ECE 4750 Computer Architecture, Fall 2015
Lab 3: Blocking Cache

School of Electrical and Computer Engineering
Cornell University

revision: 2015-10-20-18-57

In this lab, you will design two finite-state-machine (FSM) cache microarchitectures, which we will eventually compose with the processor designs you developed in the previous lab. The baseline design is a direct-mapped, write-back, write-allocate cache, and the alternative design is a two-way set associative cache that should reduce the miss rate by avoiding conflict misses. You are required to implement the baseline and alternative designs, verify the designs using an effective testing strategy, and perform an evaluation comparing the two implementations. **As with all lab assignments, the majority of your grade will be determined by the lab report. You should consult the course lab assignment assessment rubric for more information about the expectations for all lab assignments and how they will be assessed.**

This lab is designed to give you experience with:

- basic memory system design;
- complex finite-state-machine cache controllers;
- microarchitectural techniques for implementing cache associativity;
- abstraction levels including functional- and register-transfer-level modeling;
- design principles including modularity, hierarchy, and encapsulation;
- design patterns including message interfaces, control/datapath split, and FSM control;
- agile design methodologies including incremental development and test-driven development.

This handout assumes that you have read and understand the course tutorials and the lab assessment rubric. To get started, you should access the ECE computing resources and you have should have used the `ece4750-lab-admin` script to create or join a GitHub group. If you have not do so already, source the setup script and clone your lab group's remote repository from GitHub:

```
% source setup-ece4750.sh
% mkdir -p ${HOME}/ece4750
% cd ${HOME}/ece4750
% git clone git@github.com:cornell-ece4750/lab-groupXX
```

where `XX` is your group number. **You should never fork your lab group's remote repository! If you need to work in isolation then use a branch within your lab group's remote repository.** If you have already cloned your lab group's remote repository, then use `git pull` to ensure you have any recent updates before running all of the tests.You can run all of the tests in the lab like this:

```
% cd ${HOME}/ece4750/lab-groupXX
% git pull --rebase
% mkdir -p sim/build
% cd sim/build
% py.test ../lab3_mem
```

All of the tests for the provided functional-level model should pass, while the tests for the baseline and alternative cache designs should fail. For this lab, you will be working in the `lab3-mem` subproject which includes the following files:

- `BlockingCacheFL.py` – FL cache
- `BlockingCacheFL_test.py` – FL cache unit tests
- `BlockingCacheBasePRTL.py` – PyMTL blocking direct-mapped cache
- `BlockingCacheBaseCtrlPRTL.py` – PyMTL blocking direct-mapped cache's control unit
- `BlockingCacheBaseDpathPRTL.py` – PyMTL blocking direct-mapped cache's datapath
- `BlockingCacheBaseVRTL.py` – PyMTL wrapper around Verilog RTL
- `BlockingCacheBaseVRTL.v` – Verilog blocking direct-mapped cache
- `BlockingCacheBaseCtrlVRTL.v` – Verilog blocking direct-mapped cache's control unit
- `BlockingCacheBaseDpathVRTL.v` – Verilog blocking direct-mapped cache's datapath
- `BlockingCacheBaseRTL.py` – Wrapper to choose which RTL language
- `BlockingCacheBaseRTL_test.py` – Unit tests for direct-mapped cache
- `BlockingCacheAltPRTL.py` – PyMTL blocking set-associative cache
- `BlockingCacheAltCtrlPRTL.py` – PyMTL blocking set-associative cache's control unit
- `BlockingCacheAltDpathPRTL.py` – PyMTL blocking set-associative cache's datapath
- `BlockingCacheAltVRTL.py` – PyMTL wrapper around Verilog RTL
- `BlockingCacheAltVRTL.v` – Verilog blocking set-associative cache
- `BlockingCacheAltCtrlVRTL.v` – Verilog blocking set-associative cache's control unit
- `BlockingCacheAltDpathVRTL.v` – Verilog blocking set-associative cache's datapath
- `BlockingCacheAltRTL.py` – Wrapper to choose which RTL language
- `BlockingCacheAltRTL_test.py` – Unit tests for set-associative cache
- `TestCacheSink.py` – Custom test sink for cache
- `__init__.py` – Package setup

1. Introduction

Accessing main memory can require hundreds of cycles, but cache memories can significantly reduce the average memory access latency for well-structured address patterns. Caches are faster than main memory because they are smaller and are located close to the processor; but because a cache can only hold a subset of all memory locations at any one time, we must carefully manage what data we keep in the cache. A cache hit occurs when the data we are requesting is already in the cache, while a cache miss occurs when the data we are requesting is not in the cache and thus requires accessing main memory. Caches exploit spatial and temporal locality to increase the number of cache hits. In an address pattern with significant spatial locality, if we access a given address then in the near future, we are likely to access an address close to the first address. In an address pattern with significant temporal locality, if we access a given address then in the near future, we are likely to access that same address again. In this lab, you will implement and evaluate two cache microarchitectures that organize cache lines in two different ways: (1) direct-mapped where every cache line can only be placed in a single location in the cache, and (2) two-way set-associative where every cache line can be placed in one of two locations in the cache. Both caches will use a write-back, write-allocate policy for handling write misses.

We have provided you with a functional-level model of a cache, which essentially just passes all cache requests through to the memory interface, and passes all memory responses through to the cache response interface. While this might not seem useful, the functional-level model will enable us

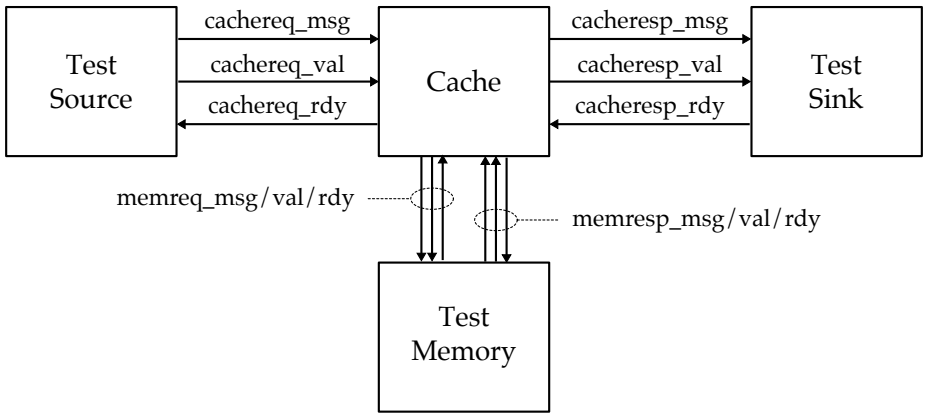

Figure 1: Memory System – The cache is integrated with a test source, test sink, and test memory for testing and evaluation.

```
                 76   74 73           66 65              34 33 32 31               0
cachreq         +-------+---------------+------------------+-----+------------------+
Message Format  | type  | opaque        | addr             | len | data             |
                +-------+---------------+------------------+-----+------------------+

                 46   44 43           36 35  34 33 32 31               0
cacheresp       +-------+---------------+------+-----+------------------+
Message Format  | type  | opaque        | test | len | data             |
                +-------+---------------+------+-----+------------------+

                 174 172 171         164 163            132 131 128 127              0
memreq          +-------+---------------+------------------+-------+------------------+
Message Format  | type  | opaque        | addr             | len   | data             |
                +-------+---------------+------------------+-------+------------------+

                 144 142 141         134 133 132 131 128 127              0
memresp         +-------+---------------+-------+-------+------------------+
Message Format  | type  | opaque        | test  | len   | data             |
                +-------+---------------+-------+-------+------------------+
```

Figure 2: Memory Request/Response Message Formats – Memory request messages are sent from the test source/sink to the cache and from the cache to/from the test memory.

to develop many of our test cases with the test memory before attempting to use these tests with the baseline and alternative designs.

Figure 1 shows a block-level diagram illustrating how the functional-level, baseline, and alternative designs are integrated with a test source, test sink, and test memory for testing and evaluation. We will load data into the test memory before resetting the cache. Once we start the execution, the test source will send memory requests into the cache, and eventually the cache will send memory responses to the test sink. If the cache needs to access main memory, then the cache will send memory requests to the test memory, and eventually the test memory will send memory responses back to the cache. We make extensive use of the latency insensitive val/rdy microprotocol in the cache interface. There are four different val/rdy channels.

- `cachereq` : from test source (processor) to cache
- `cachresp` : from cache to test sink (processor)
- `memreq` : from cache to test memory
- `memresp` : from test memory to cache

The message format for memory requests and responses are shown in Figure 2. Corresponding PyMTL `BitStructs` are defined in `pclib` here:

- `https://github.com/cornell-brg/pymtl/blob/ece4750/pclib/ifcs/MemMsg.py`

There are equivalent helper macros for Verilog located in `vc/mem-msgs.v`. The message format for the memory requests and responses are defined in `vc/mem-msgs.v` and shown in Figure 2. Memory requests use fields to encode the type (e.g., read, write, init), the address, the length of data in bytes, and the data. Memory responses use fields to encode the type (e.g., read, write), the length of data in bytes, and the data. The opaque field can be used for implementation defined behavior. You should always ensure the opaque field is correctly preserved in the response. Note that the memory messages used for the `cachereq` and `cacheresp` interfaces are for a single word (i.e., 32-bit data field and 2-bit length field), while the memory messages used for the `memreq` and `memresp` interfaces are for an entire cache line (i.e., 128-bit data field and 16-bit length field). If the length field is one then only the least significant byte of the data field (i.e., bits 7-0) is valid. If the length field is two then only the least significant two bytes of the data field (i.e., bits 15-0) are valid. If the length field is zero then all bytes are valid. We add a two-bit test field to each `cacheresp` and `memresp` message. We use the test field in `cacheresp` for testing. If a `cachereq` ends up with a cache miss, we should set the corresponding `cacheresp` message's test field to be `2'b0`. If a `cachereq` turns out to be a cache hit, we should set the corresponding `cacheresp` message's test field to be `2'b1`. By using the test field in the test harness can verify whether a cache transaction is a hit or a miss.

2. Baseline Design

The baseline design for this lab is a direct-mapped, write-back, write-allocate cache with a total capacity of 256 bytes, 16 cache lines, and 16 bytes per cache line. As with the earlier labs, we will be decomposing the baseline design into two separate modules: the datapath which has paths for moving data through various arithmetic blocks, muxes, and registers; and the control unit which is in charge of managing the movement of data through the datapath. As in the first lab, the control unit will use an FSM. Because the cache design is significantly more complicated than the first lab, we have decided to place the datapath module, control unit module, and the parent module that connects the datapath and control unit together in three different files.

The datapath for the baseline design is shown in Figure 3. The blue signals represent the control/status signals for communicating between the datapath and the control unit. Your datapath module should probably instantiate a child module for each of the blocks in the datapath diagram; in other words, you should mostly use a structural design style in the datapath. Although you are free to develop your own modules to use in the datapath, you can also use the ones provided for you in the pclib (PyMTL) or VC library (Verilog). The `repl` block takes a 32b value and simply replicates it four times to create a 128b value. The `mkaddr` block simply concatenates the tag and idx plus some zeros like this: `{ tag, idx, 4'b0000 }`. As we become more proficient, we can use our judgement about what needs to be encapsulated in a child module. For example, the `repl` and `mkaddr` blocks can probably be done directly in the datapath without encapsulating them in a module. These operations are just wiring and not really "logic". However, we strongly encourage students to use comments and a temporary signal to clearly indicate where in the code you are doing this kind of wiring so that it is still easy to connect your datapath diagram to your code. Notice that the simplify our design, we are implementing the tag and data arrays using combinational SRAMs. We provide SRAMs as a pclib component in PyMTL, see:

- `https://github.com/cornell-brg/pymtl/blob/ece4750/pclib/rtl/SRAMs.py`

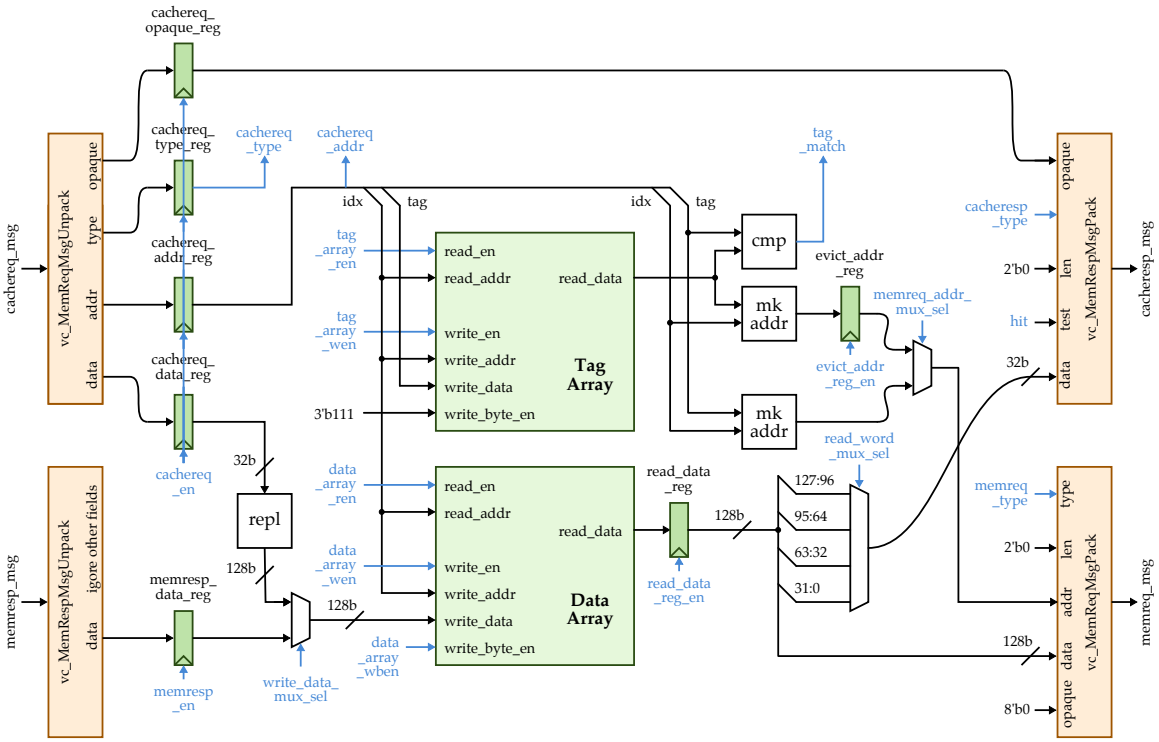

Figure 3: Baseline Design – Direct-mapped, write-back, write-allocate cache.

The equivalent Verilog component is located in `vc/srams.v` called `vc_CombinationalSRAM_1rw`. This means that after setting the read address (`raddr`) signals, the read data (`rdata`) is available that same cycle. More realistic designs would likely need to move to using synchronous SRAMs.

The FSM for the baseline design is shown in Figure 4. The control unit should also include additional valid and dirty bits to track the state of each tag entry. You should not use SRAMs for the valid and dirty bits. You should instead use standard register files. For PyMTL, register files are located in:

- `https://github.com/cornell-brg/pymtl/blob/ece4750/pclib/rtl/RegisterFile.py`

The equivalent Verilog register files are defined in `vc/regfiles.v`. The various FSM states are described in more detail below:

- `STATE_IDLE (I)` – Receive the incoming cache request and place it in the input registers
- `STATE_TAG_CHECK (TC)` – Check tag, state transition based on hit/miss, valid/dirty
- `STATE_INIT_DATA_ACCESS (IN)` – Immediately write to appropriate cache line
- `STATE_READ_DATA_ACCESS (RD)` – Read from appropriate cache line
- `STATE_WRITE_DATA_ACCESS (WD)` – Write to appropriate cache line
- `STATE_REFILL_REQUEST (RR)` – Make a request to memory for refill cache line
- `STATE_REFILL_WAIT (RW)` – Wait for memory response
- `STATE_REFILL_UPDATE (RU)` – Write the response to the cache line
- `STATE_EVICT_PREPARE (EP)` – Read tag and data, and prepare eviction message
- `STATE_EVICT_REQUEST (ER)` – Make a request to memory to write the evicted cache line
- `STATE_EVICT_WAIT (EW)` – Wait for memory response
- `STATE_WAIT (W)` – Wait while the sink (processor) is busy

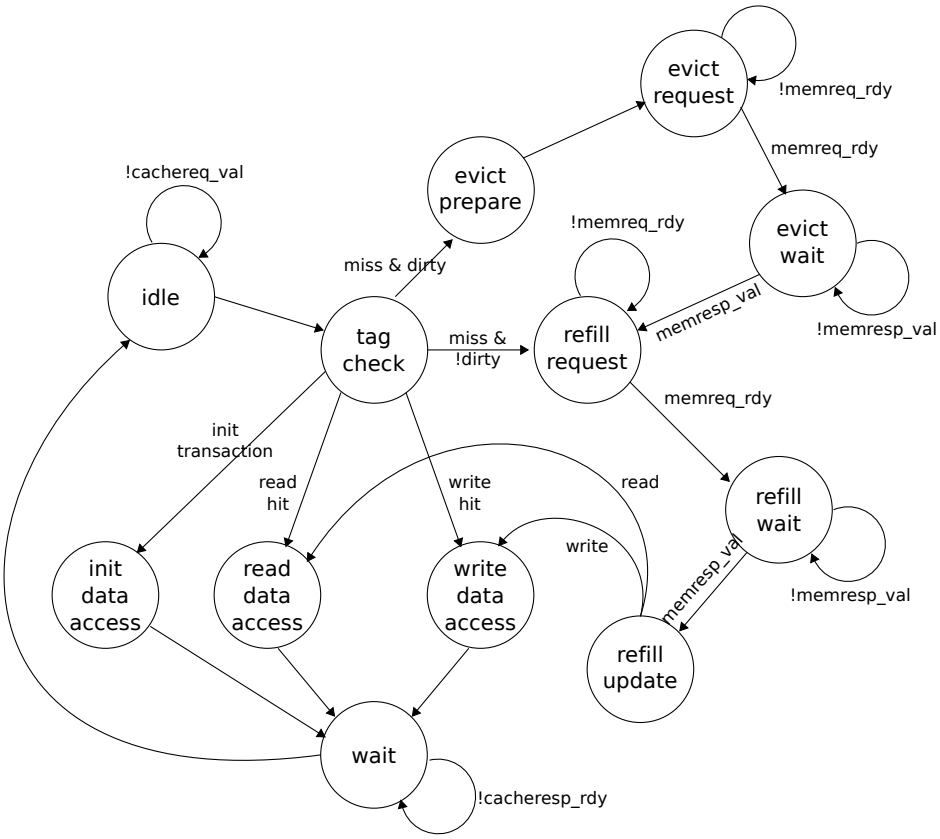

Figure 4: Baseline FSM Control Unit

This FSM control unit differs from the basic FSM discussed in lecture. This is due to the need to handle the various latency insensitive interfaces, the init transaction, and waiting for eviction responses. The FSM cache microarchitectures used in this lab will have a hit latency of four cycles (e.g., I $\rightarrow$ TC $\rightarrow$ RD $\rightarrow$ W). In the previous lab, the memory access time was usually a single cycle although you also tested your processor with random delays on the memory interface. Assuming we correctly implement the latency insensitive val/rdy microprotocol in both the processor and cache, then there should be no problem composing these two subsystems. Later optimizations can reduce the cache hit latency without the need to modify the processor.

We strongly encourage you to take an incremental design approach using the following steps:

- Implement the init transaction (I, TC, IN, W)
- Implement the read hit path (I, TC, RD, W)
- Implement the write hit path (I, TC, WD, W)
- Implement the refill path (I, TC, RR, RW, RU, RD/WD, W)
- Implement the evict path (I, TC, EP, ER, EW, RR, RW, RU, RD/WD, W)

The hit path is the simplest place to start, but in order to hit in the cache, we need valid data in the cache. The init transaction (explained further in Section 4) allows you to write data into the cache without doing a refill from main memory. This kind of transaction is an example of *design-for-test* since this transaction is only meant to simplify testing and has no real use once we have a working cache model. We recommend implementing the init transaction first, and then using this transaction

to implement and test both hit paths. Once the hit paths are working, then you can move onto the more complicated miss paths.

3. Alternative Design

The alternative design for this lab is a two-way set-associative, write-back, write-allocate cache with the same capacity (256 bytes) and cache line size (16 bytes) as the baseline cache. The general FSM for the alternative design will be very similar to that of the baseline design, except that the control signals will likely be different. Note that you will need to split the valid bits into two parts, one for each way and carefully keep track of them. You will need to AND the result of the tag match in each way with the appropriate valid bit to determine if there is a hit or miss. The control unit should use a least-recently-used (LRU) replacement policy to choose between the two ways during eviction. You should track the LRU status with separate bits in the control unit.

4. Testing Strategy

We provide you with one very basic directed test that tests the read hit path for clean lines. Although you will not need to write as many tests as in the previous lab, the tests for this lab may be more challenging since you will need to carefully craft directed tests that exercise all paths in your datapath and all states and state transitions in your FSM. As with the previous labs, you will want to initially write tests using the functional-level model. Once these tests are working on the functional-level model, you can move on to testing the baseline and alternative designs.

The following commands illustrate how to run all of tests for the entire project, how to run just the tests for this lab, and how to run just the basic test we provide on the various designs.

```
% cd ${HOME}/ece4750/lab-groupXX/sim/build
% py.test ..
% py.test ../lab3_mem
% py.test ../lab3_mem/BlockingCacheFL_test.py
% py.test ../lab3_mem/BlockingCacheBaseRTL_test.py
% py.test ../lab3_mem/BlockingCacheAltRTL_test.py
```

You will add your directed and random tests to `BlockingCacheFL_test.py`. Since this harness is shared across the functional-level model, the baseline design, and the alternative design you can write your tests once and reuse them to test all three models. You will be adding more test cases. Do not just make the given test case larger. A key challenge in writing directed tests for cache memories, is that most of the miss path must be working before you can test the hit path. The miss path is significantly more complicated than the hit path, so this lends itself more towards a monolithic design process. Most of the cache must be implemented before we can run our first directed test. We could potentially use white-box ad-hoc testing that directly initializes the tag and data arrays in the cache before starting the test, but these ad-hoc tests are fragile and difficult to develop on the functional-level model.

To address this challenge, we will use a design-for-test (DFT) approach. DFT involves adding features to a design that are solely meant to facilitate test-driven development. In this specific design, we will be adding a new init transaction to go along with our current read and write transactions. The init transaction does the following:

- Writes (overwrites) the appropriate cache line based on the index bits of the address
- Never updates main memory

- Sets the valid bit for that cache line to one
- Sets the dirty bit for that cache line to zero
- Sets the `cacheresp`'s test field to zero

Using an init transaction on a dirty cache line is undefined. Essentially, the init transaction is a special transaction that directly initializes a word in a cache line.

Figure 5 illustrates how we will be writing tests for this lab using various helper tasks which are defined in `BlockingCacheFL_test.py`. The `req` and `resp` Python helper functions will create a memory request message and the expected memory response message. We use the test field in response messages to test whether the memory request resulted in a hit or a miss. This specific example first uses an init transaction to initialize one word in the first cache line, before using a read transaction to read this same word. Note that the second transection should be a hit and we set the second `cacheresp` message's test field to be one. For each test case we define a Python function that returns a list of request-response message pairs. The request messages and response messages will be passed to the test source and the test sink respectively. In Figure 5, lines 6–11 and lines 17–22 illustrate two test cases. We also provide you a helper function to load data into the test memory before running the test. It will allow you to test the cache refill path (read miss path) without involving the cache evict path. If you want to load data into the test memory before running the test, you can create a Python function that returns a list of address-data pairs and then the test function we provide will load it to the test memory before running the test. For example lines 26–31, load the test memory with two words of data at addresses `0x00000000` and `0x00000004` before running the `read_miss_1word_msg` test case.

Once you create a new test harness, you can add it to the test case table, as shown on lines 37–41 in Figure 5. A test case table has seven columns. The first column is the name of tests, the second one is the function that generates source/sink messages, and the third one is the function that generates memory data to be loaded before running the test. If your test case does not need to load data to the test memory, simply use `None`. The last four columns are for test memory's stall probability, test memory latency, source delay, and sink delay, respectively.

Ideally, we want to use same set of tests to test the FL, baseline RTL, and alternative RTL model. We define all the test cases in `BlockingCacheFL_test.py`, test them on FL model to make sure test cases themself are correct, and we import them in the `BlockingCacheBaseRTL_test.py` and `BlockingCacheAltRTL_test.py` test scripts. We need to be careful about the test field, because the same test may have different hit/miss behavior on different cache implementations. For example, in an FL cache, every response is a miss because we simply pass `cachereq` and `cacheresp` to memory. The alternative design will have less conflict misses than the baseline design, because the alternative design is set-associative whereas the baseline design is direct-mapped. Therefore, we should expect different values for the test field in the `cacheresp` messages passed to the sink from different cache models even if we use the same source messages. To solve this issue, we provide you a custom `TestHarness` model in `BlockingCacheFL_test.py` that can optionally skip checking the test field. Lines 49–52 in Figure 5 show how to instantiate this `TestHarness` model. The first two parameters (`src_msgs` and `sink_msgs`) are source and sink messages. The `stall_prob`, `latency`, `src_delay`, and `sink_delay` parameters are the stalling probability of the test memory, latency (in cycles) of the test memory, source delay, and source delay, respectively. `CacheModel` is the model we want to test. `check_test` determines whether the test harness should check the test field in `cacheresp` messages from the model under test to the test sink. When we run a test case on the FL model, we should set `check_test` to `False` to make the harness skip checking the test field because the test field is always zero for the FL model. Otherwise the test will fail. When we test the actual RTL models, we must

```
#-------------------------------------------------------------------------
# Test Case: read hit path
#-------------------------------------------------------------------------
# The test field in the response message: 0 == MISS, 1 == HIT

def read_hit_1word_clean( base_addr ):
  return [
    #    type  opq  addr      len data               type  opq  test len data
    req( 'in', 0x0, base_addr, 0, 0xdeadbeef ), resp( 'in', 0x0, 0,   0,  0          ),
    req( 'rd', 0x1, base_addr, 0, 0          ), resp( 'rd', 0x1, 1,   0,  0xdeadbeef ),
  ]

#-------------------------------------------------------------------------
# Test Case: read miss path
#-------------------------------------------------------------------------

def read_miss_1word_msg( base_addr ):
  return [
    #    type  opq   addr      len  data              type  opq test len  data
    req( 'rd', 0x00, 0x00000000, 0, 0          ), resp('rd', 0x00, 0, 0, 0xdeadbeef ), # read word  0x00000000
    req( 'rd', 0x01, 0x00000004, 0, 0          ), resp('rd', 0x01, 1, 0, 0x00c0ffee ), # read word  0x00000004
  ]

# Data to be loaded into memory before running the test

def read_miss_1word_mem( base_addr ):
  return [
    # addr      data (in int)
    0x00000000, 0xdeadbeef,
    0x00000004, 0x00c0ffee,
  ]

#-------------------------------------------------------------------------
# Test table for generic test
#-------------------------------------------------------------------------

test_case_table_generic = mk_test_case_table([
  (                       "msg_func              mem_data_func         stall lat src sink"),
  [ "read_hit_1word",      read_hit_1word,       None,                 0.0,  0,  0,  0    ],
  [ "read_miss_1word",     read_miss_1word_msg,  read_miss_1word_mem,  0.0,  0,  0,  0    ],
])

@pytest.mark.parametrize( **test_case_table_generic )
def test_generic( test_params, dump_vcd ):
  msgs = test_params.msg_func( 0 )
  if test_params.mem_data_func != None:
    mem = test_params.mem_data_func( 0 )
  # Instantiate testharness
  harness = TestHarness( msgs[::2], msgs[1::2],
                         test_params.stall, test_params.lat,
                         test_params.src, test_params.sink,
                         BlockingCacheFL, False )
  # Load memory before the test
  if test_params.mem_data_func != None:
    harness.load( mem[::2], mem[1::2] )
  # Run the test
  run_sim( harness, dump_vcd )
```

Figure 5: Writing Directed Tests – A portion of the `BlockingCacheFL_test.py` file. We create all test cases in this file and use it to test both FL, baseline RTL and alternative RTL model.

set `check_test` to `True` in order to inspect every `cacheresp` message's test field to see verify if the request correctly hit or missed in the cache.

As mentioned above, your baseline and alternative designs will have different hit/miss behaviors, so you may need to use different sink values even for the same sequence of source messages. You should add tests designed specifically for your baseline or alternative design. For example, you should have tests that only hit in a two-way set-associative cache to make sure your alternative design is indeed two-way set-associative. We group the test cases into three test case tables. `test_case_table_generic` is shown on lines 37–41 in Figure 5 and is used to generically test both the baseline and alternative designs (i.e., tests in this table should have the same expected behavior for both the baseline and alternative design). `test_case_table_set_assoc` and `test_case_table_set_dir_mapped` are used to test only one of the designs. We provide examples which test for a very simple conflict miss in the baseline design and test for a hit with the same sequence of source messages in the alternative design.

Some suggestions for what you might want to test are listed below. Each of these would probably be a separate test case.

- Read hit path for clean lines
- Write hit path for clean lines
- Read hit path for dirty lines
- Write hit path for dirty lines
- Read miss with refill and no eviction
- Write miss with refill and no eviction
- Read miss with refill and eviction
- Write miss with refill and eviction
- Tests which stress entire cache, not just a few cache lines
- Conflict misses
- Capacity misses
- LRU replacement policy by filling up a way
- Tests specifically designed to trigger corner cases in your alternative design
- Testing all or some of the above using random source and sink delays and test memory delays

Once you have finished writing your directed tests you should move on to writing random tests. You can use the same Python-based random test generation system we used in the first lab. Some suggestions for what you might want to test are listed below. Each of these would probably be a separate test pattern, or potentially multiple test patterns with different random parameters.

- Simple address patterns, single request type, with random data
- Simple address patterns, with random request types and data
- Random address patterns, request types, and data
- Unit stride with random data
- Strided with random data
- Unit stride (high spatial locality) mixed with shared (high temporal locality)

Writing random tests for memories can actually be quite challenging. With the first lab, the correct output was trivial to calculate based on the random inputs, but with a memory system the correct output (i.e., the data we expect in a memory read response) depends on the last write to the corresponding address. To write random tests with random address patterns and/or types, you will need to keep track of a "reference memory" in your Python script. This reference memory can just be an array of words. Every time you generate a write request, you should update the reference memory in addition to printing the appropriate write request. Every time you generate a read request, you

```
cycle cachreq              state memreq                 memresp                cacheresp
  2:                        (I )                       ().                    .
  3: in:00:00001000:deadbeef (I )                      ().
  4: #                      (TC)                       ().
  5: #                      (IN)                       ().
  6: #                      (W )                       ().                    in:00:0:
  7: rd:01:00001000:        (I )                       ().
  8: .                      (TC)                       ().
  9: .                      (RD)                       ().
 10: .                      (W )                       ().                    rd:01:1:deadbeef
```

Figure 6: Line Trace for Basic Directed Test – The line trace shows two memory requests sent from the test source to the cache, the four states each transaction goes through, and then the response being sent from the cache back to the test sink.

```
cycle cachreq              state memreq                 memresp                    cacheresp
  2:                        (I )                       ().                        .
  3: rd:00:00000000:        (I )                       ().
  4: #                      (TC)                       ().
  5: #                      (RR) rd:00:00000000:       ().
  6: #                      (RW)                       ()rd:00:0:0...ba11ad0e5ca18d
  7: #                      (RU)                       ().
  8: #                      (RD)                       ().
  9: #                      (W )                       ().                        rd:00:0:0e5ca18d
 10: rd:01:00000000:        (I )                       ().
 11: #                      (TC)                       ().
 12: #                      (RD)                       ().
 13: #                      (W )                       ().                        rd:01:1:0e5ca18d
 14: rd:02:00000004:        (I )                       ().
 15: #                      (TC)                       ().
 16: #                      (RD)                       ().
 17: #                      (W )                       ().                        rd:02:1:00ba11ad
 18: wr:03:00000100:00e1de57 (I )                      ().
 19: #                      (TC)                       ().
 20: #                      (RR) rd:00:00000100:       ().
 21: #                      (RW)                       ()rd:00:0:0...00000000000000
 22: #                      (RU)                       ().
 23: #                      (WD)                       ().
 24: #                      (W )                       ().                        wr:03:0:
 25: rd:04:00000100:        (I )                       ().
 26: .                      (TC)                       ().
 27: .                      (RD)                       ().
 28: .                      (W )                       ().                        rd:04:1:00e1de57
```

Figure 7: Line Trace for More Involved Directed Test – The line trace shows several memory requests meant to trigger an eviction of a dirty line. Notice how the third request must go through a total of 10 states as the FSM does tag check, eviction, and refill. Note, only a portion of the bits for the data field for memory request/responses are shown for simplicity.

should consult your reference memory to determine what data we expect to be returned in a memory read response.

You will almost certainly want to use line tracing to visualize the execution of transactions on your baseline and alternative designs. We have provided some line tracing code for you in the test harness which traces the cache request/response and memory request/response interfaces. Figure 6 illustrates a line trace for the basic test in Figure 5 executing on the baseline design with extra annotations to indicate what the columns mean. The first columns shows when memory request messages are sent from the test source into the cache, and the last column shows when memory response messages are sent from the cache back to the test sink. The second column shows the state of the processor. This column is critical to understanding the behavior of your cache, but it is not currently

```
// loop-1d pattern
for ( i = 0; i < 100; i++ )
  result += a[i];
```

```
// loop-2d pattern
for ( i = 0; i < 5; i++ )
  for ( j = 0; j < 100; j++ )
    result += a[j];
```

```
// loop-3d pattern
for ( i = 0; i < 5; i++ )
  for ( j = 0; j < 2; j++ )
    for ( k = 0; k < 8; k++ )
      result += a[j*64 + k*4];
```

Figure 8: Evaluation Patterns – Three loops that correspond to the given evaluation patterns.

implemented in the lab harness. You will need to modify the line tracing code in your baseline and alternative designs to append a string representing the current cache state. Use the short state names as given in the state description list above (e.g., `I` for `STATE_IDLE`, `TC` for `STATE_TAG_CHECK`). The third and fourth columns show the memory request and response messages to/from the test memory.

In addition to the tests for the entire cache, you must also add additional unit tests for any datapath components you add or modify.

5. Evaluation

Once you have verified the functionality of the baseline and alternative designs, you should then use the provided simulator to evaluate these two designs. The simulator delay all responses from the test memory by 20 cycles to model a long main-memory latency. You can run the simulator to see the performance of each cache implementation as follows:

```
% cd ${HOME}/ece4750/lab-groupXX/sim/build
% ../lab3_mem/mem-sim --impl base --pattern loop-1d --stats
% ../lab3_mem/mem-sim --impl alt  --pattern loop-1d --stats
```

We provide you with three input patterns that capture common access patterns in loops. The C code for each loop that might generate the corresponding pattern is shown in Figure 8. The `loop-1d` pattern simply iterates through an array. The `loop-2d` pattern iterates through the same array five times. The `loop-3d` patterns uses a more interesting stride to iterate through an array multiple times.

The simulator will display a collection of statistics: number of cycles, number of memory and cache accesses, number of misses, miss rate, and the average memory access latency (AMAL). You should study the line traces (with the `--trace` option) and possibly the waveforms (with the `--dump-vcd` option) to understand the reason why each design performs as it does on the various patterns.

You can run simulations for all given patterns like this:

```
% cd ${HOME}/ece4750/lab-groupXX/sim/build
% python ../lab3_mem/mem_sim_eval.py
```

You will almost certainly need to add additional evaluation patterns with different amounts of spatial and temporal locality. We strongly recommend including some patterns that mix reads/writes and random patterns. We recommend maybe six or more patterns for evaluation. Obviously, these patterns need to be carefully chosen to highlight the differences between the baseline and alternative designs. You will also need to add the name of the new dataset to the `mem-sim` simulator script.

Writing an access pattern for the simulator is similar to writing a test case. Basically for each pattern you need to write a Python function that returns a list of source messages, a list of sink message, and a list of memory address-data pairs that will be loaded to the test memory before the simulation. Figure 9 shows you `loop-3d` pattern. Please keep in mind that patterns you will write in the simulator are not for testing. Instead, you need to fully test your designs using test cases and then use patterns in the simulator script as a way to evaluate your designs.

```
#-------------------------------------------------------------------------
# Pattern: loop-3d
#-------------------------------------------------------------------------

def loop_3d():
  src_msgs  = []
  sink_msgs = []

  mem_data  = []
  mem_word  = Bits( 32 )

  # Initialize memory
  addr = 0
  for i in xrange( 2 ):
    for j in xrange( 8 ):
      addr           = i*256 + j*16
      mem_word.value = i*64  + j*4
      mem_data.append( addr )
      mem_data.append( mem_word.uint() )
      addr += 4

  # Read from memory
  for i in xrange( 5 ):
    for j in xrange( 2 ):
      for k in xrange( 8 ):
        addr = j*256 + k*16
        data = j*64  + k*4
        src_msgs.append ( mk_req ( 0, addr, 0, 0 ) )
        sink_msgs.append( mk_resp( 0, 0, data    ) )

  return [ src_msgs, sink_msgs, mem_data ]
```

Figure 9: `loop-3d` **Pattern for the Simulator –** A function that returns source messages, sink messages, and a memory section for `loop-3d` access pattern.

6. PARCv3 Extensions

Each lab assignment includes additional extensions that would be required to transform the alternative design into a subsystem suitable for use in a full PARCv3 multicore. Students are free to work on these extensions, although they must implement them in a separate `lab3_mem_ext` subdirectory. Do not implement any design extensions in the main lab subdirectory! It is important that any work on design extensions not cause the tests for your baseline and alternative designs to fail. Design extensions will not be used to award extra credit, nor will they be factored into the grading of labs 1–4 in any way. However, design extensions can be factored into the grading for the baseline and alternative design section of the final lab assignment.

For this cache to be used in a reasonable PARCv3 processor, we might need to add some of the following features:

- The FSM in your baseline and alternative design can be significantly optimized. Many states can be collapsed together. Optimize the performance of your design by first focusing on collapsing the states involved on the hit path to reduce the hit latency to just two cycles. Experiment with collapsing states on the miss path as well.
- In the baseline and alternative design, we only support full word (4B) accesses. Add support for byte, and half-word reads and writes. Carefully consider how you will handle either sign extending or zero extending these subword accesses for read requests (i.e., to support `lbu`, `lb`, `lhu`, and `lh` instructions).

- In your baseline and alternative designs, you assume a fixed cache size of 256 bytes. Improve the extensibility of your design by parameterizing both the capacity and line size.
- Implement a pipelined cache microarchitecture with a two-cycle pipelined hit path. Tag check should be in the first stage and data access should be in the second stage. The miss path can still use a FSM.
- Implement a parallel read, pipelined write cache microarchitecture. Carefully consider how to resolve the structural hazard for accessing the data array.
- Add support for a simple atomic memory operation that would be suitable for implementing the `amo.add` instruction. See the PARC instruction set manual for more information on the `amo.add` instruction.
- Compose two of your cache designs in series to create a two-level cache hierarchy. Either parameterize your cache design, or duplicate the cache design so that the L2 cache can have a much larger capacity.

Acknowledgments

This lab was created by Christopher Torng, Berkin Ilbeyi, Ackerley Tng, Shreesha Srinath, Christopher Batten, and Ji Kim as part of the course ECE 4750 Computer Architecture at Cornell University.