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 MSI protocol introduced in Lecture 13.

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.
- **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→S, I→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 state.
- **Exclusive (Ex):** One of the caches is in the M state.
- **Transient states** (Ex→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 have shared or exclusive permissions for that memory location. For example, sharers = {0, 1} means that Cache 0 and Cache 1 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 *invalidate* cache block A, i.e., to send back its (dirty) data (if the cache block is in state M) and change its state to I.
<DownReq, K, A>	The directory asks cache K to *downgrade* cache block A, i.e., to send back its (dirty) data and change its state from M 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
<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 (S) 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→M and I→S notification responses. Note that to do this the directory fetches the data from memory.

Cache Hit Rules:

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

Cache-Side State Transition Diagram (without transient states):

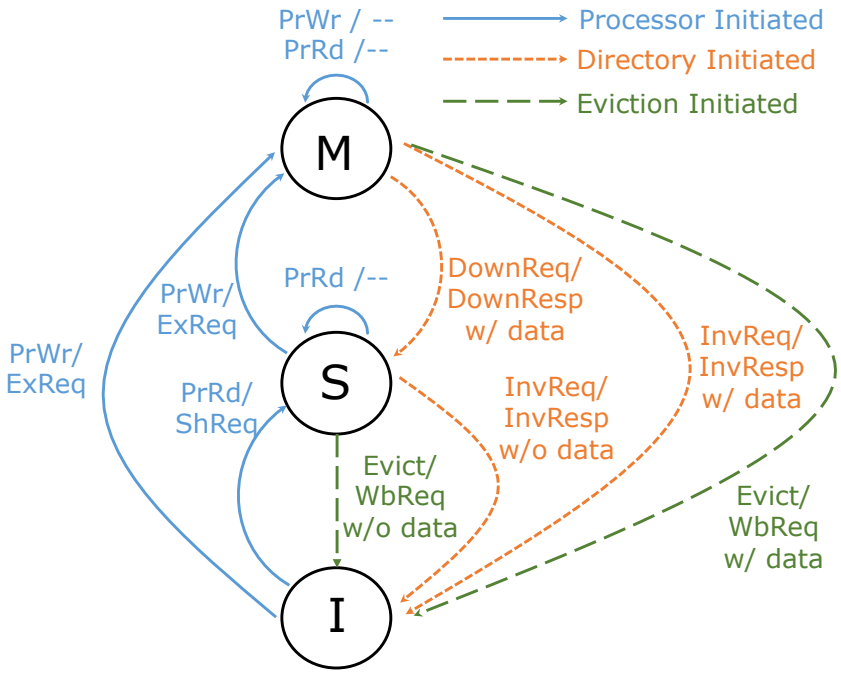

Directory-Side State Transition Diagram (without transient states):

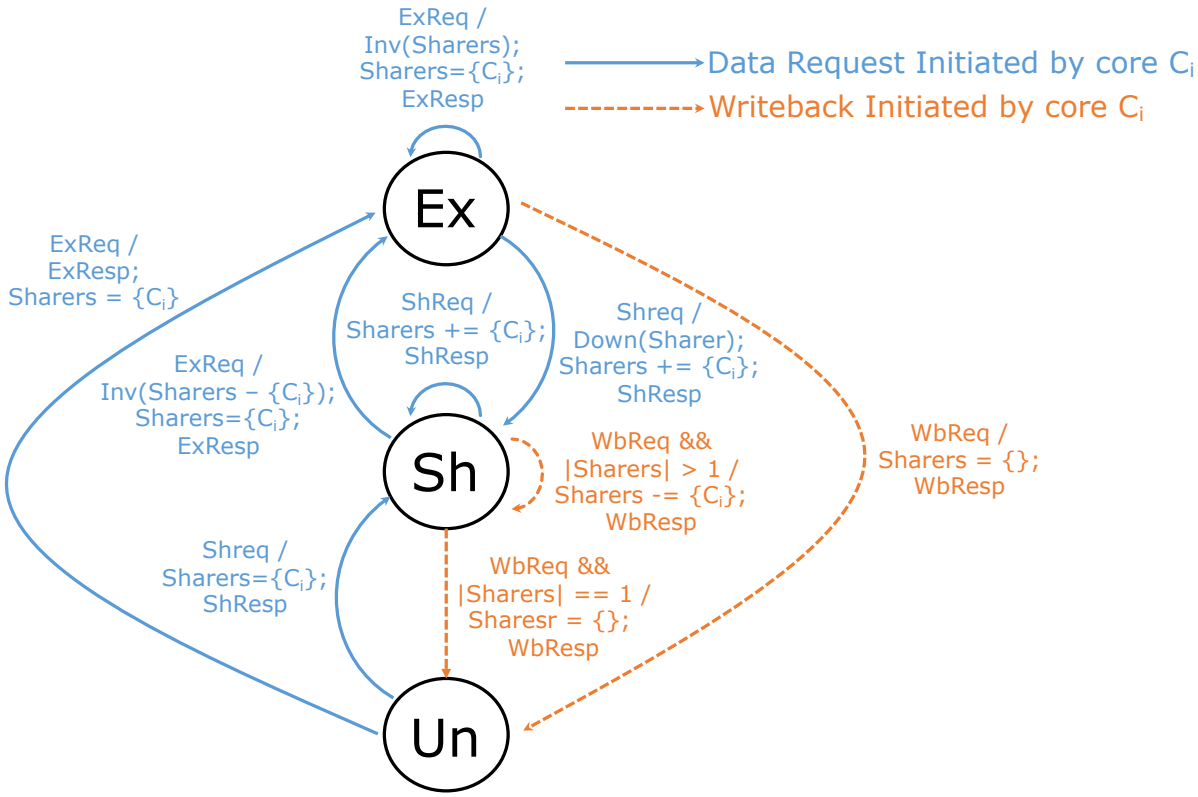

Note that the above state transition diagrams only show transitions between stable states. In reality, the coherence protocol must implement transient states when the cache or directory sends a message and waits for a response. For example, the following shows the cache waiting for a response to its ShReq within the I→S transient state: