Hardware & Enabling Software Testbeds

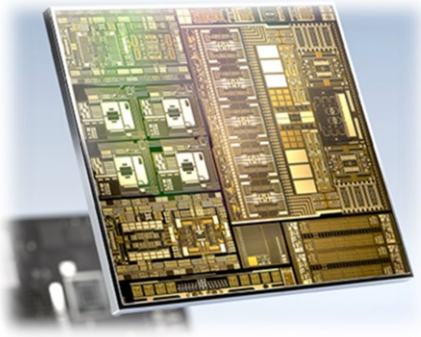

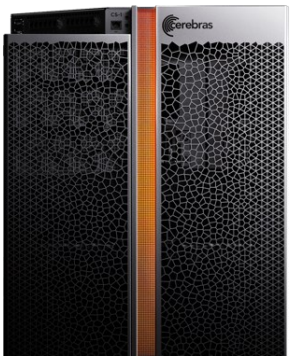

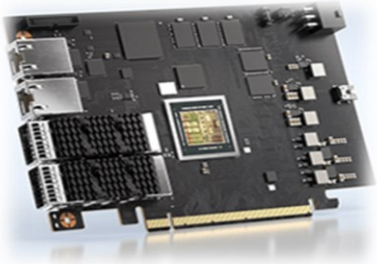

Contents

About ExCALIBUR Hardware & Enabling Software

ExCALIBUR (Exascale Computing Algorithms and Infrastructures Benefitting UK Research) is a £45.7m Strategic Priorities Fund (SPF) initiative led by the Met Office, UKAEA and UKRI that aims to deliver the next generation of high-performance simulation software for the highest-priority fields in UK research. ExCALIBUR runs from October 2019 until March 2025, redesigning high priority computer codes and algorithms to meet the demands of advancing technology and UK research

As part of ExCALIBUR, UKRI are investing £4.5m over four years to create and evaluate novel pre-Exascale Hardware and Enabling Software (H&ES) testbeds, in a co-design collaboration with technology partners and the UK research community. The H&ES programme provides grant funding for testbed projects that give UK researchers early access to novel hardware and software technologies that may play a part in future exascale systems and services.

H&ES testbeds are generally small-scale systems that are provided on an experimental basis, and clearly delineated from formal service provision to avoid interference with production workloads. Researchers can use the testbeds to validate the portability and performance of their codes, and to explore the potential of new paradigms such as workflow offload to data processing units or wafer scale compute engines.

This booklet provides brief descriptions of each of the testbeds that H&ES has supported, including completed projects, current projects (as of January 2023) and upcoming projects. In most cases the testbeds that projects have produced will continue to be available for the duration of the ExCALIBUR programme. Access arrangements vary by testbed, but testbed facilities are generally available to the whole of the UK science community. However please note that testbeds have a particular focus on the needs of ExCALIBUR projects, which will be prioritised.

For more information about the ExCALIBUR initiative and how you can engage with it, please see the programme level website: https://excalibur.ac.uk

If you have any queries for the ExCALIBUR H&ES team, please contact the H&ES Programme Office by email at: excalibur-hes@jiscmail.ac.uk

1. Adaptable Cluster Project

Status	Available / Under development
Access arrangements	Contact Dr Owain Kenway <o.kenway@ucl.ac.uk>
Organisations	University College London, ARM, NVIDIA, Lenovo
Project linkage	All ExCALIBUR projects are involved in the Benchmarking at Exascale element of the project

The Adaptable Cluster project created a testbed interconnect demonstrator consisting of two non-blocking interconnect fabrics supporting up to 60 attached nodes in a dual fabric configuration.

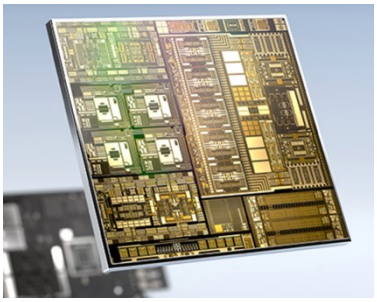

One fabric is 200 Gbps HDR Mellanox Infiniband configured so that it is possible to construct multi-hop routes between nodes. The second fabric is 100Gbps Mellanox Ethernet, with BlueField adaptors on each node. This allows us to measure the impacts of a variety of in-network technologies – doing computation at the switch level (requiring multiple hops) and looking at the possibility of using acceleration on the adaptor to off-load some of the work of the host machine (the BlueField cards). It also becomes possible to gauge the "state of the art" in using Ethernet as an Interconnect with Infiniband, to measure whether on RDMA on Converged Ethernet has reached the point where it is a performant, cost effective interconnect.

In order to understand system and application performance the Adaptable Cluster collects metrics from several sources in the system and dashboards to visualise them, which then allow focus on how to improve system design and resource usage. Alerts can be set up to draw attention to performance issues as well. The testbed uses components such as Elasticsearch, Kibana, Logstash and Prometheus to provide insights into both breadth and depth of system and application performance.

UCL is also the location of the ExCALIBUR instance of the ARM FORGE Application. This is an application that supports the debugging, profiling and optimisation of codes that use distributed resources, such as a cluster. It is both CPU and GPU enabled. UCL will support ARM FORGE for key centres in the ExCALIBUR project. It will also be available to UCL projects that are not associated with ExCALIBUR. This package enables jobs that use up to 2048 cores to be analysed in terms of code efficiency. One outcome of this project will be methodologies that enable results from Prometheus and ARM Forge to be used to improve system design, architecture performance and application performance.

2. AMD GPU testbed

Status	**Available**
Access arrangements	https://safe.epcc.ed.ac.uk/dirac (project do009) cosma-support@durham.ac.uk
Organisations	Durham University, AMD
Project linkage	ExaHype, MPPHEA, ExCALIBUR training programme

The Durham AMD GPU extends the existing AMD cluster at Durham, providing researchers with the opportunity to test their code on the AMD MI50 GPU. The testbed consists of a Gigabyte server with:

- 2 x AMD EPYC 7282 16 core 2.8GHz CPUs
- 1TB RAM
- 6 x AMD MI50 GPUs
- AMD software, ROCM, AOCC, AOMP, GCC with offload support installed

Basden, Alastair, Weinzierl, Marion, Weinzierl, Tobias, & Wylie, Brian J. N. (2021). *A novel performance analysis workshop series concept, developed at Durham University under the umbrella of the ExCALIBUR programme* (1.0). Zenodo. https://doi.org/10.5281/zenodo.5155503

Schulz, H., Gadeschi, G.B., Rudyy, O., Weinzierl, T. (2021). *Task Inefficiency Patterns for a Wave Equation Solver*. In: McIntosh-Smith, S., de Supinski, B.R., Klinkenberg, J. (eds) *OpenMP: Enabling Massive Node-Level Parallelism. IWOMP 2021*. Lecture Notes in Computer Science, vol 12870. Springer, Cham. https://doi.org/10.1007/978-3-030-85262-7_8

3. Storage and RAM as a service

Status	Available
Access arrangements	https://safe.epcc.ed.ac.uk/dirac (project do009) cosma-support@durham.ac.uk
Organisations	Durham University, NVIDIA
Project linkage	ExaHype, MPPHEA, ExCALIBUR training programme

The Durham Adaptable Memory System has distinct components that are required to investigate adaptable memory technologies that will function as a testbed and demonstrator for the ExCALIBUR H&ES programme.

The testbed provides access to BlueField-2 and Gen-Z technologies for HPC within the UK, building on prior BlueField-1 work at Durham. The testbed is integrated with the DiRAC COSMA system, allowing the existing login nodes, LDAP servers and administration consoles to be used. Researchers wishing to use the testbed should request an account through the SAFE system managed by EPCC.

BlueField-2 has significant advantages over the original BlueField-1 cards, namely increased processing power and clock rate of the embedded Arm cores. This is therefore an ideal time to realise this test cluster environment, placing the UK on the leading edge of this novel technology.

Basden, Alastair, Weinzierl, Marion, Weinzierl, Tobias, & Wylie, Brian J. N. (2021). *A novel performance analysis workshop series concept, developed at Durham University under the umbrella of the ExCALIBUR programme* (1.0). Zenodo. https://doi.org/10.5281/zenodo.5155503

Schulz, H., Gadeschi, G.B., Rudyy, O., Weinzierl, T. (2021). *Task Inefficiency Patterns for a Wave Equation Solver*. In: McIntosh-Smith, S., de Supinski, B.R., Klinkenberg, J. (eds) *OpenMP: Enabling Massive Node-Level Parallelism. IWOMP 2021*. Lecture Notes in Computer Science, vol 12870. Springer, Cham. https://doi.org/10.1007/978-3-030-85262-7_8

4. ARM+GPU Demonstrator

Status	Available
Access arrangements	Contact Professor Mark Wilkinson <miw6@leicester.ac.uk> or Dr Christopher Mountford <cjm14@leicester.ac.uk>
Organisations	University of Leicester, NVIDIA
Project linkage	UKAEA (MFEM, MOOSE, OpenMC codes), OpenQCD-Fastsum (DiRAC)

The aim of this project was to ensure that:

- ARM servers work harmoniously with accelerators such as GPUs
- Any shortcomings are understood, documented and reported to vendors
- ExCALIBUR and the wider UK research community has access to an ARM-GPU testbed

The project put into place an ARM+GPU testbed system based on the NVIDIA ARM HPC Developer Kit. This included four nodes with specifications as detailed below, incorporated into the University's existing ARM Catalyst system:

1. GIGABYTE G242-P32, 2U server
2. 1x Ampere Altra Q80-30 (Arm processor)
3. 512G DDR4 memory
4. 6TB SAS/ SATA 3.5" storage
5. 2x NVIDIA A100 GPU
6. 2x NVIDIA® BlueField®-2 E-Series DPU

The project included Research Software Engineering (RSE) effort for the porting, benchmarking and development of existing codes and creation of digital assets including progress reports, white papers and how-to documents, as well as software enhancements and modification.

5. Cerebras Wafer Scale Engine testbed

Status	Available
Access arrangements	https://www.ed.ac.uk/edinburgh-international-data-facility/services/computing/cerebras-cs-1 or contact Professor Mark Parsons <m.parsons@epcc.ed.ac.uk>
Organisations	EPCC, University of Edinburgh, Cerebras
Project linkage	ELEMENT, EXALAT, Exascale Computing for System-Level Engineering, Materials and Molecular Modelling

This project brought a Cerebras CS-1 Wafer Scale Engine system to the UK, the first such system in Europe. The CS-1 has enabled performance and usability exploration for UK academic and industrial users. The majority of the system has been funded by the University of Edinburgh, however the support from ExCALIBUR HE&S has allowed for a more general access service to be provided to researchers from across ExCALIBUR and the wider UK computational science and AI community.

Cerebras Systems have developed the world's largest processor, the Wafer Scale Engine (WSE), at over 46,000 square millimetres, with 1.2 trillion transistors, 400,000 processor cores, 18 gigabytes of SRAM, and an interconnect between processors capable of moving 100 million billion bits per second.

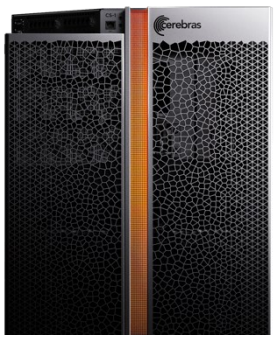

The Cerebras CS-1 system is focussed on neural network training and Cerebras Systems have integrated their hardware into common machine learning frameworks such as TensorFlow and PyTorch2, opening up the potential for easy porting of existing applications to the system. They also provide a graph compiler (CGC) and optimised library kernels, to efficiently map applications to the many processors on the WSE and ensure optimal use of the resource.

With potential for extreme performance for a wide range of machine learning training tasks, the CS-1 is an exciting new technology. However, there is currently a lack of user experience and application performance data to assess the suitability of the hardware for actual applications, and the requirements/costs for porting codes to the system. With a software environment that partially resembles standard CPU- and GPU-based systems, and partially resembles FPGA-based systems, with associated placement and routing requirements, it is important to be able evaluate both performance and usability of the CS-1 for end user applications. Such end user applications may also include more traditional numerical applications and this will be an area of exploration on the system.

6. Graphcore testbed

Status	Available
Access arrangements	Contact Dr Owain Kenway <o.kenway@ucl.ac.uk>
Organisations	University College London, STFC, Graphcore
Project linkage	Various

This testbed evaluated the Graphcore IPU-M2000 system for high performance and scientific computing applications, providing a novel architecture for the community to test and develop AI compatible codes on. The IPU (Intelligent Processing Unit) is a completely new kind of massively parallel processor, co-designed from the ground up to accelerate machine intelligence.

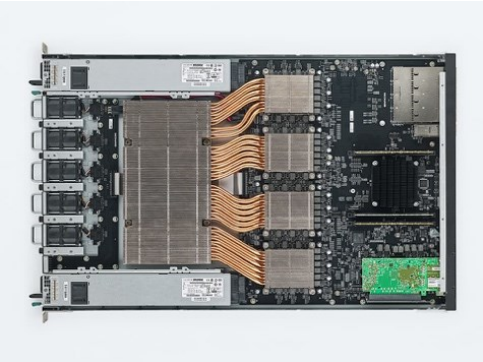

Each MK2 GC200 IPU in the IPU-M2000 unit has 1,472 processor cores, running nearly 9,000 independent parallel program threads with 900MB in processor memory and 250 TeraFlops of AI compute at FP16.16 and FP16.SR (stochastic rounding). The IPU-M2000 system has four IPUs, delivering approximately 1 PetaFlop of AI compute, and supporting ultra-low latency IPU-Fabric interconnect.

The testbed includes four IPU-M2000 systems, enabling the interconnect to be tested and characterised. The project has evaluated the Graphcore system's intended use cases around AI training and interference, and also looked at a subset of HPC codes that may be suitable for this platform. The Graphcore system is also available to the ExCALIBUR and wider UK research community with support and a training programme from the UCL team. It should be noted that codes will need to fit into small memories and must be single or half precision due to IPU requirements.

7. FPGA testbed

Status	Available
Access arrangements	https://fpga.epcc.ed.ac.uk/ fpga-testbed@mlist.is.ed.ac.uk
Organisations	EPCC, University of Edinburgh, University College London, University of Warwick
Project linkage	ExaClaw, ExCALIBUR-HEP, ELEMENT

This testbed system and associated effort for enabling software, allows researchers to port their scientific and data-science applications to Field Programmable Gate Arrays (FPGAs) and explore performance and power advantages such technology provides. Composed of next-generation hardware and software, the testbed forms an important UK resource for exploring the future role of FPGA technology in science, engineering, and the broader computational science communities.

In addition to the testbed hardware itself the project has been supported by Research Software Engineer (RSE) effort to develop the software stack to enable easier usage of FPGAs, driven by specific use cases from the Excalibur Design and Development Working Groups and other interested application communities.

Project partners EPCC, UCL, and Warwick worked in collaboration with FPGA vendor Xilinx, Inc to deliver and operate the testbed. The testbed is envisaged as the foundation of a future community and ecosystem around the role of FPGAs in HPC, data science, AI, and machine learning workloads in the UK. To this end the project has run a series of training events and workshops, and developed training material to ensure the system is accessible and usable.

The testbed provides access to next-generation Versal Adaptive Compute Acceleration Platform (ACAP) technology from Xilinx, including their revolutionary AI engines; hierarchical memory hardware provision, with high bandwidth (HBM2) and Non-Volatile (NVRAM) memory on some of the hosted hardware, providing a unique resource for software developers and algorithm designers to investigate this emerging field in computing hardware; multiple networking options including a high performance node-level network and direct FPGA to FPGA networking to enable system designers and applications developers to assess the relative merits of both approaches; and multiple families of FPGA, allowing evaluation of a range of technologies by users.

Additionally, RSE effort has been dedicated to providing an enabling software stack to significantly reduce the barrier to entry in utilising FPGAs for scientific and data-science applications.

8. NextSilicon testbed

Status	Under development
Access arrangements	TBC – please contact Principal Investigator Professor Iain Styles to discuss <i.b.styles@bham.ac.uk>
Organisations	University of Birmingham, NextSilicon, Lenovo
Project linkage	TBC

This project will create a testbed featuring novel accelerator technology from NextSilicon in collaboration with University of Birmingham HPC systems partner Lenovo, as part of a co-design partnership. The project will evaluate the performance of the main codes used by UKRI researchers, with a particular emphasis on evaluating some of the major algorithm classes used in supercomputing and data science, as well as other HPC apps.

Should the technology fulfil its promised potential in the evaluation, it will be made available to the UK HPC community as part of the recently awarded Baskerville EPSRC Tier 2 service operated by the University of Birmingham. This facility, installed in Q1 2021, includes 184 Nvidia A100 GPU accelerator cards. Other novel accelerator technologies are also planned for introduction over the lifetime of the facility, thus enabling us to assess the new technology in a heterogeneous accelerated live HPC environment.

9. Exascale Data Testbed

Status	Available / Under development
Access arrangements	Contact Dr OG Parchment <ogp21@cam.ac.uk>
Organisations	University of Cambridge, Dell, Intel, Altair
Project linkage	UKAEA MAST Experiment (visualisation of plasma simulations) Excalidata: Addressing I/O and Workflows at Exascale DiRAC: Optimising I/O for AREPO (Cosmological hydrodynamics) Excalistore, IRIS, UKSRC (SKA)

This testbed utilised HPC systems development, deployment and operational skills housed within the Cambridge Research Computing Service to build a next generation high performance PCI-Gen-4 solid state I/O testbed utilising a range of file systems including Lustre, Intel DAOS, BeeGFS and HDF5 on state-of-the-art solid state storage hardware.

The system is based on Intel PCI Gen-4 NVMe drives and Optane Data Centre Persistent Memory. The project saw the deployment of the UK's fastest HPC storage testbed delivering over 500GB/s bandwidth and over 20 million IOPS of raw I/O performance. This was deployed across applications via a range of HPC file systems such as Lustre, Intel DAOS or BeeGFS as well as other more low level direct I/O protocols.

In addition to the I/O hardware and various file system technologies the testbed is configured with comprehensive system level telemetry monitoring capability provided by the UKRI funded Scientific OpenStack middleware layer combined with a range of other more specialised application I/O profiling tools. The UK Scientific OpenStack is a HPC middleware layer developed at Cambridge and funded by over 4 years investment from STFC, EPSRC and MRC.

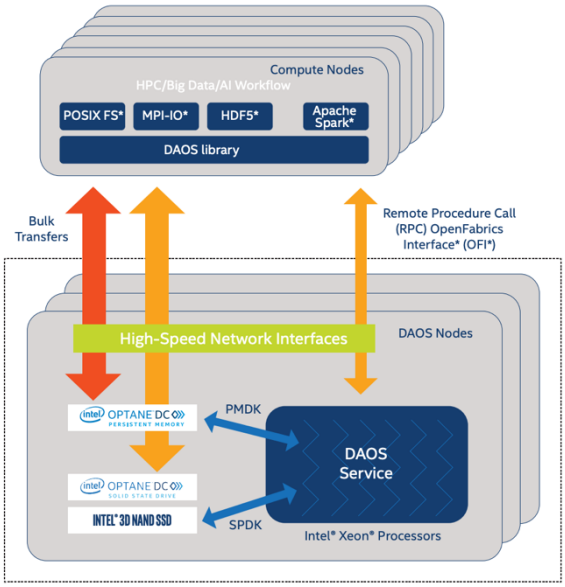

This project is supported directly by Intel in terms of hardware, staff effort and strong co-design work in collaboration with Intel engineers developing the DAOS file system.

10. In-situ Visualization

Status	Under Development
Access arrangements	Developed changes will be integrated into OSPRay (https://www.ospray.org/) and ParaView/Catalyst (https://www.paraview.org/). GRChombo example code will be made available at https://github.com/GRChombo/GRChombo.
Organisation	DAMTP, University of Cambridge, Intel
Project linkage	UKAEA: JOREK code exemplar for in-situ visualisation

This project seeks to create new in-situ visualization capabilities and to port exemplar community codes onto multi-accelerator systems using a unified programming model (oneAPI), which will be directly applicable to planned exascale systems. This programme represents a substantial development in a longstanding collaboration between Intel, TACC, Kitware and the Cosmos Intel Parallel Computing Centre and Intel Graphics and Visualization Institute of Xellence, hosted in the Cambridge Faculty of Mathematics (which has recently become an Intel oneAPI Centre of Excellence).

This project has been given early access to Intel's new accelerator systems to investigate the benefits of platform-independent oneAPI programming and the impact of architecture specific optimizations, and is also exploring performance on accelerators from other vendors. This programme builds on a strong track record of achievement by the award-winning Cosmos team.

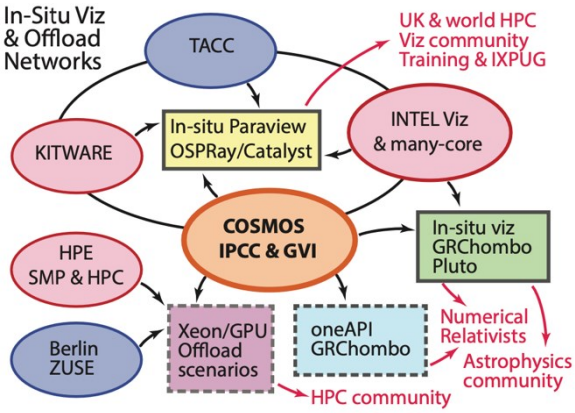

An important component of this programme is the dissemination of these in-situ visualization advances across the wider UK HPC community by implementing these capabilities in community codes, offering online training courses, and running open workshops, as well as SC and ISC presentations.

Radia, M., Sperhake, U., Drew, A., Clough, K., Figueras, P., Lim, E. A., Ripley, J., et al. (2022). Lessons for adaptive mesh refinement in numerical relativity. *Classical and Quantum Gravity* https://doi.org/10.1088/1361-6382/ac6fa9

Gerosa, Davide and Fabbri, Cecilia Maria and Sperhake, Ulrich (2022) *The irreducible mass and the horizon area of LIGO's black holes.* Classical and Quantum Gravity, 39 (17). Art. No. 175008. ISSN 0264-9381. https://doi.org/10.1088/1361-6382/ac8332

11. RISC-V Testbed

Status	Available / Under development
Access arrangements	https://riscv.epcc.ed.ac.uk riscv-testbed@mlist.is.ed.ac.uk
Organisations	EPCC, University of Edinburgh, Xilinx, Intel, RISC-V International
Project linkage	Machine Learning ExCALIBUR project, ExCALIBUR cross-cutting xDSL project

RISC-V is an open-source Instruction Set Architecture (ISA) that continues to grow rapidly since it was introduced over a decade ago. Overseen by RISC-V International, open standards are produced that the community, ranging from companies, academics, and hobbyists, are then free to implement via their own specific microarchitecture. A rich ecosystem has built up around RISC-V, with a great many implementations (many themselves open source) available. There is a growing interest in the use of RISC-V for HPC, including an official RISC-V HPC special interest group and events at leading HPC conferences.

This project is focussed on providing a RISC-V testbed system and supporting the development of enabling software to enable researchers predominantly within ExCALIBUR, but also the wider HPC community, to port their scientific and data-science applications to RISC-V CPUs and explore the performance properties such processors provide. This testbed forms an important UK resource, and one that is unique, for enabling the convenient exploring the role of this technology by application developers and researchers. We provide a mixture of physical and soft-core RISC-V CPUs, the former enabling user exploration with more mature RISC-V solutions, and the later enabling users to experiment with less mature next generation technology which we are also able to tune for select use-cases.

The project has provided the testbed itself, and Research Software Engineer (RSE) effort is being used to develop the soft-core multi-core RISC-V CPU configurations and enabled us to optimise specific HPC and data-science libraries for the architecture as well we develop tools that can aid adoption of RISC-V by the community. Access to the testbed is free and we very much welcome users to sign up and start using this exciting technology. You will find a familiar set of compilers and HPC libraries available via the module environment, and exciting hardware to experiment with. See the website for more details.

12. Rockport testbed

Status	Available
Access arrangements	Sign up at https://safe.epcc.ed.ac.uk/dirac, (project do009) cosma-support@durham.ac.uk
Organisations	Durham University, Rockport Networks, Dell
Project linkage	ExaHype, MPPHEA, Exploiting Task Parallelism

This project is based around the adaption of a large-scale HPC cluster (up to 226 nodes) to use a novel high performance 100G switchless networking technology developed by Rockport Networks.

Based on a 6D torus today, and extensible to other topologies, this network technology is highly modular and has linear scaling up to hundreds of thousands of nodes (and probably beyond), and is well suited for Exascale.

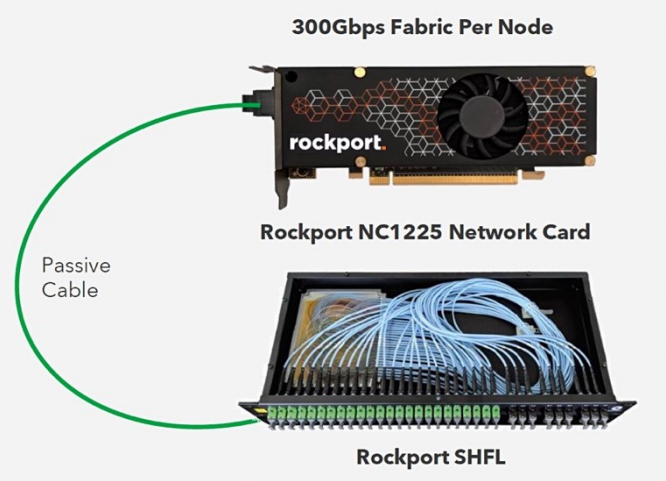

By distributing the network switching function into each device endpoint, the nodes become the network. This gives consistently low latency under load, for a predictable workload performance, on networks of any size.

For classic HPC networks, congestion can introduce significant delays for messages, significantly reducing code performance. However, with the network proposed here, packets are split into small FLITs , and so small messages are not delayed by larger messages within the pipeline, thus delivering consistently low network latency.

Switchless networking removes network complexity, replacing centralised switch fabrics with a distributed, highly reliable interconnect, providing an intelligent, adaptable and self- healing network which is simple to operate. The testbed Rockport system is at a scale large enough to be useful for scientific simulations, and will be invaluable in determining whether this technology and approach is appropriate for future Exascale class systems.

13. Benchmarking for Performance Portablity

Status	Available / Under development
Access arrangements	Contact Dr Tuomas Koskela <t.koskela@ucl.ac.uk>
Organisations	University College London, University of Reading, University of Bristol
Project linkage	All ExCALIBUR projects

This project continues the ExCALIBUR Benchmarking Initiative begun as part of the Adaptable Cluster H&ES project, bringing together and enabling ExCALIBUR projects to benchmark their applications on the diverse supercomputing platforms available in the UK's HPC ecosystem.

The team will collaborate with other ExCALIBUR teams through regular community meetings, knowledge exchange, presentations from projects and topical discussions, training sessions, and 1:1 drop-in sessions to assist application developers with contributing benchmarks to the framework. We will deploy the framework on UK Tier1/2 systems and the feasibility of deployment on H&ES testbeds will be investigated. Development of new features will be driven by feedback from users via GitHub. We will significantly expand the performance analysis in ReFrame to visualize complex performance portability (PP) data, making it convenient for ExCALIBUR application teams to measure and improve the PP of their codes.

Application benchmarking is a crucial activity in the UK's path to Exascale. It ensures future Exascale systems are understood by the community so that UK Science Applications, particularly those developed as part of ExCALIBUR, can take advantage of the scientific opportunities at Exascale. Linear solvers are a ubiquitous pattern in many ExCALIBUR codes, and this project will develop a suitable benchmark applicable to many of UK codes. Application benchmarks are part of the enabling software stack that allows the performance of hardware to be assessed for scientific workloads.

It is vital to rigorously measure the performance of benchmarks in a systematic way to enhance the transparency and enable reproducibility. International conferences in this space are now mandating reproducibility information. This project provides the tooling to make collecting and analysing performance data straightforward. This will allow application owners to develop the skills and experience to measure the performance of their code across the gamut of the UK's supercomputer facilities, and improve their code based on benchmarking as they target Exascale.

14. Coarse Grained Reconfigurable Architectures (CGRAs)

Status	Project in startup
Access arrangements	Contact Dr Nick Brown <n.brown@epcc.ed.ac.uk>
Organisations	EPCC, University of Edinburgh, Cerebras, Xilinx, Graphcore
Project linkage	ExCALIBUR xDSL and Machine Learning projects + additional ExCALIBUR projects to be identified via kick-off workshop

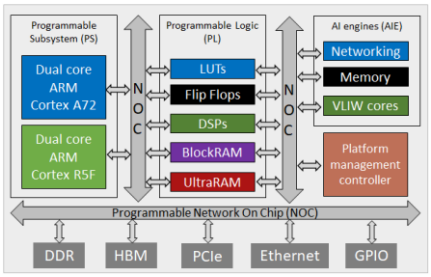

Coarse Grained Reconfigurable Architectures (CGRAs) describe a class of hardware that provides many independent Processing Elements (PEs) interconnected using a high-performance and flexible network. Each PE is effectively a simple CPU core and programmers will typically assign one or more tasks to each PE and then configure how these communicate. The raw concurrency combined with extremely high-performance networking making them very interesting for potential use as HPC accelerators at the exascale. Three of the existing ExCALIBUR H&ES testbeds provide technologies that can be described as CRGAs and will be a focus in this project:

- The Cerebras CS-1 Wafer Scale Engine
- The Versal AI engines

Whilst these architectures can provide significant raw compute performance, it is not easy for developers to easily exploit them. This is because of immaturity of the software ecosystem, where the vendor provided programming tools are low level and complex to use. This project will take advantage of existing H&ES investments by focussing on the enabling software component required to take advantage of these systems most effectively for HPC workloads and ExCALIBUR projects, effectively aiming to close the gap between the CGRA hardware and higher-level software.

These CGRA technologies demonstrate significant potential as accelerators for HPC workloads, and with AMD announcing in June 2022 plans to integrate the Versal AI engines into their future enterprise CPUs, such architectures are likely to become much more commonplace. However, whilst there is significant interest and promise from vendors, when it comes to actually using these technologies for accelerating HPC simulations, to date, they have not yet been proven for such workloads.

By exploring these CGRA architectures for HPC codes, developing best-practice, and some library support, this project will deliver the capability so that ExCALIBUR software projects and the HPC community can more widely exploit these existing testbeds in new ways, delivering benefit to them, H&ES, and the ExCALIBUR programme.

15. DWAVE Quantum Computing

Status	Project in startup
Access arrangements	Contact Dr Alastair Basden <a.g.basden@durham.ac.uk>
Organisations	Durham University, DWAVE
Project linkage	QEVEC, HPC-PAX

The project will provide researchers with access to a DWAVE quantum computing system, a high RAM HPC facility for simulating quantum systems and 5 day training course on developing applications for DWAVE's quantum annealing technology.

Quantum computing has the potential to offer an alternative route to Exascale computing for particularly suited algorithms. The QEVEC ExCALIBUR software project is investigating ways in which quantum computing can be used alongside classical computing as an accelerator or to aid the classical computations.

Currently, access to quantum computing facilities is provided for free for very limited usage. It has been the case that researchers have had to plan their usage such that their monthly limits are not used too quickly, and to allow deadlines (e.g. for publication) to be met in a timely fashion. Having to plan in this way reduces researcher effectiveness, and thus there is a community need for immediate and relatively unfettered access to quantum computing systems.

Quantum annealing provides a potential route for solution of annealing problems, including graph partitioning, and has a wide range of use cases. The main alternative approach to quantum annealing is the quantum gate model.

16. CXL technology demonstrator

Status	Project in startup
Access arrangements	https://safe.epcc.ed.ac.uk/dirac (project do009) cosma-support@durham.ac.uk
Organisations	Durham University, Liqid
Project linkage	HPC-PAX, UKSRC (SKA)

This testbed will explore technologies developed as part of the Compute Express Link (CXL) v3 standard, introducing full composable memory and cache coherent memory sharing between remote servers using the Liqid composable infrastructure platform.

CXL is a new standard for connectivity between CPUs and other components, and is largely expected to replace PCIe in the future. The most novel new feature is support for cache coherent protocols for accessing system and device memory which will open up new capabilities, new programming models, and allow heterogeneous codes to operate more efficiently, and with greater simplicity, both from programmer and execution point of view.

CXL promises to be the dominant standard for years to come, having recently accepted transfer of assets from both the Gen-Z and OpenCAPI consortiums. All major manufacturers have signed up to adopt CXL, and the use of composable RAM systems potentially offers significant cost savings for future DRI infrastructure. By equipping a server with sufficient RAM on demand, rather than a speculative purchase which may only require this RAM occasionally, RAM usage can be optimised. Additionally, further RAM can be added to systems at a later date if this becomes necessary, e.g. as problem sizes grow beyond the initial scope. The Liqid software stack is integrated with SLURM, and so it is possible to automatically compose based on job requirements.

There are several distinct communities across UKRI who require large RAM systems for Exascale, for portions of their operations. One example is the SKA post-processing teams, who need to access and process large datasets which do not fit in RAM in conventional servers. Cosmology simulation (ExCALIBUR HPC-PAX project) also require large memory systems.

Experience with composable RAM systems is necessary to assess the impact on non-local RAM on scientific codes. The non-local RAM will have higher latency and reduced bandwidth compared with node-local RAM. Therefore codes will have to be made aware of this, and utilise the RAM in the most effective manner, namely by making use of streaming datasets, with data repositioned in memory to be contiguous, and placing regularly touched datasets placed in local RAM and less frequently touched data in non-local RAM.

17. Enabling PETSc on the Cerebras Wafer Scale Engine

Status	Project in startup
Access arrangements	Contact Prof Michèle Weiland <m.weiland@epcc.ed.ac.uk>
Organisations	EPCC, University of Edinburgh, Cerebras
Project linkage	CGRAs

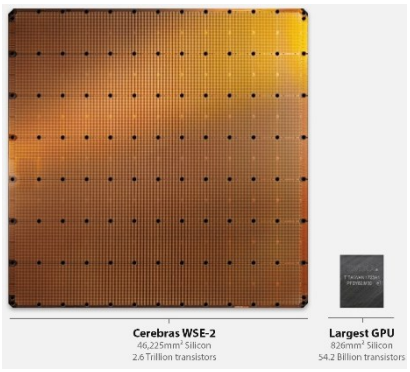

The aim of the project is to enable the widely used linear algebra library PETSc on the Cerebras Wafter Scale Engine (WSE). The Cerebras WSE has huge potential to accelerate numerical applications, however to date usage has mostly been focused on machine learning applications, primarily through frameworks such as Tensorflow and PyTorch. However, there is nothing inherent about the architecture that limits it to those use cases - what currently limits usability is the immaturity of the software stack and a lack of more general purpose HPC libraries being available for the system. Enabling a linear algebra library such as PETSc will open these systems up to a much wider user base, exposing the community to the potential benefits of this architecture for their workloads.

The project has 3 key activity areas:

1. Evaluate the feasibility of porting an established HPC library to an emerging dataflow architecture, in particular in light of a potentially immature supporting software stack;
2. Assess the performance of PETSc linear algebra kernels on the Cerebras WSE;
3. Provide a proof-of-concept implementation of PETSc for Cerebras WSE

The Cerebras WSE is programmed using a domain specific language, called CSL. Kernels that run on the WSE are written in CSL and compiled to binaries that are launched from a Python application, which runs on the host system. The host, via the Python application, performs the setup work and data marshalling (streaming data to the WSE and receiving results back). Data is streamed to the WSE in the form of NumPy objects.

PETSc2 is a widely used library for linear algebra, developed at Argonne National Laboratory. It is written in C and C++, and it provides full bindings for Fortran as well as for Python, via petsc4py. In this project, we will exploit the Python interface for PETSc to implement support for the WSE. The Python application that runs on the host will be used to set up the PETSc environment, perform the I/O and create the data structures using petsc4py. The computational kernels, e.g. axpy and spmv, will be recast using a dataflow paradigm and implemented in CSL for execution on the WSE. They will be compiled into kernel binaries that are launched from the host. The PETSc data objects will be converted to NumPy objects, thus enabling us to stream them to the binaries running on the WSE.

18. Intel Ponte Vecchio GPU Pre-Exascale Testbed

Status	Project in startup
Access arrangements	Contact Prof Paul Shellard <E.P.S.Shellard@damtp.cam.ac.uk>
Organisations	University of Cambridge, Intel, Lenovo
Project linkage	NEPTUNE (UKAEA)

This testbed will procure and co-design a multi-accelerator system based on the Intel oneAPI enabled Ponte Vecchio GPU (Xeon Data Center GPU Max) together with the 4th gen Sapphire Rapids CPU (Xeon CPU Max). The testbed will act as a demonstrator subunit or brick, with over half a petaflop of peak performance.

Project milestones:

- Co-design of Intel multiple-accelerator testbed of PVC GPUs and Sapphire Rapids CPUs on an expandable system that can also host other vendor GPUs for comparative and open standards testing
- Implementing new OSPRay/Paraview in-situ or “on the fly” visualization demonstrators using multiple accelerators, incorporating our new AMR capability and exploiting unique PVC hardware raytracing
- Development platform for porting flagship codes onto GPUs using SYCL and OpenMP, including the public GRChombo/GRAMReX numerical relativity code
- Targeting AI workloads on GPUs, we will optimize machine learning pipelines used for LIGO-Virgo gravitational-wave (GW) data analysis, notably the parameter estimation DINGO deep learning code
- We will explore enhancements from the fast Intel Xe Link GPU-GPU interconnect which enables faster peer access to cross-GPU memory, comparing with GPU-host-GPU PCIe 5.0 transfers. We will also study the performance impact of using SYCL to access the unified address space across x4 GPU memory of 0.5TB
- We will explore alternative memory hierarchies for the new high bandwidth memory (HBM) on Xeon Max, determining workload size effects on performance when configured in flat or caching modes
- The project will investigate and report on benefits of open standards programming (SYCL/OpenMP) and the impact of architecture specific optimizations

Our track record studying offload paradigms to multiple accelerators means this unique testbed will enable application development for future exascale hardware on which CPUs and GPUs are tightly integrated. Direct access to an Intel PVC GPU platform provides an important opportunity for UK ExCALIBUR researchers to assess both PVC GPU and Sapphire Rapids CPU performance and the benefits of architectural enhancements, such as fast GPU-GPU interconnect and HBM.