Multi-Device GPU-Code Generation with LLVM

Philipp Schaad

Bachelor Thesis
April - September 2017

Dr. Tobias Grosser
Prof. Dr. Torsten Hoefler

ETH zürich

Why GPU-Code Generation?

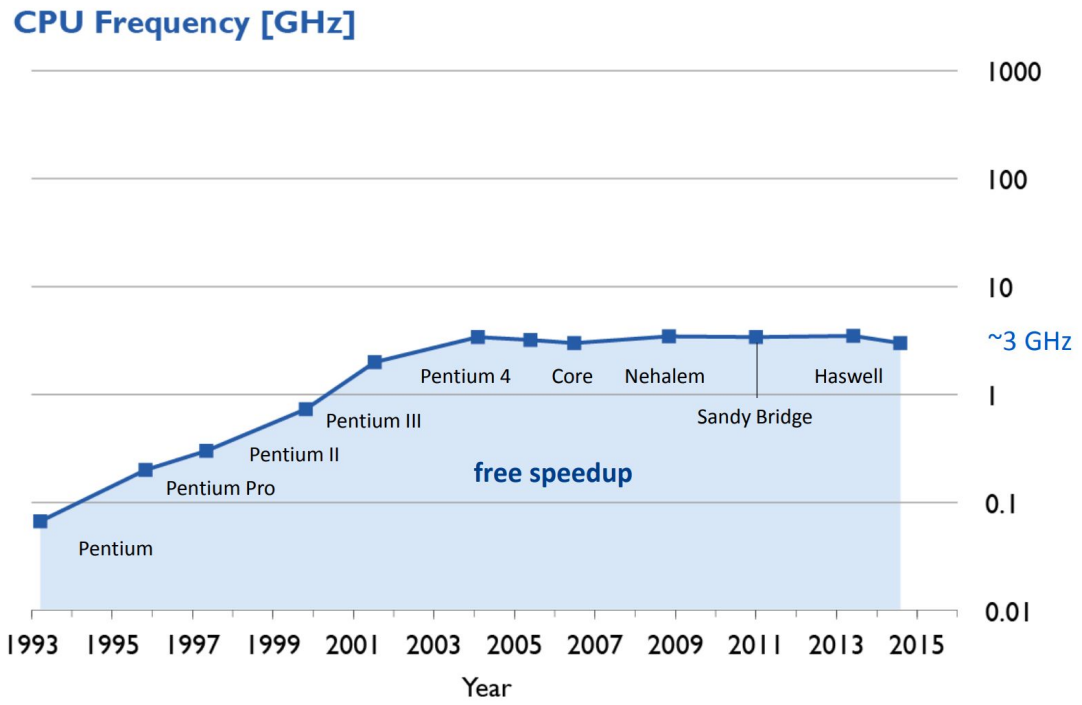

Source: DPHPC Lecture Slides FS17

Why GPU-Code Generation?

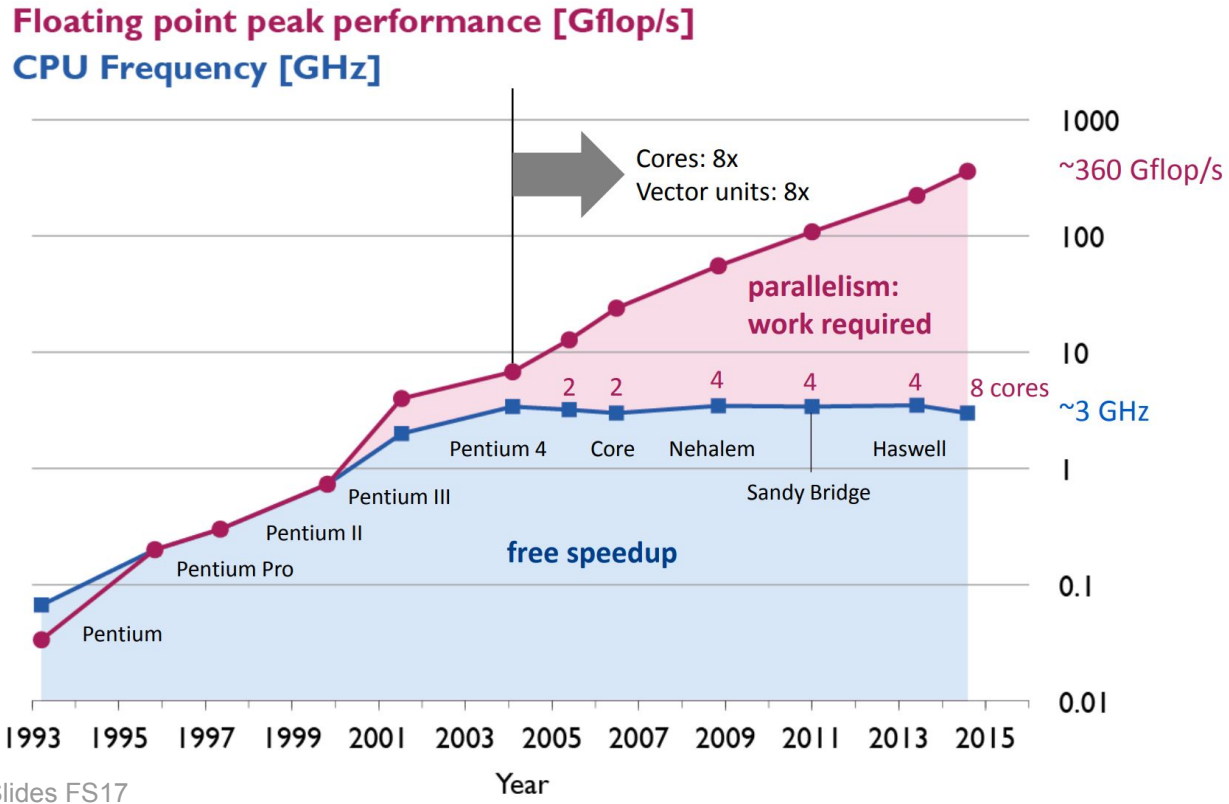

Source: DPHPC Lecture Slides FS17

Why GPU-Code Generation?

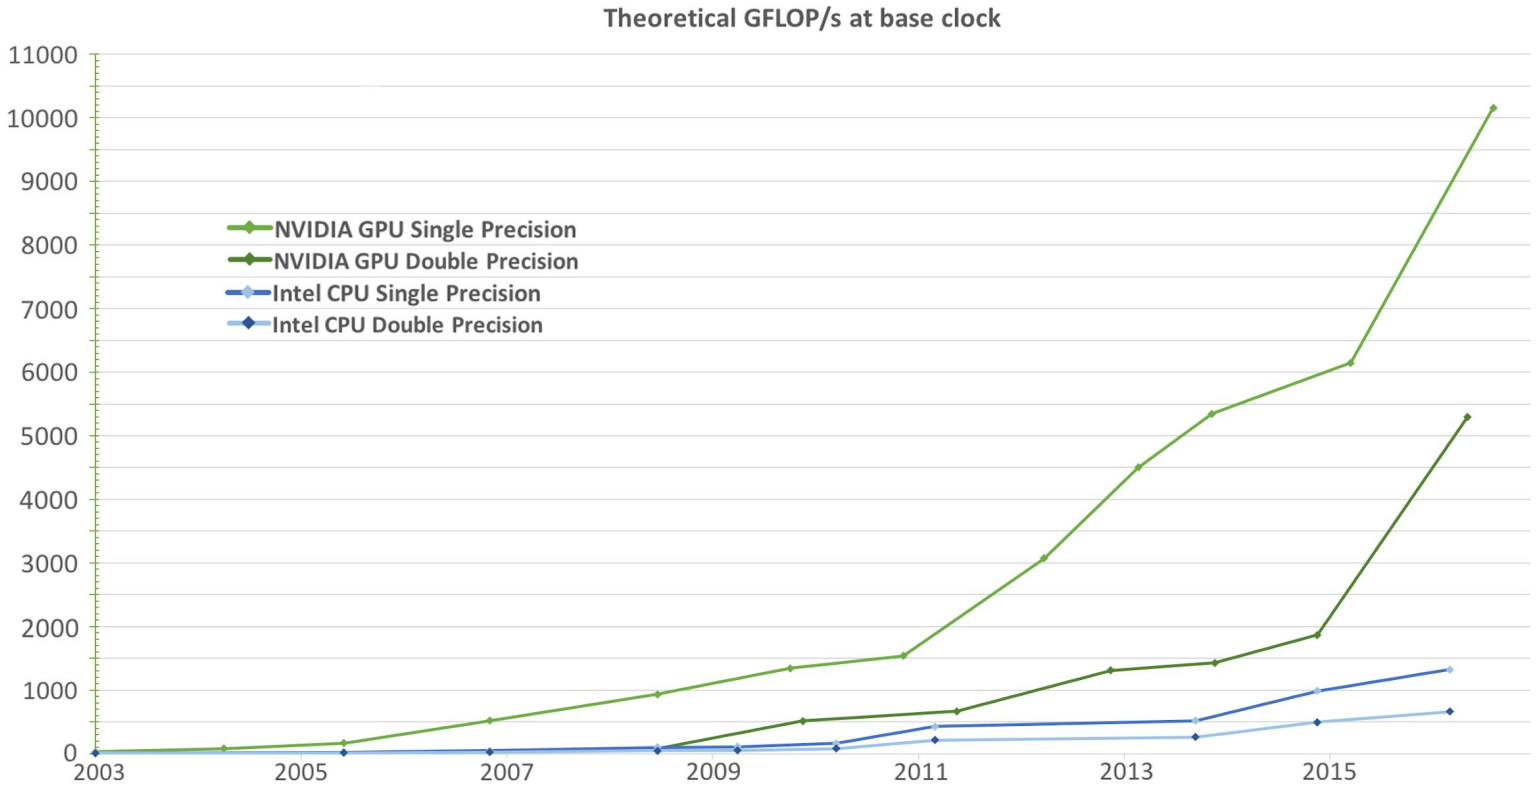

Source: NVIDIA CUDA C Programming Guide (modified)

Why GPU-Code Generation?

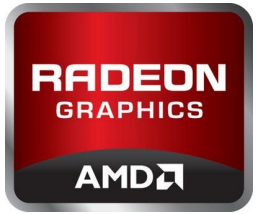

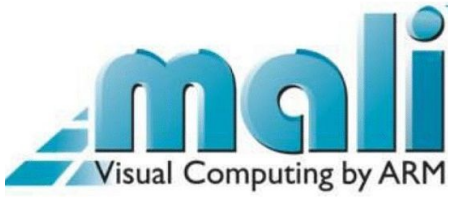

....

Why GPU-Code Generation?

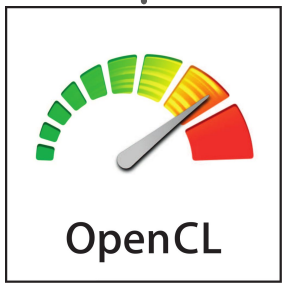

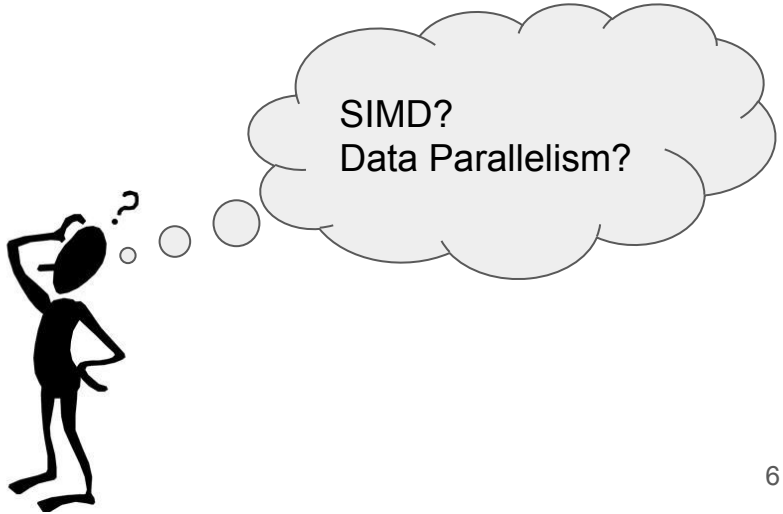

Why GPU-Code Generation?

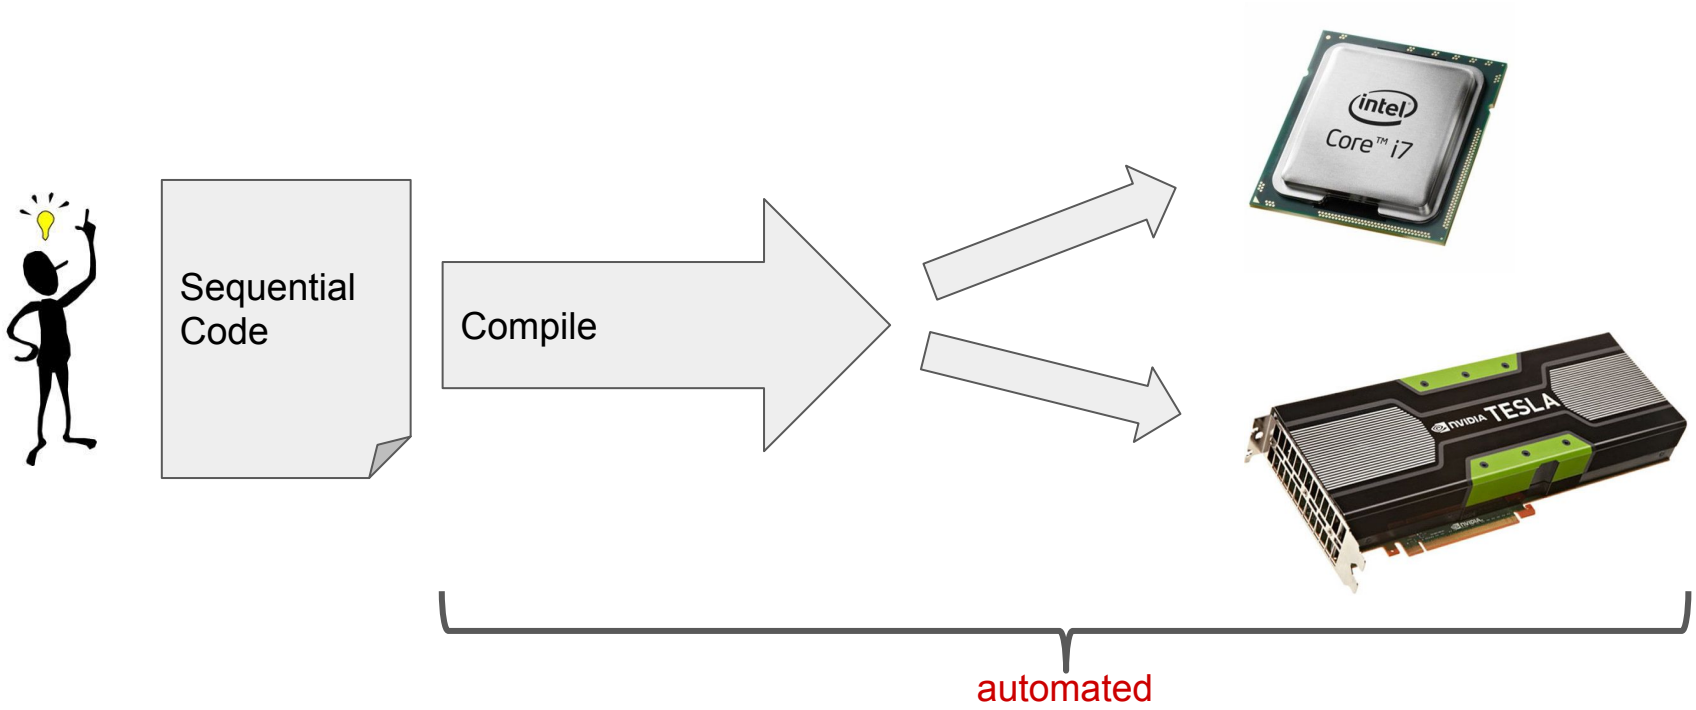

Why LLVM?

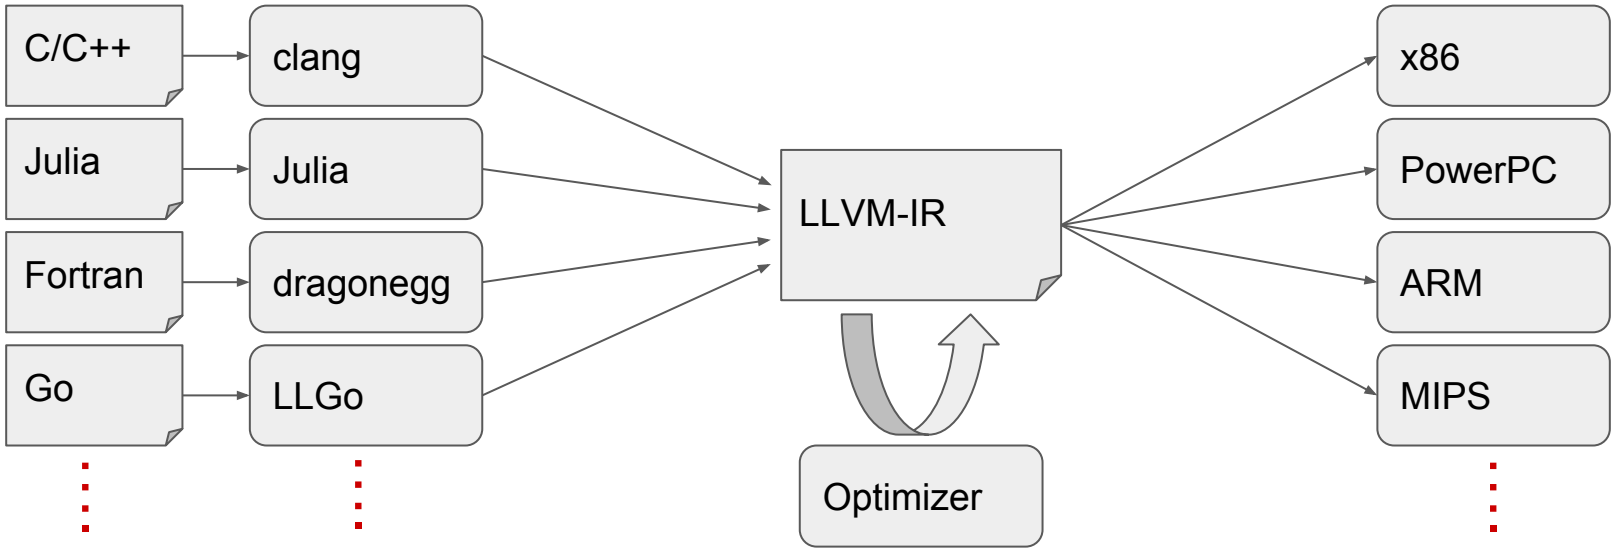
C/C++
Julia
Fortran
Go
clang
Julia
dragonegg
LLGo
LLVM-IR
Optimizer
x86
PowerPC
ARM
MIPS

Why LLVM?

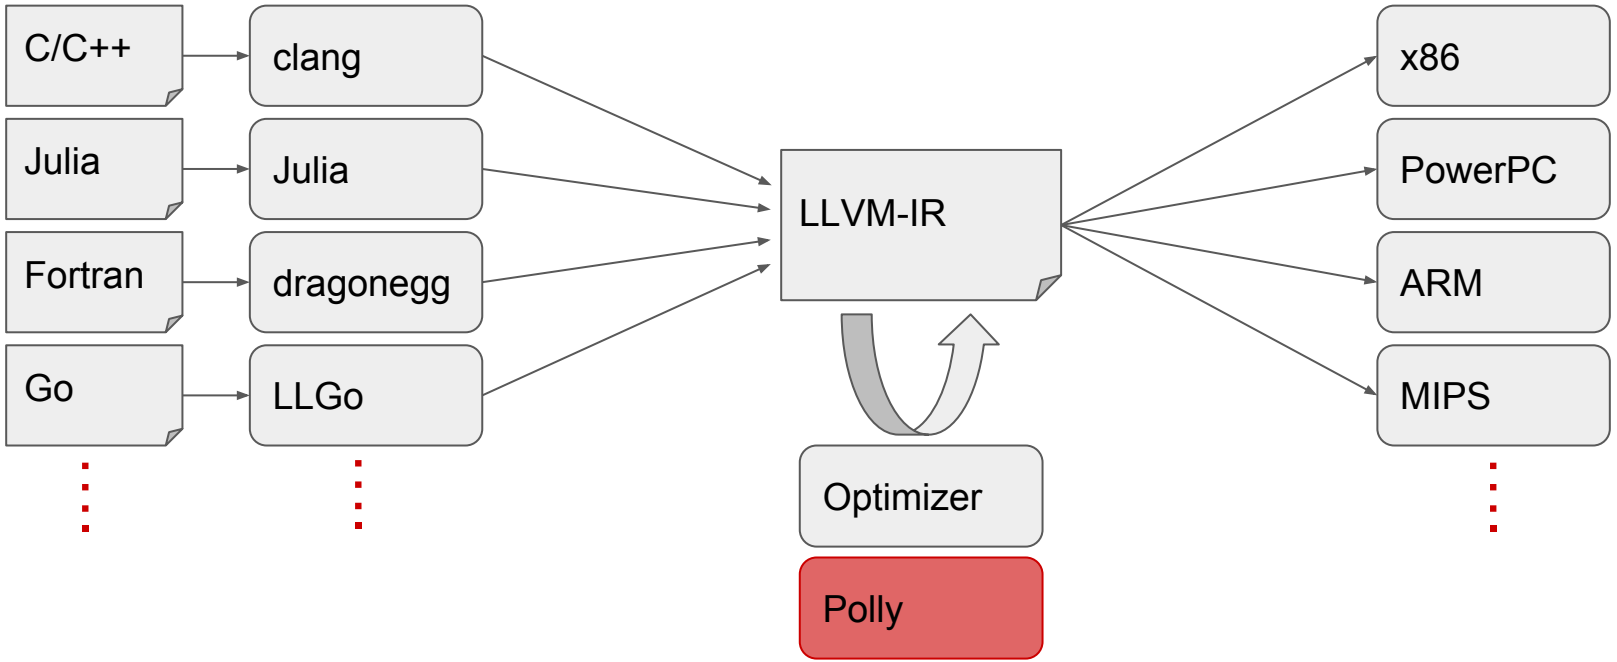
C/C++
clang
Julia
Julia
Fortran
dragonegg
Go
LLGo
LLVM-IR
Optimizer
Polly
x86
PowerPC
ARM
MIPS

Polly's Architecture

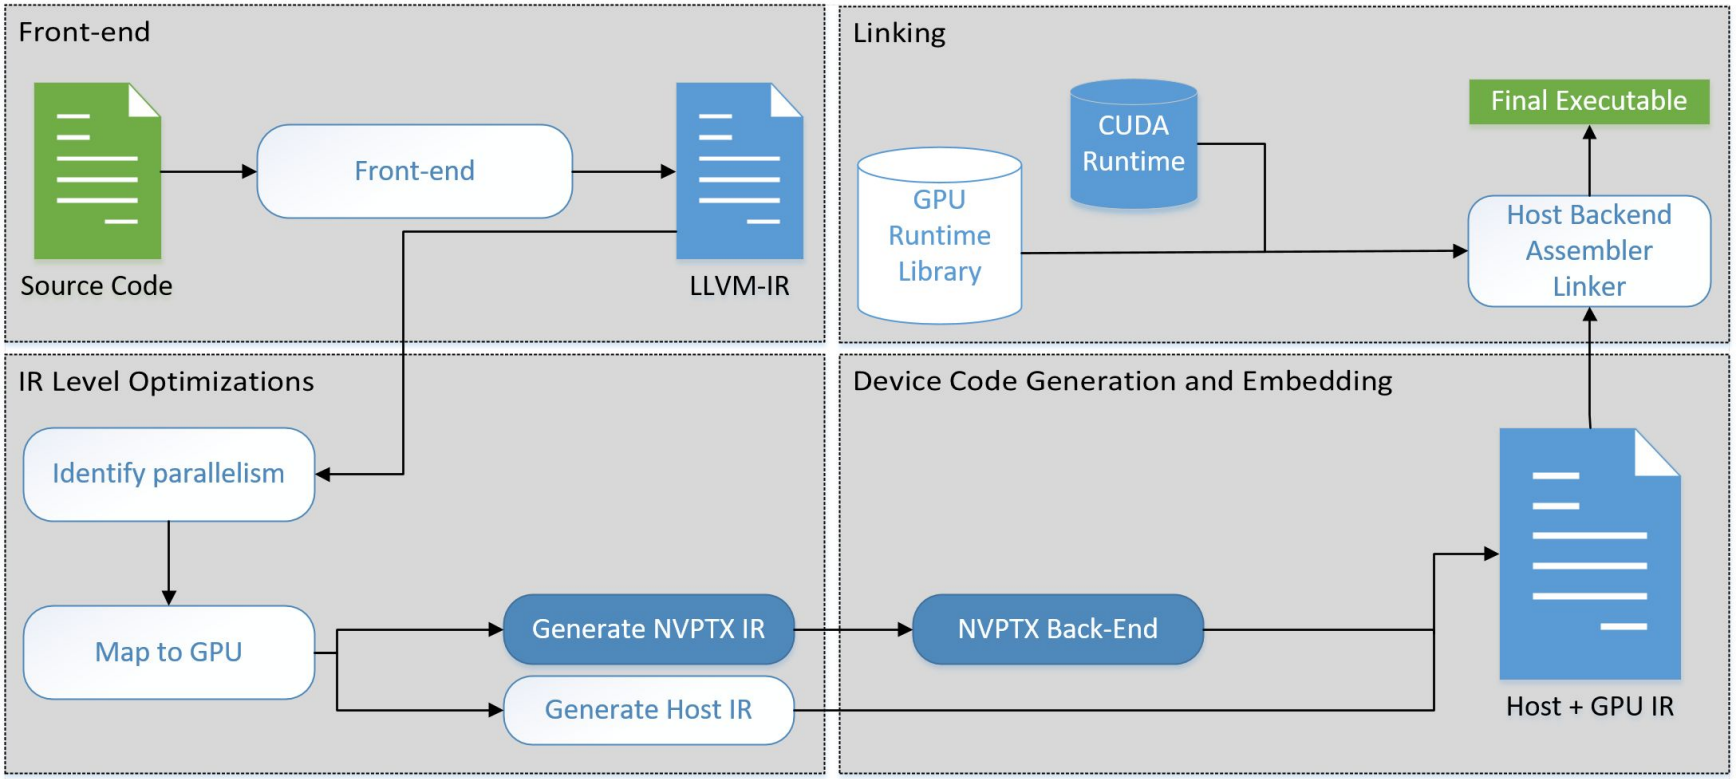

Extending Polly

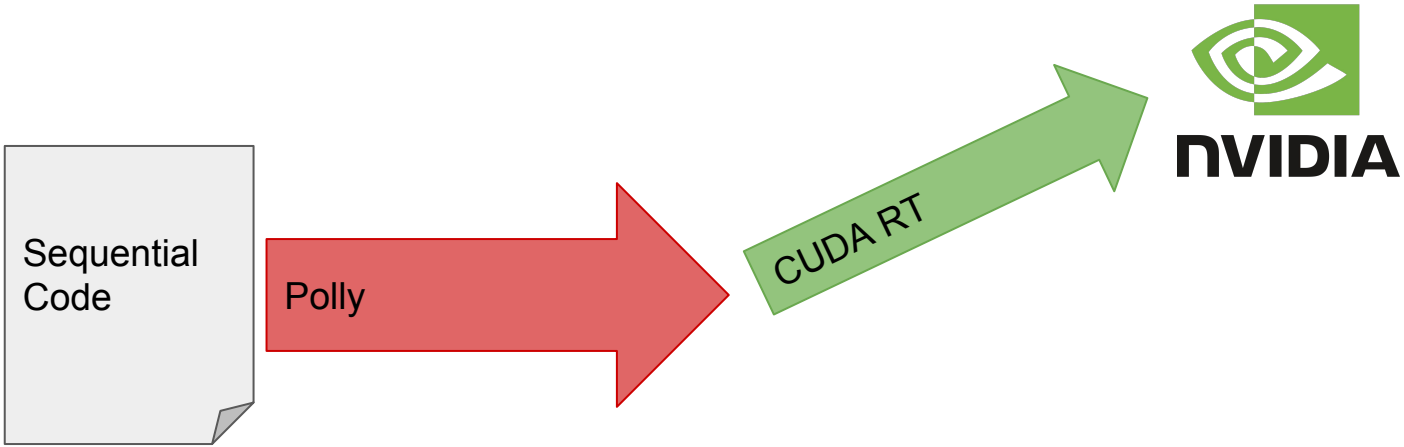

Extending Polly

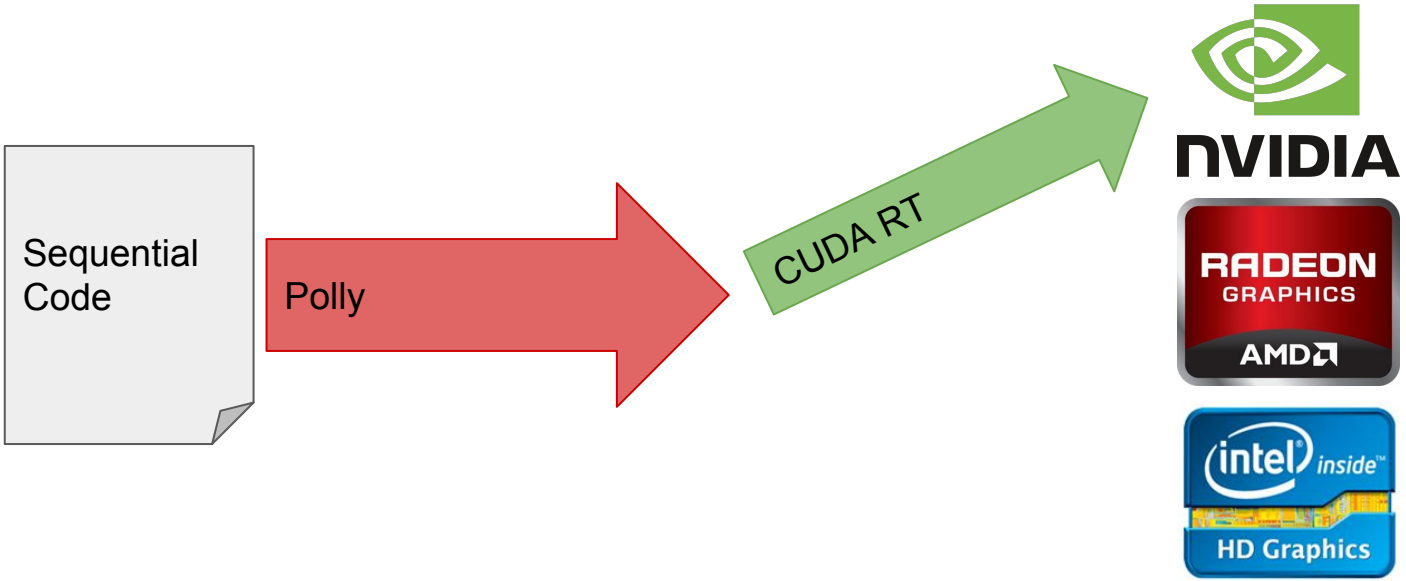

Extending Polly

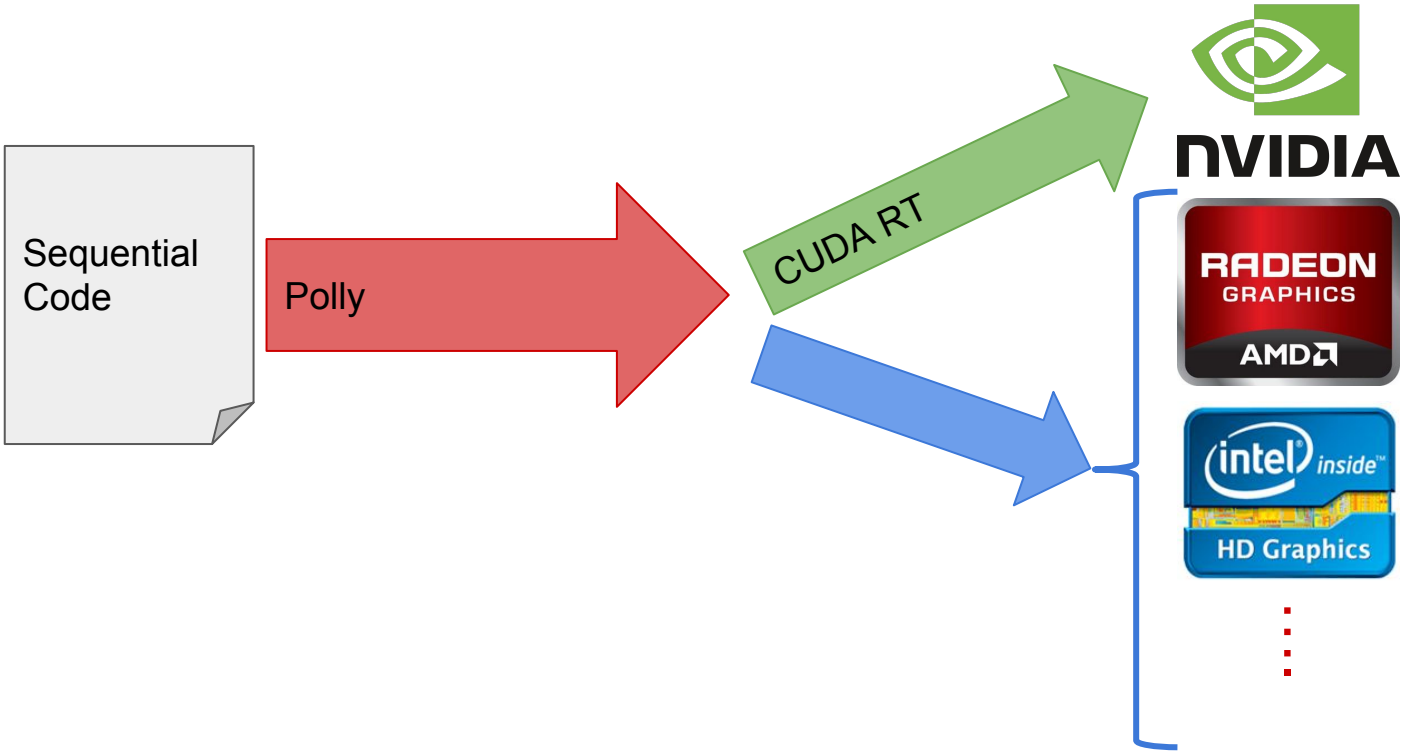

CUDA as a Limiting Factor

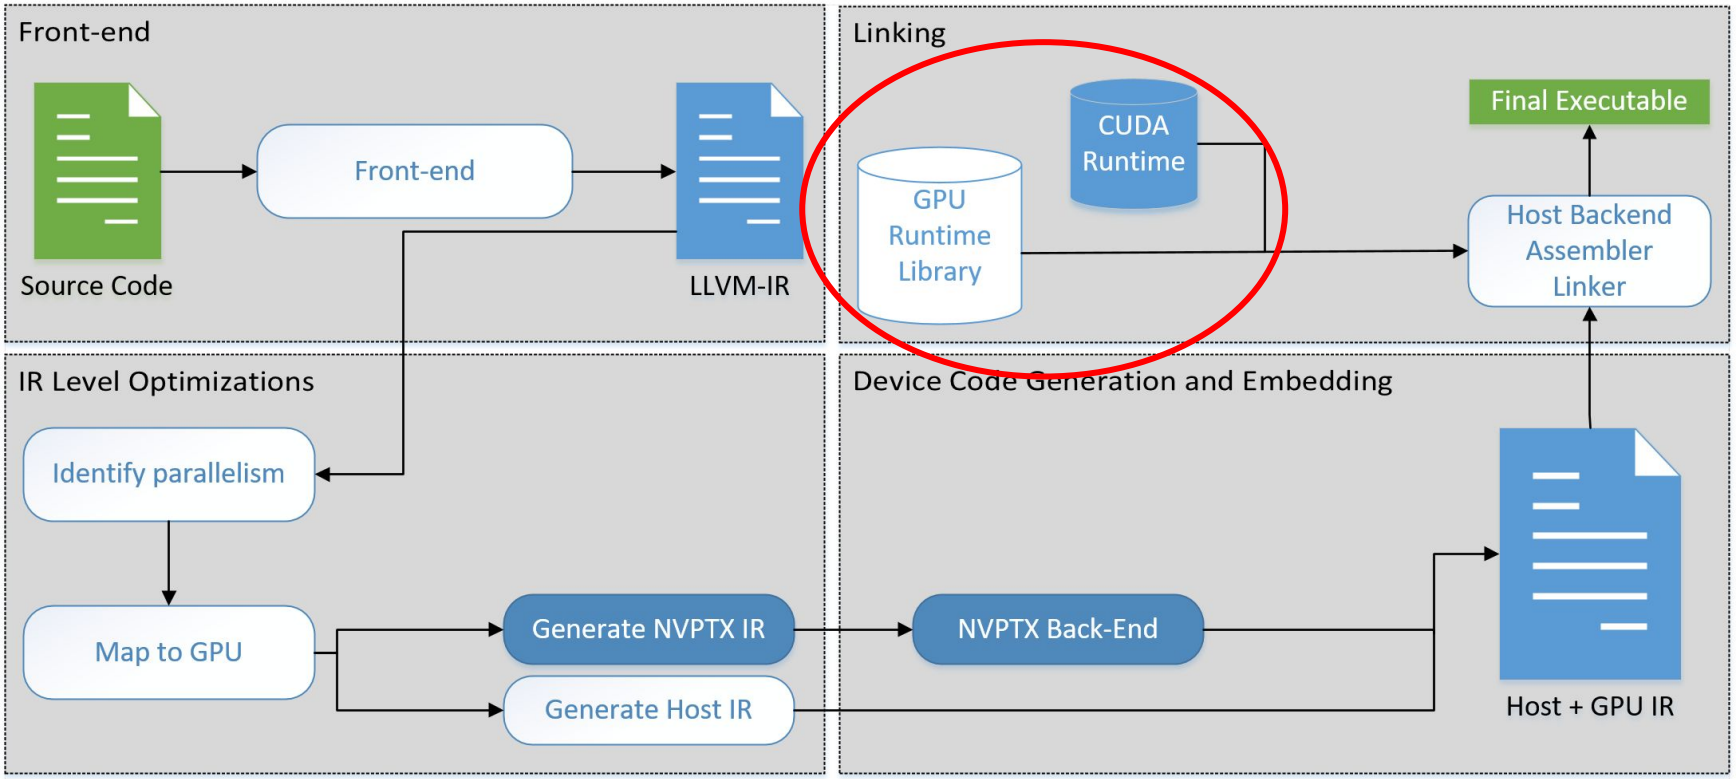

OpenCL Interface

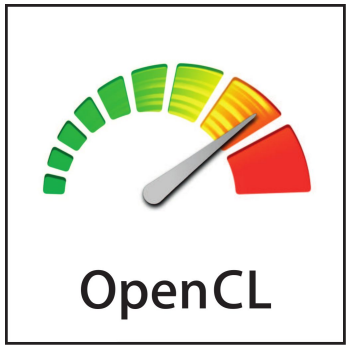

OpenCL Interface

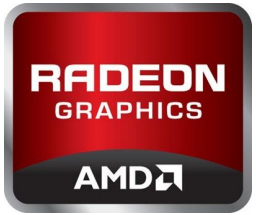

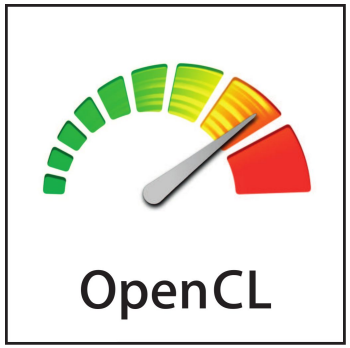

OpenCL Interface

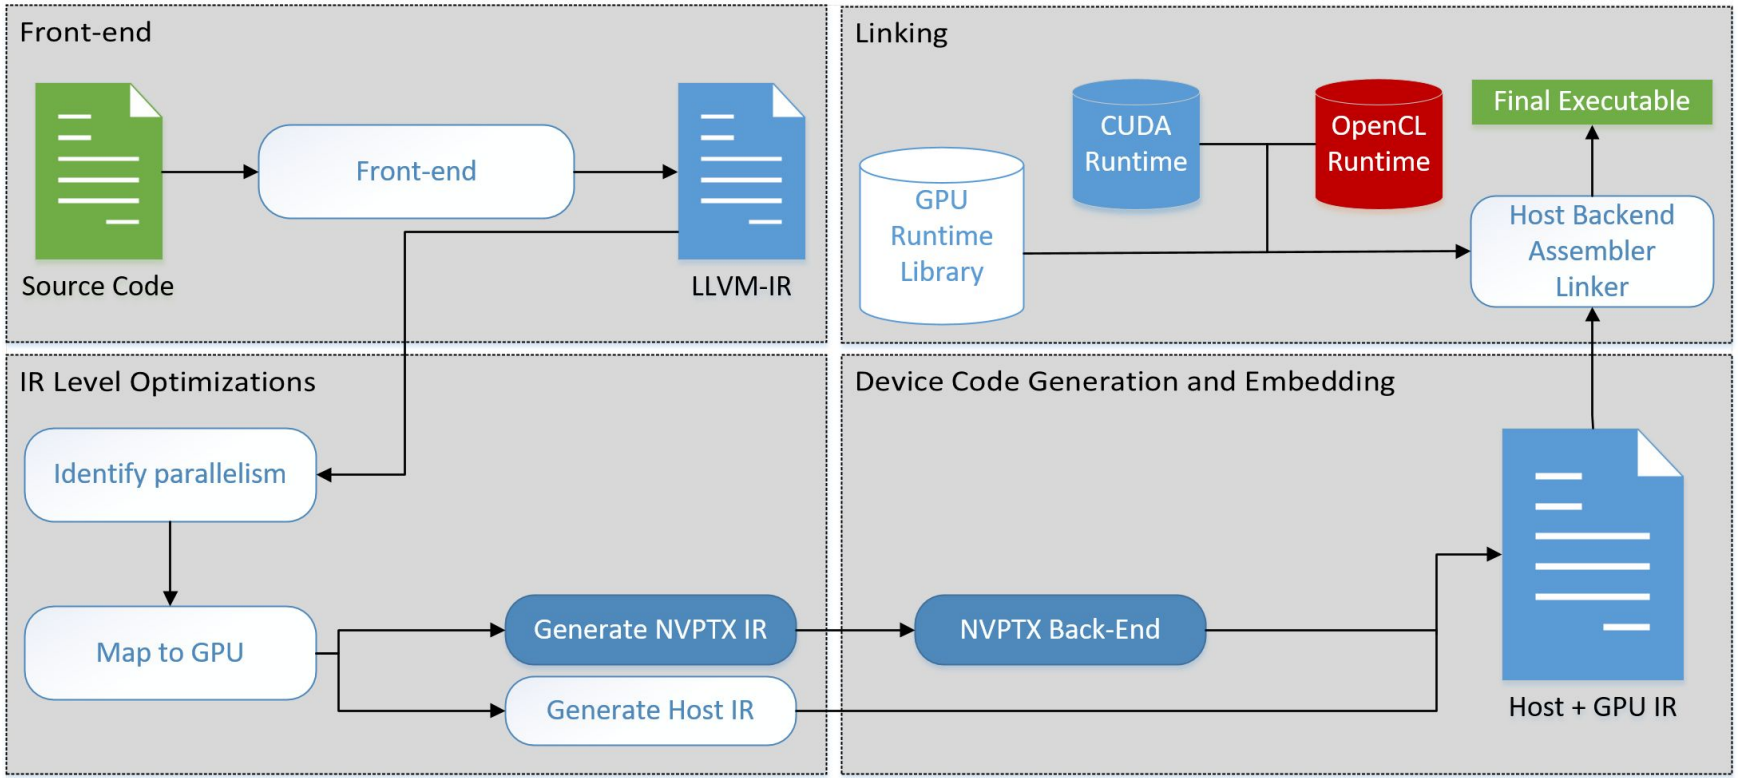

GPU Code Generation (OpenCL) Compared to -O3

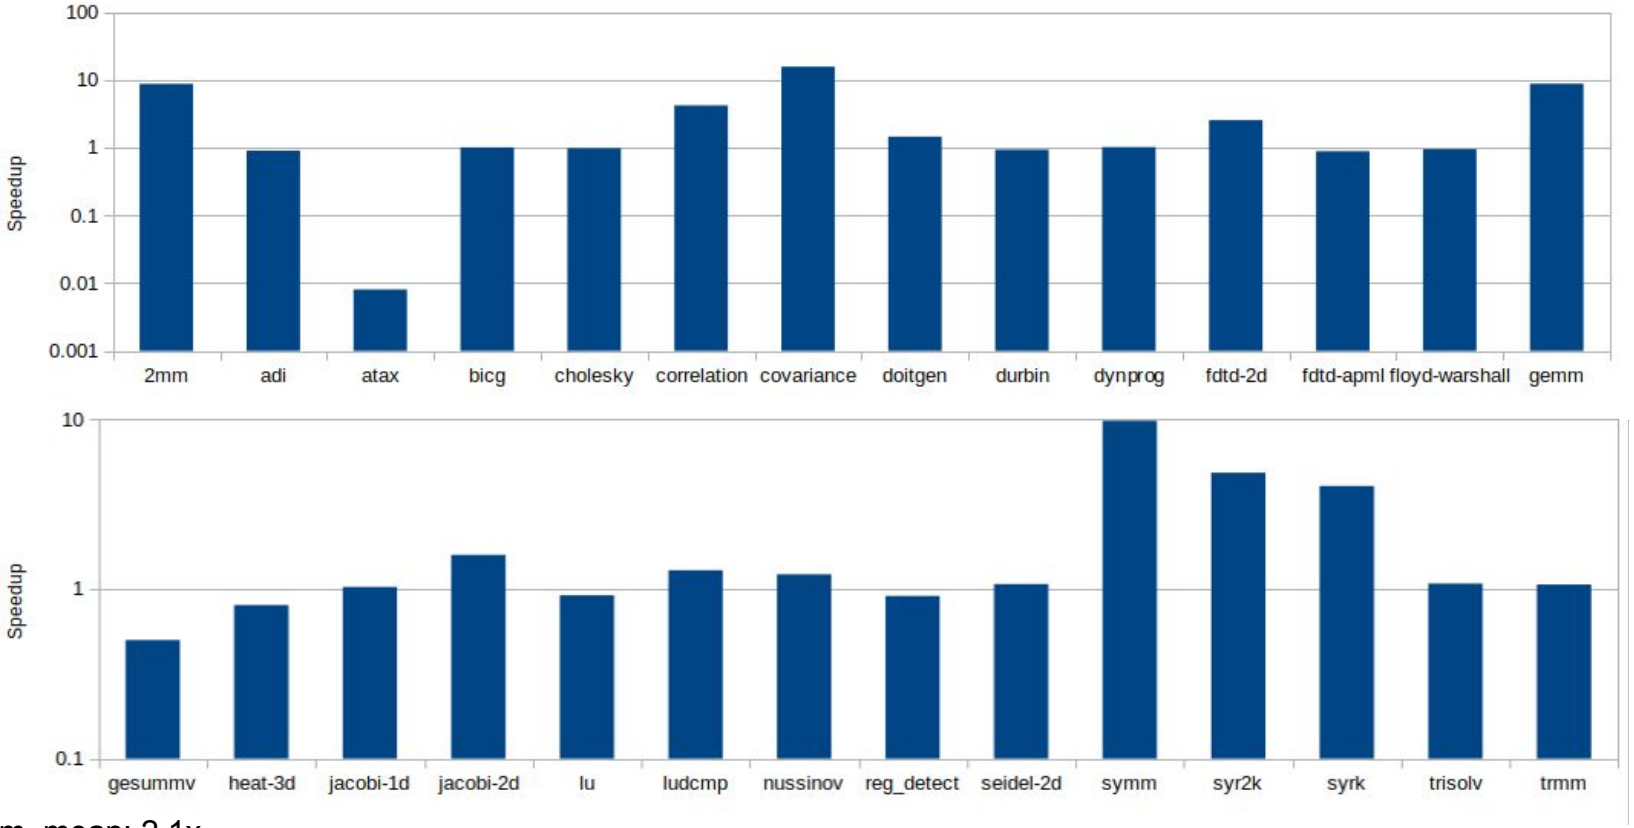

Geom. mean: 2.1x

How Does It Compare to CUDA?

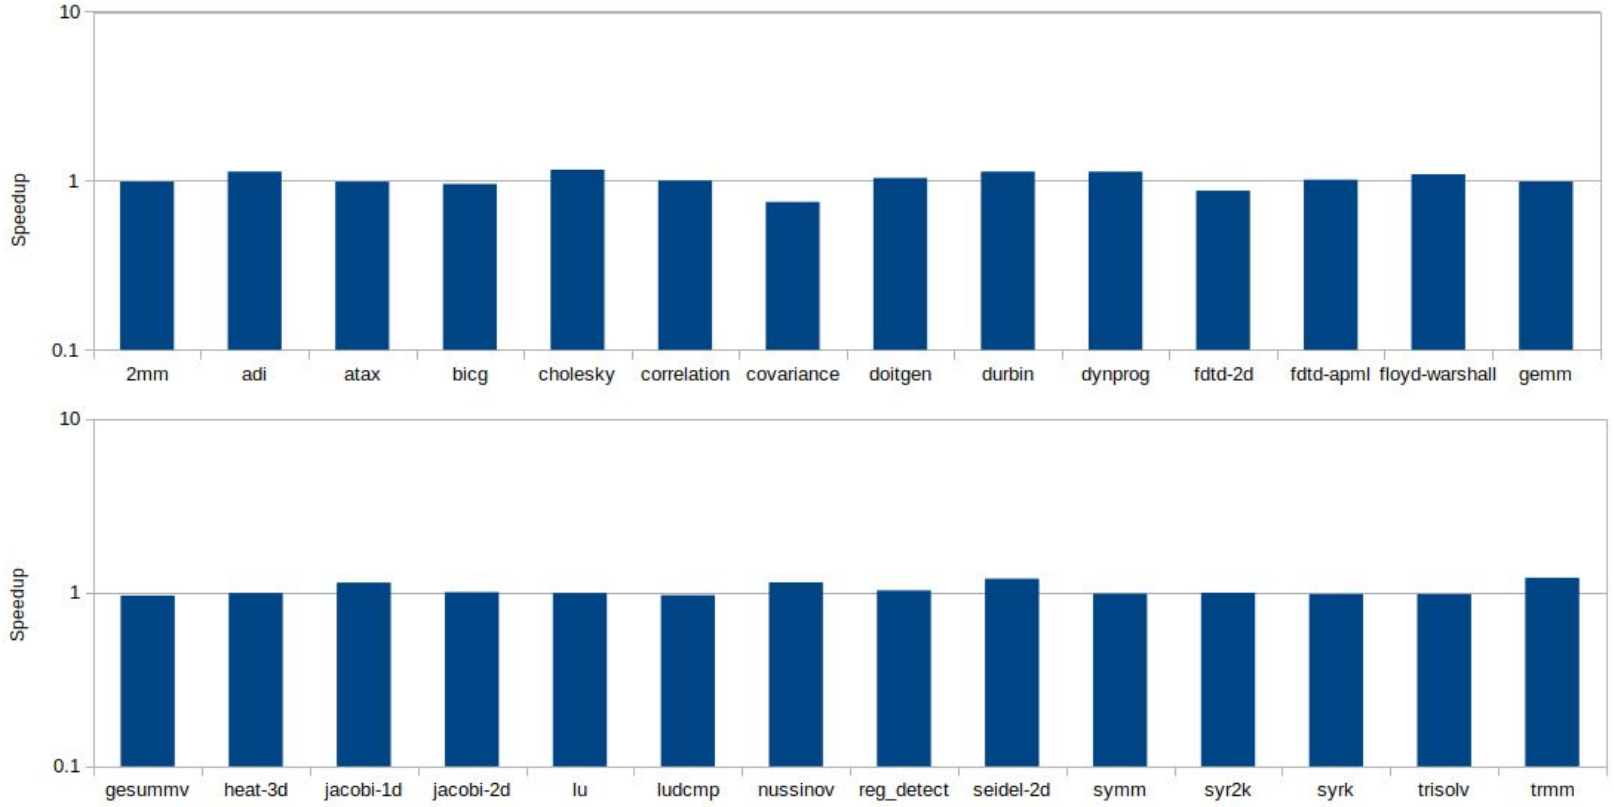

Why Intel?

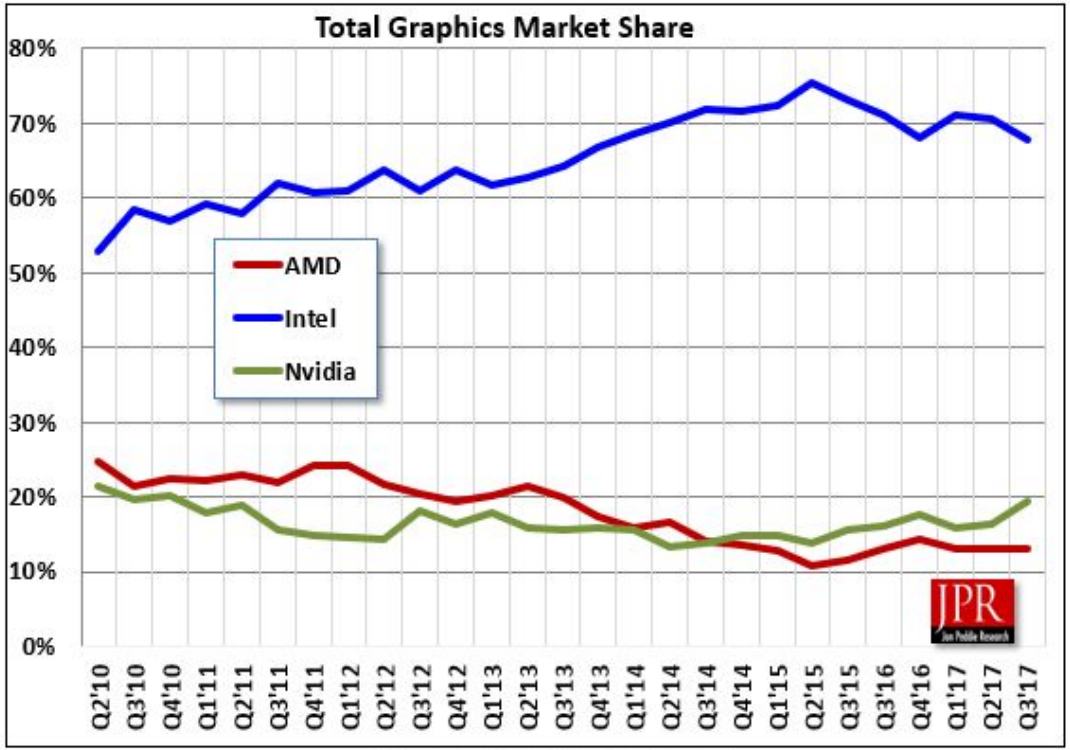

Source: Jon Peddie Research

Why Intel?

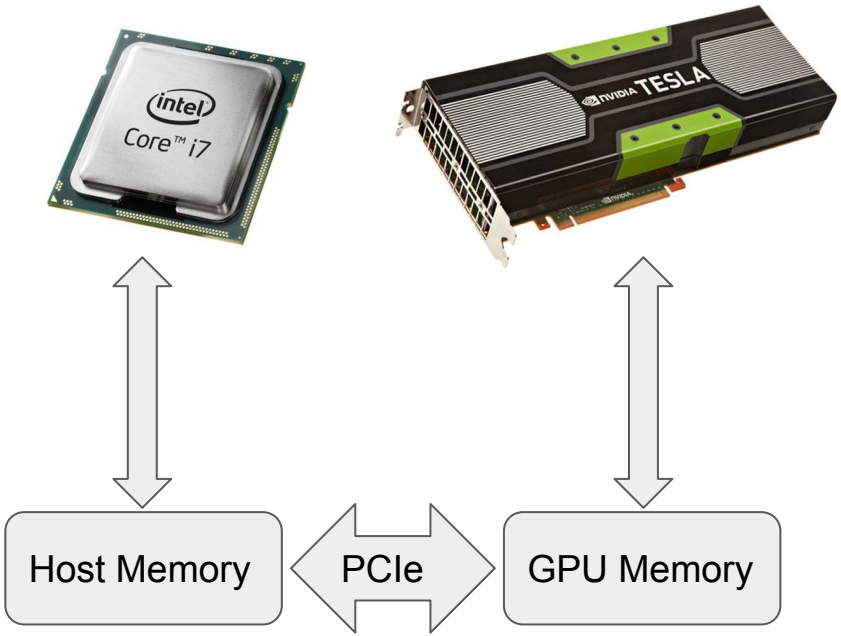

Why Intel?

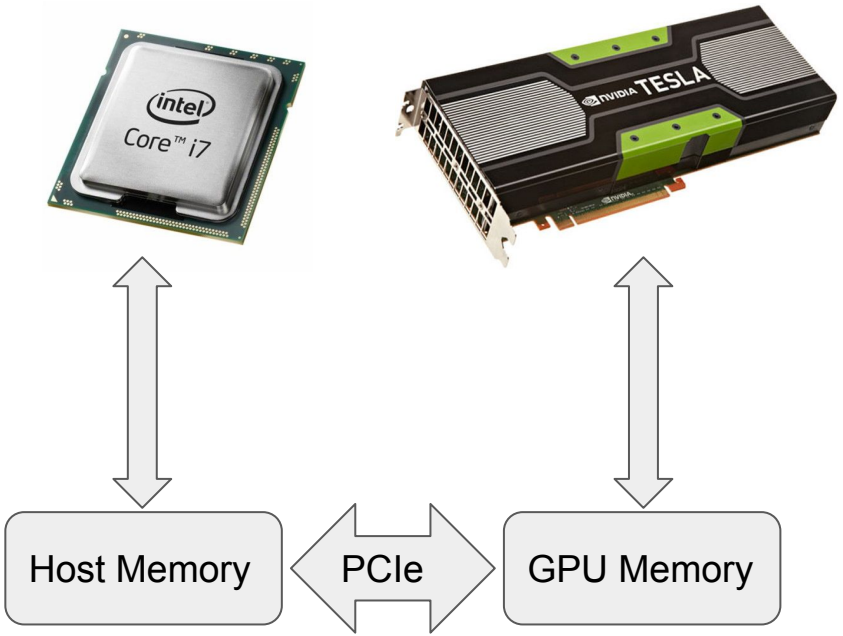

Unified Memory

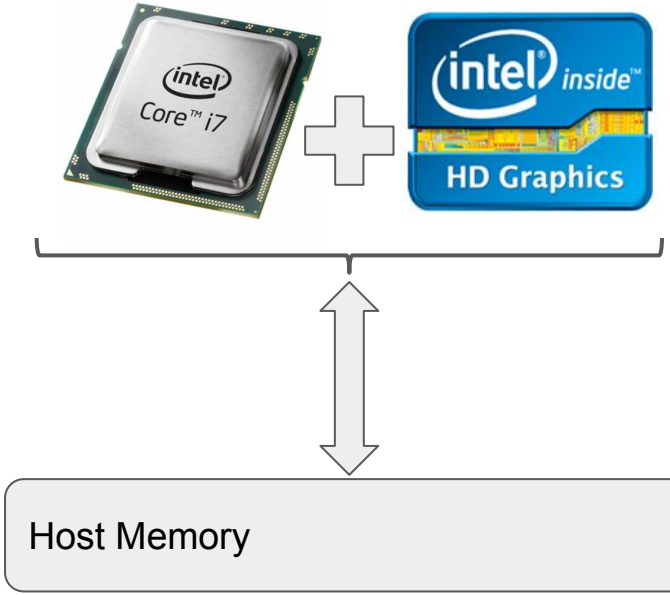

Generating Intel GPU Code

Generating Intel GPU Code

Running Intel GPU Code

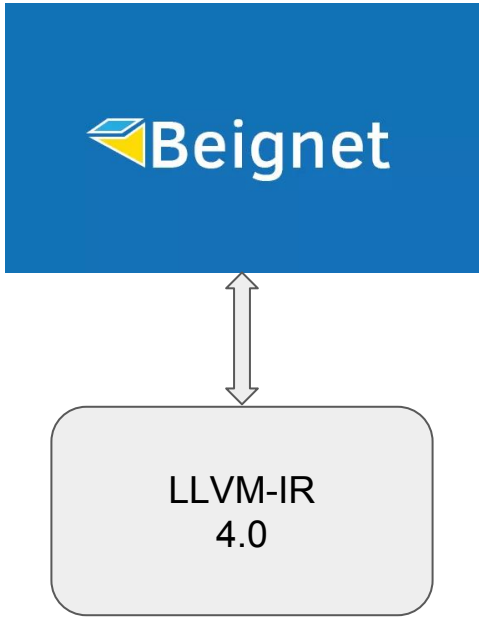

Running Intel GPU Code

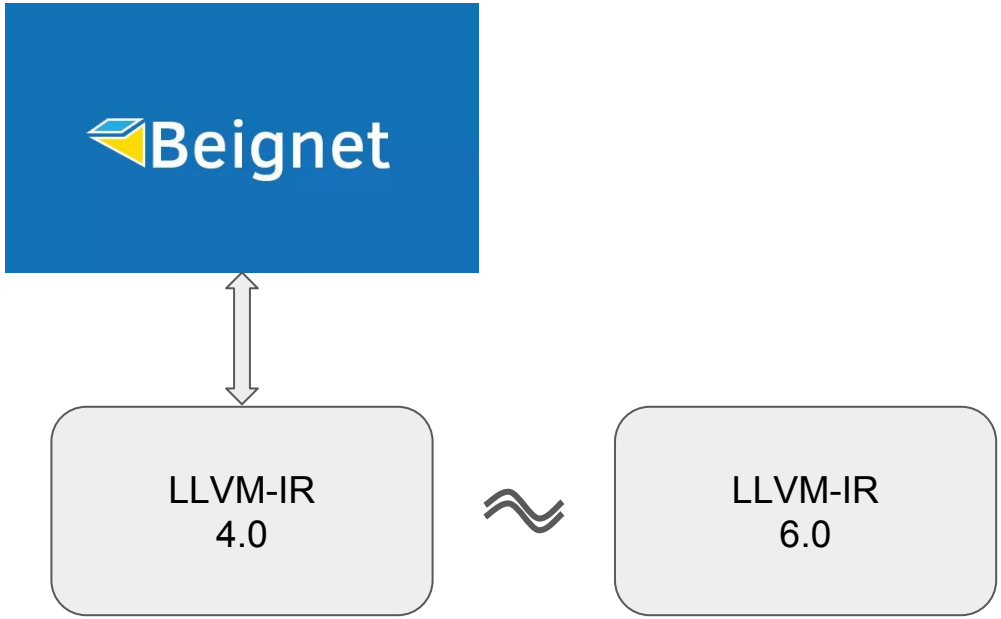

Intel Through SPIR

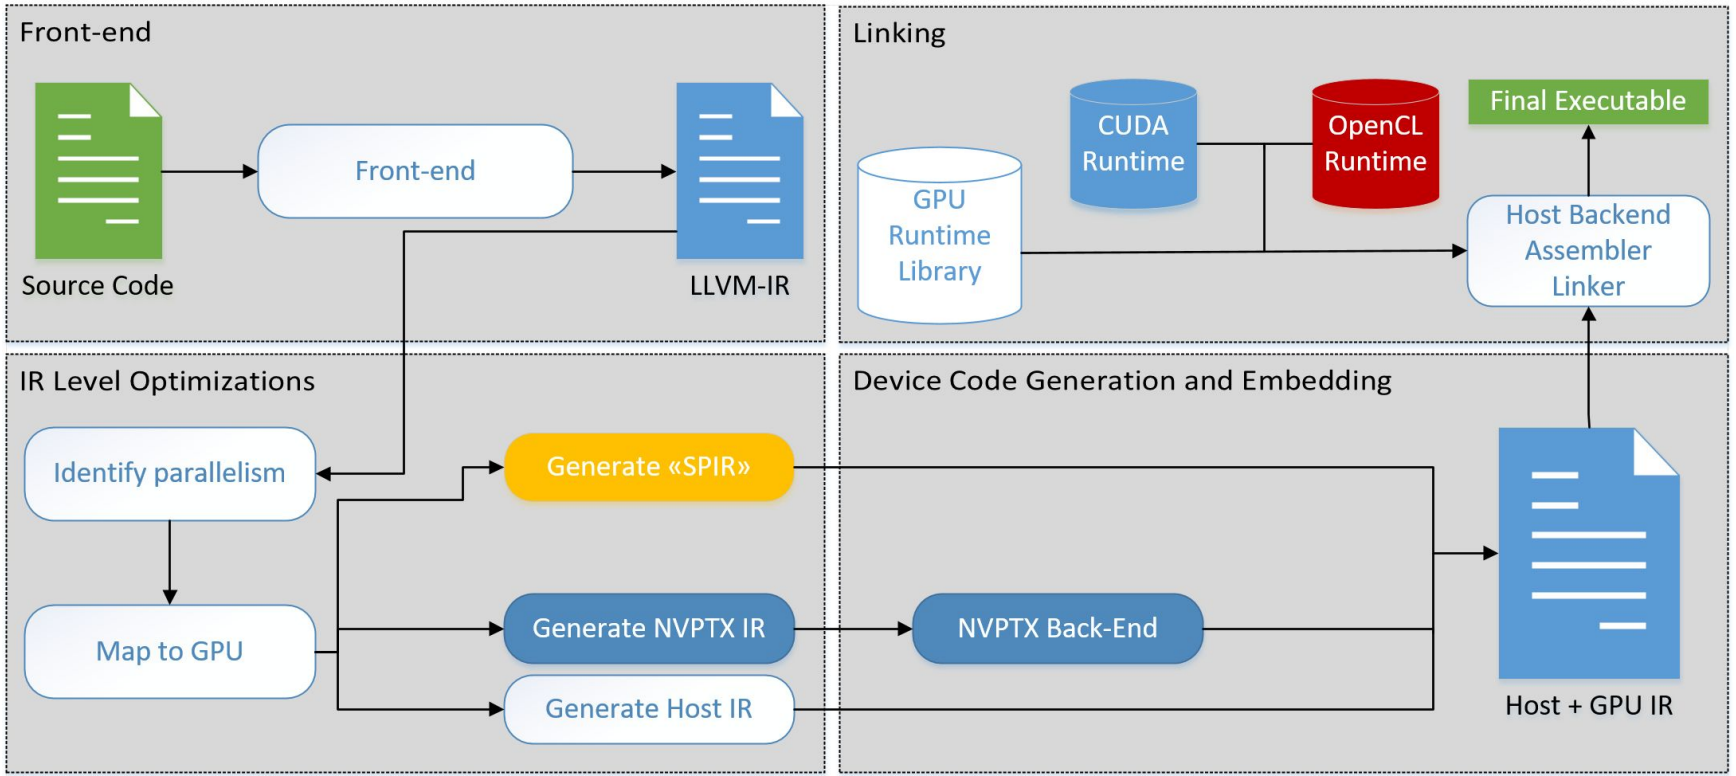

Percentage of Time Spent Copying Data

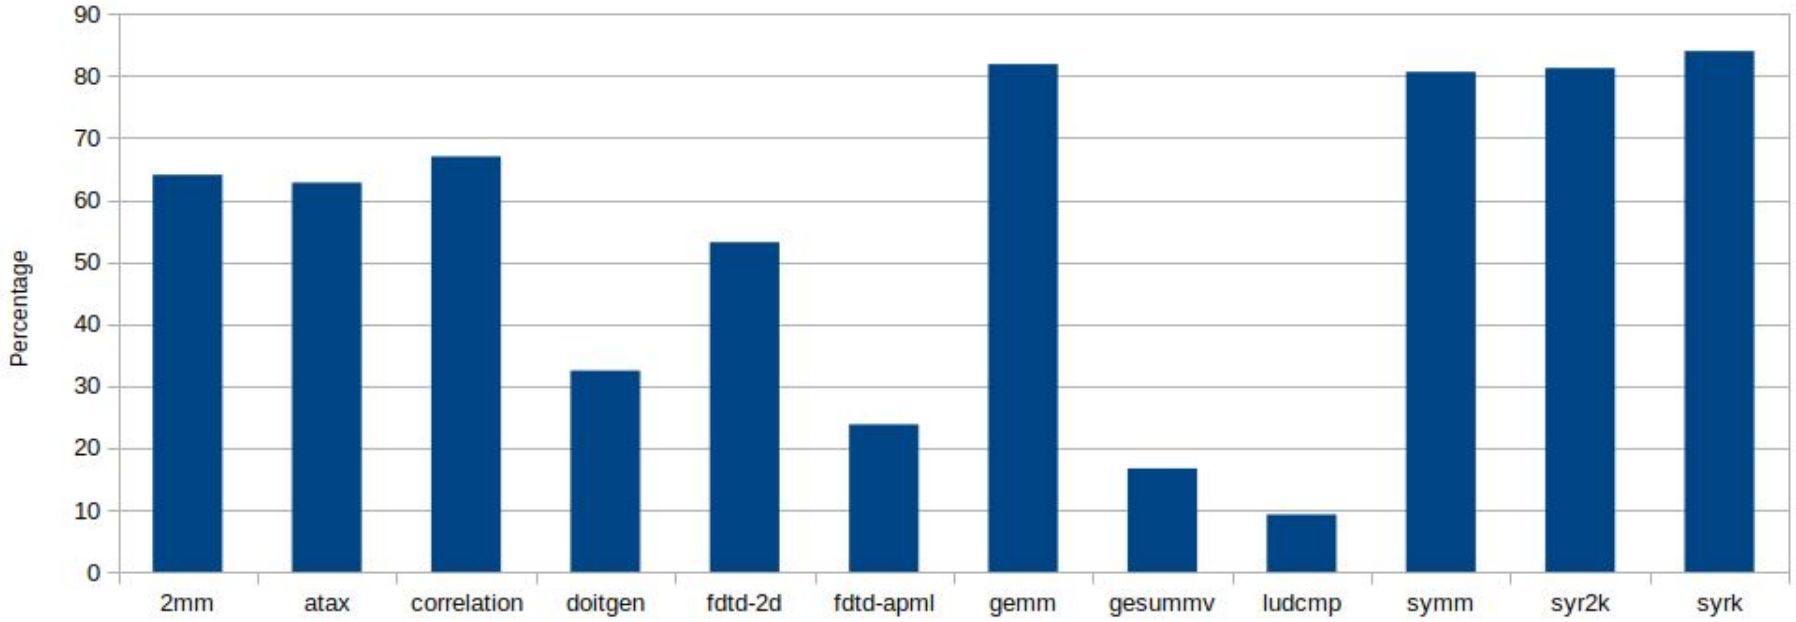

Geom. mean: 46%

This Is ‘Why Intel’!

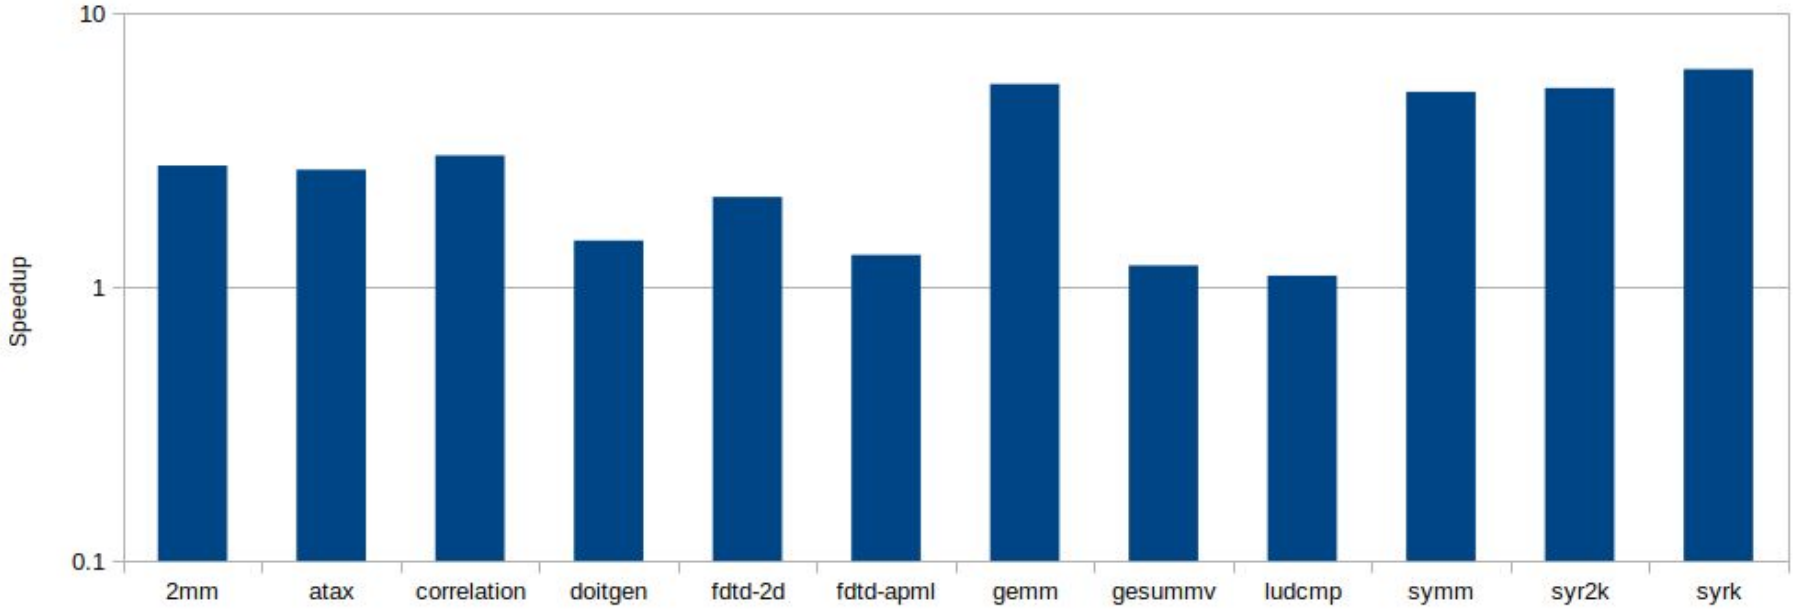

Geom. mean: 3.2x

AMD Code Generation

AMD Code Generation

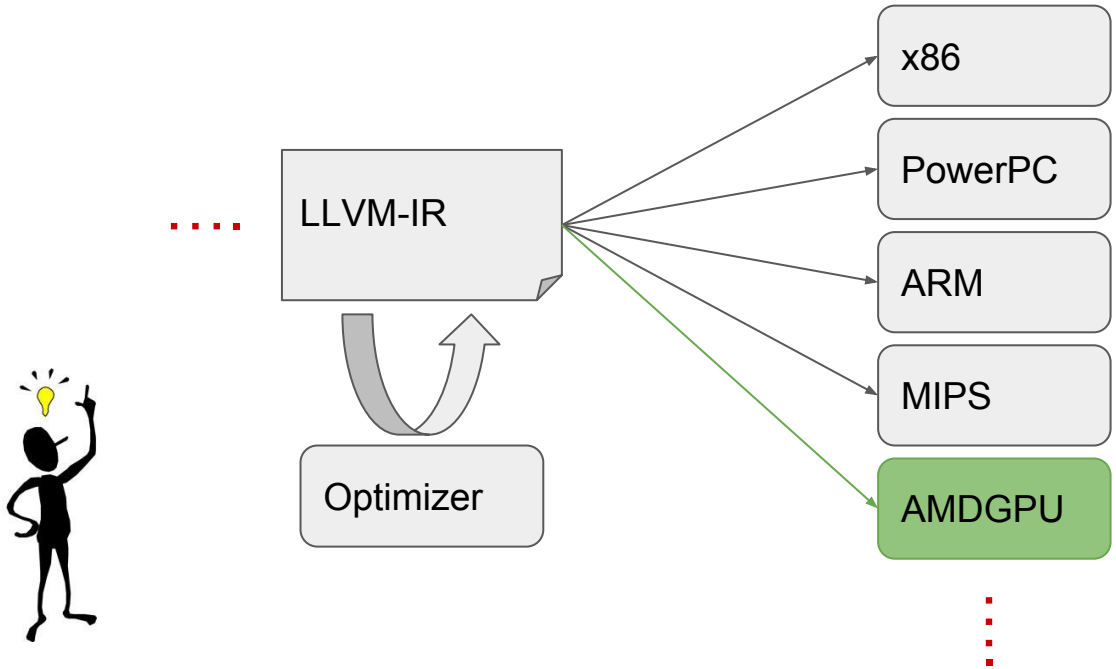

AMD Code Generation

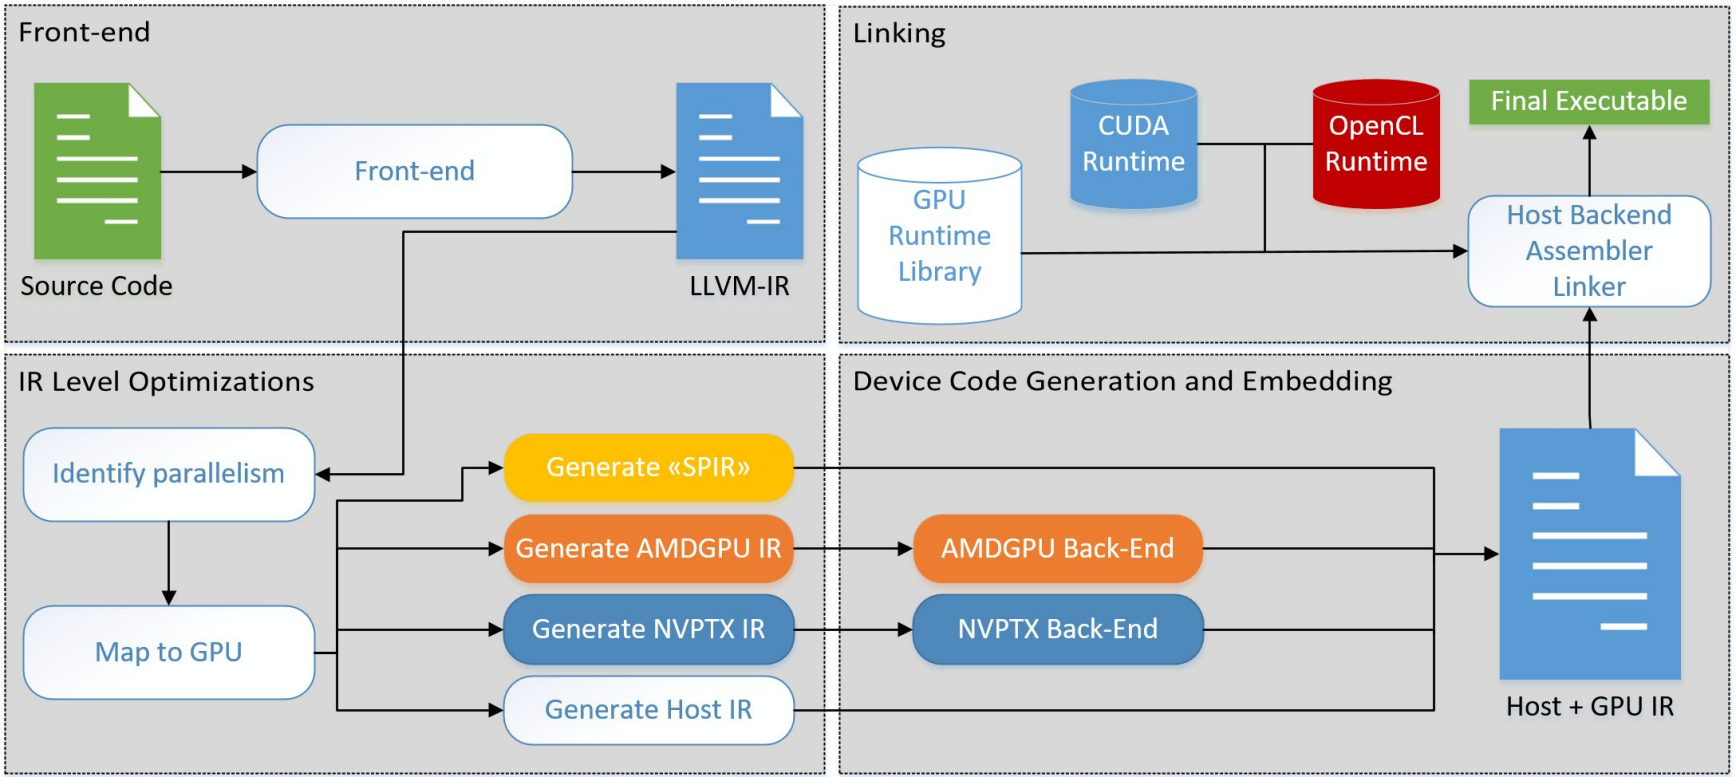

Running AMD ISA

Radeon Open Compute

Runtime Improvement by Platform over -O3

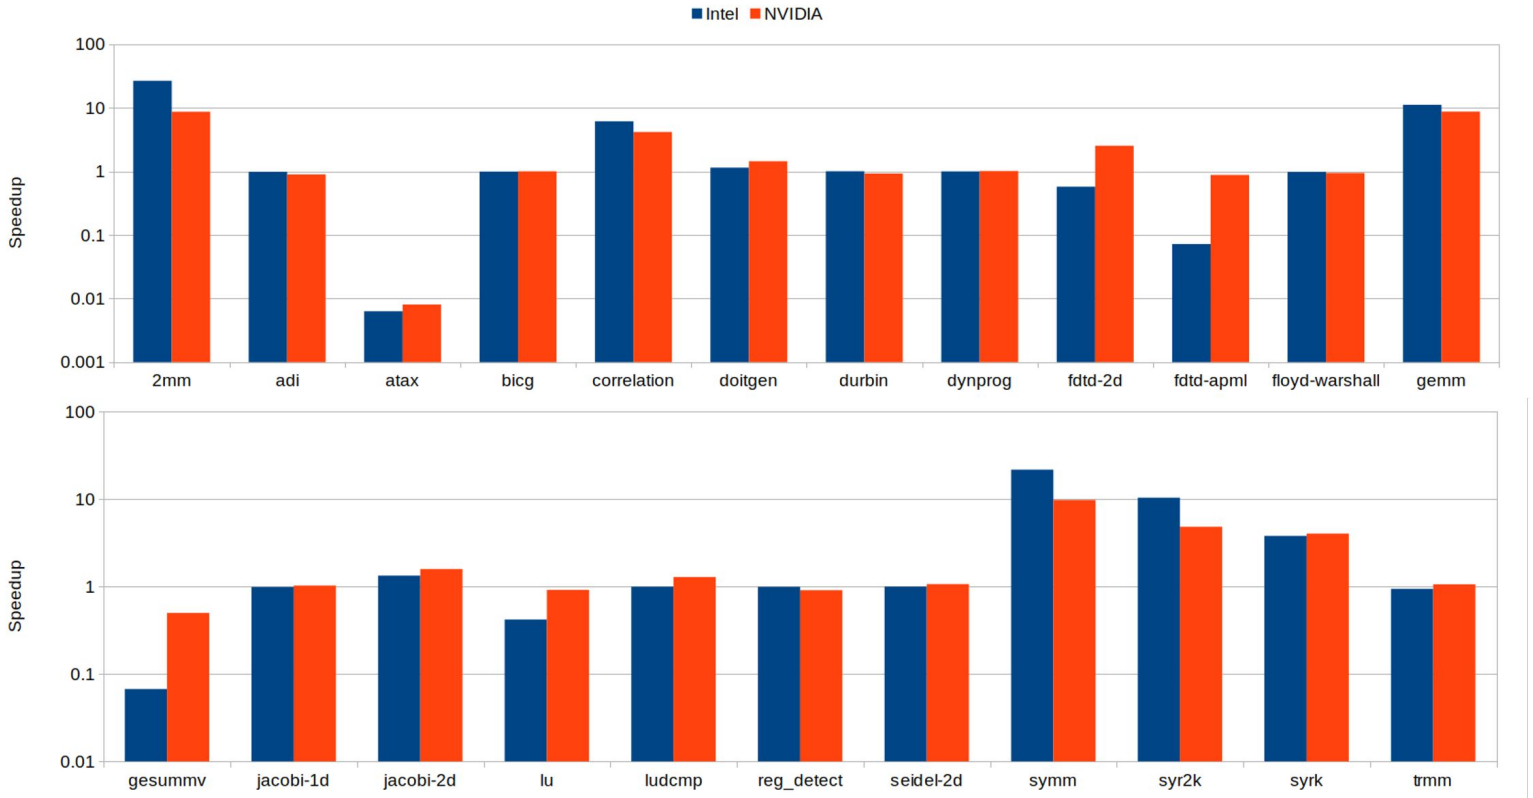

What can we do now?

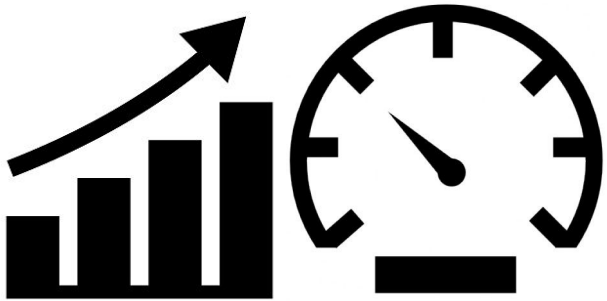

What can we do now?

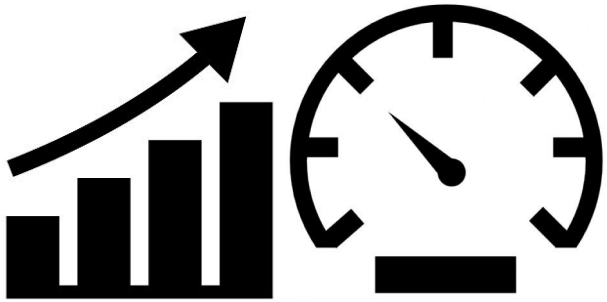

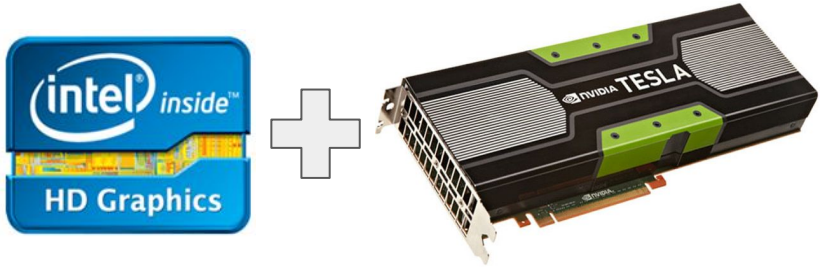

What can we do now?

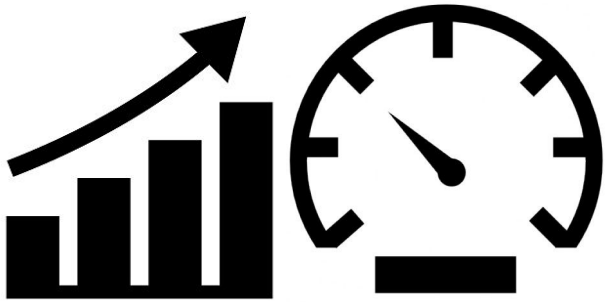

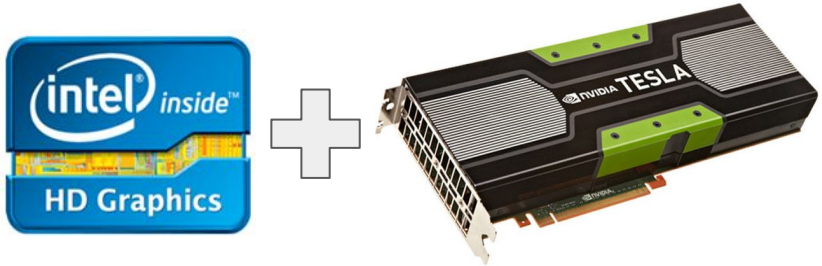

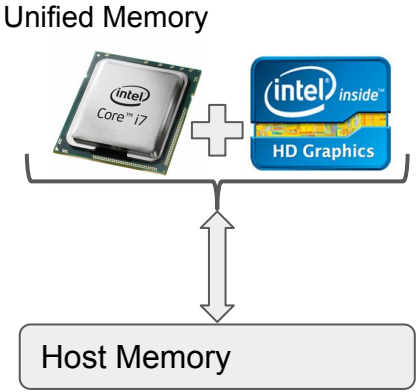

SPIR-V

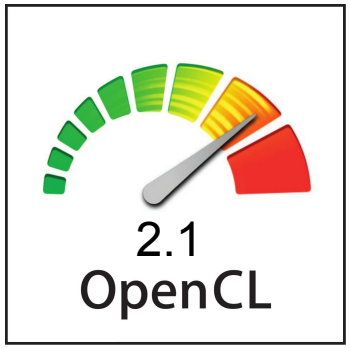

We’re One Step Closer