

## CS450/550: Parallel Programming

### Assignment 5: Improving Application Performance Using OpenMP and Algorithmic Transformations

Due: see BBLearn

## Preliminaries

You are expected to do your own work on all homework assignments. You may (and are encouraged to) engage in discussions with your classmates regarding the assignments, but specific details of a solution, including the solution itself, must always be your own work. See the academic dishonesty policy in the course syllabus.

## Submission Instructions

You should turn in an electronic archive (.zip, .tar., .tgz, etc.). The archive must contain a single top-level directory called CS450\_aX\_NAME, where “NAME” is your NAU username and “X” is the assignment number (e.g., CS450\_a1\_mg1234). Inside that directory you should have all your code (no binaries and other compiled code) and requested files, named exactly as specified in the questions below. In the event that I cannot compile your code, you may (or may not) receive an e-mail from me shortly after the assignment deadline. This depends on the nature of the compilation errors. If you do not promptly reply to the e-mail then you may receive a 0 on some of the programming components of the assignment. Because I want to avoid compilation problems, it is crucial that you use the software described in Assignment 0. Assignments need to be turned in via BBLearn.

*Turn in a single pdf document that outlines the results of each question.* For instance, screenshots that show you achieved the desired program output and a brief text explanation. If you were not able to solve a problem, please provide a brief write up (and screenshots as appropriate) that describes what you tried and why you think it does not work (or why you think it should work). You must provide this brief write up for each programming question in the assignment.

This pdf should be independent of the source code archive, but feel free to include a copy in the top level of that archive as well.

## Optional: Group Assignment

You may choose to work on this assignment in groups of 2 or 3. If you choose to work in groups, you must submit multiple *different versions* of Question 2 (see below). If your group has 2 people, you must submit 2 versions, if your group has 3 people, you must submit 3 versions. *Why a maximum of 3 people?* Because Question 2 requires creatively solving the problem in multiple ways. If the number of ways you come up with is less than the number of people in your group, then you will be in trouble.

If you choose the group assignment option, you must do the following: 1) select a group leader who will be in charge of submitting the assignment; 2) have the group leader submit by e-mail the names of the people in the group to the instructor, and CC the other group members; and 3) this must be decided *within 3 days of assignment dissemination*.

I will not entertain scenarios where there is a conflict within the group. Working in a group is optional. I will not accept assignments by multiple people within the same group. That is, I will only accept the assignment submitted by the group leader. For example, if two people decide to work together, and then decide to do their own assignments because of a conflict, I will ignore the submission by the person who was not designated the leader.




Figure 1: Example. See text for details.

## Overview

A simple operation that is used in many applications is calculating distances between points. In this assignment, we will use the Euclidean distance. For example, if  $p_1 = (x_1, y_1)$  and  $p_2 = (x_2, y_2)$ , then the distance is as follows:  $d(p_1, p_2) = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2}$ . In this assignment, you will calculate the distances between points and report the total number of distances between points that are within some distance threshold,  $\epsilon$ . Thus, you will count the total number of occurrences  $d(p_1, p_2) \leq \epsilon$ . As  $\epsilon$  increases, the total number of points within  $\epsilon$  should increase. For validation instructions see the end of the assignment.

## Illustration

Consider Figure 1 with  $p_1 = (2, 2)$ ,  $p_2 = (3, 4)$ ,  $p_3 = (5, 2)$ ,  $p_4 = (-4, 4)$ ,  $p_5 = (-2, -3)$ ,  $p_6 = (2, -1)$ . An example calculation of the distance between a few pairs of points (to a few decimal places) are as follows:

- $p_1$  and  $p_2$ :  $\sqrt{(2 - 3)^2 + (2 - 4)^2} = \sqrt{5} = 2.236$
- $p_2$  and  $p_3$ :  $\sqrt{(3 - 5)^2 + (4 - 2)^2} = \sqrt{8} = 2.828$
- $p_6$  and  $p_5$ :  $\sqrt{(2 - (-2))^2 + (-1 - (-3))^2} = \sqrt{20} = 4.472$

Consider all of the points in Figure 1.

- If  $\epsilon = 1$  or  $\epsilon = 2$ , then no pairs of points are within  $\epsilon$ .
- If  $\epsilon = 3$ :  $d(p_1, p_2) \leq \epsilon$ ,  $d(p_2, p_3) \leq \epsilon$ ,  $d(p_1, p_6) \leq \epsilon$ ,  $d(p_1, p_3) \leq \epsilon$ .
- If  $\epsilon = 100$ : all points are within  $\epsilon$  of each other.

Table 1: Measurements and metrics for the subquestions in Question 1. All response times have been averaged over 3 time trials, each using separate executions of the program. The column “No Opt.” refers to executing the program without compiler optimization. Blank lines indicate values that you need to add to the table.

Subquestion ( $\epsilon = \text{-----}$ )	#Cores	Time (s) No Opt.	Time (s) w/ -O3	Speedup	Parallel Efficiency
Sequential	1	-----		1	1
(a)	-----	-----		-----	-----
Sequential	1		-----	1	1
(b)	-----		-----	-----	-----

## Details

The Euclidean distance is reflexive, i.e.,  $d(p_1, p_2) = d(p_2, p_1)$ . You must “double count” these pairs. That is, if  $p_1$  and  $p_2$  are within  $\epsilon$  then  $p_2$  and  $p_1$  are also within  $\epsilon$ , and both must be counted towards the total number of points within  $\epsilon$  of each other. You must also add to the total count the distance between a point and itself. E.g.,  $d(p_1, p_1) = 0$  will always be within  $\epsilon$ . These requirements make the total count easier to compute, as these are considered corner cases.

If  $\epsilon = 100$  in Figure 1, then all points are within  $\epsilon$  of each other. Including the “double counting” above, and a point counting itself, this means that the total number of counts you should compute are:  $6^2 = 36$ .

If  $\epsilon = 3$  (see example above), then the total count of points within  $\epsilon$  of each other should be 14. All points are within  $\epsilon$  of themselves (6), and the 4 pairs of points within  $\epsilon$  above yield  $4 \times 2 = 8$ .

## Program Guidelines

Your program will take as input on the command line a value for  $\epsilon$  (declared as a double precision float). Your program will output the total number of point comparisons that are within  $\epsilon$ . I will provide you with starter code that passes  $\epsilon$  in on the command line.

## Question 1: Baseline with OpenMP [5 Points]

Write the brute force sequential algorithm to compute the total number of points that are within  $\epsilon$  according to the guidelines above. This algorithm compares all points to each other, calculating the distance and determining if they are within  $\epsilon$ . Hint: the algorithm is  $O(N^2)$ , where  $N$  is the input size (total number of data points). Parallelize the sequential code with OpenMP. Compare the execution time of the sequential vs. parallel algorithm. For all items below, include the number of seconds it took to run the program. Report an average of 3 time measurements.

- Without compiler optimization, report the number of cores, speedup, and parallel efficiency achieved.
- With compiler optimization (-O3), report the number of cores, speedup, and parallel efficiency achieved.
- For the items above, reason about your performance gains. Is your speedup or parallel efficiency good or bad? Explain.

Example compilation using optimization: `gcc -O3 main.c -o main`

Example of running the program with the  $\epsilon = 10$  as a parameter: `./main 10.0`

### Submission:

I provide you with `point_epsilon_starter.c`. This code generates the data, gets  $\epsilon$  from the command line, and shows you where to time your code. Download it, rename it `question1_NAME.c`, complete the question, and turn in an archive with this file. Make sure that you include responses to Questions (a)–(c) above in your pdf writeup. In addition, fill out a table that looks like that shown in Table 1.

## Question 2: Algorithmic Transformation [10 Points]

The above algorithm is  $O(N^2)$ , which is not very efficient. For instance, in Figure 1, if  $\epsilon = 3$ , then we know visually that there is no reason to compare  $p_3$  and  $p_4$ , because they are clearly far away from each other. Modify your brute force  $O(N^2)$  algorithm to be more efficient (i.e., fewer total distance calculations, fewer floating point operations, or other ways of reducing the computational load of comparing points). This will require an *algorithmic transformation*, meaning that you cannot use the brute force algorithm in Question 1. You will need to be *creative* to reduce the total computational load. Use OpenMP to parallelize the code. Remember that the algorithm should still give the correct result for a given value of  $\epsilon$ .

Report the following information:

- A description of how your new algorithm works. That is, how does it reduce the computational load in comparison to the  $O(N^2)$  algorithm? What overheads does it have? Overheads here refer to the “extra steps” that are needed to: (a) make the program utilize threads; and (b) reduce the amount of computation. Provide illustrative examples of how your new algorithm works, as appropriate.
- Without compiler optimization, report the number of seconds it took to run the program (averaged over 3 trials), the speedup you achieved, the parallel efficiency, and the number of cores.
- With compiler optimization (-O3), report the number of seconds it took to run the program (averaged over 3 trials), the speedup you achieved, the parallel efficiency, and the number of cores.

Compare the performance between your new algorithm and the parallel algorithm in Question 1 (executed using no compiler optimization and with -O3). Report the following information:

- Compare performance by reporting the ratio of the response times over the brute force implementation in Question 1. For example, you execute both algorithms in parallel, and obtain: T1- brute force response time, and T2- your new algorithm response time. Report the T1/T2 ratios.
- Answer these questions: Is your new algorithm faster? Why or why not? What optimizations worked well? What ideas did you try that did not perform well?

You must provide a detailed performance comparison. Parallel speedup gives a very coarse-grained overview of performance. But it does not actually tell us what makes the parallel program faster. Hint: you should collect these metrics separately from timing your program because collecting performance metric data may slow down your optimized program!

- Give a more detailed comparison of the performance of your algorithm in comparison to the brute force algorithm. Use other metrics to convey why your program is now faster. Examples include: reduction in floating point operations, reduction in the number of points compared, counting cache misses, and others.
- Do these metrics translate linearly to the observed reduction in response time as compared to the brute force algorithm? Why or why not?

### Submission:

Use the same starter code as above as a starting point. Download it, rename it `question2_NAME.c`, complete the question, and turn in an archive with this file. If you have  $p$  versions of the code because you worked in a group, submit  $p$  versions of the code, with different file names indicating the version number. Make sure that you include responses to Questions (a)–(g) above in your pdf writeup. In addition, fill out a table that looks like that shown in Table 2.

### Optional Group Assignment Instructions

Your group should answer  $p$  different versions of Question 2, where  $p$  is the number of people in your group. Report the following information:

- Compare the performance of your  $p$  versions of Question 2. Which algorithmic transformation worked best? Why do you think that is?

Table 2: Measurements and metrics for the subquestions in Question 2. All response times have been averaged over 3 time trials, each using separate executions of the program. “No Opt.” refers to executing the program without compiler optimization. Blank lines indicate values that you need to add to the table.

Subquestion ( $\epsilon = \text{-----}$ )	#Cores	Time (s) No Opt.	Time (s) w/ -O3	Speedup	Parallel Efficiency	T1/T2 No Opt.	T1/T2 w/ -O3
Sequential	1	-----		1	1		
(b)	-----	-----		-----	-----		
Sequential	1		-----	1	1		
(c)	-----		-----	-----	-----		
(d)						-----	-----

## Validation of all Questions

The starter program randomly generates the data to be analyzed. Do not change the seed for the random number generator. I have selected the number of points ( $N$  in the program) to be 100,000. The sequential brute force algorithm on my laptop without compiler optimization takes roughly 100 s to execute. This gives you lots of room for optimizing Question 2. It’s enough work such that the effects of *efficient* algorithmic changes should be noticeable. Since  $N = 100000$  means that you will need to wait roughly 100 s for each execution, you may want to change  $N$  for testing purposes. Below are some values of  $N$  and  $\epsilon$  and the corresponding total number of points within  $\epsilon$ . Note that I have found two different total count values which may vary slightly depending on your math library.

- You will be graded on these values of  $N$ :
- $N = 100000$ ,  $\epsilon = 10.0$ , total count: 3211140 (Ubuntu), 3213704 (OSX).
- $N = 100000$ ,  $\epsilon = 5.0$ , total count: 879688 (Ubuntu), 881058 (OSX).
- These values of  $N$  are sample ones that can be used for testing:
- $N = 1000$ ,  $\epsilon = 10.0$ , total count: 1356 (Ubuntu), 1310 (OSX).
- $N = 1000$ ,  $\epsilon = 5.0$ , total count: 1102 (Ubuntu), 1092 (OSX).

Upon submitting your programs, they should be configured to use  $N = 100000$ . Supply screenshots of your program outputting the correct values for the total number of points within  $\epsilon$  as shown above for  $N = 100000$ .

## Grading

Because you are implementing an algorithmic transformation which requires thinking about different ways to implement the problem, there are a wide range of possible solutions. This assignment will be graded based on the level of effort you put into your implementation, and the amount of thought you’ve given to effectively computing the distance between points.

The brute force algorithm is  $O(N^2)$ . However, there are solutions that can compute the distances between all points within  $O(N \log N)$ . If we assume that the brute force algorithm takes 100 seconds to compute the solution sequentially for  $N = 100000$ , then the time to perform a single distance calculation between two points is as follows:  $100 \text{ s} / (100000^2) = 10^{-8} \text{ s}$ . If the  $O(N \log N)$  algorithm were to have no overhead, then the lower bound response time would be  $100000 \log_2(100000) \cdot 10^{-8} \text{ s} = 0.0166 \text{ s}$ . This is a large reduction in the response time; however, the  $O(N \log N)$  algorithm is expected to have overhead, and you will not be able to compute the distance between all points in only 0.0166 s. Since it is possible to achieve a large reduction in the response time, there are many optimizations that can be used to reduce the time needed to compute

the distance between the points. Keep in mind that the above “back-of-the-envelope” calculation is for the sequential algorithm, and parallelizing the algorithm can achieve further performance gains. Also, note that I do not expect you to come up with an  $O(N\log N)$  solution to the problem!

Here are a few general grading guidelines/scenarios for Question 2 regarding your transformation of the brute force algorithm.

- To achieve an A: In your sequential implementation, you should have a substantial reduction in the amount of work as compared to the brute force approach. This could be a reduction in floating point operations, or number of distance comparisons (as described in Question 2). You must parallelize the algorithm as appropriate. For instance, if you have 4 cores and 4 threads, you shouldn't have 3 threads idle while one thread is performing the computation.
- To achieve a B: In your sequential implementation, you should have a moderate reduction in the amount of work as compared to the brute force approach. Like the criterion above, you must parallelize the algorithm as appropriate to your algorithm.
- To achieve a C: In your sequential implementation, your solution to reducing the number of distance calculations yields a modest reduction in the amount of work as compared to the brute force approach. The parallelization of the algorithm may not be efficient (e.g., may have too much overhead).

## Bonus 1: Beat the Professor [2 Points]

If you do not want your assignment to be considered for these bonus points, do not respond to this section in your write up.

The professor's brute force implementation as executed on a dedicated machine with 1 thread/core (using -O3) executes in 25.4 s for both  $\epsilon = 5.0$  and  $\epsilon = 10.0$ . The optimized version executes in 0.43 s ( $\epsilon = 5.0$ ) and 0.76 s ( $\epsilon = 10.0$ ). The ratio of the brute force to optimized algorithm response time is thus 59.1 and 33.4, for  $\epsilon = 5.0$  and  $\epsilon = 10.0$ , respectively.

The sequential time to execute the programs with compiler optimization (-O3) has been obtained for Questions 1 and 2. Respond to this question with your response times and ratios for  $\epsilon \in \{5, 10\}$ . I will run your program on the same machine to ensure consistency. Furthermore, I will use my brute force algorithm to make the calculation. If you beat the professor's ratio, you will get two bonus points.

Table 3: Response time ratios for the bonus question. All response times have been averaged over 3 time trials, each using separate executions of the program. The column “No Opt.” refers to executing the program without compiler optimization. Blank lines indicate values that you need to add to the table.

	T1/T2 (-O3) $\epsilon = 5.0$	T1/T2 (-O3) $\epsilon = 10.0$
Sequential:	-----	-----

## Bonus 2: Competition [5 Points]

If you do not want your assignment to be considered for these bonus points, do not respond to this section in your write up.

Bonus points will be awarded for the top three (CS550) and five (CS450) fastest algorithms, as reported in this section of the assignment as the ratio of the brute force to optimized algorithm response time using -O3 and executed sequentially. The algorithms are executed sequentially as students will have varying numbers of cores in their computers.

I reserve the right to retract these bonus points if the algorithms are all at a B level (described in the grading section). I will execute all programs on the same platform to ensure a fair comparison between approaches. Be sure to complete Table 3.

You are not eligible for the bonus points for beating the professor if you are awarded bonus points for this question.