

Performant programming for GPUs

Thursday, 13 June 2024 09:45 (1 hour)

Programming for Heterogeneous Architectures - lecture 3

- Data locality, coalesced memory accesses, tiled data processing
- GPU streams, pipelined memory transfers
- Under the hood: branchless, warps, masked execution
- Debugging and profiling a GPU application

Summary

Presenter: CAMPORA, Daniel (NVIDIA)