Live Lecture Chat Window

May 8, 2024

14:05:15 I still have capstone to take after this term (that's my only class left). If I plan to take that in the summer, I definitely cannot go through graduation this spring, correct?

You can definitely go through the EECS Graduation Celebration on the Friday of Finals Week. Ours is very informal, and our only goal is to celebrate you and your accomplishments. I thought you could also participate in the University one on Saturday, you just won't get a diploma. Check with the Registrar people to be sure.

14:12:06 Why no pivot tables for Project 4? Is it simply 4 different outputs (performance vs array size) for the various operations, all plotted together instead of using a pivot table to produce a table to plot?

In P4, there is only one independent variable – the array size. Pivot Tables are at their best when there are two or more independent variables. So, if you have an output that has a column of array size, a column of Mul speedups, and a column of MulAdd speedups, you can just sweep over that whole thing and insert a graph.

14:15:49 Was the assembly language used x64?

I think it will work for both x86 and x64.

14:16:42 I saw at&t syntax -- is there a way to use intel syntax instead?

I'm sure there is, I just don't know what it is.

14:43:51 If we want to run CUDA locally, what should we install?

The CUDA toolkit. The link is in the CUDA notes.

14:44:43 If it [module load slurm] adds nvcc to the path, do we still need to set the full path in the scripts?

module load slurm only sets the path to the slurm commands, like sbatch. It doesn't set any of the CUDA-related paths.

14:55:19 I'm going to have to rewatch this lecture like 5 more times ;)

I understand. First time I saw GPU programming, it mystified me.

14:56:55 Do all programs that utilize a GPU also use the CPU? Or is this just a CUDA thing?

Yes, all programs that use the GPU also use the CPU. You cannot directly get to the GPU (*ssh nvidiacard.engr.oregonstate.edu* won't get you anywhere. So, you need some place to get the program running, read data, send data to the GPU, invoke the kernel, get data back, print results, etc. That place is the CPU.

15:10:13 How do you send your other files to DGX?

You don't have to. The submit-* machines look just like flip. Your files are already sitting where you expect them to be.

15:12:26 What exactly constitutes a supercomputer?

Not sure there is a strict definition. The Encyclopedia Britannica says, "The term is commonly applied to the fastest high-performance systems available at any given time." Not very helpful.

15:21:58 Hypothetically, could we have the global random number arrays as --device-- variables (can't use _ here), and then use a --global-- function to fill then, avoiding passing the big arrays over to the GPU?

All random number algorithms I am aware of are sequential. That is, one random number acts as the seed to get the next one, to get the next one, and so on. There must have been attempts at generating random numbers in parallel, but I don't know of them.