

# CS 6290: High-Performance Computer Architecture

## Project 3

This project is intended to help you understand cache coherence and performance of multi-core processors. As with previous projects, for this project you will need VirtualBox and our project virtual machine. Just like in previous projects, you will put your answers in the reddish boxes in this Word document, and then submit it in T-Square (but this time the submitted file name should be PRJ3.docx).

In each answer box, you **must first provide your answer to the actual question** (e.g. a number). You can then use square brackets to provide any explanations that the question is not asking for but that you feel might help us grade your answer. E.g. answer 9.7102 may be entered as **9.7102 [Because  $9.71 + 0.0002$  is 9.7102]**. For questions that are asking “why” and/or “explain”, the correct answer is one that concisely states the cause for what the question is describing, and also states what evidence you have for that. Guesswork, even when entirely correct, will only yield 50% of the points on such questions.

Additional files to upload are specified in each part of this document. **Do not archive** (zip, rar, or anything else) the files when you submit them, except when we explicitly ask you to submit a zip file (in Part 3H). Each file we are asking for should be uploaded separately, and its name should be as specified in this assignment. You will lose up to 20 points for not following the file submission and naming guidelines. Furthermore, if it is not VERY clear which submitted file matches which requested file, we will treat the submission as missing that file. The same is true if you submit multiple files that appear to match the same requested file (e.g. several files with the same name). In short, if there is any ambiguity about which submitted file(s) should be used for grading, the grading will be done as if those ambiguous files were not submitted at all.

Most numerical answers should have **at least two decimals** of precision. Speedups should be computed to **at least 4 decimals** of precision, using the number of cycles, not the IPC (the IPC reported by report.pl is rounded to only two decimals). You lose points if you round to fewer decimals than required, or if you truncate digits instead of correctly rounding (e.g. a speedup of 3.141592 rounded to four decimals is 3.1416, not as 3.1415).

**This project can be done either individually or in groups of two students.** If doing this project as a two-student group, you can do the simulations and programming work together, but each student is responsible for his/her own project report, and each student will be graded based solely on what that student submits. Finally, **no collaboration with other students or anyone else is allowed**. If you do have a partner you **have to** provide his/her name here  (enter None here if no partner) and his/her T-Square username here . Note that this means that you cannot change partners once you begin working on the project, i.e. if you do any work with a partner you cannot “drop” your partner and submit the project as your own (or start working with someone else) because the collaboration you already had with your (original) partner then becomes unauthorized collaboration.

## Part 1 [40 points]: Running a parallel application

In this part of Project 3 we will be using the LU benchmark. We will also be using a processor with more (sixteen) cores (cmp16-noc.conf). So, for example, to simulate 4-threaded execution you would use a command like this (note the absence of spaces between -n and 512, and between -p and 4):

```
~/sesc/sesc.opt      -fAp4      -c      ~/sesc/confs/cmp16-noc.conf      -olu.out  
-elu.err lu.mipseb -n512 -p4
```

To complete this part of the project, run the lu application with 1, 4, and 16 threads. Then fill in the blanks, taking into account all the runs you were asked to do:

- A) Submit the three simulation reports: sesc\_lu.mipseb.Ap1, sesc\_lu.mipseb.Ap4, and sesc\_lu.mipseb.Ap16. You will not earn points for submitting these simulation reports, but you will lose 10 points for each missing simulation report.
- B) Fill out the execution time, parallel speedup, and parallel efficiency with 4 and 16 threads. Enter Sim Time with precision of at least three decimals, and speedup and efficiency with precision of at least two decimals.

	SimTime (in ms)	Parallel Speedup	Parallel Efficiency
-p1	<input type="text"/> ms	<input type="text"/>	
-p4	<input type="text"/> ms	<input type="text"/>	
-p16	<input type="text"/> ms	<input type="text"/>	<input type="text"/>



Note: Parallel speedup is the speedup of parallel execution over the single-thread execution with the same input size. Parallel efficiency is the parallel speedup divided by the number of threads used – ideally, the speedup would be equal to the number of threads used, so the efficiency would be 1. When computing the speedup and efficiency we cannot use IPC or Cycles that are reported for each processor, because these do not account for the cycles where that core was idle (e.g. because the thread was waiting for something to happen). So we need to use the “Sim Time” we get from report.pl because it accounts for all cycles that elapse between the start and completion of the entire benchmark.

- C) Our results indicate that parallel efficiency gets lower as we use more cores. Why do you think this is happening?
- D) When we use four threads (-p4) instead of one (-p1), the IPC achieved by Core 0 (the first processor listed) got slightly lower because
- E) Now look at the simulation reports for these simulations. Core 0 executes more than its fair share of all instructions because

## Part 2 [10 points]: Cache miss behavior

In this part of Project 3, we will be focusing on the number of read misses in the DL1 (Data L1) cache of Core 0, using the same simulations that we already did for Part 1. In the report file generated by the simulator (sesc\_lu.mipseb.something, not what you get from report.pl), the number of cache read misses that occur in each DL1 cache (one per processor core) is reported in lines that begin with “P(0)\_DL1:readMiss=”.

- F) The total number of read misses that occur in the DL1 cache of Core 0 is

Simulation	-p1	-p4	-p16
Core 0's DL1 read misses			

Your answers here should be integer numbers.

- G) The number of these misses changes this way as we go from one to two to four, etc. threads because

## Part 3 [50 points]: Identifying accesses to shared data

Your task in this part of the project is to determine how many read misses in each core's DL1 cache are compulsory (readCompMiss), replacement (capacity or conflict, the counter should be called readReplMiss), and coherence misses (readCoheMiss), and separately also classify write misses (writeCompMiss, writeReplMiss, and writeCoheMiss). Note that this classification is similar to the one you did for Project 2, except that you now we are counting different categories of misses separately for reads (load instructions) and writes (store instructions), that we are placing conflict and capacity misses in the same (replacement) category, and that we are adding a category for coherence misses that we didn't have in Project 2. To simplify classification, we will **not** follow the exact definition of coherence misses (“those misses that would have been hits were it not for coherence actions from other cores”). Instead, we will use a definition that allows much simpler implementation: a coherence miss is a miss that finds in the cache a line whose tag matches the block it wants, but that block has a coherence state that prevents such access. In the case of read misses, this means that the line was found in an “Invalid” coherence state. Note that this identification of coherence misses may not be trivial in the SESC simulator because of the way it handles tags during invalidation. If a miss is not a coherence miss, then you can classify it as either compulsory or replacement miss by checking if the block was ever in that cache. When checking whether the miss is a compulsory miss, be careful to track the “was previously in this cache” set of blocks for each cache separately.

- H) Create a Changed.zip file with any simulator source code files that you have modified in Part 3 of the project, and submit this Changed.zip file together with your project. You will not earn points for submitting this file, but you will lose 50 points if it is missing or if it does not contain all the source code modifications. Then, with your changed simulator (that now counts compulsory, replacement, and coherence read misses), re-run the simulations from Part 1 and submit the resulting simulation report files as sesc\_lu.mipseb.Hp1, sesc\_lu.mipseb.Hp4, and

sesc\_lu.mipseb.Hp16. As in Part 1, you will not earn points for these submitted simulation reports, but you will lose 10 points for each simulation that is missing.

- I) The number of all read misses, compulsory read misses, replacement read misses, and coherence read misses for the DL1 cache of Core 0 is:

	<b>Core 0's DL1 readMiss</b>	<b>Core 0's DL1 compMiss</b>	<b>Core 0's DL1 replMiss</b>	<b>Core 0's DL1 coheMiss</b>
<b>-p1</b>				
<b>-p4</b>				
<b>-p16</b>				

Note: readMiss numbers here should be the same as those you had in Part 2.

- J) The number of all write misses, compulsory write misses, replacement write misses, and coherence write misses for the DL1 cache of Core 0 is:

	<b>Core 0's DL1 writeMiss</b>	<b>Core 0's DL1 compMiss</b>	<b>Core 0's DL1 replMiss</b>	<b>Core 0's DL1 coheMiss</b>
<b>-p1</b>				
<b>-p4</b>				
<b>-p16</b>				