Exploiting OpenCL for heterogeneous computing: a case study

Simon McIntosh-Smith simonm@cs.bris.ac.uk
Head of Microelectronics Research

University of BRISTOL

Collaborators

- Richard B. Sessions, Amaurys Avila Ibarra
  - University of Bristol, Biochemistry
- James Price
  - University of Bristol, Computer Science
- Tsuyoshi Hamada, Felipe Cruz
  - University of Nagasaki, Japan

University of BRISTOL

Molecular docking

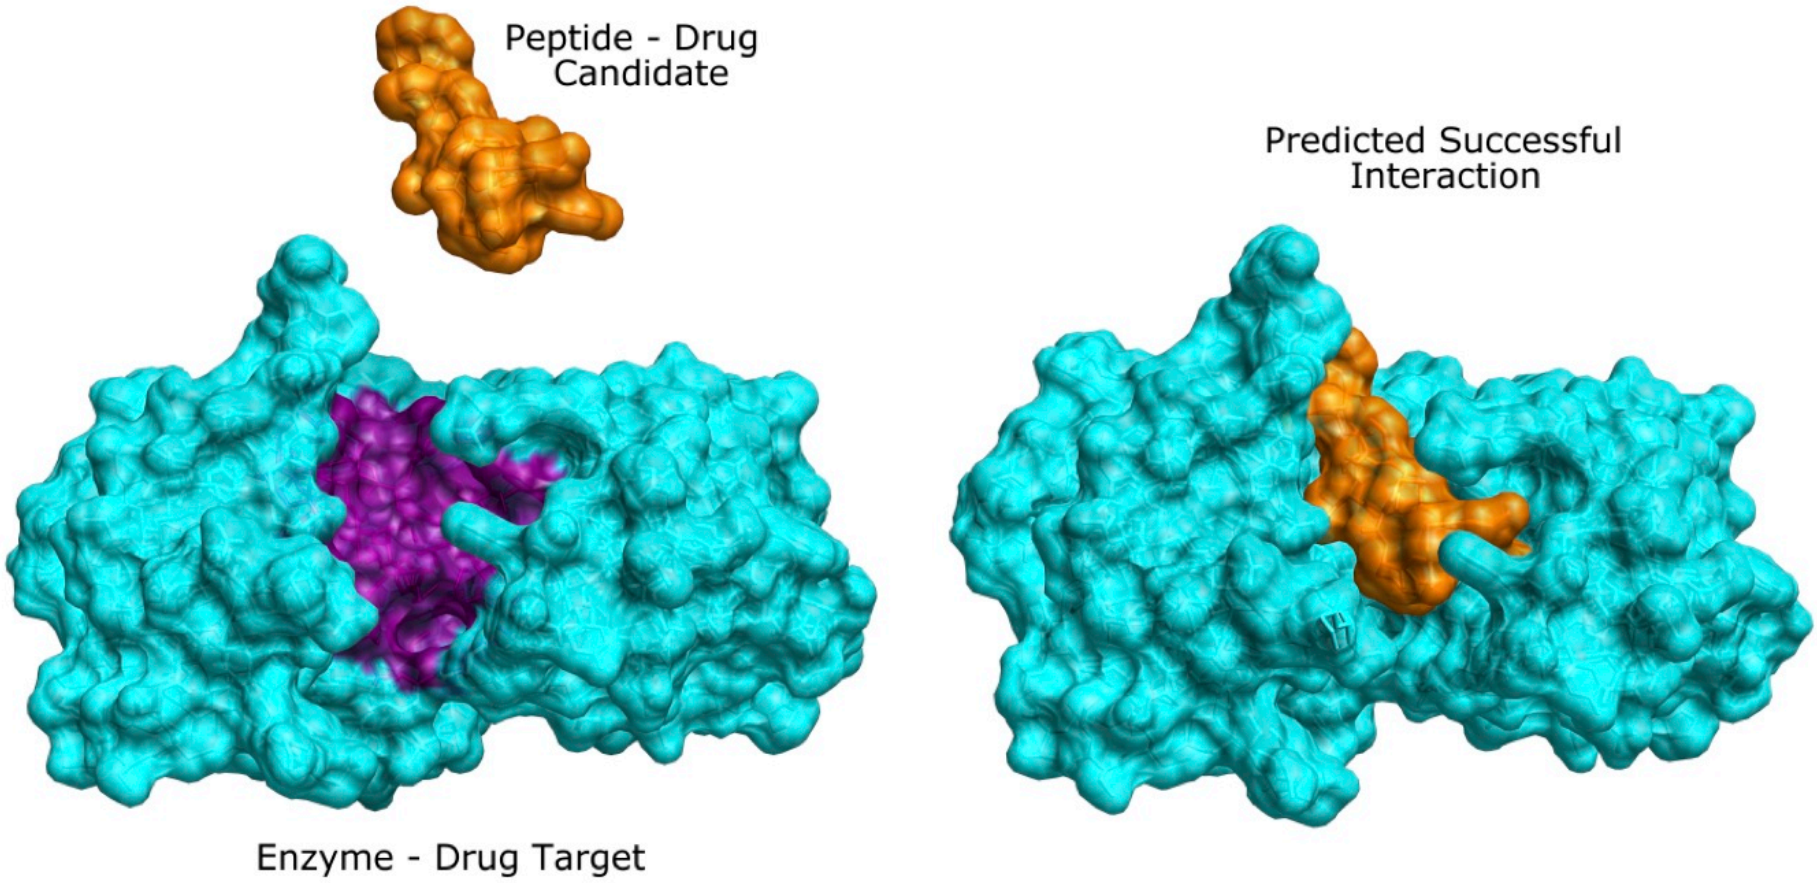

Proteins typically O(1000) atoms
Ligands typically O(100) atoms

University of BRISTOL

BUDE: Bristol University Docking Engine

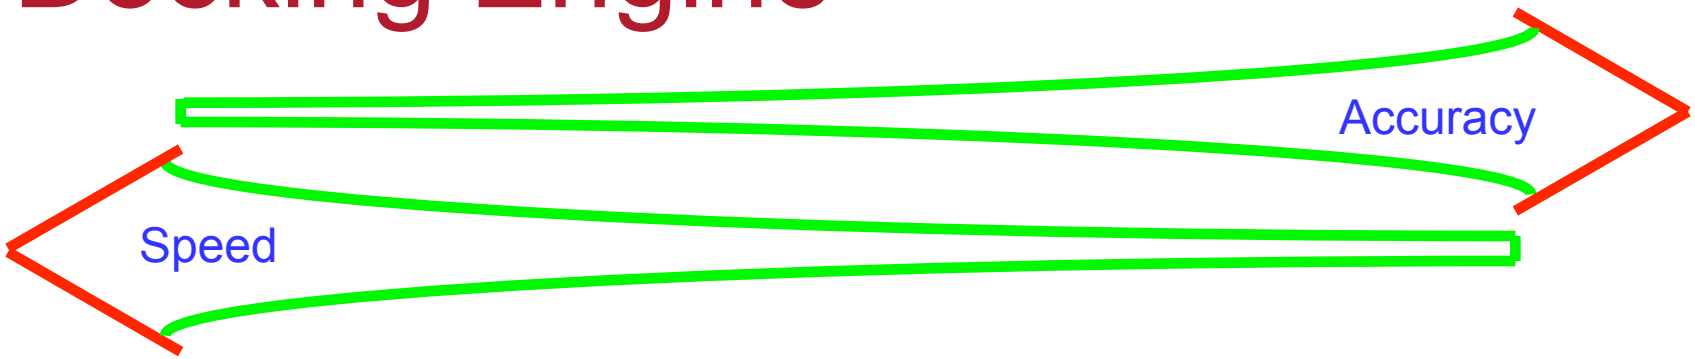

Typical docking scoring functions	Empirical Free Energy Forcefield **BUDE**	Free Energy calculations MM[1,2] QM/MM[3]
Entropy:		
solvation	**No**	**Yes**	**Yes**
configurational	**Approx**	**Approx**	**Yes**
Electrostatics	**?**	**Approx**	**Yes**
All atom	**No**	**Yes**	**Yes**
Explicit solvent	**No**	**No**	**Yes**

University of BRISTOL

1. MD Tyka, AR Clarke, RB Sessions, J. Phys. Chem. B 110 17212-20 (2006)
2. MD Tyka, RB Sessions, AR Clarke, J. Phys. Chem. B 111 9571-80 (2007)
3. CJ Woods, FR Manby, AJ Mulholland, J. Chem. Phys. 128 014109 (2008)

How BUDE's EMC works

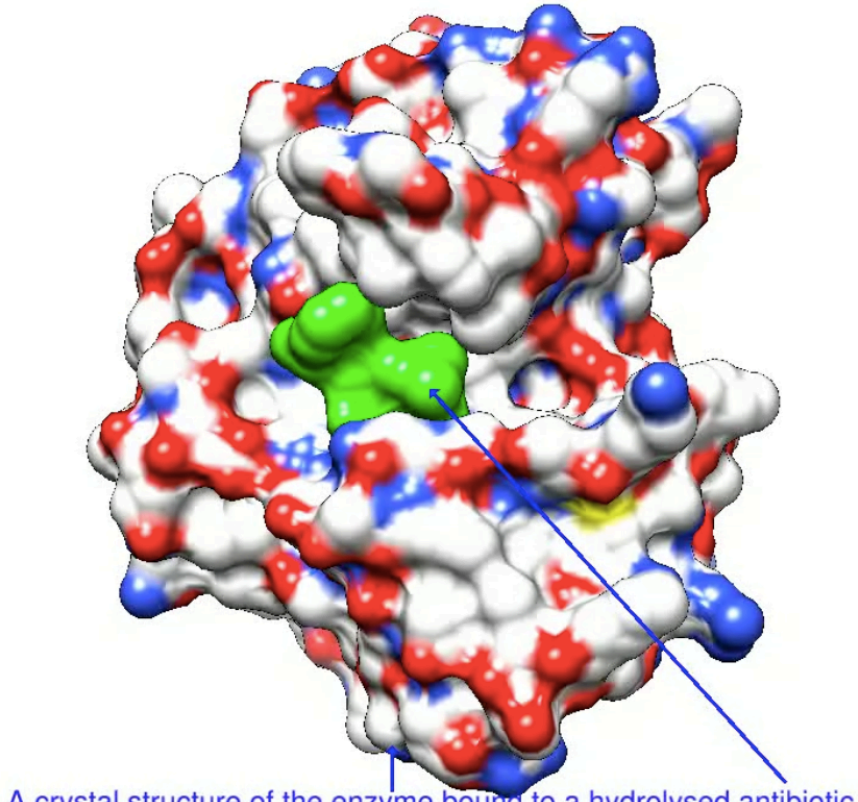

University of BRISTOL

Experimental results

Redocking into Xray Structure

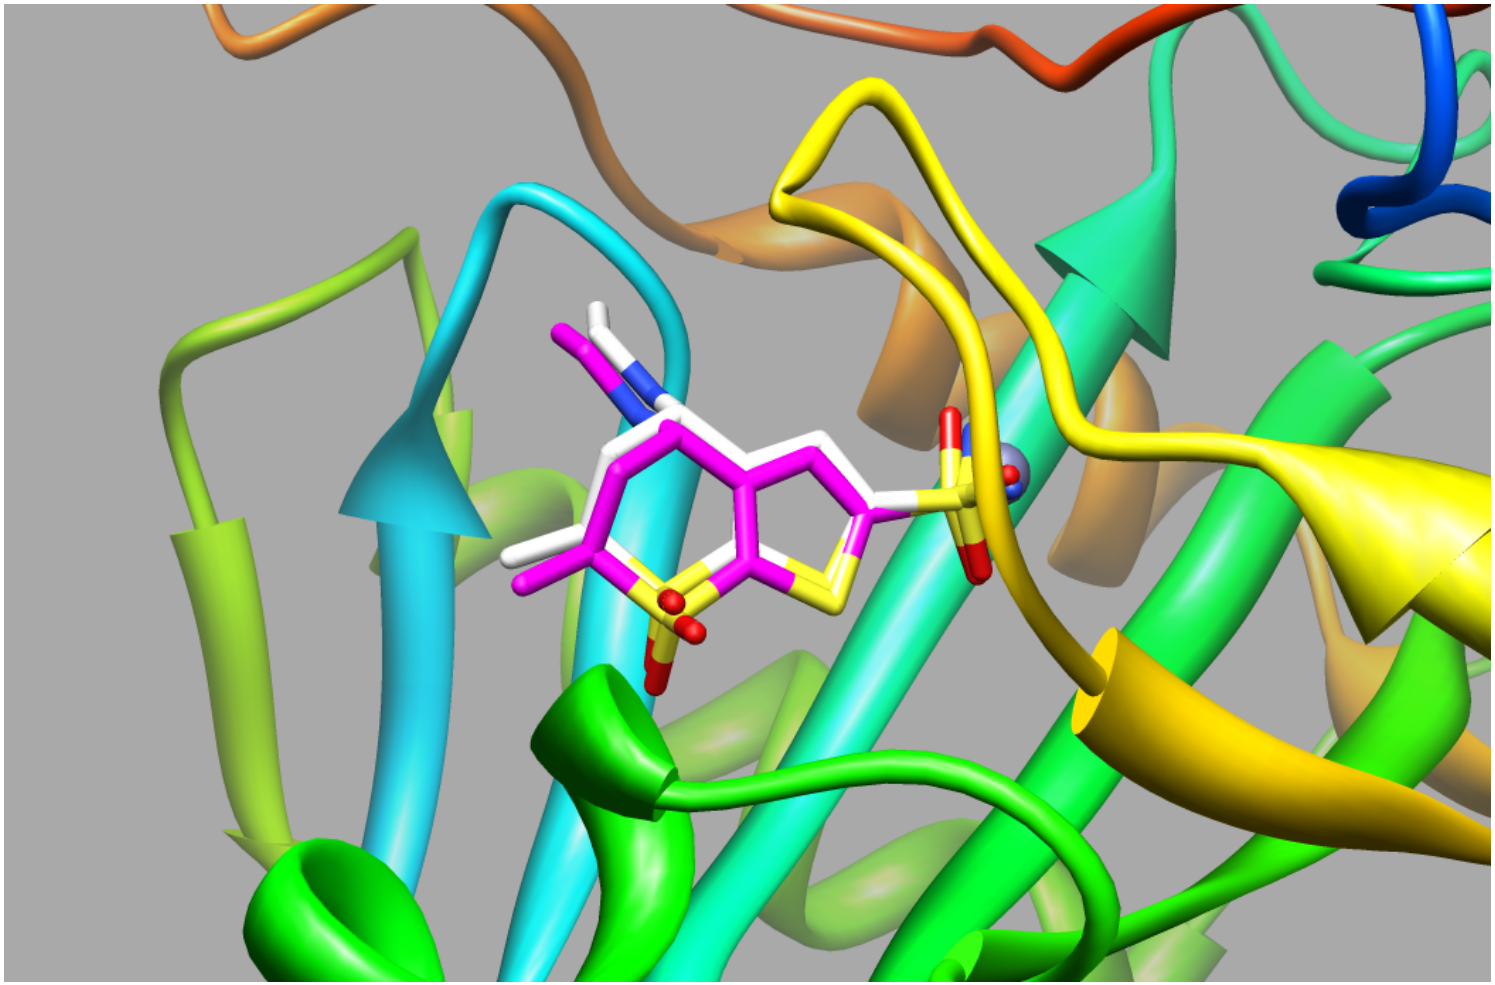

University of BRISTOL

1CIL (Human carbonic anhydrase II)

Another example

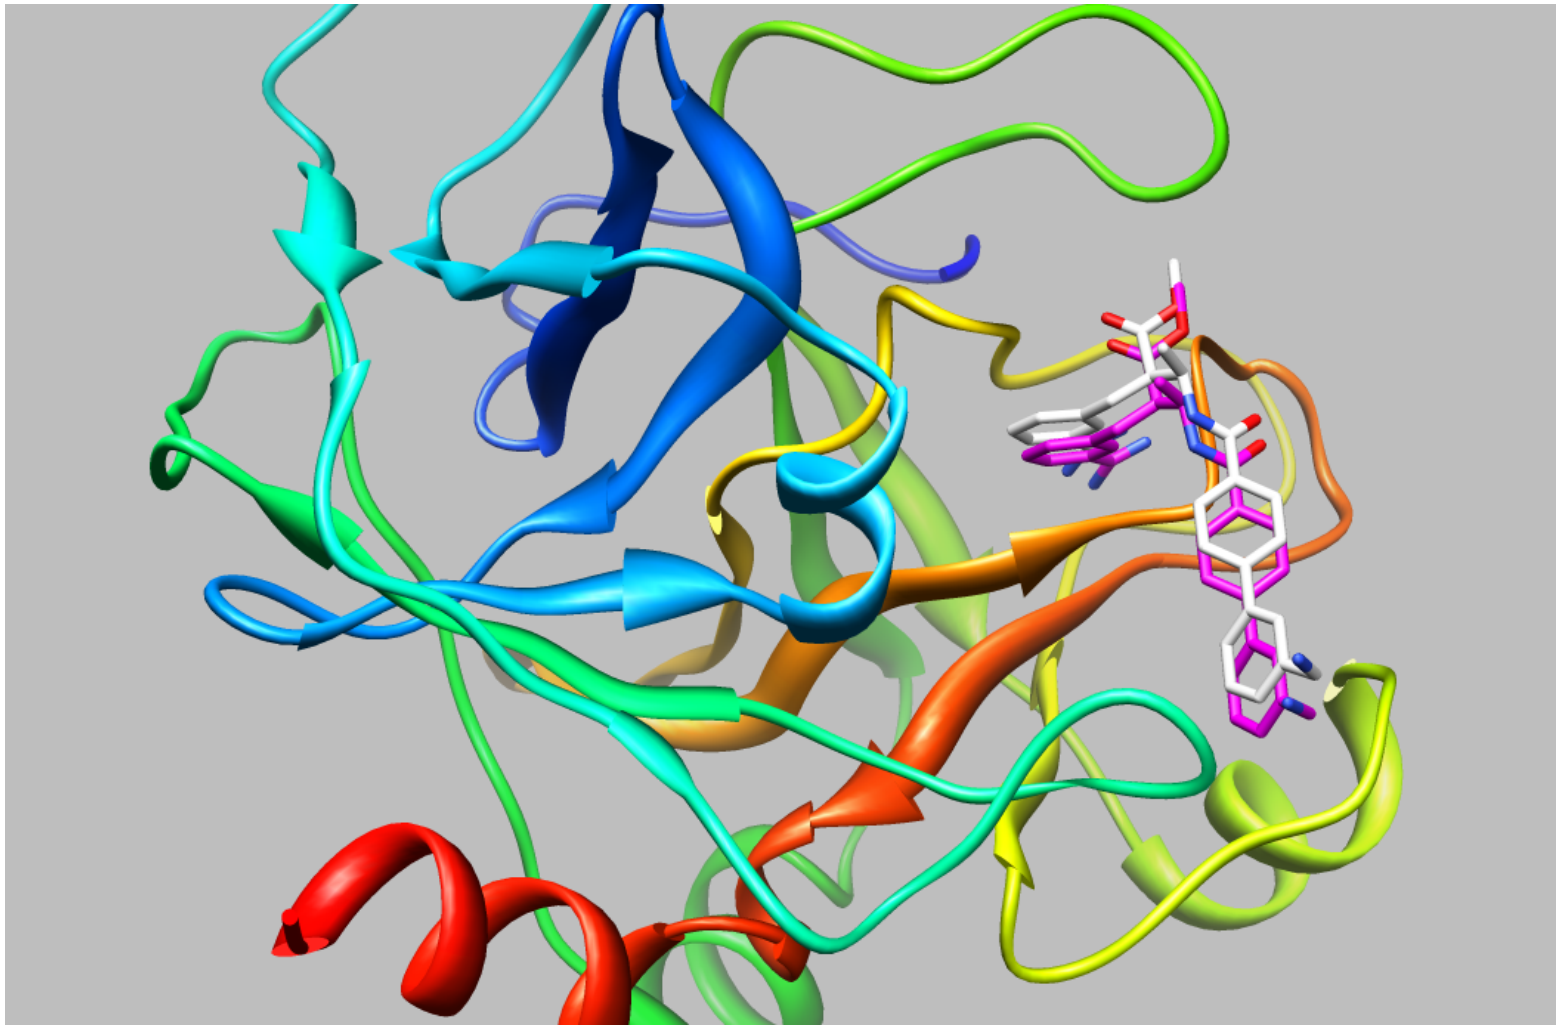

1EZQ (Human Factor XA)

OpenCL for heterogeneous computing

A modern computer includes:

- One or more CPUs
- One or more GPUs

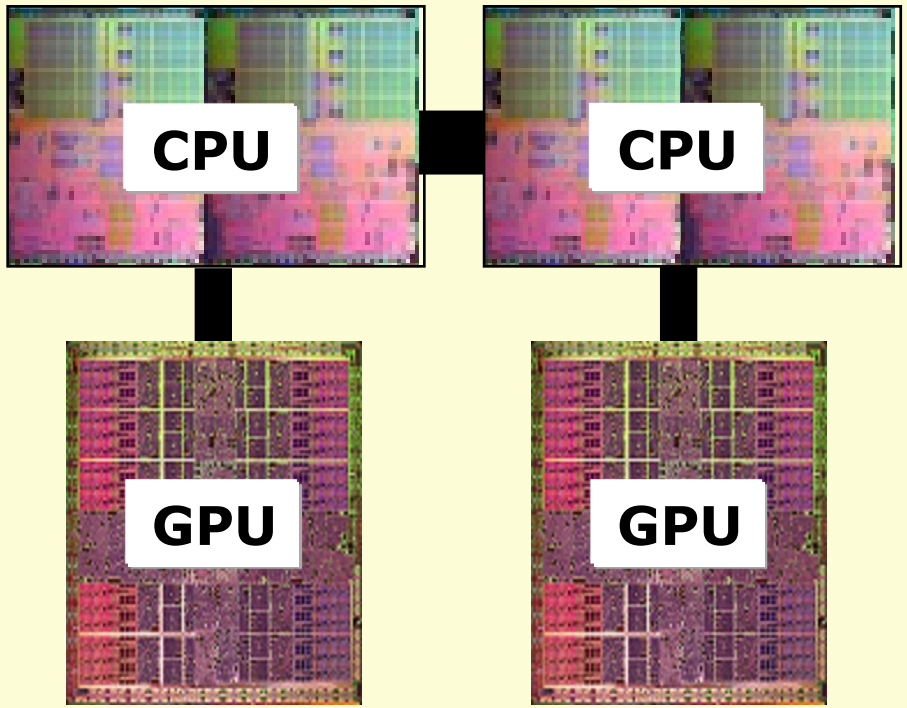

OpenCL (Open Compute Language) lets programmers write a single portable program that uses ALL resources in the heterogeneous platform

BUDE's heterogeneous approach

1. Discover all OpenCL platforms/devices, inc. both CPUs and GPUs
2. Run a *micro benchmark* on each device, ideally a short piece of real work
3. Load balance using micro benchmark results
4. Re-run micro benchmark at regular intervals in case load changes

Benchmark results

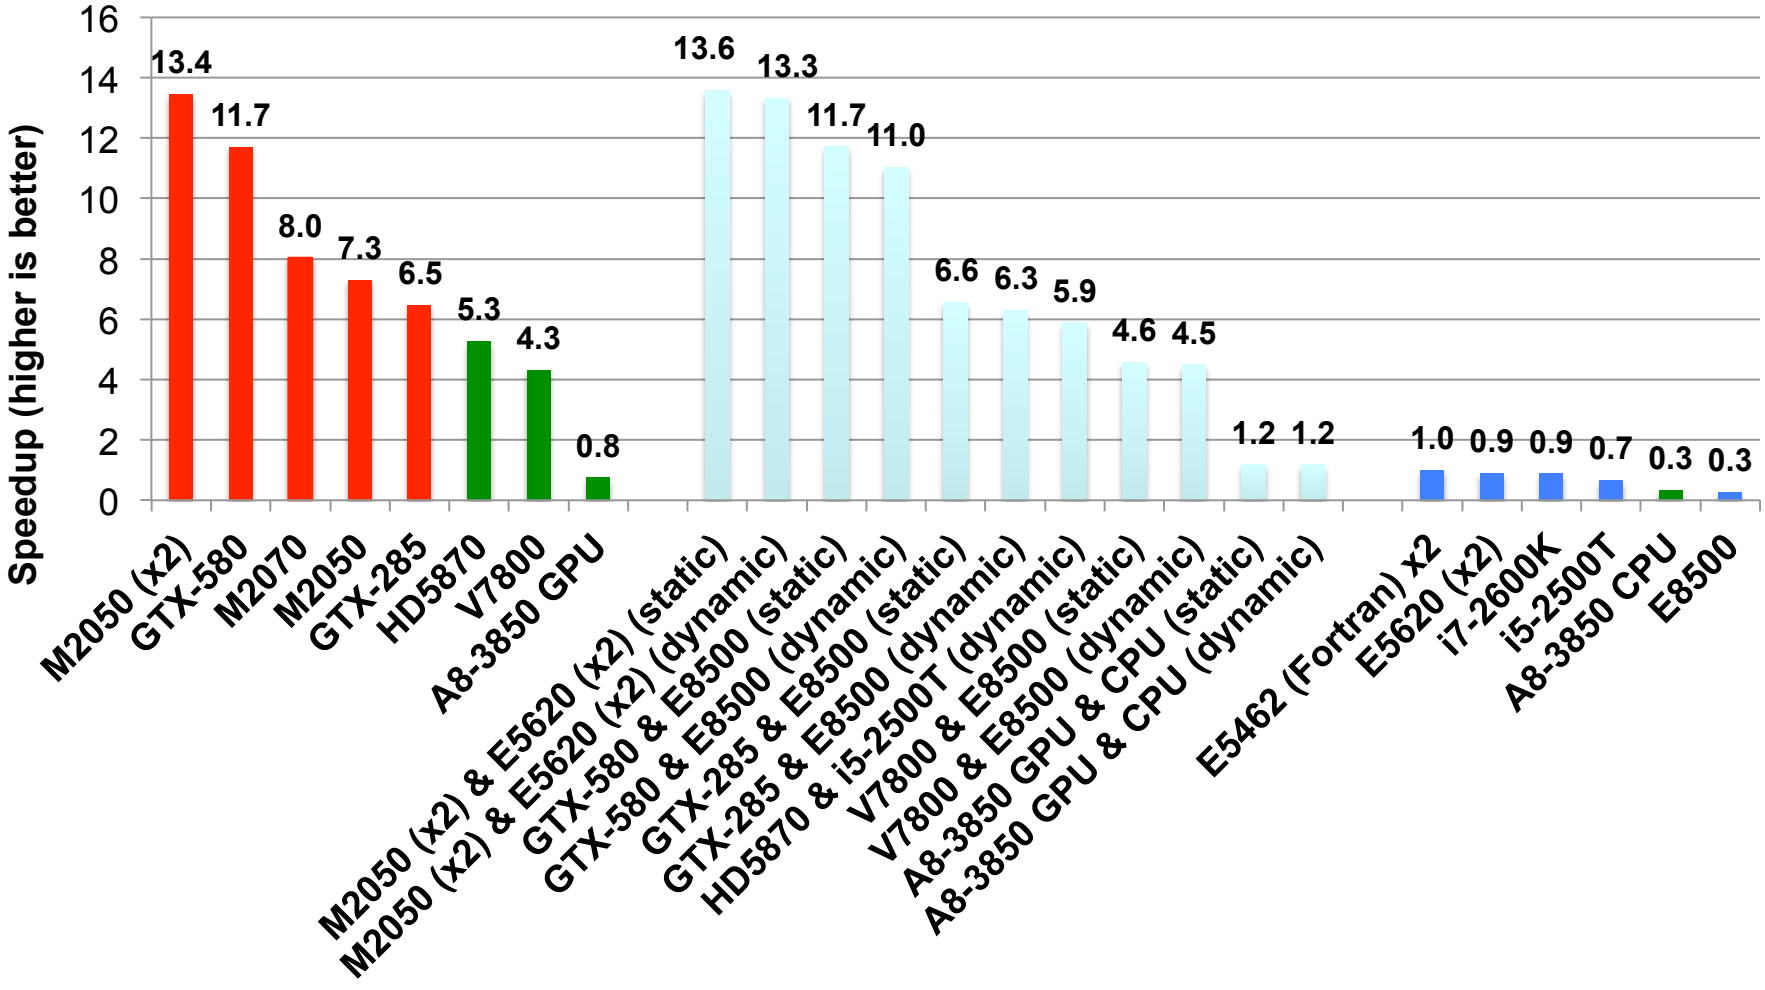

University of BRISTOL

Selected performance results

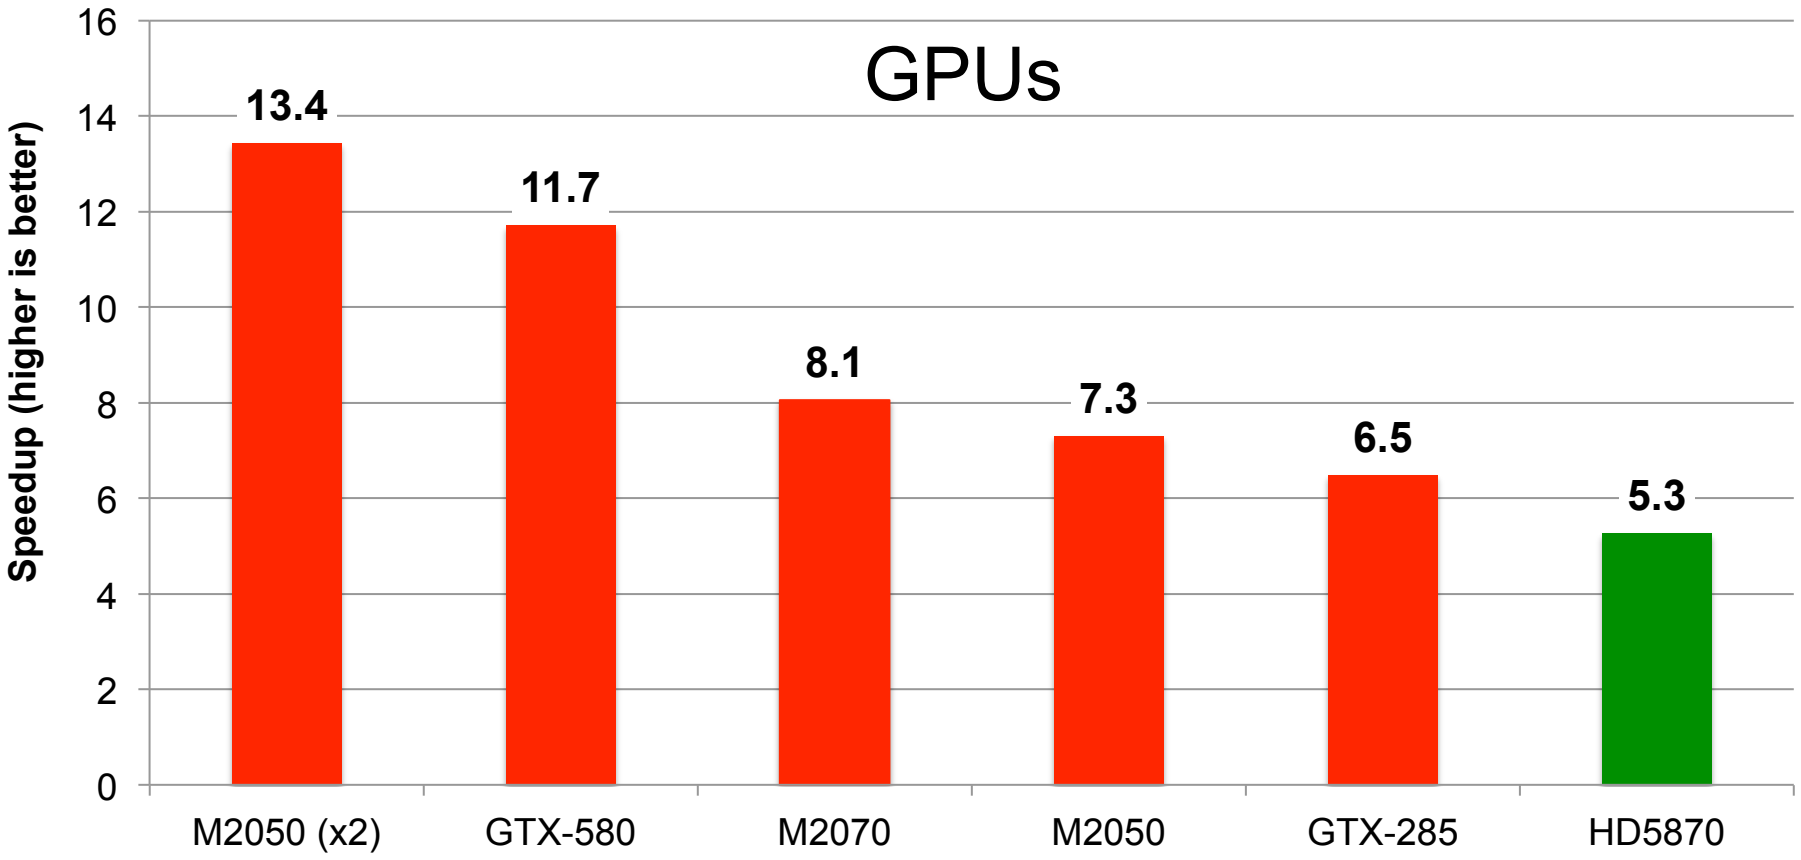

University of BRISTOL

Selected performance results

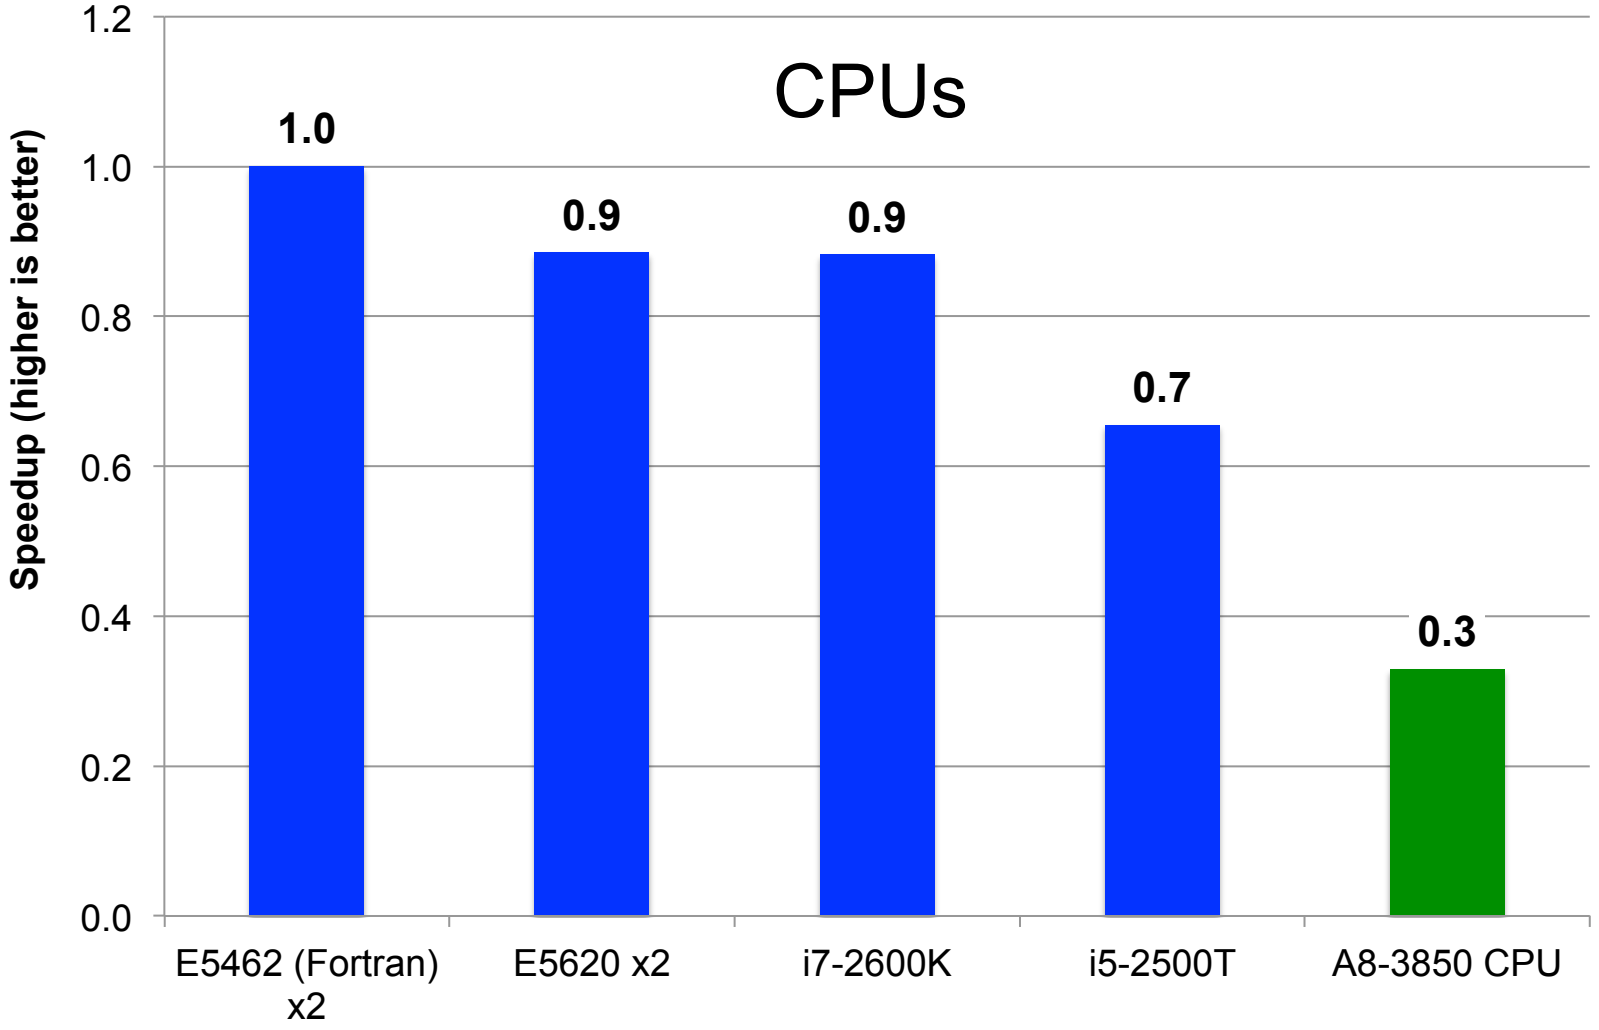

University of BRISTOL

Selected performance results

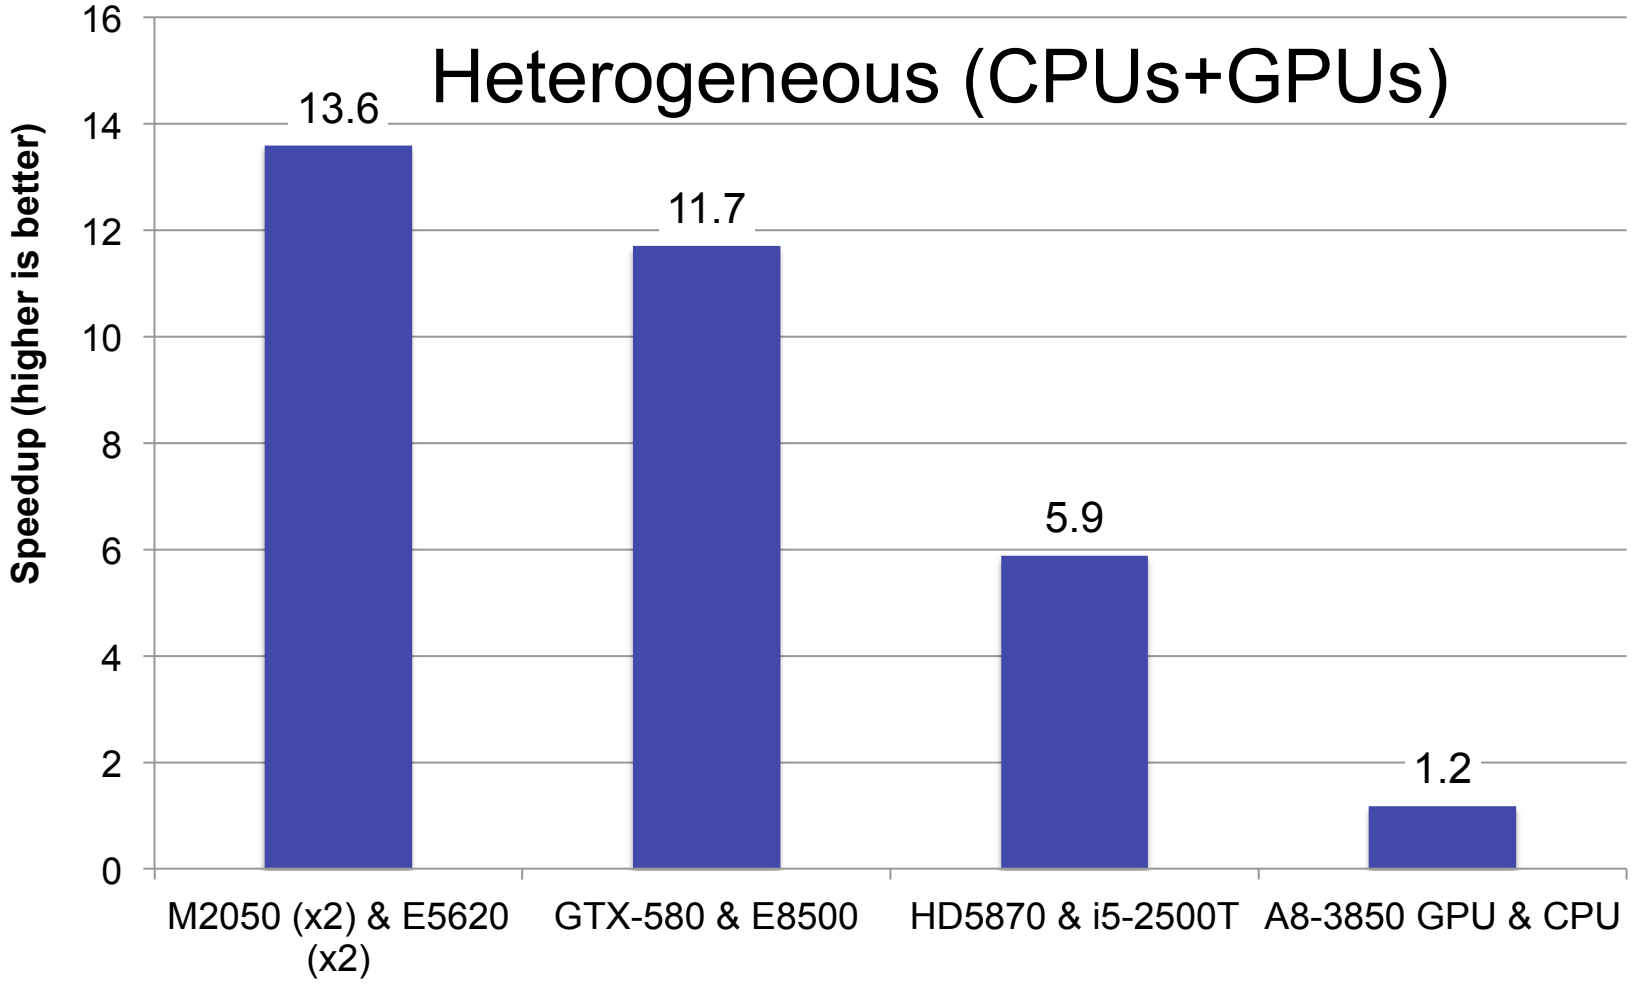

University of BRISTOL

Selected performance results

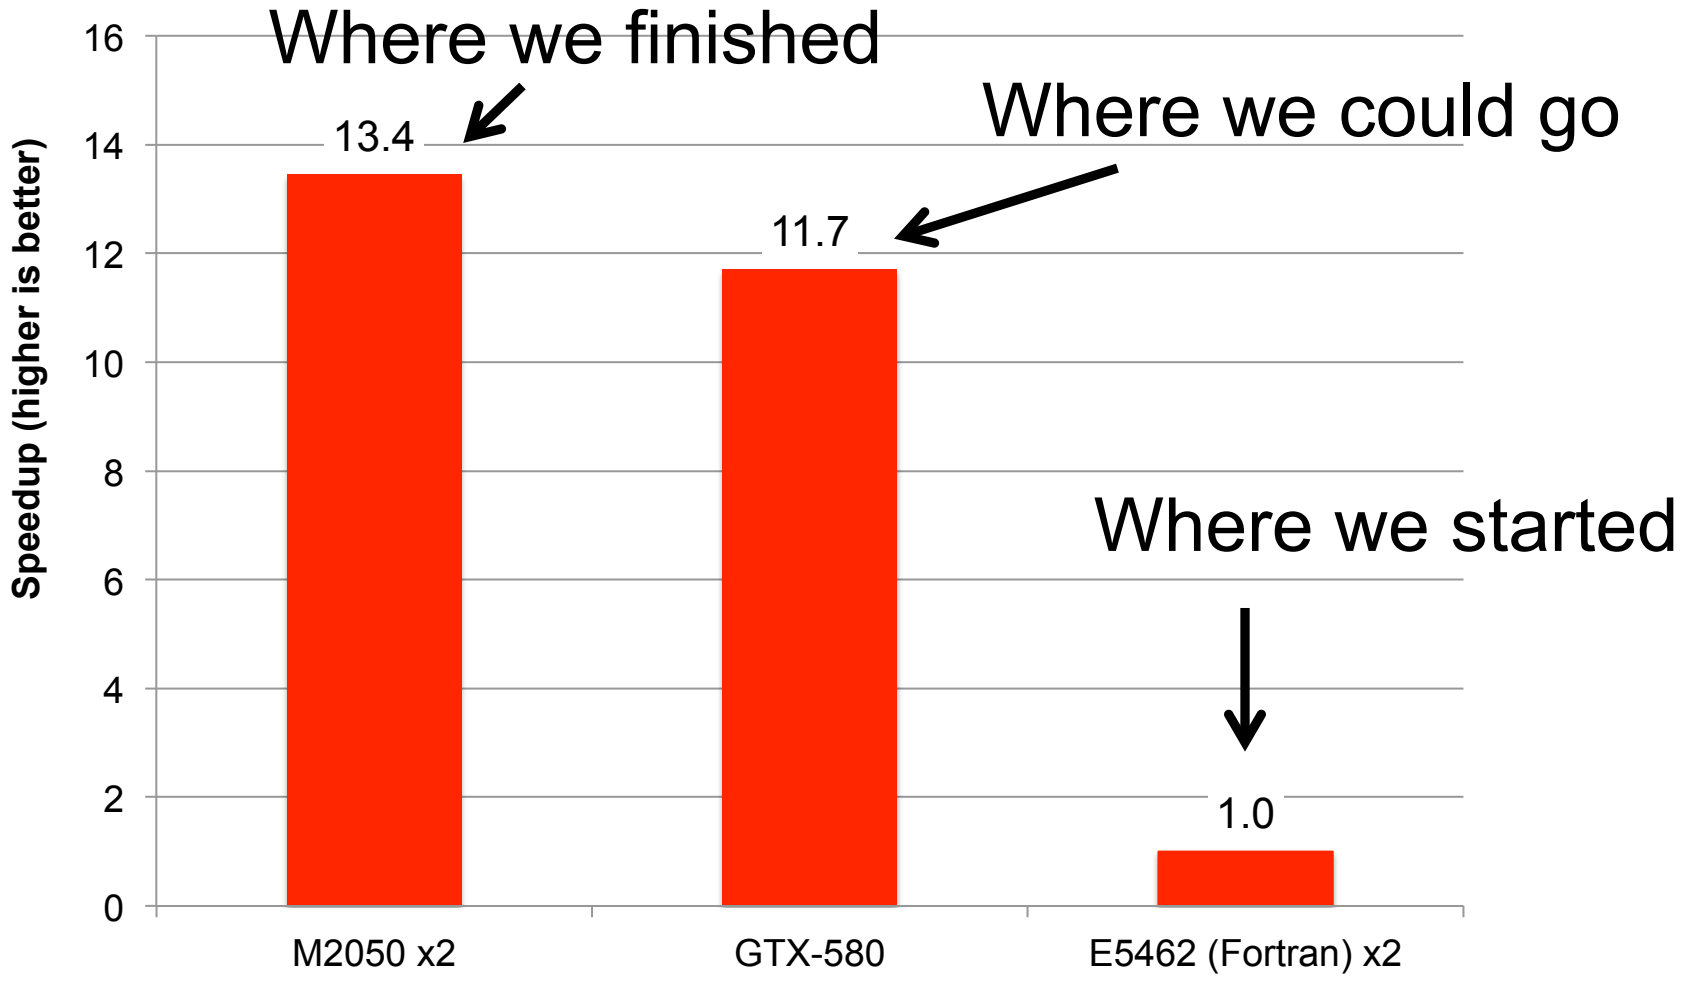

£340

University of BRISTOL

Relative energy and run-time

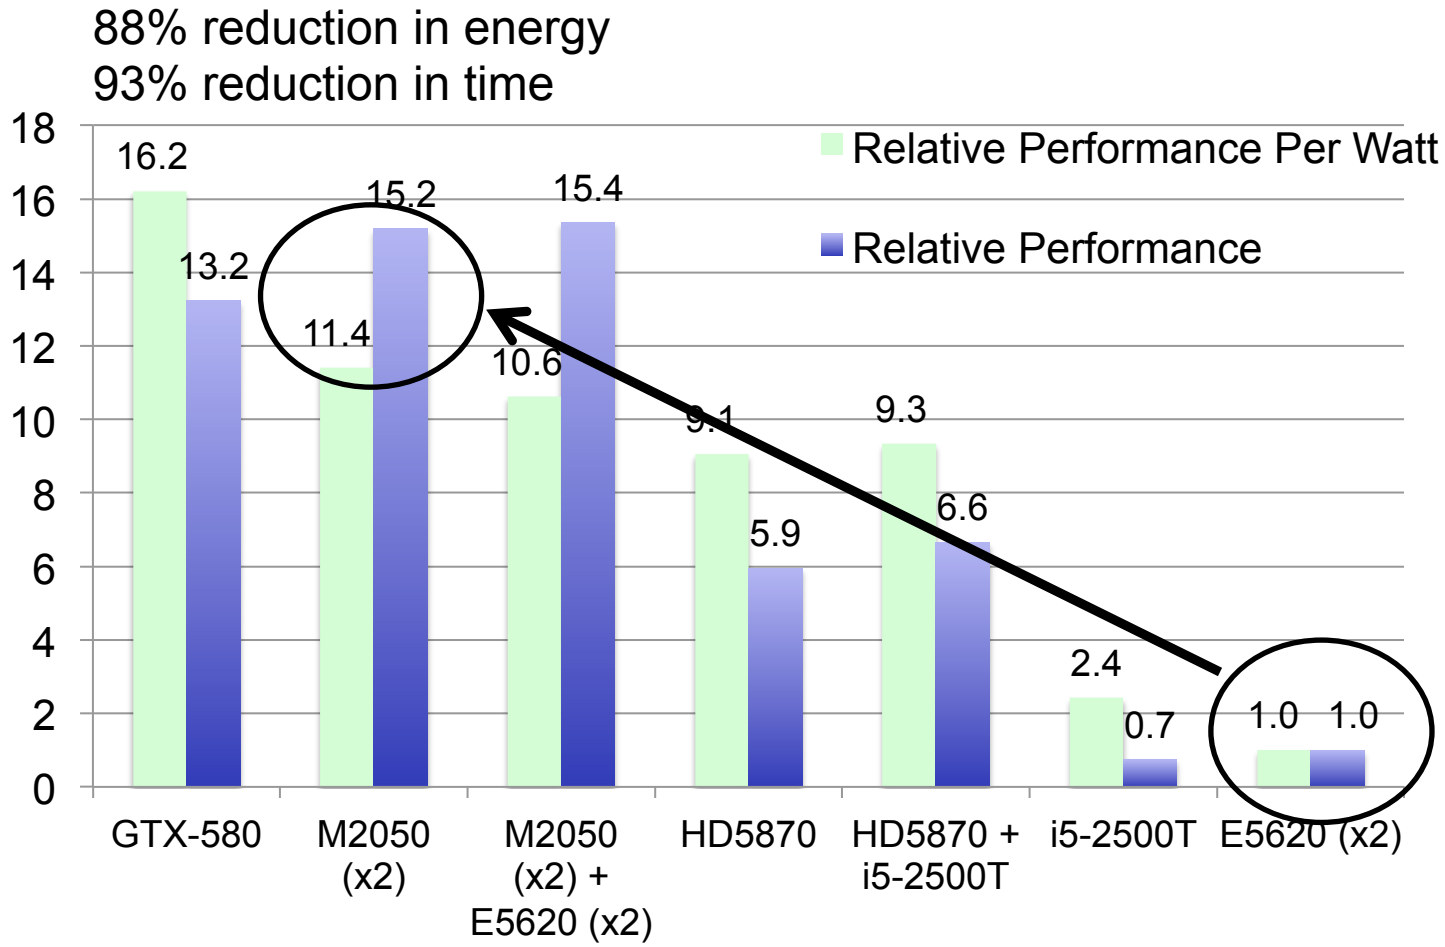

Measurements are for a constant amount of work.
Energy measurements are “at the wall” and include any idle components.

University of BRISTOL

What does this let us do?

Potentially save lives

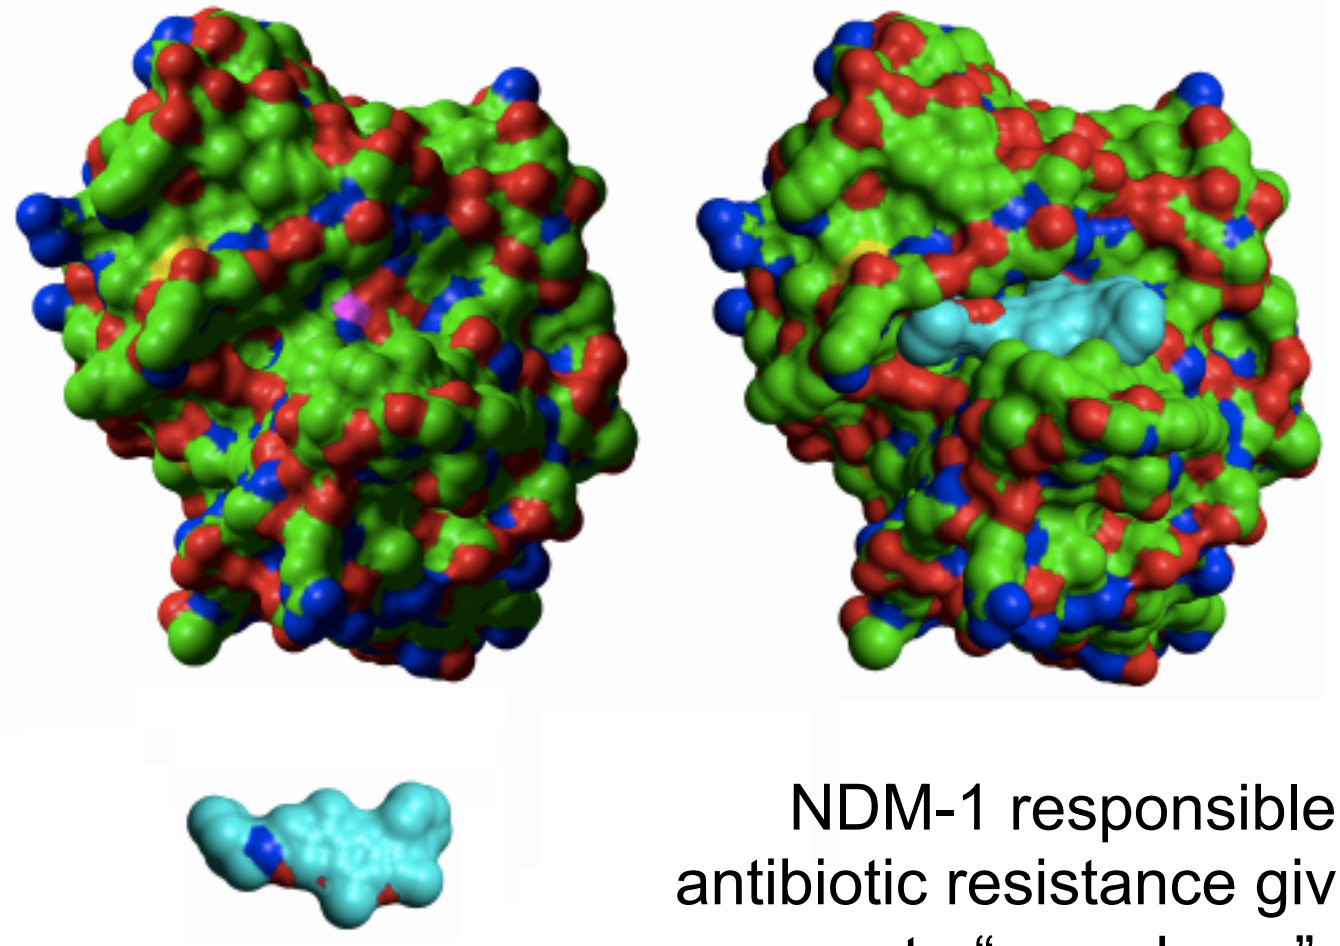

NDM-1 responsible for antibiotic resistance giving rise to “superbugs”

University of BRISTOL

GPU-system DEGIMA

- Used 144 GPUs in parallel for drug docking simulations
  - ATI Radeon HD5870 & Intel i5-2500T
- ~300 TFLOPS single precision
- Courtesy of Tsuyoshi Hamada and Felipe Cruz, Nagasaki

University of BRISTOL

NDM-1 experiment

- 1 million candidate drug molecules times 20 conformers each → 20M dockings
- 1.23×10^{17} atom-atom energies calculated
- 267 days of GPU compute time and 224 days of CPU compute time
- ~55 hours actual wall-time
- A second run with 8 million molecules, 160M conformers on >200 GPUs is running right now!

University of BRISTOL

Conclusions

- OpenCL enables truly heterogeneous computing, harnessing all hardware resources in a system
- GPUs can yield significant savings in energy costs (and equipment costs)
- OpenCL can work just as well for multi-core CPUs as it does for GPUs

It's possible to screen libraries of millions of molecules against complex targets using highly accurate, computationally-expensive methods in one weekend using equipment costing O(£100K)

For an introduction to GPUs

The GPU Computing Revolution – a Knowledge Transfer Report from the London Mathematical Society and the KTN for Industrial Mathematics

- https://ktn.innovateuk.org/web/mathsktn/articles/-/blogs/the-gpu-computing-revolution

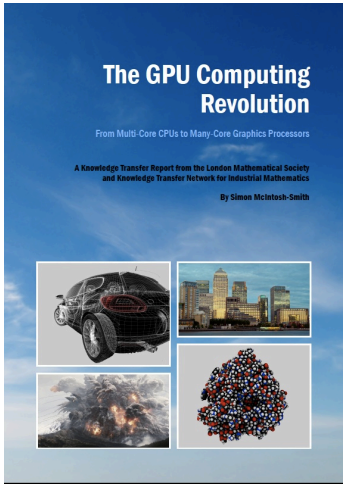

University of BRISTOL

References

- S. McIntosh-Smith, T. Wilson, A.A. Ibarra, J. Crisp and R.B. Sessions, "Benchmarking energy efficiency, power costs and carbon emissions on heterogeneous systems", The Computer Journal, September 12th 2011. DOI: 10.1093/comjnl/bxr091

- N. Gibbs, A.R. Clarke & R.B. Sessions, "Ab-initio Protein Folding using Physicochemical Potentials and a Simplified Off-Lattice Model", Proteins 43:186-202,200