Live Lecture Chat Window
April 26, 2023

15:36:34 My brother has a PhD in Huntsville, related to rocket science. He promoted parallel programming's usefulness, unaware I was taking it. Treated it like the Rosetta Stone of engineering. It's a private business trying to predict bubble fluid movement and predict catalyst movement in reactor beds for producing Polymers.

Wow, great to hear! I hope he is impressed that you are in this class.
Does he have any information on how he is doing his parallelism? Any speedup graphs we could see? That's the type of thing I would love to add to the notes!

15:42:19 Are global variables saved on the cache or in memory?

They are stored in memory, but when they are used by the program they are brought into the cache.

16:09:42 What is xmm0, xmm1 again?

They are the SIMD SSE registers. Each is 128 bits wide and thus can hold 4 ints or 4 floats or 2 doubles.

16:10:16 Are there any recommended books or resources on writing assembly in C/C++?

I don't need assembly code information very often, but when I do, my go-to reference is:
Jo Van Hoey, *Beginning x64 Assembly Programming*, Apress, 2019. If you are on-campus and want to see it, let me know.

16:19:29 Could you go over "Combine Multicore and SIMD" again?

Sure! Each core has its own dedicated SSE SIMD unit. That means that each core can do whatever it wants with its SIMD unit without any interference from the other cores. When you multithread a for-loop, each core gets only a fraction of the total for-loop passes. And, that fraction can be made to run 4x faster using SIMD. This means that the OpenMP for-loop handling combined with SIMD will give you even more speedup than just either one alone.

16:25:37 In summary, does SIMD make use of multi threads on a core and multi-registers on the thread?

Remember that threads are a software thing (that you can ask for) and cores are an execution-hardware thing (that you can’t ask for). You ask for threads, knowing that the OS will run those threads on separate cores. So, your for-loop passes can be divided up among the cores (because you asked for threads), and then the SIMD units on each core can do more than one multiple, add, etc. at one time.

16:52:15 I have a question for Proj 4. Are we supposed to include different numtries in a column by itself along with which array size?

No, you use NUMTRIES just to get a performance number that is as close to the actual peak performance as you can get. Just report peak performance.