

ECE571: Advanced Microprocessor Design – Homework 5
Fall 2024

Due: Friday 11 October 2024, 12:00pm

Create a document that contains the data and answers described in the sections below. A .pdf or .txt file is preferred but I can accept MS Office or Libreoffice format if necessary.

1. Cache parameters

A Haswell machine has a 44-bit physical address space, 32-kB L1 data cache, 8-way set associative, 64-bytes per line, write-back.

- (a) How many bits are used to calculate the offset?
- (b) How many bits are used to calculate the line?
- (c) How many bits are used for the tag?

2. Cache problem

This question assumes a 512-byte cache, 16-bytes per line, 2-way associative, 32-bit address size. (24 bits of tag, 4 bits for line, 4 bits for offset).

For this LRU=1 (least recently used is true) means it is the oldest data in the cache, LRU=0 (least recently used is false) means it is the newest.

The cache's current contents are as follows:

line	Way 0				Way 1			
	V	D	LRU	Tag	V	D	LRU	Tag
0	1	0	1	0000 00	1	0	0	0000 08
1	1	0	0	0000 00	1	1	1	0000 0a
2	0				0			
3	0				0			
4	0				0			
5	0				0			
...								
b	0				0			
c	0				0			
d	0				0			
e	0				0			
f	0				0			

For each of the following sequence of memory accesses state whether it is a cache hit or miss. If a line is evicted due to a miss, state whether the evicted data need to be written back to memory or not.

- (a) `ldrb r0, 0xffffffff0f`
- (b) `ldrb r0, 0x0000080f`
- (c) `strb r0, 0x0000001e`
- (d) `strb r0, 0xffffffffff`
- (e) `ldrb r0, 0xffffffff1e`

3. Bzip2 cache behavior on the x86 Haswell-EP Machine

For this section, log into the Haswell-EP machine just like in previous homeworks. Run the bzip2 benchmark, recall you will use a command line something like this:

```
perf stat -e instructions:u,L1-icache-load-misses:u \  
/opt/ece571/401.bzip2/bzip2 -k -f ./input.source
```

Be sure `input.source` is in your directory, you should be getting roughly 19 billion instructions.

- (a) Measure and report the L1 instruction cache miss rate.
Use `instructions:u` and `L1-icache-load-misses:u` for the events.
- (b) Measure and report the L1 data cache load miss rate.
Use `L1-dcache-loads:u` and `L1-dcache-load-misses:u`
- (c) Measure and report the L2 cache miss rate
Use `L2_rqsts.references:u` and `L2_lines_in.all:u`
- (d) Measure and report the L3 cache miss rate
Use `cache-references:u` and `cache-misses:u`

4. equake_1 cache behavior on the x86 Haswell-EP Machine

Recall that running equake looks something like this:

```
perf stat -e instructions:u,L1-icache-load-misses:u \  
/opt/ece571/equake_1.speccomp/equake_1 < \  
/opt/ece571/equake_1.speccomp/inp.in
```

- (a) Measure and report the L1 instruction cache miss rate.
Use `instructions:u` and `L1-icache-load-misses:u` for the events.
- (b) Measure and report the L1 data cache load miss rate.
Use `L1-dcache-loads:u` and `L1-dcache-load-misses:u`
- (c) Measure and report the L2 cache miss rate
Use `L2_rqsts.references:u` and `L2_lines_in.all:u`
- (d) Measure and report the L3 cache miss rate
Use `cache-references:u` and `cache-misses:u`

5. Bzip2 cache behavior on the ARM64 Ampere Machine

Now run the bzip2 benchmark on the ARM64 ampere machine. (As with the last homework, just ssh from the haswell-ep machine, `ssh ampere`).

- (a) For L1 icache rate try measuring `L1-icache-loads:u` and `L1-icache-load-misses:u`
- (b) For L1 dcache rate try measuring `L1-dcache-loads:u` and `L1-dcache-load-misses:u`
- (c) For l2 cache rate try measuring `L2d_cache:u` and `L2d_cache_refill:u`
- (d) For L3 cache try measuring `l3c0/read-hit/` and `l3c0/read-miss/` but I don't know if that's anywhere near measuring what we think it does, I haven't had the time to research this properly.

6. Short Answer Questions

To answer these questions, it might be useful to know the cache size parameters of the various machines. See if you can find the L1/L2/L3 cache sizes for the Haswell-EP (Xeon E5-2640 v3). Also for the Ampere system which is a Lenovo HR330A with an Ampere eMAG 8180 processor.

One additional piece of information, you can use "top" while a program is running to see the memory working set size. For bzip2 this is around 11MB, where equake is using 700MB.

- (a) Why is the L1-icache miss rate so small?
- (b) How does equake's L3 cache behavior differ from bzip2's on Haswell-EP?
What might be the reason for this?
- (c) How does bzip2's cache behavior on Haswell-EP differ from bzip2's cache behavior on Ampere?
What might be the reason for this?

7. Submitting your work

- Create the document containing the data as well as answers to the questions asked.
- Please make sure your name appears in the document.
- e-mail the file to me by the homework deadline.