

# The Flowing and POP Models

(supplementary material for Modelling the ARMv8 Architecture, Operationally:  
Concurrency and ISA)

## 1. The Flowing Model in Detail

### 1.1 The Storage Subsystem/Thread Interface

The model is expressed in terms of read, write, and barrier requests. Read and write requests include the kind of the memory access (e.g. exclusive, release, acquire), the ID of the issuing thread and the memory access address. Write requests also include a value. Barrier requests include the kind of the barrier (e.g. *sy*, *ld*, *st*) and the issuing thread ID.

When we refer to a write or read request without mentioning the kind of request we mean the request can be of any kind.

The storage subsystem and a thread subsystem can exchange messages through synchronous transitions:

- a **non-exclusive** write request can be passed to the storage subsystem by a thread **Commit store instruction** transition coupled with a storage subsystem **Accept request** transition;
- **[exclusive]** an **exclusive** write request can be passed to the storage subsystem by a thread **Commit store-exclusive** transition coupled with a storage subsystem **Accept a successful write-exclusive** (from segment or memory) transition;
- a (memory) barrier request can be passed to the storage subsystem by a thread **Commit barrier instruction** transition coupled with a storage subsystem **Accept request** transition;
- a read request can be passed to the storage subsystem by a thread **Issue read request** transition coupled with a storage subsystem **Accept request** transition; and
- a read response can be returned from the storage subsystem to a thread by a storage subsystem **Satisfy read from segment** or **Satisfy read from memory** transition coupled with a thread **Satisfy memory read from storage response** transition.

### 1.2 Storage Subsystem States

The Flowing storage subsystem state includes:

- *thread\_ids*, the set of thread IDs that exist in the system;
- *topology*, a data structure describing how the segments are connected to each other;
- *thread\_to\_segment*, a map from thread IDs to segments, associating each thread with its leaf segment;
- *buffers*, a map from segments to list of requests associating each segment with the requests queued in that segment;
- *reordered*, a set of request pairs that have been reordered w.r.t. each other; and
- *memory*, a map from memory addresses to the most recent write request to that address to reach memory.

### 1.3 Storage Subsystem Transitions

**Accept request** A request *r* from thread *r.tid* can be accepted if:

1. *r* has not been accepted before (i.e. *r* is not in *buffers*); and
2. *r.tid* is in *thread\_ids*.

Action: add *r* to the top of *buffers(thread\_to\_segment(r.tid))*.

**Flow request** A request *r* can flow from segment *s1* to *s2* if:

1. *r* is at the bottom of *buffers(s1)*;
2. *s1* is a child of *s2* in *topology*;
3. **[exclusive]** if *r* is a write to address *r.addr*, *s1* is not blocked (§1.4) by any write-exclusive to *r.addr*; and
4. **[exclusive]** if *r.kind* is *dmb sy* or *dmb st*, *s1* is not blocked by any write-exclusive.

Action:

1. remove *r* from *buffers(s1)*;
2. add *r* to the top of *buffers(s2)*; and
3. remove from *reordered* any pair that contains *r*.

**Reorder requests** Two requests *r\_new*, *r\_old* that appear consecutively in *buffers(s)* (*r\_new* nearer the top) can be reordered if:

1. (*r\_new*, *r\_old*) does not appear in *reordered* (i.e. they have not been reordered (in segment *s*) with each other before (preventing live lock)); and
2. *r\_new* and *r\_old* satisfy the *reorder condition* (§1.4).

Action:

1. switch the positions of *r\_new* and *r\_old* in *buffers(s)*; and
2. record the reordering of *r\_new* and *r\_old* (by adding the pair (*r\_new*, *r\_old*) to *reordered*).

**Satisfy read from segment** Two requests, *r\_read*, *r\_write*, can generate a read response to thread *r\_read.tid* if:

1. *r\_read* is a read request that has not been satisfied yet;
2. *r\_write* is a write request;
3. *r\_read*, *r\_write* appear consecutively in *buffers(s)* for some segment *s*, with *r\_read* closer to the top (newer);
4. *r\_read* and *r\_write* are to the same address; and
5. **[release/acquire]** if *r\_read.kind* is *read-acquire*, *r\_write.kind* is not *write-release*.

Action:

1. send a read response for request *r\_read* to thread *r\_read.tid*, containing *r\_write*; and
2. **[release/acquire]** if *r\_read.kind* is *read-acquire*, switch the positions of *r\_read* and *r\_write* in *buffers(s)* and mark *r\_read* as satisfied; else
3. remove *r\_read*.

**Satisfy read from memory** A read request *r\_read* from thread *r\_read.tid* can generate a read response to thread *r\_read.tid* if *r\_read* is at the bottom of *buffers(s)*, where *s* is the root segment in *topology*, and *r\_read* has not been satisfied yet. Action:

1. send a read response for request *r\_read* to thread *r\_read.tid*, containing the write stored in *memory* for the address *r\_read.addr*; and
2. remove *r\_read*.

**Flow satisfied read to memory** A satisfied read acquire request  $r\_read$  can be discarded if:  $r\_read$  is at the bottom of  $buffers(s)$ , where  $s$  is the root segment in  $topology$ . Action: remove  $r\_read$ .

**Flow write to memory** The write request  $r\_write$  can be stored into memory if:  $r\_write$  is at the bottom of  $buffers(s)$ , where  $s$  is the root segment in  $topology$ . Action:

1. update  $memory$  to map the address  $r\_write.addr$  to  $r\_write$ ; and
2. remove  $r\_write$  from  $buffers(s)$  and  $reordered$ .

**Flow barrier to memory** A barrier request  $r\_barr$  can be discarded if:  $r\_barr$  is at the bottom of  $buffers(s)$ , where  $s$  is the root segment in  $topology$ . Action: remove  $r\_barr$ .

**Accept a successful write-exclusive** (from segment) A write-exclusive request  $r\_write$  from thread  $r\_write.tid$  with an accompanying write request  $r\_write'$  that was read by a po-previous load-exclusive can be accepted and succeed if:

1. all the conditions of **Accept request** are satisfied for  $r\_write$ ;
2.  $r\_write$  and  $r\_write'$  are to the same address;
3. exists  $0 \leq i \leq n$  such that  $r\_write'$  in  $buffers(s_i)$ ;
4. for all  $i < j \leq n$ , segment  $s_j$  is not blocked by a write-exclusive to the address  $r\_write.addr$ ;
5. let  $rs$  be the set of requests  $\bigcup_{i < j \leq n} buffers(s_j)$  unioned with all the requests in  $buffers(s_i)$  above  $r\_write'$ ;
6. there is no write request in  $rs$  to the same address as  $r\_write$  from a different thread;
7. there is no dmb sy barrier request in  $rs$ ; and
8.  $[ \begin{smallmatrix} dmb\ ld/ \\ dmb\ st \end{smallmatrix} ]$  there is no dmb st barrier request in  $rs$ .

where  $s_0, s_1, \dots, s_n$  is the path of segments in  $topology$  from the root segment  $s_0$  to the leaf segment  $s_n$  associated with thread  $r\_write.tid$ . Action: the same as **Accept request**.

**Accept a successful write-exclusive** (from memory) A write-exclusive request  $r\_write$  from thread  $r\_write.tid$  with an accompanying write request  $r\_write'$  that was read by a po-previous read-exclusive can be accepted and succeed if:

1. all the conditions of **Accept request** are satisfied for  $r\_write$ ;
2.  $r\_write$  and  $r\_write'$  are to the same address;
3.  $r\_write'$  is in  $memory$ ;
4. for all  $0 < j \leq n$ , segment  $s_j$  is not blocked by a write-exclusive to the same address as  $r\_write$ ;
5. there is no write request in  $rs$  to the same address as  $r\_write$  from a different thread;
6. there is no dmb sy barrier request in  $rs$ ; and
7.  $[ \begin{smallmatrix} dmb\ ld/ \\ dmb\ st \end{smallmatrix} ]$  there is no dmb st barrier request in  $rs$ .

where  $s_0, s_1, \dots, s_n$  is the path of segments in  $topology$  from the root segment  $s_0$  to the leaf segment  $s_n$  associated with thread  $r\_write.tid$ , and let  $rs$  be the set of requests  $\bigcup_{0 \leq j \leq n} buffers(s_j)$ . Action: the same as **Accept request**.

#### 1.4 Auxiliary Definitions for Storage Subsystem

**Immediate predecessor** Intuitively, the *immediate predecessor* of a write-exclusive is the write the associated load-exclusive read from, which must be kept coherence-immediately-before the write-exclusive.

Formally we say a write  $w'$  in segment  $s_{w'}$  (or memory) is the *immediate predecessor* of a write-exclusive  $w$  in segment  $s_w$  if  $s_{w'}$  is a predecessor of  $s_w$  (in  $topology$ ),  $w'$  and  $w$  are to the same address, and all the writes queued between  $w$  and  $w'$  are to different addresses than that. Notice that a write-exclusive can have at most one immediate predecessor (in fact we will make sure it has exactly one immediate predecessor).

As write-exclusive must be the coherence-immediate-successor of its immediate predecessor, no other write to the same address

can be allowed to flow in between them. We say a write-exclusive  $w$  in segment  $s_w$  is blocking segment  $s$  if  $s_w$  is a successor of  $s$ 's parent, but is not a successor of  $s$  (or  $s$  itself), and the immediate predecessor of  $w$  is in a segment that is predecessor of  $s$ .

**Reorder condition** Two requests  $r\_new$  and  $r\_old$  are said to meet the *reorder condition* if:

1.  $[ \begin{smallmatrix} release/ \\ acquire \end{smallmatrix} ]$  if  $r\_new$  is a read-acquire and  $r\_old$  is a write-release, they are to a different address;
2.  $[ \begin{smallmatrix} release/ \\ acquire \end{smallmatrix} ]$   $r\_new$  is not a write-release;
3.  $[ \begin{smallmatrix} release/ \\ acquire \end{smallmatrix} ]$  if  $r\_old$  is a read-acquire then  $r\_new$  and  $r\_old$  originated from different threads;
4.  $[ \begin{smallmatrix} release/ \\ acquire \end{smallmatrix} ]$  if  $r\_new$  is a read-acquire and  $r\_old$  is a write-release,  $r\_new$  and  $r\_old$  originated from different threads;
5.  $[ \begin{smallmatrix} dmb\ ld/ \\ dmb\ st \end{smallmatrix} ]$  at least one of  $r\_new, r\_old$  is not a memory barrier (of any kind);
6.  $[ \begin{smallmatrix} dmb\ ld/ \\ dmb\ st \end{smallmatrix} ]$  if  $r\_new$  is dmb ld then  $r\_old$  is not a read from the same thread;
7.  $[ \begin{smallmatrix} dmb\ ld/ \\ dmb\ st \end{smallmatrix} ]$  if  $r\_old$  is dmb ld then  $r\_new$  is not a memory access from the same thread;
8.  $[ \begin{smallmatrix} dmb\ ld/ \\ dmb\ st \end{smallmatrix} ]$  if  $r\_new$  is dmb st then  $r\_old$  is not a write from the same thread;
9.  $[ \begin{smallmatrix} dmb\ ld/ \\ dmb\ st \end{smallmatrix} ]$  if  $r\_old$  is dmb st then  $r\_new$  is not a write;
10. neither one of  $r\_new, r\_old$  is a dmb sy; and
11. if both  $r\_new$  and  $r\_old$  are memory access requests, and none of them is a satisfied read, they are to different addresses.

#### 1.5 Thread States

The state of a single hardware thread includes:

- $thread\_id$ , a unique identifier of the thread;
- $register\_data$ , general information about the available registers, including name, bit width, initial bit index and direction;
- $initial\_register\_state$ , the initial values for each register;
- $initial\_fetch\_address$ , the initial fetch address for this thread;
- $instruction\_tree$ , a data structure holding the instructions that have been fetched in program order; and
- $read\_issuing\_order$ , a list of read requests in the order they were issued to the storage subsystem.

#### 1.6 Instruction State

Each instruction in the  $instruction\_tree$  has a state including:

- $program\_loc$ , the memory address where the instruction's op-code resides;
- $instruction\_kind$ , the kind of the instruction (e.g. load-acquire);
- $regs\_in$ , the input registers, for dependency calculation;
- $regs\_out$ , the output registers, for dependency calculation;
- $reg\_reads$ , accumulated register reads;
- $reg\_writes$ , accumulated register writes;
- $mem\_read$ , one of *none*, *pending a*, *requested a*, *write\_read from w*, where  $a$  is a memory address and  $w$  is a memory write;
- $mem\_write$ , one of *none*, *potential\_address a*, *pending w*, *committed w*, where  $a$  is a memory address and  $w$  is a memory write;
- $committed$ , set to *true* when the instruction can no longer affect the memory model (the Sail interpreter might still need to complete the execution of the ISA definition);
- $finished$ , set to *true* only after  $committed$  is set to *true* and the Sail interpreter has finished executing the ISA definition; and
- $micro\_op\_state$ , the meta-state of the instruction, one of *plain interpreter\_state* (ready to make a Sail interpreter transition from *interpreter\_state*), *pending\_mem\_read read\_cont* (ready to perform a read from memory) or *potential\_mem\_write write\_cont* (ready to perform a write to memory) where *read\_cont* is a Sail interpreter continuation that given the

value read from memory returns the next interpreter state, and *write\_cont* is a Sail interpreter continuation that given a boolean value (that indicates if the store was successful) returns the next interpreter state.

Instructions that have been fetched (i.e. in *instruction\_tree*) and have a *finished* value of *false* are said to be *in-flight*. Instructions that have a *committed* value of *true* are said to be *committed*. Load instructions with the value *requested* for *mem\_read* are said to have an *outstanding read request*, and if they have the value *write\_read\_from* they are said to be *satisfied*. Store instructions with *mem\_write = pending w* are said to have a *pending write*. We say instruction *i* has *fully determined address* if all instructions that write to input registers of *i* that affect memory address calculation in the ISA definition of *i* are committed. We say *i* is *fully determined* if all instructions that write to input registers of *i* are committed.

## 1.7 Thread Transitions

**Sail interpreter step** An instruction *i* in meta-state *plain interpreter\_state* can perform an interpreter step as follows:

- if *interpreter\_state* indicates a memory-read event: set *i.mem\_read* to *pending interpreter\_state.read\_addr* and update the instruction meta-state to *pending\_mem\_read interpreter\_state.read\_cont*;
- if *interpreter\_state* indicates a memory-write-address event: set *mem\_write* to *potential\_address interpreter\_state.write\_addr* and update the instruction meta-state to *plain interpreter\_state.next*;
- if *interpreter\_state* indicates a memory-write-value event: set *mem\_write* to *pending w* (where *w* is a write request with the value *interpreter\_state.write\_value* and the address that was in *mem\_write* before) and update the instruction meta-state to *potential\_mem\_write interpreter\_state.write\_cont*;
- if *interpreter\_state* indicates a register-read event: look in *instruction\_tree* for the most recent po-previous instruction, *i'* that has *interpreter\_state.reg* in *i'.regs\_out*. If *interpreter\_state.reg* is not in *i'.reg\_writes* the transition is disabled. Otherwise, add *interpreter\_state.reg* to *i.reg\_reads* and update the meta-state of *i* to *plain (interpreter\_state.reg\_count val)* (where *val* is the value written to the register by *i'*);
- if *interpreter\_state* indicates a register-write event: add *interpreter\_state.reg* to *i.reg\_writes* and update the meta-state of *i* to *plain interpreter\_state.next*; and
- if *interpreter\_state* indicates an internal step: update the meta-state of *i* to *plain interpreter\_state.next*.

**Fetch instruction** An instruction *i*, from address *loc*, can be fetched, following its program-order predecessor *i'* if:

1. *i'* is in *instruction\_tree*;
2. *loc* is a possible next fetch address for *i'* according to the ISA model;
3. none of the successors of *i'* in *instruction\_tree* are from *loc*; and
4. *i* is the instruction of the program at *loc*.

The possible next fetch addresses allow speculation past calculated jumps and conditional branches; they are defined as:

1. for a non-branch/jump instruction, the successor instruction address;
2. for a jump to constant address, that address;
3. for a conditional branch, the possible addresses for a jump<sup>1</sup> together with the successor; and
4. for a jump to an address which is not yet fully determined (i.e., where there is an uncommitted instruction with a dataflow path

<sup>1</sup>In AArch64, all the conditional branch instructions have a constant address.

to the address), any address. This is (necessarily) approximated in our implementation, c.f. §4.5.

Action: construct an initialized instruction instance, including static information available from the ISA model such as *instruction\_kind*, *regs\_in*, *regs\_out*, and add it to *instruction\_tree* as a successor of *i'*.

This is an internal action of the thread, not involving the storage subsystem, as we assume a fixed program rather than modelling fetches with memory reads; we do not model self-modifying code.

**Issue read request** An in-flight instruction *i* in meta-state *pending\_mem\_read read\_cont* can issue a read request of address *a* to the storage subsystem if:

1. *i.mem\_read* has the value *pending a*, i.e., any other reads with dataflow path to the address have been satisfied, though not necessarily committed, and any arithmetic on such a path completed;
2. all po-previous dmb sy and isb instructions are committed;
3. [ *dmb ld/ st* ] all po-previous dmb ld instructions are committed;
4. [ *release/ acquire* ] if *i* is load-acquire instruction, all po-previous store-release instructions are committed; and
5. [ *release/ acquire* ] all in-flight po-previous load-acquire instructions have outstanding read requests or were already satisfied.

Action:

1. send a read-request to the storage subsystem;
2. change *i.mem\_read* to *requested a*; and
3. update *read\_issuing\_order* to note that the read was issued last.

**Satisfy memory read from storage response** A read response with write *w*, for read request from instruction *i* in meta-state *pending\_mem\_read read\_cont* can always be received. Action:

1. if *i.mem\_read* does not have the value *requested* the response is ignored (this can happen when the read is satisfied by write forwarding while waiting for the response); else
2. if there exists a po-previous load instruction to the same address that was issued after *i* (i.e., issued out of order) and was satisfied by a write that is not *w*, set *i.mem\_read* to *pending*; else
3. for every in-flight instruction *i'* that is po-after *i* and has read from a write to the same address as *i* that is not *w* and not po-successor of *i*, restart *i'* and its data flow dependents;
4. change *i.mem\_read* to *write\_read\_from w*; and
5. update the meta-state of *i* to *plain (read\_cont w.value)*.

**Satisfy memory read by forwarding an in-flight write** A pending memory write *w* from an in-flight store instruction *i'* can be forwarded directly to a load instruction *i* in meta-state *pending\_mem\_read read\_cont* if:

1. [ *release/ acquire* ] *i* is not load-acquire;
2. [ *exclusive* ] *i* is not load-exclusive;
3. *i.mem\_read* has the value *pending w.address* or *requested w.address*; and
4. *i'* is po-before *i*, there is no other store instruction to the same address in between them, and there is no other load instruction in between them that has read from a different store instruction to the same address.

Action:

1. for every in-flight instruction *i''* that is po-after *i* and has read from a write to the same address as *i* that is not *w* and not po-successor of *i*, restart *i''* and its data flow dependents;
2. change *i.mem\_read* to *write\_read\_from w*; and
3. update the meta-state of *i* to *plain (read\_cont w.value)*.

**Commit store instruction** A store (not exclusive) instruction *i* in meta-state *potential\_mem\_write write\_cont* and with pending write *w* can be committed if:

1. *i* is fully determined;

2. all po-previous conditional branches are committed;
3. all po-previous dmb sy and isb instructions are committed;
4.  $[ \text{dmb ld/} / \text{dmb st} ]$  all po-previous dmb ld instructions are committed;
5.  $[ \text{release/} / \text{acquire} ]$  all po-previous load-acquire instructions are committed;
6.  $[ \text{release/} / \text{acquire} ]$  all po-previous store-release instructions to the same address are committed;
7.  $[ \text{release/} / \text{acquire} ]$  if  $i$  is a store-release, all po-previous memory access instructions are committed;
8.  $[ \text{exclusive} ]$  all po-previous store-exclusive to the same address are committed;
9.  $[ \text{dmb ld/} / \text{dmb st} ]$  all po-previous dmb st are committed;
10. all po-previous memory access instructions have a fully determined address; and
11. all po-previous instructions that read from the same address have either an outstanding read request or are satisfied, and cannot be restarted (see §1.8).

Action:

1. restart any in-flight loads (and their dataflow dependants) that:
  - (a) are po-after  $i$  and have read from the same address, but from a different write and where the read could not have been by forwarding an in-flight write that is po-after  $i$ ; or
  - (b) are po-after  $i$ , are to the same address, and have issued a read request that has not been satisfied yet.
2. if there is no committed po-following store to the same address, send a write request to the storage subsystem;
3. record the store as committed (i.e. set  $i.mem\_write$  to *committed*  $w$  and *committed* to *true*); and
4. update the meta-state of  $i$  to *plain* ( $write\_cont$  *true*).

**Commit failed store-exclusive instruction** A store-exclusive instruction  $i$  in meta-state *potential\_mem\_write write\_cont* can always be committed as “failed”. Action:

1. restart any in-flight loads (and their dataflow dependants) that were satisfied by a write forward from  $i$ ;
2. record the store as committed (i.e. set  $i.mem\_write$  to *none* and set *committed* to *true*); and
3. update the meta-state of  $i$  to *plain* ( $write\_cont$  *false*).

**Commit store-exclusive** A store-exclusive instruction  $i$  in meta-state *potential\_mem\_write write\_cont* with pending write  $w$  and an associated load-exclusive instruction  $i'$  with  $i'.mem\_read = write\_read\_from\ w'$ , can be committed if:

1.  $i'$  is committed and po-before  $i$ ;
2. there is no other store-exclusive or load-exclusive instructions po-between  $i'$  and  $i$ ; and
3. all the conditions of **Commit store instruction** are met.

If  $i'$  is to the same address as  $i$  the following action is applied, otherwise the action of **Commit store instruction** is applied and the destination register is set to 0. Action:

1. send a write-exclusive request (with  $w$  as the requested write and  $w'$  as the required coherence predecessor) to the storage subsystem and receive the corresponding fail/success response, setting the destination register accordingly (0 for success and 1 for fail);
2. if the write succeeded, perform the action of **Commit store instruction**; else
3. apply the action of **Commit failed store-exclusive instruction**.

**Commit barrier instruction** A barrier instruction  $i$  in meta-state *plain interpreter\_state* can be committed if:

1. *interpreter\_state* indicates a barrier event is pending;
2. all po-previous conditional branches are committed;
3.  $[ \text{dmb ld/} / \text{dmb st} ]$  if  $i$  is dmb ld, all po-previous load instructions are committed;

4.  $[ \text{dmb ld/} / \text{dmb st} ]$  if  $i$  is dmb st, all po-previous store instructions are committed;
5. all po-previous barriers (of any kind) are committed;
6. if  $i$  is isb, all po-previous memory access instructions have a fully determined address; and
7. if  $i$  is dmb sy, all po-previous memory access instructions are committed.

Action:

1. record the barrier as committed (i.e. set *committed* to *true*);
2. send a barrier request to the storage subsystem; and
3. update the meta-state of  $i$  to *plain interpreter\_state.next*.

**Finish instruction** An in-flight instruction  $i$  in meta-state *plain interpreter\_state* can be finished if:

1. *interpreter\_state* indicates the execution of the ISA definition has finished;
2. if  $i$  is a load instruction:
  - (a) all po-previous dmb sy and isb instructions are committed;
  - (b)  $[ \text{dmb ld/} / \text{dmb st} ]$  all po-previous dmb ld instructions are committed;
  - (c)  $[ \text{release/} / \text{acquire} ]$  all po-previous load-acquire instructions are committed;
  - (d) let  $i'$  be the store instruction to the same address as  $i$  that appears last in program order before  $i$ :
    - i. if  $i'$  was forwarded to *and is not a store-exclusive*,  $i'$  is fully determined, otherwise  $i'$  is committed;
    - ii. all memory access instructions, po-between  $i'$  and  $i$ , have a fully determined address; and
    - iii. all load instructions to the same address as  $i$ , that are po-between  $i'$  and  $i$ , are committed.
  - (e)  $[ \text{release/} / \text{acquire} ]$  if  $i$  is a load-acquire, all po-previous store-release instructions are committed;
3.  $i$  is fully determined; and
4. all po-previous conditional branches are committed.

Action:

1. if  $i$  is a branch instruction, abort any untaken path of execution, i.e., any in-flight instruction that are not reachable by the branch taken in *instruction\_tree*; and
2. record the instruction as finished, i.e., set *finished* (and *committed*) to *true*.

## 1.8 Auxiliary Definitions for Thread Subsystem

**Restart condition** To determine if instruction  $i$  might be restarted we use the following recursive condition:  $i$  is an in-flight instruction and at least one of the following holds,

1. there exists an in-flight store instruction  $s$  such that applying the action of the **Commit store instruction** transition to  $s$  will result in the restart of  $i$ ;
2. there exists an in-flight load instruction  $l$  such that applying the action of the **Satisfy memory read from storage response** transition to  $l$  will result in the restart of  $i$  (even if  $l$  is already satisfied);
3.  $i$  has an outstanding read request that has not been satisfied yet, and there exists a load instruction po-before  $i$  that has an outstanding read request to the same address (maybe already satisfied) after  $i$ ; or
4. there exists an in-flight instruction  $i'$  that might be restarted and an output register of  $i'$  feeds an input register of  $i$ , or  $i'$  is a load-acquire and  $i$  is a load and  $i'$  is po-before  $i$ .

## 2. The POP Model

### 2.1 The Storage Subsystem/Thread Interface

The storage subsystem and a thread subsystem can exchange messages through synchronous transitions:

- a **non-exclusive** write request can be passed to the storage subsystem by a thread **Commit store instruction** transition coupled with a storage subsystem **Accept request** transition;
- [ **exclusive** ] an **exclusive** write request can be passed to the storage subsystem by a thread **Commit store-exclusive** transition coupled with a storage subsystem **Accept a successful write-exclusive** transition;
- a (memory) barrier request can be passed to the storage subsystem by a thread **Commit barrier instruction** transition coupled with a storage subsystem **Accept request** transition;
- a read request can be passed to the storage subsystem by a thread **Issue read request** transition coupled with a storage subsystem **Accept request** transition; and
- a read response can be passed from the storage subsystem to a thread by a storage subsystem **Send read-response to thread** transition coupled with a thread **Satisfy memory read from storage response** transition.

In addition to the above, when a load instruction is restarted in the thread subsystem, all its read-requests are removed from the storage subsystem.

## 2.2 Storage Subsystem State

The POP storage subsystem state includes:

- *thread\_ids*, the set of thread IDs that exist in the system;
- *requests\_seen*, the set of requests (memory read/write requests and barrier requests) that have been seen by the subsystem;
- *order\_constraints*, the set of pairs of requests from *requests\_seen*. The pair (*r\_old*, *r\_new*) indicates that *r\_old* is before *r\_new* (*r\_old* and *r\_new* might be to different addresses and might even be of different kinds); and
- *requests\_propagated\_to*, a map from thread IDs to subsets of *requests\_seen*, associating with each thread the set of requests that has propagated (potentially visible) to it.

## 2.3 Storage Subsystem Transitions

**Accept request** A request *r\_new*, that is not store-exclusive, from thread *r\_new.tid* can be accepted if:

1. *r\_new* has not been accepted before (i.e., *r\_new* is not in *requests\_seen*); and
2. *r\_new.tid* is in *thread\_ids*.

Action:

1. add *r\_new* to *requests\_seen*;
2. add *r\_new* to *requests\_propagated\_to*(*r\_new.tid*); and
3. update *order\_constraints* to note that *r\_new* is after every request *r\_old* that has propagated to thread *r\_new.tid*, and *r\_new* and *r\_old* do not meet the Flowing *reorder condition* (see **Reorder condition**).

**Propagate request to another thread** The storage subsystem can propagate request *r* (by thread *tid*) to another thread *tid'* if:

1. *r* has been seen before (i.e., *r* is in *requests\_seen*);
2. *r* has not yet been propagated to thread *tid'*;
3. every request that is before *r* in *order\_constraints* has already been propagated to thread *tid'*;
4. [ **exclusive** ] if *r* is a write, all the write exclusive requests that have propagated to *tid'* but not to *tid*, for which the immediate predecessor write to the same address is *order\_constraints-before r*, are to a different address; and
5. [ **exclusive** ] if *r* is a dmb sy or dmb st, there are no write exclusive requests that have propagated to *tid'* but not to *tid* for which the immediate predecessor write to the same address is *order\_constraints-before r*.

Action:

1. add *r* to *requests\_propagated\_to*(*tid'*); and
2. update *order\_constraints* to note that *r* is before every request *r\_new* that has propagated to thread *tid'* but not to thread *tid*, where *r\_new* and *r* do not meet the Flowing *reorder condition* (see **Reorder condition**) and are not already ordered.

**Send read-response to thread** The storage subsystem can send a read-response for read request *r\_read* that has not been satisfied yet to thread *r\_read.tid* containing the write request *r\_write* if:

1. *r\_write* and *r\_read* are to the same address;
2. *r\_write* and *r\_read* have been propagated to (exactly) the same threads;
3. *r\_write* is *order\_constraints-before r\_read*;
4. any request that is *order\_constraints-between r\_write* and *r\_read* has been *fully-propagated* (§2.4) and is to a different address; and
5. [ **release/acquire** ] if *r\_read* is a read-acquire and *r\_write* is a write-release then *r\_write* must be *fully-propagated*.

Action:

1. send thread *r\_read.tid* a read-response for *r\_read* containing *r\_write*;
2. [ **release/acquire** ] if *r\_read.kind* is read-acquire and there are no other requests between *r\_read* and *r\_write* in *order\_constraints*, switch the positions of *r\_read* and *r\_write* in *order\_constraints* and mark *r\_read* as satisfied; else
3. remove *r\_read*.

**Accept a successful write-exclusive** A write-exclusive request *r\_write* from thread *r\_write.tid* with an accompanying write request *r\_write'* that was read by a po-previous load-exclusive can be accepted and succeed if:

1. *r\_write* meets the condition of **Accept request**;
2. *r\_write'* is in *requests\_seen*;
3. the write *r\_write* is to the same address as *r\_write'*;
4. *r\_write'* is not the immediate predecessor of a different write-exclusive;
5. every immediate predecessor write that is *order\_constraints-after r\_write'* and has propagated to thread *r\_write.tid* is to a different address than *r\_write*; and
6. for every request *r* that has been propagated to thread *r\_write.tid* and is *order\_constraints-after r\_write'*:
  - (a) if *r.tid* is different then *r\_write.tid* then *r* is not a write to the same address as *r\_write*; and
  - (b) if *r* is a dmb sy or dmb st then *r* is *fully-propagated*.

Action: the same as **Accept request**.

## 2.4 Auxiliary Definitions for Storage Subsystem

**Fully propagated** Request *r* is said to be fully-propagated if it has been propagated to all threads and so has every request that is *order\_constraints-before* it.

**Removing read request** When a read request is removed from the storage subsystem, due to restart of the instruction in the thread subsystem or satisfaction, first *order\_constraints* is restricted to exclude the request; it is then further restricted by applying the *reorder condition* to each pair of ordered requests and removing pairs that can be reordered; finally the transitive closure of the result is calculated.

**Store-exclusive immediate predecessor** The following is an important invariant that guarantees the correct behaviour of load/store-exclusive. Every write-exclusive has exactly one immediate predecessor write to the same address in *order\_constraints*; moreover, that immediate predecessor has exactly one immediate successor write to the same address in *order\_constraints* (namely the write-exclusive).