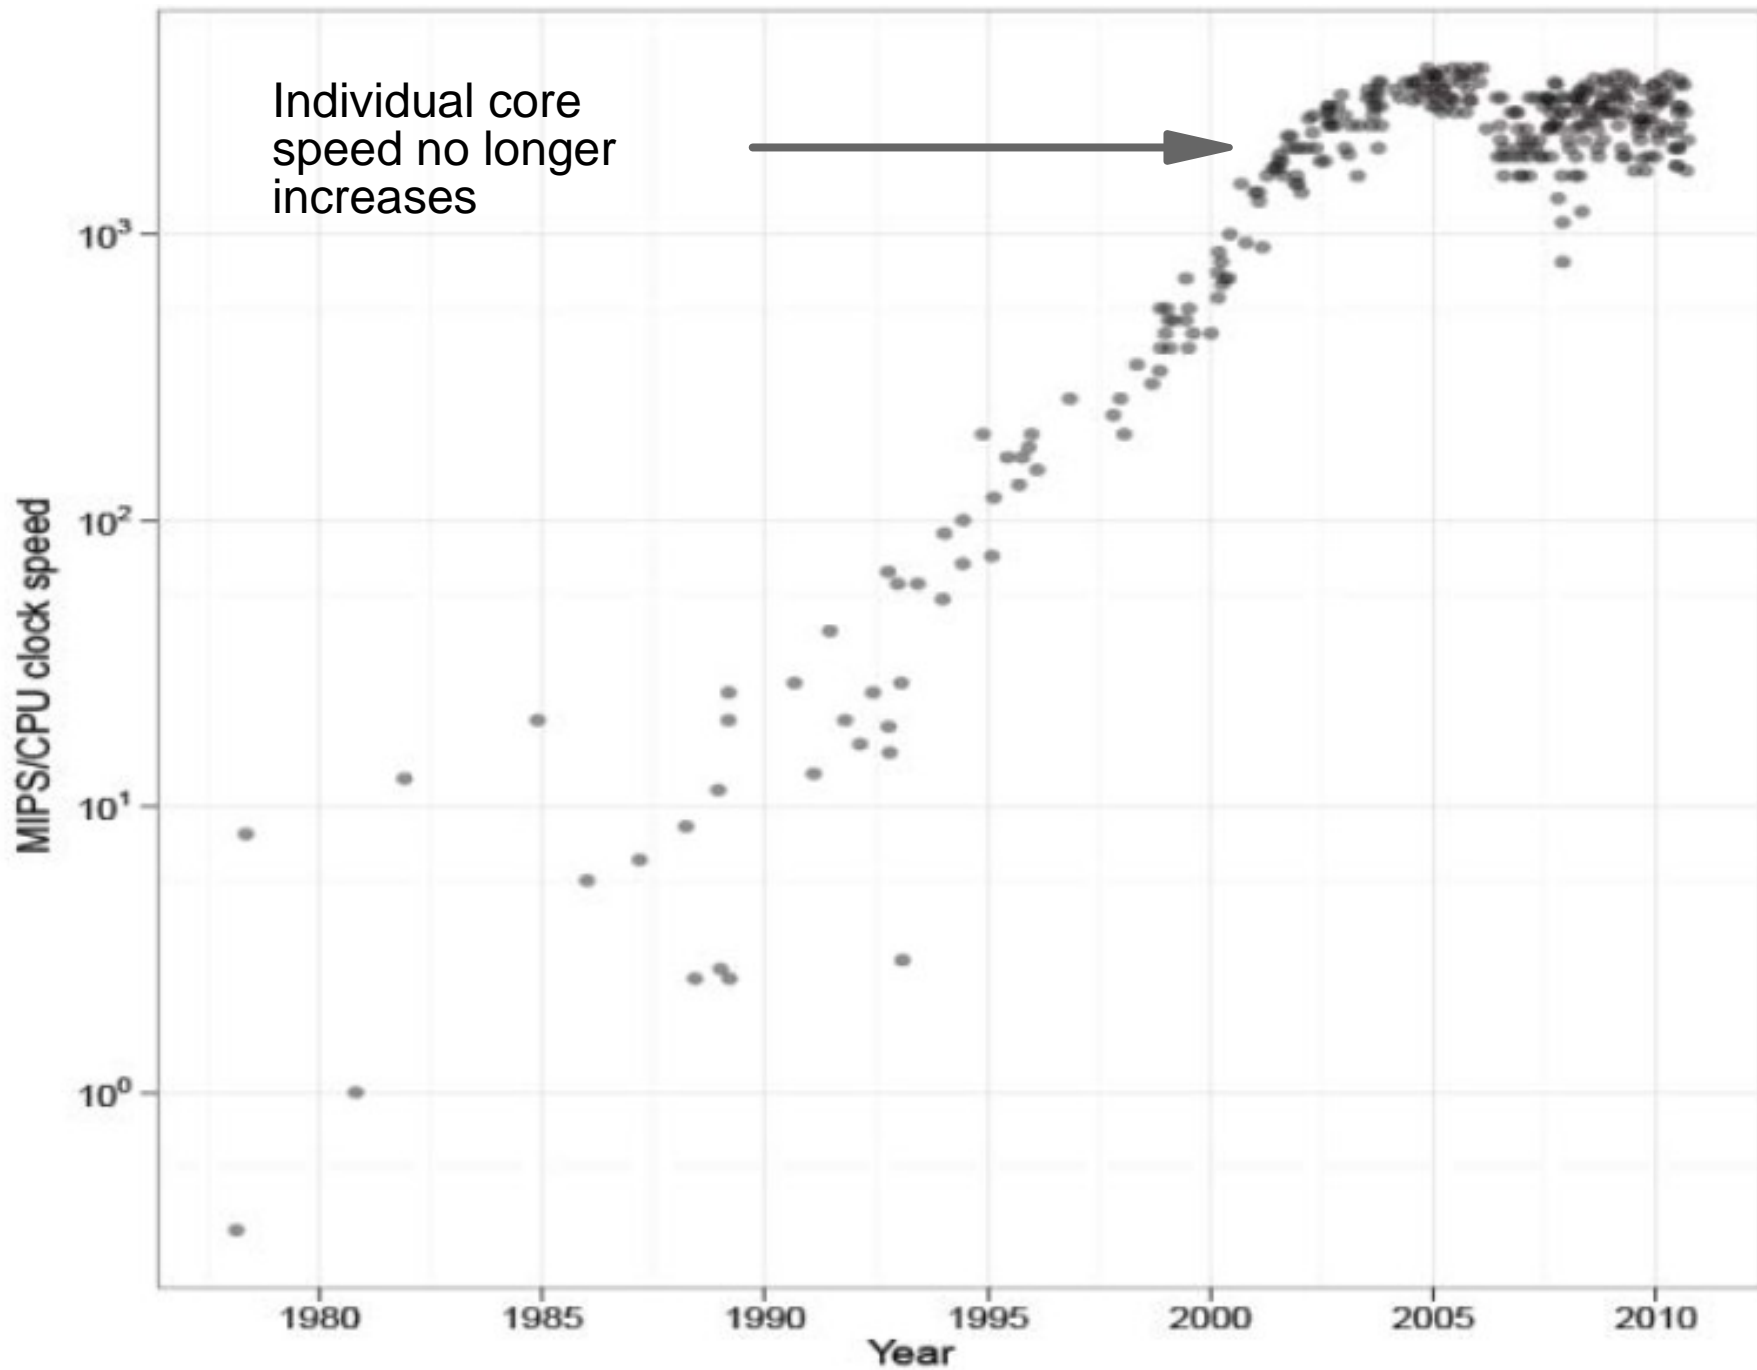


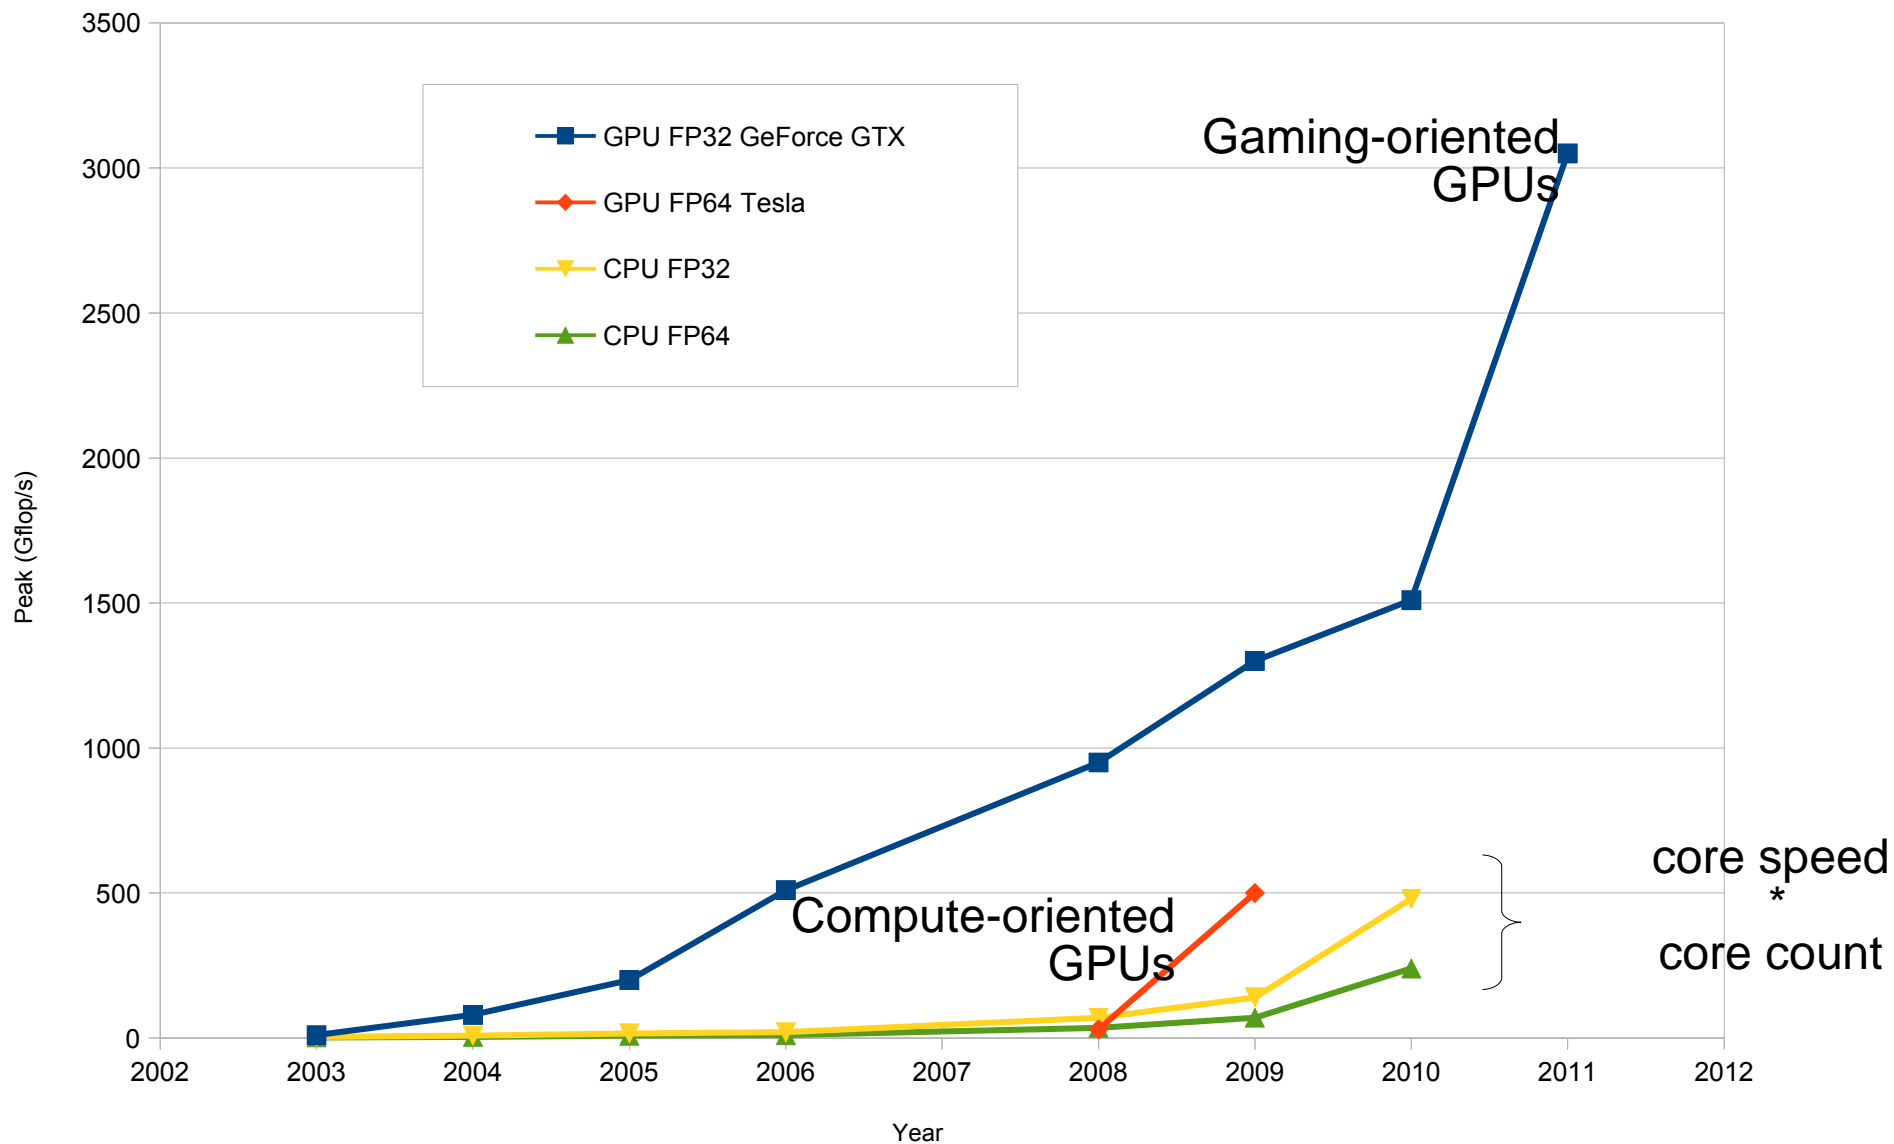
CUDA Introduction

Piotr Luszczek

Per-Core CPU Performance



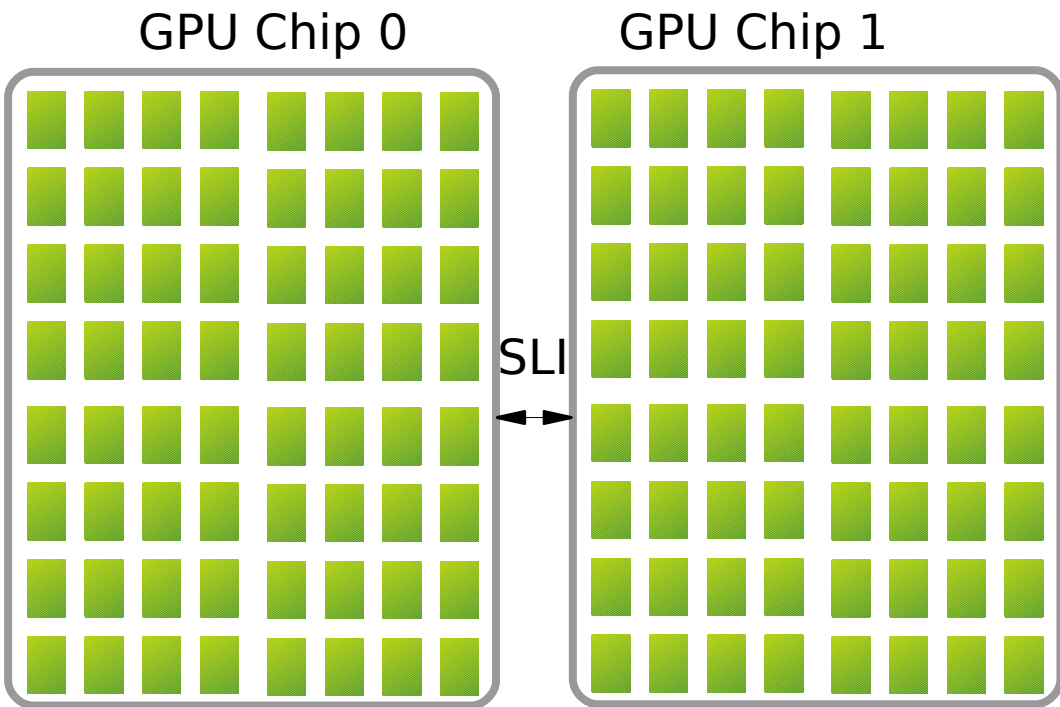
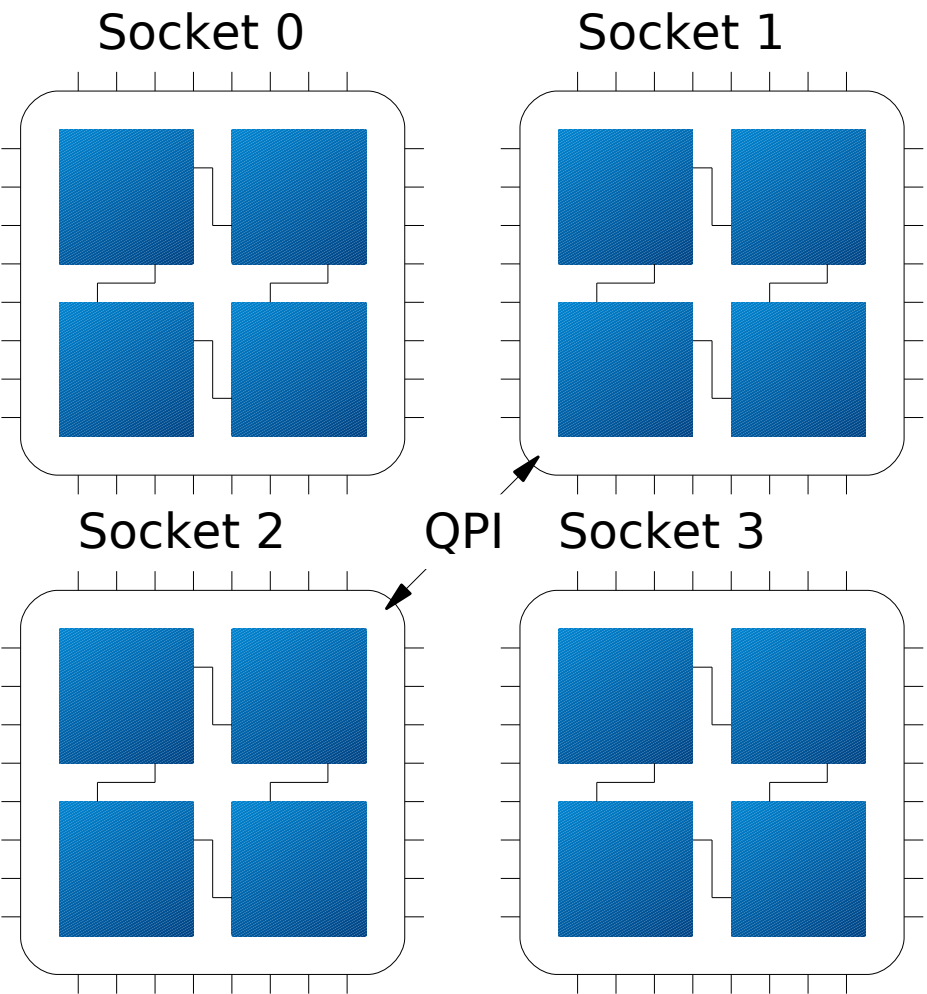
GPU vs. CPU Performance over Years



GPU and GPGPU: The Origin Story

- Programmable graphics pipeline
 - GLSL (shader language)
- Interpolation vs. dynamic range
 - Colors in graphics look better in floating-point
- Early attempts at programming
 - Cg, Brook, ...
- Modern standards or de facto standards
 - CUDA (currently 8, 9alpha)
 - Compute Unified Device Architecture
 - OpenCL (currently 2)
- High-level languages
 - OpenMP 4.5 and offload directives
 - OpenACC (like OpenMP but GPU oriented)

Hardware: CPU vs. GPU



Main Memory RAM: DDR3 or DDR4
Size: ~500 GiB
Bus width: 64 bits
Speed: ~100 GB/s

PCle express
15 GB/s

GPU Memory RAM: GDDR5
Size: ~10 GiB
Bus width: 384 bits
Speed: ~200 GB/s

Software: CPU + GPU

