

Live Lecture Chat Window Wednesday, May 4, 2022

14:31:50 Can you use CUDA on non-Nvidia GPUs, or is CUDA just for Nvidia GPUs?

CUDA is an Nvidia-only proprietary product. However, OpenCL runs everywhere.

14:32:34 Can we use the concepts we learn in CUDA in other GPUs?

Absolutely! The concepts will just be expressed differently.

15:06:28 Can we use CUDA for cluster computers? Like can we use CUDA to get a speed boost by combing multiple GPUs/computers?

Yes, modern clusters and supercomputers have clusters of CPUs, each of which has one or more GPUs (plus multicore, plus SIMD on top of that)

15:41:59 So you recommend we develop the software on rabbit and then go to the DGX on the last test run.

Yes. Your performance numbers will be ever-so-much-cooler on the DGX.

15:44:40 I'm using the WaitBarrier() on Visual Studio, and wanted to ask where InitBarrier() should be. Is it placed before omp_set_num_threads()?

Right before, or right after.

15:50:43 Will SSE SIMD code from proj 4 run on an AMD CPU?

Not sure – depends on the compiler. Try it and let me know! 😊