HECToR

Andrew Jones

nag VP HPC and **HECToR Outreach**

March 2010, HPCC-USA, Newport RI

N A G : what we do

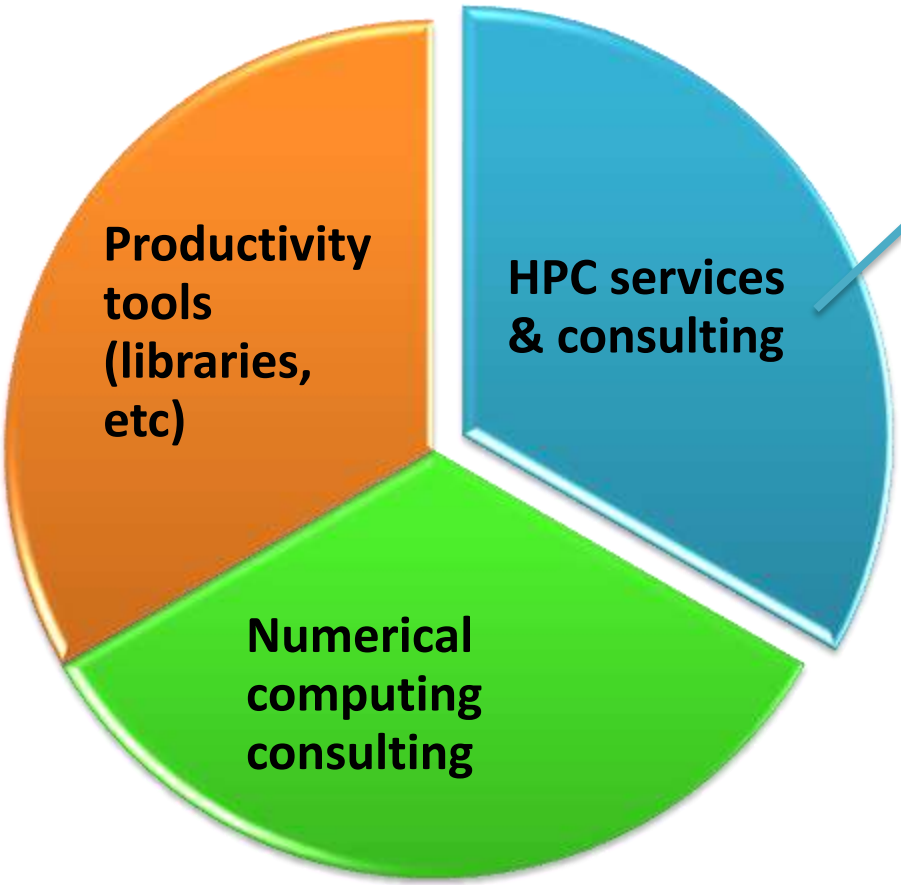

High Performance Computing Support

- Independent evaluation of HPC technology
- Application optimization and renovation
- Training and documentation

HECToR: UK national HPC service

Independent HPC advisor to procurement activity

Market survey, benchmarking, technology watch, etc

Academic user community

Helpdesk

(N A G)

Computational Science & Engineering Support & Training: 120 person years

System management & hosting (UoE HPCx)

Cray Centre of Excellence

Phase 2a: 208TF Cray XT4; [Phase 2b: ~340TF XT6]

Phase 3 (Unknown)

2007-2013, £113M

HECToR

RESEARCH COUNCILS UK

NB: No longer wholly accurate, but easier to talk through hardware with this one chart than a series of slides.

Revised HECToR Upgrade Roadmap

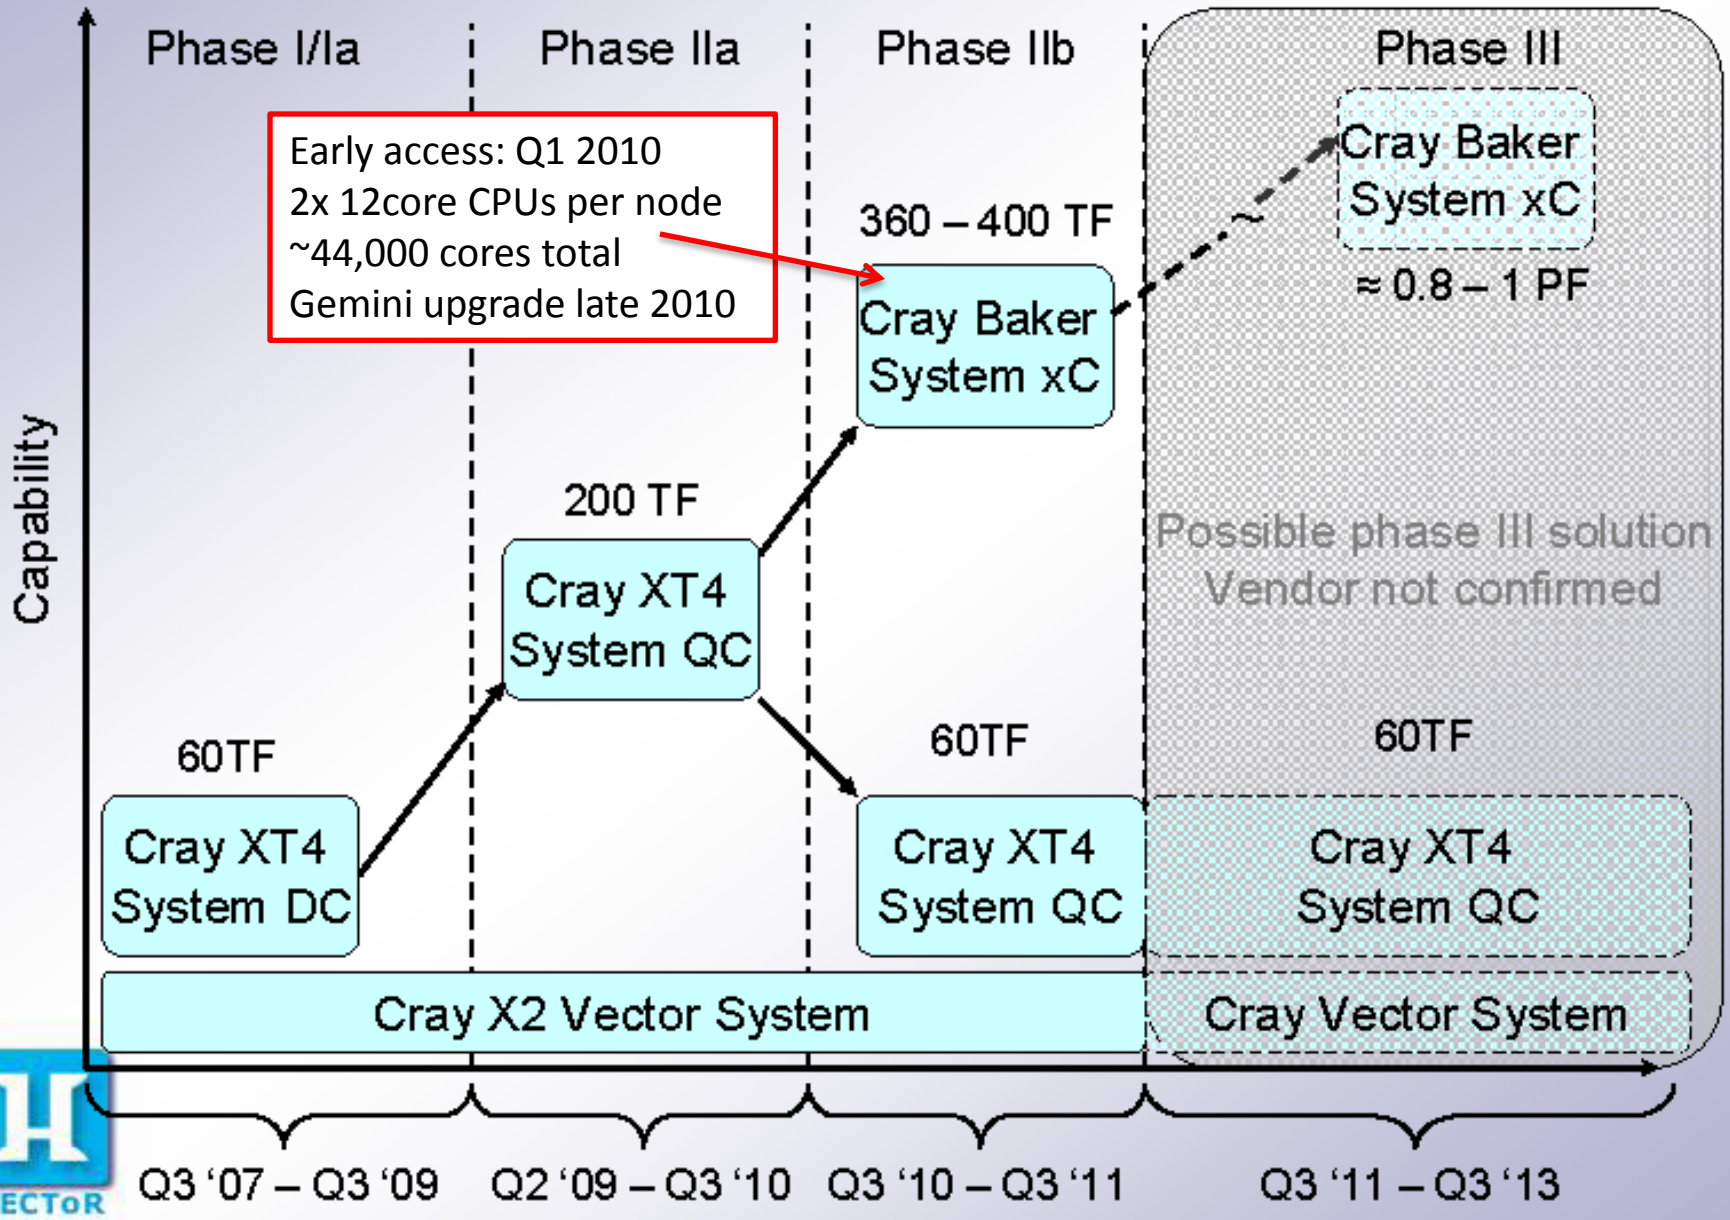

HECToR: capability CSE support

Comparable scale investment in expertise

- Computational Science & Engineering
- Capability support model

User requirements capture during HECToR procurement identified the need for strong CSE support

HECToR

RESEARCH COUNCILS UK

Users driving the need for CSE support ...

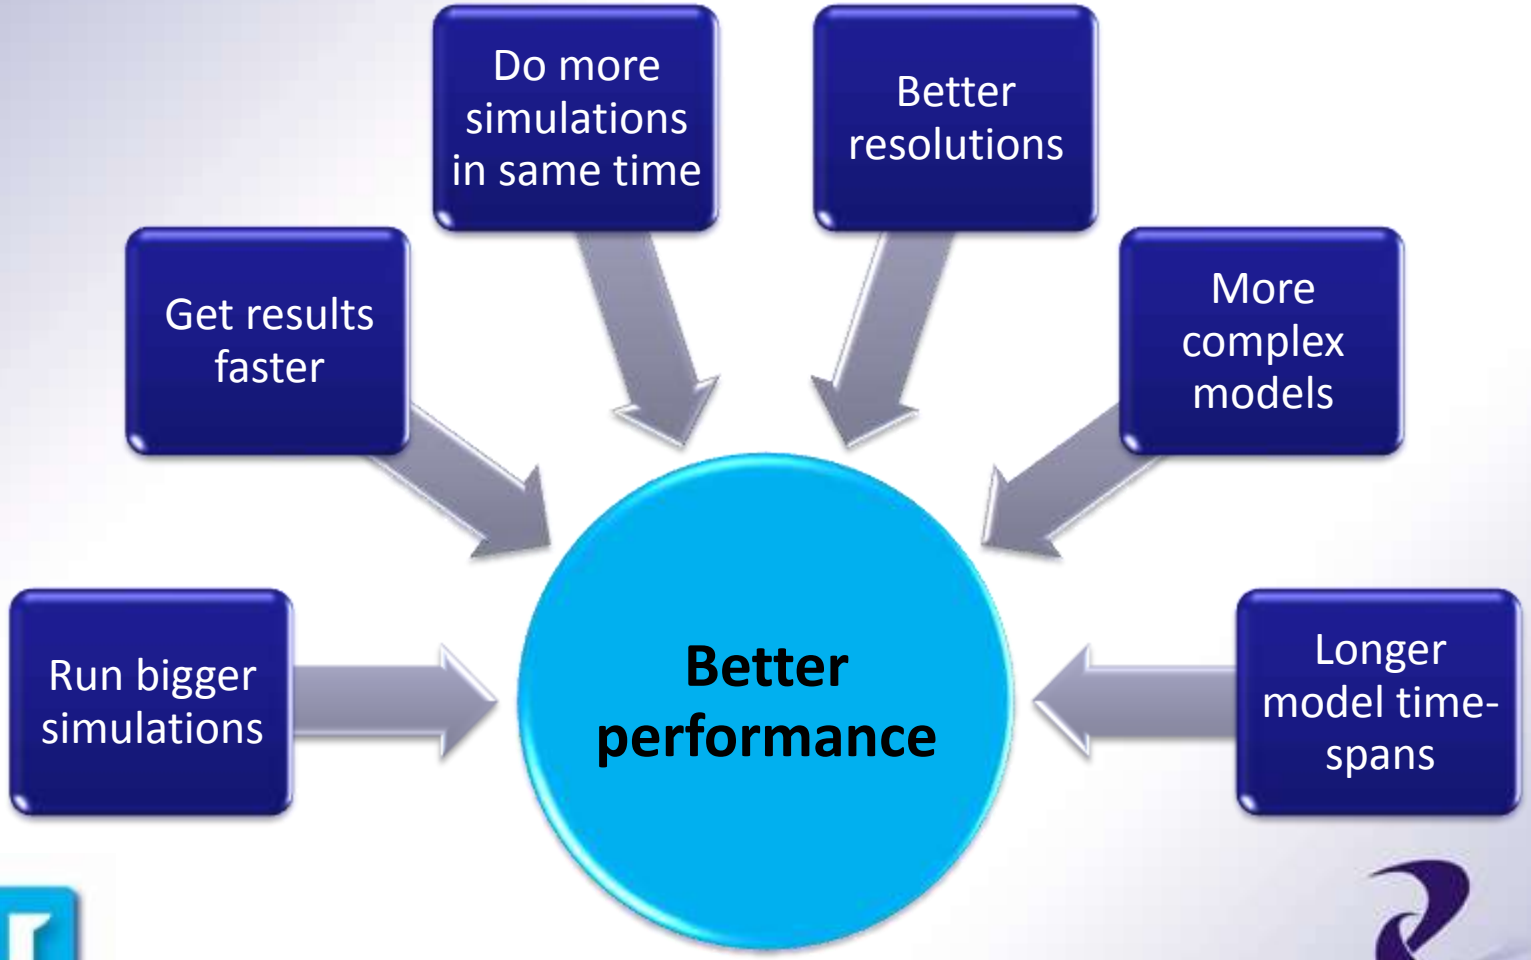

HECToR

RESEARCH COUNCILS UK

Users driving the need for CSE support ...

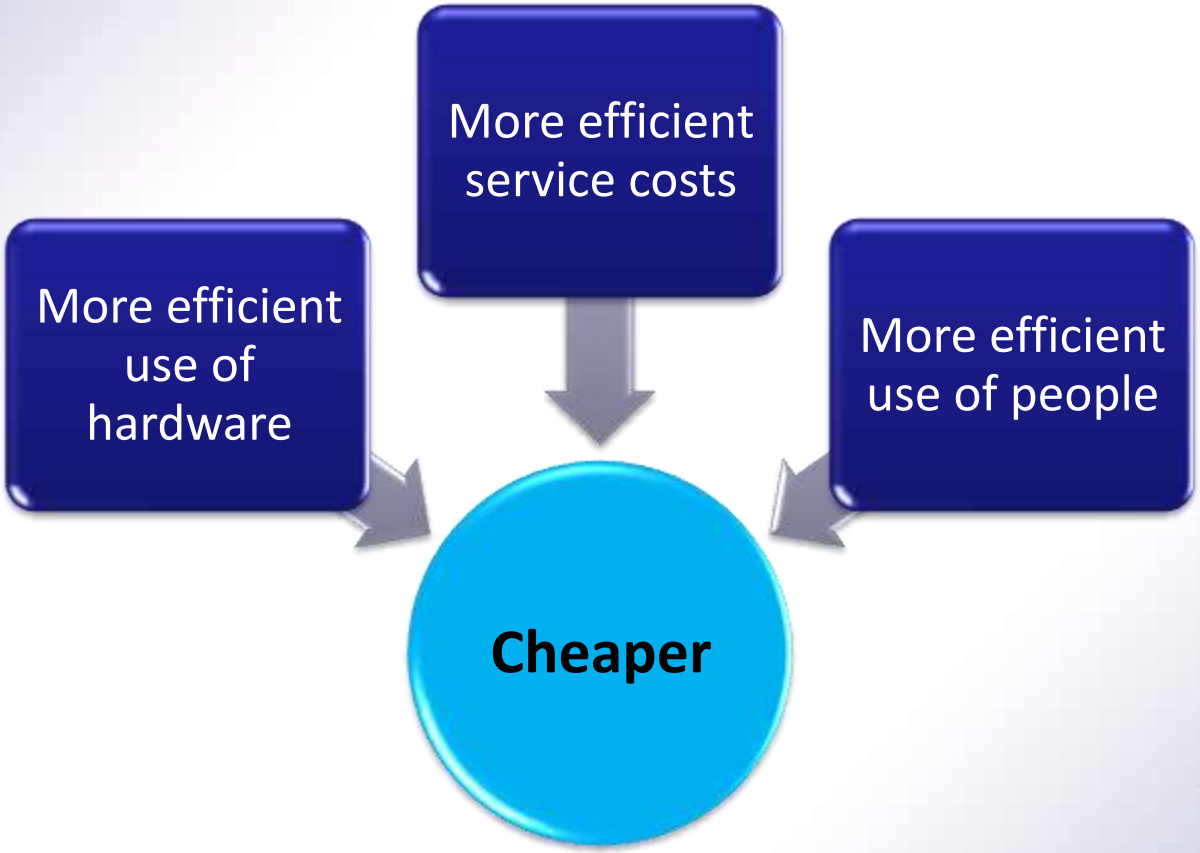

HECToR

RESEARCH COUNCILS UK

Technology driving need for CSE

There used to be 3 main types of HPC system

- Vector
- SMP
- MPP

Now we have all these rolled into one:

- SSE + multicore + MPP/cluster

Programming this technology at scale is hard

HECToR

RESEARCH COUNCILS UK

So, as well as traditional applications helpdesk ...

Support for new users

- Help with initial porting and code optimisation
- Typically 1-4 weeks effort

Training

Documentation

HECToR

RESEARCH COUNCILS UK

... HECToR has Distributed CSE support

~60% of CSE support is embedded in user community

- secondments to specific research groups
- research group staff funded by HECToR (NAG)

Facilitate new scientific research
by improving the performance of an existing code

- Optimisation or implementation of new algorithms
- Speed, scalability, methods, functionality
- 6-48 months of effort - up to 2 years duration

HECToR

RESEARCH COUNCILS UK

Some overall CSE successes so far (Feb 2010)

Code	Area	Staff	Results
CASTEP	Materials	NAG	Quad-core optimisation
CASTEP	Materials	York & STFC	4x scalability, 4x faster
CASINO	QMC	NAG	4x faster/cheaper, scaling to 40k cores
NEMO	Oceanography	EPCC	25% faster
SoFTaR	Combustion	NAG	6x faster
HYPO4D	Lattice Boltzmann	NAG	5x faster
FFT	Generic	NAG	Generic highly scalable 3D FFT routines

So far (late 2009): ~30 projects, ~20 person years

HECToR

RESEARCH COUNCILS UK

Example: CASINO

PI: Dario Alfè of University College London

Enable the CASINO Quantum Monte Carlo code to effectively use the multicore processors of HECToR

- Model more complex physical systems with greater efficiency

12 month project by NAG's Lucian Anton

HECToR

RESEARCH COUNCILS UK

Optimization and Scaling

Optimised the way CASINO used multicore nodes

- Introduced shared memory techniques
- MPI processes on a node to share common data set
- Less nodes needed for a given simulation (cheaper)

Further hierarchical parallelism with OpenMP and I/O optimisations improved the scalability of the code

- 60-80% improvement when using > 10,000 cores

CASINO now scaling up to 40,000 cores on Jaguar at ORNL

HECToR

RESEARCH COUNCILS UK

Results: CASINO

Alfè estimated that this dCSE work saved around 12M AUs for a one year research project on HECToR

- This means several millions of £ saved for his future research on HECToR and other supercomputers

"The new shared memory facility is effectively speeding up the code by a factor equal to the number of cores per node for large jobs, i.e. a fourfold increase with the current quadcores" - Alfè

HECToR

RESEARCH COUNCILS UK

Why NAG's dCSE optimizations matter

Better performance means less resources required

- Researchers can do more simulations for their budget
- More researchers can used the fixed amount of resource

In other words

Research using NAG dCSE optimized codes now costs millions of dollars less

Extra research worth millions of dollars from fixed resources is possible using NAG dCSE optimized codes

Means the supercomputers deliver the same science as a bigger machine costing millions of dollars more

HECToR

RESEARCH COUNCILS UK

In simple numbers

Assume we optimise the codes that consume about a third of the total cycles per year

If we enhance each of these codes to do 4x the science for same the allocation

Then our supercomputers costing ~$80M can now deliver the same science as $160M of supercomputers

HECToR

$(^2/_3 + 4x\ ^1/_3 = ^6/_3 = 2x)$

RESEARCH COUNCILS UK

Another example: Highly Scalable Distributed FFTs

Work by NAG expert Ning Li under HECToR dCSE project improving a CFD code at Imperial

Created a general-purpose application framework

- Based on a general 2D decomposition library
- Higher-level libraries and applications can be built on top
- Distributed multi-dimensional FFT is one example

Paper at CUG 2010

Application framework

Independent of the science application

- data movement strategy should be same for a given global data structure and domain decomposition strategy
- implementation is a software engineering issue
- proper implementation for performance is non-trivial

Recognise MPP-SMP hybrid nature of decomposition

- ALLTOALL(V) can be very expensive
- Instead, cores on a node communicate via shared buffer (System V IPC API) and only one core per node participates in global communication

HECToR

RESEARCH COUNCILS UK

FFT Library Scaling on HECToR

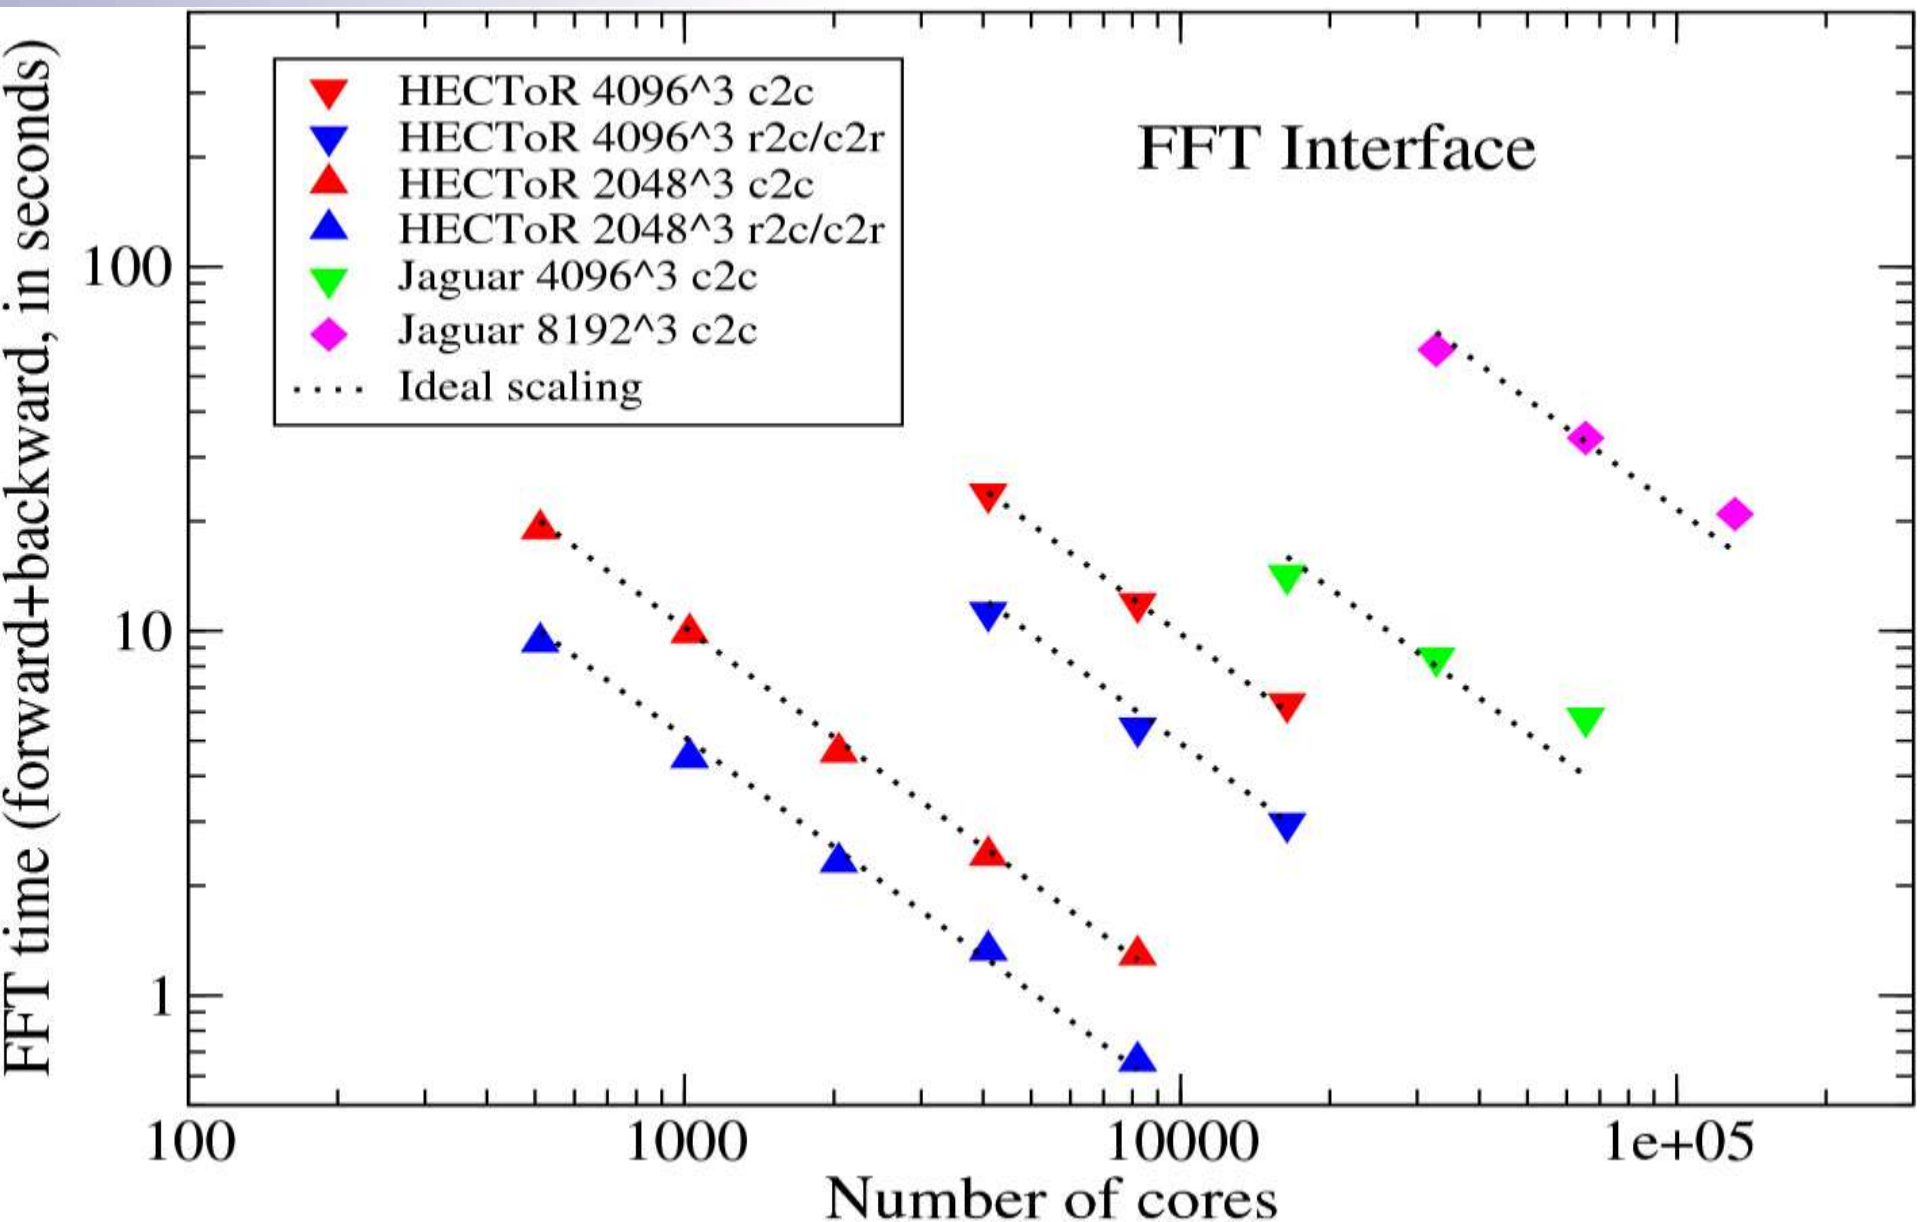

Potential Applications

Many applications using general 3D Cartesian data structures

- Spectral method
- Compact finite difference
- Multi-dimensional FFTs
- General-purpose PDF solvers

HECToR

RESEARCH COUNCILS UK

Future

Support more FFT engines

- Currently ACML, FFTW3, fftpack, generic

Additional high-level library (e.g. Poisson solver)

Port the code to other systems (e.g. IBM BG)

Hardware dependent optimisations

- Auto-tuning depending on message size on HECToR

HECToR

RESEARCH COUNCILS UK

HECToR

HIGH END COMPUTING TERASCALE RESOURCE

A Research Councils UK High End Computing Service

Capability Science. **NAG HPC expertise**.

nag®

Quad-core now, 24-core soon ...

Cost and performance benefits will be even greater as the number of cores per node increases

- 4x now, 24x by next year?

Experienced users and codes are benefiting as much as the newer ones

- Multicore really is a challenge
- even for experienced HPC user groups

HECToR

RESEARCH COUNCILS UK

Summary

world leading science

=> world leading support model

CSE investment enables better science, quicker and at lower cost

NAG's HPC expertise underpins HECToR mission

HECToR

HIGH END COMPUTING TERASCALE RESOURCE

A Research Councils UK High End Computing Service

Capability Science. **NAG HPC expertise**.

nag®