

the Exascale Report™



The voice of the emerging exascale community

<http://TheExascaleReport.com>
May 2013

Feature Report

The Path to Exascale Starts with Better Use of HPC Today

HPC Progress in the UK Adds Value and Steers the Course to Extreme Scale

By John Barr

The HPC industry is involved in a worldwide arms race to build Exascale systems. The next generation of HPC systems will be much more challenging for users, with millions of heterogeneous processor cores, complex memory hierarchies and different programming approaches. The UK is addressing these issues to keep its industry at the forefront of HPC use. In recent years the UK government has been lobbied by the HPC community to fund systems to help smooth the transition towards the next generation of HPC hardware and software. The UK e-Infrastructure has been refreshed in order to support academic use, and to increase economic output through the industrial exploitation of HPC. Using large scale HPC facilities can enable scientists and engineers to do things that were not possible before, such as adding new capabilities to an application, or increasing the fidelity of modelling used – which can in turn lead to the development of better, lighter, stronger products that are less expensive to manufacture.

The Hartree Centre

The Hartree Centre was created by the Science and Technology Facilities Council (STFC) as a research collaboration in association with IBM in 2012 at the Daresbury Science and Innovation Campus at Warrington near Manchester, with the support of a £37.5 million investment from the UK government. The Hartree

Consortium comprises STFC, IBM, Intel, NVIDIA, DDN, Mellanox, ScaleMP, Platform Computing and OCF.

STFC is a multi-disciplinary research organisation, one of the UK's seven publicly funded Research Councils responsible for supporting, co-ordinating and promoting research, innovation and skills development. IBM is a world leading supplier of both proprietary architecture and commodity component based HPC systems. Intel is the major supplier of processors for HPC systems, while NVIDIA GPUs are the most popular compute accelerator. The other members of the consortium are leaders in their own fields: DDN in HPC storage systems, Mellanox with the high performance InfiniBand interconnect, ScaleMP's middleware enables a cluster to be viewed as a large shared memory system and Platform Computing (part of IBM) is the industry leader in cluster management and workload scheduling tools. OCF provides high performance, management and storage solutions, but their main reason for being part of the Hartree team is its enCORE HPC on demand service.

The STFC Campus Centre provides a platform to build partnerships with both academia and industry. In order to bring HPC capabilities to a wider constituency the barriers to entry for the use of HPC systems need to be lowered, but this is at a time when – through heterogeneity and massive parallelism – HPC systems are becoming more difficult to use. The mission of the Hartree Centre is to help industry add value to the UK economy through the exploitation of High Performance Computing. The purpose of its many collaborations is to combine science and computer science expertise, to bring communities together and to deliver value to UK industry through accelerating the exploitation of HPC. The centre is focussed on nine sectors; life sciences, engineering, materials and chemistry, environment, nuclear science, power, data analytics, small and medium sized enterprises and government.

The Hartree Centre gives access to HPC platforms, access to IBM, and access to the many computational scientists working across STFC. The partnership with IBM brings 5,000 man days effort from experts in technical areas or business development to support projects at the

the Exascale Report™



The voice of the emerging exascale community

Hartree centre. The provision of on-ramps and easy to use interfaces can make the use of HPC facilities more accessible to scientists and engineers who are domain experts, but not HPC gurus. Working towards the development of the next generation of HPC software, the Hartree Centre works on proof of concept ideas, hosts much of this work on Hartree hardware facilities, and provides the initial capability for the enCORE Service.

Complementing the leading edge HPC hardware facilities available are the skills and experience of a 150 strong scientific computing team.

The services offered by Hartree are focussed in five distinct areas.

1. Software development: HPC systems are going through major transitions, driven by the need to reduce power consumption for the next generation of systems by two orders of magnitude. These changes will result in radical changes to the way that HPC systems are programmed.

2. Applications and optimisation: Hartree develops new approaches, software tools and algorithms to tackle these issues, as well as working to optimise existing applications and find better ways of handling big data. One aspect of its work in this area is its collaboration on the European exascale Software Initiative.

3. HPC on-demand: see section on encore

4. Collaboration: cross-fertilisation of skills and experience is an important component of the approach required to address the challenges of the next generation of HPC systems. STFC's scientific computing team, augmented by the abilities of Hartree's industrial partners, brings a unique set of skills to the table, enabling Hartree clients to quickly test new concepts and exploit the latest HPC know-how to enhance their products.

5. Training and education: Using HPC systems can be difficult, and they are becoming more complex, so training and education is ever more important. Hartree offers courses, workshops and a range of formal and informal media to educate scientists and engineers in both academia and industry.

A second tranche of funding has been provided by the UK government to explore issues relating to energy efficient computing. While it has not yet been decided how this money will be spent, it is likely that low power consuming compute accelerators will be on the agenda, including the Intel Xeon Phi, the latest Kepler GPUs from NVIDIA and Field Programmable Gate Arrays (FPGAs), an option that is very flexible, but also very difficult to program. The entire machine room will be instrumented allowing a deep understanding of power consumption for HPC systems and applications to be developed.

Systems

The Hartree Centre has two major compute facilities, and also offers access to other systems. Their flagship system is Blue Joule, a 6 rack IBM BlueGene/Q system that has 7,168 16 core Power BQC processors, uses a custom interconnect, and – delivering more than 1.2 Petaflop/s on the Linpack benchmark used by the TOP500 list – it is ranked number 16 in the world. Blue Wonder is an IBM iDataPlex that uses 1,024 8 core Intel Xeon E5-2670 processors connected by FDR InfiniBand. It delivers over 150 Teraflop/s and is ranked number 158 on the TOP500 list. Data is stored in 6 PB of disc and 15 PB of tape using IBM's GPFS General Parallel File System.

The Emerald system, hosted at STFC's Rutherford Appleton Laboratory in Oxfordshire and managed by STFC is the fastest GPU accelerated system in the UK. It is shared between the STFC and the universities of the e-Infrastructure South consortium (Bristol, Oxford, Southampton and UCL), Oxford university being a CUDA Centre of Excellence, sponsored by NVIDIA. It is based on HP SL390 nodes using 1,160 6 core Xeon E5649

the Exascale Report™



The voice of the emerging exascale community

processors, accelerated by 372 NVIDIA 2090 GPUs and, delivering 114 Teraflop/s, and is ranked 242 on the TOP500 list.

Longer term the Hartree systems will be used to support two major scientific projects. The UK Met Office and the Natural Environment Research Council (NERC) will use the Hartree systems to develop a highly scalable, next generation weather model that can efficiently use hundreds of thousands of processor cores, while vast amounts of data from the multinational Square Kilometre Array (SKA) - the world's largest radio telescope - will also be analysed at Hartree. One of the most data intensive experiments ever undertaken is the Large Hadron Collider (LHC) at CERN used in the search for the Higgs boson. It has been estimated that the SKA will generate more data every hour than the LHC generated in a year.

The OCF enCORE Service

OCF collaborates with STFC to sell HPC as a service. The main benefits of this service are that it provides an agile front end and delivery mechanism for HPC, and takes much of the pain and complexity away from the end users, enabling them to focus on solving their business problems without the need to become HPC experts. The enCORE Service can deliver a range of capabilities, but the initial service offered is based on the systems at STFC. Clients start working with the Hartree Centre on a proof of concept, exploiting the HPC and wide domain skills that the centre and its partners can deliver – before migrating to the enCORE Service for production runs. The cross fertilisation across the supply chain delivers real value.

There are many benefits of using an on-demand HPC service. A company can get access to the capability they need, when they need it, without having that expensive HPC resource sitting idle much of the time just so it is available to meet the peak demand. There is also the benefit – which is especially true for small companies – that they can outsource the tasks of buying, running, maintaining and supporting a large HPC resource. This

allows their staff to focus on their real business issues rather than becoming specialist HPC support staff.

One of the approaches that bring ease of use and flexibility to many areas of computer use is Cloud. However, Cloud is not always a simple solution for HPC. High performance is delivered for HPC applications by understanding the target architecture, and optimising the code to best exploit it. While Cloud is often about hiding the details of the target system from the user through the use of virtualisation. This conflict has stopped many HPC users targeting Cloud systems, but these two issues can be reconciled by wrapping a well understood HPC target system in a Cloud business model. It may not offer all of the flexibility of vanilla flavoured Clouds (e.g. you may need advance reservation for a large number of well connected nodes), but with careful design Cloud-like capabilities such as enCORE can deliver many of the benefits of the Cloud model as well as all of the benefits of dedicated HPC systems. An additional problem for running HPC applications in the Cloud is that of ISV application licensing. Some ISVs are enlightened and allow you to export your software licenses to the Cloud, but others are stuck in the past and only license their application on dedicated, named systems, which slows the migration to a more flexible working environment.

enCORE users pay a small annual subscription fee, and then the cost per CPU or GPU hour consumed. Unlike standard Cloud, compute nodes are not shared with other users. This is good as other applications will not interfere with the performance of your code, but it is bad in that it means that the cost of a node cannot be shared across multiple users. In order to get the best value for money users need to ensure that their codes are parallelised efficiently so they can exploit the full resource that is being paid for. Data transfer is handled by a secure web interface, or a secure shuttle service for very large data sets. Users access the enCORE facility through job queues on a web portal, there are separate job queues for workloads requiring CPU or GPU resources. Jobs are prioritised on a "fair share" basis, but reservations can be made for jobs that need a large number of dedicated nodes. enCORE makes a number of important application available to users at no additional cost, including the openFOAM CFD toolbox, the Code-Saturne open-source

the Exascale Report™



The voice of the emerging exascale community

CFD package, the TELEMATC modelling tool for free-surface flows and Ansys Fluent, one of the leading packages used for physical modelling to solve problems in areas such as aerodynamics, combustion, oil exploration and semiconductor manufacturing.

Conclusion

If UK industry is to make best use of HPC tools and technologies two issues must be addressed. One is flexible access to appropriate facilities at reasonable cost, while the other is the training, education, support and mentoring required to enable mere mortals (i.e. those that are not HPC experts) to make effective use of such HPC facilities. Together, STFC's Hartree centre and OCF's enCORE service address these two issues, adding real value to UK industry.

What do you think? We'd love to hear from you.

© 2013, The Exascale Report™

The Exascale Report

The Voice of the Emerging Exascale Community

Interested in having one of our consultants or editors host or moderate a panel at your event?

Contact us at info@TheExascaleReport.com



John Barr is a widely respected independent consultant and a contributing writer to The Exascale Report.

John can be reached at:
Johnbarr.work@hotmail.com
+43 (0) 664 35 42 667