

EXTENDS *Naturals*

CONSTANTS

KEYS

VARIABLES

database,
cache,
cacheFillStates, *cacheFillStatus*[*key*] = Fill state
invalidationQueue

INSTANCE *cacherequirements*

vars \triangleq \langle *database*, *cache*, *cacheFillStates*, *invalidationQueue* \rangle

InvalidationMessage \triangleq [*key* : *KEYS*, *version* : *DataVersion*]

CacheFillState \triangleq [*state* : { "inactive", "started", "respondedto" }, *version* : *DataVersion*]

CacheValue \triangleq *CacheMiss* \cup *CacheHit*

TypeOk \triangleq

\wedge *database* \in [*KEYS* \rightarrow *DataVersion*]

\wedge *cache* \in [*KEYS* \rightarrow *CacheValue*]

We track the cache fill state for each key. It is a multipart process

\wedge *cacheFillStates* \in [*KEYS* \rightarrow *CacheFillState*]

We model *invalidationQueue* as a set, because we cannot guarantee in-order delivery

\wedge *invalidationQueue* \in SUBSET *InvalidationMessage*

Init \triangleq

\wedge *database* = [*k* \in *KEYS* \mapsto 0]

\wedge *cache* = [*k* \in *KEYS* \mapsto [*type* \mapsto "miss"]]

Cache fill states start inactive

\wedge *cacheFillStates* = [*k* \in *KEYS* \mapsto [

state \mapsto "inactive",

Version set to earliest possible version

version \mapsto 0]

]

The invalidation queue starts empty

\wedge *invalidationQueue* = {}

DatabaseUpdate(*k*) \triangleq

LET *updatedVersion* \triangleq *database*[*k*] + 1IN

The version of that key is incremented, representing a write

$\wedge database' = [database \text{ EXCEPT } ! [k] = updatedVersion]$
 Adds invalidation message to queue.
 We don't need to model a delay in adding message as the cache can always delay handling message; to similar effect
 $\wedge invalidationQueue' = invalidationQueue \cup$
 Add updated data to invalidation message
 $\{[key \mapsto k, version \mapsto updatedVersion]\}$
 $\wedge \text{UNCHANGED } \langle cache, cacheFillStates \rangle$

Cache Fill behavior

$CacheStartReadThroughFill(k) \triangleq$
 Read through only occurs when the cache is unset for that value
 $\wedge cache[k] \in CacheMiss$
 One cache fill request at a time
 $\wedge cacheFillStates[k].state = \text{"inactive"}$
 $\wedge cacheFillStates' = [cacheFillStates \text{ EXCEPT } ! [k].state = \text{"started"}]$
 $\wedge \text{UNCHANGED } \langle database, cache, invalidationQueue \rangle$

This is the moment the database provides a value for cache fill

$DatabaseRespondToCacheFill(k) \triangleq$
 $\wedge cacheFillStates[k].state = \text{"started"}$
 $\wedge cacheFillStates' = [cacheFillStates \text{ EXCEPT } ! [k].state = \text{"respondedto"},$
 $! [k].version = database[k]$
 $]$
 $\wedge \text{UNCHANGED } \langle database, cache, invalidationQueue \rangle$

Cache fails to fill

$CacheFailFill(k) \triangleq$
 $\wedge cacheFillStates[k].state = \text{"respondedto"}$
 Cache fill state is reset, having not filled cache
 $\wedge cacheFillStates' = [cacheFillStates \text{ EXCEPT } ! [k].state = \text{"inactive"},$
 $! [k].version = 0$
 $]$
 $\wedge \text{UNCHANGED } \langle database, cache, invalidationQueue \rangle$

Cache incorporates the data

$CacheCompleteFill(k) \triangleq$
 $\wedge cacheFillStates[k].state = \text{"respondedto"}$
 Either the cache is empty for that key
 $\wedge \vee cache[k] \in CacheMiss$
 or we are filling a newer version
 $\vee \wedge cache[k] \notin CacheMiss$
 $\wedge cache[k].version < cacheFillStates[k].version$

$$\begin{aligned}
& \wedge \text{cacheFillStates}' = [\text{cacheFillStates} \text{ EXCEPT } \text{Reset to 0} \\
& \quad \quad \quad ![k].\text{state} = \text{"inactive"}, \\
& \quad \quad \quad ![k].\text{version} = 0 \\
& \quad \quad \quad] \\
& \wedge \text{cache}' = [\text{cache} \text{ EXCEPT} \\
& \quad \quad \quad ![k] = [\\
& \quad \quad \quad \quad \text{cache value is now a hit} \\
& \quad \quad \quad \quad \text{type} \mapsto \text{"hit"}, \\
& \quad \quad \quad \quad \text{set to whatever came back in response} \\
& \quad \quad \quad \quad \text{version} \mapsto \text{cacheFillStates}[k].\text{version} \\
& \quad \quad \quad \quad] \\
& \quad \quad \quad] \\
& \wedge \text{UNCHANGED} \langle \text{database}, \text{invalidationQueue} \rangle \\
\text{CacheIgnoreFill}(k) & \triangleq \\
& \wedge \text{cacheFillStates}[k].\text{state} = \text{"respondedto"} \\
& \quad \text{If we have a newer version in cache, ignore fill} \\
& \wedge \wedge \text{cache}[k] \in \text{CacheHit} \\
& \quad \wedge \text{cache}[k].\text{version} \geq \text{cacheFillStates}[k].\text{version} \\
& \wedge \text{cacheFillStates}' = [\text{cacheFillStates} \text{ EXCEPT } \text{Reset to 0} \\
& \quad \quad \quad ![k].\text{state} = \text{"inactive"}, \\
& \quad \quad \quad ![k].\text{version} = 0 \\
& \quad \quad \quad] \\
& \quad \text{Don't update cache} \\
& \wedge \text{UNCHANGED} \langle \text{cache}, \text{database}, \text{invalidationQueue} \rangle \\
\text{CacheHandleInvalidationMessage} & \triangleq \\
& \wedge \exists \text{message} \in \text{invalidationQueue} : \text{Deque invalidation queue in any order} \\
& \quad \text{Key must be in cache} \\
& \wedge \wedge \text{cache}[\text{message.key}] \in \text{CacheHit} \\
& \quad \text{Message needs to be newer than the cache} \\
& \quad \wedge \text{cache}[\text{message.key}].\text{version} < \text{message.version} \\
& \quad \text{Update item in cache} \\
& \wedge \text{cache}' = [\text{cache} \text{ EXCEPT} \\
& \quad \quad \quad ![message.key] = [\\
& \quad \quad \quad \quad \text{type} \mapsto \text{"hit"}, \\
& \quad \quad \quad \quad \text{Update to version in invalidation message} \\
& \quad \quad \quad \quad \text{version} \mapsto \text{message.version} \\
& \quad \quad \quad \quad] \\
& \quad \quad \quad] \\
& \quad \text{Remove message from queue because handled} \\
& \wedge \text{invalidationQueue}' = \text{invalidationQueue} \setminus \{\text{message}\} \\
& \wedge \text{UNCHANGED} \langle \text{cacheFillStates}, \text{database} \rangle \\
\text{CacheIgnoreInvalidationMessage} & \triangleq
\end{aligned}$$

$\wedge \exists message \in invalidationQueue :$ Dequeue invalidation queue in any order
 Ignore invalidation messages for messages not in cache
 $\wedge \vee cache[message.key] \in CacheMiss$
 Or when the cache already has the same or larger version
 $\vee \wedge cache[message.key] \notin CacheMiss$
 $\wedge cache[message.key].version \geq message.version$
 Remove message from queue to ignore
 $\wedge invalidationQueue' = invalidationQueue \setminus \{message\}$
 Don't update cache
 $\wedge UNCHANGED \langle cacheFillStates, database, cache \rangle$

Cache eviction model is unchanged
 $CacheEvict(k) \triangleq$
 $\wedge cache[k] \in CacheHit$
 $\wedge cache' = [cache \text{ EXCEPT } ![k] = [type \mapsto \text{"miss"}]]$
 $\wedge UNCHANGED \langle database, cacheFillStates, invalidationQueue \rangle$

Fairness: Normally no operation is guaranteed to happen, it just may. however that means, for example, that the cache could just stop reading forever. And so it would never update. Now that doesn't seem reasonable.

$CacheFairness \triangleq$
 $\exists k \in KEYS :$
 Cache fill process will be allowed to complete
 $\vee CacheStartReadThroughFill(k)$
 $\vee DatabaseRespondToCacheFill(k)$ Write
 $\vee CacheCompleteFill(k)$
 $\vee CacheIgnoreFill(k)$
 Cache will be allowed to process invalidation messages
 $\vee CacheHandleInvalidationMessage$
 $\vee CacheIgnoreInvalidationMessage$

Specification

$Next \triangleq$
 $\exists k \in KEYS :$
 Database states
 $\vee DatabaseUpdate(k)$
 Cache states
 $\vee CacheStartReadThroughFill(k)$
 $\vee DatabaseRespondToCacheFill(k)$
 $\vee CacheCompleteFill(k)$
 $\vee CacheIgnoreFill(k)$
 $\vee CacheHandleInvalidationMessage$

$\vee \text{CacheIgnoreInvalidationMessage}$
 $\vee \text{CacheEvict}(k)$

Cache fairness is included as part of the specification of system behavior.

This is just how the system works.

$Spec \triangleq \text{Init} \wedge \square[\text{Next}]_{vars} \wedge \text{WF}_{vars}(\text{CacheFairness})$

* Modification History
* Last modified *Wed Jun 15 12:49:34 MST 2022* by *elliotswart*
* Created *Tue Jun 14 20:36:02 MST 2022* by *elliotswart*