Towards exascale first-principles materials modelling with CASTEP

~~An eCSE~~
~~An ExCALIBUR~~
A GPU-eCSE project

Matt Smith, Arjen Tamerus and Phil Hasnip

ARCHER2 Celebration of Science
May 2025

What is CASTEP?

First-principles materials modelling package.

What is CASTEP?

First-principles materials modelling package. “First principles...?!”

What is CASTEP?

First-principles materials modelling package. “First principles...?!”

- Aim: predict materials’ behaviour from *first principles*
  i.e. no knowledge of what they’ll do beforehand
- Materials are made from atoms
- Most behaviour depends on the electrons
- We can use quantum mechanics to model electrons

The material’s behaviour should emerge from the simulation

Why materials modelling?

When we say ‘materials’, what do we mean?

Why materials modelling?

When we say ‘materials’, what do we mean?

- Metals, semiconductors, insulators

Why materials modelling?

When we say ‘materials’, what do we mean?

- Metals, semiconductors, insulators
- Superconductors

Why materials modelling?

When we say ‘materials’, what do we mean?

- Metals, semiconductors, insulators
- Superconductors
- Bone

Why materials modelling?

When we say ‘materials’, what do we mean?

- Metals, semiconductors, insulators
- Superconductors
- Bone
- Rocket fuel

Why materials modelling?

When we say ‘materials’, what do we mean?

- Metals, semiconductors, insulators
- Superconductors
- Bone
- Rocket fuel
- Jelly

Why materials modelling?

When we say ‘materials’, what do we mean?

- Metals, semiconductors, insulators
- Superconductors
- Bone
- Rocket fuel
- Jelly
- Caffeine

Why materials modelling?

When we say ‘materials’, what do we mean?

- Metals, semiconductors, insulators
- Superconductors
- Bone
- Rocket fuel
- Jelly
- Caffeine
- Neutron stars

Why materials modelling?

When we say ‘materials’, what do we mean?

- Metals, semiconductors, insulators
- Superconductors
- Bone
- Rocket fuel
- Jelly
- Caffeine
- Neutron stars

That sort of thing.

Large calculations

Quantum mechanics is quite difficult!

- Use a Nobel Prize-winning reformulation (density functional theory)
- Gives the electron density and the energy
- Unfortunately, no experiment can measure the energy...

Large calculations

Quantum mechanics is quite difficult!

- Use a Nobel Prize-winning reformulation (density functional theory)
- Gives the electron density and the energy
- Unfortunately, no experiment can measure the energy...
  ... but they do measure derivatives of the energy

Good vibrations

- How does the energy change when atoms move?

$$\frac{\partial^2 E}{\partial R_n \partial R_m}$$

- Tells you about atomic vibrations, which affect:
  - Speed of sound
  - Heat transport
  - Electrical resistance
  - Infrared spectra
- Lots of work!
- But then we have big computers... right?

We're gonna need a bigger computer

Exascale machines characterised by:

- Massive parallelism
- GPUs

CASTEP makes heavy use of

- Dense linear algebra
- 3D Fast Fourier Transforms (FFTs)

Parallel performance, 1 BP (Before PAX)

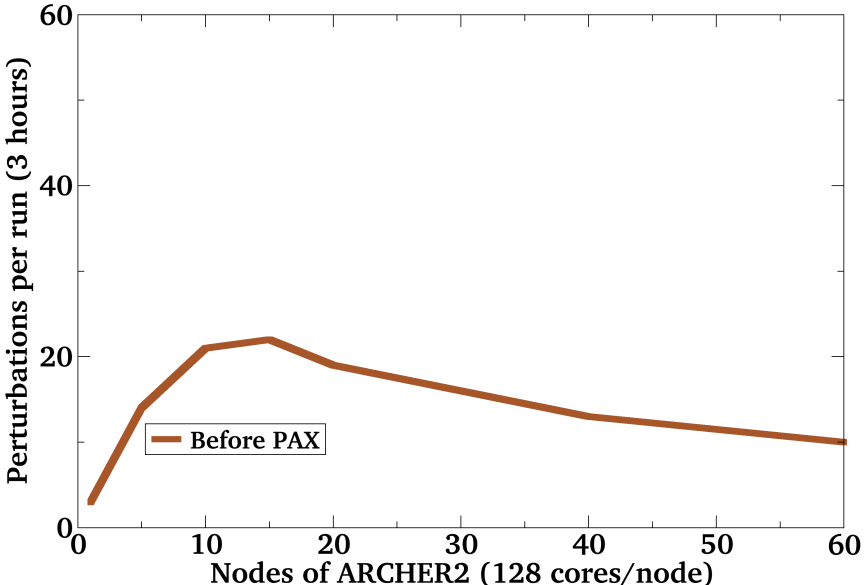

Parallelisation

- The parallel 3D FFTs limit the scaling
- FFTs need all-to-all comms...

Parallelisation

- The parallel 3D FFTs limit the scaling
- FFTs need all-to-all comms... or do they?

Parallelisation

- The parallel 3D FFTs limit the scaling
- FFTs need all-to-all comms... or do they?
- eCSE & PAX: New decomposition inspired by process grid:

Conventional decomposition

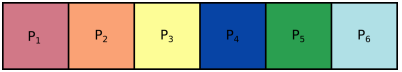

$O(P^2)$

New decomposition

$O(P)$

Parallel performance, 1 AP (After PAX)

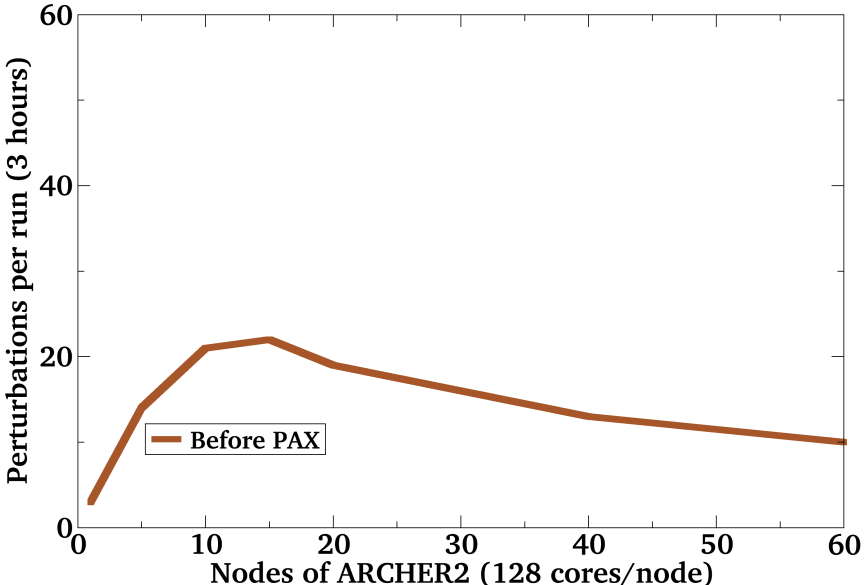

Parallel performance, 1 AP (After PAX)

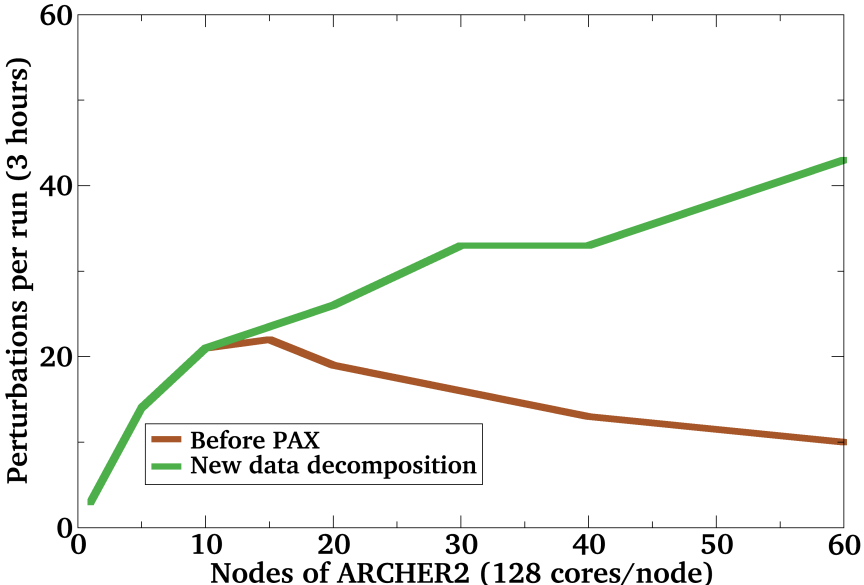

Parallel performance, 3 AP (After PAX)

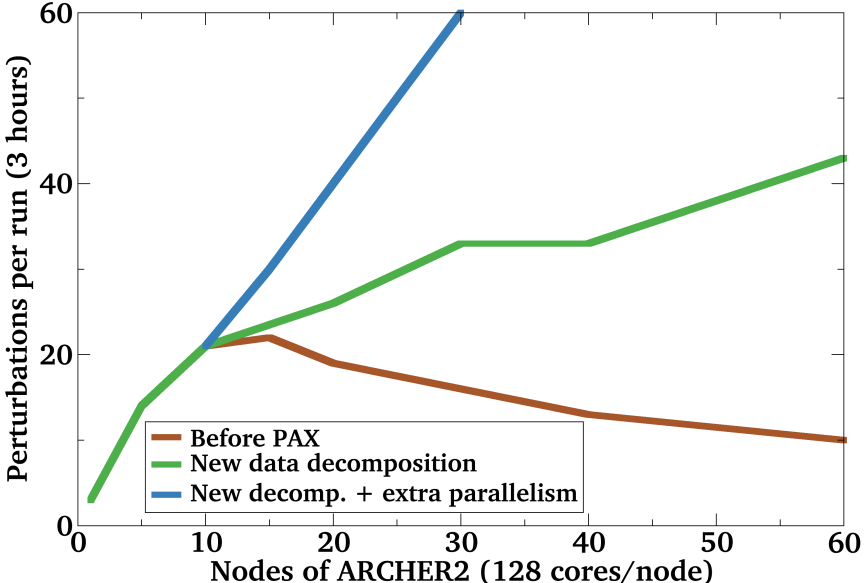

Parallel performance, 3 AP (After PAX)

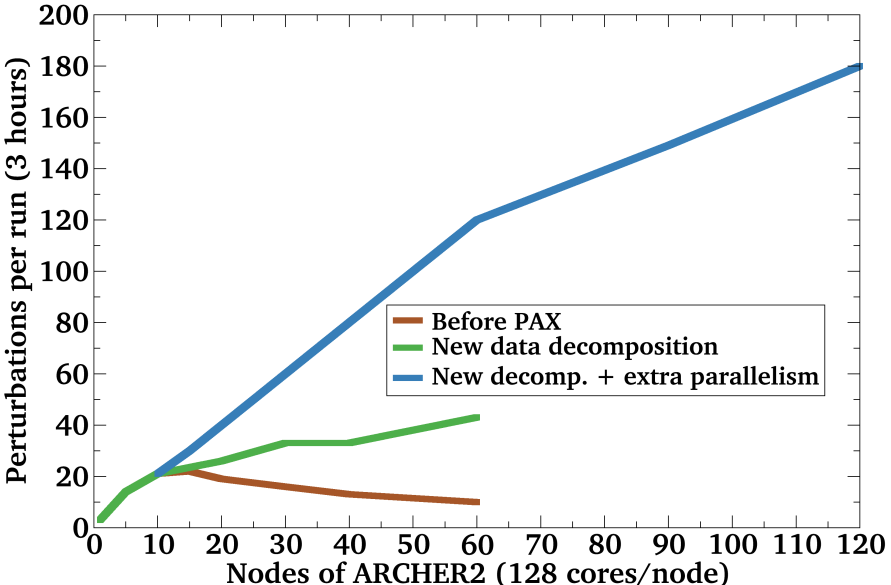

Enter the GPU

CASTEP is a big Fortran code used by many research groups and companies.

Focus on:

- Single code base
- Directives-based data movement (OpenACC)
- Use of optimised GPU libraries

CASTEP-GPU on Bede (UK Tier-2 HPC)

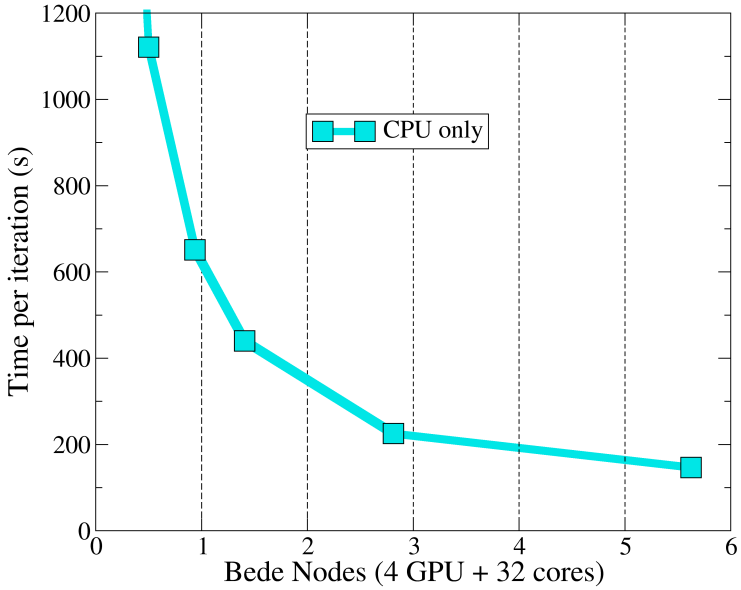

CASTEP-GPU on Bede

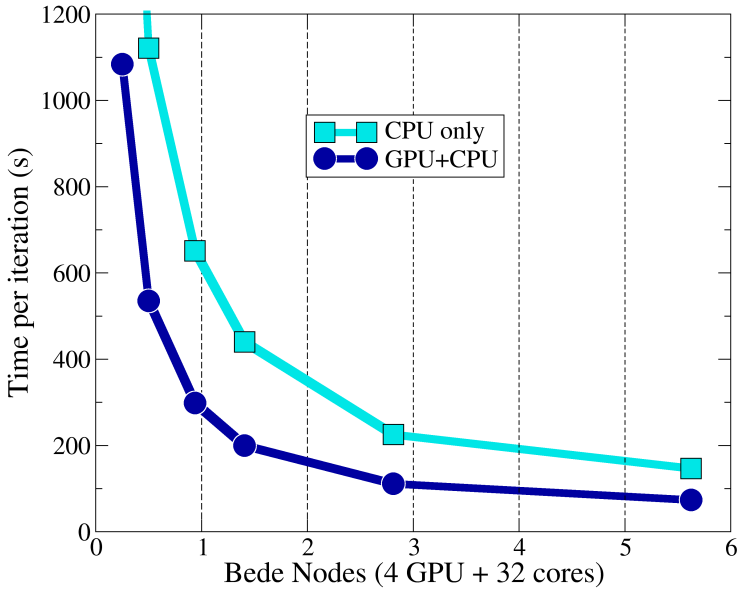

CASTEP-GPU on Bede (UK Tier-2 HPC)

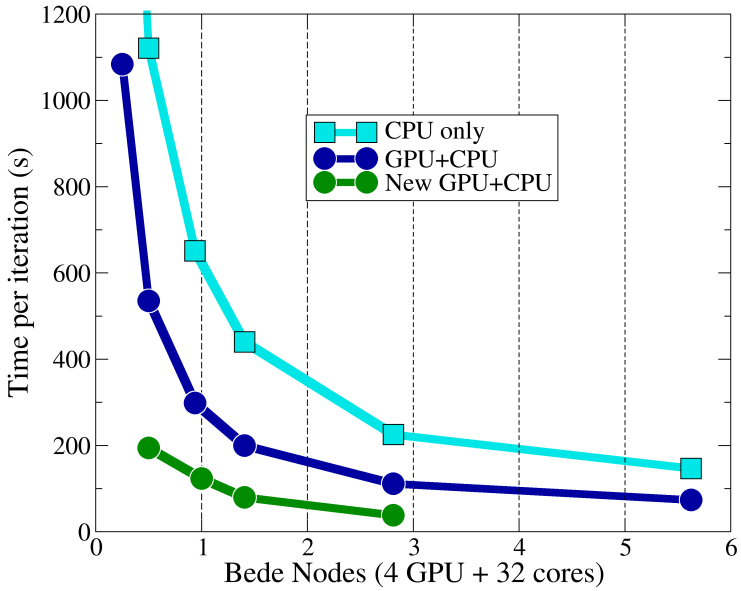

CASTEP-GPU on Bede (UK Tier-2 HPC)

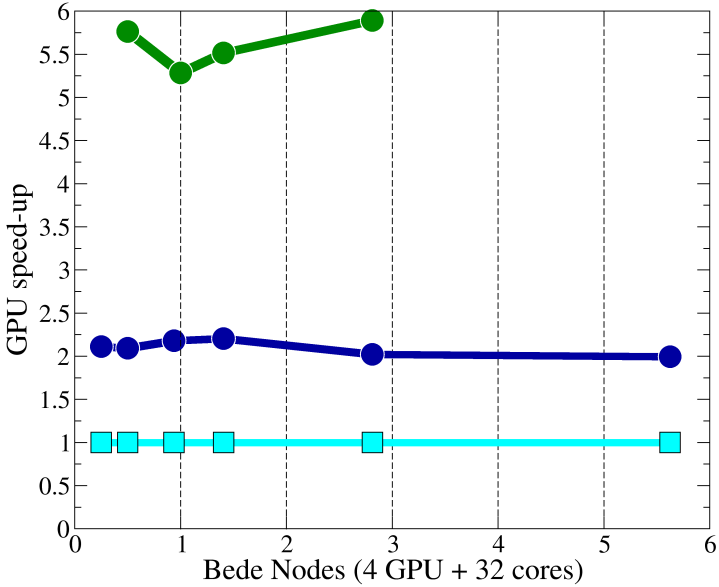

CASTEP-GPU on Bede (UK Tier-2 HPC)

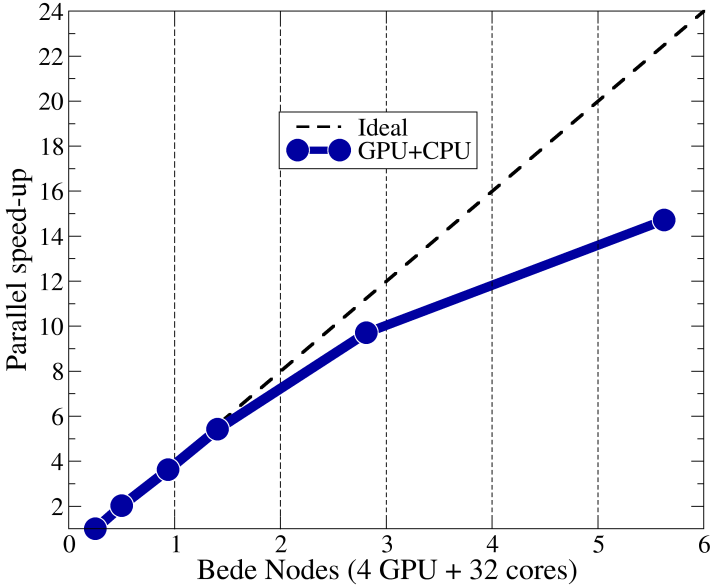

Towards Exascale

- Many challenges... but significant performance improvements
- Re-thinking parallel decompositions: 4X scaling
- Re-engineering for GPUs: 6-8X
- See Matt Smith's poster for details (poster 22)
- GPU-eCSE to combine the two!

New CASTEP parallelism:
B. Durham, M.J. Smith & P. Hasnip, eCSE Technical Report, doi: 10.5281/zenodo.14960783

CASTEP GPU:
M.J. Smith et al, *Comput. Sci. Eng.* **24**(1) 46-55 (Jan-Feb 2022); doi: 10.1109/MCSE.2022.3141714