

6.823 Computer System Architecture 6.823 Directory-based Cache Coherence Protocol

<http://csg.csail.mit.edu/6.823/>

In this handout, we describe the 4-hop directory-based **MESI** protocol with **silent evictions from S**. This means that when a cache line is evicted from S, it transitions to the I state without notifying the directory. Assume that given a sender-receiver pair, messages will be delivered to the receiver in the order they were sent.

Cache State:

- **Modified (M):** The cache has the only valid copy of the line with read and write permissions. No other cache may have a valid copy of the line.
- **Exclusive (E):** The cache has the only valid copy of the line with read-only permissions. No other cache may have a valid copy of the line.
- **Shared (S):** The cache has a shared, read-only copy of the line. Other caches may also have read-only copies.
- **Invalid (I):** Data is not present in this cache but can be present in other caches.
- **Transient states** ($I \rightarrow S$, $I \rightarrow M$, etc.): The cache has sent a request to the directory and is waiting for the response.

Directory State:

For each memory address, the directory maintains its coherence state and a sharer set:

- **Uncached (Un):** No cache has a valid copy.
- **Shared (Sh):** One or more caches are in the S or E states.
- **Exclusive (Ex):** One of the caches is in the M state (not to be confused with the E state of the cache).
- **Transient states** ($Ex \rightarrow Sh$, etc.): The directory has sent a downgrade request to a cache (or has sent an invalidate request to one or multiple caches) and is waiting for the cache response(s).
- **Sharer set:** Contains the IDs of the caches that *may* have shared or exclusive permissions for that memory location. Since we allow silent evictions, some "sharers" in the sharer set may have invalidated their copies. For example, sharers = {0, 1} means that Cache 0 and Cache 1 may potentially have shared copies. In practice, this can be implemented using a bit vector.

Messages:

- Directory to cache:

Message	Meaning
<InvReq, K, A>	The directory asks cache K to <i>invalidate</i> cache block A. If the cache block is in state M, it sends back its dirty data and transitions to I. If the cache was in E or S, it does not send back its data and transitions to I.
<DownReq, K, A>	The directory asks cache K to <i>downgrade</i> cache block. If the cache block is in M state, it sends back its dirty data and change its state from M to S. If the cache was in E, it does not send back its data and transitions to S.
<ExResp, K, A>	The directory grants cache K exclusive permission (M) for cache block A
<ShResp, K, A>	The directory grants cache K shared permission (S) for cache block A
<NoSharerShResp, K, A>	The directory grants cache K shared permission (E) for cache block A
<WbResp, K, A>	The directory acknowledges cache K's writeback request for cache block A

- Cache to directory:

Message	Meaning
<ExReq, K, A>	Cache K requests exclusive permission (M) for cache block A
<ShReq, K, A>	Cache K requests shared permission (either S or E) for cache block A
<WbReq, K, A>	Cache K sends a writeback request for cache block A, which it needs to evict
<InvResp, K, A>	Cache K notifies the directory that it has invalidated cache block A
<DownResp, K, A>	Cache K notifies the directory that it has downgraded cache block A

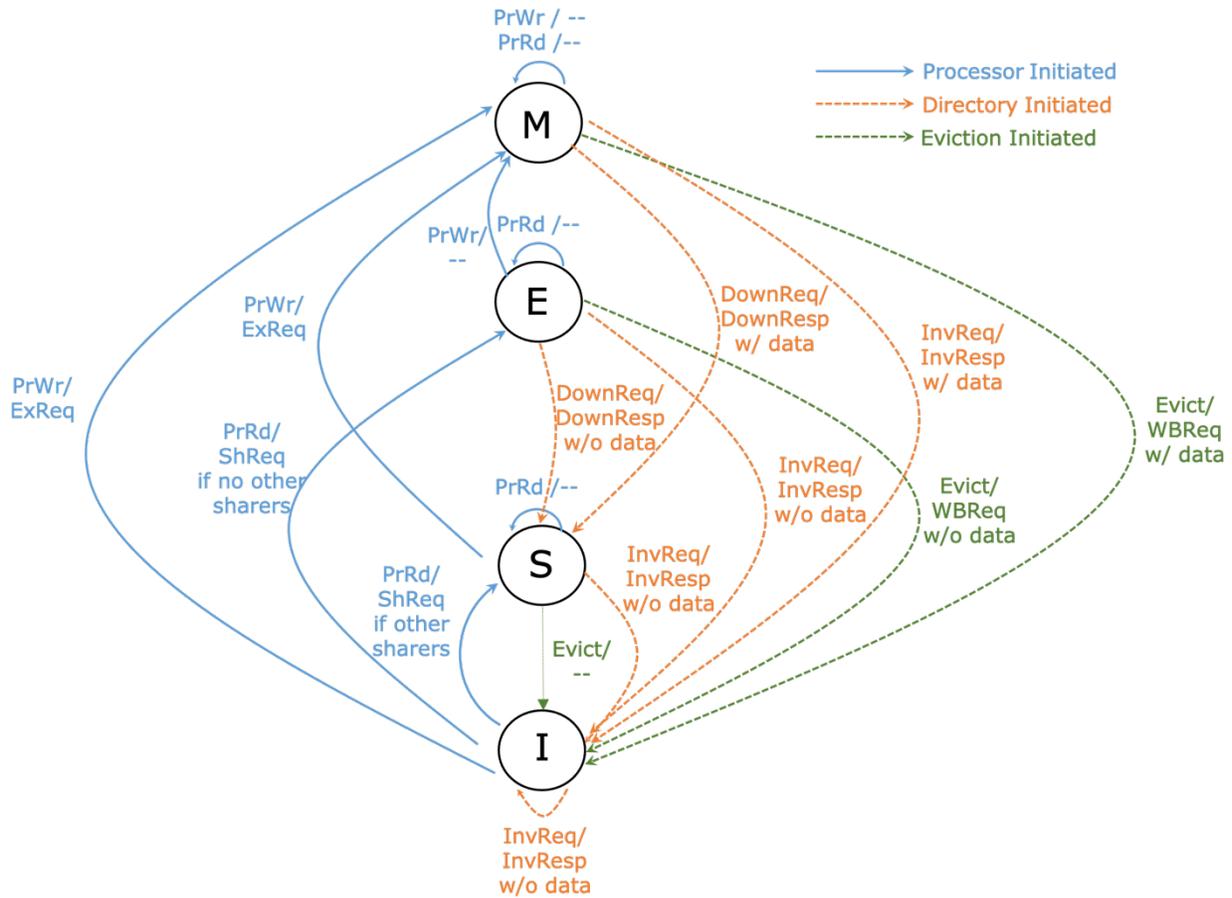
Data Transfer:

- **Cache to directory data transfer:** If a cache changes state from M to S or M to I (due to either an invalidation or a writeback), it sends the dirty data for that block to the directory.
- **Directory to cache data transfer:** The directory sends data to a cache when sending $I \rightarrow M$, $I \rightarrow S$, and $I \rightarrow E$ notification responses. Note that to do this the directory fetches the data from memory (see Lecture 14).

Cache Hit Rules:

- **Read hit:** if the cache state is E, S, or M
- **Write hit:** only if the cache state is M

Cache-Side State Transition Diagram:



Note that due to silent evictions, there are a couple differences from the usual directory-based MESI state transition diagram:

- Evictions from S state now do not send a writeback request to the directory.
- The I state must respond to an invalidation request from the directory since it could have been a prior sharer of the line that got silently evicted.