CS 140 Assignment 4: Cilkified Inner Products

Assigned February 13, 2017

Due by 11:55 pm Friday, February 17

The purpose of this assignment is to gain familiarity with Cilk constructs and tools, as well as to think about different ways of parallelizing an algorithm using Cilk. You will write a parallel routine that computes the inner product (sometimes called dot product or DDOT) of two arrays in three different ways.

1. Background

The inner product of two vectors is the sum of their elementwise products. In pure C, the inner product of two **double** arrays, each of size n, can be implemented as follows:

```
double innerprod = 0;
for(int i=0; i< n; ++i)
{
     innerprod += a[i] * b[i];
}
```

In C++, the standard library has a dedicated function for computing the inner product. It is made available by including the <numeric> header. For the purposes of this assignment (and this class), you don't have to know and use C++, but you're allowed to do so because Cilk is an extension of C++ and will accept C++ constructs.

2. What to implement

You will write three different functions to evaluate the inner product in parallel.

1. rec_cilkified(double * a, double * b, int n)

This will evaluate the inner product recursively, by splitting each array into two at each stage of the recursion and adding the partial sums at the end. You should switch to a sequential execution when the subarray sizes become smaller than a threshold called *coarseness*. Here is a picture of the calculation:



2. loop_cilkified(double * a, double * b, int n)

This function contains two nested loops: the outer loop executes *n* / *coarseness* times, and the inner loop executes *coarseness* times. You should parallelize the outer loop with *cilk_for*, and let each inner loop execution proceed sequentially. Finally, you'll combine the results of the inner loop executions by adding them all together sequentially.

3. reducer_cilkified(double * a, double * b, int n)

For this function, you'll use use a *reducer* similar to the ones shown in the Cilk programmer's guide, and let it take care of data races and combining the results.

The driver/harness code is in a file on the course GauchoSpace site. You can use this file as a template and implement the required functions, which are left blank for you to fill in. The site also includes a timing function and a Makefile.

Your program will be executed for grading as:

>> ./innerproduct [sizeofarray]

where the parameter is optional and default is 1 million. (However, larger sizes might be necessary to measure differences in performance.)

If the harness reports "incorrect" on your results, use **cilkscreen** to identify any data races that might exist in your code. (Don't forget that you must build the debug version of your code for **cilkscreen** to report line numbers where data races occur)

3. What to report/answer

- 1. Run your code with different input sizes (sizeofarray = 10^4 , 10^5 , 10^6 , 10^7 , 10^8 , 10^9) on a fixed number of cores (namely 4). Plot a graph that has input size on its x-axis (in log scale) and speedup on the y-axis. Plot all three lines (one for each parallel routine you wrote) on the same graph.
- 2. Run your code on different numbers of cores (1, 2, 3, 4, 6, 8, 12, 16) with a fixed input size (**sizeofarray** = 10^8). Plot a graph that has the number of cores on its x-axis (in normal scale) and parallel efficiency on the y-axis. Plot all three lines (one for each parallel routine) on the same graph.
- 3. (Extra Credit) Experiment with different *coarseness* values and report on the sensitivity of the performance to different values.
- 4. (Extra Credit) The harness reports the speedup of your implementations compared to the sequential inner product. It is highly likely that your routines won't enjoy linear speedup. What might be the reason? Do you think that the algorithm does not have sufficient parallelism, or do you think there is another bottleneck other than parallelism?

<u>Hint for (4)</u>: Using **cilkview**, you can empirically find out the parallelism of your code. However, you should turn the tool on only for the test of the innerproduct routine, not for the initialization and output.

You should do this assignment in a group of two. Be sure your report contains the names and perm numbers of both group members!