

Analysing Snapshot Isolation

Andrea Cerone, Imperial College London, UK
Alexey Gotsman, IMDEA Software Institute, Spain

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 running under SI.

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 [Bernstein et al., 1987] and, more recently, by transactional memory systems [Herlihy and Moss, 1993]. Ideally, programmers would like to get strong guarantees about the isolation of transactional computations, formalised by the notion of *serializability* [Bernstein et al., 1987]: 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) [Berenson et al., 1995] is one of the most popular such models, implemented by major centralised databases (e.g., MS SQL Sever and Oracle), distributed databases [Daudjee and Salem, 2006; Serrano et al., 2007; Peng and Dabek, 2010] and transactional memory systems [Litz et al., 2014; Bieniusa and Fuhrmann, 2010; Dias et al., 2011; Closure, 2016; Riegel et al., 2006].

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(f). 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

This article is a revised and expanded version of a paper that received the Best Paper Award at the 35th Annual ACM Symposium on Principles of Distributed Computing (PODC).

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 ACM 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.

© YYYY ACM. 0000-0000/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

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. [Adya, 1999; Adya et al., 2000]. 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 [Bernstein et al., 1987]. 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 [Fekete et al., 2005; Jorwekar et al., 2007; Shasha et al., 1995; Zhang et al., 2013; Cerone et al., 2015b], run-time monitoring [Cahill et al., 2009; Zellag and Kemme, 2014] and proving concurrency-control algorithms correct [Lin et al., 2009; Xie et al., 2015; Diegues and Romano, 2014]. 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 [Adya, 1999], as well as more recent proposals of consistency models [Bailis et al., 2014; Xie et al., 2015]. 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 [Adya, 1999]. 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* [Shasha et al., 1995; Xiang and Scott, 2015; Afek et al., 2011]—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 [Shasha et al., 1995] and parallel SI [Cerone et al., 2015b], a recently proposed weaker version of SI for large-scale databases [Sovran

et al., 2011; Saeida Ardekani et al., 2013a]. 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* [Terry et al., 1994; Daudjee and Salem, 2006] (aka *chains* [Zhang et al., 2013]) 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* [Fekete et al., 2005; Shasha and Snir, 1988] 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 [Fekete et al., 2005] (robustness *against SI*, §6.1). We then propose a new static analysis that checks whether an application executing under the recently proposed parallel SI [Sovran et al., 2011; Saeida Ardekani et al., 2013a] 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.

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 [Cerone et al., 2015a]. 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* $\text{Obj} = \{x, y, \dots\}$. 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 $\text{Event} = \{e, f, \dots\}$. A function $\text{op} : \text{Event} \rightarrow \text{Op}$ for

$$\text{Op} = \{\text{read}(x, n), \text{write}(x, n) \mid x \in \text{Obj}, n \in \mathbb{Z}\}$$

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 2.1. A **transaction** T, S, \dots is a pair (E, po) , where $E \subseteq \text{Event}$ is a finite, non-empty set of events and the **program order** $\text{po} \subseteq E \times E$ 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 [Adya, 1999; Doherty et al., 2013; Guerraoui and Kapalka, 2008]. We denote components of transactions and similar structures as in E_T and po_T . We write $e \xrightarrow{\text{po}} f$ and $(e, f) \in \text{po}$ interchangeably, and similarly for other relations.

To allow transaction chopping (§5), we assume that the transactional system allows its clients to group several transactions into a session [Terry et al., 1994], which establishes an ordering on the transactions. Thus, instead of classical SI and serializability, we actually define their **strong session** variants [Daudjee and Salem, 2006; Daudjee and Salem, 2004]. We represent the client-visible results of executing a set of sessions by a *history*.

Definition 2.2. A **history** is a pair $\mathcal{H} = (\mathcal{T}, \text{SO})$, where \mathcal{T} is a finite set of transactions with disjoint sets of events and the **session order** $\text{SO} \subseteq \mathcal{T} \times \mathcal{T}$ 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 2.3. An **abstract execution** (or just an execution) is a tuple

$$\mathcal{X} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO}),$$

where (\mathcal{T}, SO) is a history and the **visibility** and **commit orders** $\text{VIS}, \text{CO} \subseteq \mathcal{T} \times \mathcal{T}$ are such that $\text{VIS} \subseteq \text{CO}$ and CO is total.

For $\mathcal{H} = (\mathcal{T}, \text{SO})$ we shorten $(\mathcal{T}, \text{SO}, \text{VIS}, \text{CO})$ to $(\mathcal{H}, \text{VIS}, \text{CO})$. In terms of the SI concurrency-control algorithm sketched in §1, $T \xrightarrow{\text{VIS}} S$ means that the writes done by the transaction T are included into the snapshot taken by the transaction S ; $T \xrightarrow{\text{CO}} S$ means that T commits earlier than S . The constraint $\text{VIS} \subseteq \text{CO}$ 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 2.4. The sets of executions and histories **allowed by (strong session) SI and serializability** are:

$$\begin{aligned} \text{ExecSI} &= \{\mathcal{X} \mid \mathcal{X} \models \text{INT} \wedge \text{EXT} \wedge \text{SESSION} \wedge \text{PREFIX} \wedge \text{NOCONFLICT}\}; \\ \text{ExecSER} &= \{\mathcal{X} \mid \mathcal{X} \models \text{INT} \wedge \text{EXT} \wedge \text{SESSION} \wedge \text{TOTALVIS}\}; \\ \text{HistSI} &= \{\mathcal{H} \mid \exists \text{VIS}, \text{CO}. (\mathcal{H}, \text{VIS}, \text{CO}) \in \text{ExecSI}\}; \\ \text{HistSER} &= \{\mathcal{H} \mid \exists \text{VIS}, \text{CO}. (\mathcal{H}, \text{VIS}, \text{CO}) \in \text{ExecSER}\}. \end{aligned}$$

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 $\mathcal{R} \subseteq A \times A$, we let $\max_{\mathcal{R}}(A)$ be the element $a \in A$ such that $\forall b \in A. a = b \vee (b, a) \in \mathcal{R}$; if $A = \emptyset$, then $\max_{\mathcal{R}}(A)$ is undefined. In the following, the use of $\max_{\mathcal{R}}(A)$ in an expression implicitly assumes that it is defined. We define $\min_{\mathcal{R}}(A)$ similarly. For a relation $\mathcal{R} \subseteq A \times A$ and an element $a \in A$, we let $\mathcal{R}^{-1}(a) = \{b \mid (b, a) \in \mathcal{R}\}$. We define the sequential composition of relations \mathcal{R}_1 and \mathcal{R}_2 as

$$\mathcal{R}_1 ; \mathcal{R}_2 = \{(a, b) \mid \exists c. (a, c) \in \mathcal{R}_1 \wedge (c, b) \in \mathcal{R}_2\}.$$

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

Fig. 1. Axioms constraining an abstract execution $(\mathcal{T}, \text{SO}, \text{VIS}, \text{CO})$.

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. In particular, the axiom ensures that, if a transaction writes to an object and then reads the object, then it will observe its last write. The axiom also disallows *unrepeatable reads*: if a transaction reads an object twice without writing to it in-between, then it will read the same value in both cases.

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 $T = (E, \text{po})$, we let $T \vdash \text{write}(x, n)$ if T writes to x and the last value written is n :

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

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

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

We also let $\text{WriteTx}_x = \{T \mid T \vdash \text{write}(x, _)\}$. 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. Note that the read in T_3 in Figure 2(b) returns 25, because T_1 precedes T_2 in the commit order. The axiom EXT implies the absence of *dirty reads*: a committed transaction cannot read a value written by an aborted or ongoing transaction (as such a transaction are not present in abstract executions), and a transaction also cannot read a value that was overwritten by the transaction that wrote it (ensured by the definition of $T \vdash \text{write}(x, n)$). Finally, EXT guarantees that either all or none of writes by a transaction can be visible to another transaction. For example, EXT disallows the execution in Figure 2(c) and, in fact, any execution with the same history. This illustrates

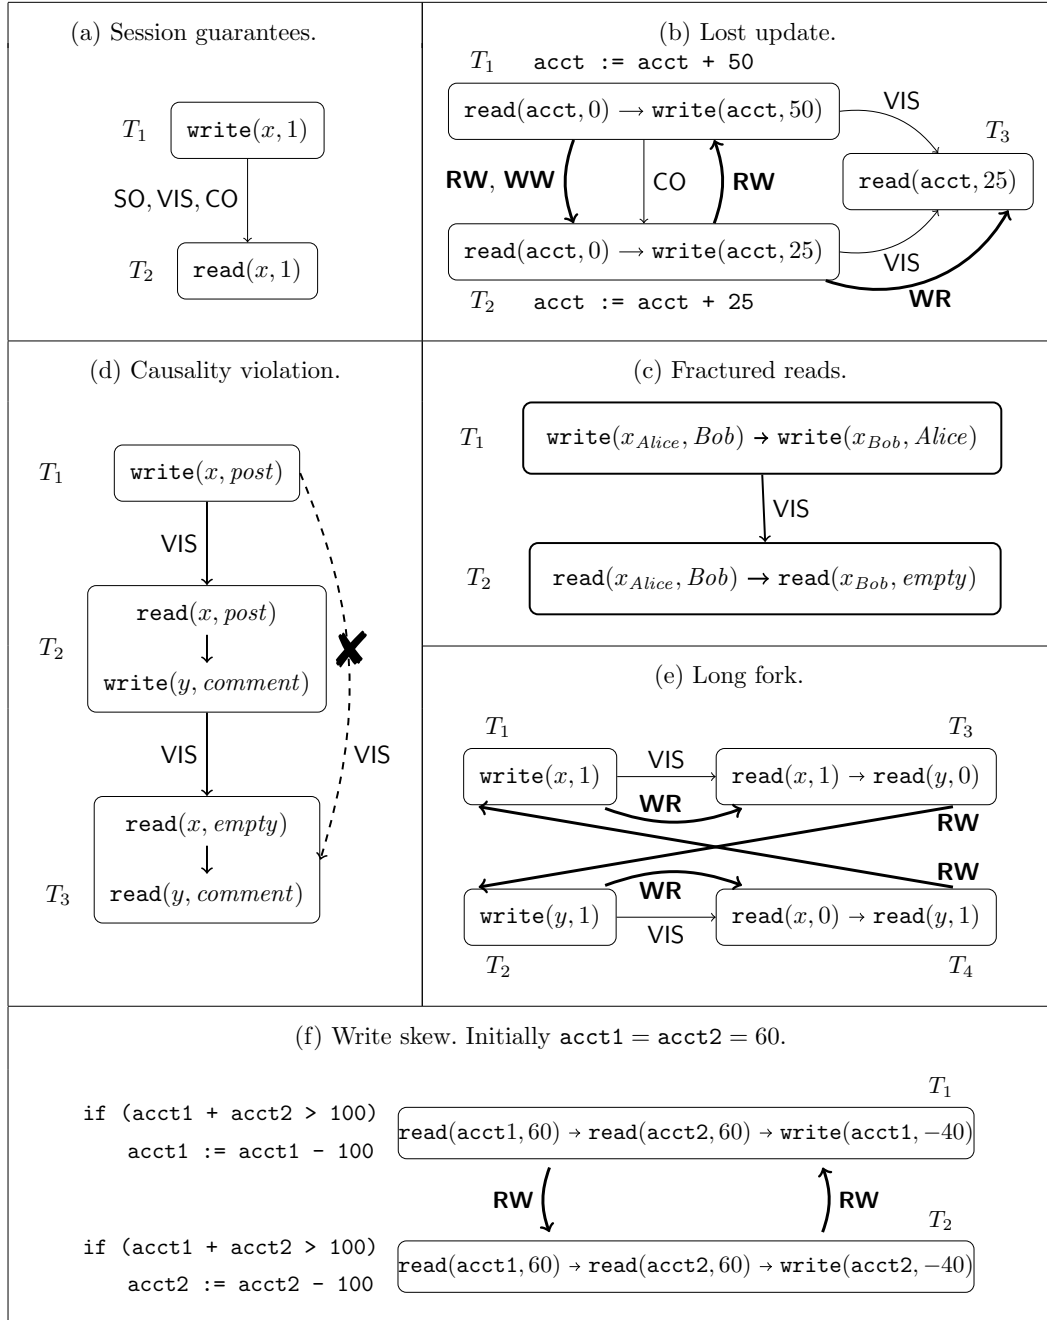


Fig. 2. Abstract executions illustrating SI and serializability. Boxes represent transactions, and arrows inside boxes represent the program order. We omit irrelevant CO edges. We also omit a special transaction that writes initial versions of all objects and precedes all the other transactions in VIS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

a **fractured reads** anomaly: T_1 makes *Alice* and *Bob* friends, but T_2 observes only one direction of the friendship relationship.

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** [Elnikety et al., 2005]. However, following strong session SI [Daudjee and Salem, 2006; Daudjee and Salem, 2004], 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 $\text{VIS} \subseteq \text{CO}$ in Definition 2.4 imply that VIS is transitive. Hence, PREFIX disallows the **causality violation** anomaly in Figure 2(d): here T_3 sees the effects of T_2 , but not the effects of T_1 , which is seen by T_2 . Since VIS is transitive, we must have a VIS edge between T_1 and T_2 ; but then by EXT, T_3 has to read *comment* from y . PREFIX also disallows the **long fork** anomaly shown in Figure 2(e), which is allowed by parallel SI [Sovran et al., 2011; Saeida Ardekani et al., 2013a], a weakening of SI. 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(e), 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 2.4) defined using the consistency axioms explained so far is exactly the one produced by the SI concurrency-control algorithm [Cerone et al., 2015a]. The axioms allow the execution in Figure 2(f) 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 $\text{HistSER} \subset \text{HistSI}$.

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. [Adya, 1999; Adya et al., 2000].

Definition 3.1. Let $\mathcal{X} = (\mathcal{H}, \text{VIS}, \text{CO})$ be an execution. For $x \in \text{Obj}$, we define the following relations on $\mathcal{T}_{\mathcal{H}}$:

— **read dependency:**

$$T \xrightarrow{\text{WR}_{\mathcal{X}}(x)} S \iff S \vdash \text{read}(x, _) \wedge T = \max_{\text{CO}}(\text{VIS}^{-1}(S) \cap \text{WriteTx}_x);$$

— **write dependency:**

$$T \xrightarrow{\text{WW}_{\mathcal{X}}(x)} S \iff T \xrightarrow{\text{CO}} S \wedge T, S \in \text{WriteTx}_x;$$

— **anti-dependency:**

$$T \xrightarrow{\text{RW}_{\mathcal{X}}(x)} S \iff T \neq S \wedge \exists T'. T' \xrightarrow{\text{WR}_{\mathcal{X}}(x)} T \wedge T' \xrightarrow{\text{WW}_{\mathcal{X}}(x)} S.$$

Informally, $T \xrightarrow{\text{WR}_{\mathcal{X}}(x)} S$ means that S reads T 's write to x (cf. the EXT axiom in Figure 1); $T \xrightarrow{\text{WW}_{\mathcal{X}}(x)} S$ means that S overwrites T 's write to x ; $T \xrightarrow{\text{RW}_{\mathcal{X}}(x)} S$ means that S overwrites the write to x read by T . For example, the dependencies of the executions in Figures 2(b), 2(e) and 2(f) 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 $\text{WR}_{\mathcal{X}}$ to also denote the relation $\bigcup_{x \in \text{Obj}} \text{WR}_{\mathcal{X}}(x) \subseteq \mathcal{T}_{\mathcal{H}} \times \mathcal{T}_{\mathcal{H}}$, and similarly for $\text{WW}_{\mathcal{X}}$ and $\text{RW}_{\mathcal{X}}$.

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 2.4. To this end, we extend histories to *dependency graphs* (aka direct serialization graphs) [Adya, 1999], which include relations representing the dependencies.

Definition 3.2. A **dependency graph** is a tuple $\mathcal{G} = (\mathcal{T}, \text{SO}, \text{WR}, \text{WW}, \text{RW})$, where (\mathcal{T}, SO) is a history and

- (1) $\text{WR} : \text{Obj} \rightarrow 2^{\mathcal{T} \times \mathcal{T}}$ is such that:
 - (a) $\forall T, S \in \mathcal{T}. \forall x. T \xrightarrow{\text{WR}(x)} S \implies \exists n. T \neq S \wedge T \vdash \text{write}(x, n) \wedge S \vdash \text{read}(x, n)$;
 - (b) $\forall S \in \mathcal{T}. \forall x. S \vdash \text{read}(x, -) \implies \exists T. T \xrightarrow{\text{WR}(x)} S$;
 - (c) $\forall T, T', S \in \mathcal{T}. \forall x. (T \xrightarrow{\text{WR}(x)} S \wedge T' \xrightarrow{\text{WR}(x)} S) \implies T = T'$.
- (2) $\text{WW} : \text{Obj} \rightarrow 2^{\mathcal{T} \times \mathcal{T}}$ is such that for every $x \in \text{Obj}$, $\text{WW}(x)$ is a total order on the set WriteTx_x ;
- (3) $\text{RW} : \text{Obj} \rightarrow 2^{\mathcal{T} \times \mathcal{T}}$ is derived from WR and WW as in Definition 3.1:

$$\forall T, S \in \mathcal{T}. \forall x. T \xrightarrow{\text{RW}(x)} S \iff T \neq S \wedge \exists T'. T' \xrightarrow{\text{WR}(x)} T \wedge T' \xrightarrow{\text{WW}(x)} S.$$

PROPOSITION 3.3. For any $\mathcal{X} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO}) \in \text{ExecSI}$,

$$\text{graph}(\mathcal{X}) = (\mathcal{T}, \text{SO}, \text{WR}_{\mathcal{X}}, \text{WW}_{\mathcal{X}}, \text{RW}_{\mathcal{X}})$$

is a dependency graph.

Note that the constraints on WR in Definition 3.2 ensure that it uniquely determines the values read by transactions. For $\mathcal{H} = (\mathcal{T}, \text{SO})$ we write $(\mathcal{H}, \text{WR}, \text{WW}, \text{RW})$ for $(\mathcal{T}, \text{SO}, \text{WR}, \text{WW}, \text{RW})$.

We write $\mathcal{T} \models \text{INT}$ if a set of transactions \mathcal{T} satisfies the internal consistency axiom INT in Figure 1. A relation $\mathcal{R} \subseteq \mathcal{T} \times \mathcal{T}$ is **acyclic** if $\mathcal{R}^+ \cap \{(T, T) \mid T \in \mathcal{T}\} = \emptyset$, where \mathcal{R}^+ is the transitive closure of \mathcal{R} . (Strong session) serializability can be characterised by the set of acyclic dependency graphs with internally consistent transactions [Adya, 1999].

THEOREM 3.4. Let

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

Then

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

For example, the histories in Figures 2(b), 2(e) and 2(f) are not serializable, and they cannot be extended to acyclic dependency graphs; in particular, the graphs $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$, shown with bold edges in figure 2(b), 2(e) and 2(f), respectively, satisfy the conditions of Definition 3.2, 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 $\mathcal{R} \subseteq \mathcal{T} \times \mathcal{T}$ let $\mathcal{R}^? = \mathcal{R} \cup \{(T, T) \mid T \in \mathcal{T}\}$. We show that (strong session) SI is characterised by dependency graphs that contain only cycles with at least two adjacent anti-dependency edges.

THEOREM 4.1. *Let*

$$\text{GraphSI} = \{\mathcal{G} \mid (\mathcal{T}_{\mathcal{G}} \models \text{INT}) \wedge (((\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}}) ; \text{RW}_{\mathcal{G}}^?) \text{ is acyclic})\}.$$

Then

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

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(f) is allowed by SI, and indeed the dependency graph \mathcal{G}_3 shown in the figure contains only cycles with two adjacent anti-dependencies (e.g., $T_1 \xrightarrow{\text{RW}} T_2 \xrightarrow{\text{RW}} T_1$). In contrast, the histories in Figures 2(b) and 2(e) 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 \mathcal{G}_1 and \mathcal{G}_2 shown in Figures 2(b) and 2(e), respectively, contain cycles without these: e.g., $T_1 \xrightarrow{\text{WW}} T_2 \xrightarrow{\text{RW}} T_1$ in Figure 2(b) and $T_1 \xrightarrow{\text{WR}} T_3 \xrightarrow{\text{RW}} T_2 \xrightarrow{\text{WR}} T_4 \xrightarrow{\text{RW}} T_1$ in Figure 2(e).

To prove Theorem 4.1, 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 4.2.

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

As we explain in §7, the easier completeness direction of this theorem actually follows from existing results [Fekete et al., 2005]. 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).

4.1. Relationships between Abstract Executions and Dependency Graphs

We start by establishing several relationships between abstract executions in **ExecSI** and dependency graphs in **GraphSI** that are useful in proving Theorem 4.2. We introduce them by example.

First, consider the dependency graph \mathcal{G}_4 in Figure 3 and an abstract execution $\mathcal{X} \in \text{ExecSI}$ such that $\text{graph}(\mathcal{X}) = \mathcal{G}_4$. We show that we cannot have $T_4 \xrightarrow{\text{VIS}_{\mathcal{X}}} T_2$. Indeed, if we had $T_4 \xrightarrow{\text{VIS}_{\mathcal{X}}} T_2$, then by Definition 3.1 we would have $T_4 \xrightarrow{\text{WR}_{\mathcal{G}_4}(x)} T_2$, contradicting the hypothesis that $T_3 \xrightarrow{\text{WR}_{\mathcal{G}_4}(x)} T_2$. Hence, $\neg(T_4 \xrightarrow{\text{VIS}_{\mathcal{X}}} T_2)$. This fact is not unique to the dependency graph \mathcal{G}_4 , as we show in the following proposition.

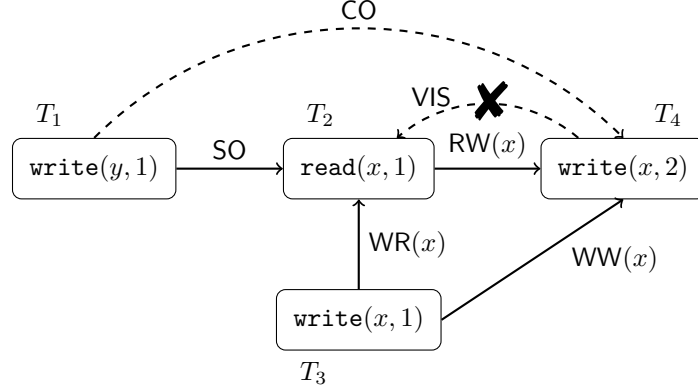


Fig. 3. An illustration of a correspondence between abstract executions and dependency graphs.




Fig. 4. An illustration of the proof of Proposition 4.3.

PROPOSITION 4.3.

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

PROOF. Fix $\mathcal{X} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO}) \in \text{ExecSI}$ and $T, S \in \mathcal{T}$. Let $\text{graph}(\mathcal{X}) = (\mathcal{T}, \text{SO}, \text{WR}, \text{WW}, \text{RW})$. We illustrate the proof in Figure 4.

“ \implies ”. Assume $S \xrightarrow{\text{RW}} T$. Then for some $x \in \text{Obj}$ and $T' \in \mathcal{T}$ we have $T' \xrightarrow{\text{WR}(x)} S$ and $T' \xrightarrow{\text{WW}(x)} T$. This implies $S \neq T$, $S \vdash \text{read}(x, -)$, $T \vdash \text{write}(x, -)$ and $T' = \max_{\text{CO}}(\text{VIS}^{-1}(S) \cap \text{WriteT}_{x,x})$. Since $T' \xrightarrow{\text{WW}(x)} T$, we have $T' \xrightarrow{\text{CO}} T$, and since $T \vdash \text{write}(x, -)$, we have $T \in \text{WriteT}_{x,x}$. Together with $T' = \max_{\text{CO}}(\text{VIS}^{-1}(S) \cap \text{WriteT}_{x,x})$, these statements imply $T \notin \text{VIS}^{-1}(S)$, i.e., $\neg(T \xrightarrow{\text{VIS}} S)$.

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

$$\left\{ \begin{array}{l} \text{SO} \cup \text{WR} \cup \text{WW} \subseteq \text{VIS} \\ \text{CO} ; \text{VIS} \subseteq \text{VIS} \\ \text{VIS} \subseteq \text{CO} \\ \text{CO} ; \text{CO} \subseteq \text{CO} \\ \text{VIS} ; \text{RW} \subseteq \text{CO} \end{array} \right. \begin{array}{l} \text{(S1)} \\ \text{(S2)} \\ \text{(S3)} \\ \text{(S4)} \\ \text{(S5)} \end{array}$$

Fig. 5. Requirements on a pre-execution $\mathcal{P} = (\mathcal{H}, \text{VIS}, \text{CO})$ constructed from a dependency graph $\mathcal{G} = (\mathcal{H}, \text{WR}, \text{WW}, \text{RW})$.

We next show that Proposition 4.3 implies an interesting relationship among anti-dependencies, visibility and the commit order. Consider again the dependency graph \mathcal{G}_4 of Figure 3 and $\mathcal{X} \in \text{ExecSI}$ such that $\text{graph}(\mathcal{X}) = \mathcal{G}_4$. Even though transactions T_1 and T_4 access different objects, we cannot choose the commit order between them arbitrarily: we must have $T_1 \xrightarrow{\text{CO}_{\mathcal{X}}} T_4$. This is because (SESSION) guarantees $T_1 \xrightarrow{\text{VIS}_{\mathcal{X}}} T_2$. Then, if we had $T_4 \xrightarrow{\text{CO}_{\mathcal{X}}} T_1$, then by (PREFIX) we would have $T_4 \xrightarrow{\text{VIS}_{\mathcal{X}}} T_2$. But this contradicts the edge $T_2 \xrightarrow{\text{RW}_{\mathcal{X}}} T_4$ and Proposition 4.3, according to which $\neg(T_4 \xrightarrow{\text{VIS}_{\mathcal{X}}} T_2)$. Again, the pattern illustrated in this example is not unique to the dependency graph \mathcal{G}_4 , as we now show.

LEMMA 4.4. $\forall \mathcal{X} \in \text{ExecSI}. \text{VIS}_{\mathcal{X}} ; \text{RW}_{\mathcal{X}} \subseteq \text{CO}_{\mathcal{X}}$.

PROOF. Consider $\mathcal{X} \in \text{ExecSI}$ and $T, S', S \in \mathcal{T}_{\mathcal{X}}$ such that $S' \xrightarrow{\text{VIS}_{\mathcal{X}}} S \xrightarrow{\text{RW}_{\mathcal{X}}} T$. If $T = S'$, then $S' \xrightarrow{\text{VIS}_{\mathcal{X}}} S \xrightarrow{\text{RW}_{\mathcal{X}}} S'$, contradicting Proposition 4.3. If $T \xrightarrow{\text{CO}_{\mathcal{X}}} S'$ then by PREFIX we get $T \xrightarrow{\text{VIS}_{\mathcal{X}}} S$, contradicting Proposition 4.3. Then, since $\text{CO}_{\mathcal{X}}$ is total, we must have $S' \xrightarrow{\text{CO}_{\mathcal{X}}} T$. \square

4.2. Pre-Executions

The main challenge in the proof of Theorem 4.2(i) 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 2.4). 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 4.5. A tuple $\mathcal{P} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO})$ is a **pre-execution** if it satisfies all the conditions of Definition 2.3, 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):

$$\text{PreExecSI} = \{ \mathcal{P} \mid \mathcal{P} \models \text{INT} \wedge \text{EXT} \wedge \text{SESSION} \wedge \text{PREFIX} \wedge \text{NoCONFLICT} \}.$$

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.

PROPOSITION 4.6. *For any $\mathcal{P} \in \text{PreExecSI}$, $\text{graph}(\mathcal{P})$ is a dependency graph.*

Towards proving Theorem 4.2(i), in this section, for a dependency graph

$$\mathcal{G} = (\mathcal{T}, \text{SO}, \text{WR}, \text{WW}, \text{RW}) \in \text{GraphSI}$$

we construct a pre-execution

$$\mathcal{P} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO}) \in \text{PreExecSI}$$

such that $\text{graph}(\mathcal{P}) = \mathcal{G}$. In §4.3 we show how to extend \mathcal{P} to a desired execution $\mathcal{X} \in \text{ExecSI}$.

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 5. First,

to ensure $\text{graph}(\mathcal{P}) = \mathcal{G}$, by Definition 3.1 at the very least we must have $\text{WR} \cup \text{WW} \subseteq \text{VIS}$. For \mathcal{P} to satisfy the `SESSION` axiom we must also have $\text{SO} \subseteq \text{VIS}$. These two observations motivate (S1). This inequality also implies that \mathcal{P} satisfies `NOCONFLICT`, since according to Definition 3.1, `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 4.5 from Definition 2.3. Inequality (S4) requires `CO` to be transitive; (S2) and (S3) ensure that so is `VIS`.

As we now show, (S5), motivated by Lemma 4.4, ensures the axiom `EXT`. More generally, the system of inequalities in Figure 5 can be used to ensure that a pre-execution $\mathcal{P} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO})$ satisfies all SI axioms save `INT`.

LEMMA 4.7. *Let $\mathcal{G} = (\mathcal{T}, \text{SO}, \text{WR}, \text{WW}, \text{RW})$ be a dependency graph such that $\mathcal{T} \models \text{INT}$ and $\text{VIS}, \text{CO} \subseteq \mathcal{T} \times \mathcal{T}$ be acyclic relations satisfying the system of inequalities in Figure 5. Then $\mathcal{P} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO})$ is a pre-execution such that $\mathcal{P} \in \text{PreExecSI}$ and $\text{graph}(\mathcal{P}) = \mathcal{G}$.*

PROOF. We only prove that $\mathcal{P} \models \text{EXT}$ and $\text{WR}_{\mathcal{P}} = \text{WR}$; discharging the other obligations is straightforward. Consider $S \in \mathcal{T}$ such that $S \vdash \text{read}(x, n)$. Then there exists a unique T' such that $T' \xrightarrow{\text{WR}(x)} S$. Let $T = \max_{\text{CO}}(\text{VIS}^{-1}(S) \cap \text{WriteTx}_x)$. This is defined because: `CO` is acyclic; by (S1) and (S3) we have $\text{WW} \subseteq \text{CO}$, so that `CO` is total over `WriteTxx`; and by (S1) we have $\text{WR} \subseteq \text{VIS}$, so that $T' \in \text{VIS}^{-1}(S) \cap \text{WriteTx}_x$. Hence, either $T = T'$ or $T' \xrightarrow{\text{CO}} T$. We now show that the latter case is impossible. Since $T \xrightarrow{\text{VIS}} S$, we cannot have $S \xrightarrow{\text{RW}(x)} T$, for in this case (S5) would imply $T \xrightarrow{\text{CO}} T$, contradicting the acyclicity of `CO`. Hence, $\neg(S \xrightarrow{\text{RW}(x)} T)$. Since $S \vdash \text{read}(x, _)$, $T \vdash \text{write}(x, _)$ and $T' \xrightarrow{\text{WR}(x)} S$, this can be the case only if $\neg(T' \xrightarrow{\text{WW}(x)} T)$, which implies $\neg(T' \xrightarrow{\text{CO}} T)$. Thus, we must have $T = T'$, which entails the required. \square

According to Lemma 4.7, to construct a desired pre-execution $\mathcal{P} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO}) \in \text{PreExecSI}$ from a dependency graph $\mathcal{G} = (\mathcal{T}, \text{SO}, \text{WR}, \text{WW}, \text{RW})$, it is sufficient to find a solution to the system of inequalities in Figure 5 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 $*$ to denote the transitive and reflexive closure of a given relation.

LEMMA 4.8. *Let $\mathcal{G} = (\mathcal{T}, \text{SO}, \text{WR}, \text{WW}, \text{RW})$ be a dependency graph. For any $R \subseteq \mathcal{T} \times \mathcal{T}$, the relations*

$$\begin{aligned} \text{VIS} &= (((\text{SO} \cup \text{WR} \cup \text{WW}) ; \text{RW}?) \cup R)^* ; (\text{SO} \cup \text{WR} \cup \text{WW}); \\ \text{CO} &= (((\text{SO} \cup \text{WR} \cup \text{WW}) ; \text{RW}?) \cup R)^+ \end{aligned} \tag{1}$$

are a solution to the system of inequalities in Figure 5. They also are the smallest solution to the system for which $\text{CO} \supseteq R$: for any other solution $(\text{VIS}', \text{CO}')$ with $\text{CO}' \supseteq R$ we have $\text{VIS} \subseteq \text{VIS}'$ and $\text{CO} \subseteq \text{CO}'$.

PROOF. We first prove that the relations in the statement of the lemma are indeed a solution to the system of inequalities in Figure 5.

(S1).

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

(S2).

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

(S3).

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

(S4).

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

(S5).

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

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

$$SO \cup WR \cup WW \subseteq VIS'. \quad (2)$$

Then by (S3) and (S5) we get

$$(SO \cup WR \cup WW) ; RW? \subseteq CO'.$$

Since $R \subseteq CO'$, this implies

$$((SO \cup WR \cup WW) ; RW?) \cup R \subseteq CO',$$

and by (S4) we get

$$(((SO \cup WR \cup WW) ; RW?) \cup R)^+ \subseteq (CO')^+ \subseteq CO'.$$

But this is exactly $CO \subseteq CO'$.

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

$$VIS = (SO \cup WR \cup WW) \cup (((SO \cup WR \cup WW) ; RW?) \cup R)^+ ; (SO \cup WR \cup WW),$$

and we prove that both parameters of the union are included in VIS' . We have already proved (2). We also have

$$(((SO \cup WR \cup WW) ; RW?) \cup R)^+ ; (SO \cup WR \cup WW) = CO ; (SO \cup WR \cup WW).$$

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

$$CO ; (SO \cup WR \cup WW) \subseteq CO' ; VIS' \subseteq VIS',$$

as required. \square

For $R = \emptyset$, Lemma 4.8 gives the smallest solution to the system of inequalities in Figure 5, which we denote by (VIS_0, CO_0) . We now show that this solution gives us the pre-execution we originally set out to construct.

COROLLARY 4.9. *If*

$$\mathcal{G} = (\mathcal{T}, SO, WR, WW, RW) \in \text{GraphSI},$$

then

$$\mathcal{P}_0 = (\mathcal{T}, SO, VIS_0, CO_0) \in \text{PreExecSI}$$

and $\text{graph}(\mathcal{P}_0) = \mathcal{G}$.

PROOF. Since $\mathcal{G} \in \text{GraphSI}$, CO_0 is acyclic. By (S3), so is VIS_0 . The required then follows from Lemma 4.7. \square

4.3. From Pre-Executions to Executions: Proof of Theorem 4.2(i)

Fix $\mathcal{G} = (\mathcal{T}, SO, WR, WW, RW) \in \text{GraphSI}$. Let $\mathcal{P}_0 = (\mathcal{T}, SO, VIS_0, CO_0)$ be the pre-execution constructed in Corollary 4.9, so that $\mathcal{P}_0 \in \text{PreExecSI}$ and $\text{graph}(\mathcal{P}_0) = \mathcal{G}$. To prove Theorem 4.2(i), we now show how to extend \mathcal{P} to a desired execution $\mathcal{X} \in \text{ExecSI}$ such that $\text{graph}(\mathcal{P}) = \mathcal{G}$. To this end, we take an incremental approach: we construct a sequence of pre-executions

$$\{\mathcal{P}_i = (\mathcal{T}, SO, VIS_i, CO_i)\}_{i=0}^n \subseteq \text{PreExecSI}$$

for some $n \geq 0$. The sequence is such that $\text{graph}(\mathcal{P}_i) = \mathcal{G}$ for $i = 0..n$, and $CO_i \subset CO_{i+1}$ for $i = 0..(n-1)$. Furthermore, CO_n is total, so that \mathcal{P}_n is an execution.

To define the above sequence, assume that \mathcal{P}_i has been defined. To construct \mathcal{P}_{i+1} , we choose two transactions (if any) T_i, S_i that are not related by CO_i and define (VIS_{i+1}, CO_{i+1}) to be the smallest solution to the system of inequalities of Figure 5 that contains $CO_i \cup \{(T_i, S_i)\}$. If the above transactions T_i, S_i do not exist, we terminate the construction and let $n = i$. To prove that the constructed sequence satisfies the properties stated above, we need to show that CO_{i+1} is acyclic. This is established using the following lemma.

LEMMA 4.10. *Let $\mathcal{P}_i = (\mathcal{T}, SO, VIS_i, CO_i) \in \text{PreExecSI}$, where (VIS_i, CO_i) is a solution to the system of inequalities in Figure 5. Assume that $T_i, S_i \in \mathcal{T}$ are two different transactions not related by CO_i , and (VIS_{i+1}, CO_{i+1}) is the smallest solution to the system of inequalities in Figure 5 such that $CO_i \cup \{(T_i, S_i)\} \subseteq CO_{i+1}$. Then CO_{i+1} is acyclic.*

PROOF. To prove the required, we exploit the fact that Lemma 4.8 can be used to express the relation CO_{i+1} as a function of CO_i and (T_i, S_i) . Namely, by Lemma 4.8 we have:

$$\begin{aligned} CO_{i+1} &= (((SO \cup WR \cup WW) ; RW?) \cup (CO_i \cup \{(T_i, S_i)\}))^+ \\ &= (((SO \cup WR \cup WW) ; RW?)^+ \cup (CO_i \cup \{(T_i, S_i)\}))^+ \\ &= (CO_0 \cup CO_i \cup \{(T_i, S_i)\})^+. \end{aligned}$$

We furthermore have $CO_0 \subseteq CO_i$, because (VIS_0, CO_0) is the smallest solution of the system in Figure 5. Then

$$CO_{i+1} = (CO_i \cup \{(T_i, S_i)\})^+.$$

We now use this equality to show that CO_{i+1} must be acyclic. Reasoning by contradiction, suppose it has a cycle. Then there exists a sequence of transactions T^0, \dots, T^m , where $m \geq 0$, such that $T^0 = T^m$, and

$$T^0 \xrightarrow{\text{CO}_i \cup \{(T_i, S_i)\}} \dots \xrightarrow{\text{CO}_i \cup \{(T_i, S_i)\}} T^m.$$

Because \mathcal{P}_i is a pre-execution, CO_i is acyclic, so in the sequence above there exists at least one index $k \in \{0, \dots, m-1\}$ such that $T^k = T_i$ and $T^{k+1} = S_i$. Let j, h be the smallest and the largest such indices, respectively (note that we can have $j = h$). Then we can convert the above cycle into the following one:

$$T^0 \xrightarrow{\text{CO}_i} \dots \xrightarrow{\text{CO}_i} T^j = T_i \xrightarrow{\{(T_i, S_i)\}} S_i = T^{h+1} \xrightarrow{\text{CO}_i} \dots \xrightarrow{\text{CO}_i} T^m,$$

where there is exactly one arrow labelled with the relation $\{(T_i, S_i)\}$, and all other arrows are labelled with CO_i . Now it remains to note that we have

$$S_i = T^{h+1} \xrightarrow{\text{CO}_i} \dots \xrightarrow{\text{CO}_i} T^m = T^0 \xrightarrow{\text{CO}_i} \dots \xrightarrow{\text{CO}_i} T^j = T_i,$$

and because CO_i is transitive by (S3), we have $S_i \xrightarrow{\text{CO}_i} T_i$. But this contradicts the assumption that S_i and T_i are not related by CO_i . Therefore, CO_{i+1} must be acyclic. \square

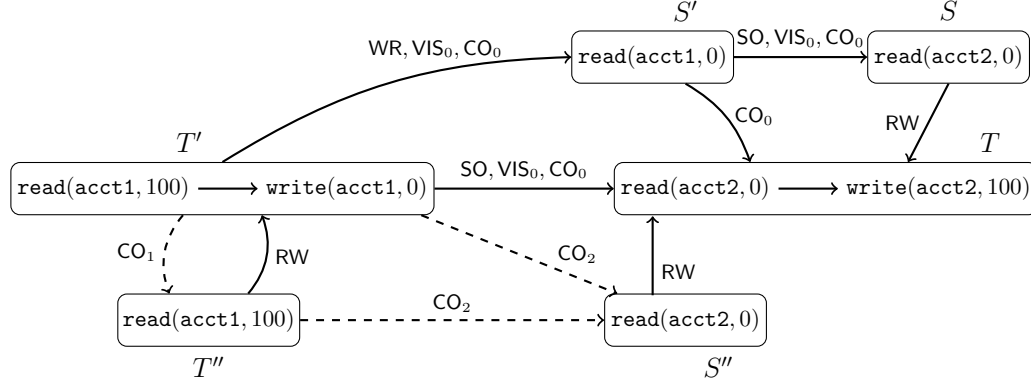
PROOF OF THEOREM 4.2(i). Consider the sequence $\{\mathcal{P}_i = (\mathcal{T}, \text{SO}, \text{VIS}_i, \text{CO}_i)\}_{i=0}^n$ constructed as described above. We prove by induction on i that $\mathcal{P}_i \in \text{PreExecSI}$ and $\text{graph}(\mathcal{P}_i) = \mathcal{G}$ for $i = 0..n$. If $i = 0$, then $\mathcal{P}_0 \in \text{PreExecSI}$ and $\text{graph}(\mathcal{P}_0) = \mathcal{G}$ by Corollary 4.9. Assume $\mathcal{P}_i \in \text{PreExecSI}$ for $0 \leq i < n$. Then by Lemma 4.10, CO_{i+1} is acyclic. Hence, by Lemma 4.7, $\mathcal{P}_{i+1} \in \text{PreExecSI}$ and $\text{graph}(\mathcal{P}_{i+1}) = \mathcal{G}$, which completes the induction. Finally, by construction CO_n is total, so that $\mathcal{X} = \mathcal{P}_n$ is an abstract execution such that $\text{graph}(\mathcal{P}_n) = \mathcal{G}$, as required. \square

We now illustrate the construction in the proof of Theorem 4.2(i) by an example. Consider the dependency graph $\mathcal{G}_5 \in \text{GraphSI}$ in Figure 6, which we also use as an example 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.

By Lemma 4.8, the smallest solution $(\text{VIS}_0, \text{CO}_0)$ of the system of inequalities in Figure 5 consists of the solid edges in the Figure 6. In particular, we have $S' \xrightarrow{\text{CO}_0} T$ because of $S' \xrightarrow{\text{SO}} S$ (so that $S' \xrightarrow{\text{VIS}} S$), $S \xrightarrow{\text{RW}} T$ and the inequality (S5). Since CO_0 is not total, we have to 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 5 such that $\text{CO}_1 \supseteq \text{CO}_0 \cup \{(T_1, S_1)\}$. By Lemma 4.8 this solution is given by (1) for $R = \text{CO}_0 \cup \{(T_1, S_1)\}$. For example, in Figure 6 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 6, then we get $\text{VIS}_1 = \text{VIS}_0$, and $\text{CO}_1 = \text{CO}_0 \cup \{(T', T'')\}$. The relation CO_1 is not yet total: for example, it 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 $R = \text{CO}_1 \cup \{(T'', S'')\}$ in (1); this corresponds to all the solid and dashed edges in Figure 6. Note that CO_2 also includes the edge (T', S'') . Since CO_2 is total in this example, the construction terminates with $n = 2$.

4.4. Proof of Theorem 4.2(ii)

Consider $\mathcal{X} = (\mathcal{T}, \text{SO}, \text{VIS}, \text{CO}) \in \text{ExecSI}$. As follows from Lemma 4.4, VIS and CO give a solution to the system of inequalities of Figure 5 for $\text{WR} = \text{WR}_{\mathcal{X}}$, $\text{WW} = \text{WW}_{\mathcal{X}}$, $\text{RW} = \text{RW}_{\mathcal{X}}$.



```

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 }

```

Fig. 6. An illustration of constructing an execution from a dependency graph (§4) and splicing an execution (§5). We omit an initialisation transaction that sets $\text{acct1} = 100$ and $\text{acct2} = 0$.

We now apply Lemma 4.8 for $R = \emptyset$; the minimality of the solution given by Lemma 4.8 implies that $((\text{SO} \cup \text{WR} \cup \text{WW}) ; \text{RW}^?)^+ \subseteq \text{CO}$. Then $((\text{SO} \cup \text{WR} \cup \text{WW}) ; \text{RW}^?)^+$ is acyclic because so is CO . This establishes $\text{graph}(\mathcal{X}) \in \text{GraphSI}$. \square

4.5. More Precise Characterisations

In this section we present corollaries of Theorem 4.1 that give more precise SI characterisations. We use these in our static analyses (§5 and §6).

For the following, it is helpful to introduce some notation. We use γ, γ', \dots to range over cycles in a dependency graph \mathcal{G} , i.e., paths of the form

$$T_1 \xrightarrow{\mathcal{R}_1} T_2 \xrightarrow{\mathcal{R}_2} \dots \xrightarrow{\mathcal{R}_{n-1}} T_n \quad (3)$$


such that $T_1 = T_n$ and for any $i = 1..n$ we have $T_i \in \mathcal{T}_{\mathcal{G}}$ and $\mathcal{R}_i \in \{\text{SO}_{\mathcal{G}}, \text{WR}_{\mathcal{G}}, \text{WW}_{\mathcal{G}}, \text{RW}_{\mathcal{G}}\}$. For a given cycle γ of the form (3), we let

$$\text{rep}(\gamma) = |\{T_i \mid \exists j. 1 \leq i, j < n \wedge i \neq j \wedge T_i = T_j\}|$$

be the number of repeated vertices in it. Note that γ has no repeated vertices if and only if $\text{rep}(\gamma) = 0$. We call such a cycle *simple*. Our first strengthening of the SI characterisation limits the scope of the acyclicity check in Theorem 4.1 to simple cycles.

THEOREM 4.11. *Let \mathcal{G} be a dependency graph. Then $\mathcal{G} \in \text{GraphSI}$ if and only if $\mathcal{G} \models \text{INT}$ and all simple cycles in \mathcal{G} contain at least two adjacent anti-dependency edges.*

PROOF. The “only if” direction trivially follows from Theorem 4.1. For the “if” direction, let \mathcal{G} be a dependency graph. We show that if \mathcal{G} contains a cycle without adjacent $\text{RW}_{\mathcal{G}}$ edges, then it also contains a *simple* cycle without adjacent $\text{RW}_{\mathcal{G}}$ edges. To this end, let γ be a cycle in \mathcal{G} without adjacent $\text{RW}_{\mathcal{G}}$ edges. If $\text{rep}(\gamma) = 0$, then there is nothing to prove, so let us assume that $\text{rep}(\gamma) > 0$. Below we show how to extract a sub-cycle γ' of γ such that

Fig. 7. A cycle with a repeated transaction T

$\text{rep}(\gamma') < \text{rep}(\gamma)$ and γ' has no adjacent $\text{RW}_{\mathcal{G}}$ edges. By applying this procedure repeatedly, we obtain a cycle γ'' with no adjacent $\text{RW}_{\mathcal{G}}$ edges and such that $\text{rep}(\gamma'') = 0$, as required.

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

$$\gamma = S \rightarrow \dots \xrightarrow{\mathcal{R}_1} T \xrightarrow{\mathcal{R}_2} \dots \xrightarrow{\mathcal{R}_3} T \xrightarrow{\mathcal{R}_4} \dots \rightarrow S$$

for some $T, S \in \mathcal{T}_{\mathcal{G}}$ and $\{\mathcal{R}_i\}_{i=1}^4 \subseteq \{\text{SO}_{\mathcal{G}}, \text{WR}_{\mathcal{G}}, \text{WW}_{\mathcal{G}}, \text{RW}_{\mathcal{G}}\}$. A graphical representation of γ is given in Figure 7. From γ we can derive the cycles

$$\gamma_1 = S \rightarrow \dots \xrightarrow{\mathcal{R}_1} T \xrightarrow{\mathcal{R}_4} \dots \rightarrow S$$

and

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

which are contained inside the dashed boxes in Figure 7.

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

Next, we show another characterisation of **GraphSI** that involves the notion of *vulnerable* anti-dependencies [Fekete et al., 2005]. In our setting, we define these as the anti-dependency edges that are not covered by the other types of edges. Formally, in a dependency graph \mathcal{G} an anti-dependency edge $T \xrightarrow{\text{RW}_{\mathcal{G}}} S$ is *vulnerable* if we do not have $T \xrightarrow{\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}}} S$.

THEOREM 4.12. *Let \mathcal{G} be a dependency graph. Then $\mathcal{G} \in \text{GraphSI}$ if and only if $\mathcal{G} \models \text{INT}$ and all simple cycles in \mathcal{G} contain at least two adjacent vulnerable anti-dependency edges.*

PROOF. First note that

$$\begin{aligned} ((\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}}) ; \text{RW}_{\mathcal{G}})^+ = \\ ((\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}}) ; (\text{RW}_{\mathcal{G}} \setminus (\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}})))^+. \end{aligned}$$

Then by Theorem 4.1 and the above equality, $\mathcal{G} \in \text{GraphSI}$ if and only if $\mathcal{G} \models \text{INT}$ and all cycles in \mathcal{G} contain at least two adjacent vulnerable anti-dependency edges. We can now proceed as in the proof of Theorem 4.13 to obtain the required. \square

So far, none of our SI characterisations has considered the objects involved in dependencies between transactions. The following theorem takes these into account.

THEOREM 4.13. *Let \mathcal{G} be a dependency graph. Then $\mathcal{G} \in \text{GraphSI}$ if and only if $\mathcal{G} \models \text{INT}$ and all simple cycles in \mathcal{G} contain at least two adjacent vulnerable anti-dependency edges over different objects.*

PROOF. The “if” part of the theorem is trivial. For the “only if” part, suppose that $\mathcal{G} \in \text{GraphSI}$. We prove that

$$\forall x. \text{RW}_{\mathcal{G}}(x) ; \text{RW}_{\mathcal{G}}(x) \subseteq \text{RW}_{\mathcal{G}}(x) ; (\text{RW}_{\mathcal{G}}(x) \cap \text{WW}_{\mathcal{G}}(x)), \quad (4)$$

This implies that, whenever we have two adjacent anti-dependencies over the same object in \mathcal{G} , then at least one of them is not vulnerable. Since $\mathcal{G} \in \text{GraphSI}$, by Theorem 4.12 the graph has only simple cycles with two adjacent vulnerable anti-dependencies, and by (4) such anti-dependencies need to be over different objects.

To prove (4), consider $x \in \text{Obj}$ and $T, T', T'' \in \mathcal{T}_{\mathcal{G}}$ such that $T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T''$. By definition, $T' \vdash \text{write}(x, _)$, $T'' \vdash \text{write}(x, _)$ and $T' \neq T''$. Therefore, either $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T''$ or $T'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$. However, we cannot have $T'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$, because $T' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T''$ and the hypothesis $\mathcal{G} \in \text{GraphSI}$ ensures that $\text{WW}_{\mathcal{G}}(x) ; \text{RW}_{\mathcal{G}}(x)$ is acyclic. Hence, we must have $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T''$ and, therefore, $T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T' \xrightarrow{\text{RW}_{\mathcal{G}}(x) \cap \text{WW}_{\mathcal{G}}(x)} T''$. \square

As an illustration of Theorem 4.13, the dependency graph of the write skew anomaly in Figure 2(f), allowed by SI, contains only cycles with at least two adjacent vulnerable anti-dependencies over different objects: e.g., $T_1 \xrightarrow{\text{RW}(\text{acct2})} T_2 \xrightarrow{\text{RW}(\text{acct1})} T_1$.

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* [Shasha et al., 1995] into sessions of smaller transactions without introducing new behaviours (the sessions are also called *chains* in this context [Zhang et al., 2013]). 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

$$\approx_{\mathcal{H}} = \text{SO}_{\mathcal{H}} \cup \text{SO}_{\mathcal{H}}^{-1} \cup \{(T, T) \mid T \in \mathcal{T}_{\mathcal{H}}\}$$

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: $\boxed{T}_{\mathcal{H}} = (E, \text{po})$, where $E = (\bigcup \{E_S \mid S \approx_{\mathcal{H}} T\})$ and

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

We let $\text{splice}(\mathcal{H})$ be the history resulting from splicing all sessions in a history \mathcal{H} :

$$\text{splice}(\mathcal{H}) = \left(\left\{ \boxed{T}_{\mathcal{H}} \mid T \in \mathcal{T}_{\mathcal{H}} \right\}, \emptyset \right).$$

A dependency graph $\mathcal{G} \in \text{GraphSI}$ is *spliceable* if there exists a dependency graph $\mathcal{G}' \in \text{GraphSI}$ such that $\mathcal{H}_{\mathcal{G}'} = \text{splice}(\mathcal{H}_{\mathcal{G}})$. 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}_5 in Figure 6 is not spliceable, because $\text{splice}(\mathcal{H}_{\mathcal{G}_5}) \notin \text{HistSI}$: informally, $\boxed{S}_{\mathcal{G}_5}$ observes the write by $\boxed{T}_{\mathcal{G}_5}$ to `acct1`, but not its write to `acct2`. On the other hand, let \mathcal{G}_6 be the graph obtained by removing the transactions S and S' from \mathcal{G}_5 . Then \mathcal{G}_6 is spliceable, as witnessed by the graph $\mathcal{G}'_6 \in \text{GraphSI}$ with $\mathcal{H}_{\mathcal{G}'_6} = \text{splice}(\mathcal{H}_{\mathcal{G}_6})$ and only the edges $\boxed{T''}_{\mathcal{G}_6} \xrightarrow{\text{RW}_{\mathcal{G}'_6}} \boxed{T}_{\mathcal{G}_6}$ and $\boxed{S''}_{\mathcal{G}_6} \xrightarrow{\text{RW}_{\mathcal{G}'_6}} \boxed{T}_{\mathcal{G}_6}$.

Our criterion for checking that a dependency graph is spliceable requires the absence of certain cycles in a variant of the graph.

Definition 5.1. Given a dependency graph \mathcal{G} , the **dynamic chopping graph** corresponding to \mathcal{G} is the graph $\text{DCG}(\mathcal{G})$ with the node set $\mathcal{T}_{\mathcal{G}}$ and the edge set obtained from the union of the following sets:

- $\text{SO}_{\mathcal{G}}$ (a *successor* edge);
- $\text{SO}_{\mathcal{G}}^{-1}$ (a *predecessor* edge);
- $\text{WR}_{\mathcal{G}} \setminus \approx_{\mathcal{G}}$ (a *read dependency* edge);
- $\text{WW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}}$ (a *write dependency* edge); and
- $\text{RW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}}$ (an *anti-dependency* edge).

We refer to the last three kinds of edges as **conflict** edges.

A cycle in a chopping graph $\text{DCG}(\mathcal{G})$ is **critical** if:

- (i) it is simple;
- (ii) it contains a fragment of three consecutive edges of the form “conflict, predecessor, conflict”; and
- (iii) any two anti-dependency edges on the cycle are separated by at least one read or write dependency edge.

THEOREM 5.2. *For $\mathcal{G} \in \text{GraphSI}$, if $\text{DCG}(\mathcal{G})$ contains no critical cycles, then \mathcal{G} is spliceable.*

For example, the above graph \mathcal{G}_6 contains no critical cycles, and the graph \mathcal{G}_5 contains a critical cycle

$$T' \xrightarrow{\text{WR}_{\mathcal{G}_5}} S' \xrightarrow{\text{SO}_{\mathcal{G}_5}} S \xrightarrow{\text{RW}_{\mathcal{G}_5}} T \xrightarrow{\text{SO}_{\mathcal{G}_5}^{-1}} T'.$$

5.1. Proof of the Dynamic Chopping Criterion under SI

To prove Theorem 5.2, we exhibit a particular dependency graph $\text{splice}(\mathcal{G})$ such that $\text{splice}(\mathcal{G}) \in \text{GraphSI}$ and $\mathcal{H}_{\text{splice}(\mathcal{G})} = \text{splice}(\mathcal{H}_{\mathcal{G}})$. We define read dependencies $\text{WR}_{\text{splice}(\mathcal{G})}$ by lifting those in $\text{WR}_{\mathcal{G}}$ to spliced transactions:

$$\forall T, S \in \mathcal{T}_{\mathcal{G}}, \forall x \in \text{Obj}. \boxed{T}_{\mathcal{G}} \xrightarrow{\text{WR}_{\text{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}} \iff \boxed{T}_{\mathcal{G}} \neq \boxed{S}_{\mathcal{G}} \wedge T \xrightarrow{\approx_{\mathcal{G}}; \text{WR}_{\mathcal{G}}(x); \approx_{\mathcal{G}}} S. \quad (5)$$

We define $\text{WW}_{\text{splice}(\mathcal{G})}$ similarly and derive $\text{RW}_{\text{splice}(\mathcal{G})}$ from $\text{WR}_{\text{splice}(\mathcal{G})}$ and $\text{WW}_{\text{splice}(\mathcal{G})}$ as in Definition 3.1. Then Theorem 5.2 is a corollary of the following theorem.

THEOREM 5.3. *For $\mathcal{G} \in \text{GraphSI}$, if $\text{DCG}(\mathcal{G})$ contains no critical cycles, then $\text{splice}(\mathcal{G})$ is a dependency graph and $\text{splice}(\mathcal{G}) \in \text{GraphSI}$.*

To prove the above theorem, we first show that $\text{RW}_{\text{splice}(\mathcal{G})}$ can be decomposed into a form similar to (5).

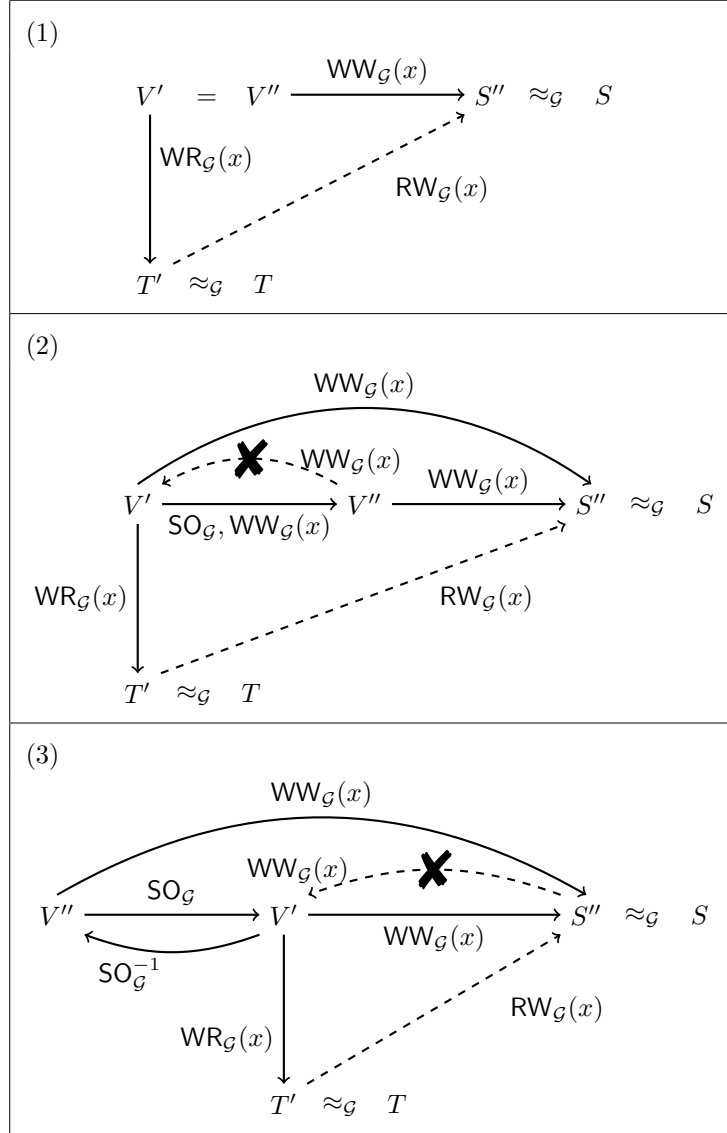


Fig. 8. Graphical representation of the different cases in the proof of Lemma 5.4.

LEMMA 5.4. *Let $\mathcal{G} \in \text{GraphSI}$ be such that $\text{DCG}(\mathcal{G})$ contains no critical cycles. Then*

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

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

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

- (1) $V' = V''$. Then $V' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} T'$ and $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$, from which $T' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S''$ follows.
- (2) $V' \xrightarrow{\text{SO}_{\mathcal{G}}} V''$. We have $V' \vdash \mathbf{write}(x, _)$ and $V'' \vdash \mathbf{write}(x, _)$. The assumption $\mathcal{G} \in \text{GraphSI}$ mandates that $\neg(V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V')$: otherwise we would have the cycle $V' \xrightarrow{\text{SO}_{\mathcal{G}}} V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$, contradicting $\mathcal{G} \in \text{GraphSI}$. Hence, $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V''$. We thus have $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$ and by transitivity of $\text{WW}_{\mathcal{G}}(x)$ we obtain $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$. We have established $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$ and $V' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} T'$, so that $T' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S''$.
- (3) $V'' \xrightarrow{\text{SO}_{\mathcal{G}}} V'$. 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 $\boxed{V'}_{\mathcal{G}} \xrightarrow{\text{WW}_{\text{splice}(\mathcal{G})}(x)} \boxed{S''}_{\mathcal{G}}$, 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 $V' \vdash \mathbf{write}(x, _)$ and $S'' \vdash \mathbf{write}(x, _)$; since $V' \neq S''$, we must have either $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$ or $S'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$. However, in the latter case we would have the critical cycle $S'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V' \xrightarrow{\text{SO}_{\mathcal{G}}^{-1}} V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$ in $\text{DCG}(\mathcal{G})$, which contradicts the hypothesis of the lemma. Therefore, we must have $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$. Since we also have $V' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} T'$, it follows that $T' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S''$.

□

We next show that $\text{splice}(\mathcal{G})$ is a dependency graph (Lemma 5.6 below) and that $\mathcal{T}_{\text{splice}(\mathcal{G})} \models \text{INT}$ (Lemma 5.7 below). The proof of the former fact relies on the following easy proposition.

PROPOSITION 5.5. *Let \mathcal{G} be a dependency graph. For any $T \in \mathcal{T}_{\mathcal{G}}$:*

- (1) $\boxed{T}_{\mathcal{G}} \vdash \mathbf{read}(x, n)$ if and only if $\text{min}_{\text{SO}_{\mathcal{G}}}\{S \mid S \approx_{\mathcal{G}} T \wedge S \vdash _.(x, _)\} \vdash \mathbf{read}(x, n)$;
- (2) $\boxed{T}_{\mathcal{G}} \vdash \mathbf{write}(x, n)$ if and only if $\text{max}_{\text{SO}_{\mathcal{G}}}\{S \mid S \approx_{\mathcal{G}} T \wedge S \vdash \mathbf{write}(x, _)\} \vdash \mathbf{write}(x, n)$.

LEMMA 5.6. *Let $\mathcal{G} \in \text{GraphSI}$ be such that $\text{DCG}(\mathcal{G})$ contains no critical cycles. Then $\text{splice}(\mathcal{G})$ is a dependency graph.*

PROOF. We prove that, if the chopping graph of \mathcal{G} contains no critical cycles, then $\text{splice}(\mathcal{G})$ satisfies all the constraints of Definition 3.2.

- (1a). Consider $x \in \text{Obj}$ and $\boxed{T}_{\mathcal{G}}, \boxed{S}_{\mathcal{G}} \in \mathcal{T}_{\text{splice}(\mathcal{G})}$ such that $\boxed{T}_{\mathcal{G}} \xrightarrow{\text{WR}_{\text{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$. 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 $T \approx_{\mathcal{G}} T' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S' \approx_{\mathcal{G}} S$. Hence, for some n we have $T' \vdash \mathbf{write}(x, n)$ and $S' \vdash \mathbf{read}(x, n)$. We now prove that (i) $\boxed{T}_{\mathcal{G}} \vdash \mathbf{write}(x, n)$ and (ii) $\boxed{S}_{\mathcal{G}} \vdash \mathbf{read}(x, n)$.

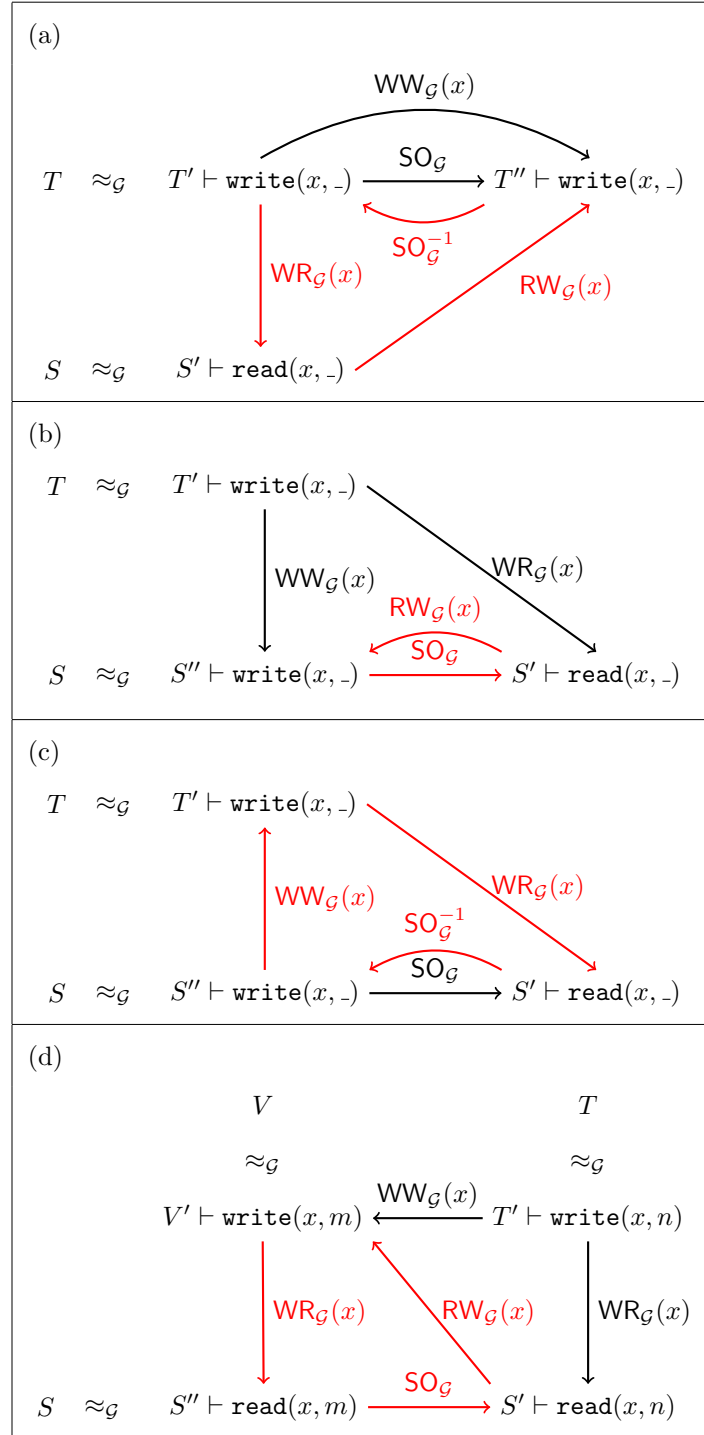


Fig. 9. Graphical representation of the different cases in the proof of Lemma 5.6.

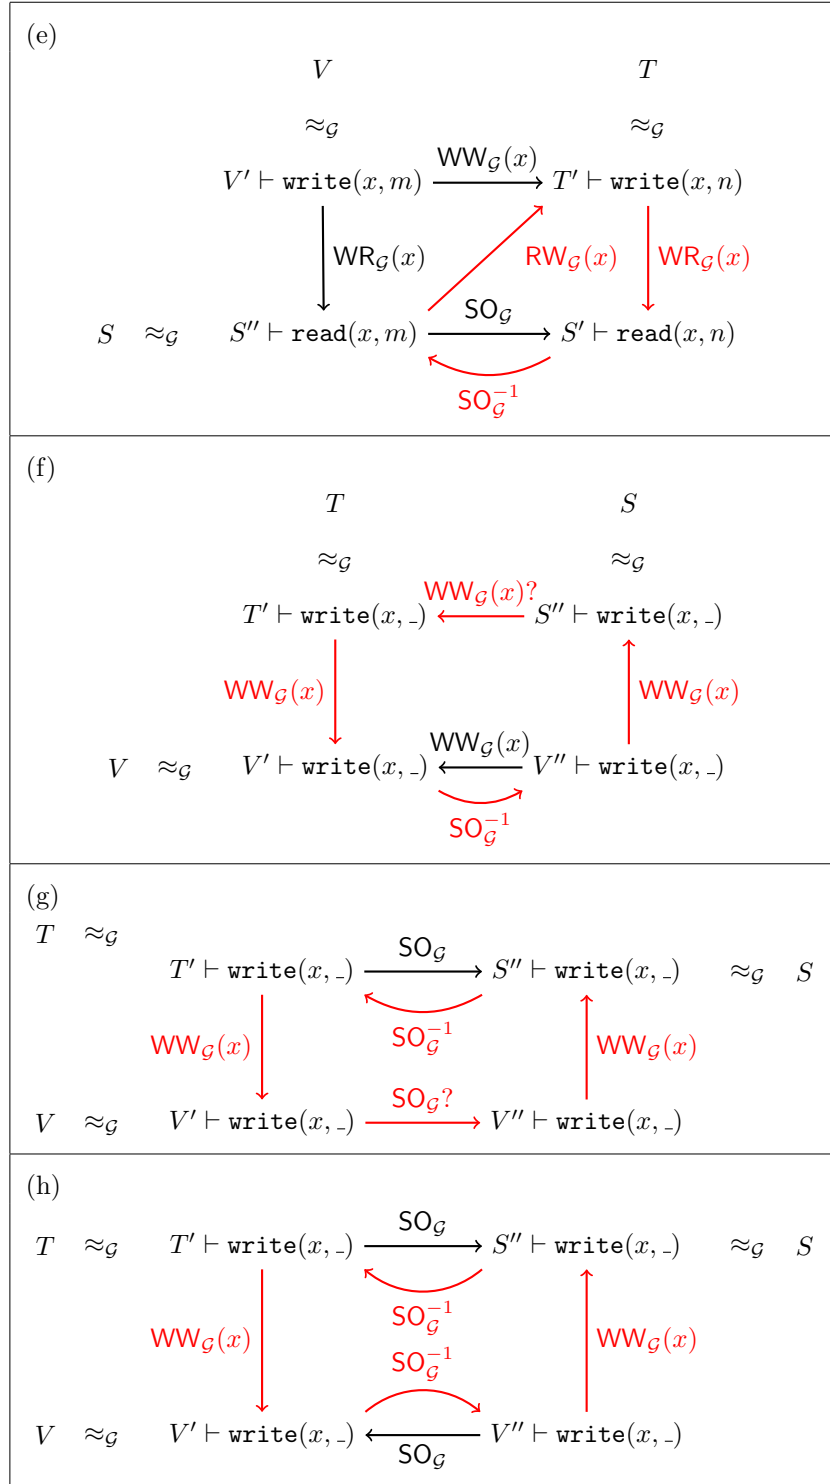


Fig. 10. Graphical representation of the different cases in the proof of Lemma 5.6.

(i). Consider an arbitrary transaction $T'' \in \mathcal{T}$ such that $T'' \vdash \mathbf{write}(x, _)$ and $T'' \approx_{\mathcal{G}} T'$ (Figure 9(a)). We show that it cannot be the case that $T' \xrightarrow{\text{SO}_{\mathcal{G}}} T''$. Then $T' = \max_{\text{SO}_{\mathcal{G}}} \{S \mid S \approx_{\mathcal{G}} T \wedge S \vdash \mathbf{write}(x, _)\}$, and by Proposition 5.5(ii) it follows that $\boxed{T}_{\mathcal{G}} \vdash \mathbf{write}(x, n)$, as required.

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

(ii). We show that, for any transaction $S'' \approx_{\mathcal{G}} S'$ such that $S'' \xrightarrow{\text{SO}_{\mathcal{G}}} S'$ we have $\neg(S'' \vdash \mathbf{write}(x, _))$, and if $S'' \vdash \mathbf{read}(x, m)$, then $m = n$. As a consequence, $\min_{\text{SO}_{\mathcal{G}}} (\{V \mid V \approx_{\mathcal{G}} S \wedge V \vdash _ (x, _)\} \vdash \mathbf{read}(x, n))$, and hence, by Proposition 5.5(i) we have $\boxed{S}_{\mathcal{G}} \vdash \mathbf{read}(x, n)$.

Let S'' be a transaction such that $S'' \xrightarrow{\text{SO}_{\mathcal{G}}} S'$; we prove that $\neg(S'' \vdash \mathbf{write}(x, _))$ by contradiction. Assume $S'' \vdash \mathbf{write}(x, _)$. Then by the definition of $\text{WW}_{\mathcal{G}}(x)$, either $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$ or $S'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$; 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 $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$ (Figure 9(b)), since together with $T' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S'$, this would imply the anti-dependency edge $S' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S''$. But then we have a cycle $S' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S'' \xrightarrow{\text{SO}_{\mathcal{G}}} S'$, contradicting $\mathcal{G} \in \text{GraphSI}$. We cannot have $S'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$ either (Figure 9(c)): in this case the chopping graph of \mathcal{G} contains the critical cycle $S'' \xrightarrow{(\text{WW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} T' \xrightarrow{(\text{WR}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} S' \xrightarrow{\text{SO}_{\mathcal{G}}^{-1}} S''$. We have thus established $\neg(S'' \vdash \mathbf{write}(x, _))$.

Suppose now that $S'' \xrightarrow{\text{SO}_{\mathcal{G}}} S'$ and $S'' \vdash \mathbf{read}(x, m)$ for some m . Then there exists a transaction $V' \in \mathcal{T}_{\mathcal{G}}$ such that $V' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S''$ and $V' \vdash \mathbf{write}(x, m)$. Since $T' \vdash \mathbf{write}(x, n)$, we have $V' = T'$, $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$ or $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$. We show that the latter two cases are impossible, so that $T' = V'$ and, hence, $m = n$, as required. If $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$, then we have the anti-dependency edge $S' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} V'$ (Figure 9(d)).

Then the graph \mathcal{G} contains the cycle $S' \xrightarrow{\text{RW}_{\mathcal{G}}} V' \xrightarrow{\text{WR}_{\mathcal{G}}} S'' \xrightarrow{\text{SO}_{\mathcal{G}}} S'$, which contradicts $\mathcal{G} \in \text{GraphSI}$. If $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$, then we have the anti-dependency edge $S'' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T'$ (Figure 10(e)). In this case $\text{DCG}(\mathcal{G})$ contains the critical cycle $S'' \xrightarrow{(\text{RW}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} T' \xrightarrow{(\text{WR}_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} S' \xrightarrow{\text{SO}_{\mathcal{G}}^{-1}} S''$, and again we obtain a contradiction.

(1b). Let $x \in \text{Obj}$ and $\boxed{S}_{\mathcal{G}} \in \mathcal{T}_{\text{splice}(\mathcal{G})}$ be such that $\boxed{S}_{\mathcal{G}} \vdash \mathbf{read}(x, _)$. By Proposition 5.5(i), there exists a transaction $S' \approx_{\mathcal{G}} S$ such that $S' \vdash \mathbf{read}(x, _)$. Since \mathcal{G} is a dependency graph, there exists a transaction $T \in \mathcal{T}_{\mathcal{G}}$ such that $T \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S'$. Then $T \neq S'$. We cannot have $T \xrightarrow{\text{SO}_{\mathcal{G}}} S'$, because $T \vdash \mathbf{write}(x, _)$ and $S' = \min_{\text{SO}} \{S \mid S \approx_{\mathcal{G}}$

$S' \wedge S \vdash _ (x, _)$ by Proposition 5.5(i). Finally, we cannot have $S' \xrightarrow{\text{SO}_{\mathcal{G}}} T$, because this would contradict the hypothesis $\mathcal{G} \in \text{GraphSI}$ due to the cycle $S' \xrightarrow{\text{SO}_{\mathcal{G}}} T \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S'$. As a consequence, it cannot be $T \approx_{\mathcal{G}} S'$. Then $\boxed{T}_{\mathcal{G}} \xrightarrow{\text{WR}_{\text{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$.

(1c). Let $x \in \text{Obj}$ and $\boxed{S}_{\mathcal{G}}, \boxed{T}_{\mathcal{G}}, \boxed{V}_{\mathcal{G}} \in \mathcal{T}_{\text{splice}(\mathcal{G})}$ be such that $\boxed{T}_{\mathcal{G}} \xrightarrow{\text{WR}_{\text{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$ and $\boxed{V}_{\mathcal{G}} \xrightarrow{\text{WR}_{\text{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$. Then $S \not\approx_{\mathcal{G}} T$, $S \not\approx_{\mathcal{G}} V$ and there exist transactions S', S'', T', V' such that $S' \approx_{\mathcal{G}} S \approx_{\mathcal{G}} S''$, $T' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S'$, $V' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S''$. 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 $S' \xrightarrow{\text{SO}_{\mathcal{G}}} S''$ or $S'' \xrightarrow{\text{SO}_{\mathcal{G}}} S'$. Without loss of generality, assume that $S'' \xrightarrow{\text{SO}_{\mathcal{G}}} S'$ (Figure 9(d)). Since $T' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S'$ and $V' \xrightarrow{\text{WR}_{\mathcal{G}}(x)} S''$, we get $T' \vdash \text{write}(x, _)$ and $V' \vdash \text{write}(x, _)$. Therefore, we must have one of the following: $T' = V'$, $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$ or $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$. However, the latter two cases are impossible. If we had $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$, then \mathcal{G} would contain the configuration shown in Figure 9(d), which contradicts $\mathcal{G} \in \text{GraphSI}$. If we had $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$, then \mathcal{G} would contain the configuration shown in Figure 10(e) and, hence, $\text{DCG}(\mathcal{G})$ would contain a critical cycle. Hence, we must have $T' = V'$ and $\boxed{T}_{\mathcal{G}} = \boxed{V}_{\mathcal{G}}$.

(2). We prove that $\text{WW}_{\text{splice}(\mathcal{G})}$ is transitive, irreflexive and total over WriteTx_x .

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

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

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

- Since $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$ and $V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$, we have $V' \vdash \text{write}(x, _)$ and $V'' \vdash \text{write}(x, _)$. Then one of the following holds: $V' = V''$, $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V''$ or $V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$. In the first two cases, the transitivity of $\text{WW}_{\mathcal{G}}(x)$ guarantees $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$, as required. Suppose now that $V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$. Then $V'' \xrightarrow{\text{SO}_{\mathcal{G}}} V'$, because $V'' \approx_{\mathcal{G}} V'$ and $\mathcal{G} \in \text{GraphSI}$. Since $T' \vdash \text{write}(x, _)$ and $S'' \vdash \text{write}(x, _)$, we have one of the following: $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$, $S'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$ 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 $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x) \approx_{\mathcal{G}}} V' \xrightarrow{\text{SO}_{\mathcal{G}}^{-1}} V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x) \approx_{\mathcal{G}}} S'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$. Therefore, it has to be $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S''$.
- 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 $V' \vdash \text{write}(x, _)$ and $V'' \vdash \text{write}(x, _)$, so that one of the following holds: $V' = V''$, $V' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V''$ or $V'' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V'$.

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

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

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

$\text{WW}_{\text{splice}(\mathcal{G})}(x)$ is total over WriteTx_x . Let $x \in \text{Obj}$ and $\boxed{T}_{\mathcal{G}}, \boxed{S}_{\mathcal{G}} \in \mathcal{T}_{\text{splice}(\mathcal{G})}$ be such that $\boxed{T}_{\mathcal{G}} \vdash \text{write}(x, _)$, $\boxed{S}_{\mathcal{G}} \vdash \text{write}(x, _)$ and $\boxed{T}_{\mathcal{G}} \neq \boxed{S}_{\mathcal{G}}$. By Proposition 5.5(ii), there exist $T', S' \in \mathcal{T}_{\mathcal{G}}$ such that $T' \approx_{\mathcal{G}} T$, $T' \vdash \text{write}(x, _)$, $S' \approx_{\mathcal{G}} S$ and $S' \vdash \text{write}(x, _)$. Also, $T' \approx_{\mathcal{G}} T \not\approx_{\mathcal{G}} S \approx_{\mathcal{G}} S'$, so that it cannot be $T' = S'$.

Since $\text{WW}_{\mathcal{G}}(x)$ is total over WriteTx_x , we must have either $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S'$ or $S' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$; without loss of generality, we assume $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S'$. We thus have $\boxed{T}_{\mathcal{G}} \neq \boxed{S}_{\mathcal{G}}$ and $T \approx_{\mathcal{G}} T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S' \approx_{\mathcal{G}} S$. Hence, $\boxed{T}_{\mathcal{G}} \xrightarrow{\text{WW}_{\text{splice}(\mathcal{G})}(x)} \boxed{S}_{\mathcal{G}}$.

(3). This follows immediately from the definition of $\text{RW}_{\text{splice}(\mathcal{G})}$.

□

LEMMA 5.7. Let $\mathcal{G} \in \text{GraphSI}$ be such that $\text{DCG}(\mathcal{G})$ contains no critical cycles. Then $\mathcal{T}_{\text{splice}(\mathcal{G})} \models \text{INT}$.

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 $\text{op}(e) = \text{read}(x, n)$ for some $x \in \text{Obj}, n \in \mathbb{N}$; and, letting $\boxed{\text{po}_T} = \text{po}_{\boxed{T}_{\mathcal{G}}}$,

$$f = \max_{\boxed{\text{po}_T}} \{e' \mid \text{op}(e') = _ (x, _) \wedge e' \xrightarrow{\boxed{\text{po}_T}} e\} \quad (6)$$

such that $\text{op}(f) = _ (x, m)$ for some $m \neq n$. Since $\mathcal{T}_{\mathcal{G}} \models \text{INT}$, we cannot have $f \in E_T$. Therefore, there exists a transaction $T' \xrightarrow{\text{SO}_{\mathcal{G}}} T$ such that $f \in E_{T'}$; see Figure 11(a). We now make a case split on whether $T' \vdash \text{write}(x, _)$.

(1) $T' \vdash \text{write}(x, _)$. Then there exists an event $g \in E_{T'}$ such that $\text{op}(g) = \text{write}(x, _)$. Without loss of generality, let g be the last write to object x in $E_{T'}$. If $f \xrightarrow{\text{po}_{T'}} g$, then $f \xrightarrow{\boxed{\text{po}_T}} g \xrightarrow{\boxed{\text{po}_T}} e$, contradicting (6). Thus, either $g = f$ or $g \xrightarrow{\text{po}_{T'}} f$. In both cases, we show that $\text{op}(g) = \text{write}(x, m)$. If $f = g$, we have $_ (x, m) = \text{op}(f) = \text{op}(g) =$

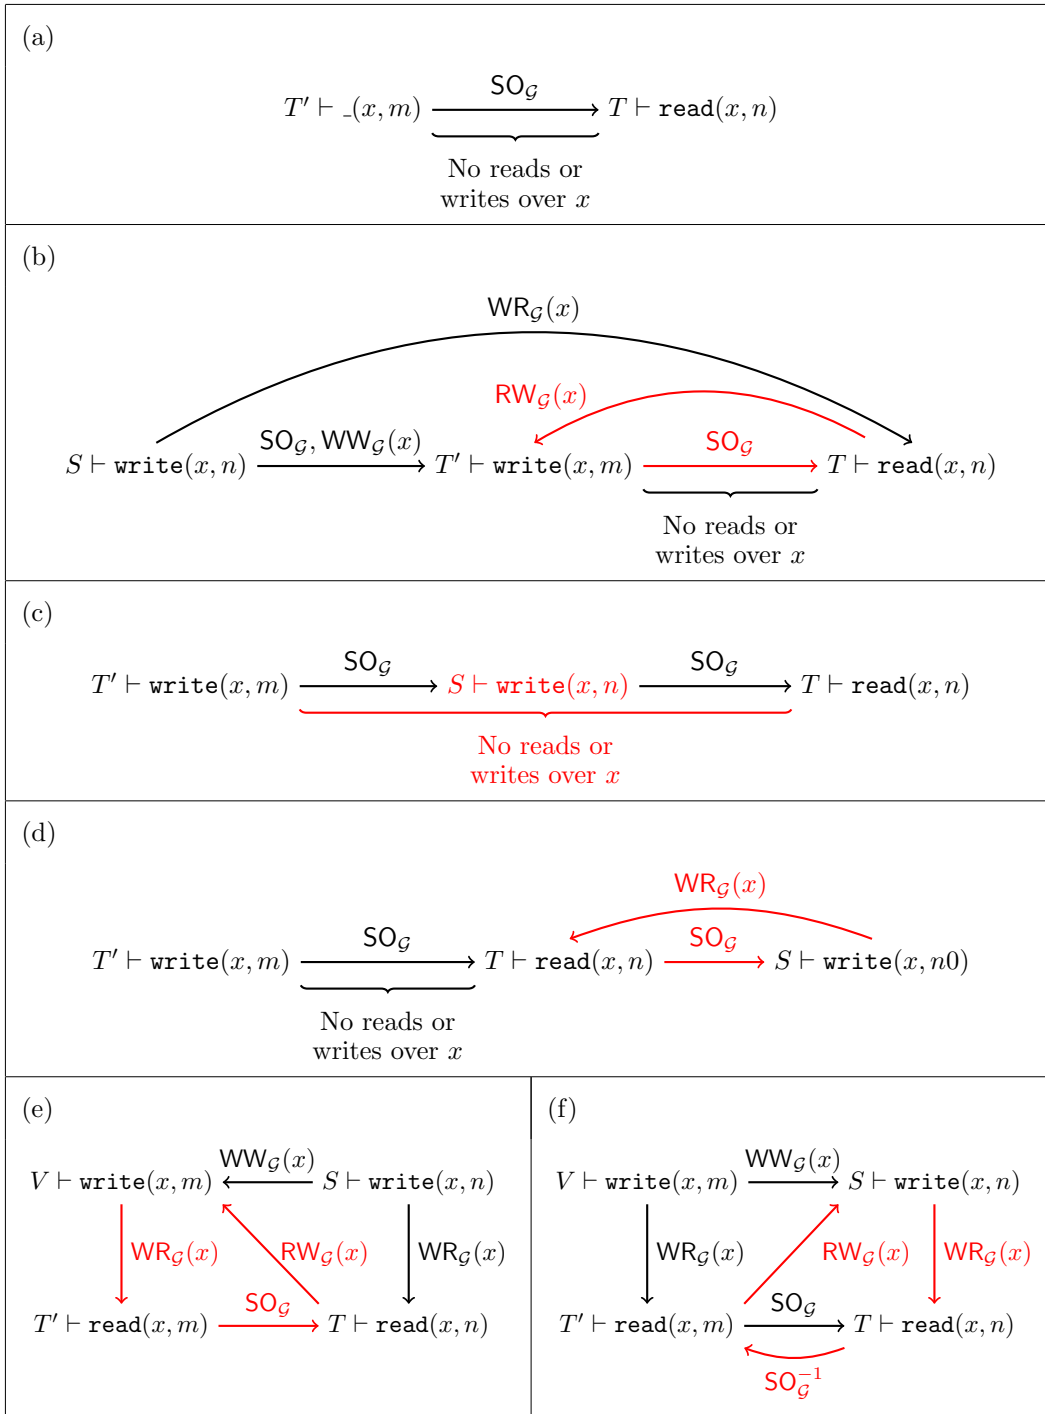


Fig. 11. Graphical representation of the different cases in the proof of Lemma 5.7.

$\text{write}(x, _)$, so that $\text{op}(g) = \text{write}(x, m)$. If $g \xrightarrow{\text{po}_{T'}} f$, then $\text{op}(f) = \text{read}(x, m)$, since g is the last write to x in T' . Also, for any other event h such that $\text{op}(h) = _.(x, _)$ and $g \xrightarrow{\text{po}_{T'}} h \xrightarrow{\text{po}_{T'}} f$, we have $\text{op}(h) = \text{read}(x, _)$. Then because $\mathcal{T}_{\mathcal{G}} \models \text{INT}$ and $\text{op}(f) = \text{read}(x, m)$, we must have $\text{op}(h) = \text{read}(x, m)$. But then $\mathcal{T}_{\mathcal{G}} \models \text{INT}$ again ensures $\text{op}(g) = \text{write}(x, m)$.

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

Next we prove that $T' \not\approx_{\mathcal{G}} S$. Assume the contrary. Since $S \neq T'$ and $T' \xrightarrow{\text{SO}_{\mathcal{G}}} T$, we must have one of the following:

$$S \xrightarrow{\text{SO}_{\mathcal{G}}} T' \xrightarrow{\text{SO}_{\mathcal{G}}} T; \quad T' \xrightarrow{\text{SO}_{\mathcal{G}}} S \xrightarrow{\text{SO}_{\mathcal{G}}} T; \quad T' \xrightarrow{\text{SO}_{\mathcal{G}}} T \xrightarrow{\text{SO}_{\mathcal{G}}} S.$$

We show that each of these cases leads to a contradiction.

— $S \xrightarrow{\text{SO}_{\mathcal{G}}} T' \xrightarrow{\text{SO}_{\mathcal{G}}} T$, Figure 11(b). Since $S, T' \vdash \text{write}(x, _)$ and $S \neq T'$, we must have either $S \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$ or $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S$. However, the last case is impossible because it would lead to a cycle $S \xrightarrow{\text{SO}_{\mathcal{G}}} T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S$, contradicting $\mathcal{G} \in \text{GraphSI}$; therefore, $S \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$. Together with $S \xrightarrow{\text{WR}_{\mathcal{G}}(x)} T$, this yields $T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T'$. But then we obtain a cycle $T' \xrightarrow{\text{SO}_{\mathcal{G}}} T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T'$, which again contradicts $\mathcal{G} \in \text{GraphSI}$.

— $T' \xrightarrow{\text{SO}_{\mathcal{G}}} S \xrightarrow{\text{SO}_{\mathcal{G}}} T$, Figure 11(c). Since $S \vdash \text{write}(x, _)$, there exists an event $h \in E_S$ such that $\text{op}(h) = \text{write}(x, _)$ and $f \xrightarrow{\boxed{\text{po}_T}} h \xrightarrow{\boxed{\text{po}_T}} e$. But this contradicts (6).

— $T' \xrightarrow{\text{SO}_{\mathcal{G}}} T \xrightarrow{\text{SO}_{\mathcal{G}}} S$, Figure 11(d). In this case we have the cycle $T \xrightarrow{\text{SO}_{\mathcal{G}}} S \xrightarrow{\text{WR}_{\mathcal{G}}(x)} T$, which contradicts $\mathcal{G} \in \text{GraphSI}$.

We have thus proved that $T' \not\approx_{\mathcal{G}} S$. Next, we observe that since $T' \vdash \text{write}(x, m)$, $S \vdash \text{write}(x, n)$ and $T' \neq S$, we must have either $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S$ or $S \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$.

We show that both of these cases lead to a contradiction. If $S \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$, then since $S \xrightarrow{\text{WR}_{\mathcal{G}}(x)} T$, we get $T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T'$, causing the cycle $T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T' \xrightarrow{\text{SO}_{\mathcal{G}}} T$; this contradicts $\mathcal{G} \in \text{GraphSI}$. On the other hand, if $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S$, then we have a critical cycle $T' \xrightarrow{(\text{WW}_{\mathcal{G}} \approx_{\mathcal{G}})} S \xrightarrow{(\text{WR}_{\mathcal{G}} \approx_{\mathcal{G}})} T \xrightarrow{\text{SO}_{\mathcal{G}}^{-1}} T'$ in $\text{DCG}(\mathcal{G})$, contradicting the hypothesis of the lemma.

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

Since \mathcal{G} is a dependency graph, there exist two transactions S, V such that $S \xrightarrow{\text{WR}_{\mathcal{G}}(x)} T$ and $V \xrightarrow{\text{WR}_{\mathcal{G}}(x)} T'$. Since $S \vdash \text{write}(x, n)$ and $V \vdash \text{write}(x, m)$, we have $S \neq V$. Since $\neg(T' \vdash \text{write}(x, _))$ and $S \vdash \text{write}(x, _)$, we have $T' \neq S$. We also have $V \neq T$, for otherwise we would have a cycle $T' \xrightarrow{\text{SO}_{\mathcal{G}}} T \xrightarrow{\text{WR}_{\mathcal{G}}} T'$, contradicting $\mathcal{G} \in \text{GraphSI}$. Hence, the transactions T', T, V, S are pairwise distinct.

We must have either $S \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V$ or $V \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S$. We show that neither of these cases is possible. If $S \xrightarrow{\text{WW}_{\mathcal{G}}(x)} V$ (Figure 11(e)), then $T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} V$. This causes a cycle $T \xrightarrow{\text{RW}_{\mathcal{G}}} V \xrightarrow{\text{WR}_{\mathcal{G}}} T' \xrightarrow{\text{SO}_{\mathcal{G}}} T$, contradicting $\mathcal{G} \in \text{GraphSI}$. If $V \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S$ (Figure 11(f)), then $T' \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S$. As in the case above, we can show that $S \not\approx_{\mathcal{G}} T$. This yields a critical

cycle $T' \xrightarrow{(RW_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} S \xrightarrow{(WR_{\mathcal{G}} \setminus \approx_{\mathcal{G}})} T \xrightarrow{SO_{\mathcal{G}}^{-1}} T'$ in $DCG(\mathcal{G})$, contradicting the assumptions of the lemma.

□

PROOF OF THEOREM 5.3. Let $\mathcal{G} \in \text{GraphSI}$ be a dependency graph such that $DCG(\mathcal{G})$ contains no critical cycles. We prove that $\text{splice}(\mathcal{G}) \in \text{GraphSI}$. First, Lemmas 5.6 and 5.7 ensure that $\text{splice}(\mathcal{G})$ is indeed a dependency graph and $\mathcal{T}_{\text{splice}(\mathcal{G})} \models \text{INT}$. Since $\text{SO}_{\text{splice}(\mathcal{G})} = \emptyset$, by Theorem 4.2 it remains to prove that the relation $((WR_{\text{splice}(\mathcal{G})} \cup WW_{\text{splice}(\mathcal{G})}) ; RW_{\text{splice}(\mathcal{G})})^?$ is acyclic. The proof goes by contradiction: we assume that this relation contains a cycle and exhibit a critical cycle in $DCG(\mathcal{G})$. Let

$$\gamma = \boxed{T_0}_{\mathcal{G}} \xrightarrow{C_0} \dots \xrightarrow{C_{n-1}} \boxed{T_n}_{\mathcal{G}} \quad (n \geq 1)$$

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

$$T_0, \dots, T_n \in \mathcal{T}_{\mathcal{G}}, \quad C_0, \dots, C_{n-1} \in \{WR_{\text{splice}(\mathcal{G})}, WW_{\text{splice}(\mathcal{G})}, RW_{\text{splice}(\mathcal{G})}\}$$

(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 $C_i = RW_{\text{splice}(\mathcal{G})}$ and $C_{(i+1) \bmod n} = RW_{\text{splice}(\mathcal{G})}$. By Theorem 4.11, 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 5.4, we can convert γ into the following path:

$$T'_0 \approx_{\mathcal{G}} T''_0 \xrightarrow{C_0^{\mathcal{G}}} T'_1 \approx_{\mathcal{G}} T''_1 \xrightarrow{C_1^{\mathcal{G}}} \dots \xrightarrow{C_{n-1}^{\mathcal{G}}} T'_n \approx_{\mathcal{G}} T''_n, \quad (7)$$

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 $i = 0..(n-1)$, $C_i^{\mathcal{G}}$ is the relation in \mathcal{G} corresponding to the relation C_i in $\text{splice}(\mathcal{G})$ (e.g., if $C_i = WR_{\text{splice}(\mathcal{G})}$, then $C_i^{\mathcal{G}} = (WR_{\mathcal{G}} \setminus \approx_{\mathcal{G}})$). We also know that

$$\neg \exists i = 0..(n-1). (C_i^{\mathcal{G}} = (RW_{\mathcal{G}} \setminus \approx_{\mathcal{G}})) \wedge (C_{(i+1) \bmod n}^{\mathcal{G}} = (RW_{\mathcal{G}} \setminus \approx_{\mathcal{G}})). \quad (8)$$

Since the cycle γ is simple, the only possibility for vertices to be repeated on the path (7) 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: $T = S$, $T \xrightarrow{SO_{\mathcal{G}}} S$ or $T \xrightarrow{SO_{\mathcal{G}}^{-1}} S$. Also, we know that $T'_n \approx_{\mathcal{G}} T_n \approx_{\mathcal{G}} T_0 \approx_{\mathcal{G}} T'_0$. Therefore, we can rewrite the path (7) as follows:

$$T'_n \xrightarrow{S_0} T''_0 \xrightarrow{C_1^{\mathcal{G}}} T'_1 \xrightarrow{S_1} T''_1 \xrightarrow{C_2^{\mathcal{G}}} \dots \xrightarrow{S_{n-1}} T''_{n-1} \xrightarrow{C_n^{\mathcal{G}}} T'_n, \quad (9)$$

where $S_0, \dots, S_{n-1} \in \{SO_{\mathcal{G}}^?, SO_{\mathcal{G}}^{-1}\}$. In this cycle repeated vertices are always adjacent and connected by an $SO_{\mathcal{G}}^?$ -edge. By removing such edges from the cycle we obtain a simple cycle, where all the occurrences of $SO_{\mathcal{G}}^?$ -edges are actually $SO_{\mathcal{G}}$ -edges; this is a cycle in $DCG(\mathcal{G})$. Due to (8), 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 $DCG(\mathcal{G})$, it remains to show that there exists an index $i = 0..(n-1)$ such that $S_i = SO_{\mathcal{G}}^{-1}$. This holds because, if we had $S_i = SO_{\mathcal{G}}$ for all $i = 0..(n-1)$, then we would obtain a cycle in $((SO_{\mathcal{G}} \cup WR_{\mathcal{G}} \cup WW_{\mathcal{G}}) ; RW_{\mathcal{G}})^+$, contradicting the assumption that $\mathcal{G} \in \text{GraphSI}$. □

Discussion. Our dependency graph characterisation of SI is instrumental in checking chopping correctness due to the ease of splicing a dependency graph (cf. (5)). Splicing abstract executions directly would be problematic. To illustrate this, consider the abstract

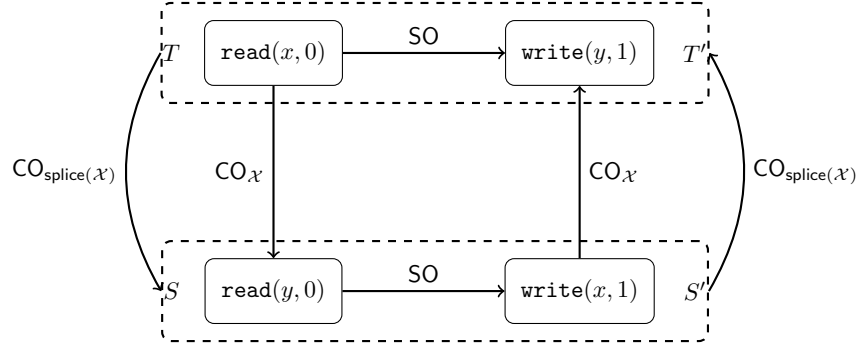


Fig. 12. An attempt to splice an execution directly.

execution \mathcal{X} in Figure 12, which is in ExecSI. A straightforward way to define $\text{splice}(\mathcal{X})$ is by letting

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

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

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

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

5.2. Static Analysis

We now derive a static analysis from Theorem 5.2. 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** P_i^j ($j = 1..k_i$), defining the code of the transactions in the sessions. We further assume that for each P_i we are given a sequence

$$(R_1^i, W_1^i) (R_2^i, W_2^i) \dots (R_{k_i}^i, W_{k_i}^i), \quad (10)$$

of **read and write sets** $R_j^i, W_j^i \subseteq \text{Obj}$, i.e., the sets of all objects that may be, respectively, read and written by the j -th piece of P_i . For example, the program **transfer** in Figure 6 consists of two pieces; the first one has the read and write sets equal to $\{\text{acct1}\}$ and the second, to $\{\text{acct2}\}$. The program **lookup1** consists of a single piece with the read set $\{\text{acct1}\}$ and the write set \emptyset .

Following Shasha et al. [Shasha et al., 1995], 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.

Definition 5.8. A history $\mathcal{H} = (\mathcal{T}, \text{SO})$ **can be produced** by programs

$$\mathcal{P} = \{(P_1^1, \dots, P_1^{k_1}), \dots, (P_i^1, \dots, P_i^{k_i})\} \quad (11)$$

with read and write sets (10) if there exists a bijective function f from transactions in \mathcal{T} to program pieces P_i^j of \mathcal{P} , such that

$$\begin{aligned} & (\forall T \in \mathcal{T}. \forall x \in \text{Obj}. \forall i, j. f(T) = P_i^j \wedge T \vdash \text{read}(x, _) \implies x \in R_i^j) \wedge \\ & (\forall T \in \mathcal{T}. \forall x \in \text{Obj}. \forall i, j. f(T) = P_i^j \wedge T \vdash \text{write}(x, _) \implies x \in W_i^j) \wedge \\ & (\forall T, S \in \mathcal{T}. \forall i, j, h. f(T) = P_i^j \wedge f(S) = P_i^h \wedge T \xrightarrow{\text{SO}} S \implies j < h). \end{aligned}$$

For example, the history in Figure 6 can be produced by the programs in the figure.

Definition 5.9. The chopping defined by the programs \mathcal{P} is **correct** if every dependency graph $\mathcal{G} \in \text{GraphSI}$, where $\mathcal{H}_{\mathcal{G}}$ can be produced by \mathcal{P} , is spliceable.

We check the correctness of \mathcal{P} by defining an analogue of the dynamic chopping graph from Definition 5.1, whose nodes are pieces of \mathcal{P} , rather than transactions in a given execution. Each piece is identified by a pair (i, j) of the number i of a program in \mathcal{P} and the piece's position j in the program.

Definition 5.10. Given a set of programs (11) with read and write sets (10), its **static chopping graph** $\text{SCG}(\mathcal{P})$ is the 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, S);
- $i_1 = i_2$ and $j_1 > j_2$ (a **predecessor** edge, P);
- $i_1 \neq i_2$ and $W_{i_1}^{j_1} \cap R_{i_2}^{j_2} \neq \emptyset$ (a **read dependency** edge, WR);
- $i_1 \neq i_2$ and $W_{i_1}^{j_1} \cap W_{i_2}^{j_2} \neq \emptyset$ (a **write dependency** edge, WW); or
- $i_1 \neq i_2$ and $R_{i_1}^{j_1} \cap W_{i_2}^{j_2} \neq \emptyset$ (an **anti-dependency** edge, RW).

The notion of a critical cycle introduced above for dynamic graphs is also applicable to static ones. The edge set of a static graph $\text{SCG}(\mathcal{P})$ over-approximates the edge sets of dynamic graphs $\text{DCG}(\mathcal{G})$ corresponding to dependency graphs \mathcal{G} produced by the programs \mathcal{P} . From this observation and Theorem 5.2 we easily get our static analysis.

COROLLARY 5.11. *The chopping defined by \mathcal{P} is correct if $\text{SCG}(\mathcal{P})$ contains no critical cycles.*

In Figure 13 we show the static chopping graph of the programs $\mathcal{P}^1 = \{\text{transfer}, \text{lookupAll}\}$, which contains a critical cycle:

$$\begin{aligned} (\text{var1} = \text{acct1}) \xrightarrow{\text{RW}} (\text{acct1} = \text{acct1} - 100) \xrightarrow{\text{S}} (\text{acct2} = \text{acct2} + 100) \xrightarrow{\text{WR}} \\ (\text{var2} = \text{acct2}) \xrightarrow{\text{P}} (\text{var1} = \text{acct1}). \quad (12) \end{aligned}$$

In fact, since the dependency graph in Figure 6 is not spliceable, the chopping defined by the above programs is incorrect. In Figure 14 we show the static chopping graph of the programs $\mathcal{P}^2 = \{\text{transfer}, \text{lookup1}, \text{lookup2}\}$. 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.

5.3. Comparison with Transaction Chopping Under Other Consistency Models

We now compare our chopping criterion for SI to criteria that have been proposed for other consistency models: serializability [Shasha et al., 1995] and parallel SI [Cerone et al., 2015b]. For clarity, in the following we refer to critical cycles of Definition 5.1 as **SI-critical**. The notion of chopping correctness in Definition 5.9 straightforwardly generalises to other consistency models.

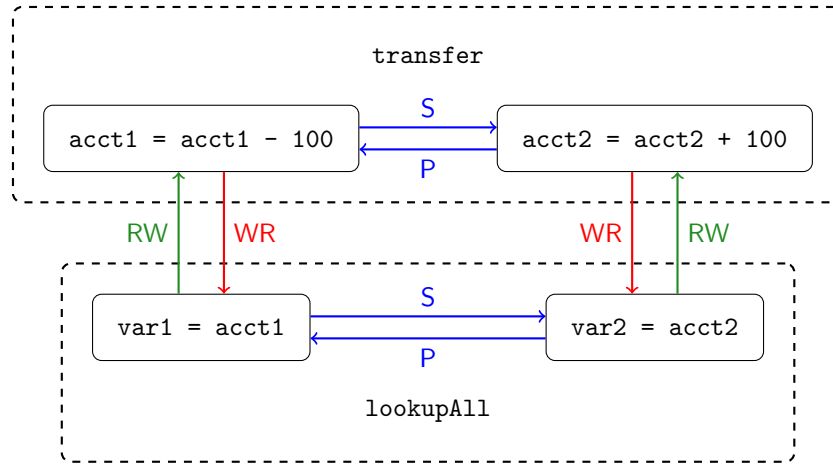


Fig. 13. The static chopping graph of the programs $\mathcal{P}^1 = \{\text{transfer}, \text{lookupAll}\}$ from Figure 6. Dashed boxes group program pieces into sessions.

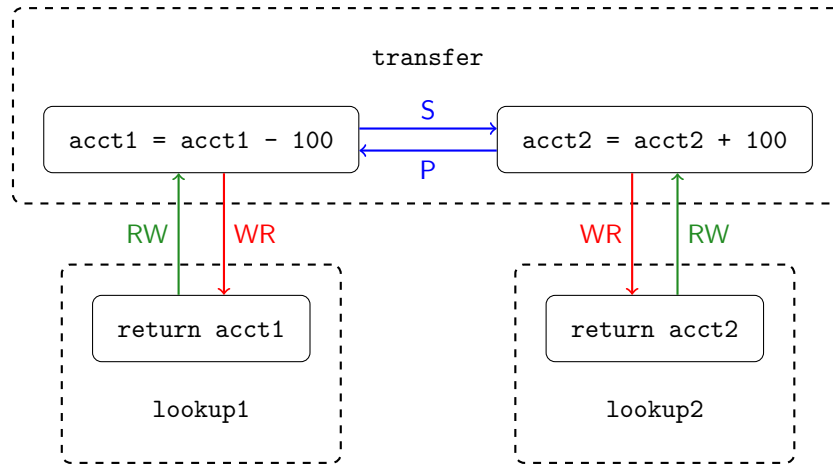


Fig. 14. The static chopping graph of the programs $\mathcal{P}^2 = \{\text{transfer}, \text{lookup1}, \text{lookup2}\}$ from Figure 6.

Transaction chopping under serializability. Following the approach in the proof of Theorem 5.2, we can easily establish the following improved version of the chopping criterion for serializability by Shasha et al. [Shasha et al., 1995].

Definition 5.12. A cycle in $\text{SCG}(\mathcal{P})$ is **SER-critical** if:

- (i) it is simple; and
- (ii) it contains a fragment of three consecutive edges of the form “conflict, predecessor, conflict”.

THEOREM 5.13. *The chopping defined by programs \mathcal{P} is correct under serializability if $\text{SCG}(\mathcal{P})$ contains no SER-critical cycles.*

For example, the chopping defined by the programs \mathcal{P}^2 considered in Figure 14 is correct under serializability. On the other hand, the chopping defined by \mathcal{P}^1 considered in Figure 13 is incorrect, and in fact $\text{SCG}(\mathcal{P}^1)$ contains a SER-critical cycle (12).

Any SI-critical cycle is also SER-critical and, thus, a chopping that is correct under serializability is also correct under SI. It follows that the classical transaction chopping analysis for serializability of Shasha et al. is also sound for SI. Note that 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.

The programs $\mathcal{P}^3 = \{\text{write1}, \text{write2}\}$ in Figure 15 define a correct chopping under SI, but not under serializability. Their static chopping graph $\text{SCG}(\mathcal{P}^3)$, shown in the figure, has only one simple cycle with three consecutive edges of the form “conflict, predecessor, conflict”:

$$(\text{var1} = x) \xrightarrow{S} (x = \text{var2}) \xrightarrow{P} (\text{var2} = y) \xrightarrow{RW} (y = \text{var1}) \xrightarrow{P} (\text{var1} = x).$$

This cycle is not SI-critical, and by Corollary 5.11, the chopping defined by \mathcal{P}^3 is correct under SI. However, it is incorrect under serializability. Indeed, consider the dependency graph $\mathcal{G}_7 \in \text{GraphSER}$ in Figure 15, whose history can be produced by \mathcal{P}^3 . Splicing the history $\mathcal{H}_{\mathcal{G}_7}$ results in a variant of the write skew anomaly, and thus $\text{splice}(\mathcal{H}_{\mathcal{G}_7}) \notin \text{HistSI}$.

Transaction chopping under parallel SI. We now compare our chopping criterion for SI to the one that we recently proposed [Cerone et al., 2015b] for parallel (aka non-monotonic) SI, a weakening of SI for large-scale databases [Sovran et al., 2011; Saeida Ardekani et al., 2013a]. 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 [Cerone et al., 2015a].

Definition 5.14. The sets of executions and histories **allowed by parallel SI** are:

$$\begin{aligned} \text{ExecPSI} &= \{\mathcal{X} \mid \mathcal{X} \models \text{INT} \wedge \text{EXT} \wedge \text{SESSION} \wedge \text{TRANSVIS} \wedge \text{NOCONFLICT}\}; \\ \text{HistPSI} &= \{\mathcal{H} \mid \exists \text{VIS}, \text{CO}. (\mathcal{H}, \text{VIS}, \text{CO}) \in \text{ExecSI}\}. \end{aligned}$$

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; in particular, it disallows the anomaly in Figure 2(d). 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(e), disallowed by the axiom PREFIX in SI.

Definition 5.15. A cycle in $\text{SCG}(\mathcal{P})$ is **PSI-critical** if:

- (i) it is simple;
- (ii) it contains a fragment of three consecutive edges of the form “conflict, predecessor, conflict”; and
- (iii) it contains at most one anti-dependency edge.

THEOREM 5.16 ([CERONE ET AL., 2015B]). *The chopping defined by programs \mathcal{P} is correct under parallel SI if $\text{SCG}(\mathcal{P})$ contains no PSI-critical cycles.*

Note that any cycle that is PSI-critical, is also SI-critical; as a consequence, the programs \mathcal{P}^2 from Figure 14 and \mathcal{P}^3 from Figure 15 considered above define a correct chopping under parallel SI. On the other hand, the programs \mathcal{P}^1 from Figure 13 are not chopped correctly under parallel SI, and the cycle (12) in $\text{SCG}(\mathcal{P}^1)$ is PSI-critical.

The programs $\mathcal{P}^4 = \{\text{write1}, \text{write2}, \text{read1}, \text{read2}\}$ in Figure 16 define a correct chopping under parallel SI, but not SI. The static chopping graph $\text{SCG}(\mathcal{P}^4)$, shown in the figure,

```

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

```

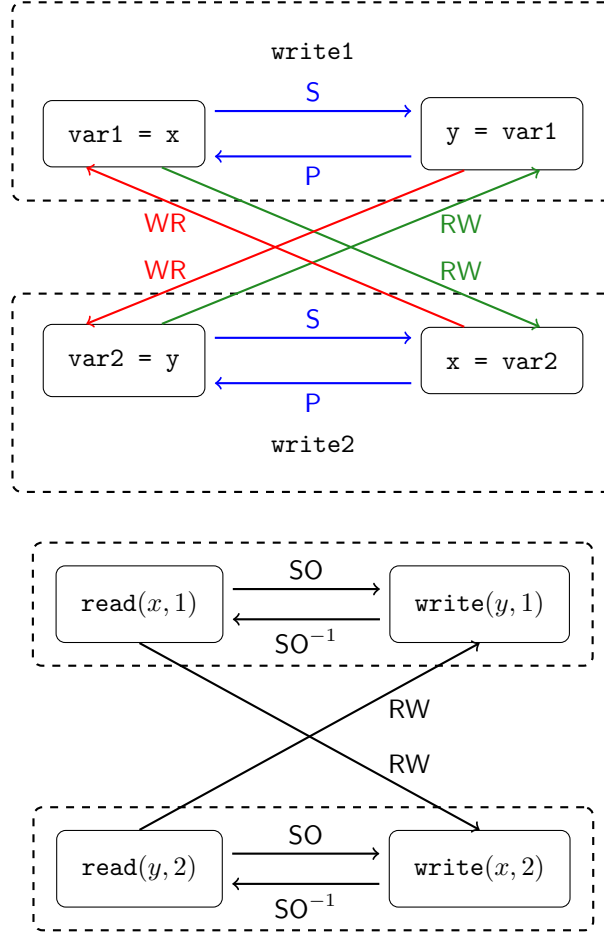


Fig. 15. Programs \mathcal{P}^3 defining a chopping correct under SI, but not under serializability.

contains exactly one simple cycle with three consecutive edges of the form “conflict, predecessor, conflict”:

$$\begin{aligned}
 (x = \text{post1}) \xrightarrow{\text{WR}} (b = x) \xrightarrow{\text{P}} (a = y) \xrightarrow{\text{RW}} (y = \text{post2}) \xrightarrow{\text{WR}} \\
 (b = y) \xrightarrow{\text{P}} (a = x) \xrightarrow{\text{RW}} (x = \text{post1}).
 \end{aligned}$$

This cycle is not PSI-critical, so that \mathcal{P}^4 indeed define a correct chopping under parallel SI. 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 $\mathcal{G}_8 \in \text{GraphSI}$ in Figure 16, whose history can be produced by \mathcal{P}^4 . Splicing the history $\mathcal{H}_{\mathcal{G}_8}$ results in a long fork anomaly, and thus $\text{splice}(\mathcal{H}_{\mathcal{G}_8}) \notin \text{HistSI}$.

```

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); }
    
```

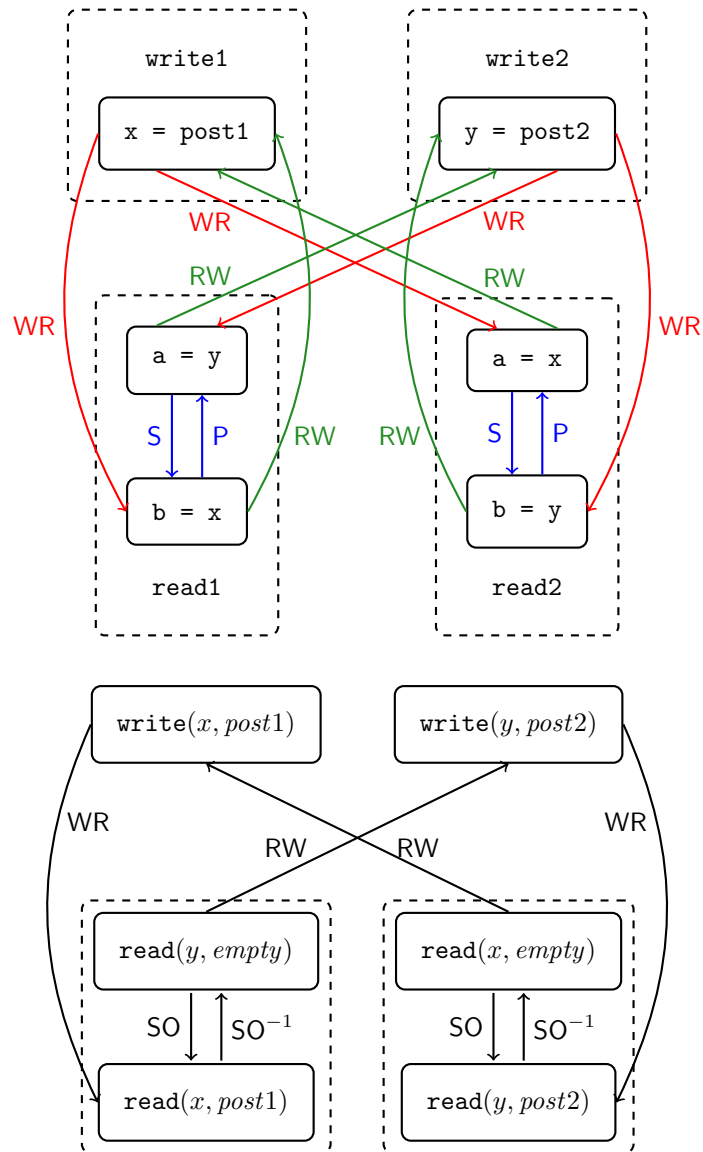


Fig. 16. Programs \mathcal{P}^4 defining a chopping correct under parallel SI, but not under SI.

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 [Fekete et al., 2005] (robustness **against SI**). For this the analysis checks that the application code may produce no histories in $\text{HistSI} \setminus \text{HistSER}$. 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 $\text{GraphSI} \setminus \text{GraphSER}$. This easily follows from Theorems 3.4 and 4.2.

THEOREM 6.1. *For any \mathcal{G} , we have $\mathcal{G} \in \text{GraphSI} \setminus \text{GraphSER}$ if and only if $\mathcal{T}_{\mathcal{G}} \models \text{INT}$, \mathcal{G} contains a cycle, and all its cycles have at least two adjacent anti-dependency edges.*

The dependency graph \mathcal{G}_3 of the write skew anomaly in Figure 2(f) contains a cycle: $T_1 \xrightarrow{\text{RW}_{\mathcal{G}_3}} T_2 \xrightarrow{\text{RW}_{\mathcal{G}_3}} T_1$. Furthermore, it is easy to see that all its cycles have two adjacent anti-dependencies, so that $\mathcal{G}_3 \in \text{GraphSI} \setminus \text{GraphSER}$. The dependency graph \mathcal{G}_5 from Figure 6 does not contain any cycles, so that $\mathcal{G}_5 \notin \text{GraphSI} \setminus \text{GraphSER}$; in fact, $\mathcal{G}_5 \in \text{GraphSER}$. Finally, the dependency graph \mathcal{G}_2 of the long fork anomaly in Figure 2(e) contains a cycle, but without two adjacent anti-dependencies; hence, $\mathcal{G}_2 \notin \text{GraphSI} \setminus \text{GraphSER}$, and in fact, $\mathcal{G}_2 \notin \text{GraphSI}$.

We can obtain the following refinement of Theorem 6.1 by using the characterisation of SI in Theorem 4.13 instead of that in Theorem 4.2.

THEOREM 6.2. *For any \mathcal{G} , we have $\mathcal{G} \in \text{GraphSI} \setminus \text{GraphSER}$ if and only if $\mathcal{T}_{\mathcal{G}} \models \text{INT}$, \mathcal{G} contains a cycle, and all its cycles have at least two adjacent vulnerable anti-dependency edges over different objects.*

Fekete et al. previously established a result roughly corresponding to the “only if” direction of the above theorem [Fekete et al., 2005]. 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.

We can derive a static analysis from Theorem 6.1 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} . It then constructs a **static dependency graph**, over-approximating possible read and write dependencies and vulnerable anti-dependencies that may exist in executions of \mathcal{P} ; in contrast with the chopping analysis, the robustness analysis also records the objects to which different edges refer to. The analysis then checks that the graph has no cycles with at least two adjacent vulnerable anti-dependency edges over different objects. By Theorem 6.2 this implies that the programs \mathcal{P} produce no histories in $\text{HistSI} \setminus \text{HistSER}$, and hence, the corresponding application is robust against SI.

The construction of the static dependency graph above can be done like in the chopping analysis, using sets of objects that may be read or written by programs in \mathcal{P} . During this construction, we can exclude vulnerable anti-dependencies using information about objects that *must* necessarily be written by certain programs [Fekete et al., 2005] using the following proposition.

PROPOSITION 6.3. *Consider an anti-dependency edge $T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S$ in a dependency graph $\mathcal{G} \in \text{GraphSI}$. If $T \vdash \text{write}(x, _)$, then the anti-dependency is not vulnerable.*

PROOF. From $T \xrightarrow{RW_{\mathcal{G}(x)}} S$ we get $T \vdash \text{write}(x, _)$ and $T \neq S$. Since we also know $T \vdash \text{write}(x, _)$, we must have either $T \xrightarrow{WW_{\mathcal{G}(x)}} S$ or $S \xrightarrow{WW_{\mathcal{G}(x)}} T$. However, the latter yields a cycle $T \xrightarrow{RW_{\mathcal{G}(x)}} S \xrightarrow{WW_{\mathcal{G}(x)}} T$, contradicting $\mathcal{G} \in \text{GraphSI}$ by Theorem 4.1. Hence, we must have $T \xrightarrow{WW_{\mathcal{G}(x)}} S$, implying the required. \square

Note that the dependency graphs characterisation of consistency models greatly facilitates deriving the above robustness 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 (§5.3) behaves the same as when executing under the classical SI (robustness *against parallel SI towards SI*). First, we give a characterisation of parallel SI in terms of dependency graphs.

THEOREM 6.4 ([CERONE ET AL., 2015B, EXTENDED VERSION, LEMMA 14]). *Let*

$$\text{GraphPSI} = \{\mathcal{G} \mid (\mathcal{T}_{\mathcal{G}} \models \text{INT}) \wedge (((\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}})^+ ; \text{RW}_{\mathcal{G}}?) \text{ is irreflexive})\}.$$

Then

$$\text{HistPSI} = \{\mathcal{H} \mid \exists \text{WR, WW, RW. } (\mathcal{H}, \text{WR, WW, RW}) \in \text{GraphPSI}\}.$$

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}_2 in Figure 2(e). It is easy to see that all its cycles contain at least two anti-dependencies, and therefore $\mathcal{G}_2 \in \text{GraphPSI}$. On the other hand, let \mathcal{G}_1 be the dependency graph in Figure 2(b).

The graph \mathcal{G}_1 contains a cycle with exactly one anti-dependency ($T_1 \xrightarrow{WW_{\mathcal{G}_1}} T_2 \xrightarrow{RW_{\mathcal{G}_1}} T_1$), and therefore $\mathcal{G}_1 \notin \text{GraphPSI}$. As a corollary of Theorems 4.1 and 6.4, we obtain a dynamic robustness criterion that checks whether a given dependency graph is in $\text{GraphPSI} \setminus \text{GraphSI}$.

THEOREM 6.5. *For any \mathcal{G} , we have $\mathcal{G} \in \text{GraphPSI} \setminus \text{GraphSI}$ if and only if $\mathcal{T}_{\mathcal{G}} \models \text{INT}$, \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}_2 of the long fork anomaly (Figure 2(e)) all cycles have at least two anti-dependencies. Furthermore, \mathcal{G}_2 also has a cycle with no adjacent anti-dependencies:

$$T_1 \xrightarrow{WR_{\mathcal{G}_2}} T_3 \xrightarrow{RW_{\mathcal{G}_2}} T_2 \xrightarrow{WR_{\mathcal{G}_2}} T_4 \xrightarrow{RW_{\mathcal{G}_2}} T_1,$$

so that $\mathcal{G}_2 \in \text{GraphPSI} \setminus \text{GraphSI}$. The dependency graph \mathcal{G}_3 of the write skew anomaly in Figure 2(f) contains only cycles with at least two adjacent anti-dependencies, so that $\mathcal{G}_3 \notin \text{GraphPSI} \setminus \text{GraphSI}$; in fact, $\mathcal{G}_3 \in \text{GraphSI}$. The dependency graph \mathcal{G}_1 of the lost update anomaly (Figure 2(b)) contains a cycle with exactly one anti-dependency, so that $\mathcal{G}_1 \notin \text{GraphPSI} \setminus \text{GraphSI}$; in fact, $\mathcal{G}_1 \notin \text{GraphPSI}$.

We can refine the characterisation in Theorem 6.4 by taking into account the objects involved in dependencies between transactions, similarly to how we did this for SI (Theorem 4.13, §4.5); for brevity, we omit a refinement that takes into account vulnerability.

THEOREM 6.6. *Let \mathcal{G} be a dependency graph. Then $\mathcal{G} \in \text{GraphPSI}$ if and only if $\mathcal{T}_{\mathcal{G}} \models \text{INT}$ and all cycles in \mathcal{G} contain at least two anti-dependency edges over different objects.*

PROOF. The “if” part of the theorem is trivial. For the “only if” part, suppose that $\mathcal{G} \in \text{GraphPSI}$. We can show that all cycles in \mathcal{G} contain at least two anti-dependency edges

over different objects using the following algebraic law:

$$\forall x. \text{RW}_{\mathcal{G}}(x) ; (\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}})^* ; \text{RW}_{\mathcal{G}}(x) \subseteq \text{RW}_{\mathcal{G}}(x) ; \text{WW}_{\mathcal{G}}(x). \quad (13)$$

Indeed, if \mathcal{G} contains a cycle where all anti-dependency edges are over the same object x , then (13) allows us to convert this cycle into one with at most one anti-dependency edge, which contradicts $\mathcal{G} \in \text{GraphPSI}$ by Theorem 6.4.

It remains to prove (13). To this end, consider $x \in \text{Obj}$ and $T, T', S', S \in \mathcal{T}_{\mathcal{G}}$ such that

$$T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T' \xrightarrow{(\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}})^*} S \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S'.$$

Then $T' \vdash \text{write}(x, _)$ and $S' \vdash \text{write}(x, _)$. Therefore, we have one of the following: $T' = S'$, $S' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$ or $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S'$. We cannot have that $T' = S'$: if this were the case, we would have the cycle

$$T' \xrightarrow{(\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}})^*} S \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T'$$

in \mathcal{G} with a single anti-dependency edge, contradicting $G \in \text{GraphPSI}$. We cannot have $S' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} T'$: in this case we would have the cycle

$$S' \xrightarrow{(\text{SO}_{\mathcal{G}} \cup \text{WR}_{\mathcal{G}} \cup \text{WW}_{\mathcal{G}})^+} S \xrightarrow{\text{RW}_{\mathcal{G}}(x)} S',$$

again disallowed because $\mathcal{G} \in \text{GraphPSI}$. We are thus left with the case $T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S$, so that $T \xrightarrow{\text{RW}_{\mathcal{G}}(x)} T' \xrightarrow{\text{WW}_{\mathcal{G}}(x)} S$. This establishes (13), as required. \square

As an illustration of Theorem 6.6, the dependency graph \mathcal{G}_2 of the long fork anomaly in Figure 2(e), allowed by parallel SI, contains only cycles with at least two anti-dependencies over different objects: e.g.,

$$T_1 \xrightarrow{\text{WR}_{\mathcal{G}_2}(x)} T_3 \xrightarrow{\text{RW}_{\mathcal{G}_2}(y)} T_2 \xrightarrow{\text{WR}_{\mathcal{G}_2}(y)} T_4 \xrightarrow{\text{RW}_{\mathcal{G}_2}(x)} T_1.$$

We can obtain the following refinement of Theorem 6.5 by using the characterisations of SI and PSI in Theorems 4.13 and 6.6.

THEOREM 6.7. *For any \mathcal{G} , we have $\mathcal{G} \in \text{GraphPSI} \setminus \text{GraphSI}$ if and only if $\mathcal{T}_{\mathcal{G}} \models \text{INT}$, \mathcal{G} contains at least one cycle with no adjacent anti-dependency edges over different objects, and all its cycles have at least two anti-dependency edges over different objects.*

From Theorem 6.7 it follows that a static analysis for robustness against parallel SI towards SI can check that the static dependency graph of an application contains no cycles where there are at least two anti-dependency edges over different objects and no two anti-dependency edges over different objects are adjacent.

7. RELATED WORK

Snapshot isolation was originally defined by an idealised algorithm formulated in terms of implementation-level concepts [Berenson et al., 1995]. Since then there have been proposals of more declarative SI specifications [Adya, 1999; Saeida Ardekani et al., 2013b; Cerone et al., 2015a], one of which [Cerone et al., 2015a] 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. [Fekete et al., 2005] 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 4.2(ii)), but they did not establish an analogue of our soundness result (Theorem 4.2(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 [Cahill et al., 2009; Zellag and Kemme, 2014] and proving the correctness of concurrency-control algorithms [Xie et al., 2015; Diegues and Romano, 2014]. 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* [Terry et al., 2013].

The constraint on dependency graphs that we use to characterise SI also arose in the work of Lin et al. [Lin et al., 2009], 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 [Zhang et al., 2013; Mu et al., 2014; Xie et al., 2015]. There have also been proposals of consistency models for transactional memory that weaken consistency guarantees in a way similar to chopping [Felber et al., 2009; Xiang and Scott, 2015; Afek et al., 2011]. Our chopping analysis enables bringing these benefits to transactional systems providing SI. We have previously proposed a chopping analysis for parallel SI [Cerone et al., 2015b], which also relies on a dependency graph characterisation of this consistency model (Theorem 6.4, §6.2). 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 the 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. The first author was also supported by the EPSRC Programme Grant REMS: Rigorous Engineering for Mainstream Systems (EP/K008528/1).

REFERENCES

- Adya, A. (1999). Weak consistency: A generalized theory and optimistic implementations for distributed transactions. PhD thesis, MIT.
- Adya, A., Liskov, B., and O'Neil, P. E. (2000). Generalized isolation level definitions. In *ICDE*.
- Afek, Y., Avni, H., and Shavit, N. (2011). Towards consistency oblivious programming. In *OPODIS*.
- Bailis, P., Fekete, A., Ghodsi, A., Hellerstein, J. M., and Stoica, I. (2014). Scalable atomic visibility with RAMP transactions. In *SIGMOD*.
- Berenson, H., Bernstein, P., Gray, J., Melton, J., O'Neil, E., and O'Neil, P. (1995). A critique of ANSI SQL isolation levels. In *SIGMOD*.
- Bernstein, P. A., Hadzilacos, V., and Goodman, N. (1987). *Concurrency Control and Recovery in Database Systems*. Addison-Wesley.
- Bieniusa, A. and Fuhrmann, T. (2010). Consistency in hindsight: A fully decentralized STM algorithm. In *IPDPS*.
- Cahill, M. J., Röhm, U., and Fekete, A. D. (2009). Serializable isolation for snapshot databases. *ACM Trans. Database Syst.*, 34(4).
- Cerone, A., Bernardi, G., and Gotsman, A. (2015a). A framework for transactional consistency models with atomic visibility. In *CONCUR*. Dagstuhl.
- Cerone, A., Gotsman, A., and Yang, H. (2015b). Transaction chopping for parallel snapshot isolation. In *DISC*. Extended version available from www.software.imdea.org/~gotsman.

- Closure (2016). Refs and transactions. <http://clojure.org/refs>.
- Daudjee, K. and Salem, K. (2004). Lazy database replication with ordering guarantees. In *ICDE*.
- Daudjee, K. and Salem, K. (2006). Lazy database replication with snapshot isolation. In *VLDB*.
- Dias, R. J., Lourenço, J. M., and Pregoça, N. (2011). Efficient and correct transactional memory programs combining snapshot isolation and static analysis. In *HotPar*.
- Diegues, N. and Romano, P. (2014). Time-warp: Lightweight abort minimization in transactional memory. In *PPoPP*.
- Doherty, S., Groves, L., Luchangco, V., and Moir, M. (2013). Towards formally specifying and verifying transactional memory. *Formal Aspects of Computing*, 25(5).
- Elnikety, S., Zwaenepoel, W., and Pedone, F. (2005). Database replication using generalized snapshot isolation. In *SRDS*.
- Fekete, A., Liarokapis, D., O’Neil, E., O’Neil, P., and Shasha, D. (2005). Making snapshot isolation serializable. *ACM Trans. Database Syst.*, 30(2).
- Felber, P., Gramoli, V., and Guerraoui, R. (2009). Elastic transactions. In *DISC*.
- Guerraoui, R. and Kapalka, M. (2008). On the correctness of transactional memory. In *PPoPP*.
- Herlihy, M. and Moss, J. E. B. (1993). Transactional memory: Architectural support for lock-free data structures. In *ISCA*.
- Jorwekar, S., Fekete, A., Ramamritham, K., and Sudarshan, S. (2007). Automating the detection of snapshot isolation anomalies. In *VLDB*.
- Lin, Y., Kemme, B., Jiménez-Peris, R., Patiño-Martínez, M., and Armendáriz-Iñigo, J. E. (2009). Snapshot isolation and integrity constraints in replicated databases. *ACM Trans. Database Syst.*, 34(2).
- Litz, H., Cheriton, D., Firoozshahian, A., Azizi, O., and Stevenson, J. P. (2014). SI-TM: Reducing transactional memory abort rates through snapshot isolation. In *ASPLOS*.
- Mu, S., Cui, Y., Zhang, Y., Lloyd, W., and Li, J. (2014). Extracting more concurrency from distributed transactions. In *OSDI*.
- Peng, D. and Dabek, F. (2010). Large-scale incremental processing using distributed transactions and notifications. In *OSDI*.
- Riegel, T., Fetzner, C., and Felber, P. (2006). Snapshot isolation for software transactional memory. In *Workshop on Languages, Compilers, and Hardware Support for Transactional Computing (TRANSACT)*.
- Saeida Ardekani, M., Sutra, P., and Shapiro, M. (2013a). Non-monotonic snapshot isolation: Scalable and strong consistency for geo-replicated transactional systems. In *SRDS*.
- Saeida Ardekani, M., Sutra, P., Shapiro, M., and Pregoça, N. (2013b). On the scalability of snapshot isolation. In *Euro-Par*.
- Serrano, D., Patiño-Martínez, M., Jiménez-Peris, R., and Kemme, B. (2007). Boosting database replication scalability through partial replication and 1-copy-snapshot-isolation. In *PRDC*.
- Shasha, D., Llirbat, F., Simon, E., and Valduriez, P. (1995). Transaction chopping: Algorithms and performance studies. *ACM Trans. Database Syst.*, 20(3).
- Shasha, D. and Snir, M. (1988). Efficient and correct execution of parallel programs that share memory. *ACM Trans. Program. Lang. Syst.*, 10(2).
- Sovran, Y., Power, R., Aguilera, M. K., and Li, J. (2011). Transactional storage for geo-replicated systems. In *SOSP*.
- Terry, D. B., Demers, A. J., Petersen, K., Spreitzer, M., Theimer, M., and Welch, B. W. (1994). Session guarantees for weakly consistent replicated data. In *PDIS*.
- Terry, D. B., Prabhakaran, V., Kotla, R., Balakrishnan, M., Aguilera, M. K., and Abu-Libdeh, H. (2013). Consistency-based service level agreements for cloud storage. In *SOSP*.
- Xiang, L. and Scott, M. L. (2015). Software partitioning of hardware transactions. In *PPoPP*.
- Xie, C., Su, C., Littley, C., Alvisi, L., Kapritsos, M., and Wang, Y. (2015). High-performance ACID via modular concurrency control. In *SOSP*.
- Zellag, K. and Kemme, B. (2014). Consistency anomalies in multi-tier architectures: Automatic detection and prevention. *The VLDB Journal*, 23(1).
- Zhang, Y., Power, R., Zhou, S., Sovran, Y., Aguilera, M., and Li, J. (2013). Transaction chains: Achieving serializability with low latency in geo-distributed storage systems. In *SOSP*.