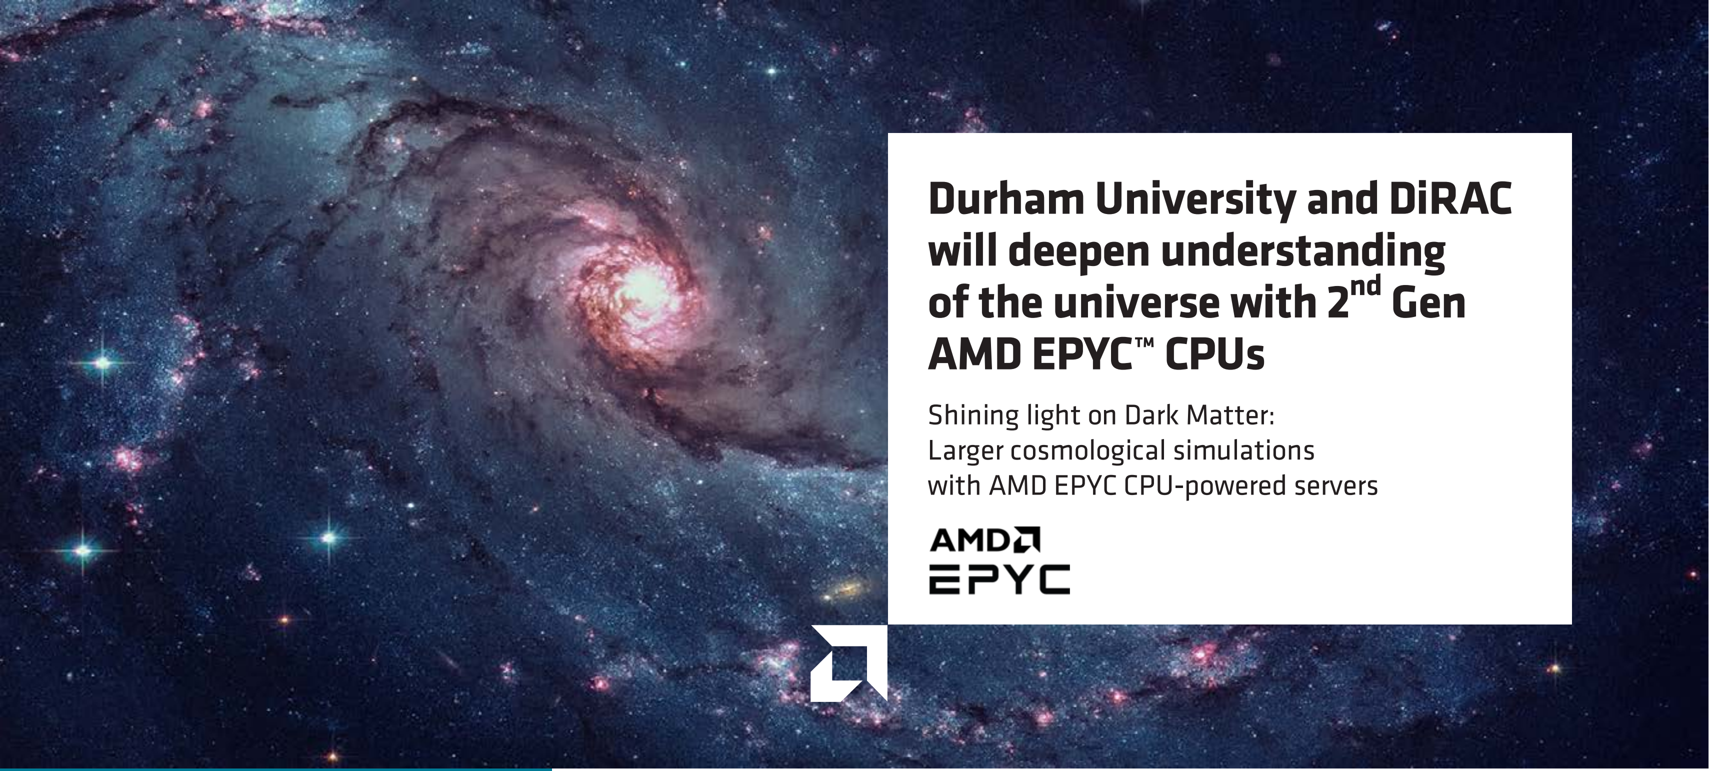

Durham University and DiRAC will deepen understanding of the universe with 2nd Gen AMD EPYC™ CPUs

Shining light on Dark Matter:
Larger cosmological simulations with AMD EPYC CPU-powered servers

AMD EPYC

CUSTOMER

Durham University
Institute for Computational Cosmology

DiRAC
High Performance Computing Facility

INDUSTRY

Cosmological research

CHALLENGES

Increasing cluster core count with high per-core memory

SOLUTION

Deploy AMD EPYC™ 7H12 processors

RESULTS

Sixfold increase in core count, providing the ability to run much larger cosmology simulations more quickly

AMD TECHNOLOGY AT A GLANCE

AMD EPYC 7H12 processors with 64 cores

TECHNOLOGY PARTNER

DELL Technologies

Researching the big cosmological questions, such as the origin of the universe, takes massive amounts of computing power. As a top 100 world university, Durham in the UK is at the forefront of this science. The university hosts the DiRAC Memory Intensive (MI) service, part of the Tier 1 DiRAC High Performance Computing (HPC) facility.

The DiRAC (Distributed Research utilizing Advanced Computing) facility has four deployments, at the Universities of Cambridge, Durham, Edinburgh and Leicester, overseen by the Project Office at University College London. DiRAC was established to provide specialized HPC services to the astrophysics, cosmology, particle and nuclear physics theory communities. HPC-based modelling is an essential tool for the exploitation and interpretation of observational and experimental data generated by astronomy and particle physics facilities. It allows scientists to test their theories and run simulations from the data gathered in experiments.

"The extra core count meant the AMD EPYC processors were significantly faster."

Alastair Basden, Technical Manager for the DiRAC Memory Intensive service, Durham University

The DiRAC MI service at Durham is focused is on workloads that require larger amounts of memory than would be typical "For our current system, we have about 18GB of RAM per core," says Alastair Basden, Technical Manager for the DiRAC MI service at Durham University. "When you compare that with typical systems of two to four gigabytes, that's a significant uplift. We do large-scale cosmological simulations of the universe, right from the start of the Big Bang up until the present day." This is a very memory intensive HPC application, so when Basden discovered how the AMD EPYC™ processor combined high core density with support for large amounts of fast RAM, it looked like it could take the DiRAC MI service to the next level.

Exploring the universe with COSMA

Given the focus on cosmological simulations at Durham, the compute clusters comprising the DiRAC MI service are named COSMA, which stands for Cosmology Machine. The first generation of COSMA came online in 2001, prior to the establishment of DiRAC. The latest iteration is COSMA 8, but successive upgrades tend to overlap, so the DiRAC MI service is actually three clusters (COSMA 6, 7 and 8), while local Durham users continue to exploit COSMA 5.

The DiRAC MI service has already been part of some literally earth-shattering science. Planetary collision simulations have been used to great effect to explain the origin of the strange rotational axis of Uranus that we observe. Furthermore, some of the simulations that led to the discovery of gravitational waves were performed on the DiRAC MI service. The COSMA system was also used to run the Eagle simulation, which led directly to the paper by Schaye et al (2015), which is the most cited astronomy paper published that year (out of more than 20000 papers).

COSMA's requirements are specific to the kind of workloads supported by the DiRAC MI service. "It's what we call a capability system," explains Basden. "This is a system that's designed to offer a capability that wouldn't otherwise be available unless you go to a much larger setup. It's a capability system primarily because of the large amounts of RAM. If you're doing large-scale cosmological simulations of the universe, you need a lot of RAM. They can have run times of months, and then, after they've produced their data, which will be snapshots of the universe at lots of different time steps and different red shifts, then years are spent in processing and analysis."

Putting COSMA's capabilities in perspective, DiRAC researchers also sometimes call upon other European facilities. "We generally find that they're not so good, because they've got much lower RAM per core. Even with a much larger allocation of CPUs on those machines, we tend to get better results on COSMA, because it's more designed for these simulations."

This balance between core count and memory is crucial for COSMA's workloads and was why the AMD EPYC processor's strengths were particularly attractive. "I first came across EPYC processors in a different role. We had four nodes on loan from Dell, before they'd gone into production. We also had a previous application that needed a very high memory bandwidth, and the 'Naples' processors fitted that bill very well."

Twice as fast with AMD EPYC CPUs

Basden was particularly interested in the 2nd Generation AMD EPYC processor (formerly codenamed 'Rome') for the next iteration of the DiRAC MI service, COSMA 8, due to its even higher core density than the 1st Generation and single NUMA memory domain. "We got access to the Dell data center in Austin so that we could do benchmarking on these processors," says Basden. A large fraction of the simulation work on COSMA 8 is performed with software called SWIFT (SPH With Inter-dependent Fine-grained Tasking). Basden used sample cosmology datasets with SWIFT to test the processors.

> *"It will have more than twice as much DRAM as previous systems. We'll be able to double the size of the simulations, but having six times as many cores means we'll be able to get through these datasets much more quickly."*
>
> *Alastair Basden, Technical Manager for the DiRAC Memory Intensive service, Durham University*

"It ran as we would hope," he says. Naturally, tests were performed against competitor processors. "In terms of core for core, it was basically the same, and when you've got more cluster cores, it's a no-brainer." Thanks to the 64 cores-per-processor with AMD EPYC, compared to the Intel Xeon Platinum Cascade Lake series' maximum of 28, much greater density was possible. "The extra core count meant the AMD EPYC processors were significantly faster." For COSMA 8, Basden has opted for dual 280-watt AMD EPYC 7H12 processors per node with a 2.6GHz base clock frequency and 64 cores, installed in Dell Cloud Service C-series chassis with a 2U form factor and custom CoolIT water cooling.

"We wanted a large number of cores per node, because then it meant we could cut down on the amount of internode communication," explains Basden. "But because there are parts of the code that don't parallelize 100 percent, we also wanted high clock rates so that the lower-threaded parts of the code stood up well, which meant the 7H12 would be the best option."

Expanding our knowledge of the universe

The initial installation of COSMA 8 consists of 32 nodes. "This was primarily for testing, for getting code up to scratch, and small-scale simulations" explains Basden. Basden specified 1TB of memory per server, giving a sizeable 8GB memory per core. The speed from having eight-channel 3,200MHz memory is invaluable too. "Because we have a large amount of data, it's important to be able to pull it through the processors quickly."

The AMD EPYC processor's support for PCI Express® 4 (PCIe4) has been another invaluable feature vs. Intel's support for PCIe3. COSMA 8 employs Mellanox® cards that allow each node to access the memory on others via RDMA. "The fastest readily available interconnect is currently 200Gbits per second, and that requires PCI Express 4.0. With PCI Express 3.0, we wouldn't be able to reach that," says Basden.

DiRAC@Durham is also exploring GPU acceleration with AMD Radeon Instinct™ graphics cards. "We have recently purchased a couple of nodes with AMD MI50 GPUs, for running AI workloads like TensorFlow. We will be porting code to them specifically to see whether they are a solution for Exascale computing." This will be aimed at investigating the benefits of this hardware for workflows across all areas of DiRAC science.

For the DiRAC MI service, whose focus is cosmology, the increased core density, larger amounts of faster DRAM per node, and faster PCI Express 4.0 connectivity together add up to a potent boost in performance. "This will help us to understand the nature of the universe, dark matter, dark energy and how the universe was formed. It's really going to help us drill down to a fundamental understanding of the world that we live in."

About Durham University

Durham University is a world top 100 higher education institution in the UK. The DiRAC Memory Intensive (MI) service clusters (COSMA) are part of the DiRAC (Distributed Research utilizing Advanced Computing) facility that provides specialized HPC resources for research into astrophysics, particle physics, cosmology and nuclear physics. COSMA is specialized on memory intensive applications, particularly cosmological simulations. The first COSMA cluster came online in 2001, and the eighth iteration is currently rolling out, using AMD EPYC processors. For more information, visit www.dur.ac.uk/cosma.

About DiRAC

The DiRAC (Distributed Research utilizing Advanced Computing) facility was established in 2009 to provide specialized High Performance Computing (HPC) services to a domain-specific theory community. 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 its systems are specifically designed to meet the different high performance computational needs within our scientific community. DiRAC provides a variety of compute resources that match machine architecture to the different algorithm design and requirements of the research problems to be solved. There are sound scientific reasons for designing the DiRAC services in this way and the methodology was adopted following a number of in-depth reviews.

About AMD

For 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies—the building blocks for gaming, immersive platforms, and the data center. Hundreds of millions of consumers, leading Fortune 500 businesses, and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit amd.com/EPYC.