Dell Technologies

Solution brief

Case study: Optimized performance and energy efficiency in cosmology using AMD Genoa CPUs

Contents

Essentials

Durham University in the UK is a leader in cosmology research. DiRAC is a UK national HPC provider for the theoretical cosmology, nuclear and particle physics communities funded by the UK Science and Technology Facilities Council (STFC). The COSMA Memory Intensive service of DiRAC is hosted by the Institute of Computational Cosmology (ICC) at Durham University, specialising in workloads with large memory requirements, including cosmology.

Business challenge

Researching big cosmological questions, such as the meaning of dark matter and energy and the origin of the universe, takes huge amounts of computing power. Cosmology is one of the most demanding supercomputing applications and requires access to the most powerful supercomputers. Quite possibly, scientists may never have a computer large enough to perform the simulations that they need to understand the mysteries of the universe.

Improving High Performance Computing (HPC) performance, particularly in terms of sustainability and cost, is key to maximizing the impact of the cosmology simulations. This is also key to many other areas. Significant research is being conducted around making HPC more sustainable and moving closer to carbon neutrality. This case study describes the energy efficiency improvements achieved from upgraded technology that were discovered during tests conducted at Durham University.

Computing power is continually developing. The current workhorse COSMA system at Durham University, COSMA8, is primarily consisting of AMD Rome processors with a Milan extension and a total of around 70,000 cores, using Dell C6525 nodes. This system also includes Dell R6525 nodes with the newest AMD Genoa and Bergamo processors.

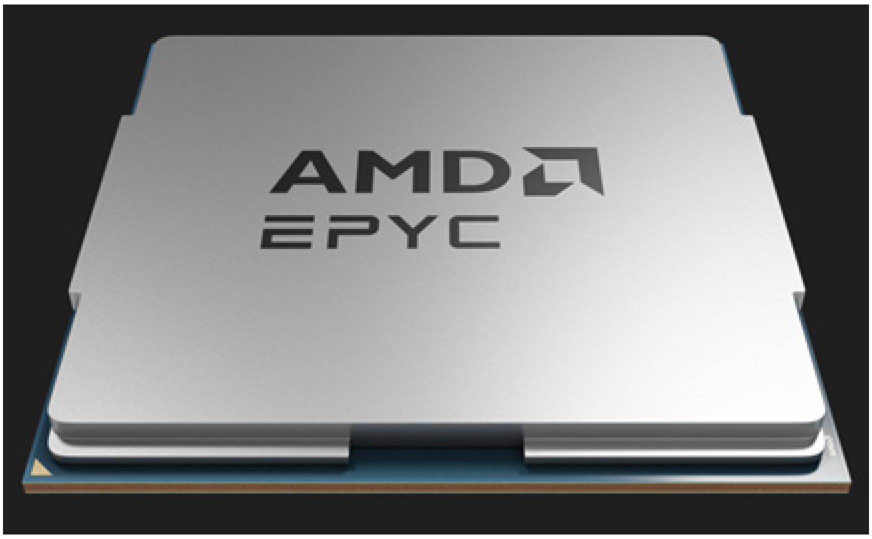

Case study overview

This case study compares the performance of these processors using two common cosmology codes, AREPO and SWIFT. AREPO is a conventional HPC code using a Message Passing Interface (MPI) paradigm with a single MPI process per core. SWIFT uses a task-parallelism-based hybrid scheme with a small number of MPI ranks each threaded with multiple tasking engines.

For this comparison, we considered the 7H12 Rome processor with 64 cores and a 2.6 GHz base frequency, the 7763 Milan processor with 64 cores and a 2.45 GHz base frequency, the 9654 Genoa processor with 96 cores and a 2.4 GHz base frequency, and the 9754 Bergamo processor with 128 cores and a 2.25 GHz base frequency. Each node contains two processors.

Our studies of several major cosmology codes show that in different situations, both Genoa and Bergamo can offer the best performance. However, one thing is clear—they offer a significant improvement, sometimes more than twice as fast than the Milan or Rome generation of processors. A key architectural feature of the newer Genoa and Bergamo processors is the availability of wide vector units (AVX512).

AREPO performance

An AREPO cosmology simulation has been run on a single node, evolving the universe from a redshift of approximately 1.8. The simulations are run for two hours, during which time they progress the evolution of the universe. The final redshift achieved provides an estimate of the compute performance. Figure 1 shows the time taken to progress the simulation by a given amount. Simulation progression is shown as a universe expansion factor rather than redshift, as this line arises in the figure. Figure 1 also demonstrates the time taken to reach a given expansion factor (redshift). The Bergamo and Genoa systems offer impressive per-node performance improvements over previous generations (with the Bergamo system being around 2.4 times faster than a Rome system).

For comparison, timings from an older Sky Lake system (COSMA7) are also demonstrated in this study, with 28 cores per node. Using 10 nodes and 256 cores shows that a single Bergamo system is around 1.76 times faster than the 10 Sky Lake nodes. We also show performance on a core to core basis, restricting Genoa and Bergamo simulations to 128 cores for a direct comparison with the Rome and Milan nodes. Here, when comparing half a Bergamo node with a full Rome node, the simulation is still around 1.67 times faster. With 128 cores, the Genoa node is around 1.8 times faster.

The performance gain from 128 to 256 Bergamo cores is around 1.4 times, primarily limited due to memory bandwidth limitations.

In Figure 1, lower values indicate better performance. Our results demonstrate a Bergamo node taking 0.42x less time to evolve the universe by a set amount than a Rome node. Unless stated in the legend, all simulations were run on a single node, making use of the number of cores given in the legend. Relative performance factors are shown

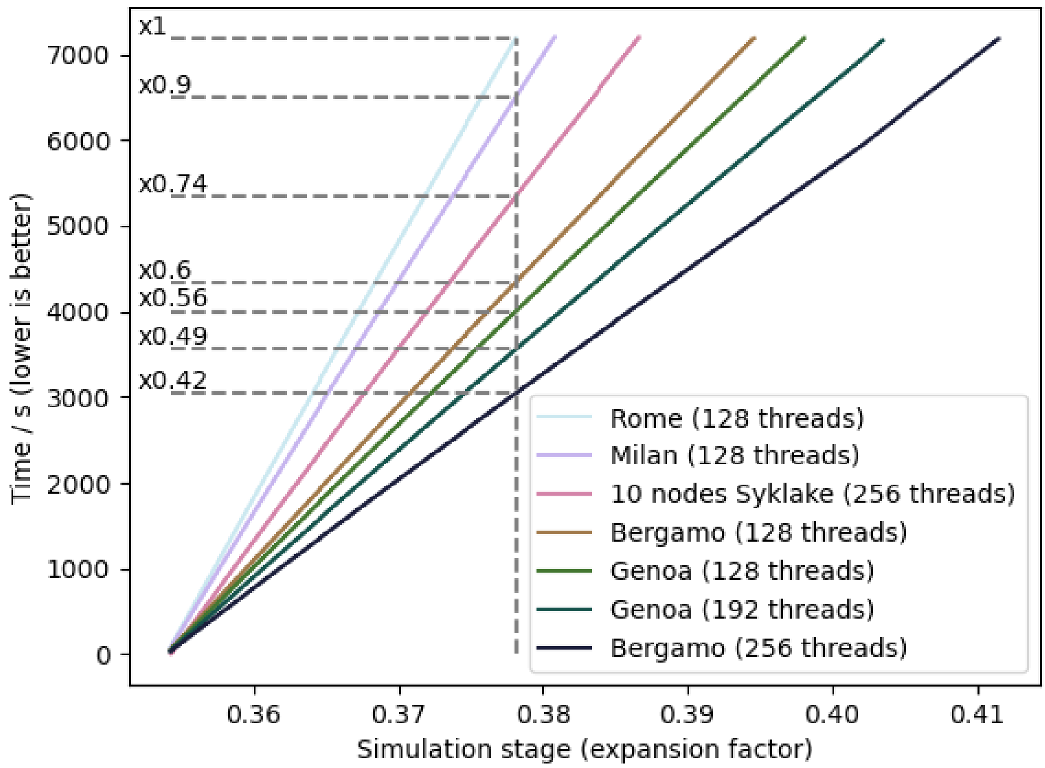

Figure 1. Simulated time as a function of expansion of the universe

We also compared the AOCC and ICX compilers, both of which are based on Clang. Figure 2 shows that there are no significant differences in performance, though AOCC generally offers a slight improvement.

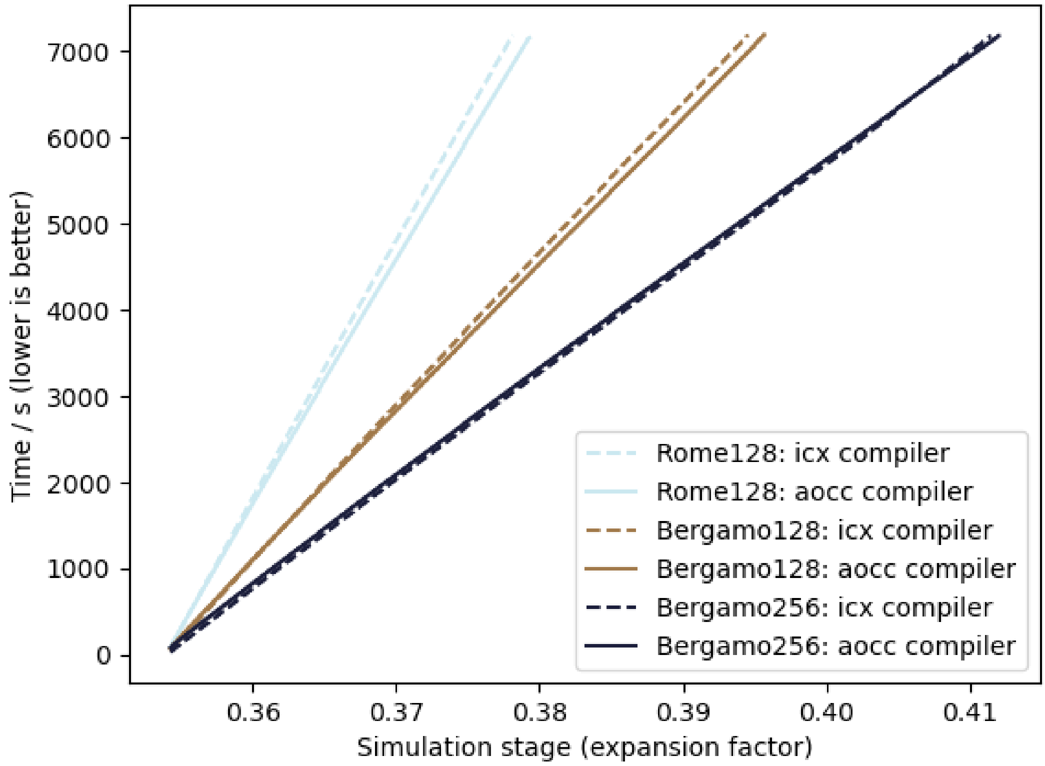

Figure 2. Comparing AOCC and ICX compiler performance for a node containing dual Rome or Bergamo processors using 128 or 256 cores

SWIFT performance

A SWIFT cosmology simulation including smoothed particle dynamics and gravity interaction. We used eight MPI Ranks per node, each associated with a NUMA domain. Each rank then has access to the 16, 24 or 32 cores within that domain. SWIFT is optimized to use AVX512 vector registers, which the Genoa and Bergamo processors contain, while the Rome and Milan processors do not.

The simulation is a scaled model of the full EAGLE simulation which is an exemplar project for the cosmology community. Our key performance metric is time to solution, or mean time per simulation iteration, which we show as a relative performance.

CPU type	Mean iteration time/ms	Cores per node	Base frequency / MHz	Max frequency/ MHz boost/MHz
Bergamo	4462	256	2.25	3.1
Genoa	4476	192	2.4	3.7
Milan	8146	128	2.45	3.5
Rome	9302	128	2.6	3.3

Figure 3 shows a significant per-node performance improvement (more than a factor of 2) for the Genoa and Bergamo processors for this workload, as compared with the Rome and Milan processors.

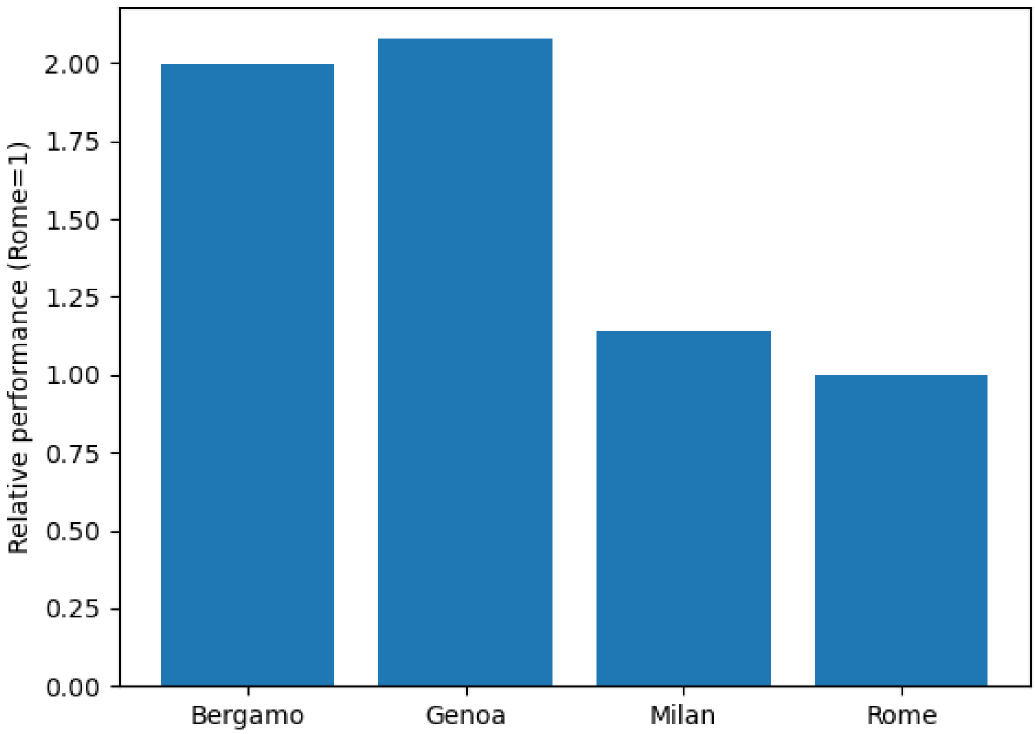

Figure 3. Relative performance (Rome time / Time taken) for a dual socket node containing these processors

If we consider energy consumption to achieve a given science output (which is directly related to CO2 emissions), our results indicate that Bergamo and Genoa show best performance, requiring less than 85 percent of the energy of the Rome processor to complete the task. This is demonstrated in Figure 4.

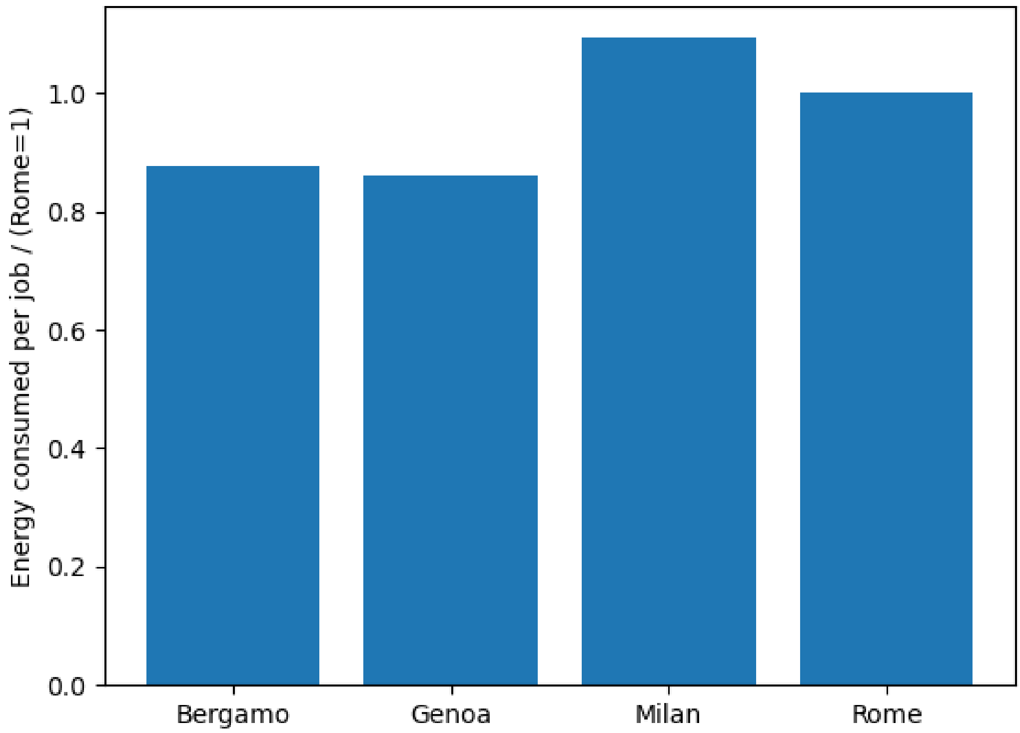

Figure 4. Relative energy use to achieve a given science goal (SWIFT simulation)

The relative performance per core and per clock MHz (such as if we scale the run times by the inverse of the number of cores and clock speeds) is also improved with Genoa and Bergamo processors compared to the Rome processor. In particular, the Genoa processor demonstrated an impressive improvement of nearly 1.5 times faster for a single core at a given frequency. The Bergamo processor has a lower level-3 cache which likely has some impact on performance.

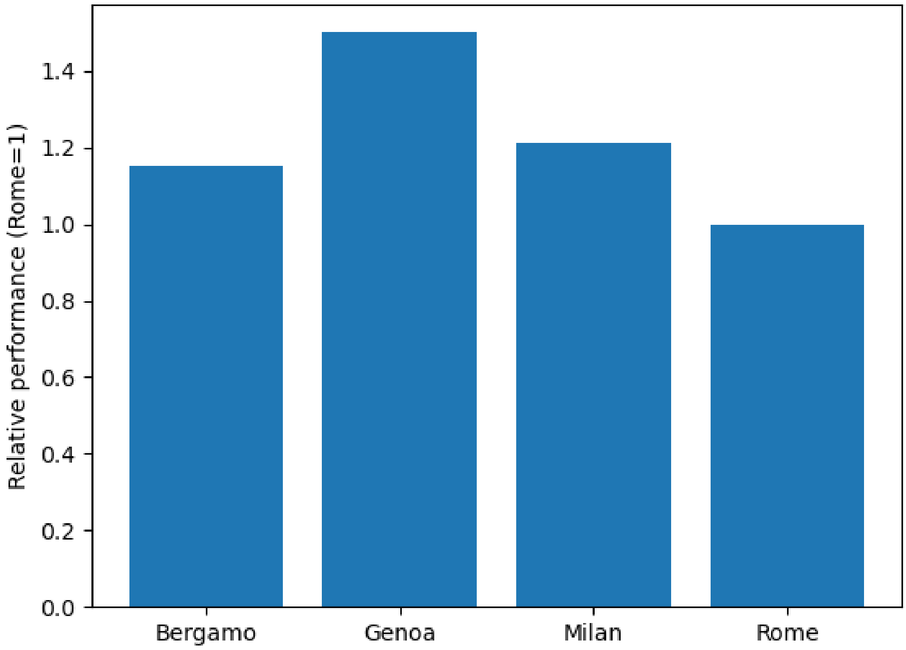

Figure 5. Relative performance (Rome=1) when scaled per core and per clock Mhz

Summary

For HPC, and particularly the memory intensive nodes of COSMA, the node cost is generally not dominated by the cost of the CPU. Therefore, significantly improving the performance of a node for small (if any) cost increases, creates major gains in cost-effectiveness.

Likewise, given that the power envelope of the newer processors is only around a factor of 30 percent greater, the improvement in science per watt significantly improved. We observed performance gains of more than 1.5 times per watt, yielding significant carbon dioxide reductions. As the need to expand computing capabilities comes up against the growing cost of required energy, and corporate and government directives around new efficiency standards, the efficiency gains this study shows offer an unprecedented opportunity for HPC systems to do more with less and contribute to ESG goals.

These new processors help solve the business challenge of the expense of HPC systems—both initial capital purchase, and ongoing operations costs—and improves sustainability efforts.

You can find a comprehensive list of documentation for this solution at the Dell Technologies Info Hub for High Performance Computing.

Dell Technologies welcomes your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by email