$$\frac{\partial u}{\partial t} = \alpha \nabla^2 u$$

Partially-Ordered Event Triggered Systems, and the Challenges of Event-Based Computing

Mark Vousden: m.vousden@soton.ac.uk
Graeme Bragg: gmb@ecs.soton.ac.uk
Alex Rast: a.d.rast@soton.ac.uk
Jonny Beaumont: j.beaumont@imperial.ac.uk
Matthew Naylor: matthew.naylor@cl.cam.ac.uk
David Thomas: d.thomas1@imperial.ac.uk
Andrey Mokhov: andrey.mokhov@newcastle.ac.uk
Andrew Brown: adb@ecs.soton.ac.uk

Describing Problems

- $\frac{\partial u}{\partial t} = \alpha \nabla^2 u$ (initial-value problem)
- Lockstep iteration vs. globally asynchronous
- Different ways (algorithms) of saying the same thing (physics).

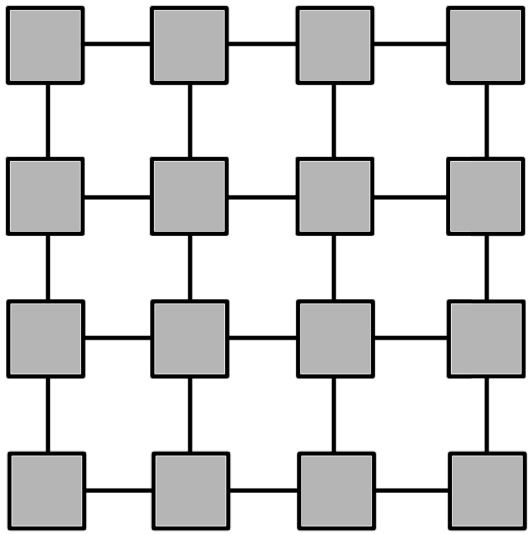

About POETS

- **Objective:** To explore the potential of event-based computing for solving certain types of massive compute problems.
- POETS exploits developments in reconfigurable platforms and the availability of cheaper cores.
- The user defines simple behaviours for the compute cores, which create emergent system dynamics.
- POETS can solve problems orders of magnitudes faster than conventional machines... but is not a general purpose compute machine.

Computers do Everything Poorly

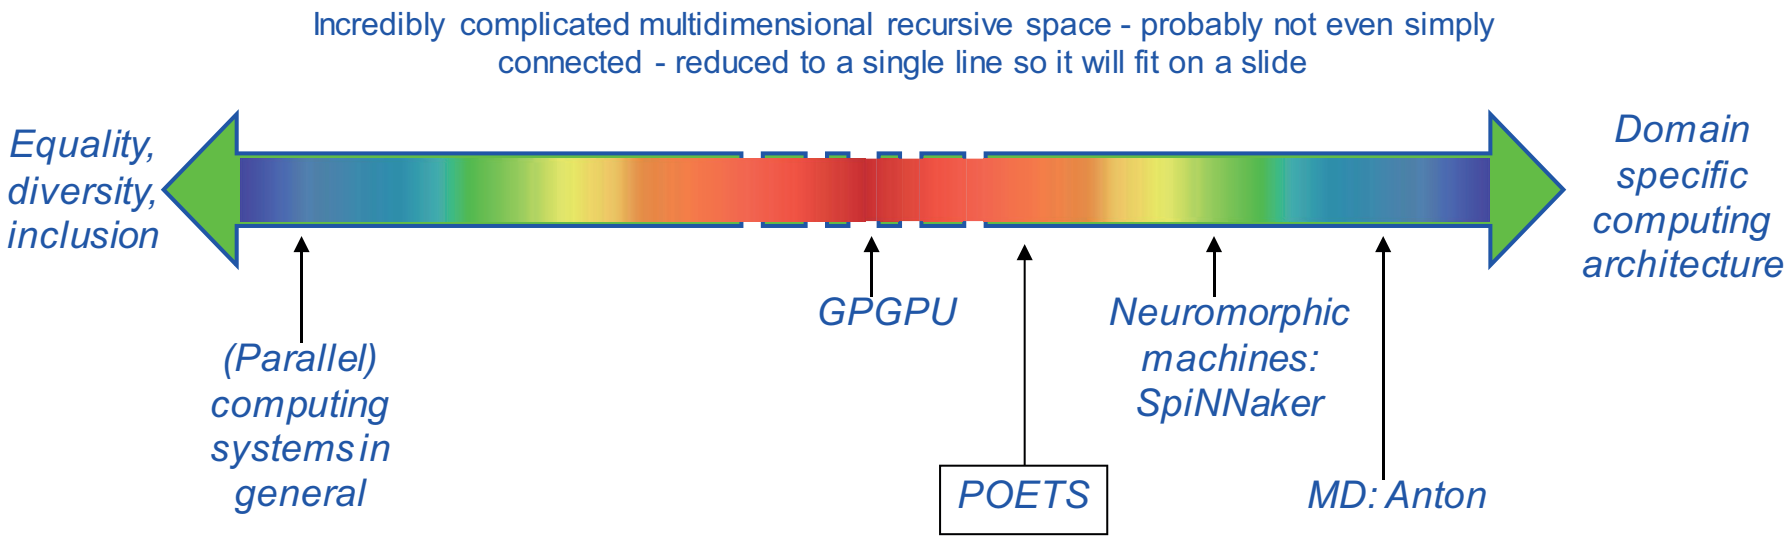

Dissipative Particle Dynamics

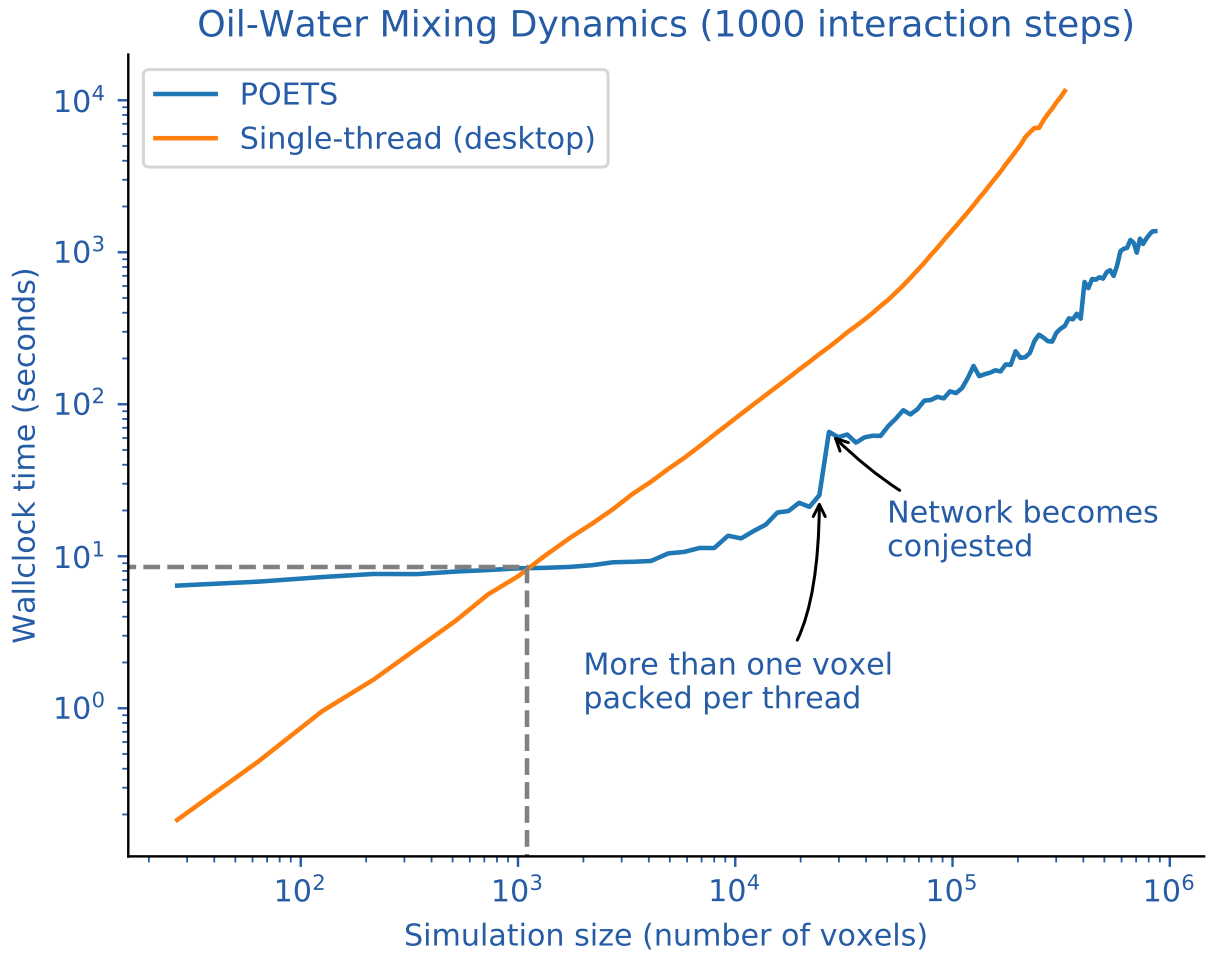

Our Partners

ARM

UNIVERSITY OF CAMBRIDGE

CCFE
CULHAM CENTRE FOR FUSION ENERGY

CU Medicine HONG KONG
香港中文大學醫學院
Faculty of Medicine
The Chinese University of Hong Kong

Imperial College London

e-Therapeutics plc
The Network Pharmacology Company

Imagination

EPSRC
Pioneering research and skills

MAXELER Technologies
MAXIMUM PERFORMANCE COMPUTING

nag®

Newcastle University

UNIVERSITY OF Southampton

nmi Semiconductors to Systems
CONNECTED COMMUNITIES

Uptake

So if this system is so great for large problems that people care about...

Why aren't all scientists using it?

Challenge: Problem Reformulation

- It's hard to leave decades of algorithm theory behind... and to justify picking up another for a slight chance at greatness.
- For some applications, it's not even possible.
- Need to convince:
  - The scientist
  - Their team
  - Paper reviewers
  - Grant panels

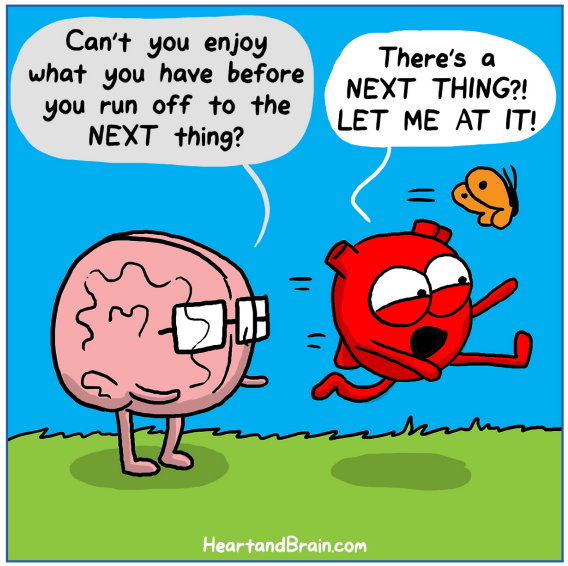

Challenge: Timeliness

- Keeping up with technology.
- The roadside of history is littered with failed bespoke machines that were overtaken by conventional compute before their design was complete.
- Strike while the iron is hot!

The solution?
Better communication!

m.vousden@soton.ac.uk

Applications of Interest

Anything you can describe as a large graph of simple processes, communicating locally and asynchronously:

- **Image processing**
- **ANN pattern recognition**
- **Neural simulation** — 10^3 neurons $\rightarrow$ Thread
- Finite differences/elements/volumes — Physics! Seismology!
- Particle and field — Volume $\rightarrow$ Thread
- Continuous, discrete simulation — Nodes $\rightarrow$ Thread
- Financial fraud detection — Accounts and transactions $\rightarrow$ Graph
- Linear algebra — Matrix-vector, matrix-matrix
- Ray tracing — Pixels $\rightarrow$ Thread
- Weather modelling — Stiff problem, chaos!

Traditional Problem Flow

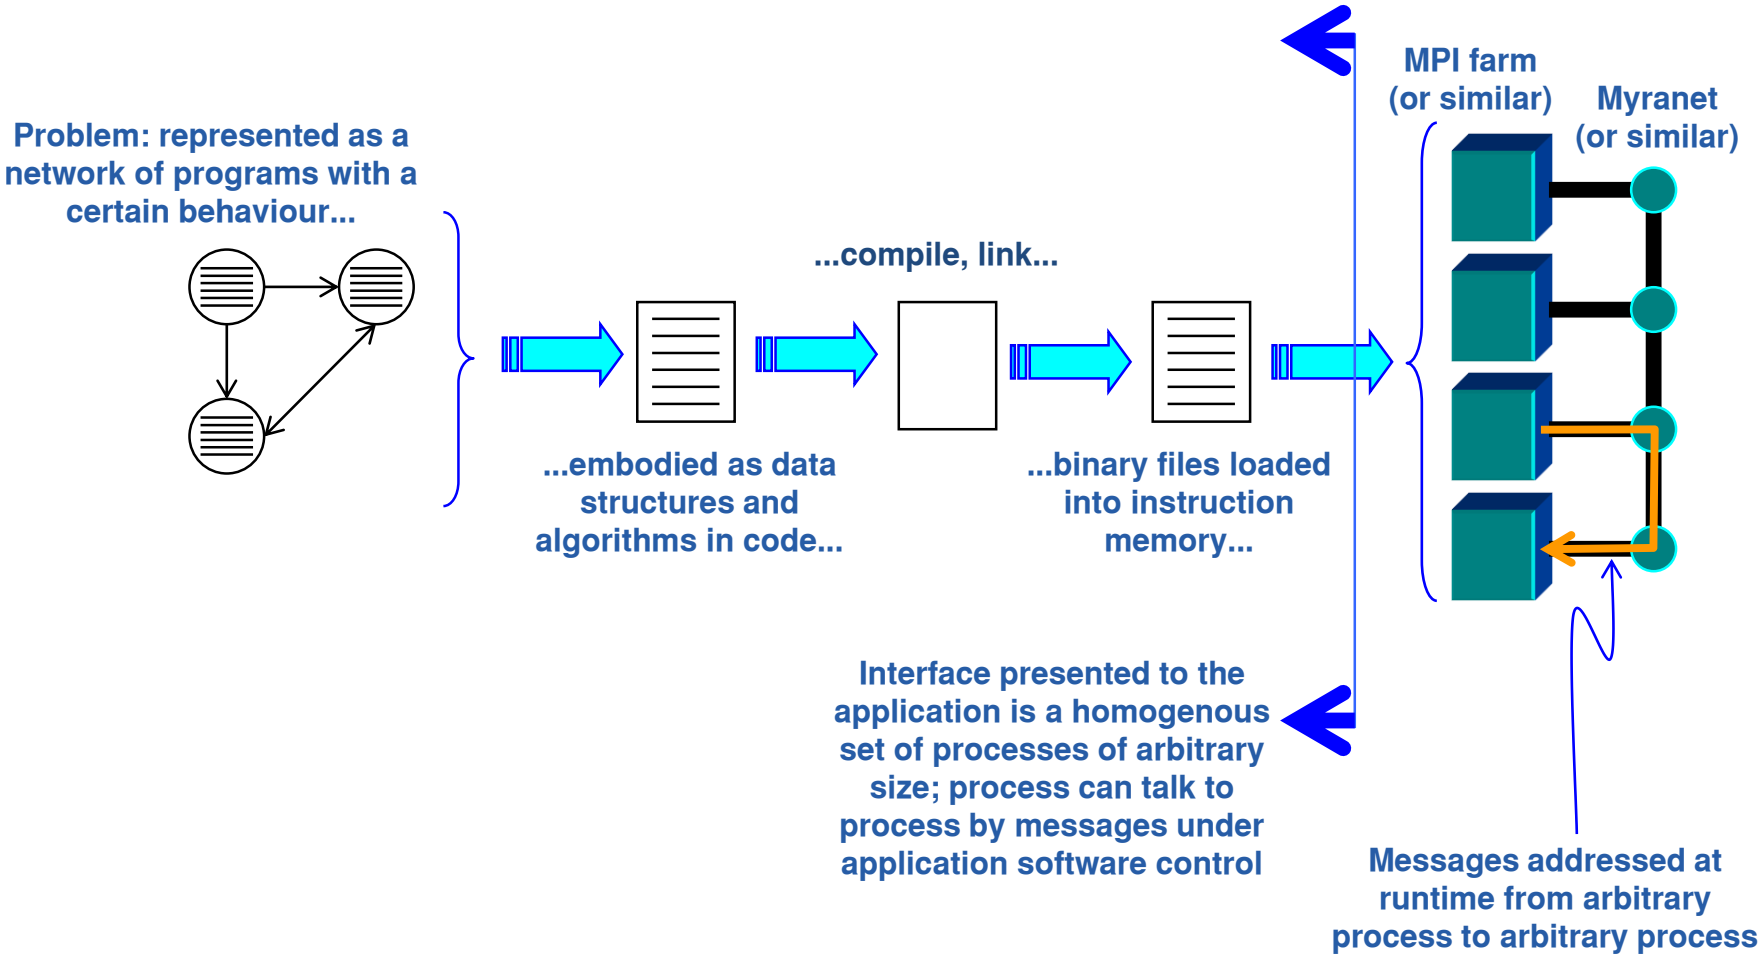

Event-Based Problem Flow

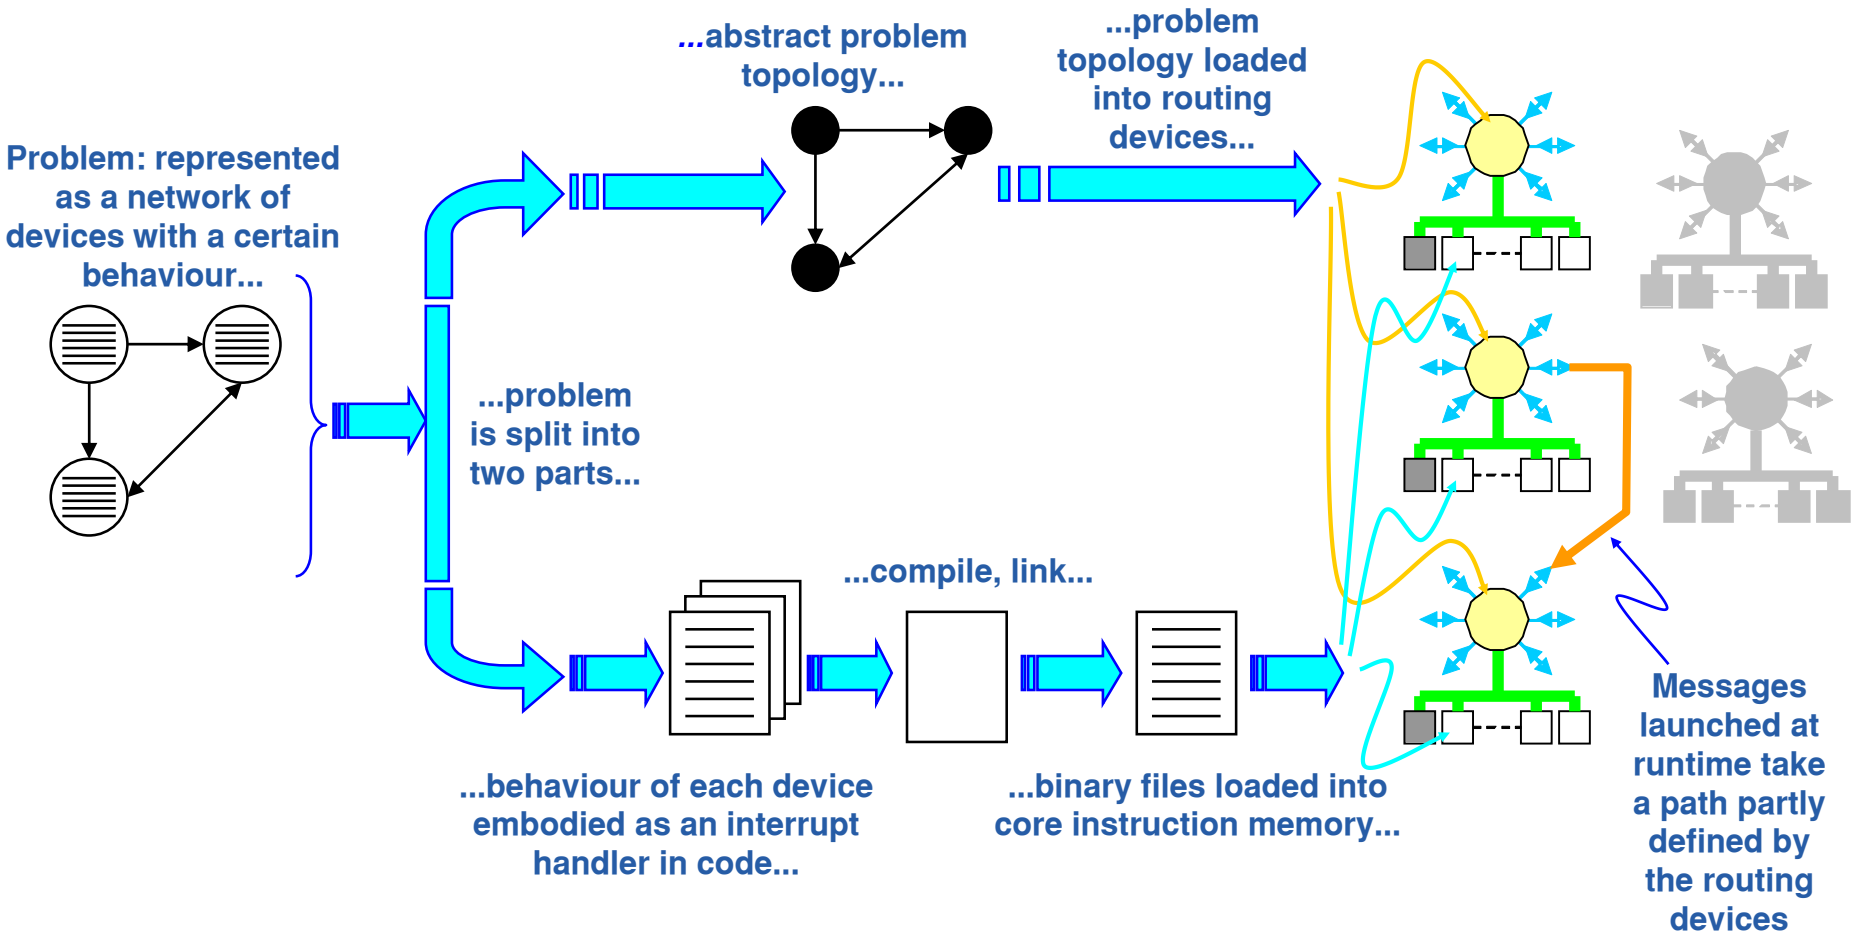

POETS Hardware (for now)

- **Eight x86 servers**, for hybrid compute and C&C.
- **56 DE5-Net FPGA boards** (seven per server, six for dedicated compute) each with:
  - 2×4GB DDR3 DRAM DIMMs
  - 4×8MB QDRII+ SRAM DIMMs
  - 4×**10Gb/s** SFP+ ports (to connect them together)
  - **1 bespoke power management board** for power switching, measurement, and health monitoring. FPGA power: **50W** (busy).
- Custom overlay (Tinsel), providing per board:
  - **64 RISC-V multithreaded (16) barrel cores**, $f_{\mathrm{MAX}} \approx 250$MHz under high loading.
  - **16 FPUs**, with maximum latency of 14 cycles at f_{MAX}.

Power Benchmarks (vs. Xeon Cluster)

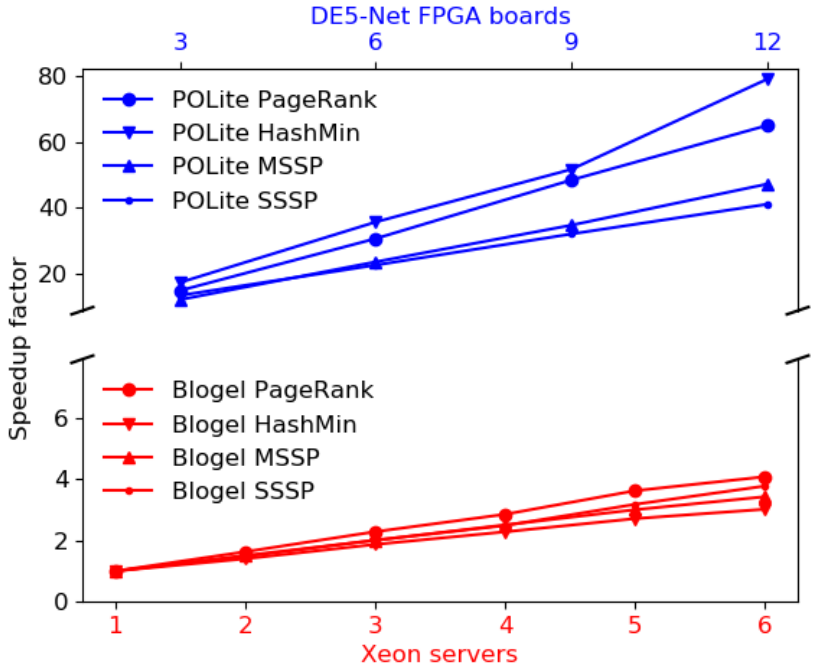

(Horizontal axis normalised by power consumption)

Naylor, M., et al. "Tinsel: a Manythread Overlay for FPGA Clusters." Field-Programmable Logic and Applications (FPL). Newcastle University, 2019.

Heated Plate on POETS

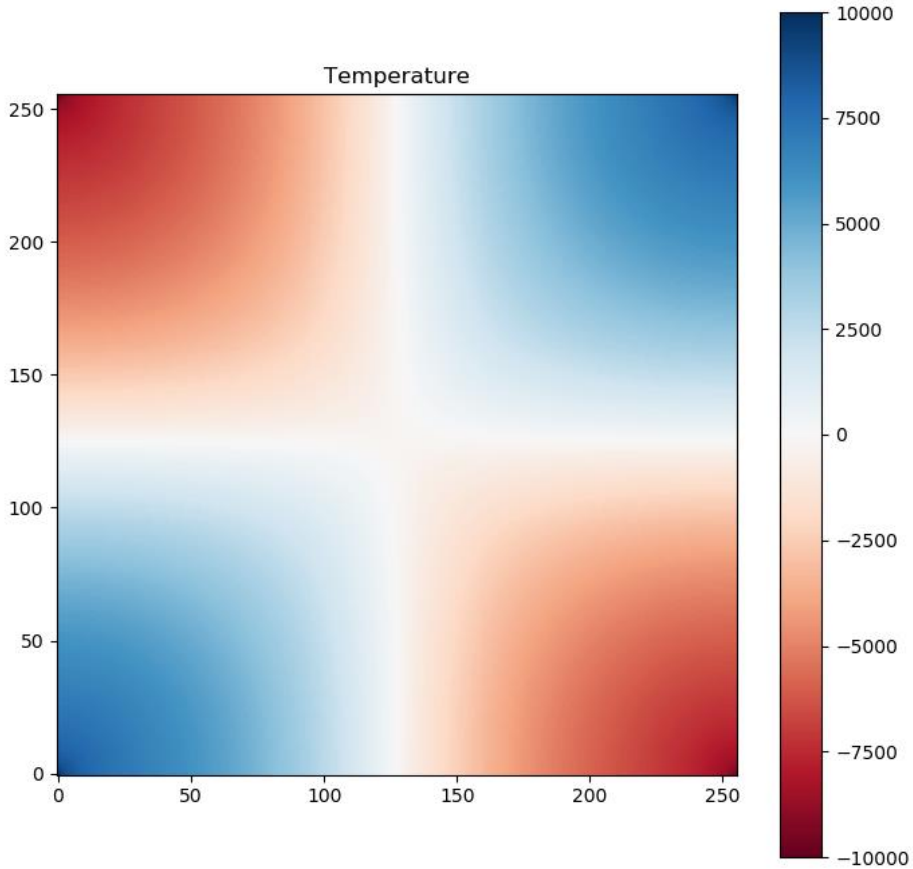

Heated Plate on POETS