

Cilk and CilkPlus

1. What guarantee does the Cilk runtime system give for the runtime of a Cilk program?
2. Compile and run the `cilk_hello.c` program on PurdueNext. What problem(s) did you run into and how did you fix them? Change the program so that Done always gets output last.
3. Examine the `fib.cpp` program and read the comments. Compile and run. Observe the runtime for different numbers of workers and interpret the results.
4. Now let's write a parallel merge sort program using Cilk.
 - a. You have been provided with a serial implementation. Modify it to create a parallel Cilk version.
 - b. Draw the execution tree for your Cilk merge sort program for a list of size 10. Find T_1 and T_∞ . In general for a list of size n , what are T_1 and T_∞ ? Given these T_1 and T_∞ , what is the asymptotic parallel runtime T_p if we use a sufficiently small number of threads?

c. Compile and run your parallel merge sort using a large list size (e.g., 40000000) and compare the runtime with the serial implementation. You may need to increase the thread stacksize. You can set the number of worker threads by setting the environment variable CILK_NWORKERS. Try to find the optimal number of workers for a given list size.

5. See <http://courses.cs.tau.ac.il/368-4064/cilk-5.3.1/examples/> for some non-trivial Cilk examples. Try some of them! You'll need to modify them to use current Cilk syntax.

6. Since Cilk is not working properly for students on PurdueNext, let's redo 4 above using Stampede. The Intel compiler (the default on Stampede) supports CilkPlus by default. So to compile the mergesort program, you can just use `icc` with no special flags. To run the program, you should use a compute node rather than a login node. You can either use `idev` to get interactive access to a compute node, or submit a batch script, to run the program. The source code, Makefile, and a sample batch script are included in the tar file provided for today.