

Accelerated graphics hardware.

Accelerated graphics hardware

- Compute Unified Device Architecture (CUDA)
 - API developed and maintained by Nvidia.
- Open Computing Language (OpenCL)
 - API originally developed Apple Inc.
 - Maintained by Khronos Group (including Nvidia).
- CUDA and OpenCL:
 - use Just In Time (JIT) run-time compilation.
 - use a C like programming language.

OpenCL: Open Computing Language

- Targeted compute hardware can be:
 - CPU (central processing unit)
 - GPU (graphics processor unit)
 - DSP (digital signal processor)
 - FPGA (field programmable gate array)

CUDA: Compute Unified Device Architecture

- Advantages:
 - Excellent parallel computation performance
 - Generally well supported in High Performance Computing applications
- Disadvantages:
 - Proprietary: Nvidia lock-in.
 - Nvidia hardware can be expensive.

OpenCL: Advantages / Disadvantages

- Advantages:
 - Excellent parallel computation performance
 - Comparable performance to CUDA
 - Hardware (possibly cheap) and ubiquitous
- Disadvantages:
 - Well supported in high performance computing applications?

GPU versus CPU

- GPU:
 - Excellent parallel computation performance
 - Compute units limited to local memory on GPU
 - GPU memory extremely fast!
- CPU:
 - Few computing cores
 - Direct access to Random Access Memory (RAM).
 - CPU main memory slow (compared to local GPU memory access).