



Case Study: Durham University's Institute for Computational Cosmology Accelerates Results with Composability from Liquid

Leading Cosmology Researchers Using Composability to
Support Modelling of the Universe on COSMA8



About

Durham University is home to the COSmology MACHine (COSMA) operated by the Institute for Computational Cosmology (ICC), and is a national supercomputing facility that is part of the UK's DiRAC, or Distributed Research utilising Advanced Computing system. DiRAC is a distributed computing facility established to provide High Performance Computing (HPC) services to the UK research community and consists of sites at four leading universities – Durham, Cambridge, Edinburgh and Leicester -- each tailored for a particular workload. Durham hosts the COSMA memory intensive site designed specifically to support the largest cosmological simulations, most notably running simulations of the universe starting with the Big Bang and propagating through the almost 14 billion years of history. COSMA, now on the eighth generation, is used by researchers across the UK and collaborators around the world. The UK ExCALIBUR programme provided funding to investigate novel hardware solutions with potential relevance for Exascale systems.

“Composing our GPUs has improved utilisation, because now we don’t have a bunch of GPUs sitting idle. When the demand for GPUs grows, we have the infrastructure in place to support more acceleration in a very flexible configuration.”

ALASTAIR BASDEN, TECHNICAL MANAGER FOR THE DIRAC MEMORY INTENSIVE SERVICE AT DURHAM UNIVERSITY

Challenge

Using large-scale simulations, COSMA8 is helping scientists tackle some of the biggest questions of humankind, including our understanding of where we came from and our place in the cosmos. Each simulation of dark matter, dark energy, black holes, galaxies and other structures in the Universe often takes months to run on COSMA8. Subsequently,

there are then long periods of data analysis during which researchers analyze simulation results. Given the episodic nature of the simulations and associated compute resources, Alastair Basden, technical manager for the DiRAC Memory Intensive Service at Durham University, is always looking for new technologies to increase the efficiency and utilisation of COSMA8. When Basden was introduced to the concept of composable infrastructure by his Dell representatives, he was inspired to learn more about its potential in his environment.

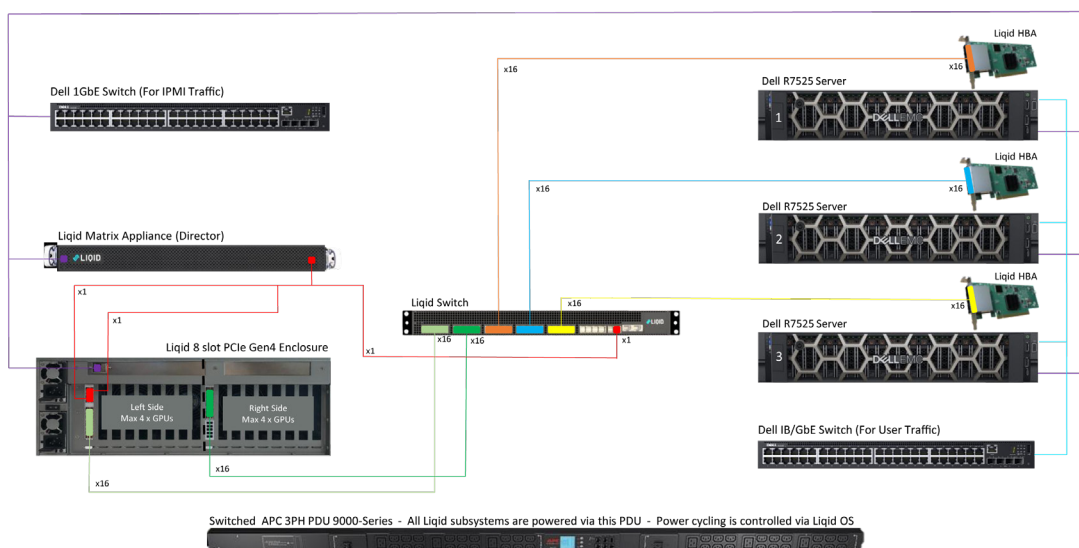




The Solution

While all simulations running on COSMA8 have high memory demands, some workloads require GPU acceleration. Since GPUs can be among the most expensive resources in an environment, it's important to ensure utilisation is maximized. To accomplish this for GPU acceleration, Basden is composing NVIDIA A100 GPUs to their servers with Liquid's composable disaggregated infrastructure (CDI) software. This then enables researchers to request and receive precise GPU quantities. "Most of our simulations don't use GPUs," according to Basden. "It would be wasteful for us to populate all our nodes with GPUs. Instead, we have some fat compute nodes and a login node, and we're able to move GPUs between those systems. Composing our GPUs gives us flexibility with a smaller number of GPUs. We can individually populate the servers with one or more GPUs as required at the click of a button."

Durham University's Liquid Composable Solution



While composing the GPUs is providing immediate flexibility for researchers using COSMA 8, what excites Basden most about composability is the potential of composing memory. Each of the main COSMA8 360 compute nodes at Durham is configured with one terabyte of RAM. "RAM is very expensive – about half the cost of our system," said Basden. "Some of our jobs don't use RAM very effectively – one simulation might need a lot of memory for a short period of time, while another simulation doesn't require much RAM. The idea of composing RAM is very attractive; our workloads could grab memory as needed."



Basden added, “When we’re doing these large simulations, some areas of the universe are more or less dense depending on how matter has collapsed. The very dense areas require more RAM to process. Composability lets resources be allocated to different workloads during these times and to share memory between the nodes. As we format the simulation and come to areas that require more RAM, we wouldn’t have to physically shift things around to process that portion of the simulation.”

Ultimately, Basden wants to compose memory over the fabric. “Rather than sending all the data over Infiniband, we could lessen congestion of the traditional fabric, plus support lower latency and higher bandwidth communication for the vast amount of memory we maintain here at Durham,” said Durham.

Results

“Durham University is using cutting-edge CDI to accelerate research, improve resource utilisation, and reduce the university’s carbon footprint,” said Basden.

“Composing our GPUs can improve utilisation, because now we don’t have a bunch of GPUs sitting idle,” said Basden. He’s seeing an increase in the number of researchers exploring artificial intelligence and expects those workloads to require more GPUs. “When the demand for GPUs grows, we have the infrastructure in place to support more acceleration in a very flexible configuration.”

Composability also has the potential to play a role in reducing the carbon footprint of COSMA. “Rather than populating all our servers with GPUs, we can compose the resources we need to each server. That reduces our carbon footprint,” said Basden.



About Liquid

Liquid's [composable infrastructure](#) software platform, Liquid Matrix™, unlocks cloud-like speed and flexibility plus higher efficiency from infrastructure at the core and edge. Now IT can configure, deploy, and scale bare-metal servers in seconds via software, and reallocate accelerator and storage resources real-time as needs evolve. Dynamically provision previously impossible systems or scale existing investments, and then redeploy resources where needed in real-time. Unlock cloud-like datacenter agility at any scale and experience new levels of resource and operational efficiency with Liquid.

About Durham University

Durham University is a world top 100 higher education institution in the UK. The COSmology MACHine (COSMA) housed at Durham is part of DiRAC (Distributed Research utilising Advanced Computing), the UK's integrated supercomputing facility for theoretical modelling and HPC-based research in particle physics, astronomy and cosmology. COSMA is specialized on memory intensive applications, particularly cosmological simulations. The first COSMA cluster came online in 2001 and is currently in its eighth generation. For more information, visit www.dur.ac.uk/cosma.

About DiRAC

DiRAC was established to provide distributed High Performance Computing (HPC) services to the STFC theory community. HPC-based modelling is an essential tool for the exploitation and interpretation of observational and experimental data generated by astronomy and particle physics facilities supported by STFC as this technology allows scientists to test their theories and run simulations from the data gathered in experiments. The UK has an extremely strong HPC community and these powerful computing facilities allow the UK science community to pursue cutting-edge research on a broad range of topics, from simulating the entire evolution of the universe, from the big bang to the present, to modelling the fundamental structure of matter. DiRAC is both an academic-led and an academic-supervised facility and our systems are specifically designed to meet the different high performance computational needs within our scientific community.



About the ICC

The Institute for Computational Cosmology is a world-leading centre for research in Cosmology and Astrophysics using large supercomputers. It hosts about 50 researchers who, with collaborators across the world, are responsible for some of the largest and most sophisticated cosmological simulations in the world. These are used to address fundamental questions regarding the nature of dark and dark energy, as well as the formation of galaxies and other large cosmological structures. See <https://www.icc.dur.ac.uk>

About ExCALIBUR H&ES

ExCALIBUR (Exascale Computing Algorithms and Infrastructures Benefitting UK Research) is a £45.7m UK Strategic Priorities Fund (SPF) programme 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 enables the redesign of high priority computer codes and algorithms to meet the demands of advancing technology and UK research. As part of the ExCALIBUR initiative, UKRI are investing £4.5m over five 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.