

CSCI 3366 (Parallel and Distributed Processing), Fall 2017

Homework 2

Credit: 55 points.

1 Reading

Be sure you have read, or at least skimmed, readings from the relevant updated appendices.

2 Honor Code Statement

Please include with each part of the assignment the Honor Code pledge or just the word “pledged”, plus one or more of the following about collaboration and help (as many as apply).¹ Text *in italics* is explanatory or something for you to fill in. For written assignments, it should go right after your name and the assignment number; for programming assignments, it should go in comments at the start of your program(s).

- This assignment is entirely my own work. (*Here, “entirely my own work” means that it’s your own work except for anything you got from the assignment itself — some programming assignments include “starter code”, for example — or from the course Web site. In particular, for programming assignments you can copy freely from anything on the “sample programs page”.*)
- I worked with *names of other students* on this assignment.
- I got help with this assignment from *source of help — ACM tutoring, another student in the course, the instructor, etc.* (*Here, “help” means significant help, beyond a little assistance with tools or compiler errors.*)
- I got help from *outside source — a book other than the textbook (give title and author), a Web site (give its URL), etc..* (*Here too, you only need to mention significant help — you don’t need to tell me that you looked up an error message on the Web, but if you found an algorithm or a code sketch, tell me about that.*)
- I provided help to *names of students* on this assignment. (*And here too, you only need to tell me about significant help.*)

3 Overview

Your mission for this assignment is to improve the programs you wrote for Homework 1, to write versions in Java and OpenCL, and to measure their performance and accuracy more systematically.

¹Credit where credit is due: I based the wording of this list on a posting to a SIGCSE mailing list. SIGCSE is the ACM’s Special Interest Group on CS Education.

4 Details

4.1 Thread-safe RNG

(5 points)

Your first step will be to write a thread-safe random number generator, i.e., one that can be called from multiple threads concurrently without ill effects. To keep this part manageable, I suggest that you just use the technique mentioned in class, LCG (Linear Congruential Generator). The [Wikipedia article](#)² has a pretty good discussion, but briefly:

This algorithm generates a pseudorandom sequence x_0, x_1, x_2, \dots from a seed S , constants a , c , and M , and a simple recurrence relation:

$$\begin{aligned}x_0 &= S \\x_i &= (ax_{i-1} + c) \pmod{M}, \text{ for } i > 0\end{aligned}$$

The Wikipedia article gives values used by many library implementation of this algorithm; to me the most attractive choice is the one cited for two POSIX functions, namely $a = 25214903917$, $c = 11$, and $M = 2^{48}$. (This seems attractive because — if I understand the discussion correctly — it will generate long sequences without duplicates (which we want), and values will be within the range of a 64-bit signed data types, which is available as `int64_t` in standard C and `Long` in Java.) Also, the `mod` part of the calculation is easily done by using bitwise and with $2^{48} - 1$. (Note that you will need to `#include <stdint.h>` to use `int64_t`.)

(If for some reason you want to try a different algorithm, check with me first — there may well be better choices, but there are probably worse choices too.)

You will need two implementations of whatever algorithm you choose, one in C and one in Java. Exactly how you package the algorithm is somewhat up to you, but you want functions analogous to `srand()` and `rand()`, and there needs to be some way to deal with the “state” of the sequence being generated (the current or next x_i) in a way that makes it possible for each thread to have its own state (rather than there being one hidden global state, as with `srand()` and `rand()`).

For C, what I think makes sense is to represent the saved state as a `int64_t` and define two functions that take a pointer to a state as a parameter:

- `void rand_set_seed(long seed, rand_state_t *state);`
- `int64_t rand_next(rand_state_t *state);`

You’ll also want to define a constant, with something such as the following:

```
const int64_t RANDMAX = (1LL << 48) - 1;
```

(Notice that this is $M - 1$.)

For Java, you’ll probably want to define a class analogous to `java.util.Random`, but much simpler, with just a `RANDMAX` constant, a constructor, and a `next` method.

²http://en.wikipedia.org/wiki/Linear_congruential_generator

4.2 Revised sequential programs

The next step is to replace the current code for generating random numbers in two starter programs, one in C and one in Java, with your RNG code:

- C program: `monte-carlo-pi.c`³. Also requires `timer.h`⁴. (This is the starter code from Homework 1, *except* it uses `long` rather than `int` for values where I think you do want at least 32 bits, which `long` guarantees but `int` does not. You can use the command `diff` to see differences between the two versions (or try `vimdiff` or `vimdiff -o`.)
- Java program: `MonteCarloPi.java`⁵. (Note that the class this defines is in package `csci3366.hw2`, so it should go in a directory named `csci3366/hw2`.)

4.2.1 Code

(5 points)

Replace the current code for generating random numbers in the two starter sequential programs with calls to your RNG. (If you didn't already test your RNG code, you might temporarily put in some debug-print statements to be sure it's generating reasonable output.) The two programs (C and Java) should now produce the same output (except for execution time).

4.2.2 Results (accuracy)

(5 points)

(You only need to do this for one of your sequential programs, since they should give the same results.) Experiment until you find a seed that seems to give reasonable results, and then measure the relationship between accuracy (difference between the computed value of π and the constant as defined in the math library) and number of samples: Generate output for at least six different values of "number of samples" (I recommend starting with a medium-size number and then repeatedly doubling it, rather than increasing by a fixed amount). Plot the results, by hand or with whatever program you like. (I use `gnuplot`. Short introduction/example below.) You can repeat this for more than one seed and plot all sets of results if you like.

4.3 Parallel programs

4.3.1 Code

(30 points)

Your mission for this step is to produce parallel programs for our four programming environments: C with OpenMP, C with MPI, Java, and C with OpenCL.

- For OpenMP and MPI, you should be able combine what you did for Homework 1 with what you did for the first step (sequential program with your own RNG).

³http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/monte-carlo-pi.c

⁴http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/timer.h

⁵http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/MonteCarloPi.java

- For Java, you'll have to figure out how to “parallelize” what you did for the first step, but you should be able to adapt the numerical integration example (on the “sample programs” page). As with the numerical integration example (as recently updated), your program should get the number of threads from an additional command-line argument.
- For OpenCL, again you'll have to figure out how to parallelize, but you should be able to adapt the numerical-integration example (on the “sample programs” page). Like that example, your program should take additional command-line arguments that let you vary what can be varied (number of work items, work group size).

So to recap, command-line arguments should be as follows:

- For OpenMPI and MPI (same as for Homework 1): number of samples, seed.
- For Java: number of samples, seed, number of threads.
- For OpenCL: number of samples, seed, number of work items, factor that lets you vary workgroup size (to me a reasonable choice here is what I do in the (latest version of the) numerical integration example, a factor by which to multiply the “preferred size”).

As we noted in class, having all UEs (processes or threads) generate points using the same RNG and seed is not useful. You have two options for dealing with this:

- Use a different seed in each UE. As noted in class, simple methods of combining a “master seed” with UE ID (adding or multiplying them) may produce overlapping sequences, but figuring out how to avoid that is somewhat beyond the scope of this assignment.
- Arrange for each UE to generate only a part of the whole “random” sequence. In principle this should be straightforward: If you want to split the above-described sequence among p UEs, you can do so by generating similar sequences in each UE, but with constants

$$\begin{aligned} a' &= a^p \pmod{M} \\ c' &= c(a^p - 1)/(a - 1) \pmod{M} \end{aligned}$$

and starting the sequence for the i -th UE at element x_i of the original sequence. (I originally found this in a paper that no longer seems to be freely accessible, but it's repeated [here](#)⁶.) To me this seems like the right way to go, but it's more work, so I'll give extra credit for trying it. Hints and partial code for a C version below.

Hints for using leapfrogging:

- What I found to make the most sense was to package things up in a slightly different way: For Java, a class still makes sense, but I think its constructor should take two more arguments, the number of “streams” (UEs for us) and which stream this object is for. You could put the code to generate the modified constants in the constructor, and I think it's fairly straightforward to do and to get right if in computing the constants you use `BigInteger` for intermediate values and only convert to `Longs` at the end (when the “ \pmod{M} ” step gives you a result you know will fit).

⁶<http://wotug.org/parallel/nhse/NHSEreview/RNG/node15.html>

For C, I thought it made sense to make the “state” a struct and introduce one more function

```
void rand_init_state(int p, int id, rand_state_t *state);
```

that computes and saves the values for the modified constants.

- Computing the modified constants — there may be some way to do this without arbitrary-precision arithmetic, but I didn’t think of one so chose to just use the GMP package, as mentioned in class. I didn’t find this so easy so am willing to share most of what I came up with — I’ve left out a few parts of the code (look for “FIXME”) to keep this from being too easy(?), but I’m also including a test program you can use to confirm that what you’re doing works:
 - [leapfrog-lcg.h](#)⁷ containing a struct and functions. (So you would use `#include "leapfrog-lcg.h"` to include this in your code.)
 - [test-leapfrog.c](#)⁸ containing a test program.
 - [Makefile](#)⁹ containing a make file that may be helpful. Notice that if you don’t use this you need to remember to compile/link with `-lgmp` to include the GMP library functions.

4.3.2 Results (accuracy)

(5 points)

(*UPDATED for OpenCL*) (You only need to do this for one of your parallel programs, since they should give the same results for the same number of units of execution, where “units of execution” is threads for OpenMP and Java, processes for MPI, and work items for OpenCL.) Experiment until you find a seed and number of samples that seem to give good results, and then measure the relationship between accuracy (difference between the computed value of π and the constant as defined in the math library) and number of UEs. Generate output for at least six different values of “number of UEs” (I recommend powers of two, starting with one). (Since for OpenCL the number of work items has to be a multiple of the minimum work-group size, it might be interesting to make a second plot showing that minimum value and then several multiples of it.) Plot the results, again by hand or with whatever program you like.

4.3.3 Results (performance)

(5 points)

For the values of seed and number of samples you used above, measure execution times for both sequential programs and all three parallel programs. For the parallel programs, measure execution times using different numbers of UEs (start with one and double until you notice that execution time is no longer decreasing). I strongly encourage you to do this on the machines that to me seem most suitable in terms of being able to “scale up” to interesting numbers of UEs: For OpenMP and Java, that would be Dione, for MPI, the Pandora cluster, and for OpenCL, Deimos or one of the Atlas machines. You should do each measurement more than once; if you get wildly different

⁷http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/leapfrog-starter/leapfrog-lcg.h

⁸http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/leapfrog-starter/test-leapfrog.c

⁹http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/leapfrog-starter/Makefile

results it probably means you are competing with other work on the machine and should try again another time or using another machine or machines.

Plot the results, again by hand or with whatever program you like:

- For the OpenMP, MPI, and Java programs, plot execution time versus number of UEs, and also show execution time for the sequential program in the same base language (C or Java).
- (*UPDATED for OpenCL*) For the OpenCL program, do as for the others, but also show at least two sets of values for different work-group sizes.

4.4 Hints and tips

- Feel free to borrow code from any of the sample programs linked from the [course sample programs page](#)¹⁰. This page also contains links to my writeups about compiling and running programs on the lab machines. The [course “useful links” page](#)¹¹ has pointers to documentation on all four environments.
- You can develop your programs on any system that provides the needed functionality, but I will test them on the department’s Linux classroom/lab machines, so you should probably make sure they work in that environment before turning them in.

5 What to turn in and how

Turn in the following:

- All source code (your two RNG implementations, revised sequential programs, and parallel programs). Call them whatever you like, as long as it’s clear what’s what, but please have them get input from command-line arguments as discussed above.
- Plots (accuracy of sequential program(s), accuracy of parallel program(s), and performance of parallel programs).
- Input data for plots. A text file or text files is fine for this. Also say which machines you used for the performance measurements.

Submit your program source code by sending mail to bmassing@cs.trinity.edu. Send program source as attachments. You can turn in your plots and input data as hardcopy or by e-mail; I have a slight preference for e-mail and a definite preference for something easily readable on one of our Linux machines — so, PDF or PNG or the like (in the past I think some students have sent me Excel spreadsheets, which — I’d rather you didn’t). Please use a subject line that mentions the course number and the assignment (e.g., “csci 3366 homework 2”).

6 A very little about gnuplot

I talked about the plotting tool `gnuplot` in class one day (9/25). Here are files for a simple example along the lines of what you need to do for this assignment (plot parallel times as a function of UEs, also showing sequential time):

¹⁰http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/SamplePrograms/index.html

¹¹http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/HTML/links.html

- Plot input file `par.plotin`¹².
- Data files `seq-times.dat`¹³, `par-1-times.dat`¹⁴, `par-2-times.dat`¹⁵.

With all these files in a directory, the command `gnuplot < par.plotin` will generate a file `par-times.png` with the plot.

¹²http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/par.plotin

¹³http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/seq-times.dat

¹⁴http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/par-1-times.dat

¹⁵http://www.cs.trinity.edu/~bmassing/Classes/CS3366_2017fall/Homeworks/HW02/Problems/par-2-times.dat