

EXTENDS *Integers*

The set of all keys.

CONSTANTS *KEY*

The sets of optimistic clients and pessimistic clients.

CONSTANTS *OPTIMISTIC_CLIENT*, *PESSIMISTIC_CLIENT*
 $CLIENT \triangleq PESSIMISTIC_CLIENT \cup OPTIMISTIC_CLIENT$

$CLIENT_KEY$ is a set of $[Client \rightarrow SUBSET\ KEY]$
 representing the involved keys of each client.

CONSTANTS *CLIENT_KEY*

ASSUME $\forall c \in CLIENT : CLIENT_KEY[c] \subseteq KEY$

$CLIENT_PRIMARY$ is the primary key of each client.

CONSTANTS *CLIENT_PRIMARY*

ASSUME $\forall c \in CLIENT : CLIENT_PRIMARY[c] \in CLIENT_KEY[c]$

Timestamp of transactions.

$Ts \triangleq Nat \setminus \{0\}$

$NoneTs \triangleq 0$

The algorithm is easier to understand in terms of the set of *msgs* of all messages that have ever been sent. A more accurate model would use one or more variables to represent the messages actually in transit, and it would include actions representing message loss and duplication as well as message receipt.

In the current spec, there is no need to model message loss because we are mainly concerned with the algorithm's safety property. The safety part of the spec says only what messages may be received and does not assert that any message actually is received. Thus, there is no difference between a lost message and one that is never received.

VARIABLES *req_msgs*

VARIABLES *resp_msgs*

$key_data[k]$ is the set of multi-version data of the key. Since we don't care about the concrete value of data, a *strat_ts* is sufficient to represent one data version.

VARIABLES *key_data*

$key_lock[k]$ is the set of lock (zero or one element). A lock is of a record of $[ts: start_ts, primary: key, type: lock_type]$. If primary equals to k , it is a primary lock, otherwise secondary lock. *lock_type* is one of {"prewrite_optimistic", "prewrite_pessimistic", "lock_key"}. *lock_key* denotes the pessimistic lock performed by *ServerLockKey* action, the *prewrite_pessimistic* denotes percolator optimistic lock

who is transformed from a *lock_key* lock by action *ServerPrewritePessimistic*, and *prewrite_optimistic* denotes the classic optimistic lock.

In *TiKV*, *key_lock* has an additional *for_update_ts* field and the *LockType* is of four variants: $\{“PUT”, “DELETE”, “LOCK”, “PESSIMISTIC”\}$.

In the spec, we abstract them by:

- (1) $LockType \in \{“PUT”, “DELETE”, “LOCK”\} \wedge for_update_ts = 0 \equiv type = “prewrite_optimistic”$
- (2) $LockType \in \{“PUT”, “DELETE”\} \wedge for_update_ts > 0 \equiv type = “prewrite_pessimistic”$
- (3) $LockType = “PESSIMISTIC” \equiv type = “lock_key”$

VARIABLES *key_lock*

key_write[*k*] is a sequence of commit or rollback record of the key. It's a record of [*ts*, *start_ts*, type, [protected]]. type can be either “write” or “rollback”. *ts* represents the *commit_ts* of “write” record. Otherwise, *ts* equals to *start_ts* on “rollback” record. “rollback” record has an additional protected field. protected signifies the rollback record would not be collapsed.

VARIABLES *key_write*

client_state[*c*] indicates the current transaction stage of client *c*.

VARIABLES *client_state*

client_ts[*c*] is a record of [*start_ts*, *commit_ts*, *for_update_ts*]. Fields are all initialized to *NoneTs*.

VARIABLES *client_ts*

client_key[*c*] is a record of [locking: {*key*}, prewriting: {*key*}]. Hereby, “locking” denotes the keys whose pessimistic locks haven't been acquired, “prewriting” denotes the keys that are pending for prewrite.

VARIABLES *client_key*

next_ts is a globally monotonically increasing integer, representing the virtual clock of transactions. In practice, the variable is maintained by *PD*, the time oracle of a cluster.

VARIABLES *next_ts*

$msg_vars \triangleq \langle req_msgs, resp_msgs \rangle$
 $client_vars \triangleq \langle client_state, client_ts, client_key \rangle$
 $key_vars \triangleq \langle key_data, key_lock, key_write \rangle$
 $vars \triangleq \langle msg_vars, client_vars, key_vars, next_ts \rangle$

$SendReqs(msgs) \triangleq req_msgs' = req_msgs \cup msgs$
 $SendResp(msg) \triangleq resp_msgs' = resp_msgs \cup \{msg\}$

Type Definitions

$ReqMessages \triangleq$

- $[start_ts : Ts, primary : KEY, type : \{ "lock_key" \}, key : KEY,$
 $for_update_ts : Ts]$
- $\cup [start_ts : Ts, primary : KEY, type : \{ "prewrite_optimistic" \}, key : KEY]$
- $\cup [start_ts : Ts, primary : KEY, type : \{ "prewrite_pessimistic" \}, key : KEY]$
- $\cup [start_ts : Ts, primary : KEY, type : \{ "commit" \}, commit_ts : Ts]$
- $\cup [start_ts : Ts, primary : KEY, type : \{ "cleanup" \}]$
- $\cup [start_ts : Ts, primary : KEY, type : \{ "resolve_rollbacked" \}]$
- $\cup [start_ts : Ts, primary : KEY, type : \{ "resolve_committed" \}, commit_ts : Ts]$

$RespMessages \triangleq$

- $[start_ts : Ts, type : \{ "prewrited", "locked_key" \}, key : KEY]$
- $\cup [start_ts : Ts, type : \{ "lock_failed" \}, key : KEY, latest_commit_ts : Ts]$
- $\cup [start_ts : Ts, type : \{ "committed",$
 $"commit_aborted",$
 $"prewrite_aborted",$
 $"lock_key_aborted" \}]$

$TypeOK \triangleq$

- $\wedge req_msgs \in \text{SUBSET } ReqMessages$
 - $\wedge resp_msgs \in \text{SUBSET } RespMessages$
 - $\wedge key_data \in [KEY \rightarrow \text{SUBSET } Ts]$
 - $\wedge key_lock \in [KEY \rightarrow \text{SUBSET } [ts : Ts,$
 $primary : KEY,$
 $type : \{ "prewrite_optimistic",$
 $"prewrite_pessimistic",$
 $"lock_key" \}]]]$
 - $\wedge \forall k \in KEY :$
 $\text{At most one lock in } key_lock[k]$
 $\forall l, l2 \in key_lock[k] :$
 $l = l2$
 - $\wedge key_write \in [KEY \rightarrow \text{SUBSET } ($
 $[ts : Ts, start_ts : Ts, type : \{ "write" \}]$
 $\cup [ts : Ts, start_ts : Ts, type : \{ "rollback" \}, protected : \text{BOOLEAN }])]]$
 - $\wedge client_state \in [CLIENT \rightarrow \{ "init", "locking", "prewriting", "committing" \}]$
 - $\wedge client_ts \in [CLIENT \rightarrow [start_ts : Ts \cup \{ NoneTs \},$
 $commit_ts : Ts \cup \{ NoneTs \},$
 $for_update_ts : Ts \cup \{ NoneTs \}]]]$
 - $\wedge client_key \in [CLIENT \rightarrow [locking : \text{SUBSET } KEY, prewriting : \text{SUBSET } KEY]]]$
 - $\wedge next_ts \in Ts$
-

Client Actions

$ClientLockKey(c) \triangleq$

$$\begin{aligned}
& \wedge \text{client_state}[c] = \text{"init"} \\
& \wedge \text{client_state}' = [\text{client_state} \text{ EXCEPT } ![c] = \text{"locking"}] \\
& \wedge \text{client_ts}' = [\text{client_ts} \text{ EXCEPT } ![c].\text{start_ts} = \text{next_ts}, ![c].\text{for_update_ts} = \text{next_ts}] \\
& \wedge \text{next_ts}' = \text{next_ts} + 1 \\
& \text{Assume we need to acquire pessimistic locks for all keys} \\
& \wedge \text{client_key}' = [\text{client_key} \text{ EXCEPT } ![c].\text{locking} = \text{CLIENT_KEY}[c]] \\
& \wedge \text{SendReqs}(\{[type \mapsto \text{"lock_key"}, \\
& \quad \text{start_ts} \mapsto \text{client_ts}'[c].\text{start_ts}, \\
& \quad \text{primary} \mapsto \text{CLIENT_PRIMARY}[c], \\
& \quad \text{key} \mapsto k, \\
& \quad \text{for_update_ts} \mapsto \text{client_ts}'[c].\text{for_update_ts}] : k \in \text{CLIENT_KEY}[c]\}) \\
& \wedge \text{UNCHANGED} \langle \text{resp_msgs}, \text{key_vars} \rangle \\
\\
\text{ClientLockedKey}(c) & \triangleq \\
& \wedge \text{client_state}[c] = \text{"locking"} \\
& \wedge \exists \text{resp} \in \text{resp_msgs} : \\
& \quad \wedge \text{resp.type} = \text{"locked_key"} \\
& \quad \wedge \text{resp.start_ts} = \text{client_ts}[c].\text{start_ts} \\
& \quad \wedge \text{resp.key} \in \text{client_key}[c].\text{locking} \\
& \quad \wedge \text{client_key}' = [\text{client_key} \text{ EXCEPT } ![c].\text{locking} = @ \setminus \{\text{resp.key}\}] \\
& \quad \wedge \text{UNCHANGED} \langle \text{msg_vars}, \text{key_vars}, \text{client_ts}, \text{client_state}, \text{next_ts} \rangle \\
\\
\text{ClientRetryLockKey}(c) & \triangleq \\
& \wedge \text{client_state}[c] = \text{"locking"} \\
& \wedge \exists \text{resp} \in \text{resp_msgs} : \\
& \quad \wedge \text{resp.type} = \text{"lock_failed"} \\
& \quad \wedge \text{resp.start_ts} = \text{client_ts}[c].\text{start_ts} \\
& \quad \wedge \text{resp.latest_commit_ts} > \text{client_ts}[c].\text{for_update_ts} \\
& \quad \wedge \text{client_ts}' = [\text{client_ts} \text{ EXCEPT } ![c].\text{for_update_ts} = \text{resp.latest_commit_ts}] \\
& \quad \wedge \text{SendReqs}(\{[type \mapsto \text{"lock_key"}, \\
& \quad \text{start_ts} \mapsto \text{client_ts}'[c].\text{start_ts}, \\
& \quad \text{primary} \mapsto \text{CLIENT_PRIMARY}[c], \\
& \quad \text{key} \mapsto \text{resp.key}, \\
& \quad \text{for_update_ts} \mapsto \text{client_ts}'[c].\text{for_update_ts}]\}) \\
& \quad \wedge \text{UNCHANGED} \langle \text{resp_msgs}, \text{key_vars}, \text{client_state}, \text{client_key}, \text{next_ts} \rangle \\
\\
\text{ClientPrewritePessimistic}(c) & \triangleq \\
& \wedge \text{client_state}[c] = \text{"locking"} \\
& \wedge \text{client_key}[c].\text{locking} = \{\} \\
& \wedge \text{client_state}' = [\text{client_state} \text{ EXCEPT } ![c] = \text{"prewriting"}] \\
& \wedge \text{client_key}' = [\text{client_key} \text{ EXCEPT } ![c].\text{prewriting} = \text{CLIENT_KEY}[c]] \\
& \wedge \text{SendReqs}(\{[type \mapsto \text{"prewrite_pessimistic"}, \\
& \quad \text{start_ts} \mapsto \text{client_ts}[c].\text{start_ts}, \\
& \quad \text{primary} \mapsto \text{CLIENT_PRIMARY}[c], \\
& \quad \text{key} \mapsto k] : k \in \text{CLIENT_KEY}[c]\}) \\
& \wedge \text{UNCHANGED} \langle \text{resp_msgs}, \text{key_vars}, \text{client_ts}, \text{next_ts} \rangle
\end{aligned}$$

$$\begin{aligned}
& \text{ClientPrewriteOptimistic}(c) \triangleq \\
& \wedge \text{client_state}[c] = \text{"init"} \\
& \wedge \text{client_state}' = [\text{client_state} \text{ EXCEPT } ![c] = \text{"prewriting"}] \\
& \wedge \text{client_ts}' = [\text{client_ts} \text{ EXCEPT } ![c].\text{start_ts} = \text{next_ts}] \\
& \wedge \text{next_ts}' = \text{next_ts} + 1 \\
& \wedge \text{client_key}' = [\text{client_key} \text{ EXCEPT } ![c].\text{prewriting} = \text{CLIENT_KEY}[c]] \\
& \wedge \text{SendReqs}(\{[type \mapsto \text{"prewrite_optimistic"}, \\
& \quad \text{start_ts} \mapsto \text{client_ts}'[c].\text{start_ts}, \\
& \quad \text{primary} \mapsto \text{CLIENT_PRIMARY}[c], \\
& \quad \text{key} \mapsto k] : k \in \text{CLIENT_KEY}[c]\}) \\
& \wedge \text{UNCHANGED} \langle \text{resp_msgs}, \text{key_vars} \rangle
\end{aligned}$$

$$\begin{aligned}
& \text{ClientPrewritten}(c) \triangleq \\
& \wedge \text{client_state}[c] = \text{"prewriting"} \\
& \wedge \text{client_key}[c].\text{locking} = \{\} \\
& \wedge \exists \text{resp} \in \text{resp_msgs} : \\
& \quad \wedge \text{resp.type} = \text{"prewritten"} \\
& \quad \wedge \text{resp.start_ts} = \text{client_ts}[c].\text{start_ts} \\
& \quad \wedge \text{resp.key} \in \text{client_key}[c].\text{prewriting} \\
& \quad \wedge \text{client_key}' = [\text{client_key} \text{ EXCEPT } ![c].\text{prewriting} = @ \setminus \{\text{resp.key}\}] \\
& \quad \wedge \text{UNCHANGED} \langle \text{msg_vars}, \text{key_vars}, \text{client_ts}, \text{client_state}, \text{next_ts} \rangle
\end{aligned}$$

$$\begin{aligned}
& \text{ClientCommit}(c) \triangleq \\
& \wedge \text{client_state}[c] = \text{"prewriting"} \\
& \wedge \text{client_key}[c].\text{prewriting} = \{\} \\
& \wedge \text{client_state}' = [\text{client_state} \text{ EXCEPT } ![c] = \text{"committing"}] \\
& \wedge \text{client_ts}' = [\text{client_ts} \text{ EXCEPT } ![c].\text{commit_ts} = \text{next_ts}] \\
& \wedge \text{next_ts}' = \text{next_ts} + 1 \\
& \wedge \text{SendReqs}(\{[type \mapsto \text{"commit"}, \\
& \quad \text{start_ts} \mapsto \text{client_ts}'[c].\text{start_ts}, \\
& \quad \text{primary} \mapsto \text{CLIENT_PRIMARY}[c], \\
& \quad \text{commit_ts} \mapsto \text{client_ts}'[c].\text{commit_ts}]\}) \\
& \wedge \text{UNCHANGED} \langle \text{resp_msgs}, \text{key_vars}, \text{client_key} \rangle
\end{aligned}$$

Server Actions

Write the write column and unlock the lock iff the lock exists.

$$\begin{aligned}
& \text{commit}(pk, \text{start_ts}, \text{commit_ts}) \triangleq \\
& \exists l \in \text{key_lock}[pk] : \\
& \quad \wedge l.\text{ts} = \text{start_ts} \\
& \quad \wedge \text{key_lock}' = [\text{key_lock} \text{ EXCEPT } ![pk] = \{\}] \\
& \quad \wedge \text{key_write}' = [\text{key_write} \text{ EXCEPT } ![pk] = @ \cup \{[ts \mapsto \text{commit_ts}, \\
& \quad \text{type} \mapsto \text{"write"}, \\
& \quad \text{start_ts} \mapsto \text{start_ts}]\}]
\end{aligned}$$

Rollback the transaction that starts at start_ts on key k .

$rollback(k, start_ts) \triangleq$
 LET
 Rollback record on the primary key of a pessimistic transaction needs to be protected from being collapsed. If we can't decide whether it suffices that because the lock is missing or mismatched, it should also be protected.
 $protected \triangleq \vee \exists l \in key_lock[k] :$
 $\quad \wedge l.ts = start_ts$
 $\quad \wedge l.primary = k$
 $\quad \wedge l.type \in \{ "lock_key", "prewrite_pessimistic" \}$
 $\vee \exists l \in key_lock[k] : l.ts \neq start_ts$
 $\vee key_lock[k] = \{ \}$
 IN
 If a lock exists and has the same ts , unlock it.
 \wedge IF $\exists l \in key_lock[k] : l.ts = start_ts$
 THEN $key_lock' = [key_lock \text{ EXCEPT } ![k] = \{ \}]$
 ELSE UNCHANGED key_lock
 \wedge $key_data' = [key_data \text{ EXCEPT } ![k] = @ \setminus \{ start_ts \}]$
 \wedge IF
 $\quad \wedge \neg \exists w \in key_write[k] : w.ts = start_ts$
 THEN
 $\quad key_write' = [key_write \text{ EXCEPT } ![k] =$
 $\quad \quad \text{collapse rollback}$
 $\quad \quad (@ \setminus \{ w \in @ : w.type = "rollback" \wedge \neg w.protected \wedge w.ts < start_ts \})$
 $\quad \quad \text{write rollback record}$
 $\quad \quad \cup \{ [ts \mapsto start_ts,$
 $\quad \quad \quad start_ts \mapsto start_ts,$
 $\quad \quad \quad type \mapsto "rollback",$
 $\quad \quad \quad protected \mapsto protected] \}$
 ELSE
 UNCHANGED $\langle key_write \rangle$
 $ServerLockKey \triangleq$
 $\exists req \in req_msgs :$
 $\quad \wedge req.type = "lock_key"$
 $\quad \wedge$ LET
 $\quad \quad k \triangleq req.key$
 $\quad \quad start_ts \triangleq req.start_ts$
 IN
 Pessimistic lock is allowed only if no stale lock exists. If there is one, wait until $ServerCleanupStaleLock$ to clean it up.
 $\quad \wedge key_lock[k] = \{ \}$
 $\quad \wedge$ LET
 $\quad \quad latest_write \triangleq \{ w \in key_write[k] : \forall w2 \in key_write[k] : w.ts \geq w2.ts \}$

$all_commits \triangleq \{w \in key_write[k] : w.type = \text{"write"}\}$
 $latest_commit \triangleq \{w \in all_commits : \forall w2 \in all_commits : w.ts \geq w2.ts\}$

IN
 IF $\exists w \in key_write[k] : w.start_ts = start_ts \wedge w.type = \text{"rollback"}$
 THEN
 If corresponding rollback record is found, which indicates that the transaction is rolled back, abort the transaction.
 $\wedge SendResp([start_ts \mapsto start_ts, type \mapsto \text{"lock_key_aborted"}])$
 $\wedge UNCHANGED \langle req_msgs, client_vars, key_vars, next_ts \rangle$
 ELSE
 Acquire pessimistic lock only if for_update_ts of req is greater or equal to the latest "write" record. Because if the latest record is "write", it means that a new version is committed after for_update_ts , which violates Read Committed guarantee.
 $\vee \wedge \neg \exists w \in latest_commit : w.ts > req.for_update_ts$
 $\wedge key_lock' = [key_lock \text{ EXCEPT } ![k] = \{[ts \mapsto start_ts,$
 $primary \mapsto req.primary,$
 $type \mapsto \text{"lock_key"}]\}]$
 $\wedge SendResp([start_ts \mapsto start_ts, type \mapsto \text{"locked_key"}, key \mapsto k])$
 $\wedge UNCHANGED \langle req_msgs, client_vars, key_data, key_write, next_ts \rangle$
 Otherwise, reject the request and let client to retry with new for_update_ts .
 $\vee \exists w \in latest_commit :$
 $\wedge w.ts > req.for_update_ts$
 $\wedge SendResp([start_ts \mapsto start_ts,$
 $type \mapsto \text{"lock_failed"},$
 $key \mapsto k,$
 $latest_commit_ts \mapsto w.ts])$
 $\wedge UNCHANGED \langle req_msgs, client_vars, key_vars, next_ts \rangle$

$ServerPrewritePessimistic \triangleq$

$\exists req \in req_msgs :$

$\wedge req.type = \text{"prewrite_pessimistic"}$

$\wedge LET$

$k \triangleq req.key$

$start_ts \triangleq req.start_ts$

IN

Pessimistic prewrite is allowed only if pessimistic lock is acquired, otherwise abort the transaction.

$\wedge IF \exists l \in key_lock[k] : l.ts = start_ts$

THEN

$\wedge key_lock' = [key_lock \text{ EXCEPT } ![k] = \{[ts \mapsto start_ts,$
 $primary \mapsto req.primary,$

$$\begin{aligned}
& \text{type} \mapsto \text{"prewrite_pessimistic"} \}}] \\
& \wedge \text{key_data}' = [\text{key_data} \text{ EXCEPT } ![k] = @ \cup \{start_ts\}] \\
& \wedge \text{SendResp}([start_ts \mapsto start_ts, \text{type} \mapsto \text{"prewrited"}, \text{key} \mapsto k]) \\
& \wedge \text{UNCHANGED} \langle \text{req_msgs}, \text{client_vars}, \text{key_write}, \text{next_ts} \rangle \\
\text{ELSE} \\
& \wedge \text{SendResp}([start_ts \mapsto start_ts, \text{type} \mapsto \text{"prewrite_aborted"}]) \\
& \wedge \text{UNCHANGED} \langle \text{req_msgs}, \text{client_vars}, \text{key_vars}, \text{next_ts} \rangle \\
\text{ServerPrewriteOptimistic} \triangleq \\
& \exists \text{req} \in \text{req_msgs} : \\
& \wedge \text{req.type} = \text{"prewrite_optimistic"} \\
& \wedge \text{LET} \\
& \quad k \triangleq \text{req.key} \\
& \quad start_ts \triangleq \text{req.start_ts} \\
& \text{IN} \\
& \wedge \text{IF } \exists w \in \text{key_write}[k] : w.ts \geq start_ts \\
& \quad \text{THEN} \\
& \quad \wedge \text{SendResp}([start_ts \mapsto start_ts, \text{type} \mapsto \text{"prewrite_aborted"}]) \\
& \quad \wedge \text{UNCHANGED} \langle \text{req_msgs}, \text{client_vars}, \text{key_vars}, \text{next_ts} \rangle \\
& \quad \text{ELSE} \\
& \quad \text{Optimistic prewrite is allowed only if no stale lock exists. If} \\
& \quad \text{there is one, wait until } \text{ServerCleanupStaleLock} \text{ to clean it up.} \\
& \quad \wedge \forall \text{key_lock}[k] = \{\} \\
& \quad \quad \vee \exists l \in \text{key_lock}[k] : l.ts = start_ts \\
& \quad \wedge \text{key_lock}' = [\text{key_lock} \text{ EXCEPT } ![k] = \{[ts \mapsto start_ts, \\
& \quad \quad \quad \text{primary} \mapsto \text{req.primary}, \\
& \quad \quad \quad \text{type} \mapsto \text{"prewrite_optimistic"} \}}] \\
& \quad \wedge \text{key_data}' = [\text{key_data} \text{ EXCEPT } ![k] = @ \cup \{start_ts\}] \\
& \quad \wedge \text{SendResp}([start_ts \mapsto start_ts, \text{type} \mapsto \text{"prewrited"}, \text{key} \mapsto k]) \\
& \quad \wedge \text{UNCHANGED} \langle \text{req_msgs}, \text{client_vars}, \text{key_write}, \text{next_ts} \rangle \\
\text{ServerCommit} \triangleq \\
& \exists \text{req} \in \text{req_msgs} : \\
& \wedge \text{req.type} = \text{"commit"} \\
& \wedge \text{LET} \\
& \quad pk \triangleq \text{req.primary} \\
& \quad start_ts \triangleq \text{req.start_ts} \\
& \text{IN} \\
& \text{IF } \exists w \in \text{key_write}[pk] : w.start_ts = start_ts \wedge w.type = \text{"write"} \\
& \quad \text{THEN} \\
& \quad \text{Key has already been committed. Do nothing.} \\
& \quad \wedge \text{SendResp}([start_ts \mapsto start_ts, \text{type} \mapsto \text{"committed"}]) \\
& \quad \wedge \text{UNCHANGED} \langle \text{req_msgs}, \text{client_vars}, \text{key_vars}, \text{next_ts} \rangle \\
& \quad \text{ELSE} \\
& \quad \text{IF } \exists l \in \text{key_lock}[pk] : l.ts = start_ts
\end{aligned}$$

THEN
 Commit the key only if the prewrite lock exists.
 $\wedge \text{commit}(pk, start_ts, req.commit_ts)$
 $\wedge \text{SendResp}([start_ts \mapsto start_ts, type \mapsto \text{"committed"}])$
 $\wedge \text{UNCHANGED} \langle req_msgs, client_vars, key_data, next_ts \rangle$
 ELSE
 Otherwise, abort the transaction.
 $\wedge \text{SendResp}([start_ts \mapsto start_ts, type \mapsto \text{"commit_aborted"}])$
 $\wedge \text{UNCHANGED} \langle req_msgs, client_vars, key_vars, next_ts \rangle$

In the spec, the primary key with a lock may clean up itself spontaneously. There is no need to model a client to request clean up because there is no difference between a optimistic client trying to read a key that has lock timeouted and the key trying to unlock itself.

$ServerCleanupStaleLock \triangleq$

$\exists k \in KEY :$
 $\exists l \in key_lock[k] :$
 $\wedge \text{SendReqs}(\{[type \mapsto \text{"cleanup"},$
 $start_ts \mapsto l.ts,$
 $primary \mapsto l.primary]\})$
 $\wedge \text{UNCHANGED} \langle resp_msgs, client_vars, key_vars, next_ts \rangle$

Clean up stale locks by checking the status of the primary key. Commit the secondary keys if primary key is committed; otherwise rollback the transaction by rolling-back the primary key, and then also rollback the secondaries.

$ServerCleanup \triangleq$

$\exists req \in req_msgs :$
 $\wedge req.type = \text{"cleanup"}$
 $\wedge \text{LET}$
 $pk \triangleq req.primary$
 $start_ts \triangleq req.start_ts$
 $committed \triangleq \{w \in key_write[pk] : w.start_ts = start_ts \wedge w.type = \text{"write"}\}$
 IN
 IF $committed \neq \{\}$
 THEN
 $\wedge \text{SendReqs}(\{[type \mapsto \text{"resolve_committed"},$
 $start_ts \mapsto start_ts,$
 $primary \mapsto pk,$
 $commit_ts \mapsto w.ts] : w \in committed\})$
 $\wedge \text{UNCHANGED} \langle resp_msgs, client_vars, key_vars, next_ts \rangle$
 ELSE
 $\wedge \text{rollback}(pk, start_ts)$
 $\wedge \text{SendReqs}(\{[type \mapsto \text{"resolve_rollbacked"},$
 $start_ts \mapsto start_ts,$

$$\wedge \text{UNCHANGED } \langle \text{primary} \mapsto pk \rangle \rangle$$

$$\wedge \text{UNCHANGED } \langle \text{resp_msgs}, \text{client_vars}, \text{next_ts} \rangle$$

ServerResolveCommitted \triangleq

$$\begin{aligned} & \exists \text{req} \in \text{req_msgs} : \\ & \wedge \text{req.type} = \text{"resolve_committed"} \\ & \wedge \text{LET} \\ & \quad \text{start_ts} \triangleq \text{req.start_ts} \\ & \text{IN} \\ & \quad \exists k \in \text{KEY} : \\ & \quad \quad \exists l \in \text{key_lock}[k] : \\ & \quad \quad \quad \wedge l.\text{primary} = \text{req.primary} \\ & \quad \quad \quad \wedge l.\text{ts} = \text{start_ts} \\ & \quad \quad \quad \wedge \text{commit}(k, \text{start_ts}, \text{req.commit_ts}) \\ & \quad \quad \wedge \text{UNCHANGED } \langle \text{msg_vars}, \text{client_vars}, \text{key_data}, \text{next_ts} \rangle \end{aligned}$$

ServerResolveRollbacked \triangleq

$$\begin{aligned} & \exists \text{req} \in \text{req_msgs} : \\ & \wedge \text{req.type} = \text{"resolve_rollbacked"} \\ & \wedge \text{LET} \\ & \quad \text{start_ts} \triangleq \text{req.start_ts} \\ & \text{IN} \\ & \quad \exists k \in \text{KEY} : \\ & \quad \quad \exists l \in \text{key_lock}[k] : \\ & \quad \quad \quad \wedge l.\text{primary} = \text{req.primary} \\ & \quad \quad \quad \wedge l.\text{ts} = \text{start_ts} \\ & \quad \quad \quad \wedge \text{rollback}(k, \text{start_ts}) \\ & \quad \quad \wedge \text{UNCHANGED } \langle \text{msg_vars}, \text{client_vars}, \text{next_ts} \rangle \end{aligned}$$

Specification

Init \triangleq

$$\begin{aligned} & \wedge \text{next_ts} = 1 \\ & \wedge \text{req_msgs} = \{\} \\ & \wedge \text{resp_msgs} = \{\} \\ & \wedge \text{client_state} = [c \in \text{CLIENT} \mapsto \text{"init"}] \\ & \wedge \text{client_key} = [c \in \text{CLIENT} \mapsto [\text{locking} \mapsto \{\}, \text{prewriting} \mapsto \{\}]] \\ & \wedge \text{client_ts} = [c \in \text{CLIENT} \mapsto [\text{start_ts} \mapsto \text{NoneTs}, \\ & \quad \quad \quad \text{commit_ts} \mapsto \text{NoneTs}, \\ & \quad \quad \quad \text{for_update_ts} \mapsto \text{NoneTs}]] \\ & \wedge \text{key_lock} = [k \in \text{KEY} \mapsto \{\}] \\ & \wedge \text{key_data} = [k \in \text{KEY} \mapsto \{\}] \\ & \wedge \text{key_write} = [k \in \text{KEY} \mapsto \{\}] \end{aligned}$$

Next \triangleq

$$\vee \exists c \in \text{OPTIMISTIC_CLIENT} :$$

$\vee \text{ClientPrewriteOptimistic}(c)$
 $\vee \text{ClientPrewritten}(c)$
 $\vee \text{ClientCommit}(c)$
 $\vee \exists c \in \text{PESSIMISTIC_CLIENT} :$
 $\vee \text{ClientLockKey}(c)$
 $\vee \text{ClientLockedKey}(c)$
 $\vee \text{ClientRetryLockKey}(c)$
 $\vee \text{ClientPrewritePessimistic}(c)$
 $\vee \text{ClientPrewritten}(c)$
 $\vee \text{ClientCommit}(c)$
 $\vee \text{ServerLockKey}$
 $\vee \text{ServerPrewritePessimistic}$
 $\vee \text{ServerPrewriteOptimistic}$
 $\vee \text{ServerCommit}$
 $\vee \text{ServerCleanupStaleLock}$
 $\vee \text{ServerCleanup}$
 $\vee \text{ServerResolveCommitted}$
 $\vee \text{ServerResolveRollbacked}$

$\text{Spec} \triangleq \text{Init} \wedge \square[\text{Next}]_{\text{vars}}$

Consistency Invariants

Check whether there is a “write” record in $\text{key_write}[k]$ corresponding to start_ts .

$\text{keyCommitted}(k, \text{start_ts}) \triangleq$
 $\exists w \in \text{key_write}[k] :$
 $\wedge w.\text{start_ts} = \text{start_ts}$
 $\wedge w.\text{type} = \text{“write”}$

A transaction can’t be both committed and aborted.

$\text{UniqueCommitOrAbort} \triangleq$
 $\forall \text{resp}, \text{resp2} \in \text{resp_msgs} :$
 $(\text{resp.type} = \text{“committed”}) \wedge (\text{resp2.type} = \text{“commit_aborted”}) \Rightarrow$
 $\text{resp.start_ts} \neq \text{resp2.start_ts}$

If a transaction is committed, the primary key must be committed and the secondary keys of the same transaction must be either committed or locked.

$\text{CommitConsistency} \triangleq$
 $\forall \text{resp} \in \text{resp_msgs} :$
 $(\text{resp.type} = \text{“committed”}) \Rightarrow$
 $\exists c \in \text{CLIENT} :$
 $\wedge \text{client_ts}[c].\text{start_ts} = \text{resp.start_ts}$
 $\text{Primary key must be committed}$
 $\wedge \text{keyCommitted}(\text{CLIENT_PRIMARY}[c], \text{resp.start_ts})$

Secondary key must be either committed or locked by the $start_ts$ of the transaction.

$$\wedge \forall k \in CLIENT_KEY[c] : \\ (\neg \exists l \in key_lock[k] : l.ts = resp.start_ts) = \\ keyCommitted(k, resp.start_ts)$$

If a transaction is aborted, all key of that transaction must be not committed.

$$AbortConsistency \triangleq \\ \forall resp \in resp_msgs : \\ (resp.type = \text{"commit_aborted"}) \Rightarrow \\ \forall c \in CLIENT : \\ (client_ts[c].start_ts = resp.start_ts) \Rightarrow \\ \neg keyCommitted(CLIENT_PRIMARY[c], resp.start_ts)$$

For each write, the $commit_ts$ should be strictly greater than the $start_ts$ and have data written into $key_data[k]$. For each rollback, the $commit_ts$ should equals to the $start_ts$.

$$WriteConsistency \triangleq \\ \forall k \in KEY : \\ \forall w \in key_write[k] : \\ \vee \wedge w.type = \text{"write"} \\ \wedge w.ts > w.start_ts \\ \wedge w.start_ts \in key_data[k] \\ \vee \wedge w.type = \text{"rollback"} \\ \wedge w.ts = w.start_ts$$

When the lock exists, there can't be a corresponding commit record, vice versa.

$$UniqueLockOrWrite \triangleq \\ \forall k \in KEY : \\ \forall l \in key_lock[k] : \\ \forall w \in key_write[k] : \\ w.start_ts \neq l.ts$$

For each key, each record in write column should have a unique $start_ts$.

$$UniqueWrite \triangleq \\ \forall k \in KEY : \\ \forall w, w2 \in key_write[k] : \\ (w.start_ts = w2.start_ts) \Rightarrow (w = w2)$$

Snapshot Isolation

Asserts that $next_ts$ is monotonically increasing.

$$NextTsMonotonicity \triangleq \\ \forall ts \in Ts :$$

$$(ts \leq next_ts) \Rightarrow \Box(ts \leq next_ts)$$

Asserts that no *msg* would be deleted once sent.

$$\begin{aligned} & \text{MsgMonotonicity} \triangleq \\ & \wedge \forall req \in ReqMessages : \\ & \quad req \in req_msgs \Rightarrow \Box(req \in req_msgs) \\ & \wedge \forall resp \in RespMessages : \\ & \quad resp \in resp_msgs \Rightarrow \Box(resp \in resp_msgs) \end{aligned}$$

Asserts that all messages sent should have *ts* less than *next_ts*.

$$\begin{aligned} & \text{MsgTsConsistency} \triangleq \\ & \wedge \forall req \in req_msgs : \\ & \quad \wedge req.start_ts \leq next_ts \\ & \quad \wedge req.type \in \{ \text{"commit"}, \text{"resolve_committed"} \} \Rightarrow \\ & \quad \quad req.commit_ts \leq next_ts \\ & \wedge \forall resp \in resp_msgs : resp.start_ts \leq next_ts \end{aligned}$$

SnapshotIsolation is implied from the following assumptions (but is not necessary), because *SnapshotIsolation* means that:

(1) Once a transaction is committed, all keys of the transaction should be always readable or have lock on secondary *keys(eventually readable)*.

PROOF BY *CommitConsistency*, *MsgConsistency*

(2) For a given transaction, all transaction that commits after that transaction should have greater *commit_ts* than the *next_ts* at the time that the given transaction commits, so as to be able to distinguish the transactions that commits before and after from all transactions that preserved by (1).

PROOF BY *NextTsConsistency*, *MsgTsConsistency*

(3) All aborted transactions would be always not readable.

PROOF BY *AbortConsistency*, *MsgConsistency*

$$\begin{aligned} \text{SnapshotIsolation} \triangleq & \wedge \text{CommitConsistency} \\ & \wedge \text{AbortConsistency} \\ & \wedge \text{NextTsMonotonicity} \\ & \wedge \text{MsgMonotonicity} \\ & \wedge \text{MsgTsConsistency} \end{aligned}$$

THEOREM *Safety* \triangleq

$$\begin{aligned} \text{Spec} \Rightarrow \Box(& \wedge \text{TypeOK} \\ & \wedge \text{UniqueCommitOrAbort} \\ & \wedge \text{CommitConsistency} \\ & \wedge \text{AbortConsistency} \\ & \wedge \text{WriteConsistency} \\ & \wedge \text{UniqueLockOrWrite} \\ & \wedge \text{UniqueWrite} \\ & \wedge \text{SnapshotIsolation}) \end{aligned}$$