


1


Parallel Programming Project Notes




**Oregon State
University**
Mike Bailey
mjb@cs.oregonstate.edu

**PILOT CAR
FOLLOW ME**

To a Better Grade



project_notes.pptx



mjb - April 27, 2022

Why Are These Notes Here?

2

These notes are here to:

1. Help you setup and run your projects
2. Help you get the data you collect in the right format for submission
3. Help you get a **better grade** by doing all of this correctly!



better grade!

better grade!

better grade!

better grade!

better grade!

mjb - April 27, 2022

Project Notes, I

3



- Feel free to run your projects on whatever systems you have access to.
- If you don't have access to your own systems, then you can use what we have at OSU. **On-campus users** will have access to Windows and Linux systems here. **Ecampus users** will have remote access to our Linux systems, such as **flip**.
- Most of the projects will require timing to determine performance. Use the OpenMP timing functions. They give decent answers, and this will make the timing consistent across projects and across people. The OpenMP call:

```
double prec = omp_get_wtick( );
```

tells you the precision of the clock in seconds. I get 10⁻⁹ seconds on the systems I've been using. (I really doubt that this is true.) The OpenMP call:

```
double time0 = omp_get_wtime( );
```

samples the clock right now. It gives you wall-clock time in seconds. In parallel computing, memory latency and thread-idle time are part of the equation, so wall clock time is what you want.

mjb - April 27, 2022

How We Will Be Doing Timing



4

In this class, we don't want to just *implement* – we want to *characterize* performance. What speed-ups do we see, and why do we see them? How do we generalize that to other types of problems? What insights does this give us?

So, as part of your project assignments, you will be doing a lot of timing to determine program speed-ups.

```
#include <omp.h>
double time0 = omp_get_wtime( );           // seconds
...
double time1 = omp_get_wtime( );           // seconds
printf( stderr, "Elapsed time = %10.2f microseconds\n", 1000000. * ( time1 - time0 ) );
```

%10.2f is a good way to print doubles ("long float")

mjb - April 27, 2022

How Reliable is the Timing?

5

This way of timing measures **wall-clock time**, which is really what we want to know in a parallel environment, *not CPU time*.

However, this puts you at the mercy of the other users on the system. If you are on one of our public systems (e.g., flip), I advise you to check the system load to see how much off your wall-clock time measurement will be due to the competition from other users. Use the Linux **uptime** command::

flip01 34% uptime

11:13:37 up 96 days, 11:52, 23 users, load average: 3.56, 3.08, 2.82

These three numbers represent total CPU load averages for the last 1, 5, and 15 minutes respectively. If the CPU load average is greater than the number of CPUs, then each CPU is over-burdened.

Clearly you want these numbers, especially the 1-minute one, to be as small as possible when you run your test. If they are "big", you might want to ssh to other systems (flip01, flip02, flip03, ...) to see if you can find a better place to run, or try again later.



mjb - April 27, 2022

How Reliable is the Timing? A Useful Trick!

6

I like to check the consistency of the timing by computing both peak speed and average speed and seeing how close they are:

```
double maxmflops = 0.;
double summflops = 0.;

for( int t = 0; t < NUMTRIES; t++ )
{
    double time0 = omp_get_wtime();

    #pragma omp parallel for
    for( int i = 0; i < ARRAYSIZE; i++ )
    {
        C[ i ] = A[ i ] * B[ i ];
    }

    double time1 = omp_get_wtime();
    double mflops = (double)ARRAYSIZE/(time1-time0)/1000000.;
    summflops += mflops;
    if( mflops > maxmflops )
        maxmflops = mflops;
}

printf( " Peak Performance = %8.2f MFLOPS\n", maxmflops );
printf( "Average Performance = %8.2f MFLOPS\n", summflops/(double)NUMTRIES );
```

You should record the peak performance value. This gives you as close to the best answer that you will get. But, compare that with the average performance value. That will tell you how reliable that peak value is.

This is a reliable result:

Peak Performance = 1183.31 MFLOPS
Average Performance = 1141.41 MFLOPS

This is an unreliable result:

Peak Performance = 627.39 MFLOPS
Average Performance = 294.86 MFLOPS

Project Notes, II

7

If you are on Linux and have access to the Intel compiler, *icc*, don't use it unless we tell you to! (icc is so good that it often does optimizations that undermine the very things you are testing.)

Use g++. The compilation sequences are:

On Linux, the typical compile sequence for files that use OpenMP is:

```
g++ -o proj proj.cpp -O3 -lm -fopenmp
```

```
icc -o proj proj.cpp -O3 -lm -openmp -align -qopt-report=3 -qopt-report-phase=vec
```

Note that OpenMP should always be included because we are using OpenMP calls for timing.

Note that the second character in the 3-character sequence "-lm" is an ell, i.e., a lower-case L. This is how you link in the **math** library.



mjb - April 27, 2022

Project Notes, III

8

• Most of these projects will require you to submit graphs. You can prepare the graphs any way you want, except for drawing them by hand. (The Excel Scatter-with-Smooth-Lines-and-Markers works well.) So that we can easily look at each other's graphs, please follow the convention that **up is faster**. That is, do not plot seconds on the Y axis because then "up" would mean "slower". Instead, plot something like Speedup or MFLOPS or frames-per-second.

• I expect the graphs to show off your **scientific literacy** -- that is, I expect axes with numbers, labels, and units. If there are multiple curves on the same set of axes, I expect to be able to easily tell which curve goes with which quantity. After all, there is a reason this major is called Computer **Science**. Not doing this makes your project unacceptable for grading.

You lose points if you don't do it this way.



mjb - April 27, 2022

Diverting to a File

9

We all have a tendency to want to write our performance results out using *printf* (or *cout*) so that we can see them on the screen. That's fine. But, then we want to get those results into a file. You could mess with file I/O, or you could use a *divert* on the command line.

If you are currently running your program like this:

```
./proj01
```

and it prints to the standard output screen via *printf* or *cout*, then running it like this:

```
./proj01 > output.csv
```

will write your results into the file *output.csv*

(If you do it a second time, you will probably have to remove the previous *output.csv* first.)

You can also divert the *entire* output (standard out and standard error) of a looping script:

```
loop.bash >& output
```



mjb - April 27, 2022

Importing into Excel - csv Files

10

csv stands for **comma-separated values**. It is a file format where you write your numbers out as text with commas between them. The great part is that Excel recognizes csv files and will read them in automatically.

Say you are using a printf that looks like this:

```
printf( "%2d threads ; %8d trials ; probability = %6.2f%% ; megatrials/sec = %6.2fn",
        NUMT, NUMTRIALS, 100.*currentProb, maxPerformance);
```

You probably did this because it looks really nice on your screen as you use this output to debug your program. But, now you want to change it to get the numbers into Excel quickly and painlessly. Comment out the old way and change it to this:

```
//printf( "%2d threads ; %8d trials ; probability = %6.2f%% ; megatrials/sec = %6.2fn",
        //NUMT, NUMTRIALS, 100.*currentProb, maxPerformance);
```

```
printf( "%2d, %8d, %6.2fn", NUMT, NUMTRIALS, maxPerformance );
```

This will now be printing just what you need into CSV format. You could divert it like this:

```
./proj01 > OUT.csv
```

or

```
loop.bash > OUT.csv
```

Which would then let you read the OUT.csv file right into Excel.

mjb - April 27, 2022

Importing into Excel - csv Files

11

```
flip3 154% cat OUT.csv
```

```
1, 1, 1.44
1, 10, 3.99
1, 100, 8.07
1, 1000, 9.33
1, 10000, 23.40
1, 100000, 25.13
1, 500000, 25.97
2, 1, 0.23
2, 10, 4.62
2, 100, 19.26
2, 1000, 17.91
2, 10000, 34.34
2, 100000, 49.83
2, 500000, 49.27
4, 1, 0.34
4, 10, 0.259
4, 100, 16.7
4, 1000, 38.66
4, 10000, 82.39
4, 100000, 91.09
4, 500000, 91.49
8, 1, 0.26
8, 10, 2.39
8, 100, 16.21
8, 1000, 48.49
8, 10000, 137.59
8, 100000, 166.17
8, 500000, 181.62
```



mjb - April 27, 2022

	A	B	C
1	1	1	1.44
2	1	10	3.99
3	1	100	8.07
4	1	1000	9.33
5	1	10000	23.4
6	1	100000	25.13
7	1	500000	25.97
8	2	1	0.23
9	2	10	4.62
10	2	100	19.26
11	2	1000	17.91
12	2	10000	34.34
13	2	100000	49.83
14	2	500000	49.27
15	4	1	0.34
16	4	10	0.259
17	4	100	16.7
18	4	1000	38.66
19	4	10000	82.39
20	4	100000	91.09
21	4	500000	91.49
22	8	1	0.26
23	8	10	2.39
24	8	100	16.21
25	8	1000	48.49
26	8	10000	137.59
27	8	100000	166.17
28	8	500000	181.62

Making Graphs

12

In Excel, I have had the most success with creating graphs from tables that look like this:

	A	B	C
1	1	1	1.44
2	1	10	3.99
3	1	100	8.07
4	1	1000	9.33
5	1	10000	23.4
6	1	100000	25.13
7	1	500000	25.97
8	2	1	0.23
9	2	10	4.62
10	2	100	19.26
11	2	1000	17.91
12	2	10000	34.34
13	2	100000	49.83
14	2	500000	49.27
15	4	1	0.34
16	4	10	0.259
17	4	100	16.7
18	4	1000	38.66
19	4	10000	82.39
20	4	100000	91.09
21	4	500000	91.49
22	8	1	0.26
23	8	10	2.39
24	8	100	16.21
25	8	1000	48.49
26	8	10000	137.59
27	8	100000	166.17
28	8	500000	181.62

	E	F	G	H	I	J	K	L	M
		1	10	100	1000	10000	100000	500000	
1		1.44	3.99	8.07	9.33	23.4	25.13	25.97	
2		0.23	4.62	19.26	17.91	34.34	49.83	49.27	
4		0.34	0.259	16.7	38.66	82.39	91.09	91.49	
8		0.26	2.39	16.21	48.49	137.59	166.17	181.62	

where the 1,2,4,8 rows are holding the number of threads constant, and the 1, 10, 100, 1000, etc. columns are holding the dataset size constant. The cells are holding performance numbers, with higher numbers representing faster performance.

You will need to do some copying and pasting to get the linear format (from the previous page) into this 2D format, but it will be worth it!



You can avoid *all* that copying and pasting by using an Excel feature called *Pivot Tables*! See the Pivot Table noteset.

mjb - April 27, 2022

Making Graphs

13

Then, sweep over the entire table

	E	F	G	H	I	J	K	L	M
		1	10	100	1000	10000	100000	500000	
1	1.44	3.99	8.07	9.33	23.4	25.13	25.97		
2	0.23	4.62	19.26	17.91	34.34	49.83	49.27		
4	0.34	0.259	16.7	38.66	82.99	91.09	91.49		
8	0.26	2.39	16.21	48.49	137.59	166.17	181.62		



mjb - April 27, 2022

Making Graphs

14

Then, select Copy, and then insert it into one of the scatterplot options.

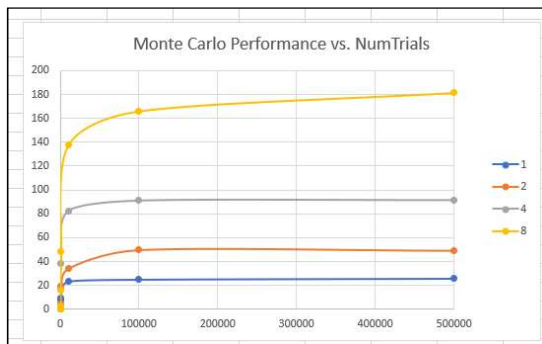
I like this one:



mjb - April 27, 2022

Here's the Graph that Excel Produced from this Data

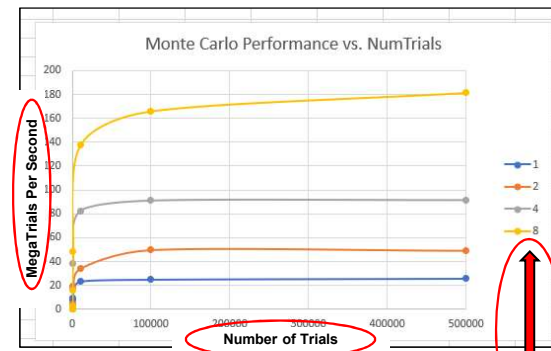
15



mjb - April 27, 2022

Adding Annotation

16



I like using PowerPoint for this, but you need to do it somewhere to get the proper graph into your project report.



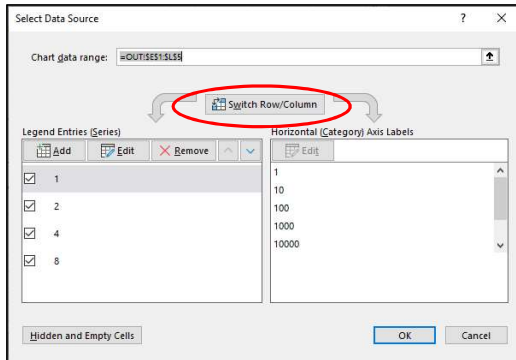
mjb - April 27, 2022

Transposing the Graph

17

To transpose the sense of the graph (which you also need to do), right-click on the border of the graph and then click on "Select Data".

Then click on "Switch Row/Column".



Here's the Graph that Excel Produced from this Data

18



You might have to adjust the X axis number maximum and minimum.

Adding Annotation

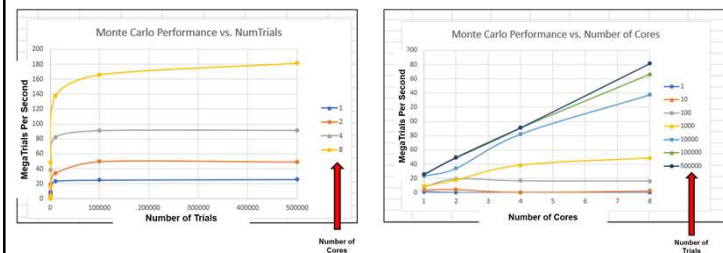
19

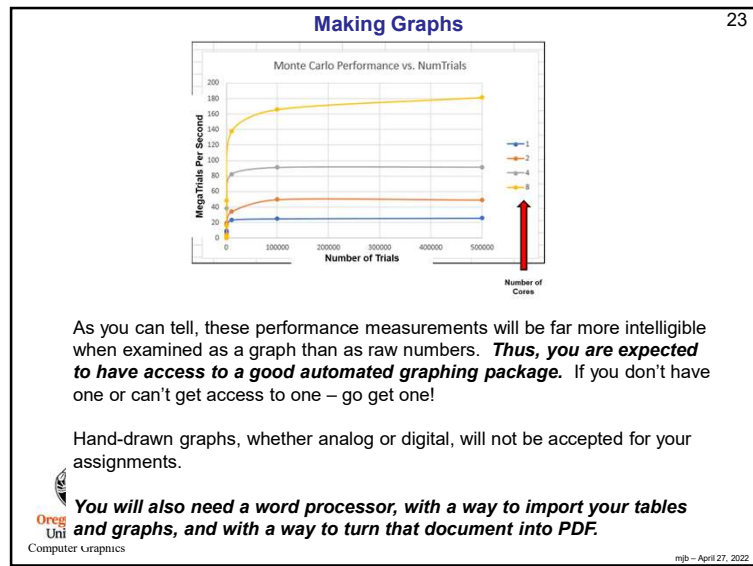
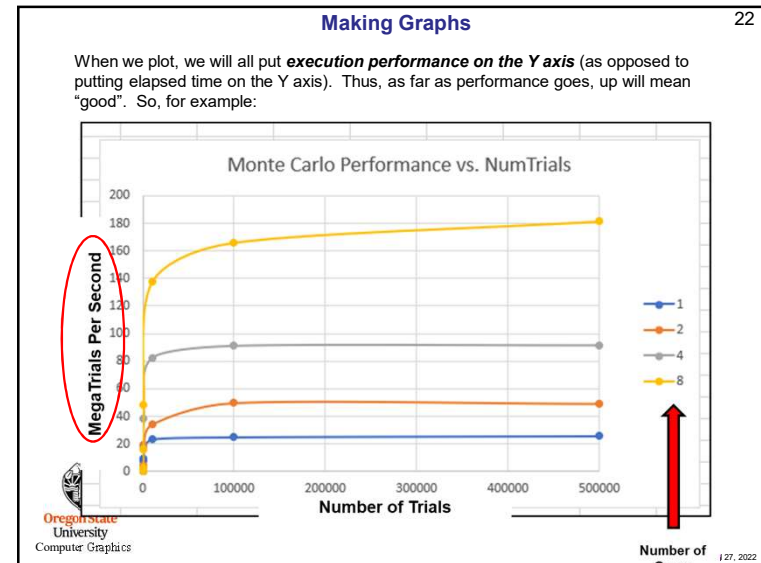
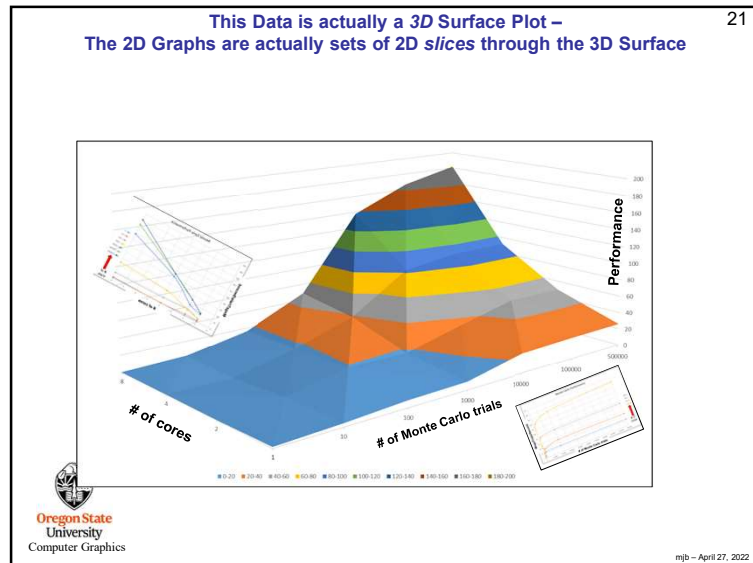


I like using PowerPoint for this, but you need to do it *somewhere* to get the proper graph into your project report.

It's the Same Data, but Each Graph Gives You a Different Insight into what the Data is Telling You

20





Setting up Your Benchmarks to run from Scripts:
#1 -- the #define Approach

There are always advantages to not hardcoding constants into the middle of your program and, instead, setting them with a #define at the top where you can find that value and change it easily, like this:

```
#include <stdio.h>
#include <math.h>

#ifndef NUMT
#define NUMT      2
#endif

#ifndef NUMS
#define NUMS     32
#endif
```

Then, in the C or C++ program, all you have to do is use NUMT to set the number of threads, like this:

```
omp_set_num_threads( NUMT );
```

But, the use of the #ifndef/#endif construct has other advantages. It lets you either run this as a standalone program or run many occurrences of the program from a script.

Oregon State University
Computer Graphics

mjb – April 27, 2022

Setting up Your Benchmarks to run from Scripts: #1 -- the #define Approach

25

In our project assignments, you will run benchmarks, that is, you will try your application using several different combinations of parameters. Setting these combinations by hand inside your program one-by-one is a time-consuming pain.

Your time is more valuable than that. Try doing it from a script.

In most C and C++ compilers, there is some mechanism to set a **#define** from outside the program. Many of them use the **-D** construct on the command line:

```
#!/bin/bash

#number of threads:
for t in 1 2 4 6 8
do
  echo NUMT = $t
  g++ -DNUMT=$t prog.cpp -o prog -lm -fopenmp
  ./prog
done
```

Then, in the C or C++ program, all you have to do is use NUMT. For example:

```
omp_set_num_threads( NUMT );
```

This lets you automatically run your program 5 times with 1, 2, 4, 6, and 8 threads.

http://pages.cse.oregonstate.edu/~mjb/

Setting up Your Benchmarks to run from Scripts: Method #1 -- the #define Approach

26

You can also test multiple parameters from the same script by nesting the loops. This one is done using **Bash Shell (bash)**:

```
#!/bin/bash

# number of threads:
for t in 1 2 4 6 8
do
  echo NUMT = $t
  # number of subdivisions:
  for s in 2 4 8 16 32 64 128 256 512 1024 2048 3072 4096
  do
    echo NUMS = $s
    g++ -DNUMS=$s -DNUMT=$t prog.cpp -o prog -lm -fopenmp
    ./prog
  done
done
```



http://teespring.com

mjb - April 27, 2022

Or, in C Shell (csh)...

27

```
#!/bin/csh

# number of threads:
foreach t ( 1 2 4 6 8 )
  echo NUMT = $t
  # number of subdivisions:
  foreach s ( 2 4 8 16 32 64 128 256 512 1024 2048 3072 4096 )
    echo NUMS = $s
    g++ -DNUMS=$s -DNUMT=$t prog.cpp -o prog -lm -fopenmp
    ./prog
  end
end
```



mjb - April 27, 2022

Or, in Python...

28

```
import os

for t in [ 1, 2, 4, 6, 8 ]:
  print "NUMT = %d" % t
  for s in [ 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 3072, 4096 ]:
    print "NUMS = %d" % s
    cmd = "g++ -DNUMS=%d -DNUMT=%d prog.cpp -o prog -lm -fopenmp % ( s, t )"
    os.system( cmd )
    cmd = "./prog"
    os.system( cmd )
```



mjb - April 27, 2022

Setting up Your Benchmarks to run from Scripts: Method #2 -- the Command Line Arguments Approach

29

Instead of this:

```
#include <stdio.h>
#include <math.h>

#ifdef NUMT
#define NUMT 8
#endif

#ifdef NUMS
#define NUMS 32
#endif
```

Do this:

```
#include <stdio.h>
#include <math.h>

int NUMT = 8;
int NUMS = 32
```

Then, in the C or C++ program, all you have to do is use NUMT to set the number of threads, like this:

```
omp_set_num_threads( NUMT );
```

But, the use of the `#ifdef/#endif` construct has other advantages. It lets you either run this as a standalone program or run many occurrences of the program from a *script*.

argc and argv

30

When you write in C or C++, your *main* program, which is really a special function call, looks like this:

```
int main( int argc, char *argv[ ] )
{
    ...
}
```

These arguments describe what was entered on the command line used to run the program.

The **argc** is the number of arguments (the arg **C**ount)

The **argv** is a list of argc character strings that were typed (the arg **V**ector).

The name of the program counts as the 0th argv (i.e., argv[0])

So, for example, when you type

```
ls -l
```

in a shell, the *ls* program sees argc and argv filled like this:

```
argc = 2
argv[0] = "ls"
argv[1] = "-l"
```



argc and argv

31

So, if NUMT and NUMS are global int variables:

```
int NUMT = 2;
int NUMS = 32;
```

and you want to set them from the command line, like this:

```
./prog 1 64
```

Then, *inside your main program*, you would say this:

```
if( argc >= 2 )
    NUMT = atoi( argv[1] );
```

```
if( argc >= 3 )
    NUMS = atoi( argv[2] );
```

The if-statements guarantee that nothing bad happens if you forget to type values on the command line.

The *atoi* function converts a string into an integer ("ascii-to-integer"). If you ever need it, there is also an *atof* function for floating-point.



shared() in the #pragma omp Line

32

Also remember, if you use Method #2, then NUMS is a *variable*, and it needs to be declared as *shared* in the *#pragma omp* line:

```
#pragma omp parallel for default(none) shared(NUMS,xcs,ycs,rs,tn) reduction(+:numHits)
```

NUMT does not need to be declared in this way because it is not used in the for-loop that has the *#pragma omp* in front of it.



Setting up Your Benchmarks to run from Scripts: Method #2 -- the Command Line Arguments Approach

33

In our project assignments, you will run benchmarks, that is, you will try your application using several different combinations of parameters. Setting these combinations by hand inside your program one-by-one is a time-consuming pain. Your time is more valuable than that. Try doing it from a script.

```
#!/bin/bash

g++ prog.cpp -o prog -lm -fopenmp

#number of threads:
for t in 1 2 4 6 8
do
    echo NUMT = $t
    ./prog $t
done
```

Then, in the C or C++ program, all you have to do is use NUMT. For example:

```
omp_set_num_threads( NUMT );
```

This lets you automatically run your program 5 times with 1, 2, 4, 6, and 8 threads.

http://pages.cs.oregonstate.edu/~mjb/ - April 27, 2022

Setting up Your Benchmarks to run from Scripts: #2 -- the Command Line Arguments Approach

34

You can also test multiple parameters from the same script by nesting the loops. This one is done using **Bash Shell (bash)**:

```
#!/bin/bash

g++ prog.cpp -o prog -lm -fopenmp

# number of threads:
for t in 1 2 4 6 8
do
    echo NUMT = $t
    # number of subdivisions:
    for s in 2 4 8 16 32 64 128 256 512 1024 2048 3072 4096
    do
        echo NUMS = $s
        ./prog $t $s
    done
done
```



http://teespring.com

mjb - April 27, 2022

Or, in csh (C Shell) ...

35

```
#!/bin/csh

g++ prog.cpp -o prog -lm -fopenmp

# number of threads:
foreach t ( 1 2 4 6 8 )
    echo NUMT = $t
    # number of subdivisions:
    foreach s ( 2 4 8 16 32 64 128 256 512 1024 2048 3072 4096 )
        echo NUMS = $s
        ./prog $t $s
    end
end
```



mjb - April 27, 2022

Or, in Python...

36

```
import os

cmd = "g++ prog.cpp -o prog -lm -fopenmp"
os.system(cmd)

for t in [ 1, 2, 4, 6, 8 ]:
    print "NUMT = %d" % t
    for s in [ 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 3072, 4096 ]:
        print "NUMS = %d" % s
        cmd = "./prog %d %d" % ( s, t )
        os.system( cmd )
```



mjb - April 27, 2022

I Don't Recommend That You Put These Loops in the .cpp Program! 37

I know what you're thinking.

You're thinking:

Those scripts are a pain, and I've never done them before. So, I'll just build the iterations through all the parameters into for-loops in the program.

Don't! I have seen evidence that the first time OpenMP does anything, it also does some one-time setups. This will mess up your timing because your first test will seem slower than it should be and the others will seem faster by comparison.

I recommend you run the program *separately* for each combination of parameters. (The script code in the previous pages shows that.)



mjb - April 27, 2022

Computing Performance 38

When computing performance, be sure that the **numerator is amount of work done** and the **denominator is the amount of time it took to do that work**. For example, in the Bezier surface example, computing one height is the work done at each node and you have NUMS*NUMS total nodes, so (NUMS*NUMS)/dt is one good way to measure performance.

NUMS, NUMS*NUMS, 1./dt, and NUMS/dt are not good ways to measure performance as they don't reflect the true amount of **work done per time**.

If you are using ridiculously high values for NUMS, the quantity NUMS*NUMS might overflow a normal 32-bit int. You can use a long int, or just float each one separately. Instead of (float)(NUMS*NUMS)/dt, you could say (float)NUMS*(float)NUMS/dt

If you are squaring a size number, and are using signed ints, the largest NUMS you can use is:

$$\sqrt{2,147,483,647} = 46,340$$

If you are squaring a size number, and are using unsigned ints, the largest NUMS you can use is:

$$\sqrt{4,294,967,295} = 65,535$$



mjb - April 27, 2022

Project Turn-in Procedures 39

Your project turnins will all be electronic. For *all* projects:

Your project turnins will be done at <http://enr.oregonstate.edu/teach> and will consist of:

1. Source files of everything (.cpp, .cl, .cuda)
2. A report in PDF format. **You can .zip everything else if you want, but please leave the PDF as a separate file.**

Electronic submissions are due at 23:59:59 on the listed due date.

Your PDF report will include:

1. **A title area, including your name, email, project number, and project name.**
2. Any tables or graphs to show your results.
3. An explanation of what you did and *why it worked the way it did*. Your submission will not be considered valid unless you at least attempt to explain why it works the way it does.

Your project will be graded and the score posted to the class web page. Any loss of points will be explained in a Canvas grade comment.

Don't forget that title area! Because I concatenate all the class PDFs together for grading, lack of a name on your project will make it difficult to know who to give this grade to. It will be **5 points off** if you forget.



mjb - April 27, 2022

Bonus Days 40

Projects are due at 23:59:59 on the listed due date, with the following exception:

Each of you has been granted **five** Bonus Days, which are no-questions-asked 24-hour project extensions which may be applied to any project, subject to the following rules:

1. If the project handout says "No Bonus Days", then no Bonus Days can be used on it. It is due exactly at 23:59:59 on the given due date, no later.
2. No more than **2** Bonus Days may be applied to *any one project*
3. Bonus Days cannot be applied to quizzes or tests



mjb - April 27, 2022

Bonus Days

41

To use one or more Bonus Days on a given project:

- You don't need to let me know ahead of time.
- Turn-in promptness is measured **by date**. Don't worry if *teach* tells you it's late because it is between 23:30:01 and 23:59:59. But, *after* 23:59:59 on the posted due date, **it's late!**
- For any project for which Bonus Days can be used, *teach* has been instructed to accept your turn-in, no matter when you do it.
- I will run a script to identify the projects that will have Bonus Days deducted
- If you have used Bonus Days on a given project, I will send you an email (well, a Python script will, actually) telling you how many you used and how many you have left.



mjb - April 27, 2022

A Warning About Virtual Machines

42

Virtual machines are, apparently, not automatically setup to do multithreading.

If you are running on your own virtual machine, and are getting performance numbers that make absolutely no sense, try using one of the OSU machines.



mjb - April 27, 2022

A Warning about Editing on Windows and Running on Linux

43

Some of you will end up having strange, unexplainable problems with your csh scripts or .cpp programs. This could be because you are typing your code in on Windows (using Notepad or Wordpad or Word) and then running it on Linux. Windows likes to insert an extra carriage return ('\r') at the end of each line, which Linux interprets as a garbage character.

You can confirm this by typing the Linux command:

```
od -c loop.csh
```

which will show you all the characters, even the '\r' (carriage returns, which you don't want) and the '\n' (newlines, which you do want).

To get rid of the carriage returns, enter the Linux command:

```
tr -d '\r' < loop.csh > loop1.csh
```

Then run loop1.csh

This works too:

```
sed -i -e 's/\r$//' loop.csh
```

Or, on some systems, there is a utility called *dos2unix* which does this for you:

```
dos2unix < loop.csh > loop1.csh
```

Sorry about this. Unfortunately, this is a fact of life when you mix Windows and Linux.



mjb - April 27, 2022