

Checking Student Programs

Since I haven't given you much work yet, I thought I'd send out a little "extra" lab today. The goal of this lab is to occupy as much of your time as possible (*i.e.*, all of it) so as to encourage you to feel that you must be learning something. Unfortunately, I didn't feel that I still had the time to identify any objectives that actually involve learning. So. There's that. As you probably noticed, I set the deadline nominally for mid-June. I don't think you'll be able to finish this lab and get a good grade on it unless you're really willing to dedicate yourself, so you may want to go ahead and cancel that internship, prestigious summer activity, research plan, *etc.*, so as to focus on this assignment. Maybe drop your other classes, too.

1. Execution Time

(a) Given a C++ program of no more than 100,000 lines and an input data set of no more than 1 GB, use PIN to determine the program's execution time. You need not be exact: an answer within a factor of two will earn some credit.

If you complete this problem quickly, you may wish to apply your solution to the subset of programs and relevant input sets available through SourceForge and let the programs' authors know the results.

(b) Once you have your solution working, have the proofs checked by a mathematician—try looking in Altgeld hall if you are not sure where to find a mathematician. A CS theorist will suffice in a pinch—these can be found in the Siebel Center.

(c) Write your solution up cleanly, put my name at the top (your name is not necessary... I will ... deduce it from your style), and submit it to a STOC, FOCS, or some other venue of interest to the theorist or mathematician from part (b). They may encourage you to add their name to the submitted version. Add them as an "acknowledgement" if desired. (Your and/or their receiving a Turing Award *and* credit for the assignment is an unethical conflict to my interests.) Acceptance for publication is necessary for full credit.

2. Available Parallelism

Choose a large open source program—at least 100,000 lines, but preferably somewhat larger. Write a specification for the program. Do not use a human language for this purpose. As you know, human languages are ambiguous and lead to errors. Unfortunately, the state-of-the-art does not provide a sufficiently exact and easily checkable specification language for arbitrary programs, so you may have to invent something that suits your needs. Do so. Then write the specification for the program. Check that the implementation matches the execution for all possible inputs, and deliver any exceptions along with clear examples to the program's authors.

Once you have a working specification, you can use it to identify all sources of parallelism and to estimate the possible speedup for execution on a range of platforms. Be sure to consider any architectural, microarchitectural, or interference effects when estimating achieved speedup. Choose at least 20 target platforms, including the EWS lab in Everitt as one of them. For the lab case, be sure to include the interference effects due to other students' using the machines, an aspect that is both time-of-day and time-of-semester dependent. Identify the best choice of machine as a function of algorithm choice(s), input data set, time of day, day of week, which professor is teaching which course that semester, and the current prices of bulk commodities in world markets.

3. Work Efficiency

Given an assignment specification, use PIN to determine the work efficiency of the assignment. Work efficiency is defined as the ratio of engineering learning achieved through completion of the assignment relative to the learning that might have been achieved through some other activity, such as drinking. Since drinking can also be done during assignment execution, you will need to consider all interleavings of other activities with assignment execution when calculating work efficiency. Please be complete. By drinking, of course, I meant something like Dew or coffee. Please try to focus.