

# Live Lecture Chat Window

## May 3, 2023

**15:20:00 Why were the per node computations always in the Fragment shader vs Vertex (or other) shaders?**

Remember how our temperature equation had to look at the values around it? Through textures, pixels can look at other pixels. But, vertices cannot look at other vertices. Also, pixels are already arranged in a nice grid, just like the diagram for the 2D differential equation. Vertices aren't.

**15:39:46 Do GPU's still conform to Von Neumann architecture?**

Pretty much. The pieces are all there. Things like the CUDA Cores are different.

**15:41:00 Do CPUs with "Integrated Graphics" include a mini GPU or just incorporate some of the GPU architecture to their chip?**

They have both a CPU and a GPU in the same chip.

**15:48:44 Do CPU and GPU always have separate memory?**

For the most part, yes, but it wouldn't surprise me if some of the integrated chips do some sharing of memory.

**15:52:54 I think Apple's ARM chip has shared memory between CPU and GPU**

**16:01:51 Quick question re: Moore's Law and GPU chips - is this why we're seeing more specialization for stuff like tensor cores?**

Yes, the ability to fit more transistors onto a chip has allowed designers to be more creative.

**16:01:52 Do GPU's have globals, heap, stack?**

No stack, but lots of registers. The globals and "heap" live in the on-board memory. We will talk about that when we talk about CUDA next week.

**16:05:22 Are GPUs popular for crypto-mining because of the CUDA core count?**

That's what I would assume. Lots of data parallelism.

**16:06:00 So the GPU can't function without a CPU?**

That is my assumption, that GPUs have a "runtime" but not a full operating system. You need to first login to a system with an OS. However, when Nvidia bundles an ARM chip on the same board as a GPU chip, maybe it looks like you are logging in to the GPU?

**16:06:07 What is your take/opinion on Intel Arc?**

I'm hoping to get my hands on one this summer. It would be fun to run all our benchmark programs there.

**16:07:42 Why isn't there fused multiply add on the CPU yet?**

Not sure. FMA seems like obviously-useful to me. Maybe it is on CPUs and I just haven't heard about it.

**16:09:20 This might be slightly off topic, but I've been wondering how parallelism is used most often in web development.**

You will use OpenCL for Project #6. There is a JavaScript version of it called WebCL, meaning you can get GPU-parallelism in a web program. I've heard it used for things like data mining.

**16:39:05 If I want to experiment with parallel computing on my own hardware, are there any risks of overheating under normal circumstances, or is that pretty unlikely?**

Pretty unlikely I think. These systems are designed to protect themselves. For example, many systems measure chip temperature and crank up the fan or lower the clock speed if it gets too hot.