ELEMENT

Element Project Workshop

Welcome and Introduction

Professor Mark Parsons, EPCC Director

ExCALIBUR and ELEMENT

- **E**xascale **C**omputing: **A**lgorithms & **I**nfrastructures **B**enefitting **U**K **R**esearch
- Funded by UK Government Strategic Priorities Fund
- £50m over 4 years
- One of a small number of Exascale Projects funded to date
- ExCALIBUR is focussed on Exascale software challenges faced by
  - UK Met Office
  - UK Atomic Energy Authority (represented by Culham Centre for Fusion Energy)
- Project led jointly by the UK Met Office and EPSRC
- ELEMENT is one of 10 High Priority Use Case projects funded in early 2020

ExCALIBUR High Priority Use Cases

1. ELEMENT: Exascale Mesh Network
2. Materials And Molecular Modelling Exascale Design And Development Working Group
3. Gen X: ExCALIBUR working group on Exascale continuum mechanics through code generation
4. Exascale Computing for System-Level Engineering: Design, Optimisation and Resilience
5. Massively Parallel Particle Hydrodynamics for Engineering and Astrophysics
6. BASE: Benchmarking for AI for Science at Exascale
7. EXA-LAT: Lattice Field Theory at the Exascale Frontier
8. ExaClaw: Clawpack-enabled ExaHyPE for heterogeneous hardware
9. ExCALIBUR-HEP: ExCALIBUR and High Energy Physics
10. Turbulent Flow Simulations at the Exascale: Application to Wind Energy and Green Aviation

ELEMENT

- Led by EPCC, The University of Edinburgh in partnership with
  - University of Cambridge, Imperial College London, University of Exeter and Swansea University
- Focusses on the high priority use case o
  - Developing highly scalable solutions to crea
  - Partitioning efficiently to minimise load im
  - Ensuring meshes are of sufficient quality to
- Objectives
  - To build a community around meshing practi … ve network
  - Undertake a small number of proof on concept studies
  - Publish a Vision Paper which will inform a Strategic Research Agenda

> Strategic Research Agenda will cover the full meshing workflow at the Exascale including mesh generation, adaptation, partitioning and visualisation

Workshop Day 1

Morning Session: Introduction to ELEMENT & Exascale System Technologies		
10:30-11:00	Mark Parsons, EPCC	Introduction to ELEMENT & The UK Exascale Project
11:00-11:20	Simon McIntosh-Smith, University of Bristol	The evolution of computer architecture and its implications for meshing
11:20-11:40	Bernhard Homoelle, SVA	Memory technologies, what comes next?
11:40-12:00	Nic Dube, HPE	Exascale and Beyond: Supercomputing Heterogeneity
12:00-12:50	Breakout groups & discussion	
12:50-13:00	Summary of breakout groups	

ELEMENT

Workshop Day 1

Afternoon Session: Parallel mesh generation		
14:00-14:20	Trevor Robinson, Queen's University Belfast	Applying Simulation Intent to Parallel Mesh Generation
14:20-14:40	Christos Tsolakis, Polykarpos Thomadakis and Nikos Chrisochoides, Old Dominion University	Exascale-Era Parallel Adaptive Mesh Generation and Runtime Software System Activities at the Center for Real-Time Computing
14:40-15:00	Christophe Geuzaine, University of Liege	Towards (very) large scale finite element mesh generation with Gmsh
15:00-15:20	Tzanio Kolev, Lawrence Livermore National Laboratory	Large-scale Finite Element Applications on High-Order Meshes
15:20-15:40	ELEMENT project talk	Meshing towards the Exascale
15:40-16:30	Breakout groups & discussion	
16:30-16:45	Summary of breakout groups	

Workshop Day 2

Morning Session: End user stories		
10:30-10:50	Paul Cusdin, Renault F1	Practical CFD and Meshing. An Inconvenient Truth
10:50-11:10	Paolo Adami, Rolls-Royce	A view from Rolls Royce
11:10-11:30	Carolyn Woeber, Pointwise	A Mesh Generation Perspective on Exascale CFD
11:30-11:50	ELEMENT project talk	Translating high order spectral/hp element methods from academia to industry
11:50-12:45	Breakout groups & discussion	
12:45-13:00	Summary of breakout groups	

ELEMENT

Workshop Day 2

Afternoon Session: Geometry definition, CAD interaction and mesh adaptivity		
14:00-14:20	Henry Bucklow, ITI	Geometry for mesh generation
14:20-14:40	Xevi Roca, Barcelona Supercomputing Centre	Meshing from CAD vision: curved adaption to geometry and solution
14:40-15:00	Adrien Loseille, INRIA	Parallel anisotropic mesh adaptation in complex geometries and extreme anisotropy
15:00-15:20	Bob Haimes, MIT	A lightweight geometry kernel for distributed mesh generation and adaptation
15:20-15:40	ELEMENT project talk	Mesh Adaptation towards the Exascale
15:40-16:30	Breakout groups & discussion	
16:30-16:45	Summary of breakout groups	
16:45-17:00	**Conclusion - Summary of Workshop**	

ELEMENT

The UK Exascale Project

ELEMENT Workshop – October 2020

Professor Mark Parsons

EPCC Director
EPSRC Director of Research Computing

The Exascale era – international context

Country or Region	Timescale	Detail
China	2020 / 2021	Little known at present – updated CPU plus accelerator as per Sunway
Japan	2020	Fugaku : based on A64FX Arm processor
USA	2021/2	Frontier : based on AMD EPYC CPU + AMD GPU Aurora : based on Intel A21 CPU + Intel GPU
Europe	2020 2023/4	Pre-Exascale hosting sites chosen (Finland / Spain / Italy) Future Exascale systems will use Europe's own CPU

Building the case for Exascale computing in the UK

- Europe Union is investing heavily in Exascale through EuroHPC
- UK scientists mustn’t be left behind
- In mid-2018 UK Government decided it needed a strategy
- Established Exascale Project Working Group to develop Business Case for investment
- Parallel review of e-Infrastructures by UKRI led by EPSRC
- Supercomputing Science Case developed to understand scientific needs

epcc

Supercomputing Science Case themes

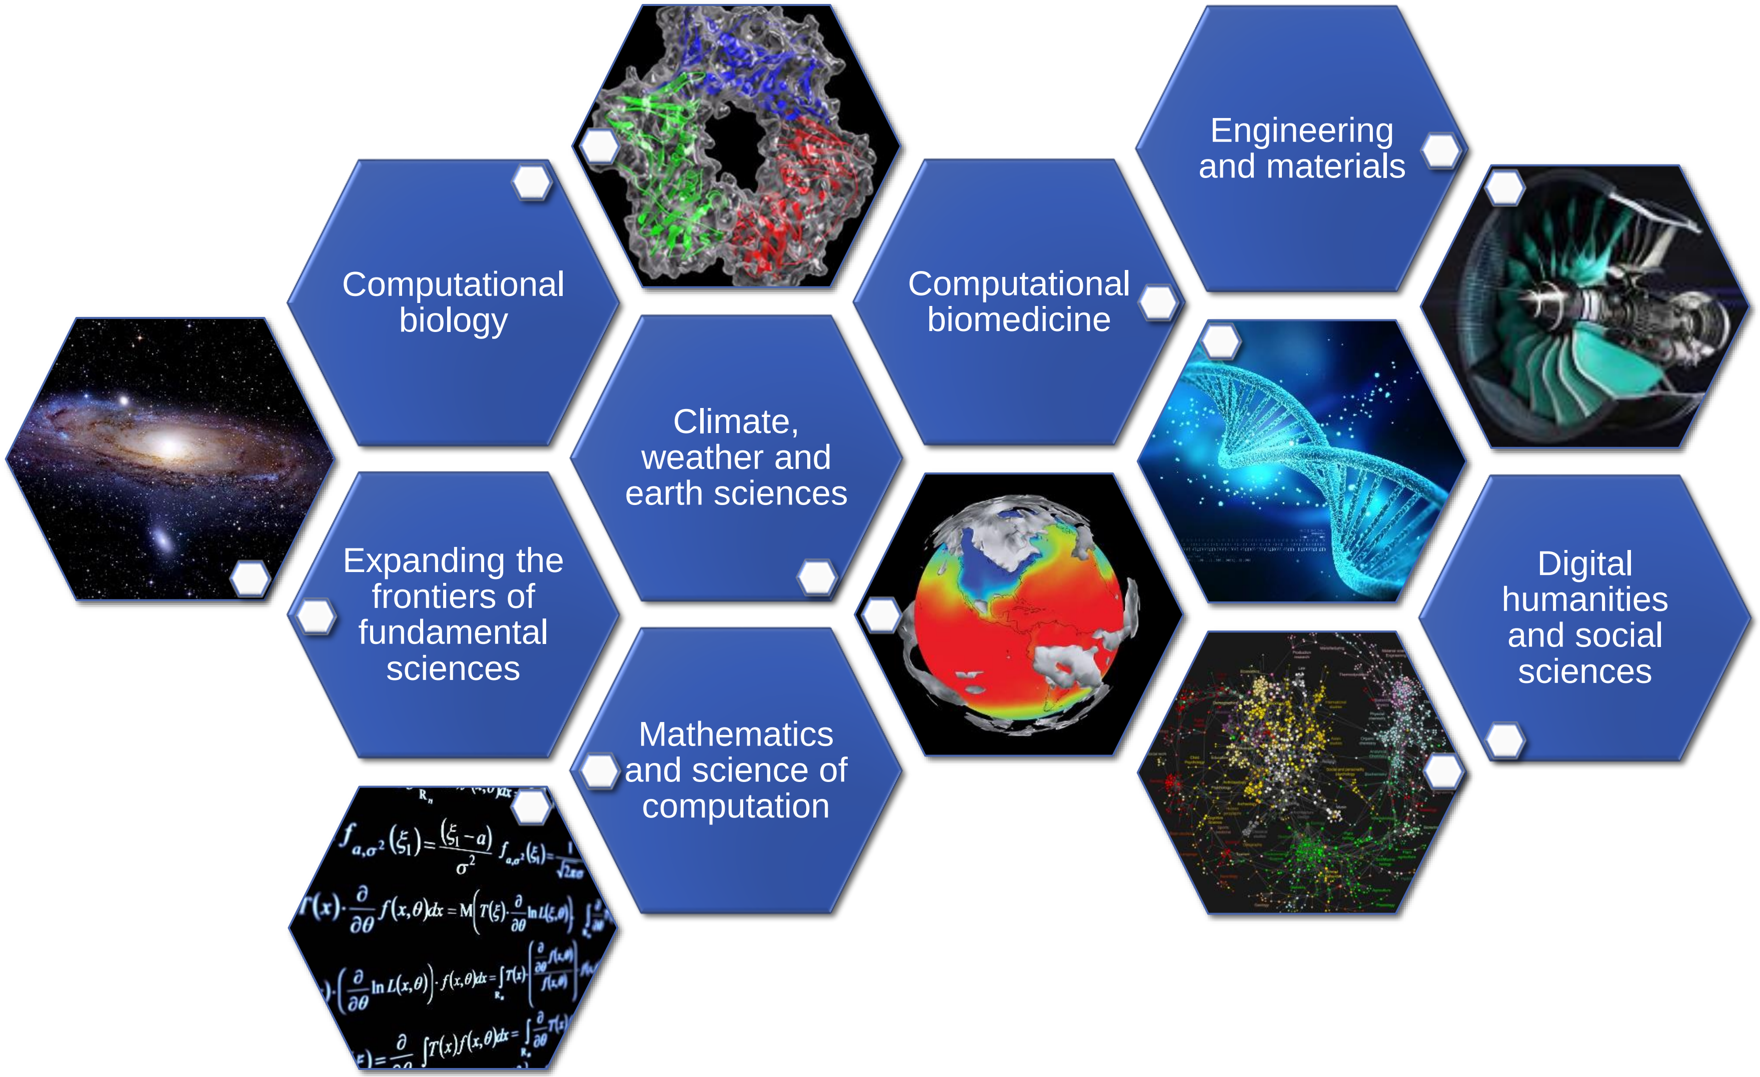

Exascale Project Specific Requirements from Government

- System should support both **traditional Modelling & Simulation** and **Artificial Intelligence / Deep Learning** applications
  - Technology choices may be impacted by this
  - But future technologies blur the distinction
- System should support both **scientific user communities** and **industry users**
  - A greater focus is proposed with regard to industry use for research
  - Pay-per-use production access will be supported
  - Specific support for SMEs
- System should be **operational by 2023**

epcc

Infrastructure takes time and money

£20m – New computer room
£8m – 30MW additional power

Opening Dec 2020

epcc

Aiming for Net Zero

Bilston Glen Colliery, 670m, 15.0C, Minewater
Monktonhall, 866m, 25.5C, Rock
Lady Victoria, 768m, 18C, Minewater

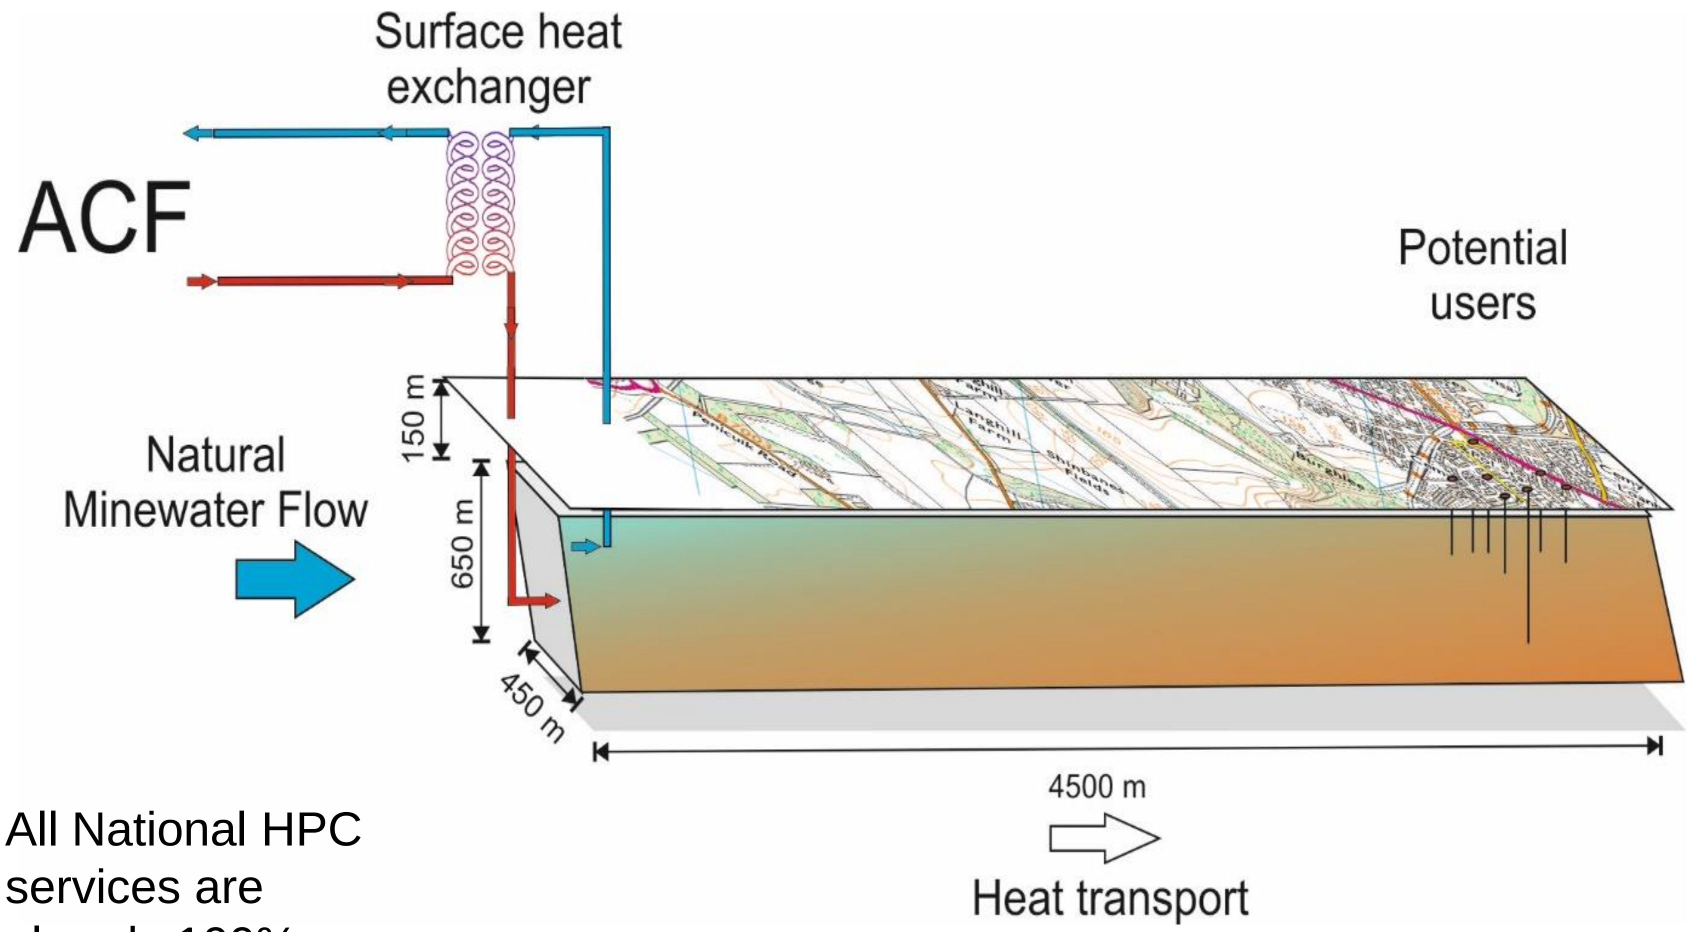

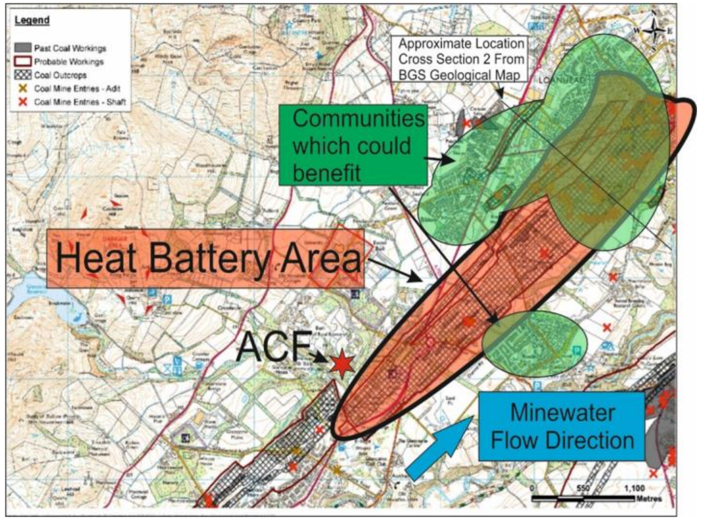

All National HPC services are already 100% Green Electricity

epcc

On the road to Exascale …

- USA's SUMMIT system was the world's fastest supercomputer from June 2018 – June 2020 according to Top500 HPL benchmark
- 2,414,592 CPU cores and 27,000 GPUs
- R_{peak} = 201 Petaflop/s
- Power consumption of 13 Megawatts
- To reach the Exascale with this technology
  - 12 million CPU cores + 68,000 GPUs
  - 65 Megawatts
- … very high levels of parallelism

… but at ISC in June 2020 Japan's Fugaku system took the crown

epcc

… Fugaku takes the crown

- Fugaku became the world's fastest supercomputer in June 2020 with a cores-only approach based on the Fujitsu A64FX Arm CPU
- Processor developed in long-term co-design (10 years) with Japanese computational science community led by Riken CCS
- 7,299,072 Arm CPU cores
- 4.866 Petabytes of RAM
- R_{peak} = 513.9 Petaflop/s
- Power = 28.3 Megawatts
- Single precision > 1 Exaflop

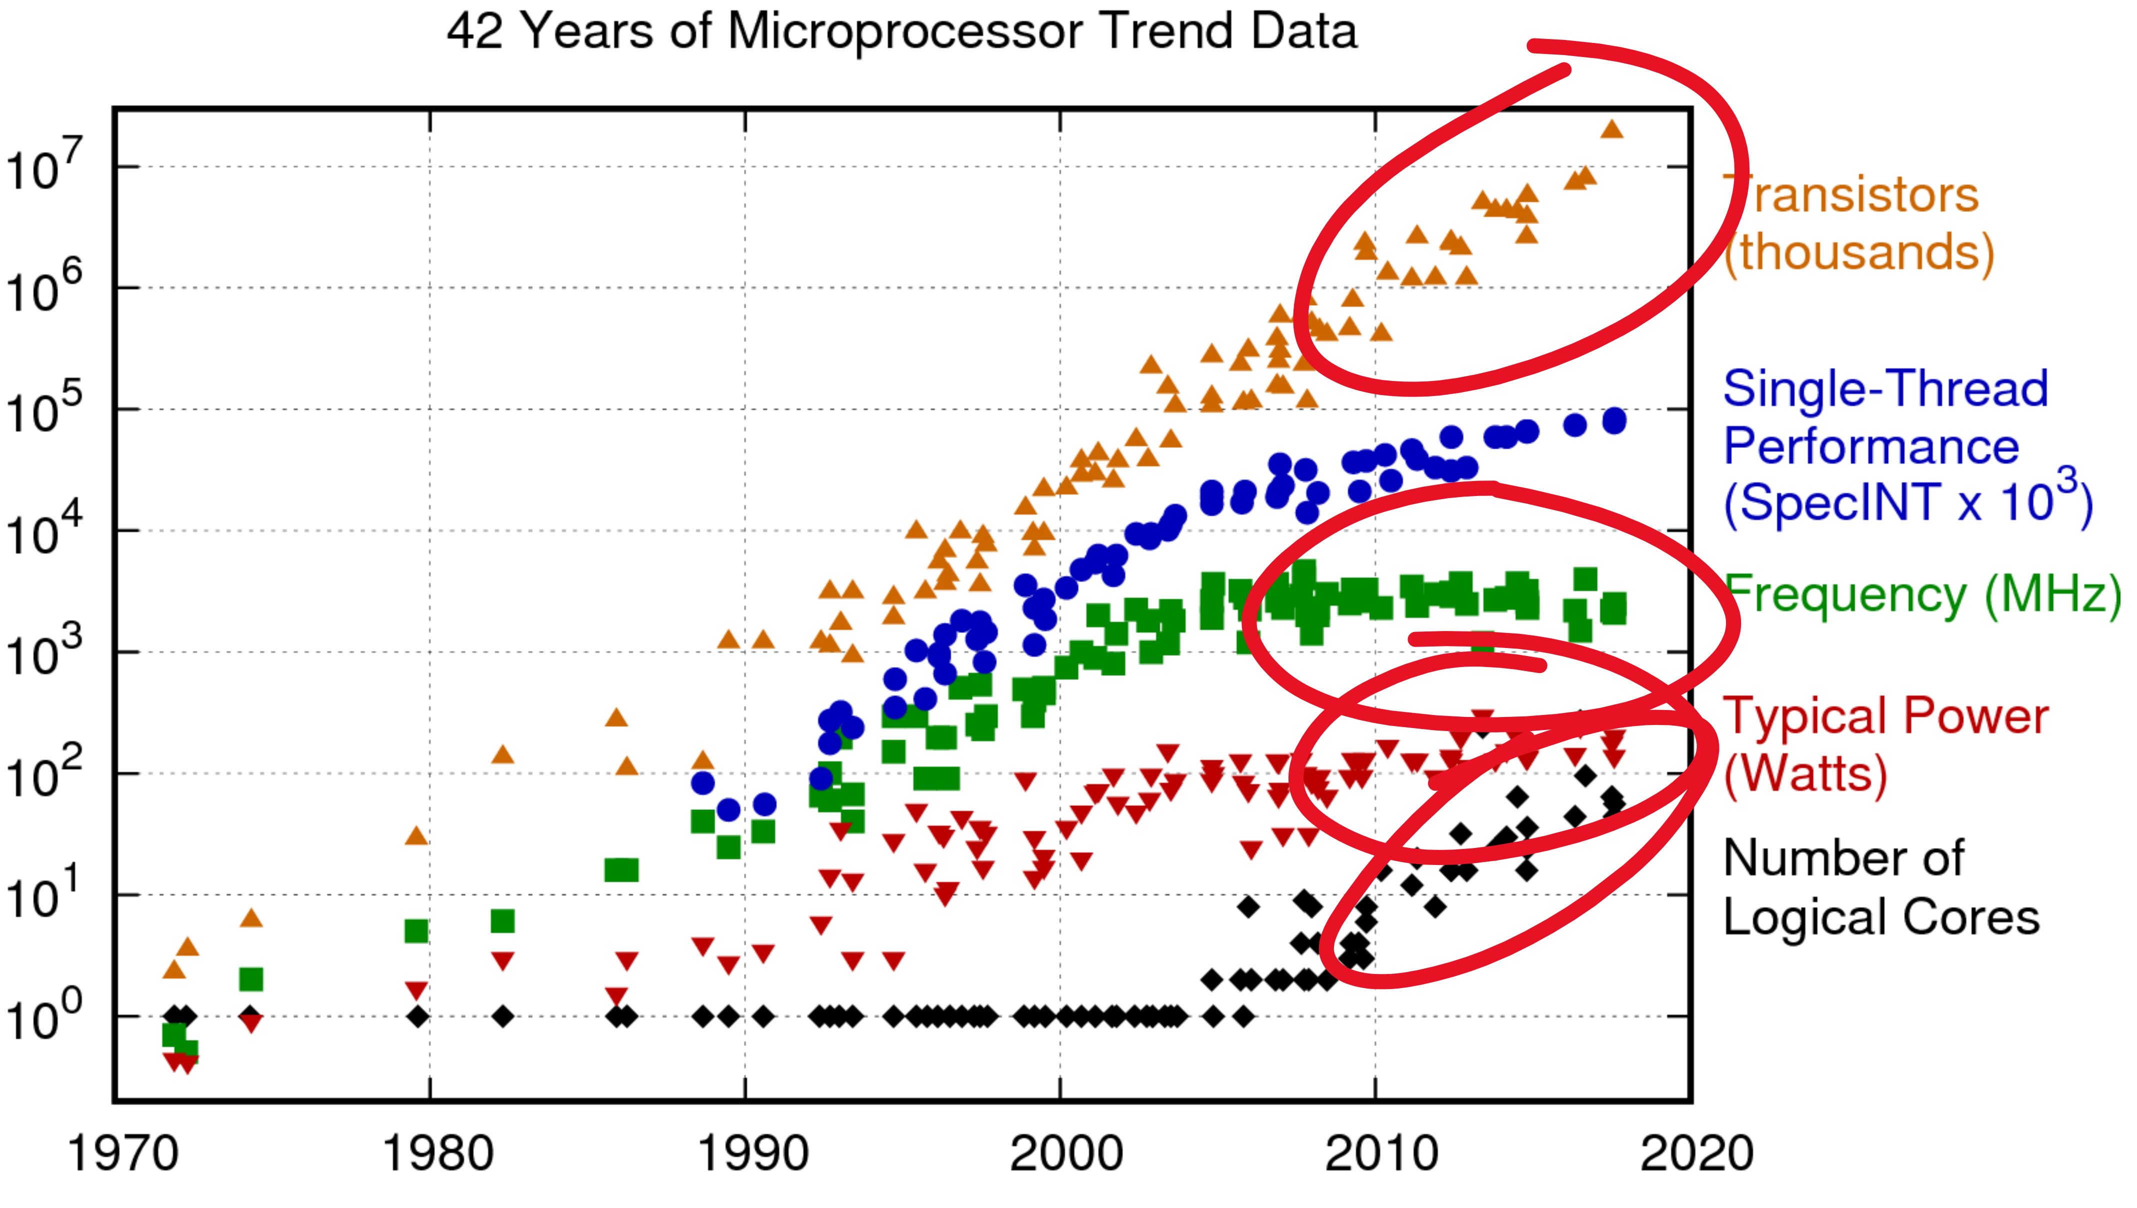

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Why GPUs are good at AI – and how CPUs are catching up

- Keys operation in computer graphics are matrix multiplications
- All GPUs support General Matrix Multiplication (GEMM) operations of the form **D = (A x B) + C**
- For computer graphics these are generally low precision FP16 calculations
- It turns out that for many AI Deep Learning algorithms – which use GEMM operations – low precision is good enough
- It’s the ability to do lots of calculations in parallel that is key
- CPUs focus on excellent FP64 arithmetic – although many designs have now added 16-bit (often the BFloat16 format) and GEMM operations (often called MMA)

Technology – two routes to the Exascale

- Cores-only route
  - Favoured by Japan
  - Evidence this approach leads to longer lifetimes of systems
  - Hardware support for AI increasing – main focus traditional simulation
  - Larger power requirements and physical dimensions
- Cores plus accelerator route
  - Favoured by the USA
  - Traditional multi-core processors coupled to accelerator
  - Sweet spot seems to 10 cores per GPU – pushes towards 1 socket + 4 GPU blades at the Exascale
  - Strong AI performance – traditional simulation more challenging

epcc

Technology – Japanese versus American model

- Japanese model has attractions but difficult to sell to Government
- Lower peak performance but much longer science lifetime

HPCG Results (June 2019 and June 2020)

Rank	Site	Computer	Cores	HPL Rpeak	HPL Rmax (Pflop/s)	Top500 Rank	HPCG (Pflop/s)	Rmax to Rpeak	Fraction of Rmax
1	Riken CCS JAPAN	Fugaku	7,299,072	513.9	415.5	1	133.7	80.9%	32.17%
2	DOE/SC/ORNL USA	Summit	2,414,592	200.8	148.6	2	2.9	74.0%	1.97%
5	Riken CCS JAPAN	K-Computer	705,024	11.3	10.5	22	0.6	93.2%	5.73%

epcc

Technology – key insights from recent Exascale vendor briefings

- High Bandwidth Memory is coming
  - Many Exascale blades include HBM
  - Some designs have no DRAM at all
- Four-way competition for CPUs and/or GPUs
  - Intel versus AMD versus Arm versus NVIDIA
- GPUs are getting ever more powerful
  - We're already seeing the market broaden
- Cabinet energy densities are rocketing
  - Today's 80-100KW cabinets will be eclipsed by cabinets at 300KW+
  - Density of blades is a key battleground
- Multicore CPUs are also getting AI Deep Learning features

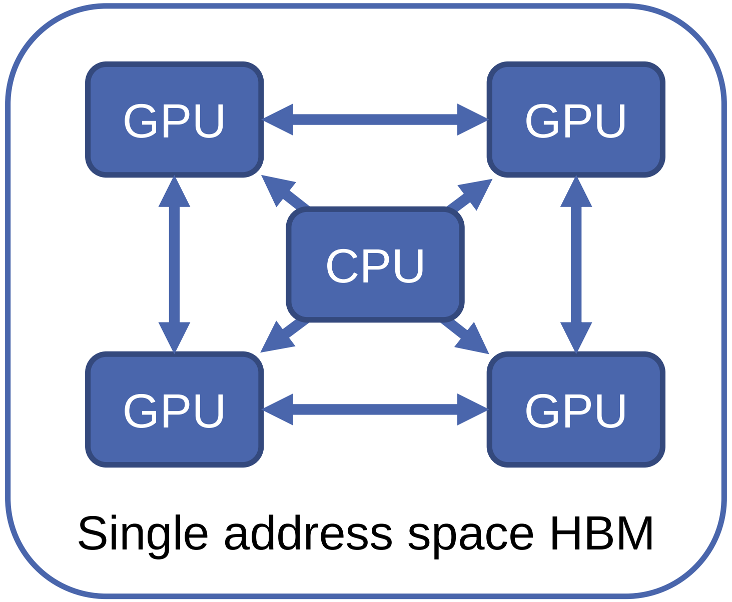

Exascale systems – How parallel? How large?

- For a cores only approach it should be possible to get to 1 Exaflop theoretical peak with 4-5 million cores
- A cores plus GPU approach will reduce the number of cores but overall parallelism will increase as GPUs have much higher parallelism – better for AI less so for simulation
- 1 Exaflop power requirements range from 20MW to 160MW
- Size of systems is highly dependent on density of blades and cabinet design
  - Number of cabinets ranges from circa. 60 to over 800!
- Key metric is always usefulness for both science and industry

Conclusions

- The next generation of blade designs focus on moving data to and from the processor faster than ever before
- This is just as relevant for modelling and simulation applications as it is for AI deep learning applications
- Supercomputing **and data science** computing are converging
- Need to focus on data processing performance **not flops** in future
- Exascale is driving much of this convergence **but so are AI applications that use large amounts of data**
- For the most demanding problems **the Cloud Hyperscaler world and Supercomputing are converging**