CRESTA

Collaborative **R**esearch into **E**xascale **S**ystemware, **T**ools and **A**pplications

Dr Mark Parsons

EPCC Executive Director
Associate Dean for e-Research

The University of Edinburgh

CRESTA

- **C**ollaborative **R**esearch into **E**xascale **S**ystemware, **T**ools and **A**pplications
- Developing techniques and solutions which address the most difficult challenges that computing at the exascale can provide
- Focus is predominately on software not hardware.
- European Commission funded project
  - FP7 project
  - Projects started 1st October 2011, three year project
  - 13 partners, EPCC project coordinator
  - €12 million costs, €8.57 million funding

www.cresta-project.eu

CRESTA

Partnership

- Consortium
  - Leading Eu … owners and
    - EPCC –
    - HLRS –
    - CSC – E … sity – Abo,
    - PDC – S … – Jyvaskyla,
  - A world lea … ondon –
    - Cray UK … UK
  - World lead … – Paris, France
    - Technisc … many
      (Vampir)
    - Allinea Lt

CREST

Key Principles

- Two strand project
  - Building and exploring appropriate *systemware* for exascale platforms
  - Enabling a set of key *co-design* applications for exascale
- Co-design is at the heart of the project. Co-design applications:
  - provide guidance and feedback to the systemware development process
  - integrate and benefit from this development in a cyclical process
- Employing both incremental and disruptive solutions
  - Exascale requires both approaches
  - Particularly true for applications at the limit of scaling today
  - Solutions will also help codes scale at the peta- and tera-scales
- World leading hardware vendor as project partner – Cray
- Committed to open source interfaces, standards and new software

CREST

Choosing what to study

- As we all know the Exascale research domain is enormous
- A €12million project can only tackle a small part of the overall problem
- We decided very early not to focus on hardware but on the problems of massive heterogeneous parallelism
- Research selection process
  - Capture and categorisation of problems
  - Each partner asked to indicate their interests
  - Topics selected based on
    - Amount of interest
    - Complementary skills
    - Viability within funding constraints

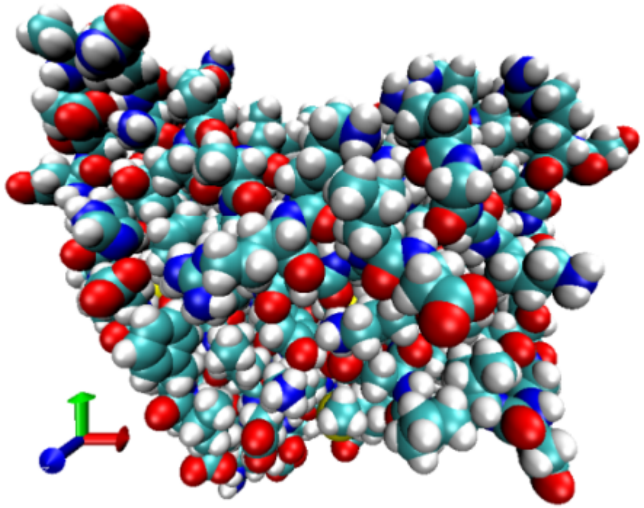

- In supercomputing it is too easy to tie a ribbon around new hardware and much more difficult to tie a ribbon around people

CREST

Co-design Applications

- Exceptional group of six applications used by academia and industry to solve critical grand challenge issues
- Applications are either developed in Europe or have a large European user base
- Enabling Europe to be at the forefront of solving world-class science challenges

Application	Grand challenge	Partner responsible
GROMACS	Biomolecular systems	KTH (Sweden)
ELMFIRE	Fusion energy	ABO/ JYU (Finland)
HemeLB	Virtual Physiological Human	UCL (UK)
IFS	Numerical weather prediction	ECMWF (International)
OpenFOAM	Engineering	EPCC / HLRS / ECP
Nek5000	Engineering	KTH (Sweden)

CREST

Example of incremental and disruptive approaches

- FFTs are a challenge at Exascale because
  - Very large number of HPC applications use them
  - Distributed memory parallel FFT is already a major performance issue today – we accept some FFTs will not scale further
- Two approaches:

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

CRESTA

Systemware

- Software components required for grand challenge applications to exploit future exascale platforms
- Underpinning and cross cutting technologies
  - Operating systems, fault tolerance, energy, performance optimisation
- Development environment
  - Runtime systems, compilers, programming models and languages including domain specific
- Algorithms and libraries
  - Key numerical algorithms and libraries for exascale
- Debugging and Application performance tools
  - World leader’s in Allinea’s DDT, TUD’s Vampir and KTH’s perfminer
- Pre- and post- processing of data resulting from simulations
  - Often neglected, hugely important at Exascale

CRESTA

Systemware: integrated set of software components

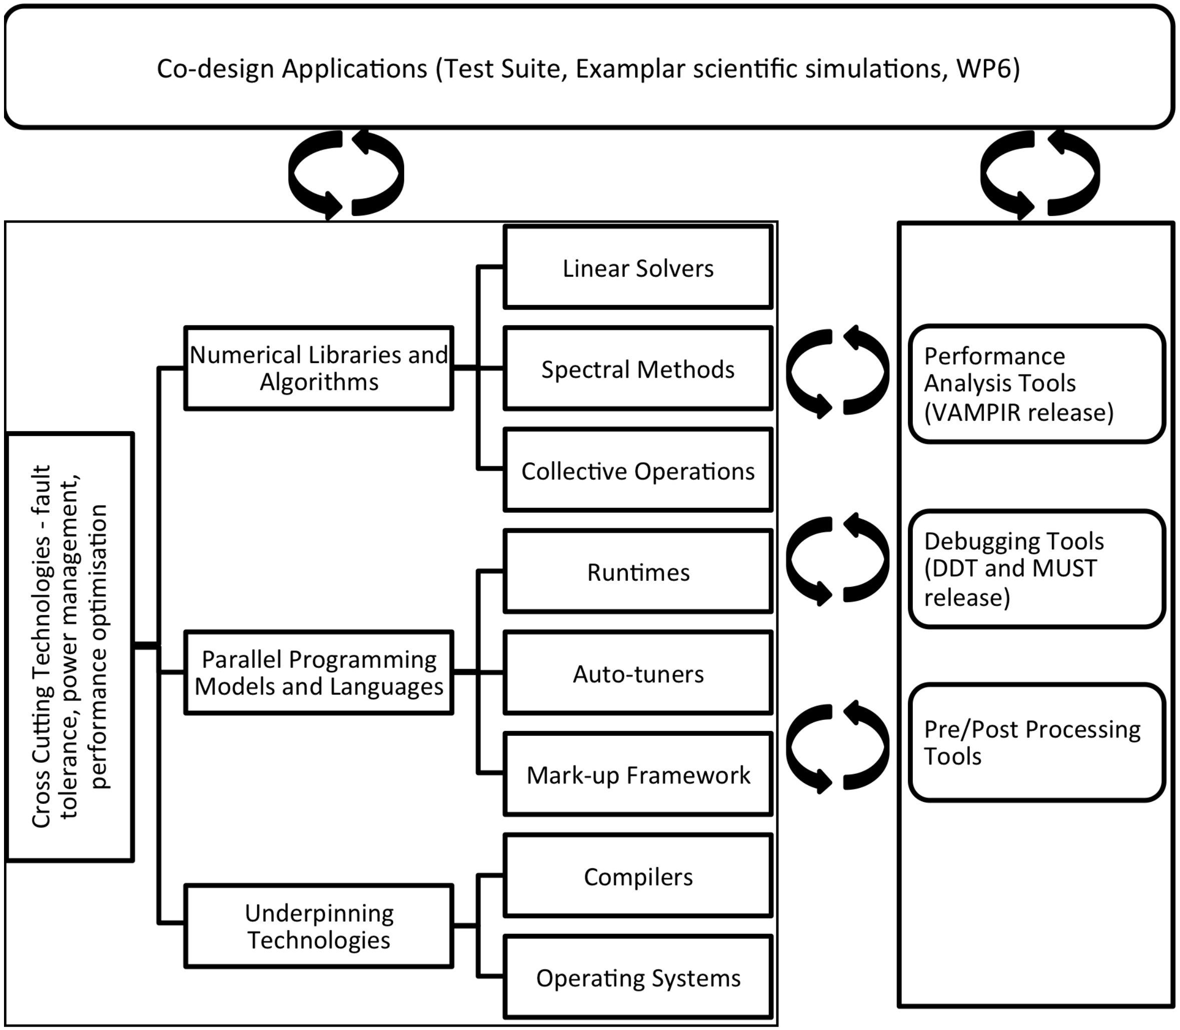

CREST

Conclusion

- CRESTA focuses on software not hardware
- Far too little money is being spent on software worldwide at present
- We need both incremental and disruptive approaches
- Exascale computers challenge our basic understanding of how we model and simulate numerical problems
  - CRESTA can only look at a small number of problems
  - There needs to be much more engagement from the mathematics and computer science communities
- Exascale should enable previously impossible simulations
  - There are a huge number of opportunities - but we need the right tools
- By focussing on co-design of applications and systemware we expect CRESTA to play a key role in the exascale roadmap

www.cresta-project.eu