

Live Lecture Chat Window

Wednesday, May 18, 2022

14:34:06 Yeah, it [the performance boost when moving from rabbit to the DGX] is around 5x. It's hard to wrap my head around the numbers and even the example on mile stacks of paper done in an instant. But it's nice to know what is available.

A typical Montecarlo simulation code is more complicated than our simple snowball example. So what this really means is that you could make the code 10x more complicated and still get performance in the billions of trials/second.

14:38:27 Is this [Project #7A particle system] going to look similar to the one we did [CPU particle system in CS 491] in the fall?

Yes, in the same way that a Ferrari is similar to my Prius. 😊 Same effect – but way, way, way faster. Keeping all the data on the GPU and computing using GPU-parallelism is what makes the difference.

15:27:33 Why does the performance increase to busier you keep the GPU?

Think of it in terms of the yellow robot. If the robot can use its 5000 grippers to move 5000 grapefruit/second, but you decide to only use 1000 of the grippers, you will only get a performance of 1000 grapefruit/second. Leaving some of the capacity unused results in a lower performance number.

15:30:53 Could you use OpenGL with CUDA?

Yes, CUDA has its own set of OpenGL compatibility functions.

15:30:53 The difference between OpenGL and OpenCL is that OpenGL is used for graphics programming and OpenCL is used for the computing, right?

Correct! The G stands for Graphics, the C stands for Compute.

15:31:32 They are all just libraries for c (or maybe its C++)

Correct, but just for C. Both OpenGL and OpenCL decided not to exclude C-based legacy code by requiring C++ in the API.