CRESTA

APPLICATION SOARS ABOVE PETASCALE AFTER TOOLS COLLABORATION

The HemeLB research group at University College London (UCL) has an exciting vision for HPC, one that will change the way neurosurgeons operate in the future and improve outcomes for patients. The group develops software to model intercranial blood flow – and collaboration with Allinea Software is helping them to address the challenges of application development at scale.

HemeLB is a Lattice-Boltzmann (LB) simulation code - the LB method is widely used in Computational Fluid Dynamics (CFD) in fields such as Aerospace or Oil and Gas reservoir modelling – but the UCL group are applying the model of fluid interactions to understanding how blood flows around cerebral vessels. Using a model of a patient's blood vessels generated by an MRI scan, simulation calculates pressure near points of weakness such as aneurysms.

The group aims to make the software perform well on larger machines – whereas presently the MRI scans are small reduced models, as large scale supercomputing grows in affordability and capability they envisage simulations being of clinical value – helping to decide the best clinical option within the right time frame for the patient in need.

Many of today's large supercomputers have around 100,000 CPU cores or more – and yet, of all the parallel applications that can run on the machines, precious few can use even close to that number of cores.

Without development, codes will run slower or may not run at all and crash.

HemeLB is one of the codes in the CRESTA project – www.cresta-project.eu CRESTA brings together four of Europe's leading supercomputing centres, with one of the world's major equipment vendors, two leading programming tools providers and six application and problem owners to explore how the Exascale challenge can be met.

The Scalable Software Challenge

Each new generation of supercomputers brings more parallelism – more CPU cores – and that brings greater challenges for application developers as they try to exploit the available computational power.

This problem is focussing minds. The supercomputers set to become reality early in the next decade – Exascale computers – capable of 10^{18} floating point operations (an Exaflop) – will be 50 times faster than today's largest machines, and have many millions of cores.

CRESTA is at the forefront of preparing applications such as HemeLB and tools such as Allinea's MAP and DDT for Exascale systems. The co-design process within CRESTA has allowed HemeLB to utilise significant fractions of the UK's flagship HPC system, ARCHER. Equally, this has facilitated the development of Allinea MAP and DDT, further extending their ability to profile and debug real applications at very large core counts", Dr Lorna Smith, EPCC and CRESTA project manager.

Taking HemeLB to the next level

Although a modern C++ code, with well-designed test and verification processes, development still has challenges. The CRESTA project is enabling the developers to work hand in hand with Allinea Software, to address those challenges. Software challenges become apparent when moving from the desktop to the clusters and supercomputers that HemeLB was designed for.

After learning about Allinea's tools for developers of parallel software during the co-design project, the team decided to explore what the tools could do for them.

The debugging and profiling tools, Allinea DDT and Allinea MAP, initially helped with development on the local systems – on the Linux desktops where most of the code is written and one of UCL's own clusters.

Research Officer, Sebastian Schmieschek quickly saw benefits. "Previously, we were developing on the desktop and bug fixing was mostly hard labour, while performance profiling was mostly done with a self-coded tool", says Schmieschek.

"In one case, we were faced with a crash and couldn't see why it was happening. We considered that a third party library might have a bug or it might have been a side-effect of another error. Allinea DDT helped us identify it."

Allinea DDT brought unique insight – integration of version control information helped to spot differences between previously reliable code and the current version. Schmieschek spotted an inconsistent vector being used due to a change elsewhere in control flow.

> "It was solved in about half an hour rather than the days it might have taken us without Allinea DDT"
>
> – Research Associate, Sebastian Schmieschek

In performance profiling, Allinea MAP soon detected that vectorization was being missed, and the impact of recent changes on run time.

"Allinea MAP allows us to get into the code quickly and identify performance problems. If we're using older, default input patterns, I can see how that impedes performance – and compare against new implementations that we're trialling."

Early Results

The true goal of the CRESTA project is to prepare applications for higher scale – and so the next step meant looking to a much larger system: the ARCHER UK National Supercomputing Service (http://www.archer.ac.uk). ARCHER is a Cray XC30 – its 3,008 nodes each hold two Intel ES-2697 Ivy Bridge CPUs with 12 cores, making a system total of 72,192 cores. It is one of the world's Top 20 machines by capability.

The machine is managed by the Edinburgh Parallel Computing Centre on behalf of EPSRC and other UK research councils.

While UCL's software had previously performed well at 32,000 cores, application tests did not perform well at 512 cores and it was hard to see where the problem was.

Allinea MAP explained the performance at a glance – and a simple adjustment in HemeLB's input parameters avoided an I/O bottleneck, enabling the application to scale successfully – improving performance on those cases by over 25%.

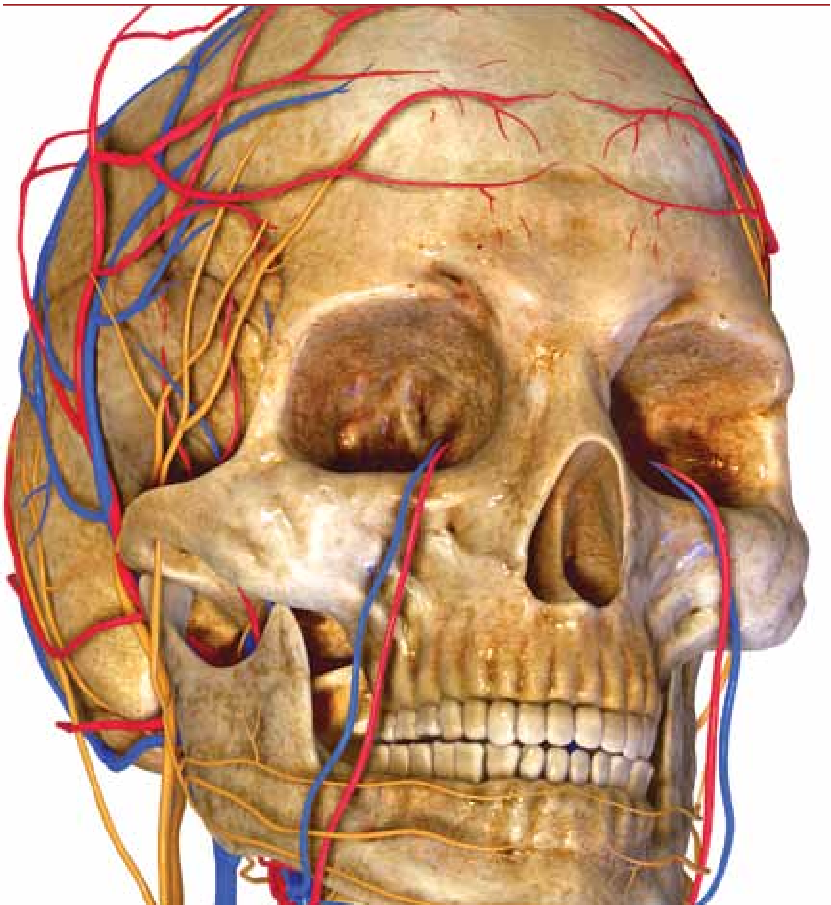

Heading for Scale

The team was ready for an even larger test case – one that would exercise the software, and ARCHER!

At 49,152 cores, the test run was set to let the developers understand how the application performed.

"We'd never been able to look at this many cores – and get a clear view of how the time was being used – we were keen to see it in Allinea MAP." Post-doctoral Researcher, Derek Groen, added.

But something happened – the application crashed! Running fine at 24,768 cores, it was natural to assume this would work at 49,152 too.

"I just had no idea. The crash was totally unexpected. I didn't know how I would diagnose or fix it at that scale – it was beyond anything I had tried to do before." says Groen. "I knew it was beyond printf. Allinea helped us straight away – they knew that if we could run the simulation with the debugger, we would find the problem."

Allinea DDT is the only parallel debugger that can handle that scale of problem – with incredibly fast responsiveness at extreme core counts and a user interface that has been designed to handle that scale in its stride.

At first, running in "offline mode" – where Allinea DDT runs unattended and records pertinent information about the crash in a simple HTML log – gave the team the head start they needed.

A follow-up interactive debugging session at the full size of the problem explored the findings in more depth and confirmed suspicions quickly. The debugger's automatic detection of outliers and consolidation of values and state across the whole application makes the task as straightforward as debugging on far smaller core counts.

"The partitioning library we were using was crashing – it's widely used but not many teams have taken it to beyond Petascale," added Groen. "We couldn't have been as confident about the source of the problem with anything else."

"We'd never been able to look at this many cores – and get a clear view of how the time was being used – we were keen to see it in Allinea MAP."

– Post-doctoral Researcher, Derek Groen

The Cycle of Tools

Now the bug has been resolved, it's back to Allinea MAP to look at the scalability again – as the team continue their drive to understand performance and scale their application to the largest machines available. The team are even looking at closer integration of debugging and profiling – as they explore performance artefacts identified with Allinea MAP inside Allinea DDT to explore how their data is being distributed and how the application workload is impacted by the partitioning strategies. This in turn is feeding ideas into the Allinea tools.

"The collaboration has enabled everyone to focus on the challenge of large-scale parallel scientific software development. Co-design with real application developers is key to ensuring that development tools are able to give the insight that those developers need to address their ambitions." says David Lecomber, CEO of Allinea Software, "We see at first-hand the impact on progress that the challenges make. Tools designed to handle scale are changing the way developers think about those hard challenges."

> "Scalable tools have been very helpful. The performance insight of Allinea MAP has been valuable – and Allinea DDT made problem solving a lot easier."
>
> – Research Associate, Sebastian Schmieschek

> "Getting HemeLB to scale to 50,000 ARCHER cores is a real achievement: we are thankful for the productive collaborations we enjoy with Allinea in the CRESTA project that have allowed us to reach these intoxicating heights, which are enabling us to study hemodynamics within the Circle of Willis for the first time."
>
> – Professor Peter Coveney, Director, Centre for Computational Science, UCL

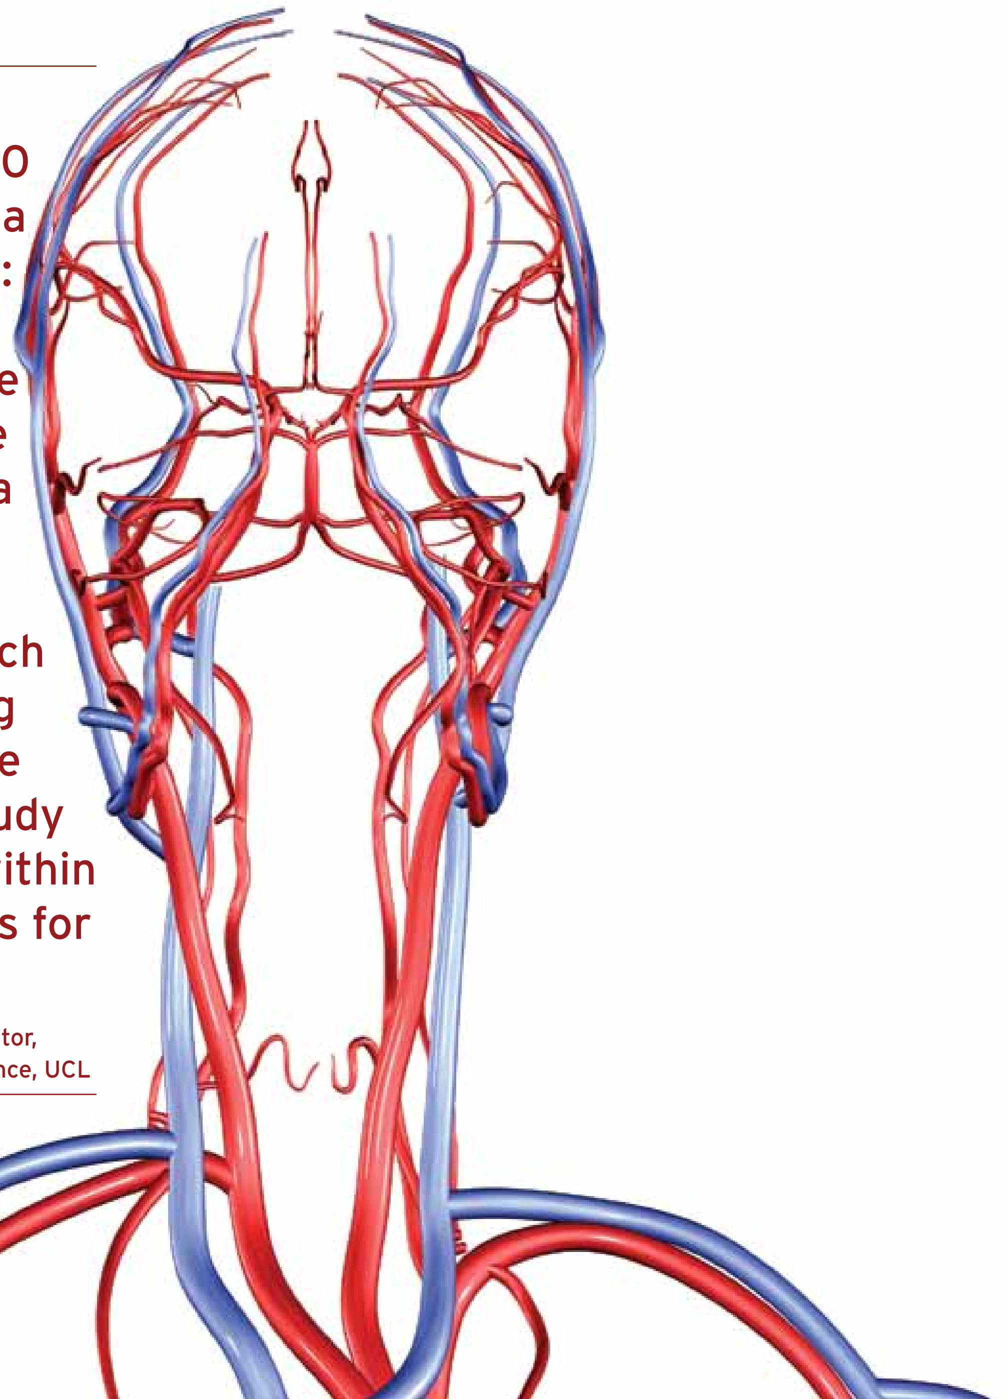

The CRESTA project receives funding from the EU's Seventh Framework Programme (FP7/2007-2013) under the grant agreements nr. 287703.