Analysing Snapshot Isolation

Andrea Cerone
IMDEA Software Institute

Alexey Gotsman
IMDEA Software Institute

ABSTRACT

Snapshot isolation (SI) is a widely used consistency model for transaction processing, implemented by most major databases and some of transactional memory systems. Unfortunately, its classical definition is given in a low-level operational way, by an idealised concurrency-control algorithm, and this complicates reasoning about the behaviour of applications running under SI. We give an alternative specification to SI that characterises it in terms of transactional dependency graphs of Adya et al., generalising serialization graphs. Unlike previous work, our characterisation does not require adding additional information to dependency graphs about start and commit points of transactions. We then exploit our specification to obtain two kinds of static analyses. The first one checks when a set of transactions running under SI can be chopped into smaller pieces without introducing new behaviours, to improve performance. The other analysis checks whether a set of transactions running under a weakening of SI behaves the same as when it running under SI.

Keywords

Snapshot isolation; transaction chopping; robustness

1. INTRODUCTION

Transactions simplify concurrent programming by enabling computations on shared data that are isolated from other concurrent computations and resilient to failures. They are commonly provided by databases [7] and, more recently, by transactional memory systems [21]. Ideally, programmers would like to get strong guarantees about the isolation of transactional computations, formalised by the notion of *serializability* [7]: the results of concurrently executing a set transactions could be obtained if these transactions executed atomically in some order. Unfortunately, ensuring serializability carries a significant performance penalty. For this reason, transactional systems often provide weaker guarantees about transaction processing, formalised by *weak consistency models*. Snapshot isolation (SI) [6] is one of the most popular such models, implemented by major centralised databases (e.g., MS SQL Sever, Oracle), distributed databases [13, 26, 28] and transactional memory systems [1, 8, 14, 24].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PODC'16, July 25 - 28, 2016, Chicago, IL, USA

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ISBN 978-1-4503-3964-3/16/07...$15.00

DOI: http://dx.doi.org/10.1145/2933057.2933096

Informally, SI is defined by a multi-version concurrency control algorithm as follows. A transaction T reads values of shared objects from a snapshot taken at its start. The transaction commits only if it passes a *write-conflict* detection check: since T started, no other committed transaction has written to any object that T also wrote to. If the check fails, T aborts. Once T commits, its changes become visible to all transactions that take a snapshot afterwards. This concurrency-control algorithm allows unserializable behaviours, called *anomalies*. One of them, *write skew*, is graphically illustrated in Figure 2(d). Each of the transactions T_1 and T_2 checks that the combined balance of two accounts exceeds 100 and, if so, withdraws 100 from one of them. Under SI, both transactions may pass the checks and make the withdrawals from different accounts, resulting in the combined balance going negative. This outcome cannot occur under serializability. Given such anomalies, reasoning about the behaviour of applications executing under SI is far from trivial. This task is further complicated by the fact that the specification of SI is given in a low-level operational way, by a concurrency control algorithm. To facilitate reasoning about applications using SI and establishing useful results about this consistency model, we need a more declarative specification that abstracts from implementation-level details as much as possible.

An approach that yields such consistency model specifications was proposed by Adya et al. [2, 3]. In this approach, an execution of a set of transactions is described by three kinds of *dependencies* between pairs of transactions T_1 and T_2: *read dependencies* record when T_1 reads the value of an object written by T_2; *write dependencies* record when T_1 overwrites the value of an object written by T_2; finally, *anti-dependencies* are derived from read and write dependencies in a certain way (§3). A set of transactions and dependencies between them form a *dependency graph*, generalising classical serialization graphs [7]. Then the set of executions allowed by a given consistency model is defined by those dependency graphs that lack certain cycles; in particular, serializable executions are characterised by acyclic dependency graphs. This way of specifying consistency models has been shown to be particularly appropriate for designing static analyses [11, 18, 22, 29, 37], run-time monitoring [9, 36] and proving concurrency-control algorithms correct [15, 23, 35]. In particular, specifications in terms of dependency graphs facilitate exploring possible program executions in a static analysis, because the analysis can determine which dependencies can possibly exist at run time by looking for pairs of read or write accesses to the same object in the code of different transactions. In contrast, it is hard to predict statically more low-level information about transaction execution, such as the order in which transactions commit.

Specifications in terms of dependency graphs have been proposed for ANSI isolation levels such as serializability, Read Committed and Repeatable Read [2], as well as more recent proposals of consistency models [5, 35]. But surprisingly, there is no such specification of SI. This is not for the want of trying: Adya did propose a definition of SI that refers to dependency graphs [2]. However, to capture the subtle semantics of SI, this definition extends the graphs by a relation describing low-level information about transaction execution, which negates their benefits.

In this paper we propose the first characterisation of SI solely in terms of dependency graphs (§4) and apply it to develop new static analyses (§5 and §6). Namely, we show that SI allows exactly the executions represented by dependency graphs that contain only cycles with at least two adjacent anti-dependency edges. The proof of this fact is highly non-trivial and represents a key technical contribution of this paper. It requires showing that, given a dependency graph satisfying the above acyclicity condition, we can construct certain relations describing how the transactions can be processed by the SI concurrency control, e.g., the order in which transactions commit. Constructing these relations from transactional dependencies is challenging, and the main insight of our proof is given by a procedure for this construction, based on solving certain kinds of inequalities over relations.

To illustrate the benefits of our dependency graph characterisation of SI, we exploit it to develop two kinds of static analyses. First, we propose a new static analysis for the classical problem of *transaction chopping* [4, 29, 34]—checking when transactions in an application can be chopped into smaller pieces without introducing new behaviours (§5). When applied to long-running transactions executing under SI, chopping can improve performance, because the longer an SI transaction runs, the higher the chances are that it will abort due to a write conflict. There are analyses for transaction chopping under serializability [29] and parallel SI [11], a recently-proposed weaker version of SI for large-scale databases [31]. However, there has been no such analysis under SI, despite the widespread use of this consistency model.

Our dependency graph characterisation of SI is instrumental in deriving the static analysis for transaction chopping, and not only due to the feasibility of determining possible dependencies statically. In more detail, chopping transforms transactions in a program into *sessions* [13, 32] (aka *chains* [37]) of smaller transactions, which ensure that the transactions will be executed in the order given, but provide no isolation guarantees. A chopping is correct if each SI execution of the resulting program can be *spliced* into an SI execution that has the same operations as the original one, but where all operations from each session are executed inside a single transaction. Showing the existence of the spliced execution is challenging on SI because it is non-trivial to pick the order in which its transactions should commit. Our characterisation of SI in terms of transactional dependencies avoids this complication, because unlike low-level aspects of an execution, these dependencies do not change significantly during splicing, and this makes it easy to construct the spliced execution.

The other kind of static analyses that we consider checks whether an application is *robust* [18, 30] against weakening consistency: it behaves the same regardless of whether it uses a database providing a weak consistency model or a database proving a stronger model (§6). When this is the case, the application programmer can reap the performance benefits of using the weaker model, yet can reason about the correctness of the application assuming the stronger one. We first show that our SI characterisation allows easily deriving a variant of an existing analysis that checks whether an application executing under SI behaves the same as when executing under serializability [18] (robustness *against SI*, §6.1). We then propose a new static analysis that checks whether an application executing under the recently-proposed parallel SI [31] behaves the same as when executing under the stronger classical SI (robustness *against parallel SI towards SI*, §6.2). To derive this static analysis, we formulate a dependency graph characterisation of parallel SI, which can be given more easily than for classical SI. Again, our characterisations of consistency models in terms of dependency graphs greatly facilitate deriving the above robustness analyses, since the characterisations allow us to easily map between executions on different models.

Due to space constraints, we defer some of the proofs to §A.

2. SNAPSHOT ISOLATION

We start by formally defining snapshot isolation (SI), as well as serializability. Rather than using the classical definition of SI by a concurrency-control algorithm (§1), it is technically convenient for us to build on a more declarative specification that we previously proposed and proved equivalent to the standard one [10]. Even though this specification is stated in terms of lower-level relations than transactional dependencies, it avoids referring explicitly to times at which a transaction takes a snapshot in the SI concurrency-control algorithm. We first introduce mathematical structures that represent transaction execution in the specification.

We consider a transactional system managing a set of integer-valued ***objects*** $\mathsf{Obj} = \{x, y, \ldots\}$. Transactions read and write the objects, and in our representation of executions, we denote each invocation of such an operation by an ***event*** from a set $\mathsf{Event} = \{e, f, \ldots\}$. A function $\mathsf{op} : \mathsf{Event} \to \mathsf{Op}$ for $\mathsf{Op} = \{\mathtt{read}(x, n), \mathtt{write}(x, n) \mid x \in \mathsf{Obj}, n \in \mathbb{Z}\}$ determines the operation a given event denotes: reading a value n from an object x or writing n to x. We call a binary relation a ***strict partial order*** if it is transitive and irreflexive. We call it a ***total order*** if it additionally relates any pair of distinct elements one way or another. We represent an execution of a single transaction by the following structure, recording a set of operations and the order in which they were invoked.

DEFINITION 1. *A **transaction** $T, S, \ldots$ is a pair (E, po), where $E \subseteq \mathsf{Event}$ is a finite, non-empty set of events and the **program order** $\mathsf{po} \subseteq E \times E$ is a total order.*

For simplicity, all transactions in this paper are assumed to be committed: our specifications do not constrain values read inside aborted or ongoing transactions; this limitation could be lifted following [2, 16, 20]. We denote components of transactions and similar structures as in E_T and po_T.

To allow transaction chopping (§5), we assume that the transactional system allows its clients to group several transactions into a session [32], which establishes an ordering on the transactions. Thus, instead of classical SI and serializability, we actually define their ***strong session*** variants [12, 13]. We represent the client-visible results of an execution of a set of sessions by a *history*.

DEFINITION 2. *A **history** is a pair $\mathcal{H} = (\mathcal{T}, \mathsf{SO})$, where $\mathcal{T}$ is a finite set of transactions with disjoint sets of events and the **session order** $\mathsf{SO} \subseteq \mathcal{T} \times \mathcal{T}$ is a union of total orders defined on disjoint subsets of $\mathcal{T}$, which correspond to transactions in different sessions.*

For simplicity, we elide the treatment of infinite computations, and thus histories are always finite. A consistency model, such as SI or serializability, is specified by a set of histories. To define this set, we extend histories with two relations, declaratively describing how the transactional system processes transactions.

DEFINITION 3. *An **abstract execution** (or just an execution) is a tuple $\mathcal{X} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO})$, where $(\mathcal{T}, \mathsf{SO})$ is a history and the **visibility** and **commit orders** $\mathsf{VIS}, \mathsf{CO} \subseteq \mathcal{T} \times \mathcal{T}$ are such that $\mathsf{VIS} \subseteq \mathsf{CO}$ and CO is total.*

$\forall (E, \mathsf{po}) \in \mathcal{T}.\, \forall e \in E.\, \forall x, n.\, \mathsf{op}(e) = \mathtt{read}(x, n) \wedge \{f \mid \mathsf{op}(f) = \_(x, \_) \wedge f \xrightarrow{\mathsf{po}} e\} \neq \emptyset \implies$ $\mathsf{op}(\max_{\mathsf{po}}\{f \mid \mathsf{op}(f) = \_(x, \_) \wedge f \xrightarrow{\mathsf{po}} e\}) = \_(x, n)$ (INT)	$\mathsf{SO} \subseteq \mathsf{VIS}$ (SESSION)
	$\mathsf{CO} \;;\, \mathsf{VIS} \subseteq \mathsf{VIS}$ (PREFIX)
$\forall T \in \mathcal{T}.\, \forall x, n.\, T \vdash \mathtt{read}(x, n) \implies \max_{\mathsf{CO}}(\mathsf{VIS}^{-1}(T) \cap \mathsf{WriteTx}_x) \vdash \mathtt{write}(x, n)$ (EXT)	$\mathsf{CO} = \mathsf{VIS}$ (TOTALVIS)
$\forall T, S \in \mathcal{T}.\, \forall x.\, (T, S \in \mathsf{WriteTx}_x \wedge T \neq S) \implies (T \xrightarrow{\mathsf{VIS}} S \vee S \xrightarrow{\mathsf{VIS}} T)$ (NOCONFLICT)	

Figure 1: Axioms constraining an abstract execution $(\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO})$.

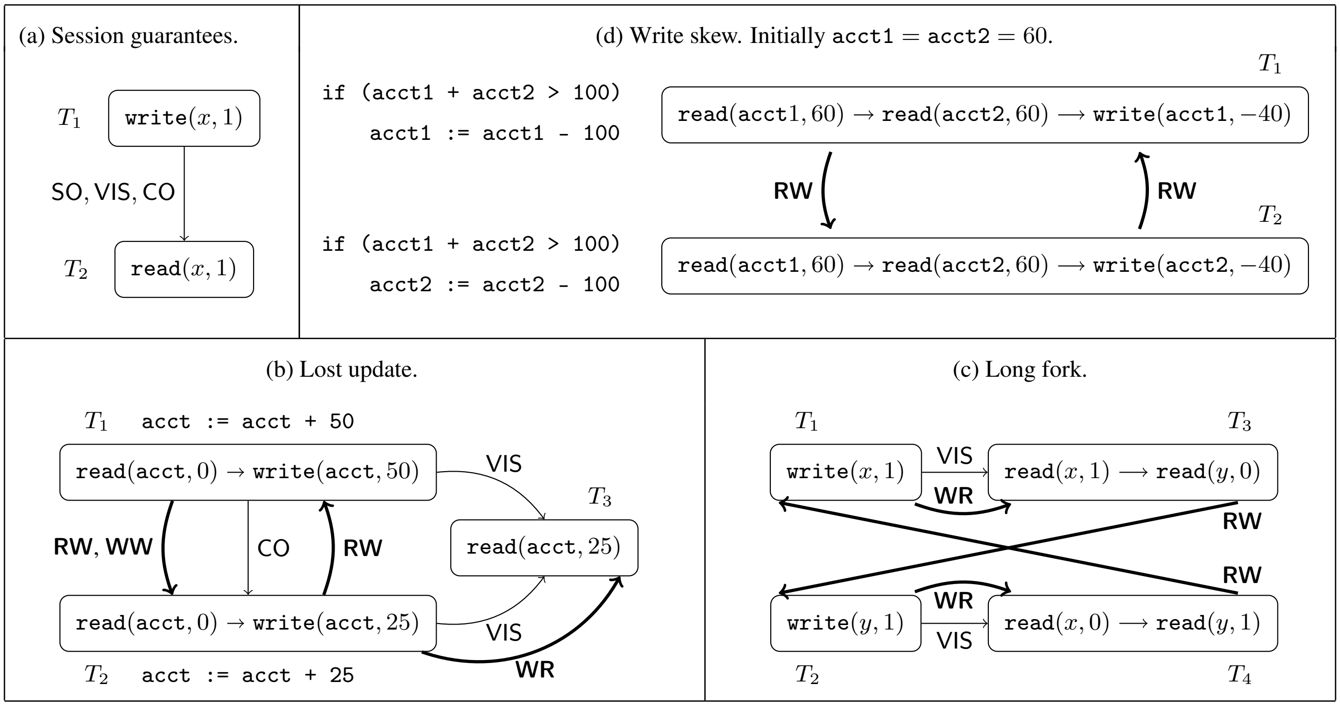

Figure 2: Abstract executions illustrating SI and serializability. Boxes represent transactions, and arrows inside boxes represent the program order. We omit irrelevant CO edges. We also omit a special transaction that writes initial versions of all objects and precedes all the other transactions in VIS and CO. The bold edges are explained in §3.

We write $T \xrightarrow{\mathsf{VIS}} S$ and $(T, S) \in \mathsf{VIS}$ interchangeably, and similarly for other relations. For $\mathcal{H} = (\mathcal{T}, \mathsf{SO})$ we shorten $(\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO})$ to $(\mathcal{H}, \mathsf{VIS}, \mathsf{CO})$. In terms of the SI concurrency-control algorithm sketched in §1, $T \xrightarrow{\mathsf{VIS}} S$ means that the writes done by the transaction T are included into the snapshot taken by the transaction S; $T \xrightarrow{\mathsf{CO}} S$ means that T commits earlier than S. The constraint $\mathsf{VIS} \subseteq \mathsf{CO}$ ensures that the snapshot taken by a transaction may only include previously committed transactions. SI or serializability allow those histories that can be extended to an abstract execution satisfying certain ***consistency axioms*** from Figure 1, which specify the corresponding guarantees about transaction processing.

DEFINITION 4. *The sets of executions and histories* ***allowed by (strong session) SI and serializability*** *are:*

$$\begin{aligned}
\mathsf{ExecSI} &= \{\mathcal{X} \mid \mathcal{X} \models \textsc{Int} \wedge \textsc{Ext} \wedge \textsc{Session} \wedge \\
&\qquad\qquad \textsc{Prefix} \wedge \textsc{NoConflict}\};\\
\mathsf{ExecSER} &= \{\mathcal{X} \mid \mathcal{X} \models \textsc{Int} \wedge \textsc{Ext} \wedge \textsc{Session} \wedge \textsc{TotalVis}\};\\
\mathsf{HistSI} &= \{\mathcal{H} \mid \exists \mathsf{VIS}, \mathsf{CO}.\, (\mathcal{H}, \mathsf{VIS}, \mathsf{CO}) \in \mathsf{ExecSI}\};\\
\mathsf{HistSER} &= \{\mathcal{H} \mid \exists \mathsf{VIS}, \mathsf{CO}.\, (\mathcal{H}, \mathsf{VIS}, \mathsf{CO}) \in \mathsf{ExecSER}\}.
\end{aligned}$$

We now explain the axioms in Figure 1, as well as anomalies that SI allows or disallows; the latter are summarised in Figure 2. We use the following notation. For a set A and a total order $R \subseteq A \times A$, we let $\max_R(A)$ be the element $a \in A$ such that $\forall b \in A.\, a = b \vee (b, a) \in R$; if $A = \emptyset$, then $\max_R(A)$ is undefined. In the following, the use of $\max_R(A)$ in an expression implicitly assumes that it is defined. We define $\min_R(A)$ similarly. For a relation $R \subseteq A \times A$ and an element $a \in A$, we let $R^{-1}(a) = \{b \mid (b, a) \in R\}$. We define the sequential composition of relations R_1 and R_2 as

$$R_1 \;;\, R_2 = \{(a, b) \mid \exists c.\, (a, c) \in R_1 \wedge (c, b) \in R_2\}.$$

We write _ for a value that is irrelevant and implicitly existentially quantified.

The INT and EXT axioms in Figure 1 ensure that a transaction reads from a snapshot of object states and its own writes. The ***internal consistency axiom*** INT ensures that a read event e on an object x returns the same value as the last write to or a read from x preceding e in the same transaction. If a read is not preceded in the same transaction by an operation on the same object, then its value is determined in terms of writes by other transactions using the ***external consistency axiom*** EXT. For $T = (E, \mathsf{po})$, we let $T \vdash \mathtt{write}(x, n)$ if T writes to x and the last value written is n:

$$\mathsf{op}(\max_{\mathsf{po}}\{e \mid \mathsf{op}(e) = \mathtt{write}(x, \_)\}) = \mathtt{write}(x, n).$$

We let $T \vdash \mathtt{read}(x, n)$ if T reads from x before writing to it and n is the value returned by the first such read:

$$\mathsf{op}(\min_{\mathsf{po}}\{e \mid \mathsf{op}(e) = \_(x, \_)\}) = \mathtt{read}(x, n).$$

We also let $\mathsf{WriteTx}_x = \{T \mid T \vdash \mathtt{write}(x, \_)\}$. Then EXT ensures that, if a transaction T reads an object x before writing to it, then the value read is determined by the transactions that are included into T's snapshot according to VIS and that wrote to x; T reads the value written by the transaction from this set that committed last according to CO. For simplicity, we consider only executions where the above set is always non-empty; this can be ensured by introducing a special transaction that writes initial values of all objects. The executions in Figures 2(a) and 2(b) satisfy EXT.

Our specification determines the snapshot that a transaction reads from based on an arbitrary visibility relation and does not require the snapshot to be "latest"; this is similar to so-called ***generalised SI*** [17]. However, following strong session SI [12, 13], the SESSION axiom requires the snapshot to include the effects of all preceding transactions in the same session. For example, in the execution in Figure 2(a), the session order between T_1 and T_2 induces a visibility edge according to SESSION.

The PREFIX axiom ensures that, if the snapshot taken by a transaction T includes a (committed) transaction S, then this snapshot also includes all transactions that committed before S. Note that PREFIX and the property $\mathsf{VIS} \subseteq \mathsf{CO}$ in Definition 4 imply that VIS is transitive. PREFIX disallows the ***long fork*** anomaly shown in Figure 2(c), which is allowed by some weakening of SI (such as parallel SI [31]). There transactions T_1 and T_2 concurrently write to objects x and y. Transaction T_3 sees the write by T_1, but not the write by T_2; conversely, transaction T_4 sees the write by T_2, but not the write by T_1. Thus, from the perspectives of T_3 and T_4, the writes of T_1 and T_2 happen in different orders. PREFIX disallows any execution with the history in Figure 2(c), because in such an execution T_1 and T_2 have to be related by CO one way or another; but then by PREFIX, either T_4 has to observe the write to x or T_3 has to observe the write to y.

The axioms explained so far do not prevent the ***lost update*** anomaly, illustrated by the execution in Figure 2(b). This execution could arise from the code in the figure that uses transactions T_1 and T_2 to make deposits into an account. The two transactions read the initial balance of the account and concurrently modify it, resulting in one deposit getting lost. This anomaly is disallowed by the NOCONFLICT axiom: if two distinct transactions write to the same object, then one of them has to be aware of the other. This axiom rules out any execution with the history in Figure 2(b): it forces T_1 and T_2 to be ordered by VIS, so that they cannot both read 0 from `acct`. In the SI concurrency control this is ensured by the write-conflict detection check (§1).

The set HistSI (Definition 4) defined using the consistency axioms explained so far is exactly the one produced by the SI concurrency-control algorithm [10]. The axioms allow the execution in Figure 2(d) with the characteristic SI anomaly of ***write skew*** (§1), disallowed by serializability. We formalise the latter by the axiom TOTALVIS, which requires visibility to totally order all transactions. Then the axioms INT and EXT ensure that the transactions are processed according to the usual sequential semantics. We thus have $\mathsf{HistSER} \subset \mathsf{HistSI}$.

3. DEPENDENCY GRAPHS

From an abstract execution we can extract several kinds of dependencies between its transactions, which are used in consistency model specifications in the style of Adya et al. [2, 3].

DEFINITION 5. *Let $\mathcal{X} = (\mathcal{H}, \mathsf{VIS}, \mathsf{CO})$ be an execution. For $x \in \mathsf{Obj}$, we define the following relations on $\mathcal{T}_{\mathcal{H}}$:*

- ***read dependency**: $T \xrightarrow{\mathsf{WR}_{\mathcal{X}}(x)} S \iff$*
  $$S \vdash \mathtt{read}(x, \_) \wedge T = \max_{\mathsf{CO}}(\mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x);$$
- ***write dependency**: $T \xrightarrow{\mathsf{WW}_{\mathcal{X}}(x)} S \iff$*
  $$T \xrightarrow{\mathsf{CO}} S \wedge T, S \in \mathsf{WriteTx}_x;$$
- ***anti-dependency**: $T \xrightarrow{\mathsf{RW}_{\mathcal{X}}(x)} S \iff$*
  $$T \neq S \wedge \exists T'.\, T' \xrightarrow{\mathsf{WR}_{\mathcal{X}}(x)} T \wedge T' \xrightarrow{\mathsf{WW}_{\mathcal{X}}(x)} S.$$

Informally, $T \xrightarrow{\mathsf{WR}_{\mathcal{X}}(x)} S$ means that S reads T's write to x (cf. the EXT axiom in Figure 1); $T \xrightarrow{\mathsf{WW}_{\mathcal{X}}(x)} S$ means that S overwrites T's write to x; $T \xrightarrow{\mathsf{RW}_{\mathcal{X}}(x)} S$ means that S overwrites the write to x read by T. For example, the dependencies of the executions in Figures 2(b), 2(d) and 2(c) are shown there with bold arrows (keep in mind that the pictures omit a special initialisation transaction). We often abuse notation and use the symbol $\mathsf{WR}_{\mathcal{X}}$ to also denote the relation $\bigcup_{x \in \mathsf{Obj}} \mathsf{WR}_{\mathcal{X}}(x) \subseteq \mathcal{T}_{\mathcal{H}} \times \mathcal{T}_{\mathcal{H}}$, and similarly for $\mathsf{WW}_{\mathcal{X}}$ and $\mathsf{RW}_{\mathcal{X}}$.

A key goal of this paper is to characterise SI solely in terms of dependencies: we want to determine whether SI allows a given history by looking for appropriate dependencies between its transactions rather than visibility and commit orders, as in Definition 4. To this end, we extend histories to *dependency graphs* (aka direct serialization graphs) [2], which include relations representing the dependencies.

DEFINITION 6. *A **dependency graph** is a tuple $\mathcal{G} = (\mathcal{T}, \mathsf{SO}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$, where $(\mathcal{T}, \mathsf{SO})$ is a history and*

- *$\mathsf{WR} : \mathsf{Obj} \to 2^{\mathcal{T} \times \mathcal{T}}$ is such that:*
  - $\forall T, S \in \mathcal{T}.\, \forall x.\, T \xrightarrow{\mathsf{WR}(x)} S \implies$
    $\exists n.\, T \neq S \wedge T \vdash \mathtt{write}(x, n) \wedge S \vdash \mathtt{read}(x, n)$*;*
  - $\forall S \in \mathcal{T}.\, \forall x.\, S \vdash \mathtt{read}(x, \_) \implies \exists T.\, T \xrightarrow{\mathsf{WR}(x)} S$*;*
  - $\forall T, T', S \in \mathcal{T}.\, \forall x.\, (T \xrightarrow{\mathsf{WR}(x)} S \wedge T' \xrightarrow{\mathsf{WR}(x)} S) \implies$
    $T = T'$*.*
- *$\mathsf{WW} : \mathsf{Obj} \to 2^{\mathcal{T} \times \mathcal{T}}$ is such that for every $x \in \mathsf{Obj}$, $\mathsf{WW}(x)$ is a total order on the set $\mathsf{WriteTx}_x$;*
- *$\mathsf{RW} : \mathsf{Obj} \to 2^{\mathcal{T} \times \mathcal{T}}$ is derived from WR and WW as in Definition 5.*

PROPOSITION 7. *For any $\mathcal{X} \in \mathsf{ExecSI}$, $\mathsf{graph}(\mathcal{X}) = (\mathcal{T}_{\mathcal{X}}, \mathsf{SO}_{\mathcal{X}}, \mathsf{WR}_{\mathcal{X}}, \mathsf{WW}_{\mathcal{X}}, \mathsf{RW}_{\mathcal{X}})$ is a dependency graph.*

Note that the constraints on WR in Definition 6 ensure that it uniquely determines the values read by transactions. For $\mathcal{H} = (\mathcal{T}, \mathsf{SO})$ we write $(\mathcal{H}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$ for $(\mathcal{T}, \mathsf{SO}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$.

We write $\mathcal{T} \models \text{INT}$ if a set of transactions $\mathcal{T}$ satisfies the internal consistency axiom INT in Figure 1. (Strong session) serializability can be characterised by the set of acyclic dependency graphs with internally consistent transactions [2].

THEOREM 8. *Let*

$$\mathsf{GraphSER} = \{\mathcal{G} \mid (\mathcal{T}_{\mathcal{G}} \models \text{INT}) \wedge ((\mathsf{SO}_{\mathcal{G}} \cup \mathsf{WR}_{\mathcal{G}} \cup \mathsf{WW}_{\mathcal{G}} \cup \mathsf{RW}_{\mathcal{G}}) \textit{ is acyclic})\}.$$

Then

$$\mathsf{HistSER} = \{\mathcal{H} \mid \exists \mathsf{WR}, \mathsf{WW}, \mathsf{RW}.\ (\mathcal{H}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW}) \in \mathsf{GraphSER}\}.$$

For example, the histories in Figures 2(b), 2(d) and 2(c) are not serializable, and they cannot be extended to acyclic dependency graphs; in particular, the graphs shown in the figures with bold edges satisfy the conditions of Definition 6, but contain cycles. We now set out to find a characterisation of the above form for SI.

4. SI CHARACTERISATION

For a set $\mathcal{T}$ and a relation $R \subseteq \mathcal{T} \times \mathcal{T}$ let $R? = R \cup \{(T, T) \mid T \in \mathcal{T}\}$. We show that (strong session) SI is characterised by dependency graphs that contain only cycles with at least two adjacent anti-dependency edges.

THEOREM 9. *Let*

$$\mathsf{GraphSI} = \{\mathcal{G} \mid (\mathcal{T}_\mathcal{G} \models \textsc{Int}) \wedge (((\mathsf{SO}_\mathcal{G} \cup \mathsf{WR}_\mathcal{G} \cup \mathsf{WW}_\mathcal{G}) \ ; \mathsf{RW}_\mathcal{G}?) \textit{ is acyclic})\}.$$

Then

$$\mathsf{HistSI} = \{\mathcal{H} \mid \exists \mathsf{WR}, \mathsf{WW}, \mathsf{RW}.\ (\mathcal{H}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW}) \in \mathsf{GraphSI}\}.$$

According to the theorem, to determine whether a particular history is allowed by SI, we can look for dependencies that extend it to a graph in GraphSI. As we demonstrate in §5 and §6, this way of defining SI is particularly suitable for developing static analyses for this consistency model. The history in Figure 2(d) is allowed by SI, and indeed the dependency graph shown in the figure contains only cycles with two adjacent anti-dependencies (e.g., $T_1 \xrightarrow{\mathsf{RW}} T_2 \xrightarrow{\mathsf{RW}} T_1$). In contrast, the histories in Figures 2(b) and 2(c) are not allowed by SI, and they cannot be extended to graphs where every cycle has at least two adjacent anti-dependencies. In particular, the graphs shown in the figures contains cycles without these: e.g., $T_1 \xrightarrow{\mathsf{WW}} T_2 \xrightarrow{\mathsf{RW}} T_1$ in Figure 2(b) and $T_1 \xrightarrow{\mathsf{WR}} T_3 \xrightarrow{\mathsf{RW}} T_2 \xrightarrow{\mathsf{WR}} T_4 \xrightarrow{\mathsf{RW}} T_1$ in Figure 2(c).

To prove Theorem 9, we prove a slightly stronger result, showing that we can establish a correspondence between executions in ExecSI and graphs in GraphSI that preserves histories and dependencies.

THEOREM 10.
(i) ***Soundness:*** $\forall \mathcal{G} \in \mathsf{GraphSI}.\ \exists \mathcal{X} \in \mathsf{ExecSI}.\ \mathsf{graph}(\mathcal{X}) = \mathcal{G}$.
(ii) ***Completeness:*** $\forall \mathcal{X} \in \mathsf{ExecSI}.\ \mathsf{graph}(\mathcal{X}) \in \mathsf{GraphSI}$.

As we explain in §7, the easier completeness direction of this theorem actually follows from existing results [18]. Our main technical contribution is the more challenging proof of the soundness direction, which is required for the static analyses that we propose (§5 and §6). We present this proof first.

The main challenge is to construct a total commit order in the desired execution $\mathcal{X}$ from the dependencies given by $\mathcal{G}$ while satisfying the SI axioms (Definition 4). We do this incrementally; at intermediate stages of the construction we get structures similar to abstract executions, but where the commit order can be partial.

DEFINITION 11. *A tuple* $\mathcal{P} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO})$ *is a* ***pre-execution*** *if it satisfies all the conditions of Definition 3, except* CO *is a strict partial order that may not be total. We let* PreExecSI *be the set of pre-executions satisfying the SI axioms (Figure 1):*

$$\mathsf{PreExecSI} = \{\mathcal{P} \mid \mathcal{P} \models \textsc{Int} \wedge \textsc{Ext} \wedge \textsc{Session} \wedge \textsc{Prefix} \wedge \textsc{NoConflict}\}.$$

$$\begin{cases} \mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW} \subseteq \mathsf{VIS} & \text{(S1)} \\ \mathsf{CO} \ ; \mathsf{VIS} \subseteq \mathsf{VIS} & \text{(S2)} \\ \mathsf{VIS} \subseteq \mathsf{CO} & \text{(S3)} \\ \mathsf{CO} \ ; \mathsf{CO} \subseteq \mathsf{CO} & \text{(S4)} \\ \mathsf{VIS} \ ; \mathsf{RW} \subseteq \mathsf{CO} & \text{(S5)} \end{cases}$$

Figure 3: Requirements on a pre-execution $\mathcal{P} = (\mathcal{H}, \mathsf{VIS}, \mathsf{CO})$ constructed from a dependency graph $\mathcal{G} = (\mathcal{H}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$.

Thus, an execution is a pre-execution whose commit order is total. In the following, we apply the graph function of §3 also to pre-executions; for $\mathcal{P} \in \mathsf{PreExecSI}$, $\mathsf{graph}(\mathcal{P})$ is indeed a dependency graph.

We first obtain auxiliary results that, given a dependency graph $\mathcal{G} = (\mathcal{T}, \mathsf{SO}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW}) \in \mathsf{GraphSI}$, allow us to construct a pre-execution $\mathcal{P} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO}) \in \mathsf{PreExecSI}$ such that $\mathsf{graph}(\mathcal{P}) = \mathcal{G}$. Later we show how to extend $\mathcal{P}$ to a desired execution $\mathcal{X} \in \mathsf{ExecSI}$. We start by restating the requirements on the pre-execution $\mathcal{P}$ in a way more suitable for guiding its construction; these are given by the system of inequalities in Figure 3. First, to ensure $\mathsf{graph}(\mathcal{P}) = \mathcal{G}$, by Definition 5 at the very least we must have $\mathsf{WR} \cup \mathsf{WW} \subseteq \mathsf{VIS}$. For $\mathcal{P}$ to satisfy the SESSION axiom we must also have $\mathsf{SO} \subseteq \mathsf{VIS}$. These two observations motivate (S1). This inequality also implies that $\mathcal{P}$ satisfies NOCONFLICT, since according to Definition 5, WW is total over transactions that write to a given object. Inequality (S2) is equivalent to PREFIX, and inequality (S3) states a relationship between VIS and CO inherited by Definition 11 from Definition 3. Inequality (S4) requires CO to be transitive; (S2) and (S3) ensure that so is VIS.

As we now explain, (S5) ensures the axiom EXT. Consider the dependency graph $\mathcal{G} \in \mathsf{GraphSI}$ in Figure 4, which we use as our running example in this section and in §5. Its transactions could arise from the programs, also shown in the figure, that make a transfer between two accounts and query their balances or the sum thereof. The transactions arising from `lookup1` and `lookup2` see the initial state of the database, while the transactions arising from `lookupAll` see its state in the middle of a transfer. The VIS and CO relations shown by the solid arrows give a pre-execution satisfying inequalities (S1)-(S4) and, in fact, all the SI axioms. Suppose we want to construct a pre-execution with a bigger CO by adding an edge $T \xrightarrow{\mathsf{CO}} S'$ (shown dotted). Since $S' \xrightarrow{\mathsf{VIS}} S$, for the resulting pre-execution to satisfy PREFIX we also need to add an edge $T \xrightarrow{\mathsf{VIS}} S$. But then the pre-execution violates EXT: S sees the write to `acct2` by T, but reads the value from the initialisation transaction (elided from Figure 4) that writes 0 to `acct2` and precedes T in CO and WW; the latter fact is witnessed by the edge $S \xrightarrow{\mathsf{RW}} T$. On the other hand, adding an edge $S' \xrightarrow{\mathsf{CO}} T$ (shown dashed), which belongs to VIS ; RW, does not violate EXT. This example illustrates a general pattern. First, as the following lemma shows, (S5) must hold in any SI execution.

LEMMA 12. $\forall \mathcal{X} \in \mathsf{ExecSI}.\ \mathsf{VIS}_\mathcal{X} \ ; \mathsf{RW}_\mathcal{X} \subseteq \mathsf{CO}_\mathcal{X}$.

Conversely, the system of inequalities in Figure 3 can be used to ensure that a pre-execution $\mathcal{P} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO})$ satisfies EXT and has other desired properties.

LEMMA 13. *Let* $\mathcal{G} = (\mathcal{T}, \mathsf{SO}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$ *be a dependency graph such that* $\mathcal{T} \models \textsc{Int}$ *and* $\mathsf{VIS}, \mathsf{CO} \subseteq \mathcal{T} \times \mathcal{T}$ *be acyclic relations satisfying the system of inequalities in Figure*

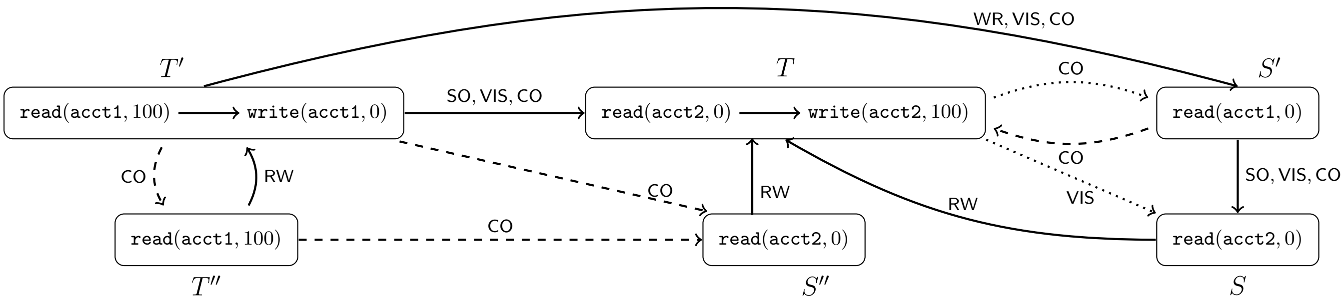

T', T: `session transfer  { tx { acct1 = acct1 - 100 }; tx { acct2 = acct2 + 100 } }`
T'': `session lookup1   { tx { return acct1 } }`
S'': `session lookup2   { tx { return acct2 } }`
S', S: `session lookupAll { tx { var1 = acct1 }; tx { var2 = acct2 }; return var1 + var2 }`

Figure 4: An illustration of constructing an execution from a dependency graph (§4) and splicing an execution (§5). We omit an initialisation transaction that sets $\texttt{acct1} = 100$ and $\texttt{acct2} = 0$.

3. Then $\mathcal{P} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO})$ is a pre-execution such that $\mathcal{P} \in \mathsf{PreExecSI}$ and $\mathsf{graph}(\mathcal{P}) = \mathcal{G}$.

The proof of Lemma 12 depends on the following characterisation of anti-dependencies in terms of visibility edges.

PROPOSITION 14.
$$\forall \mathcal{X} \in \mathsf{ExecSI}.\, \forall T, S \in \mathcal{T}_{\mathcal{X}}.\, S \xrightarrow{\mathsf{RW}_{\mathcal{X}}} T \iff S \neq T \wedge$$
$$\exists x.\, S \vdash \mathtt{read}(x, \_) \wedge T \vdash \mathtt{write}(x, \_) \wedge \neg(T \xrightarrow{\mathsf{VIS}_{\mathcal{X}}} S).$$

Informally, if we had $S \xrightarrow{\mathsf{RW}_{\mathcal{X}}} T$ and $T \xrightarrow{\mathsf{VIS}_{\mathcal{X}}} S$, then S would have to read a value of x at least as up-to-date as that written by T, contradicting the definition of $\mathsf{RW}_{\mathcal{X}}$.

PROOF OF LEMMA 12. Consider $\mathcal{X} \in \mathsf{ExecSI}$ and $T, S', S \in \mathcal{T}_{\mathcal{X}}$ such that $S' \xrightarrow{\mathsf{VIS}_{\mathcal{X}}} S \xrightarrow{\mathsf{RW}_{\mathcal{X}}} T$. If $T = S'$, then $S' \xrightarrow{\mathsf{VIS}_{\mathcal{X}}} S \xrightarrow{\mathsf{RW}_{\mathcal{X}}} S'$, contradicting Proposition 14. If $T \xrightarrow{\mathsf{CO}_{\mathcal{X}}} S'$ (see Figure 4), then by PREFIX we get $T \xrightarrow{\mathsf{VIS}_{\mathcal{X}}} S$, contradicting Proposition 14. Then, since $\mathsf{CO}_{\mathcal{X}}$ is total, we must have $S' \xrightarrow{\mathsf{CO}_{\mathcal{X}}} T$. □

PROOF OF LEMMA 13. We only prove that $\mathcal{P} \models \text{EXT}$ and $\mathsf{WR}_{\mathcal{P}} = \mathsf{WR}$; discharging the other obligations is straightforward. Consider $S \in \mathcal{T}$ such that $S \vdash \mathtt{read}(x, n)$. Then there exists a unique T' such that $T' \xrightarrow{\mathsf{WR}(x)} S$. Let $T = \max_{\mathsf{CO}}(\mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x)$. This is defined because: CO is acyclic; by (S1) and (S3) we have $\mathsf{WW} \subseteq \mathsf{CO}$, so that CO is total over $\mathsf{WriteTx}_x$; and by (S1) we have $\mathsf{WR} \subseteq \mathsf{VIS}$, so that $T' \in \mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x$. We now show that $T = T'$, which entails the required.

Assume the contrary: $T \neq T'$. We have $T, T' \in \mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x$. Hence, T and T' are related by WW. Since $\mathsf{WW} \subseteq \mathsf{CO}$, they are related in the same way by the acyclic $\mathsf{CO}(x)$. Then by the definition of T we must have $T' \xrightarrow{\mathsf{CO}(x)} T$ and $T' \xrightarrow{\mathsf{WW}} T$. From the latter and $T' \xrightarrow{\mathsf{WR}(x)} S$ we get $S \xrightarrow{\mathsf{RW}} T$. But $T \xrightarrow{\mathsf{VIS}} S$, so from (S5) we get $T \xrightarrow{\mathsf{CO}} T$. This contradicts the assumption that CO is acyclic. Hence, we must have $T = T'$. □

According to Lemma 13, to construct a desired pre-execution $\mathcal{P} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO}) \in \mathsf{PreExecSI}$ from a dependency graph $\mathcal{G} = (\mathcal{T}, \mathsf{SO}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$, it is sufficient to find a solution to the system of inequalities in Figure 3 in terms of *acyclic* relations VIS and CO. This is not completely trivial because of the recursive nature of the inequalities: according to them, adding more edges into VIS forces adding more edges into CO and vice versa, increasing the risk of tying a cycle. Our insight is to look for the solution that is *smallest* and, hence, least likely to contain cycles. The following lemma gives a closed form for this solution. In anticipation of using the lemma when extending a pre-execution to an execution, we state it in a generalised form that gives the smallest solution where CO contains at least a given set of edges R. We use $^+$ and * to denote the transitive closure and the transitive and reflexive closure of a given relation.

LEMMA 15. *Let $\mathcal{G} = (\mathcal{T}, \mathsf{SO}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$ be a dependency graph. For any $R \subseteq \mathcal{T} \times \mathcal{T}$, the relations*
$$\begin{aligned} \mathsf{VIS} &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW})\, ;\, \mathsf{RW}?) \cup R)^*\, ; \\ &\quad (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}); \\ \mathsf{CO} &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW})\, ;\, \mathsf{RW}?) \cup R)^+ \end{aligned} \tag{1}$$
are a solution to the system of inequalities in Figure 3. They also are the smallest solution to the system for which $\mathsf{CO} \supseteq R$: for any other solution $(\mathsf{VIS}', \mathsf{CO}')$ with $\mathsf{CO}' \supseteq R$ we have $\mathsf{VIS} \subseteq \mathsf{VIS}'$ and $\mathsf{CO} \subseteq \mathsf{CO}'$.

In particular, for $R = \emptyset$, Lemma 15 gives the smallest solution $(\mathsf{VIS}_0, \mathsf{CO}_0)$ to the system of inequalities in Figure 3.

If $\mathcal{G} \in \mathsf{GraphSI}$, then CO_0 is acyclic, and by (S3), so is VIS_0 (in fact, Lemma 15 is our motivation for defining $\mathsf{GraphSI}$ the way we did). Hence, by Lemma 13, $\mathcal{P}_0 = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}_0, \mathsf{CO}_0)$ is a pre-execution such that $\mathcal{P}_0 \in \mathsf{PreExecSI}$ and $\mathsf{graph}(\mathcal{P}_0) = \mathcal{G}$, which is what we originally set out to construct. We now proceed to prove Theorem 10(i) by extending the pre-execution $\mathcal{P}_0$ to an execution $\mathcal{X} \in \mathsf{ExecSI}$.

PROOF OF THEOREM 10(I). Assume $\mathcal{G} = (\mathcal{T}, \mathsf{SO}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW}) \in \mathsf{GraphSI}$. To construct $\mathcal{X}$, we define a sequence of pre-executions $\{\mathcal{P}_i = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}_i, \mathsf{CO}_i)\}_{i=0}^{n}$ for some $n \geq 0$ and let $\mathcal{X} = \mathcal{P}_n$. The sequence is such that $\mathsf{VIS}_i \subseteq \mathsf{VIS}_{i+1}$ and $\mathsf{CO}_i \subset \mathsf{CO}_{i+1}$ for $i = 0..(n-1)$; furthermore, CO_n is total, so that $\mathcal{P}_n$ is an execution. That is, on every step of our construction we add edges to the commit order until it becomes total. Each pair $(\mathsf{VIS}_i, \mathsf{CO}_i)$ gives an acyclic solution to the system in Figure 3. Then by Lemma 13, $\mathcal{P}_i \in \mathsf{PreExecSI}$ and $\mathsf{graph}(\mathcal{P}_i) = \mathcal{G}$. In particular, $\mathcal{P}_n \in \mathsf{ExecSI}$ and $\mathsf{graph}(\mathcal{P}_n) = \mathcal{G}$, as required.

We start the construction of the sequence by taking as $\mathcal{P}_0$ the pre-execution that we constructed above. For example, for the dependency graph in Figure 4, VIS_0 and CO_0 consist of the solid edges in the figure and the dashed edge $S' \xrightarrow{\mathsf{CO}} T$. If the relation CO_0 is not total, then we pick an arbitrary pair of transactions (T_1, S_1) unrelated by CO_0 and construct VIS_1 and CO_1 as the smallest solution to the system of inequalities in Figure 3 such that $\mathsf{CO}_1 \supseteq \{(T_1, S_1)\}$. By Lemma 15 this solution is given by (1) for $R = \{(T_1, S_1)\}$. For example, in Figure 3 the transactions T' and T'' are unrelated by the commit order. If we pick as (T_1, S_1) the pair (T', T'') in Figure 4, then we get $\mathsf{VIS}_1 = \mathsf{VIS}_0$, and $\mathsf{CO}_1 = \mathsf{CO}_0 \cup \{(T', T'')\}$. In general, the construction continues in the same way: while CO_i is not total, we pick an arbitrary pair of transactions (T_i, S_i) unrelated by CO_i and force CO_{i+1} to include it. In our example, CO_1 does not relate the transactions T'' and S''. By picking as (T_2, S_2) the pair (T'', S''), we construct VIS_2 and CO_2 by letting $R = \{(T', T''), (T'', S'')\}$ in (1); this corresponds to all the solid and dashed edges in Figure 4. Note that CO_2 also includes the edge (T', S''). Since CO_2 is total in this example, the construction terminates: $\mathcal{X} = \mathcal{P}_2$.

Formally, in addition to VIS_i and CO_i we construct sets $R_i = \{(T_k, S_k) \mid k = 1..i\}$, $i = 0..n$ that accumulate the edges enforced in the commit order at every step. The relations VIS_i and CO_i and the sets R_i are defined recursively as follows: we let VIS_i and CO_i be defined by (1) for $R = R_i$; we let $R_0 = \emptyset$ and $R_{i+1} = R_i \cup \{(T_i, S_i)\}$, where (T_i, S_i) is an arbitrary pair of transactions unrelated by CO_i; such a pair must exist if CO_i is not total. By Lemma 15, each $(\mathsf{VIS}_i, \mathsf{CO}_i)$ is a solution to the system of inequalities in Figure 3. It is easy to check that $\mathsf{CO}_{i+1} = (\mathsf{CO}_i \cup \{(T_i, S_i)\})^+$, $i = 0..(n-1)$. Since CO_0 is acyclic, by the choice of the edges (T_i, S_i) it follows that CO_i, $i = 1..n$ are acyclic as well. Hence, the sequence $\{\mathcal{P}_i\}_{i=0}^n$ constructed above satisfies the properties stated at the beginning of the proof, as required. $\square$

PROOF OF THEOREM 10(II). Consider $\mathcal{X} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO}) \in \mathsf{ExecSI}$. As follows from Lemma 12, VIS and CO give a solution to the system of inequalities of Figure 3 for $\mathsf{WR} = \mathsf{WR}_\mathcal{X}$, $\mathsf{WW} = \mathsf{WW}_\mathcal{X}$, $\mathsf{RW} = \mathsf{RW}_\mathcal{X}$. We now apply Lemma 15 for $R = \emptyset$; the minimality of the solution given by Lemma 15 implies that $((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW})\ ;\ \mathsf{RW}?)^+ \subseteq \mathsf{CO}$. Then $((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW})\ ;\ \mathsf{RW}?)^+$ is acyclic because so is CO. This establishes $\mathsf{graph}(\mathcal{X}) \in \mathsf{GraphSI}$. $\square$

5. TRANSACTION CHOPPING UNDER SI

In this section, we exploit our characterisation of SI in terms of dependency graphs to derive a static analysis that checks when transactions in an application executing under SI can be ***chopped*** [29] into sessions of smaller transactions without introducing new behaviours (the sessions are also called ***chains*** in this context [37]). To this end, the analysis must check that any SI execution of the application with chopped transactions can be *spliced* into an SI execution that has the same operations as the original one, but where all operations from each session are executed inside a single transaction. We first establish a *dynamic chopping criterion* that checks whether a single SI execution, represented by a dependency graph, is spliceable. From this we then derive a static analysis that checks whether this is the case for all executions produced by a given chopped application.

For a history $\mathcal{H}$, let $\approx_\mathcal{H} = \mathsf{SO}_\mathcal{H} \cup \mathsf{SO}_\mathcal{H}^{-1} \cup \{(T, T) \mid T \in \mathcal{T}_\mathcal{H}\}$ be the equivalence relation grouping transactions from the same session. We let $\boxed{T}_\mathcal{H}$ be the result of splicing all transactions in the session to which T belongs in $\mathcal{H}$ into a single transaction: $\boxed{T}_\mathcal{H} = (E, \mathsf{po})$, where $E = (\bigcup\{E_S \mid S \approx_\mathcal{H} T\})$ and

$$\mathsf{po} = \{(e, f) \mid (\exists S.\, e, f \in E_S \wedge e \xrightarrow{\mathsf{po}_S} f \wedge S \approx_\mathcal{H} T) \vee (\exists S, S'.\, e \in E_S \wedge f \in E_{S'} \wedge S \xrightarrow{\mathsf{SO}_\mathcal{H}} S' \wedge S' \approx_\mathcal{H} T)\}.$$

We let $\mathsf{splice}(\mathcal{H})$ be the history resulting from splicing all sessions in a history $\mathcal{H}$: $\mathsf{splice}(\mathcal{H}) = \left(\left\{\boxed{T}_\mathcal{H} \mid T \in \mathcal{T}_\mathcal{H}\right\}, \emptyset\right)$. A dependency graph $\mathcal{G} \in \mathsf{GraphSI}$ is ***spliceable*** if there exists a dependency graph $\mathcal{G}' \in \mathsf{GraphSI}$ such that $\mathcal{H}_{\mathcal{G}'} = \mathsf{splice}(\mathcal{H}_\mathcal{G})$. For a dependency graph $\mathcal{G}$, we let $\approx_\mathcal{G} = \approx_{\mathcal{H}_\mathcal{G}}$ and $\boxed{T}_\mathcal{G} = \boxed{T}_{\mathcal{H}_\mathcal{G}}$.

For example, the graph $\mathcal{G}_1$ in Figure 4 is not spliceable, because $\mathsf{splice}(\mathcal{H}_{\mathcal{G}_1}) \notin \mathsf{HistSI}$: informally, $\boxed{S}_{\mathcal{G}_1}$ observes the write by $\boxed{T}_{\mathcal{G}_1}$ to `acct1`, but not its write to `acct2`. On the other hand, let $\mathcal{G}_2$ be the graph obtained by removing the transactions S and S' from $\mathcal{G}_1$. Then $\mathcal{G}_2$ is spliceable, as witnessed by the graph $\mathcal{G}_2 \in \mathsf{GraphSI}$ with $\mathcal{H}_{\mathcal{G}_2} = \mathsf{splice}(\mathcal{H}_{\mathcal{G}_2})$ and only the edges $\boxed{T''}_{\mathcal{G}_2} \xrightarrow{\mathsf{RW}_{\mathcal{G}_2}} \boxed{T}_{\mathcal{G}_2}$ and $\boxed{S''}_{\mathcal{G}_2} \xrightarrow{\mathsf{RW}_{\mathcal{G}_2}} \boxed{T}_{\mathcal{G}_2}$.

Given a dependency graph $\mathcal{G}$, we let the ***dynamic chopping graph*** corresponding to $\mathcal{G}$ be the graph $\mathsf{DCG}(\mathcal{G})$ obtained by removing $\mathsf{WR}_\mathcal{G}, \mathsf{WW}_\mathcal{G}$ and $\mathsf{RW}_\mathcal{G}$ edges between transactions related by $\approx_\mathcal{G}$, and by extending $\mathcal{G}$ with edges in the reverse of the session order: $\mathsf{SO}_\mathcal{G}^{-1}$. We refer to the latter edges as ***predecessor*** edges, to those in $\mathsf{SO}_\mathcal{G}$ as ***successor*** edges, and to those in $(\mathsf{WR}_\mathcal{G} \cup \mathsf{WW}_\mathcal{G} \cup \mathsf{RW}_\mathcal{G}) \setminus \approx_\mathcal{G}$ as ***conflict*** edges. A cycle in a chopping graph $\mathsf{DCG}(\mathcal{G})$ is ***critical*** if: *(i)* it does not contain two occurrences of the same vertex; *(ii)* it contains a fragment of three consecutive edges of the form "conflict, predecessor, conflict"; and *(iii)* any two anti-dependency edges ($\mathsf{RW}_\mathcal{G} \setminus \approx_\mathcal{G}$) on the cycle are separated by at least one read ($\mathsf{WR}_\mathcal{G} \setminus \approx_\mathcal{G}$) or write ($\mathsf{WW}_\mathcal{G} \setminus \approx_\mathcal{G}$) dependency edge. Our ***dynamic chopping criterion*** is as follows.

THEOREM 16. *For $\mathcal{G} \in \mathsf{GraphSI}$, if $\mathsf{DCG}(\mathcal{G})$ contains no critical cycles, then $\mathcal{G}$ is spliceable.*

For example, the above graph $\mathcal{G}_2$ (Figure 4) contains no critical cycles, and the graph $\mathcal{G}_1$ contains a critical cycle

$$T' \xrightarrow{\mathsf{WR}_{\mathcal{G}_1}} S' \xrightarrow{\mathsf{SO}_{\mathcal{G}_1}} S \xrightarrow{\mathsf{RW}_{\mathcal{G}_1}} T \xrightarrow{\mathsf{SO}_{\mathcal{G}_1}^{-1}} T'.$$

To prove Theorem 16, we exhibit a particular dependency graph $\mathsf{splice}(\mathcal{G})$ such that $\mathsf{splice}(\mathcal{G}) \in \mathsf{GraphSI}$ and $\mathcal{H}_{\mathsf{splice}(\mathcal{G})} = \mathsf{splice}(\mathcal{H}_\mathcal{G})$. We define read dependencies $\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}$ by lifting those in $\mathsf{WR}_\mathcal{G}$ to spliced transactions:

$$\forall T, S \in \mathcal{T}_\mathcal{G}.\, \forall x \in \mathsf{Obj}.\, \boxed{T}_\mathcal{G} \xrightarrow{\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_\mathcal{G} \iff \boxed{T}_\mathcal{G} \neq \boxed{S}_\mathcal{G} \wedge T \xrightarrow{\approx_\mathcal{G}\,;\,\mathsf{WR}_\mathcal{G}(x)\,;\,\approx_\mathcal{G}} S. \quad (2)$$

We define $\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}$ similarly and derive $\mathsf{RW}_{\mathsf{splice}(\mathcal{G})}$ from $\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}$ and $\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}$ as in Definition 5. As the following lemma shows, $\mathsf{RW}_{\mathsf{splice}(\mathcal{G})}$ defined in this way can be decomposed into a form similar to (2).

LEMMA 17. *Let $\mathcal{G} \in \mathsf{GraphSI}$ be such that $\mathsf{DCG}(\mathcal{G})$ contains no critical cycles. Then*

$$\forall T, S \in \mathcal{T}_\mathcal{G}.\, \forall x \in \mathsf{Obj}.\, \boxed{T}_\mathcal{G} \xrightarrow{\mathsf{RW}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_\mathcal{G} \implies \boxed{T}_\mathcal{G} \neq \boxed{S}_\mathcal{G} \wedge T \xrightarrow{\approx_\mathcal{G}\,;\,\mathsf{RW}_\mathcal{G}(x)\,;\,\approx_\mathcal{G}} S.$$

To prove Theorem 16, we assume $\mathsf{splice}(\mathcal{G}) \notin \mathsf{GraphSI}$ and use Theorem 9 to obtain a cycle in $(\mathsf{WR}_{\mathsf{splice}(\mathcal{G})} \cup \mathsf{WW}_{\mathsf{splice}(\mathcal{G})})\ ;$

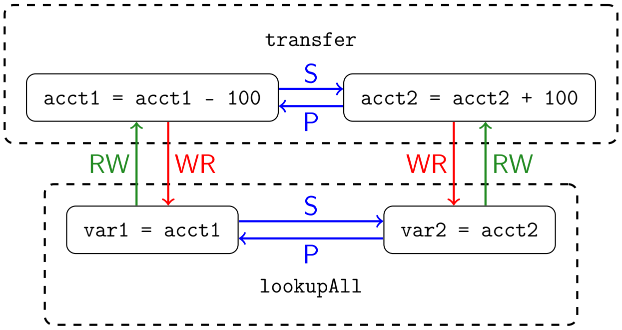

Figure 5: The static chopping graph of the programs $\{\texttt{transfer}, \texttt{lookupAll}\}$ from Figure 4. Dashed boxes group program pieces into sessions.

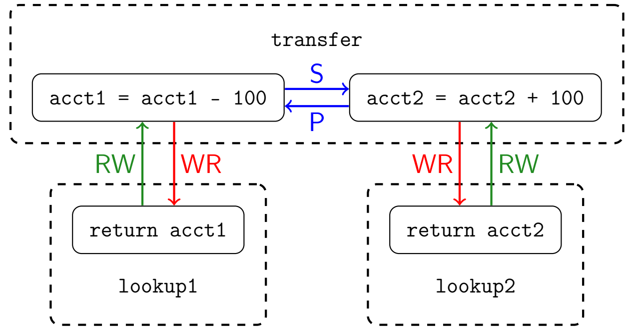

Figure 6: The static chopping graph of the programs $\{\texttt{transfer}, \texttt{lookup1}, \texttt{lookup2}\}$ from Figure 4.

$\mathsf{RW}_{\mathsf{splice}(\mathcal{G})}$?. We then use (2) and Lemma 17 to decompose this cycle into a critical cycle in $\mathsf{DCG}(\mathcal{G})$, yielding a contradiction.

We now derive a static analysis from Theorem 16. Assume a set of ***programs*** $\mathcal{P} = \{P_1, P_2, \dots\}$, each defining the code of sessions resulting from chopping the code of a single transaction. We leave the precise syntax of the programs unspecified, but assume that each P_i consists of k_i ***program pieces***, defining the code of the transactions in the sessions. We further assume that we are given the sets R^i_j and W^i_j of all objects that can respectively be read and written by the j-th piece of P_i. For example, the program `transfer` in Figure 4 consists of two pieces; the first one has the read and write sets equal to $\{\texttt{acct1}\}$ and the second, to $\{\texttt{acct2}\}$. The program `lookup1` consists of a single piece with the read set $\{\texttt{acct1}\}$ and the write set $\emptyset$.

Following Shasha et al. [29], we make certain assumptions about the way clients execute programs. We assume that, if a transaction initiated by a program piece aborts, it will be resubmitted repeatedly until it commits, and, if a piece is aborted due to system failure, it will be restarted. We also assume that the client does not abort transactions explicitly.

A history $\mathcal{H}$ ***can be produced*** by the programs $\mathcal{P}$, if there is a one-to-one correspondence between every session in $\mathcal{H}$ and a program $P_i \in \mathcal{P}$ whose read and write sets cover the sets of objects read or written by the corresponding transactions in the session. For example, the history in Figure 4 can be produced by the programs in the figure. The chopping defined by the programs $\mathcal{P}$ is ***correct*** if every dependency graph $\mathcal{G} \in \mathsf{GraphSI}$, where $\mathcal{H}_{\mathcal{G}}$ can be produced by $\mathcal{P}$, is spliceable.

We check the correctness of $\mathcal{P}$ using its ***static chopping graph*** $\mathsf{SCG}(\mathcal{P})$. It is a directed graph whose nodes are pairs of indices identifying the pieces in $\mathcal{P}$: $\{(i, j) \mid i = 1..|\mathcal{P}|,\ j = 1..k_i\}$. We have an edge $((i_1, j_1), (i_2, j_2))$ if and only if one of the following holds: $i_1 = i_2$ and $j_1 < j_2$ (a ***successor*** edge); $i_1 = i_2$ and $j_1 > j_2$ (a ***predecessor*** edge); $i_1 \neq i_2$ and $W^{i_1}_{j_1} \cap R^{i_2}_{j_2} \neq \emptyset$ (a ***read dependency*** edge); $i_1 \neq i_2$ and $W^{i_1}_{j_1} \cap W^{i_2}_{j_2} \neq \emptyset$ (a ***write dependency*** edge); or $i_1 \neq i_2$ and $R^{i_1}_{j_1} \cap W^{i_2}_{j_2} \neq \emptyset$ (an ***anti-dependency*** edge). The notion of a critical cycle introduced above for dynamic graphs is also applicable to static ones. The edge set of a static graph $\mathsf{SCG}(\mathcal{P})$ over-approximates the edge sets of dynamic graphs $\mathsf{DCG}(\mathcal{G})$ corresponding to dependency graphs $\mathcal{G}$ produced by the programs $\mathcal{P}$. From this observation and Theorem 16 we easily get our static analysis.

COROLLARY 18. *The chopping defined by $\mathcal{P}$ is correct if $\mathsf{SCG}(\mathcal{P})$ contains no critical cycles.*

In Figure 5 we show the static chopping graph of the programs $\{\texttt{transfer}, \texttt{lookupAll}\}$, which contains a critical cycle:

$$\begin{aligned}
&(\texttt{var1} = \texttt{acct1}) \xrightarrow{\mathsf{RW}} (\texttt{acct1} = \texttt{acct1} - 100) \xrightarrow{\mathsf{S}} \\
&(\texttt{acct2} = \texttt{acct2} + 100) \xrightarrow{\mathsf{WR}} \\
&(\texttt{var2} = \texttt{acct2}) \xrightarrow{\mathsf{P}} (\texttt{var1} = \texttt{acct1}).
\end{aligned}$$

In fact, since the dependency graph in Figure 4 is not spliceable, the chopping defined by the above programs is incorrect. In Figure 6 we show the static chopping graph of the programs $\{\texttt{transfer}, \texttt{lookup1}, \texttt{lookup2}\}$. This graph contains no critical cycles, and hence, the chopping defined by these programs is correct: they behave the same as when `transfer` is implemented by a single transaction.

Our SI characterisation is instrumental in deriving the above static analysis due to the ease of splicing a dependency graph (cf. (2)). As we explain in §B, splicing abstract executions directly would be problematic.

In §B, we also show that the conditions on chopping required by Corollary 18 are laxer than those of the analysis for serializability [29], but stricter than those for parallel SI [11]. In particular, this implies that the classical transaction chopping analysis for serializability is also sound for SI. This result is non-trivial: the correctness of a chopping requires that the set of histories produced by the chopped program be included into the set of histories produced by the original program. Enlarging both sets when switching from serializability to SI may not preserve the inclusion.

6. ROBUSTNESS CRITERIA FOR SI

We now consider another type of a static analysis that checks whether an application is ***robust*** against weakening consistency: executing it under a weak consistency model produces the same client-observable behaviour as executing it under a stronger one.

6.1 Robustness against SI

We first show that our SI characterisation allows deriving a variant of an existing analysis that checks whether an application executing under SI behaves the same as when executing under serializability [18][1] (robustness ***against SI***). For this the analysis checks that the application code may produce no histories in $\mathsf{HistSI} \setminus \mathsf{HistSER}$. Like for transaction chopping (§5), we first establish a ***dynamic robustness criterion*** that checks whether a single execution, represented by a dependency graph, is in $\mathsf{GraphSI} \setminus \mathsf{GraphSER}$. This easily follows from Theorems 8 and 10.

[1]Our analysis does not take into account the notion of a *vulnerable* dependency from [18].

THEOREM 19. *For any $\mathcal{G}$, we have $\mathcal{G} \in \mathsf{GraphSI} \setminus \mathsf{GraphSER}$ if and only if $\mathcal{T}_\mathcal{G} \models$ INT, $\mathcal{G}$ contains a cycle, and all its cycles have at least two adjacent anti-dependency edges.*

Fekete et al. previously established a result corresponding to the "only if" direction of the above theorem [18]. The "if" direction strengthens their result by showing that the criterion in the theorem is complete for checking whether a given dependency graph is admitted by SI, but not serializability.

The dependency graph $\mathcal{G}_3$ of the write skew anomaly in Figure 2(d) contains a cycle: $T_1 \xrightarrow{\mathsf{RW}_{\mathcal{G}_3}} T_2 \xrightarrow{\mathsf{RW}_{\mathcal{G}_3}} T_1$. Furthermore, it is easy to see that all its cycles have two adjacent anti-dependencies, so that $\mathcal{G}_3 \in \mathsf{GraphSI} \setminus \mathsf{GraphSER}$. The dependency graph $\mathcal{G}_1$ from Figure 4 does not contain any cycles, so that $\mathcal{G}_1 \notin \mathsf{GraphSI} \setminus \mathsf{GraphSER}$; in fact, $\mathcal{G}_1 \in \mathsf{GraphSER}$. Finally, the dependency graph $\mathcal{G}_4$ of the long fork anomaly in Figure 4 contains a cycle, but without two adjacent anti-dependencies; hence, $\mathcal{G}_4 \notin \mathsf{GraphSI} \setminus \mathsf{GraphSER}$, and in fact, $\mathcal{G}_4 \notin \mathsf{GraphSI}$.

We can derive a static analysis from Theorem 19 similarly to how it was done in §5. Namely, the analysis assumes that the code of transactions in an application is defined by a set of programs $\mathcal{P}$ with given read and write sets. Based on these sets, it constructs a ***static dependency graph***, over-approximating possible dependencies that can exist in executions of the programs $\mathcal{P}$. The analysis then checks that the graph has no cycles with at least two adjacent anti-dependency edges. By Theorem 19 this implies that the programs $\mathcal{P}$ produce no histories in $\mathsf{HistSI} \setminus \mathsf{HistSER}$, and hence, the corresponding application is robust against SI. Note that the dependency graphs characterisation of consistency models greatly facilitates deriving the above static analysis, since the characterisations allow us to easily establish correspondences between executions on different models with the same histories.

6.2 Robustness against parallel SI towards SI

We now use our SI characterisation to derive a static analysis that checks whether an application executing under parallel SI [31] behaves the same as when executing under the classical SI (robustness ***against parallel SI towards SI***). To specify parallel SI in the framework of §2, we drop the axiom PREFIX, while still requiring visibility to be transitive, a property that we refer to as TRANSVIS [10].

DEFINITION 20. *The sets of executions and histories **allowed by parallel SI** are:*

$$\begin{aligned} \mathsf{ExecPSI} = \{\mathcal{X} \mid \mathcal{X} \models \textsc{Int} \wedge \textsc{Ext} \wedge \textsc{Session} \wedge \\ \textsc{TransVis} \wedge \textsc{NoConflict}\}; \\ \mathsf{HistPSI} = \{\mathcal{H} \mid \exists \mathsf{VIS}, \mathsf{CO}.\, (\mathcal{H}, \mathsf{VIS}, \mathsf{CO}) \in \mathsf{ExecSI}\}. \end{aligned}$$

Note that this specification essentially does not use the commit order CO: according to NOCONFLICT, its edges used in EXT are uniquely determined by VIS.

The axiom TRANSVIS ensures that transactions ordered by VIS are observed by others in this order. However, it allows two transactions unrelated by VIS to be observed in different orders; in particular, parallel SI allows the long fork anomaly of Figure 2(c), disallowed by the axiom PREFIX in SI.

Following [11, extended version, Lemma 14], we can give a characterisation of parallel SI in terms of dependency graphs.

THEOREM 21. *Let*

$$\begin{aligned} \mathsf{GraphPSI} = \{\mathcal{G} \mid (\mathcal{T}_\mathcal{G} \models \textsc{Int}) \wedge \\ (((\mathsf{SO}_\mathcal{G} \cup \mathsf{WR}_\mathcal{G} \cup \mathsf{WW}_\mathcal{G})^+ \,;\, \mathsf{RW}_\mathcal{G}?) \text{ is irreflexive})\}. \end{aligned}$$

Then

$$\begin{aligned} \mathsf{HistPSI} = \{\mathcal{H} \mid \exists \mathsf{WR}, \mathsf{WW}, \mathsf{RW}. \\ (\mathcal{H}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW}) \in \mathsf{GraphPSI}\}. \end{aligned}$$

Thus, parallel SI is characterised by dependency graphs that contain only cycles with at least two anti-dependency edges. For example, consider the dependency graph $\mathcal{G}_4$ in Figure 2(c). It is easy to see that all its cycles contain at least two anti-dependencies, and therefore $\mathcal{G}_4 \in \mathsf{GraphPSI}$. On the other hand, let $\mathcal{G}_5$ be the dependency graph in Figure 2(b). The graph $\mathcal{G}_5$ contains a cycle with exactly one anti-dependency ($T_1 \xrightarrow{\mathsf{WW}_{\mathcal{G}_5}} T_2 \xrightarrow{\mathsf{RW}_{\mathcal{G}_5}} T_1$), and therefore $\mathcal{G}_5 \notin \mathsf{GraphPSI}$. As a corollary of Theorems 9 and 21, we obtain a dynamic robustness criterion that checks whether a given dependency graph is in $\mathsf{GraphPSI} \setminus \mathsf{GraphSI}$.

THEOREM 22. *For any $\mathcal{G}$, we have $\mathcal{G} \in \mathsf{GraphPSI} \setminus \mathsf{GraphSI}$ if and only if $\mathcal{T}_\mathcal{G} \models$ INT, $\mathcal{G}$ contains at least one cycle with no adjacent anti-dependency edges, and all its cycles have at least two anti-dependency edges.*

For example, we have already noted that in the dependency graph $\mathcal{G}_4$ of the long fork anomaly all cycles have at least two anti-dependencies. Furthermore, $\mathcal{G}_4$ also has a cycle with no adjacent anti-dependencies: $T_1 \xrightarrow{\mathsf{WR}_{\mathcal{G}_4}} T_3 \xrightarrow{\mathsf{RW}_{\mathcal{G}_4}} T_2 \xrightarrow{\mathsf{WR}_{\mathcal{G}_4}} T_4 \xrightarrow{\mathsf{RW}_{\mathcal{G}_4}} T_1$, so that $\mathcal{G}_4 \in \mathsf{GraphPSI} \setminus \mathsf{GraphSI}$. The dependency graph $\mathcal{G}_3$ of the write skew anomaly in Figure 2(d) contains only cycles with at least two adjacent anti-dependencies, so that $\mathcal{G}_3 \notin \mathsf{GraphPSI} \setminus \mathsf{GraphSI}$; in fact, $\mathcal{G}_3 \in \mathsf{GraphSI}$. The dependency graph $\mathcal{G}_5$ of the lost update anomaly contains a cycle with exactly one anti-dependency, so that $\mathcal{G}_5 \notin \mathsf{GraphPSI} \setminus \mathsf{GraphSI}$; in fact, $\mathcal{G}_5 \notin \mathsf{GraphPSI}$.

From Theorem 22 it follows that the desired static analysis can check that the static dependency graph of an application contains no cycles where there are at least two anti-dependency edges and no two anti-dependency edges are adjacent.

7. RELATED WORK

Snapshot isolation was originally defined by an idealised algorithm formulated in terms of implementation-level concepts [6]. Since then there have been proposals of more declarative SI specifications [2, 10, 27], one of which [10] was our starting point (§2). However, these specifications are stated in terms of relations which make it challenging to obtain results such as transaction chopping and robustness analyses.

Fekete et al. [18] proposed the analysis for robustness against SI that we considered in §6.1. To this end, they have proved a fact roughly equivalent to our completeness result (Theorem 10(ii)), but they did not establish an analogue of our soundness result (Theorem 10(i)). The latter more challenging result is the one that is needed to obtain analyses for transaction chopping under SI and for robustness against parallel SI towards SI: both require proving that an execution with a particular dependency graph is in SI, rather than the other way round. We also hope that our specification of SI will be beneficial in other domains where dependency graphs have been useful, such as run-time monitoring [9, 36] and proving the correctness of concurrency-control algorithms [15, 35]. Finally, we expect that the approach to constructing a total commit order from transactional dependencies in the proof of our soundness theorem can be used to give dependency graph characterisations to other consistency models whose formulation includes similar total orders, such as *prefix consistency* [33].

The constraint on dependency graphs that we use to characterise SI also arose in the work of Lin et al. [23], who used it to formulate conditions under which a replicated database guarantees SI provided every one of its replicas does so. In comparison to them, we solve a more general problem of characterising SI regardless of how it is implemented and handle a variant of SI that does not require transactions to see the latest snapshot.

Transaction chopping has recently received a lot of attention. In particular, researchers have demonstrated that transactions arising in web applications can be chopped in a way that drastically improves their performance when executed under serializability [25, 35, 37]. There have also been proposals of consistency models for transactional memory that weaken consistency guarantees in a way similar to chopping [4, 19, 34]. Our chopping analysis enables bringing these benefits to transactional systems providing SI. We have previously proposed a chopping analysis for parallel SI [11], which also relies on a dependency graph characterisation of this consistency model (Theorem 21). But since parallel SI can be formulated without using an analogue of SI's commit order, its dependency graph characterisation did not present the challenges that we had to deal with when establishing our soundness theorem.

Acknowledgements. We thank Hongseok Yang, who participated in early stages of this work. We also thank Hagit Attiya, Ricardo Jiménez-Peris and Pierre Sutra for comments that helped improve the paper. This work was supported by an EU project ADVENT.

References

[1] The Clojure language: Refs and transactions. http://clojure.org/refs.
[2] A. Adya. Weak consistency: A generalized theory and optimistic implementations for distributed transactions. PhD thesis, MIT, 1999.
[3] A. Adya, B. Liskov, and P. E. O'Neil. Generalized isolation level definitions. In *ICDE*, 2000.
[4] Y. Afek, H. Avni, and N. Shavit. Towards consistency oblivious programming. In *OPODIS*, 2011.
[5] P. Bailis, A. Fekete, A. Ghodsi, J. M. Hellerstein, and I. Stoica. Scalable atomic visibility with RAMP transactions. In *SIGMOD*, 2014.
[6] H. Berenson, P. Bernstein, J. Gray, J. Melton, E. O'Neil, and P. O'Neil. A critique of ANSI SQL isolation levels. In *SIGMOD*, 1995.
[7] P. A. Bernstein, V. Hadzilacos, and N. Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987.
[8] A. Bieniusa and T. Fuhrmann. Consistency in hindsight: A fully decentralized STM algorithm. In *IPDPS*, 2010.
[9] M. J. Cahill, U. Röhm, and A. D. Fekete. Serializable isolation for snapshot databases. *ACM Trans. Database Syst.*, 34(4), 2009.
[10] A. Cerone, G. Bernardi, and A. Gotsman. A framework for transactional consistency models with atomic visibility. In *CONCUR*. Dagstuhl, 2015.
[11] A. Cerone, A. Gotsman, and H. Yang. Transaction chopping for parallel snapshot isolation. In *DISC*, 2015. Extended version available from www.software.imdea.org/~gotsman.
[12] K. Daudjee and K. Salem. Lazy database replication with ordering guarantees. In *ICDE*, 2004.
[13] K. Daudjee and K. Salem. Lazy database replication with snapshot isolation. In *VLDB*, 2006.
[14] R. J. Dias, J. M. Lourenço, and N. Preguiça. Efficient and correct transactional memory programs combining snapshot isolation and static analysis. In *HotPar*, 2011.
[15] N. Diegues and P. Romano. Time-warp: Lightweight abort minimization in transactional memory. In *PPoPP*, 2014.
[16] S. Doherty, L. Groves, V. Luchangco, and M. Moir. Towards formally specifying and verifying transactional memory. *Formal Aspects of Computing*, 25(5), 2013.
[17] S. Elnikety, W. Zwaenepoel, and F. Pedone. Database replication using generalized snapshot isolation. In *SRDS*, 2005.
[18] A. Fekete, D. Liarokapis, E. O'Neil, P. O'Neil, and D. Shasha. Making snapshot isolation serializable. *ACM Trans. Database Syst.*, 30(2), 2005.
[19] P. Felber, V. Gramoli, and R. Guerraoui. Elastic transactions. In *DISC*, 2009.
[20] R. Guerraoui and M. Kapalka. On the correctness of transactional memory. In *PPoPP*, 2008.
[21] M. Herlihy and J. E. B. Moss. Transactional memory: Architectural support for lock-free data structures. In *ISCA*, 1993.
[22] S. Jorwekar, A. Fekete, K. Ramamritham, and S. Sudarshan. Automating the detection of snapshot isolation anomalies. In *VLDB*, 2007.
[23] Y. Lin, B. Kemme, R. Jiménez-Peris, M. Patiño-Martínez, and J. E. Armendáriz-Iñigo. Snapshot isolation and integrity constraints in replicated databases. *ACM Trans. Database Syst.*, 34(2), 2009.
[24] H. Litz, D. Cheriton, A. Firoozshahian, O. Azizi, and J. P. Stevenson. SI-TM: Reducing transactional memory abort rates through snapshot isolation. In *ASPLOS*, 2014.
[25] S. Mu, Y. Cui, Y. Zhang, W. Lloyd, and J. Li. Extracting more concurrency from distributed transactions. In *OSDI*, 2014.
[26] D. Peng and F. Dabek. Large-scale incremental processing using distributed transactions and notifications. In *OSDI*, 2010.
[27] M. Saeida Ardekani, P. Sutra, M. Shapiro, and N. Preguiça. On the scalability of snapshot isolation. In *Euro-Par*, 2013.
[28] D. Serrano, M. Patiño-Martínez, R. Jiménez-Peris, and B. Kemme. Boosting database replication scalability through partial replication and 1-copy-snapshot-isolation. In *PRDC*, 2007.
[29] D. Shasha, F. Llirbat, E. Simon, and P. Valduriez. Transaction chopping: Algorithms and performance studies. *ACM Trans. Database Syst.*, 20(3), 1995.
[30] D. Shasha and M. Snir. Efficient and correct execution of parallel programs that share memory. *ACM Trans. Program. Lang. Syst.*, 10(2), 1988.
[31] Y. Sovran, R. Power, M. K. Aguilera, and J. Li. Transactional storage for geo-replicated systems. In *SOSP*, 2011.
[32] D. B. Terry, A. J. Demers, K. Petersen, M. Spreitzer, M. Theimer, and B. W. Welch. Session guarantees for weakly consistent replicated data. In *PDIS*, 1994.
[33] D. B. Terry, V. Prabhakaran, R. Kotla, M. Balakrishnan, M. K. Aguilera, and H. Abu-Libdeh. Consistency-based service level agreements for cloud storage. In *SOSP*, 2013.
[34] L. Xiang and M. L. Scott. Software partitioning of hardware transactions. In *PPoPP*, 2015.
[35] C. Xie, C. Su, C. Littley, L. Alvisi, M. Kapritsos, and Y. Wang. High-performance ACID via modular concurrency control. In *SOSP*, 2015.
[36] K. Zellag and B. Kemme. Consistency anomalies in multi-tier architectures: Automatic detection and prevention. *The VLDB Journal*, 23(1), 2014.
[37] Y. Zhang, R. Power, S. Zhou, Y. Sovran, M. Aguilera, and J. Li. Transaction chains: Achieving serializability with low latency in geo-distributed storage systems. In *SOSP*, 2013.

APPENDIX

A. ADDITIONAL PROOFS

A.1 Proof of Proposition 7

We prove a more general result, from which Proposition 7 follows immediately.

PROPOSITION 23. *Let $\mathcal{X} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO})$ be an execution such that $\mathcal{X} \models$ EXT. Then $\mathcal{G} = (\mathcal{T}, \mathsf{SO}, \mathsf{WR}_{\mathcal{X}}, \mathsf{WW}_{\mathcal{X}}, \mathsf{RW}_{\mathcal{X}})$ is a dependency graph.*

PROOF. It suffices to show that $\mathcal{G}$ satisfies all the constraints imposed by Definition 6:

- $\forall T, S \in \mathcal{T}. \forall x. T \xrightarrow{\mathsf{WR}_{\mathcal{X}}(x)} S \implies \exists n. T \neq S \wedge T \vdash \mathtt{write}(x, n) \wedge S \vdash \mathtt{read}(x, n)$. Let T, S, x be such that $T \xrightarrow{\mathsf{WR}_{\mathcal{X}}(x)} S$. By Definition 5, $S \vdash \mathtt{read}(x, n)$ for some $n \in \mathbb{N}$ and $T = \max_{\mathsf{CO}}(\mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x)$. Then $T \neq S$. Since $\mathcal{X} \models$ EXT, we have $T \vdash \mathtt{write}(x, n)$.

- $\forall S \in \mathcal{T}. \forall x. S \vdash \mathtt{read}(x, \_) \implies \exists T. T \xrightarrow{\mathsf{WR}_{\mathcal{X}}(x)} S$. Suppose that $S \vdash \mathtt{read}(x, \_)$. Since $\mathcal{X} \models$ EXT, the set $\mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x$ is non-empty, so that $T = \max_{\mathsf{CO}}(\mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x)$ is defined. By Definition 5, $T \xrightarrow{\mathsf{WR}_{\mathcal{X}}(x)} S$.

- $\forall T, T', S \in \mathcal{T}. \forall x. (T \xrightarrow{\mathsf{WR}(x)} S \wedge T' \xrightarrow{\mathsf{WR}(x)} S) \implies T = T'$. This holds because by Definition 5 we have $T = \max_{\mathsf{CO}}(\mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x) = T'$.

- for any x, $\mathsf{WW}_{\mathcal{G}}(x)$ is a total order over $\mathsf{WriteTx}_x$. Fix an object x and recall that $T \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S$ if and only if $T, S \in \mathsf{WriteTx}_x$ and $T \xrightarrow{\mathsf{CO}} S$. Since CO is transitive and irreflexive, so is $\mathsf{WW}_{\mathcal{G}}(x)$. Since CO is total over $\mathcal{T} \times \mathcal{T}$, so is $\mathsf{WW}_{\mathcal{G}}(x)$ over $\mathcal{T} \cap \mathsf{WriteTx}_x$.

- $T \xrightarrow{\mathsf{RW}_{\mathcal{X}}(x)} S \iff T \neq S \wedge \exists T'. T' \xrightarrow{\mathsf{WR}_{\mathcal{X}}(x)} T \wedge T' \xrightarrow{\mathsf{WW}_{\mathcal{X}}(x)} S$. This follows directly from Definition 5.

□

A.2 Proof of Proposition 14

Fix $\mathcal{X} = (\mathcal{T}, \mathsf{SO}, \mathsf{VIS}, \mathsf{CO}) \in \mathsf{ExecSI}$ and $T, S \in \mathcal{T}$. Let $\mathsf{graph}(\mathcal{X}) = (\mathcal{T}, \mathsf{SO}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$. We illustrate the proof in Figure 7.

"$\Longrightarrow$". Assume $S \xrightarrow{\mathsf{RW}} T$. Then for some $T' \in \mathcal{T}$ we have $T' \xrightarrow{\mathsf{WR}(x)} S$ and $T' \xrightarrow{\mathsf{WW}(x)} T$. This implies $S \neq T$, $S \vdash \mathtt{read}(x, \_)$ and $T \vdash \mathtt{write}(x, \_)$. Assume that $T \xrightarrow{\mathsf{VIS}} S$. Then $T \in \mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x$. Since $T' \xrightarrow{\mathsf{WW}(x)} T$, we have $T' \xrightarrow{\mathsf{CO}} T$. But then $T' \neq \max_{\mathsf{CO}}(\mathsf{VIS}^{-1}(S) \cap \mathsf{WriteTx}_x)$, contradicting $T' \xrightarrow{\mathsf{WR}(x)} S$. Hence, we cannot have $T \xrightarrow{\mathsf{VIS}} S$.

"$\Longleftarrow$". Assume $S \vdash \mathtt{read}(x, \_)$, $T \vdash \mathtt{write}(x, \_)$ and $\neg(T \xrightarrow{\mathsf{VIS}} S)$ for some $x \in \mathsf{Obj}$. Let T' be the unique transaction such that $T' \xrightarrow{\mathsf{WR}(x)} S$; then $T' \vdash \mathtt{write}(x, \_)$. Since CO is total, we must have one of $T' = T$, $T \xrightarrow{\mathsf{CO}} T'$ or $T' \xrightarrow{\mathsf{CO}} T$. We cannot have $T' = T$, since then we would get $T \xrightarrow{\mathsf{VIS}} S$ from $T' \xrightarrow{\mathsf{WR}(x)} S$. We also cannot have $T \xrightarrow{\mathsf{CO}} T'$, since then we would get $T \xrightarrow{\mathsf{VIS}} S$ by PREFIX. Therefore, $T' \xrightarrow{\mathsf{CO}} T$ and, hence, $T' \xrightarrow{\mathsf{WW}(x)} T$. But then $S \xrightarrow{\mathsf{RW}(x)} T$. □

A.3 Proof of Lemma 15

We first prove that the relations in the statement of the lemma are indeed a solution to the system of inequalities in Figure 3:

1.

$$\begin{aligned} \mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW} &= (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \\ &\subseteq (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \\ &= \mathsf{VIS}. \end{aligned}$$

2.

$$\begin{aligned} \mathsf{CO} \cdot \mathsf{VIS} &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \cdot \\ &\quad\; (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \\ &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \\ &\subseteq (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \\ &= \mathsf{VIS}. \end{aligned}$$

3.

$$\begin{aligned} \mathsf{VIS} &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \\ &\subseteq (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot ((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \\ &\subseteq (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R) \\ &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \\ &= \mathsf{CO}. \end{aligned}$$

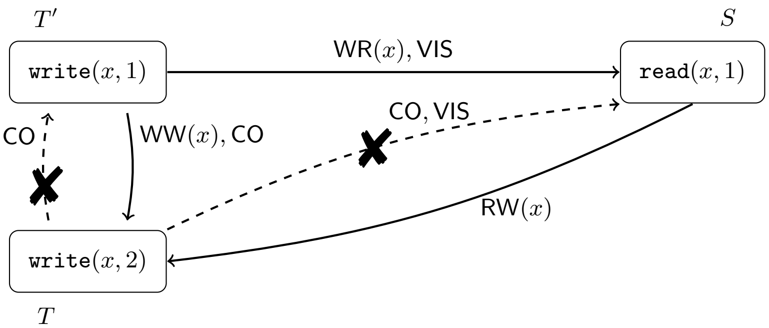

Figure 7: An illustration of the proof of Proposition 14.

4.

$$\begin{aligned} \mathsf{CO} \cdot \mathsf{CO} &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \cdot (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \\ &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \\ &= \mathsf{CO}. \end{aligned}$$

5.

$$\begin{aligned} \mathsf{VIS} \cdot \mathsf{RW} &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW} \\ &\subseteq (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot ((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \\ &\subseteq (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^* \cdot (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R) \\ &= (((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \\ &= \mathsf{CO}. \end{aligned}$$

Next, consider VIS' and CO' such that they satisfy (S1)-(S5) and $R \subseteq \mathsf{CO}'$. We first show that $\mathsf{CO} \subseteq \mathsf{CO}'$. By (S1) we have

$$\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW} \subseteq \mathsf{VIS}'. \tag{3}$$

Then by (S3) and (S5) we get

$$(\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}? \subseteq \mathsf{CO}'.$$

Since $R \subseteq \mathsf{CO}'$, this implies

$$((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R \subseteq \mathsf{CO}',$$

and by (S4) we get

$$(((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \subseteq (\mathsf{CO}')^+ \subseteq \mathsf{CO}'.$$

But this is exactly $\mathsf{CO} \subseteq \mathsf{CO}'$.

To prove that $\mathsf{VIS} \subseteq \mathsf{VIS}'$, we rewrite VIS as

$$\mathsf{VIS} = (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cup ((((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW})),$$

and we prove that both parameters of the union are included in VIS'. This establishes $\mathsf{VIS} \subseteq \mathsf{VIS}'$ because, according to (S2) and (S3), VIS' is transitive. We have already proved (3). We also have

$$(((\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \cdot \mathsf{RW}?) \cup R)^+ \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) = \mathsf{CO} \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}).$$

Since $\mathsf{CO} \subseteq \mathsf{CO}'$ and $\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW} \subseteq \mathsf{VIS}'$, by (S2) we get

$$\mathsf{CO} \cdot (\mathsf{SO} \cup \mathsf{WR} \cup \mathsf{WW}) \subseteq \mathsf{CO}' \cdot \mathsf{VIS}' \subseteq \mathsf{VIS}',$$

as required. $\square$

A.4 Proof of Lemma 17

Let $\mathcal{G} \in \mathsf{GraphSI}$ and suppose that $\boxed{T}_{\mathcal{G}} \xrightarrow{\mathsf{RW}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$ for some $T, S \in \mathcal{T}_{\mathcal{G}}$ and $x \in \mathsf{Obj}$. By definition, $\boxed{T}_{\mathcal{G}} \neq \boxed{S}_{\mathcal{G}}$ and there exists $V \in \mathcal{T}_{\mathcal{G}}$ such that $\boxed{V}_{\mathcal{G}} \xrightarrow{\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{T}_{\mathcal{G}}$, $\boxed{V}_{\mathcal{G}} \xrightarrow{\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$. That is, there exist $T' \approx_{\mathcal{G}} T, V' \approx_{\mathcal{G}} V, V'' \approx_{\mathcal{G}} V$ and $S'' \approx_{\mathcal{G}} S$ such that $V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} T'$, $V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$. We show that $T' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} S''$, so that $T \xrightarrow{\approx_{\mathcal{G}}\, ;\, \mathsf{RW}_{\mathcal{G}}(x)\, ;\, \approx_{\mathcal{G}}} S$.

First note that from $V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} T'$ and $V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$ we can infer $V' \vdash \mathtt{write}(x, \_)$, $V'' \vdash \mathtt{write}(x, \_)$ and $S'' \vdash \mathtt{write}(x, \_)$. Since $V' \approx_{\mathcal{G}} V \approx_{\mathcal{G}} V''$, one of the following must be true: $V' = V''$, $V' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} V''$ or $V'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} V'$. We handle these three cases separately, as illustrated in Figure 8.

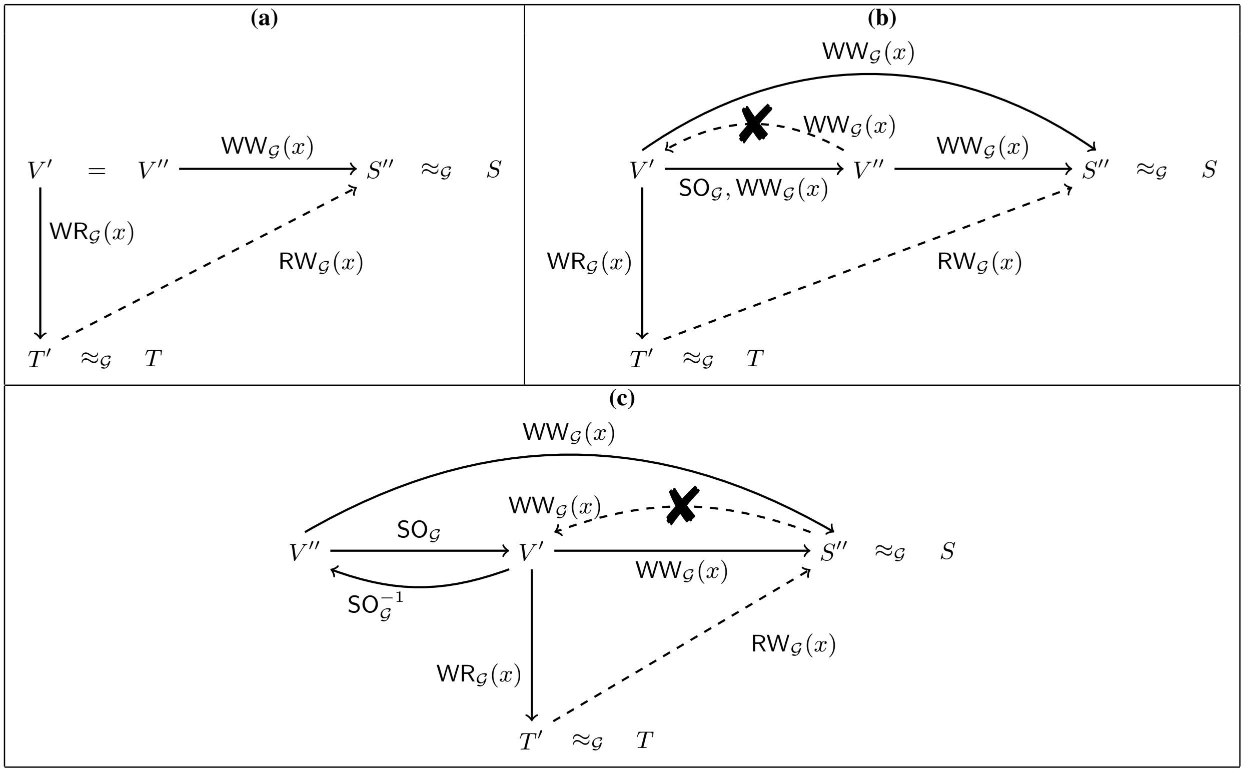

Figure 8: Graphical representation of different cases in the proof of Lemma 17.

(a) $V' = V''$. Then $V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} T'$ and $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$, from which $T' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} S''$ follows.

(b) $V' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} V''$. We have $V' \vdash \mathtt{write}(x, \_)$ and $V'' \vdash \mathtt{write}(x, \_)$. The assumption $\mathcal{G} \in \mathsf{GraphSI}$ mandates that $\neg(V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V')$: otherwise we would have the cycle $V' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$, contradicting $\mathcal{G} \in \mathsf{GraphSI}$. Hence, $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V''$. We thus have $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$ and by transitivity of $\mathsf{WW}_{\mathcal{G}}(x)$ we obtain $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$. We have established $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$ and $V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} T'$, so that $T' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} S''$.

(c) $V'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} V'$. Since $V' \approx_{\mathcal{G}} V$, we have $\boxed{V'}_{\mathcal{G}} = \boxed{V}_{\mathcal{G}}$; similarly, $\boxed{S''}_{\mathcal{G}} = \boxed{S}_{\mathcal{G}}$. Therefore we have $\boxed{V'}_{\mathcal{G}} \xrightarrow{\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S''}_{\mathcal{G}}$, which implies $\boxed{V'}_{\mathcal{G}} \neq \boxed{S''}_{\mathcal{G}}$, or equivalently $V' \not\approx_{\mathcal{G}} S''$. In particular $V' \neq S''$. We have already observed that $V' \vdash \mathtt{write}(x, \_)$ and $S'' \vdash \mathtt{write}(x, \_)$; since $V' \neq S''$, we must have either $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$ or $S'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$. However, in the latter case we would have the critical cycle $S'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V' \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$ in $\mathsf{DCG}(\mathcal{G})$, which contradicts the hypothesis of the lemma. Therefore, we must have $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$. Since we also have $V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} T'$, it follows that $T' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} S''$.

A.5 Proof of Theorem 16

Before proving Theorem 16, we need to establish several technical results.

We use $\gamma, \gamma', \dots$ to range over cycles in a dependency graph $\mathcal{G}$, i.e., paths of the form $T_1 \xrightarrow{\mathcal{R}_1} T_2 \xrightarrow{\mathcal{R}_2} \cdots \xrightarrow{\mathcal{R}_{n-1}} T_n$ such that $T_1 = T_n$ and for any $i = 1, \dots, n$ we have $T_i \in \mathcal{T}_{\mathcal{G}}$ and $\mathcal{R}_i \in \{\mathsf{SO}_{\mathcal{G}}, \mathsf{WR}_{\mathcal{G}}, \mathsf{WW}_{\mathcal{G}}, \mathsf{RW}_{\mathcal{G}}\}$. For a given cycle $\gamma = T_1 \xrightarrow{\mathcal{R}_1} T_2 \xrightarrow{\mathcal{R}_2} \cdots \xrightarrow{\mathcal{R}_{n-1}} T_n$, we let $\mathsf{rep}(\gamma) = |\{T_i \mid \exists j.\, 1 \leq i, j < n \wedge i \neq j \wedge T_i = T_j\}|$ be the number of repeated vertices in it. Note that γ has no repeated vertices if and only if $\mathsf{rep}(\gamma) = 0$. We call such a cycle ***simple***.

LEMMA 24. *Let $\mathcal{G}$ be a dependency graph such that the relation* $((\mathsf{SO}_{\mathcal{G}} \cup \mathsf{WR}_{\mathcal{G}} \cup \mathsf{WW}_{\mathcal{G}}) \,;\, \mathsf{RW}_{\mathcal{G}}?)$ *contains a cycle. Then this relation contains a simple cycle.*

PROOF. Let γ be a cycle in $((\mathsf{SO}_{\mathcal{G}} \cup \mathsf{WR}_{\mathcal{G}} \cup \mathsf{WW}_{\mathcal{G}}) \,;\, \mathsf{RW}_{\mathcal{G}}?)$, i.e., a cycle without adjacent $\mathsf{RW}_{\mathcal{G}}$ edges. If $\mathsf{rep}(\gamma) = 0$, then there is nothing to prove, so let us assume that $\mathsf{rep}(\gamma) > 0$. Below we show how to extract a sub-cycle γ' of γ such that $\mathsf{rep}(\gamma') < \mathsf{rep}(\gamma)$ and γ' has no adjacent $\mathsf{RW}_{\mathcal{G}}$ edges. By applying this procedure repeatedly, we obtain a cycle γ'' with no adjacent $\mathsf{RW}_{\mathcal{G}}$ edges and such that $\mathsf{rep}(\gamma'') = 0$, as required.

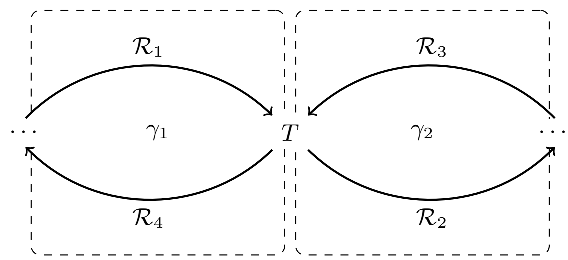

Figure 9: A cycle with a repeated transaction T

To construct γ' from γ, we proceed as follows: since $\mathsf{rep}(\gamma) > 0$, then

$$\gamma = S \dashrightarrow \ldots \xrightarrow{\mathcal{R}_1} T \xrightarrow{\mathcal{R}_2} \ldots \xrightarrow{\mathcal{R}_3} T \xrightarrow{\mathcal{R}_4} \ldots \dashrightarrow S$$

for some $T, S \in \mathcal{T}_{\mathcal{G}}$ and $\{\mathcal{R}_i\}_{i=1}^{4} \subseteq \{\mathsf{SO}_{\mathcal{G}}, \mathsf{WR}_{\mathcal{G}}, \mathsf{WW}_{\mathcal{G}}, \mathsf{RW}_{\mathcal{G}}\}$. A graphical representation of γ is given in Figure 9.

From γ we can derive the cycles

$$\gamma_1 = S \dashrightarrow \cdots \xrightarrow{\mathcal{R}_1} T \xrightarrow{\mathcal{R}_4} \cdots \dashrightarrow S$$

and

$$\gamma_2 = T \xrightarrow{\mathcal{R}_2} \cdots \xrightarrow{\mathcal{R}_3} T,$$

which are contained inside the dashed boxes in the picture to the right.

Since γ contains no adjacent $\mathsf{RW}_{\mathcal{G}}$ edges, we can have two adjacent $\mathsf{RW}_{\mathcal{G}}$ edges in γ_1 only if $\mathcal{R}_1 = \mathsf{RW}_{\mathcal{G}}$ and $\mathcal{R}_4 = \mathsf{RW}_{\mathcal{G}}$; similarly, we have two adjacent $\mathsf{RW}_{\mathcal{G}}$ edges in γ_2 only if $\mathcal{R}_2 = \mathsf{RW}_{\mathcal{G}}$ and $\mathcal{R}_3 = \mathsf{RW}_{\mathcal{G}}$. Therefore, if either $\mathcal{R}_1 \neq \mathsf{RW}_{\mathcal{G}}$ or $\mathcal{R}_2 \neq \mathsf{RW}_{\mathcal{G}}$, then we know that γ_1 has no adjacent $\mathsf{RW}_{\mathcal{G}}$ edges, and we choose $\gamma' = \gamma_1$. Otherwise $\mathcal{R}_1 = \mathsf{RW}_{\mathcal{G}}$; since γ contains no adjacent $\mathsf{RW}_{\mathcal{G}}$ edges, it follows that $\mathcal{R}_2 \neq \mathsf{RW}_{\mathcal{G}}$; therefore γ_2 contains no adjacent $\mathsf{RW}_{\mathcal{G}}$ edges, and we choose $\gamma' = \gamma_2$. $\square$

PROPOSITION 25. *Let $\mathcal{G}$ be a dependency graph. For any $T \in \mathcal{T}_{\mathcal{G}}$:*

(i) $\boxed{T}_{\mathcal{G}} \vdash \mathtt{read}(x, n)$ *if and only if* $\min_{\mathsf{SO}_{\mathcal{G}}}\{S \mid S \approx_{\mathcal{G}} T \wedge S \vdash \_(x, \_)\} \vdash \mathtt{read}(x, n)$*;*

(ii) $\boxed{T}_{\mathcal{G}} \vdash \mathtt{write}(x, n)$ *if and only if* $\max_{\mathsf{SO}_{\mathcal{G}}}\{S \mid S \approx_{\mathcal{G}} T \wedge S \vdash \mathtt{write}(x, \_)\} \vdash \mathtt{write}(x, n)$.

PROOF. We only prove Statement i; the proof for Statement ii is analogous. To maintain the notation easy, we let $\boxed{T}_{\mathcal{G}} = \left(\boxed{E_T}_{\mathcal{G}}, \boxed{\mathsf{po}_T}_{\mathcal{G}}\right)$.

Suppose that $\boxed{T}_{\mathcal{G}} \vdash \mathtt{read}(x, n)$; by definition, the event $e_{\mathsf{rd}} = \min_{\boxed{\mathsf{po}}_{\mathcal{G}}}\{e \in \boxed{E_T}_{\mathcal{G}} \mid \mathsf{op}(e) = \_(x, \_)\}$ exists, and $\mathsf{op}(e_{\mathsf{rd}}) = \mathtt{read}(x, n)$. Recall that $\boxed{E_T}_{\mathcal{G}} = (\bigcup\{E_S \mid S \approx_{\mathcal{G}} T\})$, from which it follows that $e_{\mathsf{rd}} \in E_{T_{\mathsf{rd}}}$ for some $T_{\mathsf{rd}} \approx_{\mathcal{G}} T$. Also, $\boxed{\mathsf{po}_T}_{\mathcal{G}} = \{(e, f) \mid \exists S.\, S \approx_{\mathcal{G}} T \wedge (e, f \in S \wedge e \xrightarrow{\mathsf{po}_S} f) \vee (\exists S'.\, e \in E_S \wedge f \in E_{S'} \wedge S \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S')\}$. A first consequence of this fact is that, for any event $f \in E_{\mathsf{rd}}$ such that $\mathsf{op}(f) = \_(x, \_)$, $e_{\mathsf{rd}} \xrightarrow{\mathsf{po}_{T_{\mathsf{rd}}}} f$, hence $T_{\mathsf{rd}} \vdash \mathtt{read}(x, n)$. A second consequence is that, for any transaction $S \approx_{\mathcal{G}} T_{\mathsf{rd}}$ with $S \neq T_{\mathsf{rd}}$, and for any event $f \in E_S$ such that $\mathsf{op}(f) = \_(x, \_)$, then $T_{\mathsf{rd}} \xrightarrow{\mathsf{SO}}_{\mathcal{G}} S$. Alternatively, we can write that if $S \approx_{\mathcal{G}} T_{\mathsf{rd}}$, $S \neq T_{\mathsf{rd}}$, and either $S \vdash \mathtt{read}(x, \_)$ or $S \vdash \mathtt{write}(x, \_)$, then $T_{\mathsf{rd}} \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S$; therefore, $T_{\mathsf{rd}} = \min_{\mathsf{SO}_{\mathcal{G}}}\{S \mid S \approx_{\mathcal{G}} T \wedge (S \vdash \mathtt{read}(x, \_) \vee S \vdash \mathtt{write}(x, \_))\}$.

Next, suppose that $T_{\mathsf{rd}} = \min_{\mathsf{SO}_{\mathcal{G}}}\{S \mid S \approx_{\mathcal{G}} T \wedge S \vdash \mathtt{read}(x, \_) \vee S \vdash \mathtt{write}(x, \_)\}$ is defined, and $T_{\mathsf{rd}} \vdash \mathtt{read}(x, n)$. By definition, there exists an event $e_{\mathsf{rd}} \in E_{T_{\mathsf{rd}}}$ such that $\mathsf{op}(e_{\mathsf{rd}}) = \mathtt{read}(x, n)$, and for any $f \in E_{T_{\mathsf{rd}}}$ such that $\mathsf{op}(f) = \_(x, \_)$ we have that $e_{\mathsf{rd}} \xrightarrow{\mathsf{po}_{T_{\mathsf{rd}}}} f$. Note that $e_{\mathsf{rd}} \in \boxed{E_T}_{\mathcal{G}}$. We prove that, for any event $f \in \boxed{E_T}_{\mathcal{G}}$, if $\mathsf{op}(f) = \_(x, \_)$ then $e_{\mathsf{rd}} \xrightarrow{\boxed{\mathsf{po}_T}_{\mathcal{G}}} f$, from which it follows that $\boxed{T}_{\mathcal{G}} \vdash \mathtt{read}(x, n)$. First note that we have already shown that if $\mathsf{op}(f) = \_(x, \_)$ and $f \in E_{T_{\mathsf{rd}}}$, then $e_{\mathsf{rd}} \xrightarrow{\mathsf{po}_{T_{\mathsf{rd}}}} f$, from which $e_{\mathsf{rd}} \xrightarrow{\boxed{\mathsf{po}_T}_{\mathcal{G}}} f$ follows. Suppose then that $f \in E_S$ for some $S \neq T_{\mathsf{rd}}$, $S \approx_{\mathcal{G}} T_{\mathsf{rd}}$. If $\mathsf{op}(f) = \mathtt{read}(x, \_)$ then either $S \vdash \mathtt{read}(x, \_)$ (if there exists no event $f' \in E_S$ such that $\mathsf{op}(f') = \mathtt{write}(x, \_)$) or $S \vdash \mathtt{write}(x, \_)$ (otherwise); If $\mathsf{op}(f) = \mathtt{write}(x, \_)$, then $S \vdash \mathtt{write}(x, \_)$. We have proved that if $\mathsf{op}(f) = \_(x, \_)$, then $(S \vdash \mathtt{read}(x, \_) \vee S \vdash \mathtt{write}(x, \_))$, so that it cannot be $S \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T_{\mathsf{rd}}$; therefore, it has to be the case that $T_{\mathsf{rd}} \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S$, from which $e_{\mathsf{rd}} \xrightarrow{\boxed{\mathsf{po}_T}_{\mathcal{G}}} f$ follows, and there is nothing left to prove. $\square$

LEMMA 26. *Let $\mathcal{G} \in \mathsf{GraphSI}$ such that $\mathsf{DCG}(\mathcal{G})$ contains no critical cycles. Then $\mathsf{splice}(\mathcal{G})$ is a dependency graph.*

PROOF. We prove that, if the chopping graph of $\mathcal{G}$ contains no critical cycles, then $\mathsf{splice}(\mathcal{G})$ satisfies all the constraints of Definition 6.

(a)

$$\begin{tikzcd}
T & \approx_{\mathcal{G}} & T' \vdash \mathtt{write}(x, _) \arrow[r, "\mathsf{SO}_{\mathcal{G}}"] \arrow[r, bend left=45, "\mathsf{WW}_{\mathcal{G}}(x)"] \arrow[d, "\mathsf{WR}_{\mathcal{G}}(x)"] & T'' \vdash \mathtt{write}(x, _) \arrow[l, bend left, "\mathsf{SO}^{-1}_{\mathcal{G}}(x)"] \\
S & \approx_{\mathcal{G}} & S' \vdash \mathtt{read}(x, _) \arrow[ur, "\mathsf{RW}_{\mathcal{G}}(x)"'] &
\end{tikzcd}$$

(b)

$$\begin{tikzcd}
T & \approx_{\mathcal{G}} & T' \vdash \mathtt{write}(x, _) \arrow[d, "\mathsf{WW}_{\mathcal{G}}(x)"] \arrow[dr, "\mathsf{WR}_{\mathcal{G}}(x)"] & \\
S & \approx_{\mathcal{G}} & S'' \vdash \mathtt{write}(x, _) \arrow[r, "\mathsf{SO}_{\mathcal{G}}"] & S' \vdash \mathtt{read}(x, _) \arrow[l, bend right, "\mathsf{RW}_{\mathcal{G}}(x)"']
\end{tikzcd}$$

(c)

$$\begin{tikzcd}
T & \approx_{\mathcal{G}} & T' \vdash \mathtt{write}(x, _) \arrow[dr, "\mathsf{WR}_{\mathcal{G}}(x)"] & \\
S & \approx_{\mathcal{G}} & S'' \vdash \mathtt{write}(x, _) \arrow[u, "\mathsf{WW}_{\mathcal{G}}(x)"] \arrow[r, "\mathsf{SO}_{\mathcal{G}}"] & S' \vdash \mathtt{read}(x, _) \arrow[l, bend right, "\mathsf{SO}^{-1}_{\mathcal{G}}"']
\end{tikzcd}$$

(d)

$$\begin{tikzcd}
 & & V & T \\
 & & \approx_{\mathcal{G}} & \approx_{\mathcal{G}} \\
 & & V' \vdash \mathtt{write}(x, m) \arrow[d, "\mathsf{WR}_{\mathcal{G}}(x)"] & T' \vdash \mathtt{write}(x, n) \arrow[l, "\mathsf{WW}_{\mathcal{G}}(x)"'] \arrow[d, "\mathsf{WR}_{\mathcal{G}}(x)"] \\
S & \approx_{\mathcal{G}} & S'' \vdash \mathtt{read}(x, m) \arrow[r, "\mathsf{SO}_{\mathcal{G}}"] & S' \vdash \mathtt{read}(x, n) \arrow[ul, "\mathsf{RW}_{\mathcal{G}}(x)"]
\end{tikzcd}$$

(e)

$$\begin{tikzcd}
 & & V & T \\
 & & \approx_{\mathcal{G}} & \approx_{\mathcal{G}} \\
 & & V' \vdash \mathtt{write}(x, m) \arrow[r, "\mathsf{WW}_{\mathcal{G}}(x)"] \arrow[d, "\mathsf{WR}_{\mathcal{G}}(x)"] & T' \vdash \mathtt{write}(x, n) \arrow[d, "\mathsf{WR}_{\mathcal{G}}(x)"] \\
S & \approx_{\mathcal{G}} & S'' \vdash \mathtt{read}(x, m) \arrow[r, "\mathsf{SO}_{\mathcal{G}}"] \arrow[ur, "\mathsf{RW}_{\mathcal{G}}(x)"'] & S' \vdash \mathtt{read}(x, n) \arrow[l, bend left, "\mathsf{SO}^{-1}_{\mathcal{G}}"]
\end{tikzcd}$$

(f)

$$\begin{tikzcd}
 & & T & S \\
 & & \approx_{\mathcal{G}} & \approx_{\mathcal{G}} \\
 & & T' \vdash \mathtt{write}(x, _) \arrow[d, "\mathsf{WW}_{\mathcal{G}}(x)"'] & S'' \vdash \mathtt{write}(x, _) \arrow[l, "\mathsf{WW}_{\mathcal{G}}(x)?"'] \\
V & \approx_{\mathcal{G}} & V' \vdash \mathtt{write}(x, _) \arrow[r, bend right, "\mathsf{SO}^{-1}_{\mathcal{G}}"'] & V'' \vdash \mathtt{write}(x, _) \arrow[l, "\mathsf{WW}_{\mathcal{G}}(x)"'] \arrow[u, "\mathsf{WW}_{\mathcal{G}}(x)"']
\end{tikzcd}$$

(g)

$$\begin{tikzcd}
T \approx_{\mathcal{G}} S & \approx_{\mathcal{G}} & T' \vdash \mathtt{write}(x, _) \arrow[r, "\mathsf{SO}_{\mathcal{G}}"] \arrow[d, "\mathsf{WW}_{\mathcal{G}}(x)"'] & S'' \vdash \mathtt{write}(x, _) \arrow[l, bend left, "\mathsf{SO}^{-1}_{\mathcal{G}}"] \\
V & \approx_{\mathcal{G}} & V' \vdash \mathtt{write}(x, _) \arrow[r, "\mathsf{SO}_{\mathcal{G}}?"] & V'' \vdash \mathtt{write}(x, _) \arrow[u, "\mathsf{WW}_{\mathcal{G}}(x)"']
\end{tikzcd}$$

(h)

$$\begin{tikzcd}
T & \approx_{\mathcal{G}} & T' \vdash \mathtt{write}(x, _) \arrow[r, "\mathsf{SO}_{\mathcal{G}}"] \arrow[d, "\mathsf{WW}_{\mathcal{G}}(x)"'] & S'' \vdash \mathtt{write}(x, _) \arrow[l, bend left, "\mathsf{SO}^{-1}_{\mathcal{G}}"] \\
V & \approx_{\mathcal{G}} & V' \vdash \mathtt{write}(x, _) \arrow[r, bend left, "\mathsf{SO}^{-1}_{\mathcal{G}}"] & V'' \vdash \mathtt{write}(x, _) \arrow[l, "\mathsf{SO}_{\mathcal{G}}"] \arrow[u, "\mathsf{WW}_{\mathcal{G}}(x)"']
\end{tikzcd}$$

Figure 10: Graphical representation of different cases in the proof of Lemma 26.

- Consider $\boxed{T}_{\mathcal{G}}, \boxed{S}_{\mathcal{G}} \in \mathcal{T}_{\mathsf{splice}(\mathcal{G})}$ such that $\boxed{T}_{\mathcal{G}} \xrightarrow{\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$. By definition $\boxed{T}_{\mathcal{G}} \neq \boxed{S}_{\mathcal{G}}$ and there exist two transactions $T', S' \in \mathcal{T}_{\mathcal{G}}$ such that $T \approx_{\mathcal{G}} T' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S' \approx_{\mathcal{G}} S$. Hence, for some n we have $T' \vdash \mathtt{write}(x, n)$ and $S' \vdash \mathtt{read}(x, n)$. We now prove that *(i)* $\boxed{T}_{\mathcal{G}} \vdash \mathtt{write}(x, n)$ and *(ii)* $\boxed{S}_{\mathcal{G}} \vdash \mathtt{read}(x, n)$.

  (i) Consider an arbitrary transaction $T'' \in \mathcal{T}$ such that $T'' \vdash \mathtt{write}(x, \_)$ and $T'' \approx_{\mathcal{G}} T'$ (Figure 10(a)). We show that it cannot be the case that $T' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T''$. Then $T' = \max_{\mathsf{SO}_{\mathcal{G}}}\{S \mid S \approx_{\mathcal{G}} T \wedge S \vdash \mathtt{write}(x, \_)\}$, and by Proposition 25(ii) it follows that $\boxed{T}_{\mathcal{G}} \vdash \mathtt{write}(x, n)$, as required.

  Assume $T' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T''$; then $T' \neq T''$. Since $T', T'' \vdash \mathtt{write}(x, \_)$, either $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T''$ or $T'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$. However, the latter case would lead to the cycle $T'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}} T' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T''$, which cannot exist because $\mathcal{G} \in \mathsf{GraphSI}$. Therefore $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T''$. Together with $T' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S'$, this yields the anti-dependency $S' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} T''$. Since $S' \approx_{\mathcal{G}} S \not\approx_{\mathcal{G}} T \approx_{\mathcal{G}} T''$, this implies $S' \xrightarrow{(\mathsf{RW}_{\mathcal{G}}(x) \setminus \approx_{\mathcal{G}})} T''$. We have thus obtained the cycle $T' \xrightarrow{(\mathsf{WR}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} S' \xrightarrow{(\mathsf{RW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} T'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} T'$ in $\mathsf{DCG}(\mathcal{G})$, which is critical. This contradicts the assumption of the lemma.

  (ii) We show that, for any transaction $S'' \approx_{\mathcal{G}} S'$ such that $S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S'$ we have $\neg(S'' \vdash \mathtt{write}(x, \_))$, and if $S'' \vdash \mathtt{read}(x, m)$, then $m = n$. As a consequence, $\min_{\mathsf{SO}_{\mathcal{G}}}(\{V \mid V \approx_{\mathcal{G}} S \wedge V \vdash \_(x, \_)\} \vdash \mathtt{read}(x, n))$, and hence, by Proposition 25(i) we have $\boxed{S}_{\mathcal{G}} \vdash \mathtt{read}(x, n)$.

  Let S'' be a transaction such that $S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S'$; we prove that $\neg(S'' \vdash \mathtt{write}(x, \_))$ by contradiction. Assume $S'' \vdash \mathtt{write}(x, \_)$. Then by the definition of $\mathsf{WW}_{\mathcal{G}}(x)$, either $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$ or $S'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$; the case $S'' = T'$ is ruled out because $S'' \approx_{\mathcal{G}} S \not\approx_{\mathcal{G}} T \approx_{\mathcal{G}} T'$.

  We cannot have $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$ (Figure 10(b)), since together with $T' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S'$, this would imply the anti-dependency edge $S' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} S''$. But then we have a cycle $S' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S'$, contradicting $\mathcal{G} \in \mathsf{GraphSI}$. We cannot have $S'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$ either (Figure 10(c)): in this case the chopping graph of $\mathcal{G}$ contains the critical cycle $S'' \xrightarrow{(\mathsf{WW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} T' \xrightarrow{(\mathsf{WR}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} S' \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} S''$. We have thus established $\neg(S'' \vdash \mathtt{write}(x, \_))$.

  Suppose now that $S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S'$ and $S'' \vdash \mathtt{read}(x, m)$ for some m. Then there exists a transaction $V' \in \mathcal{T}_{\mathcal{G}}$ such that $V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S''$ and $V' \vdash \mathtt{write}(x, m)$. Since $T' \vdash \mathtt{write}(x, n)$, we have $V' = T'$, $V \xrightarrow{\mathsf{WW}_{\mathcal{G}(x)}} T'$ or $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}(x)}} V$. We show that the latter two cases are impossible, so that $T' = V'$ and, hence, $m = n$, as required. If $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$, then we have the anti-dependency edge $S' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} V'$ (Figure 10(d)). Then the graph $\mathcal{G}$ contains the cycle $S' \xrightarrow{\mathsf{RW}_{\mathcal{G}}} V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}} S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S'$, which contradicts $\mathcal{G} \in \mathsf{GraphSI}$. If $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$, then we have the anti-dependency edge $S'' \xrightarrow{\mathsf{RW}_{\mathcal{G}}(x)} T'$ (Figure 10(e)). In this case $\mathsf{DCG}(\mathcal{G})$ contains the critical cycle $S'' \xrightarrow{(\mathsf{RW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} T' \xrightarrow{(\mathsf{WR}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} S' \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} S''$, and again we obtain a contradiction.

- Let $\boxed{S}_{\mathcal{G}}$ be a transaction such that $\boxed{S}_{\mathcal{G}} \vdash \mathtt{read}(x, \_)$. By Proposition 25(i), there exists a transaction $S' \approx_{\mathcal{G}} S$ such that $S' \vdash \mathtt{read}(x, \_)$. Since $\mathcal{G}$ is a dependency graph, there exists a transaction $T \in \mathcal{T}_{\mathcal{G}}$ such that $T \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S'$. Then $T \neq S'$. We cannot have $T \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S'$, because $T \vdash \mathtt{write}(x, \_)$ and $S' = \min_{\mathsf{SO}}\{S \mid S \approx_{\mathcal{G}} S' \wedge S \vdash \_(x, \_)\}$ by Proposition 25(i). Finally, we cannot have $S' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T$, because this would contradict the hypothesis $\mathcal{G} \in \mathsf{GraphSI}$ due to the cycle $S' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S'$. As a consequence, it cannot be $T \approx_{\mathcal{G}} S'$. Then $\boxed{T}_{\mathcal{G}} \xrightarrow{\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$.

- Let $\boxed{S}_{\mathcal{G}}, \boxed{T}_{\mathcal{G}}, \boxed{V}_{\mathcal{G}} \in \mathcal{T}_{\mathsf{splice}(\mathcal{G})}$ be such that $\boxed{T}_{\mathcal{G}} \xrightarrow{\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$ and $\boxed{V}_{\mathcal{G}} \xrightarrow{\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$. Then $S \not\approx_{\mathcal{G}} T$, $S \not\approx_{\mathcal{G}} V$ and there exist transactions S', S'', T', V' such that $S' \approx_{\mathcal{G}} S \approx_{\mathcal{G}} S''$, $T' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S'$, $V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S''$. Note that if $S' = S''$, then $T' = V'$, because $\mathcal{G}$ is a dependency graph; hence $\boxed{T}_{\mathcal{G}} = \boxed{T'}_{\mathcal{G}} = \boxed{V'}_{\mathcal{G}} = \boxed{V}_{\mathcal{G}}$, and there is nothing left to prove.

  It remains to analyse the case when $S' \neq S''$, so that either $S' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S''$ or $S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S'$. Without loss of generality, assume that $S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S'$ (Figure 10(d)). Since $T' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S'$ and $V' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} S''$, we get $T' \vdash \mathtt{write}(x, \_)$ and $V' \vdash \mathtt{write}(x, \_)$. Therefore, we must have one of the following: $T' = V'$, $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$ or $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$. However, the latter two cases are impossible. If we had $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$, then $\mathcal{G}$ would contain the configuration shown in Figure 10(d), which contradicts $\mathcal{G} \in \mathsf{GraphSI}$. If we had $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$, then $\mathcal{G}$ would contain the configuration shown in Figure 10(d) and, hence, $\mathsf{DCG}(\mathcal{G})$ would contain a critical cycle. Hence, we must have $T' = V'$ and $\boxed{T}_{\mathcal{G}} = \boxed{V}_{\mathcal{G}}$.

- We prove that $\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}$ is transitive, irreflexive, and whenever $\boxed{T}_{\mathcal{G}} \vdash \mathtt{write}(x, \_)$ and is total over $\mathsf{WriteTx}_x$.

$\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}$ **is transitive.** Let $\boxed{T}_{\mathcal{G}}, \boxed{V}_{\mathcal{G}}, \boxed{S}_{\mathcal{G}}$ be such that $\boxed{T}_{\mathcal{G}} \xrightarrow{\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{V}_{\mathcal{G}} \xrightarrow{\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$. By definition there exist transactions $T', V', V'', S'' \in \mathcal{T}_{\mathcal{G}}$ such that

$$T \approx_{\mathcal{G}} T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V' \approx_{\mathcal{G}} V \approx_{\mathcal{G}} V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S'' \approx_{\mathcal{G}} S.$$

To prove that $\boxed{T}_{\mathcal{G}} \xrightarrow{\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$, it suffices to prove that $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$ and $T \not\approx_{\mathcal{G}} S$. We prove these two statements separately.

- Since $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$ and $V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$, we have $V' \vdash \mathtt{write}(x, \_)$ and $V'' \vdash \mathtt{write}(x, \_)$. Then one of the following holds: $V' = V''$, $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V''$ or $V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$. In the first two cases, the transitivity of $\mathsf{WW}_{\mathcal{G}}(x)$ guarantees $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$, as required. Suppose now that $V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$. Then $V'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} V'$, because $V'' \approx_{\mathcal{G}} V'$ and $\mathcal{G} \in \mathsf{GraphSI}$. Since $T' \vdash \mathtt{write}(x, \_)$ and $S'' \vdash \mathtt{write}(x, \_)$, we have one of the following: $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$, $S'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$ or $T' = S''$. However, in the latter two cases we would end up with the chopping graph of Figure 10(f), which contains the critical cycle $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x) \setminus \approx_{\mathcal{G}}} V' \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x) \setminus \approx_{\mathcal{G}}} S'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$. Therefore, it has to be $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$.
- We prove $T \not\approx_{\mathcal{G}} S$ by contradiction. Suppose that $T \approx_{\mathcal{G}} S$ and, hence, $T' \approx_{\mathcal{G}} S''$. We have $V' \vdash \mathtt{write}(x, \_)$ and $V'' \vdash \mathtt{write}(x, \_)$, so that one of the following holds: $V' = V''$, $V' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V''$ or $V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$.

  In the first two cases, $\mathcal{G} \in \mathsf{GraphSI}$ guarantees $V' \xrightarrow{\mathsf{SO}_{\mathcal{G}}(x)?} V''$ and the transitivity of $\mathsf{WW}_{\mathcal{G}}(x)$ guarantees $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S''$. Since $\mathsf{WW}_{\mathcal{G}}(x)$ is irreflexive, we cannot have $T' = S''$, and since $\mathcal{G} \in \mathsf{GraphSI}$, we cannot have $S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T'$. Therefore, $T' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S''$. Then, as illustrated in Figure 10(g), $\mathsf{DCG}(\mathcal{G})$ contains the critical cycle $T' \xrightarrow{(\mathsf{WW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} V' \xrightarrow{\mathsf{SO}_{\mathcal{G}}?} V'' \xrightarrow{(\mathsf{WW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} T'$, yielding a contradiction.

  It remains to consider the case when $V'' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} V'$. Then $V'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} V'$, because $V'' \approx_{\mathcal{G}} V'$ and $\mathcal{G} \in \mathsf{GraphSI}$. Since we are assuming $T' \approx_{\mathcal{G}} S''$, one of the following holds: $T' = S''$, $T' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S''$ or $S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T'$. In all cases there is a critical cycle in $\mathsf{DCG}(\mathcal{G})$. For example, if $T' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S''$, then $\mathcal{G}$ has a configuration shown in Figure 10(h), and $\mathsf{DCG}(\mathcal{G})$ contains the critical cycle $T' \xrightarrow{(\mathsf{WW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} V' \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} V'' \xrightarrow{(\mathsf{WW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} S'' \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} T'$.

$\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)$ **is irreflexive.** this follows immediately from the definition of $\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)$.

$\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)$ **is total over** $\mathsf{WriteTx}_x$**.** Let $\boxed{T}_{\mathcal{G}}, \boxed{S}_{\mathcal{G}}$ be such that $\boxed{T}_{\mathcal{G}} \vdash \mathtt{write}(x, \_)$, $\boxed{S}_{\mathcal{G}} \vdash \mathtt{write}(x, \_)$ and $\boxed{T}_{\mathcal{G}} \neq \boxed{S}_{\mathcal{G}}$. By Proposition 25(ii), there exist $T', S' \in \mathcal{T}_{\mathcal{G}}$ such that $T' \approx_{\mathcal{G}} T$, $T' \vdash \mathtt{write}(x, \_)$, $S' \approx_{\mathcal{G}} S$ and $S' \vdash \mathtt{write}(x, \_)$. Also, $T' \approx_{\mathcal{G}} T \not\approx_{\mathcal{G}} S \approx_{\mathcal{G}} S'$, so that it cannot be $T' = S'$. Since $\mathsf{WW}_{\mathcal{G}}(x)$ is total over $\mathsf{WriteTx}_x$, we must have either $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S'$ or $S' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} T'$; without loss of generality, we assume $T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S'$. We thus have $\boxed{T}_{\mathcal{G}} \neq \boxed{S}_{\mathcal{G}}$ and $T \approx_{\mathcal{G}} T' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(x)} S' \approx_{\mathcal{G}} S$. Hence, $\boxed{T}_{\mathcal{G}} \xrightarrow{\mathsf{WW}_{\mathsf{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$.

$\square$

LEMMA 27. *Let $\mathcal{G} \in \mathsf{GraphSI}$ be a dependency graph whose chopping graph has no critical cycle. Then $\mathcal{T}_{\mathsf{splice}(\mathcal{G})} \models$ INT.*

PROOF. Proceeding by contradiction, let us assume that INT is violated. Then there exist: $T \in \mathcal{T}_{\mathcal{G}}$; $e \in E_T$ such that $\mathsf{op}(e) = \mathtt{read}(x, n)$ for some $x \in \mathsf{Obj}$, $n \in \mathbb{N}$; and, letting $\boxed{\mathsf{po}_T} = \mathsf{po}_{\boxed{T}_{\mathcal{G}}}$,

$$f = \max_{\boxed{\mathsf{po}_T}}\{e' \mid \mathsf{op}(e') = \_(x, \_) \land e' \xrightarrow{\boxed{\mathsf{po}_T}} e\} \tag{4}$$

such that $\mathsf{op}(f) = \_(x, m)$ for some $m \neq n$. Since $\mathcal{T}_{\mathcal{G}} \models$ INT, we cannot have $f \in E_T$. Therefore, there exists a transaction $T' \xrightarrow{\mathsf{SO}_{\mathcal{G}}} T$ such that $f \in E_{T'}$. We now make a case split on whether $T' \vdash \mathtt{write}(x, \_)$.

1. $T' \vdash \mathtt{write}(x, \_)$. Then there exists an event $g \in E_{T'}$ such that $\mathsf{op}(g) = \mathtt{write}(x, \_)$. Without loss of generality, let g be the last write to object x in $E_{T'}$.

   If $f \xrightarrow{\mathsf{po}_{T'}} g$, then $f \xrightarrow{\boxed{\mathsf{po}_T}} g \xrightarrow{\boxed{\mathsf{po}_T}} e$, contradicting (4). Thus, either $g = f$ or $g \xrightarrow{\mathsf{po}_{T'}} f$. In both cases, we show that $\mathsf{op}(g) = \mathtt{write}(x, m)$. If $f = g$, we have $\_(x, m) = \mathsf{op}(f) = \mathsf{op}(g) = \mathtt{write}(x, \_)$, so that $\mathsf{op}(g) = \mathtt{write}(x, m)$. If $g \xrightarrow{\mathsf{po}_{T'}} f$, then $\mathsf{op}(f) = \mathtt{read}(x, m)$, since g is the last write to x in T'. Also, for any other event event h such that $\mathsf{op}(h) = \_(x, \_)$ and $g \xrightarrow{\mathsf{po}_{T'}} h \xrightarrow{\mathsf{po}_{T'}} f$, we have $\mathsf{op}(h) = \mathtt{read}(x, \_)$. Then because $\mathcal{T}_{\mathcal{G}} \models$ INT and $\mathsf{op}(f) = \mathtt{read}(x, m)$, we must have $\mathsf{op}(h) = \mathtt{read}(x, m)$. But then $\mathcal{T}_{\mathcal{G}} \models$ INT again ensures $\mathsf{op}(g) = \mathtt{write}(x, m)$.

   We have proved that $T' \vdash \mathtt{write}(x, m)$. By hypothesis $T \vdash \mathtt{read}(x, n)$ for some $n \neq m$, so that there exists a transaction $S \neq T'$ such that $S \vdash \mathtt{write}(x, n)$ and $S \xrightarrow{\mathsf{WR}_{\mathcal{G}}(x)} T$.

Next we prove that it cannot be either $S \xrightarrow{\mathsf{SO}_\mathcal{G}} T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T$, nor $T' \xrightarrow{\mathsf{SO}_\mathcal{G}} S \xrightarrow{\mathsf{SO}_\mathcal{G}} T$, nor $T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T \xrightarrow{\mathsf{SO}_\mathcal{G}} S$. Since we have already proved that $S \neq T'$ and $T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T$, these imply $T' \not\approx_\mathcal{G} S$.

- If $S \xrightarrow{\mathsf{SO}_\mathcal{G}} T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T$, then since $S, T' \vdash \mathtt{write}_\mathcal{G}(x, \_)$ and $S \neq T'$, it has to be either $S \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} T'$ or $T' \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} S$. However, the last case is impossible because it would lead to a cycle $S \xrightarrow{\mathsf{SO}_\mathcal{G}} T' \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} S$, contradicting $\mathcal{G} \in \mathsf{GraphSI}$; therefore, $S \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} T'$. Together with $S \xrightarrow{\mathsf{WR}_\mathcal{G}(x)} T$, this yields $T \xrightarrow{\mathsf{RW}_\mathcal{G}(x)} T'$. But then we obtain a cycle $T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T \xrightarrow{\mathsf{RW}_\mathcal{G}(x)} T'$, which contradicts $\mathcal{G} \in \mathsf{GraphSI}$.
- If $T' \xrightarrow{\mathsf{SO}_\mathcal{G}} S \xrightarrow{\mathsf{SO}_\mathcal{G}} T$, since $S \vdash \mathtt{write}(x, \_)$, this implies that there exists an event $h \in E_S$ such that $\mathsf{op}(h) = \mathtt{write}(x, \_)$; furthermore, in this case we have that $f \xrightarrow{\boxed{\mathsf{po}_T}} h \xrightarrow{\boxed{\mathsf{po}_T}} e$, contradicting the hypothesis that $f = \max_{\boxed{\mathsf{po}_T}}\{f \mid \mathsf{op}(f) = \_(x, \_) \wedge f \xrightarrow{\boxed{\mathsf{po}_T}} e\}$.
- If $T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T \xrightarrow{\mathsf{SO}_\mathcal{G}} S$, then we have the cycle $T \xrightarrow{\mathsf{SO}_\mathcal{G}} S \xrightarrow{\mathsf{WR}_\mathcal{G}(x)} T$, which is not allowed because we are assuming that $\mathcal{G} \in \mathsf{GraphSI}$.

We have proved that $T' \not\approx_\mathcal{G} S$. Next, we observe that since $T' \vdash \mathtt{write}(x, m)$, $S \vdash \mathtt{write}(x, n)$ and $T' \neq S$, we must have either $T' \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} S$ or $S \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} T'$. We show that both of these cases lead to a contradiction. If $S \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} T'$, then since $S \xrightarrow{\mathsf{WR}_\mathcal{G}(x)} T$, we get $T \xrightarrow{\mathsf{RW}_\mathcal{G}(x)} T'$, causing the cycle $T \xrightarrow{\mathsf{RW}_\mathcal{G}(x)} T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T$; this contradicts $\mathcal{G} \in \mathsf{GraphSI}$. On the other hand, if $T' \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} S$, then we have a critical cycle $T' \xrightarrow{(\mathsf{WW}_\mathcal{G} \setminus \approx_\mathcal{G})} S \xrightarrow{(\mathsf{WR}_\mathcal{G} \setminus \approx_\mathcal{G})} T \xrightarrow{\mathsf{SO}_\mathcal{G}^{-1}} T'$ in $\mathsf{DCG}(\mathcal{G})$, contradicting the hypothesis of the lemma.

2. $\neg(T' \vdash \mathtt{write}(x, \_))$. In this case there exists no event $g \in E_{T'}$ such that $\mathsf{op}(g) = \mathtt{write}(x, \_)$. Using the fact that $\mathcal{T}_\mathcal{G} \models \textsc{Int}$, we can easily show that for any $g \in E_{T'}$ such that $\mathsf{op}(g) = \mathtt{read}(x, \_)$, we have $\mathsf{op}(g) = \mathtt{read}(x, m)$. Then $T' \vdash \mathtt{read}(x, m)$.

   Since $\mathcal{G}$ is a dependency graph, there exist two transactions S, V such that $S \xrightarrow{\mathsf{WR}_\mathcal{G}(x)} T$ and $V \xrightarrow{\mathsf{WR}_\mathcal{G}(x)} T'$. We have $S \vdash \mathtt{write}(x, n)$ and $V \vdash \mathtt{write}(x, m)$, so that $S \neq V$. Then either $S \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} V$ or $V \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} S$. We show that neither of these cases is possible.

   If $S \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} V$, then $T \xrightarrow{\mathsf{RW}_\mathcal{G}(x)} V$. This causes a cycle $T \xrightarrow{\mathsf{RW}_\mathcal{G}} V \xrightarrow{\mathsf{WR}_\mathcal{G}} T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T$, contradicting $\mathcal{G} \in \mathsf{GraphSI}$.

   On the other hand, if $V \xrightarrow{\mathsf{WW}_\mathcal{G}(x)} S$, then observe that $T' \neq S$, since $\neg(T' \vdash \mathtt{write}(x, \_))$ and $S \vdash \mathtt{write}(x, \_)$. Therefore $T' \xrightarrow{\mathsf{RW}_\mathcal{G}(x)} S$. We can show that $S \not\approx_\mathcal{G} T$ in a way similar to the one above, proving that neither of the cases $S \xrightarrow{\mathsf{SO}_\mathcal{G}} T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T$, $T' \xrightarrow{\mathsf{SO}_\mathcal{G}} S \xrightarrow{\mathsf{SO}_\mathcal{G}} T$ or $T' \xrightarrow{\mathsf{SO}_\mathcal{G}} T \xrightarrow{\mathsf{SO}_\mathcal{G}} S$ are possible. This leads to a critical cycle $T' \xrightarrow{(\mathsf{RW}_\mathcal{G} \setminus \approx_\mathcal{G})} S \xrightarrow{(\mathsf{WR}_\mathcal{G} \setminus \approx_\mathcal{G})} T \xrightarrow{\mathsf{SO}_\mathcal{G}^{-1}} T'$ in $\mathsf{DCG}(\mathcal{G})$, contradicting the assumptions of the lemma.

□

PROOF OF THEOREM 16. Let $\mathcal{G} \in \mathsf{GraphSI}$ be a dependency graph such that $\mathsf{DCG}(\mathcal{G})$ contains no critical cycle. We prove that $\mathsf{splice}(\mathcal{G}) \in \mathsf{GraphSI}$. First, Lemmas 26 and 27 ensure that $\mathsf{splice}(\mathcal{G})$ is indeed a dependency graph and $\mathcal{T}_{\mathsf{splice}(\mathcal{G})} \models \textsc{Int}$. Since $\mathsf{SO}_{\mathsf{splice}(\mathcal{G})} = \emptyset$, by Theorem 10 it remains to prove that the relation $((\mathsf{WR}_{\mathsf{splice}(\mathcal{G})} \cup \mathsf{WW}_{\mathsf{splice}(\mathcal{G})})\ ;\ \mathsf{RW}_{\mathsf{splice}(\mathcal{G})}?)$ is acyclic. The proof goes by contradiction: we assume that this relation contains a cycle and exhibit a critical cycle in $\mathsf{DCG}(\mathcal{G})$. Let

$$\gamma = \boxed{T_0}_\mathcal{G} \xrightarrow{\mathcal{C}_0} \ldots \xrightarrow{\mathcal{C}_{n-1}} \boxed{T_n}_\mathcal{G} \quad (n \geq 1)$$

be a cycle in $\mathsf{splice}(\mathcal{G})$, where

$$T_0, \ldots, T_n \in \mathcal{T}_\mathcal{G}, \quad \mathcal{C}_0, \ldots \mathcal{C}_{n-1} \in \{\mathsf{WR}_{\mathsf{splice}(\mathcal{G})}, \mathsf{WW}_{\mathsf{splice}(\mathcal{G})}, \mathsf{RW}_{\mathsf{splice}(\mathcal{G})}\}$$

(the letter $\mathcal{C}$ stands for *conflict*), $\boxed{T_n}_\mathcal{G} = \boxed{T_0}_\mathcal{G}$ (in particular, $T_n \approx_\mathcal{G} T_0$), and there is no index $i = 0..(n-1)$ such that $\mathcal{C}_i = \mathsf{RW}_{\mathsf{splice}(\mathcal{G})}$ and $\mathcal{C}_{(i+1) \mod n} = \mathsf{RW}_{\mathsf{splice}(\mathcal{G})}$. By Lemma 24, we can assume that γ is simple. Thus, for any $i, j = 0..(n-1)$ we have $\boxed{T_i}_\mathcal{G} = \boxed{T_j}_\mathcal{G}$ (equivalently, $T_i \approx_\mathcal{G} T_j$) only if $i = j$.

By applying Lemma 17, we can convert γ into the following path:

$$T_0' \approx_\mathcal{G} T_0'' \xrightarrow{\mathcal{C}_0^\mathcal{G}} T_1' \approx_\mathcal{G} T_1'' \xrightarrow{\mathcal{C}_2^\mathcal{G}} \ldots \xrightarrow{\mathcal{C}_{n-1}^\mathcal{G}} T_n' \approx_\mathcal{G} T_n'', \tag{5}$$

where for any $i = 0..n$, $T_i' \approx_\mathcal{G} T_i \approx_\mathcal{G} T_i''$ (note that because $T_n \approx_\mathcal{G} T_0$, this implies $T_n'' \approx_\mathcal{G} T_0'$), and for any $i = 0..(n-1)$, $\mathcal{C}_i^\mathcal{G}$ is the relation in $\mathcal{G}$ corresponding to the relation $\mathcal{C}_i$ in $\mathsf{splice}(\mathcal{G})$ (e.g., if $\mathcal{C}_i = \mathsf{WR}_{\mathsf{splice}(\mathcal{G})}$, then $\mathcal{C}_i^\mathcal{G} = (\mathsf{WR}_\mathcal{G} \setminus \approx_\mathcal{G})$). We also know that

$$\neg\exists i = 0..(n-1).\, \mathcal{C}_i^\mathcal{G} = (\mathsf{RW}_\mathcal{G} \setminus \approx_\mathcal{G}) \wedge \mathcal{C}_{(i+1) \mod n}^\mathcal{G} = (\mathsf{RW}_\mathcal{G} \setminus \approx_\mathcal{G}). \tag{6}$$

Since the cycle γ is simple, the only possibility for vertices to be repeated on the path (5) is when they are adjacent: $T'_i = T''_i$ for some $i = 0..(n-1)$. Recall that whenever $T \approx_{\mathcal{G}} S$, for some transaction $T, S \in \mathcal{T}_{\mathcal{G}}$, then one of the following holds: $T = S$, $T \xrightarrow{\mathsf{SO}_{\mathcal{G}}} S$ or $T \xrightarrow{\mathsf{SO}_{\mathcal{G}}^{-1}} S$. Also, we know that $T'_n \approx T_n \approx T_0 \approx T''_0$. Therefore, we can rewrite the path (5) as follows:

$$T'_n \xrightarrow{\mathcal{S}_0} T''_0 \xrightarrow{\mathcal{C}_1^{\mathcal{G}}} T'_1 \xrightarrow{\mathcal{S}_1} T''_1 \xrightarrow{\mathcal{C}_2^{\mathcal{G}}} \dots \xrightarrow{\mathcal{S}_{n-1}} T''_{n-1} \xrightarrow{\mathcal{C}_{n-1}^{\mathcal{G}}} T'_n, \tag{7}$$

where $\mathcal{S}_0, \dots \mathcal{S}_{n-1} \in \{\mathsf{SO}_{\mathcal{G}}?, \mathsf{SO}_{\mathcal{G}}^{-1}\}$ (the letter $\mathcal{S}$ stands for *siblings*). In this cycle repeated vertices are always adjacent and connected by a $\mathsf{SO}_{\mathcal{G}}?$-edge. By removing such edges from the cycle we obtain a simple cycle, where all the occurrences of $\mathsf{SO}_{\mathcal{G}}?$-edges are actually $\mathsf{SO}_{\mathcal{G}}$-edges; this is a cycle in $\mathsf{DCG}(\mathcal{G})$. Due to (6), in this cycle any two anti-dependency edges are separated by a read- or write-dependency edge. To prove this cycle yields a critical cycle in $\mathsf{DCG}(\mathcal{G})$, it remains to show that there exists an index $i = 0..(n-1)$ such that $\mathcal{S}_i = \mathsf{SO}_{\mathcal{G}}^{-1}$. This holds because, if we had $\mathcal{S}_i = \mathsf{SO}_{\mathcal{G}}$ for all $i = 0..(n-1)$, then we would obtain a cycle in $((\mathsf{SO}_{\mathcal{G}} \cup \mathsf{WR}_{\mathcal{G}} \cup \mathsf{WW}_{\mathcal{G}}) \; ; \; \mathsf{RW}_{\mathcal{G}}?)^+$, contradicting the assumption that $\mathcal{G} \in \mathsf{GraphSI}$. $\square$

B. ADDITIONAL MATERIAL ON TRANSACTION CHOPPING UNDER SI

In this section we compare our chopping criterion for SI to criteria that have been proposed for other consistency models: serializability (§B.1) and parallel SI (§B.2). We also discuss why we defined splicing over dependency graphs, rather than over executions (§B.3).

B.1 Comparison with Transaction Chopping under Serializability

We now compare our chopping criterion for SI to the one previously proposed for serializability. For clarity, we refer to critical cycles of §5 as ***SI-critical***. The following is an improved version of the chopping criterion by Shasha et al. [29].

DEFINITION 28. *A cycle in $\mathsf{SCG}(\mathcal{P})$ is **SER-critical** if: (i) it does not contain two occurrences of the same vertex; and (ii) it contains three consecutive edges of the form "conflict, predecessor, conflict".*

THEOREM 29. *The chopping defined by programs $\mathcal{P}$ is correct under serializability if $\mathsf{SCG}(\mathcal{P})$ contains no SER-critical cycles.*

Note that any SI-critical cycle is also SER-critical. As a consequence, any set of programs that is chopped correctly under SI is also chopped correctly under serializability. In particular, the chopping defined by the programs $\mathcal{P}_2$ considered in Figure 6 is correct under serializability. On the other hand, the chopping defined by $\mathcal{P}_1$ from Figure 5 is incorrect, and in fact $\mathsf{SCG}(\mathcal{P}_1)$ contains a SER-critical cycle:

$$(\mathtt{var1} = \mathtt{acct1}) \xrightarrow{\mathsf{RW}} (\mathtt{acct1} = \mathtt{acct1} - 100) \xrightarrow{\mathsf{S}} (\mathtt{acct2} = \mathtt{acct2} + 100) \xrightarrow{\mathsf{WR}} (\mathtt{var2} = \mathtt{acct2}) \xrightarrow{\mathsf{P}} (\mathtt{var1} = \mathtt{acct1}). \tag{8}$$

The programs $\mathcal{P}_3 = \{\mathtt{write1}, \mathtt{write2}\}$ in Figure 11 define a correct chopping under SI, but not under serializability. Their static chopping graph $\mathsf{SCG}(\mathcal{P}_3)$ is depicted on the bottom left of the same figure. There is only one cycle that contains no repeated vertices and three consecutive edges of the form "conflict, predecessor, conflict":

$$(\mathtt{var1} = \mathtt{x}) \xrightarrow{\mathsf{S}} (\mathtt{x} = \mathtt{var2}) \xrightarrow{\mathsf{P}} (\mathtt{var2} = \mathtt{y}) \xrightarrow{\mathsf{RW}} (\mathtt{y} = \mathtt{var1}) \xrightarrow{\mathsf{P}} (\mathtt{var1} = \mathtt{x}). \tag{9}$$

This cycle is not SI-critical, and by Corollary 18, the chopping defined by $\mathcal{P}_3$ is correct under SI. However, it is incorrect under serializability. Indeed, consider the execution $\mathcal{G}_6 \in \mathsf{GraphSER}$ depicted in Figure 11, which can be produced by $\mathcal{P}_3$. It is immediate to observe that $\mathsf{splice}(\mathcal{H}_{\mathcal{G}_6})$ is a write skew: $\mathsf{splice}(\mathcal{H}_{\mathcal{G}_6}) \notin \mathsf{HistSER}$.

B.2 Comparison with Transaction Chopping under Parallel SI

Next, we compare our chopping criterion for SI to the one that we recently proposed for PSI [11].

DEFINITION 30. *Let $\mathcal{P}$ be a transactional application. A cycle in $\mathsf{SCG}(\mathcal{P})$ is **PSI-critical** if: (i) it does not contain two instances of the same vertex; (ii) it contains three consecutive edges of the form "conflict, predecessor, conflict"; and (iii) it contains at most one anti-dependency edge.*

THEOREM 31. *The chopping defined by programs $\mathcal{P}$ is correct under PSI if $\mathsf{SCG}(\mathcal{P})$ contains no PSI-critical cycles.*

Note that any cycle that is PSI-critical, is also SI-critical; as a consequence, the sets of programs $\mathcal{P}_2$ and $\mathcal{P}_3$ considered in this section define a correct chopping under PSI. On the other hand, it is easy to see that $\mathcal{P}_1$ cannot be chopped correctly under PSI: the cycle (8) for $\mathsf{SCG}(\mathcal{P}_1)$ is PSI-critical.

The programs $\mathcal{P}_4 = \{\mathtt{write1}, \mathtt{write2}, \mathtt{read1}, \mathtt{read2}\}$ in Figure 12 define a correct chopping under PSI, but not SI. The static chopping graph $\mathsf{SCG}(\mathcal{P}_4)$ is depicted on the bottom left of the same figure. The graph contains exactly one cycle with no repeated vertices, and three consecutive vertices of the form "conflict, predecessor, conflict":

$$(\mathtt{x} = \mathtt{post1}) \xrightarrow{\mathsf{WR}} (\mathtt{b} = \mathtt{x}) \xrightarrow{\mathsf{P}} (\mathtt{a} = \mathtt{y}) \xrightarrow{\mathsf{RW}} (\mathtt{y} = \mathtt{post2}) \xrightarrow{\mathsf{WR}} (\mathtt{b} = \mathtt{y}) \xrightarrow{\mathsf{P}} (\mathtt{a} = \mathtt{x}) \xrightarrow{\mathsf{RW}} (\mathtt{x} = \mathtt{post1}). \tag{10}$$

This cycle is not PSI-critical, so that $\mathcal{P}_4$ indeed define a correct chopping under PSI. On the other hand, this cycle is SI-critical and $\mathcal{P}_4$ do not define a correct chopping under SI. Indeed, consider the dependency graph $\mathcal{G}_7 \in \mathsf{GraphSI}$ in Figure 12, which can be produced by $\mathcal{P}_4$. Splicing the history $\mathcal{H}_{\mathcal{G}_7}$ results in a long fork anomaly: $\mathsf{splice}(\mathcal{H}_{\mathcal{G}_7}) \notin \mathsf{HistSI}$. It follows that $\mathcal{P}_4$ does not define a correct chopping under SI.

```
session write1 { tx { var1 = x } ; tx { y = var1 } }
session write2 { tx { var2 = y } ; tx { x =var2 } }
```

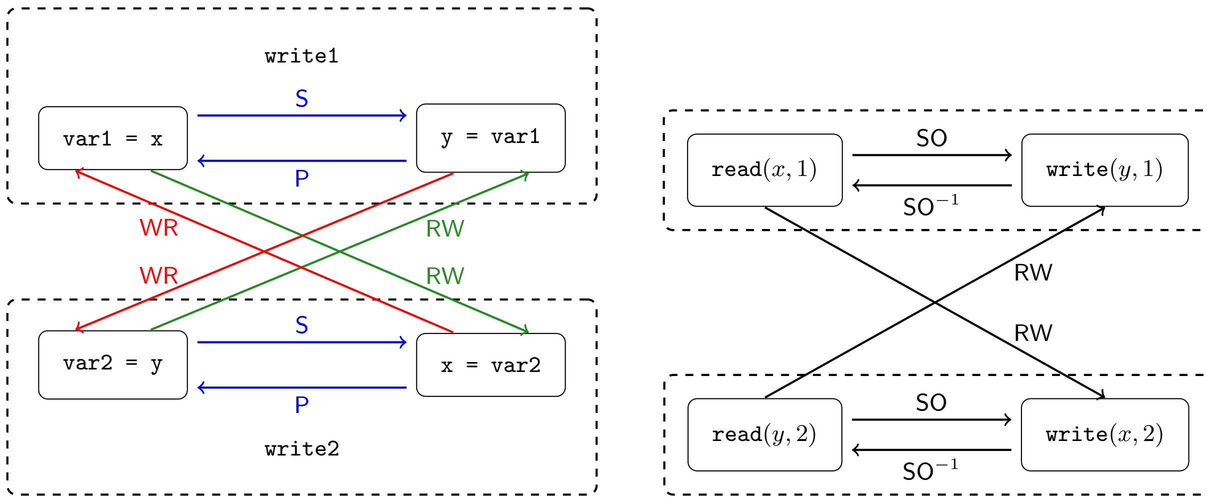

Figure 11: Example of a chopping correct under SI, but not under SER.

```
session write1 { tx { x = post1 } }
session write2 { tx { y = post2 } }
session read1 { tx { a = y }; tx { b = x }; return (a, b); }
session read2 { tx { a = x }; tx { b = y }; return (a, b); }
```

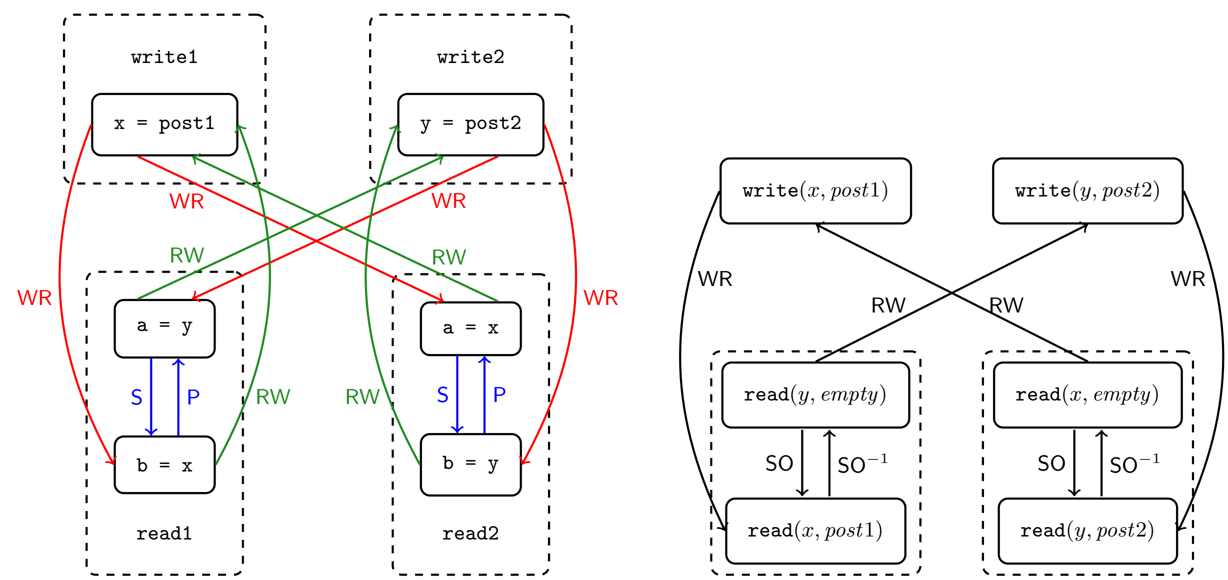

Figure 12: Example of a chopping correct under PSI, but not under SI.

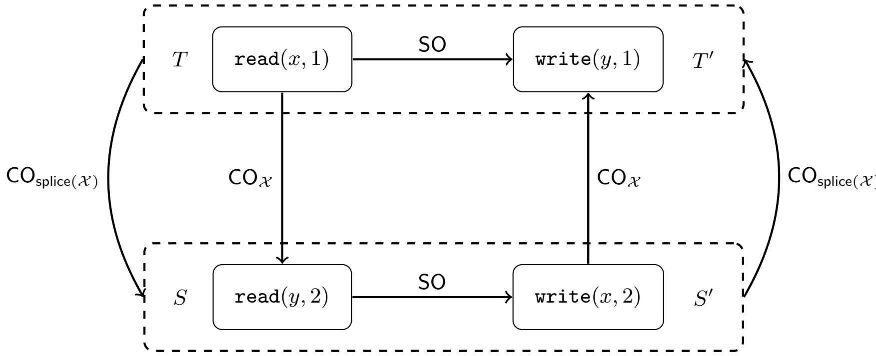

Figure 13: An attempt to splice an execution directly.

B.3 Challenges in Splicing Executions

In §5 we claimed that splicing abstract executions, rather than dependency graphs, is challenging. Here we illustrate why via a simple example. Consider the abstract execution $\mathcal{X}$ of Figure 13, which is in ExecSI. A straightforward way to define $\mathsf{splice}(\mathcal{X})$ is by letting

$$\boxed{T}_{\mathcal{X}} \xrightarrow{\mathsf{CO}_{\mathsf{splice}(\mathcal{X})}} \boxed{S}_{\mathcal{X}} \iff \exists T', S'.\, T \approx_{\mathcal{H}_{\mathcal{X}}} T' \xrightarrow{\mathsf{CO}_{\mathcal{X}}} S' \approx_{\mathcal{H}_{\mathcal{X}}} S,$$

and similarly for $\mathsf{VIS}_{\mathsf{splice}(\mathcal{X})}$. In this case we would have

$$\boxed{T}_{\mathcal{X}} \xrightarrow{\mathsf{CO}_{\mathsf{splice}(\mathcal{X})}} \boxed{S}_{\mathcal{X}} \xrightarrow{\mathsf{CO}_{\mathsf{splice}(\mathcal{X})}} \boxed{T}_{\mathcal{X}},$$

so that $\mathsf{CO}_{\mathsf{splice}(\mathcal{X})}$ is not irreflexive. Hence $\mathsf{splice}(\mathcal{X})$ is not a valid execution. On the other hand, by extracting a dependency graph $\mathcal{G}$ from $\mathcal{X}$ and computing $\mathsf{splice}(\mathcal{G})$, we easily obtain a dependency graph in $\mathsf{GraphSI}$. This allows us to construct an execution $\mathcal{X}'$ with the dependency graph $\mathsf{splice}(\mathcal{G})$ such that $\mathcal{X}' \in \mathsf{ExecSI}$.