each application.

ELMFIRE: In order to optimize the structure of fusion reactors, understanding of turbulent plasma transport is required. With exemplar scientific simulations in ELMFIRE, the scientists seek to better understand confinement properties of the fusion plasma. For computations at a small scale, an excellent match between the experiments and computational results has been observed. For modelling ITER-sized problems, exascale computing is needed.

GROMACS: With GROMACS, two different types of exemplar simulations are described. The first type aims at a better understanding of liquid/surface interactions at an atomic level and the second type considers the modelling of state transfers in biomolecular systems. Both type of simulations have direct applications in medicine and in drug design. Exascale computational resources are needed for modelling of realistic systems up to the continuum limit.

HemeLB: Exemplar simulations in HemeLB consider cerebrovascular blood flow in an authentic geometry. Since some aspects of such flows, such as wall-shear stresses, are difficult to obtain or reproduce in a clinical setting, simulations allow increased understanding and better treatment of conditions that can ultimately cause health issues, financial cost and ultimately even the death of a patient. Exascale computational resources are required in order to run the simulations within prespecified clinical time limits.

IFS: Enabling more detailed grid resolutions having added physical interactions is known to increase the accuracy of numerical weather forecasts. This can help save lives as well as result in capital savings.

With exemplar scientific simulations of IFS, by using actual data from the hurricane Sandy, the meteorologists aim to understand how the approximation with better grid resolution and additional physics relate to the accuracy of the numerical model. Another objective is to study if it is feasible to do computations with increased resolution and complexity within the ECMWFs operational requirement of computing 240 weather forecast days per day.

Nek5000: Simulation of incompressible fluid flow in pipes with a circular cross-section is important in many applications such as gas and oil pipelines. However, in order to solve the resulting problems at scale with turbulence, exascale computational resources are required. To this end, by using OpenACC directives, Nek5000 has been enabled to exploit GPUs.

OpenFOAM: Pump turbines are commonly used in hydro power plants to generate electricity. The exemplar simulation with OpenFOAM is to model a whole Francis pump turbine. This enables the engineers to study the behaviour of the flow vortex and the resulting pressure pulsations in order to improve the turbine design. Since a long time evolution combined with a fine timescale is needed, exascale computational resources are required.

Copyright © CRESTA Consortium Partners 2014

2 Exemplar scientific simulations

CRESTA applications, ELMFIRE, GROMACS, HemeLB, IFS, Nek5000 and OpenFOAM, all have a scientific need for computations at an exascale. In the case of these applications, the need for large-scale computations arises mainly from the complexity of the physics and the sheer size of the physical models to be simulated.

Although simple benchmarks are commonly used to evaluate the performance of scientific applications, the conclusions drawn from them can be occasionally quite misleading. We stress that the simulations described, apart from the cases where scalability has been studied, should not be mistaken for such simple benchmark runs. Each exemplar simulation is motivated by an actual scientific need to better understand and model the phenomena under study. This has been reflected already in the deliverable "D6.2 Needs analysis", in which, for each application, the critical components in need of modification for enabling successful scaling up to an exascale level have been documented.

In this document, an exemplar simulation with potential exascale computational need is presented for each CRESTA application. At the end of the project these simulations, alongside with standard application benchmarks, will be used to show the improved scalability of the applications. After the project has ended and the development of the applications continues, the simulations can be used as development tools when scalability and performance evaluation is being done.

None of the simulations described in this document require exascale resources in order to be computed. They are designed in an expandable way in order to be exploitable once the first exascale systems become available.

Exemplar simulations for different applications are covered in separate sections. In the beginning of each section, a very brief introduction to the application is given. The first subsection, "scientific use case", presents the exemplar scientific use case as well as its computational requirements. The second subsection, "exemplar simulation", describes one or more example exemplar simulation runs in detail: the timescales for the simulations, list of people who performed the simulations, the computing environment in which the simulations were performed in as well as scalability data about the execution of the simulations. The third subsection, "scientific impact", describes the scientific impact of the exemplar simulations and highlights gains achievable with exascale computations. References are listed in the fourth and the final subsection.

We note that for some applications the results given in the second subsection, "exemplar simulation", have been computed with a version of the code containing modifications implemented during the project. A more detailed description of the impact of the project to the performance and scalability of the applications will however be discussed in the deliverables "D6.1.2 Roadmap to exascale (Update 2)" and "D6.5 Peta to exascale enabled applications".

2.1 Glossary of Acronyms

API	Application Programming Interface
CAF	Coarray Fortran
CFD	Computational Fluid Dynamics
CMC	Canadian Meteorological Centre
CoW	Circle of Willis
CPU	Central Processing Unit
DNS	Direct Numerical Simulation
DOPC	Di-Oleoyl-PhosphatidylCholine
ECMWF	European Centre for Medium-Range Weather Forecasts
FD/D	Forecast Day/Day

Copyright © CRESTA Consortium Partners 2014

GAM	Geodesic Acoustic Mode
GB	Gibabyte
GGI	General Grid Interface
GLIC	Gloeobacter violacaeus Ligand-gated Ion Channel
GPU	Graphics Processing Unit
HPC	High-performance Computing
IDDES	Improved Delayed Detached Eddy Simulation
ITER	International Thermonuclear Experimental Reactor
LES	Large Eddy Simulation
MCA	Middle Cerebral Artery
MPI	Message Passing Interface
MSLP	Mean sea level pressure
NCEP	National Centers for Environmental Prediction
NDBC	National Data Buoy Center
NWP	Numerical Weather Prediction
OpenACC	Open Accelerators
OpenMP	Open Multiprocessing
ORNL	Oak Ridge National Laboratory
PF	Petaflop
PME	Particle Mesh Ewald
PRACE	Partnership for Advanced Computing in Europe
RAM	Random Access Memory
RANS	Reynolds Averaged Navier-Stokes
SOL	Scrape-off Layer
SST	Shear-Stress Transport
TEXTOR	Tokamak EXperiment for Technology Oriented Research
UKMO	United Kingdom Meteorological Office
UTC	Coordinated Universal Time

Copyright © CRESTA Consortium Partners 2014

3 ELMFIRE

ELMFIRE is a particle-in-cell code that simulates the movement and interaction between extended gyrokinetic particles moving at high speed in a torus-shaped geometry on a three dimensional grid. The particles are held together by a strong external magnetic field.

ELMFIRE approximates the Coulomb interaction between particles by solving a global electrostatic field on a grid, using the particle charges as sources. ELMFIRE then advances particles in time by free streaming along the magnetic field line and particle drift perpendicular to the magnetic field. Typically, time steps correspond to 30-50ns of real time.

Application	Identification (e.g. Version No)
ELMFIRE	ELMFIRE, version v.12b (1D-decomposition, quasi-ballooning coordinates, without SOL extension and magnetic axis)

Table 3.1 ELMFIRE - version information

3.1 ELMFIRE: scientific use case

3.1.1 Scientific use case: Gyrokinetic full-f simulation of GAMs at Textor plasma edge

Understanding turbulent transport is needed for further optimization of fusion reactors. Unfortunately realistic transport time scale simulations of plasma turbulence are computationally very demanding. The aim of the recent ELMFIRE simulations is to increase the understanding of the mechanisms behind the sudden improvement in confinement observed in experiments by investigating the possible role of the radial derivative of a time-varying electric field in triggering transition. The ELMFIRE turbulence simulation code investigates these phenomena with a so-called first principal computer model. This model tracks individual particles providing information on the complex interplay between the magnetic field, the electric field and the particle trajectories. Using the 30 million CPUh granted from the 4th PRACE call, a scan over local parameters like temperature and density starting from experimental Textor parameters was carried out starting from Textor low confinement mode case for which strong oscillation was observed in the simulations [1].

3.1.2 Computational requirements

ELMFIRE is even more computationally demanding than most of the other turbulence codes as it does not assume a Maxwellian distribution of particles, but instead simulates the full distribution of electrons and ions. While this approach enables more realistic simulations especially near the plasma edge, a very large number of simulation particles are needed in order to have good statistical properties. Previously excellent results have been obtained for a small FT-2 tokamak (R=0.55 m), whereas the current simulations are about a one third of a middle-sized Textor tokamak (R=1.75 m).

Simulations for a whole volume of a real tokamak reactors like ITER (R=6.2 m) require exascale computing. An ITER-sized simulation requires electrostatic grids up to 3000x4000x16. The version of ELMFIRE originally provided for CRESTA does not implement any spatial decomposition, which effectively prevents scaling to grids larger than approximately 120x150x8 regardless of the number of cores available due to data duplication. For an ITER simulation approximately 640000 cores would be needed consisting of about 590 billion particles. With the original version of the code this would require approximately 28TB memory per core and is thus completely unfeasible.

Within CRESTA, a version of the code implementing 3D domain decomposition has been implemented. Performance benchmarks for the scientific test case will be described in the final WP6 deliverable, D6.1.2 Roadmap to exascale (update 2).

Copyright © CRESTA Consortium Partners 2014

3.2 ELMFIRE: exemplar simulation

Simulation	Description
GAMbasecase	ELMFIRE v12.b. 100X210x4, 940 million ions and electrons (11200 per cell). In the series of runs, physics parameters (density, temperature, magnetic field and isotope mass) were varied (about 20 runs altogether).
GAMNscan	ELMFIRE v12b.Testing the effect of noise in GAMbasecase. Test was done in Sisu by reducing number of test particles by factor of 4, 16 and 64.

Table 3.2 ELMFIRE - exemplar simulations

3.2.1 Timescales

Start date:	01/09/12
End date:	31/7/2013
Total duration:	25 million CPU hours (including test runs)

Table 3.3 ELMFIRE - timescales

3.2.2 People

Name	Role
Timo Kiviniemi	Senior research associate, Principle investigator in PRACE project

Table 3.4 ELMFIRE - people

3.2.3 Environment

Platform	Platform specs
SuperMUC	iDataPlex Intel Sandy Bridge-EP (Xeon E5-2680) 2.7GHz, infiniband
Sisu	Cray XC30 (2.6 GHz E5-2670 processors)

Table 3.5 ELMFIRE - environment

3.2.4 Exemplar simulation runs

Simulation	Description of run			
	Nodes	**Total cores**	**Wall time**	**Total time**
GAMbasecase (N/1)	256	4096	82.1 h	336281 h

Table 3.6 ELMFIRE - exemplar simulation runs on SuperMUC

Simulation	Description of run			
	Nodes	**Total cores**	**Wall time**	**Total core time**
GAMNscan (N/4)	32	512	51.2 h	26214 h
GAMNscan (N/16)	32	512	17.8 h	9113 h
GAMNscan (N/64)	16	256	9.1 h	2329 h

Table 3.7 ELMFIRE - exemplar simulation runs on Sisu

Copyright © CRESTA Consortium Partners 2014

3.3 ELMFIRE: Scientific impact

ITER is a large-scale scientific experiment that aims to demonstrate that it is possible to produce commercial energy from fusion. The construction work on ITER is expected to come to an end in 2019 and the total price is expected to be in excess of 15 billion euros. Although the scaling laws from experimental databases of present smaller tokamaks can give some clue of the performance, the numerical simulations are the only possible way to investigate the behaviour of full scale reactors.

Understanding plasma turbulence is very important for the success of ITER concerning especially the physics of plasma turbulence which is a key issue for predicting plasma confinement and performance for future fusion devices. Of particular importance in the study of plasma confinement is the transition between low and high confinement states, still not explained by theory although experimentally observed in 1982. Here, the most advanced approach is the full distribution gyrokinetic turbulence simulation approach which is not yet computationally possible for ITER and, in present work, the work on middle-sized tokamak Textor is reported.

A series of gyrokinetic first principles simulations of plasma turbulence for TEXTOR tokamak was carried out using PRACE resources. The main effort was a parameter scan where temperature, density, scale lengths, magnetic field and isotope was varied starting from Textor L-mode case with strong GAMs. From this data, we have shown radial propagation velocity of GAMs with similar decrease as a function of temperature as in recent experiments. A clear correlation between Er and transport coefficients was found and ExB shear due to phase shift of GAMs were found to reach values which may affect the transport levels. The radial wavelengths of GAMs show good agreement analytic theory where radial mode is Airy function with certain characteristic scale. Part of the PRACE resources were also used to finish the GAM work for FT-2 which is published and also included in doctoral thesis of S.Leerink [4]. For videos of results, thesis and for list of scientific references over last 7 years see ELMFIRE web site [5].

3.4 ELMFIRE: References

[1] T.P. Kiviniemi, S. Leerink, J.A. Heikkinen, S.J. Janhunen, T. Korpilo, “Gyrokinetic Simulations of the Edge Pedestal in the TEXTOR Tokamak”, Contributions to Plasma Physics, **52,** 406 (2012).

[2] E Z Gusakov, A B Altukhov, V V Bulanin, A D Gurchenko, J A Heikkinen, S J Janhunen, S Leerink, L A Esipov, M Yu Kantor, T P Kiviniemi, T Korpilo, D V Kouprienko, S I Lashkul, A V Petrov and N V Teplova, "Anomalous transport and multi-scale drift turbulence dynamics in tokamak ohmic discharge as measured by high resolution diagnostics and modeled by full-f gyrokinetic code", Plasma Phys. Control. Fusion **55**, 124034 (2013).

[3] T. Korpilo, J.A. Heikkinen, S.J. Janhunen, T.P. Kiviniemi, S. Leerink, F. Ogando, "Numerically stable method for kinetic electrons in gyrokinetic particle-in-cell simulation of toroidal plasmas", Journal of Computational Physics, **239**, 22 (2013).

[4] S. Leerink, V. V. Bulanin, A. D. Gurchenko, E. Z. Gusakov, J. A. Heikkinen, S. J. Janhunen, S. I. Lashkul, A. B. Altukhov, L. A. Esipov, M. Yu. Kantor, T. P. Kiviniemi, T. Korpilo, D. V. Kuprienko, and A. V. Petrov, "Multiscale Investigations of Drift-Wave Turbulence and Plasma Flows: Measurements and Total-Distribution-Function Gyrokinetic Simulations" Phys. Rev. Lett. **109**, 165001 (2012).

[5] ELMFIRE web site. http://www.elmfire.eu. (2014). Last visited 2014-03-20.

Copyright © CRESTA Consortium Partners 2014

4 GROMACS

Molecular simulation in general, and biomolecular simulation in particular, has extensive potential to use very large amounts of CPU-cycles since the problem is dominated by sampling. GROMACS is focused on molecular dynamics (as exemplified by the 2013 chemistry Nobel prize), i.e. the evaluation of forces on all atoms in a system using classical interaction models, which makes it possible to take a very small time step (femtoseconds) and compute how atoms move.

By repeating billions or trillions of these steps, it is possible to predict both dynamic properties of molecules (i.e., how they move) as well as the statistical distribution of states visited by very complex molecules. The latter enables us to calculate properties such as free energies that correspond directly to solvation or binding energies for small molecules, which is critical e.g. in the development of medical drugs.

Historically it has been a huge challenge to achieve sufficient scaling for molecular dynamics, in particular for systems of sizes relevant to molecular applications, rather than artificially inflated ones. The work on GROMACS in CRESTA has focused both on improving scaling and performance of individual simulations, and to enable automatic ensemble simulation using the related Copernicus toolkit developed in the same lab.

The needs analysis in deliverable 6.2 identified several important targets for GROMACS software development in CRESTA. Some of those targets have been incorporated into the code that executed the simulations reported here. Specifically, these simulations used the portable future-proof SIMD-friendly non-bonded kernels, and the new ensemble-level simulation-management software called Copernicus. Attempts to exploit lower-level communication modes have not yet succeeded; in particular, an attempt to use shmem proved unsuccessful because GROMACS requires MPMD parallelism at scale. Re-engineering for better task parallelism in GROMACS is ongoing, and will be coupled with the non-blocking collectives from WP4 for use with FFTs in GROMACS for the final simulations. Improvements for efficient parallel I/O are also in progress.

Application	Identification (e.g. Version No)
GROMACS	GROMACS, version 4.6, patch level 5 (4.6.5)
GROMACS/Copernicus	Copernicus, version 1.0 (and 2.0 beta)

Table 4.1 GROMACS - version information

4.1 GROMACS: Scientific use case

4.1.1 Scientific use case: Droplet wetting & membrane protein free energy calculations

For many time-dependent systems, weak scaling is not very relevant since time scales for transitions grow exponentially with system size. One notable exception to this rule is flow, where the timescales usually scale from linearly to quadratically with the system size, for instance the poorly understood concept of liquid wetting at atomic level. Existing models assume continuum behavior, which inevitably means the liquid/surface interactions are described poorly, and the models make poor predictions at small scales where experimental observations are also very difficult. Simulations can help fill this need by doing direct observation of wetting phenomena in atomistic detail.

When going from a mixed viscous/inertial regime to a pure inertial regime, we expect a scaling between linear (viscous) and inertial (power 1.5). Thus going from 100 nm radius, which needs about 20 ns simulation time, we expect about 400 ns for a radius

Copyright © CRESTA Consortium Partners 2014

of one micron. On a million Intel Sandy Bridge cores with perfect linear scaling, that would take about 5 days of wall-clock time. Since these wetting simulations do not contain long-range electrostatics interactions, scaling is linear. We demonstrated linear scaling of GROMACS (without long-range electrostatics) to 150000 cores already in 2009. So scaling to a million cores should certainly be possible. But it remains to be seen if the network connecting a million cores can keep up with a millisecond iteration rate and if the dynamic load balancing in GROMACS can balance such an inhomogeneous system over so many cores.

The second major application class described in the introduction concerns free energy calculations. This is a critically important technique for many biomolecular systems where it is possible to calculate the work required to convert a state of the system into a different state, for instance moving a molecule from vacuum to water, or moving a small compound from water into a specific binding site of medical relevance. Membrane proteins constitute a particularly important class of molecules for these studies since they constitute most receptors on our cells, for instance in the central nervous system. When a small molecule (such as a neurotransmitter) is released from the end of one nerve cell it will bind to receptors on the surface on the next. These neurotransmitters are central molecules themselves, but an even more important recent discovery is that it is possible to amplify or reduce the response of nerves by introducing separate compounds, so-called “allosteric modulators” that bind to other parts of the receptor. Some of the most important examples of this are anesthetics and alcohol that bind to the transmembrane domain of ligand-gated ion channels, which is the way these chemicals influence and affect our brains. Surprisingly, very small differences in these molecules can lead to completely different behavior, and it is first with state-of-the-art molecular simulations it has been possible to predict both where molecules will bind and what effect the will have – will a particular compound facilitate or reduce opening of a specific channel?

This type of calculation is typically performed as a whole series of independent simulations where the compound is gradually turned into a ghost particle by slowly removing the interactions with the receptor in ~40 separate simulations, and each of these simulations can in turn be split into 10-100 different runs to improve sampling. This all needs to be repeated for several different molecules and mutants to draw biological conclusions. For this, we use the Copernicus toolkit on top of GROMACS to automate ensemble simulations of a very large number of independent simulations.

4.1.2 Computational requirements: Droplet wetting & membrane protein free energy

A cylindrical droplet of diameter 100 nm wetting a surface can be modeled in around 1.9 million atoms, however such a model is only 2.3 nm thick, which limits the accuracy with which macroscopic fluid flow effects can be studied. To achieve more realistic systems, we have also repeated this slice to make it a factor 2, 4, or 8 thicker. Achieving strong scaling on such simulations are challenging for GROMACS (and other programs), because the distribution of atoms over space is heterogeneous, and the dynamic load-balance algorithm is tailored for much lower heterogeneity. For these types of systems, GROMACS exhibits virtually perfect weak scaling, so the scaling work in CRESTA has focused on improving the *strong* scaling, i.e. for a system of given size we want to scale down to as few atoms as possible per node. Without substantial re-engineering, for these kinds of simulations, a single CPU core can efficiently handle down to roughly a thousand atoms, in contrast to simulations of condensed phases where the efficiency of the implementation is higher and so the calculation scales to use more hardware, so that the scaling is efficient down to where a single CPU core handles only around 150 atoms. However, even with this limitation, running on an exascale-era machine with one million cores is still only a billion atoms. This would provide 10 times larger radius and a 5 times larger thickness. This will increase the Reynolds number by a factor of 10, which should probably just get us from a mixed viscous-inertial regime to an inertial dominated regime. This might just be enough to reach the continuum limit, beyond which larger atomistic simulations will be

Copyright © CRESTA Consortium Partners 2014

less useful. Reaching this limit should provide enough information to study and understand the physics at the atomistic scale, which will provide a better understanding of the fundamentals of wetting and provide better boundary conditions for continuum models that can be used on larger scales. Figure 4.1 shows an illustration of a small part of the droplet slice.

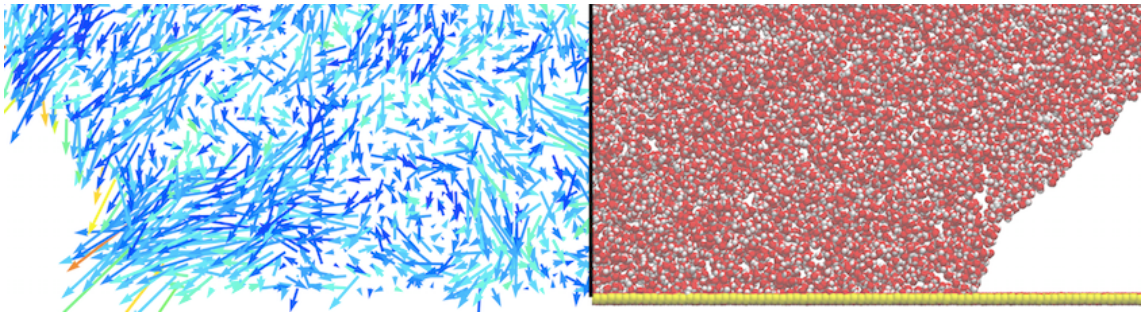

Figure 4.1 A small part of a thin slice of a water droplet used to simulate surface wetting effects. Left side shows the flow representation while the right side shows the atomic detail, both close to the surface (bottom). Already at 2.3nm thickness this system re requires close to 2 million atoms, and increasing the thickness further will enable us to bridge atomic-scale interaction models with continuum mechanics.

For the membrane protein free energy calculation project, the system itself is of much more modest size: 10-20nm in all directions, and not more than 143,000 atoms. This includes a complete membrane protein (GLIC, with both a transmembrane part and an extracellular domain) in a DOPC bilayer surrounded by water, as illustrated in Figure 4.2. This is an actual critical application area where it historically has been extremely challenging to use large numbers of cores. The GROMACS versions developed in CRESTA now scale all the way to roughly 100-200 atoms per core for these systems, but even this limits single individual simulations to using 1000-1500 cores for the small systems. However, since the 20-40 free energy steps can be evaluated in parallel with ensemble-level communication in Copernicus, this directly makes it possible to use ~50,000 cores at very high efficiency. When necessary, each such step can be split into ~10 separate runs for ensemble sampling (which increases the maximum number of cores another order of magnitude), and since biological applications typically involve more than one compound or mutations in the protein, this is a highly viable option for life science simulations to use many millions of cores efficiently in the near future. One particularly important aspect is that individual simulations can be quite short, frequently completing in mere hours given enough resources. For this example, the ensemble level statistics was used to improve accuracy, but it can also be configured to minimize the total number of CPU hours or minimize the wall-clock time.

Copyright © CRESTA Consortium Partners 2014

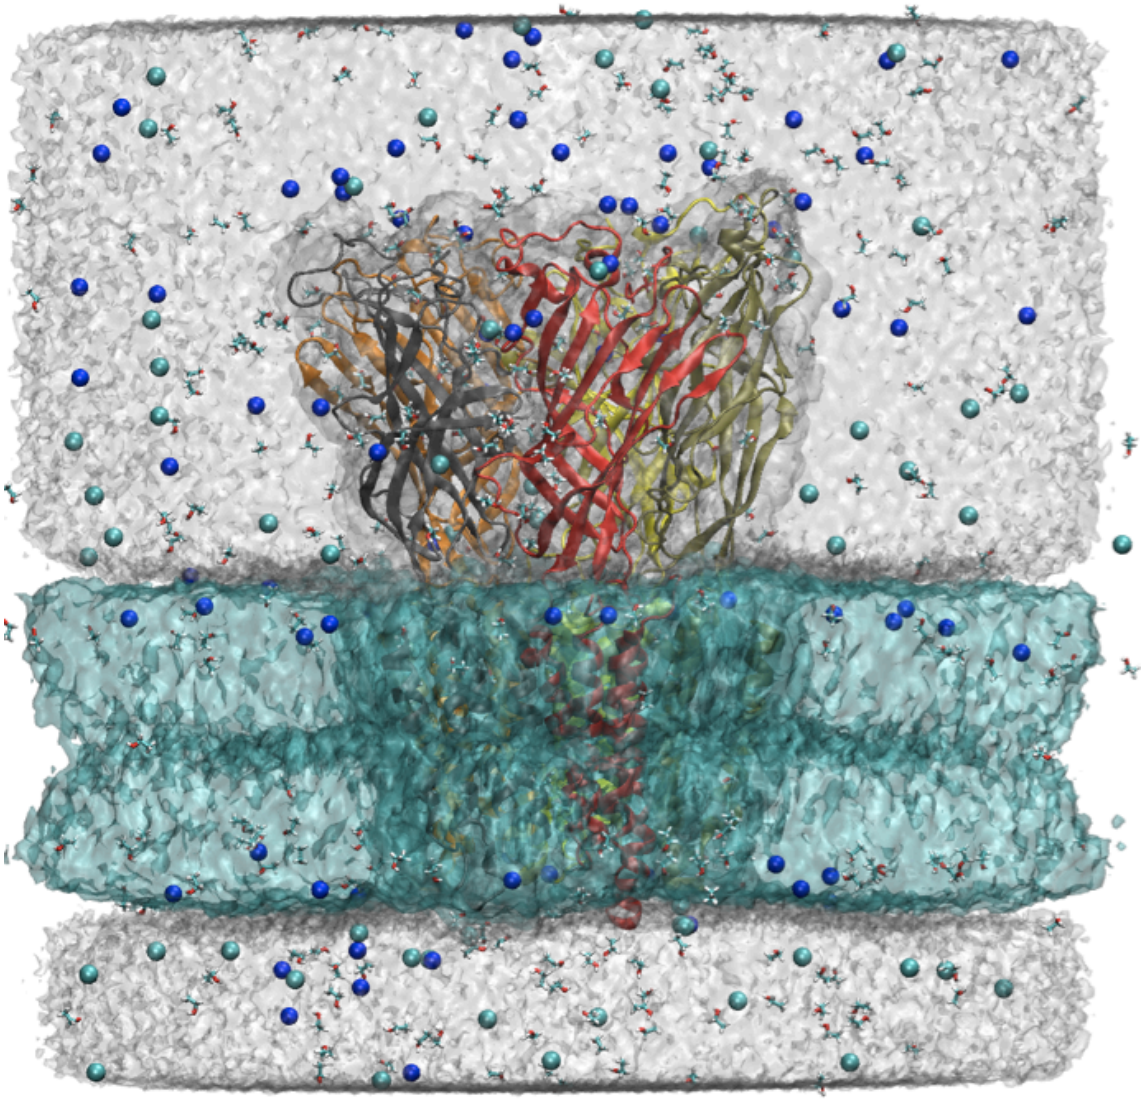

Figure 4.2 The prokaryotic ligand-gated ion channel GLIC is a central model system for understanding the effects of anesthetics and alcohol on the nervous system. Free energy calculations of the binding energies for allosteric ligands (e.g. anesthetics) are enabli enabling us to understand the action sites, and will help design new pairs of anesthetics to control sedation much better in the operating room.

4.2 GROMACS: exemplar simulation

. Simulation	Description
Droplet wetting	GROMACS-4.6, From 1.9 million to 15.2 million atoms in total. A slice of a (large) water cylinder hitting a surface. Reaction-field electrostatics and shifted Van der Waals interactions, 1.0nm cutoffs. 4 fs time steps.
Free energy calc. of anesthetics binding to ligand-gated ion channels	GROMACS-4.6, 143,000 atoms, membrane protein in DOPC bilayer surrounded by water, atomic detail, 2fs time step, PME electrostatics evaluated every step (grid spacing 0.12nm) and van der Waals interactions switched off from 0.8nm. Cutoff at 1.0nm. The ligands (Desflurane and Chloroform) are decoupled from the system with 21 simulations each used for decoupling electrostatics and Van der Waals interactions. Soft-core interactions used for free energy decoupling.

Table 4.2 GROMACS - exemplar simulations

Copyright © CRESTA Consortium Partners 2014

4.2.1 Timescales

Start date:	01/07/2012
End date:	28/02/2014
Total duration:	Approximately 40 million core-hours in total for both systems including development and setup.

Table 4.3 GROMACS - timescales

4.2.2 People

Name	**Role**
Mark Abraham	Running of benchmarks, data assembly, GROMACS scaling
Berk Hess	Water droplet scientific lead
Petter Johansson	System preparation, scientific analysis of droplets
Erik Lindahl	Membrane protein scientific lead
Magnus Lundborg	Ensemble simulation execution & development
Iman Pouya	Ensemble simulation analysis

Table 4.4 GROMACS - people

4.2.3 Environment

Platform	**Platform specs**
Cray XE6, Stockholm	2.1GHz AMD Magny-cours, 12 cores/CPU, Cray Gemini 3D torus (Used for first project part, and to achieve a large number of CPU hours). Free energy calculations.
Cray XC30, Edinburgh	2.7GHz Intel E5-2697v2, 12 cores/CPU, Cray Aries Dragonfly topology interconnect. Droplet wetting.

Table 4.5 GROMACS - environment

4.2.4 Exemplar simulation runs

Simulation	**Description of run**			
	Nodes	**CPUs**	**Wall time (h)**	**CPU time**
Droplet wetting 1x thick (10ns simulated time)	512	12288	1.57	19300
Droplet wetting 2x thick (10ns simulated time)	1024	24576	1.93	47000
Droplet wetting 4x thick (10ns simulated time)	1024	24576	2.78	68300
Droplet wetting 4x thick (10ns simulated time)	1024	24576	4.43	109000
Droplet wetting 8x thick (10ns simulated time)	2048	49152	3.93	193000

Copyright © CRESTA Consortium Partners 2014

Free energy calculation for 1 out of 40 steps, for one compound in one site, for one mutant	56	1344 672 MPI ranks, 2 OpenMP threads per rank	9.348h per point	12563 per point
Automatic ensemble simulation of free energy calculation (all steps) for one compound in one site.	Average 500 (Can use up to 2352)	Average 12000, will vary during project depending on queue load.	Completed in 47h. With enough nodes, wall-clock <10h is possible.	527646

Table 4.6 : GROMACS - exemplar simulation runs on XC30 (Droplet wetting) (see also Figure 4.3) and XE6 (Free energy).

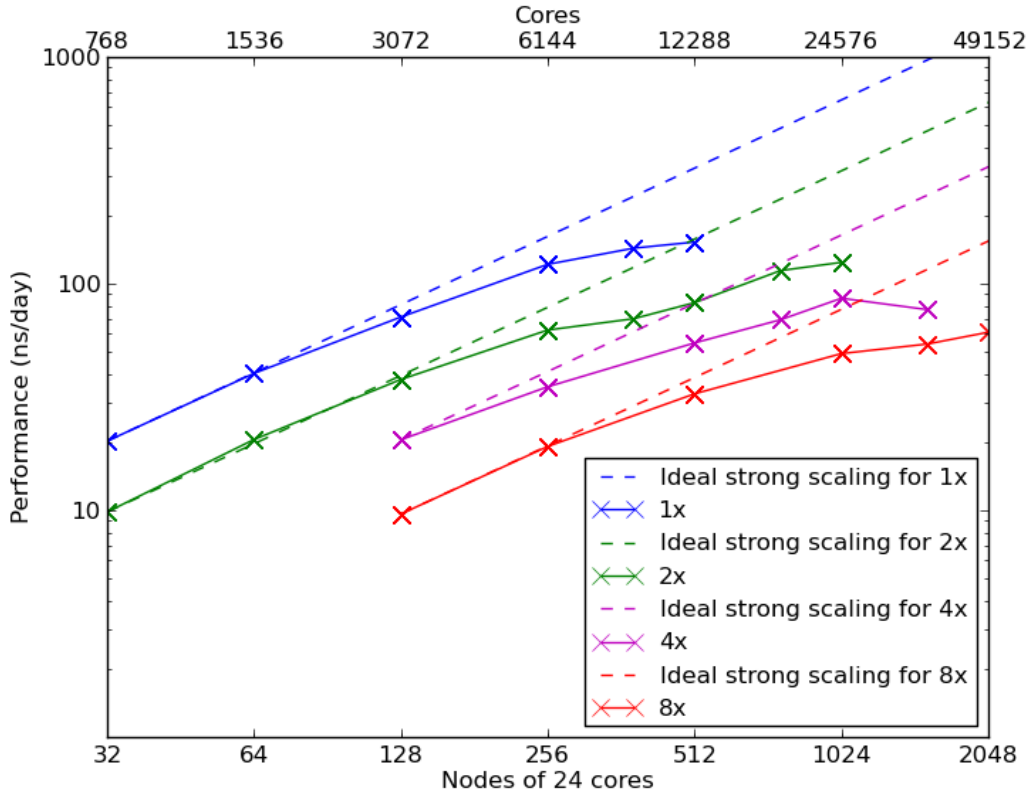

Figure 4.3 Parallel scaling of the thin water droplet slice to simulate wetting. Performance reported as ns/day of production simulation with 4fs timesteps when running on a Cray XC30.

4.3 GROMACS: Scientific impact

The increase of computing power has made it possible to include more physical effects and study more realistic fluids in larger configurations than before. Currently, we can study realistic behavior of water droplets of sub-micrometer diameter. Exploring a wide range of length scales offers a direct view of how the important molecular effects around the interfaces couple to macroscopic phenomena. Large systems also greatly increase the availability of statistical data, reducing thermal noise to the point in which detailed information about fluid flow during dynamic processes can be obtained. Additionally, the nature of wetting changes fundamentally when going from the viscous dominated to the inertia-dominated regime. The transition is linked to the Reynolds number, which, given a liquid and substrate, only depends on the droplet radius. This implies that large droplets are needed to cover all the regimes. The work that enabled

Copyright © CRESTA Consortium Partners 2014

this scaling has been reported in [7][8], and the scientific results are currently being written up for publication in [10].

The computational work on ligand-gated ion channel has had very large application impact, and the methods are now being used by several groups in combination with experimental electrophysiology work – since it is simply not possible to understand the systems and predict which modulatory site will have highest affinity without simulations and free energy calculations. The whole concept of dual allosteric modulation was first reported in [6] as a result of simulations on this very system, and with the performance improvements reported in [7][8] as well as previous CRESTA reports, reference [10] was the first-ever work that predicted altered allosteric modulation purely by computational methods, and later confirming them experimentally. This project needed to investigate a total of 4 different binding sites with 2-3 subposes of the ligands in each, for two separate mutants of the protein. This work proved that single-site mutations can salter both desflurane and chloroform (common anesthetics) from inhibitors to potentiators, which is an extremely important fundamental science step towards designing pairs of anesthetics for very fine-grained control of sedation during surgery.

4.4 GROMACS: References

[6] Murail S, Howard RJ, Broemstrup T, Bertaccini EJ, Harris RA, Trudell JR, Lindahl E, *Molecular mechanism for the dual alcohol modulation of Cys-loop receptors*, PLoS computational biology **8**(10), e1002710 (2012)

[7] Pronk S, Pall S, Schulz R, Larsson P, Bjelkmar P, Apostolov R, Shirts MR, Smith JC, Kasson PM, van der Spoel D, Hess B, Lindahl E, *GROMACS 4.5: a high-throughput and highly parallel open source molecular simulation toolkit.* Bioinformatics 29(7), 845-54 (2013)

[8] Pall S, Abraham MJ, Kutzner C, Hess B, Lindahl E, *Tackling exascale software challenges in molecular dynamics simulations with GROMACS*, Exascale Applications and Software conference EASC14, *in press* (2014)

[9] Broemstrup T, Howard RJ, Trudell JR, Harris RA, Lindahl E, *Inhibition versus potentiation of ligand-gated ion channels can be altered by a single mutation that moves ligands between intra- and intersubunit sites,* Structure 21, 1307-16 (2013)

[10] Johansson P, Hess B, *Importance of atomistic topology on dynamic wetting of water,* manuscript (2014)

Copyright © CRESTA Consortium Partners 2014

5 HemeLB

HemeLB is a tool for fluid flows in complex sparse geometries. Its main focus is simulating blood flow in parts of the cerebral arterial network. HemeLB employs an implementation of the lattice Boltzmann (LB) algorithm which, due to its locality, is intrinsically easy to parallelize

Application	Identification (e.g. Version No)
HemeLB	HemeLB, Branch default, Changeset 2883:1717125ee61d

Table 7: HemeLB - version information

5.1 HemeLB: scientific use case

5.1.1 Scientific use case: The Circle of Willis

The Circle of Willis (CoW) is a ring-like structure of blood vessels at the base of the brain. It was first described by the British physician Thomas Willis in the 17th century [11][12].

Connecting the internal carotid to the cerebral arteries, the CoW is the main distribution system of oxygen-rich blood to the brain. Its circular structure introduces redundancy into the cerebral circulation and allows supply to the whole of the brain even if the blood flow on one side is limited or obstructed [13][14][15].

Many studies have however found a large variation in the actual form of the CoW in individuals, ranging from differences in the ratio of diameters of vessels to missing parts of the circle, which result in large variations in the flow patterns [15][16].

Variations and defects of the CoW and adjacent vessels have been related to different problems ranging from Alzheimer's disease to aneurysm formation and infarcts. Furthermore anatomical variations here can affect the efficiency of treatment approaches as well [17][18][19][20][21][26].

Motivated by this, several numerical studies have implemented different models of flow in the CoW [22][23][24][25][26]. Validations here have typically been made against flow measurements by transcranial Doppler ultrasound or phase-contrast contrast magnetic resonance angiography [21][27][28].

It has been shown that 1d Poiseuille models do not intrinsically capture individual features of the flow and have to be calibrated by 3d models [24]. As the dynamics of artery diameters in the cerebrovascular system are negligible and we consider vessels of sufficiently large diameter to assume Newtonian fluid behaviour of the blood our 3d lattice Boltzmann implementation is well suited for simulation in this domain [29].

Figure 5.1 shows a ray-traced representation of the simulation domain employed in our exemplar simulation runs. In addition to the above mentioned vessels, the geometry contains the basilar artery and the medium cerebral arteries not normally considered to be part of the Circle of Willis as well. The front of the brain is here oriented in positive y-direction. The data was kindly provided by J. Figueroa et al.

Copyright © CRESTA Consortium Partners 2014

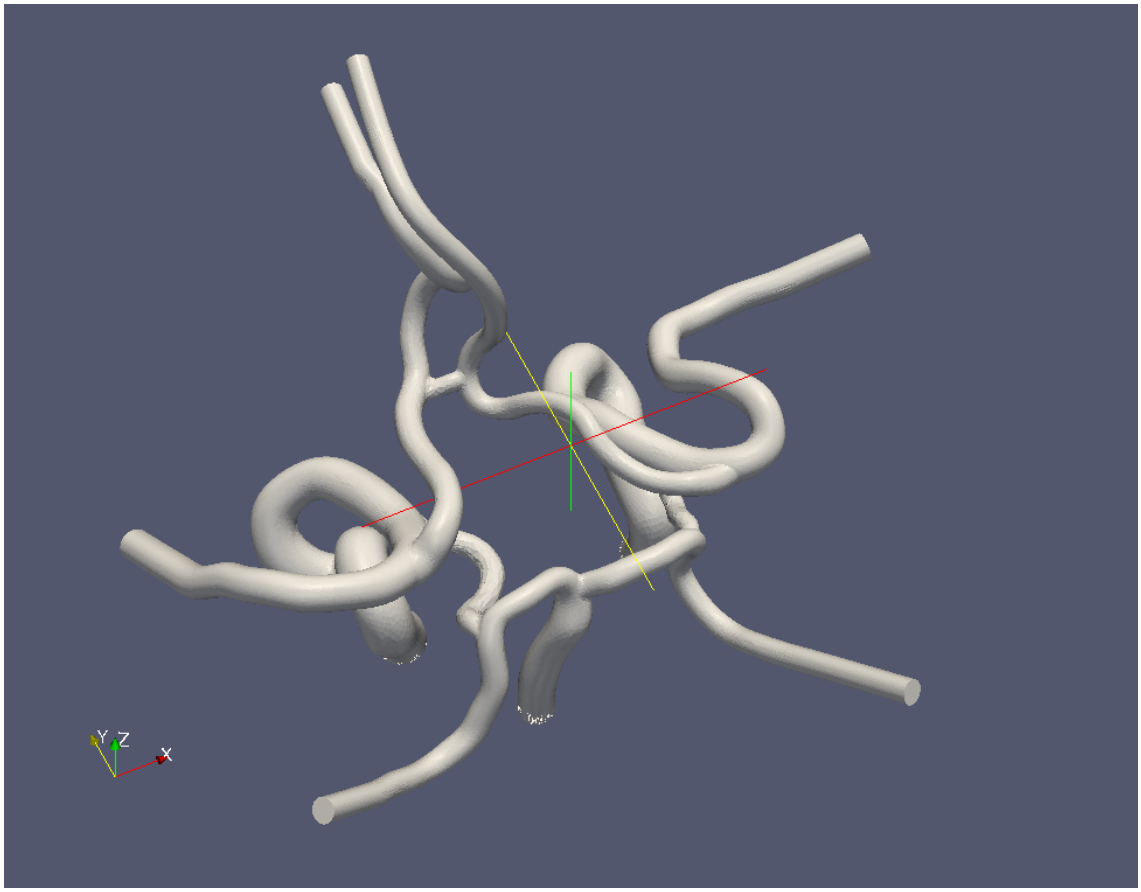

Figure 5.1 Rendering of the simulation domain. Front of the brain in positive y-direction. Depicted is the Circle of Willis and in addition the Basilar artery and medium cerebral artery. Data kindly provided by J. Figueroa et al.

5.1.2 Computational requirements

To accurately resolve the flow properties of blood in a section of the Middle Cerebral Artery (MCA), we require discretised geometries using up to 10 million voxels (or lattice sites). We have run these geometries frequently, and they are ideally run using 2048-4096 cores.

In this application, however, we seek to model a considerably larger system, namely the full Circle of Willis. Based on our estimates for the MCA, we believe that this system will need to contain up to 160 million lattice sites (we demonstrate a geometry with approximately 70 million lattice sites as a test). These systems are well-suited for execution on a petascale architecture, as they require around 64k Intel Ivy Bridge cores to be executed efficiently.

Individual simulations can only be scaled up further to a limited extent (most likely by placing the outlets further down the arterial tree), as we expect a full model of the Circle of Willis to already provide best effort flow estimates without incorporating additional models (e.g., for vascular elasticity, arterial wall deterioration, or for cerebral auto-regulation).

However, it is instructive to analyse the flow properties under a range of flow conditions, and possibly geometry distortions. To do so effectively we would require a set of concurrent simulations for each patient-specific geometry, each initialized with a different velocity profile (e.g., ones associated with blood pressures ranging from those obtained at rest to those obtained during vigorous exercise). Using such an approach, which will require at least multi-petascale resources, would provide clinicians with adapted flow predictions. These predictions can then account for patient activities that are difficult to reproduce in clinical sessions.

Copyright © CRESTA Consortium Partners 2014

5.2 HemeLB: exemplar simulation

Simulation	Description
Circle of Willis	HemeLB, 73M lattice points resolution, Exemplary large run over 2 pulses (24,576 cores, 800,000 time steps), Scaling tests with varying core counts over 10,000 time steps.

Table 8: HemeLB - exemplar simulations

5.2.1 Timescales

Start date:	28/02/2014
End date:	03/03/2014
Total duration:	15700 CPUh

Table 9: HemeLB - timescales

5.2.2 People

Name	Role
Sebastian Schmieschek	Developer
Derek Groen	Developer

Table 10: HemeLB - people

5.2.3 Environment

Platform	Platform specs
ARCHER	ARCHER - Cray XC30, Intel Xeon E5 v2 12C 2.700GHz, Aries interconnect, 76192 cores.

Table 11: HemeLB - environment

5.2.4 Exemplar simulation runs

Simulation	Description of run			
	Nodes	**Total cores**	**Wall time**	**Total core time**
Circle of Willis Scaling test, 10k time steps, no I/O	32	768	235.0	180480
Circle of Willis Scaling test, 10k time steps, no I/O	64	1536	132.0	202752
Circle of Willis Scaling test,10k time steps, no I/O	128	3072	68.7	211046.4
Circle of Willis Scaling test,10k time steps, no I/O	256	6144	37.1	227942.4

Copyright © CRESTA Consortium Partners 2014

Circle of Willis Scaling test,10k time steps, no I/O	512	12288	23.2	285081.6
Circle of Willis Scaling test,10k time steps, no I/O	1024	24576	25.3	621772.8
Circle of Willis Production test, 800k time steps, I/O every 10k time steps	1024	24576	2270.0	55787520

Table 12: HemeLB - exemplar simulation runs on ARCHER

5.3 HemeLB: Scientific impact

Cerebrovascular disorders, including brain aneurysms, are both common and devastating in the UK. They cause close to 10% of all deaths [34] and place a great financial burden on society. It has been suggested that haemodynamics interact with vessel structure in complex manners leading to aneurysm development and rupture. Understanding the characteristics of brain blood flow can help neurosurgeons to better manage these disorders in individual patients, and to better assess the need for interventional cerebrovascular procedures.

In the Centre for Computational Science, University College London we have developed a simulation environment, named HemeLB [26][37], which allows us to model blood flow in a variety of three-dimensional geometries [36]. We have already applied HemeLB to a variety of problems, including models of the middle cerebral artery [30][32] and of retinal blood flow [31].

Within this study we use a three-dimensional, patient-specific model of the Circle of Willis and the directly connected arteries. This model has been kindly provided to us by Figueroa et al. [26]. When we apply this model to HemeLB, we are able to make predictions of the flow properties within key arteries in the cerebrovasculature. These include for example the left and right middle cerebral arteries, locations where brain aneurysms are commonly observed.

With HemeLB, we are able to predict the velocities in the branching arteries, as well as the wall-shear stress, the latter of which is very difficult to obtain in a clinical environment. In addition, we are able to vary the inflow properties, something which is often undesirable to attempt in-vivo, and make rough predictions of the flow properties in the cerebrovasculature in cases where the blood pressure is elevated (for example, during exercise) or reduced.

5.4 HemeLB: References

[11] O'Connor, J., Thomas Willis and the background to Cerebri Anatome, J R Soc Med 96, 139 (2003)

[12] [1] Willis, T. Cerebri Anatome: cui accessit nervorum descriptio et usus, London (1664)

[13] The Lancet, The Circle of Willis 278 (7201), 536 (1961)

[14] Ipers, B.J. et al., Anatomical Studies of the Circle of Willis in Normal Brain, AMA Arch NeurPsych. 81 (4), 409 (1959)

[15] Riggs, H.E. et al., Variation in Form of Circle of Willis, Arch Neurol., 8 (1), 8 (1963)

[16] Bergman, R.A et al. Circle of Willis in Illustrated Encyclopedia of Human Anatomic Variation, http://www.anatomyatlases.org/AnatomicVariants/Cardiovascular/Test/Arteries/CircleofWillis.shtml, last visited 2014-03-02

Copyright © CRESTA Consortium Partners 2014

[17] Roher, A.E. et al., Circle of Willis Atherosclerosis is a Risk Factor for Sporadic Alzheimer's Disease, Art Thromb Vasc Bio. 23, 2055 (2003)
[18] Hashimoto, N. et al., Experimentally induced cerebral aneurysms in rats: Part V. Relation of hemodynamics in the circle of Willis to formation in aneurysms, Surgical Neurology 13 (1), 41 (1980)
[19] Lord, R.S.A. et al., Contribution of the circle of Willis to the Subclavian Steal Syndrome, Circulation 40, 871 (1969)
[20] Ringelstein, B.E. et al., Cerebral vasomotor reactivity is significantly reduced in low-flow as compared to thromboembolic infarctions: The key role of the circle of Willis, J Neur Sc. 121(1), 103 (1994)
[21] Klingelhoefer, J. et al., Transcranial Doppler ultrasonography of carotid-basilar collateral circulation in subclavian steal, Stroke 19, 1036 (1988).
[22] Alastruey, J. et al., Modelling the circle of Willis to assess the effects of anatomical variations and occlusions on cerebral flow, J BioMech. 40 (8), 1794 (2006)
[23] Oshima, M. et al., Modelling of inflow boundary conditions for image-based simulation of cerebrovascular flow, Num Meth Fluid. 47(6), 603 (2005)
[24] Moore, S.M. et al., One-dimensional and three dimensional models of cerebrovascular flow, J Biomech Eng. 127(3), 440 (2004)
[25] De Vault, K. et al., Blood flow in the Circle of Willis: Modeling and Calibration, Multiscale Model Simul. 7 (2), 888 (2008)
[26] Coogan, J.S. et al., Computational simulations of hemodynamic changes with thoratic, coronary and cerebral arteries following early wall remodeling in response to distal aortic coarctation, Biomech Model Mechanobiol. 12, 79 (2013)
[27] Hendrikse, J. et al., Distribution of cerebral blood flow in the circle of Willis, Radiology 235(1), 031799 (2005)
[28] Bishop, C.C. et al., Transcranial Doppler measurement of middle cerebral artery blood flow velocity: a validation study, Stroke 17, 913 (1986)
[29] Dempere-Marco L. et al., CFD Analysis Incorporating the Influence of Wall Motion: Application to intracranial Aneurysms, Lect Notes in Comp Sc. 4191, 438 (2006)
[30] Bernabeu et al., “Impact of blood rheology on wall shear stress in a model of the middle cerebral artery”, Interface Focus 3(2), 20120094, 2013.
[31] Bernabeu et al., “Computer simulations reveal complex distribution of haemodynamic forces in a mouse retina model of angiogenesis” arXiv:1311.1640, 2013.
[32] Groen et al., “Flexible composition and execution of high performance, high fidelity multiscale biomedical simulations”, Interface Focus 3(2), 20120087, 2013.
[33] Poster presented at the CLMS Symposium: http://www.slideshare.net/DerekGroen/multiscale-modelling-of-brain-bloodflow, last visited 2014-03-02
[34] http://www.who.int/healthinfo/global_burden_disease/estimates_country/en/index.html, last visited 2014-03-02
[35] http://www.newscientist.com/article/mg21729066.000-virtual-body-double-gets-ill-so-you-dont-have-to.html, last visited 2014-03-02
[36] Nash et al., “Choice of boundary condition for lattice-Boltzmann simulation of moderate-Reynolds-number flow in complex domains”, Physics Review E 89, 023033, 2014.
[37] HemeLB web site, http://ccs.chem.ucl.ac.uk/hemelb, last visited 2014-03-02

Copyright © CRESTA Consortium Partners 2014

6 IFS

ECMWF's Integrated Forecasting System (IFS) is the world leading application for the purpose of global analysis and forecast in the field of Numerical Weather Prediction. ECMWF's operational activities and wide-ranging programmes of research and development continue to play a pioneering role in the remarkable advancement of data assimilation and weather forecasting systems. The comprehensive Earth-system model encompasses the interactive coupling of ocean, ocean-waves, and complex land-surface processes. IFS forms the basis of all the data assimilation and forecasting activities and the main applications required are available through one integrated computer software system (a set of computer programs mainly written in Fortran). The increasing detail and sophistication of the forecast and data assimilation system makes huge demands on the computational resources. The strong increase in computer power experienced in the last 20 years, together with major increases in computational efficiency provided by the continual development of the numerical algorithms used to solve the forecast equations, has enabled ECMWF to greatly increase the resolution and the complexity of its forecasting system. The long-term strategy of ECMWF envisages this trend of increasing resolution and complexity of the forecasting system to continue: to improve the realism of the near surface interactions and to resolve the relevant physical forcings, perhaps most importantly the vertical re-distribution of energy via convection (at resolutions 1km or less), that is currently emulated by means of subgrid-scale parameterization. As the operational IFS model is currently based on a 16 km horizontal resolution and as the cost of the computations scale approximately like the cube of the resolution increase, the requirement leads the IFS into the exascale computing domain, in particular when considering that any resolution increase must always meet the operational (speed) requirement of approximately 240 forecast days per day.

Application	Identification (e.g. Version No)
IFS	IFS, CY38R2, RAPS13 (benchmark version) for scaling runs on TITAN IFS, CY38R2, for SANDY Exemplar runs on IBM Power7

Table 6.1 IFS - version information

6.1 IFS: scientific use case

6.1.1 Scientific use case: Evaluation of medium-range forecasts for hurricane Sandy

On 30 October 2012 hurricane Sandy made landfall on the U.S. East coast with a devastating impact. Here, the performance of the ECMWF forecasts (both high-resolution and ensemble) are evaluated together with ensemble forecasts from other numerical weather prediction centres. The sensitivity to model resolution for the ECMWF forecasts is explored.

6.1.2 Computational requirements

It is well understood that increasing the resolution of weather forecast models is an important factor in improving their skill. However, this by itself is only part of the picture; it must be accompanied by increases in the use of observations and the resolution of the data assimilation system used to provide the initial conditions for the forecast model. The latter is not addressed in CRESTA but it is understood (and worked on independently) that the scalability of the assimilation part is of equal importance. Of course, improvements in model physics, dynamics and the interactions with the oceans all have an important part to play in continually improving forecast skill.

The highest model resolution used as part of the hurricane Sandy simulations was T_L3999 (5 km global grid) with 91 levels. This required 512 nodes (about 2/3 of the available nodes on one of ECMWF's two Power7 clusters). A single 10 day forecast took about 2 hours. This has been repeated for different start dates leading up to

Copyright © CRESTA Consortium Partners 2014

landfall (the forecast integrations being shorter up to landfall), so 512 Power7 nodes were used for about 10 hours for this resolution.

Scaling studies of T_L3999, T_C1999 (both 5 km global grids) and T_L2047, T_C1023 (both 10 km global grids) required about 2-3 million core hours on TITAN. Runs were also performed with and without the use of Fortran coarrays for the scaling studies.

Increasing the existing resolution, for example to a 2.5km global grid, is scientifically extremely interesting, since it moves towards resolving deep convection and thus resolving rather than parameterizing the aforementioned energy re-distribution. The spatial uncertainty associated with these processes also likely necessitates the use of ensembles of simulations at these very high resolutions.

The benefits of improved weather forecasts and in particular of severe weather events can be measured in many ways, in lives saved and 10's of billions of Euros in savings from providing early warnings to government agencies.

6.2 IFS: exemplar simulation

IFS exemplar simulations were based on IFS internal cycle CY38R2. These simulations were conducted in two parts, with scaling runs on TITAN using the RAPS13 benchmark version, and Sandy simulations being run on one of ECMWF's IBM Power7 clusters. The reason to run the Sandy simulations at ECMWF is mainly due the need to use analysis/visualization tools for the large volume of data generated by the simulations (terabytes) and collaboration of a number of research and operations staff in this work. The operational IFS internal cycle at the time of Sandy was CY38R1.

Simulation	Description
Scaling runs on TITAN	IFS 10km global grid and 5km global grid cases have been run at scale using a linear grid and a cubic grid. In addition, these cases were run with and without the use of Fortran 2008 coarray optimizations developed in the CRESTA project. For the 10km case, runs used up to 98,304 cores (6,144 nodes), while the 5km case runs used up to 212,992 cores (13,312 nodes).
Sandy simulations on IBM Power7 at ECMWF	These simulations used linear grid resolutions, T_L159 (160 km), T_L319 (64 km), T_L639 (32 km), T_L1279 (16km) and T_L3999 (5km).

Table 6.2: IFS - exemplar simulations

6.2.1 Timescales (scaling runs on TITAN)

Start date:	11/2013
End date:	02/2014
Total duration:	3 million core hours on TITAN.

Table 6.3: IFS - timescales

6.2.2 People

Name	Role
George Mozdzynski	Scaling runs on TITAN
Linus Magnusson et al.	Evaluation of forecasts for hurricane Sandy on ECMWF IBM Power7 cluster [38]

Table 6.4: IFS - people

Copyright © CRESTA Consortium Partners 2014

6.2.3 Environment

Platform	Platform specs
TITAN	At ORNL, Interlagos cores (16 cores per node), NVIDIA K20X GPU (1 per node - not used for IFS scaling runs), CRAY Gemini interconnect, 18,688 nodes, 27.1 PF peak
IBM Power7	At ECMWF, IBM Power7 cores, 32 cores per node, 768 nodes, IBM proprietary interconnect, 754 TF peak per cluster, 2 clusters
CRAY XC-30	At ECMWF, Intel Ivy bridge cores, 24 cores per node, 3500 nodes, CRAY Aries interconnect, 1.74 PF peak per cluster, 2 clusters

Table 6.5: IFS - environment

6.2.4 Exemplar simulation runs (Scaling runs on TITAN)

Figure 6.1 shows the performance of the IFS model at 10 km global resolution on up to 98,304 cores of TITAN. The use of the new cubic grid (denoted by subscript $_C$) ([39]) gives a significant improvement in performance over the linear grid (denoted by subscript $_L$). The cubic grid has the same number of grid points as the linear grid (5,447,118 at each of the 137 atmospheric levels, total grid points 746×10^6) but half the spectral resolution; in this case the cubic grid has a spectral resolution of T_C1023 whereas the linear grid has a spectral resolution of T_L2047. Besides the reduction in the spectral truncation from 2047 to 1023 (and corresponding reduction in spectral transforms and associated transpositions), the cubic grid allows a longer time-step to be used from 450 seconds for the linear grid to 600 seconds for the cubic grid which is a major factor in the performance improvement. The 10 km IFS model is planned to enter operations in 2015, and is required to run in excess of 240 Forecast Days per Day (FD/D), which on TITAN can be achieved with 12K cores and on ECMWF's new XC-30 system with just 6K Ivy bridge cores. It is further encouraging that the use of Fortran2008 coarrays in IFS developed in the CRESTA project [40] continues to deliver a significant improvement in performance at scale.

Copyright © CRESTA Consortium Partners 2014

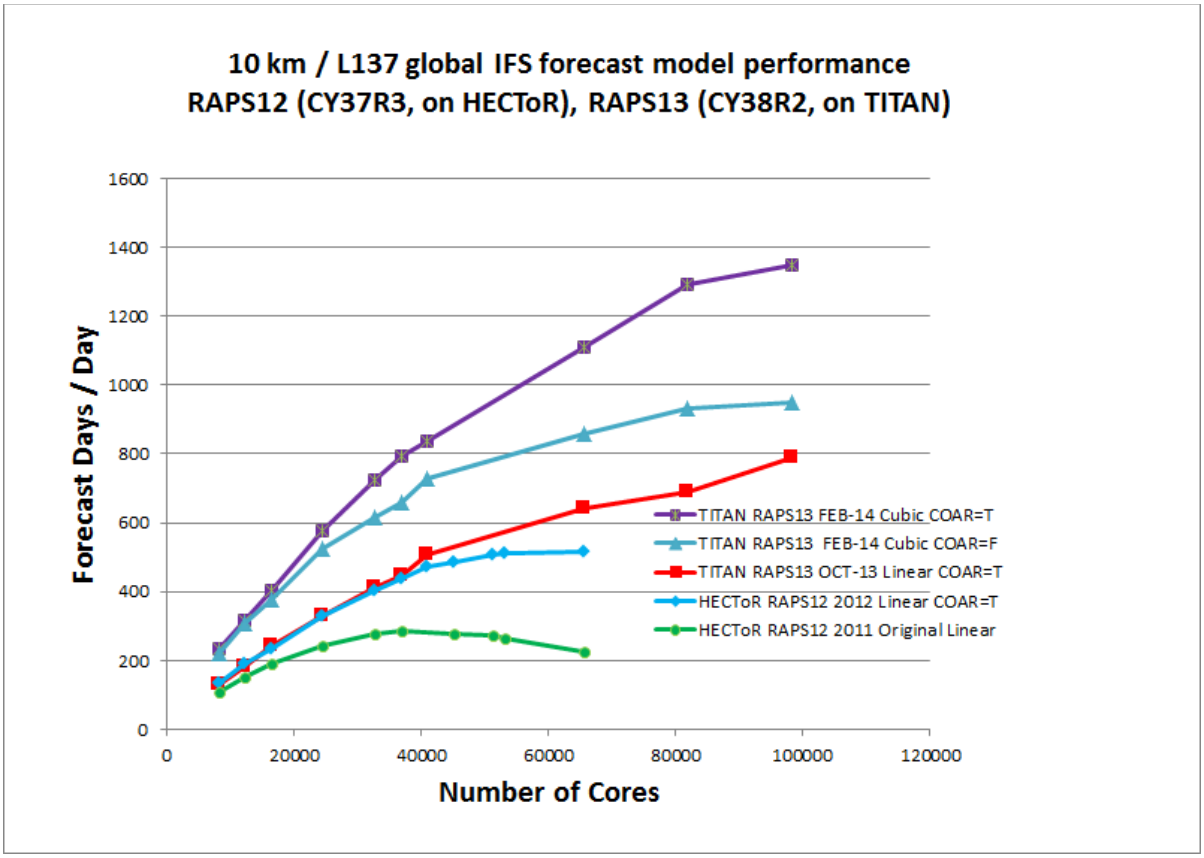

Figure 6.1: IFS 10 km global model performance on TITAN

Figure 6.2 shows the performance of the IFS model at 5 km global resolution on up to 212,992 cores of TITAN. Following the trend of the past 30 years, This model resolution is envisaged for operational implementation in 2023-2024. The runs performed on TITAN show for the first time performance in excess of the 240 FD/D operational requirement at about 100K cores. This is with the new 5km cubic grid with a spectral resolution of T_C1999, whereas the 5 km linear grid has a spectral resolution of T_L3999. The time-step used for the linear grid was 240 seconds and 400 seconds for the cubic grid. Both these grids have 20,696,844 grid points at each of the 137 atmospheric levels (total grid points 2.8×10^9). Again, the use of Fortran2008 coarrays provides a marked improvement in performance. This 5 km cubic grid model case has also been run on ECMWF's XC-30 (without using coarrays), and this achieved 240 FD/D at about 25K Ivybridge cores, one quarter of the cores for a comparable run on TITAN.

Copyright © CRESTA Consortium Partners 2014

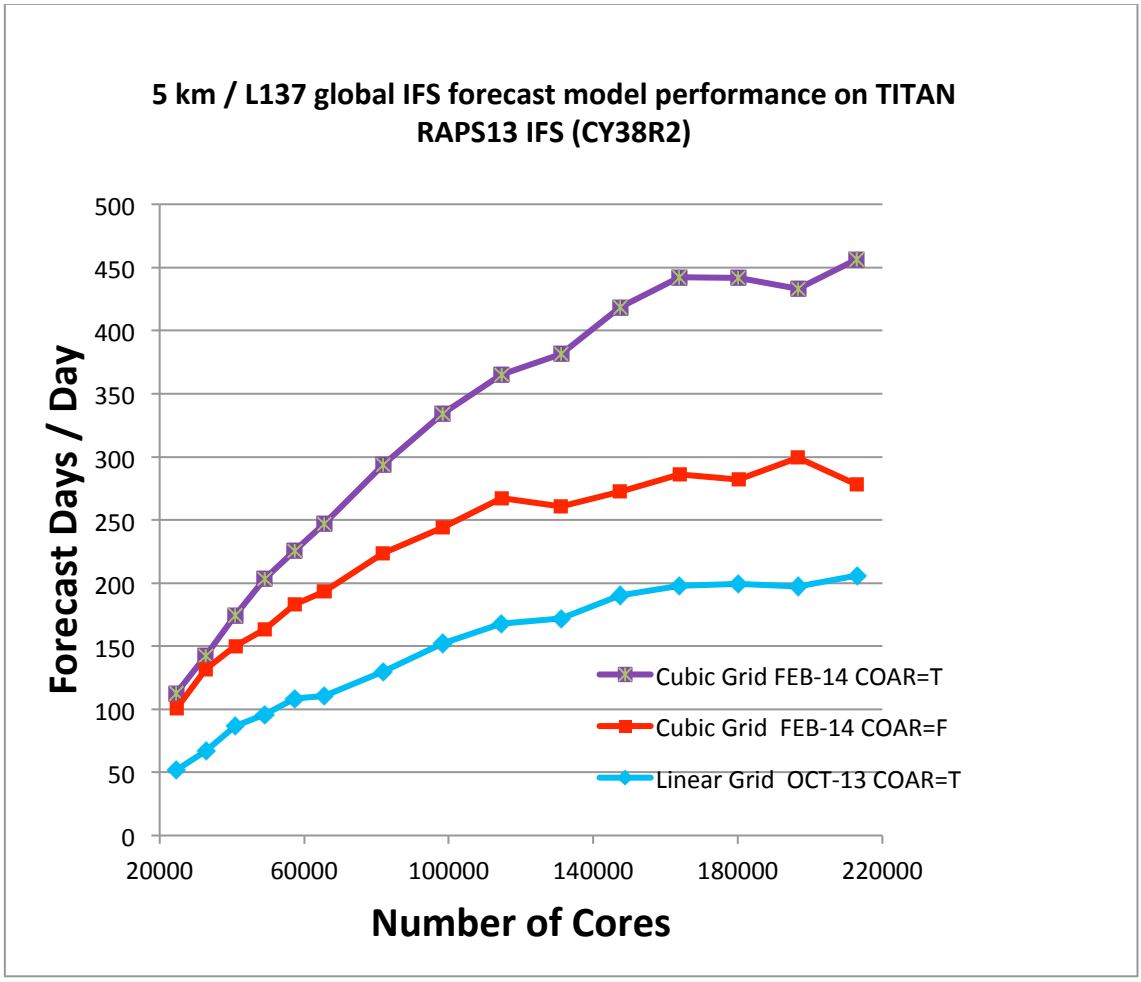

Figure 6.2 : IFS 5 km global model performance on TITAN

6.3 IFS: Scientific impact

A full evaluation of IFS model forecasts for hurricane Sandy is presented in [38] and only some highlights are presented here which are taken from the paper.

The results show that the ECMWF forecasts provided a clear indication of the landfall location from 7 days in advance.

The forecast provided by ECMWF was with the then operational model which was a 16 km global model (T_L1279) with 91 atmospheric levels, which was upgraded in 2013 to 137 atmospheric levels. It is important to note that the forecasts of such severe weather were provided by an operational suite that is run twice a day at ECMWF, 365 days a year.

Figure 6.3 shows the cyclone track and Mean Sea Level Pressure (MSLP) for 9 days prior to Sandy landfall (note the blue line is the forecast, while red line is the analysis - or the best indicator of the actual track that Sandy took). At 9 days ahead, even ECMWF's then operational model showed Sandy veering east out into the Atlantic, but at 8 days ahead (and subsequent days) ECMWF's forecast correctly predicted the path that Sandy took.

Figure 6.4 shows the MSLP from the analysis at 30 October 00 UTC (a) and forecasts from 5 different model resolutions, issued 5 days earlier (25 October 00 UTC). In the plots the cyclone tracks for the analysis (red) and forecasts (blue) are included. The T_L1279 forecast has a position similar to the T_L3999 and the analysis. For the lower resolutions, the cyclone is still east of the coast, indicating a slower movement in these forecasts. For the T_L319 and T_L159 forecasts, the curve of the cyclone track to the west

Copyright © CRESTA Consortium Partners 2014

is not as sharp as in the analysis, while the two highest resolutions turn westward somewhat too early.

In addition to MSLP, the evaluation shows that the T_L3999 forecast (compared to the other resolutions) is the most accurate in relation to hourly observations (MSLP, wind speed and significant wave height, from 27 October 00 UTC) from two buoys located close to New York Harbour and Long Island, the centre of the hurricane passing south of the two buoys. Figure 6.5 shows the significant wave height over this period for the New York Harbour buoy, with the T_L3999 forecast providing a near perfect match to the observations compared to other resolutions.

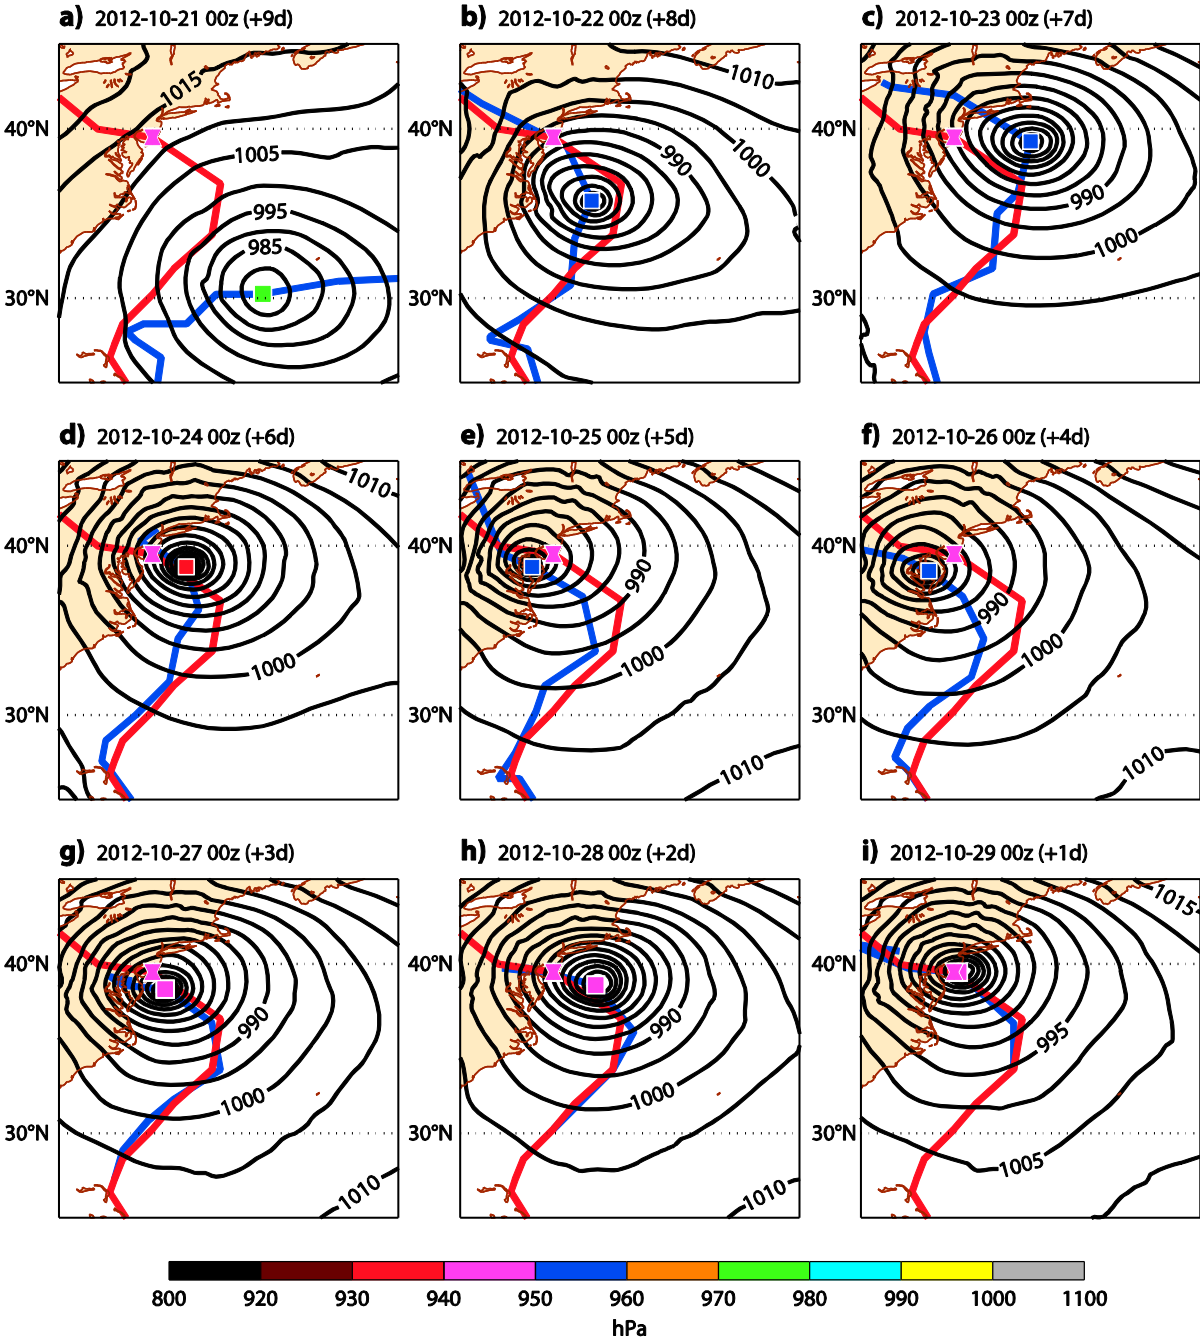

Figure 6.3 : MSLP from IFS $T_L1279L91$ forecasts valid at 30 October 00 UTC and the cyclone track for the forecasts (blue) and the analysis (red). The position of the cyclone centre in the analysis at 30 October 00 UTC is marked as an hourglass and the forecast with a square. The colour of the symbol represents the depth of the cyclone centre.

Copyright © CRESTA Consortium Partners 2014

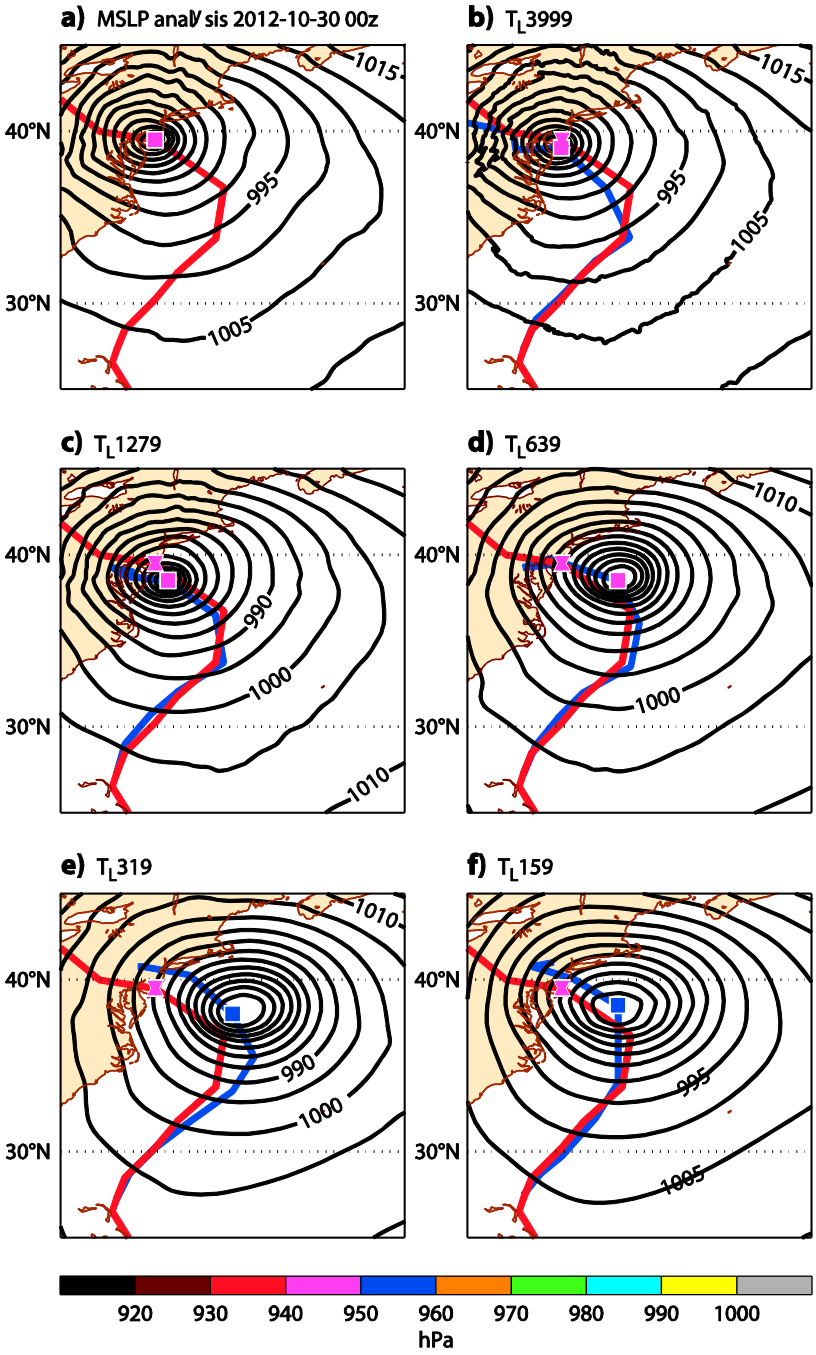

Figure 6.4 : MSLP forecasts from 25 October 00 UTC valid 30 October 00 UTC (5 day forecasts) from different resolutions (starting from the same T_L1279 analysis) including the cyclone track (analysis - red, forecast -blue). Again, the position of the cyclone centre in the analysis at 30 October 00 UTC is marked as an hourglass and the forecast with a square. The colour of the symbol represents the depth of the cyclone centre.

The largest impact in New York City was due to the storm surge connected to Sandy. While a storm surge model is currently not included in the ECMWF forecasting system, our results suggest great potential in including this in the future. We have not investigated the impact of resolution in the data assimilation system and all forecasts were initialised from the same T_L1279 analysis. Our results do not exclude the possibility that the resolution and the quality of the analysis played a major role for the forecast success.

Finally, the evaluation gives a good example for the use of ensemble forecasts for such a severe weather event, comparing results from 4 forecast centres, ECMWF, UKMO, NCEP (USA) and CMC (Canada).

Copyright © CRESTA Consortium Partners 2014

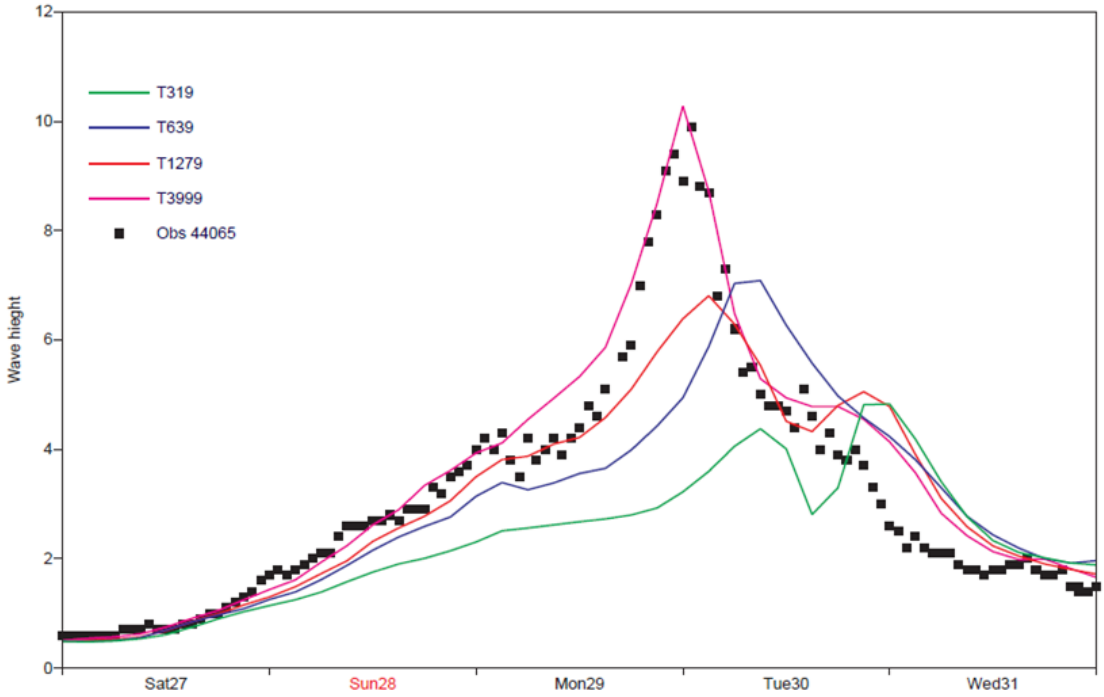

Figure 6.5 : Significant wave height (in metres) for NDBC New York Harbour entrance buoy

Sandy was a severe weather event that led to lives lost and an estimated $68 billion of damage for the USA. The evaluation of hurricane Sandy [38] shows that the ECMWF IFS operational model provided a world leading forecast of this hurricane up to 7 days in advance that likely saved lives. Achieving such a forecast is a combination of many parts, the assimilation of observations from many sources (mainly satellite observations) to produce the best possible initial conditions for the model and then for the forecast model to produce a forecast for up to 10 days or more ahead. Ensemble forecasts importantly provide a quantitative measure of the uncertainty of each individual prediction.

Increasing the resolution of the forecast model further will lead to improved forecasting skill together with a proportional increase in the data assimilation system used to produce the initial conditions. The T_L3999L91 resolution model used in the evaluation should ideally have been initialised from a higher resolution data assimilation system, which will only become available in the period leading up to the operational implementation of such a higher resolution model.

It is important to assess early the computational performance of higher-than-operational resolution models, and this has been successfully accomplished with scaling runs of a 5 km model on TITAN. Using the newly formulated cubic grid [39] new pathways have been explored to selectively increase the resolution for some but not all parts of the global IFS NWP system while exploiting the associated higher efficiency.

6.4 IFS: References

[38] Magnusson, L., J. Bidlot, S. Lang, A. Thorpe, N. Wedi, and M. Yamaguchi, 2014: Evaluation of medium-range forecasts for hurricane Sandy. Mon. Wea. Rev. doi:10.1175/MWR-D-13-00228.1, in press. http://journals.ametsoc.org/doi/pdf/10.1175/MWR-D-13-00228.1

[39] Nils P. Wedi, 2014: Increasing horizontal resolution in NWP and climate simulations - illusion or panacea? Submitted to rspa.royalsocietypublishing.org, accepted.

[40] Roadmap to Exascale (Update 1), D6.1.2.

Copyright © CRESTA Consortium Partners 2014

7 Nek5000

Nek5000 [45] is an open-source code for the simulation of incompressible flow in complex geometries. OpenACC enables existing HPC application codes to run on accelerators with minimal source code changes. This is done with compiler directives and API calls, with the compiler being responsible for generating the GPU optimized code and the user guiding performance only where necessary.

The Nek5000 discretization scheme is based on the spectral-element method. In the approach, the incompressible Navier-Stokes equations are discretized in space by using high-order weighted residual techniques employing tensor-product bases. The tensor-product-based operator evolution can be implemented as matrix-matrix products. This makes it is possible to accelerate the most compute intensive parts of the simulation in a straightforward manner with a GPU.

Application	Identification (e.g. Version No)
Nek5000 with OpenACC	Nekbone_OpenACC (r141), Nek5000_OpenACC (20140211)

Table 7.1 Nek5000 - version information

7.1 Nek5000: scientific use case

7.1.1 Scientific use case: Pipe Flow Simulations

The study of pipe flows is closely related to finding the relation between the average flow velocity to the friction coefficient. The flow of fluid in pipes with circular cross-section is frequently encountered in a variety of environmental, technical and even biological applications. Typical examples of pipe flows can be found in urban drainage system transport of natural gas or oil in the energy sector (i.e. pipelines) or the flow of blood in veins and arteries. Simulations can aid to find methods to solve problems related to pipe flow, such as drag reduction. Thus, the understanding of flow physics in pipes has a direct and rather substantial impact on everyday life.

7.1.2 Computational requirements

In the CFD code Nek5000, Direct numerical simulation (DNS) is employed to numerically solve the Navier-Stokes equations. In order to capture all the features of an eddy in the flow field, the computational domain should be larger than the structures in flow to resolve the smallest scale in the turbulent eddies. Consequently, the computational cost of performance a DNS including all scales grows by Re^3. The highest Reynolds number Re for a pipe flow can reach 5.3x10^5. With such a high computational costs, only a substantial increase in the computing power makes it possible to fully resolve numerical solutions of a truly turbulent flow.

In order to validate the simulation results, simulations and experiments from the KTH wind tunnel have been compared. The conclusion has been that the simulations have replicated the experiments down to the smallest of scales. Thus, conducting realistic experiments in a virtual wind tunnel is feasible.

7.2 Nek5000: exemplar simulation

Simulation	Description
Nekbone mini applications	Run one template test with a set number of elements (16-2048) per CPU/GPU and a set polynomial order (8-20). Furthermore run without the multigrid preconditioner and without user-provided decompositions of the processor counts and the elements.

Copyright © CRESTA Consortium Partners 2014

Pipe Flow Simulation	Laminar pipe flow simulation. P_N-P_N algorithm, 36480 elements, Order of polynomial N=8,10,12,14,16. The total points are around 18.7, 36.5, 63.0, 100.1, and 149.4 M. Typically these cases can be run on 512-8192 CPU/GPUs.
Turbulent Pipe Flow Simulation	Turbulent pipe flow simulation. P_N-P_N-2 algorithm, 1264032 elements, Order of polynomial N=10,12,14,16. The total points are around 1260.0, 2184.2, 3468.5 and 5177.5 M. Typically these cases can be run on 10 240 – 204 800 CPU/GPUs

Table 7.2: Nek5000 - exemplar simulations

7.2.1 Timescales

Start date:	01/07/2012
End date:	01/12/2014
Total duration:	3 000 000 CPU hours

Table 7.3: Nek5000 - timescales

7.2.2 People

Name	**Role**
Jing Gong	Develop and debug the OpenACC version of code. Run all cases on Titan
Philipp Schlatter	Create the input and restarting files. Run pilot tests on the other platforms
Dan Henningson	Create the turbulent modelling
Alistair Hart	Optimize the OpenACC derivatives. Run pilot tests on the other platforms

Table 7.4: Nek5000 - people

7.2.3 Environment

Platform	**Platform specs**
Raven (only for code development and debugging)	Cray XE/XK series system. The system contains 8 XK7 compute nodes with a total of 8 AMD Interlagos Opteron processors and 8 NVIDIA Tesla K20 GPUs. An XK7 compute node has one Opteron processor (2.1GHZ) with a total of 16 processor cores and 2 NUMA nodes with a total of 16 GB of DDR3-1600 main memory. Each GPU has 6 GB of GDDR5 memory.
Titan	Cray XK7 serial system. The system contains 18,688 physical compute nodes, each with a processor, physical memory, and a connection to the Cray custom high-speed interconnect. Each compute node contains 16-core 2.2GHZ AMD Opteron 6274 (Interlagos) processor and 32GB of RAM. Two node shared one Gemini high-speed interconnect router. All of Titan's 18,688 physical compute nodes contain an NVIDIA Kepler GPU with 6 GB of GDDR5 memory.

Table 7.5: Nek5000 - environment

7.2.4 Exemplar simulation runs

Simulation	**Description of run**

Copyright © CRESTA Consortium Partners 2014

	Nodes	Total cores	Wall time	Total core time
Nekbone mini applications	4	4 (GPUs)	25.5 s	102 s
	16	16 (GPUs)	24.8 s	398 s
	64	64 (GPUs)	26.7 s	1706 s
	256	256 (GPUs)	27.0 s	6913 s
	1024	1024 (GPUs)	29.2 s	29912 s
	4096	4096 (GPUs)	31.4 s	128629 s
	8192	8192 (GPUs)	33.0 s	270268 s
	16384	16384 (GPUs)	36.0 s	590539 s
Pipe Flow Simulation	64	1024	2 h	2048 h
	640	10240	2 h	20480 h
	128	128 (GPUs)	1.5 h	192 h
	256	256 (GPUs)	1.5 h	384 h
	512	512 (GPUs)	1.5 h	768 h
	1024	1024 (GPUs)	1.5 h	1536 h
	2048	2048(GPUs)	0.5 h	3072 h
Turbulence Pipe Flow Simulation	512	8192	3 h	24576 h
	1024	16384	1 h	8192 h
	4096	65536	1 h	65536 h
	8192	131072	0.5 h	65536 h
	10240	163840	0.5 h	81920 h

Table 7.6: Nek5000 - exemplar simulation runs on Titan

7.3 Nek5000: Scientific impact

CFD code Nek5000 is being prepared for the exascale era. Our focus has been on solving software challenges of using hybrid computer architectures with accelerators in order to enable simulations at an exascale. To our knowledge, Nek5000 is one of the first examples where a production application has been enabled to use GPUs by using OpenACC.

We have executed the GPU enabling work with several steps. First, we have investigated how to enable running a simplified version of the Nek5000 code on a parallel hybrid architecture [42]. Second, we have investigated the strategies to speed up the full Nek5000 code when running on a similar system. Our work provides an example of how to port real-world application massively parallel hybrid architecture.

We also assess the effectiveness of OpenACC [43] and the compiler support for GPU programming with a view to future hybrid exascale simulations. This is importance in order to avoid the inefficiencies of rewriting large portions of a code in a low-level language, thus ending up with multiple versions of the source code to maintain.

Copyright © CRESTA Consortium Partners 2014

OpenACC enables existing HPC application codes to run on accelerators with minimal source code changes [44].

7.4 Nek5000: References

[41] El Khoury, G., Schlatter, P., Noorani, A., Fischer, P.F., Brethouwer, G., Johansson, A.V., Direct numerical simulation of turbulent pipe flow at moderately high Reynolds numbers. Flow, Turbulence and Combustion, 91:475–495, 2013.

[42] Markidis S., Gong J., Schliephake M, Laure E., Hart A., Henty D., Heisey K, Fischer P.F., OpenACC Acceleration of Nek5000, Spectral Element Code, submitted to Advances in Engineering Software Journal.

[43] Gong J., Markidis S., Schliephake M., Laure E., Henningson D., Schlatter P., Peplinski A., Hart A., Doleschal J., Henty D., Fischer P.F. Nek5000 with OpenACC, accepted by the EASC2014 Solving Software Challenges for Exascale, April 02-04, 2014, Stockholm, Sweden

[44] Henningson D.,Schlatter P, Peplinski A., Markidis S., Schliephake M., Laure E., Hart A., Henty D., Fischer P.F., Heisey K., Toward Exascale Simulations with the CFD code Nek5000, accepted by 11th World congress on Computational Mechanics (WCCM XI) , July 20-25, 2014, Barcelona, Spain

[45] P. Fischer, J. Lottes, S. Kerkemeier, nek5000 Web page (2008). http://nek5000.mcs.anl.gov

Copyright © CRESTA Consortium Partners 2014

8 OpenFOAM

The open source CFD toolbox OpenFOAM® is widely used in engineering including turbo machinery, both in academia and industry. One of its outstanding features is its capability of simulating an entire hydraulic turbine.

Application	Identification (e.g. Version No)
OpenFOAM® Extend	OpenFOAM-1.6-ext

Table 8.1 OpenFOAM® - version information

8.1 OpenFOAM: scientific use case

8.1.1 Scientific use case: Francis pump turbine

In the last years hydro power plants have taken over the task of power-frequency control for the electrical grid. Therefore, turbines in storage power plants often operate outside their optimum in off-design conditions such as e.g. part load. At this operating point the discharge is reduced and strong swirling flow occurs in the draft tube. This is why the phenomenon of the vortex rope occurs which induces strong pressure pulsations into the hydraulic machine. It is very difficult to capture these pressure pulsations and transient phenomena by flow simulations with classical RANS (Reynolds averaged Navier-Stokes) turbulence models. The LES (Large Eddy Simulation) turbulence model approach resolves almost all turbulent structures (expect for the very smallest) and can therefore capture the vortex rope phenomenon much better. Due to the mesh requirements, especially in the boundary layer, such a simulation is not feasible with the available computational resources and with the performance of OpenFOAM-1.6-ext. For this reason a hybrid RANS-LES turbulence model approach is chosen which uses RANS in the boundary layer and LES in the core flow.

8.1.2 Computational requirements

The flow simulation of a pump turbine in part load operating conditions needs large computational simulations as the meshes must be quite large. The used mesh size for the hybrid RANS-LES turbulence model used in this project (10 and 20 million elements) is at the lower required border of what is necessary in terms of correct turbulence and physical modelling. As the time step is quite small (1°rev/time step) and up to 40 turbine revolutions must be simulated to capture some vortex rope rotations, roughly 15000 time steps are needed. If we wanted to run a LES simulation, a mesh which is approximately 50 times larger and an at least 10 times smaller time step would be required.

The latter simulation would be scientifically interesting as in hybrid models the boundary layer is still modelled in RANS and hybrid RANS-LES turbulence models are using blending functions between both approaches which can induce a modelling error in the area where both approaches are blended because they have different physical background.

8.2 OpenFOAM: exemplar simulation

The flow simulation contains all relevant components for capturing the vortex rope, e.g. the spiral casing, 20 stay and guide vanes, 7 runner blades and a straight draft tube. The simulation uses GGI (General Grid Interface) for coupling different meshes and a rotating mesh approach for the turbine. Two mesh sizes are used: one with 10 million elements (10M) and one with 20 million elements (20M). As RANS turbulence model the common SST model and as hybrid RANS-LES model a SST-based IDDES approach are used.

Simulation	Description

Copyright © CRESTA Consortium Partners 2014

RANS 10M	10 million elements grid size, RANS-SST turbulence model, 1°rev / time step, 15000 time steps, 24 iterations per time step, 224 cores
Hybrid 10M	10 million elements grid size, IDDES-SST turbulence model, 1°re v / time step, 15000 time steps, 24 iterations per time step, 224 cores
Hybrid 20M	20 million elements grid size, IDDES-SST turbulence model, 1°rev / time step, 15000 time steps, 24 iterations per time step, 448 cores

Table 8.2: OpenFOAM® - exemplar simulations

8.2.1 Timescales

Start date:	27/01/2014
End date:	28/02/2014
Total duration:	9856 CPU days / 236544 CPU hours

Table 8.3: OpenFOAM® - timescales

8.2.2 People

Name	**Role**
Timo Krappel	Pre-Processing, Simulations, Post-Processing
Stephen Sachs	Compiling

Table 8.4: OpenFOAM® - people

8.2.3 Environment

The Cray XE6 "HECTOR" at EPCC at the University of Edinburgh was used for the simulation runs. HECTOR has 2816 compute nodes, each with two 16-core AMD Opteron 2.3GHz Interlagos processors. This gives a total of 90112 cores. Each 16-core processor has 16 GB of shared memory. The nodes are coupled by a Cray Gemini network with 5 GB/s or more MPI point-to-point bandwidth and 1-1.5μs latency between two nodes. The theoretical peak performance of HECTOR is over 800 Tflops.

Platform	**Platform specs**
HECTOR	Cray XE6

Table 8.5: OpenFOAM® - environment

8.2.4 Exemplar simulation runs

CPU time = nodes*CPUs*wall time

Simulation	**Description of run**			
	Nodes	**Total cores**	**Wall time**	**Total core time**
RANS 10M	14	224	240 h	53760 h
Hybrid 10M	14	224	264 h	59136 h
Hybrid 20M	28	448	276 h	123648 h

Table 8.6: OpenFOAM® - exemplar simulation runs on HECTOR

8.3 OpenFOAM: Scientific impact

As progress in the development of hybrid RANS-LES turbulence models has been achieved within the last few years, their use is not common. Flow simulations with

Copyright © CRESTA Consortium Partners 2014

hybrid RANS-LES turbulence models need large computational resources, especially for Francis pump turbines having a high Reynolds number. For getting good statistical results, large meshes and a very large number of small time steps (i.e., simulation over a long time interval), is needed. Validation of the simulation results is performed with a large set of measurement data.

Increased knowledge of pressure pulsations in the draft tube can help in the early design phase of hydraulic machinery. As pump turbines are often located in power plants inside concrete dams, computer simulations can be of aid in developing a better and a more stable construction of the turbine.

Flow simulations with a hybrid RANS-LES turbulence model can resolve turbulent flow structures more detailed than RANS turbulence models (as can be seen in Figure 8.1). This leads to a better resolution of the vortex rope phenomenon and the pressure fluctuations in the draft tube. Results of flow simulations performed within the CRESTA-project are used for an abstract in [46].

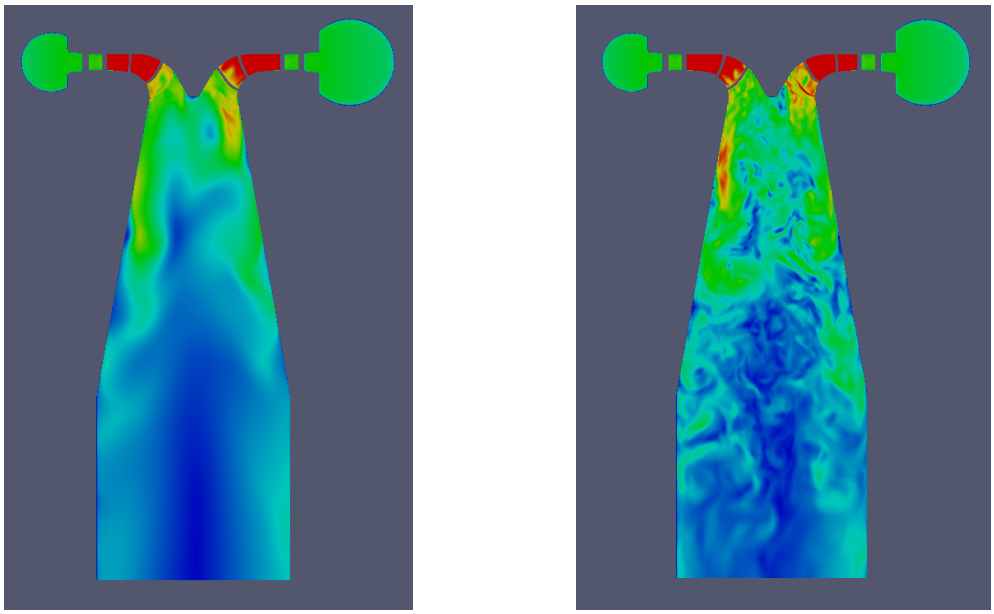

Figure 8.1 Francis turbine flow simulation using OpenFOAM®; left: RANS turbulence model, right: hybrid RANS-LES turbulence model

8.4 OpenFOAM: References

[46] Krappel T., Kuhlmann H., Kirschner O., Ruprecht A., Riedelbauch S. "Validation of an IDDES-type Turbulence Model and Application to a Francis Pump Turbine Flow Simulation", 10th International ERCOFTAC Symposium on Engineering Turbulence Modelling and Measurements, Marbella, Spain, 17-19 September 2014

Copyright © CRESTA Consortium Partners 2014