

# Practical 2: Cache Coherence Protocols

CS4/MSc Parallel Architectures

Issued: Tuesday, February 3, 2015
Due: Friday, February 27, 2015 at 4:00 PM

## 1 Introduction

This assignment is the second practical of the Parallel Architectures module for CS4/MSc students. This practical will contribute 20% of the final mark for this course.

It consists of a programming exercise culminating in a brief written report. Assessment of this practical will be based on the correctness and the clarity of the solution. This practical is to be solved individually to assess your competence on the subject. Please bear in mind the School of Informatics guidelines on plagiarism. You must return your solutions to the ITO before the due date shown above.

## 2 Overview of the Practical

The goal of this practical is to develop a cache coherence protocol simulator and evaluate it on 2 memory access traces. The trace files will be given to you and their format is described in §3.2. To accomplish this goal, you will need to undertake the following subtasks, which will be marked separately; however, bearing in mind, that 1-3 can only be awarded fully if these subtasks are documented sufficiently in the report.

1. **Implementation of the simulator (30 marks).** In order to simulate coherent caches, first you will have to write a simple cache simulator. For simplicity, assume a direct mapped cache with write back policy for your experiments. The size of the lines and the number of lines in the cache, however, should be parametrise and your final experiments should be performed for different values of such parameters. Please refer to [1] or the CS3 Computer Architecture notes for details on how addresses (word addresses in the case of this practical) are broken into index and tag. Note that the cache simulator code should be modular as you will need multiple instances of it, one for each processor in the system.

On top of this very simple cache simulator you will have to build a simple coherence protocol simulator. For simplicity, you can assume a system with only 4 processors and with a single level of cache per processor (all given trace files will have only 4 processors). At a minimum, you will have to implement a simple MSI bus based

snooping coherence protocol (*20 marks*), but marks (*10 marks*) will be given for the addition of a MESI protocol implementation. Again, for simplicity you can assume that all transactions on the bus are instantaneous and you *do not* have to model transient states nor split transactions.

Additionally, you will have to augment your simulator with mechanisms to collect key statistics, such as miss rate<sup>1</sup>, number of invalidations<sup>2</sup>, and anything else that will help you perform a quantitative analysis of the results you obtain from your experiments (see below).

2. **Validation of the simulator (*10 marks*).** In order to have high confidence in the validity of your results, you should provide a clearly written description of the structure and operation of your simulator and a set of small, artificial examples which use the various output options to demonstrate that your simulator is working correctly. The examples you choose are entirely up to you, but you should make sure you cover important cases, such as reads and writes to a cache-line/block that is in states **Modified** or **Shared** in another cache.
3. **Design and execution of experiments (*30 marks*).** The aim of the experiments with the given traces is to investigate the behaviour of both the caches (miss rate, etc.) and programs. Obtain statistics regarding sharing pattern of each cache-line/block, the type of permissions required for each memory access, miss rate for different types of accesses etc. Think of other statistics that will provide more information about the cache usage and the access patterns of the parallel application. Try and answer these questions:
  - a) Vary the size of the cache and cache-line/block size to iterate over the above experiments. How do differing cache-line/block sizes and the number of cache-lines affect miss/hit rates and invalidations?
  - b) In terms of memory accesses, what is the distribution of accesses to private data (hits in state **Modified**, and also **Exclusive** for MESI), shared data (hits in **Shared**).
  - c) If you manage to implement a MESI protocol, then you should also perform experiments to compare the two protocols and to identify the main advantages of it, if any.
4. **Report and discussion (*30 marks*).** You must report and discuss the results of your experiments. The written report should be clear and concise, and you should present your results in a suitable form (e.g., with plots). You should carefully plot the results to highlight the impact of various parameters. At the end of the report you should provide a summary of the results along with your inferences and conclusions.

*You are strongly advised to work in stages and iteratively. For instance, try to validate each new feature of the simulator and describe it in the report as you develop it. Also, try to perform the main experiments in stages with a small set of parameters at a time and alternating with the writing of the report.*

---

<sup>1</sup>For consistency purposes assume that a tag match on a line cached in shared state upon a write also counts as a “miss”.

<sup>2</sup>You should count both the number of invalidation broadcasts and the number of lines invalidated at each broadcast.

## 3 Baseline specification

Any assumptions you make must be stated clearly in your report. Below are the given parameters of the system you must implement.

### 3.1 Architecture

Our imaginary architecture has 4 processors and a very simple memory and bus model. Memory is addressed at the word level and data addresses start from 0. There is no virtual memory. Thus, address 0 indexes the first word in memory, address 1 the second word, and so on. The cache-line/block size will correspond to some multiple of 2 number of words, which is a parameter that should be parametrised in your simulator in a suitable manner (e.g., command line argument). The caches are direct mapped with a write back policy.

In the default configuration, private cache of each processor has 1024 cache-lines and the cache-line/block size is 4 words. In the report, results for all other configurations should be analysed relative to the default configuration.

The main point of this simple model is that your simulator does not have to concern itself with byte addressing within words, word alignment and so on. The bus is non-split-transaction and you can assume that transactions are instantaneous. You can also assume that all state changes in the caches are instantaneous.

### 3.2 Trace File Format

Your simulator *must* accept its input from trace files which have the format described here. Every line in the trace describes either a memory operation or a command for your simulator to produce some output.

**Memory operations:** A memory access line in the trace consists of the processor identifier (e.g., P1), followed by a space, followed by the type of operation (R or W, for read or write, respectively), followed by another space, followed by the word address, and ending with a newline. Thus for instance, the line:

```
P2 R 12
```

indicates a read by processor 2 to word 12. Memory operations must appear in the bus in the order specified in the trace. The actual timing of the operations is not relevant and is not taken into account.

*Note that neither the actual data value being written nor the register being used are specified. Thus, your simulator cannot keep track of real data. All that matters in calculating hit rate and invalidations is the sequence of addresses used.*

**Output commands:** An output command line in the trace consists of one of the following:

- **v:** indicates that full line-by-line explanation should be toggled, i.e. switched on (if it is currently off) or off (if it is currently on). The default is that it is off.
- **p:** indicates that the complete content of the cache should be output in some suitable format (see below).

- **h**: indicates that the hit-rate achieved so far should be output.
- **i**: indicates the that total number of invalidations seen so far should be output.

You are free to choose the actual wording of the explanations and other output. As a suggestion you can use: “A read by processor P2 to word 17 looked for tag 0 in cache-line/block 2, was found in state Invalid (cache miss) in this cache and found in state Modified in the cache of P1.”

For example, the meaning of the following trace file is shown in parenthesis, to the right of the commands (these comments are not actually present in the trace):

```
v          (switch on line by line explanation)
P0 R 17    (a read by processor 0 to word 17)
P1 R 18    (a read by processor 1 to word 18)
P2 W 17    (a write by processor 2 to word 17)
v          (switch off line by line explanation)
P2 R 25    (a read by processor 2 to word 25)
P0 R 35    (a read by processor 0 to word 35)
p          (print out the cache contents)
h          (print out the hit rate)
i          (print out the number of invalidations)
```

## 4 Format of your submission

Your submission should clearly indicate which parts of the coherence protocol you have simulated completely and which you have only partially completed. Submit an archive of your simulator *and* a soft copy of your report as follows, using the DICE environment:

```
$ submit pa 2 <cache-simulator.tar.gz> <report.pdf>
```

Ensure that your source code is well documented. Your report should be self-sufficient, but also ensure your simulator source code is well documented. The report should contain the following:

1. Written description of the internal structure and workings of your simulator (*1-2 pages, single spaced*, depending upon complexity), explaining clearly which parts are complete and which incomplete.
2. Extracts of simple trace files used for validation and a description of expected and actual results on these, in order to provide evidence that your simulations are accurate (a few pages including examples).
3. A report detailing the experiments you ran and giving a critical summary of the results (*2-4 pages, single spaced*, including results). You should provide enough information to make the experiments repeatable. Present your data clearly and concisely. Also present any inferences and conclusions you have obtained from the results of these experiments.

## 5 Reporting Problems

Send an email to [arpit.joshi@ed.ac.uk](mailto:arpit.joshi@ed.ac.uk) for any issues regarding the assignment.

## References

- [1] J. L. Hennessy and D. A. Patterson. *Computer Architecture - A Quantitative Approach* (5. ed.). Morgan Kaufmann, 2011.