

Assignment 2: Understanding Data Cache Prefetching

Computer Architecture

Due: Thursday, March 22, 2018 at 4:00 PM

This assignment represents the second practical component of the Computer Architecture module. This practical contributes 12.5% of the overall mark for the module. It consists of a programming exercise culminating in a written report. Assessment of this assignment will be based on the correctness of the code and the clarity of the report as explained below. The practical is to be solved individually. Please bear in mind the School of Informatics guidelines on Academic Misconduct. You must submit your solution (code and a report) by the due date shown above. Follow the instructions provided in Section 6 for submission details.

In this assignment, you are required to explore data cache prefetching techniques using the Intel Pin simulation tool. You are strongly advised to commence working on the simulator as soon as possible.

1 Overview

In this assignment you will evaluate the hit ratio of a data cache when different hardware prefetch mechanisms are used. Towards this end, you will compare the following data prefetch mechanisms:

1. *Next-N-Line Prefetcher* – provided to you
2. *Stride Prefetcher* – for you to implement
3. *Distance Prefetcher* – for you to implement

2 Prefetching

Cache prefetching is a technique that reduces cache miss rate by fetching data from memory to a cache, ideally before the data has been demanded from the processor.

2.1 Sequential Prefetcher

The simplest hardware prefetcher is a Next-N-Line Prefetcher, which brings one or several cache blocks adjacent to the one that was not found in a cache. If the next block(s)

are already in a cache, they are not prefetched. The number of next blocks to prefetch (N) is called *aggressiveness*. The aggressiveness of the Next-N-Line Prefetcher can vary across different implementations. The simplest Next-N-Line Prefetcher requests just the next block ($N=1$). However, a more aggressive prefetcher can prefetch more next blocks.

2.2 Stride Prefetcher

Stride prefetching is a more advanced hardware prefetching technique, which is particularly beneficial for array traversals. A *strided access* with stride M means that every M th cache block is touched. Once a strided access pattern is detected, upon a load miss, the prefetcher fetches one or more blocks according to the pattern. Similarly to the Next-N-Line prefetcher, the number of blocks prefetched is called *aggressiveness*. If the block(s) are already in the cache, they are not prefetched.

The difference between addresses referenced by the same load instruction when it misses in the cache is called a *stride*. If a stride is the same for several consecutive misses of a load instruction, then the load access pattern is detected, and prefetching is enabled for this load instruction.

The information, required for detecting a strided access pattern is accumulated in a *Reference Prediction Table* (RPT). RPT entries are indexed by the address (i.e. Program Counter – PC) of a load instruction. An RPT entry contains:

- The PC of the load instruction, which acts as a tag.
- The previous address accessed by this load instruction on a cache miss.
- The stride for those entries which have established a pattern.
- A two-bit state.

The RPT structure is shown in Figure 1.

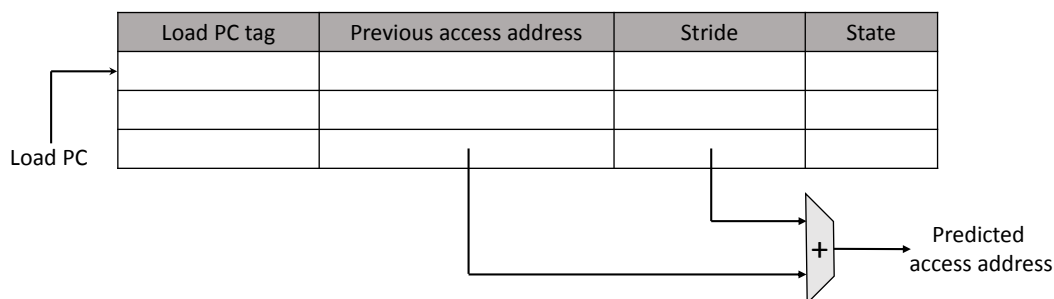


Figure 1: RPT structure.

The state diagram for an RPT entry is given in Figure 2. The four states, captured by the 2-bit state field, are:

- *Initial*: set on allocation of an entry in the RPT or after the entry experienced an incorrect prediction from steady state.

- *Transient*: corresponds to the case when the system is not sure whether it should prefetch or not.
- *Steady*: indicates that a pattern has been detected.
- *No prediction*: no pattern was detected for this entry.

An RPT entry is accessed when a load instruction experiences a cache miss. If an entry corresponding to a load instruction address is not found in the RPT, a new entry is allocated (after replacing an entry according to RPT replacement policy). The stride field of a new entry is initialized to zero. The previous address of a new entry is initialized to the memory access address. If there is a hit to RPT, the load address accessed by the instruction is predicted to be

$$access_addr = prev_access_addr + stride_value$$

The prediction is correct if the predicted address matches the access address. Note that the prediction is made only if an entry is found in the RPT. On a hit to the RPT, the state of an entry is updated depending on the prediction's correctness, according to the state transition diagram shown in Figure 2. In addition, the previous address field is updated with the memory access address. In all states except steady, in case of an incorrect prediction, a stride field is updated with the last stride. If the prediction is incorrect and the state is steady, the stride value remains unchanged. A prefetch occurs only if the corresponding RPT entry is in the state steady and the prediction is correct. In case of a prefetch, the previous address field is updated with the address of the last prefetched block.

For more information on the Stride Prefetcher, see the original Chen and Baer's paper¹.

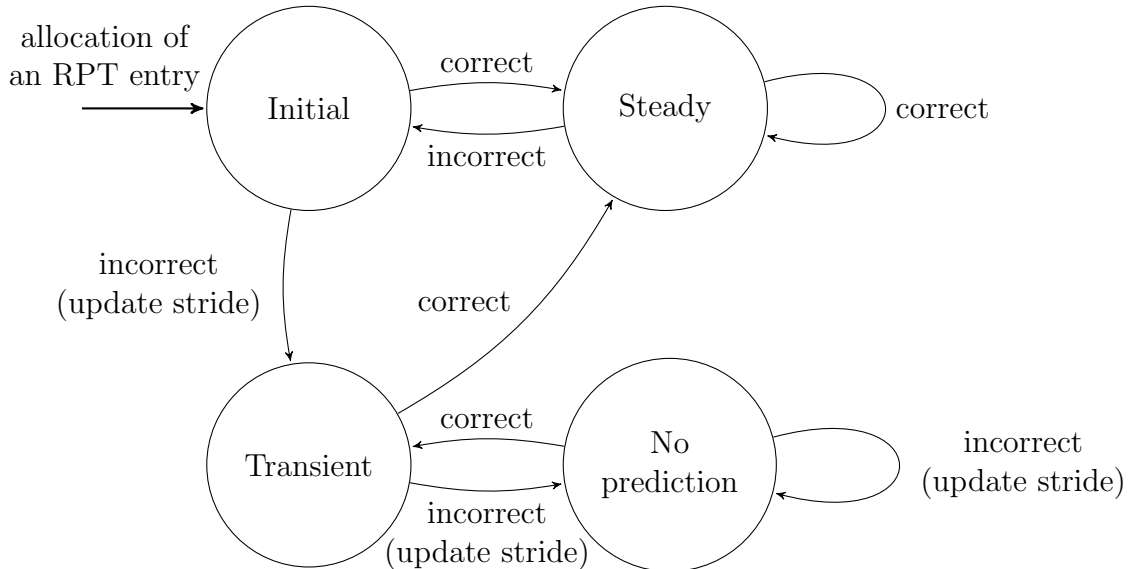


Figure 2: State transition diagram.

¹Jean-Loup Baer and Tien-Fu Chen, "An effective on-chip preloading scheme to reduce data access penalty"

2.3 Distance Prefetcher

The *distance* between two memory accesses is defined as the numerical difference between two addresses. The Distance Prefetcher (DP) observes distances between pairs of consecutive misses, stores sequences of the observed distances, and makes predictions based on these sequences. For instance, after seeing a stream of global misses to addresses 10, 11, 13, 14, it observes that a distance of 1 (11 - 10) is followed by a distance of 2 (13 - 11), and that a distance of 2 is followed by a distance of 1 (14 - 13). Based on the first observation, the last miss (i.e. on address 14) will trigger a prefetch for address 16. Note that distances can take both positive and negative values.

As shown in Figure 3, the DP stores the following state:

- The previous miss address
- The previous distance
- An RPT with entries that consist of:
 - A tag of the entry, indicating the distance to which the entry refers to
 - A number of predicted distances that have been observed to follow the distance that the entry refers to. The number of the stored predicted distances determines the aggressiveness of the DP.

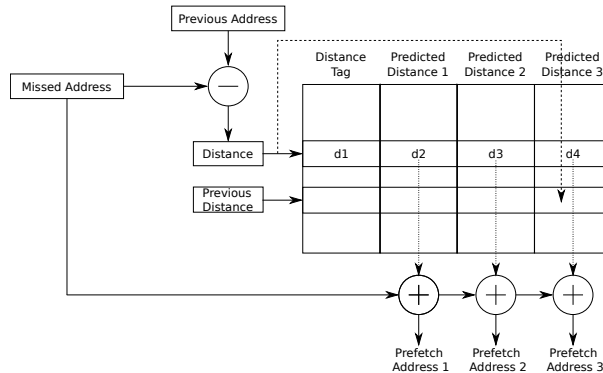


Figure 3: The Distance Prefetcher with aggressiveness equal to three

When the DP receives a new miss address, it compares it to the previous miss address to calculate the distance. Then, it searches the RPT for an entry that corresponds to that distance. If such an entry exists, then DP calculates the prefetch addresses by adding the predicted distances to the miss address and issues prefetches for all calculated addresses. Alternatively, if no entry corresponds to the distance, then an entry is allocated for the distance; if no unused entry exists, then an entry is evicted according to the *RPT replacement policy*.

The DP must be trained with the newly observed distance. As discussed, each entry has a number of slots that contain the predicted distances. In order to train the DP, the newly observed distance must be added as a predicted distance to the RPT entry that refers to the previous distance; if that entry has all of its predicted distances filled, then one of the predicted distances must be evicted to make way for the new distance.

The policy that determines which slot of the entry should be evicted is called the *entry replacement policy*. Finally, the previous distance and previous address state must be updated with the newly observed distance and address, respectively.

Let's examine the DP's functionality through an example. Assume an RPT with aggressiveness equal to 3, where there is an RPT entry with a distance tag equal to 2 and predicted distances 1, 3 and 6. The previous address field stores the address 23 and the previous distance field stores distance 5. Then assume that the DP receives a miss to address 25. Firstly, it will calculate a distance of 2 and will search the RPT for an entry with tag equal to 2. When the entry is found, DP will compute the prefetch addresses by adding the predicted distances to the miss address, and thus, it will issue prefetches for addresses 26, 28 and 31. Then, DP will perform the training: it will locate the entry with tag equal to 5, that refers to the previous distance and it will add the distance 2 to one of its predicted distances. If all the predicted address slots are in use, then one of them will be evicted, subject to the entry replacement policy. Finally, the previous address will be replaced with the address 25, and the previous distance will be replaced with 2.

3 Simulator infrastructure

Pin is a dynamic binary instrumentation engine that enables the creation of dynamic analysis tools (e.g. architectural simulators). Pin intercepts program execution and provides an API that allows you to execute high-level C++ code in response to runtime events like loads, stores, and branches. A Pin tool is a program which uses the API and specifies what has to happen in response. In this assignment, you are given an example Pin tool that intercepts data memory requests (loads and stores) and simulates a data cache with a Next-N-Line prefetcher. You will extend this tool to simulate a Stride Prefetcher and a Distance Prefetcher.

The directions below apply to DICE machines.

The Pin tool and the benchmark are available in a directory that can be accessed through a DICE machine. Inside that directory there is a README text file that contains guidelines on how to run the simulator with the benchmark.

To access the guidelines within the directory, type the following commands in the terminal:

1. `CAR_ASSGN2_PATH=/afs/inf.ed.ac.uk/group/teaching/car/assignment2`
2. `cd $CAR_ASSGN2_PATH`
3. `gedit README`

The directory that contains the skeleton source code for a data cache (`dcache_for_prefetcher.hpp`) and Next-N-Line prefetcher (`prefetcher_example.cpp`) is further down that path, at:

`$CAR_ASSGN2_PATH/tools/pin-3.5-97503-gac534ca30-gcc-linux/source/tools/PrefetchExample/`

When the Pin tool runs `prefetcher_example.cpp`, it requires a benchmark program (along with its arguments) to be given as the last command line argument. The Pin tool runs the benchmark and while running it, passes all memory instructions to the simulated data cache.

A Next N-Line Prefetcher is implemented inside `prefetcher_example.cpp` in a function called `prefetchNextNLines`.

Once Pin exits, it will generate a set of statistics for the simulated data cache and prefetcher in a file. The file's name will be given as an argument. The default name is `data`. The cache simulator, in the file `prefetcher_example.cpp`, allows for 4 different command line arguments:

1. `-pref_type`: Denotes which Prefetcher should be used; the choices are: none, distance, stride and next_n_lines. For example, in order to enable the Distance Prefetcher, use this command line argument `-pref_type distance`
2. `-aggr`: Sets aggressiveness. For example, when using the Next N-Line Prefetcher, in order to prefetch next line only, use this command line argument `-aggr 1`
3. `-o`: Sets the name of the output file to be generated. To generate an output file called `MyOutput.out`, give this command line argument `-o MyOutput.out`

Note that in the given example implementation, trying to use the Distance or Stride Prefetcher will have no impact on the simulation, as these Prefetchers are not implemented. The provided code contains comments on how to use the arguments that are necessary for your implementation (i.e. `-pref_type`, `-aggr`). These arguments will be used when testing the submitted code, so they must be used as described. Although it is not required for the assignment, there are also arguments available to change the dimensions of the data cache (`-b` for block size, `-a` for associativity and `-sets` for the number of sets).

Inside the skeleton code directory² there is a folder called `benchmark` that contains a simple micro-benchmark program, `microBench.exe`, along with its source code (in c++ and x86 assembly). The given micro-benchmark streams through several arrays, exhibiting a behaviour that can be benefited by the Next-N-Line Prefetcher. The purpose of the micro-benchmark is to act as an example of the kind of micro-benchmarks you can build to test your prefetchers. Note that the micro-benchmark has been compiled with all optimizations disabled (i.e. with `-O0`) to eliminate compiler interference. You are not required to submit your results with the micro-benchmark (or any other benchmark); rather, the given micro-benchmark is meant to serve as an example for you to use. The guidelines inside the `README` text file contain more information about how exactly to run the prefetcher simulator and how to use the command line arguments.

IMPORTANT: The skeleton code `prefetcher_example.cpp` includes counters that you **MUST** increment in your code. The counters are: `prefetches`, `accesses` and `hits`. The comments inside the skeleton code describe the counters and the source code demonstrates when to increment them.

4 Prefetcher Parameters

Your task is to implement a Stride Prefetcher and a Distance Prefetcher. The skeleton code, found in `prefetcher_example.cpp`, includes comments and guidelines on how you can use the existing classes and their member functions. As an example, a Next-N-Line Prefetcher is already implemented. The parameters of the system are as follows:

- Data cache (provided - no changes need to be made):
 - size: 512B

²`$(CAR_ASSGN2_PATH)/tools/pin-3.5-97503-gac534ca30-gcc-linux/source/tools/PrefetchExample/`

- block size: 4 bytes
- associativity: 2
- replacement policy: LRU
- prefetched blocks are allocated to LRU position
- Stride Prefetcher:
 - RPT size: 64 entries
 - associativity: fully associative
 - RPT replacement policy: random
- Distance Prefetcher:
 - RPT size: 64 entries
 - entry size: equal to the aggressiveness (i.e. each RPT entry has as many predicted distances as specified by the command line argument “aggr”)
 - associativity: fully associative
 - RPT replacement policy: random
 - entry replacement policy: random

5 Marking Scheme

Correctness (80 marks). Includes code functionality and quality for the Stride and Distance Prefetchers.

Report (20 marks). You should write a report (max 1 page). In the report you should answer the following questions:

- What are the advantages and disadvantages of the Next-N-Line prefetcher?
- What are the advantages and disadvantages of the Stride prefetcher?
- What are the advantages and disadvantages of the Distance prefetcher?
- Why and how, varying aggressiveness impacts (or not) the hit ratio?
- Which predictor would you choose and why?

Note: no graphs or measurements are required for the report.

6 Submission

Your submission should clearly indicate (in both the report and the code) which prefetching techniques you have simulated completely and, in case you did not finish, which you have only partially completed. Make sure all code written by you is well-commented to make it easy for the markers to understand. Remember that your simulator will be compiled/executed on a DICE machine, so you must ensure that it works on DICE.

Submit your source code, along with a soft copy of your report as follows, using the DICE environment:

```
$ submit car cw2 prefetcher_example.cpp report.pdf
```

7 Similarity Checking and Academic Misconduct

You must submit your own work. Any code that is not written by you must be clearly identified and explained through comments at the top of your files. Failure to do so is plagiarism. Note that you are required to take reasonable measures to protect your assessed work from unauthorized access. Detailed guidelines on what constitutes plagiarism and other issues of academic misconduct can be found at:

<http://web.inf.ed.ac.uk/infweb/admin/policies/academic-misconduct>

All submitted code is checked for similarity with other submissions using software tools such as MOSS. These tools have been very effective in the past at finding similarities and are not fooled by name changes and reordering of code blocks.

8 Reporting Problems

Send an email to Antonios.Katsarakis@ed.ac.uk **and** Vasilis.Gavrielatos@ed.ac.uk if you have any issues regarding the assignment.