

# Formalization of (Conflict-)Serializability and Strict Two-Phase Locking

Dmitriy Traytel

February 11, 2025

## Abstract

Concurrency control is an essential component of any transactional database management system, which is responsible for providing isolation (the “I” in ACID) to transactions. Formally, concurrency control aims to achieve serializability: a way to rearrange the actions of concurrently executing transactions that eliminates concurrency while leaves the database modifications unchanged. In this small entry, we define serializability, a syntactic over-approximation called conflict-serializability, and characterize schedules generated by the frequently used concurrency control mechanism of strict two-phase locking (S2PL). We also prove two inclusions: S2PL implies conflict-serializability, which in turn implies serializability. The formalization is based on standard material from an advanced database systems course [1, Chapter 17].

## 1 Transactions

We work with a rather abstract model of transactions comprised of read/write actions.

Read/written values are natural numbers.

**type-alias**  $val = nat$

Transactions  $'xid$  may read from/write two addresses  $'addr$ .

**datatype**  $( 'xid, 'addr ) action = isRead: Read (xid-of: \langle 'xid \rangle) (addr-of: 'addr) | isWrite: Write (xid-of: \langle 'xid \rangle) (addr-of: 'addr)$

A schedule is a sequence of actions.

**type-synonym**  $( 'xid, 'addr ) schedule = \langle ( 'xid, 'addr ) action list \rangle$

A database, which is being modified by the read/write actions, maps addresses to values.

**type-synonym**  $'addr db = \langle 'addr \Rightarrow val \rangle$

Each transaction has a local state, which is represented as the list of previously read values (and the addresses they have been read from).

**type-synonym**  $'addr\ xstate = \langle ('addr \times val)\ list \rangle$

The values written by a transaction are given by a higher-order parameter and may depend on the previously read values.

**context fixes**  $write\ logic :: \langle 'xid \Rightarrow 'addr\ xstate \Rightarrow 'addr \Rightarrow val \rangle$  **begin**

Read values are recorded in the transaction's local state; writes modify the database.

**fun**  $action\ effect :: \langle ('xid, 'addr)\ action \Rightarrow ('xid \Rightarrow 'addr\ xstate) \times 'addr\ db \Rightarrow ('xid \Rightarrow 'addr\ xstate) \times 'addr\ db \rangle$  **where**  
 $\langle action\ effect\ (Read\ xid\ addr)\ (xst, db) = (xst(xid := (addr, db\ addr)\ \# xst\ xid), db) \rangle$   
 $| \langle action\ effect\ (Write\ xid\ addr)\ (xst, db) = (xst, db(addr := write\ logic\ xid\ (xst\ xid)\ addr)) \rangle$

We are interested in how a schedule modifies the database (local state changes are discarded at the end).

**definition**  $schedule\ effect :: \langle ('xid, 'addr)\ schedule \Rightarrow 'addr\ db \Rightarrow 'addr\ db \rangle$  **where**  
 $\langle schedule\ effect\ s\ db = snd\ (fold\ action\ effect\ s\ (\lambda\_. [], db)) \rangle$

**end**

Actions that belong to the same transaction.

**definition**  $eq\ xid$  **where**  
 $\langle eq\ xid\ a\ b = (xid\ of\ a = xid\ of\ b) \rangle$

## 2 Serial and Serializable Schedules

**declare**  $length\ dropWhile\ le[termination\ simp]$

A serial schedule does not interleave actions of different transactions.

**fun**  $serial :: \langle ('xid, 'addr)\ schedule \Rightarrow bool \rangle$  **where**  
 $\langle serial\ [] = True \rangle$   
 $| \langle serial\ (a\ \#\ as) = (let\ bs = dropWhile\ (\lambda b. eq\ xid\ a\ b)\ as\ in\ serial\ bs \wedge xid\ of\ a \notin xid\ of\ 'set\ bs) \rangle$

A schedule  $s$  can be rearranged into schedule  $t$  by a permutation  $\pi$ , which preserves the relative order of actions related by  $eq$ .

**definition**  $permutes\ upto$  **where**  
 $\langle permutes\ upto\ eq\ \pi\ s\ t =$   
 $(bij\ betw\ \pi\ \{..\langle length\ s \rangle\}\ \{..\langle length\ t \rangle\}) \wedge$   
 $(\forall i < length\ s. s\ !\ i = t\ !\ \pi\ i) \wedge$   
 $(\forall i < length\ s. \forall j < length\ s. i < j \wedge eq\ (s\ !\ i)\ (s\ !\ j) \longrightarrow \pi\ i < \pi\ j) \rangle$

**lemma** *permutes-upto-Nil[simp]*:  $\langle \text{permutes-upto } R \ \pi \ \square \ \square \rangle$   
 $\langle \text{proof} \rangle$

Two schedules are equivalent if one can be rearranged into another without rearranging the actions of each transaction and in addition they have the same effect on any database for any fixed write logic.

**abbreviation** *equivalent* ::  $\langle ('xid, 'addr) \text{ schedule} \Rightarrow ('xid, 'addr) \text{ schedule} \Rightarrow \text{bool} \rangle$   
**where**

$\langle \text{equivalent } s \ t \equiv (\exists \pi. \text{permutes-upto } eq\text{-xid } \pi \ s \ t \wedge (\forall \text{write-logic } db. \text{schedule-effect } \text{write-logic } s \ db = \text{schedule-effect } \text{write-logic } t \ db)) \rangle$

A schedule is serializable if it is equivalent to some serial schedule. Serializable schedules thus provide isolation: even though actions of different transactions may be interleaved, the effect from the point of view of each transaction is as if the transaction was the only one executing in the system (as is the case in serial schedules).

**definition** *serializable* ::  $\langle ('xid, 'addr) \text{ schedule} \Rightarrow \text{bool} \rangle$  **where**  
 $\langle \text{serializable } s = (\exists t. \text{serial } t \wedge \text{equivalent } s \ t) \rangle$

### 3 Conflict Serializable Schedules

Two actions of different transactions are conflicting if they access the same address and at least one of them is a write.

**definition** *conflict* **where**

$\langle \text{conflict } a \ b = (\text{xid-of } a \neq \text{xid-of } b \wedge \text{addr-of } a = \text{addr-of } b \wedge (\text{isWrite } a \vee \text{isWrite } b)) \rangle$

Two schedules are conflict-equivalent if one can be rearranged into another without rearranging conflicting actions or actions of one transaction. Note that unlike equivalence, the conflict-equivalence notion is purely syntactic, i.e., not talking about databases and action/schedule effects.

**abbreviation** *conflict-equivalent* ::  $\langle ('xid, 'addr) \text{ schedule} \Rightarrow ('xid, 'addr) \text{ schedule} \Rightarrow \text{bool} \rangle$  **where**

$\langle \text{conflict-equivalent } s \ t \equiv (\exists \pi. \text{permutes-upto } (\text{sup } eq\text{-xid } \text{conflict}) \ \pi \ s \ t) \rangle$

A schedule is conflict-serializable if it is conflict equivalent to some serial schedule.

**definition** *conflict-serializable* ::  $\langle ('xid, 'addr) \text{ schedule} \Rightarrow \text{bool} \rangle$  **where**  
 $\langle \text{conflict-serializable } s = (\exists t. \text{serial } t \wedge \text{conflict-equivalent } s \ t) \rangle$

### 4 Conflict-Serializability Implies Serializability

In the following, we prove that the syntactic notion implies the semantic one. The key observation is that swapping non-conflicting actions of different transactions preserves the overall effect on the database.

**lemma** *swap-actions*:  $\langle \neg \text{conflict } a \ b \implies \neg \text{eq-xid } a \ b \implies$   
 $\text{action-effect } wl \ a \ (\text{action-effect } wl \ b \ st) = \text{action-effect } wl \ b \ (\text{action-effect } wl \ a \ st) \rangle$   
 $\langle \text{proof} \rangle$

**lemma** *swap-many-actions*:  $\langle \forall i < \text{length } p. \neg \text{conflict } a \ (p \ ! \ i) \wedge \neg \text{eq-xid } a \ (p \ ! \ i) \implies$   
 $\text{action-effect } wl \ a \ (\text{fold } (\text{action-effect } wl) \ p \ st) = \text{fold } (\text{action-effect } wl) \ p \ (\text{action-effect } wl \ a \ st) \rangle$   
 $\langle \text{proof} \rangle$

**lemma** *fold-action-effect-eq*:  
**assumes**  $\langle t = p \ @ \ a \ \# \ u \rangle$   
**shows**  
 $\langle \text{fold } (\text{action-effect } wl) \ s \ (\text{action-effect } wl \ a \ st) =$   
 $\text{fold } (\text{action-effect } wl) \ (p \ @ \ u) \ (\text{action-effect } wl \ a \ st) \implies$   
 $\forall i < \text{length } p. \neg \text{conflict } a \ (p \ ! \ i) \wedge \neg \text{eq-xid } a \ (p \ ! \ i) \implies$   
 $\text{fold } (\text{action-effect } wl) \ (a \ \# \ s) \ st = \text{fold } (\text{action-effect } wl) \ t \ st \rangle$   
 $\langle \text{proof} \rangle$

**definition** *shift where*  
 $\langle \text{shift } \pi = (\lambda i. \text{if } i < \pi \ 0 \ \text{then } i \ \text{else } i - 1) \ o \ \pi \ o \ \text{Suc} \rangle$

**lemma** *bij-betw-remove*:  $\langle \text{bij-betw } f \ A \ B \implies x \in A \implies \text{bij-betw } f \ (A - \{x\}) \ (B - \{f \ x\}) \rangle$   
 $\langle \text{proof} \rangle$

**lemma** *permutes-upto-shift*:  
**assumes**  $\langle \text{permutes-upto } eq \ \pi \ (a \ \# \ s) \ t \rangle$   
**shows**  $\langle \text{permutes-upto } eq \ (\text{shift } \pi) \ s \ (\text{take } (\pi \ 0) \ t \ @ \ \text{drop } (\text{Suc } (\pi \ 0)) \ t) \rangle$   
 $\langle \text{proof} \rangle$

**lemma** *permutes-upto-prefix-upto*:  
**assumes**  $\langle \text{permutes-upto } eq \ \pi \ (t \ ! \ \pi \ 0 \ \# \ s) \ t \ \langle i < \pi \ 0 \rangle$   
**shows**  $\langle \neg \text{eq } (t \ ! \ \pi \ 0) \ (t \ ! \ i) \rangle$   
 $\langle \text{proof} \rangle$

**lemma** *conflict-equivalent-imp-equivalent*:  
**assumes**  $\langle \text{conflict-equivalent } s \ t \rangle$   
**shows**  $\langle \text{equivalent } s \ t \rangle$   
 $\langle \text{proof} \rangle$

**theorem** *conflict-serializable-imp-serializable*:  $\langle \text{conflict-serializable } s \implies \text{serializable } s \rangle$   
 $\langle \text{proof} \rangle$

## 5 Schedules Generated by Strict Two-Phase Locking (S2PL).

To enforce conflict-serializability database management systems use locks. Locks come in two kinds: shared locks for reads and exclusive locks for writes. An address can be accessed in a reading fashion by multiple transactions, each holding a shared lock. If one transaction however holds an exclusive locks to write to an address, then no other transaction can hold a lock (neither shared nor exclusive) for the same address.

**datatype** *'addr lock* = *S (addr-of: 'addr) | X (addr-of: 'addr)*

**fun** *lock-for where*

*⟨lock-for (Read - addr) = S addr⟩*  
*| ⟨lock-for (Write - addr) = X addr⟩*

**definition** *valid-locks where*

*⟨valid-locks locks = (∀ addr xid1 xid2. X addr ∈ locks xid1 →*  
*X addr ∈ locks xid2 ∨ S addr ∈ locks xid2 → xid1 = xid2)⟩*

A frequently used lock strategy is strict two phase locking (S2PL) in which transactions attempt to acquire locks gradually (whenever they want to execute an action that needs a particular lock) and release them all at once at the end of each transaction.

The following predicate checks whether a schedule could have been generated using the S2PL strategy. To this end, the predicate checks for each action, whether the corresponding lock could have been acquired by the transaction executing the action. We also allow lock upgrades (from shared to exclusive), i.e., one transaction can hold both a shared and an exclusive lock

As in our model there is no explicit transaction end marker (commit), we treat each transaction as finished immediately when it has executed its last action in the given schedule. This is the moment, when the transaction's locks are released.

**fun** *s2pl :: ⟨('xid ⇒ 'addr lock set) ⇒ ('xid, 'addr) schedule ⇒ bool⟩ where*

*⟨s2pl locks [] = True⟩*  
*| ⟨s2pl locks (a # s) =*  
*(let xid = xid-of a; addr = action.addr-of a*  
*in if ∃ xid'. xid' ≠ xid ∧ (X addr ∈ locks xid' ∨ isWrite a ∧ S addr ∈ locks*  
*xid')*  
*then False*  
*else s2pl (locks(xid := if xid ∉ xid-of ' set s then {} else locks xid ∪ {lock-for*  
*a})) s)⟩*

We prove in the following that S2PL schedules are conflict-serializable (and thus also serializable). The proof proceeds by induction on the number of transactions in a schedule. To construct the conflict-equivalent serial

schedule we always move the actions of the transaction that finished first in our S2PL schedule to the front. To do so we show that these actions are not conflicting with any preceding actions (due to the acquired/held locks).

**lemma** *conflict-equivalent-trans*:

$\langle \text{conflict-equivalent } s \ t \implies \text{conflict-equivalent } t \ u \implies \text{conflict-equivalent } s \ u \rangle$   
 $\langle \text{proof} \rangle$

**lemma** *conflict-equivalent-append*:  $\langle \text{conflict-equivalent } s \ t \implies \text{conflict-equivalent } (u \ @ \ s) \ (u \ @ \ t) \rangle$

$\langle \text{proof} \rangle$

**lemma** *conflict-equivalent-Cons*:  $\langle \text{conflict-equivalent } s \ t \implies \text{conflict-equivalent } (a \ # \ s) \ (a \ # \ t) \rangle$

$\langle \text{proof} \rangle$

**lemma** *conflict-equivalent-rearrange*:

**assumes**  $\langle \bigwedge i \ j. \text{xid-of } (s \ ! \ i) = \text{xid} \implies j < i \implies i < \text{length } s \implies \neg \text{conflict } (s \ ! \ j) \ (s \ ! \ i) \rangle$

**shows**  $\langle \text{conflict-equivalent } s \ (\text{filter } ((=) \ \text{xid} \circ \ \text{xid-of}) \ s \ @ \ \text{filter } (\text{Not} \circ \ (=) \ \text{xid} \circ \ \text{xid-of}) \ s) \rangle$

$\langle \text{is } \langle \text{conflict-equivalent } s \ (\ ?\text{filter } s) \rangle \rangle$

$\langle \text{proof} \rangle$

**lemma** *serial-append*:

$\langle \text{serial } s \implies \text{serial } t \implies \text{xid-of } \text{'set } s \cap \text{xid-of } \text{'set } t = \{\} \implies \text{serial } (s \ @ \ t) \rangle$   
 $\langle \text{proof} \rangle$

**lemma** *serial-same-xid*:  $\langle \forall x \in \text{set } s. \text{xid-of } x = \text{xid} \implies \text{serial } s \rangle$

$\langle \text{proof} \rangle$

**lemma** *conflict-equivalent-same-set*:  $\langle \text{conflict-equivalent } s \ t \implies \text{set } s = \text{set } t \rangle$

$\langle \text{proof} \rangle$

**lemma** *s2pl-filter*:

$\langle \text{s2pl } \text{locks } s \implies \text{s2pl } (\text{locks}(\text{xid} := \{\})) \ (\text{filter } (\text{Not} \circ \ (=) \ \text{xid} \circ \ \text{xid-of}) \ s) \rangle$

$\langle \text{proof} \rangle$

**lemma** *valid-locks-grab[simp]*:  $\langle \text{valid-locks } \text{locks} \implies$

$\neg (\exists \text{xid}'. \text{xid}' \neq \text{xid-of } a \wedge$

$(X \ (\text{action.addr-of } a) \in \text{locks } \text{xid}' \vee \text{isWrite } a \wedge S \ (\text{action.addr-of } a) \in \text{locks } \text{xid}') \implies$

$\text{valid-locks } (\text{locks}(\text{xid-of } a := \text{insert } (\text{lock-for } a) \ (\text{locks } (\text{xid-of } a)))) \rangle$

$\langle \text{proof} \rangle$

**lemma** *s2pl-suffix*:  $\langle \text{valid-locks } \text{locks} \implies \text{s2pl } \text{locks } (s \ @ \ t) \implies$

$\forall a \in \text{set } s. \exists b \in \text{set } t. \text{eq-xid } a \ b \implies$

$\exists \text{locks}'. \text{valid-locks } \text{locks}' \wedge (\forall \text{xid}. \text{locks } \text{xid} \subseteq \text{locks}' \ \text{xid}) \wedge \text{s2pl } \text{locks}' \ t \rangle$

$\langle \text{proof} \rangle$

**lemma** *set-drop*:  $\langle l \leq \text{length } xs \implies \text{set } (\text{drop } l \text{ } xs) = \text{nth } xs \text{ } \{l..<\text{length } xs\} \rangle$   
 $\langle \text{proof} \rangle$

**lemma** *drop-eq-Cons*:  $\langle i < \text{length } xs \implies \text{drop } i \text{ } xs = xs ! i \# \text{drop } (\text{Suc } i) \text{ } xs \rangle$   
 $\langle \text{proof} \rangle$

**theorem** *s2pl-conflict-serializable*:  $\langle s2pl (\lambda-. \{\}) s \implies \text{conflict-serializable } s \rangle$   
 $\langle \text{proof} \rangle$

**corollary** *s2pl-serializable*:  $\langle s2pl (\lambda-. \{\}) s \implies \text{serializable } s \rangle$   
 $\langle \text{proof} \rangle$

## 6 Example Executing S2PL

To make the S2PL check executable regardless of the transaction id type, we restrict the quantification to transaction ids that are occurring in the schedule.

**fun** *s2pl-code* ::  $\langle ('xid \Rightarrow 'addr \text{ lock set}) \Rightarrow ('xid, 'addr) \text{ schedule} \Rightarrow \text{bool} \rangle$  **where**  
 $\langle s2pl\text{-code } locks [] = \text{True} \rangle$   
 $| \langle s2pl\text{-code } locks (a \# s) =$   
    $(\text{let } xid = xid\text{-of } a; \text{addr} = \text{action.addr-of } a$   
    $\text{in if } \exists xid' \in xid\text{-of ' set } s. xid' \neq xid \wedge (X \text{ addr} \in \text{locks } xid' \vee \text{isWrite } a \wedge S$   
    $\text{addr} \in \text{locks } xid')$   
    $\text{then False}$   
    $\text{else } s2pl\text{-code } (\text{locks}(xid := \text{if } xid \notin xid\text{-of ' set } s \text{ then } \{\} \text{ else } \text{locks } xid \cup$   
    $\{\text{lock-for } a\})) s \rangle$

**lemma** *s2pl-code-cong*:  $\langle \bigwedge xid. xid \in xid\text{-of ' set } s \implies f \text{ } xid = g \text{ } xid \implies$   
 $\langle \bigwedge xid. xid \notin xid\text{-of ' set } s \implies f \text{ } xid = \{\} \implies$   
 $s2pl \text{ } f \text{ } s = s2pl\text{-code } g \text{ } s$   
 $\langle \text{proof} \rangle$

**lemma** *s2pl-code[code-unfold]*:  $s2pl (\lambda-. \{\}) s = s2pl\text{-code } (\lambda-. \{\}) s$   
 $\langle \text{proof} \rangle$

**definition** *TB* =  $(0 :: \text{nat})$

**definition** *TA* =  $(1 :: \text{nat})$

**definition** *TC* =  $(2 :: \text{nat})$

**definition** *AX* =  $(0 :: \text{nat})$

**definition** *AY* =  $(1 :: \text{nat})$

**definition** *AZ* =  $(2 :: \text{nat})$

Good example involving a lock upgrade by *TA* and *TB*

**lemma**  $\langle s2pl (\lambda-. \{\})$

$[ \text{Write } TB \text{ } AZ, \text{Read } TA \text{ } AX, \text{Read } TC \text{ } AX, \text{Write } TB \text{ } AY, \text{Write } TC \text{ } AY, \text{Write } TA \text{ } AX, \text{Write } TA \text{ } AY ] \rangle$

$\langle \text{proof} \rangle$

Bad example: *TC* cannot acquire exclusive lock for *AY*, which is already held by *TA*

**lemma**  $\langle \neg \text{spl} (\lambda-. \{\}) \rangle$

$\langle \text{Read } TA \text{ } AX, \text{Read } TB \text{ } AX, \text{Read } TC \text{ } AX, \text{Write } TA \text{ } AY, \text{Write } TC \text{ } AY, \text{Write } TB \text{ } AY, \text{Write } TA \text{ } AZ \rangle$

$\langle \text{proof} \rangle$

**hide-const** *TB TA TC AX AY AZ*

**hide-fact** *TB-def TA-def TC-def AX-def AY-def AZ-def*

## References

- [1] R. Ramakrishnan and J. Gehrke. *Database Management Systems*. McGraw-Hill Higher Education, 3rd edition, January 2003.