THE CRESTA PROJECT

Professor Mark Parsons

EPCC Executive Director
Associate Dean for e-Research

Alison Kennedy

EPCC Executive Director

CRESTA epcc

Collaborative Research into Exascale Systemware, Tools and Applications

- Exascale Sc
- 3.5 years - € 2014
  - Leading E wners and
    - EPCC –
    - HLRS – ity – Abo, Finland
    - CSC – E - Jyvaskyla,
    - KTH – S
  - A world lea ndon –UK
    - Cray UK UK
  - World lead Paris, France
    - Technisc any
      (Vampir) eden
    - Allinea Lt Germany

CRESTA epcc

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 emb t get to 20MW
- The Exascale expo data flow hierarchies inside ll balanced comp
- The solution wil processors with communications
- But these solutions **INC** of parallelism
  - Today's leader scales to 92 million cores and 526MW at the Exascale
- Slower better balanced cores means parallelism at the 500 million – 1 billion thread scale

Hardware is leaving software behind – many codes scale badly today

CRESTA epcc

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 much components are custom-built for eac
  - … but the years o e and users are reluctant t g takes place
- Strongly believe
  - Without fundar y areas will be limited … and n ems
  - But it's not just a case s – it's much more difficult
- This doesn't just apply to Exascale
  - It's apparent at the Petascale if you look
- And we have a huge skills and tools gap …

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

CRESTA epcc

Key principles behind CRESTA

- Two strand project
  - Building and exploring appropriate sy…
  - Enabling a set of key *co-design* appli…
- Co-design is at the heart of the proje…
  - provide guidance and feedback to the…
  - integrate and benefit from this development … cyclical process
- Employing both incremental and disruptive solutions
  - Exascale requires bot… …proaches
  - Particularly true for app… …limit of scaling today
  - Solutions will …
- Committed to …

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

CRESTA epcc

Co-design applications

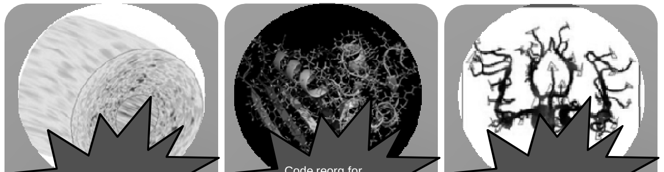

Exascale 3D decomposition and visualisation

Elm…
Gy… …bulent fu…
Simula… …asma behavior in large s… …ale fusion reactors

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

Code reorg for Exascale + ensemble engine + FMM electrostatics

G…
… …ological
- …omputational 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
Med… …simulations to help surg…ries
Brain aneurysm simulation

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

CRESTA epcc

Co-design applications

Parallel radiation co-model + cubic grid + OmpSs experiments

IFS … …ction
Th… …weather … …ve lives
Simu…ng the trajectory of hurricane Sandy

Acceleration
Task-graph based parallelization
New communication models

GPGPU engine + AMR + Exascale mesh partitioner

N… …n Mira
S… …lant cooling simu…ons

Adaptive mesh refinement acceleration

Development stopped – OpenFOAM is **NOT** an Exascale code

… fine … …ace and time
Wi… …urbines, hydroelectric power plants
Francis pump turbine simulation

Linear solver optimization

CRESTA epcc

What has CRESTA achieved?

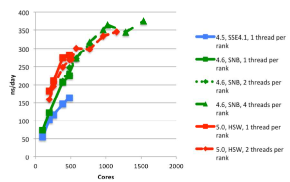

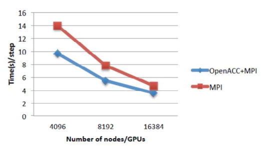

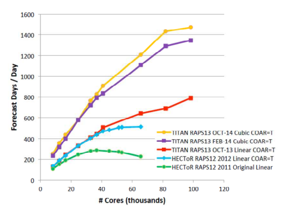

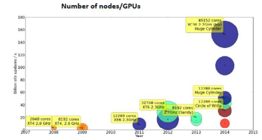

- Lots of scaling results ….
- But that wasn't really the point of CRESTA

CRESTA epcc

Achievements

- Shown 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
- Key result has been with regard to software development planning
  - Partners now understand the true scale of the Exascale challenge
  - One partner has created an in-house scaling team – direct result of CRESTA

Conclusions

- CRESTA has shown the scale of the Exascale application challenge
- It has moved some codes towards this challenge and published the results for others to learn from
- Far too few projects like CRESTA exist
- It is clear that many DISRUPTIVE INNOVATIONS will be needed to model and simulate on Exascale systems
- We still need to re-think many algorithms – and continue to build engagement with mathematics community
- There is a clear need for many more CRESTA-like projects … but who will be brave enough to fund them?

Thank you