CRESTA

Collaborative Research into Exascale Systemware, Tools and Applications

Autumn 2012

Welcome to the 1st CRESTA Project Newsletter. I hope you find it interesting and informative. CRESTA's focus on the software challenges of exascale computing makes it a unique project. While a number of projects worldwide are studying hardware aspects of the race to perform 10^{18} calculations per second, no other project is focussing on the exascale software stack in the way that we are.

By limiting our work to a small set of representative applications we hope to develop key insights into the necessary changes to applications and system software required to compute at this scale. In its first year, the project has made a very good start but we don't underestimate the challenges ahead.

When studying how to compute at the exascale it is very easy to slip into a comfort zone where incremental improvements to applications eventually deliver the necessary performance. In CRESTA, we recognise that incremental improvements are simply not enough and we need to look at disruptive changes to the HPC software stack from the operating system, through tools and libraries to the applications themselves. From the mid-1990s to the end of the last decade, HPC systems have remained remarkably similar (with performance increases being delivered largely through the increase in microprocessor speeds). Today, at the petascale, we are already in an era of massive parallelism with many systems containing several hundred thousand cores. At the exascale, HPC systems may have tens of millions of cores. We simply don't know how to compute with such high levels of parallelism.

CRESTA is studying these issues and identifying a huge range of challenges. With the first exascale systems expected in the early 2020s, we need to prepare now for the software challenges we face which, we believe, greatly outnumber the corresponding hardware challenges. It's a very exciting time to be involved in such a project.

I hope you enjoy reading this newsletter. If you have any comments, queries or suggestions for future editions please let us know.

Dr Mark Parsons
CRESTA Project Coordinator

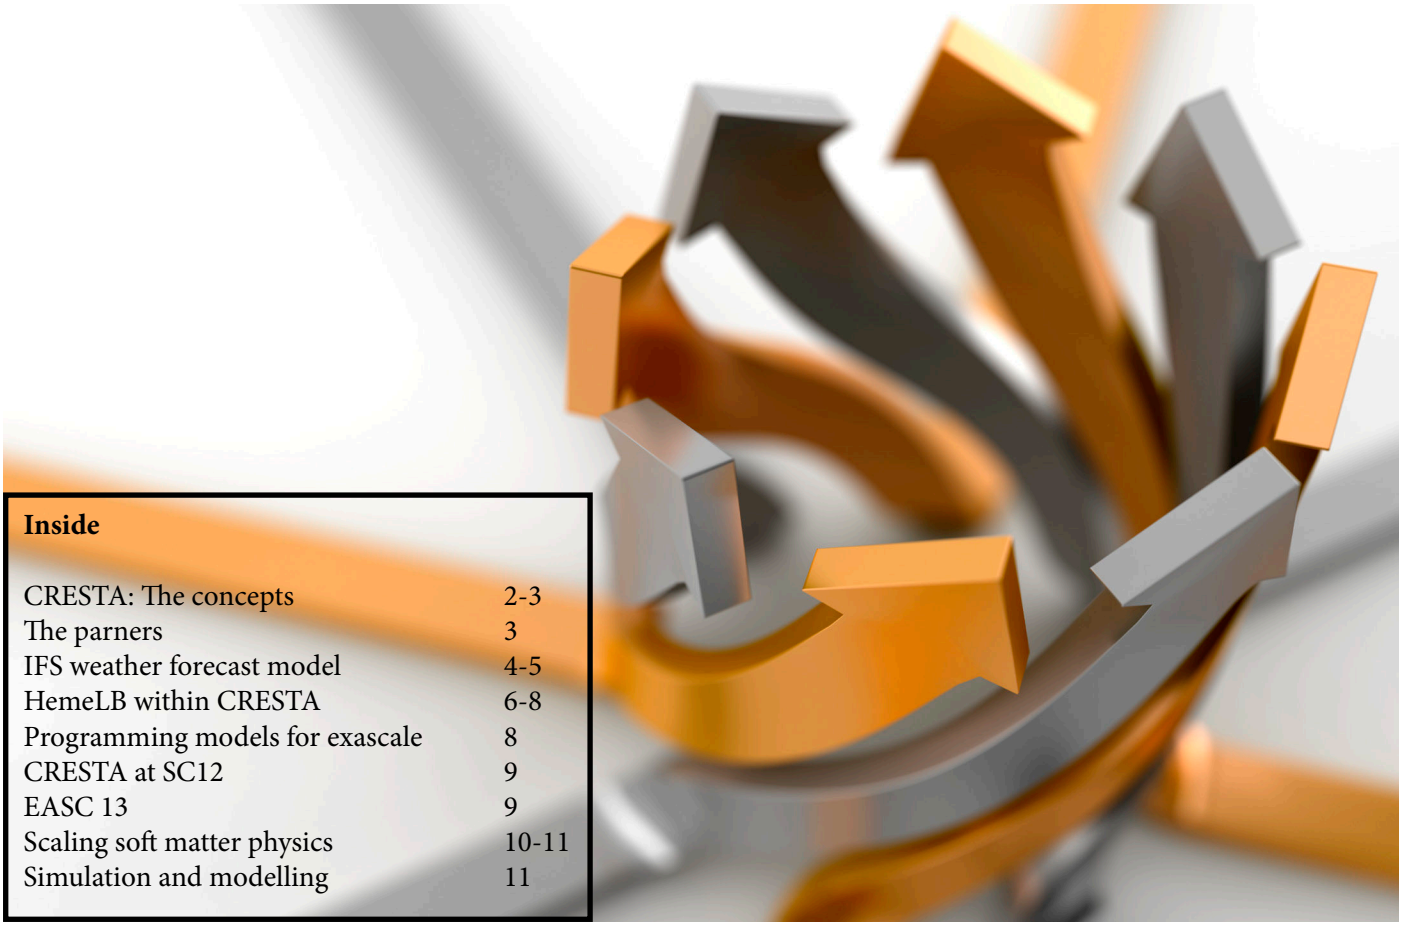

Inside

CRESTA: the concepts, facts and figures

Lorna Smith, EPCC

'The main goal of CRESTA is to develop techniques and solutions which address the most difficult challenges that computing at the exascale can present.'

CRESTA has recently reached the end of its first year. During this time each of the project's partners has been working closely on research topics across the exascale domain and have been publicising the project to the community. The CRESTA project is a large and complex one, so in this article I hope to provide some facts and figures and to introduce some of the topics and challenges we are looking to address. Coupled with the different articles within this newsletter we hope to give you a real feel for the activities being undertaken within CRESTA.

Concepts

CRESTA has two integrated strands: one focused on enabling a key set of applications for exascale (the co-design applications), the other focused on building and exploring appropriate 'systemware for exascale platforms.' Associated with this is one of our key project concepts: the co-design process, with the co-design applications providing guidance and feedback to the exascale software development process, and integrating and benefiting from this development in a cyclical process.

The second important concept within CRESTA is the use of dual pathways to exascale solutions. Many problems in HPC hardware and software have been solved over the years using an incremental approach. Most of todays systems have been developed incrementally, growing larger and more powerful with each product release. However, we know that some issues at the exascale, particularly on the software side, will require a completely new, disruptive approach. CRESTA is employing both incremental and disruptive approaches to achieve technical innovation.

Co-design applications

The six co-design vehicles within CRESTA represent an exceptional group of applications used by European academia and industry to solve critical grand challenge issues. CRESTA will enable these applications to prepare for, and exploit, exascale technology, enabling Europe to be at the forefront of solving world-class science challenges. These challenges include:

- Biomolecular systems. Molecular Dynamics simulation is one the most rapidly growing methods for modelling in general, and for life science in particular.
- Global Numerical Weather Prediction. ECMWF leads the world in this advanced computer observation-analysis modelling technique used to predict the weather.
- The Virtual Physiological Human initiative is a collection of European projects which undertake research in biomedical modelling and simulation of the human body. This will improve our ability to predict, diagnose and treat disease, and will have a dramatic impact on the future of healthcare, the pharmaceutical, and the medical device industries.
- Fusion Energy. The decision to construct the first power-plant scale fusion reactor, ITER, in Cadarache, France has dramatically increased the need to solve the remaining technical challenges of fusion reactor design. ITER is planned to be operational within approximately a decade and the budget is counted in tens of billions of euros. Before ITER is operational, numerical simulations are the only possible way to investigate the behaviour of full scale reactors.
- Engineering. Fluid dynamic simulations are widely used within the automotive, aeronautical and heavy engineering industries. Of particular interest to CRESTA is the ability to model the complete hydraulic system in hydro power plants.

Systemware for exascale

Systemware represents any additional enabling software required for the applications. For CRESTA, this involves

five distinct areas with the aim of producing an integrated suite including open source interfaces, standards and new software. These areas are:

- The development environment, led by KTH. In order to cope with the challenges of exascale computing, application developers need support in all phases of the application lifecycle, including programming models that allow the construction of efficient, yet portable, applications, and advanced compilation techniques and adaptive runtime environments. The development environment also encompasses debugging and Application Performance Tools. CRESTA boasts an impressive range of European tool partners for debugging and performance analysis. This includes Allinea's DDT debugger, KTH's perfminer and TUD's Vampir tool-suite and MUST runtime error detection tool (developed in collaboration with LLNL and ASC Trt-Labs).
- Key numerical algorithms and libraries, led by the University of Stuttgart. Many of the most common underlying algorithms are already performance limited on current tera and petascale platforms. These limitations will only increase, and in some cases become untenable, at exascale. For example, achieving a reasonable computation/communication balance without a major increase in problem size, coping with increasing memory latencies, avoiding global synchronisation points and load imbalance are all important considerations.
- Pre- and post- processing tools, led by DLR. Application domains, such as fluid dynamics, meteorology, nuclear physics, or material science, heavily rely on numerical simulations on HPC resources. A simulation and analysis process is typically composed of three steps: the first step is the domain decomposition by means of partitioning and mesh creation; the second step is the numerical computation of the simulation; visualisation and analysis of the resulting data is the third step. The performance of extreme-scale simulations on exascale computing systems will depend on the efficiency of massively parallel numerical computations and their scalability. However, pre- and post-processing, including rendering, are also crucial factors for the success of a simulation. CRESTA is focused on these two steps.
- Underpinning and cross cutting technologies, led by EPCC. While hardware advances towards exascale are outwith the scope of CRESTA, these advances will significantly influence software and systemware developments. Hence, CRESTA is tracking underpinning exascale technology and future architecture developments towards exascale and assessing their impact on exascale software developments within CRESTA.

As you can see, CRESTA covers a broad range of interesting topics and we hope that you enjoy reading about some of these areas in more detail in this newsletter.

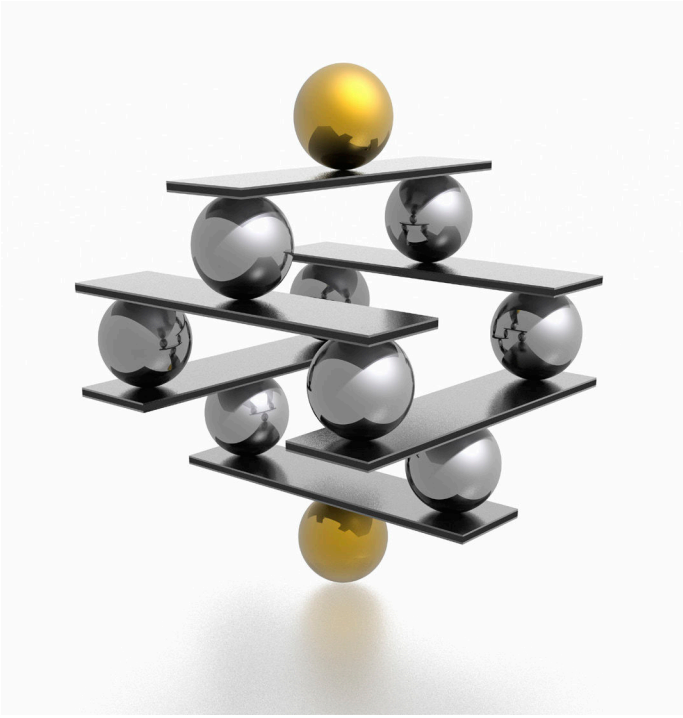

The CRESTA Partners:

Åbo Akademi University, Finland

Allinea Software, UK

Cray UK Limited, UK

Ecole Centrale Paris, France

EPCC, The University of Edinburgh, UK

German Aerospace Centre (DLR), Germany

Kungliga Tekniska högskolan (KTH), Sweden

The European Centre for Medium-Range Weather Forecasts, International

The University of Jyvaskyla, Finland

The University of Stuttgart, Germany

Tieteen Tietotekniikan Keskus OY (CSC)'s IT Centre for Science Ltd, Finland

University College London, UK

ZIH, Technische Universität Dresden, Germany

Preparing ECMWF's Integrated Forecast System (IFS) weather forecast model for Exascale

George Mozdzynski, ECMWF. Mats Hamrud, ECMWF. Nils Wedi, ECMWF.

Jens Doleschal, TUD. Harvey Richardson, Cray.

Today, the European Centre for Medium-Range Weather Forecasts (ECMWF) runs its world leading 16 km global T1279 operational weather forecast model using 1,536 cores of an IBM Power6 system located at its headquarters in Reading, UK.

By 2025 ECMWF expects to be running a 2.5 km global forecast model on an exascale system that should be available, and hopefully affordable, by then. To achieve this would require IFS to run efficiently on about 1000 times the number of cores it uses today. Nothing like a challenge!

But can this be done? Well the good news is that ECMWF, working within the CRESTA project, has already demonstrated IFS running a 10 km global model efficiently on over 50,000 cores of HECToR, the UK's national HPC resourse and a CRAY XE6 based at EPCC, UK (see charts). Of course, getting to over a million cores remains a formidable challenge, and many scalability improvements have yet to be implemented.

CRESTA has been essential for the above IFS achievement as it brings together partners such as CRAY, with their experience in PGAS languages, and Technische

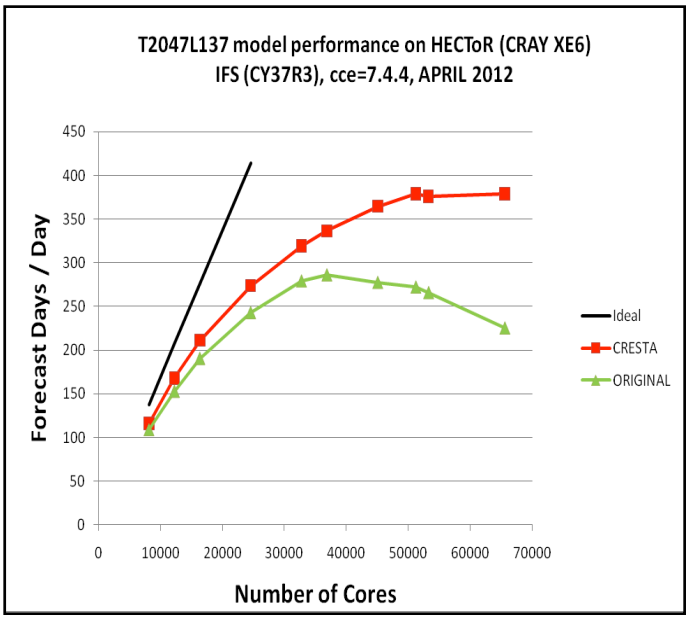

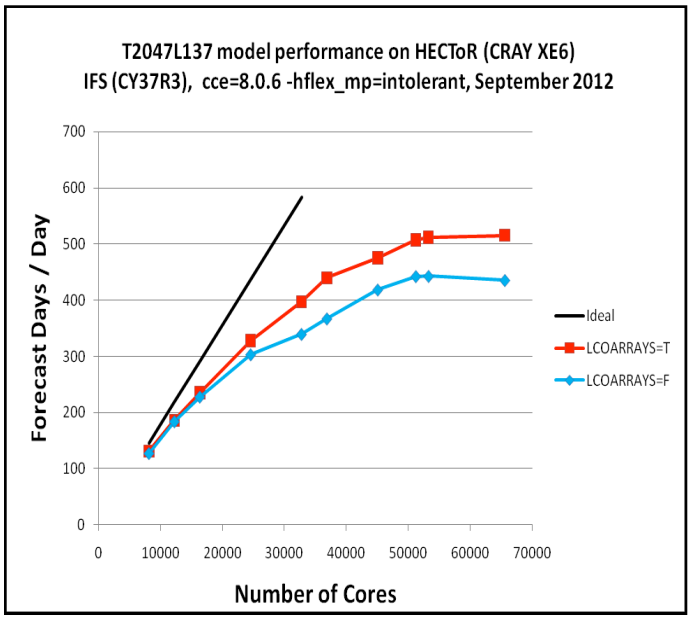

Universität Dresden, providing exascale enhancements to their profiling tool vampir, just to mention two of the partners in the project.

Now the technical bit. Within CRESTA, ECMWF is exploring the use of Fortran2008 coarrays; in particular, it is possibly the first time that coarrays have been used in a world leading production application within the context of OpenMP parallel regions. The purpose of these optimisations is primarily to allow the overlap of computation and communication, and further, in the semi-Lagrangian advection scheme, to reduce the volume of the data communicated by removing the need for a constant width halo for computing the trajectory of particles of air backwards in time. The importance of this research is such that if these developments are successful, then the IFS model can continue to use the spectral method to 2025 and beyond for the currently planned model resolutions on an exascale sized system. This research is further significant as the techniques used should be applicable to other hybrid MPI/OpenMP codes with the potential to overlap computation and communication.

In a nutshell, IFS is a spectral, semi-implicit, semi-Lagrangian weather prediction code, where model data exists in three spaces, namely, grid-point, Fourier and spectral space. In a single time-step, data is transposed between these spaces so that the respective grid-point, Fourier and spectral computations are independent over two of the three co-ordinate directions in each space. Fourier transforms are performed between grid-point and Fourier space, and Legendre transforms are performed between Fourier and spectral space.

At ECMWF, this same model is used in an Ensemble Prediction System (EPS) suite where today 51 models are run at lower resolution with perturbed input conditions to provide probabilistic information to complement the accuracy of the high resolution deterministic forecast.

The EPS suite is a perfect candidate to run on future exascale systems, with each ensemble member being independent of other such jobs. Increase the number of members and their resolution and we have the potential to fill an exascale system having an estimated 100 million cores. Hopefully by 2025, a more affordable ($20M) system will have 10 million cores. Even with these huge core counts there will always be a need for a high resolution deterministic forecast which is more challenging to scale than an EPS suite and the reason for ECMWF's focus in the CRESTA project.

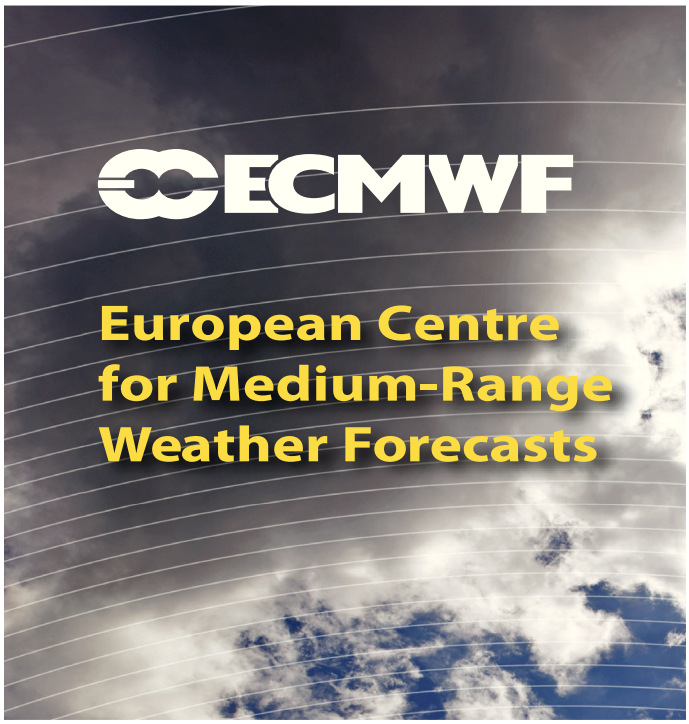

HemeLB within CRESTA: computational haemodynamics en route to patient-specific treatment planning using the exascale

Timm Krueger, UCL

It is estimated that about 1 - 5% of the entire population have aneurysms in their blood vessels; these are pathological balloon-like bulges of blood vessels which can form following the weakening of the vessel wall. With a cranial aneurysm, there is an increased risk of artery rupture and internal haemorrhage, which is a life-threatening situation.

Aneurysms are often found during routine investigations as they are usually asymptomatic. Most patients with a cranial aneurysm never suffer consequences, but in some cases, aneurysms rupture with severe, if not fatal, outcomes (e.g., stroke). Once a rupture has occurred, physicians must plan and begin treatment rapidly. The choice of treatment strongly depends on patient-specific factors (e.g., location of aneurysm, vessel geometry, blood pressure). Therefore, the ultimate goal is the development of a real-time predictive tool accounting for the short time-scale of clinical decision-making.

But even in the case of a non-ruptured aneurysm it is not directly obvious whether it might turn into a life-threatening situation in the future and whether intervention is necessary at all. Although there are means to stabilise or remove aneurysms (e.g., inserting stents into the artery to modify the flow field or clipping the aneurysm from the outside), these interventions are not without risks and should be applied only when required. Unfortunately, the factors affecting the potential danger of a given aneurysm are not yet known. It is thus desired to provide methods predicting the risk of a future vascular malfunction in individual patients.

There is strong evidence that the blood flow pattern plays a role in the formation process of aneurysms. This brings us to the question as to whether it may be possible to find significant correlations between specific features of the flow near the aneurysm and how it eventually develops.

Computational fluid dynamics (CFD) comprises a wide class of numerical algorithms designed to tackle problems related to the flow of fluids. In recent years, larger and larger volumes have been simulated with higher and higher resolution. This development calls for more efficient large-scale simulations. Therefore, nowadays, many CFD applications belong in the realm of high performance computing (HPC). Current supercomputers are capable of achieving peak performance of a few petaflops per second (1 Pflops/s = 10^{15} floating point operations per second), whereas modern laptop computers typically reach of the order of 10^{10} flops/s. The aim of CRESTA is to reach the exascale (10^{18} flops/s) by 2018. This endeavour strongly relies on designing software solutions which have to run with hundreds of thousands, if not millions, of parallel processes. As blood flow is essentially a fluid dynamics problem, it seems natural to take advantage of recent developments, in both the CFD and HPC communities, in order to better understand the effects of blood flow on vessel topologies. Eventually, it will be possible to provide a simulation tool assessing the personal risk of aneurysm formation and rupture since each patient's circulatory system is different.

HemeLB is designed as a specialised predictive tool for fluid flows in complex geometries. Its main focus is simulating blood flow in parts of the cerebral arterial network. HemeLB employs a highly efficient implementation of the lattice Boltzmann (LB) algorithm which found its way into the CFD community around two decades ago. Due to its locality, the LB algorithm is intrinsically easy to parallelise and therefore a method of choice for HPC. A sophisticated communication approach based on the Message Passing Interface (MPI) allows HemeLB to be run on large supercomputers. Therefore, the HemeLB project fits perfectly into CRESTA, as it stimulates the development of technologies leading to the next generation of supercomputers on the exascale.

The typical workflow of a HemeLB simulation has three main steps: (1) obtain patient-specific data from a medical scan, (2) pre-process the scan to extract geometry data suitable for a simulation, and (3) simulation, visualisation, and data assessment by running HemeLB. At the beginning, an X-ray rotational angiography scan is performed to obtain patient-specific data. It provides the highest resolution of the available imaging methods. The dataset provided by the scan typically contains between 10^5 and 10^8 voxels whose brightness carries data about their locations – inside or outside the blood vessels. In the next step, the data has to be processed in order to construct a geometrical representation of the blood vessel network. At this point, the intervention of the user is required: the voxel data is converted into a surface representation which defines the blood vessel boundaries. There is ongoing work at the University of Jyväskylä in Finland to automate this currently labour intensive procedure. Once the surface data is available, the region

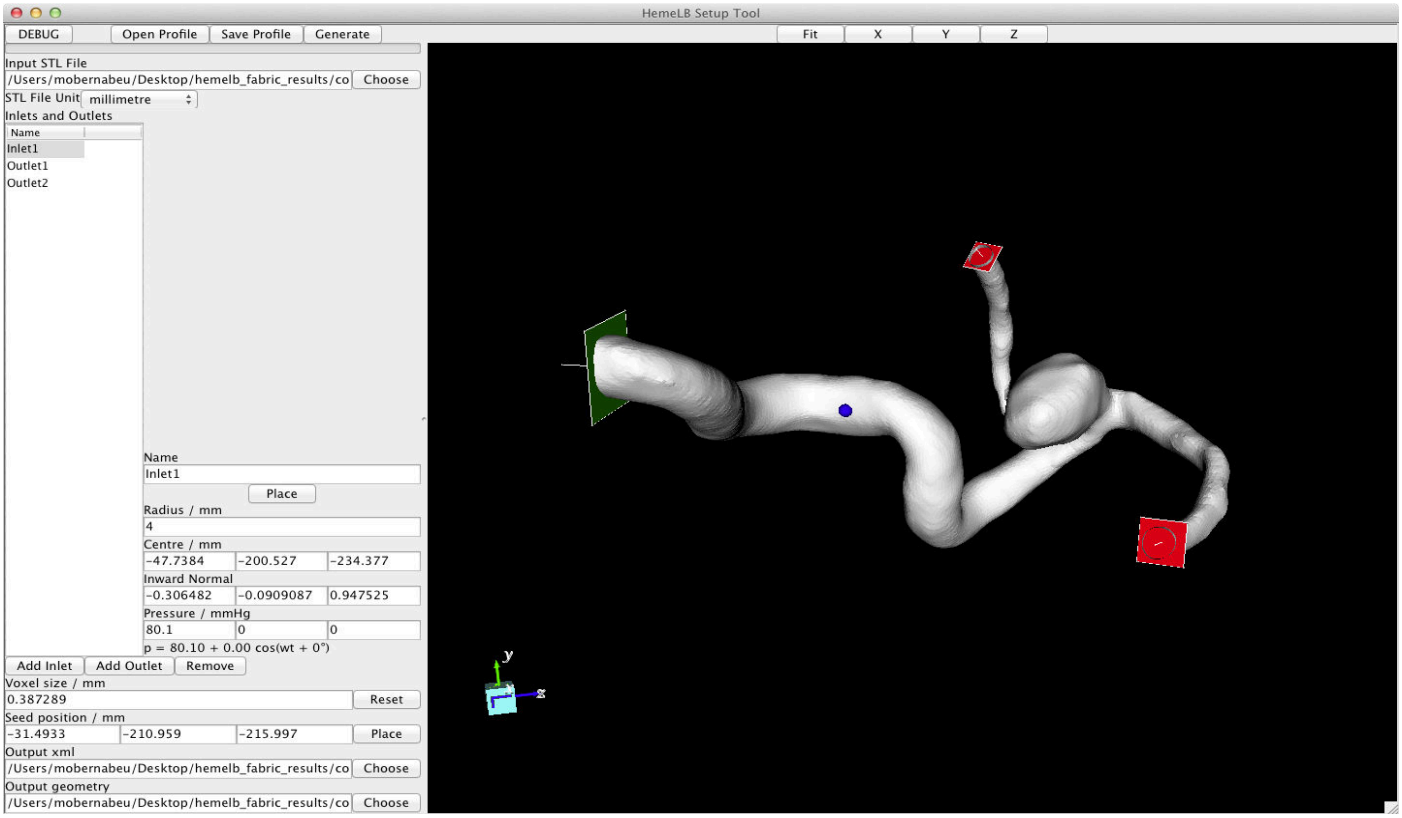

Figure 1: HemeLB Setup Tool and a processed aneurysm geometry taken from an X-ray angiography scan provided by the Universitat Pompeu Fabra in Barcelona.

of interest (ROI, containing the aneurysm and connected blood vessels) has to be selected employing the graphical user interface of the HemeLB Setup Tool as shown in figure 1. The user has to specify the simulation resolution and the so-called inlet and outlet planes through which the blood enters and leaves the network.
Only now does the main HemeLB application come into play. The prepared geometry file is loaded together with an additional input file. The latter contains information about the pressure values applied at the inlets and outlets. The geometry is mapped onto a regular lattice where each lattice site can either be fluid (within the blood vessel) or tissue (outside). Due to the sparseness of the arterial network, usually only 5 - 10% of the volume consists of fluid sites. This underlines how important it is to use software optimised for sparse geometries, like HemeLB. Since the simulation is running on a highly parallelised computer, the total fluid volume has to be subdivided into domains of approximately equal size to ensure computational load-balancing. This is performed at runtime by the parallel partitioner ParMETIS which eventually provides the final partition of fluid sites between processes.

After reading the geometry and boundary conditions, the flow simulation is executed. For large-scale simulations, it is impractical to write complete snapshots to the hard disks: a single snapshot can easily require several gigabytes. Due to the ever-growing size and complexity of computer simulations, new data-handling approaches are necessary. Indeed, the German Aerospace Centre (DLR) institutes in Cologne and Braunschweig explore innovative data-processing ideas within CRESTA. The fundamental idea is to enable the interactive inspection of a coarse data subset at runtime, which drastically decreases the amount of data to be transferred. In the future, the user will be able to select regions of particular interest and request highly resolved snapshots of these. The HemeLB Steering Client for remote connections receives real-time visualisations of the simulation and allows the user to define which data to visualise (e.g., streamlines, wall stresses, or pressure distributions, cf. figure 2). An advantage of this approach is that both the computation and visualisation are executed on the same machine and share the same data and resources, producing data quickly and in real-time.

There are several ideas for enhancing the efficiency and physical relevance of HemeLB in the future. In order to increase the computational performance and to eventually approach exascale capability, EPCC is working on a hybrid parallelisation using GPUs and OpenMP. It is also possible to apply steering not only to the visualisation pipeline, but also to the simulation itself. This way, simulation parameters could be changed at runtime to investigate different scenarios without performing a restart. There are additional model components planned, including magnetic microparticles (targeted drug delivery for cancer treatment), vascular remodelling (predicting effects of flow on aneurysm formation), and a coupling to coarse-grained circulation models (improving estimates of the inlet/outlet pressure conditions). The latter multiscale

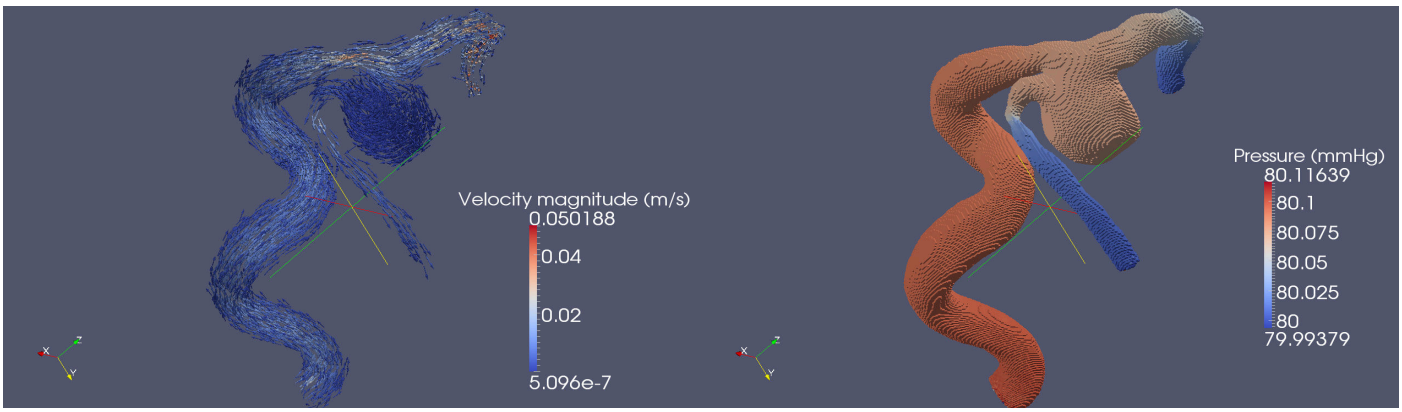

Figure 2: Velocity (left) and pressure (right) distributions in an aneurysm and connected blood vessels as obtained from an example HemeLB simulation.

project overlaps with the scope of the MAPPER (Multiscale APPlications on European e-infRastructures) project. The ultimate goal of the HemeLB project is to reach a state where the entire workflow between patient scan and assessment of the simulation results is automatised and takes no longer than one hour which is the typical time-scale on which physicians have to make decisions about to procedures with the patient on the operating table. CRESTA provides the optimal environment to approach this ambitious aim.

Programming models for the exascale era

Stefano Markidis, KTH

Programming models for exascale computing are an active research area in CRESTA. We are studying new features of traditional parallel programming models, and the advantages of novel emerging programming models for CRESTA's co-design applications, numerical libraries, and pre- and post-processing tools.

On the road to exascale, programming models face three main challenges. First, programming models should enhance the scalability of the co-design vehicle applications and allow efficient and fast parallelisation. Second, programming models should make it easier to program effectively hybrid parallel supercomputers with, for example, Graphics Processing Units (GPU) and Intel Many Integrated Cores (MIC), as accelerators are becoming more and more common on massively parallel computers. Third, programming languages should provide support for automatic load balancing and simplify the implementation of fault-tolerant algorithms.

In CRESTA, two approaches are being utilised to achieve scalability on exascale machines. First, the new features of MPI-3 and OpenMP 4 are being considered. In particular, non-blocking collectives, the increased support for one-sided communication in MPI-3 and thread affinity in OpenMP 4, can be used to make co-design vehicles scale on a higher number of cores. Second, CRESTA focuses on exploiting the features of Partitioned Global Address Space (PGAS) languages, such as Coarray Fortran, UPC and Chapel, and on coupling PGAS languages and MPI to increase the scalability of existing MPI-based codes. Coarray Fortran is currently used in the IFS code in combination with MPI and OpenMP, resulting in an increase of IFS parallel speed-up.

OpenACC, a programming model based on compiler directives for GPUs, is CRESTA's main approach to addressing the difficulty of program GPU systems effectively. Part of the computational kernels of Nek5000, and of the HemeLB and GROMACS analysis tools, are being ported to GPU systems using OpenACC. This requires very few modifications to the code. However, the challenge is still in achieving the performance of lower level approaches, such as CUDA.

Task-based programing languages, such as Cilk, enable automatic load balancing and are the most promising approach to implementing fault-tolerant algorithms. A skeleton version of the Particle-in-Cell method, used in the Elmfire code, is being developed in Cilk to study the potential of task-based approaches for exascale computing.

All of these improved and innovative programming models are being tested against CRESTA's co-design applications, as well as CRESTA's numerical libraries and pre- and post-processing tools. These investigations are expected to provide important indicators of the programming models that will be used in future exascale supercomputers.

CRESTA at SC12

Katie Urquhart, EPCC

CRESTA will be represented in a number of ways at this years SC conference being held Salt Lake City, Utah from 10th – 16th November.

Firstly, congratulations to our TUD partners for becoming a finalist for best paper at SC12. The paper, 'MPI Runtime Error Detection with MUST', details TUD's work within CRESTA.

We are hosting a workshop at the conference on co-design for exascale, with both invited and submitted talks. The workshop will showcase the work of the project provide an opportunity for attendees to learn about state of the art in co-design from international speakers. 'Preparing Applications for Exascale through Co-design' takes place on Friday 16th Nov from 8.30am – 12.30pm.

Two of our CRESTA partners, Allinea Software and Technische Universität Dresden, are contributing to a tutorial on Allinea's flagship debugger DDT and TUD's MUST. The turorial, 'Debugging MPI and Hybrid-Heterogenous Applications at Scale', will take place on Sunday 11th Nov from 8.30am – 5pm.

Alistair Hart, one of the key CRESTA staff members at Cray, will be presenting a tutorial on OpenACC with his Cray colleagues Luiz DeRose, James Beyer, and Heidi Poxon. The tutorial will not only introduce the OpenACC parallel programming model for accelerator based systems, but will also demonstrate the full development cycle when porting applications to OpenACC, covering high level compilers, libraries, and tools for OpenACC. This tutorial, 'Productive, Portable Performance on Accelerators Using OpenACC Compilers and Tools' takes place on Sunday 11th Nov from 8.30am – 5pm.

Finally, a number of our partners are exhibiting this year, these include: Allinea Software, Cray, CSC, EPCC, TUD and the University of Stuttgart. All are keen to see you at their booths so please do drop by.

We look forward to discussing our project with you at one, or more, of these events at SC12.

EASC 2013

EASC 2013 is a conference organised by EPCC, University of Edinburgh, in partnership with the CRESTA, NAIS and NuFuse projects. The aim of the conference is to bring together all of the stakeholders involved in solving the software challenges of the exascale – from application developers, through numerical library experts, programming model developers and integrators, to tools designers.

Invited talks from Satoshi Matsuoka, Vladimir Voevodin, Bill Tang, George Mozdzynski, Peter Coveney, and Jack Dongarra.

Authors are invited to submit novel research and experience in all areas associated with developing applications for exascale. Further information can be found at:

www.easc2013.org.uk

Scaling soft matter physics to a thousand GPUs and beyond

Alan Gray, EPCC. Alistair Hart, Cray

Simulations based around fluid dynamics offer a powerful way to study, predict and ultimately improve the behaviour of 'soft matter': everyday materials such as paints, engine lubricants, foodstuffs, and cosmetic items; hi-tech items such as liquid crystal displays (LCDs) and even biological fluids within the body.

Over the last 12 years, a collaboration between the University of Edinburgh's Soft Condensed Matter Physics group and EPCC has developed simulations of soft matter systems using the lattice Boltzmann method and the parallel computing code 'Ludwig' to accurately capture the physics of systems such as mixtures, suspensions and liquid crystals. Understanding and controlling the phase separation of liquid mixtures, for example, can improve the shelf-life of foodstuffs, one of the many practical applications of the research.

Ludwig performs excellently on 'traditional' CPU-based supercomputers, and is able to take advantage of many thousands of compute cores in parallel. But recent work at EPCC has taken this a step further by engineering the code for exploitation of large-scale GPU accelerated architectures, such as the cutting-edge Cray XK6™ supercomputer. Such machines offer an impressively high ratio of performance to power consumption and can be thought of as a possible template for future exascale systems (for which power consumption will be a key issue). This work is supported by the EU-funded, EPCC-led FP7 research network Collaborative Research into Exascale Systemware, Tools and Applications (CRESTA) project, in which EPCC and Cray collaborate with other institutions and companies across Europe, looking forward towards the next generation of supercomputers.

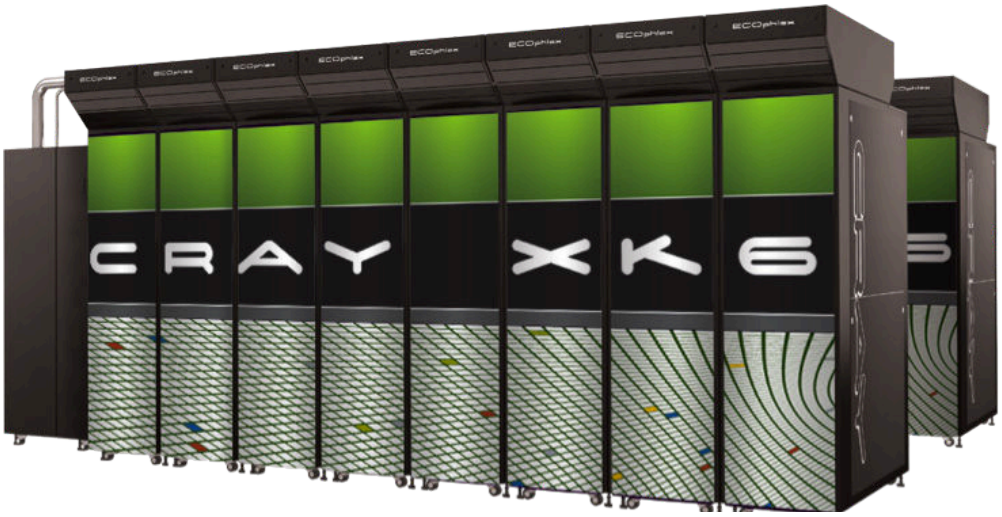

The Cray XK6 Supercomputer

"Large-scale computer simulation has become a central tool for research in materials physics," says University of Edinburgh professor and head of the Soft Matter group Michael Cates. "And thanks to the Cray XK6 system and other similar innovations, the rate of increase in computational capability is breathtaking."

The main challenge in this work is to maintain good scalability whilst exchanging data between many GPUs. To ensure optimal performance scaling, use was made of NVIDIA® CUDA™ stream technology and was combined with MPI to minimise communication overheads through overlapping. The new code is observed to scale excellently such that it can exploit a thousand GPUs in parallel on the Cray XK6, with a view to even larger simulations as bigger machines become available. "A crucial factor in obtaining accurate results is the size of the physical system that we can simulate," says lead Ludwig author Kevin Stratford (EPCC). "To model complex problems in large systems, we need efficient, scalable application performance over large numbers of nodes."

Further work is underway to enable advanced functionality for the GPU architecture such as the ability to include, in

the simulation, particle suspensions together with liquid crystals: this is of particular interest for new generations of LCD display technology.

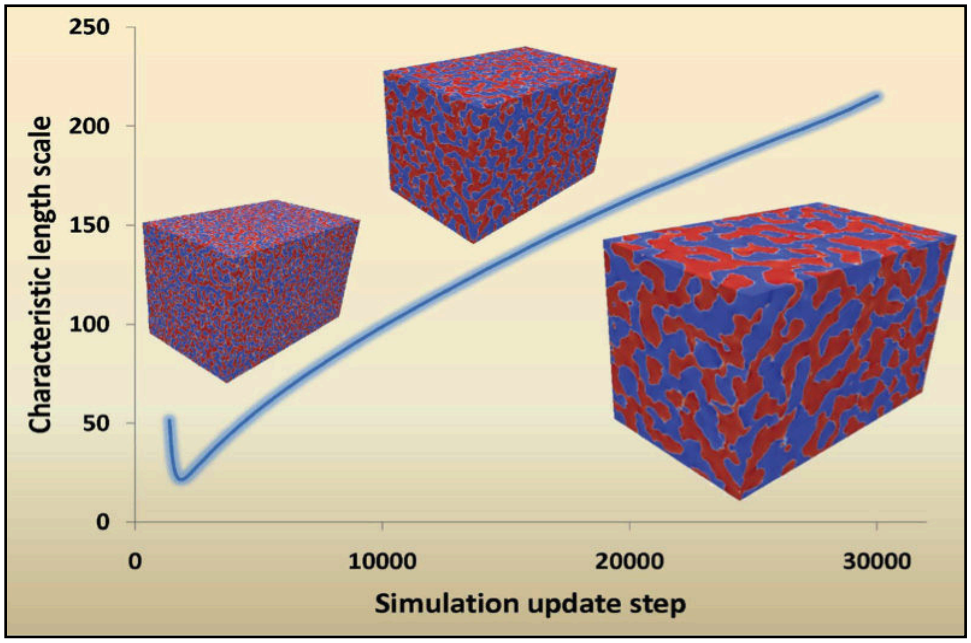

LIQUID INSIGHT: Plot of the two components of a binary fluid separating during a six-hour portion of a 12-hour run on 936 GPU nodes of a 10-cabinet Cray XK6 system. After initial equilibration, the characteristic length scale of teh system grows as predicted theoretically. (Credit: K. Stratford, EPCC)

Simulation and modelling technologies

Stephen Booth, EPCC

Simulation and modelling are important tools in the development of exascale systems. There are very few other mechanisms for evaluating our designs for exascale hardware and software other than developing models of their behaviour and simulating these models in a computer. The behaviour of both hardware and software needs to be modelled. EPCC has recently been investigating the modelling requirements of the CRESTA project and assessing the most appropriate technologies.

In the early stages of the design process these models need to be quite simple and abstract. This allows us to develop and evolve our designs quickly and efficiently. If we attempt to use overly complex models in these early design stages then we will waste time and resources performing overly detailed simulations of design choices that will be abandoned before the final system is built. Current thinking about exascale hardware design is that these designs will be highly constrained by system power consumption. To keep power consumption within acceptable levels exascale systems will need to utilize very high degrees of parallelism and the performance of their communication systems may also have to be limited. This implies that we should be using software models that explicitly capture the available parallelism and the communication requirements of an algorithm. One way of capturing this information is to consider modelling the parallelisable sections of the algorithm as directed acyclic graphs.

Application behaviour can be simulated at a high level by simulating the communication pattern of the application. This allows the behaviour of the application to be extrapolated to different (possibly theoretical) hardware platforms. This allows us to explore the behaviour of applications on exascale hardware well before such hardware becomes available. Various different approaches to application simulation exist. One particularly interesting approach is the use of simple skeleton applications to drive the simulation. These are lightweight simple codes intended to capture the essential behaviour of larger much more complex applications. They provide a mechanism of exploring the behaviour of new designs without the cost of first developing the design into a fully functional application. As this approach aims to capture the communication pattern rather than the details of the computational sections it provides a mechanism to develop a directed acyclic graph model into a form that can be simulated.

Our research has found a number of simulation platforms that are suitable for this kind of simulation. These can give useful insights into the limitations that the network imposes on application performance. Many of these platforms are explicitly targeting the development of exascale systems. Though these simulation tools are useful they are fairly complex research tools and can be quite difficult to use.

www.cresta-project.eu

The CRESTA project has received funding from the European Community's Seventh Framework Programme (ICT-2011.9.13) under Grant Agreement no. 287703