

# Live Lecture Chat Window

April 3, 2024

**14:29:55 For this graphing project does it matter if we are running the executable on our local machine or a school Linux server like flip?**

Probably won't matter. From what I have seen, the general shape of the curves will be about the same between machines. It's the absolute numbers that change. And it is the *shape* of the curves that I most want you to see and learn from.

**14:31:00 In the example, wouldn't we still need to change the chart titles in Excel?**

Yes, or just erase the chart title in Excel and fill it in inside PowerPoint.

**14:35:11 Is Linux/Flip required to do the projects? Or can everything in the class be done on Windows with a GPU?**

Projects #0, #1, #2, and #3 can be done anywhere with a C++/OpenMP compiler. Project #4 probably needs to be done on flip because we are using special code to get access to SIMD. Project #5 needs to be done somewhere with an Nvidia GPU, Project #6 needs to be done somewhere with an OpenCL-compatible GPU, and Project #7 needs to be done on Linux unless you have access to a compute cluster or supercomputer of your own.

**14:35:30 Is it recommended to finish the projects first before the quiz?**

The quizzes are more based on the notes than on the projects.

**14:35:38 If you don't have a GPU locally, we must use OSU servers?**

For Project #5 and #6, yes.

**14:38:38 To clarify, most of the projects are run on the CPU not the GPU until we get into CUDA and the final project. So, we won't need to use the servers until those points correct?**

See the answer above. And there is no final project.

**14:58:55 What note handouts should we make sure to read before we do Project 1?**

I would read OpenMP, Amdahl, and Trapezoid.

**15:07:15 Question - what version of g++ does flip use? Simple OpenMP says its default is 4.8.5, this slide set says 11.4**

It's 11.4 – the notes have been corrected.

**15:07:18** For the question on the Project #0 that asked why we think the 1 thread to 4 thread speed up should be less than 4, are there additional reasons we should include besides that we are dividing the workload across 4 processing units and which would be a 4 times speed up maximum in ideal perfect conditions but since we need to account for OpenMP overhead, we are going to be under the 4X upper threshold

Yes, that's basically it. There are overhead parts of the program that need to be there but cannot be parallelized, like setting up the for-loop, firing off the threads, etc.

**15:35:06** How would you know how to improve your thread allocation? Would you start and stop a timer for each thread and figure out when each thread is completing and then adjust to balance their individual runtimes as best as possible, or is there an easier way to figure out which thread or threads are lagging?

No, it really all comes down to you the domain expert understanding something about the algorithm that is being parallelized. Also, I believe Intel has a "thread-watcher" utility that could help decide if threads were lagging.

**15:38:42** Question: what was the default thread allocation option? Is it static with chunksize 1 or something else?

It is "static,1"

**15:39:18** What would be the reasoning for ever using a higher chunk size if using a chunk size of 1 handles variability better?

Can't think of a good use case right now. Possibly if an algorithm wanted to process adjacent data values.

**15:39:46** What if you have three threads, and every 3<sup>rd</sup> thread takes a long time?

I would probably use a chunksize of 3 so that each thread gets two of the easy thread and one hard one, thus balancing the load among the threads.

**15:44:57** Just wanted to confirm: after handing in project 0 on Teach, we don't need to leave a note on canvas or anything, right? We just need to hand in the cpp file and the pdf on teach and we're all set.

That's correct. Just turn it in. My script will go looking for it after the due date.

**15:49:33** I noticed that the performance is better when I plug in my [laptop] machine, I think the vendor is trying to save my battery life.

Yes, a lot of laptops underclock the laptop when it is not plugged in to save battery life.

**16:00:02** To access flip/rabbit server, it must be done through Linux?

You should be able to VScode-access to flip or rabbit from anything. You should be able to ssh to flip or rabbit from anything. It's not so much where you are coming from but where you are connecting to.