University of Hull
High Performance Computing Newsletter
Spring 2018 Edition 3

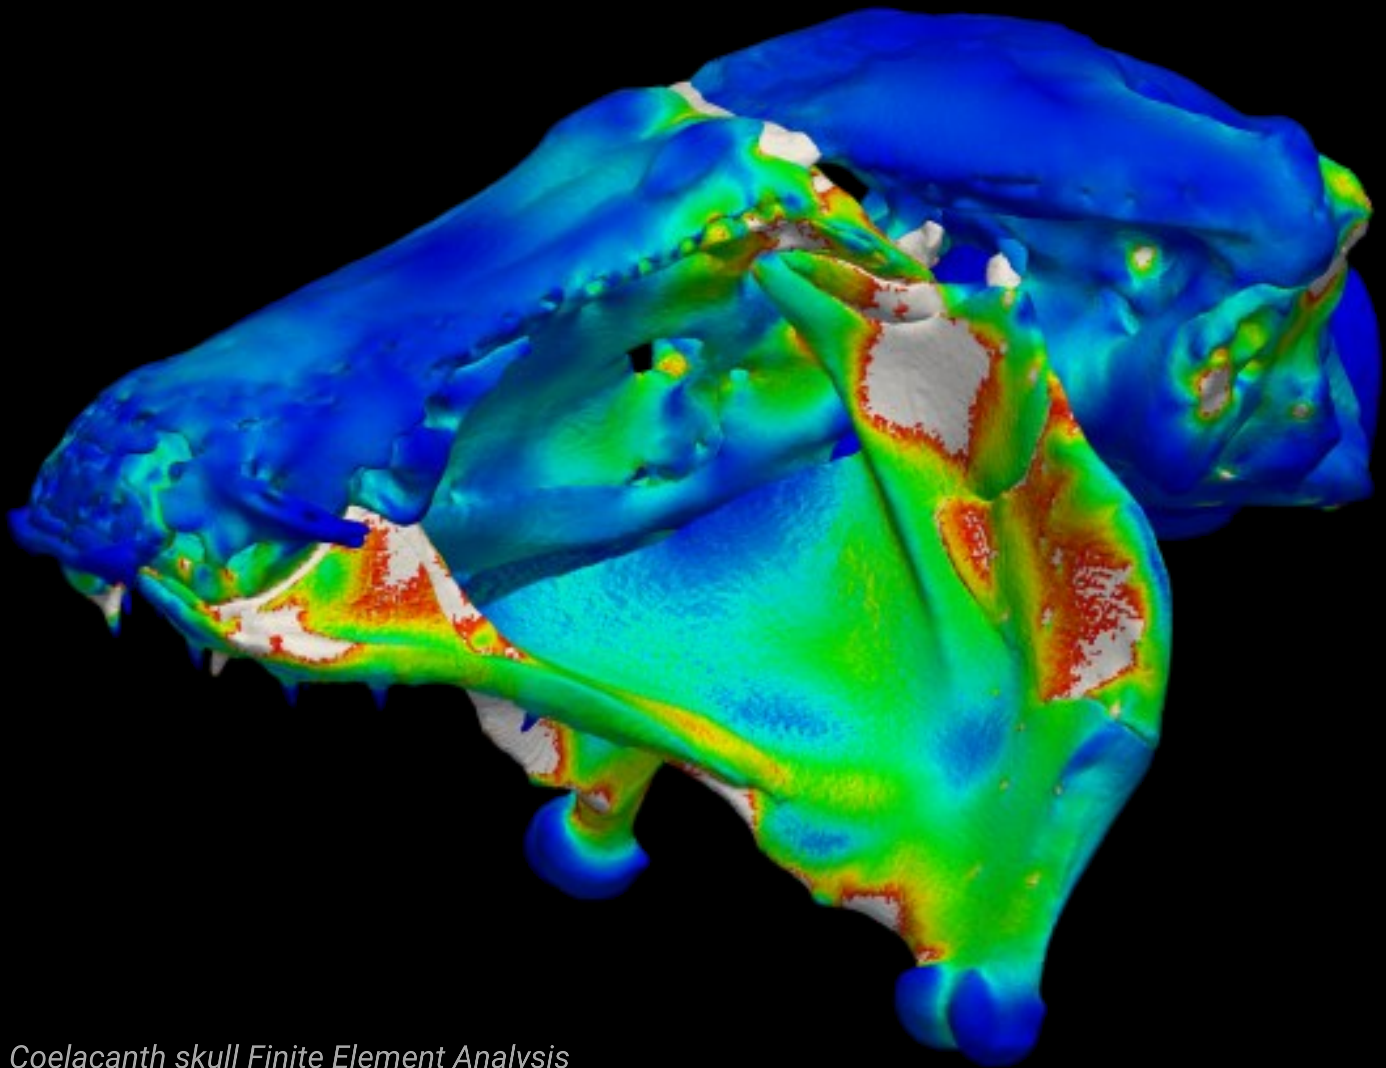

Coelacanth skull Finite Element Analysis 1000με by Dr Hugo Dutel

In this edition:

Looking Ahead to HPC-2

With Viper having established itself as a key tool for many researchers across the university, we are already looking at the next HPC to meet the growing needs of the research community. Chris Downing of Red Oak Consulting has been back at the University looking into this and explains the requirements gathering work being carried out

It is almost three years since work started on the process to build a HPC service in Hull. Back in April 2015, myself and colleagues from Red Oak had our first introductions to the researchers who were keen to have access to a local scientific computing platform. The process of gathering information on what people actually required and then buying it took just over a year; in that period, the university primary data centre had to be completely re-organised to make room for Viper (though we didn't have a name for it yet), and new infrastructure was added to the roof of the building to extract the significant extra heat that would be produced. By January 2016, the University had decided on a supplier and made a good start on putting together the team that would run the service. After an extensive period of testing in May and June, the service went live for all users without any major hiccups, and it has grown in usage ever since.

Viper has been active for more than 18 months now - for the cutting-edge hardware used in HPC, that represents about half of its "prime" lifespan. Because of that, it's now time to start thinking about HPC-2 and what might be required the next time round. In December, we started the process of talking to Viper users, the HPC steering group, and other university staff who have an interest in aspects of computational research such as business engagement and data preservation.

Viper in its home in Applied Science 3.
HPC-2 will be bigger and better

The process of gathering input from the wider community is now drawing to a close, and so we've produced a report setting out what the university will need to consider when buying the next system, and some options for how to proceed. That information should make its way to you via the HPC team and the steering group soon.

As you may have already heard from the HPC team or as part of one of the requirements capture meetings, the plan is to have HPC-2 ready to go in the second half of 2019. We are certainly expecting the system to be bigger.

Although HPC-2 will become the main research system by the end of 2019, Viper is unlikely to go away entirely. The transition will mean keeping both systems online for a while in order for people to move their data, after which Viper may be re-purposed for less intensive workloads such as use as part of teaching activities.

Although the university now has experience with buying a HPC system, there is still likely to be a significant set of hurdles to deal with this time. Firstly, the plan to increase the scale of the system when moving to HPC-2 means that putting it in AS3 isn't ideal. Fortunately, the university has a project underway to build a new data centre, which will house HPC-2 and potentially future systems as well. Additionally, those of you who follow the technology market may have noticed that things are a lot more "interesting" now than they were in 2015/16; back then, Intel were the only CPU supplier who were really competitive in HPC. Now, AMD and various manufacturers of ARM hardware have stepped up their presence in the industry. While these complications will make it harder to design the system, they should also improve competition and allow the university to get even better value for money.

Viper has provided a massive productivity boost to some researchers in Hull, and has been completely transformational to others. Over the next 18 months, we will all be working hard to make sure the next system is just as impressive!

In February, Dr Chris Downing returned to Hull to report back to the HPC User Forum on the findings of the requirements gathering exercise. The forum was attended by PVC for Research and Enterprise Dr Dave Richards, members of the HPC Steering Group, HPC users and some of those involved in the requirements gathering exercise not directly using HPC at this time. The forum, chaired by Dr David Benoit, also featured feedback on the Deep Learning Winter School hosted by the University in January, a discussion led by Dr Richards on the expectations and ambitions for HPC within the context of the University's research strategy, and a demonstration of Worktribe to show how users can tag research projects and outputs as utilising HPC.

In his session, Chris Downing summarised the output of the exercise, looking at what is likely to be required, what can be done to help make users more effective and what to do with Viper, once it is no longer needed.

Review

- Viper has been a great success for the university
- Compute utilisation has steadily grown, and the machine is now regularly at full capacity
- Storage capacity is steadily being consumed, but not to the point where quotas have been needed
- Other components (visualisation, GPUs (Graphics Processing Units), high-memory nodes) are seeing good, steady levels of use
- There is therefore a good case to be made for the university increasing its investment for HPC-2

Summary of requirements

- Expect HPC-2 to have more cores and more storage, but be broadly familiar
- No expectation of a wholesale shift to POWER or ARM CPU (too much risk porting applications), so x86 CPUs will form the majority of the system – could be Intel or AMD, with no real preference
- Accelerators will almost certainly be more GPUs – most effective solution will be to have combined high-density GPU high-memory systems along with low-density GPU hybrid standard nodes
- Making best use of the GPUs will require some software engineering – Viper offered a big boost over previous compute resources, but unlikely to get the same gain again without some work as CPU performance gains are running out of steam
- Separate /home from /scratch - /home will be small, and backed up as part of the HPC service
- Scratch requirements likely to increase somewhat – estimating 1PB storage system being the best choice
- Software will remain largely unchanged – part of the tender requirements will be for the supplier to ensure compatibility and the same feature set
- Plan to have support for containerised applications (Singularity) to improve reproducibility and ease-of-use

External and Teaching Use of HPC

- There is growing interest in having collaborators and external partners use the HPC system in the future. It is our view that HPC-2 should remain primarily a tool for university researchers, with commercial use fitting in as part of collaborative projects (as is the case already)
- Like Viper is currently, HPC-2 will be a research-focused system. Viper will be 3-4 years old when HPC-2 goes live – no longer fit for high-end research, but could be re-purposed for teaching if there is sufficient demand. After HPC-2 goes live, research will no longer be supported on Viper

PhD Research Case Studies

About half of the HPC users at the University of Hull are PhD students, with Viper being a critical part of their research. In 2018 the University of Hull is investing in 50 PhD studentships, many of which will utilise the University's High Performance Computing facilities. These include PhD clusters in Astrophysical Data Science; High Throughput Approaches to Decipher the Role of Human Endothelial Cells in Chronic Disease; Offshore Wind Energy Operations and Maintenance, and Molecular Stress in Changing Aquatic Environments.

Below we hear from PhD students currently using Viper as part of their studies, about their research and what impact Viper has had on their work.

Conducting Wave Energy Conversion Simulations Using Viper

Siya Jin
School of Engineering
Supervisor Professor Ron Patton

The potential wave power in global open sea has been estimated to be approximate 10^{13} W, which is comparable to the world's current power consumption. Therefore, massive efforts are being made to optimise a wave energy converter (WEC) device for effective wave energy extraction.

To capture wave energy, various WEC devices have been proposed. Among these, Heaving Point Absorber Wave Energy Converter (PAWEC) is the most promising one. Siya Jin [1] on opposite page shows the experimental platform of a 1/50 scale cylindrical PAWEC system built in the Hull University wave tank. The PAWEC obtains kinetic energy by interacting with the wave. Then it drives the connected power take-off (PTO) system to generate electricity. Thus, to capture wave energy, it is of fundamental importance to study the interaction performances among the PAWEC, the wave and the PTO system.

A series of tests has been conducted on the designed 1/50 scale PAWEC system. The tests demonstrate useful PAWEC hydrodynamics under varying wave conditions. However, due to the existence of wave reflection, while involving PTO system, it is difficult to achieve stable results. Moreover, it is interesting to understand the structure optimisation on the PAWEC performance maximisation.

Therefore, computational fluid dynamic (CFD) method (used software package, ANSYS/LS-DYNA) is applied to simulate PAWEC system. Two topics are focused (1) validation of the developed CFD model by experimental data, (2) study on the structure and PTO optimisation for power efficiency maximisation.

To reproduce the physical wave tank, a numerical wave tank is 13 m in length, 1.5 m in width and 0.9 m both in water and air is built (see 3D NWT built in ANSYS/LS-DYNA). It takes a long computation time in my own PC for running jobs. Therefore, Viper is applied, which reduces the computation time from three days to one and a half days running with one core. Presently, our lab has purchased a license for another 4 cores for LS-DYNA. In the future, we would like to use Viper to run jobs with cores in parallel. This will improve the computation efficiency drastically. The work related to the CFD model validation and PAWEC optimisation are being prepared for journal papers and will be published soon.

Direct Reaction of Anion-Deficient Solids with CO_2

Matthew J. Bennett
School of Mathematics and Physical Science-Chemistry
Supervisors Dr David Benoit and Dr Grazia Francesconi

Rising global temperatures are occurring as a result of man-made climate change from increasing carbon dioxide (CO_2) emissions. One of the ways to reduce the amount of CO_2 being released is through the utilisation of this greenhouse gas for preparing new materials. My PhD

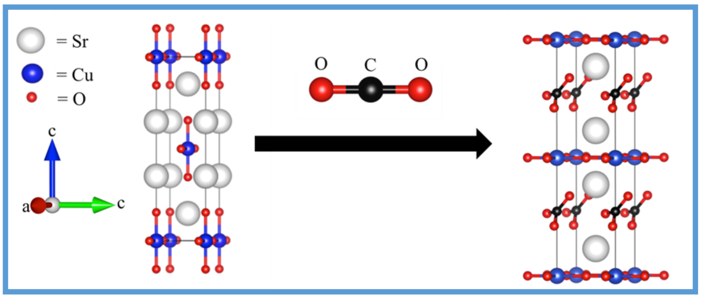

Matthew J. Bennett: Direct Reaction of Anion-Deficient Solids with CO_2: Sr_2CuO_3 (left) reacting with CO_2 to prepare $Sr_2CuO_2(CO_3)$ (right). This reaction has been studied using DFT methods within Viper.

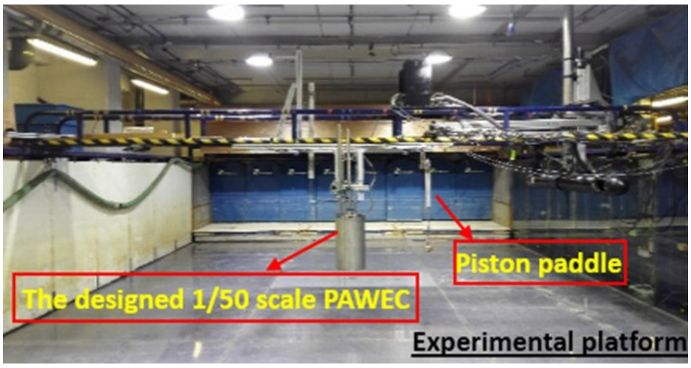

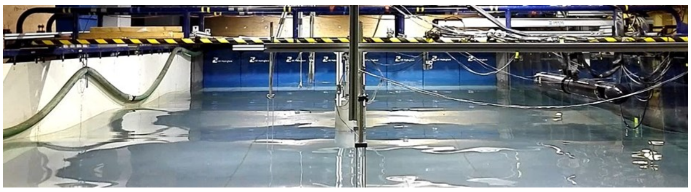

Siya Jin [1] - Physical scenario in the laboratory. The cylindrical buoy installed in the centred area is the experimental 1/50 scale PAWEC device

Siya Jin [2] - 3D view of the NWT built in ANSYS/LS-DYNA. Blue, pink, brown, black and green parts are water, air, sloped beach, PTO system and optimised PAWEC device, respectively

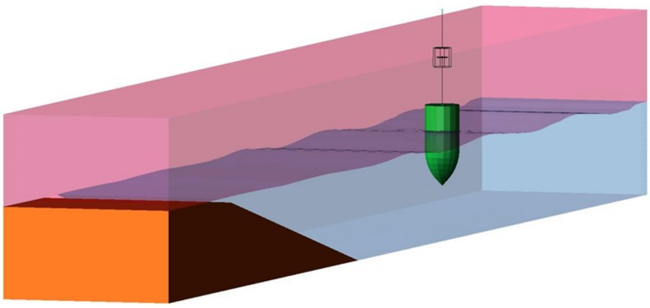

project is involved in studying the reaction of compounds with CO_2 for preparing new materials, using a combination of experimental and theoretical chemistry. This is not only to find new ways of making materials with CO_2, but also to try and understand the chemical reaction process-taking place. In my experiments I look at directly reacting CO_2 with a variety of mixed-metal oxide based compounds, A_2BO_3 (A = Sr, Ca and Ba and B = Cu and Pd), to prepare oxide carbonate $A_2BO_2(CO_3)$ materials. My theoretical work then utilises density functional theory (DFT) methods for understanding these chemical reaction processes. This is to try and explain how CO_2 reacts within these solids and form a carbonate (CO_3) group.

The use of Viper allows simultaneous calculations of CO_2 with solid materials to be carried out in real time. What would have taken hours to days using conventional computer systems has now been reduced to minutes. This level of high speed performance is essential in providing an in-depth study of a chemical reaction.

Through theoretical methods, a thermodynamic study of the direct reaction between CO_2 and the mixed-metal oxide Sr_2CuO_3 for preparing $Sr_2CuO_2(CO_3)$ (see figure to the left) has been carried out. This information is essential for providing a possible reaction pathway that would be difficult to study via a purely experimental approach. At present the results from these calculations are being prepared for publication alongside the experimental values.

Molecular diet analysis to investigate the impacts of invasive non-native species on invertebrate communities in the UK

Marco Benucci
School of Environmental Sciences
Supervisor Dr Lori Lawson-Handley

Invasive Non-Native Species (INNS) are a massive issue worldwide. Species that get established outside of their native range end up causing drastic damages to infrastructures, biodiversity, and in some cases posing a risk to human health. The detection of damages to biodiversity, in particular those through active predation, can be challenging especially in invertebrate communities due to their elusiveness. For these cases we can now use DNA analysis and modern High-Throughput Sequencing, that allow us to generate millions of DNA strands extracted from the gut contents of individual animals to detect potential preys.

In our project we are focusing on the impacts that two invasive species in the UK are causing to the native invertebrate communities through active predation using exactly this DNA technology. In order to go from DNA to a species name we need to match each DNA we obtain against worldwide DNA databases so that we can identify each prey. This is not an easy task when we have millions of DNA reads from 100s of samples. For this intensive task we use the HPC to automatically clean our DNA information, sort them in a more manageable format, cluster them together to identify communalities and then in the end match them against the DNA database that we use as reference. All of these processes can take up to a few days per step and it can be very computational heavy; however the HPC allow us to do that efficiently, accurately and in a way that we can replicate our methods over and over again.

Currently our achievements spanned broadly from gut contents analysis of 300 individuals to description of the distribution of species into 4 water bodies, data that we can match with our data from the gut contents to cross verify presence and absence of species.

Processing and Visualisation of Graphic-Intensive Synchrotron X-ray Tomography Data

Billy Koe

School of Engineering & Computer Science and Diamond Light Source
Supervisors Prof Jiawei Mi (University of Hull and Dr Thomas Connolley (Diamond Light Source)

My research involves the studies of intermetallic particles in metallic alloys. These particles, which form when molten alloy freeze during casting, can significantly affect the properties of the cast products and anything made from the castings during subsequent processing (e.g. rolling and forging).

Typically, intermetallics have length scales from a few hundreds of micrometres to a few tens of nanometres. The understanding of how these intermetallics form is crucial in current material processing technologies. I am working with colleagues at Hull and collaborators from other research institutions on experimental studies of molten metal processing, with the aim of improving materials properties. Various molten metal processing techniques (e.g. ultrasonic treatment and applying electromagnetic fields) are being investigated as methods of refining or removing these intermetallics to enhance the properties of alloys. We use advanced synchrotron X-ray facilities in the UK and abroad to study in real-time the solidification and semi-solid processing of these alloys, using high-speed imaging and micro/nano-tomography techniques.

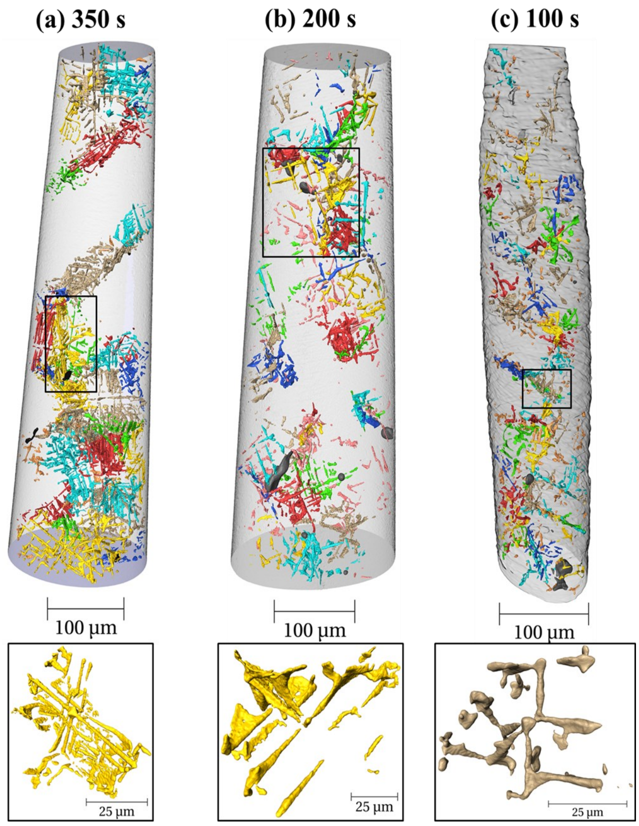

Billy Koe: High resolution 3D rendering of the intermetallics (metal carbides) of a widely-used INCONEL 713LC Nickel superalloy under different solidification time.

During the past two years of my PhD research project, I have been fortunate enough to be given multiple opportunities to carry out very advanced experiments at these facilities, which include Diamond Light Source (UK), BESSY (Germany), Swiss Light Source (Switzerland), and European Synchrotron Radiation Facility (France). The accumulated real-time datasets collected have surpassed the 10 TB mark, and this tremendous amount of data can easily overwhelm the current capabilities of a workstation due to the limited graphics, memory, and storage capabilities. These real-time datasets are able to provide novel information to further understand the underlying mechanism of how a specific alloy is formed. My usage on Viper is slightly different in nature compared to other users. I heavily rely on the visualisation aspect of Viper to study these imaging and tomography datasets. The datasets, in particular the 3D tomography scans, are extremely graphic-intensive and require a powerful graphics performance for smooth visualisation.

The figure illustrates the high resolution 3D rendering on Viper of the intermetallics (metal carbides) of a widely used INCONEL 713LC Nickel superalloy under different solidification time. With Viper, these hurdles were overcome without the need to down-sample the datasets, maintaining its high fidelity. In addition, these computing capabilities now allow us to perform full statistical analyses revealing much richer information from these datasets. This would not have been possible without the use of Viper's HPC services. We currently have manuscripts in preparation for submission to the top material science journals.

The Effect of Progenitor Metallicity on Nucleosynthesis in Multidimensional Type Ia Supernova Models

James Keegans
PhD
E.A. Milne Center for Astrophysics
Supervisor Dr Marco Pignatari

My research centers around providing detailed isotopic and elemental yields from a variety of astrophysical conditions. The main focus of my project is on Type Ia and core collapse supernovae, hugely powerful explosions which contribute significant material to their environments, although the suite of tools I use to investigate these can also be used in any number of astrophysical environments.

Collaborators around the world perform hydrodynamical simulations of different environments - for example tracking the shockwave which propagates through an exploding Type Ia supernova. From these simulations, temperatures and densities are extracted and used as the conditions under which nuclear reactions occur. You may well ask, why don't these hydrodynamics models also follow nuclear reactions? The answer is a practical one - computing costs and timescales would be far too large to make this method viable. Post processing allows these problems to become tractable, and for us to provide other areas of astrophysical and nuclear research with the data they need to inform further investigation.

My work has focused on making these investigations possible, by helping to test and implement a new version of the post processing code to handle the huge numbers of particles necessary to follow the nuclear reactions happening in a supernova explosion. Currently, 100,000 different trajectories are being investigated as part of the 2 dimensional Type Ia supernova models. Once development of the code is completed we can increase this number to millions or tens of millions of particles, allowing us to investigate models in 3 dimensions. These models are the next step in providing elemental and isotopic yields to the galactic chemical evolution community, yields which have been in high demand for over a decade.

Without access to a HPC facility, and VIPER in particular, this work would simply not be possible. The large number of particles to process, the time involved and the data storage requirements of this and related projects would simply be too great. Without HPC, these problems in nuclear astrophysics would remain unanswered, and other areas of physics research which rely on this information - experimental nuclear astrophysics for example - would also suffer.

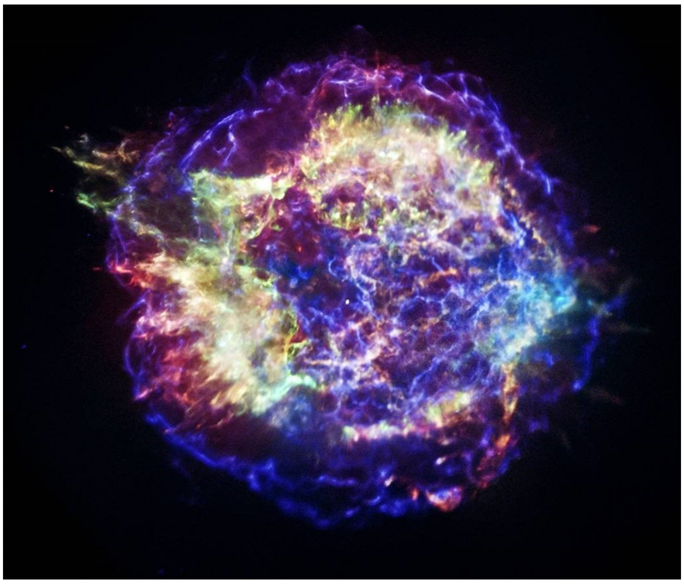

James Keegans: Cassiopeia A, a supernova remnant. In order to accurately model the isotopic abundances from events such as this, it is necessary to consider the 3D nature of the supernova events.
Image credit: NASA/CXC/SAO

If you are using Viper for your PhD and would like to contribute a short case study of your work for future newsletters, we would love to hear from you.

If you are interested in finding out more about how Viper could support your PhD research, please contact viper@hull.ac.uk

HPC Symposium

In September, the University of Hull held a one day HPC Symposium to celebrate research carried out on Viper by researchers, and showcase the most exciting research in HPC at the minute. Graeme Murphy, Head of Research and Enterprise ICT Services reports on the day

On the 8th September 2017 the University hosted its first HPC symposium by way of celebrating the first year of successful operations for its own institutional HPC, Viper. The event was very well attended by both internal and external guests and speakers with presentations covering a range of subjects and disciplines.

Dr Nina Dethlefs from the University's Department of Computer Science opened proceedings by describing how her interest in machine learning came about and how she is using it to identify patterns and relationships in large datasets.

Nina was followed in quick succession by Dr Michel Steuwer (University of Glasgow) and Dr Tim Bellerby (University of Hull Geography department) who both talked about changing the way we approach programming on a HPC. Michel discussed 'primitives' and 're-write rules' that he was using to improve performance and portability whilst Tim explained how his own PM language makes programming easier by managing some of the parallel aspects of coding.

Following a short break the fourth session of the morning saw the first of the industry contributions with ANSYS describing how using a HPC enables product designers to meet regulatory, environmental, quality and efficiency targets and how this can get new products to market faster. ANSYS explained that using a HPC enables product developers to make better design decisions because they are able to simulate larger and more complex models which then give them earlier insights into issues and problems which in turn allow for earlier interventions.

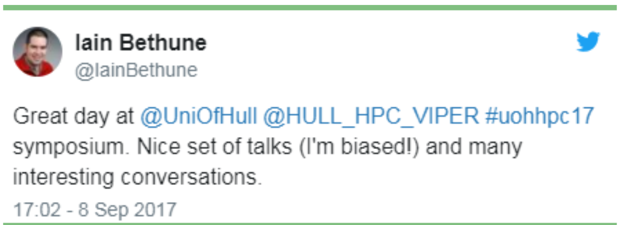

Iain Bethune from the Hartree Centre was next up. Iain described the role of the Hartree Centre in supporting industrial research and talked about how they are looking to get 'HPCs to work smart not hard'. Iain revealed how manufacturers use HPCs to help them decide why particular ingredients work well together in, for example, bath soaks. Iain also caused a certain amount of amusement when he talked of some of their new technologies and referred to an oil-cooled HPC as 'The Deep Fat Fryer'; little realising that the company who installed and commissioned it were also the supplier of the University HPC Viper and were sat in the audience. Lunch probably came at the right time…

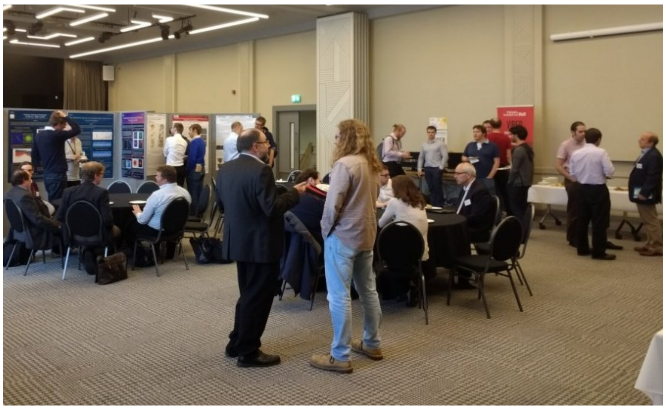

HPC Symposium lunch: food, research posters and lots of HPC debate

Over lunch posters highlighting the range of research carried out on Viper were on display. Posters covered subjects ranging from the 'Benefits to Positron Emission Tomography from Ultra fast Time-of-Flight detectors' to 'The evolutionary history of hermaphroditism in fish', with simulating cosmic structures and merging galaxy clusters, the use of HPC in processing and visualisation of Synchotron X-ray Tomography data, biomechanics and modelling the evolution of animal forms and studies of anharmonic vibrations of molecules in between. A competition for the best PhD poster would be judged later in the day.

Dr David Benoit from the University Chemistry department took to the stage after lunch and he gave an interesting talk on how difficult 'accurate' quantum chemistry is to do on a HPC and how scalability is restored by breaking down the problem into smaller units.

Lively debate around the Human Genome project then followed as Professor Peter Coveney from UCL talked about the 'Virtual Human', a personal 'digital doppelganger' that can help medics and clinicians target treatments based on the individual patients' physiology and needs.

Professor Coveney was followed by our very own Dr Hugo Dutel who described the work taking place in Hull using finite element analysis to study the evolution and diversity of animal body forms. Hugo showed how 3D virtual models, multi-body dynamic analysis, in vivo data dissection and materials properties testing, alongside finite element analysis, is being used to analyse the stresses and strain levels being imposed on bone structures.

> **Dave Lunt**
> @davelunt
>
> Today is @UniOfHull High Performance Compute symposium- lots of impressive talks and really interesting approaches @HULL_HPC_VIPER
>
> 11:10 - 8 Sep 2017

Next up was Mark Wilkinson from DiRAC who talked about the work being done to understand how the Universe works and how particle physics can help us comprehend the fundamental laws of nature. Mark took us on a glorious journey from the birth of the Universe into the inside of molecules, atoms, protons and electrons – from the humongously huge to the infinitesimally small.

Finally it was left to Dr Gareth Few to continue the astronomical theme and to educate us all in the work of the E.A. Milne Centre and how they are looking back towards the birth of the Universe in order to expand our understanding of how stars and galaxies are born and evolve.

The event was followed by a drinks reception in Canham-Turner where the student poster competition was judged by members of the HPC Steering Group and Laura Redfern from the University's HPC partner ClusterVision. The eventual winner was Alex Sheardown, PhD student in the E. A. Milne Centre for Astrophysics, with a poster titled "Simulating the Largest Cosmic Structures: The Infall of N1404", taking away an NVidia Shield donated by ClusterVision. After the competition, Alex worked with ClusterVision on a case study to showcase his research.

> **Brad Gibson**
> @profbradgibson
>
> Congrats to @EAMilneCentre student Alex Sheardown awarded ClusterVision HPC Poster Prize! #uohhpc17 @UniOfHull @clustervhpc @HULL_HPC_VIPER
>
> 23:59 - 10 Sep 2017

The general consensus was that it had been an interesting and thought provoking day and that time had flown by – almost as if it were being driven by our own HPC, Viper, itself. Work is already underway to host a similar event earmarked for September 2018.

For more information on any of the speakers that took part in the 2017 HPC Symposium, access to the presentations themselves, information on accessing HPC resources at Hull or for more information on HPC Symposium 2018 please contact Chris Collins at chris.collins@hull.ac.uk.

Laura Redfern from ClusterVision and Viper's Chris Collins with student poster competition winner Alex Sheardown

ClusterVision case study of Alex Sheardown's research: Exploring the growth of galaxy clusters using HPC is available at http://clustervision.com/case-study-university-of-hull-viper/

HPC Update

High Performance Computing means pushing what can be done with the latest technologies. The HPC community includes hardware and software vendors, integrators and administrators. Here we look at how we keep up to date with the latest news and trends in High Performance Computing

Chris Collins, HPC Systems Manager attended Supercomputing '17

SC17 Denver, CO hpc connects.

November sees one of the highlights of the HPC calendar with the holding of the Supercomputing conference, the premier international conference showcasing work in HPC, networking, storage and analysis by the international HPC community. This year, over 12,000 people from across the globe descended on Denver, Colorado to participate in SC17, representing HPC users, researchers and professionals from across industry and academia.

During my week in Denver I had the chance to attend a vast array of sessions on the management and application of HPC resources. From workshops on application performance, panels on software sustainability and energy efficiency to excellent Birds-of-a-Feather sessions on HPC training and HPC outreach. The packed schedule was diverse but always interesting.

Researchers at the University of Wyoming and software company Intelligent Light showcased their work in simulating the Lillgrund Wind Plant, Sweden's largest offshore wind farm. This work represents the most highly resolved simulation of a wind farm to date, performed over more than 32,000 compute cores and covering the smallest scale (air interacting on the surface of the blade) to the large scale (turbulence between windmills).

Denver's 40ft tall big blue bear keeps an eye on the latest technology updates at SC17

On the exhibition floor there were great opportunities to talk with some of our existing HPC partners such as Intel, BeeGFS and Slurm, as well as new companies, projects and products who could play a part in future HPC services here at Hull. Stands of particular interest included software to improve remote visualisation, tools to allow better insight into the vast data involved in HPC and better metadata management and technology for cooling HPC data centre hardware, something that is always a consideration in HPC.

It was also interesting to see other academic institutes from across the globe exhibiting their work in HPC and carried out using HPC. SC17 also saw AMD and ARM joining IBM's Power 9 as creditable competitors to Intel in the CPU market for HPC, which is promising for the future of HPC.

The annual supercomputing conference also includes the November release of the latest Top500 list of the fastest supercomputers in the world. There was no significant changes at the top of the list since the June update, with Sunway TaihuLight at China's National Supercomputing Center still the top ranked system in the world. For the first time China overtook the US in terms of the number of systems on the Top500 list, with 40% of the list compared to 29% for the US. The UK was 6th in the list with 15 entries (3%), with the UK Meteorological Office Cray system remaining as the top UK system, 15th on the list.

The next conference, SC18, takes place in Dallas, Texas in November 2018.

Research Software Engineer Darren Bird reports from GTC

GPU TECHNOLOGY CONFERENCE

NVidia Roborace autonomous racing cars - deep learning touches every research area

In October 2017, we visited NVidia's showcase conference GTC 2017 (Europe) in Munich, Germany. GTC focuses on GPU technologies for developers, data scientists and senior decision makers. The event brings the most vital work in the computing industry today, including deep learning and AI, big data analytics, virtual reality, self-driving cars and a vast number of new areas where this technology is starting to be exploited.

NVidia have aligned themselves into the AI market completely from originally developing their range of GPU cards for the PC market to now producing cards that have thousands of streaming processors in one chip. This makes them ideal for running artificial neural networks upon them. Although the theory has been around since the 1950s the emergence of this type of technology is now making it possible to implement these networks. Our own HPC Viper has four separate accelerator units which have four of these cards per unit as well. At the moment these are used mainly for deep learning although their flexibility allows them to be used in a host of different areas too, such as big data analysis.

The conference defined eight significant areas of research they intended to follow. These are Deep Learning & AI, autonomous vehicles, supercomputing & HPC, VR & AR, pro-visualisation, healthcare, finance and autonomous machines. The attendees here ranged from company CTO, IT managers, Data Scientists, GPU developers and academics who presented a range of research from big data mining, classifiers, language comprehension, autonomous machines and vehicles.

University of Hull's Graeme Murphy takes on role of Vice Chair of UK High Performance Computing Special Interest Group

At its meeting in Bristol in February, the UK High Performance Computing Special Interest Group (HPC-SIG) elected University of Hull's Graeme Murphy, Head of Research and Enterprise ICT, as their new Deputy Chair.

The HPC-SIG draws its membership from UK Universities that have, or aspire to have, HPC facilities and aims to promote the use of High Performance Computing (HPC) in academia and industry whilst also providing support and resources to the community through a range of HPC focussed meetings, events, and workshops. Group members include Directors/Heads of Research Computing Services, HPC managers, systems administrators and research software engineers amongst others.

Hull hosted a two day HPC-SIG event in September 2017 and has won the respect of the UK HPC community with what it has achieved with Viper. Having Graeme Murphy hold the position he now does within this group clearly shows how far the University has come in so short a period and helps to maintain our progress towards our vision for HPC at Hull of providing access to technically advanced and innovative research computing systems that will be an enabler of research excellence at the University of Hull whilst also being acknowledged as a leader in the field of quality HPC services and support.

HPC Update

Dr David Benoit, Senior Lecturer in Physical Chemistry and member of the HPC Steering Group, had the opportunity to visit the Swiss National Computing Centre's Piz Daint facility, one of the world's top supercomputers. Here he tells us about his visit.

In September, I was lucky enough to visit one of my former workplaces, the CSCS (Swiss National Supercomputing Centre) home to the third largest supercomputer in the world - Piz Daint (number 3 on top500.org). This is one of the largest PRACE (Partnership for Advanced Computing in Europe) Tier-0 system in Europe.

Swiss National supercomputing centre building in Lugano. Offices on the left connected to the computer building on the right by a footbridge.

I was curious to see what the 3rd largest supercomputer in the world looks like and also what type of infrastructure and maintenance was necessary for such a large facility. The CSCS has a split design that spans two buildings: one office building and a separate windowless concrete monolith for that houses the supercomputers (yes, they have more than one) and also hosts resources for most Swiss universities. Both buildings are linked by a covered foot bridge. The computer building is truly designed to be both energy efficient, secure and has space for expansion. Indeed, my first surprise was that Piz Daint is not as big as one would imagine, it only takes about an eighth of the space available.

It would be remiss not to mention the technological advances that are at the core of the CSCS building: first of all, instead of the traditional raised floor that we find in most computing centres to house cables, cooling pipes, etc, the CSCS has decided early on that they would reserve an entire floor (resource deck) to those, thus making both maintenance, installation and access much easier. Therefore, the computer floor is actually on the first floor, with a technical stage on the ground floor. The next novelty is that the CSCS cooling system uses water from nearly lake Lugano. In order to do this, the computing centre collaborated with the town of Lugano, where it is based, to lay 2.8 km of pipelines between the lake and the computer building. The pipes are buried 45 meters deep and provide a steady supply of water at 8-9˚C to the CSCSC at 460L/s. The remaining water capacity (300L/s) is used by the town's main utilities contractor. Once the water has cooled the supercomputers, it passes through a heat exchanger to feed a secondary air-based cooling circuit, used for file-storage appliances and computer housing. Overall this technique provides a cooling power of 21MW (14 MW primary and 7MW secondary). The used water is then piped back to the lake and mixed back to ensure its temperature does not exceed 25C.

Piz Daint is an impressive machine, despite its surprisingly small footprint of 3 rows of hybrid nodes and 1 row of standard CPU nodes. This is achieved by using high-density racks, similar to those used at Archer. The hardware was provided by Cray (XC40/XC50) as the CSCS has a 13 years working relationship with the supercomputing company. Cray has two hardware engineers based at CSCS and has twice altered their technology road-map to accommodate the needs of the CSCS. One of the main feature of Piz Daint is its very large number of hybrid nodes (5320 nodes) compared to the remaining 1431 "standard" nodes in the system. Hybrid nodes are equipped with 12-cores intel Haswell CPUs and Nvidia P100 cards, while the CPU nodes have 18-cores Intel Broadwell cpus (similar to those on VIPER). The peak performance (maximal LINPACK performance on hybrid nodes) is 19,590 TFlop/s which is equivalent to 114 times that of VIPER. According to Sadaf Alam – CSCS's Chief Architect and Head of HPC Operations – the Cray's Aries networking system is able to essentially remove any communication bottleneck between racks so that the entire machine operates like a huge 361,760-core system. Interestingly, the physical hybrid nodes are much longer than standard nodes and need to be extracted with a special trolley system or they would bend out of shape due to their own weight. The huge lustre storage system (nearly 9 PB) dedicated to the Piz Daint facility is located a few feet away from the Cray racks in a separate air-cooled island.

Large cooling pipes at lake Lugano's pumping station.

In terms of operations, the user base is quite broad: most of them are based at Switzerland's universities and research centres but there is also a large proportion of international users who access this resource through the PRACE scheme. Most users run highly parallel MPI tasks and the GPU usage on the machine is similar to that of the CPU usage. This implies that most codes are GPU-enabled and the CSCS has been working with users over the years to facilitate the move to hybrid systems. The CSCS also hosts other types of accelerators, such as Xeon Phi (KNL) cards and has a number of accelerated systems that can be accessed separately for testing or production. One of the main objective at CSCS is to develop extremely performant and cost-effective HPC but also to minimise environmental impact. This was one of the main driver for the development of their truly unique lake-based cooling system, for example. Thus, while delivering world-class performance, Piz Daint is also one of the greenest of its class (number 10 on the Green500 list).

Piz Daint - [Image creditCSCS]

University of Hull Deep Learning Winter School

Dr Nina Dethlefs reports on the 1st Deep Learning Winter School which took place at the University on the 23rd-24th of January 2018, organised jointly by Computer Science, Chemistry and the Viper HPC team

Funded by the Ferens Education Trust, our aim was to provide an introduction to the theory and practice of Deep Learning in a focused crash course of two days. Deep Learning is the main technique behind recent breakthroughs in AI including autonomous cars, machine translation and personal assistants such as Alexa or Siri. It is therefore increasingly important for our own research, collaborations with others, funding opportunities and is also an increasingly sought after skill for graduates both in academia and in industry.

The schedule combined lectures on Deep Learning with lab sessions, an invited talk and social events, such as lunches and a dinner. In particular, we looked at the details of the maths and theory behind Deep Learning and explored the most important libraries and frameworks in hands-on lab sessions training variants of neural networks on a handwritten digit recognition task. Dr Yannis Konstas from Heriot-Watt University told us about using Deep Learning for natural language generation. The second day allowed participants to apply their knowledge and work in small groups on projects that they chose themselves. We had four groups working on projects on song identification from lyrics, Titanic survival chances, flag identification from images and poker hand classification. The Winter School finished with a Viper tour for all interested participants.

We had 45 participants overall including seven PhD students from Hull, four sixth formers and the remainder undergraduate students from the University of Hull, including Computer Science (majority) but also Chemistry / Physics, and Business.

From the Desktop to Viper

Applications such as Matlab, Stata, R and COMSOL are common on the desktop, but users of these tools can also benefit from using Viper

There is a common misconception that HPC clusters such as Viper can only benefit either those programming specifically for the HPC, or those running large parallel models or simulations that run over hundreds of computers. While these are use cases seen on Viper, other examples of common workflows include many applications which are more than familiar on the desktop.

As an example, both **Matlab** and **Stata** can provide a familiar graphical user environment to that on a Windows or Mac desktop. This makes transitioning to using Viper a simpler experience and helps open up the benefits of using HPC for your research.

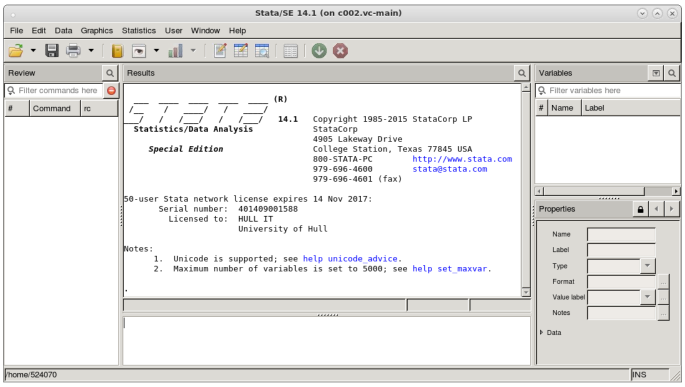

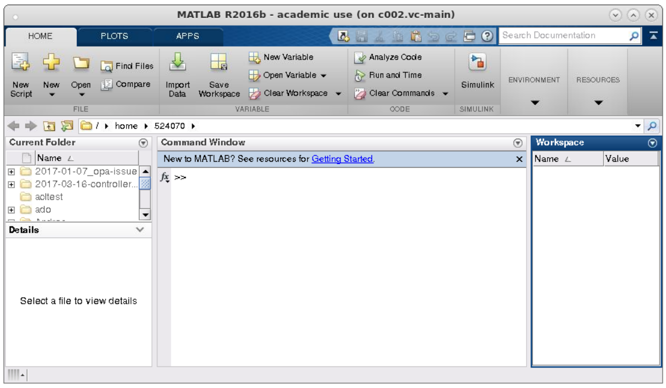

Matlab and Stata user interface running on Viper. A familiar working environment for those familiar with the desktop versions

Each of Viper's 180 standard compute nodes have 128GB of memory and 28 CPU cores, and with the addition of dedicated high memory and GPU systems this provides over 5,500 cores. With the University of Hull having a **Matlab** Total Academic Headcount license, even if individual tasks have modest resource requirements it is possible to run large numbers of jobs concurrently. Both **Matlab** and **Stata** can also see significant performance improvements with Viper with minimal changes to jobs taken from the desktop.

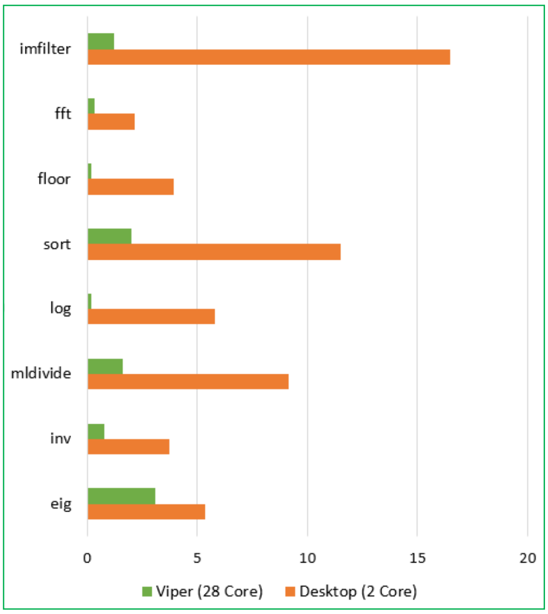

Matlab performance, in seconds, for certain standard functions on various sized inputs. Run times comparing processing on a 2 core desktop (orange) and a single 28 core Viper compute node (green)

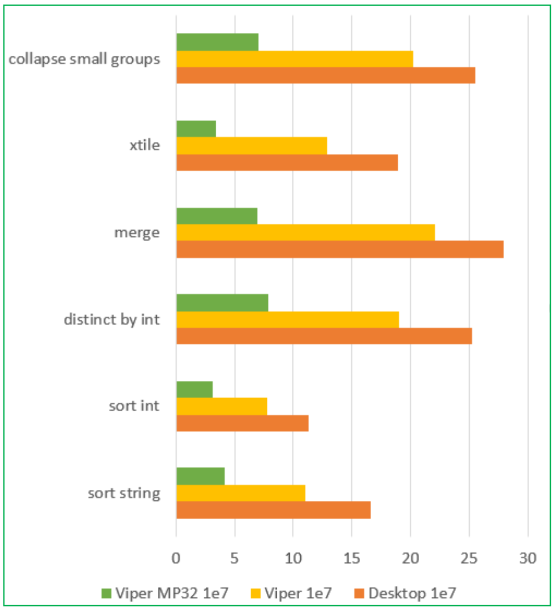

Stata performance, in seconds, for certain standard functions on various sized inputs. Comparison of Stata/SE single core on a desktop (orange) and a Viper compute node (yellow) and a trial version of Stata/MP32 running on 28 cores (green)

Other tools available on Viper that are commonly used on the desktop include **COMSOL** (depending on license availability), **Python** and **R**. All these applications can make use of the multiple cores, multiple nodes and increased memory available on Viper, often with just small updates to the original task.

Delivering Research Outputs

Viper is helping contribute to important research being undertaken at the University. Below are some of the outputs Viper has helped contribute to since the last HPC newsletter

Roggatz C. C., Lorch M., Benoit D. M., 2018-03-22, **Influence of Solvent Representation on Nuclear Shielding Calculations of Protonation States of Small Biological Molecules**, The Journal of Chemical Theory and Computation, doi: 10.1021/acs.jctc.7b01020

Zhao Y., Du W., Koe B., Connolley T., Irvine S., Allan P. K.. Schlepütz C. M., Zhang W., Wang F., Eskin D. G. and Mi J., 2018-03-15, **3D characterisation of the Fe-rich intermetallic phases in recycled Al alloys by synchrotron X-ray microtomography and skeletonisation**, Scripta Materialia, doi: 10.1016/j.scriptamat.2017.12.010

Wang B., Tan D., Lee, T. L., Khong, J. C., Wang F., Eskin D., Connolley T., Fezzaa K., and Mi J., 2018-02-01, **Ultrafast synchrotron X-ray imaging studies of microstructure fragmentation in solidification under ultrasound**, Acta Materialia, doi: 10.1016/j.actamat.2017.10.067

Gottschalk H. C. *et al* including Benoit D. M., 2018-01-01, **The furan microsolvation blind challenge for quantum chemical methods: First steps**, The Journal of Chemical Physics, doi: 10.1063/1.5009011

Vorotnikova, N. A. *et al* including Benoit, D. M., 2017-12-22, **23-Electron Octahedral Molybdenum Cluster Complex [{Mo6I8}Cl6]–**, Inorganic Chemistry, doi: 10.1021/acs.inorgchem.7b02760

Ruiz-Lara, T., Few, C. G., Florido, E., Gibson, B. K., Pérez, I., Sánchez-Blázquez, P., 2017-12-15, **The role of stellar radial motions in shaping galaxy surface brightness profiles**, Astronomy & Astrophysics, doi: 10.1051/0004-6361/201731485

Szitenberg A., Salazar-Jaramillo L., Blok V. C., Laetsch D. R., Joseph S., Williamson V. M., et al., 2017-09-25, **Comparative genomics of apomictic root-knot nematodes: hybridization, ploidy, and dynamic genome change.**, Genome Biology and Evolution, doi: 10.1093/gbe/evx201

Thompson, B. B. *et al* including Few, C. G. and Gibson, B. K., 2017-09-11, **The Gaia-ESO Survey: Matching Chemo-Dynamical Simulations to Observations of the Milky Way**, Monthly Notices of the Royal Astronomical Society, doi: 10.1093/mnras/stx2316

deBoer R. J., Görres J., Wiescher M. , Azuma R. E., Best A., Brune C. R., Fields C. E., Jones S., Pignatari, M., Sayre D., Smith K., Timmes F. X. and E. Uberseder, 2017-09-07, **The 12C(alpha,gamma)16O reaction and its implications for stellar helium burning**, Reviews of Modern Physics, doi: 10.1103/RevModPhys.89.035007

Hahn C., Genner M. J., Turner, G. F., Joyce, D. A., 2017-08-29, **The genomic basis of cichlid fish adaptation within the deepwater "twilight zone" of Lake Malawi**, Evolution Letters, doi: 10.1002/evl3.20

Dethlefs, N., 2017-07-18, **Domain Transfer for Deep Natural Language Generation from Abstract Meaning Representations**, IEEE Computational Intelligence Magazine, doi: 10.1109/MCI.2017.2708558

Liu N., Nittler L. R., Pignatari M., O'D. Alexander C. M., Wang J., 2017-06-06, **Stellar Origin of 15 N-rich Presolar SiC Grains of Type AB: Supernovae with Explosive Hydrogen Burning**, The Astrophysical Journal Letters, doi: 10.3847/2041-8213/aa74e5

A full list of outputs can be found at http://hpc.wordpress.hull.ac.uk/research-outputs/

We are always interested to know of any output that features work carried out on Viper, please contact viper@hull.ac.uk

Viper Support Services

Viper is supported by a dedicated support team, with a wide range of expertise and experience from Linux and HPC system administration to code development and software engineering.

As well as the day to day running of the HPC facility, we also provide the following services to the research community to help them make the most of the resources available including:

- Training on topics such as Linux, High Performance Computing and software development
- Support and advice on areas such as migrating tasks from the desktop to HPC and HPC workflows
- Software engineering resource to work with researchers and research groups to develop and optimise code to run on the HPC
- Advice and assistance on the inclusion of HPC in research proposals

We regularly hold open surgery sessions which give users the opportunity to talk directly to team members about possible workflows and HPC issues or even go through some informal training. Come along to the ICT Consultation Room by the IT Service Desk, **Wednesday 2-3pm**. The Viper team are happy to meet with any researchers interested in looking at how Viper may be used in their research, and can also organise group sessions or school seminars focussed on the usage of HPC to specific research topics.

Getting Information

The main HPC website is at http://www.hpc.hull.ac.uk and includes information about:

- The research carried out on Viper, including case studies, research areas and research outputs
- The support team and the HPC Steering Group
- The support FAQ
- How to apply for an account

Our documentation and technical support wiki can be found at http://www.hpc.hull.ac.uk/wiki which includes:

- General pages with information on Viper, HPC concepts and basic usage
- Training pages including guides on getting connected to Viper, working remotely, and introductions to Linux
- Application pages with detailed information on the hundreds of applications, tools and libraries available on Viper, including typical usage instructions

To contact the team, please use one of the following email addresses:

- For support enquiries, please raise a ticket via the University of Hull Service Desk by emailing help@hull.ac.uk
- For queries regarding the use of Viper, please email the team directly on viper@hull.ac.uk. Please feel free to contact the HPC Systems Manager, Chris Collins, directly on chris.collins@hull.ac.uk

@hull_hpc_viper   YouTube Viper HPC