

Live Lecture Chat Window

May 22, 2024

14:00:33 Just to confirm, at the bottom of the Project 5 description, when it says "benchmark combinations", is it just talking about running the combinations of NumTrials and BlockSize?

That's correct.

14:04:43 It's okay for the submission if we just run that bash script on rabbit correct? We don't have to include a rabbit vs dgx comparison?

Sort of like King Kong vs Godzilla? No, just taking numbers from rabbit alone is fine.

14:09:42 Should we report M and B by array size, or are they expected to be more or less the same thing regardless of array size?

In Project #6, just report the M and B values for the run with the largest DATASIZE.

14:11:55 I have a GPU, and I want to figure out how many Processing Units it has per Compute Unit, NVidia GPU. That printf stuff only tells me Compute Units

Newer Nvidia cards seem to all have 128. Older ones seem to have 32.

14:16:38 So when does OpenCL execute out of order? When clEnqueueNDRangeKernel() doesn't have the event in its arguments?

That is my understanding.

15:04:54 When doing FMA, is there a loss in accuracy for the result?

There is some loss in accuracy for *every* floating-point operation. But my experiments showed that the FMA floating-point produced a *different* loss of accuracy than the "intermediate value" approach. Just goes to confirm that one should never test for FP equality with ==. Instead, always test $fabs(f1-f2)$ against a tolerance.

15:36:35 Is anyone else having problems with Nvidia GeForce Experience? Specifically error 0x0003 right out of the box (when I try to open it on my desktop)?

It's not the GeForce Experience you want – it's the CUDA Toolkit:

https://developer.nvidia.com/cuda-downloads?target_os=Windows&target_arch=x86_64&target_version=11&target_type=exe_local