Towards Green Computing: A Survey of Performance and Energy Efficiency of Different Platforms using OpenCL

Philip Heinisch, Katharina Ostaszewski, Hendrik Ranocha

Full paper available on arXiv:
arxiv.org/abs/2003.03794

GitHub:
github.com/IANW-Projects/toolkitICL

When considering hardware platforms, not just time-to-solution can be of importance but also the energy necessary to reach it. This is not only the case with battery powered mobile devices but also with HPC cluster systems due to financial and practical limits on power consumption and cooling. With a variety of hardware options available, the question arises which combination of devices is best suited for a given problem. The answer depends not only on the runtime but also on the energy-to-solution and the price of the hardware. The energy required to reach a solution becomes increasingly important as battery powered systems have to handle computationally intensive tasks e.g. image processing or machine learning. Even for data centers or HPC facilities the energy cost over the lifetime of the systems can be higher than the acquisition cost. To showcase the differences and give a basic outlook on the applicability of different architectures, devices ranging from ARM systems to server CPUs and GPUs have been used with diverse benchmarking test cases taken from applied research applications.

Test Cases

- Based on ToolkitICL profiling & benchmarking tool
- Power measurement using power profiling API or RMS multimeter
- 4 different OpenCL test cases:
  - 2D Median Filter (nonlinear edge preserving filter e.g. image noise)
  - Dot Product (typical microbenchmark, requires data reduction)
  - Cross-Correlation (measure for data similarity e.g. machine learning)
  - Runge-Kutta Solver (iterative, quasi standard for differential equations)

Results

- No hardware specific OpenCL code changes required for any device
- OpenCL performance comparable to OpenMP, MATLAB & C++ AMP
- Energy efficiency improves by 1.22x every two years for CPUs
- Energy efficiency improves by 1.50x every two years for GPUs
- GPUs between 2x and 20x faster than CPUs
- GPUs approx. order of magnitude more energy efficient than CPUs
- Integrated GPUs up to 3x slower than corresponding CPU
- Integrated GPUs comparable to CPUs in energy for most tasks
- Integrated GPUs only better (~2x) in energy with single pre. floating point
- ARM CPUs comparable in required energy to CPUs
- ARM Mali better in energy relative to CPUs, almost comparable to GPUs
- ARM with OpenCL possible alternative for applications constrained by electrical power or in mobile applications

Methods

- Only single nodes
- Tested on MacOS, Windows and Linux
- x86 CPUs with POCL, Intel and AMD OCL framework
- ARM CPUs based on POCL
- Mali GPUs using ARM drivers
- Averaged over thousands of runs with same data
- No biased device specific tuning or optimizations

Hardware

- 22 discrete GPUs & 3 integrated GPUs & 2 Mali GPUs
- 14 server and workstation CPUs & 4 ARM CPUs

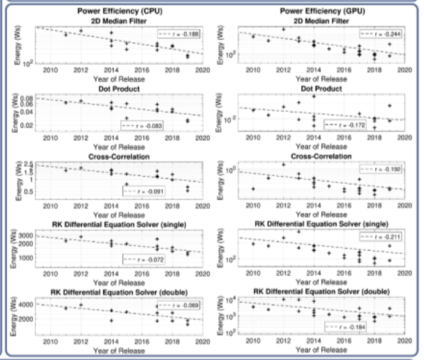

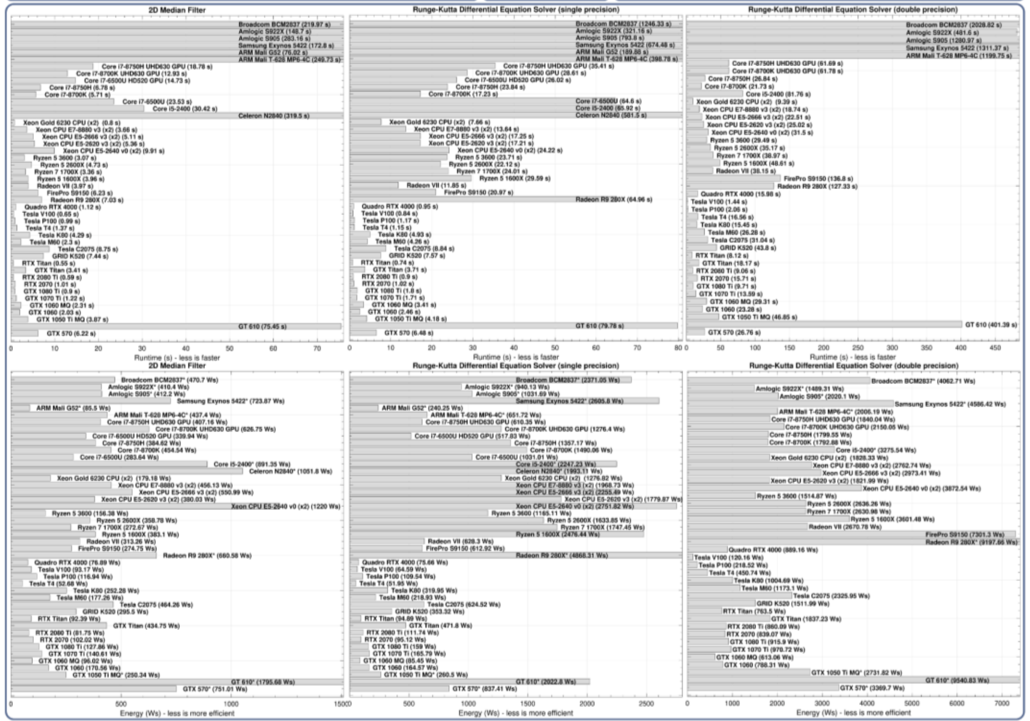

Dr. Philip Heinisch
p.heinisch@tu-bs.de

Technische Universität Braunschweig

IANW Institut für angewandte numerische Wissenschaft e. V.

Towards Green Computing: A Survey of Performance and Energy Efficiency of Different Platforms using OpenCL

When considering hardware platforms, not just time-to-solution can be of importance but also the energy necessary to reach it. This is not only the case with battery powered mobile devices but also with HPC cluster systems due to financial and practical limits on power consumption and cooling. With a variety of hardware options available, the question arises which combination of devices is best suited for a given problem. The answer depends not only on the runtime but also on the energy-to-solution and the price of the hardware. The energy required to reach a solution becomes increasingly important as battery powered systems have to handle computationally intensive tasks e.g. image processing or machine learning. Even for data centers or HPC facilities the energy cost over the lifetime of the systems can be higher than the acquisition cost. To showcase the differences and give a basic outlook on the applicability of different architectures, devices ranging from ARM systems to server CPUs and GPUs have been used with diverse benchmarking test cases taken from applied research applications.

Test Cases

- Based on ToolkitICL profiling & benchmarking tool
- Power measurement using power profiling API or RMS multimeter
- 4 different OpenCL test cases:
  - 2D Median Filter (nonlinear edge preserving filter e.g. image noise)
  - Dot Product (typical microbenchmark, requires data reduction)
  - Cross-Correlation (measure for data similarity e.g. machine learning)
  - Runge-Kutta Solver (iterative, quasi standard for differential equations)

Hardware

- 22 discrete GPUs & 3 integrated GPUs & 2 Mali GPUs
- 14 server and workstation CPUs & 4 ARM CPUs

Methods

- Only single nodes
- Tested on MacOS, Windows and Linux
- x86 CPUs with POCL, Intel and AMD OCL framework
- Nvidia and AMD GPUs using corresponding proprietary drivers
- ARM CPUs based on POCL
- Mali GPUs using official ARM drivers
- Averaged over thousands of runs with same data
- General code optimizations compared against OpenMP and CUDA
- No biased device specific tuning or optimizations

Technische Universität Braunschweig

IANW Institut für angewandte numerische Wissenschaft e. V.

Towards Green Computing: A Survey of Performance and Energy Efficiency of Different Platforms using OpenCL

2D Median Filter

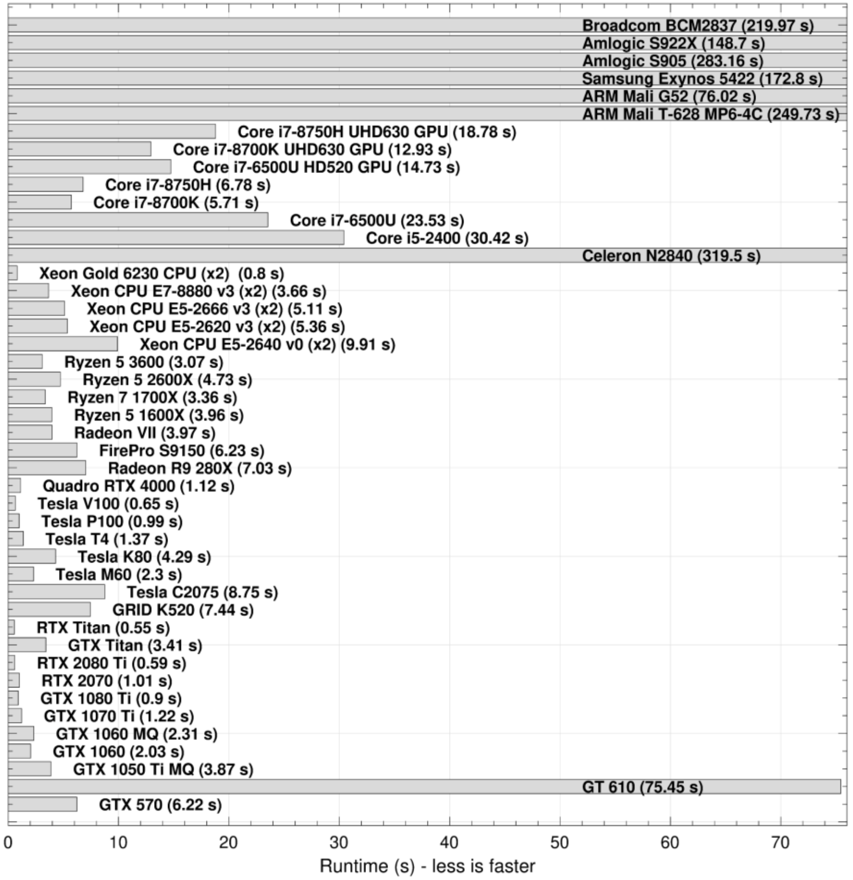

Runge-Kutta Differential Equation Solver (single precision)

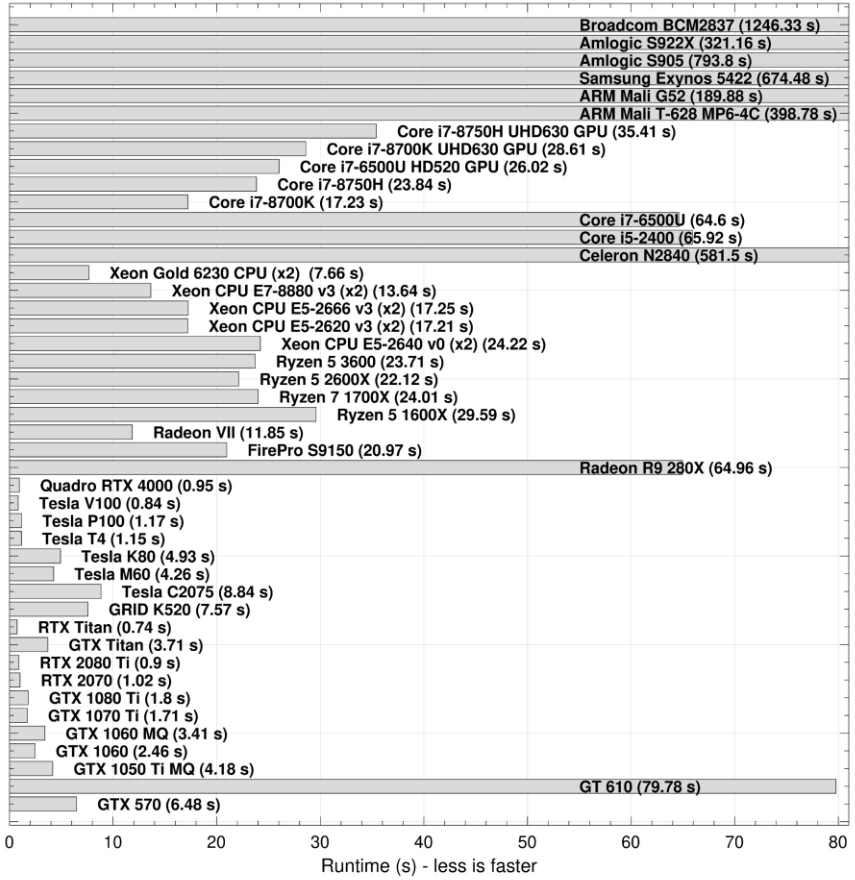

Runge-Kutta Differential Equation Solver (double precision)

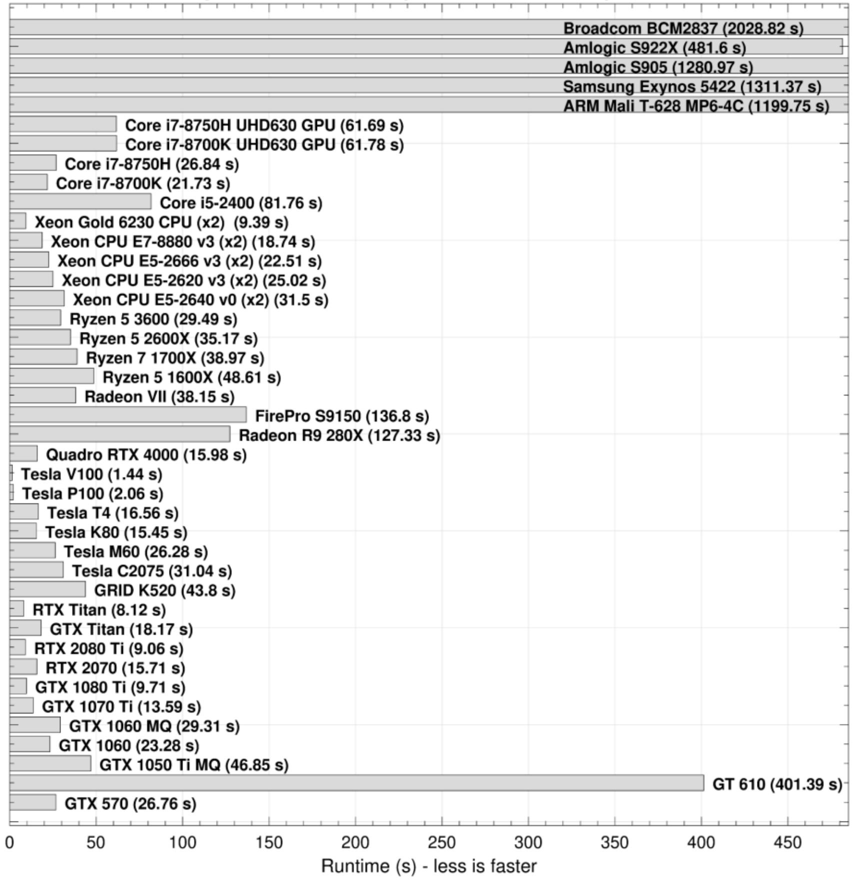

Technische Universität Braunschweig

IANW Institut für angewandte numerische Wissenschaft e. V.

Towards Green Computing: A Survey of Performance and Energy Efficiency of Different Platforms using OpenCL

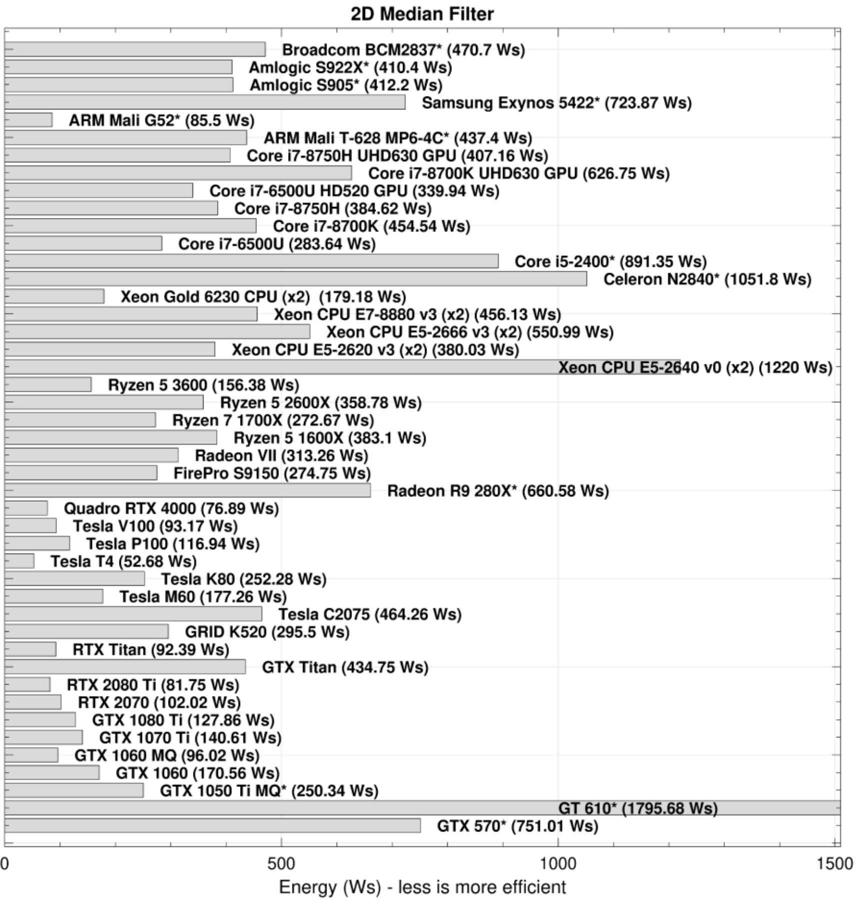

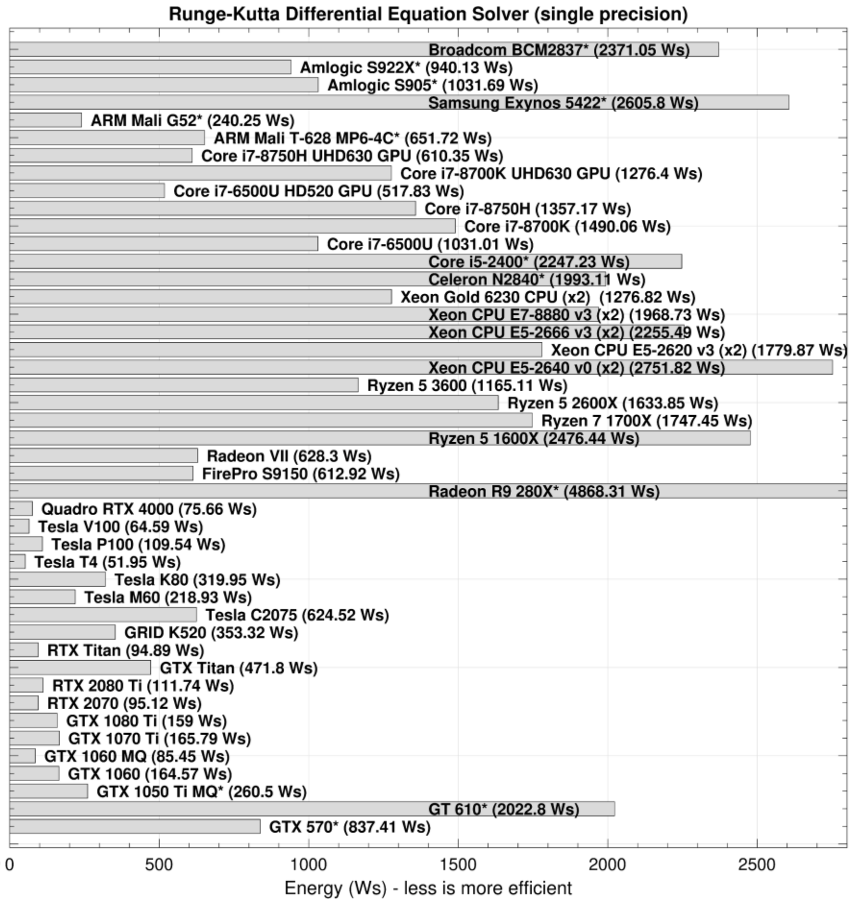

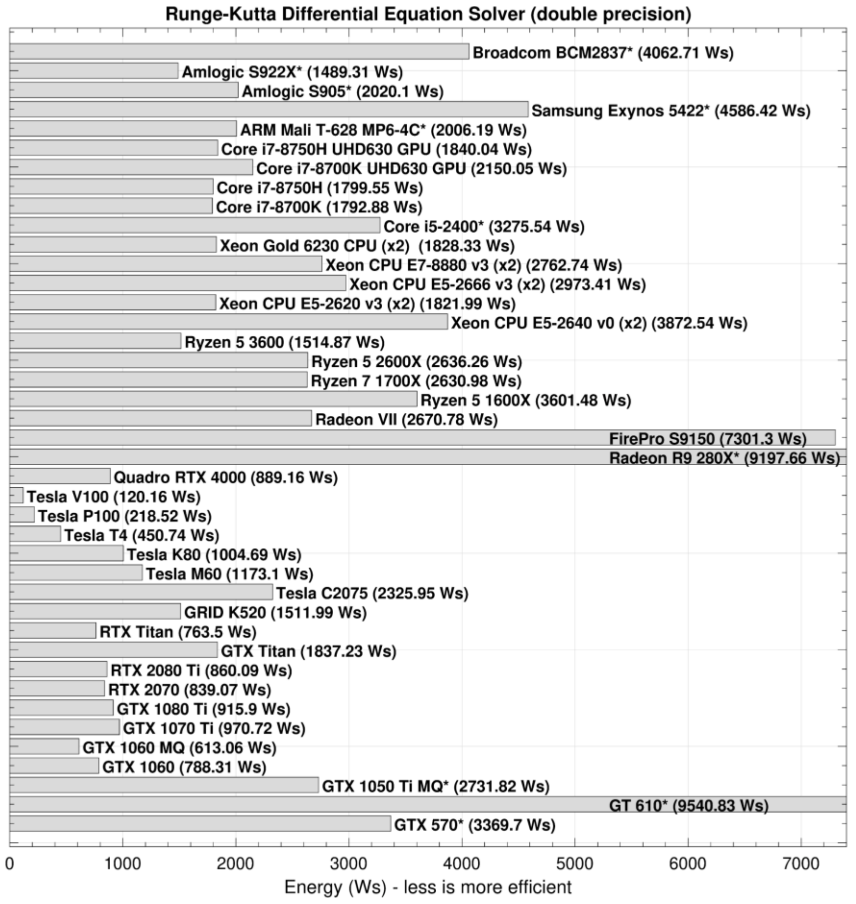

Technische Universität Braunschweig

IANW Institut für angewandte numerische Wissenschaft e. V.

Towards Green Computing: A Survey of Performance and Energy Efficiency of Different Platforms using OpenCL

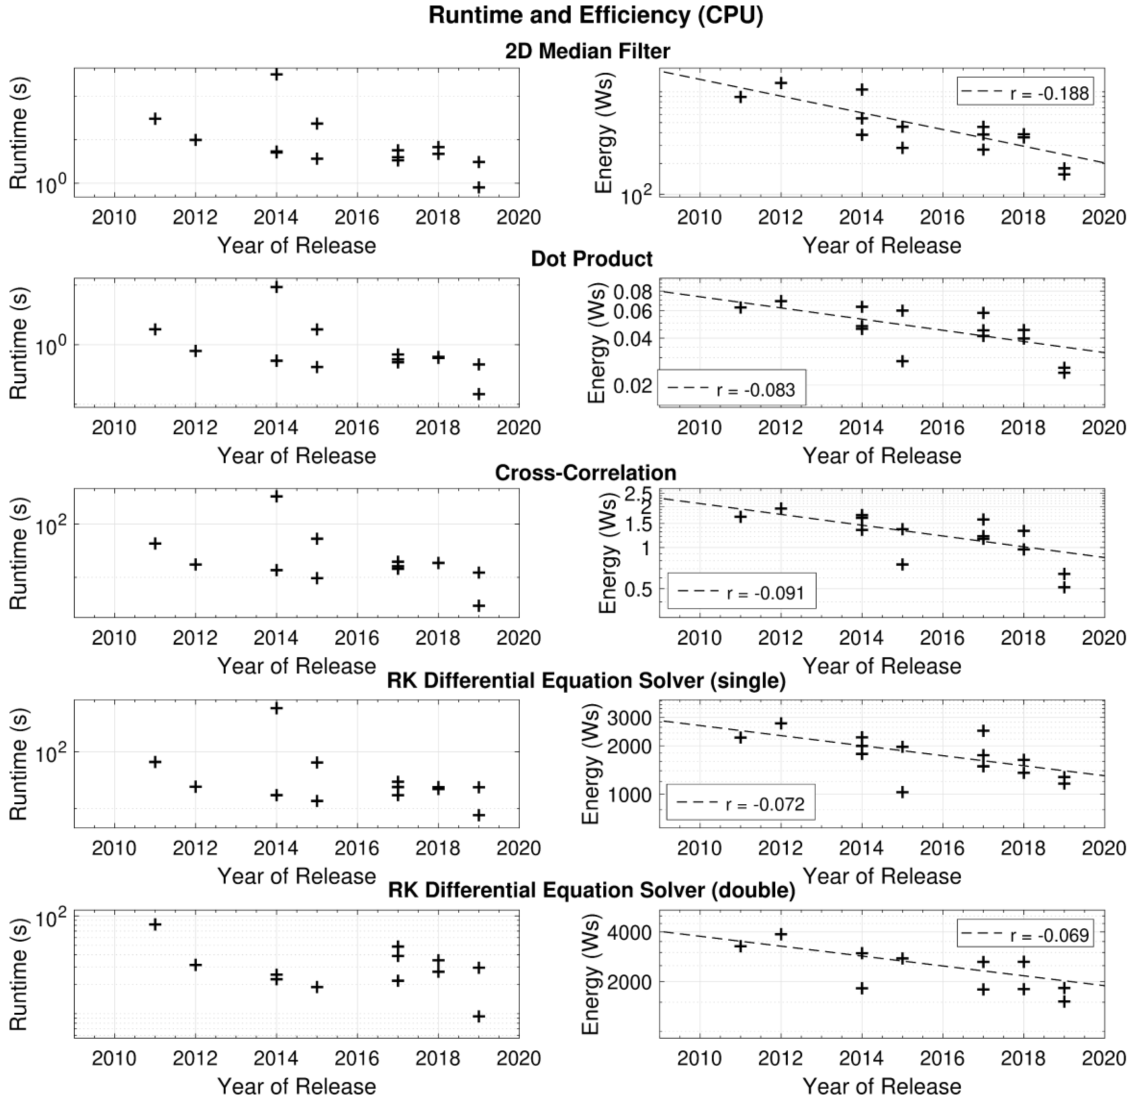

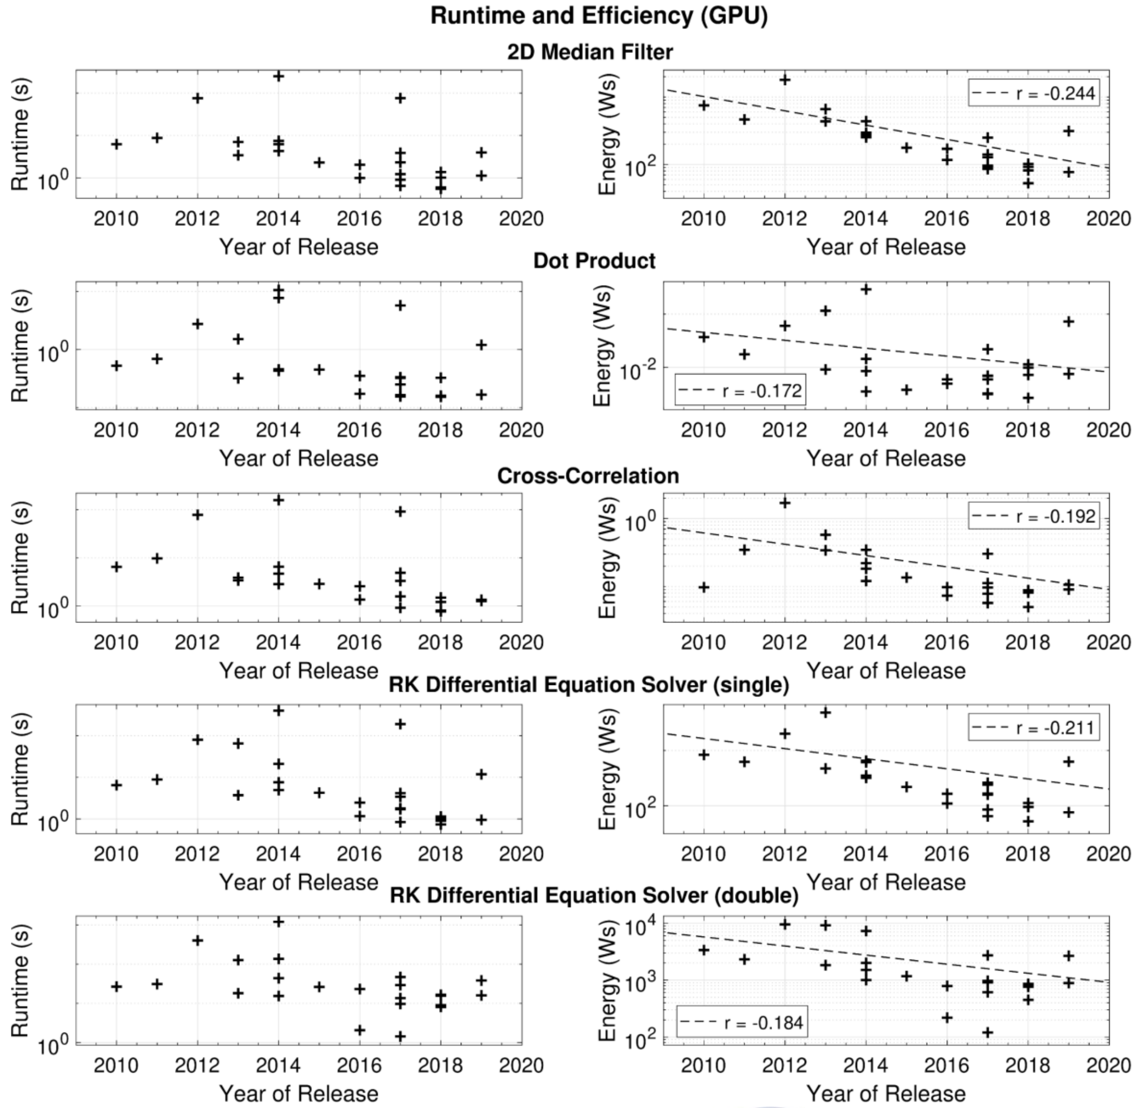

Technische Universität Braunschweig

IANW Institut für angewandte numerische Wissenschaft e. V.

Towards Green Computing: A Survey of Performance and Energy Efficiency of Different Platforms using OpenCL

Results

- No hardware specific OpenCL code changes required for any device
- OpenCL performance comparable to OpenMP, MATLAB & C++ AMP
- Energy efficiency improves by 1.22x every two years for CPUs
- Energy efficiency improves by 1.50x every two years for GPUs
- GPUs between 2x and 20x faster than CPUs
- GPUs approx. order of magnitude more energy efficient than CPUs
- Integrated GPUs up to 3x slower than corresponding CPU
- Integrated GPUs comparable to CPUs in energy for most tasks
- Integrated GPUs only better (~2x) in energy with single pre. floating point
- ARM CPUs comparable in required energy to CPUs
- ARM Mali better in energy relative to CPUs, almost comparable to GPUs
- ARM with OpenCL possible alternative for applications constrained by electrical power or in mobile applications

Full paper available on arXiv:
arxiv.org/abs/2003.03794

GitHub:
github.com/IANW-Projects/toolkitICL

Technische Universität Braunschweig

IANW Institut für angewandte numerische Wissenschaft e. V.