Role of heterogeneous computing on the expressway to net zero compute

Jamie Quinn ✉️ *jamie.quinn@ucl.ac.uk* 🐦 @jimjonquinn

University College London • 19-10-2022

Overview

What is heterogeneous computing?

The ENERGETIC project: challenges and early results

How can we effectively use heterogeneous computing to improve energy efficiency?

What is heterogeneous computing?

Hardware

- CPU + co-processor(s)
- E.g. GPU, DPU, FPGA, etc
- Co-processors match algorithmic needs
- E.g. GPU <-> DNNs
- Most large systems are now heterogeneous

Software

Different models:

- Distribute similar work amongst all processing units
- Distribute algorithmic components to approp. processing units

The ENERGETIC project

ENERGETIC

Which classes of algorithms are most energy efficient on which architectures?

ENERGETIC

Secondary questions:

- What is a good, portable way to run energy benchmarks?
- How well are HPC systems set up to measure energy?
- Can we incorporate energy measurements into existing benchmarking frameworks (e.g. ExCALIBUR benchmarking)

ENERGETIC: Benchmarks

Benchmarks chosen to represent large UK HPC workloads:

- HPC-Challenge (including HPL)
- HPCG
- Rodinia (representing 7 dwarves)
- MiniMD (molecular dynamics mini-app)

In total ~10 benchmarks

Lesson: Finding FPGA benchmarks is tricky

ENERGETIC: Energy measurement

Power/energy recorded using:

- Intel RAPL
- nvidia-smi/NVML
- Xilinx Power Analyzer
- AMD (??)
- Intel FPGA (??)

Lesson: Is coarse grained energy measurement better handled by scheduler?

HPL Results (Ns=128000)

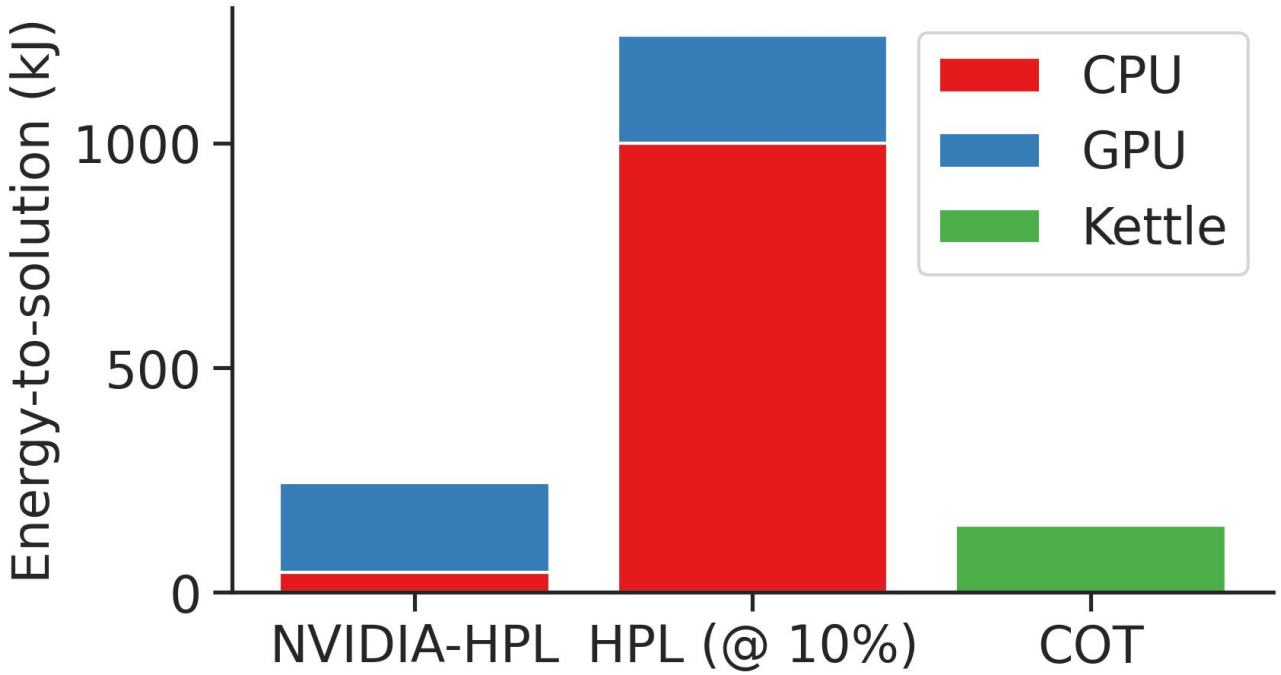

Image Processing Results

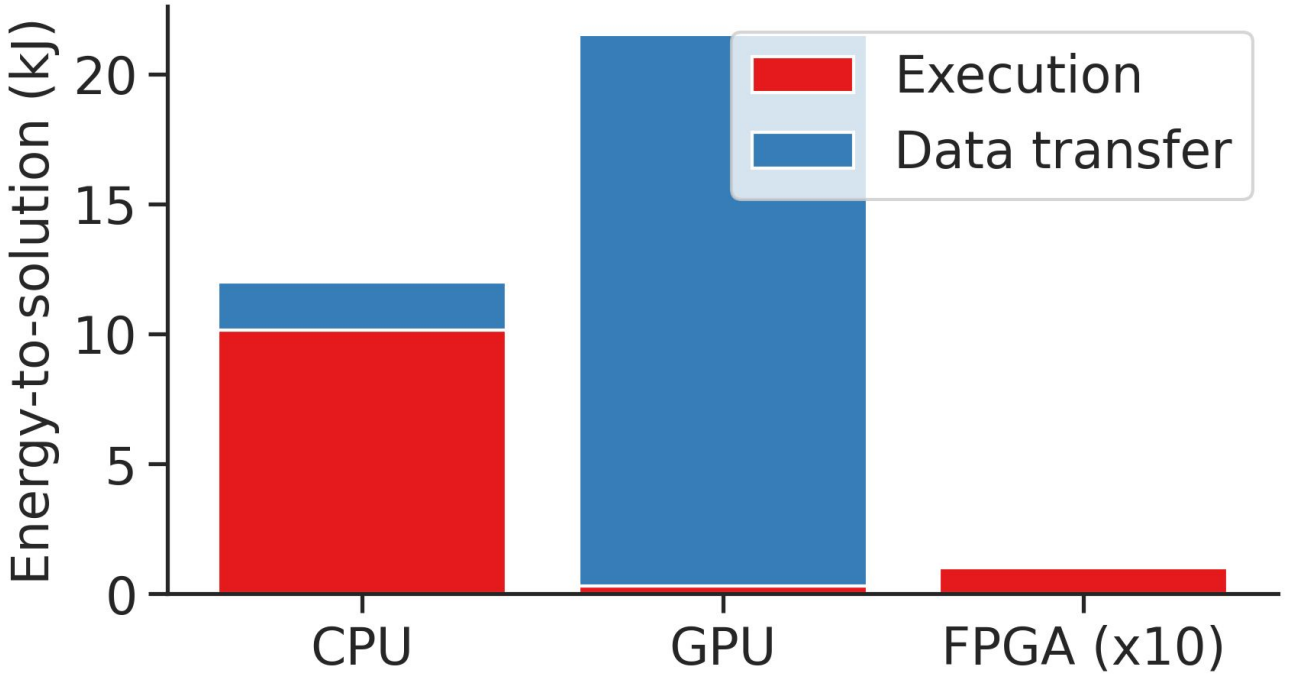

ENERGETIC: Challenges

- There are few HPC benchmarks/mini-apps for FPGAs
- We could not find any “truly” heterogeneous benchmarks
- Energy reporting can require root access
- Every vendor reports energy/power use through different mechanisms
- Whole-node/rack energy reporting is usually unavailable to end user

~~Conclusions~~ questions

- **FPGA HLS is still early days**
  - Is energy efficiency worth excruciating developer experience?
  - More benchmarks required to understand appropriate workloads
- **Heterogeneous code is complex**
  - What is the impact of bugs on energy use?
  - Is it worth the potential energy efficiency?
- **Idle power seems significant**
  - Should all code be as heterogeneous as possible?
  - Or should idle hardware power down further?
- **Userspace power monitoring is tough**
  - Are schedulers the more appropriate place for job energy stats?

How can we use heterogeneous computing to improve energy efficiency?

- Software should **target most energy-efficient hardware** for particular algorithms
- HPC systems should **target specific workloads**
- HPC systems should **expose energy measurements**
- **Benchmarking should inform** porting and procurement decisions