

ROBUSTNESS AGAINST TRANSACTIONAL CAUSAL CONSISTENCY

SIDI MOHAMED BEILLAHI, AHMED BOUAJJANI, AND CONSTANTIN ENEA

Université de Paris, IRIF, CNRS, F-75013 Paris, France

e-mail address: beillahi@irif.fr

e-mail address: abou@irif.fr

e-mail address: cenea@irif.fr

ABSTRACT. Distributed storage systems and databases are widely used by various types of applications. Transactional access to these storage systems is an important abstraction allowing application programmers to consider blocks of actions (i.e., transactions) as executing atomically. For performance reasons, the consistency models implemented by modern databases are weaker than the standard serializability model, which corresponds to the atomicity abstraction of transactions executing over a sequentially consistent memory. Causal consistency for instance is one such model that is widely used in practice.

In this paper, we investigate application-specific relationships between several variations of causal consistency and we address the issue of verifying automatically if a given transactional program is robust against causal consistency, i.e., all its behaviors when executed over an arbitrary causally consistent database are serializable. We show that programs without write-write races have the same set of behaviors under all these variations, and we show that checking robustness is polynomial time reducible to a state reachability problem in transactional programs over a sequentially consistent shared memory. A surprising corollary of the latter result is that causal consistency variations which admit incomparable sets of behaviors admit comparable sets of robust programs. This reduction also opens the door to leveraging existing methods and tools for the verification of concurrent programs (assuming sequential consistency) for reasoning about programs running over causally consistent databases. Furthermore, it allows to establish that the problem of checking robustness is decidable when the programs executed at different sites are finite-state.

1. INTRODUCTION

Distribution and replication are widely adopted in order to implement storage systems and databases offering performant and available services. The implementations of these systems must ensure consistency guarantees allowing to reason about their behaviors in an abstract and simple way. Ideally, programmers of applications using such systems would like to have strong consistency guarantees, i.e., all updates occurring anywhere in the system are seen immediately and executed in the same order by all sites. Moreover, application

Key words and phrases: Distributed Databases, Causal Consistency, Model Checking.

This work is supported in part by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 678177).

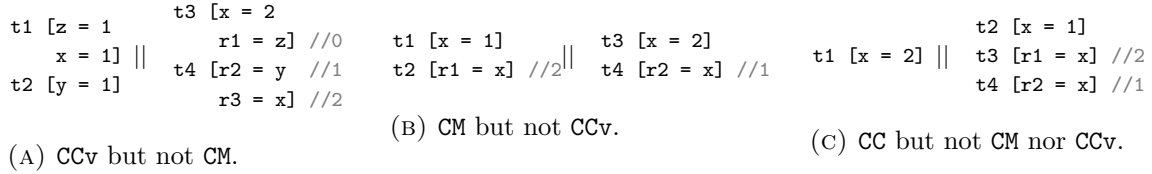


FIGURE 1. Program computations showing the relationship between CC, CCv and CM. Transactions are delimited using brackets and the transactions issued on the same site are aligned vertically. The values read in a transaction are given in comments.

programmers also need an abstract mechanism such as transactions, ensuring that blocks of actions (writes and reads) of a site can be considered as executing atomically without interferences from actions of other sites. For transactional programs, the consistency model offering strong consistency is *serializability* [37], i.e., every computation of a program is equivalent to another one where transactions are executed serially one after another without interference. In the non-transactional case this model corresponds to *sequential consistency* (SC) [31]. However, while serializability and SC are easier to apprehend by application programmers, their enforcement (by storage systems implementors) requires the use of global synchronization between all sites, which is hard to achieve while ensuring availability and acceptable performances [23, 24]. For this reason, modern storage systems ensure weaker consistency guarantees. In this paper, we are interested in studying *causal consistency* [30].

Causal consistency is a fundamental consistency model implemented in several production databases, e.g., AntidoteDB, CockroachDB, and MongoDB, and extensively studied in the literature [7, 22, 33, 34, 39]. Basically, when defined at the level of actions, it guarantees that every two causally related actions, say a_1 is causally before (i.e., it has an influence on) a_2 , are executed in that same order, i.e., a_1 before a_2 , by all sites. The sets of updates visible to different sites may differ and read actions may return values that cannot be obtained in SC executions. The definition of causal consistency can be lifted to the level of transactions, assuming that transactions are visible to a site in their entirety (i.e., all their updates are visible at the same time), and they are executed by a site in isolation without interference from other transactions. In comparison to serializability, causal consistency allows that conflicting transactions, i.e., which read or write to a common location, be executed in different orders by different sites as long as they are not causally related. Actually, we consider three variations of causal consistency introduced in the literature, weak causal consistency (CC) [38, 13], causal memory (CM) [3, 38], and causal convergence (CCv) [17].

The weakest variation of causal consistency, namely CC, allows speculative executions and roll-backs of transactions which are not causally related (concurrent). For instance, the computation in Fig. 1c is only feasible under CC: the site on the right applies t_2 after t_1 before executing t_3 and roll-backs t_2 before executing t_4 . CCv and CM offer more guarantees. CCv enforces a total *arbitration* order between all transactions which defines the order in which delivered concurrent transactions are executed by *every* site. This guarantees that all sites reach the same state when all transactions are delivered. CM ensures that *all* values read by a site can be explained by an interleaving of transactions consistent with the causal order, enforcing thus PRAM consistency [32] on top of CC. Contrary to CCv, CM allows that two sites diverge on the ordering of concurrent transactions, but both models do not

allow roll-backs of concurrent transactions. Thus, CCv and CM are incomparable in terms of computations they admit. The computation in Fig. 1a is not admitted by CM because there is no interleaving of those transactions that explains the values read by the site on the right: reading 0 from z implies that the transactions on the left must be applied after t_3 while reading 1 from y implies that both t_1 and t_2 are applied before t_4 which contradicts reading 2 from x . However, this computation is possible under CCv because t_1 can be delivered to the right after executing t_3 but arbitrated before t_3 , which implies that the write to x in t_1 will be lost. The CM computation in Fig. 1b is not possible under CCv because there is no arbitration order that could explain both reads from x .

As a first contribution of our paper, we show that the three causal consistency models coincide for transactional programs containing no *write-write races*, i.e., concurrent transactions writing on a common variable. We also show that if a transactional program has a write-write race under one of these models, then it must have a write-write race under any of the other two models. This property is rather counter-intuitive since CC is strictly weaker than both CCv and CM , and CCv and CM are incomparable (in terms of admitted behaviors). Notice that each of the computations in Figures 1a, 1b, and 1c contains a write-write race which explains why none of these computations is possible under all three models.

Then, we investigate the problem of checking *robustness* of application programs against causal consistency relaxations: Given a program P and a causal consistency variation X , we say that P is robust against X if the set of computations of P when running under X is the same as its set of computations when running under *serializability*. This means that it is possible to reason about the behaviors of P assuming the simpler serializability model and no additional synchronization is required when P runs under X such that it maintains all the properties satisfied under serializability. Checking robustness is not trivial, it can be seen as a form of checking program equivalence. However, the equivalence to check is between two versions of the same program, obtained using two different semantics, one more permissive than the other one. The goal is to check that this permissiveness has actually no effect on the particular program under consideration. The difficulty in checking robustness is to apprehend the extra behaviors due to the reorderings introduced by the relaxed consistency model w.r.t. serializability. This requires a priori reasoning about complex order constraints between operations in arbitrarily long computations, which may need maintaining unbounded ordered structures, and make the problem of checking robustness hard or even undecidable.

We show that verifying robustness of transactional programs against causal consistency can be reduced in polynomial time to the reachability problem in concurrent programs over SC. This allows to reason about distributed applications running on causally consistent storage systems using the existing verification technology and it implies that the robustness problem is decidable for finite-state programs; the problem is PSPACE-complete when the number of sites is fixed, and EXPSpace-complete otherwise. This is the first result on the decidability and complexity of verifying robustness against causal consistency. In fact, the problem of verifying robustness has been considered in the literature for several consistency models of distributed systems, including causal consistency [11, 15, 16, 19, 35]. These works provide (over- or under-)approximate analyses for checking robustness, but none of them provides precise (sound and complete) algorithmic verification methods for solving this problem, nor addresses its decidability and complexity.

The approach we adopt for tackling this verification problem is based on a precise characterization of the set of robustness violations, i.e., executions that are causally consistent

$$\begin{array}{ll}
\langle prog \rangle ::= \text{program } \langle process \rangle^* & \langle linst \rangle ::= \langle label \rangle : \langle inst \rangle ; \text{goto } \langle label \rangle ; \\
\langle process \rangle ::= \text{process } \langle pid \rangle \text{ regs } \langle reg \rangle^* & \langle inst \rangle ::= \langle reg \rangle := \langle var \rangle \\
& \langle ltxn \rangle^* & | \langle var \rangle := \langle reg\text{-expr} \rangle \\
\langle ltxn \rangle ::= \langle binst \rangle \langle linst \rangle^* \langle einst \rangle & | \text{assume } \langle bepr \rangle \\
\langle binst \rangle ::= \langle label \rangle : \text{begin} ; \text{goto } \langle label \rangle ; & \\
\langle einst \rangle ::= \langle label \rangle : \text{end} ; \text{goto } \langle label \rangle ; &
\end{array}$$

FIGURE 2. Program syntax. a^* indicates zero or more occurrences of a . $\langle pid \rangle$, $\langle reg \rangle$, $\langle label \rangle$, and $\langle var \rangle$ represent a process identifier, a register, a label, and a shared variable respectively. $\langle reg\text{-expr} \rangle$ is an expression over registers while $\langle bepr \rangle$ is a Boolean expression over registers.

but not serializable. For both CCv and CM , we show that it is sufficient to search for a special type of robustness violations, that can be simulated by serial (SC) computations of an instrumentation of the original program. These computations maintain the information needed to recognize the pattern of a violation that would have occurred in the original program under a causally consistent semantics (executing the same set of operations). A surprising consequence of these results is that a program is robust against CM iff it is robust against CC , and robustness against CM implies robustness against CCv . This shows that the causal consistency variations we investigate can be incomparable in terms of the admitted behaviors, but comparable in terms of the robust applications they support.

2. CAUSAL CONSISTENCY

2.1. Program syntax. We consider a simple programming language where a program is parallel composition of *processes* distinguished using a set of identifiers \mathbb{P} . Our simple programming language syntax is given in Fig. 2. Each process is a sequence of *transactions* and each transaction is a sequence of *labeled instructions*. Each transaction starts with a **begin** instruction and finishes with an **end** instruction. Each other instruction is either an assignment to a process-local *register* from a set \mathbb{R} or to a *shared variable* from a set \mathbb{V} , or an **assume** statement. The assignments use values from a data domain \mathbb{D} . An assignment to a register $\langle reg \rangle := \langle var \rangle$ is called a *read* of $\langle var \rangle$ and an assignment to a shared variable $\langle var \rangle := \langle reg\text{-expr} \rangle$ is called a *write* to $\langle var \rangle$ ($\langle reg\text{-expr} \rangle$ is an expression over registers). The statement **assume** $\langle bepr \rangle$ blocks the process if the Boolean expression $\langle bepr \rangle$ over registers is false. Each instruction is followed by a **goto** statement which defines the evolution of the program counter. Multiple instructions can be associated with the same label which allows us to write non-deterministic programs and multiple **goto** statements can direct the control to the same label which allows us to mimic imperative constructs like loops and conditionals. We assume that the control cannot pass from one transaction to another without going as expected through **begin** and **end** instructions.

2.2. Program Semantics Under Causal Memory. Informally, the semantics of a program under causal memory is defined as follows. The shared variables are replicated across each process, each process maintaining its own local valuation of these variables. During the execution of a transaction in a process, the shared-variable writes are stored in a *transaction log* which is visible only to the process executing the transaction and which is broadcasted

to all the processes at the end of the transaction¹. To read a shared variable x , a process p first accesses its transaction log and takes the last written value on x , if any, and then its own valuation of the shared variables, if x was not written during the current transaction. Transaction logs are delivered to every process in an order consistent with the *causal delivery* relation between transactions, i.e., the transitive closure of the union of the *program order* (the order in which transactions are executed by a process), and the *delivered-before* relation (a transaction t_1 is delivered-before a transaction t_2 iff the log of t_1 has been delivered at the process executing t_2 before t_2 starts). By an abuse of terminology, we call this property *causal delivery*. Once a transaction log is delivered, it is immediately applied on the shared-variable valuation of the receiving process. Also, no transaction log can be delivered to a process p while p is executing another transaction, we call this property *transaction isolation*.

Formally, a program configuration is a triple $\mathbf{gs} = (\mathbf{ls}, \mathbf{msgs})$ where $\mathbf{ls} : \mathbb{P} \rightarrow \mathbb{S}$ associates a local state in \mathbb{S} to each process in \mathbb{P} , and \mathbf{msgs} is a set of messages in transit. A local state is a tuple $\langle \mathbf{pc}, \mathbf{store}, \mathbf{rval}, \mathbf{log} \rangle$ where $\mathbf{pc} \in \mathbb{Lab}$ is the program counter, i.e., the label of the next instruction to be executed, $\mathbf{store} : \mathbb{V} \rightarrow \mathbb{D}$ is the local valuation of the shared variables, $\mathbf{rval} : \mathbb{R} \rightarrow \mathbb{D}$ is the valuation of the local registers, and $\mathbf{log} \in (\mathbb{V} \times \mathbb{D})^*$ is the transaction log, i.e., a list of variable-value pairs. For a local state s , we use $s.\mathbf{pc}$ to denote the program counter component of s , and similarly for all the other components of s . A message $m = \langle t, \mathbf{log} \rangle$ is a transaction identifier t from a set \mathbb{T} together with a transaction log $\mathbf{log} \in (\mathbb{V} \times \mathbb{D})^*$. We let \mathbb{M} denote the set of messages.

Then, the semantics of a program \mathcal{P} under causal memory is defined using a labeled transition system (LTS) $[\mathcal{P}]_{\text{CM}} = (\mathbb{C}, \mathbb{Ev}, \mathbf{gs}_0, \rightarrow)$ where \mathbb{C} is the set of program configurations, \mathbb{Ev} is a set of transition labels called *events*, \mathbf{gs}_0 is the initial configuration, and $\rightarrow \subseteq \mathbb{C} \times \mathbb{Ev} \times \mathbb{C}$ is the transition relation. As it will be explained later in this section, the executions of \mathcal{P} under causal memory are a subset of those generated by $[\mathcal{P}]_{\text{CM}}$. The set of events is defined by:

$$\mathbb{Ev} = \{ \mathbf{begin}(p, t), \mathbf{ld}(p, t, x, v), \mathbf{isu}(p, t, x, v), \mathbf{del}(p, t), \mathbf{end}(p, t) : p \in \mathbb{P}, t \in \mathbb{T}, x \in \mathbb{V}, v \in \mathbb{D} \}$$

where \mathbf{begin} and \mathbf{end} label transitions corresponding to the start, resp., the end of a transaction, \mathbf{isu} and \mathbf{ld} label transitions corresponding to writing, resp., reading, a shared variable during some transaction, and \mathbf{del} labels transitions corresponding to applying a transaction log to the local state of the process issuing the transaction or to the state of another process that received the log. An event \mathbf{isu} is called an *issue* while an event \mathbf{del} is called a *store*.

The transition relation \rightarrow is partially defined in Fig. 3 (we will present additional constraints later in this section). The events labeling a transition are written on top of \rightarrow . A \mathbf{begin} transition will just reset the transaction log while an \mathbf{end} transition will add the transaction log together with the transaction identifier to the set \mathbf{msgs} of messages in transit. An \mathbf{ld} transition will read the value of a shared-variable looking first at the transaction log \mathbf{log} and then, at the shared-variable valuation \mathbf{store} , while an \mathbf{isu} transition will add a new write to the transaction log. Finally, a \mathbf{del} transition represents the delivery of a transaction log that was in transit which is applied immediately on the shared-variable valuation \mathbf{store} .

We say that an execution ρ satisfies *transaction isolation* if no transaction log is delivered to a process p while p is executing a transaction, i.e., if an event $ev = \mathbf{del}(p, t)$ occurs in ρ before an event $ev' = \mathbf{end}(p, t')$ with $t' \neq t$, then ρ contains an event $ev'' = \mathbf{begin}(p, t')$ between

¹For simplicity, we assume that every transaction commits. The effects of aborted transactions shouldn't be visible to any process.

$$\begin{array}{c}
\frac{\text{begin} \in \text{inst}(\text{ls}(p).\text{pc}) \quad s = \text{ls}(p)[\text{log} \mapsto \epsilon, \text{pc} \mapsto \text{next}(\text{pc})]}{(\text{ls}, \text{msgs}) \xrightarrow{\text{begin}(p, t)} (\text{ls}[p \mapsto s], \text{msgs})} \\
\frac{r := x \in \text{inst}(\text{ls}(p).\text{pc}) \quad \text{eval}(\text{ls}(p), x) = v \quad \text{rval} = \text{ls}(p).\text{rval}[r \mapsto v] \quad s = \text{ls}(p)[\text{rval} \mapsto \text{rval}, \text{pc} \mapsto \text{next}(\text{pc})]}{(\text{ls}, \text{msgs}) \xrightarrow{\text{id}(p, t, x, v)} (\text{ls}[p \mapsto s], \text{msgs})} \\
\frac{x := v \in \text{inst}(\text{ls}(p).\text{pc}) \quad \text{log} = (\text{ls}(p).\text{log}) \cdot (x, v) \quad s = \text{ls}(p)[\text{log} \mapsto \text{log}, \text{pc} \mapsto \text{next}(\text{pc})]}{(\text{ls}, \text{msgs}) \xrightarrow{\text{isu}(p, t, x, v)} (\text{ls}[p \mapsto s], \text{msgs})} \\
\frac{\text{end} \in \text{inst}(\text{ls}(p).\text{pc}) \quad s = \text{ls}(p)[\text{pc} \mapsto \text{next}(\text{pc})]}{(\text{ls}, \text{msgs}) \xrightarrow{\text{end}(p, t)} (\text{ls}[p \mapsto s], \text{msgs} \cup \{(t, \text{ls}(p).\text{log})\})} \\
\frac{\langle t, \text{log} \rangle \in \text{msgs} \quad \text{store} = \text{ls}(p).\text{store}[x \mapsto \text{last}(\text{log}, x) : x \in \mathbb{V}, \text{last}(\text{log}, x) \neq \perp] \quad s = \text{ls}(p)[\text{store} \mapsto \text{store}]}{(\text{ls}, \text{msgs}) \xrightarrow{\text{del}(p, t)} (\text{ls}[p \mapsto s], \text{msgs})}
\end{array}$$

FIGURE 3. The set of transition rules defining the causal memory semantics. We assume that all the events which come from the same transaction use a unique transaction identifier t . For a function f , we use $f[a \mapsto b]$ to denote a function g such that $g(c) = f(c)$ for all $c \neq a$ and $g(a) = b$. The function inst returns the set of instructions labeled by some given label while next gives the next instruction to execute. We use \cdot to denote sequence concatenation. The function $\text{eval}(\text{ls}(p), x)$ returns the value of x in the local state $\text{ls}(p)$: (1) if $\text{ls}(p).\text{log}$ contains a pair (x, v) , for some v , then $\text{eval}(\text{ls}(p), x)$ returns the value of the last such pair in $\text{ls}(p).\text{log}$, and (2) $\text{eval}(\text{ls}(p), x)$ returns $\text{ls}(p).\text{store}(x)$, otherwise. Also, $\text{last}(\text{log}, x)$ returns the value v in the last pair (x, v) in log , and \perp , if such a pair does not exist.

ev and ev' . For an execution ρ satisfying transaction isolation, we assume w.l.o.g. that transactions executed by different processes do not interleave, i.e., if an event ev associated to a transaction t (an event of the process executing t or the delivery of the transaction log of t) occurs in ρ before $ev' = \text{end}(p', t')$, then ρ contains an event $ev'' = \text{begin}(p', t')$ between ev and ev' . Formally, we say that an execution ρ satisfies *causal delivery* if the following hold:

- for any event $\text{begin}(p, t)$, and for any process p' , ρ contains at most one event $\text{del}(p', t)$,
- for any two events $\text{begin}(p, t)$ and $\text{begin}(p, t')$, if $\text{begin}(p, t)$ occurs in ρ before $\text{begin}(p, t')$, then the event $\text{del}(p, t)$ occurs before $\text{begin}(p, t')$ in ρ . This ensures that when p issues t it must store the writes of t in its local state before issuing another transaction t' ;
- for any events $ev_1 \in \{\text{del}(p, t_1), \text{end}(p, t_1)\}$, $ev_2 = \text{begin}(p, t_2)$, and $ev'_2 = \text{del}(p', t_2)$ with $p \neq p'$, if ev_1 occurs in ρ before ev_2 , then there exists $ev'_1 = \text{del}(p', t_1)$ such that ev'_1 occurs before ev'_2 in ρ .

An execution ρ satisfies *causal memory* if it satisfies transaction isolation and causal delivery. The set of executions of \mathcal{P} under causal memory, denoted by $\mathbb{E}_{\text{CM}}(\mathcal{P})$, is the set of executions of $[\mathcal{P}]_{\text{CM}}$ satisfying causal memory.

Fig. 4a shows an execution under CM. This execution satisfies transaction isolation since no transaction is delivered while another transaction is executing.

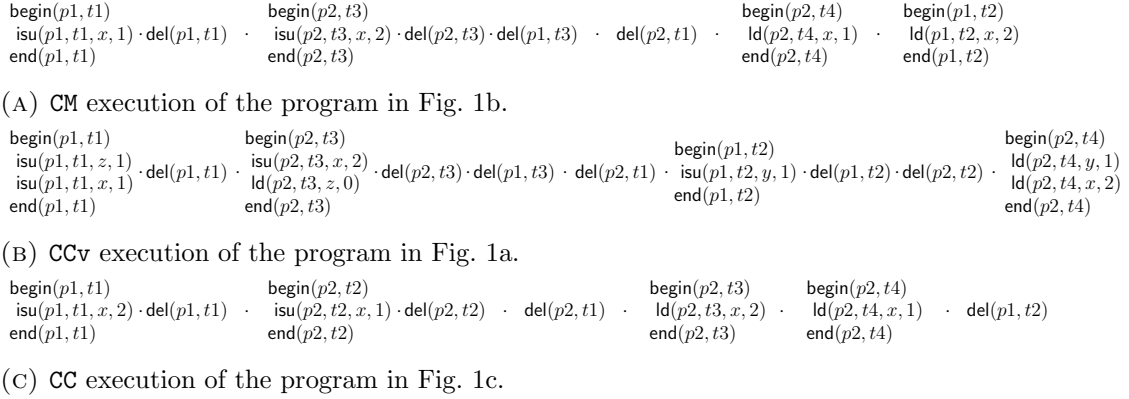


FIGURE 4. For readability, the sub-sequences of events delimited by **begin** and **end** are aligned vertically, the execution-flow advancing from left to right and top to bottom.

2.3. Program Semantics Under Causal Convergence. Compared to causal memory, causal convergence ensures eventual consistency of process-local copies of the shared variables. Each transaction log is associated with a timestamp and a process applies a write on some variable x from a transaction log only if it has a timestamp larger than the timestamps of all the transaction logs it has already applied and that wrote the same variable x . For simplicity, we assume that the transaction identifiers play the role of timestamps, which are totally ordered according to some relation $<$. CCv satisfies both *causal delivery* and *transaction isolation* as well. Assuming that transactions are constituted of either a read alone or a write alone, CCv is equivalent to Strong Release-Acquire (SRA), a strengthening of the standard Release-Acquire of the C11 memory model [28]².

Formally, we define a variation of the LTS $[\mathcal{P}]_{\text{CM}}$, denoted by $[\mathcal{P}]_{\text{CCv}}$, where essentially, the transition identifiers play the role of timestamps and are ordered by a total order $<$, each process-local state contains an additional component **tstamp** storing the largest timestamp the process has seen for each variable, and a write on a variable x from a transaction log is applied on the local valuation **store** only if it has a timestamp larger than **tstamp**(x). Also, a **begin**(p, t) transition will choose a transaction identifier t greater than those in the image of the **tstamp** component of p 's local state. The transition rules of $[\mathcal{P}]_{\text{CCv}}$ that change w.r.t. those of $[\mathcal{P}]_{\text{CM}}$ are given in Fig. 5.

The set of executions of \mathcal{P} under causal convergence, denoted by $\mathbb{E}x_{\text{CCv}}(\mathcal{P})$, is the set of executions of $[\mathcal{P}]_{\text{CCv}}$ satisfying transaction isolation, causal delivery, and the fact that every process p generates monotonically increasing transaction identifiers.

The execution in Fig. 4a is not possible under causal convergence since $t4$ and $t2$ read 2 and 1 from x , respectively. This is possible only if $t1$ and $t3$ write x at $p2$ and $p1$, respectively, which contradicts the definition of **del** transition given in Fig. 5 where we cannot have both $t1 < t3$ and $t3 < t1$ at the same time. Fig. 4b shows an execution under CCv (we assume $t_1 < t_2 < t_3 < t_4$). Notice that **del**($p2, t1$) did not result in an update of x because the timestamp t_1 is smaller than the timestamp of the last transaction that wrote x at p_2 ,

²This equivalence excludes the atomic read-modify-write (also know as compare-and-swap) operation which is not provided by CCv.

$$\begin{array}{c}
\frac{\Phi_1 \quad \text{img}(\text{ls}(p).\text{tstamp}) < t}{(\text{ls}, \text{lk}, \text{msgs}) \xrightarrow{\text{begin}(p, t)} (\text{ls}[p \mapsto s], \text{lk}, \text{msgs})} \\
\frac{\langle t, \text{log} \rangle \in \text{msgs} \quad \text{store} = \text{ls}(p).\text{store}[x \mapsto \text{last}(\text{log}, x) : x \in \mathbb{V}, \text{last}(\text{log}, x) \neq \perp, \text{tstamp}(x) < t] \\
\text{tstamp} = \text{ls}(p).\text{tstamp}[x \mapsto t : x \in \mathbb{V}, \text{last}(\text{log}, x) \neq \perp, \text{tstamp}(x) < t] \quad s = \text{ls}(p)[\text{store} \mapsto \text{store}, \text{tstamp} \mapsto \text{tstamp}]}{(\text{ls}, \text{lk}, \text{msgs}) \xrightarrow{\text{del}(p, t)} (\text{ls}[p \mapsto s], \text{lk}, \text{msgs})}
\end{array}$$

FIGURE 5. Transition rules for defining causal convergence. Φ_1 is the hypothesis of the $\text{begin}(p, t)$ transition rule in Fig. 3, and img denotes the image of a function.

namely t_3 , a behavior that is not possible under CM . The two processes converge and store the same shared variable copy at the end of the execution.

2.4. Program Semantics Under Weak Causal Consistency. Compared to the previous semantics, CC allows that reads of the same process observe concurrent writes as executing in different orders. Each process maintains a *set* of values for each shared variable, and a read returns any one of these values non-deterministically. Transaction logs are associated with *vector clocks* [30] which represent the causal delivery relation, i.e., a transaction t_1 is before t_2 in causal-delivery iff the vector clock of t_1 is smaller than the vector clock of t_2 . We assume that transactions identifiers play the role of vector clocks, which are partially ordered according to some relation $<$. In applying the log of a transaction t on the local state of the receiving process p , the final *set* of values for each shared variable in p will be constituted of the value in the log of t and the values that were written by *concurrent* transactions (not related by causal delivery to t). CC satisfies both *causal delivery* and *transaction isolation*.

Formally, in CC semantics, the local valuation of the shared variables $\text{store} : \mathbb{V} \rightarrow (\mathbb{D} \times \mathbb{T})^*$ is a map that accepts a shared variable and returns a set of pairs. The pairs are constituted of values that were written concurrently and identifiers of the transactions that wrote those values. When applying a transaction log on the local valuation store, we keep the values that were written by transactions that are concurrent with the current transaction. Additionally, in the CC semantics, the local state of a process has an additional component $\text{snapshot} : \mathbb{V} \rightarrow (\mathbb{D} \times \mathbb{T})$ that maps each shared variable to a single pair. snapshot is obtained by taking a “consistent” snapshot from store when a new transaction starts. Such a snapshot corresponds to a linearization of the transactions that were delivered to the process, which is consistent with the vector clock order. The snapshot associates to each variable the last value written in this linearization. When a process does a read from a shared variable x , it looks first at the transaction log log and then, at the variable valuation snapshot . In Fig. 6, we provide the transition rules of $[\mathcal{P}]_{\text{CC}}$ that change w.r.t. those of $[\mathcal{P}]_{\text{CCv}}$ and $[\mathcal{P}]_{\text{CM}}$.

The set of executions of \mathcal{P} under weak causal consistency model, denoted by $\mathbb{E}_{\text{XCC}}(\mathcal{P})$, is the set of executions of $[\mathcal{P}]_{\text{CC}}$ satisfying transaction isolation and causal delivery. We denote by $\text{Tr}(\mathcal{P})_{\text{CC}}$ the set of traces of executions of a program \mathcal{P} under weak causal consistency.

Fig. 4c shows an execution under CC , which is not possible under CCv and CM because t_3 and t_4 read 2 and 1, respectively. Since the transactions t_1 and t_2 are concurrent, p_2 stores both values 2 and 1 written by these transactions. A read of x can return any of these two values.

$$\begin{array}{c}
\text{begin} \in \text{inst}(\text{ls}(p).\text{pc}) \quad \text{img}(\text{ls}(p).\text{tstamp}) < t \\
\hline
s = \text{ls}(p)[\text{log} \mapsto \epsilon, \text{snapshot} \mapsto \text{buildSnapshot}(\text{store}), \text{pc} \mapsto \text{next}(\text{pc})] \\
(\text{ls}, \text{msgs}) \xrightarrow{\text{begin}(p, t)} (\text{ls}[p \mapsto s], \text{msgs}) \\
r := x \in \text{inst}(\text{ls}(p).\text{pc}) \quad \text{cceval}(\text{ls}(p), x) = (v, t') \quad \text{rval} = \text{ls}(p).\text{rval}[r \mapsto v] \\
\hline
s = \text{ls}(p)[\text{rval} \mapsto \text{rval}, \text{pc} \mapsto \text{next}(\text{pc})] \\
(\text{ls}, \text{msgs}) \xrightarrow{\text{ld}(p, t, x, v)} (\text{ls}[p \mapsto s], \text{msgs}) \\
\text{end} \in \text{inst}(\text{ls}(p).\text{pc}) \quad s = \text{ls}(p)[\text{snapshot} \mapsto \epsilon, \text{pc} \mapsto \text{next}(\text{pc})] \\
\hline
(\text{ls}, \text{msgs}) \xrightarrow{\text{end}(p, t)} (\text{ls}[p \mapsto s], \text{msgs} \cup \{(t, \text{ls}(p).\text{log})\}) \\
\langle t, \text{log} \rangle \in \text{msgs} \quad \text{store} = \text{ls}(p).\text{store}[x \mapsto \text{update}(\text{ls}(p), x, t, \text{last}(\text{log}, x)) : x \in \mathbb{V}] \\
\hline
s = \text{ls}(p)[\text{store} \mapsto \text{store}, \text{pc} \mapsto \text{next}(\text{pc})] \\
(\text{ls}, \text{msgs}) \xrightarrow{\text{del}(p, t)} (\text{ls}[p \mapsto s], \text{msgs})
\end{array}$$

FIGURE 6. Transition rules for defining weak causal consistency semantics: $\text{buildSnapshot}(\text{store})$ returns a consistent snapshot of store . $\text{cceval}(\text{ls}(p), x)$ returns the pair $(\text{last}(\text{log}, x), t)$ if $\text{last}(\text{log}, x) \neq \perp$, and returns the pair (v, t') in $\text{ls}(p).\text{snapshot}(x)$, otherwise. $\text{update}(\text{ls}(p), x, t, \text{last}(\text{log}, x))$ returns the result of appending the pair $(\text{last}(\text{log}, x), t)$ to the set $\text{ls}(p).\text{store}(x)$ after removing all pairs that contain values overwritten by t .

2.5. Execution Summary. Let ρ be an execution under $X \in \{\text{CCv}, \text{CM}, \text{CC}\}$, a sequence τ of events $\text{isu}(p, t)$ and $\text{del}(p, t)$ with $p \in \mathbb{P}$ and $t \in \mathbb{T}$ is called a *summary of ρ* if it is obtained from ρ by substituting every sub-sequence of transitions in ρ delimited by a **begin** and an **end** transition, with a single “macro-event” $\text{isu}(p, t)$. For example, $\text{isu}(p1, t1) \cdot \text{isu}(p2, t3) \cdot \text{del}(p1, t3) \cdot \text{del}(p2, t1) \cdot \text{isu}(p2, t4) \cdot \text{isu}(p1, t2)$ is a summary of the execution in Fig. 4a.

We say that a transaction t in ρ performs an *external read* of a variable x if ρ contains an event $\text{ld}(p, t, x, v)$ which is not preceded by a write on x of t , i.e., an event $\text{isu}(p, t, x, v)$. Under **CM** and **CC**, a transaction t *writes* a variable x if ρ contains an event $\text{isu}(p, t, x, v)$, for some v . In Fig. 4a, both $t2$ and $t4$ perform external reads and $t2$ writes to y . A transaction t executed by a process p *writes x at process p'* if t writes x and ρ contains an event $\text{del}(p', t)$ (e.g., in Fig. 4a, $t1$ writes x at $p2$). Under **CCv**, we say that a transaction t executed by a process p *writes x at process p'* if t writes x and ρ contains an event $\text{del}(p', t)$ which is not preceded by an event $\text{del}(p', t')$ with $t < t'$ and t' writing x (if it would be preceded by such an event then the write to x of t will be discarded). For example, in Fig. 4b, $t1$ does *not* write x at $p2$.

2.6. Trace. We define an abstract representation of executions that satisfy transaction isolation³, called *trace*. Essentially, a trace contains the summary of an execution (it forgets the order in which shared-variables are accessed inside a transaction) and several happens-before relations between events in its summary which record control-flow dependencies, the order between transactions issued in the same process, and data-flow dependencies, e.g. which transaction wrote the value read by another transaction.

More precisely, the *trace* of an execution ρ is a tuple $\text{tr}(\rho) = (\tau, \text{PO}, \text{WR}, \text{WW}, \text{RW}, \text{STO})$ where τ is the summary of ρ , **PO** is the *program order*, which relates any two issue events

³We refer collectively to executions in $[\mathcal{P}]_X$ with $X \in \{\text{CCv}, \text{CM}, \text{CC}\}$.

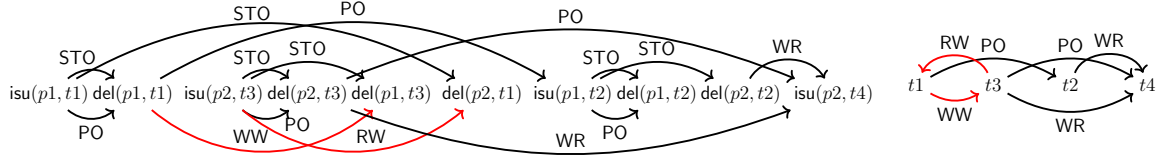


FIGURE 7. The trace of the execution in Fig. 4b and its transactional happens-before.

$isu(p, t)$ and $isu(p, t')$ that occur in this order in τ , WR is the *write-read* relation (also called *read-from*), which relates events of two transactions t and t' such that t writes a value that t' reads, WW is the *write-write* order (also called *store-order*), which relates events of two transactions that write to the same variable, RW is the *read-write* relation (also called *conflict*), which relates events of two transactions t and t' such that t reads a value overwritten by t' , and STO is the *same-transaction* relation, which relates events of the same transaction.

Definition 2.1 (Trace). Formally, the *trace* of an execution ρ satisfying transaction isolation is $\text{tr}(\rho) = (\tau, PO, WR, WW, RW, STO)$ where τ is a summary of ρ , and

PO: relates the issue and store events $isu(p, t)$ and $del(p, t)$ of t and subsequently, the event $del(p, t)$ with any issue event $isu(p, t')$ that occurs after it in τ .

WR: relates any store and issue events $ev_1 = del(p, t)$ and $ev_2 = isu(p, t')$ that occur in this order in τ such that t' performs an external read of x , and ev_1 is the last event in τ before ev_2 such that t writes x at p . To make the shared variable x explicit, we may use $WR(x)$ to name the relation between ev_1 and ev_2 .

WW: relates events of two transactions that write to the same variable. More precisely, WW relates any two store events $ev_1 = del(p, t_1)$ and $ev_2 = del(p, t_2)$ that occur in this order in τ provided that t_1 and t_2 both write the same variable x , and if ρ is an execution under causal convergence, then t_1 and t_2 writes x at p , and $t_1 < t_2$. To make the shared variable x explicit, we may use $WW(x)$ to name the relation between ev_1 and ev_2 .

RW: relates events of two distinct transactions t and t' such that t reads a value that is overwritten by t' . Formally, $RW(x) = WR^{-1}(x); WW(x)$ (we use $;$ to denote the standard composition of relations) and $RW = \bigcup_{x \in \mathbb{V}} RW(x)$. If a transaction t reads the initial value of x then $RW(x)$ relates $isu(p, t)$ with every event $del(p', t')$ with $p' \in \mathbb{P}$ of any other transaction t' that writes to x at p' .

STO: relates issue events with store events of the same transaction. More precisely, STO relates every event $isu(p, t)$ with every event $del(p', t)$ with $p' \in \mathbb{P}$.

The following result states an important property of the store order relation WW that is enforced by the CCv semantics. It holds because the writes in different transactions are applied by different processes in the same order given by their timestamps, when visible (delivered) to those processes.

Lemma 2.2. *Let $\tau \in \text{Tr}_{\text{CCv}}(\mathcal{P})$ be a trace. If $(del(p_0, t_0), del(p_0, t_1)) \in WW(x)$, then for every other process p , $(del(p, t_1), del(p, t_0)) \notin WW(x)$.*

We define the *happens-before* relation HB as the transitive closure of the union of all the relations in the trace, i.e., $HB = (PO \cup WR \cup WW \cup RW \cup STO)^+$. Since we reason

about only one trace at a time, we may say that a trace is simply a summary τ , keeping the relations implicit. The trace of the CCv execution in Fig. 4b is shown on the left of Fig. 7. $\text{Tr}(\mathcal{P})_{\mathbb{X}}$ denotes the set of traces of executions of a program \mathcal{P} under $\mathbb{X} \in \{\text{CCv}, \text{CM}, \text{CC}\}$.

For readability, we write $ev_1 \rightarrow_{\text{HB}} ev_2$ instead of $(ev_1, ev_2) \in \text{HB}$ and ev_1 and ev_2 can be either $\text{isu}(p, t)$ or $\text{del}(p, t)$. We use the notation $ev_1 \rightarrow_{\text{HB}^1} ev_2$ (resp., $(ev_1, ev_2) \in \text{HB}^1$) to denote $(ev_1, ev_2) \in (\text{PO} \cup \text{WW} \cup \text{WR} \cup \text{STO} \cup \text{RW})$.

The *causal order* CO of a trace $tr = (\tau, \text{PO}, \text{WR}, \text{WW}, \text{RW}, \text{STO})$ is the transitive closure of the union of the program order, write-read relation, and the same-transaction relation, i.e., $\text{CO} = (\text{PO} \cup \text{WR} \cup \text{STO})^+$. For readability, we write $ev_1 \rightarrow_{\text{CO}} ev_2$ instead of $(ev_1, ev_2) \in \text{CO}$.

Let t_1 and t_2 be two transactions issued in a trace tr that originate from two different processes p_1 and p_2 , respectively. If $(\text{isu}(p_1, t_1), \text{isu}(p_2, t_2)) \notin \text{CO}$ and $(\text{isu}(p_2, t_2), \text{isu}(p_1, t_1)) \notin \text{CO}$, then t_1 and t_2 are called *concurrent* transactions.

The happens-before relation between events is extended to transactions as follows: a transaction t_1 *happens-before* another transaction $t_2 \neq t_1$ if the trace tr contains an event of transaction t_1 which happens-before an event of t_2 . The happens-before relation between transactions is denoted by HB_t and called *transactional happens-before* (an example is given on the right of Fig. 7).

Remark 2.3. The operational models of causal consistency we described are equivalent to the axiomatic models defined in [13]. These axiomatic models are defined as a set of constraints on abstractions of executions, called *histories*, that consist of a set of read and write operations along with a program order, denoted by PO' , and a read-from relation, denoted by WR' : PO' relates operations in the same process and WR' associates every read operation to the write operation which wrote the read value. For instance, the axiomatic model of CC requires that the union of PO' and WR' (denoted CO') is acyclic⁴, and its composition with a variation of the conflict relation, denoted by RW' , $((a, b) \in \text{RW}'^5 \text{ iff } \exists c. (c, b) \in \text{CO}' \wedge (c, a) \in \text{WR}'$) is irreflexive⁶. These models can be extended easily to histories that contain transactions instead of operations by adapting the above relations. Note that every program trace (cf. Definition 2.1) can be “projected” to a history where issue and store events from the same transaction in the trace are mapped to a single transaction in the history. Also, the read-from and the program order between trace events are mapped to the WR' and PO' of the history.

To show equivalence between these models, it is sufficient to show that (1) every history corresponding to a trace in the operational model satisfies the constraints of the axiomatic model, and (2) every history that is valid under the axiomatic model is the “projection” of a trace of the operational model. For instance, for CC , it is easy to see that the relation $\text{CO}' = \text{PO}' \cup \text{WR}'$ in a history that is the projection of a trace $\tau \in \text{Tr}_{\text{CC}}(\mathcal{P})$ is acyclic because the causal order CO in τ is. Also, the proof that $\text{CO}'; \text{RW}'$ is irreflexive can be derived easily by contradiction (for instance, if $(a, b) \in \text{RW}'$ and $(b, a) \in \text{CO}'$, then there exists c such that $(c, b) \in \text{CO}'$ which means that by causal delivery, a can never read the value written by c).

3. WRITE-WRITE RACE FREEDOM

We say that an execution ρ has a *write-write race* on a shared variable x if there exist two concurrent transactions t_1 and t_2 that were issued in ρ and each transaction contains a write

⁴This constraint corresponds to the absence of the CyclicCO bad pattern in [13].

⁵ b is overwriting the value a is reading.

⁶This constraint corresponds to the absence of the WriteCORead bad pattern in [13].

to the variable x . We call ρ write-write race free if there is no variable x such that ρ has a write-write race on x . Also, we say a program \mathcal{P} is *write-write race free* under a consistency semantics $X \in \{\text{CCv}, \text{CM}, \text{CC}\}$ iff for every $\rho \in \mathbb{E}_{X\mathcal{X}}(\mathcal{P})$, ρ is write-write race free.

We show that if a given program has a write-write race under one of the three causal consistency models then it must have a write-write race under the remaining two. The intuition behind this is that the three models coincide for programs without write-write races. Indeed, without concurrent transactions that write to the same variable, every process local valuation of a shared variable will be a singleton set under CC and no process will ever discard a write when applying an incoming transaction log under CCv .

Theorem 3.1. *Given a program \mathcal{P} and two consistency semantics $X, Y \in \{\text{CCv}, \text{CM}, \text{CC}\}$, \mathcal{P} has a write-write race under X iff \mathcal{P} has a write-write race under Y .*

Proof. Since CC is weaker than both CCv and CM , it is sufficient to prove the following two cases: (1) if \mathcal{P} has a write-write race under CC , then \mathcal{P} has a write-write race under CCv and (2) if \mathcal{P} has a write-write race under CC , then \mathcal{P} has a write-write race under CM .

We prove the first case by induction on the number of transactions in \mathcal{P} . The second case can be proved in a similar way.

Base case: \mathcal{P} is constituted of two transactions t_1 and t_2 . Assume that \mathcal{P} has a write-write race under CC then the transactions t_1 and t_2 must originate from different processes. Thus, in any trace τ of \mathcal{P} under CCv where the transactions t_1 and t_2 are executed concurrently we will have a write-write race between these two transactions. Thus, \mathcal{P} has a write-write race under CCv .

Induction step: If $n > 2$ is the number of transactions in \mathcal{P} , we assume that for any program \mathcal{P}' with $n' < n$ transactions, if \mathcal{P}' has a write-write race under CC , then \mathcal{P}' has a write-write race under CCv . Assume that \mathcal{P} has a write-write race under CC . Let τ be a trace of \mathcal{P} under CC where we have a write-write race between two transactions t_1 and t_2 that were issued by processes p_1 and p_2 , respectively. Executing t_1 and t_2 concurrently while writing to a common variable is not possible under CCv only if the writes were enabled by some events that occurred before t_1 and t_2 under CC and are not possible under CCv . However, based on the semantic models of both CC and CCv , if all the transactions that write to common variables are causally related then such events cannot occur under CC but not CCv . Thus, we must have two other transactions t'_1 and t'_2 of \mathcal{P} that were executed concurrently in τ under CC and occurred before t_1 (or t_2 or both) which write to a common variable. Without loss of generality, let \mathcal{P}_1 be the program resulting from removing the transaction t_1 from \mathcal{P} . We know that \mathcal{P}_1 admits a trace τ_1 under CC where the transactions t'_1 and t'_2 are involved in a data race. Also, the size of \mathcal{P}_1 is $n - 1 < n$. Thus, from the induction hypothesis we get that \mathcal{P}_1 has a write-write race under CCv . Because adding a new transaction to \mathcal{P}_1 will not eliminate existing data races, \mathcal{P} has a write-write race under CCv as well. \square

The following result shows that indeed, the three causal consistency models coincide for programs which are write-write race free under any one of these three models.

Theorem 3.2. *Let \mathcal{P} be a program. Then, $\mathbb{E}_{X\text{CC}}(\mathcal{P}) = \mathbb{E}_{X\text{CCv}}(\mathcal{P}) = \mathbb{E}_{X\text{CM}}(\mathcal{P})$ iff \mathcal{P} has no write-write race under neither CC , CM , and CCv .*

Proof. Left-to-right direction: By Theorem 3.1, it is sufficient to prove that \mathcal{P} has no write-write race under CM . Suppose by contradiction that \mathcal{P} has a write-write race under CM . Then, there must exist a trace $\tau \in \text{Tr}_{\text{CC}}(\mathcal{P})$ such that we have two concurrent transactions t_1 and t_2 that are issued in τ and write to a variable x . Assume w.l.o.g that the issue event of

t_1 occurs before the issue event of t_2 in τ . Since t_1 and t_2 are concurrent in τ , the issue event of t_1 and the store events of t_2 are commutative, and the issue event of t_2 and the store events of t_1 are commutative. Then, $\tau' = \alpha \cdot \text{isu}(p_1, t_1) \cdot \text{del}(p_1, t_1) \cdot \beta \cdot \text{isu}(p_2, t_2) \cdot \text{del}(p_2, t_2) \cdot \text{del}(p_1, t_2) \cdot \text{del}(p_2, t_1)$ where α and β are sequences of events in τ that t_1 and t_2 causally depend on (since we are not interested in other events)⁷, is a trace of \mathcal{P} under **CM**. In τ' , both store events $\text{del}(p_2, t_1)$ and $\text{del}(p_1, t_2)$ do not discard any writes (guaranteed under **CM**). Therefore, $(\text{del}(p_1, t_1), \text{del}(p_1, t_2)) \in \text{WW}(x)$ and $(\text{del}(p_2, t_2), \text{del}(p_2, t_1)) \in \text{WW}(x)$ since both t_1 and t_2 write to x . However, it is impossible to obtain τ' under **CCv** as we cannot have $(\text{del}(p_2, t_2), \text{del}(p_2, t_1)) \in \text{WW}(x)$ if $(\text{del}(p_1, t_1), \text{del}(p_1, t_2)) \in \text{WW}(x)$ which leads to a contradiction (\mathcal{P} has different sets of traces under **CM** and **CCv**).

Right-to-left direction: It is sufficient to prove the following two cases: if τ has no write-write race under **CC** then $\tau \in \text{Tr}_{\text{CC}}$ implies $\tau \in \text{Tr}_{\text{CM}}$ and $\tau \in \text{Tr}_{\text{CCv}}$ ($\text{Tr}_{\text{CCv}}(\mathcal{P}) \subseteq \text{Tr}_{\text{CC}}(\mathcal{P})$ and $\text{Tr}_{\text{CM}}(\mathcal{P}) \subseteq \text{Tr}_{\text{CC}}(\mathcal{P})$ hold by definition).

Let $\tau \in \text{Tr}_{\text{CC}}$ be a trace under **CC**. Then, τ satisfies transactions isolation and causal delivery. It is important to notice that if τ has no write-write race then the contents of *store* at a given variable will contain a single value at any time during τ . This implies that *store* can be simulated by a single value memory which does not discard writes. Thus, we obtain a program semantics that is the same as the one for **CM**. Thus, τ is also a trace of \mathcal{P} under **CM**. To prove that $\tau \in \text{Tr}_{\text{CCv}}$, we also need to ensure that the transitive closure of store order in τ is acyclic which is enough to guarantee the existence of a total arbitration between transactions which is ensured by **CCv** semantics. Suppose by contradiction that the transitive closure of store order is cyclic then there must exist a sequence of events $ev_1 \cdot ev_2 \cdot \dots \cdot ev_n$ in τ such that $(ev_i, ev_{i+1}) \in \text{WW}$, for all $1 \leq i \leq n-1$ and $(ev_n, ev_1) \in \text{WW}$. Since τ has no write-write races then $(ev_i, ev_{i+1}) \in \text{WW}$ implies that the issue events corresponding to ev_i and ev_{i+1} must be related by causal ordered (since the corresponding transactions must be causally related to prevent concurrency which will lead to write-write races for transactions that write to a common variable). For all i s.t. $1 \leq i \leq n-1$, let ev'_i and ev'_{i+1} denote these issue events then $(ev'_i, ev'_{i+1}) \in \text{CO}$ which implies that the causal order **CO** is cyclic. This is a contradiction since it is not possible under **CC**. Thus, there exists a total order between transactions in τ that includes both the causal order and the transitive closure of store order. Thus, τ is also a trace of \mathcal{P} under **CCv**. \square

4. PROGRAM ROBUSTNESS

4.1. Program Semantics Under Serializability. The semantics of a program under serializability [37] can be defined using a transition system where the configurations keep a single shared-variable valuation (accessed by all processes) with the standard interpretation of read or write statements. Each transaction executes in isolation. Alternatively, the serializability semantics can be defined as a restriction of $[\mathcal{P}]_{\mathcal{X}}$, $\mathcal{X} \in \{\text{CCv}, \text{CM}, \text{CC}\}$, to the set of executions where each transaction is *immediately* delivered to all processes, i.e., each event $\text{end}(p, t)$ is immediately followed by all $\text{del}(p', t)$ with $p' \in \mathbb{P}$. Such executions are called *serializable* and the set of serializable executions of a program \mathcal{P} is denoted by $\text{EX}_{\text{SER}}(\mathcal{P})$. The

⁷Note that other cases such as $\tau' = \alpha \cdot \text{isu}(p_1, t_1) \cdot \beta \cdot \text{isu}(p_2, t_2) \cdot \text{del}(p_2, t_2) \cdot \text{del}(p_1, t_2) \cdot \text{del}(p_1, t_1)$ implies that $\tau'' = \alpha \cdot \text{isu}(p_1, t_1) \cdot \text{del}(p_1, t_1) \cdot \beta \cdot \text{isu}(p_2, t_2) \cdot \text{del}(p_2, t_2) \cdot \text{del}(p_1, t_2) \cdot \text{del}(p_2, t_1)$ is a trace of \mathcal{P} as well since all events in β are not causally dependent on t_1 .

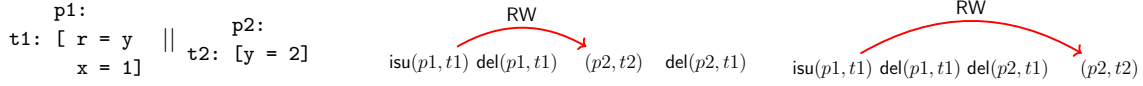


FIGURE 8. Two executions of the same serializable trace.

latter definition is easier to reason about when relating executions under causal consistency and serializability, respectively.

Given a trace $tr = (\tau, \text{PO}, \text{WR}, \text{WW}, \text{RW}, \text{STO})$ of a serializable execution, we have that every event $\text{isu}(p, t)$ in τ is immediately followed by all $\text{del}(p', t)$ with $p' \in \mathbb{P}$. For simplicity, we write τ as a sequence of “atomic macro-events” (p, t) where (p, t) denotes a sequence $\text{isu}(p, t) \cdot \text{del}(p, t) \cdot \text{del}(p_1, t) \cdot \dots \cdot \text{del}(p_n, t)$ with $\mathbb{P} = \{p, p_1, \dots, p_n\}$. We say that t is *atomic*. In Fig. 7, $t3$ is atomic and we can use $(p2, t3)$ instead of $\text{isu}(p2, t3) \cdot \text{del}(p2, t3) \cdot \text{del}(p1, t3)$.

The following result characterizes traces of serializable executions, and follows from previous works [2, 41] that considered a notion of history/trace that corresponds to our notion of transactional happens-before. The transactional happens-before of any trace under SER is acyclic, and conversely, any trace obtained under a weaker semantics $\mathbb{X} \in \{\text{CCv}, \text{CM}, \text{CC}\}$ with an acyclic transactional happens-before can be transformed into a trace under SER by successive swaps of consecutive events in its summary, which are not related by happens-before (the happens-before relations remain the same). Indeed, note that multiple executions/traces can have the same (transactional) happens-before (an example for traces is given in Fig. 8). In particular, it is possible that a trace tr produced by a variation of causal consistency has an acyclic transactional happens-before even though $\text{isu}(p, t)$ events are not immediately followed by the corresponding $\text{del}(p', t)$ events. However, tr would be equivalent, up to reordering of consecutive summary events that are not related by happens-before to a serializable trace.

Theorem 4.1 [2, 41]. *For any trace $tr \in \text{Tr}_{\text{SER}}(\mathcal{P})$, the transactional happens-before of tr is acyclic. Moreover, for any trace $tr = (\tau, \text{PO}, \text{WR}, \text{WW}, \text{RW}, \text{STO}) \in \text{Tr}_{\mathbb{X}}(\mathcal{P})$ with $\mathbb{X} \in \{\text{CCv}, \text{CM}, \text{CC}\}$, if the transactional happens-before of tr is acyclic, then there exists a permutation τ' of τ such that $(\tau', \text{PO}, \text{WR}, \text{WW}, \text{RW}, \text{STO}) \in \text{Tr}_{\text{SER}}(\mathcal{P})$.*

As a consequence of Theorem 4.1, we define a trace tr to be *serializable* if it has the same happens-before relations as a trace of a serializable execution. Let $\text{Tr}_{\text{SER}}(\mathcal{P})$ denote the set of serializable traces of a program \mathcal{P} .

4.2. Robustness Problem. We consider the problem of checking whether the causally-consistent semantics of a program produces only serializable traces (it produces all serializable traces because every issue event can be immediately followed by all the corresponding store events).

Definition 4.2. A program \mathcal{P} is called *robust* against a semantics $\mathbb{X} \in \{\text{CCv}, \text{CM}, \text{CC}\}$ iff $\text{Tr}_{\mathbb{X}}(\mathcal{P}) = \text{Tr}_{\text{SER}}(\mathcal{P})$.

A trace $tr \in \text{Tr}_{\mathbb{X}}(\mathcal{P}) \setminus \text{Tr}_{\text{SER}}(\mathcal{P})$ is called a *robustness violation* (or *violation*, for short). By Theorem 4.1, the transactional happens-before HB_t of tr is cyclic.

We discuss several examples of programs which are (non-) robust against both CM and CCv or only one of them. Robustness violations are presented in terms of “observable” behaviors, tuples of values that can be read in the different transactions and that are not possible

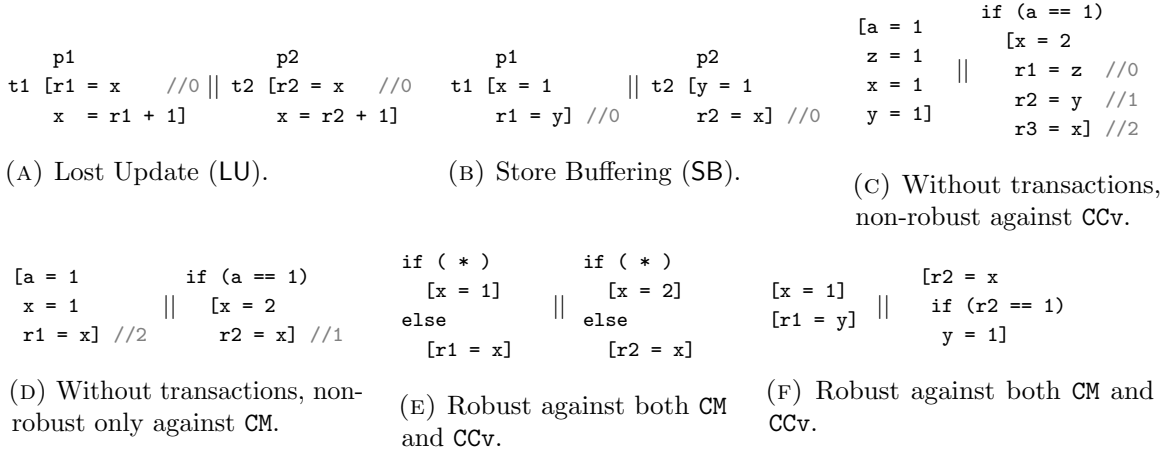


FIGURE 9. (Non-)robust programs. For non-robust programs, the read instructions are commented with the values they return in robustness violations. The condition of `if-else` is checked inside a transaction whose demarcation is omitted for readability (* denotes non-deterministic choice).

under the serializability semantics (they correspond to traces with acyclic transactional happens-before). Fig. 9a and Fig. 9b show examples of programs that are *not* robust against both CM and CCv, which have also been discussed in the literature on weak memory models, e.g. [6]. The execution of Lost Update under both CM and CCv allows that the two reads of `x` in transactions `t1` and `t2` return 0 although this cannot happen under serializability. Also, executing Store Buffering under both CM and CCv allows that the reads of `x` and `y` return 0 although this would not be possible under serializability. These values are possible because the transaction in each of the processes may not be delivered to the other process.

Assuming for the moment that each instruction in Fig. 9c and Fig. 9d forms a different transaction, the values we give in comments show that the program in Fig. 9c, resp., Fig. 9d, is not robust against CCv, resp., CM. The values in Fig. 9c are possible assuming that the timestamp of the transaction `[x = 1]` is smaller than the timestamp of `[x = 2]` (which means that if the former is delivered after the second process executes `[x = 2]`, then it will be discarded). Moreover, enlarging the transactions as shown in Fig. 9c, the program becomes robust against CCv. The values in Fig. 9d are possible under CM because different processes do not need to agree on the order in which to apply transactions, each process applying the transaction received from the other process last. However, under CCv this behavior is not possible, the program being actually robust against CCv. As in the previous case, enlarging the transactions as shown in the figure leads to a robust program against CM.

We end the discussion with several examples of programs that are robust against both CM and CCv. These are simplified models of real applications reported in [27]. The program in Fig. 9e can be understood as the parallel execution of two processes that either create a new user of some service, represented abstractly as a write on a variable `x` or check its credentials, represented as a read of `x` (the non-deterministic choice abstracts some code that checks whether the user exists). Clearly this program is robust against both CM and CCv since each process does a single access to the shared variable. Although we considered simple transactions that access a single shared-variable this would hold even for “bigger”

transactions that access an arbitrary number of variables. The program in Fig. 9f can be thought of as a process creating a new user of some service and reading some additional data in parallel to a process that updates that data only if the user exists. It is rather easy to see that it is also robust against both **CM** and **CCv**.

5. MINIMAL VIOLATIONS

We define a class of robustness violations called *minimal* violations. The particular shapes of these violations, that we determine through a series of results in this section, Section 6, and Section 7, enables a polynomial-time reduction of robustness checking to a reachability problem in a program running under serializability.

For simplicity, we use “atomic macro-events” (p, t) even in traces obtained under causal consistency (recall that this notation was introduced to simplify serializable traces), i.e., we assume that any sequence of events formed of an issue $\text{isu}(p, t)$ followed immediately by all the store events $\text{del}(p', t)$ is replaced by (p, t) . Then, all the relations that held between an event ev of such a sequence and another event ev' , e.g., $(ev, ev') \in \text{PO}$, are defined to hold as well between the corresponding macro-event (p, t) and ev' , e.g., $((p, t), ev') \in \text{PO}$.

5.1. Happens-Before Through Relation. To decide if two events in a trace are “independent” (or commutative) we use the information about the existence of a happens-before relation between the events. If two events are not related by happens-before then they can be swapped while preserving the same happens-before. Thus, we extend the happens-before relation to obtain the *happens-before through* relation as follows:

Definition 5.1. Let $\tau = \alpha \cdot a \cdot \beta \cdot b \cdot \gamma$ be a trace where a and b are events (or atomic macro events), and α , β , and γ are sequences of events (or atomic macro events) under a semantics $X \in \{\text{CCv}, \text{CM}\}$. We say that a *happens-before through* b if there is a non empty sub-sequence $c_1 \cdots c_n$ of β that satisfies:

$$c_i \rightarrow_{\text{HB}^1} c_{i+1} \quad \text{for all } i \in [0, n]$$

where $c_0 = a$, $c_{n+1} = b$.

The following result shows that any two events in a trace which are not related via the happens-before through relation can be reordered without affecting the happens-before or they can be placed one immediately after the other.

Lemma 5.2. *Let τ be a trace of a program \mathcal{P} under a semantics $X \in \{\text{CCv}, \text{CM}\}$, and a and b be two events such that $\tau = \alpha \cdot a \cdot \beta \cdot b \cdot \gamma$. Then, one of the following holds:*

- (1) a happens-before through b ;
- (2) $\tau' = \alpha \cdot \beta_1 \cdot a \cdot b \cdot \beta_2 \cdot \gamma \in \text{Tr}_X(\mathcal{P})$ where $(a, b) \in \text{HB}^1$ has the same happens-before as τ ;
- (3) $\tau' = \alpha \cdot \beta_1 \cdot b \cdot a \cdot \beta_2 \cdot \gamma \in \text{Tr}_X(\mathcal{P})$ has the same happens-before as τ .

Proof. We prove that $\neg(1) \Rightarrow ((2) \text{ or } (3))$ using induction on the size of β .

Base case: If $|\beta| = 0$, then $\tau = \alpha \cdot a \cdot b \cdot \gamma$, which implies that a does not happen-before b through β (by definition, β cannot be empty). Thus, either a and b are HB^1 -related, which corresponds to (2), or a and b are not HB^1 -related, which implies that b can move to the left of a producing the trace $\tau' = \alpha \cdot b \cdot a \cdot \gamma$ that has the same happens-before as τ and that corresponds to (3).

Induction step: We assume that the lemma holds for $|\beta| \leq n$. Consider $\tau_{n+1} = \alpha \cdot a \cdot \beta \cdot b \cdot \gamma$ with $|\beta| = n + 1$. Consider c the last event in the sequence $\beta = \beta_1 \cdot c$. If a does not happen before b through β , then either a does not happen before c through β_1 and a and c are not HB^1 -related, or c and b are not HB^1 -related.

First case: suppose that a does not happen before c through β_1 and a and c are not HB^1 -related. Using the induction hypothesis over τ_{n+1} with respect to a and c (since $|\beta_1| \leq n$) results in $\tau'_{n+1} = \alpha \cdot \beta_{11} \cdot c \cdot a \cdot \beta_{12} \cdot b \cdot \gamma$ that has the same happens-before as τ_{n+1} . We know that if a happens-before b through β_{12} then a happens-before b through β because β_{12} is a subset of β . Therefore, a does not happen-before b through β_{12} . Since $|\beta_{12}| \leq |\beta_1| \leq n$, then we can apply the induction hypothesis to τ'_{n+1} with respect to a and b which yields either $\tau''_{n+1} = \alpha \cdot \beta_{11} \cdot c \cdot \beta_{121} \cdot b \cdot a \cdot \beta_{122} \cdot \gamma$ which has the same happens-before as τ'_{n+1} , if a and b are not HB^1 -related, or $\tau''_{n+1} = \alpha \cdot \beta_{11} \cdot c \cdot \beta_{121} \cdot a \cdot b \cdot \beta_{122} \cdot \gamma$ which has the same happens-before as τ'_{n+1} , otherwise.

Second case: suppose c and b are not HB^1 -related. We apply the induction hypothesis to τ_{n+1} with respect to c and b , and we get $\tau'_{n+1} = \alpha \cdot a \cdot \beta_1 \cdot b \cdot c \cdot \gamma$ with the same happens-before as τ_{n+1} . As we already know that a does not happen before b through β then a does not happen before b through β_1 . Subsequently by using the induction hypothesis over τ'_{n+1} with respect to a and b , we obtain $\tau''_{n+1} = \alpha \cdot \beta_{11} \cdot b \cdot a \cdot \beta_{12} \cdot c \cdot \gamma$ where τ''_{n+1} has the same happens-before as τ'_{n+1} , if a and b are not HB^1 -related, or $\tau''_{n+1} = \alpha \cdot \beta_{11} \cdot a \cdot b \cdot \beta_{12} \cdot c \cdot \gamma$ where τ''_{n+1} has the same happens-before as τ'_{n+1} , otherwise. \square

We show next that a robustness violation should contain at least an issue and a store event of the same transaction that are separated by another event that occurs after the issue and before the store and which is related to both via the happens-before relation. Otherwise, since any two events which are not related by happens-before could be swapped in order to derive a trace with the same happens-before, every store event could be swapped until it immediately follows the corresponding issue and the trace would be serializable.

Lemma 5.3. *Given a violation τ , there must exist a transaction t such that $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot \text{del}(p_0, t) \cdot \gamma$ and $\text{isu}(p, t)$ happens-before $\text{del}(p_0, t)$ through β .*

Proof. Assume by contradiction that the lemma does not hold. For every transaction t of τ suppose there exist $p' \in \mathbb{P}$ such that $\text{del}(p', t)$ does not occur immediately after $\text{isu}(p, t)$. Thus, $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot \text{del}(p', t) \cdot \gamma$, and $(\text{isu}(p, t), \text{del}(p', t)) \in \text{STO} \subset \text{HB}^1$. From Lemma 5.2, $\tau' = \alpha \cdot \beta_1 \cdot \text{isu}(p, t) \cdot \text{del}(p', t) \cdot \beta_2 \cdot \gamma$ has the same happens-before as τ (since $\text{isu}(p, t)$ does not happens-before $\text{del}(p', t)$ through β). Then, the trace τ^* where for every transaction t of τ the store events occur immediately after the issue event has the same happens-before as τ . Thus, τ^* is serializable which means that its HB_t is acyclic which contradicts the fact that τ is a violation. \square

The transaction t in the trace τ above is called a *delayed* transaction. The happens-before constraints imply that t belongs to a transactional happens-before cycle in the trace. In the remainder of the paper, when given a violation $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot \text{del}(p_0, t) \cdot \gamma$, we assume that t is the *first* delayed transaction in τ .

5.2. Minimal Violations. Given a trace $\tau = \alpha \cdot b \cdot \beta \cdot c \cdot \omega$ containing two events $b = \text{isu}(p, t)$ and c , the *distance* between b and c , denoted by $d_\tau(b, c)$, is the number of events in β that are causally related to b , excluding events that correspond to the delivery of t , i.e., $d_\tau(b, c) = |\{d \in \beta \mid (b, d) \in \text{CO} \wedge d \neq \text{del}(p', t) \text{ for every } p' \in \mathbb{P}\}|$

The *number of delays* $\#(\tau)$ in a trace τ is the sum of all distances between issue and store events that originate from the same transaction:

$$\#(\tau) = \sum_{\text{isu}(p,t), \text{del}(p',t) \in \tau} d_{\tau}(\text{isu}(p,t), \text{del}(p',t))$$

Definition 5.4 (Minimal Violation). A robustness violation τ is called *minimal* if it has the least number of delays among all robustness violations (for a given program \mathcal{P} and semantics $\mathsf{X} \in \{\text{CC}, \text{CCv}, \text{CM}\}$).

Remark 5.5. It is important to note that a non-robust program can admit multiple minimal violations with different happens-before relations. For instance, Fig. 10 pictures two minimal violations that do not have the same happens-before and both traces have 0 delays. In the trace in Fig. 10b a single transaction is delayed while in the trace in Fig. 10c two transactions are delayed and are not causally related. For the trace τ_1 in Fig. 10b, we have that $\#(\tau_1) = d_{\tau_1}(\text{isu}(p2, t2), \text{del}(p2, t2)) + d_{\tau_1}(\text{isu}(p2, t2), \text{del}(p3, t2)) = 0$. For the trace τ_2 in Fig. 10c, we have that $\#(\tau_2) = d_{\tau_2}(\text{isu}(p1, t1), \text{del}(p1, t1)) + d_{\tau_2}(\text{isu}(p1, t1), \text{del}(p3, t1)) + d_{\tau_2}(\text{isu}(p2, t2), \text{del}(p3, t2)) = 0$. Hence, the number of delays for both cases is 0.

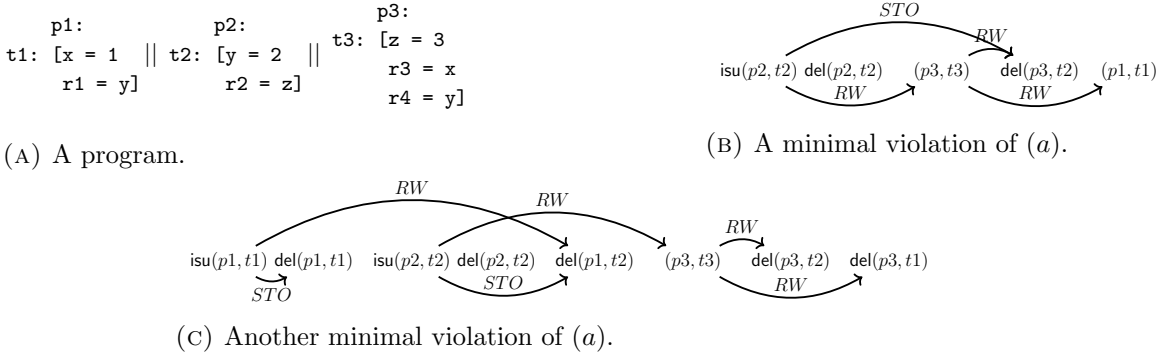


FIGURE 10. Example of two minimal violation traces that do not have the same happens-before relation (possible under both CCv and CM). Both traces have the same number of delays which is equal to 0. The minimal violation in (b) contains a single delayed transaction ($t2$), and the minimal violation in (c) contains two delayed transactions ($t1$ and $t2$). For readability, we do not show all PO and STO transitions.

Given a minimal violation $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot \text{del}(p_0, t) \cdot \gamma$, the following lemma shows that we can assume w.l.o.g. that γ contains only store events from transactions that were issued before $\text{del}(p_0, t)$ in τ .

Lemma 5.6. *Let $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot \text{del}(p_0, t) \cdot \gamma$ be a minimal violation such that $\text{isu}(p, t)$ happens-before $\text{del}(p_0, t)$ through β . Then, $\tau' = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot \text{del}(p_0, t) \cdot \gamma'$, such that γ' contains only store events from transactions that were issued before $\text{del}(p_0, t)$ in τ , is also a minimal violation.*

Proof. The prefix $\alpha \cdot \text{isu}(p, t) \cdot \beta \cdot \text{del}(p_0, t)$ has a cyclic transactional happens-before and it is already a minimal violation independently of whether γ contains additional transactions. \square

The following result shows that for every minimal violation, we can extract another minimal violation of the shape $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma$ such that $(\text{isu}(p, t), (p', t')) \in \text{HB}$, and $((p', t'), \text{del}(p', t)) \in \text{HB}^1$.

Lemma 5.7. *If \mathcal{P} is a program that is not robust against some $X \in \{\text{CCv}, \text{CM}, \text{CC}\}$, then its set of traces under the semantics X must admit a minimal violation of the shape $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma$ such that $(\text{isu}(p, t), (p', t')) \in \text{HB}$ and $((p', t'), \text{del}(p', t)) \in \text{HB}^1$.*

Proof. Let $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta \cdot \text{del}(p_0, t) \cdot \gamma$ be a minimal violation of \mathcal{P} , such that $\text{isu}(p, t)$ happens-before $\text{del}(p_0, t)$ through β . By Lemma 5.6, we assume that γ contains only store events. We prove by induction on the size of β that \mathcal{P} admits another minimal violation against X of the form $\tau' = \alpha' \cdot \text{isu}(p_1, t_1) \cdot \beta' \cdot (p_2, t_2) \cdot \text{del}(p_2, t_1) \cdot \gamma'$ such that $(\text{isu}(p_1, t_1), (p_2, t_2)) \in \text{HB}$, $((p_2, t_2), \text{del}(p_2, t_1)) \in \text{HB}^1$, and τ' is a permutation of a subsequence of τ .

Note that $\text{isu}(p, t)$ happens-before $\text{del}(p_0, t)$ through β implies that there exists a subsequence $c_1 \cdots c_n$ of β that satisfies: $c_i \rightarrow_{\text{HB}^1} c_{i+1}$ for all $i \in [0, n]$ where $c_0 = \text{isu}(p, t)$, $c_{n+1} = \text{del}(p_0, t)$. Then, we have three possibilities for c_n : (p', t') , $\text{isu}(p', t')$, or $\text{del}(p_0, t')$.

Base case: $|\beta| = 1$ implies that $\beta = c_n$. If $c_n = (p', t')$ then τ is a minimal violation s.t. $\text{isu}(p, t) \rightarrow_{\text{HB}} (p', t')$ and $(p', t') \rightarrow_{\text{HB}^1} \text{del}(p_0, t)$. If $c_n = \text{isu}(p', t')$ then we regroup together the issue event $\text{isu}(p', t')$ with its store events obtaining $\tau' = \alpha \cdot \text{isu}(p, t) \cdot (p', t') \cdot \text{del}(p_0, t) \cdot \gamma'$ to be a minimal violation as well (since the transactional happens-before of the trace resulting from reordering store events in $\text{del}(p_0, t) \cdot \gamma'$ will always be cyclic). Since $(p', t') \rightarrow_{\text{HB}^1} \text{del}(p_0, t)$ implies that $(p', t') \rightarrow_{\text{HB}^1} \text{del}(p', t) \in \gamma$, then $\tau'' = \alpha \cdot \text{isu}(p, t) \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma''$, where the two store events $\text{del}(p', t)$ and $\text{del}(p', t)$ are reordered, is a minimal violation. $c_n = \text{del}(p_0, t')$ is not possible since t is the first delayed transaction in τ .

Induction step: We assume that the induction hypothesis holds for $|\beta| \leq m$. The case $c_n = (p', t')$ is trivial. If $c_n = \text{isu}(p', t')$ then removing the issue events that occur after c_n will not impact the happens-before. Thus, we remove every issue and atomic marco event that occurs after $\text{isu}(p', t')$ with all their store events and regroup together the event $\text{isu}(p', t')$ with its store events obtaining $\tau' = \alpha \cdot \text{isu}(p, t) \cdot \beta' \cdot (p', t') \cdot \text{del}(p_0, t) \cdot \gamma'$ to be a minimal violation. Similar to before, $\tau'' = \alpha \cdot \text{isu}(p, t) \cdot \beta' \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma''$ is a minimal violation.

If $c_n = \text{del}(p_0, t')$, then the corresponding issue event $\text{isu}(p', t')$ must occur in β (α contains only atomic macro events because t is the first delayed transaction). If $\text{isu}(p', t')$ does not happen before $\text{del}(p_0, t')$ (or any store event of t' in $\beta \cdot \text{del}(p_0, t) \cdot \gamma$) through a subsequence of β (resp., $\beta \cdot \text{del}(p_0, t) \cdot \gamma$) then we can regroup together the issue and store events of t' and get that $\tau' = \alpha \cdot \text{isu}(p, t) \cdot \beta' \cdot (p', t') \cdot \beta'' \cdot \text{del}(p_0, t) \cdot \gamma'$ is a minimal violation. Otherwise, if $\text{isu}(p', t')$ happens-before $\text{del}(p_0, t')$ through a subsequence of β , then τ can be written as $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p', t') \cdot \beta_2 \cdot \text{del}(p_0, t') \cdot \beta_3 \cdot \text{del}(p_0, t) \cdot \gamma$. Note that if there exists an issue event $\text{isu}(p_1, t_1)$ in $\beta_1 \cdot \text{isu}(p', t') \cdot \beta_2$ s.t. $(\text{isu}(p_1, t_1), \text{del}(p_0, t)) \in \text{RW}$ (or $(\text{isu}(p_1, t_1), \text{del}(p_1, t)) \in \text{RW}$) then similar to before the following trace $\tau' = \alpha \cdot \text{isu}(p, t) \cdot \beta' \cdot (p_1, t_1) \cdot \text{del}(p_0, t) \cdot \gamma'$ (resp., $\tau' = \alpha \cdot \text{isu}(p, t) \cdot \beta' \cdot (p_1, t_1) \cdot \text{del}(p_1, t) \cdot \gamma'$) is a minimal violation. Assume now that there does not exist an issue event $\text{isu}(p_1, t_1)$. Then, let $\text{isu}(p_2, t_2)$ be the first issue event in $\text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p', t')$ s.t. $\tau = \alpha \cdot \text{isu}(p, t) \cdot \beta'_1 \cdot \text{isu}(p_2, t_2) \cdot \beta'_2 \cdot \text{del}(p_3, t_2) \cdot \beta'_3 \cdot \gamma$ and $\text{isu}(p_2, t_2)$ happens-before $\text{del}(p_3, t_2)$ through β'_2 and s.t. for every issue event in $\text{isu}(p, t) \cdot \beta'_1$ of a transaction t_4 there does not exist an event in $\beta'_1 \cdot \text{isu}(p_2, t_2) \cdot \beta'_2$ that reads from a variable that t_4 overwrites. We can remove every issue event and atomic marco event which occur after $\text{del}(p_3, t_2)$ with all related stores: $\tau' = \alpha \cdot \text{isu}(p, t) \cdot \beta'_1 \cdot \text{isu}(p_2, t_2) \cdot \beta'_2 \cdot \text{del}(p_3, t_2) \cdot \gamma'$ where γ' contains only store events is a minimal violation. Then, not delaying the transactions

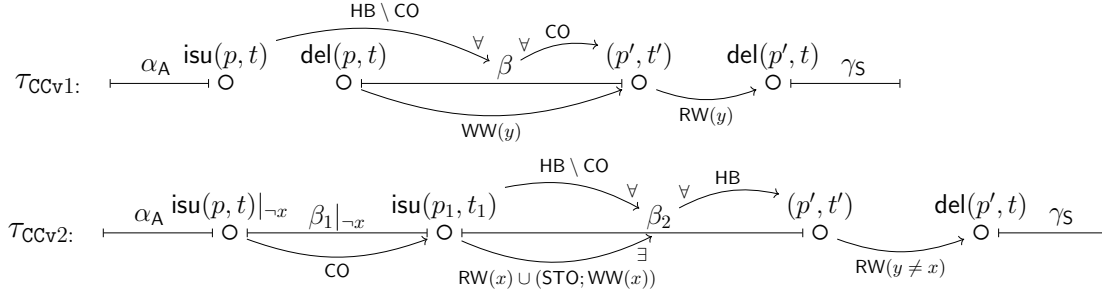
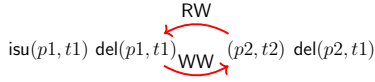
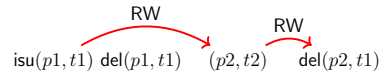


FIGURE 11. Robustness violation patterns under CCv . We use $a \xrightarrow{R} \forall \beta$ to denote $\forall b \in \beta. (a, b) \in R$. We use $\beta_1|_{\neg x}$ to say that all delayed transactions in β_1 do not access x . For violation τ_{CCv1} , t is the only delayed transaction. For τ_{CCv2} , all delayed transactions are in $\text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1)$ and they form a causality chain that starts at $\text{isu}(p, t)$ and ends at $\text{isu}(p_1, t_1)$.



(A) Violation of LU program in Fig. 9a.



(B) Violation of SB program in Fig. 9b.

FIGURE 12. (a) A τ_{CCv1} violation where $\beta_2 = \epsilon$, $\gamma_S = \epsilon$, and t and t' correspond to t_1 and t_2 . (b) A τ_{CCv2} (resp., τ_{CM2}) violation where t and t_1 coincide and correspond to t_1 . Also, $\beta_1 = \epsilon$, $\beta_2 = (p_2, t_2)$, $\gamma_S = \epsilon$, such that $(\text{isu}(p_1, t_1), (p_2, t_2)) \in \text{RW}(y)$ and $((p_2, t_2), \text{del}(p_2, t_1)) \in \text{RW}(x)$. In all traces, we show only the relations that are part of the happens-before cycle.

in $\text{isu}(p, t) \cdot \beta'_1$ does not affect the reads in $\beta'_1 \cdot \text{isu}(p_2, t_2) \cdot \beta'_2$, and thus, we get that $\tau'' = \alpha \cdot (p, t) \cdot \beta''_1 \cdot \text{isu}(p_2, t_2) \cdot \beta''_2 \cdot \text{del}(p_3, t_2) \cdot \gamma''$, where t_2 is the first delayed transaction in τ'' and $\text{isu}(p_2, t_2)$ happens-before $\text{del}(p_3, t_2)$ through β''_2 , is a minimal violation. Note that $|\beta''_2| < |\beta| = m + 1$, and we can apply the induction hypothesis to τ'' and conclude the proof. \square

Next, we show that a program which is not robust against CCv or CM admits violations of particular shapes. For the remainder of the paper, we write a minimal violation in the shape $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ to say that all the events in the sequence α_A are atomic macro events and all the events in the sequence γ_S are store events. As before, we assume that t is the first delayed transaction in τ , and by Lemma 5.7, we assume that $(\text{isu}(p, t), (p', t')) \in \text{HB}$ and $((p', t'), \text{del}(p', t)) \in \text{HB}^1$.

6. ROBUSTNESS VIOLATIONS UNDER CAUSAL CONVERGENCE

In this section, we present a precise characterization of minimal violations under CCv . In particular, we show that in these violations, the first delayed transaction (which must exist by Lemma 5.3) is followed by a possibly-empty sequence of delayed transactions that form a “causality chain”, i.e., the issue of every new delayed transaction is causally ordered after the issue of the first delayed transaction. Also, we show that the issue event of the last delayed transaction happens-before an event of another transaction that reads a variable updated

by the first delayed transaction (which implies a cycle in the transactional happens-before). This characterization will allow us to build a monitor for detecting the existence of robustness violations that is linear in the size of the input program.

Next, we give a precise definition of the “causality chain”. It consists of a sequence of issue events such that the first issue is causally ordered before every other issue event and every issued transaction is delivered to the process executing the next issue event in the chain, before this issue event executes.

Definition 6.1. We say that a sequence of issue events $ev_1 \cdot ev_2 \cdot \dots \cdot ev_n$ forms a *causality chain* that starts with ev_1 and ends at ev_n in a trace τ if the followings hold:

- (1) $(ev_1, ev_i) \in \text{CO}$, for all $2 \leq i \leq n$
- (2) for all $1 \leq i \leq n - 1$ such that $ev_i = \text{isu}(p_i, t_i)$, $ev_{i+1} = \text{isu}(p_{i+1}, t_{i+1})$, the store event $\text{del}(p_{i+1}, t_i)$ occurs before the issue event ev_{i+1} in τ .

The characterization of robustness violations under CCv is stated in the following theorem and pictured in Fig. 11.

Theorem 6.2. *A program \mathcal{P} is not robust under CCv iff there exists a minimal violation in $\text{Tr}(\mathcal{P})_{\text{CCv}}$ of one of the following forms:*

- (1) $\tau_{\text{CCv1}} = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ where:
 - (a) $\text{isu}(p, t)$ is the issue of the first and only delayed transaction (Lemma 6.3);
 - (b) $\exists y. \text{s.t. } (\text{del}(p, t), (p', t')) \in \text{WW}(y)$ and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$ (Lemma 6.3);
 - (c) $\forall a \in \beta_2. (\text{isu}(p, t), a) \in \text{HB} \setminus \text{CO}$ and $(a, (p', t')) \in \text{CO}$ (Lemma 6.3).
- (2) $\tau_{\text{CCv2}} = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ where:
 - (a) $\text{isu}(p, t)$ and $\text{isu}(p_1, t_1)$ are the issues of the first and last delayed transactions (Lemmas 6.3 and 6.4);
 - (b) the issues of all delayed transactions are in β_1 and are included in a causality chain that starts with $\text{isu}(p, t)$ and ends at $\text{isu}(p_1, t_1)$ (Lemma 6.4);
 - (c) for every $a \in \beta_2$, we have that $(\text{isu}(p_1, t_1), a) \in \text{HB} \setminus \text{CO}$ and $(a, (p', t')) \in \text{HB}$ (Lemma 6.3);
 - (d) there exist $a \in \beta_2 \cdot (p', t')$, x , and y s.t. $x \neq y$, $(\text{isu}(p_1, t_1), a) \in \text{RW}(x) \cup (\text{STO}; \text{WW}(x))$, $(a, (p', t')) \in \text{HB}^{?8}$, and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$ (Lemma 6.3);
 - (e) all delayed transactions in $\text{isu}(p, t) \cdot \beta_1$ do not access the variable x (Lemma 6.6).

Above, τ_{CCv1} contains a single delayed transaction while τ_{CCv2} may contain arbitrarily many delayed transactions. In τ_{CCv1} the store event $\text{del}(p, t)$ of the only delayed transaction happens before (p', t') which is conflicting with t , thus resulting in a cycle in the transactional happens-before. In τ_{CCv2} the issue event of the last delayed transaction t_1 , which is causally ordered after the issue of the first delayed transaction t , happens before (p', t') which is conflicting with t , thus resulting in a cycle in the transactional happens-before as well. The theorem above allows $\alpha_A = \epsilon$, $\beta_1 = \epsilon$, $\beta_2 = \epsilon$, $\beta = \epsilon$, $\gamma_S = \epsilon$, $p = p_1$, $t = t_1$, and t_1 to be a read-only transaction. Fig. 12a and Fig. 12b show two violations under CCv where such equalities hold. If t_1 is a read-only transaction then $\text{isu}(p_1, t')$ has the same effect as (p_1, t_1) since t_1 does not contain writes.

The minimality of the violation enforces the constraints stated above. For example, in the context of τ_{CCv2} , the delayed transactions in β_1 cannot create a cycle in the transactional

⁸HB? is the reflexive closure of HB.

happens-before (otherwise, there exists a sequence of store events γ'_S such that $\alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta_1 \cdot \text{del}(p_0, t) \cdot \gamma'_S$ is a violation with a smaller measure, which contradicts minimality). Moreover, (c) implies that β_2 contains no stores of delayed transactions from β_1 . If this were the case, then these stores can either be reordered after $\text{del}(p', t)$ or if this is not possible due to happens-before constraints, then there would exist an issue event which is after such a store in the happens-before order and thus causally after $\text{isu}(p, t)$, which would contradict the fact that $\text{isu}(p_1, t_1)$ is the last issue event in τ that is causally ordered after $\text{isu}(p, t)$. Also, if it were to have a delayed transaction t_2 in β_2 (resp., β for τ_{CCv1}), then it is possible to remove some transaction (the issue and all its store events) from the original trace and obtain a new violation trace with a smaller number of delays. For instance, in the case of β_2 , if $t_1 \neq t$, then we can remove the events of the last delayed transaction (i.e., t_1), that is causally related to $\text{isu}(p, t)$, since all events in $\beta_2 \cdot \text{del}(p_0, t) \cdot \gamma_S$ neither read from the writes of t_1 nor are issued by the same process as t_1 (because of the $\text{HB} \setminus \text{CO}$ relation between events β_2 and $\text{isu}(p_1, t_1)$). The resulting trace is still a robustness violation (because of the transactional happens-before cycle involving t_2 since it is delayed in β_2) but with a smaller measure. Note that all processes that delayed transactions, stop executing new transactions in β_2 (resp., β) because of the relation $\text{HB} \setminus \text{CO}$, shown in Fig. 11, between the delayed transaction t_1 (resp., t) and events in β_2 (resp., β).

In the following we give a series of lemmas that collectively imply Theorem 6.2. Next lemma gives the decomposition of minimal violations under CCv into two possible patterns. It also characterizes the nature of the happens-before dependencies in these traces. For instance, we show that the last dependency in the happens-before cycle is always a conflict dependency. The lemma proof starts with a minimal violation as characterized in Lemma 5.7 and uses induction to show that we can always obtain a minimal violation which follows one of the two patterns. The induction is based on the size of the sequence of events between the issue and delayed store events of the first delayed transaction (the sequence β in Lemma 5.7).

Lemma 6.3. *If \mathcal{P} is a program that is not robust under CCv , then it must admit a minimal violation τ that satisfies one of the following:*

- (1) $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ where:
 - (a) $\exists y$. s.t. $(\text{del}(p, t), (p', t')) \in \text{WW}(y)$ and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$;
 - (b) $\forall a \in \beta$. $(\text{isu}(p, t), a) \in \text{HB} \setminus \text{CO}$ and $(a, (p', t')) \in \text{CO}$.
- (2) $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ where:
 - (a) $\text{isu}(p_1, t_1)$ is the last issue event in $\{c \in \beta \mid (\text{isu}(p, t), c) \in \text{CO}\}$;
 - (b) $\exists x, y$, and $a \in \beta_2 \cdot (p', t')$ s.t. $(\text{isu}(p_1, t_1), a) \in \text{RW}(x) \cup (\text{STO}; \text{WW}(x))$, $(a, (p', t')) \in \text{HB}?$, and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$;
 - (c) $\forall a \in \beta_2$. $(\text{isu}(p_1, t_1), a) \in \text{HB} \setminus \text{CO}$ and $(a, (p', t')) \in \text{HB}$.

Proof. Let $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ be a minimal violation under CCv (cf. Lemma 5.7). We prove by induction on the size of β that there exists a minimal violation trace τ' that satisfies (1) or (2) and τ' is obtained from τ . By the definition of the happens-before $((p', t'), \text{del}(p', t)) \in \text{HB}^1$ implies that $((p', t'), \text{del}(p', t)) \in \text{RW} \cup \text{WW}$. Since t' was issued after t in τ , then based on the total order of timestamps under CCv , we cannot have $((p', t'), \text{del}(p', t)) \in \text{WW}$. Then, there must exist y s.t. $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$.

Base case: $|\beta| = 0$. Since $(\text{isu}(p, t), (p', t')) \in \text{HB}$ then from the definition of the happens-before the only possible relation is $(\text{isu}(p, t), (p', t')) \in \text{RW}$. Thus, there must exist x s.t. $(\text{isu}(p, t), (p', t')) \in \text{RW}(x)$. If $x = y$ then both t and t' write to x . Thus, by reordering

the store event $\text{del}(p, t) \in \gamma_S$ to occur just after the corresponding issue event we get $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma'_S$ is also a minimal violation where $(\text{del}(p, t), (p', t')) \in \text{WW}(x)$ (since t was issued before t' and both write to x) and $((p', t'), \text{del}(p', t)) \in \text{RW}(x)$. τ' satisfies the first case of the lemma. If $x \neq y$ then $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ where there exist x and y s.t. $x \neq y$, $(\text{isu}(p, t), (p', t')) \in \text{RW}(x)$, and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$ satisfies the second case of the lemma where t and t_1 coincide and a corresponds to (p', t') .

Induction step: We assume the induction hypothesis holds for $|\beta| \leq m$. Let $\sigma = \{c \in \beta \mid (\text{isu}(p, t), c) \in \text{CO}\}$, we will consider the following three possible cases:

First, assume that σ is empty. Since $(\text{isu}(p, t), (p', t')) \in \text{HB}$ then there must exist $a \in \beta \cdot (p', t')$ s.t. $(\text{isu}(p, t), a) \in \text{HB}^1$ and $(a, (p', t')) \in \text{HB}^?$. σ is empty implies that β does not contain events that are related to $\text{isu}(p, t)$ through CO (which includes $\text{PO} \cup \text{WR} \cup \text{STO}$), therefore, $(\text{isu}(p, t), a) \in \text{WW} \cup \text{RW}$. It is impossible to have $(\text{isu}(p, t), a) \in \text{WW}$ since $\text{isu}(p, t)$ does not contain writes. Thus, there must exist x s.t. $(\text{isu}(p, t), a) \in \text{RW}(x)$. If $x = y$ then both the transaction of the event a , denoted t_2 , and t write to x . We consider the two cases of $(a, (p', t')) \in \text{HB}^?$: i) $a = (p', t')$ (i.e., $t_2 = t'$), and ii) $(a, (p', t')) \in \text{HB}$. Assume $a = (p', t')$ then by reordering the store event $\text{del}(p, t) \in \tau$ to occur just after the corresponding issue event (since the events in β are not causally related to $\text{isu}(p, t)$) we get $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma'_S$ is also a minimal violation where $(\text{del}(p, t), (p', t')) \in \text{WW}(x)$ (since t was issued before t' and both write to x) and $((p', t'), \text{del}(p', t)) \in \text{RW}(x)$. In τ' we remove all events in β that are not causally ordered before (p', t') since they do not contribute to the happens-before cycle. We obtain a new violation trace that satisfies the first case of the lemma. Assume now that $(a, (p', t')) \in \text{HB}$. This implies that $\text{isu}(p_2, t_2) \in \beta$ happens-before (p', t') (since a is an event t_2). Since both t_2 and t write to x and t occurs before t_2 in τ then from the definition of store and conflict relations $((p', t'), \text{del}(p', t)) \in \text{RW}(x)$ implies that $((p', t'), \text{del}(p', t_2)) \in \text{RW}(x)$. Also, since in β we do not have events that are causally related to $\text{isu}(p, t)$ then let τ' be the trace resulting from removing all events of t in τ : $\tau' = \alpha_A \cdot \text{isu}(p_2, t_2) \cdot \beta' \cdot (p', t') \cdot \text{del}(p', t_2) \cdot \gamma'_S$ where τ' is a subsequence of τ and β' is a subsequence of β . τ' is a minimal violation as well since it was obtained from τ by just removing events and $(\text{isu}(p_2, t_2), (p', t')) \in \text{HB}$ and $((p', t'), \text{del}(p', t_2)) \in \text{RW}(x)$. Since $|\beta'| \leq m$ then we can apply the induction hypothesis on τ' . If $x \neq y$ we get that in τ , $(\text{isu}(p, t), a) \in \text{RW}(x)$ and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$ which satisfies the second case of the lemma.

Second, assume that σ is not empty and all the elements of σ are store events. Since t is the first delayed transaction in τ then all stores in σ are stores of t . Then, following the same analogy as before there must exist x and an event $a \in \beta \cdot (p', t')$ that is not a store event of t s.t. $(\text{isu}(p, t), a) \in (\text{STO}; \text{WW}(x)) \cup \text{RW}(x)$ and $(a, (p', t')) \in \text{HB}^?$. Similar to before we consider the two cases $x = y$ and $x \neq y$ and apply the induction hypothesis in the first case.

Third, assume that σ is not empty and $\text{isu}(p_1, t_1)$ is the last issue event in σ , i.e., $\beta = \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2$ where all the events in β_2 are either stores of transactions that are causally related to $\text{isu}(p, t)$ (we can reorder these stores to be part of γ_S except the store $\text{del}(p_1, t_1)$) or other events that are not causally related to $\text{isu}(p, t)$. We also have that $\text{isu}(p, t)$ is causally ordered before $\text{isu}(p_1, t_1)$. Since $(\text{isu}(p, t), (p', t')) \in \text{HB}$ then $(\text{isu}(p_1, t_1), (p', t')) \in \text{HB}$, otherwise, we remove $\text{isu}(p_1, t_1)$ and all related store events from τ and the resulting trace is still a violation and it has less delays since $\text{isu}(p, t)$ was not delayed after $\text{isu}(p_1, t_1)$ in the trace. Thus, $(\text{isu}(p_1, t_1), (p', t')) \in \text{HB}$. Similar to before we obtain that there exist x and an event $a \in \beta_2 \cdot (p', t')$ that is not a store event of t_1 s.t. $(\text{isu}(p_1, t_1), a) \in (\text{STO}; \text{WW}(x)) \cup \text{RW}(x)$ and $(a, (p', t')) \in \text{HB}^?$. If $x = y$ then both the transaction of the event a , denoted t_2 , and

t write to x . Thus, $(\text{isu}(p, t), a) \in \text{STO}; \text{WW}(x)$. Then, since the events in $\beta_2 \cdot (p', t')$ do not causally depend on $\text{isu}(p_1, t_1)$ then we can remove the events of t_1 and obtain τ' where $(\text{isu}(p, t), a) \in \text{STO}; \text{WW}(x)$, $(a, (p', t')) \in \text{HB}?$, and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$ where t was not delayed after $\text{isu}(p_1, t_1)$ in the trace, which means that τ' has less delays than τ (a contradiction to τ being a minimal violation). Therefore, we must have $x \neq y$ s.t. $(\text{isu}(p_1, t_1), a) \in (\text{STO}; \text{WW}(x)) \cup \text{RW}(x)$ and $(a, (p', t')) \in \text{HB}?$ and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$ which satisfies the second case of the lemma. \square

We use $\mathbb{T}\text{ccv1}$ and $\mathbb{T}\text{ccv2}$ to denote the class of minimal violations that satisfy the first and second case in Lemma 6.3, respectively. The following lemma shows that we can always obtain a minimal violation trace in either $\mathbb{T}\text{ccv1}$ or $\mathbb{T}\text{ccv2}$ where β and β_2 contain no delayed transactions, respectively. We distinguish two cases in the proof: i) a minimal violation in $\mathbb{T}\text{ccv2}$ where t and t_1 are distinct transactions, and ii) a minimal violation in $\mathbb{T}\text{ccv1}$ or in $\mathbb{T}\text{ccv2}$ where t and t_1 coincide. In the first case, we show that if it were to have a delayed transaction in β_2 , then it is possible to remove some transaction from τ that is causally dependent on the first delayed transaction in τ , and obtain a new violation with a smaller number of delays (which contradicts the minimality assumption). The second case is proved by induction on the size of β (note that if t and t_1 coincide, then $\beta = \beta_2$) where the base case is trivial (i.e., $\beta = \epsilon$), and in the induction step, we show that if it were to have a delayed transaction in β then we can remove one of the delayed transactions in the trace and obtain another violation with the same number of delays as the original violation and for which we can apply the induction hypothesis.

Lemma 6.4. *Let τ be a minimal violation in $\mathbb{T}\text{ccv1}$ or $\mathbb{T}\text{ccv2}$. Then, there exist a violation τ_1 in $\mathbb{T}\text{ccv1}$ where β contains no delayed transactions or a violation τ_2 in $\mathbb{T}\text{ccv2}$ where β_2 contains no delayed transactions.*

Proof. We consider two cases: i) τ in $\mathbb{T}\text{ccv2}$ where t_1 and t are two distinct transactions, ii) τ in $\mathbb{T}\text{ccv1}$ or τ in $\mathbb{T}\text{ccv2}$ where t_1 and t coincide. We prove the first case by contradiction and the second case by induction on the size β (we abused terminology here and considered $\beta_2 = \beta$ since $\beta_1 = \epsilon$ in the second case).

First case: let $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ and suppose by contradiction that β_2 contains a delayed transaction t_0 issued by a process $q \neq p$. W.l.o.g., we assume that the delayed store events of t_0 occur in β_2 . Thus, $\beta_2 = \beta_{21} \cdot \text{isu}(q, t_0) \cdot \beta_{22} \cdot \text{del}(q', t_0) \cdot \beta_{23}$ and $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_{21} \cdot \text{isu}(q, t_0) \cdot \beta_{22} \cdot \text{del}(q', t_0) \cdot \beta_{23} \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$. In τ , $\text{isu}(q, t_0)$ happens-before $\text{del}(q', t_0)$ through β_{22} . Hence, we deduce that we can get a robustness violation when the event $\text{del}(q', t_0)$ is executed, thus we can remove all issued transactions from $\beta_{23} \cdot (p', t')$ except stores of already issued transactions and we obtain: $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_{21} \cdot \text{isu}(q, t_0) \cdot \beta_{22} \cdot \text{del}(q', t_0) \cdot \beta'_{23} \cdot \text{del}(p', t) \cdot \gamma_S$ which is a minimal violation because $\text{isu}(q, t_0)$ happens-before $\text{del}(q', t_0)$ through β_{22} and its number of delays is less or equal to the one of τ . We know that in $\beta_{21} \cdot \text{isu}(q, t_0) \cdot \beta_{22} \cdot \text{del}(q', t_0) \cdot \beta'_{23} \cdot \text{del}(p', t) \cdot \gamma_S$ there are no transactions from the process p_1 or that see the effect of transactions from p_1 (because of the $\text{HB} \setminus \text{CO}$ relation between events β_2 and $\text{isu}(p_1, t_1)$). Therefore, $\text{isu}(p_1, t_1)$ is the last issued transaction from p_1 and we do not have any transaction in τ' that depends on it. Thus, we can remove $\text{isu}(p_1, t_1)$ and we obtain the following trace: $\tau'' = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \beta_{21} \cdot \text{isu}(q, t_0) \cdot \beta_{22} \cdot \text{del}(q', t_0) \cdot \beta'_{23} \cdot \text{del}(p', t) \cdot \gamma'_S$, which is a robustness violation because $\text{isu}(q, t_0)$ happens-before $\text{del}(q', t_0)$ through β_{22} . τ'' has less delays than τ' ($\text{del}(p', t)$ was not delayed after $\text{isu}(p_1, t_1)$ which was removed), which is a contradiction to the fact that τ is a minimal violation.

Second case: let $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$. We show by induction that we can construct either τ_1 in $\mathbb{Tccv1}$ where β of τ_1 contains no delayed transactions or τ_2 in $\mathbb{Tccv2}$ where β_2 of τ_2 contains no delayed transactions.

Base case: $|\beta| = 0$ is trivial.

Induction step: We assume that the induction hypothesis holds for $|\beta| \leq m$. Let t_0 be the first delayed transaction in β . Similar to before, we assume w.l.o.g. that the delayed store events of t_0 occurs in β . Thus, $\beta = \beta_{01} \cdot \text{isu}(q, t_0) \cdot \beta_{02} \cdot \text{del}(q', t_0) \cdot \beta_{03}$ and $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_{01} \cdot \text{isu}(q, t_0) \cdot \beta_{02} \cdot \text{del}(q', t_0) \cdot \beta_{03} \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ where $\text{isu}(q, t_0)$ happens-before $\text{del}(q', t_0)$ through β_{02} . Using the same arguments as before, we can remove the event $\text{isu}(p, t)$, its related stores in τ , and all issued transactions in $\beta_{03} \cdot (p', t')$. We obtain: $\tau' = \alpha'_A \cdot \text{isu}(q, t_0) \cdot \beta_{02} \cdot \text{del}(q', t_0) \cdot \gamma'_S$ where $\alpha'_A = \alpha_A \cdot \beta_{01}$ and $\text{isu}(q, t_0)$ happens-before $\text{del}(q', t_0)$ through β_{02} . τ' is a robustness violation, and it has the same number of delays as τ . We now consider two possible case of τ' : i) τ' is in $\mathbb{Tccv2}$ where t_0 and t_{01} , the last delayed transaction causally dependent on $\text{isu}(q, t_0)$ in τ' , are two distinct transactions, or ii) τ in $\mathbb{Tccv1}$ or τ in $\mathbb{Tccv2}$ where t_0 and t_{01} coincide. From the first part of the proof, it is guaranteed that in the first case there are no delayed transactions after t_{01} . For the second case, we use the induction hypothesis since $|\beta_{02}| \leq m$ (β_{02} is a strict subsequence of β). \square

We have now showed all the necessary characterizations for minimal violations that fall under the first pattern (i.e., $\mathbb{Tccv1}$). In the rest of this section, we focus on minimal violations that fall under the second pattern (i.e., $\mathbb{Tccv2}$). In particular, we look at minimal violations in $\mathbb{Tccv2}$ where t and t_1 are distinct transactions. In the following lemma, we show that for these minimal violations the issue events of delayed transactions in $\text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1)$ constitute a causality chain. Our proof can be decomposed to two parts. In the first part, we show that we cannot have an issue event of a delayed transaction in $\beta_1 \cdot \text{isu}(p_1, t_1)$ that is not causally dependent on $\text{isu}(p, t)$. We prove this by showing that if this were possible then we can remove a transaction that is causally dependent on one of the two delayed transactions and obtain a new violation trace with less delays than the original violation (which contradicts the minimality assumption). For the second part, we show that for a given minimal violation, we can construct a happens-before equivalent trace where for every two successive issue events of delayed transactions in $\text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1)$, the transaction in the first issue is delivered to the process executing the second issue before this event happens.

Lemma 6.5. *Let $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ be a minimal violation in $\mathbb{Tccv2}$ s.t $t \neq t_1$ and β_2 contains no delayed transactions (cf. Lemma 6.4). Then, there exists a violation $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \beta'_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma'_S$ obtained from τ where $\beta'_1 \cdot \gamma'_S$ is a subsequence of $\beta_1 \cdot \gamma_S$ and the sequence of issue events of delayed transactions forms a causality chain that starts with $\text{isu}(p, t)$ and ends at $\text{isu}(p_1, t_1)$.*

Proof. First, we show that we can obtain a violation τ' from τ where all delayed transactions in $\beta'_1 \cdot \text{isu}(p_1, t_1)$ are causally dependent on $\text{isu}(p, t)$. From the definition of t_1 in Lemma 6.3, we already have that $(\text{isu}(p, t), \text{isu}(p_1, t_1)) \in \text{CO}$. In the proof, we assume w.l.o.g that in $\beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2$ there is no event a that reads a value that t overwrites, otherwise, we can shortcut the trace by removing (p', t') and instead using the conflict relation between a and a store event of t to build the transactional happens-before cycle. Now, assume that β_1 contains a delayed transaction t_0 from another process $q \neq p$ that is not causally dependent on $\text{isu}(p, t)$. We show that we either can obtain a contradiction or we can remove all events of t_0 and obtain a new violation trace τ' that has the same number of delays as τ . We have

three possible cases based on whether the delayed store event $\text{del}(q', t_0)$ of t_0 occurs in β_1 , β_2 or γ_S . Hence, we get that τ can be one of the following:

- (a) $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_{11} \cdot \text{isu}(q, t_0) \cdot \beta_{12} \cdot \text{del}(q', t_0) \cdot \beta_{13} \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$
- (b) $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_{11} \cdot \text{isu}(q, t_0) \cdot \beta_{12} \cdot \text{isu}(p_1, t_1) \cdot \beta_{21} \cdot \text{del}(q', t_0) \cdot \beta_{22} \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$
- (c) $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_{11} \cdot \text{isu}(q, t_0) \cdot \beta_{12} \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S^1 \cdot \text{del}(q', t_0) \cdot \gamma_S^2$

In case (a) (resp., (b)) we can notice that since $\text{isu}(q, t_0)$ happens-before $\text{del}(q', t_0)$ through β_{12} (resp., $\beta_{12} \cdot \text{isu}(p_1, t_1) \cdot \beta_{21}$) then after executing $\text{del}(q', t_0)$ we obtain a cycle in the transactional happens-before. Thus, we can remove (p', t') from both traces and still obtain a robustness violation. Let τ' be the resulting trace. τ' has the same number of delays as τ . In τ' , we do not have events that read values that t overwrites. Therefore, we do not need to delay the transaction t to ensure that the trace is a violation. Let τ'' be the resulting trace where the transaction t executes atomically. In τ'' , the transaction t was not delayed after the issue event of t_1 which means that τ'' has less delays than τ . This contradicts the fact that τ is a minimal violation.

Case (c): we assume that $\text{del}(q', t_0)$ happens-before after (p', t') , otherwise, we can reorder it before (p', t') and get case (b). Since γ_S^1 contains only store events, then by the happens-before definition, $\text{del}(q', t_0)$ must be a store event executed by p' which means that $q' = p'$. Let e_1 and e_2 be the read/write actions t' that are the source of the conflict between (p', t') and $\text{del}(p', t)$ and the happens-before between (p', t') and $\text{del}(p', t_0)$, respectively. Similar to before we assume w.l.o.g that there is no event in $\beta_{12} \cdot \text{isu}(p_1, t_1) \cdot \beta_2$ that reads a value that t_0 overwrites. We consider the two cases: i) e_2 occurs before e_1 in t' or the two coincide, and ii) e_1 occurs before e_2 in t' . In the first case we can obtain a new violation where we do not delay the transaction t which will not affect the action e_2 that is the source of the happens-before between (p', t') and $\text{del}(p', t_0)$ (since e_1 occurs after e_2 then it cannot disable it). The new trace τ' is a violation since the store event $\text{del}(p', t_0)$ is delayed. Also, since the store event $\text{del}(p', t)$ of t was not delayed after $\text{isu}(p_1, t_1)$ then τ' has less delays than τ , which contradicts the fact that τ is a minimal violation. In the second case, if in $\beta_{12} \cdot \text{isu}(p_1, t_1) \cdot \beta_2$ we do not have any event that is causally dependent on $\text{isu}(q, t_0)$ other than the store events of t_0 , then we can remove all events of t_0 from τ without affecting the happens before between $\text{isu}(p, t)$ and $\text{del}(p', t)$ through $\beta_{12} \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t')$. Let $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1' \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S'$ be the resulting trace which has the same number of delays as τ . Otherwise, if in $\beta_{12} \cdot \text{isu}(p_1, t_1) \cdot \beta_2$ we have an event a that is causally dependent on $\text{isu}(q, t_0)$ that is not a store event of t_0 , then the new trace t' resulting from not delaying t_0 is a violation. This is because the store event $\text{del}(p', t)$ is delayed. τ' has less delays than τ since the store event $\text{del}(p', t_0)$ of t_0 was not delayed after a . This contradicts the fact that τ is a minimal violation.

Now, we show that for every two successive issue events of delayed transactions in τ' , we can deliver the first to the process of the second before the second is issued. Let $ev_i = \text{isu}(p_i, t_i)$ and $ev_j = \text{isu}(p_j, t_j)$ be two successive issue events of delayed transactions in $\beta_1 \cdot \text{isu}(p_1, t_1)$ s.t. either $(ev_i, ev_j) \in \text{HB}$ or $(ev_i, \text{del}(p_j, t_j)) \in \text{HB}$. Note that the only case where the store event $\text{del}(p_j, t_j)$ cannot be moved to occur before ev_j in β_1 is when the two events are related by a happens-before relation, i.e., $(ev_j, \text{del}(p_j, t_j)) \in \text{HB}$. In this case, we get that the transactions t_i and t_j are involved in a cycle in the transactional happens-before in τ' which means that $\tau'' = \alpha_A \cdot (p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \gamma_S$ is a violation which has less delays than τ (since t was not delayed after $\text{isu}(p_1, t_1)$). Therefore, the trace τ'' where the store event $\text{del}(p_j, t_j)$ occurs before ev_2 is happens-before equivalent to τ' . Similarly, when the

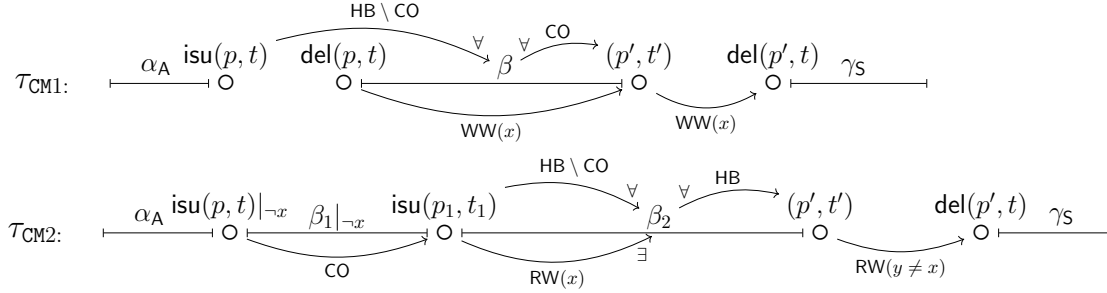


FIGURE 13. Robustness violation patterns under CM. For violation τ_{CM1} , t is the only delayed transaction. For τ_{CM2} , all delayed transactions are in $\text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1)$ and they form a causality chain that starts at $\text{isu}(p, t)$ and ends at $\text{isu}(p_1, t_1)$.

two events are concurrent, the trace τ'' where the store event $\text{del}(p_j, t_i)$ occurs before ev_j is happens-before equivalent to τ' . Thus, given the sequence of issue events $ev_1 \cdot ev_2 \cdot \dots \cdot ev_n$ of delayed transactions in τ' s.t. $ev_1 = \text{isu}(p, t)$ and $ev_n = \text{isu}(p_1, t_1)$, the trace τ'' where for every $1 \leq k \leq n-1$ s.t. $ev_k = \text{isu}(p_k, t_k)$ and $ev_{k+1} = \text{isu}(p_{k+1}, t_{k+1})$, we have the store event $\text{del}(p_{k+1}, t_k)$ occurs before the issue event ev_{k+1} is happens-before equivalent to τ' . Also, in τ'' for every $2 \leq k \leq n$, we have that ev_k is causally dependent on $ev_1 = \text{isu}(p, t)$. Thus, in τ'' the sequence of issue events $ev_1 \cdot ev_2 \cdot \dots \cdot ev_n$ of delayed transactions forms a causality chain. \square

Next, we show that for minimal violations in $\mathbb{T}\text{ccv2}$ where t and t_1 are distinct transactions, all delayed transactions in $\text{isu}(p, t) \cdot \beta_1$ do not access the shared variable x that starts the happens-before path in β_2 (Lemma 6.3) between $\text{isu}(p_1, t_1)$ and (p', t') . If this were not the case, then the events of t_1 can be removed and we still guarantee a happens-before path to $\text{del}(p', t)$ (starting in the delayed transaction accessing the variable x), thus obtaining a new robustness violation trace with less delays (since $\text{del}(p', t)$ was not delayed after $\text{isu}(p_1, t_1)$), which contradicts the minimality assumption.

Lemma 6.6. *Let τ be a minimal violation in $\mathbb{T}\text{ccv2}$ where t_1 and t are two distinct transactions. Then, all the delayed transactions in $\text{isu}(p, t) \cdot \beta_1$ do not access the variable x from Lemma 6.3.*

Proof. Suppose by contradiction that we have an issue event $\text{isu}(p_2, t_2)$ in $\text{isu}(p, t) \cdot \beta_1$ (i.e., $\text{isu}(p, t) \cdot \beta_1 = \text{isu}(p, t) \cdot \beta_{11} \cdot \text{isu}(p_2, t_2) \cdot \beta_{12}$) which accesses the shared variable x with either a read or a write instruction. Then, since there exists an event $a \in \beta_2$ s.t. $(\text{isu}(p_1, t_1), a) \in \text{WW}(x) \cup (\text{STO}; \text{RW}(x))$, we have that $(\text{isu}(p_2, t_2), a) \in \text{WW}(x) \cup (\text{STO}; \text{RW}(x))$. Moreover, because $\beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ does not contain any transaction that causally depends on $\text{isu}(p_1, t_1)$, we get that $\text{isu}(p_1, t_1)$ is the issue event by the process p_1 and we can remove it together with all the related stores in γ_S to obtain: $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_{11} \cdot \text{isu}(p_2, t_2) \cdot \beta_{12} \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma'_S$ which is a violation because $\text{isu}(p, t)$ happens-before $\text{del}(p', t)$ through $\beta_{11} \cdot \text{isu}(p_2, t_2) \cdot \beta_{12} \cdot \beta_2 \cdot (p', t')$. Furthermore, τ' has less delays than τ since $\text{del}(p', t)$ was not delayed after $\text{isu}(p_1, t_1)$. This contradicts the fact that τ is a minimal violation. \square

7. ROBUSTNESS VIOLATIONS UNDER CAUSAL MEMORY

The characterization of robustness violations under CM is at some level similar to that of robustness violations under CCv . However, some instance of the violation pattern under CCv is not possible under CM and CM admits some class of violations that is not possible under CCv . This reflects the fact that these consistency models are incomparable in general.

The following theorem gives the characterization of minimal violations under CM which is pictured in Fig. 13. Roughly, a program is not robust iff it admits a violation that either contains two concurrent transactions that write to the same variable, or it is a restriction of the pattern $\tau_{\text{CCv}2}$ admitted by CCv where the last delayed transaction is related only by RW to future transactions. The first pattern is not admitted by CCv because the writes to each variable are executed according to the timestamp order (CM does not satisfy the CCv property stated in Lemma 2.2).

Theorem 7.1. *A program \mathcal{P} is not robust under CM iff there exists a minimal violation in $\text{Tr}(\mathcal{P})_{\text{CM}}$ of one of the following forms:*

- (1) $\tau_{\text{CM}1} = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$, where:
 - (a) $\text{isu}(p, t)$ is the issue of the first and only delayed transaction;
 - (b) $\exists y. \text{ s.t. } (\text{del}(p, t), (p', t')) \in \text{WW}(y) \text{ and } ((p', t'), \text{del}(p', t)) \in \text{WW}(y)$ (Lemma 7.3);
 - (c) $\forall a \in \beta. (\text{isu}(p, t), a) \in \text{HB} \setminus \text{CO}$ and $(a, (p', t')) \in \text{CO}$ (Lemma 7.3).
- (2) $\tau_{\text{CM}2} = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$, where
 - (a) $\text{isu}(p, t)$ and $\text{isu}(p_1, t_1)$ are the issues of the first and last delayed transactions (Lemma 7.3);
 - (b) the issues of all delayed transactions are in β_1 are included in a causality chain that starts with $\text{isu}(p, t)$ and ends with $\text{isu}(p_1, t_1)$;
 - (c) for every $a \in \beta_2$, we have that $(\text{isu}(p_1, t_1), a) \in \text{HB} \setminus \text{CO}$ and $(a, (p', t')) \in \text{HB}$ (Lemma 7.3);
 - (d) there exist $a \in \beta_2 \cdot (p', t')$, x , and y s.t. $x \neq y$, $(\text{isu}(p_1, t_1), a) \in \text{RW}(x)$, $(a, (p', t')) \in \text{HB}?$, and $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$ (Lemma 7.3);
 - (e) all delayed transactions in $\text{isu}(p, t) \cdot \beta_1$ do not access the variable x .

The violation pattern $\tau_{\text{CM}2}$ is a restriction of the pattern $\tau_{\text{CCv}2}$ under CCv . For instance, the trace in Fig.12b is a valid minimal violation of the SB program under CM . The violation pattern $\tau_{\text{CM}1}$ implies the existence of a write-write race under CM . Fig.14 shows a minimal violation under CM that corresponds to a write-write race in the LU program. Conversely, if a program \mathcal{P} admits a trace τ which contains a write-write race under CM , then \mathcal{P} also admits a trace τ' where the two transactions t_1 and t_2

that caused the write-write race form a cycle in the store order (the store events of t_1 and t_2 on the two processes p_1 and p_2 that issued them can be reordered to occur in opposite orders, i.e., $\text{del}(p_1, t_1)$ before $\text{del}(p_1, t_2)$ and $\text{del}(p_2, t_2)$ before $\text{del}(p_2, t_1)$, which implies that are also in opposite orders w.r.t. the store order). Thus, \mathcal{P} has a trace τ' with a cycle in the transactional happens-before which means that \mathcal{P} is not robust against CM . Therefore, a program which is robust against CM is also write-write race free under CM . Since without write-write races, the CM and the CCv semantics coincide, we get the following the result.

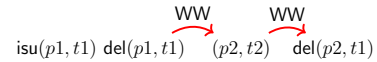


FIGURE 14. Violation of LU program in Fig. 9a. A $\tau_{\text{CM}1}$ violation where $\beta_2 = \gamma_S = \epsilon$, and t and t' correspond to t_1 and t_2 .

Lemma 7.2. *If a program \mathcal{P} is robust against CM, then \mathcal{P} is robust against CCv.*

Next, we discuss the proof of Theorem 7.1. The following lemma reveals the two possible minimal violation patterns under causal memory. The characterization of the patterns in this lemma can be refined further using arguments similar to the case of CCv (see the discussion at the end of this section).

Lemma 7.3. *If \mathcal{P} is a program that is not robust under CM, then it must admit a minimal violation τ that satisfies one of the following:*

- (1) $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ where:
 - (a) $\exists y$. s.t. $(\text{del}(p, t), (p', t')) \in \text{WW}(y)$ and $((p', t'), \text{del}(p', t)) \in \text{WW}(y)$;
 - (b) $\forall a \in \beta$. $(\text{isu}(p, t), a) \in \text{HB} \setminus \text{CO}$ and $(a, (p', t')) \in \text{CO}$.
- (2) $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ where:
 - (a) $\text{isu}(p_1, t_1)$ is the last issue event from $\{c \in \beta \mid (\text{isu}(p, t), c) \in \text{CO}\}$ in τ ;
 - (b) there exist two variables $x \neq y$, a in $\beta_2 \cdot (p', t')$, and $b = (p', t')$ such that $(\text{isu}(p_1, t_1), a) \in \text{RW}(x)$, $(b, \text{del}(p', t)) \in \text{RW}(y)$, and $(a, b) \in \text{HB}?$;
 - (c) $\forall a \in \beta_2$. $(\text{isu}(p_1, t_1), a) \in \text{HB} \setminus \text{CO}$ and $(a, (p', t')) \in \text{HB}$.

Proof. The proof will contain many arguments which are similar to those used in the proof of Lemma 6.3. Let $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ be a minimal violation under CM (cf. Lemma 5.7). We prove that there exists a minimal violation trace τ' obtained from τ that satisfies (1) or (2). Similar to Lemma 6.3, we get that there must exist y s.t. $((p', t'), \text{del}(p', t)) \in \text{RW}(y) \cup \text{WW}(y)$.

We consider two cases: i) $((p', t'), \text{del}(p', t)) \in \text{WW}(y)$, and ii) $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$. If $((p', t'), \text{del}(p', t)) \in \text{WW}(y)$, then by reordering the store event $\text{del}(p, t) \in \gamma_S$ to occur just after the corresponding issue and removing all events in β (and all related stores in γ_S) that are not causally ordered before (p', t') (since they do not contribute to the transactional happens-before cycle) we obtain a trace $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta' \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma'_S$ that is also a minimal violation and where $(\text{del}(p, t), (p', t')) \in \text{WW}(y)$ and $((p', t'), \text{del}(p', t)) \in \text{WW}(y)$. The trace τ' satisfies the first case of the lemma.

Now assume that $((p', t'), \text{del}(p', t)) \in \text{RW}(y)$, and let $\sigma = \{c \in \beta \mid (\text{isu}(p, t), c) \in \text{CO}\}$. We consider the following three cases.

First, assume that σ is empty. As in the proof of Lemma 6.3, we obtain that there exist $a \in \beta \cdot (p', t')$ and x s.t. $(\text{isu}(p, t), a) \in \text{RW}(x)$ and $(a, (p', t')) \in \text{HB}?$. If $x = y$ then both t and the transaction t_2 by a process p_2 of the event a write to x . Similar to before we can reorder the store event $\text{del}(p, t) \in \gamma_S$ to occur just after the corresponding issue and remove all issue events in $\beta \cdot (p', t')$ that occur after the issue event of t_2 and all their related stores. Also, we remove all events in β that are not causally ordered before the issue event of t_2 . We obtain $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta'(p_2, t_2) \cdot \text{del}(p_2, t) \cdot \gamma'_S$. In τ' the events of t_2 are assembled together, $\text{del}(p_2, t) \in \gamma_S$ is reordered to occur just after (p_2, t_2) , and $(\text{del}(p, t), (p_2, t_2)) \in \text{WW}(y)$ and $((p_2, t_2), \text{del}(p_2, t)) \in \text{WW}(y)$. Thus, τ' is a minimal violation and it satisfies the first case of the lemma. If $x \neq y$ then we get the second case of the lemma.

Second, assume that σ is not empty and all the elements of σ are store events. As in the proof of Lemma 6.3, we obtain that there exist x and an event $a \in \beta \cdot (p', t')$ that is not a store event of t s.t. $(\text{isu}(p, t), a) \in (\text{STO}; \text{WW}(x)) \cup \text{RW}(x)$ and $(a, (p', t')) \in \text{HB}?$. If $(\text{isu}(p, t), a) \in (\text{STO}; \text{WW}(x))$ or $x = y$ then both t and the transaction t_2 by a process p_2

of the event a write to x . Using the same procedure as in the previous paragraph we can obtain $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta' \cdot (p_2, t_2) \cdot \text{del}(p_2, t) \cdot \gamma'_S$ that satisfies the first case of the lemma. Similarly, if $(\text{isu}(p, t), a) \in \text{RW}(x)$ and $x \neq y$ then we get the second case of the lemma.

Third, assume that σ is not empty and $\text{isu}(p_1, t_1)$ is the last issue event in σ , i.e., $\beta = \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t')$. As in the proof of Lemma 6.3, we obtain that there exist x and an event $a \in \beta_2 \cdot (p', t')$ that is not a store event of t_1 s.t. $(\text{isu}(p_1, t_1), a) \in (\text{STO}; \text{WW}(x)) \cup \text{RW}(x)$ and $(a, (p', t')) \in \text{HB}?$. If $x = y$ then both t and the transaction t_2 by a process p_2 of the event a write to x . Using the same procedure as before we can obtain a trace $\tau' = \alpha_A \cdot \text{isu}(p, t) \cdot \text{del}(p, t) \cdot \beta'_1 \cdot \beta'_2 \cdot (p_2, t_2) \cdot \text{del}(p_2, t) \cdot \gamma'_S$ that is a minimal violation. τ' has less delays than τ since the store of t was not delayed after $\text{isu}(p_1, t_1)$. This contradicts the fact that τ is a minimal violation. Assume now that $x \neq y$. We assume w.l.o.g. that all events in β_2 do not read values that any transaction with an issue event in $\text{isu}(p, t) \cdot \beta_{t_1} \cdot \text{isu}(p_1, t_1)$ overwrites. If $(\text{isu}(p_1, t_1), a) \in (\text{STO}; \text{WW}(x))$ and $a \neq (p', t')$ then we can remove all issue events in $\beta_2 \cdot (p', t')$ that occur after the issue event of t_2 including (p', t') and assemble together the events of t_2 . We obtain that $(\text{del}(p_1, t_1), (p_2, t_2)) \in \text{WW}(x)$ and $((p_2, t_2), \text{del}(p_2, t_1)) \in \text{WW}(x)$ where we do not need to delay the transaction t and obtain $\tau' = \alpha_A \cdot (p, t) \cdot \beta'_1 \cdot \text{isu}(p_1, t_1) \cdot \text{del}(p_1, t_1) \cdot \beta'_2 \cdot (p_2, t_2) \cdot \text{del}(p_2, t_1) \cdot \gamma'_S$ that is a violation and has less delays than τ . This contradicts the fact that τ is a minimal violation. If $(\text{isu}(p_1, t_1), a) \in (\text{STO}; \text{WW}(x))$ and $a = (p', t')$ (i.e., $t' = t_2$) then we construct τ' such that all transactions that have issue events in σ and t are executed atomically after all the events in $(\beta_1 \setminus \sigma) \cdot \beta_2 \cdot \text{isu}(p', t') \cdot \text{del}(p', t')$ are executed first, i.e., $\tau' = \alpha_A \cdot \beta_{t_1} \cdot \beta_2 \cdot \text{isu}(p', t') \cdot \text{del}(p', t') \cdot (p, t) \cdot \beta_{t_2} \cdot \beta' \cdot (p_1, t_1) \cdot \text{del}(p_1, t_1) \cdot \gamma'_S$. τ' is a robustness violation since $(\text{del}(p', t'), (p_1, t_1)) \in \text{WW}(x)$ and $((p_1, t_1), \text{del}(p_1, t_1)) \in \text{WW}(x)$. Also, τ' has less delays than τ since t' was not delayed after a causally dependent event other than its store events and t is no longer delayed after the issue event of t_1 . This contradicts the fact that τ is a minimal violation. Finally, the only remaining possibility is $(\text{isu}(p_1, t_1), a) \in \text{RW}(x)$ where $x \neq y$ which corresponds to the second case of the lemma. \square

We use $\mathbb{T}\text{cm1}$ and $\mathbb{T}\text{cm2}$ to denote the class of minimal violations that satisfy the first and second case in Lemma 7.3, respectively. To show that for a non robust program, we can always find a minimal violation in either $\mathbb{T}\text{cm1}$ or $\mathbb{T}\text{cm2}$ where β and β_2 do not contain delayed transactions we can use the same proof arguments as in Lemma 6.4. For minimal violations in $\mathbb{T}\text{cm2}$ where t and t_1 are distinct transactions, the two properties that issue events of all delayed transactions form a causality chain and that delayed transactions in $\text{isu}(p, t) \cdot \beta_1$ do not access the shared variable x can also be proved in the same manner as in Lemmas 6.5 and 6.6, respectively.

8. ROBUSTNESS VIOLATIONS UNDER WEAK CAUSAL CONSISTENCY

If a program is robust against CM , then it must not contain a write-write race under CM (note that this is not true for CCv). Therefore, by Theorem 3.2, a program which is robust against CM has the same set of traces under both CM and CC , which implies that it is also robust against CC . Conversely, since CC is weaker than CM (i.e., $\text{Tr}_{\text{CM}}(\mathcal{P}) \subseteq \text{Tr}_{\text{CC}}(\mathcal{P})$ for any \mathcal{P}), if a program is robust against CC then it is robust against CM . Thus, we obtain the following result.

Theorem 8.1. *A program \mathcal{P} is robust against CC iff it is robust against CM .*

9. REDUCTION TO SC REACHABILITY

We describe a reduction of robustness checking to a reachability problem in a program executing under the serializability semantics, which can be simulated on top of standard sequential consistency (SC) by considering that each transaction is an atomic section (guarded by a fixed global lock). Essentially, given a program \mathcal{P} and a semantics $X \in \{\text{CCv}, \text{CM}, \text{CC}\}$, we define an instrumentation of \mathcal{P} such that \mathcal{P} is not robust against X iff the instrumentation reaches an error state under the serializability semantics. The instrumentation uses auxiliary variables in order to simulate the robustness violations (in particular, the delayed transactions) satisfying the patterns given in Fig. 11 and Fig. 13. We will focus our presentation on the second violation pattern of CCv (which is similar to the second violation pattern of CM): $\tau_{\text{CCv}2} = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$.

The process p that delayed the first transaction t is called the *Attacker*. The other processes delaying transactions in $\beta_1 \cdot \text{isu}(p_1, t_1)$ are called *Visibility Helpers*. Recall that all the delayed transactions must be causally ordered after $\text{isu}(p, t)$. The processes that execute transactions in $\beta_2 \cdot (p', t')$ and contribute to the happens-before path between $\text{isu}(p_1, t_1)$ and $\text{del}(p', t)$ are called *Happens-Before Helpers*. A happens-before helper cannot be the attacker or a visibility helper since this would contradict the causal delivery guarantee provided by causal consistency (a transaction of a happens-before helper is not delayed, so visible immediately to all processes, and it cannot follow a delayed transaction). γ_S contains the stores of the delayed transactions in $\text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1)$. It is important to notice that we may have $t = t_1$. In this case, $\beta_1 = \epsilon$ and the only delayed transaction is t . Also, all delayed transactions in β_1 including t_1 may be issued by the same process as t . In all of these cases, the set of Visibility Helpers is empty.

The instrumentation uses two copies of the set of shared variables in the original program. We use primed variables x' to denote the second copy. When a process becomes the attacker or a visibility helper, it will write only to the second copy that is visible only to these processes (and remains invisible to the other processes including the happens-before helpers). The writes made by other processes including the happens-before helpers are made visible to all processes, i.e., they are applied on both copies of every shared variable.

To establish the causality chains of the delayed transactions issued by the attacker and the visibility helpers, we look whether a transaction can extend the causality chain started by the first delayed transaction issued by the attacker. This is to ensure that all such transactions are causally related to the first delayed transaction (of the attacker). In order for a transaction to “join” the causality chain, it has to satisfy one of the following conditions:

- the transaction is issued by a process that has already another transaction in the causality chain. Thus, we ensure the continuity of the causality chain through program order;
- the transaction is reading from a variable that was updated by a previous transaction in the causality chain. Hence, we ensure the continuity of the causality chain through the write-read relation.

We introduce a flag for each shared variable to mark the fact that it was updated by a previous transaction in the causality chain. These flags are used by the instrumentation to establish whether a transaction “joins” a causality chain. Enforcing a happens-before path starting in the last delayed transaction, using transactions of the happens-before helpers, can be done in the same way. Compared to causality chains, there are two more cases in which a transaction can extend a happens-before path:

- the transaction writes to a shared variable that was read by a previous transaction in the happens-before path. Hence, we ensure the continuity of the happens-before path through the read-write relation;
- the transaction writes to a shared variable that was updated by a previous transaction in the happens-before path. Hence, we ensure the continuity of the happens-before path through write-write order.

Thus, we extend the shared variables flags used for causality chains in order to record if a variable was read or written by a previous transaction (in this case, a previous transaction in the happens-before path). Overall, the instrumentation uses a flag $x.event$ or $x'.event$ for each (copy of a) shared variable, that stores the type of the last access (read or write) to the variable. Initially, these flags and other flags used by the instrumentation as explained below are initialized to null (\perp).

In general, whether a process is an attacker, visibility helper, or happens-before helper is not enforced syntactically by the instrumentation, and can vary from execution to execution. The role of a process in an execution is set *non-deterministically* during the execution using some additional process-local flags. Thus, during an execution, each process chooses to set to `true` at most one of the flags $p.a$, $p.vh$, and $p.hbh$, implying that the process becomes an attacker, visibility helper, or happens-before helper, respectively. At most one process can be an attacker, i.e., set $p.a$ to `true`.

9.1. Instrumentation of the Attacker. We provide in Fig. 15, the instrumentation of the instructions for the attacker process. Such a process passes through an initial phase where it executes transactions that are visible immediately to all the other processes (i.e., they are not delayed), and then non-deterministically it can choose to delay a transaction. When the attacker randomly chooses the first transaction to start delaying of transactions, it sets a *global* flag a_{tr_A} to `true` in the instruction `begin` (line (9.3)). Then, it sets the flag $p.a$ to `true` to indicate that the current process is the attacker. During the first delayed transaction, the attacker non-deterministically chooses a write instruction to a shared variable y and stores the name of this variable in the flag a_{st_A} (line (9.8)). The values written during delayed transactions are stored in the primed variables and are visible only to the attacker and the visibility helpers. For example, given a variable z , all the writes to z from the original program are transformed into writes to the primed version z' (line (9.6)). Each time the attacker writes to a variable z' , it sets the flag $z'.event$ to `st` (line (9.7)) which will allow other processes that read the same variable to join the set of visibility helpers and start delaying their transactions. Once the attacker delays a transaction, it will read only from the primed variables (i.e., z').

To start the happens-before path, the attacker has to execute a transaction that either reads or writes to a shared variable x that was not accessed by a delayed transaction (i.e., $x'.event = \perp$). In this case, it sets the variable `HB` to `true` (lines (9.4) and (9.9)) to mark the start of the happens before path and the end of the visibility chains, and it sets the flag $x.event$ to `ld` (lines (9.5) and (9.10)). We set $x.event$ to `ld` even in the case of a write to x in order to simplify the instrumentation of the happens-before helpers (to check that this transaction is related to a transaction of a happens-before helper p through $WW(x)$ or $RW(x)$ it is enough that p writes to x and it “observers” the same value `ld` in $x.event$). When the flag `HB` is set to `true` the attacker stops executing new transactions. We can

<pre> [[l1: begin; goto l2;]]_A^{P2} = // Typical execution of begin l1: assume HB = ⊥ ∧ (p.a ≠ ⊥ ∨ a_{tr_A} = ⊥); goto l_{x1};(9.1) l_{x1}: begin; goto l2; // Begin of first delayed transaction l1: assume HB = ⊥ ∧ a_{tr_A} = ⊥; goto l_{x2}; (9.2) l_{x2}: begin; goto l_{x3}; l_{x3}: p.a := true; goto l_{x4}; l_{x4}: Foreach x ∈ V. x' := x; goto l_{x5}; l_{x5}: a_{tr_A} := true; goto l2; (9.3) </pre>	<pre> [[l1: x := e; goto l2;]]_A^{P2} = // Write before delaying transactions l1: assume a_{tr_A} = ⊥; goto l_{x1}; l_{x1}: x := e; goto l2; // Write in delayed transactions l1: assume a_{tr_A} ≠ ⊥ ∧ p.a ≠ ⊥; goto l_{x2}; (9.6) l_{x2}: x' := e; goto l_{x3}; (9.7) l_{x3}: x'.event := st; goto l2; (9.7) // Special write in first delayed transaction l1: assume a_{st_A} = x.event = ⊥ ∧ p.a ≠ ⊥; goto l_{x4}; l_{x4}: x' := e; goto l_{x5}; (9.8) l_{x5}: a_{st_A} := 'x'; goto l_{x6}; (9.8) l_{x6}: x'.event := st; goto l2; // Special write in last delayed transaction l1: assume x'.event = ⊥ ∧ p.a ≠ ⊥; goto l_{x7}; l_{x7}: x' := e; goto l_{x8}; l_{x8}: HB := true; goto l_{x9}; (9.9) l_{x9}: x.event := ld; goto l2; (9.10) </pre>
<pre> [[l1: r := x; goto l2;]]_A^{P2} = // Read before delaying transactions l1: assume a_{tr_A} = ⊥; goto l_{x1}; l_{x1}: r := x; goto l2; // Read in delayed transactions l1: assume a_{tr_A} ≠ ⊥ ∧ p.a ≠ ⊥; goto l_{x2}; l_{x2}: r := x'; goto l_{x3}; l_{x3}: x'.event := ld; goto l2; // Special read in last delayed transaction l1: assume x'.event = ⊥ ∧ p.a ≠ ⊥; goto l_{x4}; l_{x4}: r := x'; goto l_{x5}; l_{x5}: HB := true; goto l_{x6}; (9.4) l_{x6}: x.event := ld; goto l2; (9.5) </pre>	<pre> [[l1: end; goto l2;]]_A^{P2} = l1: assume p.a ≠ ⊥ ∧ a_{st_A} = ⊥; assume false; l1: end; goto l2; </pre>

FIGURE 15. Instrumentation of the Attacker. We use ‘ x ’ to denote the name of the shared variable x .

notice that when the HB is set to true, we can no longer execute new transactions from the attacker (all conditions in lines (9.1) and (9.2) become false).

9.2. Instrumentation of the Visibility Helpers. Fig. 16 lists the instrumentation of the instructions of a process that belongs to the set of visibility helpers. Such a process passes through an initial phase where it executes the original code instructions (lines (9.18) and (9.13)) until the flag a_{tr_A} is set to true by the attacker. Then, it continues the execution of its original instructions but, whenever it stores a value it writes it to both the shared variable z and the primed variable z' so it is visible to all processes. Non deterministically it chooses a first transaction to delay, at which point it joins the set of visibility helpers. It sets the flag $p.vh$ to false signaling its desire to join the visibility helpers, and it chooses a transaction (the begin of this transaction is shown in line (9.12)) through which the process will join the set of visibility helpers. The process directly starts delaying its writes, i.e., writing to primed variables, and reading only from delayed writes, i.e., from primed variables, and behaving the same as the attacker. In order to check that it can extend the sequence of causal dependencies (required by the causal chain definition), it takes a snapshot of the $_.event$ fields at the beginning of the transaction and stores it to $_.event'$ fields (line l_{x4} in the instrumentation of begin). This snapshot is necessary to check that it reads from writes

<pre> [[l₁: begin; goto l₂];]]_{VH}^{P2} = // Before joining visibility helpers l₁: assume HB = ⊥ ∧ (a_{trA} = ⊥ ∨ p.vh = ⊥); goto l_{x1}; l_{x1}: begin; goto l₂; // Joining visibility helpers l₁: assume HB = p.vh = p.a = ⊥ ∧ a_{trA} ≠ ⊥; goto l_{x2}; l_{x2}: begin; goto l_{x3}; l_{x3}: p.vh := false; goto l_{x4}; l_{x4}: Foreach x' ∈ V. x'.event' := x'.event; goto l₂; // After joining visibility helpers l₁: assume HB = ⊥ ∧ a_{trA} ≠ ⊥ ∧ p.vh; goto l_{x5}; l_{x5}: begin; goto l₂; [[l₁: r := x; goto l₂];]]_{VH}^{P2} = // Before joining visibility helpers l₁: assume a_{trA} = ⊥ ∨ (p.vh = p.a = ⊥); goto l_{x1}; l_{x1}: r := x; goto l₂; // After joining visibility helpers l₁: assume p.vh ≠ ⊥; goto l_{x2}; l_{x2}: r := x'; goto l_{x3}; l_{x3}: assume x'.event' = st ∧ ¬p.vh; goto l_{x4}; l_{x4}: p.vh := true; goto l₂; l_{x3}: assume x'.event' ≠ st ∨ p.vh; goto l₂; // Last delayed transaction l₁: assume x'.event = ⊥ ∧ p.vh ≠ ⊥; goto l_{x5}; l_{x5}: HB := true; goto l_{x6}; l_{x6}: x.event := ld; goto l_{x7}; l_{x7}: r := x'; goto l₂; </pre>	<pre> [[l₁: x := e; goto l₂];]]_{VH}^{P2} = // Before attacker delays transactions l₁: assume a_{trA} = ⊥; goto l_{x1}; l_{x1}: x := e; goto l₂; // Before joining visibility helpers l₁: assume a_{trA} ≠ ⊥ ∧ p.vh = p.a = ⊥; goto l_{x2}; l_{x2}: x' := e; goto l_{x3}; l_{x3}: x := e; goto l₂; // After joining visibility helpers l₁: assume p.vh ≠ ⊥; goto l_{x4}; l_{x4}: x' := e; goto l_{x5}; l_{x5}: x'.event := st; goto l_{x6}; l_{x6}: x'.event' := ⊥; goto l₂; // Last delayed transaction l₁: assume x'.event = ⊥ ∧ p.vh ≠ ⊥; goto l_{x7}; l_{x7}: HB := true; goto l_{x8}; l_{x8}: x.event := ld; goto l_{x9}; l_{x9}: x' := e; goto l₂; [[l₁: end; goto l₂];]]_{VH}^{P2} = // Before joining visibility helpers l₁: assume a_{trA} = ⊥ ∨ (a_{trA} ≠ ⊥ ∧ p.vh = ⊥); goto l_{x1}; l_{x1}: end; goto l₂; // After joining visibility helpers l₁: assume a_{trA} ≠ ⊥ ∧ p.vh; goto l_{x2}; l_{x2}: end; goto l₂; // Failed to join visibility helpers l₁: assume a_{trA} ≠ ⊥ ∧ ¬p.vh; assume false; </pre>
--	---

FIGURE 16. Instrumentation of the Visibility Helpers.

made in other transactions (ignoring the writes in the current transaction). When a process chooses a first transaction to delay (during the **begin** instruction), it has made a pledge that during this transaction it will read from a variable that was updated by another delayed transaction from either the attacker or some other visibility helper. This is to ensure that this transaction extends the visibility chain. Hence, the local process flag $p.vh$ will be set to **true** when the process meets its pledge (line (9.15)). If the process does not keep its pledge (i.e., $p.vh$ is equal to **false**) at the end of the transaction (i.e., during the **end** instruction) we block the execution. Thus, when executing the **end** instruction of the underlying transaction we check whether the flag $p.vh$ is null, if so we block the execution (line (9.23)).

When a process joins the visibility helpers, it delays all writes and reads only from the primed variables (lines (9.19) and (9.14)). Similar to the attacker, a process in the visibility helpers delays a write to a shared variable z by writing to z' , it sets the flag $z'.event$ to **st** (line (9.20)). In order for a process in the visibility helpers to start the happens-before path, it has to either read or write a shared variable x that was not accessed by a delayed transaction (i.e., $x'.event = \perp$). In this case we set the flag **HB** to **true** (lines (9.21) and (9.16)) to mark the start of the happens before path and the end of the visibility chains

<pre> [[l₁: begin; goto l₂;]]_{HbH}^{P2} = // Before joining happens-before helpers l₁: assume HB = p.vh = p.a = ⊥ ; goto l_{x1}; l_{x1}: begin; goto l₂; // Joining happens-before helpers l₁: assume HB ≠ ⊥ ∧ p.hbh = p.vh = p.a = ⊥ ; goto l_{x2}; l_{x2}: begin; goto l_{x3}; (9.24) l_{x3}: Foreach x ∈ V. x.event' := x.event; goto l₂; // After joining happens-before helpers l₁: assume HB ≠ ⊥ ∧ p.hbh ≠ ⊥ ; goto l_{x4}; l_{x4}: begin; goto l₂; [[l₁: x := e; goto l₂;]]_{HbH}^{P2} = // Before the first delayed transaction l₁: assume HB = ⊥ ∧ a_{trA} = ⊥ ; goto l_{x1}; l_{x1}: x := e; goto l₂; // After the first delayed transaction l₁: assume HB = p.vh = p.a = ⊥ ∧ a_{trA} ≠ ⊥ ; goto l_{x2}; l_{x2}: x' := e; goto l_{x3}; (9.25) l_{x3}: x := e; goto l₂; (9.26) // After the last delayed transaction l₁: assume HB ≠ ⊥ ∧ p.vh = p.a = ⊥ ; goto l_{x4}; l_{x4}: x := e; goto l_{x5}; l_{x5}: x.event := st; goto l_{x6}; (9.27) l_{x6}: assume x.event' ≠ ⊥ ∧ p.hbh = ⊥ ; goto l_{x7}; l_{x7}: p.hbh := true; goto l₂; (9.28) l_{x6}: assume x.event' = ⊥ ∨ p.hbh ≠ ⊥ ; goto l₂; </pre>	<pre> [[l₁: r := x; goto l₂;]]_{HbH}^{P2} = // Before the last delayed transaction l₁: assume HB = ⊥ ∧ p.vh = p.a = ⊥ ; goto l_{x1}; l_{x1}: r := x; goto l₂; (9.29) // After the last delayed transaction l₁: assume HB ≠ ⊥ ∧ p.vh = p.a = ⊥ ; goto l_{x2}; l_{x2}: r := x; goto l_{x3}; l_{x3}: assume x.event' = st ∧ p.hbh = ⊥ ; goto l_{x4}; l_{x4}: p.hbh := true; goto l₂; (9.30) l_{x3}: assume x.event = ⊥ ; goto l_{x5}; l_{x5}: x.event := ld; goto l₂; (9.31) l_{x3}: assume x.event ≠ ⊥ ∨ p.hbh ≠ ⊥ ; goto l₂; [[l₁: end; goto l₂;]]_{HbH}^{P2} = // Before joining happens-before helpers l₁: assume HB = p.vh = p.a = ⊥ ; goto l_{x1}; l_{x1}: end; goto l₂; // After joining happens-before helpers l₁: assume HB ≠ ⊥ ∧ p.hbh ≠ ⊥ ; goto l_{x2}; l_{x2}: end; goto l_{x3}; l_{x3}: r̃ := a_{stA}; goto l_{x4}; (9.32) l_{x4}: r̃ := r̃.event; goto l_{x5}; (9.33) l_{x5}: assume r̃ ≠ ⊥ ; assert false; (9.34) l_{x5}: assume r̃ = ⊥ ; goto l₂; // Failed to join happens-before helpers l₁: assume HB ≠ ⊥ ∧ p.hbh = p.vh = p.a = ⊥ ; assume false; (9.35) </pre>
--	---

FIGURE 17. Instrumentation of Happens-Before Helpers.

and set the flag $x.event$ to `ld` (lines (9.22) and (9.17)). When the flag `HB` is set to `true`, all processes in the set of visibility helpers stop issuing new transactions because all conditions for executing the `begin` instruction become false.

9.3. Instrumentation of the Happens-Before Helpers. The remaining processes, which are not the attacker or a visibility helper, can become happens-before helpers. Fig. 17 lists the instrumentation of the instructions of a happens-before helper process. Similar to above, when the flag a_{trA} is set to `true` by the attacker, other processes enter a phase where they continue executing their instructions, however, when they store a value they write it in both the shared variable z and the primed variable z' (lines (9.25) and (9.26)). However, they only read from the original shared variables (line (9.29)). Once the flag `HB` is set to `true`, a process that cannot be the attacker (i.e., the flag $p.a$ is null) or a visibility helper (i.e., the flag $p.vh$ is null) chooses non-deterministically a transaction t (the `begin` of this transaction is shown in line (9.24)) through which it wants to join the set of happens-before helpers, i.e., continue the happens-before path created by the existing happens-before helpers. Similar to visibility helpers, when a process chooses the transaction t , it makes a pledge (while executing

the `begin` instruction) that during this transaction it will either read a variable updated by another happens-before helper or write to a variable that was accessed (read or written) by another happens-before helper (every process that executes a transaction after `HB` is set to `true` makes this pledge). When the pledge is met, the process sets the flag $p.hbh$ to `true` (lines (9.30) and (9.28)). The execution is blocked if a process does not keep its pledge (i.e., the flag $p.hbh$ is null) at the end of the transaction (line (9.35)). We use a flag $x.event$ for each variable x to record the type (read `ld` or write `st`) of the last access made by a happens-before helper (lines (9.31) and (9.27)). Moreover, once `HB` is set to `true` (i.e., there are no more delayed transactions), the process can write and read only the original shared variables, since the primed versions are no longer in use. A particular case is when the transaction t is from the first process trying to join the happens-before helpers, in which the transaction must contain a read accessing the variable x that was read or written to by a transaction from the attacker of a visibility helper.

The happens-before helpers continue executing their instructions, until one of them reads from the shared variable y whose name was stored in a_{st_A} . This establishes a happens-before path between the last delayed transaction and a “fictitious” store event corresponding to the first delayed transaction that could be executed just after this read of y . The execution does not have to contain this store event explicitly since it is always enabled. Therefore, at the end of every transaction, the instrumentation checks whether the transaction read y . If it is the case, then the execution stops and goes to an error state to indicate that this is a robustness violation. The happens-before helpers processes continue executing their instructions, until one of them executes a load that reads from the shared variable y that was stored in a_{st_A} which implies the existence of a happens-before cycle. Thus, when executing the instruction `end` at the end of every transaction, we have a conditional check to detect if we have a load or a write accessing the variable y (lines (9.32), (9.33), and (9.34)). When the check detects that the variable y was accessed, the execution goes to an error state (line (9.34)) to indicate that it has produced a robustness violation.

In Fig. 18, we show an excerpt of the instrumentations of the two transactions of the `SB` program. In particular, we only give the instructions of the instrumented `SB` that are reached during the execution that leads to an error state. The attacker instrumentation is applied to the transaction $t1$ of $p1$ and the happens-before helpers instrumentation is applied to the transaction $t2$ of $p2$. The first conflict order from $t1$ to $t2$ (shown in Fig. 11) is simulated by the fact that at line 9.39, $y.event' = ld$ (see lines 9.37 and 9.38). Also, the second conflict order from $t2$ to $t1$ is simulated by the fact that at line 9.41 we reach the error state where $a_{st_A}.event = x.event = ld$ (see lines 9.36 and 9.40).

9.4. Correctness. As we have already mentioned, the role of a process in an execution is chosen non-deterministically at runtime. Therefore, the final instrumentation of a given program \mathcal{P} , denoted by $\llbracket \mathcal{P} \rrbracket^{P2}$, is obtained by replacing each labeled instruction $\langle linst \rangle$ with the concatenation of the instrumentations corresponding to the attacker, the visibility helpers, and the happens-before helpers, i.e., $\llbracket \langle linst \rangle \rrbracket^{P2} ::= \llbracket \langle linst \rangle \rrbracket_A^{P2} \llbracket \langle linst \rangle \rrbracket_{VH}^{P2} \llbracket \langle linst \rangle \rrbracket_{HBH}^{P2}$. The instrumented program $\llbracket \mathcal{P} \rrbracket^{P2}$ reaches the error state iff \mathcal{P} admits a violation of the pattern τ_{CCv2} . Let $\llbracket \mathcal{P} \rrbracket^{P1}$ be the instrumented program that reaches an error state iff \mathcal{P} admits a violation of the pattern τ_{CCv1} . The instrumentation $\llbracket \rrbracket^{P1}$ does not include the visibility helpers since only a single transaction is delayed in τ_{CCv1} , and it can be obtained in the same manner as $\llbracket \rrbracket^{P2}$. The following theorem states the correctness of the instrumentation.

$$\begin{array}{l}
\llbracket l_1: \text{begin}; \text{goto } l_2; \rrbracket_A^{P2} = \\
l_1: \text{assume } \text{HB} = \perp \wedge a_{\text{tr}_A} = \perp; \text{goto } l_{b2}; \\
l_{b2}: \text{begin}; \text{goto } l_{b3}; \\
l_{b3}: p1.a := \text{true}; \text{goto } l_{b4}; \\
l_{b4}: \text{Foreach } z \in \mathbb{V}. z' := z; \text{goto } l_{b5}; \\
l_{b5}: a_{\text{tr}_A} := \text{true}; \text{goto } l_2; \\
\llbracket l_2: x := 1; \text{goto } l_3; \rrbracket_A^{P2} = \\
l_2: \text{assume } a_{\text{st}_A} = x.\text{event} = \perp \wedge p1.a \neq \perp; \text{goto } l_{s4}; \\
l_{s4}: x' := 1; \text{goto } l_{s5}; \\
l_{s5}: a_{\text{st}_A} := x'; \text{goto } l_{s6}; \\
l_{s6}: x'.\text{event} := \text{st}; \text{goto } l_3; \\
\llbracket l_3: r1 := y; \text{goto } l_4; \rrbracket_A^{P2} = \\
l_3: \text{assume } y'.\text{event} = \perp \wedge a_{\text{tr}_A} \neq \perp; \text{goto } l_{l4}; \\
l_{l4}: r1 := y'; \text{goto } l_{l5}; \\
l_{l5}: \text{HB} := \text{true}; \text{goto } l_{l6}; \\
l_{l6}: y.\text{event} := \text{ld}; \text{goto } l_4; \\
\llbracket l_4: \text{end}; \text{goto } l_5; \rrbracket_A^{P2} = \\
l_4: \text{end}; \text{goto } l_5;
\end{array}$$

$$\begin{array}{l}
\llbracket l_1: \text{begin}; \text{goto } l_2; \rrbracket_{\text{HbH}}^{P2} = \\
l_1: \text{assume } \text{HB} \neq \perp \wedge p2.\text{hbh} = p2.\text{vh} = p2.a = \perp; \text{goto } l_{b2}; \\
l_{b2}: \text{begin}; \text{goto } l_{b3}; \\
l_{b3}: x.\text{event}' := x.\text{event}; y.\text{event}' := y.\text{event}; \text{goto } l_2; \quad (9.38) \\
\llbracket l_3: y := 1; \text{goto } l_4; \rrbracket_{\text{HbH}}^{P2} = \\
l_3: \text{assume } \text{HB} \neq \perp \wedge p2.\text{vh} = p2.a = \perp; \text{goto } l_{s4}; \\
l_{s4}: y := 1; \text{goto } l_{s5}; \\
l_{s5}: y.\text{event} := \text{st}; \text{goto } l_{s6}; \\
l_{s6}: \text{assume } y.\text{event}' \neq \perp \wedge p2.\text{hbh} = \perp; \text{goto } l_{s7}; \quad (9.39) \\
l_{s7}: p2.\text{hbh} := \text{true}; \text{goto } l_4; \\
(9.36) \llbracket l_2: r2 := x; \text{goto } l_3; \rrbracket_{\text{HbH}}^{P2} = \\
l_2: \text{assume } \text{HB} \neq \perp \wedge p.\text{vh} = p.a = \perp; \text{goto } l_{l2}; \\
l_{l2}: r2 := x; \text{goto } l_{l3}; \\
l_{l3}: \text{assume } x.\text{event} = \perp; \text{goto } l_{l5}; \\
l_{l5}: x.\text{event} := \text{ld}; \text{goto } l_3; \quad (9.40) \\
\llbracket l_4: \text{end}; \text{goto } l_5; \rrbracket_{\text{HbH}}^{P2} = \\
(9.37) l_4: \text{assume } \text{HB} \neq \perp \wedge p2.\text{hbh} \neq \perp; \text{goto } l_{e2}; \\
l_{e2}: \text{end}; \text{goto } l_{e3}; \\
l_{e3}: \tilde{r} := a_{\text{st}_A}; \text{goto } l_{e4}; \\
l_{e4}: \tilde{r} := \tilde{r}.\text{event}; \text{goto } l_{e5}; \\
l_{e5}: \text{assume } \tilde{r} \neq \perp; \text{assert false}; \quad (9.41)
\end{array}$$

FIGURE 18. Instrumentation of SB program in Fig. 9b.

Theorem 9.1. *A program \mathcal{P} is not robust against CCv iff either $\llbracket \mathcal{P} \rrbracket^{P1}$ or $\llbracket \mathcal{P} \rrbracket^{P2}$ reaches the error state.*

The proof of this theorem relies on the explanations given above. One can define a bijection between executions of the instrumentation that reach an error state and executions of the original program that satisfy the constraints in one of the two violation patterns. The former can be rewritten to the latter by roughly, removing all accesses to the auxiliary variables used by the instrumentation, replacing the writes to shared variable copies by writes to the original variables, delivering delayed transactions only to visibility helpers, and appending store events for all the delayed transactions. For the reverse, given a robustness violation $\tau = \alpha_A \cdot \text{isu}(p, t) \cdot \beta_1 \cdot \text{isu}(p_1, t_1) \cdot \beta_2 \cdot (p', t') \cdot \text{del}(p', t) \cdot \gamma_S$ of type τ_{CCv2} , we can build an execution of the instrumentation that reaches an error state, where p is the attacker, the processes delaying transactions in $\beta_1 \cdot \text{isu}(p_1, t_1)$ are visibility helpers, and the processes that issue transactions between $\text{isu}(p_1, t_1)$ and $\text{del}(p', t)$ and that are part of the happens-before path between these two events are the happens-before helpers.

The following result states the complexity of checking robustness for finite-state programs⁹ against one of the three variations of causal consistency considered in this work (we use causal consistency as a generic name to refer to all of them). The upper bound is a direct consequence of Theorem 9.1 and of previous results concerning the reachability problem in concurrent programs running over SC, with a fixed [26] or parametric number of

⁹That is, programs where the number of variables and the data domain are bounded.

processes [40]. For the lower bound, given an instance of the reachability problem under sequential consistency, denoted by (\mathcal{P}, ℓ) ¹⁰, we construct a program \mathcal{P}' where each statement s of \mathcal{P} is a different transaction (guarded by a global lock), and where reaching the location ℓ enables the execution of a “gadget” that corresponds to the SB program in Figure 9b. Executing each statement under a global lock ensures that every execution of \mathcal{P}' under causal consistency is serializable, and faithfully represents an execution of the original \mathcal{P} under sequential consistency. Moreover, \mathcal{P} reaches ℓ iff \mathcal{P}' contains a robustness violation, which is due to the execution of SB.

Corollary 9.2. *Checking robustness of finite-state programs against causal consistency is PSPACE-complete when the number of processes is fixed and EXPSPACE-complete, otherwise.*

Remark 9.3. The reduction to reachability does not manipulate transaction identifiers and it is insensitive to the number of transactions executed by one process. Thus, all our results extend to processes that include unbounded loops of transactions. This includes programs where each process can call a statically known set of transactions (with parameters) an arbitrary number of times.

10. RELATED WORK

Causal consistency is one of the oldest consistency models for distributed systems [30]. Formal definitions of several variants of causal consistency, suitable for different types of applications, have been introduced recently [18, 17, 38, 13]. The definitions in this paper are inspired from these works and coincide with those given in [13]. In that paper, the authors address the decidability and the complexity of verifying that an implementation of a storage system is causally consistent (i.e., all its computations, for every client, are causally consistent).

While our paper focuses on *trace-based* robustness, *state-based robustness* requires that a program is robust if the set of all its reachable states under the weak semantics is the same as its set of reachable states under the strong semantics. While state-robustness is the necessary and sufficient concept for preserving state-invariants, its verification, which amounts in computing the set of reachable states under the weak semantics, is in general a hard problem. The decidability and the complexity of this problem has been investigated in the context of relaxed memory models such as TSO and Power, and it has been shown that it is either decidable but highly complex (non-primitive recursive), or undecidable [8, 9]. As far as we know, the decidability and complexity of this problem has not been investigated for causal consistency. Automatic procedures for approximate reachability/invariant checking have been proposed using either abstractions or bounded analyses, e.g., [10, 5, 20, 1]. Proof methods have also been developed for verifying invariants in the context of weakly consistent models such as [29, 25, 36, 4]. These methods, however, do not provide decision procedures.

Decidability and complexity of trace-based robustness has been investigated for the TSO and Power memory models [14, 12, 21]. The work we present in this paper borrows the idea of using minimal violation characterizations for building an instrumentation allowing to obtain a reduction of the robustness checking problem to the reachability checking problem over SC. However, applying this approach to the case of causal consistency is not straightforward and

¹⁰That is, whether the program \mathcal{P} reaches the control location ℓ under SC.

requires different proof techniques. Dealing with causal consistency is far more tricky and difficult than dealing with TSO, and requires coming up with radically different arguments and proofs, for (1) characterizing in a finite manner the set of violations, (2) showing that this characterization is sound and complete, and (3) using effectively this characterization in the definition of the reduction to the reachability problem.

As far as we know, our work is the first one that establishes results on the decidability and complexity issues of the robustness problem in the context of causal consistency, and taking into account transactions. The existing work on the verification of robustness for distributed systems consider essentially trace-based concepts of robustness and provide either over- or under-approximate analyses for checking it. In [11, 15, 16, 19], static analysis techniques are proposed based on computing an abstraction of the set of computations that is used in searching for robustness violations. These approaches may return false alarms due to the abstractions they consider. In particular, [11] shows that a trace under causal convergence is not admitted by the serializability semantics iff it contains a (transactional) happens-before cycle with a RW dependency, and another RW or WW dependency. This characterization alone is not sufficient to prove our result concerning robustness checking. Our result relies on a characterization of more refined robustness violations and relies on different proof arguments. In [35] a sound (but not complete) bounded analysis for detecting robustness violation is proposed. Our approach is technically different, is precise, and provides a decision procedure for checking robustness when the program is finite-state.

11. CONCLUSION

We have studied three variations of transactional causal consistency, showing that they are equivalent for programs without write-write races. We have shown that the problem of verifying that a transactional program is robust against causal consistency can be reduced, modulo a linear-size instrumentation, to a reachability problem in a transactional program running over a sequentially consistent shared memory. This reduction leads to the first decidability result concerning the problem of checking robustness against a weak transactional consistency model. Furthermore, this reduction opens the door to the use of existing methods and tools for the analysis and verification of SC concurrent programs, in order to reason about weakly-consistent transactional programs. It can be used for the design of a large spectrum of static/dynamic tools for testing/verifying robustness against causal consistency.

Our notion of robustness relies on a particular interpretation of behaviors as traces recording all happens-before dependencies. This is stronger than a more immediate notion of *state-based* robustness that requires equality of sets of reachable states, which means that it could produce false alarms, i.e., robustness violations that are not also violations of the intended program specification. This trade-off is similar in spirit to data races being used as an approximation of concurrency errors (since data races are easier to detect, compared to violations of arbitrary specifications).

An interesting direction for future work is looking at the robustness problem in the context of *hybrid* consistency models where some of the transactions in the program can be declared serializable. These models include synchronization primitives similar to lock acquire/release which allow to enforce a serialization order between some transactions. Such mechanisms can be used as a “repair” mechanism in order to make programs robust.

REFERENCES

- [1] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Ahmed Bouajjani, and Tuan Phong Ngo. Context-bounded analysis for POWER. In Axel Legay and Tiziana Margaria, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 23rd International Conference, TACAS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Part II*, volume 10206 of *Lecture Notes in Computer Science*, pages 56–74, 2017.
- [2] Atul Adya. *Weak consistency: A generalized theory and optimistic implementations for distributed transactions*. PhD thesis, 1999.
- [3] Mustaque Ahamad, Gil Neiger, James E. Burns, Prince Kohli, and Phillip W. Hutto. Causal memory: Definitions, implementation, and programming. *Distributed Comput.*, 9(1):37–49, 1995.
- [4] Jade Alglave and Patrick Cousot. OGRE and pythia: an invariance proof method for weak consistency models. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 3–18. ACM, 2017.
- [5] Jade Alglave, Daniel Kroening, and Michael Tautschnig. Partial orders for efficient bounded model checking of concurrent software. In Natasha Sharygina and Helmut Veith, editors, *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *Lecture Notes in Computer Science*, pages 141–157. Springer, 2013.
- [6] Jade Alglave, Luc Maranget, and Michael Tautschnig. Herding cats: Modelling, simulation, testing, and data mining for weak memory. *ACM Trans. Program. Lang. Syst.*, 36(2):7:1–7:74, 2014.
- [7] Sérgio Almeida, João Leitão, and Luís E. T. Rodrigues. Chainreaction: a causal+ consistent datastore based on chain replication. In Zdenek Hanzálek, Hermann Härtig, Miguel Castro, and M. Frans Kaashoek, editors, *Eighth EuroSys Conference 2013, EuroSys '13, Prague, Czech Republic, April 14-17, 2013*, pages 85–98. ACM, 2013.
- [8] Mohamed Faouzi Atig, Ahmed Bouajjani, Sebastian Burckhardt, and Madanlal Musuvathi. On the verification problem for weak memory models. In Manuel V. Hermenegildo and Jens Palsberg, editors, *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010*, pages 7–18. ACM, 2010.
- [9] Mohamed Faouzi Atig, Ahmed Bouajjani, Sebastian Burckhardt, and Madanlal Musuvathi. What’s decidable about weak memory models? In Helmut Seidl, editor, *Programming Languages and Systems - 21st European Symposium on Programming, ESOP 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012. Proceedings*, volume 7211 of *Lecture Notes in Computer Science*, pages 26–46. Springer, 2012.
- [10] Mohamed Faouzi Atig, Ahmed Bouajjani, and Gennaro Parlato. Getting rid of store-buffers in TSO analysis. In Ganesh Gopalakrishnan and Shaz Qadeer, editors, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *Lecture Notes in Computer Science*, pages 99–115. Springer, 2011.
- [11] Giovanni Bernardi and Alexey Gotsman. Robustness against consistency models with atomic visibility. In Josée Desharnais and Radha Jagadeesan, editors, *27th International Conference on Concurrency Theory, CONCUR 2016, August 23-26, 2016, Québec City, Canada*, volume 59 of *LIPICs*, pages 7:1–7:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [12] Ahmed Bouajjani, Egor Derevenetc, and Roland Meyer. Checking and enforcing robustness against TSO. In Matthias Felleisen and Philippa Gardner, editors, *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*, volume 7792 of *Lecture Notes in Computer Science*, pages 533–553. Springer, 2013.
- [13] Ahmed Bouajjani, Constantin Enea, Rachid Guerraoui, and Jad Hamza. On verifying causal consistency. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 626–638. ACM, 2017.
- [14] Ahmed Bouajjani, Roland Meyer, and Eike Möhlmann. Deciding robustness against total store ordering. In Luca Aceto, Monika Henzinger, and Jirí Sgall, editors, *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part II*, volume 6756 of *Lecture Notes in Computer Science*, pages 428–440. Springer, 2011.

- [15] Lucas Brutschy, Dimitar I. Dimitrov, Peter Müller, and Martin T. Vechev. Serializability for eventual consistency: criterion, analysis, and applications. In Giuseppe Castagna and Andrew D. Gordon, editors, *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, pages 458–472. ACM, 2017.
- [16] Lucas Brutschy, Dimitar I. Dimitrov, Peter Müller, and Martin T. Vechev. Static serializability analysis for causal consistency. In Jeffrey S. Foster and Dan Grossman, editors, *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2018, Philadelphia, PA, USA, June 18-22, 2018*, pages 90–104. ACM, 2018.
- [17] Sebastian Burckhardt. Principles of eventual consistency. *Found. Trends Program. Lang.*, 1(1-2):1–150, 2014.
- [18] Sebastian Burckhardt, Alexey Gotsman, Hongseok Yang, and Marek Zawirski. Replicated data types: specification, verification, optimality. In Suresh Jagannathan and Peter Sewell, editors, *The 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '14, San Diego, CA, USA, January 20-21, 2014*, pages 271–284. ACM, 2014.
- [19] Andrea Cerone and Alexey Gotsman. Analysing snapshot isolation. *J. ACM*, 65(2):11:1–11:41, 2018.
- [20] Andrei Marian Dan, Yuri Meshman, Martin T. Vechev, and Eran Yahav. Effective abstractions for verification under relaxed memory models. *Comput. Lang. Syst. Struct.*, 47:62–76, 2017.
- [21] Egor Derevenetc and Roland Meyer. Robustness against power is pspace-complete. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II*, volume 8573 of *Lecture Notes in Computer Science*, pages 158–170. Springer, 2014.
- [22] Jiaqing Du, Sameh Elnikety, Amitabha Roy, and Willy Zwaenepoel. Orbe: scalable causal consistency using dependency matrices and physical clocks. In Guy M. Lohman, editor, *ACM Symposium on Cloud Computing, SOCC '13, Santa Clara, CA, USA, October 1-3, 2013*, pages 11:1–11:14. ACM, 2013.
- [23] Michael J. Fischer, Nancy A. Lynch, and Mike Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.
- [24] Seth Gilbert and Nancy A. Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, 2002.
- [25] Alexey Gotsman, Hongseok Yang, Carla Ferreira, Mahsa Najafzadeh, and Marc Shapiro. ‘cause i’m strong enough: reasoning about consistency choices in distributed systems. In Rastislav Bodík and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 371–384. ACM, 2016.
- [26] Dexter Kozen. Lower bounds for natural proof systems. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, pages 254–266. IEEE Computer Society, 1977.
- [27] Arthur Kurath. Analyzing Serializability of Cassandra Applications. Master’s thesis, ETH Zurich, Switzerland, 2017.
- [28] Ori Lahav, Nick Giannarakis, and Viktor Vafeiadis. Taming release-acquire consistency. In Rastislav Bodík and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 649–662. ACM, 2016.
- [29] Ori Lahav and Viktor Vafeiadis. Owicki-gries reasoning for weak memory models. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, volume 9135 of *Lecture Notes in Computer Science*, pages 311–323. Springer, 2015.
- [30] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978.
- [31] Leslie Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Computers*, 28(9):690–691, 1979.
- [32] Richard J Lipton and Jonathan S Sandberg. PRAM: A scalable shared memory. Technical Report TR-180-88, Princeton University, Department of Computer Science, August 1988.
- [33] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Don’t settle for eventual: scalable causal consistency for wide-area storage with COPS. In Ted Wobber and Peter Druschel, editors,

- Proceedings of the 23rd ACM Symposium on Operating Systems Principles 2011, SOSP 2011, Cascais, Portugal, October 23-26, 2011*, pages 401–416. ACM, 2011.
- [34] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Stronger semantics for low-latency geo-replicated storage. In Nick Feamster and Jeffrey C. Mogul, editors, *Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2013, Lombard, IL, USA, April 2-5, 2013*, pages 313–328. USENIX Association, 2013.
- [35] Kartik Nagar and Suresh Jagannathan. Automated detection of serializability violations under weak consistency. In Sven Schewe and Lijun Zhang, editors, *29th International Conference on Concurrency Theory, CONCUR 2018, September 4-7, 2018, Beijing, China*, volume 118 of *LIPICs*, pages 41:1–41:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.
- [36] Mahsa Najafzadeh, Alexey Gotsman, Hongseok Yang, Carla Ferreira, and Marc Shapiro. The CISE tool: proving weakly-consistent applications correct. In Peter Alvaro and Alysson Bessani, editors, *Proceedings of the 2nd Workshop on the Principles and Practice of Consistency for Distributed Data, PaPoC@EuroSys 2016, London, United Kingdom, April 18, 2016*, pages 2:1–2:3. ACM, 2016.
- [37] Christos H. Papadimitriou. The serializability of concurrent database updates. *J. ACM*, 26(4):631–653, 1979.
- [38] Matthieu Perrin, Achour Mostéfaoui, and Claude Jard. Causal consistency: beyond memory. In Rafael Asenjo and Tim Harris, editors, *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2016, Barcelona, Spain, March 12-16, 2016*, pages 26:1–26:12. ACM, 2016.
- [39] Nuno M. Pregoça, Marek Zawirski, Annette Bieniusa, Sérgio Duarte, Valter Balegas, Carlos Baquero, and Marc Shapiro. Swiftcloud: Fault-tolerant geo-replication integrated all the way to the client machine. In *33rd IEEE International Symposium on Reliable Distributed Systems Workshops, SRDS Workshops 2014, Nara, Japan, October 6-9, 2014*, pages 30–33. IEEE Computer Society, 2014.
- [40] Charles Rackoff. The covering and boundedness problems for vector addition systems. *Theor. Comput. Sci.*, 6:223–231, 1978.
- [41] Dennis E. Shasha and Marc Snir. Efficient and correct execution of parallel programs that share memory. *ACM Trans. Program. Lang. Syst.*, 10(2):282–312, 1988.