

Live Lecture Chat Window

April 30, 2025

16:01:04 Thinking about the inverse of the ARM cards — what does the gpu portion of cpus with integrated graphics look like, that allow for computer graphics without the presence of a graphics card?

The integrated graphics chips make up for lack of a supporting graphics card by sharing resources, such as memory, with the CPU.

16:03:42 I was wondering if you had a highly parallel program that used a very large data set would the GPU be able to stream the data out before loading more or would it run out of memory? (Are you limited by VRAM)

The path to/from GPU memory is very fast, so some amount of streaming is possible. Not as good as having it all in GPU memory. This is one of the reasons that many graphics cards have so much memory (e.g., the Nvidia A6000s with 48 GB).

16:06:39 I was hoping to do project 3 in C — can I show the structs, the strings are the ones I am not sure I am changing correctly.

Should work fine except for the `std::string` construct. That will need to be a `char *`.

16:08:55 Do GPUs also read/write from VRAM in 64 byte cache lines?

Looks like the newer Nvidia GPUs use 128 bytes per cache line. Looks like AMD is transitioning from 64 to 128. Looks like Intel GPUs use 64.

16:15:08 For Project 3 (K-means), could/should we use parallel programming with OpenMP for the extra credit too? Specifically, over capitals? Or is the amount of work per capital too small to benefit from it?

I remember thinking about this for my version of the project. I think I decided that any performance gains wouldn't be worth it. Try it both ways and let me know what you find out!

16:32:24 From Wong, Becky to Everyone:

For project 4, is it expected to see speedup increase from small array sizes to moderate array sizes, then plateau as the array size gets larger (6M - 8M)

Yes. And then, especially if you do the Extra Credit, the performance will drop due to congestion on the path from the CPU to memory.