

Modeling Database Recovery

An Exercise in Stepwise Refinement

Egon Börger

Dipartimento di Informatica, Università di Pisa

<http://www.di.unipi.it/~boerger>

For details see Chapter 3.2 (Incremental Design by Refinements) of:

E. Börger, R. Stärk

Abstract State Machines

A Method for High-Level System Design and Analysis

Springer-Verlag 2003

For update info see AsmBook web page:

<http://www.di.unipi.it/AsmBook>

The Problem of Recovery from Failure in Databases

- **Goal:** design and verification by stepwise refinement of a typical recovery algorithm for a multiple-user, concurrently accessed database, guaranteeing despite of system failures
 - **atomicity** of the effect of durative database transactions (either committing all the transaction updates or aborting all of them)
 - **durability** of the effect of committed transactions
- **Basic method:**
 - Construction of a high-level database model with proof of atomicity and durability of transactions under appropriate **run constraints**, using **abstract stable, volatile, committed storage** with abstract operations READ, WRITE, COMMIT, ABORT, RECOVER, FAIL, FLUSH and abstract auxiliary fcts
 - Provably correct refinement of volatile and committed storage and related operations by stable and **cache memory with log management**
 - Provably correct refinement by
 - run-time **computation of auxiliary functions** (e.g. last committed value)
 - **sequentialization of parallelism** (e.g. log scanning during abort/recovery)
 - **computing run constraints** (e.g. on caching policy for log records)

Database Transactions and Failure

- A multiple user database is accessed by TRANSACTIONS
 - a (static or monitored) set consisting of sequences of OPERATIONS with static fcts
 - issuer: OPERATION \rightarrow TRANSACTION yields the issuing user
 - type: OPERATION \rightarrow {read, write, commit, abort}
 - loc: OPERATION \rightarrow LOC yields the read/write location
 - val: OPERATION \rightarrow VALUE yields the write value
 - atomicity: each transaction t completes normally or abnormally with one of
 - commit: values of all writes of t have to remain in db (durability) until the next overwrite by a transaction
 - abort: values of all writes of t must be replaced with the previous values which remain until the next overwrite by a transaction
 - a system failure when t is active aborts t
- so that db values after a failure reflect only updates made by transactions committed before the failure
- at the time of recovery, recovery information has to be in stable db

Recovery Goal: install committed Db as current Db

- Dynamic fcts **stableDb**, **currDb**, **commDb** : $LOC \rightarrow VALUE$ for
 - stable storage
 - current - most recent - values (over volatile or stable storage)
 - last committed values
- Auxiliary dynamic functions:
 - **writeSet** : $TRANSACTION \rightarrow 2^{LOC}$ yields the set of locations for which the given transaction has issued a write operation
 - **currOp** monitored function yielding the currently issued operation
 - **fail** : $BOOL$ monitored fct indicating the presence/absence of a system failure
 - treated as event which is consumed by firing the corresponding FAIL rule below
 - **cacheFlush** : $LOC \rightarrow BOOL$ monitored function indicating when to copy the cache value in the given location to stable storage
- **Initially** **stableDb**, **currDb**, **commDb** everywhere undefined, **writeSet** yields everywhere \emptyset .

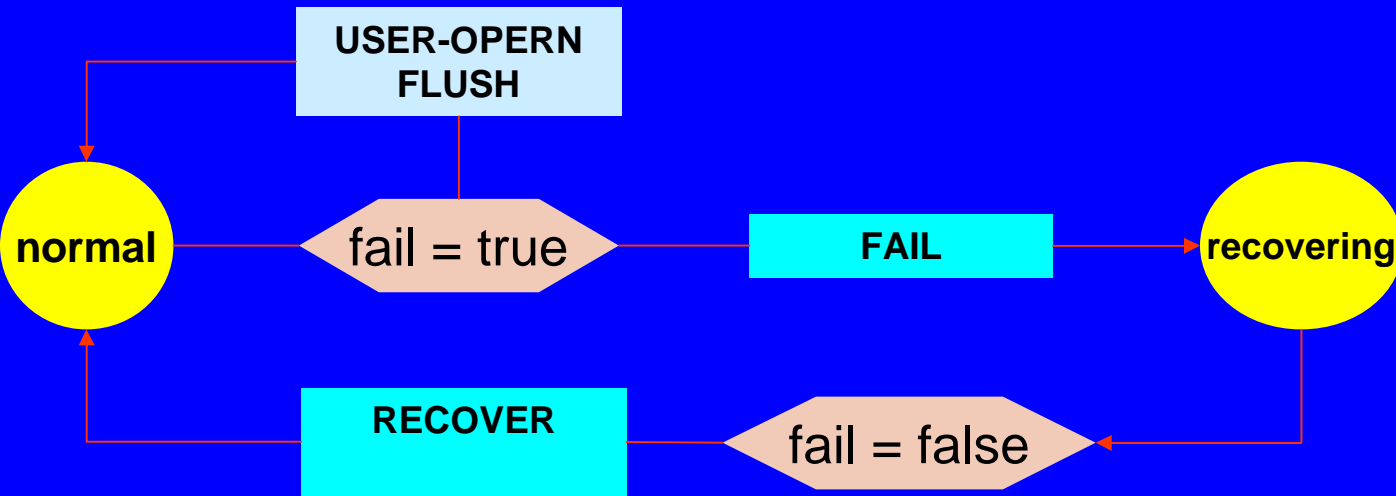
Run Constraint: Strictness of Operation Schedule

- Runs are supposed to be **strict** , i.e.
 - operations are issued only by **active** transactions
 - if transaction t writes to location l in state S and transaction t' writes to l in state $T \succ S$, then t is not active any more in state T

Run Constraint: Strictness of Operation Schedule

- Lemma. In strict runs
 - transactions only read values written by committed transactions
 - per location l there is at most one active writer at any time (so that undoing a write to l restores the value written by the committed transaction that last wrote to l)
- Def. t is **active** in state S if since its first read/write, in some $T \leq S$, neither does it commit or abort nor does the system recover in (T, S)
 - t does o iff $\text{type}(\text{currOp}) = o \ \& \ \text{issuer}(\text{currOp}) = t$
 $\quad \quad \quad \& \ \text{ctlstate} = \text{normal} \ \& \ \text{fail} = \text{false}$
 - t **commits/aborts a write to l** iff t commits/aborts $\& \ l \in \text{writeSet}(t)$
 - t **encounters a failure** in S iff in S t is active $\& \ \text{fail} = \text{true}$ holds
 - t **terminates** in S iff in S t commits or aborts or encounters a failure

High-Level Db Recovery Machine M_1



at this level of abstraction, READ has no recovery effect

M_1 is running iff fail = false

at recovery time, recovery info must be in stable storage

FAIL \circ forall $l \in \text{LOC}$ restore stable value(l)

restore stable value(l) \circ currDb(l):=stableDb(l)

FLUSH \circ forall l to be flushed Flush(l)

l to be flushed \circ cacheFlush(l) = true

RECOVER \circ forall $l \in \text{LOC}$ restore comm value(l)

Flush(l) \circ stableDb(l):=currDb(l)

COMMIT \circ commit(issuer(currOp))

restore comm value(l) \circ currDb(l):=commDb(l)

ABORT \circ abort(issuer(currOp))

commit(t) \circ forall $l \in \text{writeSet}(t)$ commDb(l):= currDb(l)

WRITE \circ write into volatile loc
record loc

abort(t) \circ forall $l \in \text{writeSet}(t)$ currDb(l):= commDb(l)

write into volatile loc \circ currDb(loc(currOp)):=val(currOp)

record loc \circ writeSet(issuer(currOp)):= writeSet(issuer(currOp)) \cup {loc(currOp)}

Proof of Durability for commDb Values in Machine M_1

- **Durability Proposition.** When a write of v to location l is committed in a state, then from the next state $\text{commDb}(l) = v$ holds until some transaction commits a new write to l .
- **Proof.**
 - WRITE by t of v to l in state R yields $\text{currDb}(l)=v$, $l \in \text{writeSet}(t)$
 - COMMIT of t in state $S \succ R$ yields $\text{commDb}(l) = \text{currDb}(l)$
- By strictness, from state R (excluded) until state S (included) no transaction WRITES to l or ABORTs a write to l and there is no system failure followed by an application of RECOVER. Therefore $\text{currDb}(l)$ is not updated in $(R, S]$ so that $\text{commDb}(l) = v$.
 - If for some T no transaction commits a write to l in $(S, T]$, $\text{commDb}(l) = v$ continues to hold in $(S, T]$.

Proof of Atomicity for currDb Values in Machine M_1

- **Atomicity Proposition.** If transaction t in state R writes value v to location l and thereafter terminates, in every following normal state T until which no other WRITE is issued the value $\text{currDb}(l)_T$ of $\text{currDb}(l)$ is
 - the lastly written new value v if t terminates by a COMMIT
 - the old value $\text{currDb}(l)_R$ before the last writing if t terminates by an ABORT or a system FAILURE
- **Def.** for $c=\text{normal}$, recovering:
 S is c iff $\text{fail}=\text{false}$ & $\text{ctl} = c$ hold in S
- **Proof.** After t issued a WRITE in state R , resulting in $\text{currDb}(l)=v$ and $l \in \text{writeSet}(t)$, there are two cases upon termination of t in a state $S \in (R, T)$:
 - Case 1. t commits in S
 - Case 2. t aborts or encounters a failure in S

Proof of Atomicity for currDb Values: Case 1

- **Case 1.** t commits in S
 - By strictness, (*) $\text{currDb}(I)$ is not updated in $(R, S]$ since from state R (excluded) until state S (included) no transaction WRITES to I or ABORTs a write to I and there is no system failure followed by an application of RECOVER.
 - COMMIT in S yields $\text{currDb}(I)_{S+1} = \text{currDb}(I)_S = \text{commDb}(I)_{S+1}$
 - This implies $\text{currDb}(I)_T = v$ by Lemma 1 and (*).
- **Lemma 1** (Preservation of committed values in $\text{currDb}(I)$ between normal states without WRITES to I) If $\text{currDb}(I)_S = \text{commDb}(I)_S$ in a normal state S and T is a later normal state until which no transaction writes to I , then $\text{currDb}(I)_T = \text{currDb}(I)_S$.

Proof of Atomicity for currDb Values: Case 2

- **Case 2.** t aborts or encounters a failure in $S \in (R, T)$:
 - By strictness, (*) $\text{commDb}(l)$ is not updated in $[R, S]$ since no transaction commits a write to l in $[R, T]$.
 - By Lemma 2, (#) $\text{currDb}(l)_R = \text{commDb}(l)_R$.
 - Subcase 1. t ABORTs in S implying $\text{currDb}(l)_T =_{\text{by Lemma 1}} \text{currDb}(l)_{S+1} = \text{commDb}(l)_S =_{\text{by (*)}} \text{commDb}(l)_R =_{\text{by (#)}} \text{currDb}(l)_R$
 - Subcase 2. Let S' be the first non-failure state in (S, T) (which does exist to reach normal T) reached by RECOVER, resulting in $\text{currDb}(l)_{S'+1} = \text{commDb}(l)_{S'} =_{\text{by (*)}} \text{commDb}(l)_R =_{\text{by (#)}} \text{currDb}(l)_R =_{\text{by Lemma 1}} \text{currDb}(l)_T$
- **Lemma 2.** In a state where WRITE to l is issued, $\text{commDb}(l) = \text{currDb}(l)$:
If t writes to l in state S , then

$$\text{commDb}(l)_S = \text{currDb}(l)_S.$$

Proof of Lemma 1

- **Lemma 1.** If $\text{currDb}(l)_S = \text{commDb}(l)_S$ in a normal state S and T is a later normal state such that in $[S, T)$ no transaction writes to l , then $\text{currDb}(l)_T = \text{currDb}(l)_S$.
- **Proof.** Induction($|[S, T)|$).
 - Case 1. $\text{fail} = \text{false}$ in $[S, T)$. Then $\text{currDb}(l) = \text{commDb}(l)$ in $[S, T)$ (since in $[S, T)$ no t writes to l) so that $\text{commDb}(l)$ does not change its value in $[S, T)$. Thus $\text{currDb}(l)_S = \text{currDb}(l)_{T-1} = \text{currDb}(l)_T$.
 - Case 2. Let S' be the last normal state with $\text{fail} = \text{true}$, with next non-failure (recovering) state S'' such that $\text{fail} = \text{false}$ in $[S''+1, T)$ and $S''+1$ is a normal state.
 - via RECOVER $\text{currDb}(l)_{S''+1} = \text{commDb}(l)_{S''}$
 - $\text{commDb}(l)_{S'} = \text{commDb}(l)_{S''} = \text{commDb}(l)_{S''+1}$ (FAIL, RECOVER do not update commDB)
 - Therefore $\text{currDb}(l)_{S''+1} = \text{commDb}(l)_{S''+1}$ so that
 - $\text{currDb}(l)_T = \text{currDb}(l)_{S''+1}$ by Case 1 since $\text{fail} = \text{false}$ in $[S''+1, T)$
 - $\text{currDb}(l)_{S''+1} = \text{currDb}(l)_S$ by ind.hyp.

Proof of Lemma 2

- **Lemma 2.** If t writes to l in state S , then
$$\text{commDb}(l)_S = \text{currDb}(l)_S.$$
- **Proof.**
 - Fact 1. $\text{commDb}(l)_R = \text{currDb}(l)_R$ and $\text{commDb}(l)$ is unchanged in $[R, S]$ for the state R reached by the last commit before S of a write to l (if case of no such commit: R =initial state, $\text{commDb}(l) = \text{currDb}(l) = \text{undef}$).
 - Fact 2. In $[R, S]$, $\text{currDb}(l)$ can be updated to a val $\neq \text{commDb}(l)$ only by a WRITE or a FAIL.
 - Case 1: a transaction t writes to l . Then t aborts before S (by strictness of runs and definition of R). ABORT reestablishes $\text{currDb}(l) = \text{commDb}(l)$.
 - Case 2: fail = true at some $R' \in [R, S)$. Then RECOVER fires at some $R'' \in (R', S)$ since S is normal, thereby reestablishing $\text{currDb}(l) = \text{commDb}(l)$.

Refinement by decomposing currDb into stableDb and cache

- **Refinement goal**
 - compute currDb as composed from stableDb and volatile **cache** : $\text{LOC} \rightarrow \text{VALUE}$
 - $\text{currDb}(l) = \text{stableDb}(l)$ if $\text{cache}(l) = \text{undef}$
= $\text{cache}(l)$ otherwise
 - compute the recording of commDb in stable storage
- **Problem**: upon system failure, in the cache
 - writes by uncommitted transactions (with values possibly flushed to stable storage) must be **undone**
 - writes of committed transactions (whose values might reside only in the cache) must be **redone** (reinstalled)

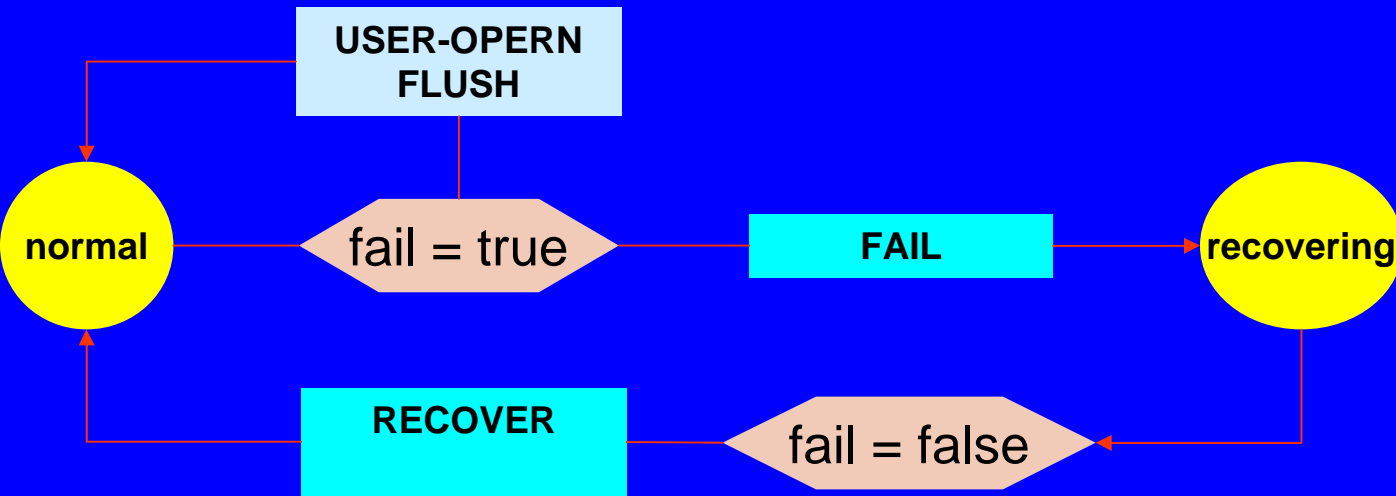
Idea for keeping track of writes and commits

- **Refinement idea:** keeping track of writes/commits via
 - a **dynamic subset** $\text{log} \subseteq \text{LOG}$ of an ordered set of **records of all current writes**, equipped with field accessing fcts
 - **issuer**: $\text{LOG} \rightarrow \text{TRANSACTION}$
 - **loc**: $\text{LOG} \rightarrow \text{Loc}$
 - **afterImage**: $\text{LOG} \rightarrow \text{VAL}$
 - log stored in stable or volatile memory, with monitored set **stableLog** $\subseteq \text{log}$, satisfying two run constraints:
 - RCS1. Upon system failure, records of last committed values must be in stable storage: **lastCommRcds** $\subseteq \text{stableLog} \subseteq \text{log}$
 - RCS2. If a location has no record in the stable log portion, then its value in stableDb remains undef and therefore may not be flushed to stable storage: **if not $\exists r \in \text{stableLog}$ s.t. $\text{loc}(r) = l$, then $\text{stableDb}(l) = \text{undef}$**
 - dynamic list of **committed** transactions, residing in stable storage

Additional signature to compute commDb

- **cacheRemove**: $\text{LOC} \rightarrow \text{BOOL}$ monitored fct to be used for possible removal of cache values upon flushing
- auxiliary derived functions for computing
$$\text{commDb}(l) = \text{afterImage}(\text{lastRcd}(l, \text{commRcds}))$$
 - **lastRcd**: $\text{LOC} \times 2^{\text{LOG}} \rightarrow \text{LOG}$
 - $\text{lastRcd}(l, L) = \max\{r \in L \mid \text{loc}(r) = l\}$
 - **commRcds** = $\{r \in \text{log} \mid \text{committed}(\text{issuer}(r)) = \text{true}\}$
 - **lastCommRcds** =
$$\{r \in \text{commRcds} \mid r = \text{lastRcd}(\text{loc}(r), \text{commRcds})\}$$
 - **undoRcd**(r) = $\max\{r' \in \text{commRcds} \mid r' < r \ \& \ \text{loc}(r') = \text{loc}(r)\}$
(undef if - initially - no such r' exists)
- standard order related fcts **next**, **lastLog**, **max** on LOG
 - **RCS3**. **cacheRemove**(l) = false upon reading/writing to l and upon undoing/redoing an l -record

Refined Db Recovery Machine M_2 : Operation Refinement



FAIL \circ forall $l \in \text{LOC}$ restore stable value(l)
 $\text{log} := \text{stableLog}$

restore stable value(l) \circ currDb(l) := stableDb(l)

restore stable value(l) \circ cache(l) := undef

FLUSH \circ forall l to be flushed Flush(l)

if cacheRemove(l)
 then cache(l) :=
 undef

l to be flushed \circ cacheFlush(l) = true

& cache(l)
 $\neq \text{undef}$

Flush(l) \circ stableDb(l) := currDb(l)

Flush(l) \circ stableDb(l) := cache(l)

Refined Db Recovery Machine M_2 : Operation Refinement (Cont'd)

RECOVER \circ

forall $l \in \text{LOC}$
restore comm value(l)

restore comm value(l) \circ $\text{currDb}(l) := \text{commDb}(l)$

let $r = \text{lastRcd}(l, \text{log})$ if $r \neq \text{undef}$ then
if $\text{committed}(\text{issuer}(r))$ then $\text{REDO}(r)$
else $\text{UNDO}(r)$

$\text{commit}(t)$ \circ forall $l \in \text{writeSet}(t)$ $\text{commDb}(l) := \text{currDb}(l)$

$\text{committed}(t) := \text{true}$

COMMIT \circ $\text{commit}(\text{issuer}(\text{currOp}))$

ABORT \circ $\text{abort}(\text{issuer}(\text{currOp}))$

$\text{abort}(t)$ \circ forall $l \in \text{writeSet}(t)$ $\text{currDb}(l) := \text{commDb}(l)$

forall $r \in \text{log}$ s.t. $\text{issuer}(r) = t$ $\text{UNDO}(r)$

WRITE \circ write into volatile loc
record loc

write into volatile loc \circ $\text{currDb}(\text{loc}(\text{currOp})) := \text{val}(\text{currOp})$

$\text{cash}(\text{loc}(\text{currOp})) := \text{val}(\text{currOp})$

record loc \circ $\text{writeSet}(\text{issuer}(\text{currOp})) := \text{writeSet}(\text{issuer}(\text{currOp})) \cup \{\text{loc}(\text{currOp})\}$

$\text{WRITELOG}(\text{next}(\text{lastLog}))$

READ \circ if $\text{cache}(\text{loc}(\text{currOp})) = \text{undef}$ then $\text{cache}(\text{loc}(\text{currOp})) := \text{stableDb } \text{loc}(\text{currOp})$

WRITELOG(r) \circ

$\text{issuer}(r) := \text{issuer}(\text{currOp})$
 $\text{loc}(r) := \text{loc}(\text{currOp})$
 $\text{afterImage}(r) := \text{val}(\text{currOp})$
 $\text{log} := \text{log} \cup \{r\}$

UNDO(r) \circ

$\text{cache}(\text{loc}(r)) := \text{afterImage}(\text{undoRcd}(r))$

REDO(r) \circ

$\text{cache}(\text{loc}(r)) := \text{afterImage}(r)$

Refinement diagram: definition of equivalence

- corresponding **locations of interest** must be identical via the following definitions for the correspondence:
 - $\text{currDb}(l)_1 \quad \text{currDb}(l)_2 = \text{stableDb}(l) \quad \text{if } \text{cache}(l) = \text{undef}$
 $\quad \quad \quad = \text{cache}(l) \quad \quad \text{otherwise}$
 - $\text{writeSet}(t)_1$
 $\quad \quad \text{writeSet}(t)_2 = \{l \in \text{LOC} \mid \exists r \in \text{log}: \text{issuer}(r) = t \ \& \ \text{loc}(r) = l\}$
 - $\text{commDb}(l)_1 \quad \text{commDb}(l)_2 =$
 $\quad \quad \text{afterImage}(\text{lastRcd}(l, \text{commRcds}))$
 - **homonymous functions** stableDb , etc.
- **one-to-one correspondence of rule applications** via homonymy of rules

Correctness of the Refinement

- **Equivalence Theorem.** In runs of M_1 and M_2 which are started in equivalent initial states, in every pair of corresponding states corresponding locations are equivalent.
- **Proof:** Follows by run induction from the following equivalence of the effect of every M_2 -operation macro to the effect of the homonymous M_1 -macro, for every state S reached by M_2 :
 - FLUSH: if $\text{cacheFlush}(l)=\text{true}$ in S , then $\text{stableDb}(l)_{S+1} = \text{currDb}(l)_S$
 - WRITE: if t writes v to l in S , then $\text{currDb}(l)_{S+1} = v$ and $l \in \text{writeSet}(t)_{S+1}$
 - COMMIT: if t commits in S & $l \in \text{writeSet}(t)_S$, then $\text{commDb}(l)_{S+1} = \text{currDb}(l)_S$
 - ABORT: if t aborts in S & $l \in \text{writeSet}(t)_S$, then $\text{currDb}(l)_{S+1} = \text{commDb}(l)_S$
 - FAIL: if $\text{fail}=\text{true}$ in S , then $\text{currDb}(l)_{S+1} = \text{stableDb}(l)_S$ for every l
 - RECOVER: if the system recovers in S , then $\text{currDb}(l)_{S+1} = \text{commDb}(l)_S$ for every l

Equivalence Proof for FLUSH and WRITE

- To show for FLUSH: if $\text{cacheFlush}(l)_S = \text{true}$, then
$$\text{stableDb}(l)_{S+1} = \text{currDb}(l)_S$$
- Proof. By the definition of $\text{currDb}(l)$ in M_2 there are two cases to distinguish (in state S):
 - Case 1. $\text{cache}(l)_S = \text{undef}$. Then by definition $\text{currDb}(l)_S = \text{stableDb}(l)_S$ and by applying FLUSH $\text{stableDb}(l)$ is not changed. Therefore $\text{stableDb}(l)_{S+1} = \text{stableDb}(l)_S = \text{currDb}(l)_S$.
 - Case 2. Otherwise. Then by definition $\text{currDb}(l)_S = \text{cache}(l)_S$ and by firing FLUSH $\text{stableDb}(l)_{S+1} = \text{cache}(l)_S = \text{currDb}(l)_S$.
- To show for WRITE : if t writes v to l in S , then
$$\text{currDb}(l)_{S+1} = v \text{ and } l \in \text{writeSet}(t)_{S+1}$$
- Proof. By WRITE of v to l in state S
 - $\text{cache}(l)_{S+1} = v$, so that by definition $\text{currDb}(l)_{S+1} = v$.
 - \log has a new l -record with issuer t , so that $l \in \text{writeSet}(t)_{S+1}$.

Equivalence Proof for FAIL and RECOVER

- To show for FAIL : if fail=true in S, then $\text{currDb}(l)_{S+1} = \text{stableDb}(l)_S$
- **Proof.** By FAIL in S, $\text{cache}(l)_{S+1} = \text{undef}$ for every l, so that by definition $\text{currDb}(l)_{S+1} = \text{stableDb}(l)_S$.
- To show for RECOVER: if the system recovers in S, then for every l $\text{currDb}(l)_{S+1} = \text{commDb}(l)_S$
- **Proof.** Case 1: in S there is no l-record in log. Then there has been no committed write to l, so that $\text{currDb}(l)_{S+1} =_{\text{RECOVER}} \text{currDb}(l)_S =_{\text{FAIL in S-1}} \text{stableDb}(l)_S =_{\text{FAIL in S-1}} \text{stableDb}(l)_{S-1} =_{\text{initialization}} \text{undef}$ and $\text{commDb}(l)_S = \text{undef}$ by definition.
- **Proof.** Case 2: $r = \text{lastRcd}(l, \text{log})$ is defined. Let $t = \text{issuer}(r)$.
 - Case 2.1. $\text{committed}(t)_S = \text{true}$. Then by REDO and definition $\text{cache}(l)_{S+1} = \text{afterImage}(r)_S =_{\text{case 2.1 assumption}} \text{commDb}(l)_S$.
 - Case 2.2. otherwise, by UNDO $\text{currDb}(l)_{S+1} = \text{cache}(l)_{S+1} = \text{afterImage}(r')_S$ for the last committed l-record $r' < r$ in log_S . Since r is the last l-record in log_S though not committed, r' is the last committed l-record in log_S . Thus $\text{afterImage}(r')_S = \text{commDb}(l)_S$.

Equivalence Proof for COMMIT and ABORT

- To show for COMMIT : if t commits in S & $l \in \text{writeSet}(t)_S$, then $\text{commDb}(l)_{S+1} = \text{currDb}(l)_S$
- **Proof.** By $l \in \text{writeSet}(t)_S$ there is a record in log with issuer t & location l , so that t writes a value v to l in some state $R < S$. By strictness, t is active in S . By Lemma below this implies for the last l -record $r \in \log_S$ that $\text{currDb}(l)_S = v = \text{afterImage}(r)_S$. By applying COMMIT in S $\text{committed}(t)_{S+1} = \text{true}$. Therefore $\text{currDb}(l)_S = \text{afterImage}(r)_S = \text{commDb}(l)_{S+1}$ by definition.
- To show for ABORT : if t aborts in S & $l \in \text{writeSet}(t)_S$, then $\text{currDb}(l)_{S+1} = \text{commDb}(l)_S$.
- **Proof.** By $l \in \text{writeSet}(t)_S$ there is an l -record $r \in \log_S$ with issuer t . By strictness, t is active in S . By the Lemma below r is the last l -record $r \in \log_S$. By UNDO in S , $\text{cache}(l)_{S+1} = \text{afterImage}(r')_S$ for the last committed l -record $r' < r$ in \log_S . Since r is the last l -record in \log_S , r' is the last committed l -record in \log_S . Thus $\text{afterImage}(r')_S = \text{commDb}(l)_S$. Therefore $\text{currDb}(l)_{S+1} = \text{cache}(l)_{S+1} = \text{commDb}(l)_S$.

Lemma

- **Lemma.** (Persistence of write effects of active transactions) If t writes v to I in state R and is still active in state $S \succ R$, then
 - the last I -record in \log_S is the record r written at R
 - $\text{currDb}(I)_S = v$
- **Proof: Exercise.**
 - Hint: use induction on $|\{S \mid S \succ R\}|$. See Lemma 5 in Gurevich, Soparkar, Wallace op.cit.

Refinement Exercises

- Refine M_2 to M_3 by **computing derived functions**, e.g. `afterImage(undoRcd(r))`. Formulate the correctness of this refinement and prove it.
 - Hint. Refine WRITELOG by logging also the before-image of the written location (fetched from stable or cache memory), to be used in UNDO.
 - See Sect.4.1 in the paper by Gurevich, Soparkar, Wallace
- Refine M_3 to M_4 by **sequentializing “forall”** as iterated log scanning in FAIL, ABORT, RECOVER. Formulate and prove the correctness of this refinement.
 - Hint. Refine FAIL to initialize an iterator `thisRec` for scanning all records in `stableLog` for recovery. Refine ABORT to first initialize the iterator `thisRec` and then scanning all log records of the considered transaction. Refine WRITELOG by linking the new log to the current transaction’s last log record. See Sect.4.2 in the paper by Gurevich, Soparkar, Wallace

Refinement Exercise

- Refine M_4 to M_5 by implementing run conditions, e.g. on cache policy for log:
 - RCS1. The records of all last committed writes must be in stable storage (for reinstallation during recovery), formally: $\text{lastCommRcds} \subseteq \text{stableLog} \subseteq \text{log}$
 - Hint: Refine COMMIT by flushing log to stableLog
 - RCS2. If there is no stableLog record of a write to l , then $\text{stableDb}(l) = \text{undef}$
 - Hint: maintain for each l the last log record with l (refining WRITE); constrain flushing from cache to stableDb by a check of the index of the last record for the location of the value to be flushed against the index of the last record in stable storage.
- Formulate and prove the correctness of this refinement. See Sect.4.3 in the paper by Gurevich, Soparkar, Wallace

Reference

- Y. Gurevich and N. Soparkar and C. Wallace:
Formalizing Database Recovery
 - In: J. of Universal Computer Science 3 (4) 1997
- **E. Börger, R. Stärk:** Abstract State Machines. A
Method for High-Level System Design and Analysis
Springer-Verlag 2003, see
<http://www.di.unipi.it/AsmBook>