

CSE471 Homework Assignment #2: Cache Coherence

Due Tuesday, May 17th, before lecture

Milestone 3: Evaluating the 4-State Protocol

In the second milestone, you implemented your 4-state protocol in the MultiCacheSim simulator. In this milestone, you will be doing a quantitative analysis of the 4-state protocol. You will then write a report summarizing your work, and the conclusions you were able to draw from your quantitative analysis.

Experiments

To better understand the impact of the 4-state protocol, you will conduct a series of experiments.

First, you need to run each benchmark in simulation, using PARSEC's `simdev` input (using the `-i simdev` option for `parsecmgmt`). During development, you were using the `test` input, which is only good for making sure your code runs correctly, but isn't good for experiments. The `simdev` input is larger, and will give you a better approximation of real system behavior.

When you run your simulations, you should run your 4-state protocol implementation alongside the provided 3-state protocol implementation. By doing so, you will get numbers that came from the same execution trace, and can, therefore, be compared directly. Recall that you can specify a comma-separated list of protocol plugins to MultiCacheSim to simulate multiple protocols on the same execution.

Part of your analysis will be to evaluate the relative scalability of the 4-state protocol to the 3-state protocol. In order to do this analysis, you will have to run your experiments three times – once with 2 threads, once with 4 threads, and once with 8 threads. Take note to change the thread count option that you are providing to `parsecmgmt` (e.g., `-n 8`) and the one you are providing to MultiCacheSim (e.g., `-nCaches 8`).

One goal of this assignment is to give you experience dealing with the heaps of semi-structured data produced by architectural simulators. It is up to you to keep your data organized. Keeping things neat will help you stay focused on the interesting parts of the experiments.

The Analysis

After completing your experiments, you will analyze the data collected from these experiments.

The first step to your analysis should be to identify the simulation outputs that vary between the 3-state and 4-state protocols. You should try to be as inclusive as possible during this step – don't limit yourself to analyzing only the outputs that you expected to vary. It will probably be helpful during this step to focus on a subset of your data. For example, to try to decide what the variables of interest are, you might want to look at only one or two benchmarks at first. You also probably want to restrict your initial attention to a single thread-count. Plotting the data you've collected will probably help you to see trends in your data.

Once you have decided which are your variables of interest, you should figure out why you are seeing the difference you are seeing. Be sure not to jump to conclusions – if you see something that doesn't make sense, given the protocol and the data, don't try to manufacture an explanation – there might still be a bug lurking in your code.

Next, you should look at how the differences in these variables of interest change as you vary the number of threads in the simulation. For example, does the 4-state protocol lead to a *bigger* difference in some output

value when there are more threads? What conclusions can you draw regarding the performance scalability of these two protocols from the data you've collected?

The Report

You should write up a report that summarizes the work you've done, and the conclusions you are able to draw from your results. Your report should include the following:

- A brief introduction and problem statement
- A description of your 4-state protocol design, and the shortcomings of the 3-state design that it addresses. This section can be kept brief, as your milestone 1 report should have elaborated on your design decisions in detail.
- A description of your evaluation. You should describe what you intend to show with your experiments, describe what experiments you performed, and discuss your experimental methodology.
- A discussion of your results. This should be a write up of the quantitative analysis. You should include at least one plot, showing at least one simulator output that varies between protocols. The plot should show how that output varies between protocols with two, four, and eight threads. You should discuss the implications of the 4-state protocol on the performance scalability of systems. You should also use this section to describe any other interesting properties of your data, such as surprising differences between the protocols, benchmark-specific characterization (if you decide to do any), etc.

The report is a large fraction of your grade for this milestone. Please keep that in mind when you are allocating your time. It is important that you communicate clearly – converting your analysis from raw data to English prose should be a top-level concern when you are budgeting your time, not just an afterthought.