

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49

The Flat Operational Model

CHRISTOPHER PULTE, University of Cambridge
SHAKED FLUR, University of Cambridge
WILL DEACON, ARM Ltd.
JON FRENCH, University of Cambridge
SUSMIT SARKAR, University of St. Andrews
PETER SEWELL, University of Cambridge

This document gives a prose description of the Flat operational model, as formally defined in its Lem definition. This is part of the supplementary material for “Simplifying ARM Concurrency: Multicopy-atomic Axiomatic and Operational Models for ARMv8”.

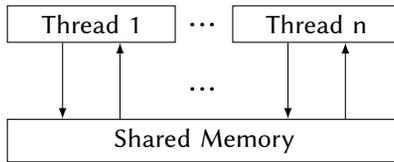
1 AN OPERATIONAL MODEL FOR MCA ARMV8

To help reading this document we have colour-coded some text as follows:

- [release/acquire] Release/Acquire instructions
- [exclusive] Exclusive instructions
- [dmb ld/dmb st] dmb ld and dmb st instructions

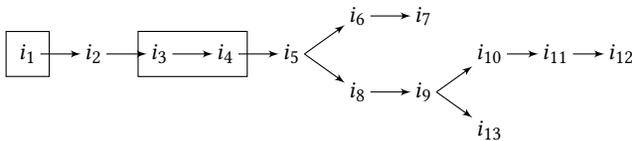
The operational model is expressed as a state machine, with states that are an abstract representation of hardware machine states. We first introduce the model states and transitions informally.

Model states A model state consists just of a shared memory and a tuple of thread model states:



The shared memory state effectively just records the most recent write to each location. To handle load/store-exclusives, the memory is extended with a map (the exclusives map) from read requests to sets of write slices, that associates a read request of a load-exclusive with the write slices it read from (excluding writes that have been forwarded to the read and have not reached memory yet).

Each thread model state consists principally of a list or tree of instruction instances, some of which have been finished, and some of which have not. For example, below we show a thread model state with instruction instances i_1, \dots, i_{13} , and the program-order-successor relation between them. Three of those (i_1 , i_3 , and i_4 , boxed) have been finished; the remainder are non-finished.



Authors’ addresses: Christopher Pulte, University of Cambridge, first.last@cl.cam.ac.uk; Shaked Flur, University of Cambridge, first.last@cl.cam.ac.uk; Will Deacon, ARM Ltd. first.last@arm.com; Jon French, University of Cambridge, first.last@cl.cam.ac.uk; Susmit Sarkar, University of St. Andrews, ss265@st-andrews.ac.uk; Peter Sewell, University of Cambridge, first.last@cl.cam.ac.uk.

2017. XXXX-XXXX/2017/10-ART \$15.00
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

50 Non-finished instruction instances can be subject to restart, e.g. if they depend on an out-of-order or
51 speculative read that turns out to be unsound. The finished instances are not necessarily contiguous:
52 in the example, i_3 and i_4 are finished even though i_2 is not, which can only happen if they are
53 sufficiently independent. Instruction instances i_5 and i_9 are conditional branches for which the
54 thread has fetched multiple possible successors. When a conditional branch is finished, any un-
55 taken alternative paths are discarded, and instruction instances that follow (in program order) a
56 non-finished conditional branch cannot be finished until that conditional branch is. One can choose
57 whether or not to allow simultaneous exploration of multiple successors of a conditional branch
58 (as shown above); this does not affect the set of allowed outcomes.

59 The intra-instruction behaviour of a single instruction can largely be treated as sequential
60 (but not atomic) execution of its ASL/Sail pseudocode. Each instruction instance state includes a
61 pseudocode execution state, which one can think of as a representation of the pseudocode control
62 state, pseudocode call stack, and local variable values. An instruction instance state also includes
63 information, detailed below, about the instruction instance's memory and register footprints, its
64 register and memory reads and writes, whether it is finished, etc.

65 **Model transitions** For any state, the model defines the set of allowed transitions, each of which is
66 a single atomic step to a new abstract machine state. Each transition arises from the next step of a
67 single instruction instance; it will change the state of that instance, and it may depend on or change
68 the rest of its thread state and/or the shared memory state. Instructions cannot be treated as atomic
69 units: complete execution of a single instruction instance may involve many transitions, which
70 can be interleaved with those of other instances in the same or other threads, and some of this is
71 programmer-visible. The transitions are introduced below and defined in §1.4, with a precondition
72 and a construction of the post-transition model state for each. The transitions labelled \circ can always
73 be taken eagerly, as soon as they are enabled, without excluding other behaviour; the \bullet cannot.

74 Transitions for all instructions:

- 75 \bullet **Fetch instruction**: This transition represents a fetch and decode of a new instruction instance,
76 as a program-order successor of a previously fetched instruction instance, or at the initial
77 fetch address for a thread.
- 78 \circ **Register read**: This is a read of a register value from the most recent program-order predecessor
79 instruction instance that writes to that register.
- 80 \circ **Register write**
- 81 \circ **Pseudocode internal step**: this covers ASL/Sail internal computation, function calls, etc.
- 82 \circ **Finish instruction**: At this point the instruction pseudocode is done, the instruction cannot
83 be restarted or discarded, and all memory effects have taken place. For a conditional branch,
84 any non-taken po-successor branches are discarded.

85 Load instructions:

- 86 \circ **Initiate memory reads of load instruction**: At this point the memory footprint of the load is
87 provisionally known and its individual reads can start being satisfied.
- 88 \bullet **Satisfy memory read by forwarding from writes**: This partially or entirely satisfies a single
89 read by forwarding from its po-previous writes.
- 90 \bullet **Satisfy memory read from memory**: This entirely satisfies the outstanding slices of a single
91 read, from memory.
- 92 \circ **Complete load instruction (when all its reads are entirely satisfied)**: At this point all the
93 reads of the load have been entirely satisfied and the instruction pseudocode can continue
94 execution. A load instruction can be subject to being restarted until the **Finish instruction**
95 transition. In some cases it is possible to tell that a load instruction will not be restarted or
96
97
98

99 discarded before that, e.g. when all the instructions po-before the load instruction are finished.
 100 The **Restart condition** over-approximates the set of instructions that might be restarted.

101 Store instructions:

- 102 ○ **Initiate memory writes of store instruction, with their footprints:** At this point the memory
- 103 footprint of the store is provisionally known.
- 104 ○ **Instantiate memory write values of store instruction:** At this point the writes have their
- 105 values and program-order-subsequent reads can be satisfied by forwarding from them.
- 106 ○ **Commit store instruction:** At this point the store is guaranteed to happen (it cannot be
- 107 restarted or discarded), and the writes can start being propagated to memory.
- 108 ● **Propagate memory write:** This propagates a single write to memory.
- 109 ○ **Complete store instruction (when its writes are all propagated):** At this point all writes have
- 110 been propagated to memory, and the instruction pseudocode can continue execution.

112 **Store-exclusive instructions:**

- 113 ● **Guarantee the success of store-exclusive: This guarantees the success of the store-exclusive.**
- 114 ● **Make a store-exclusive fail: This makes the store-exclusive fail.**

116 Barrier instructions:

- 117 ○ **Commit barrier**

119 1.1 Intra-instruction Pseudocode Execution

120 To link the model transitions introduced above to the execution of the instructions an interface
 121 is needed between Sail and the rest of the concurrency model. For each instruction instance this
 122 intra-instruction semantics is expressed as a state machine, essentially running the instruction
 123 pseudocode, where each pseudocode execution state is a request of one of the following forms:

124	READ_MEM(<i>read_kind</i> , <i>address</i> , <i>size</i> , <i>read_continuation</i>)	Read request
125	EXCL_RES(<i>res_continuation</i>)	Store-exclusive result
126	WRITE_EA(<i>write_kind</i> , <i>address</i> , <i>size</i> , <i>next_state</i>)	Write effective address
127	WRITE_MEMV(<i>memory_value</i> , <i>write_continuation</i>)	Write value
128	BARRIER(<i>barrier_kind</i> , <i>next_state</i>)	Barrier
129	READ_REG(<i>reg_name</i> , <i>read_continuation</i>)	Register read request
130	WRITE_REG(<i>reg_name</i> , <i>register_value</i> , <i>next_state</i>)	Write register
131	INTERNAL(<i>next_state</i>)	Pseudocode internal step
132	DONE	End of pseudocode

134 Each of these states is a suspended computation with a request for an action or input from the
 135 concurrency model and, except in the case of DONE, a continuation for the remaining execution.

136 Here memory values are lists of bytes, addresses are 64-bit numbers, read and write kinds identify
 137 whether they are regular, exclusive, and/or release/acquire operations, register names identify
 138 a register and slice thereof (start and end bit indices), and the continuations describe how the
 139 instruction instance will continue for any value that might be provided by the surrounding memory
 140 model. This largely follows Gray et al. [2015, §2.2], except that memory writes are split into two
 141 steps, WRITE_EA and WRITE_MEMV. We ensure these are paired in the pseudocode, but there may
 142 be other steps between them: it is observable that the WRITE_EA can occur before the value to
 143 be written is determined, because the potential memory footprint of the instruction becomes
 144 provisionally known then.

145 We ensure that each instruction has at most one memory read, memory write, or barrier step,
 146 by rewriting the pseudocode to coalesce multiple reads or writes, which are then split apart into

147

the architecturally atomic units by the thread semantics; this gives a single commit point for all memory writes of an instruction.

Each bit of a register read should be satisfied from a register write by the most recent (in program order) instruction instance that can write that bit, or from the thread’s initial register state if there is no such. That instance may not have executed its register write yet, in which case the register read should block. The semantics therefore has to know the register write footprint of each instruction instance, which it calculates when the instruction instance is created. We ensure in the pseudocode that each instruction does exactly one register write to each bit of its register footprint, and also that instructions do not do register reads from their own register writes. In some cases, but not in the fragment of ARM that we cover at present, register write footprints need to be dynamically recalculated, when the actual footprint only becomes known during pseudocode execution.

Data-flow dependencies in the model emerge from the fact that a register read has to wait for the appropriate register write to be executed (as described above). This has to be carefully handled in order not to create unintentional strength. First, for some instructions we need to ensure that the pseudocode is in the maximally liberal order, e.g. to allow early computed-address register writebacks before the corresponding memory write. Leaving load-pair aside (which we do not cover), and the treatment of the multiple reads or writes that can be associated with a single load or store instruction (which we do), we have not so far needed other intra-instruction concurrency. Second, the model has to be able to know when a register read value can no longer change (i.e. due to instruction restart). We approximate that by recording, for each register write, the set of register and memory reads the instruction instance has performed at the point of executing the write. This information is then used as follows to determine whether a register read value is final: if the instruction instance that performed the register write from which the register reads from is finished, the value is final; otherwise check that the recorded reads for the register write do not include memory reads, and continue recursively with the recorded register reads. For the instructions we cover this approximation is exact.

We express the pseudocode execution semantics in two ways: a definitional interpreter for Sail [Gray et al. 2015], with an exhaustive symbolic mode to (re)calculate an instruction’s memory and register footprints, and as a shallow embedding, translating Sail into directly executable code, with separate hand-written definitions of the footprint functions. The two are essentially equivalent: the first lets one small-step through the pseudocode interactively, while the second is more efficient and should be more convenient for proof.

1.2 Instruction Instance States

Each instruction instance i has a state comprising:

- *program_loc*, the memory address from which the instruction was fetched;
- *instruction_kind*, identifying whether this is a load, store, or barrier instruction, each with the associated kind; or a conditional branch; or a ‘simple’ instruction.
- *regs_in*, the set of input *reg_names*, as statically determined;
- *regs_out*, the output *reg_names*, as statically determined;
- *pseudocode_state* (or sometimes just ‘state’ for short), one of
 - *PLAIN_next_state*, ready to make a pseudocode transition;
 - *PENDING_MEM_READS read_cont*, performing the read(s) from memory of a load; or
 - *PENDING_MEM_WRITES write_cont*, performing the write(s) to memory of a store;
- *reg_reads*, the accumulated register reads, including their sources and values, of this instance’s execution so far;

- *reg_writes*, the accumulated register writes, including dependency information to identify the register reads and memory reads (by this instruction) that might have affected each;
- *mem_reads*, a set of memory read requests. Each request includes a memory footprint (an address and size) and, if the request has already been satisfied, the set of write slices (each consisting of a write and a set of its byte indices) that satisfied it.
- *mem_writes*, a set of memory write requests. Each request includes a memory footprint and, when available, the memory value to be written. In addition, each write has a flag that indicates whether the write has been propagated (passed to the memory) or not.
- **[exclusive] *successful_exclusive*, for store-exclusives, indicates whether it was previously guaranteed to succeed or made to fail.**
- information recording whether the instance is committed, finished, etc.

Read requests include their read kind and their memory footprint (their address and size), the as-yet-unsatisfied slices (the byte indices that have not been satisfied), and, for the satisfied slices, information about the write(s) that they were satisfied from. Write requests include their write kind, their memory footprint, and their value. When we refer to a write or read request without mentioning the kind of request we mean the request can be of any kind. A load instruction which has initiated (so its read request list *mem_reads* is not empty) and for which all its read requests are satisfied (i.e. there are no unsatisfied slices) is said to be *entirely satisfied*. **A load-exclusive is called *successful* if the first po-following store-exclusive that has not been made to fail has been guaranteed to succeed (as opposed to does not exist or has not been guaranteed to succeed or made to fail). The successful load-exclusive and the successful store-exclusive are said to be *paired*. If a successful load-exclusive has a read request that is mapped, in the exclusives map, to a write slice *ws*, we say the load-exclusive has an outstanding lock on *ws*.**

1.3 Thread States

The model state of a single hardware thread includes:

- *thread_id*, a unique identifier of the thread;
- *register_data*, the name, bit width, and start bit index for each register;
- *initial_register_state*, the initial register value for each register;
- *initial_fetch_address*, the initial fetch address for this thread;
- *instruction_tree*, a tree or list of the instruction instances that have been fetched (and not discarded), in program order.

1.4 Model Transitions

Fetch instruction A possible program-order successor of instruction instance *i* can be fetched from address *loc* if:

- (1) it has not already been fetched, i.e., none of the immediate successors of *i* in the thread's *instruction_tree* are from *loc*;
- (2) *loc* is a possible next fetch address for *i*:
 - (a) for a non-branch/jump instruction, the successor instruction address (*i.program_loc+4*);
 - (b) for an instruction that has performed a write to the program counter register (*_PC*), the value that was written;
 - (c) for a conditional branch, either the successor address or the branch target address¹; or
 - (d) for a jump to an address which is not yet determined, any address (this is approximated in our tool implementation, necessarily); and

¹In AArch64, all the conditional branch instructions have statically determined addresses.

(3) there is a decodable instruction in program memory at *loc*.

Note that this allows speculation past conditional branches and calculated jumps.

Action: construct a freshly initialized instruction instance *i'* for the instruction in the program memory at *loc*, including the static information available from the ISA model such as its *instruction_kind*, *regs_in*, and *regs_out*, and add *i'* to the thread's *instruction_tree* as a successor of *i*.

This involves only the thread, not the storage subsystem, as we assume a fixed program rather than modelling fetches with memory reads; we do not model self-modifying code.

Initiate memory reads of load instruction An instruction instance *i* with next state *READ_MEM(read_kind, address, size, read_cont)* can initiate the corresponding memory reads. Action:

- (1) Construct the appropriate read requests *rrs*:
 - if *address* is aligned to *size* then *rrs* is a single read request of *size* bytes from *address*;
 - otherwise, *rrs* is a set of *size* read requests, each of one byte, from the addresses *address*. . . *address+size-1*.
- (2) set *i.mem_reads* to *rrs*; and
- (3) update the state of *i* to *PENDING_MEM_READS read_cont*.

Satisfy memory read by forwarding from writes For a load instruction instance *i* in state *PENDING_MEM_READS read_cont*, and a read request, *r* in *i.mem_reads* that has unsatisfied slices, the read request can be partially or entirely satisfied by forwarding from unpropagated writes by store instruction instances that are po-before *i*, if the *read-request-condition* predicate holds. This is if:

- (1) all po-previous dmb sy and isb instructions are finished;
- (2) [dmb ld/ dmb st] all po-previous dmb ld instructions are finished;
- (3) [release/ acquire] if *i* is a load-acquire, all po-previous store-releases are finished; and
- (4) [release/ acquire] all non-finished po-previous load-acquire instructions are entirely satisfied.

Let *wss* be the maximal set of unpropagated write slices from store instruction instances po-before *i* (if *i* is a load-acquire, exclude store-exclusive writes), that overlap with the unsatisfied slices of *r*, and which are not superseded by intervening stores that are either propagated or read from by this thread. That last condition requires, for each write slice *ws* in *wss* from instruction *i'*:

- that there is no store instruction po-between *i* and *i'* with a write overlapping *ws*, and
- that there is no load instruction po-between *i* and *i'* that was satisfied from an overlapping write slice from a different thread.

Action:

- (1) update *r* to indicate that it was satisfied by *wss*; and
- (2) restart any speculative instructions which have violated coherence as a result of this, i.e., for every non-finished instruction *i'* that is a po-successor of *i*, and every read request *r'* of *i'* that was satisfied from *wss'*, if there exists a write slice *ws'* in *wss'*, and an overlapping write slice from a different write in *wss*, and *ws'* is not from an instruction that is a po-successor of *i*, restart *i'* and its data-flow dependents (including po-successors of load-acquire instructions).

Note that store-release writes cannot be forwarded to load-acquires: a load-acquire instruction cannot be satisfied before all po-previous store-release instructions are finished, and *wss* does not include writes from finished stores (as those must be propagated).

Satisfy memory read from memory For a load instruction instance *i* in state *PENDING_MEM_READS read_cont*, and a read request *r* in *i.mem_reads*, that has unsatisfied slices, the read request can be satisfied from memory if *i* is not a successful load-exclusive or no other successful load-exclusive from a different thread has an outstanding lock on the writes *r* is trying to read from.

If: the read-request-condition holds (see previous transition).

295 Action: let wss be the write slices from memory covering the unsatisfied slices of r , and apply
 296 the action of **Satisfy memory read by forwarding from writes**. In addition, if i is a successful
 297 load-exclusive, union wss with the set of write slices r is mapped to in the exclusives map.

298 Note that **Satisfy memory read by forwarding from writes** might leave some slices of the read
 299 request unsatisfied. **Satisfy memory read from memory**, on the other hand, will always satisfy all
 300 the unsatisfied slices of the read request.
 301

302 **Complete load instruction (when all its reads are entirely satisfied)** A load instruction in-
 303 stance i in state $PENDING_MEM_READS$ ($read_cont$) can be completed (not to be confused with finished)
 304 if all the read requests $i.mem_reads$ are entirely satisfied (i.e., there are no unsatisfied slices).

305 Action: update the state of i to $PLAIN$ ($read_cont$ ($memory_value$)), where $memory_value$ is assem-
 306 bled from all the write slices that satisfied $i.mem_reads$.

307 **Guarantee the success of store-exclusive** A store-exclusive instruction instance i with next
 308 state $EXCL_RES(res_cont)$ can be guaranteed to succeed if:

- 310 (1) the store-exclusive has not been made to fail (as recorded in $i.successful_exclusive$);
- 311 (2) assuming i is successful, it can be paired with a load-exclusive i' (see §1.2); and
- 312 (3) if i' has already been satisfied (not necessarily entirely), let wss be the set of propagated write
 313 slices i' has read from, then, no slice in wss has been overwritten (in memory) by a write from
 314 this thread, and no other successful load-exclusive from a different thread has an outstanding
 315 lock on a write slice from wss .

316 Action:

- 317 (1) record in $i.successful_exclusive$ that the store-exclusive will be successful;
- 318 (2) if i' has already been satisfied, union wss with the set of write slices the read request of i' is
 319 mapped to in the exclusives map, where wss is as above; and
- 320 (3) update the state of i to $PLAIN$ (res_cont ($true$)).
 321

322 **Make a store-exclusive fail** A store-exclusive instruction instance i with next state $EXCL_RES(res_con-$
 323 $tinuation)$ can be made to fail if the store-exclusive has not been guaranteed to succeed (as recorded
 324 in $i.successful_exclusive$) Action:

- 325 (1) record in $i.successful_exclusive$ that the store-exclusive was made to fail; and
- 326 (2) update the state of i to $PLAIN$ (res_cont ($false$)).
 327

328 Note the promise-success transition is enabled before the store-exclusive commits, and we do not
 329 require it to have a fully-determined address or to be non-restartable. As a result, a store-exclusive
 330 that has already promised its success might be restarted. Since other instructions may rely on
 331 its promise, the restart will not affect the value of $i.successful_exclusive$. Instead, when the store-
 332 exclusive is restarted it will take the same promise/failure transition as before its restart — based
 333 on the value of $i.successful_exclusive$.

334 **Initiate memory writes of store instruction, with their footprints** An instruction instance
 335 i with next state $WRITE_EA(write_kind, address, size, next_state')$ can announce its pending write
 336 footprint. Action:

- 337 (1) construct the appropriate write requests:
 - 338 • if $address$ is aligned to $size$ then ws is a single write request of $size$ bytes to $address$;
 - 339 • otherwise ws is a set of $size$ write requests, each of one byte size, to the addresses $ad-$
 340 $dress \dots address+size-1$.
- 341 (2) set $i.mem_writes$ to ws ; and
- 342 (3) update the state of i to $PLAIN$ $next_state'$.
 343

Note that at this point the write requests do not yet have their values. This state allows non-overlapping po-following writes to propagate.

Instantiate memory write values of store instruction An instruction instance i with next state $WRITE_MEMV(memory_value, write_cont)$ can initiate the corresponding memory writes. Action:

- (1) split $memory_value$ between the write requests $i.mem_writes$; and
- (2) update the state of i to $PENDING_MEM_WRITES write_cont$.

Commit store instruction For an uncommitted store instruction i in state $PENDING_MEM_WRITES write_cont$, i can commit if:

- (1) i has fully determined data (i.e., the register reads cannot change, see §1.5);
- (2) all po-previous conditional branch instructions are finished;
- (3) all po-previous dmb sy and isb instructions are finished;
- (4) $[\begin{smallmatrix} \text{dmb ld/} \\ \text{dmb st} \end{smallmatrix}]$ all po-previous dmb ld instructions are finished;
- (5) $[\begin{smallmatrix} \text{release/} \\ \text{acquire} \end{smallmatrix}]$ all po-previous load-acquire instructions are finished;
- (6) all po-previous store instructions, **except for store-exclusives that failed**, have initiated and so have non-empty mem_writes ;
- (7) $[\begin{smallmatrix} \text{release/} \\ \text{acquire} \end{smallmatrix}]$ if i is a store-release, all po-previous memory access instructions are finished;
- (8) $[\begin{smallmatrix} \text{dmb ld/} \\ \text{dmb st} \end{smallmatrix}]$ all po-previous dmb st instructions are finished;
- (9) all po-previous memory access instructions have a fully determined memory footprint; and
- (10) all po-previous load instructions have initiated and so have non-empty mem_reads .

Action: record i as committed.

Propagate memory write For an instruction i in state $PENDING_MEM_WRITES write_cont$, and an unpropagated write, w in $i.mem_writes$, the write can be propagated if:

- (1) all memory writes of po-previous store instructions that overlap w have already propagated
- (2) all read requests of po-previous load instructions that overlap with w have already been satisfied, and the load instruction is non-restartable (see §1.5);
- (3) all read requests satisfied by forwarding w are entirely satisfied; **and**
- (4) $[\text{exclusive}]$ **no successful load-exclusive from a different thread has an outstanding lock on a write slice that overlaps with w .**

Action:

- (1) restart any speculative instructions which have violated coherence as a result of this, i.e., for every non-finished instruction i' po-after i and every read request r' of i' that was satisfied from wss' , if there exists a write slice ws' in wss' that overlaps with w and is not from w , and ws' is not from a po-successor of i , restart i' and its data-flow dependents;
- (2) record w as propagated;
- (3) update the memory with w ; **and**
- (4) $[\text{exclusive}]$ **for every successful load-exclusive that has read from w (by forwarding), add the slices of w this load-exclusive read from to the set of write slices the read request of the load-exclusive is mapped to in the exclusives map.**

Complete store instruction (when its writes are all propagated) A store instruction i in state $PENDING_MEM_WRITES write_cont$, for which all the memory writes in $i.mem_writes$ have been propagated, can be completed. Action: update the state of i to $PLAIN(write_cont(true))$.

Commit barrier A barrier instruction i in state $PLAIN next_state$ where $next_state$ is $BARRIER(barrier_kind, next_state')$ can be committed if:

- (1) all po-previous conditional branch instructions are finished;

- 393 (2) $\left[\begin{smallmatrix} \text{dmb ld/} \\ \text{dmb st} \end{smallmatrix} \right]$ if i is a dmb ld instruction, all po-previous load instructions are finished;
 394 (3) $\left[\begin{smallmatrix} \text{dmb ld/} \\ \text{dmb st} \end{smallmatrix} \right]$ if i is a dmb st instruction, all po-previous store instructions are finished;
 395 (4) all po-previous dmb sy barriers are finished;
 396 (5) if i is an isb instruction, all po-previous memory access instructions have fully determined
 397 memory footprints; and
 398 (6) if i is a dmb sy instruction, all po-previous memory access instructions and barriers are
 399 finished.

400 Note that this differs from the previous Flowing and POP models: there, barriers committed in
 401 program-order and potentially re-ordered in the storage subsystem. Here the thread subsystem is
 402 weakened to subsume the re-ordering of Flowing's (and POP's) storage subsystem.

403 Action: update the state of i to *PLAIN next_state'*.

404
 405 **Register read** An instruction instance i with next state *READ_REG*(reg_name , $read_cont$) can do a
 406 register read if every instruction instance that it needs to read from has already performed the
 407 expected register write.

408 Let $read_sources$ include, for each bit of reg_name , the write to that bit by the most recent (in
 409 program order) instruction instance that can write to that bit, if any. If there is no such instruction,
 410 the source is the initial register value from *initial_register_state*. Let $register_value$ be the assembled
 411 value from $read_sources$. Action:

- 412 (1) add reg_name to $i.reg_reads$ with $read_sources$ and $register_value$; and
 413 (2) update the state of i to *PLAIN* ($read_cont(register_value)$).

414
 415 **Register write** An instruction instance i with next state *WRITE_REG*(reg_name , $register_value$,
 416 $next_state'$) can do the register write. Action:

- 417 (1) add reg_name to $i.reg_writes$ with $write_deps$ and $register_value$; and
 418 (2) update the state of i to *PLAIN next_state'*.

419 where $write_deps$ is the set of all $read_sources$ from $i.reg_reads$ and a flag that is set to true if i is a
 420 load instruction that has already been entirely satisfied.

421 **Pseudocode internal step** An instruction instance i with next state *INTERNAL*($next_state'$) can do
 422 that pseudocode-internal step. Action: update the state of i to *PLAIN next_state'*.

423
 424 **Finish instruction** A non-finished instruction i with next state *DONE* can be finished if:

- 425 (1) if i is a load instruction:
 426 (a) all po-previous dmb sy and isb instructions are finished;
 427 (b) $\left[\begin{smallmatrix} \text{dmb ld/} \\ \text{dmb st} \end{smallmatrix} \right]$ all po-previous dmb ld instructions are finished;
 428 (c) $\left[\begin{smallmatrix} \text{release/} \\ \text{acquire} \end{smallmatrix} \right]$ all po-previous load-acquire instructions are finished;
 429 (d) it is guaranteed that the values read by the read requests of i will not cause coherence
 430 violations, i.e., for any po-previous instruction instance i' , let cfp be the combined footprint
 431 of propagated writes from store instructions po-between i and i' and fixed writes that were
 432 forwarded to i from store instructions po-between i and i' including i' , and let cfp' be the
 433 complement of cfp in the memory footprint of i . If cfp' is not empty:
 434 (i) i' has a fully determined memory footprint;
 435 (ii) i' has no unpropagated memory write that overlaps with cfp' ; and
 436 (iii) If i' is a load with a memory footprint that overlaps with cfp' , then all the read requests
 437 of i' that overlap with cfp' are satisfied and i' can not be restarted (see §1.5).

438 Here a memory write is called fixed if it is the write of a non-exclusive-store instruction
 439 that has fully determined data.

- 440 (e) $\left[\begin{smallmatrix} \text{release/} \\ \text{acquire} \end{smallmatrix} \right]$ if i is a load-acquire, all po-previous store-release instructions are finished;
 441

- 442 (2) i has fully determined data; and
- 443 (3) all po-previous conditional branches are finished.

444 Action:

- 445 (1) if i is a branch instruction, discard any untaken path of execution, i.e., remove any (non-
- 446 finished) instructions that are not reachable by the branch taken in *instruction_tree*; and
- 447 (2) record the instruction as finished, i.e., set *finished* to *true*.

449 1.5 Auxiliary Definitions

450 **Fully determined** An instruction is said to have fully determined footprint if the memory reads
451 feeding into its footprint are finished: A register write w , of instruction i , with the associated
452 *write_deps* from *i.reg_writes* is said to be *fully determined* if one of the following conditions hold:

- 453 (1) i is finished; or
- 454 (2) the load flag in *write_deps* is *false* and every register write in *write_deps* is fully determined.

455 An instruction i is said to have *fully determined data* if all the register writes of *read_sources* in
456 *i.reg_reads* are fully determined. An instruction i is said to have a *fully determined memory footprint*
457 if all the register writes of *read_sources* in *i.reg_reads* that are associated with registers that feed
458 into i 's memory access footprint are fully determined.

459 **Restart condition** To determine if instruction i might be restarted we use the following recursive
460 condition: i is a non-finished instruction and at least one of the following holds,

- 461 (1) there exists an unpropagated write w such that applying the action of the **Propagate memory**
462 **write** transition to s will result in the restart of i ;
- 463 (2) there exists a non-finished load instruction l such that applying the action of the **Satisfy**
464 **memory read from memory** transition to l will result in the restart of i (even if l is already
465 entirely satisfied); or
- 466 (3) there exists a non-finished instruction i' that might be restarted and i is in its data-flow
467 dependents (**including po-successors of load-acquire instructions**).

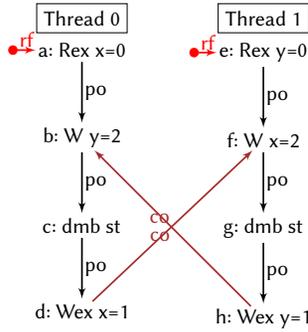
472 1.6 Remarks about load/store exclusive instructions

473 The MCA ARMv8 architecture intends that the success bit of store exclusives does not introduce
474 dependencies, to allow (e.g.) hardware optimisations that dynamically replace load/store exclusive
475 pairs by atomic read-modify-write operations that can execute in the memory subsystem and there-
476 fore be guaranteed to succeed. The ARMv8-axiomatic definition assumes all address/data/control
477 dependencies to be from reads, not writes. In the operational model, matching this weakness has
478 proved to be difficult: it means the operational model must be able to promise the success or failure
479 of a store-exclusive instruction even before any of its registers reads/writes have been done, so
480 before the store-exclusive's address and data are available. The early success promises are the
481 source of deadlocks in the operational model. To illustrate this consider, for example, the following
482 litmus test and a state where both a and e are satisfied and finished, and where b and f are not
483 propagated. Then d can promise its success, locking memory location x , and h can promise its
484 success, locking location y . But now there is a deadlock:

- 485 • For d to propagate c has to be committed and hence b propagated.
486 But b cannot propagate since y is locked.
- 487 • For h to propagate g has to be committed and hence f propagated.
488 But f cannot propagate since x is locked.

490

491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539



Similar situations arise from cases where there are other barriers or release/acquire instructions in-between the load and the store exclusive, or if the store exclusive has additional dependencies that the load exclusive does not have. These are cases that are not really intended to be supported by the architecture.

The model can also currently deadlock if a load and a store-exclusive are paired successfully but later turn out to have different addresses: if the store-exclusive promises its success before its address is known it locks the matched load-exclusive's memory location; when they later turns out to be to a different addresses it never unlocks it. This issue can be fixed, but it is currently still being clarified what exactly the architecturally allowed behaviour should be.

REFERENCES

Kathryn E. Gray, Gabriel Kerneis, Dominic Mulligan, Christopher Pulte, Susmit Sarkar, and Peter Sewell. 2015. An integrated concurrency and core-ISA architectural envelope definition, and test oracle, for IBM POWER multiprocessors. In *Proc. MICRO-48, the 48th Annual IEEE/ACM International Symposium on Microarchitecture*.