

**There is no excuse why your software is not running on
CPU and GPU!**

al**aka**

Who am I?

- René Widera
- Institute for Radiation Physics / Laser Particle Acceleration Division
- Application developer / full stack HPC developer
- Maintainer/ developer of
 - PIconGPU
 - alpaka/ cupla
 - mallocMC
 - ISAAC
 - cuda_memtest
 - redGrapes

The Mission and Reality

- Providing best re-usable research tools for our scientists
- Growing algorithm complexity
- Limited number of developers
- Large amount of new architectures



What is my Framework?

- alpaka - Abstraction Library for Parallel Kernel Acceleration
- C++17 and newer
- CMake build system, stand alone headers and single header
- Open source and open development
- Supports: NVIDIA GPUs, AMD GPUs, INTEL GPUs, x86, ARM, OpenPower, RISC-V and FPGAs

Features

- Zero overhead abstraction of vendor API's
- Support kernel specialization for a single architecture (write assembler code if needed)

What are our Challenges?

- Finding a common feature set supported by all platforms/ vendor API's
- Describe many-core parallelism that is understandable by the compiler
- How to scale the development and testing for the large zoo of platforms
- Provide an easy to use but explicit programming model



Contribute to the Community

- Use alpaka in you projects
- Open issues
- Provide features and bug fixes
- Bring in new ideas

<https://github.com/alpaka-group>



WE WANT YOU!