



THE CRESTA PROJECT

Professor Mark Parsons

Project Coordinator
EPCC Executive Director



Collaborative Research into Exascale Systemware, Tools and Applications - CRESTA

- Exascale Software
- 39 months -
- Leading European research centres
 - EPCC – Edinburgh
 - HLRS – Stuttgart
 - CSC – Edinburgh
 - KTH – Stockholm
- A world leading research centre
 - Cray UK – London
- World leading research centres
 - Technische Universität München (Munich)
 - Allinea Ltd – London



2014

owners and

ity – Abo, Finland

– Jyvaskyla,

ndon –UK

UK

Paris, France

many

Sweden

Germany

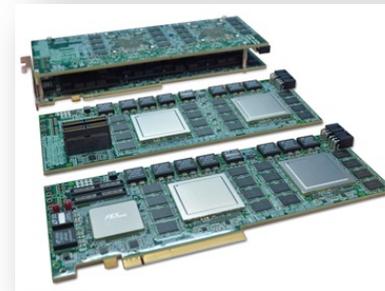


A hardware *and* software challenge

- Based on current technology roadmaps Exascale systems will be impossible to build below 50MW
 - GPUs and Xeon Phi plus traditional multi-core microprocessors, memory hierarchies and even with embedded interconnects cannot get to 20MW
- The Exascale experiments are showing that modern data flow hierarchies inside a well balanced computer are not enough
- The solution is to use more parallelism, but we don't know how to use such parallelism
- But these solutions are not enough
 - Today's leader scales to 92 million threads and 526MW at the Exascale
- Slower better balanced cores mean parallelism at the 500 million – 1 billion thread scale

Lowering the energy requirements

- Most energy efficient supercomputer today is Japanese Shoubo system from RIKEN
 - 7.031 GigaFLOPS per watt
- Would require
 - 3.3 **BILLION** cores to deliver an ExaFLOP
 - 142 Megawatts
- Current target by 2025
 - 20 Megawatts
 - 200-500 million cores
- BUT ... we have no idea how to use such systems for modelling and simulation
- Exascale is a **SOFTWARE** challenge



At the Exascale software leaves algorithms behind

- Few mathematical algorithms are designed with parallelism in mind
 - ... parallelism is “just a matter of implementation” is the mind-set
- This approach generates many more components as components are custom-built for each application
 - ... but the years of development and testing are reluctant to change and users are reluctant to write new code as writing takes place
- HPC is at a “fork in the road”
 - Without fundamental changes many areas will be limited ... and systems will be limited
 - But it's not just a case of “more” – it's much more difficult
- This doesn't just apply to Exascale – it's apparent at the Petascale too
- CRESTA tackled this challenge head on

Software and how we
model and simulate
remain the key
Exascale challenges

Key principles behind CRESTA

- Two strand project
 - Building and exploring appropriate
 - Enabling a set of key *co-design* ap
- Co-design was at the heart of the
 - provided guidance and feedback to
 - integrated and benefited from this develop
- Employed both incremental and disruptive solutions
 - Exascale requires
 - Particularly true for
 - Solutions
- Project has
 - and new sc
 - and case s

Disruptive approach

- Work with co-design applications to consider alternative algorithms
- Crucial we understand maximum performance before very major application redesigns undertaken

Incremental approach

- Through optimisations, performance modelling and co-design application feedback
- Look to achieve maximum performance at Exascale and understand limitations e.g. through sub-domains, overlap of compute and comms

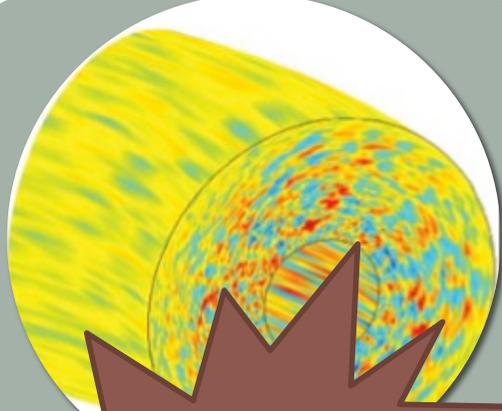
Where are we on the road to Exascale?

- The call text from 2010 said:

The target is to develop a small number of advanced computing platforms with potential for extreme performance (100 petaflop/s in 2014 with potential for exascale by 2020), as well as optimised application codes driven by the computational needs of science and engineering and of today's grand challenges such as climate change, energy, industrial design and manufacturing, systems biology. These platforms should rely on vendors' proprietary hardware or on COTS.

- The fastest machine in Europe is 6.3 Petaflops Linpack today
- The fastest machine in the World is 33.9 Petaflops Linpack today

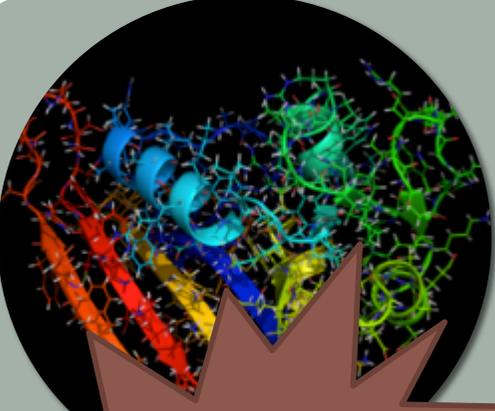
Co-design applications



Exascale 3D decomposition and visualisation

Elm
Gy
fur
Simul
plasma behavior in
large scale fusion reactors

An almost complete code restructuring
Radical reduction of memory consumption per core

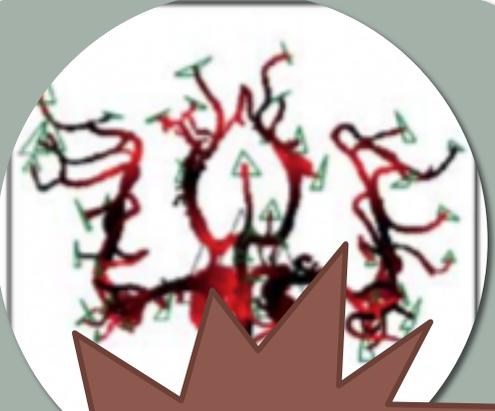


Improved implementations + code reorg for task parallelism + ensemble engine

G
biological

- computational material and drug design
- 10M atom simulation

Coupling strong scaling techniques with ensemble scaling



Physics for Exascale + performance / scaling of LB

H
S
ovascular
LB

Medical simulations to help surgeries
Brain aneurysm simulation

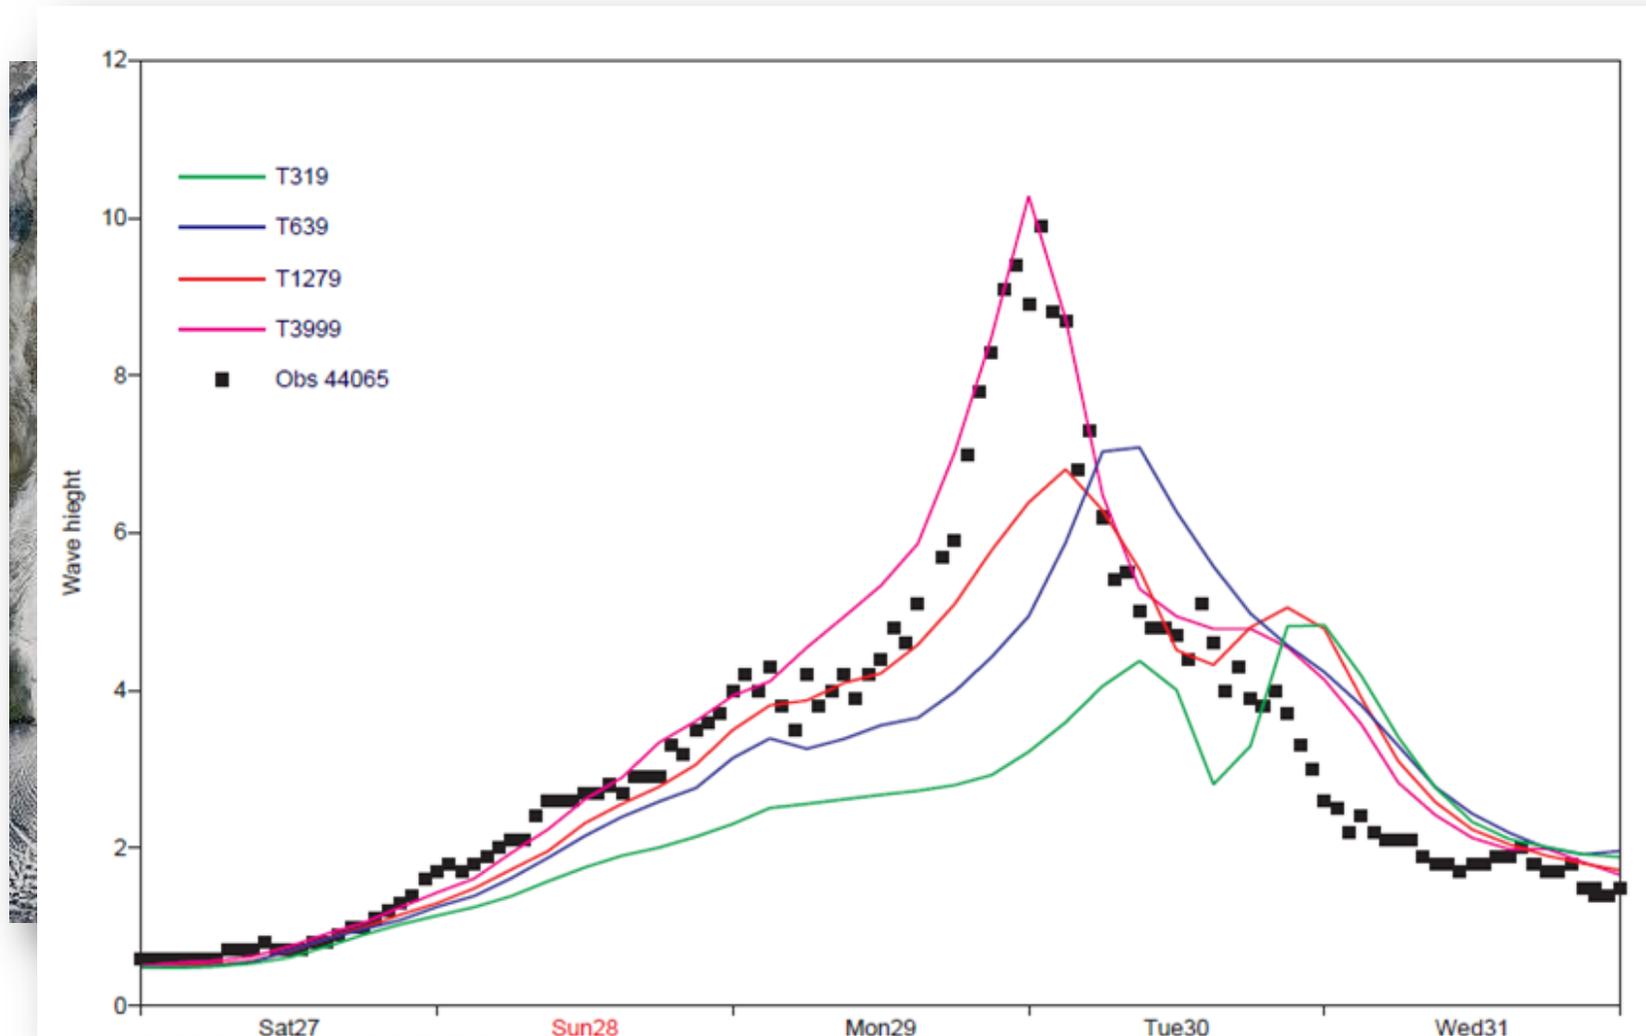
Pre- and post-processing and load balancing
Hybridisation, restructuring

Co-design applications

The image consists of three vertical panels, each with a circular image at the top and a starburst callout in the middle. The panels are set against a light grey background with rounded corners. A double-headed arrow is at the bottom of the panels.

- Panel 1 (Left):**
 - Image: A circular image showing a satellite view of Earth with a grid overlay. Text in the image reads "European Centre for Medium-Range Weather Forecasts".
 - Starburst: "PGAS approaches + cubic grid + OmpSs experiments"
 - Text below: "IES...ction", "The... weather", "ive lives", "Simulating the trajectory of hurricane Sandy"
 - Text at bottom: "Acceleration", "Task-graph based parallelization", "New communication models"
- Panel 2 (Middle):**
 - Image: A circular image showing a 3D visualization of a globe with a color-coded atmospheric simulation.
 - Starburst: "GPGPU engine + AMR + Exascale mesh partitioner"
 - Text below: "S...n Mira", "plant cooling", "simulations"
 - Text at bottom: "Adaptive mesh refinement", "acceleration"
- Panel 3 (Right):**
 - Image: A circular image showing a cross-section of a turbine or flow simulation with a color gradient from blue to red.
 - Starburst: "Development stopped – OpenFOAM is NOT an Exascale code"
 - Text below: "fine", "space and time", "Wind turbines, hydroelectric power plants", "Francis pump turbine simulation"
 - Text at bottom: "Linear solver optimization"

Why high resolution modelling is important



East Coast of ...

October 2012

Achievements

- We showed how software co-design can work
 - Driven by a general understanding of the scale of parallelism that Exascale hardware will deliver
- Identified many challenges – not just with parallelism but also I/O performance, tools, libraries – software and systemware
- Made tools advances which also benefit Petascale
- Shown that some codes e.g. OpenFoam will never run at the Exascale in their present form
- Given code owners the space to explore the Exascale and plan how to respond to it e.g. the ECMWF scaling team
- A key success was been to create awareness of the challenges – so that all of us can properly plan – **SOFTWARE IS THE KEY**