Data Consistency in Transactional Storage Systems: A Centralised Approach

SHALE XIONG, Imperial College London, UK
ANDREA CERONE, Football Radar, UK
AZALEA RAAD, MPI-SWS, Germany
PHILIPPA GARDNER, Imperial College London, UK

We introduce an interleaving operational semantics for describing the client-observable behaviour of atomic transactions on distributed key-value stores. Our semantics builds on abstract states comprising centralised, global key-value stores and partial client views. We provide operational definitions of consistency models for our key-value stores which are shown to be equivalent to the well-known declarative definitions of consistency model for execution graphs. We explore two immediate applications of our semantics: specific protocols of geo-replicated databases (e.g. COPS) and partitioned databases (e.g. Clock-SI) can be shown to be correct for a specific consistency model by embedding them in our centralised semantics; programs can be directly shown to have invariant properties such as robustness results against a weak consistency model.

1 INTRODUCTION

Transactions are the *de facto* synchronisation mechanism in modern distributed databases. To achieve scalability and performance, distributed databases often use weak transactional consistency guarantees. These weak guarantees pose several challenges: the formalisation of client-observable behaviour; and the verification of database protocols and client applications. Much work has been done to formalise the semantics of such consistency guarantees, both declaratively and operationally. On the declarative side, several general formalisms have been proposed, such as dependency graphs [Adya 1999] and abstract executions [Burckhardt et al. 2012], to provide a unified semantics for formulating different consistency models. On the operational side, the semantics of *specific* consistency models have been captured using reference implementations [Berenson et al. 1995; Raad et al. 2018; Sovran et al. 2011]. However, unlike declarative approaches, there has been little work on general operational semantics for describing a range of consistency models.

We summarise briefly the existing work on general operational semantics; further discussion can be found in §7. Kaki et al. [2017] propose an operational semantics for SQL transaction programs under the consistency models given by the standard ANSI/SQL isolation levels [Berenson et al. 1995]. Their operational semantics accesses a global, centralised store, and is used to develop a program logic and prototype tool for reasoning about client programs. They capture consistency models such as snapshot isolation (SI) [Berenson et al. 1995], but not weaker ones such as parallel snapshot isolation (PSI) [Sovran et al. 2011] and causal consistency (CC) [Lloyd et al. 2011] which are important for distributed databases. Following this work, Nagar and Jagannathan [2018] propose an operational semantics over abstract execution graphs, rather than a concrete centralised store, in order to prove the robustness of applications against a given consistency model. They are able to capture weaker consistency models such as PSI and CC . However, although they focus on consistency models with snapshot property[1], their semantics allows for the fine-grained

---

[1] *Snapshot property*, also known as *atomic visibility*, means that a transaction reads from an atomic snapshot of the database, and commits atomically.

interleaving of operations in different transactions. We believe that this results in an unnecessarily complicated semantics. Crooks et al. [2017] provide a trace semantics over a global centralised store, where the behaviour of clients is formalised by the observations they make on the totally-ordered history of states. They use their semantics to demonstrate the equivalence of several implementation-specific definitions of consistency model. However, the usefulness of their approach for analysing client programs is not clear, since observations made by their clients involve information that is not generally available to real-world clients, such as the total order in which transactions commit.

We introduce an interleaving operational semantics for describing the client-observable behaviour of atomic transactions on distributed key-value stores (§ 2, § 3), successfully abstracting from the internal details of protocols of geo-replicated and partitioned databases. In our semantics, transactions execute atomically, preventing fine-grained interleaving of the operations they perform. Our semantics comprises a global, centralised key-value store (kv-store) with *multi-versioning*, which records all the versions of a key, and partial *client views*, which let clients see only a subset of the versions. Our approach is partly inspired by the views in the C11 operational semantics in [Kang et al. 2017]. Our operational semantics is parametric in the notion of *execution test*, which determines if a client with a given view is allowed to commit a transaction. Just as with standard operational semantics, the next transaction step just depends on the current kv-store and client view. Our execution tests resemble an approach taken in [Crooks et al. 2017], except that the next step requires an analysis of the whole trace. We capture most of the well-known consistency models in a uniform way (§ 4): e.g, CC, PSI, SI and serialisability (SER). We also identify a new consistency model that sits between PSI and SI and retains good properties of both. We call this new consistency model *weak snapshot isolation* (WSI). Since we focus on snapshot property, we are not able to capture popular consistency models such as *read committed*. Because our focus is on protocols and applications employed by distributed databases, most of which guarantee snapshot property, we do not find this constraint to be a severe limitation.

We prove that our operational definitions of consistency models for kv-stores are equivalent to the well-known declarative definitions of consistency model for execution graphs. Such results have not been given in [Kaki et al. 2017], but have been given in [Crooks et al. 2017]. We provide a general proof technique which captures the correspondence between our execution tests and the axiomatic specifications of consistency models for abstract executions (§ 5 and § E). Using this technique, we prove that our definitions of consistency model for kv-stores are equivalent to the declarative definitions of consistency model for abstract executions [Cerone et al. 2015a], and hence for dependency graphs [Adya 1999; Cerone et al. 2017].

We explore two immediate applications of our operational semantics: the establishment of invariant properties such as robustness for simple client applications; and the correctness of specific distributed protocols. By contrast, these tasks tend to be carried out in different declarative formalisms: clients are analysed using dependency graphs [Bernardi and Gotsman 2016; Cerone and Gotsman 2016; Cerone et al. 2015b; Fekete et al. 2005; Nagar and Jagannathan 2018]; protocols are verified using abstract execution graphs [Burckhardt et al. 2014; Cerone et al. 2015a]; and equivalence results are used to move between the two [Cerone et al. 2017]. We prove the robustness of a single counter against PSI and the robustness of multiple counters against our new model WSI (§ 6). Alomari et al. [2008] present a banking example that is robust against SI. We show that it is also robust against WSI (§6). To our knowledge, our robustness results are the first to take into account client sessions. With sessions, we demonstrate that multiple counters are not robust against PSI. Interestingly, without sessions, it can be shown that multiple counters are robust against PSI using static-analysis techniques [Bernardi and Gotsman 2016]; these techniques are known not to be applicable to sessions. We also establish the correctness of two database protocols against their

$k \mapsto \boxed{0 \; \frac{t_0}{\emptyset}}$

(a) Initial state

$k \mapsto \boxed{0 \; \frac{t_0}{\{t\}} \; 1 \; \frac{t}{\emptyset}}$

(b) After t

$k \mapsto \boxed{0 \; \frac{t_0}{\{t\}} \; 1 \; \frac{t}{\emptyset}}$

(c) A possible view of cl_2

$k \mapsto \boxed{0 \; \frac{t_0}{\{t,t'\}} \; 1 \; \frac{t}{\emptyset} \; 1 \; \frac{t'}{\emptyset}}$
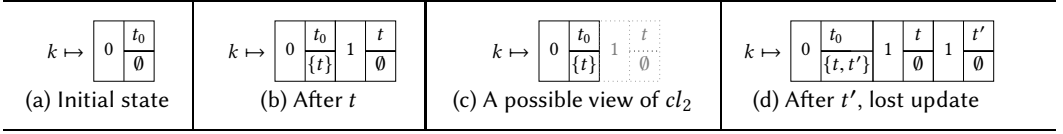
(d) After t', lost update

Fig. 0. Example key-value stores (a, b, d); a client view (c)

consistency models, demonstrating that they can be accurately expressed in our centralised semantics: the COPS protocol for the fully replicated kv-stores [Lloyd et al. 2011] which satisfies CC (§6); and the Clock-SI protocol for partitioned kv-stores [Du et al. 2013] which satisfies SI (§H).

2 OVERVIEW

We motivate our key ideas, centralised kv-stores, partial client views and execution tests, via an intuitive example. We show that our interleaving semantics is an ideal middle point for proving invariant properties such as robustness, and verifying distributed protocols.

Example. We use a simple transactional library, Counter(k), to introduce our operational semantics. Clients of this counter library can manipulate the value of key k via two transactions:

$$\text{inc}(k) \triangleq [\mathsf{x} := [k]; \; [k] := \mathsf{x} + 1] \qquad \text{read}(k) \triangleq [\mathsf{x} := [k]]$$

Command $\mathsf{x} := [k]$ reads the value of key k to local variable x; command $[k] := \mathsf{x} + 1$ writes the value of $\mathsf{x} + 1$ to key k. The code of each operation is wrapped in square brackets, denoting that it must be executed *atomically* as a transaction.

Consider a replicated database where a client only interacts with one replica. For such a database, the correctness of atomic transactions is subtle, depending heavily on the particular consistency model under consideration. Consider the client program $\mathsf{P}_{\mathsf{LU}} = (cl_1 : \text{inc}(k) \; || \; cl_2 : \text{inc}(k))$, where we assume that the clients cl_1 and cl_2 work on different replicas and the k initially holds value 0 in all replicas. Intuitively, since transactions are executed atomically, after both calls to inc(k) have terminated, the counter should hold the value 2. Indeed, this is the only outcome allowed under SER where transactions appear to execute in a sequential (serial) order, one after another. The implementation of SER in distributed kv-stores comes at a significant performance cost. Therefore, implementers are content with weaker consistency models [Ardekani et al. 2014; Bailis et al. 2014; Binnig et al. 2014; Du et al. 2013; Li et al. 2012; Liu et al. 2018; Lloyd et al. 2011; Saeida Ardekani et al. 2013; Sovran et al. 2011; Spirovska et al. 2018]. For example, if the replicas provide no synchronisation mechanism for transactions, then it is possible for both clients to read the same initial value 0 for k at their distinct replicas, update them to 1, and eventually propagate their updates to other replicas. Consequently, both sites are unchanged with value 1 for k. This weak behaviour is known as the *lost update* anomaly, which is allowed under the consistency model called *causal consistency* [Li et al. 2012; Lloyd et al. 2011; Spirovska et al. 2018].

Centralised Operational Semantics. A well-known declarative approach for providing general reasoning about clients of distributed kv-stores is to use execution graphs [Adya 1999; Adya et al. 2000; Burckhardt et al. 2012; Cerone et al. 2015a], where nodes are atomic transactions and edges describe the known dependencies between transactions. The graphs capture the behaviour of the whole program, with different consistency models corresponding to different sets of axioms constraining the graphs. However, execution graphs provide little information about how the state of a kv-store evolves throughout the execution of a program. By contrast, we provide an interleaving operational semantics based on an abstract centralised state. The centralised state comprises a centralised, multi-versioned kv-store, which is *global* in the sense that it contains all the versions written by clients, and client views of the store, which are *partial* in the sense that clients may see

different subsets of the versions in the kv-store. Each update is given by either a simple primitive command or an atomic transaction. The atomic transaction steps are subject to an *execution test* which analyses the state to determine whether the update is allowed by the associated consistency model.

Let us introduce our global kv-stores and partial client views by showing that we can reproduce the lost update anomaly given by P_{LU}. Our kv-stores are functions mapping keys to lists of versions, where the versions record all the values written to each key together with the meta-data of the transactions that access it. In the P_{LU} example, the initial kv-store comprises a single key k, with only one possible version $(0, t_0, \emptyset)$, stating that k holds value 0, that the version *writer* is the initialising transaction t_0 (this version was written by t_0), and that the version *reader set* is empty (no transaction has read this version as of yet). Fig. 0a depicts this initial kv-store, with the version represented as a box sub-divided in three sections: the value 0; the writer t_0; and the reader set $\emptyset$.

First, suppose that cl_1 invokes inc on Fig. 0a. It does this by choosing a fresh transaction identifier, t, and then proceeds with inc(k). It reads the initial version of k with value 0 and then writes a new value 1 for k. The resulting kv-store is depicted in Fig. 0b, where the initial version of k has been updated to reflect that it has been read by t.

Second, client cl_2 invokes inc on Fig. 0b. As there are now two versions available for k, we must determine the version from which cl_2 fetches its value, before running inc(k). This is where *client views* come into play. Intuitively, a view of client cl_2 comprises those versions in the kv-store that are *visible* to cl_2, i.e. those that can be read by cl_2. If more than one version is visible, then the newest (right-most) version is selected, modelling the *last-writer-wins* resolution policy used by many distributed kv-stores [Vogels 2009]. In our example, there are two view candidates for cl_2 when running inc(k) on Fig. 0b: (1) one containing only the initial version of k; (2) the other containing both versions of k.[2] For (1), the view is depicted in Fig. 0c. Client cl_2 chooses a fresh transaction identifier t', reads the initial value 0 and writes a new version with value 1, as depicted in Fig. 0d. Such a kv-store does not contain a version with value 2, despite two increments on k, producing the lost update anomaly. For (2), client cl_2 reads the newest value 1 and writes a new version with value 2.

To avoid undesirable behaviour, such as the lost update anomaly, we use an *execution test* which restricts the possible update at the point of the transaction commit. One such test is to enforce a client to commit a transaction writing to k if and only if its view contains all versions available in the global state for k. This prevents cl_2 from running inc(k) on Fig. 0b if its view only contains the initial version of k. Instead, the cl_2 view must contain both versions of k, thus enforcing cl_2 to write a version with value 2 after running inc(k). This particular test corresponds to *write-conflict-freedom* of distributed kv-stores: at most one concurrent transaction can write to a key at any one time. In §4 we give many examples of execution tests and their associated consistency models on kv-stores. In §5, we develop a proof technique, which we use in §F to show the equivalence of our operational definitions of consistency models and the declarative ones based on execution graphs.

General Robustness Results. The first application of our operational semantics is to prove general robustness results for clients with respect to specific consistency models (§ 6.1). Using our general operational semantics, we can prove invariant properties (e.g. robustness) of a program P under weak consistency models. That is, the invariant obtained by executing P under a weak consistency model can also be obtained under serialisability. To demonstrate this, we prove the robustness of the single counter library discussed above against PSI, and the robustness of a multi-counter library and the banking library of Alomari et al. [2008] against our new proposed model WSI and all stronger models such as SI. The latter is done through general conditions on invariant

---

[2] As we explain in §3.1, we always require the view of a client to include the initial version of each key.

which guarantees robustness against WSI. Thanks to our operational semantics, our invariant-based approaches only need to work with single program steps rather than whole program traces.

Verifying Implementation Protocols. The second application of our operational semantics is for showing that implementations of distributed kv-stores satisfy certain consistency models. Kv-stores and views provide a faithful abstraction of geo-replicated and partitioned databases, and execution tests provide a powerful abstraction of the synchronisation mechanisms enforced by these databases when committing a transaction. This then allows us to use our formalism to verify the correctness of distributed database protocols. To demonstrate this, we show that the COPS protocol [Lloyd et al. 2011] for implementing a replicated database satisfies causal consistency (§6.2), and the Clock-SI protocol [Du et al. 2013] for implementing a partitioned database satisfies snapshot isolation (§H.2).

3 OPERATIONAL MODEL

We define an interleaving operational semantics for atomic transactions over global, centralised kv-stores and partial client views.

3.1 Key-Value Stores and Client Views

Our global, centralised key-value stores (kv-store) and partial client views provide the abstract machine states for our operational semantics. A kv-store comprises key-indexed lists of versions which record the history of the key with values and meta-data of the transactions that accessed it: the writer and readers.

We assume a countably infinite set of *client identifiers*[3], CLIENT $\ni cl$. The set of *transaction identifiers*, TRANSID $\ni t$, is defined by TRANSID $\triangleq \{t_0\} \uplus \{t_{cl}^n \mid cl \in \text{CLIENT} \wedge n \geq 0\}$, where t_0 denotes the *initialisation transaction* and t_{cl}^n identifies a transaction committed by client cl with n determining the client session order: that is, SO $\triangleq \{(t, t') \mid \exists cl, n, m. \ t = t_{cl}^n \wedge t' = t_{cl}^m \wedge n < m\}$. Subsets of TRANSID are ranged over by $\mathcal{T}, \mathcal{T}', \cdots$. We let TRANSID$_0 \triangleq$ TRANSID $\setminus \{t_0\}$.

Definition 3.1 (Kv-stores). Assume a countably infinite set of *keys*, KEY $\ni k$, and a countably infinite set of *values*, VAL $\ni v$, which includes the keys and an *initialisation value* v_0. The set of *versions*, VERSION $\ni \nu$, is defined by VERSION $\triangleq$ VAL$\times$TRANSID$\times\mathcal{P}$ (TRANSID$_0$). A *kv-store* is a function $\mathcal{K} :$ KEY $\rightarrow$ List(VERSION), where List(VERSION) $\ni \mathcal{V}$ is the set of lists of versions.

Each version has the form $\nu=(v, t, \mathcal{T})$, where v is a value, the *writer* t identifies the transaction that wrote v, and the *reader set* $\mathcal{T}$ identifies the transactions that read v. We use the notation val(ν), w(ν) and rs(ν) to project the individual components of ν. Given a kv-store $\mathcal{K}$ and a transaction t, we write $t \in \mathcal{K}$ if t is either the writer or one of the readers of a version included in $\mathcal{K}$, $|\mathcal{K}(k)|$ for the length of the version list $\mathcal{K}(k)$, and write $\mathcal{K}(k, i)$ for the ith version of k, with $0 \leq i < |\mathcal{K}(k)|$.

We focus on kv-stores whose consistency model satisfies the *snapshot property*, ensuring that a transaction reads and writes at most one version for each key. This is a normal assumption for distributed databases, e.g. in [Ardekani et al. 2014; Bailis et al. 2014; Binnig et al. 2014; Du et al. 2013; Li et al. 2012; Liu et al. 2018; Lloyd et al. 2011; Saeida Ardekani et al. 2013; Sovran et al. 2011; Spirovska et al. 2018]. We also assume that the version list for each key has an initial version carrying the initialisation value v_0, written by the initialisation transaction t_0 with an initial empty reader set. Finally, we assume that the kv-store agrees with the session order of clients: a client cannot read a version of a key that has been written by a future transaction within the same session; and the order in which versions are written by a client must agree with its session order. A kv-store

---

[3] We use the notation A $\ni a$ to denote that elements of A are ranged over by a and its variants such as a', a_1, $\cdots$.

is *well-formed* if it satisfies these three assumptions, defined formally in Def. A.1. Henceforth, we assume kv-stores are well-formed, and write KVS to denote the set of well-formed kv-stores.

A global kv-store provides an abstract centralised description of updates associated with distributed kv-stores that is *complete* in that no update has been lost in the description. By contrast, in both replicated and partitioned distributed databases, a client may have incomplete information about updates distributed between machines. We model this incomplete information by defining a *view* of the kv-store which provides a *partial* record of the updates observed by a client. We require that a client view be *atomic* in that it can see either all or none of the updates of a transaction.

Definition 3.2 (Views). A *view* of a kv-store $\mathcal{K} \in$ KVS is a function $u \in \text{Views}(\mathcal{K}) \triangleq \text{Key} \rightarrow \mathcal{P}(\mathbb{N})$ such that, for all i, i', k, k':

$$0 \in u(k) \land (i \in u(k) \Rightarrow 0 \leq i < |\mathcal{K}(k)|) \qquad \text{(well-formed)}$$

$$i \in u(k) \land \text{w}(\mathcal{K}(k, i)) = \text{w}(\mathcal{K}(k', i')) \Rightarrow i' \in u(k') \qquad \text{(atomic)}$$

Given two views $u, u' \in \text{Views}(\mathcal{K})$, the order between them is defined by $u \sqsubseteq u' \overset{\text{def}}{\Leftrightarrow} \forall k \in$ dom$(\mathcal{K})$. $u(k) \subseteq u'(k)$. The set of views is $\text{Views} \triangleq \bigcup_{\mathcal{K} \in \text{KVS}} \text{Views}(\mathcal{K})$. The *initial view*, u_0, is defined by $u_0(k) = \{0\}$ for every $k \in$ Key.

Our operational semantics updates *configurations*, which are pairs comprising a kv-store and a function describing the views of a finite set of clients.

Definition 3.3 (Configurations). A *configuration*, $\Gamma \in$ Conf, is a pair $(\mathcal{K}, \mathcal{U})$ with $\mathcal{K} \in$ KVS and $\mathcal{U} : \text{Client} \xrightarrow{\text{fin}} \text{Views}(\mathcal{K})$. The set of *initial configurations*, $\text{Conf}_0 \subseteq$ Conf, contains configurations of the form $(\mathcal{K}_0, \mathcal{U}_0)$, where $\mathcal{K}_0$ is the initial kv-store defined by $\mathcal{K}_0(k) \triangleq (v_0, t_0, \emptyset)$ for all $k \in$ Key.

Given a configuration $(\mathcal{K}, \mathcal{U})$ and a client cl, if $u = \mathcal{U}(cl)$ is defined then, for each k, the configuration determines the sub-list of versions in $\mathcal{K}$ that cl sees. If $i, j \in u(k)$ and $i < j$, then cl sees the values carried by versions $\mathcal{K}(k, i)$ and $\mathcal{K}(k, j)$, and it also sees that the version $\mathcal{K}(k, j)$ is more up-to-date than $\mathcal{K}(k, i)$. It is therefore possible to associate a *snapshot* with the view u, which identifies, for each key k, the last version included in the view. This definition assumes that the database satisfies the *last-write-wins* resolution policy, employed by many distributed kv-stores. However, our formalism can be adapted straightforwardly to capture other resolution policies.

Definition 3.4 (View Snapshots). Given $\mathcal{K} \in$ KVS and $u \in \text{Views}(\mathcal{K})$, the *snapshot* of u in $\mathcal{K}$ is a function, snapshot$(\mathcal{K}, u) : \text{Key} \rightarrow \text{Val}$, defined by snapshot$(\mathcal{K}, u) \triangleq \lambda k.\, \text{val}(\mathcal{K}(k, \max_<(u(k))))$, where $\max_<(u(k))$ is the maximum element in $u(k)$ with respect to the natural order $<$ over $\mathbb{N}$.

3.2 Operational Semantics

Programming Language. A *program* P comprises a finite number of clients, where each client is associated with a unique identifier $cl \in$ Client, and executes a sequential *command* C, given by the following grammar:

$$\text{C} ::= \text{skip} \,|\, \text{C}_p \,|\, [\text{T}] \,|\, \text{C}; \text{C} \,|\, \text{C} + \text{C} \,|\, \text{C}^* \qquad \text{C}_p ::= \text{x} := \text{E} \,|\, \text{assume(E)}$$

$$\text{T} ::= \text{skip} \,|\, \text{T}_p \,|\, \text{T}; \text{T} \,|\, \text{T} + \text{T} \,|\, \text{T}^* \qquad \text{T}_p ::= \text{C}_p \,|\, \text{x} := [\text{E}] \,|\, [\text{E}] := \text{E}$$

Sequential commands C comprise skip, primitive commands C_p, atomic transactions [T], and standard compound constructs: sequential composition (;), non-deterministic choice (+) and iteration (∗). Primitive commands include variable assignment x := E and assume statements assume(E) which can be used to encode conditionals. They are used for computations based on client-local

variables and can hence be invoked without restriction. Transactional commands T comprise skip, primitive transactional commands T_p, and the standard compound constructs. Primitive transactional commands comprise primitive commands, lookup x := [E] and mutation [E] := E used for reading and writing to kv-stores respectively, which can only be invoked as part of an atomic transaction.

A *program* is a finite partial function from client identifiers to sequential commands. For clarity, we often write $C_1 \parallel \dots \parallel C_n$ as syntactic sugar for a program P with n clients associated with identifiers $cl_1 \dots cl_n$, where each client cl_i executes C_i. Each client cl_i is associated with its own client-local *stack*, $s_i \in \text{STACK} \triangleq \text{VARS} \rightarrow \text{VAL}$, mapping program variables (ranged over by x, y, $\cdots$) to values. We assume a language of expressions built from values v and program variables x: E ::= $v \mid x \mid E + E \mid \cdots$. The *evaluation* $[\![E]\!]_s$ of expression E is parametric in the client-local stack s:

$$[\![v]\!]_s \triangleq v \quad [\![x]\!]_s \triangleq s(x) \quad [\![E_1 + E_2]\!]_s \triangleq [\![E_1]\!]_s + [\![E_2]\!]_s \quad \dots$$

Transactional Semantics. In our operational semantics, transactions are executed *atomically*. It is still possible for an underlying implementation, such as COPS, to update the distributed kv-store while the transaction is in progress. It just means that, given the abstractions captured by our global kv-stores and partial views, such an update is modelled as an instantaneous atomic update. Intuitively, given a configuration $\Gamma = (\mathcal{K}, \mathcal{U})$, when a client cl executes a transaction [T], it performs the following steps: (1) it constructs an initial *snapshot ss* of $\mathcal{K}$ using its view $\mathcal{U}(cl)$ as defined in Def. 3.4; (2) it executes T in isolation over *ss* accumulating the effects (the reads and writes) of executing T; and (3) it commits T by incorporating these effects into $\mathcal{K}$.

Definition 3.5 (Transactional Snapshots). A *transactional snapshot, ss* $\in \text{SNAPSHOT} \triangleq \text{KEY} \rightarrow \text{VAL}$, is a function from keys to values. When the meaning is clear, it is just called a *snapshot*.

The rules for transactional commands will be defined using an arbitrary transactional snapshot. The rules for sequential commands and programs will be defined using a transactional snapshot given by a view snapshot. To capture the effects of executing a transaction T on a snapshot *ss* of kv-store $\mathcal{K}$, we identify a *fingerprint* of T on *ss* which captures the values T reads from *ss*, and the values T writes to *ss* and intends to commit to $\mathcal{K}$. Execution of a transaction in a given configuration may result in more than one fingerprint due to non-determinism (non-deterministic choice).

Definition 3.6 (Fingerprints). Let Ops denote the set of read (r) and write (w) *operations* defined by Ops $\triangleq \{(l, k, v) \mid l \in \{r, w\} \land k \in \text{KEY} \land v \in \text{VAL}\}$. A *fingerprint* $\mathcal{F}$ is a set of operations, $\mathcal{F} \subseteq \text{Ops}$, such that: $\forall k \in \text{KEY}, l \in \{r, w\} . (l, k, v_1), (l, k, v_2) \in \mathcal{F} \Rightarrow v_1 = v_2$.

According to Def. 3.6, a fingerprint contains at most one read operation and at most one write operation. This reflects our assumption regarding transactions that satisfy the snapshot property: reads are taken from a single snapshot of the kv-store; and only the last write of a transaction to each key is committed to the kv-store.

Fig. 1 presents the rules for transactional commands. The only non-standard rule is TPRIMITIVE, which updates the snapshot and the fingerprint of a transaction: the premise $(s, ss) \xrightarrow{T_p} (s', ss')$ describes how executing T_p affects the local state (the client stack and the snapshot) of a transaction; and the premise $o = \text{op}(s, ss, T_p)$ identifies the operation on the kv-store associated with T_p.

TPrimitive

$$\frac{(s, ss) \overset{\mathsf{T}_p}{\rightsquigarrow} (s', ss') \qquad o = \mathsf{op}(s, ss, \mathsf{T}_p)}{(s, ss, \mathcal{F}), \mathsf{T}_p \to (s', ss', \mathcal{F} \lll o), \mathsf{skip}}$$

$$\mathcal{F} \lll (\mathsf{r}, k, v) \triangleq \begin{cases} \mathcal{F} \cup \{(\mathsf{r}, k, v)\} & \text{if } \forall l, v'. \ (l, k, v') \notin \mathcal{F} \\ \mathcal{F} & \text{otherwise} \end{cases}$$

$$\mathcal{F} \lll (\mathsf{w}, k, v) \triangleq (\mathcal{F} \setminus \{(\mathsf{w}, k, v') \mid v' \in \mathrm{Val}\}) \cup \{(\mathsf{w}, k, v)\}$$

$$\mathcal{F} \lll \epsilon \triangleq \mathcal{F}$$

TChoice

$$\frac{i \in \{1, 2\}}{(s, ss, \mathcal{F}), \mathsf{T}_1 + \mathsf{T}_2 \to (s, ss, \mathcal{F}), \mathsf{T}_i}$$

TIter

$$\frac{}{(s, ss, \mathcal{F}), \mathsf{T}^* \to (s, ss, \mathcal{F}), \mathsf{skip} + (\mathsf{T}; \mathsf{T}^*)}$$

TSeqSkip

$$\frac{}{(s, ss, \mathcal{F}), \mathsf{skip}; \mathsf{T} \to (s, ss, \mathcal{F}), \mathsf{T}}$$

TSeq

$$\frac{(s, ss, \mathcal{F}), \mathsf{T}_1 \to (s', ss', \mathcal{F}'), \mathsf{T}_1'}{(s, ss, \mathcal{F}), \mathsf{T}_1; \mathsf{T}_2 \to (s', ss', \mathcal{F}'), \mathsf{T}_1'; \mathsf{T}_2}$$

Fig. 1. Rules for transactional commands.

Definition 3.7. The relation $\overset{\mathsf{T}_p}{\rightsquigarrow} \subseteq (\mathrm{Stack} \times \mathrm{Snapshot}) \times (\mathrm{Stack} \times \mathrm{Snapshot})$ is defined by[4] :

$$(s, ss) \overset{\mathsf{x} := \mathsf{E}}{\rightsquigarrow} (s[\mathsf{x} \mapsto [\![\mathsf{E}]\!]_s], ss) \qquad (s, ss) \overset{\mathsf{assume}(\mathsf{E})}{\rightsquigarrow} (s, ss) \text{ where } [\![\mathsf{E}]\!]_s \neq 0$$

$$(s, ss) \overset{\mathsf{x} := [\mathsf{E}]}{\rightsquigarrow} (s[\mathsf{x} \mapsto ss([\![\mathsf{E}]\!]_s)], ss) \qquad (s, ss) \overset{[\mathsf{E}_1] := \mathsf{E}_2}{\rightsquigarrow} (s, ss[[\![\mathsf{E}_1]\!]_s \mapsto [\![\mathsf{E}_2]\!]_s])$$

The function op, computing the fingerprint of primitive transactional commands, is defined by:

$$\mathsf{op}(s, ss, \mathsf{x} := \mathsf{E}) \triangleq \epsilon \qquad\qquad \mathsf{op}(s, ss, \mathsf{assume}(\mathsf{E})) \triangleq \epsilon$$

$$\mathsf{op}(s, ss, \mathsf{x} := [\mathsf{E}]) \triangleq (\mathsf{r}, [\![\mathsf{E}]\!]_s, ss([\![\mathsf{E}]\!]_s)) \qquad \mathsf{op}(s, ss, [\mathsf{E}_1] := \mathsf{E}_2) \triangleq (\mathsf{w}, [\![\mathsf{E}_1]\!]_s, [\![\mathsf{E}_2]\!]_s)$$

The empty operation ϵ is used for those primitive commands that do not contribute to the fingerprint.

The conclusion of the TPrimitive rule uses the *combination operator* $\lll : \mathcal{P}(\mathrm{Ops}) \times (\mathrm{Ops} \uplus \{\epsilon\}) \to \mathcal{P}(\mathrm{Ops})$, defined in Fig. 1, to extend the fingerprint $\mathcal{F}$ accumulated with operation o associated with T_p, as appropriate: it adds a read from k if $\mathcal{F}$ does not already contain an entry for k; and it always updates the write for k to $\mathcal{F}$, removing previous writes to k.

Definition 3.8 (Fingerprint Set). Given a client stack s and a snapshot ss, the *fingerprint set* of a transaction T is: $\mathrm{Fp}(\mathsf{T}) \triangleq \{\mathcal{F} \mid \exists s, s', ss, ss'. \ (s, ss, \emptyset), \mathsf{T} \to^* (s', ss', \mathcal{F}), \mathsf{skip}\}$ where $\to^*$ denotes the reflexive, transitive closure of $\to$ given in Fig. 1. A set $\mathcal{F} \in \mathrm{Fp}(\mathsf{T})$ is called a *final fingerprint* of T.

Operational Semantics. We give the operational semantics of commands and programs in Fig. 2. The command semantics describes transitions of the form $cl \vdash (\mathcal{K}, u, s), \mathsf{C} \overset{\lambda}{\to}_{\mathrm{ET}} (\mathcal{K}', u', s'), \mathsf{C}'$, stating that given the kv-store $\mathcal{K}$, view u and stack s, a client cl may execute command C for one step, updating the kv-store to $\mathcal{K}'$, the stack to s', and the command to its continuation C'. The label λ is either of the form (cl, ι) denoting that cl executed a primitive command that required no access to $\mathcal{K}$, or $(cl, u'', \mathcal{F})$ denoting that cl committed an atomic transaction with final fingerprint $\mathcal{F}$ under the view u''. The semantics is parametric in the choice of *execution test* ET for kv-stores, which is used to generate the *consistency model* on kv-stores under which a transaction can execute. In §4, we give many examples of execution tests for well-known consistency models. In §5 and §F, we prove that our execution tests (using kv-stores) generate consistency models that are equivalent to existing definitions of consistency models (using execution graphs).

The rules for the compound constructs are straightforward and given in §A. The rule for primitive commands, CPrimitive, depends on the transition system $\overset{\mathsf{C}_p}{\rightsquigarrow} \subseteq \mathrm{Stack} \times \mathrm{Stack}$ which simply

---

[4] For any function f, the new function f[d ↦ r] is defined by f[d ↦ r](d) = r, and f[d ↦ r](d') = f(d') if d' ≠ d.

CPrimitive

$$\frac{s \overset{C_p}{\leadsto} s'}{cl \vdash (\mathcal{K}, u, s), C_p \xrightarrow{(cl,\iota)}_{ET} (\mathcal{K}, u, s'), \text{skip}}$$

$$s \xrightarrow{\text{x:=E}} s[\text{x} \mapsto [\![E]\!]_s]$$

$$s \overset{\text{assume(E)}}{\leadsto} s \text{ where } [\![E]\!]_s \neq 0$$

CAtomicTrans

$$u \sqsubseteq u''$$

$$\frac{\begin{array}{cccc} ss = \text{snapshot}(\mathcal{K}, u'') & (s, ss, \emptyset), T \rightarrow^* (s', \_, \mathcal{F}), \text{skip} & \text{can-commit}_{ET}(\mathcal{K}, u'', \mathcal{F}) \\ t \in \text{nextTid}(cl, \mathcal{K}) & \mathcal{K}' = \text{update}(\mathcal{K}, u'', \mathcal{F}, t) & \text{vshift}_{ET}(\mathcal{K}, u'', \mathcal{K}', u') \end{array}}{cl \vdash (\mathcal{K}, u, s), [T] \xrightarrow{(cl, u'', \mathcal{F})}_{ET} (\mathcal{K}', u', s'), \text{skip}}$$

PProg

$$\frac{u = \mathcal{U}(cl) \qquad s = \mathcal{E}(cl) \qquad C = P(cl) \qquad cl \vdash (\mathcal{K}, u, s), C \xrightarrow{\lambda}_{ET} (\mathcal{K}', u', s'), C'}{\vdash (\mathcal{K}, \mathcal{U}, \mathcal{E}), P \xrightarrow{\lambda}_{ET} (\mathcal{K}', \mathcal{U}[cl \mapsto u'], \mathcal{E}[cl \mapsto s']), P[cl \mapsto C'])}$$

Fig. 2. Rules for sequential commands and programs.

describes how the primitive command C_p affects the local state of a client. The rule CAtomicTrans describes the execution of an atomic transaction under the execution test ET.

We describe the CAtomicTrans rule in detail. The first premise states that the current view u of the executing command may be advanced to a newer view u'' (see Def. 3.2). Given the new view u'', the transaction obtains a snapshot ss of the kv-store $\mathcal{K}$, and executes T locally to completion (skip), updating the stack to s', while accumulating the fingerprint $\mathcal{F}$; this behaviour is modelled in the second and third premises of CAtomicTrans. Note that the resulting snapshot is ignored as the effect of the transaction is recorded in the fingerprint $\mathcal{F}$. The can-commit$_{ET}(\mathcal{K}, u'', \mathcal{F})$ premise ensures that under the execution test ET, the final fingerprint $\mathcal{F}$ of the transaction is compatible with the (original) kv-store $\mathcal{K}$ and the client view u'', and thus the transaction *can commit*. Note that the can-commit check is parametric in the execution test ET. This is because the conditions checked upon committing depend on the consistency model under which the transaction is to commit. In §4, we define can-commit for several execution tests associated with well-known consistency models.

Now we are ready for client cl to commit the transaction resulting in the kv-store $\mathcal{K}'$ with the client view u'' *shifting* to a new view u': pick a fresh transaction identifier $t \in \text{nextTid}(cl, \mathcal{K})$; compute the new kv-store via $\mathcal{K}' = \text{update}(\mathcal{K}, u'', \mathcal{F}, t)$; and check if the *view shift* is permitted under execution test ET using vshift$_{ET}(\mathcal{K}, u'', \mathcal{K}', u')$. Observe that as with can-commit, the vshift check is parametric in the execution test ET. Once again this is because the conditions checked for shifting the client view depend on the consistency model. In §4 we define vshift for several execution tests associated with well-known consistency models. The set nextTid$(cl, \mathcal{K})$ is defined by nextTid$(cl, \mathcal{K}) \triangleq \{ t_{cl}^n \mid \forall m. \ t_{cl}^m \in \mathcal{K} \Rightarrow m < n \}$. The function update$(\mathcal{K}, u, \mathcal{F}, t)$ describes how the fingerprint $\mathcal{F}$ of transaction t executed under view u updates kv-store $\mathcal{K}$: for each read $(r, k, v) \in \mathcal{F}$, it adds t to the reader set of the last version of k in u; for each write (w, k, v), it appends a new version $(v, t, \emptyset)$ to $\mathcal{K}(k)$.

Definition 3.9 (Transactional update). The function update$(\mathcal{K}, u, \mathcal{F}, t)$, is defined by:

$$\begin{aligned} \text{update}(\mathcal{K}, u, \emptyset, t) &\triangleq \mathcal{K} \\ \text{update}(\mathcal{K}, u, \{(r, k, v)\} \uplus \mathcal{F}, t) &\triangleq \text{let } i = \max_{<}(u(k)), (v, t', \mathcal{T}) = \mathcal{K}(k, i) \text{ in} \\ &\qquad \text{update}(\mathcal{K}[k \mapsto \mathcal{K}(k)[i \mapsto (v, t', \mathcal{T} \uplus \{t\})]], u, \mathcal{F}, t) \\ \text{update}(\mathcal{K}, u, \{(w, k, v)\} \uplus \mathcal{F}, t) &\triangleq \text{let } \mathcal{K}' = \mathcal{K}[k \mapsto \mathcal{K}(k)::(v, t, \emptyset)] \text{ in } \text{update}(\mathcal{K}', u, \mathcal{F}, t) \end{aligned}$$

where, given a list of versions $\mathcal{V} = v_0::\cdots::v_n$ and an index $i : 0 \leq i \leq n$, then $\mathcal{V}[i \mapsto v] \triangleq v_0::\cdots::v_{i-1}::v::v_{i+1}\cdots v_n$,

The last rule, PProg, captures the execution of a program step given a *client environment* $\mathcal{E} \in$ CEnv. A client environment $\mathcal{E}$ is a function from client identifiers to variable stacks, associating each client with its stack. We assume that the domain of client environments contains the domain of the program throughout the execution: $\mathrm{dom}(\mathsf{P}) \subseteq \mathrm{dom}(\mathcal{E})$. Program transitions are simply defined in terms of the transitions of their constituent client commands. This yields an interleaving semantics for transactions of different clients: a client executes a transaction in an atomic step without interference from the other clients.

4 CONSISTENCY MODELS: KV-STORES

We define what it means for a kv-store to be in a consistent state. Many different consistency models for distributed databases have been proposed in the literature [Berenson et al. 1995; Burckhardt 2014; Li et al. 2012; Liu et al. 2018; Sovran et al. 2011], capturing different trade-offs between performance and application correctness: examples range from *serialisability*, a strong consistency model which only allows kv-stores obtained from a serial execution of transactions with inevitable performance drawbacks, to *eventual consistency*, a weak consistency model which imposes few conditions on the structure of a kv-store leading to good performance but anomalous behaviour. We define consistency models for our kv-stores, by introducing the notion of *execution test* which specifies whether a client is allowed to commit a transaction in a given kv-store. Each execution test induces a consistency model as the set of kv-stores obtained by having clients non-deterministically commit transactions so long as the constraints imposed by the execution test are satisfied. We explore a range of execution tests associated with well-known consistency models in the literature. In §5, we demonstrate that our operational formulation of consistency models over kv-stores using execution tests are equivalent to the established declarative definitions of consistency models over abstract executions [Burckhardt et al. 2012; Cerone et al. 2015a] and dependency graphs [Adya 1999].

Definition 4.1 (Execution tests). An *execution test* is a set of tuples $\mathsf{ET} \subseteq \mathrm{KVS} \times \mathrm{Views} \times \mathrm{Fp} \times \mathrm{KVS} \times \mathrm{Views}$ such that, for all $(\mathcal{K}, u, \mathcal{F}, \mathcal{K}', u') \in \mathsf{ET}$: (1) $u \in \mathrm{Views}(\mathcal{K})$ and $u' \in \mathrm{Views}(\mathcal{K}')$; (2) can-commit$_{\mathsf{ET}}(\mathcal{K}, u, \mathcal{F})$ and (4) for all $k \in \mathcal{K}$ and $v \in \mathrm{Val}$, if $(\mathsf{r}, k, v) \in \mathcal{F}$ then $\mathcal{K}(k, \max_{<}(u(k)))=v$.

Intuitively, $(\mathcal{K}, u, \mathcal{F}, \mathcal{K}', u') \in \mathsf{ET}$ means that, under the execution test ET, a client with initial view u over a kv-store $\mathcal{K}$ can commit a transaction with fingerprint $\mathcal{F}$ to obtain the resulting kv-store $\mathcal{K}'$ (Def. 3.9) while shifting its view to u'. We adopt the notation $\mathsf{ET} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$ to capture this intuition. Note that the last condition in Def. 4.1 enforces the last-write-wins policy [Vogels 2009]: a transaction always reads the most recent writes from the initial view u.

The largest execution test is denoted by $\mathsf{ET}_\top$, where for all $\mathcal{K}, \mathcal{K}', u, u, \mathcal{F}$:

$$\mathrm{can\text{-}commit}_{\mathsf{ET}_\top}(\mathcal{K}, u, \mathcal{F}) \overset{\mathrm{def}}{\Leftrightarrow} \mathrm{true} \qquad \text{and} \qquad \mathrm{vshift}_{\mathsf{ET}_\top}(\mathcal{K}, u, \mathcal{K}', u') \overset{\mathrm{def}}{\Leftrightarrow} \mathrm{true}$$

In §5, we show that the consistency model induced by $\mathsf{ET}_\top$ corresponds to the *Read Atomic* model [Bailis et al. 2014], a variant of *Eventual Consistency* [Burckhardt et al. 2012] for atomic transactions.

In §4.1, we give many examples of execution tests. Here, we explain how an execution test induces a consistency model. Given an execution test ET, we define a ET-reduction as a labelled transition which captures how client *cl* updates a configuration: either *cl* shifts its view to a more up-to-date one with the label ε denoting that there was no access to the kv-store; or *cl* commits a transaction with the label $\mathcal{F}$ denoting the fingerprint of the committed transaction.

ET	can-commit$_{\text{ET}}(\mathcal{K}, u, \mathcal{F})$	Closure Relation (where applicable)	vshift$_{\text{ET}}(\mathcal{K}, u, \mathcal{K}', u')$
MR	true		$u \sqsubseteq u'$
RYW	true		$\forall t \in \mathcal{K}' \setminus \mathcal{K}. \ \forall k, i.$ $\text{w}(\mathcal{K}'(k, i)) \xrightarrow{\text{SO}^?} t \Rightarrow i \in u'(k)$
CC	closed$(\mathcal{K}, u, R_{\text{CC}})$	$R_{\text{CC}} \triangleq \text{SO} \cup \text{WR}_{\mathcal{K}}$	vshift$_{\text{MR} \cap \text{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
UA	closed$(\mathcal{K}, u, R_{\text{UA}})$	$R_{\text{UA}} \triangleq \bigcup_{(\text{w}, k, \_) \in \mathcal{F}} \text{WW}_{\mathcal{K}}^{-1}(k)$	true
PSI	closed$(\mathcal{K}, u, R_{\text{PSI}})$	$R_{\text{PSI}} \triangleq R_{\text{UA}} \cup R_{\text{CC}} \cup \text{WW}_{\mathcal{K}}$	vshift$_{\text{MR} \cap \text{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
CP	closed$(\mathcal{K}, u, R_{\text{CP}})$	$R_{\text{CP}} \triangleq \text{SO}; \text{RW}_{\mathcal{K}}^? \cup \text{WR}_{\mathcal{K}}; \text{RW}_{\mathcal{K}}^? \cup \text{WW}_{\mathcal{K}}$	vshift$_{\text{MR} \cap \text{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
SI	closed$(\mathcal{K}, u, R_{\text{SI}})$	$R_{\text{SI}} \triangleq R_{\text{UA}} \cup R_{\text{CP}} \cup (\text{WW}_{\mathcal{K}}; \text{RW}_{\mathcal{K}})$	vshift$_{\text{MR} \cap \text{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
SER	closed$(\mathcal{K}, u, R_{\text{SER}})$	$R_{\text{SER}} \triangleq \text{WW}^{-1}$	true

Fig. 3. Execution tests of well-known consistency models, where SER* denotes an alternative equivalent SER specification and SO is as given in §3.1.

Definition 4.2 (ET-reduction). An ET-*reduction*, $(\mathcal{K}, \mathcal{U}) \xrightarrow{(cl, \alpha)}_{\text{ET}} (\mathcal{K}', \mathcal{U}')$, is defined by:

(1) either $\alpha = \varepsilon$, $\mathcal{K}' = \mathcal{K}$ and $\mathcal{U}' = \mathcal{U}[cl \mapsto u]$ for some u such that $\mathcal{U}(cl) \sqsubseteq u$; or

(2) $\alpha = \mathcal{F}$ for some $\mathcal{F}$, and ET $\vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$, where $\mathcal{K}' = \text{update}(\mathcal{K}, u, \mathcal{F}, t)$ for some $t \in \text{nextTid}(cl, \mathcal{K})$, $\mathcal{U}' = \mathcal{U}[cl \mapsto u']$.

A finite sequence of ET-reductions starting in an initial configuration Γ_0 is called an ET-*trace*.

Each ET-trace starting with an initial configuration (Def. 3.3) terminates in a configuration $(\mathcal{K}, \_)$ where $\mathcal{K}$ is obtained as a result of several clients committing transactions under the execution test ET. The consistency model induced by ET, written CM(ET), is the set of all such terminal kv-stores.

Definition 4.3 (Consistency Model). The *consistency model* induced by an execution test ET is defined as CM(ET) $\triangleq \big\{ \mathcal{K} \mid \exists \Gamma_0 \in \text{CONF}_0. \ \Gamma_0 \dashrightarrow^*_{\text{ET}} (\mathcal{K}, \_) \big\}$.

Note that in the definition of ET-traces, the view-shifts and transaction commits are decoupled. This is in contrast to the operational semantics (§3, Fig. 2), where view-shifts and transaction commits are combined in a single transition of programs (CAtomicTrans). The reason for this mismatch is best understood when looking at the intended applications. ET-traces are useful for proving that a distributed transactional protocol implements a given consistency model: in this case, it is convenient to separate shifting a view from committing a transaction, as these two steps often take place separately in distributed protocols. The operational semantics is particularly useful for reasoning about transactional programs: in this case, the treatment of the view-shifts and transaction commits as a single transition reduces the number of interleavings in programs. The ET-traces and operational semantics are equally expressive as the following theorem states.

THEOREM 4.4. *Let* $[\![P]\!]_{\text{ET}}$ *be the set of kv-stores reachable by executing* P *under the execution test* ET. *Then for all* ET, CM(ET) $= \bigcup_P [\![P]\!]_{\text{ET}}$.

4.1 Example Execution Tests

We give several examples of execution tests which give rise to consistency models on kv-stores. Recall that the snapshot property and the last write wins are hard-wired into our model. This means that we can only define consistency models that satisfy these two constraints. Although this forbids us to express interesting consistency models such as *Read Committed*, we are able to express a large variety of consistency models employed by distributed kv-stores. We proceed with a summary of our notational conventions.

Notation. Given relations $r, r' \subseteq A \times A$, we write: $r^?$, r^+ and r^* for its reflexive, transitive and reflexive-transitive closures of r, respectively; r^{-1} for its inverse; $a_1 \xrightarrow{r} a_2$ for $(a_1, a_2) \in r$; and $r; r'$ for $\{(a_1, a_2) \mid \exists a. \ (a_1, a) \in r \land (a, a_2) \in r'\}$.

Recall from Def. 4.1 that an execution test ET comprises tuples of the form $(\mathcal{K}, u, \mathcal{F}, \mathcal{K}', u')$ where can-commit$_{\mathsf{ET}}(\mathcal{K}, u, \mathcal{F})$ and vshift$_{\mathsf{ET}}(\mathcal{K}, u, \mathcal{K}', u')$. We define can-commit and vshift for several consistency models, using a couple of auxiliary definitions.

Prefix Closure. Given a kv-store $\mathcal{K}$ and a view u, the *set of visible transactions* is:

$$\mathsf{visTx}(\mathcal{K}, u) \triangleq \{\mathsf{w}(\mathcal{K}(k, i)) \mid i \in u(k)\}$$

Given a binary relation on transactions, $R \subseteq \textsc{TransID} \times \textsc{TransID}$, we say that a view u is closed with respect to a kv-store $\mathcal{K}$ and R, written $\mathsf{closed}(\mathcal{K}, u, R)$, iff:

$$\mathsf{closed}(\mathcal{K}, u, R) \overset{\text{def}}{\Leftrightarrow} \mathsf{visTx}(\mathcal{K}, u) = \left((R^*)^{-1}\mathsf{visTx}(\mathcal{K}, u)\right) \setminus \{t \mid \forall k, i. \ t \neq \mathsf{w}(\mathcal{K}(k, i))\}$$

That is, if transaction t is visible in u, then all transactions that are R^*-before t are also visible in u.

Dependency Relations. We define transaction dependency relations for kv-stores, inspired by analogous relations for dependency graphs due to Adya [1999]. Given a kv-store $\mathcal{K}$, a key k and indexes i, j such that $0 \leq i < j < |\mathcal{K}(k)|$, if there exists $t_i, \mathcal{T}_i, t$ such that $\mathcal{K}(k, i) = (\_, t_i, \mathcal{T}_i)$, $\mathcal{K}(k, j) = (\_, t_j, \_)$ and $t \in \mathcal{T}_i$, then we say that there is:

(1) a *Write-Read* dependency over k from t_i to t, written $(t_i, t) \in \mathsf{WR}_{\mathcal{K}}(k)$;
(2) a *Write-Write* dependency over k from t_i to t_j, written $(t_i, t_j) \in \mathsf{WW}_{\mathcal{K}}(k)$; and
(3) a *Read-Write* anti-dependency from t to t_j, provided that $t \neq t_j$, written $(t, t_j) \in \mathsf{RW}_{\mathcal{K}}(k)$.

Fig. 3a illustrates an example kv-store and its transaction dependency relations.

We give several definitions of execution tests using vshift and can-commit in Fig. 3. In §5, we demonstrate that the associated consistency models on kv-stores correspond to well-known consistency models on execution graphs. We anticipate these results, by labelling the execution test with their well-known consistency models.

4.1.1 Monotonic Reads (MR). This consistency model states that when committing, a client cannot lose information in that it can only see increasingly more up-to-date versions from a kv-store. This prevents, for example, the kv-store of Fig. 3b, since client cl first reads the latest version of k in t_{cl}^1, and then reads the older, initial version of k in t_{cl}^2. As such, the vshift$_{\mathsf{MR}}$ predicate in Fig. 3 ensures that clients can only extend their views. When this is the case, clients can then *always* commit their transactions, and thus can-commit$_{\mathsf{MR}}$ is simply defined as true.

4.1.2 Read Your Writes (RYW). This consistency model states that a client must always see all the versions written by the client itself. The vshift$_{\mathsf{RYW}}$ predicate thus states that after executing a transaction, a client contains all the versions it wrote in its view. This ensures that such versions will be included in the view of the client when committing future transactions. Note that under RYW the kv-store in Fig. 3c is prohibited as the initial version of k holds value v_0 and client cl tries to increment the value of k twice. For its first transaction t_{cl}^1, it reads the initial value v_0 and then writes a new version with value v_1. For its second transaction t_{cl}^2, it reads the initial value v_0 again and write a new version with value v_1. The vshift$_{\mathsf{RYW}}$ predicate rules out this example by requiring that the client view, after it commits the transaction t_{cl}^1, includes the version it wrote. When this is the case, clients can always commit their transactions, and thus can-commit$_{\mathsf{RYW}}$ is simply true.
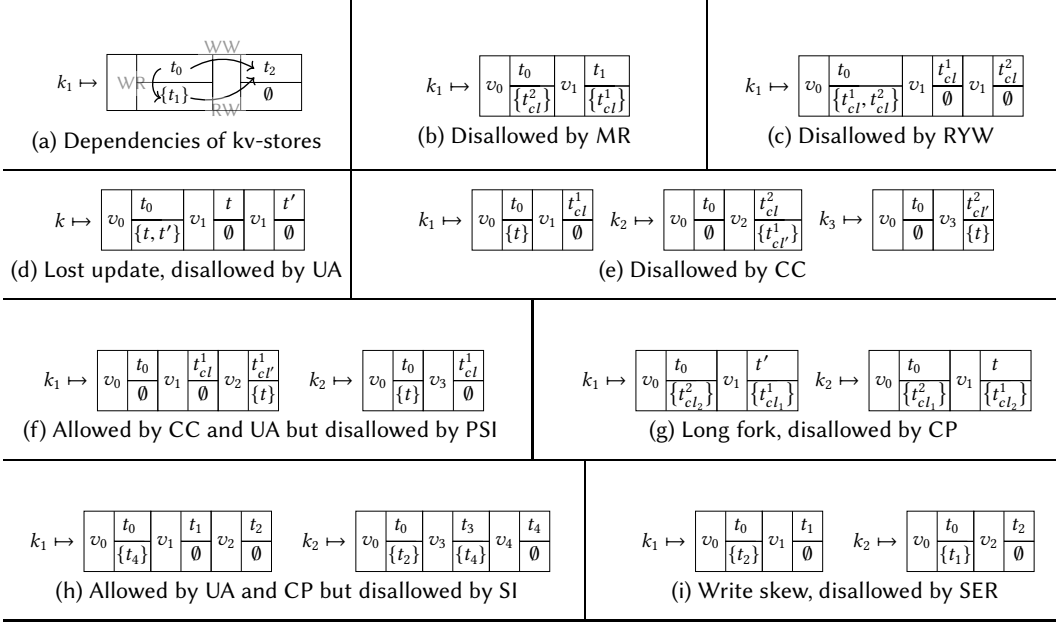
Fig. 3. Behaviours disallowed under different consistency models. Sub-figure 3a shows the dependencies of transactions in kv-stores, where values of versions have been removed for simplicity.

The MR and RYW models together with *monotonic writes* (MW) and *write follows reads* (WFR) models are collectively known as *session guarantees*. Due to space constraints, we omit the definitions associated with MW and WFR, and refer the reader to §A.

4.1.3 Causal Consistency (CC). Causal consistency subsumes the four session guarantees discussed above. As such, the vshift$_{CC}$ predicate is defined as the *conjunction* of their associated vshift predicates. However, as shown in Fig. 3, it is sufficient to define vshift$_{CC}$ as the conjunction of the MR and RYW session guarantees alone, where for brevity we write vshift$_{MR \cap RYW}$ for vshift$_{MR}$ ∧ vshift$_{RYW}$. This is because as we demonstrate in §A, the vshift$_{MW}$ and vshift$_{WFR}$ are defined simply as $\texttt{true}$, allowing us to remove them from vshift$_{CC}$.

Additionally, CC strengthens the session guarantees by requiring that if a client sees a version v prior to committing a transaction, then it must also see the versions on which v depends. If t is the writer of v, then v clearly depends on all versions that t reads. Moreover, if v is, or it depends on, a version v' accessed by a client cl, then it also depends on all versions that were previously read or written by cl. This is captured by the can-commit$_{CC}$ predicate in Fig. 3, defined as $\texttt{closed}(\mathcal{K}, u, R_{CC})$ with $R_{CC} \triangleq SO \cup WR_{\mathcal{K}}$. For example, the kv-store of Fig. 3e is disallowed by CC: the version of key k_3 carrying value v_3 depends on the version of key k_1 carrying value v_1. However, transaction t must have been committed by a client whose view included v_3, but not v_1.

4.1.4 Update Atomic (UA). This consistency model has been proposed by Cerone et al. [2015a] and implemented by Liu et al. [2018]. UA disallows concurrent transactions writing to the same key, a property known as *write-conflict freedom*: when two transactions write to the same key, one must see the version written by the other. Write-conflict freedom is enforced by can-commit$_{UA}$ which allows a client to write to key k only if its view includes all versions of k; i.e. its view is closed with respect to the $WW^{-1}(k)$ relation for all keys k written in the fingerprint $\mathcal{F}$. This prevents the kv-store of Fig. 3d, as t and t' concurrently increment the initial version of k by 1.

As client views must include the initial versions, once t commits a new version v with value v_1 to k, then t' must include v in its view as there is a WW edge from the initial version to v. As such, when t' subsequently increments k, it must read from v, and not the initial version as depicted in Fig. 3d.

4.1.5 Parallel Snapshot Isolation (PSI). This consistency model is defined as the conjunction of the guarantees provided by CC and UA [Cerone et al. 2015a]. As such, the vshift$_\mathsf{PSI}$ predicate is defined as the conjunction of the vshift predicates for CC and UA. However, we cannot simply define can-commit$_\mathsf{PSI}$ as the conjunction of the can-commit predicates for CC and UA. This is for two reasons. First, their conjunction would only mandate that u be closed with respect to R_CC and R_UA *individually*, but *not* with respect to their *union* (recall that closure is defined in terms of the transitive closure of a given relation and thus the closure of R_CC and R_UA is smaller than the closure of $R_\mathsf{CC} \cup R_\mathsf{UA}$). As such, we define can-commit$_\mathsf{PSI}$ as closure with respect to R_PSI which must include $R_\mathsf{CC} \cup R_\mathsf{UA}$. Second, recall that can-commit$_\mathsf{UA}$ requires that a transaction writing to a key k must be able to see all previous versions of k, i.e. all versions of k. That is, when write-conflict freedom is enforced, a version v of k depends on all previous versions of k. This observation leads us to include write-write dependencies (WW$_\mathcal{K}$) in R_PSI. Observe that the kv-store in Fig. 3f shows an example kv-store that satisfies can-commit$_\mathsf{CC}$ $\wedge$ can-commit$_\mathsf{UA}$, but not can-commit$_\mathsf{PSI}$.

4.1.6 Consistent Prefix (CP). If the total order in which transactions commit is known, CP can be described as a strengthening of CC: if a client sees the versions written by a transaction t, then it must also see all versions written by transactions that commit before t. Although kv-stores only provide *partial* information about the transaction commit order via the dependency relations, this is sufficient to formalise *Consistent Prefix* [Cerone et al. 2017].

In practice, we approximate the order in which transactions commit in an ET-trace that terminates in a configuration $(\mathcal{K}, \_)$ via the WR$_\mathcal{K}$, WW$_\mathcal{K}$, RW$_\mathcal{K}$ and SO relations. This approximation is best understood in terms of an idealised implementation of CP on a centralised system, where the snapshot of a transaction is determined at its start point and its effects are made visible to future transactions at its commit point. With respect to this implementation, if $(t, t') \in$ WR, then t must commit before t' starts, and hence before t' commits. Similarly, if $(t, t') \in$ SO, then t commits before t' starts, and thus before t' commits. Recall that $(t'', t') \in$ RW denotes that t'' reads a version that is later overwritten by t'. That is, t'' cannot see the write of t', and thus t'' must starts before t' commits. As such, if t commits before t'' starts $((t, t'') \in$ WR or $(t, t'') \in$ SO), and $(t'', t') \in$ RW, then t must commit before t' commits. In other words, if $(t, t') \in$ WR; RW or $(t, t') \in$ SO; RW, then t commits before t'. Finally, if $(t, t') \in$ WW, then t must commit before t'. We therefore define $R_\mathsf{CP} \triangleq (\mathsf{WR}_\mathcal{K}; \mathsf{RW}_\mathcal{K}^? \cup \mathsf{SO}; \mathsf{RW}_\mathcal{K}^? \cup \mathsf{WW})$, approximating the order in which transactions commit. Cerone et al. [2017] show that the set $(R_\mathsf{CP}^{-1})^+(t)$ contains all transactions that must be observed by t under CP. We define can-commit$_\mathsf{CP}$ by requiring that the client view be closed with respect to R_CP.

Consistent prefix disallows the *long fork anomaly* shown in Fig. 3g, where clients cl_1 and cl_2 observe the updates to k_1 and k_2 in different orders. Assuming without loss of generality that $t_{cl_1}^2$ commits before $t_{cl_2}^2$, then prior to committing its transaction cl_2 sees the version of k_1 with value v_0. However, since $t \xrightarrow{\mathsf{WR}_\mathcal{K}} t_{cl_1}^1 \xrightarrow{\mathsf{SO}} t_{cl_1}^2 \xrightarrow{\mathsf{RW}} t' \xrightarrow{\mathsf{WR}} t_{cl_2}^1 \xrightarrow{\mathsf{SO}} t_{cl_2}^2$, then cl_2 should also see the version of k_1 with value v_2, leading to a contradiction.

4.1.7 Snapshot Isolation (SI). When the total order in which transactions commit is known, SI can be defined compositionally from CP and UA. As such, vshift$_\mathsf{SI}$ is defined as the conjunction of their associated vshift predicates. However, as with PSI, we cannot define can-commit$_\mathsf{SI}$ as the

conjunction of their associated can-commit predicates. Rather, we define can-commit$_{SI}$ as closure with respect to R_{SI}, which includes $R_{CP} \cup R_{UA}$. Observe that the kv-store in Fig. 3h shows an example kv-store that satisfies can-commit$_{UA} \wedge$ can-commit$_{CP}$, but not can-commit$_{SI}$. Additionally, we include WW; RW in R_{SI}. This is because when the centralised CP implementation (discussed above) is strengthened with write-conflict freedom, then a write-write dependency between two transactions t and t' does not only mandate that t commits before t' commits but also before t' starts. Consequently, if $(t, t') \in$ WW; RW, then t must commit before t' commit.

4.1.8 (Strict) serialisability (SER). Serialisability is the strongest consistency model in the literature, requiring that transactions execute in a total sequential order. The can-commit$_{SER}$ thus allows clients to commit transactions only when their view of the kv-store is complete, i.e. the client view is closed with respect to WW^{-1}. This requirement prevents the kv-store in Fig. 3i: without loss of generality, suppose that t_1 commits before t_2. Then the client committing t_2 must see the version of k_1 written by t_1, and thus cannot read the outdated value v_0 for k_1. This example is allowed by all other execution tests in Fig. 3.

4.1.9 Weak Snapshot Isolation (WSI): *A New Consistency Model.* Kv-stores and execution tests are useful for investigating new consistency models. One example is the consistency model induced by combining CP and UA, which we refer to as *Weak Snapshot Isolation* (WSI). To justify this consistency model in full, it would be useful to explore its implementations. Here we focus on the benefits of implementing WSI. Because WSI is stronger than CP and UA by definition, it forbids all the anomalies forbidden by these consistency models, e.g. the long fork (Fig. 3g) and the lost update (Fig. 3d). Moreover, WSI is strictly weaker than SI. As such, WSI allows all SI anomalies, e.g. the write skew (Fig. 3i), and allows behaviour not allowed under SI such as that in Fig. 3h. We can construct a (ET$_{CP} \cap$ ET$_{UA}$)-trace terminating in $(\mathcal{K}, \_)$ by executing transactions t_1, t_2, t_3 and t_4 in this order. In particular, t_4 is executed using $u=[k_1 \mapsto \{0\}, k_2 \mapsto \{0, 1\}]$. However, the same trace is not a valid ET$_{SI}$-trace. Under SI transaction t_4 cannot be executed using u: t_4 reads the version of k_2 written by t_3, meaning that u must include the version written by t_3. Since $(t_2, t_3) \in$ RW and $(t_1, t_2) \in$ WW, then u should contain the version of k_1 written by t_1, contradicting the fact that t_4 reads the initial version of k_1.

As WSI is a weaker consistency model than SI, we believe that WSI implementations would outperform known SI implementations. Nevertheless, the two consistency models are very similar in that many applications that are correct under SI are also correct under WSI. We give an example of such an application in §6.

5 CONSISTENCY MODELS: DEPENDENCY GRAPHS AND ABSTRACT EXECUTIONS

We demonstrate that our consistency models for kv-stores are equivalent to the declarative consistency models for dependency graphs [Adya 1999] and abstract executions [Burckhardt et al. 2012; Cerone et al. 2015a]. We outline our results here, and refer the reader to §D to F for the full details.

5.1 Relating KV-Stores and Dependency Graphs

Dependency graphs [Adya 1999; Adya et al. 2000] provide perhaps the most well-known formalism used for specifying transactional consistency models. A dependency graph $\mathcal{G}$ is a directed, labelled graph whose nodes denote transactions, and whose edges denote *dependencies* between transactions. More specifically, nodes are labelled with a transaction identifier and the fingerprint associated with the transaction. Edges are labelled with a dependency relation SO, WR, WW, RW, in the same spirit of dependencies of transactions in kv-stores in §4. An example of dependency graph is given in Fig. 4a. We give the formal definition of dependency graphs in §B.

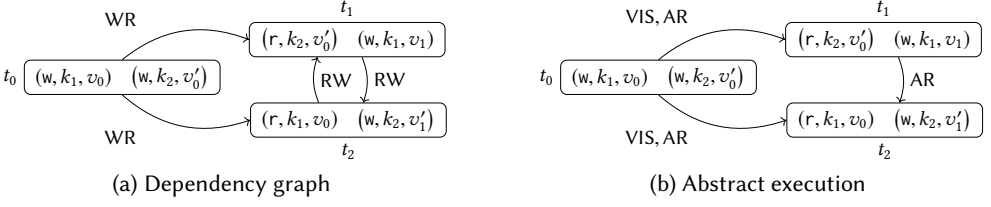(a) Dependency graph   (b) Abstract execution

Fig. 4. The dependency graph (a) and abstract execution graph (b) associated with the kv-store in Fig. 3i

We can always extract a dependency graph from a kv-store, and vice-versa. For example, Fig. 4a corresponds to the dependency graph extracted from the kv-store in Fig. 3i.

THEOREM 5.1. *Dependency graphs are isomorphic to kv-stores.*

Consistency models using dependency graphs can be specified by constraining the shape of the graphs, typically by requiring the absence of certain cycles. For example, strict serialisability is defined as the set of dependency graphs with no cycles. We can immediately use such constraints to define execution tests on kv-stores, and hence consistency models for kv-stores. However, to show that our consistency models over kv-stores given in Fig. 3 are equivalent to existing consistency model definitions using dependency graphs, we first prove that our models are equivalent to existing definitions using abstract executions, and then appeal to the results of Cerone et al. [2017] showing the equivalence between definitions using dependency graphs and those using abstract executions.

5.2 Relating KV-Stores and Abstract Executions

We compare our consistency model specifications using execution tests over kv-stores with an alternative, axiomatic specification style based on abstract executions [Cerone et al. 2015a], defined shortly. Our main contribution here is the development of a general proof technique for proving the equivalence of our execution-test-based specifications and abstract-execution-based specifications. Our proof technique keeps the proof obligations (conditions that must be satisfied by its user) to a minimum. In particular, the user only needs to show that the constraints on client views in execution tests relate to analogous constraints on visibility edges in abstract executions. We then provide a mapping between the ET-traces to $\mathcal{K}$, to a set of abstract executions that satisfy the axiomatic specification corresponding to ET. Here we use our proof technique to prove that the execution tests for serialisability (SER) are equivalent to their axiomatic specifications. In §F we apply our proof technique to show that all the execution tests from Fig. 3 are equivalent to their respective axiomatic specifications.

5.2.1 Abstract Executions. Abstract executions are labelled graphs whose nodes comprise transaction identifiers and their fingerprints. These nodes may be connected either by a *visibility edge*, $t \xrightarrow{\text{VIS}} t'$, when t' sees the updates of t; or an *arbitration edge*, $t \xrightarrow{\text{AR}} t'$, when t' updates overwrite t updates. Fig. 4b depicts an example abstract execution.

Definition 5.2 (Abstract executions). An *abstract execution* is a triple $X = (\mathcal{T}, \text{VIS}, \text{AR})$, where $\mathcal{T} : \text{TransID} \rightharpoonup \mathcal{P}(\text{Ops})$ is a partial function mapping transaction identifiers to fingerprints, with $\mathcal{T}(t_0) = \{(\text{w}, k, v_0) \mid k \in \text{Key}\}$, $\text{VIS} \subseteq \text{dom}(\mathcal{T}) \times \text{dom}(\mathcal{T})$ is an irreflexive relation such that, for any $t \in \text{dom}(\mathcal{T})$, $t_0 \xrightarrow{\text{VIS}} t$ for the initial transaction t_0, and $\text{AR} \subseteq \text{dom}(\mathcal{T}) \times \text{dom}(\mathcal{T})$ is a strict, total order such that $\text{VIS} \subseteq \text{AR}$, $\min_{\text{AR}}(\text{dom}(\mathcal{T})) = t_0$ and $t_{cl}^n \xrightarrow{\text{AR}} t_{cl}^m$ only if $n < m$.

Given an abstract execution $X = (\mathcal{T}, \text{VIS}, \text{AR})$, we let $\mathcal{T}_X = \mathcal{T}$, $\text{VIS}_X = \text{VIS}$ and $\text{AR}_X = \text{AR}$. We write $(l, k, v) \in_X t$ as a shorthand for $(l, k, v) \in \mathcal{T}_X(t)$. For $t \in \mathcal{T}_X$, we define its *visible writes* in X as $\text{visibleWrites}_X(k, t) \triangleq \text{VIS}_X^{-1}(t) \cap \{t' \mid (\text{w}, k, \_) \in_X t'\}$. An abstract execution X satisfies the *last-writer-wins* policy: if a transaction t reads key k, it must read from the latest transaction in the arbitration order that is visible to t and wrote to key k, i.e. $\forall t \in \mathcal{T}_X. (\text{r}, k, v) \in_X t \Rightarrow (\text{w}, k, v) \in \max_{\text{AR}_X}(\text{visibleWrites}_X(k, t))$. then $\forall k, v. (\text{r}, k, v) \in_X t \Rightarrow (\text{w}, k, v) \in \max_{\text{AR}_X}(\text{visibleWrites}_X(k, t))$. Henceforth we assume that abstract executions satisfy the last-writer-wins policy, and we let AbsExecs be the set of all such abstract executions.

Abstract-execution-based specifications of consistency models constrain the overall structure of abstract executions. For most consistency models [Cerone et al. 2015a, 2017; Nagar and Jagannathan 2018], such constraints are over the set of transactions that **must** be seen by other transactions. For example, monotonic reads is specified by requiring that if a transaction t follows another transaction t' in the session order, then t must see all transactions that are seen by t'. Serialisability can be specified by requiring that a transaction t see all transactions preceding t in the arbitration order.

Formally, an axiomatic specification $\mathcal{A}$ is a set of *axioms* $\text{A} : \text{AbsExecs} \rightarrow \mathcal{P}(\text{TransID} \times \text{TransID})$, where $\forall X. \text{A}(X) \subseteq \text{AR}_X$. We write $X \models \text{A}$ when $\text{A}(X) \subseteq \text{VIS}_X$. We refer the reader to §C for details about abstract executions.

Returning to the monotonic reads (MR) example, we define $\mathcal{A}_{\text{MR}} \triangleq \{\text{A}_{\text{MR}}\}$, where $\text{A}_{\text{MR}}(X) \triangleq \text{VIS}_X; \text{SO}_{\text{VIS}}$. By definition, for a given X, $X \models \text{A}_{\text{MR}}$ if and only if $\text{VIS}_X; \text{SO}_X \subseteq \text{VIS}_X$. That is, whenever $t'' \xrightarrow{\text{VIS}_X} t' \xrightarrow{\text{SO}_X} t$, then $t'' \xrightarrow{\text{VIS}_X} t$. Similarly, for serialisability (SER) we define $\mathcal{A}_{\text{SER}} \triangleq \{\text{A}_{\text{SER}}\}$, where $\text{A}_{\text{SER}}(X) \triangleq \text{AR}_X$, captures the constraint that a transaction t must see all transactions preceding it in the arbitration order.

Any abstract executions X can be mapped into an equivalent dependency graph $\mathcal{G}_X$ (Cerone et al. [2017]), and hence into a kv-store $\mathcal{K}_X$ (Theorem 5.1). We can then use this construction to define the consistency model induced by an abstract-execution-based specification $\text{CM}(\mathcal{A})$ by projecting abstract executions that satisfy the axioms in $\mathcal{A}$ to kv-stores: $\text{CM}(\mathcal{A}) \triangleq \{\mathcal{K}_X \mid \forall \text{A} \in \mathcal{A}.X \models \text{A}\}$.

In the remainder of this section we develop a proof techniques for showing that an execution test ET and an axiomatic specification $\mathcal{A}$ induce the same consistency model, i.e. $\text{CM}(\text{ET})=\text{CM}(\mathcal{A})$. Due to space constraints, we focus only on *soundness*, i.e. on proving the left-to-right inclusion: $\text{CM}(\text{ET}) \subseteq \text{CM}(\mathcal{A})$; we describe the other direction in full in §E. The core of our proof technique lies in the soundness of the most permissive execution test, $\text{CM}(\text{ET}_\top)$, with respect to the weakest axiomatic specification, given by the empty set of axioms, which we prove next.

5.2.2 Equivalence of Read Atomic and $\text{CM}(\text{ET}_\top)$. The weakest axiomatic specification, given by the empty set of axioms, corresponds to the *Read Atomic* consistency model [Bailis et al. 2014]. To prove that the most permissive execution test $\text{ET}_\top$ is sound with respect to the weakest axiomatic specification, we map $\text{ET}_\top$-traces terminating in a configuration of the form $(\mathcal{K}, \_)$, into the set of abstract executions whose underlying kv-store is $\mathcal{K}$.

THEOREM 5.3. *Given an* $\text{ET}_\top$*-trace* τ *terminating in* $(\mathcal{K}, \_)$*, there exists a non-empty set of abstract executions* $\text{execs}(\tau)$ *such that:* $\forall X \in \text{execs}(\tau). \mathcal{K}_X = \mathcal{K}$*, and the order in which transactions are executed in* τ *is consistent with* AR_X*.*

The proof of Theorem 5.3 is highly non-trivial, and relies on the following intuition that drives the construction of the set $\text{execs}(\tau)$: whenever a client cl in τ with view u commits a transaction t, then in all abstract executions included in $\text{execs}(\tau)$, transaction t must see the writers of the versions included in u, and it never sees the writers of versions not included in u (Fig. 5). These are defined by the set $\text{visTx}(\mathcal{K}, u)$ (defined in §4.1). Furthermore, abstract executions in $\text{execs}(\tau)$ differ only in
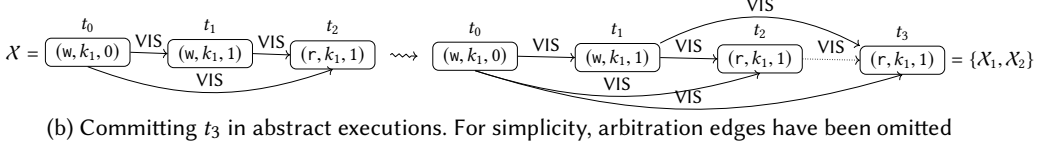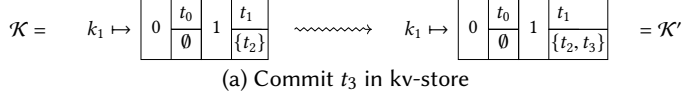
(a) Commit t_3 in kv-store

(b) Committing t_3 in abstract executions. For simplicity, arbitration edges have been omitted

Fig. 5. Correspondence between committing t_3 in kv-stores and abstract executions. The figure to the right denotes a set of abstract executions, which differ in the presence of the dashed visibility edge.

the set of read-only transactions (i.e. those with no write operations) that transactions see. While mapping an $ET_\top$-trace into multiple abstract executions is not strictly necessary for proving the soundness of $ET_\top$ with respect to the weakest axiomatic specification, it plays a significant role when proving the soundness of an arbitrary execution test ET with respect to its counterpart in axiomatic specifications.

The definition of $execs(\tau)$ is by induction on the length of the $ET_\top$-trace τ. For the base case with τ_0 consisting of a single configuration $(\mathcal{K}_0, \_)$, we define $execs(\tau_0)$ to contain a single abstract execution with a single transaction t_0 that initialises all the keys to the initial value v_0: $execs(\tau_0) \triangleq \{([t_0 \mapsto \{(w, k, v_0 \mid k \in \text{KEY})\}], \emptyset, \emptyset\}$. For the inductive case with $\tau = \tau' \xrightarrow{(cl, \alpha)}_{ET}$ $(\mathcal{K}', \mathcal{U}')$, let $(\mathcal{K}, \mathcal{U})$ be the last configuration appearing in τ'. If $\alpha = \varepsilon$, then $execs(\tau) \triangleq execs(\tau')$. If $\alpha = \mathcal{F}$ for some $\mathcal{F}$, we first determine the transaction identifier t' that was used to commit $\mathcal{F}$ in $\mathcal{K}'$, the view $u' = \mathcal{U}'(cl)$ of the client cl when committing t', the set of transactions that cl must see when committing t', given by $visTx(\mathcal{K}', u')$, and the set of read-only transactions $\mathcal{T}_{rd}$ in $\mathcal{K}'$: the latter are those transactions that never appear as writers. Then, for all abstract execution $\mathcal{X}' \in execs(\tau')$, we define $extend(\mathcal{X}', t', visTx(\mathcal{K}', u'), \mathcal{F})$ as the largest set such that, whenever $\mathcal{X} \in extend(\mathcal{X}', t', visTx(\mathcal{K}', u'), \mathcal{F})$, then (1) $\mathcal{T}_\mathcal{X} = \mathcal{T}_{\mathcal{X}'}[t' \mapsto \mathcal{F}]$; (2) for all $t' \in \mathcal{T}_\mathcal{X}$, $t' \xrightarrow{AR_{\mathcal{X}'}} t$; and (3) if $t' \xrightarrow{VIS_\mathcal{X}} t$, then either $t \in visTx(\mathcal{K}', u')$, or $t \in \mathcal{T}_{rd}$. Finally, we define $execs(\tau) \triangleq \bigcup_{\mathcal{X} \in execs(\tau')} extend(\mathcal{X}, t, visTx(\mathcal{K}', u'), \mathcal{T}_{rd}, \mathcal{F})$. In §D.2 and D.3 we give the full details. To understand the construction outlined above, we illustrate one use of the function extend. The abstract execution $\mathcal{X}$ to the left of Fig. 5b has underlying kv-store $\mathcal{K}$, depicted to the left of Fig. 5a. If a client commits a transaction t_3 that reads the last version of k_1, then the resulting kv-store $\mathcal{K}'$ would be the one to the right of Fig. 5a. This commit is simulated by the function $extend(\mathcal{X}', t_3, visTx(\mathcal{K}, u), \mathcal{F})$, where $u, \mathcal{F}$ are the view and fingerprint used to update $\mathcal{K}$ to $\mathcal{K}'$: the result of this function consists of two abstract executions $\mathcal{X}_1, \mathcal{X}_2$, that only differ in read-only transactions (the right of Fig. 5a). The visibility edges of $\mathcal{X}_1$ are exactly the concrete edges in Fig. 5b; however, $\mathcal{X}_2$ has the extra dashed visibility edge $t_2 \xrightarrow{VIS} t_3$. Note that $\mathcal{K}_{\mathcal{X}_2} = \mathcal{K}_{\mathcal{X}_3} = \mathcal{K}'$.

THEOREM 5.4. *Given an abstract execution $\mathcal{X}$, there exists a non-empty set of $ET_\top$-traces $\{\tau_i\}_{i \in I}$ such that, for each $i \in I$, the last configuration of τ_i is $(\mathcal{K}_\mathcal{X}, \_)$, and τ_i executes transactions in the order established by $AR_\mathcal{X}$.*

The proof of Theorem 5.4 is given in §D.2 and D.3. Theorems 5.3 and 5.4 together establish the equivalence of $ET_\top$ with the weakest axiomatic specification.

5.2.3 Equivalence of axiomatic specifications and execution tests. We are now ready to present our proof technique for proving the soundness of an execution test ET with respect to an axiomatic specification $\mathcal{A}$. It can be summarised as follows: the user considers an arbitrary $\mathcal{X} : \mathcal{X} \models \mathcal{A}$

, and a tuple of the form $\mathsf{ET} \vdash (\mathcal{K}_X, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$. Then, it constructs a non-empty subset of $\mathsf{extend}(X, t, \mathsf{visTx}(\mathcal{K}, u), \mathcal{F})$ whose elements satisfy the axioms $\mathcal{A}$. Note that, because abstract executions in $\mathsf{extend}(X, t, \mathsf{visTx}(\mathcal{K}, u), \mathcal{F})$ differ only in visibility edges of the form $t_{\mathrm{rd}} \xrightarrow{\mathsf{VIS}} t$, where t_{rd} is a read-only transaction, then constructing the set above amounts to identifying a suitable set of read-only transactions in X.

To see why our proof technique guarantees the soundness of ET with respect to $\mathcal{A}$ (Theorem 5.6), we apply an inductive argument over the number of ET-reductions n in a ET-trace τ: first, if $n = 0$ then $\tau = (\mathcal{K}_0, \mathcal{U}_0)$, and $X_0 \in \mathsf{execs}(\tau)$ trivially satisfies the axioms $\mathcal{A}$: $\forall \mathsf{A} \in \mathcal{A}.\mathsf{A}(X_0) \subseteq \mathsf{AR}_{X_0} = \emptyset \subseteq \mathsf{VIS}_{X_0}$. Otherwise, if $\tau = \tau' \xrightarrow{(cl, \mathcal{F})} (\mathcal{K}, u)$, suppose that there exists an abstract execution $X' \in \mathsf{execs}(\tau')$ that satisfies the axioms $\mathcal{A}$. If the proof obligations of our proof technique are satisfied, we can construct an abstract execution $X \in \mathsf{extend}(X, t, \mathsf{visTx}(\mathcal{K}, u), \mathcal{F}) \subseteq X(\tau)$ that satisfies $\mathcal{A}$; furthermore $\mathcal{K}_X = \mathcal{K}$.

In practice, our proof techniques allows defining a per-client invariant on the visibility relation of abstract executions, which must be proved to be preserved by ET-reductions (Def. 5.5). This invariant carries client-specific information that links to vshift (defined in §4) in execution tests. Defining the right invariant is crucial for proving the soundness of several execution-test-based specifications (see §F).

Definition 5.5. An execution test ET is sound with respect to an axiomatic specification $\mathcal{A}$ if and only if there exists an invariant condition I such that, for any $cl, t, u, u', \mathcal{K}, u', \mathcal{F}, X$, if:

- $\mathsf{ET} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$, where $\mathcal{K}' = \mathsf{update}(\mathcal{K}, u, \mathcal{F}, t)$
- $\mathcal{K}_X = \mathcal{K}$ and $I(X, cl) \subseteq \mathsf{visTx}(\mathcal{K}, u)$,

then there exists a non-empty subset of $\mathscr{X} \subseteq \mathsf{extend}(X, t, \mathsf{visTx}(\mathcal{K}, u), \mathcal{F})$ such that, for any $X' in \mathscr{X}$, $X \models \mathcal{A}$ and $I(X', cl) \subseteq \mathsf{visTx}(\mathcal{K}', u')$.

THEOREM 5.6. *If* ET *is sound with respect to* $(\mathcal{A})$, *then* $\mathsf{CM}(\mathsf{ET}) \subseteq \mathsf{CM}(\mathcal{A})$.

We conclude this section by outlining how our proof technique can be applied to show that the execution test $\mathsf{ET}_{\mathsf{SER}}$ defined for serialisability is sound with respect to the axiomatic specification $\mathcal{A}_{\mathsf{SER}}$. Let X be an abstract execution such that $X \models \mathcal{A}_{\mathsf{SER}}$, and suppose that the underlying kv-store $\mathcal{K}_X$ is in $\mathsf{CM}(\mathsf{ET}_{\mathsf{SER}})$. An example is the abstract execution X and kv-store $\mathcal{K}$ to the left of Fig. 5. We pick an invariant I_{SER} that is always empty since $\mathsf{vshift}_{\mathsf{SER}}$ is always true. When a client cl commits a transaction t with fingerprint $\mathcal{F}$ in $\mathcal{K}_X$ under $\mathsf{ET}_{\mathsf{SER}}$, then its view u must contain all the versions stored in $\mathcal{K}_X$ ($\mathsf{can\text{-}commit}_{\mathsf{SER}}(\mathcal{K}_X, u, \mathcal{F})$). This means that all the abstract executions X'' in $\mathsf{extend}(X, t, \mathsf{visTx}(\mathcal{K}_X, u), \mathcal{F})$ are such that there is a visibility edge $t' \xrightarrow{\mathsf{VIS}_{X''}} t$ where t' is a writer transaction in $\mathcal{K}_X$. It is easy to see there exists an singleton set $\{X'\}$ that is a subset of $\mathsf{extend}(X, t, \mathsf{visTx}(\mathcal{K}, u), \mathcal{F})$ and $X' \models \mathcal{A}_{\mathsf{SER}}$; in particular, there is an edge $t' \xrightarrow{\mathsf{VIS}_{X'}} t$ for any transaction $t' \in \mathcal{T}_{X'}$. For example, in Fig. 5, the possible abstract executions are X_1, X_2. We pick $\{X_2\}$, because the new transaction t_3 in X_2 sees all previous transactions including the visibility edge $t_2 \xrightarrow{\mathsf{VIS}} t_3$. Last, the invariant $I_{\mathsf{SER}}(X', cl) \subseteq \mathsf{visTx}(\mathcal{K}', u)$ given $I_{\mathsf{SER}}(X', cl) = \emptyset$.

6 APPLICATIONS

To demonstrate the applications of our operational semantics, in §6.1 we use our formalism to prove the robustness of several transactional libraries. In §6.2 we then use our formalism to verify several distributed protocols.
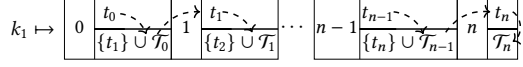
Fig. 6. The canonical structure of a kv-store in $\mathrm{CM}(\mathrm{ET_{PSI}}, \mathrm{Counter})$, where each t_i is the result of a $\mathrm{inc}(k)$ operation, and each transaction in $\mathcal{T}_i$ is the result of a $\mathrm{read}(k)$ operation.

6.1 Application: Robustness of Transactional Libraries

A transactional library, $L = \{[\mathsf{T}]_i\}_{i \in I}$, provides a set of transactional operations which can be used by its clients to access the underlying kv-store[5]. For instance, the set of operations of the counter library on key k in §2 is: $\mathrm{Counter}(k) \triangleq \{\mathrm{inc}(k), \mathrm{read}(k)\}$. A program P is a *client program* of L if the only transactional calls in P are to operations of L. Let $\mathrm{CM}(\mathrm{ET}, L)$ denote the set of kv-stores obtained by running L clients under ET. A library L is *robust* against an execution test ET if for all clients of L, the kv-stores obtained under ET can also be obtained under SER; i.e. $\mathrm{CM}(\mathrm{ET}, L) \subseteq \mathrm{CM}(\mathrm{ET_{SER}})$.

THEOREM 6.1. *For all kv-stores* $\mathcal{K}$, $\mathcal{K} \in \mathrm{CM}(\mathrm{ET_{SER}})$ *iff* $(\mathrm{SO} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}} \cup \mathrm{RW}_{\mathcal{K}})^+$ *is irreflexive.*

This theorem is an adaptation of a well-known result on dependency graphs [Adya 1999] stating that $\mathcal{K} \in \mathrm{CM}(\mathrm{ET_{SER}})$ if and only if its associated dependency graph is acyclic. Using this theorem, we prove the robustness of a single counter against PSI. As discussed in §2, the multi-counter library is not robust against PSI. We thus prove the robustness of the multi-counter library and the banking library of Alomari et al. [2008] against WSI instead. While previous work on checking robustness [Bernardi and Gotsman 2016; Cerone and Gotsman 2016; Cerone et al. 2017; Nagar and Jagannathan 2018] uses static-analysis techniques that cannot be extended to support client sessions, we give the first robustness proofs that take client sessions into account.

6.1.1 Robustness of a Single Counter against PSI. In the single-counter library $\mathrm{Counter}(k)$, a client reads from k by calling $\mathrm{read}(k)$, and writes to k by calling $\mathrm{inc}(k)$ which first reads the value of k and subsequently increments it by one. Pick an arbitrary key k and a kv-store $\mathcal{K} \in \mathrm{CM}(\mathrm{ET_{PSI}}, \mathrm{Counter}(k))$. As PSI enforces write conflict freedom (UA), we know that if a transaction t updates k (by calling $\mathrm{inc}(k)$) and writes version v to k, then it must have read the version of k immediately preceding v: $\forall t, i > 0 . t = \mathrm{w}(\mathcal{K}(k, i)) \Rightarrow t \in \mathrm{rs}(\mathcal{K}(k, i{-}1))$. Moreover, as PSI enforces monotonic reads (MR), the order in which clients observe the versions of k (by calling $\mathrm{read}(k)$) is consistent with the order of versions in $\mathcal{K}(k)$. As such, the kv-stores in $\mathrm{CM}(\mathrm{ET_{PSI}}, \mathrm{Counter}(k))$ have the canonical structure depicted in Fig. 6 and defined below, where :: denotes list concatenation, and $\{t_i\}_{i=1}^{n}$ and $\bigcup_{i=0}^{n} \mathcal{T}_i$ denote disjoint sets of transactions calling $\mathrm{inc}(k)$ and $\mathrm{read}(k)$, respectively:

$$\mathcal{K}(k) = (0, t_0, \mathcal{T}_0 \cup \{t_1\}) :: (1, t_1, \mathcal{T}_1 \cup \{t_2\}) :: \cdots :: (n-1, t_{n-1}, \mathcal{T}_{n-1} \cup \{t_n\}) :: (n, t_n, \mathcal{T}_n)$$

We define the $\dashrightarrow$ relation depicted in Fig. 6 by extending $\mathrm{SO} \cup \{(t_i, t_j) \mid t_j \in \mathcal{T}_i \vee (t_i \in \mathcal{T}_i \wedge j = i+1)\}$ to a strict total order (i.e. is $\dashrightarrow$ irreflexive and transitive). Note that $\dashrightarrow$ contains $\mathrm{SO} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}} \cup \mathrm{RW}_{\mathcal{K}}$ and thus $(\mathrm{SO} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}} \cup \mathrm{RW}_{\mathcal{K}})^+$ is irreflexive; i.e. $\mathrm{Counter}(k)$ is robust against PSI.

Recall from §2 that unlike in SER and SI, under PSI clients can observe the increments on different keys in different orders (see Fig. 3g). As such, multiple counters are not robust against PSI.

6.1.2 Robustness Conditions against WSI. Several libraries in the literature that are robust against SI [Alomari et al. 2008; Bernardi and Gotsman 2016] are also robust against WSI. The operations

---

[5]For simplicity, we model each operation as a single transaction; it is straightforward to extend this to multiple transactions.

of these libraries all yield kv-stores that adhere to a particular pattern captured by the following definition.

Definition 6.2 (WSI-safe). A kv-store $\mathcal{K}$ is WSI-*safe* if it is reachable by executing a program P from an initial configuration Γ_0 under WSI (i.e. $\Gamma_0, \mathsf{P} \rightarrow_{\mathsf{ET}_{\mathsf{WSI}}} (\mathcal{K}, \_), \_$), and for all t, k, i:

$$t \in \mathsf{rs}(\mathcal{K}(k,i)) \wedge t \neq \mathsf{w}(\mathcal{K}(k,i)) \Rightarrow \forall k', j. \ t \neq \mathsf{w}(\mathcal{K}(k',j)) \tag{6.1}$$

$$t \neq t_0 \wedge t = \mathsf{w}(\mathcal{K}(k,i)) \Rightarrow \exists j. \ t \in \mathsf{rs}(\mathcal{K}(k,j)) \tag{6.2}$$

$$t \neq t_0 \wedge t = \mathsf{w}(\mathcal{K}(k,i)) \wedge \exists k', j. \ t \in \mathsf{rs}(\mathcal{K}(k',j)) \Rightarrow t = \mathsf{w}(\mathcal{K}(k',j)) \tag{6.3}$$

This definition states that a kv-store $\mathcal{K}$ is WSI-safe if for each transaction t: (1) if t reads from k without writing to it then t must be a read-only transaction (6.1); (2) if t writes to k, then it must also read from it (6.2), a property known as *strictly-no-blind writes*; and (3) if t writes to k, then it must also write to all keys it reads (6.3). It is straightforward to see that the version j read by t in (6.2) must be written immediately before the version i written by t, i.e. $i=j+1$.

THEOREM 6.3 (WSI ROBUSTNESS). *If a kv-store $\mathcal{K}$ is WSI-safe, then it is robust against WSI.*

From Theorem 6.1 it suffices to prove that $(\mathsf{SO} \cup \mathsf{WR}_{\mathcal{K}} \cup \mathsf{WW}_{\mathcal{K}} \cup \mathsf{RW}_{\mathcal{K}})^+$ is irreflexive. We proceed by contradiction and assume that there exists t_1 such that $t_1 \xrightarrow{(\mathsf{SO} \cup \mathsf{WR}_{\mathcal{K}} \cup \mathsf{WW}_{\mathcal{K}} \cup \mathsf{RW}_{\mathcal{K}})^+} t_1$. Since $\mathcal{K}$ is reachable under WSI and thus also reachable under CC, this cycle is of the form:

$$t_1 \xrightarrow{R^*} t_2 \xrightarrow{\mathsf{RW}} t_3 \xrightarrow{R^*} \cdots \xrightarrow{R^*} t_{n-2} \xrightarrow{\mathsf{RW}} t_{n-1} \xrightarrow{R^*} t_n = t_1$$

where $R \triangleq \mathsf{WR} \cup \mathsf{SO} \cup \mathsf{WW}$. From Eqs. (6.2) and (6.3) we know that an RW edge from a writing transaction can be replaced by a WW edge. Moreover, WW edges can be replaced by WR^* edges since $\mathcal{K}$ is reachable under UA. We thus have:

$$t_1 \xrightarrow{R'^*} t_2' \xrightarrow{\mathsf{RW}} t_3' \xrightarrow{R'^+} \cdots \xrightarrow{R'^+} t_{m-2}' \xrightarrow{\mathsf{RW}} t_{m-1}' \xrightarrow{R'^*} t_m' = t_n = t_1$$

where $R' \triangleq \mathsf{WR} \cup \mathsf{SO}$. That is, $t_1 \xrightarrow{((\mathsf{WR} \cup \mathsf{SO}); \mathsf{RW}^?)^*} t_n$. This however leads to a contradiction as $(\mathsf{WR} \cup \mathsf{SO}); \mathsf{RW}^? \subseteq R_{\mathsf{CP}}$ and WSI requires views to be closed under R_{CP} (see Fig. 3).

6.1.3 Robustness of Multiple Counters against WSI. A multi-counter library on a set of keys K is: $\mathsf{Counters(K)} \triangleq \bigcup_{k \in K} \mathsf{Counter}(k)$. We next show that a multi-counter library is WSI-safe, and is therefore robust against WSI and all stronger models such as SI.

THEOREM 6.4. *A multi-counter library* $\mathsf{Counters(K)}$ *is robust against WSI.*

It is sufficient to show that a kv-store obtained by executing arbitrary transactional calls from the library are WSI-safe. We proceed by induction on the length of traces. Let $\Gamma_0 = (\mathcal{K}_0, \mathcal{U}_0)$ be an initial configuration and P_0 be a program such that $\mathsf{dom}(\mathsf{P}_0) \subseteq \mathsf{dom}(\mathcal{U}_0)$. The initial kv-store trivially satisfies (6.1), (6.2) and (6.3) above. Let $\mathcal{K}_i$ be the resulting kv-store after i steps of execution under WSI. The next transaction t_{i+1} may then be a call to either inc(k) or read(k). If t_{i+1} is a read(k) call, then the resulting kv-store is:

$$\mathcal{K}_{i+1} = \mathcal{K}_i[\mathsf{k} \mapsto \mathcal{K}_i(\mathsf{k})[j \mapsto (v, t, \mathcal{T} \uplus \{t_{i+1}\})]]$$

for some j and $\mathcal{K}_i(\mathsf{k}, j)=(v, t, \mathcal{T})$. Since t_{i+1} is a read-only transaction, then (6.1), (6.2) and (6.3) immediately hold. On the other hand, if t_{i+1} is an inc(k) call, then the resulting kv-store is:

$$\mathcal{K}_{i+1} = \mathcal{K}_i[\mathsf{k} \mapsto (\mathcal{K}_i(\mathsf{k})[j \mapsto (v, t, \mathcal{T} \uplus \{t_{i+1}\})])::(v+1, t_{i+1}, \emptyset)]$$

where $j=|\mathcal{K}_i(\mathsf{k})|$ and $\mathcal{K}_i(\mathsf{k}, j)=(v, t, \mathcal{T})$. As t_{i+1} reads the latest version of k and writes a new version to k, the new kv-store $\mathcal{K}_{i+1}$ satisfies (6.1), (6.2) and (6.3).

6.1.4 Robustness of a Banking Library against WSI. Alomari et al. [2008] present a banking library which is robust against SI. Here, we show that this library is also robust against WSI. The banking example is based on relational databases and has three tables: account, saving and checking. The account table maps customer names to customer IDs (Account(Name, CID)); the saving table maps customer IDs to their saving balances (Saving(CID, Balance)); and the checking table maps customer IDs to their checking balances (Checking(CID, Balance)).

For simplicity, we encode the saving and checking tables as a kv-store, and forgo the account table as it is an immutable lookup table. We model a customer ID as an integer $n \in \mathbb{N}$, and assume that balances are integer values. We then define the key associated with customer n in the checking table as $n_c \triangleq 2n$; and define the key associated with n in the saving table as $n_s \triangleq 2n+1$. That is, $\text{KEY} \triangleq \bigcup_{n \in \mathbb{N}} \{n_c, n_s\}$. Moreover, if n identifies a customer, i.e. $(\_, n) \in \text{Account}(\underline{\text{Name}}, \text{CID})$, then:

$$(n, \text{val}(\mathcal{K}(n_s, |\mathcal{K}(n_s)|))) \in \text{Saving}(\underline{\text{CID}}, \text{Balance}) \qquad (n, \text{val}(\mathcal{K}(n_s, |\mathcal{K}(n_c)|))) \in \text{Checking}(\underline{\text{CID}}, \text{Balance})$$

The banking library provides five transactional operations for accessing the database:

$$\text{balance(n)} \triangleq [\text{x} := [n_c]; \text{ y} := [n_s]; \text{ ret} := \text{x} + \text{y}]$$

$$\text{depositChecking(n, v)} \triangleq [\text{if}(\text{v} \geq 0)\{ \text{ x} := [n_c]; \text{ } [n_c] := \text{x} + \text{v}; \}]$$

$$\text{transactSaving(n, v)} \triangleq [\text{x} := [n_s]; \text{ if}(\text{v} + \text{x} \geq 0)\{ \text{ } [n_s] := \text{x} + \text{v}; \}]$$

$$\text{amalgamate(n, n')} \triangleq [\text{x} := [n_s]; \text{ y} := [n_c]; \text{ z} := [n'_c]; \text{ } [n_s] := 0; \text{ } [n_c] := 0; \text{ } [n'_c] := \text{x} + \text{y} + \text{z};]$$

$$\text{writeCheck(n, v)} \triangleq \begin{bmatrix} \text{x} := [n_s]; \text{ y} := [n_c]; \\ \text{if}(\text{x} + \text{y} < \text{v})\{ \text{ } [n_c] := \text{y} - \text{v} - 1; \}\text{else}\{ \text{ } [n_c] := \text{y} - \text{v}; \} \text{ } [n_s] := \text{x}; \end{bmatrix}$$

The balance(n) operation returns the total balance of customer n in ret. The depositChecking(n, v) operation deposits v to the checking account of customer n when v is non-negative; otherwise the checking account remains unchanged. When $\text{v} \geq 0$, transactSaving(n, v) deposits v to the saving account of n. When $\text{v} < 0$, transactSaving(n, v) withdraws v from the saving account of n only if the resulting balance is non-negative; otherwise the saving account remains unchanged. The amalgamate(n, n') operation moves the combined checking and saving funds of consumer n to the checking account of customer n'. Lastly, writeCheck(n, v) cashes a cheque of customer n in the amount v by deducting v from its checking account. If n does not hold sufficient funds (i.e. the combined checking and saving balance is less than v), customer n is penalised by deducting one additional pound. Alomari et al. [2008] argue that to make the banking library robust against SI, the writeCheck(n, v) operation must be strengthened by writing back the balance to the saving account (via $[n_s] := \text{x}$), even though the saving balance is unchanged.

The banking library is more complex than the multi-counter library discussed in §6.1.3. Nevertheless, all banking transactions are either read-only or satisfy the strictly-no-blind-writes property; i.e. the banking library is WSI-safe. As such, we can prove its robustness against WSI in a similar fashion to that of the multi-counter library. More concretely, given a WSI-safe kv-store $\mathcal{K}$, we show that the kv-store resulting from executing a banking operation on $\mathcal{K}$ remains WSI-safe.

As balance(n) is read-only, it immediately satisfies (6.1), (6.2) and (6.3). When $\text{v} \geq 0$, then depositChecking(n, v) both reads and writes n_c, and thus preserves (6.1), (6.2) and (6.3). When $\text{v} < 0$, then depositChecking(n, v) leaves the kv-store unchanged and thus (6.1), (6.2) and (6.3) are trivially preserved. Lastly, the transactSaving(n, v), amalgamate(n, n') and writeCheck(n, v) operations always read and write the keys they access, thus satisfying (6.1), (6.2) and (6.3).

6.2 Verifying Database Protocols

Kv-stores and views faithfully abstract the state of geo-replicated and partitioned databases, and execution tests provide a powerful abstraction of the synchronisation mechanisms enforced by

r_1

$(k_1, v_0, (t_0, r_0), \emptyset)$ $(k_2, v_0, (t_0, r_0), \emptyset)$

r_2

$(k_1, v_0, (t_0, r_0), \emptyset)$ $(k_2, v_0, (t_0, r_0), \emptyset)$

(a) An initial COPS state with two replicas (r_1, r_2); each replica contains two keys (k_1, k_2) with initial versions.

r_1

$(k_1, v_0, (t_0, r_0), \emptyset)$ $(k_2, v_0, (t_0, r_0), \emptyset)$ $(k_1, v_1, (t_1, r_1), \emptyset)$

r_2

$(k_1, v_0, (t_0, r_0), \emptyset)$ $(k_2, v_0, (t_0, r_0), \emptyset)$ $(k_1, v_1', (t_1, r_2), \emptyset)$

$(k_2, v_2, (t_2, r_2), \{(k_1, t_1, r_2)\})$

(b) Client cl_1 commits a new version of k_1 carrying value v_1 to replica r_1; other clients commit versions to r_2. The new versions in r_1 and r_2 have not yet been propagated to each other.

r_1

$(k_1, v_0, (t_0, r_0), \emptyset)$ $(k_2, v_0, (t_0, r_0), \emptyset)$ $(k_1, v_1, (t_1, r_1), \emptyset)$

$(k_1, v_1', (t_1, r_2), \emptyset)$ $(k_2, v_2', (t_2, r_2), \{(k_1, t_1, r_2)\})$

(c) Replica r_1 optimistically fetches the newest version for k_1, k_2 one by one, during which it receives synchronisation messages from r_2.

r_1

$(k_1, v_0, (t_0, r_0), \emptyset)$ $(k_2, v_0, (t_0, r_0), \emptyset)$ $(k_1, v_1, (t_1, r_1), \emptyset)$

$(k_1, v_1', (t_1, r_2), \emptyset)$ $(k_2, v_2', (t_2, r_2), \{(k_1, t_1, r_2)\})$

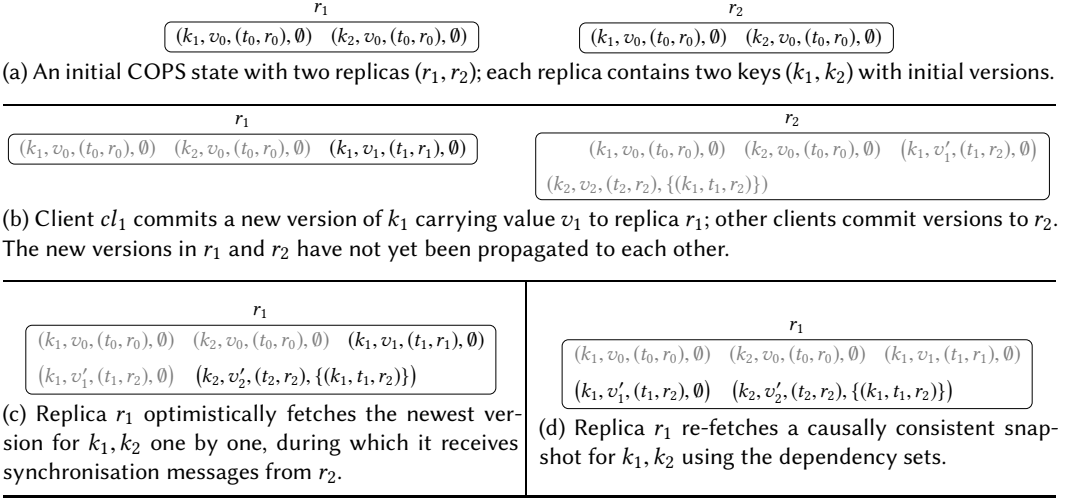(d) Replica r_1 re-fetches a causally consistent snapshot for k_1, k_2 using the dependency sets.

Fig. 7. COPS protocol

these databases when committing a transaction. This makes it possible to use our semantics to verify the correctness of distributed database protocols. We demonstrate this by showing that the replicated database COPS [Lloyd et al. 2011] satisfies causal consistency (this section and §H.1) and the partitioned database Clock-SI [Du et al. 2013] satisfies snapshot isolation (§H.2);

COPS is a fully replicated database, with each replica storing multiple versions of each key as illustrated in Figs. 7a and 7b. Each COPS version v, e.g. $(k_1, v_1, (t_1, r_1), \emptyset)$ in Fig. 7b, contains a key (e.g. k_1), a value (e.g. v_1), a time-stamp denoting when a client first wrote the version to the replica (e.g. (t_1, r_1)), and a set of dependencies, written $\text{deps}(v)$ (e.g. $\emptyset$). A time-stamp associated with a version v is of the form (t, r), where r identifies the replica that committed v, and t denotes the local, real time when r committed v. Each dependency in $\text{deps}(v)$ comprises a key and the time-stamp of the version of that key on which v directly depends. COPS assumes a total order among replica identifiers. Thus, time-stamps can be totally ordered lexicographically over time-stamps.

The COPS API provides two operations: one for writing to a single key; and another for atomically reading from a set of keys. Each call to a COPS operation is processed by a single replica. Each client maintains a *context*, which is a set of dependencies tracking the versions the client observes.

We demonstrate how a client cl interacts with a replica through the following example:

$$cl : [[k_1 := v_1;\,]\,;\; [x := [k_1];\; y := [k_2];\,]] \qquad \text{(cops-cl)}$$

For brevity, we assume that there are two keys (k_1 and k_2) and two replicas (r_1 and r_2) as shown in Fig. 7a, where $r_1 < r_2$. Initially, client cl connects to replica r_1 and initialises its local context as $ctx = \emptyset$. To execute its first single-write transaction, client cl requests to write v_1 to k_1 by sending the message (k_1, v_1, ctx) to its associated replica r_1. It then awaits a reply from r_1. Upon receiving the message, r_1 produces a monotonically increasing local time t_1, and uses it to install a new version $v = (k_1, v_1, (t_1, r_1), ctx)$, as shown in Fig. 7b. Note that the dependency set of v is the cl context ($ctx = \emptyset$). Replica r_1 then sends the time-stamp (t_1, r_1) back to client cl_1, and cl_1 in turn incorporates (k_1, t_1, r_1) in its local context; i.e. cl observes its own write. Finally, replica r_1 propagates the written version to other replicas asynchronously by sending a *synchronisation message* using *causal delivery* as follows. In the general case, when a replica r' receives a version v' from another replica r, it first waits for all dependencies of v to be committed to r', and then commits v. As such, the set of versions contained in each replica is closed with respect to causal dependencies. In the example
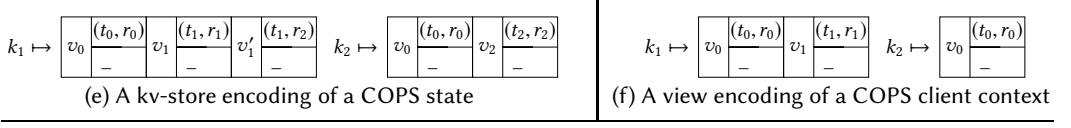
(e) A kv-store encoding of a COPS state

(f) A view encoding of a COPS client context

Fig. 7. COPS encoding

above, when other replicas receive version v from r_1, they can immediately commit v as its dependency set is empty. Note that replicas may accept new versions from different clients in parallel (see Fig. 7b).

To execute its second multi-read transaction, client cl requests to read from the k_1, k_2 keys by sending the message $\{k_1, k_2\}$ to replica r_1 and awaits a reply. Upon receiving this message, replica r_1 builds a causally consistent snapshot (a mapping from $\{k_1, k_2\}$ to values) in two rounds as follows. First, r_1 optimistically fetches the most recent versions for k_1 and k_2, one at a time. This process may be interleaved with other writes and synchronisation messages (propagated from other replicas). For instance, Fig. 7c depicts a scenario where r_1 (1) first fetches $(k_1, v_1, (t_1, r_1), \emptyset)$ for k_1 (highlighted); (2) then receives two synchronisation messages from r_2, containing versions $(k_1, v_1', (t_1, r_2), \emptyset)$ and $(k_2, v_2', (t_2, r_2), \{(k_1, t_1, r_2)\})$; and (3) finally fetches $(k_2, v_2', (t_2, r_2), \{(k_1, t_1, r_2)\})$ for k_2 (highlighted). As such, the versions fetched for $\{k_1, k_2\}$ are not causally consistent: $(k_2, v_2', (t_2, r_2), \{(k_1, t_1, r_2), \}$, depends on a k_1 version with time-stamp (t_1, r_2) which is bigger than that fetched for k_1, namely (t_1, r_1). To remedy this, after the first round of optimistic reads, r_1 collects all dependency sets and uses them as an lower-bound in the second round to re-fetch the most recent version (with the biggest time-stamp) of each key from the dependency sets. For instance, in Fig. 7c replica r_1 re-fetches the newer version $(k_1, v_1', (t_1, r_2), \emptyset)$ for k_1, as depicted in Fig. 7d. The snapshot obtained after the second round is thus causally consistent. Finally, the snapshot and the dependencies of each version read are sent to the client, and subsequently added to the client context.

To prove that COPS satisfies causal consistency, we encode the state of the system (comprising the state of all replicas and clients) as a configuration in our operational semantics. As each replica stores a set of versions (for each key in COPS) and their dependencies, we can project the state of COPS replicas into a kv-store by mapping COPS versions into our kv-store versions. The writer of a mapped version is uniquely determined by the time-stamp of the corresponding COPS version. The reader set of the mapped version can be recovered by annotating read-only transactions. For example, the COPS state in Fig. 7b can be encoded as the kv-store depicted in Fig. 7e. Similarly, as the context of a client cl identifies the set of COPS versions that cl sees, we can project COPS client contexts to our client views over kv-stores. For example, the context of cl after committing its first single-write transaction in (cops-cl) is encoded as the client view depicted in Fig. 7f.

We next map the execution of a COPS transaction into an ET_{CC} reduction between the configurations obtained from encoding the COPS states before and after executing the transaction. Note that existing verification techniques [Cerone et al. 2015a; Crooks et al. 2017] require examining the *entire* sequence of operations of a protocol to show that it implements a consistency model. By contrast, we only need to look at how the system evolves after a *single* transaction is executed. In particular, we check the client views (obtained from COPS client contexts) over kv-stores after executing a single transaction as follows. Intuitively, we observe that when a COPS client cl executes a transaction then: (1) the cl context grows, and thus we obtain a more up-to-date view of the associated kv-store; i.e. vshift$_{MR}$ holds; (2) the cl context always includes the time-stamp of the versions written by itself, and thus the corresponding client view always includes the versions cl has written; i.e. vshift$_{RYW}$ holds; and (3) the cl context always contains the dependencies corresponding to versions it has either written or read from other transactions, and thus the corresponding client view is always closed-down with respect to the relation $SO \cup WR_{\mathcal{K}}$; i.e. closed$(\mathcal{K}, u, R_{CC})$)

holds. As such, from the definition of CC in Fig. 3 we know that COPS satisfies causal consistency (CC). We refer the reader to §H.1 for further details and the full soundness proof.

7 CONCLUSIONS AND RELATED WORK

We have introduced a simple interleaving semantics for atomic transactions, based on a global, centralised kv-stores and partial client views. It is expressive enough to capture the anomalous behaviour of many weak consistency models. We have demonstrated that our semantics can be used to both verify protocols of distributed databases and analyse client programs.

We have defined a large variety of consistency models for kv-stores based on execution tests, and have shown these models to be equivalent to well-known declarative consistency models for dependency graphs and abstract executions. We do not know of an appropriate consistency model that we cannot express with our semantics, bearing in mind the constraints that our transactions satisfy snapshot property and the last-write-wins policy. We have identified a new consistency model, called weak snapshot isolation, which lies between PSI and SI and inherits many of the good properties of SI. We have shown that examples are robust against WSI. We would need to provide an implementation of this model to justify it in full in the future. We have proved the correctness of two real-world protocols employed by distributed databases: COPS [Lloyd et al. 2011], a protocol for replicated databases that satisfies causal consistency; and Clock-SI [Du et al. 2013], a protocol for partitioned databases that satisfies snapshot isolation. We have also demonstrated the usefulness of our framework for proving invariant properties: the robustness of simple transactional libraries against different consistency models.

In future, we aim to extend our framework to handle other weak consistency models. For example, we believe that, by introducing promises in the style of [Kang et al. 2017], we can capture consistency models such as *Read Committed*. We also plan to validate further the usefulness of our framework by: verifying other well-known protocols of distributed databases, e.g. Eiger [Lloyd et al. 2013], Wren [Spirovska et al. 2018] and Red-Blue [Li et al. 2012]; exploring robustness results for OLTP workloads such as TPC-C [TPCC 1992] and RUBiS [RUBIS 2008]; and exploring other program analysis techniques such as transaction chopping [Cerone et al. 2015b; Shasha et al. 1995], invariant checking [Gotsman et al. 2016; Zeller 2017] and program logics [Kaki et al. 2017].

Related Work. In the introduction (§ 1), we highlight three general operational semantics for distributed transactional databases. We discuss these semantics in more detail here, and also give some additional related work on program analysis.

Kaki et al. [2017] propose an operational semantics of SQL transactional programs under the consistency models given by the standard ANSI/SQL isolation levels [Berenson et al. 1995]. In their framework, transactions work on a local copy of the global state of the system, and the local effects of a transaction are committed to the system state when it terminates. Because state changes are made immediately available to all clients of a system, this model is not suitable to capture weak consistency models such as PSI or CC. They introduce a program logic and prototype verification tool for reasoning about client programs. However, their definitions of consistency models are not validated against previously known formal definitions.

Nagar and Jagannathan [2018] propose an operational semantics for weak consistency based on abstract executions. Their semantics is parametric in the declarative definition of a consistency model. They introduce a tool for checking the robustness of transactional libraries. They focus on consistency models with snapshot property, but confusingly allow the interleaving of fine-grained operations between transactions. This results in an unnecessary explosion of the space of traces generated by the program. In our semantics, the interleaving is between transactions.

Crooks et al. [2017] propose a state-based formal framework for weak consistency models that employs concepts similar to execution tests and views, called commit tests and read states respectively. They prove that consistency models previously thought to be different are in fact equivalent in their semantics. They capture a wide range of consistency models including read committed which we cannot do. In their semantics, one-step trace reduction is determined by the whole previous history of the trace. In contrast, our reduction step only depends on the current configuration (kv-store and view). They do not consider program analysis. Their notion of commit tests and read states requires the knowledge of information that is not known to clients of the system, i.e. the total order of system changes that happened in the database prior to committing a transaction. For this reason, we believe that their framework is not suitable for the development of techniques for analysing client programs.

Doherty et al. [2019] develop an operational semantics for release-acquire fragment of C11 memory model, an variant of causal consistency. Their semantics is based on an variant of dependency graph where nodes and edges are tailored for C11 operations. They introduce per-thread observations, and they must be compatible for executing next operations; this is similar to our views and execution tests. We believe we can also model release-acquire fragment of C11.

Several other works have focused on program analysis for transactional systems. Dias et al. [2012] developed a separation logic for the robustness of applications against SI. Fekete et al. [2005] derived a static analysis check for SI based on dependency graph. Bernardi and Gotsman [2016] developed a static analysis check for several consistency models with snapshot property. Beillahi et al. [2019] developed a tool based on Lipton's reduction [Lipton 1975] for checking robustness against SI. Cerone et al. [2017] investigated the relationship between abstract executions and dependency graphs from an algebraic perspective, and applied it to infer robustness checks for several consistency models.

REFERENCES

Atul Adya. 1999. *Weak Consistency: A Generalized Theory and Optimistic Implementations for Distributed Transactions*. Ph.D. Dissertation. Massachusetts Institute of Technology, Cambridge, MA, USA.

Atul Adya, Barbara Liskov, and Patrick E. O'Neil. 2000. Generalized Isolation Level Definitions. In *ICDE*.

M. Alomari, M. Cahill, A. Fekete, and U. Rohm. 2008. The Cost of Serializability on Platforms That Use Snapshot Isolation. In *2008 IEEE 24th International Conference on Data Engineering*. 576–585. https://doi.org/10.1109/ICDE.2008.4497466

Masoud Saeida Ardekani, Pierre Sutra, and Marc Shapiro. 2014. G-DUR: A Middleware for Assembling, Analyzing, and Improving Transactional Protocols. In *Proceedings of the 15th International Middleware Conference (Middleware'14)*. ACM, New York, NY, USA, 13–24. https://doi.org/10.1145/2663165.2663336

Peter Bailis, Alan Fekete, Ali Ghodsi, Joseph M. Hellerstein, and Ion Stoica. 2014. Scalable Atomic Visibility with RAMP Transactions. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*. 27–38.

Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. 2019. Checking Robustness Against Snapshot Isolation. *CoRR* abs/1905.08406 (2019). arXiv:1905.08406 http://arxiv.org/abs/1905.08406

Hal Berenson, Phil Bernstein, Jim Gray, Jim Melton, Elizabeth O'Neil, and Patrick O'Neil. 1995. A Critique of ANSI SQL Isolation Levels. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data (SIGMOD'95)*. ACM, 1–10. https://doi.org/10.1145/223784.223785

Giovanni Bernardi and Alexey Gotsman. 2016. Robustness against Consistency Models with Atomic Visibility. In *Proceedings of the 27th International Conference on Concurrency Theory*. 7:1–7:15. https://doi.org/10.4230/LIPIcs.CONCUR.2016.7

Carsten Binnig, Stefan Hildenbrand, Franz Färber, Donald Kossmann, Juchang Lee, and Norman May. 2014. Distributed Snapshot Isolation: Global Transactions Pay Globally, Local Transactions Pay Locally. *The VLDB Journal* 23, 6 (December 2014), 987–1011. https://doi.org/10.1007/s00778-014-0359-9

Sebastian Burckhardt. 2014. Principles of Eventual Consistency. *Found. Trends Program. Lang.* 1, 1-2 (October 2014), 1–150. https://doi.org/10.1561/2500000011

Sebastian Burckhardt, Manuel Fahndrich, Daan Leijen, and Mooly Sagiv. 2012. Eventually Consistent Transactions. In *Proceedings of the 21nd European Symposium on Programming*. Springer.

Sebastian Burckhardt, Alexey Gotsman, Hongseok Yang, and Marek Zawirski. 2014. Replicated Data Types: Specification, Verification, Optimality. In *Proceedings of the 41st Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'14)*. ACM, 271–284.

Andrea Cerone, Giovanni Bernardi, and Alexey Gotsman. 2015a. A Framework for Transactional Consistency Models with Atomic Visibility. In *Proceedings of the 26th International Conference on Concurrency Theory (Leibniz International Proceedings in Informatics (LIPIcs))*, Luca Aceto and David de Frutos-Escrig (Eds.), Vol. 42. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 58–71. https://doi.org/10.4230/LIPIcs.CONCUR.2015.58

Andrea Cerone and Alexey Gotsman. 2016. Analysing Snapshot Isolation. In *Proceedings of the 2016 ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC'16)*. ACM, 55–64. https://doi.org/10.1145/2933057.2933096

Andrea Cerone, Alexey Gotsman, and Hongseok Yang. 2015b. Transaction Chopping for Parallel Snapshot Isolation. In *Proceedings of the 29th International Symposium on Distributed Computing*. 388–404.

Andrea Cerone, Alexey Gotsman, and Hongseok Yang. 2017. Algebraic Laws for Weak Consistency. In *Proceedings of the 27th International Conference on Concurrency Theory (Leibniz International Proceedings in Informatics (LIPIcs))*, Roland Meyer and Uwe Nestmann (Eds.), Vol. 85. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 26:1–26:18. https://doi.org/10.4230/LIPIcs.CONCUR.2017.26

Natacha Crooks, Youer Pu, Lorenzo Alvisi, and Allen Clement. 2017. Seeing is Believing: A Client-Centric Specification of Database Isolation. In *Proceedings of the 2017 ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC'17)*. ACM, New York, NY, USA, 73–82. https://doi.org/10.1145/3087801.3087802

Ricardo J. Dias, Dino Distefano, João Costa Seco, and João M. Lourenço. 2012. Verification of Snapshot Isolation in Transactional Memory Java Programs. In *ECOOP 2012 – Object-Oriented Programming*, James Noble (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 640–664.

Simon Doherty, Brijesh Dongol, Heike Wehrheim, and John Derrick. 2019. Verifying C11 Programs Operationally. In *Proceedings of the 24th Symposium on Principles and Practice of Parallel Programming (PPoPP '19)*. ACM, New York, NY, USA, 355–365. https://doi.org/10.1145/3293883.3295702

Jiaqing Du, Sameh Elnikety, and Willy Zwaenepoel. 2013. Clock-SI: Snapshot Isolation for Partitioned Data Stores Using Loosely Synchronized Clocks. In *Proceedings of the 32nd Leibniz International Proceedings in Informatics (LIPIcs) (SRDS'13)*. IEEE Computer Society, Washington, DC, USA, 173–184. https://doi.org/10.1109/SRDS.2013.26

Alan Fekete, Dimitrios Liarokapis, Elizabeth O'Neil, Patrick O'Neil, and Dennis Shasha. 2005. Making Snapshot Isolation Serializable. *ACM Transactions on Database Systems* 30, 2 (June 2005), 492–528. https://doi.org/10.1145/1071610.1071615

Alexey Gotsman, Hongseok Yang, Carla Ferreira, Mahsa Najafzadeh, and Marc Shapiro. 2016. 'Cause I'm Strong Enough: Reasoning about Consistency Choices in Distributed Systems. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Rastislav Bodik and Rupak Majumdar (Eds.). ACM, 371–384.

Gowtham Kaki, Kartik Nagar, Mahsa Najafzadeh, and Suresh Jagannathan. 2017. Alone Together: Compositional Reasoning and Inference for Weak Isolation. *Proceedings of the ACM on Programming Languages* 2, POPL, Article 27 (December 2017), 34 pages. https://doi.org/10.1145/3158115

Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A Promising Semantics for Relaxed-memory Concurrency. In *Proceedings of the 44th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'17)*. ACM, New York, NY, USA, 175–189. https://doi.org/10.1145/3009837.3009850

Cheng Li, Daniel Porto, Allen Clement, Johannes Gehrke, Nuno Preguiça, and Rodrigo Rodrigues. 2012. Making Geo-Replicated Systems Fast as Possible, Consistent when Necessary. In *Proceedings of the 10th Symposium on Operating Systems Design and Implementation*. 265–278.

Richard J. Lipton. 1975. Reduction: A Method of Proving Properties of Parallel Programs. *Commun. ACM* 18, 12 (December 1975), 717–721. https://doi.org/10.1145/361227.361234

Si Liu, Peter Csaba Ölveczky, Keshav Santhanam, Qi Wang, Indranil Gupta, and José Meseguer. 2018. ROLA: A New Distributed Transaction Protocol and Its Formal Analysis. In *Fundamental Approaches to Software Engineering*, Alessandra Russo and Andy Schürr (Eds.). Springer, Cham, 77–93.

Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. 2011. Don't Settle for Eventual: Scalable Causal Consistency for Wide-area Storage with COPS. In *Proceedings of the 23rd ACM Symposium on Operating Systems Principles (SOSP'11)*. ACM, 401–416.

Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. 2013. Stronger Semantics for Low-Latency Geo-Replicated Storage. In *Presented as part of the 10th USENIX Symposium on Networked Systems Design and Implementation*. USENIX, Lombard, IL, 313–328. https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/lloyd

Kartik Nagar and Suresh Jagannathan. 2018. Automated Detection of Serializability Violations Under Weak Consistency. In *Proceedings of the 29th International Conference on Concurrency Theory*. 41:1–41:18. https://doi.org/10.4230/LIPIcs.CONCUR.2018.41

Azalea Raad, Ori Lahav, and Viktor Vafeiadis. 2018. On Parallel Snapshot Isolation and Release/Acquire Consistency. In *Proceedings of the 27th European Symposium on Programming*, Amal Ahmed (Ed.). Lecture Notes in Computer Science,

Cham, 940–967.

RUBiS 2008. The RUBiS benchmark. https://rubis.ow2.org/index.html.

M. Saeida Ardekani, P. Sutra, and M. Shapiro. 2013. Non-monotonic Snapshot Isolation: Scalable and Strong Consistency for Geo-replicated Transactional Systems. In *Proceedings of the 32^{nd} Leibniz International Proceedings in Informatics (LIPIcs)*. 163–172.

Dennis Shasha, Francois Llirbat, Eric Simon, and Patrick Valduriez. 1995. Transaction Chopping: Algorithms and Performance Studies. *ACM Transactions on Database Systems* 20, 3 (September 1995), 325–363. https://doi.org/10.1145/211414.211427

KC Sivaramakrishnan, Gowtham Kaki, and Suresh Jagannathan. 2015. Declarative Programming over Eventually Consistent Data Stores. In *Proceedings of the 36^{th} ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '15)*. ACM, New York, NY, USA, 413–424. https://doi.org/10.1145/2737924.2737981

Yair Sovran, Russell Power, Marcos K. Aguilera, and Jinyang Li. 2011. Transactional Storage for Geo-replicated Systems. In *Proceedings of the 23^{rd} ACM Symposium on Operating Systems Principles (SOSP'11)*. ACM, New York, NY, USA, 385–400. https://doi.org/10.1145/2043556.2043592

Kristina Spirovska, Diego Didona, and Willy Zwaenepoel. 2018. Wren: Nonblocking Reads in a Partitioned Transactional Causally Consistent Data Store. In *Proceedings of the 48^{th} Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'18)*. 1–12. https://doi.org/10.1109/DSN.2018.00014

TPCC 1992. The TPC-C benchmark. http://www.tpc.org/tpcc/.

Werner Vogels. 2009. Eventually Consistent. *Commun. ACM* 52, 1 (January 2009), 40–44.

Peter Zeller. 2017. Testing Properties of Weakly Consistent Programs with Repliss. In *Proceedings of the 3^{rd} International Workshop on Principles and Practice of Consistency for Distributed Data (PaPoC'17)*. ACM, New York, NY, USA, Article 3, 5 pages. https://doi.org/10.1145/3064889.3064893

A OPERATIONAL SEMANTICS ON KV-STORES

Definition A.1 (Multi-version Key-value Stores). Assume a countably infinite set of *keys* KEY $\ni k$, and a countably infinite set of *values* VAL $\ni v$, including an *initialisation value* v_0. The set of *versions*, VERSION $\ni v$, is: VERSION $\triangleq$ VAL $\times$ TRANSID $\times$ $\mathcal{P}$ (TRANSID$_0$). A *kv-store* is a function $\mathcal{K}$: KEY $\rightarrow$ List(VERSION), where List(VERSION) $\ni \mathcal{V}$ is the set of lists of versions VERSION. Well-formed key-values store satisfy:

$$\forall k, i, j. \ \mathsf{rs}(\mathcal{K}(k, i)) \cap \mathsf{rs}(\mathcal{K}(k, j)) \neq \emptyset \vee \mathsf{w}(\mathcal{K}(k, i)) = \mathsf{w}(\mathcal{K}(k, j)) \Rightarrow i = j \tag{1.1}$$

$$\forall k. \ \mathcal{K}(k, 0) = (v_0, t_0, \_) \tag{1.2}$$

$$\forall k, cl, i, j, n, m. \ t_{cl}^n = \mathsf{w}(\mathcal{K}(k, i)) \wedge t_{cl}^m \in \{\mathsf{w}(\mathcal{K}(k, j))\} \cup \mathsf{rs}(\mathcal{K}(k, i)) \Rightarrow n < m \tag{1.3}$$

The full semantics is in Fig. 8 and the full definition of consistency models is in Fig. 9.

$$\rightarrow : ((\textsc{Stack} \times \textsc{Snapshot} \times \textsc{Fp}) \times \textsc{Trans}) \times ((\textsc{Stack} \times \textsc{Snapshot} \times \textsc{Fp}) \times \textsc{Trans})$$

TPrimitive
$$\frac{(s, ss) \overset{\mathsf{T}_p}{\rightsquigarrow} (s', ss') \qquad o = \mathsf{op}(s, ss, \mathsf{T}_p)}{(s, ss, \mathcal{F}), \mathsf{T}_p \rightarrow (s', ss', \mathcal{F} \lessdot o), \mathsf{skip}}$$

TChoice
$$\frac{i \in \{1, 2\}}{(s, ss, \mathcal{F}), \mathsf{T}_1 + \mathsf{T}_2 \rightarrow (s, ss, \mathcal{F}), \mathsf{T}_i}$$

TIter
$$\frac{}{(s, ss, \mathcal{F}), \mathsf{T}^* \rightarrow (s, ss, \mathcal{F}), \mathsf{skip} + (\mathsf{T}; \mathsf{T}^*)}$$

TSeqSkip
$$\frac{}{(s, ss, \mathcal{F}), \mathsf{skip}; \mathsf{T} \rightarrow (s, ss, \mathcal{F}), \mathsf{T}}$$

TSeq
$$\frac{(s, ss, \mathcal{F}), \mathsf{T}_1 \rightarrow (s', ss', \mathcal{F}'), \mathsf{T}_1'}{(s, ss, \mathcal{F}), \mathsf{T}_1; \mathsf{T}_2 \rightarrow (s', ss', \mathcal{F}'), \mathsf{T}_1'; \mathsf{T}_2}$$

$$\rightarrow : \textsc{Client} \times ((\textsc{Kvs} \times \textsc{Views} \times \textsc{Stack}) \times \textsc{Cmd}) \times \text{ET} \times \textsc{Labels} \times ((\textsc{Kvs} \times \textsc{Views} \times \textsc{Stack}) \times \textsc{Cmd})$$

CAtomicTrans
$$\frac{\begin{array}{c} u \sqsubseteq u'' \qquad ss = \mathsf{snapshot}(\mathcal{K}, u'') \qquad (s, ss, \emptyset), \mathsf{T} \rightarrow^* (s', \_, \mathcal{F}), \mathsf{skip} \qquad t \in \mathsf{nextTid}(cl, \mathcal{K}) \\ \mathcal{K}' = \mathsf{update}(\mathcal{K}, u'', \mathcal{F}, t) \qquad \mathsf{can\text{-}commit}_{\text{ET}}(\mathcal{K}, u'', \mathcal{F}) \qquad \mathsf{vshift}_{\text{ET}}(\mathcal{K}, u'', \mathcal{K}', u') \end{array}}{cl \vdash (\mathcal{K}, u, s), [\mathsf{T}] \xrightarrow{(cl, u'', \mathcal{F})}_{\text{ET}} (\mathcal{K}', u', s'), \mathsf{skip}}$$

CPrimitive
$$\frac{s \overset{\mathsf{C}_p}{\rightsquigarrow} s'}{cl \vdash (\mathcal{K}, u, s), \mathsf{C}_p \xrightarrow{(cl, \iota)}_{\text{ET}} (\mathcal{K}, u, s'), \mathsf{skip}}$$

CChoice
$$\frac{i \in \{1, 2\}}{cl \vdash (\mathcal{K}, u, s), \mathsf{C}_1 + \mathsf{C}_2 \xrightarrow{(cl, \iota)}_{\text{ET}} (\mathcal{K}, u, s), \mathsf{C}_i}$$

CIter
$$\frac{}{cl \vdash (\mathcal{K}, u, s), \mathsf{C}^* \xrightarrow{(cl, \iota)}_{\text{ET}} (\mathcal{K}, u, s), \mathsf{skip} + (\mathsf{C}; \mathsf{C}^*)}$$

CSeqSkip
$$\frac{}{cl \vdash (\mathcal{K}, u, s), \mathsf{skip}; \mathsf{C} \xrightarrow{(cl, \iota)}_{\text{ET}} (\mathcal{K}, u, s), \mathsf{C}}$$

CSeq
$$\frac{cl \vdash (\mathcal{K}, u, s), \mathsf{C}_1 \xrightarrow{(cl, \iota)}_{\text{ET}} (\mathcal{K}, u', s'), \mathsf{C}_1'}{cl \vdash (\mathcal{K}, u, s), \mathsf{C}_1; \mathsf{C}_2 \xrightarrow{(cl, \iota)}_{\text{ET}} (\mathcal{K}, u', s'), \mathsf{C}_1'; \mathsf{C}_2}$$

$$\rightarrow : (\textsc{Conf} \times \textsc{CEnv} \times \textsc{Prog}) \times \text{ET} \times \textsc{Label} \times (\textsc{Conf} \times \textsc{CEnv} \times \textsc{Prog})$$

PProg
$$\frac{cl \vdash (\mathcal{K}, \mathcal{U}(cl), \mathcal{E}(cl)), \mathsf{P}(cl), \xrightarrow{\lambda}_{\text{ET}} (\mathcal{K}', u', s'), \mathsf{C}'}{(\mathcal{K}, \mathcal{U}, \mathcal{E}), \mathsf{P} \xrightarrow{\lambda}_{\text{ET}} (\mathcal{K}', \mathcal{U}[cl \mapsto u'], \mathcal{E}[cl \mapsto s']), \mathsf{P}[cl \mapsto \mathsf{C}']}$$

Fig. 8. Operational Semantics on Key-value Store

ET	can-commit$_{\mathrm{ET}}(\mathcal{K}, u, \mathcal{F})$	Closure Relation (where applicable)	vshift$_{\mathrm{ET}}(\mathcal{K}, u, \mathcal{K}', u')$
MR	true		$u \sqsubseteq u'$
RYW	true		$\forall t \in \mathcal{K}' \setminus \mathcal{K}.\ \forall k, i.$ $\mathrm{w}(\mathcal{K}'(k,i)) \xrightarrow{\mathrm{SO}^?} t \Rightarrow i \in u'(k)$
MW	closed$(\mathcal{K}, u, R_{\mathrm{MW}})$	$R_{\mathrm{MW}} \triangleq \mathrm{SO} \cap \mathrm{WW}_{\mathcal{K}}$	true
WFR	closed$(\mathcal{K}, u, R_{\mathrm{WFR}})$	$R_{\mathrm{WFR}} \triangleq \mathrm{WR}_{\mathcal{K}}; (\mathrm{SO} \cup \mathrm{RW}_{\mathcal{K}})^?$	true
CC	closed$(\mathcal{K}, u, R_{\mathrm{CC}})$	$R_{\mathrm{CC}} \triangleq \mathrm{SO} \cup \mathrm{WR}_{\mathcal{K}}$	vshift$_{\mathrm{MR} \cap \mathrm{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
UA	closed$(\mathcal{K}, u, R_{\mathrm{UA}})$	$R_{\mathrm{UA}} \triangleq \bigcup_{(w,k,\_) \in \mathcal{F}} \mathrm{WW}_{\mathcal{K}}^{-1}(k)$	true
PSI	closed$(\mathcal{K}, u, R_{\mathrm{PSI}})$	$R_{\mathrm{PSI}} \triangleq R_{\mathrm{UA}} \cup R_{\mathrm{CC}} \cup \mathrm{WW}_{\mathcal{K}}$	vshift$_{\mathrm{MR} \cap \mathrm{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
CP	closed$(\mathcal{K}, u, R_{\mathrm{CP}})$	$R_{\mathrm{CP}} \triangleq \mathrm{SO}; \mathrm{RW}_{\mathcal{K}}^? \cup \mathrm{WR}_{\mathcal{K}}; \mathrm{RW}_{\mathcal{K}}^? \cup \mathrm{WW}_{\mathcal{K}}$	vshift$_{\mathrm{MR} \cap \mathrm{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
WSI	closed$(\mathcal{K}, u, R_{\mathrm{WSI}})$	$R_{\mathrm{SI}} \triangleq R_{\mathrm{UA}} \cup R_{\mathrm{CP}}$	vshift$_{\mathrm{MR} \cap \mathrm{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
SI	closed$(\mathcal{K}, u, R_{\mathrm{SI}})$	$R_{\mathrm{SI}} \triangleq R_{\mathrm{UA}} \cup R_{\mathrm{CP}} \cup (\mathrm{WW}_{\mathcal{K}}; \mathrm{RW}_{\mathcal{K}})$	vshift$_{\mathrm{MR} \cap \mathrm{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$
SER	closed$(\mathcal{K}, u, R_{\mathrm{SER}})$	$R_{\mathrm{SER}} \triangleq \mathrm{WW}^{-1}$	true
SER*	closed$(\mathcal{K}, u, R_{\mathrm{SER}^*})$	$R_{\mathrm{SER}^*} \triangleq R_{\mathrm{UA}} \cup \mathrm{SO} \cup \mathrm{WW}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{RW}_{\mathcal{K}}$	vshift$_{\mathrm{MR} \cap \mathrm{RYW}}(\mathcal{K}, u, \mathcal{K}', u')$

Fig. 9. Execution tests of well-known consistency models, where SER* denotes an alternative equivalent SER specification and SO is as given in §3.1.

B RELATIONS TO DEPENDENCY GRAPHS

Dependency graphs were introduced by Adya to define consistency models of transactional databases [Adya 1999]. They are directed graphs consisting of transactions as nodes, each of which is labelled with transaction identifier and a set of read and write operations, and labelled edges between transactions for describing how information flows between nodes. Specifically, a transaction t reads a version for a key k that has been written by another transaction t' (*write-read dependency* WR), overwrites a version of k written by t' (*write-write dependency* WW), or reads a version of k that is later overwritten by t' (*read-write anti-dependency* RW). Note that we have named dependencies in kv-stores after the labelled edges of dependency graph. The main result of this Section shows that kv-stores are in fact isomorphic to dependency graphs, and dependencies in a kv-store naturally translates into a labelled edge in the associated dependency graph.

Definition B.1. A *dependency graph* is a quadruple $\mathcal{G} = (\mathcal{T}, \text{WR}, \text{WW}, \text{RW})$, where

- $\mathcal{T} : \textsc{TransID} \rightharpoonup \mathcal{P}(\textsc{Ops})$ is a partial mapping from transaction identifiers to the set of operations, where there are at most one read operation and one write operation per key, and such that $\mathcal{T}(t_0) = \{(\text{w}, k, v_0 \mid k \in \textsc{Key}\}$; furthermore, $t_0 \in \text{dom}(\mathcal{T})$, and $\mathcal{T}(t_0) = \{(\text{w}, k, v_0) \mid k \in \textsc{Key}\}$,
- $\text{WR} : \textsc{Key} \rightarrow \mathcal{P}(\text{dom}(\mathcal{T}) \times \text{dom}(\mathcal{T}))$ is a function that maps each key k into a relation between transactions, such that for any t, t_1, t_2, k, cl, m, n:
  - if $(\text{r}, k, v) \in \mathcal{T}(t)$, there exists $t' \neq t$ such that $(\text{w}, k, v) \in \mathcal{T}(t')$, and $t' \xrightarrow{\text{WR}(k)} t$,
  - if $t_1 \xrightarrow{\text{WR}(k)} t$ and $t_2 \xrightarrow{\text{WR}(k)} t$, then $t_1 = t_2$.
  - if $t_{cl}^m \xrightarrow{\text{WR}(k)} t_{cl}^n$, then $m < n$.
- $\text{WW} : \textsc{Key} \rightarrow \mathcal{P}(\text{dom}(\mathcal{T}) \times \text{dom}(\mathcal{T}))$ is a function that maps each key into an irreflexive relation between transactions, such that for any t, t', k, cl, m, n,
  - if $t \xrightarrow{\text{WW}(k)} t'$, then $(\text{w}, k, \_) \in \mathcal{T}(t), (\text{w}, k, \_) \in \mathcal{T}(t')$,
  - if $(\text{w}, k, \_) \in \mathcal{T}(t), (\text{w}, k, \_) \in \mathcal{T}(t')$, then either $t = t'$, $t \xrightarrow{\text{WW}(k)} t'$, or $t' \xrightarrow{\text{WW}(k)} t$; furthermore, if $t = t_0$, then it must be the case that $t \xrightarrow{(\text{WW}(k))} t'$,
  - if $t_{cl}^m \xrightarrow{\text{WW}(k)} t_{cl}^n$, then $m < n$,
- $\text{RW} : \textsc{Key} \rightarrow \mathcal{P}(\text{dom}(\mathcal{T}) \times \text{dom}(\mathcal{T}))$ is defined by letting $t \xrightarrow{\text{RW}(k)} t'$ if and only if $t'' \xrightarrow{\text{WR}(k)} t, t'' \xrightarrow{\text{WW}(k)} t'$ for some t''.

Let $\textsc{Dgraphs}$ be the set of all dependency graphs.

Given a dependency graph $\mathcal{G} = (\mathcal{T}, \text{WR}, \text{WW}, \text{RW})$, we let $\text{WR}_\mathcal{G} = \text{WR}$, and similarly for WW and RW. We also let $\mathcal{T}_\mathcal{G} = \text{dom}(\mathcal{T})$, and write $(l, k, v) \in_\mathcal{G} t$ if $(;, k, v) \in \mathcal{G}(t)$. We often commit an abuse of notation and use WR to denote the relation $\bigcup_{k \in \textsc{Key}} \text{WR}(k)$; a similar notation is adopted for WW, RW. It will always be clear from the context whether the symbol WR refers to a function from keys to relations, or to a relation between transactions.

As stated above, kv-stores are isomorphic to dependency graphs. The proof of this result is the topic of this Section.

THEOREM B.2. *There is a one-to-one map between kv-stores and dependency graphs.*

The proof structure of Theorem B.2 is standard in its nature. We first how to encode a kv-store into a dependency graph. Then we show how to encode a dependency graph into a kv-store. Finally, we prove that the two constructions are one the inverse of the other: if we convert a kv-store $\mathcal{K}$ into a dependency graph $\mathcal{G}_\mathcal{K}$, then back to a kv-store $\mathcal{K}_{\mathcal{G}_\mathcal{K}}$, we obtain the initial kv-store.

To convert a kv-store $\mathcal{K}$ into a dependency graph, we first define how to extract a fingerprint of a transaction identifier t appearing in $\mathcal{K}$:

Definition B.3. Let $\mathcal{K}$ be a kv-store. For any transaction identifier t, we define $\mathcal{F}_{\mathcal{K}}(t)$ to be the smallest set such that whenever $\mathcal{K}(k, \_) = (v, t, \_)$ then $(\mathsf{w}, k, v) \in t$, and whenever $\mathcal{K}(k, \_) = (v, \_, \{t\} \cup \_)$, then $(\mathsf{r}, k, v) \in t$.

PROPOSITION B.4. *For any $\mathcal{K}, t$, the fingerprint $\mathcal{F}_{\mathcal{K}}(t)$ is well defined. That is, whenever $(\mathsf{w}, k, v_1), (\mathsf{w}, k, v_2) \in \mathcal{F}_{\mathcal{K}}(t)$, then $v_1 = v_2$, and whenever $(\mathsf{r}, k, v_1), (\mathsf{r}, k, v_2) \in \mathcal{F}_{\mathcal{K}}(t)$, then $v_1 = v_2$.*

PROOF. Suppose that $(\mathsf{w}, k, v_1), (\mathsf{w}, k, v_2) \in \mathcal{F}_{\mathcal{K}}(t)$ for some *key*, v_1, v_2. That is, there exist two indexes i_1, i_2 such that $\mathcal{K}(k, i_1) = (v_1, t, \_)$, and $\mathcal{K}(k, i_2) = (v_2, t, \_)$. That is, $\mathsf{w}(\mathcal{K}(k, i_1)) = \mathsf{w}(\mathcal{K}(k, i_2))$, and it follows from Def. A.1 that $i_1 = i_2$. In particular, this implies that $v_1 = v_2$.

A similar argument can be used to prove that if $(\mathsf{r}, k, v_1), (\mathsf{r}, k, v_2) \in \mathcal{F}_{\mathcal{K}}(t)$, then $v_1 = v_2$. In this case, in fact, we have that there exist two indexes i_1, i_2 such that $\mathcal{K}(k, i_1) = (v_1, \_, \{t\} \cup \_)$, and $\mathcal{K}(k, i_2) = (v_2, \_, \{t\} \cup \_)$. Equivalently, $t \in \mathsf{rs}(\mathcal{K}(t, i_1)) \cap \mathsf{rs}(\mathcal{K}(t, i_2))$, and from Def. A.1 it must be the case that $i_1 = i_2$, hence $v_1 = v_2$. □

Using Def. B.3, conerting a kv-store $\mathcal{K}$ into a dependency graph is immediate, as the following definition shows:

Definition B.5. Given a kv-store $\mathcal{K}$, the *dependency graph* $\mathcal{G}_{\mathcal{K}} = (\mathscr{T}_{\mathcal{K}}, \mathsf{WR}_{\mathcal{K}}, \mathsf{WW}_{\mathcal{K}}, \mathsf{RW}_{\mathcal{K}})$ is defined by letting $\mathscr{T}_{\mathcal{K}}(t)$ be defined if and only if $\mathcal{F}_{\mathcal{K}}(t) \neq \emptyset$, in which case we let $\mathscr{T}_{\mathcal{K}}(t) = \mathcal{F}_{\mathcal{K}}(t)$. The relations $\mathsf{WR}_{\mathcal{K}}, \mathsf{WW}_{\mathcal{K}}, \mathsf{RW}_{\mathcal{K}}$ are inherited directly from the transactional dependencies defined for $\mathcal{K}$.

Definition B.6. Given a dependency graph $\mathcal{G} = (\mathscr{T}, \mathsf{WR}, \mathsf{WW}, \mathsf{RW})$, we define the kv-store $\mathcal{K}_{\mathcal{G}}$ as follows:

(1) for any transaction $t \in \mathrm{dom}(\mathscr{T})$ such that $(\mathsf{w}, k, v) \in \mathscr{T}(t)$, let $\mathcal{T} = \left\{ t' \;\middle|\; t \xrightarrow{\mathsf{WR}(k)} t' \right\}$, and let
$$v(t, k) = (v, t, \mathcal{T}),$$

(2) For each key k, let $v_k^0 = (v_0, t_0, \mathcal{T}_k^0)$, where $\mathcal{T}_k^0 = \left\{ t \;\middle|\; (\mathsf{r}, k, \_) \in \mathscr{T}(t) \wedge \forall t'.\ \neg(t' \xrightarrow{\mathsf{WR}(k)} t) \right\}$. Let also $\left\{ v_k^i \right\}_{i=1}^n$ be the ordered set of versions such that, for any $i = 1, \cdots, n$, $v_k^i = v(t, k)$ for some t such that $(\mathsf{w}, k, \_) \in \mathscr{T}(t)$, and such that for any $i, j : 1 \le i < j \le n$, $\mathsf{w}(v_k^i) \xrightarrow{\mathsf{WW}(k)} \mathsf{w}(v_k^j)$. Then we let $\mathcal{K}_{\mathcal{G}} = \lambda k.\ \prod_{i=0}^n v_k^i$.

PROPOSITION B.7. *Let $\mathcal{K}$ be a well-formed kv-store. Then $\mathcal{G}_{\mathcal{K}}$ is a well-formed dependency graph.*

PROOF. Let $\mathcal{K}$ be a (well-formed) kv-store. We need to show that $\mathcal{G}_{\mathcal{K}} = (\mathscr{T}_{\mathcal{K}}, \mathsf{WR}_{\mathcal{K}}, \mathsf{WW}_{\mathcal{K}}, \mathsf{RW}_{\mathcal{K}})$ is a dependency graph. As a first step, we show that $\mathcal{G}_{\mathcal{K}}$ is a dependency graph, i.e. it satisfies all the constraints placed by Def. B.1.

- Let $t \in \mathrm{dom}(\mathscr{T}_{\mathcal{K}})$, and suppose that $(\mathsf{r}, k, v) \in \mathscr{T}_{\mathcal{K}}(t)$. We need to prove that there exists a transaction $t' \in \mathrm{dom}(\mathscr{T}_{\mathcal{K}})$ such Because $(\mathsf{r}, k, v) \in \mathscr{T}_{\mathcal{K}}(t)$, there must exist an index $i : 0 \le i < |\mathcal{K}(k)|$ such that $\mathcal{K}(k, i) = (v, t', \{t\} \cup \_)$ for some $t' \in \mathrm{TRANSID}$. In this case we have that $t' \xrightarrow{\mathsf{WR}_{\mathcal{K}}(k)} t$, and by Def. B.3 we have that $(\mathsf{w}, k, v) \in_{\mathcal{G}_{\mathcal{K}}} t'$.

- Let $t \in \mathrm{dom}(\mathscr{T}_{\mathcal{K}})$, and suppose that there exist t_1, t_2 such that $t_1 \xrightarrow{\mathsf{WR}_k(\mathcal{K})} t$, $t_2 \xrightarrow{\mathsf{WR}_k(\mathcal{K})} t$. By Def. B.5, there exist two indexes $i, j : 0 \le i, j < |\mathcal{K}(k)|$, such that $\mathcal{K}(k, i) = (\_, t_1, \{t\} \cup \_)$, $\mathcal{K}(k, j) = (\_, t_2, \{t\} \cup \_)$. We have that $t \in \mathsf{rs}(\mathcal{K}(k, i)) \cap \mathsf{rs}(\mathcal{K}(k, j))$, i.e. $\mathsf{rs}(\mathcal{K}(k, i)) \cap \mathsf{rs}(\mathcal{K}(k, j)) \neq \emptyset$. Because we are assuming that $\mathcal{K}$ is well-formed, then it must be the case that $i = j$. This implies that $t_1 = t_2$.

- Let $cl \in \text{Client}$, $m, n \in \mathbb{N}$ and $k \in \text{Key}$ be such that $t_{cl}^n \xrightarrow{\text{WR}_{\mathcal{K}}(k)} t_{cl}^m$. We prove that $n < m$. By Def. B.5, it must be the case that there exists an index $i : 0 \le i < |\mathcal{K}(k)|$ such that $\mathcal{K}(k, i) = (\_, t_{cl}^n, \{t_{cl}^m\} \cup \_)$. Because $\mathcal{K}$ is well-formed, it must be the case that $n < m$.

- Let $t \in \text{dom}(\mathscr{T}_{\mathcal{K}})$. We show that $\neg(t \xrightarrow{\text{WW}_{\mathcal{K}}} t)$. We prove this fact by contradiction: suppose that $t \xrightarrow{\text{WW}_{\mathcal{K}}(k)} t$ for some key k. By Def. B.5, there must exist two indexes $i, j : 0 \le i < j < |\mathcal{K}(k)|$ such that $t = \text{w}(\mathcal{K}(k, i))$ and $t = \text{w}(\mathcal{K}(k, j))$. Because we are assuming that $\mathcal{K}$ is well-formed, then it must be the case that $i = j$, contradicting the statement that $i < j$.

- Let t, t' be such that $t' \xrightarrow{\text{WW}_k(\mathcal{K})} t$. We must show that $(\text{w}, k, \_) \in \mathscr{T}_{\mathcal{K}}(t')$, and $(\text{w}, k, \_) \in \mathscr{T}_{\mathcal{K}}(t)$. By Def. B.5, there exist $i, j : 0 \le i, j < |\mathcal{K}(k)|$ such that $\mathcal{K}(k, i) = (v', t', \_)$ and $\mathcal{K}(k, j) = (v, t, \_)$, for some $v, v' \in \text{Val}$. Def. B.5 also ensures that $(\text{w}, k, v') \in \mathscr{T}_{\mathcal{K}}(t')$, and $(\text{w}, k, v) \in \mathscr{T}_{\mathcal{K}}(t)$.

- Let t, t' be such that $(\text{w}, k, \_) \in \mathscr{T}_{\mathcal{K}}(t)$ and $(\text{w}, k, \_) \in \mathscr{T}_{\mathcal{K}}(t')$. We need to prove that either $t = t'$, $t \xrightarrow{\text{WW}_{\mathcal{K}}(k)} t'$, or $t' \xrightarrow{\text{WW}_{\mathcal{K}}(k)} t$. By Def. B.5 there exist two indexes $i, j : 0 < i, j < |\mathcal{K}(k)|$ such that $\mathcal{K}(k, i) = (\_, t, \_)$ and $\mathcal{K}(k, j) = (\_, t', \_)$. If $i = j$, then $t = t'$ and there is nothing left to prove. Otherwise, suppose without loss of generality that $i < j$. Then Def. B.5 ensures that $t \xrightarrow{\text{WW}_{\mathcal{K}}(k)} t'$.

- Suppose that $t_{cl}^m \xrightarrow{\text{WW}_{\mathcal{K}}(k)} t_{cl}^n$ for some $cl \in \text{Client}$ and $m, n \in \mathbb{N}$. We need to show that $m < n$. By Def. B.5, because $t_{cl}^m \xrightarrow{\text{WW}_{\mathcal{K}}(k)} t_{cl}^n$ there exist two indexes $i, j : 0 < i, j < |\mathcal{K}(k)|$ such that $\text{w}(\mathcal{K}(k, i)) = t_{cl}^m$ and $\text{w}(\mathcal{K}(k, j)) = t_{cl}^n$. From the assumption that $\mathcal{K}$ is well-formed, it follows that $n < m$.

<div style="text-align: right">□</div>

Next, we show how to convert a dependency graph $\mathcal{G}$ into a kv-store $\mathcal{K}$. The main idea is that any transaction $t \in \mathcal{T}_{\mathcal{G}}$ induces a set of versions, and for each key k, the write-write-dependency order $\text{WW}_{\mathcal{G}}(k)$ determines the order of these versions in $\mathcal{K}_{\mathcal{G}}$.

Definition B.8. Let $\mathcal{G}$ be a dependency graph. Given a key k, let n_k, $\{v_i^k\}_{i=0}^{n_k}$ $\{t_i^k\}_{i=0}^{n_k}$ be such that $\{t_i^k\}_{i=0}^{n_k} = \{t \mid (\text{w}, k, v_i^k) \in_{\mathcal{G}} t\}$, where the index set $\{1, \cdots, n_k\}$ is chosen to be consistent with $\text{WW}_{\mathcal{G}}(k)$: that is, $t_i \xrightarrow{\text{WW}(k)} t_j$ if and only if $i < j$. Given a key k and an index $i = 1, \cdots, n_k$, we also let $\mathcal{T}_i^k = \{t \mid t_i^k \xrightarrow{\text{WR}(k)} \} t$. Note that this set is possibly empty. Finally, we let $\mathcal{K}_{\mathcal{G}}$ be such that, for any $k \in \text{Key}$, $|\mathcal{K}_{\mathcal{G}}(k)| = n_k$, and for any $i = 0, \cdots, n$, $\mathcal{K}_{\mathcal{G}}(k, i) = (v_i^k, t_i^k, \mathcal{T}_i^k)$.

PROPOSITION B.9. *For any dependency graph $\mathcal{G}$, $\mathcal{K}_{\mathcal{G}}$ is a (well-formed) kv-store.*

PROOF. We show that $\mathcal{K}_{\mathcal{G}}$ satisfies all the constraints fromf Def. A.1. Throughout the proof, we adopt the same notation of Def. B.8.

Let $k \in \text{Key}$, and let i, j be such that $\text{rs}(\mathcal{K}_{\mathcal{G}}(k, i)) \cap \text{rs}(\mathcal{K}_{\mathcal{G}}(k, j)) \ne \emptyset$, that is there exists a transaction $t \in \text{rs}(\mathcal{K}_{\mathcal{G}}(k, i)) \cap \text{rs}(\mathcal{K}_{\mathcal{G}}(k, j))$. We show that $i = j$. By definition, $\text{rs}(\mathcal{K}_{\mathcal{G}}(k, i)) = \mathcal{T}_k^i$, and $\text{rs}(\mathcal{K}_{\mathcal{G}}(k, j)) = \mathcal{T}_k^j$. Def. B.8 ensures that $t_i^k \xrightarrow{\text{WR}_{\mathcal{G}}(k)} t$, and $t_j^k \xrightarrow{\text{WR}_{\mathcal{G}}(k)} t$. By definition of dependency graph, it must be the case that $t_i^k = t_j^k$, and because the order of writers transactions in versions in $\mathcal{K}_{\mathcal{G}}(k)$ is defined to be consistent with $\text{WW}_{\mathcal{G}}(k)$, then it must also be the case that $i = j$.

Suppose no that k, i, j are such that $\text{w}(\mathcal{K}_{\mathcal{G}}(k, i)) = \text{w}(\mathcal{K}_{\mathcal{G}}(k, j))$. By definition $\text{w}(\mathcal{K}_{\mathcal{G}}(k, i)) = t_i^k$, and $\text{w}(\mathcal{K}_{\mathcal{G}}(k, j) = t_j^k$. That is, $t_i^k = t_j^k$. Because the order of writer transactions in $\mathcal{K}_{\mathcal{G}}(k)$ is consistent with $\text{WW}_{\mathcal{G}}(k)$, we also have that $i = j$.

Next, note that for any key k, $t_0^k = t_0$. In fact, because $t_0 \in_{\mathcal{G}} (\mathsf{w}, k, v_0)$, we have that $t_0 = t_i^k$ for some $i = 0, \cdots, n_k$. Also, because whenever t is such that $(\mathsf{w}, k, \_) \in_{\mathcal{G}} t$, then it must be the case that $t_0 \xrightarrow{\mathsf{WW}(k)} t$, then it must be the case that $i = 0$. It follows that, for any $k \in \text{Key}$, $\mathcal{K}_{\mathcal{G}}(k, 0) = (v_0, t_0, \_)$.

Finally, suppose that $t_{cl}^n = \mathsf{w}(\mathcal{K}_{\mathcal{G}}(k, i))$, $t_{cl}^m = \mathsf{w}(\mathcal{K}_{\mathcal{G}}(k, j))$ for some i, j such that $i < j$. In this case we have that $t_{cl}^n = t_i^k$, $t_{cl}^m = t_j^k$, and because $i < j$ it must be the case that $t_{cl}^n \xrightarrow{\mathsf{WW}_{\mathcal{G}}(k)} t_{cl}^m$. The definition of dependency graph ensures then that it must $n < m$. A similar argument shows that, if $t_{cl}^n \in \mathsf{w}(\mathcal{K}_{\mathcal{G}}(k, i))$, $t_{cl}^m \in \mathsf{rs}(\mathcal{K}_{\mathcal{G}}(k, i))$, then it must be the case that $t_{cl}^n \xrightarrow{\mathsf{WR}_{\mathcal{G}}(k)} t_{cl}^m$, and therefore $n < m$. □

Finally, we need to show that the two constructions outlined in Def. B.8 and Def. B.6 are one the inverse of the other.

PROPOSITION B.10. *For any kv-store $\mathcal{K}$, $\mathcal{K}_{\mathcal{G}_{\mathcal{K}}} = \mathcal{K}$.*

PROOF. We prove that for any $k \in \text{Key}$, $\mathcal{K}(k) = \mathcal{K}_{\mathcal{G}_{\mathcal{K}}}(k)$.

Let then $k \in \text{Key}$, and suppose that $\mathcal{K}(k) = (v_0, t_0, \mathcal{T}_0) \cdots (v_n, t_n, \mathcal{T}_n)$. By construction, in $(\mathsf{w}, k, v_i) \in_{\mathcal{G}_{\mathcal{K}}} t_i$, and whenever there is a transaction t such that $(\mathsf{w}, k, t) \in_{\mathcal{G}_{\mathcal{K}}} t$, then $t = t_i$ for some $i = 0, \cdots, n$. In particular, we have that $t_0 \xrightarrow{\mathsf{WW}_{\mathcal{K}}(k)} \cdots \xrightarrow{\mathsf{WW}_{\mathcal{K}}(k)} t_n$ completely characterises the write-write-dependency relation $\mathsf{WW}_{\mathcal{K}}(k)$ over $\mathcal{K}_{\mathcal{G}}$ (recall that, by Def. B.8, $\mathsf{WR}_{\mathcal{G}_{\mathcal{K}}} = \mathsf{WR}_{\mathcal{K}}$). By definition, we have that $\mathcal{K}_{\mathcal{G}_{\mathcal{K}}} = (v_0, t_0, \mathcal{T}_0') \cdots (v_n, t_n, \mathcal{T}_n')$.

It remains to prove that, for any $i = 0, \cdots, n$, $\mathcal{T}_i' = \mathcal{T}_i$. For any $i = 0, \cdots, n$, and transaction $t \in \mathcal{T}_i$, Def. B.8 ensures that $t_i \xrightarrow{\mathsf{WR}_{\mathcal{K}}} t$, and by Def. B.6 it must be the case that $t \in \mathcal{T}_i'$. Furthermore, if $t' \in \mathcal{T}_i'$, then from Def. B.6 it must be the case that $t_i \xrightarrow{\mathsf{WR}_{\mathcal{G}_{mkvs}}(k)} t'$, or equivalently $t_i \xrightarrow{\mathsf{WR}_{\mathcal{K}}(k)} t_i'$. (Def. B.8). Then it must be the case that $t' \in \mathcal{T}_i$. □

PROPOSITION B.11. *For any dependency graph $\mathcal{G}$, $\mathcal{G}_{\mathcal{K}_{\mathcal{G}}} = \mathcal{G}$.*

PROOF. The proof of this claim is similar to Prop. B.10, and therefore omitted. □

C OPERATIONAL SEMANTICS OF ABSTRACT EXECUTIONS

Abstract executions are a framework originally introduced in [Burckhardt et al. 2012] to capture the run-time behaviour of clients interacting with a database. In abstract execution, two relations between transactions are introduced: the *visibility* relation establishes when a transaction observes the effects of another transaction; and the *arbitration* relation helps to determine the value of a key k read by a transaction, in the case that the transaction observes multiple updates to k performed by different transactions.

Definition C.1. An abstract execution is a triple $X = (\mathcal{T}, \text{VIS}, \text{AR})$, where

- $\mathcal{T} : \text{TransID} \rightharpoonup \mathcal{P}(\text{Ops})$ is a partial, finite function mapping transaction identifiers to the set of operations that they perform, with $\mathcal{T}(t_0) = \{(\mathsf{w}, k, v_0 \mid k \in \text{Key}\}$,
- $\text{VIS} \subseteq \text{dom}(\mathcal{T}) \times \text{dom}(\mathcal{T})$ is an irreflexive relation, called *visibility*,
- $\text{AR} \subseteq \text{dom}(\mathcal{T}) \times \text{dom}(\mathcal{T})$ is a strict, total order such that $\text{VIS} \subseteq \text{AR}$, and whenever $t_{cl}^n \xrightarrow{\text{AR}} t_{cl}^m$, then $n < m$.

The set of abstract executions is denoted by absExec.

Given an abstract execution $X = (\mathcal{T}, \text{VIS}, \text{AR})$, the notation $\mathcal{T}_X = \mathcal{T}$, $\mathcal{T}_X = \text{dom}(\mathcal{T})$, $\text{VIS}_X = \text{VIS}$ and $\text{AR}_X = \text{AR}$. The session order for a client $\text{SO}_X(cl)$ and then the overall session order SO_X are defined as the following:

$$\text{SO}_X(cl) = \left\{ (t_{cl}^n, t_{cl}^m) \mid cl \in \text{Client} \wedge t_{cl}^n \in \mathcal{T}_X \wedge t_{cl}^m \in \mathcal{T}_X \wedge n < m \right\}$$

and

$$\text{SO}_X = \bigcup_{cl \in \text{Client}} \text{SO}_X(cl)$$

The notation $(\mathsf{r}, k, v) \in_X t$ denotes $(\mathsf{r}, k, v) \in \mathcal{T}_X(t)$, and similarly for write operations $(\mathsf{w}, k, v) \in_X t$. Given an abstract execution X, a transaction $t \in \mathcal{T}_X$, and a key k, the visible writers set $\text{visibleWrites}_X(k, t) \triangleq \left\{ t' \mid t' \xrightarrow{\text{VIS}_X} t \wedge (\mathsf{w}, k, \_) \in_X t' \right\}$.

The operational semantics on abstract executions (Fig. 10) is parametrised in the axiomatic definition $(\text{RP}, \mathcal{A})$ of a consistency model: transitions take the form $(X, \text{CEnv}, \text{P}) \rightarrow_{(\text{RP}, \mathcal{A})} (X', \text{CEnv}', \text{P}')$. An axiomatic definition of a consistency model is given by a pair $(\text{RP}, \mathcal{A})$, where RP is a resolution policy (Def. C.2) and $\mathcal{A}$ is a set of axioms for visibility relation (Def. C.3). An abstract execution X satisfies the consistency model, written $X \models (\text{RP}, \mathcal{A})$ if it satisfies its individual components. The set of abstract executions induced by an axiomatic definition is given by $\text{CM}(\text{RP}, \mathcal{A}) = \{X \mid X \models (\text{RP}, \mathcal{A})\}$.

We first introduce a notation of two abstract executions *agree*. Given two abstract executions $X_1, X_2 \in \text{absExec}$ and set of transactions $\mathcal{T} \subseteq \mathcal{T}_{X_1} \cap \mathcal{T}_{X_2}$, X_1 and X_2 *agree* on $\mathcal{T}$ if and only if for any transactions $t\ t'$ in $\mathcal{T}$:

$$\mathcal{T}_{X_1}(t) = \mathcal{T}_{X_2}(t) \wedge ((t \xrightarrow{\text{VIS}_{X_1}} t') \iff (t \xrightarrow{\text{VIS}_{X_2}} t')) \wedge ((t \xrightarrow{\text{AR}_{X_1}} t') \iff (t \xrightarrow{\text{AR}_{X_2}} t'))$$

Definition C.2. A resolution policy RP is a function $\text{RP} : \text{absExec} \times \mathcal{P}(\text{TransID}) \rightarrow \mathcal{P}(\text{Snapshot})$ such that, for any X_1, X_2 that agree on a subset of transactions $\mathcal{T}$, then $\text{RP}(X_1, \mathcal{T}) = \text{RP}(X_2, \mathcal{T})$. An abstract execution X satisfies the execution policy RP if,

$$\forall t \in \mathcal{T}_X. \ \exists ss \in \text{RP}(X, \text{VIS}_X^{-1}(t)). \ \forall k, v. \ (\mathsf{r}, k, v) \in_X t \Rightarrow ss(k) = v$$

Definition C.3. An axiom A is a function from abstract executions to relations between transactions, $\text{A} : \text{absExec} \rightarrow \mathcal{P}(\text{TransID} \times \text{TransID})$, such that whenever X_1, X_2 agree on a subset of transactions $\mathcal{T}$, then $\text{A}(X_1) \cap (\mathcal{T} \times \mathcal{T}) \subseteq \text{A}(X_2)$.

$\rightarrow$: Client $\times$ ((AbsExecs $\times$ Stack) $\times$ Cmd) $\times$ ET $\times$ Label $\times$ ((AbsExecs $\times$ Stack) $\times$ Cmd)

AAtomicTrans
$$\frac{\mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}} \qquad ss \in \mathsf{RP}(\mathcal{X}, \mathcal{T}) \qquad (s, ss, \emptyset), \mathsf{T} \rightarrow^* (s', \_, \mathcal{F}), \mathtt{skip}}{t \in \mathsf{nextTid}(\mathcal{T}_{\mathcal{X}}, cl) \qquad \mathcal{X}' = \mathsf{extend}(\mathcal{X}, t, \mathcal{T}, \mathcal{F}) \qquad \forall A \in \mathcal{A}. \ \{t' \mid (t', t) \in \mathsf{A}(\mathcal{X}')\} \subseteq \mathcal{T}}{cl \vdash (\mathcal{X}, s), [\mathsf{T}] \xrightarrow{(cl, \mathcal{T}, \mathcal{F})}_{(\mathsf{RP}, \mathcal{A})} (\mathcal{X}', s'), \mathtt{skip}}$$

APrimitive
$$\frac{s \xrightarrow{\mathsf{C}_p} s'}{cl \vdash (\mathcal{X}, s), \mathsf{C}_p \xrightarrow{(cl, \iota)}_{\mathsf{ET}} (\mathcal{X}, s'), \mathtt{skip}}$$

AChoice
$$\frac{i \in \{1, 2\}}{cl \vdash (\mathcal{X}, s), \mathsf{C}_1 + \mathsf{C}_2 \xrightarrow{(cl, \iota)}_{\mathsf{ET}} (\mathcal{X}, s), \mathsf{C}_i}$$

AIter
$$cl \vdash (\mathcal{X}, s), \mathsf{C}^* \xrightarrow{(cl, \iota)}_{\mathsf{ET}} (\mathcal{X}, s), \mathtt{skip} + (\mathsf{C}; \mathsf{C}^*)$$

ASeqSkip
$$cl \vdash (\mathcal{X}, s), \mathtt{skip}; \mathsf{C} \xrightarrow{(cl, \iota)}_{\mathsf{ET}} (\mathcal{X}, s), \mathsf{C}$$

ASeq
$$\frac{cl \vdash (\mathcal{X}, s), \mathsf{C}_1 \xrightarrow{(cl, \iota)}_{\mathsf{ET}} (\mathcal{X}, s'), \mathsf{C}_1'}{cl \vdash (\mathcal{X}, s), \mathsf{C}_1; \mathsf{C}_2 \xrightarrow{(cl, \iota)}_{\mathsf{ET}} (\mathcal{X}, s'), \mathsf{C}_1'; \mathsf{C}_2}$$

$\rightarrow$: (AbsExecs $\times$ CEnv $\times$ Prog) $\times$ ET $\times$ Label $\times$ (AbsExecs $\times$ CEnv $\times$ Prog)

AProg
$$\frac{cl \vdash (\mathcal{X}, \mathcal{E}(cl)), \mathsf{P}(cl) \xrightarrow{\lambda}_{(\mathsf{RP}, \mathcal{A})} (\mathcal{X}', s'), \mathsf{C}'}{(\mathcal{X}, \mathcal{E}), \mathsf{P} \xrightarrow{\lambda}_{(\mathsf{RP}, \mathcal{A})} (\mathcal{X}', \mathcal{E}[cl \mapsto s']), \mathsf{P}[cl \mapsto \mathsf{C}'])}$$

Fig. 10. Operational Semantics on Abstract Executions

Axioms of a consistency model are constraints of the form $\mathsf{A}(\mathcal{X}) \subseteq \mathsf{VIS}_{\mathcal{X}}$. For example, if we require $A(\mathcal{X}) = \mathsf{AR}_{\mathcal{X}}$, then the corresponding axiom is given by $\mathsf{AR}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}$, thus capturing the serialisability of transactions, i.e. this axiom is equivalent to require that $\mathsf{VIS}_{\mathcal{X}}$ is a total order. The requirement on subsets of transactions on which abstract executions agree will be needed later, when we define an operational semantics of transactions where clients can append a new transaction t at the tail of an abstract execution $\mathcal{X}$, which satisfies an axiom A. This requirement ensures that we only need to check that the axiom is A is satisfied by the pre-visibility and pre-arbitration relation of the transaction t in $\mathcal{X}'$. In fact, the resulting abstract execution $\mathcal{X}'$ agrees with $\mathcal{X}$ on the set $\mathcal{T}_{\mathcal{X}}$: in this case we'll note that we can rewrite $\mathsf{A}(\mathcal{X}') = \mathsf{A}(\mathcal{X}') \cap ((\mathcal{T}_{\mathcal{X}} \times \mathcal{T}_{\mathcal{X}})) \cup (\mathcal{T}_{\mathcal{X}} \times \{t\}))$. Then $\mathsf{A}(\mathcal{X}') \cap ((\mathcal{T}_{\mathcal{X}} \times \mathcal{T}_{\mathcal{X}})) \subseteq \mathsf{A}(\mathcal{X}) \cap (\mathcal{T}_{\mathcal{X}} \times \mathcal{T}_{\mathcal{X}}) \subseteq \mathsf{VIS}_{\mathcal{X}} \cap (\mathcal{T}_{\mathcal{X}} \times \mathcal{T}_{\mathcal{X}}) \subseteq \mathsf{VIS}_{\mathcal{X}'}$, hence we only need to check that $\mathsf{A}(\mathcal{X}') \cap (\mathcal{T}_{\mathcal{X}} \times \{t\}) \subseteq \mathsf{VIS}_{\mathcal{X}'}$.

We say that an abstract execution $\mathcal{X}$ satisfies an axiom A, if $\mathsf{A}(\mathcal{X}) \subseteq \mathsf{VIS}_{\mathcal{X}}$. An abstract execution $\mathcal{X}$ satisfies $(\mathsf{RP}, \mathcal{A})$, written $\mathcal{X} \models (\mathsf{RP}, \mathcal{A})$, if the abstract execution $\mathcal{X}$ satisfies RP and $\mathcal{A}$.

Definition C.4 (Abstract executions induced by axiomatic definition). The set of all abstract executions induced by an axiomatic definition, $\mathsf{CM}(\mathsf{RP}, \mathcal{A})$ is defined as $\mathsf{CM}(\mathsf{RP}, \mathcal{A}) \triangleq \{\mathcal{X} \mid \mathcal{X} \models (\mathsf{RP}, \mathcal{A})\}$.

The Fig. 10 presents all rules of the operational semantics of programs based on abstract executions. TheACommit rule is the abstract execution counterpart of rule PCommit for kv-stores. TheACommit models how an abstract execution X evolves when a client wants to execute a transaction whose code is [T]. In the rule, $\mathcal{T}$ is the set of transactions of X that are visible to the client cl that wishes to execute [T]. Such a set of transactions is used to determine a snapshot $ss \in \mathrm{RP}(X, \mathcal{T})$ that the client cl uses to execute the code [T], and obtain a fingerprint $\mathcal{F}$. This fingerprint is then used to extend abstract execution X with a transaction from the set $\mathrm{nextTid}(\mathcal{T}_X, cl)$. Similar PProg rule, the AProg rule in Fig. 10 models multi-clients concurrency in an interleaving fashion. All the rest rules of the abstract operational semantics in Fig. 10 have a similar counterpart in the kv-store semantics.

Note that AAtomicTrans is more general than Rule PAtomicTrans in the kv-store semantics. In the latter, the snapshot of a transaction is uniquely determined from a view of the client, in a way that roughly corresponds to the last write wins policy in the abstract execution framework. In contrast, the snapshot of a transaction used in AAtomicTrans is chosen non-deterministically from those made available to the client by the resolution policy RP, which may not necessarily be last-write-win.

Throughout this report we will work mainly with the *Last Write Wins* resolution policy (Def. C.5). When discussing the operational semantics of transactional programs, we will also introduce the *Anarchic* resolution policy.

Definition C.5. The Last Write Wins resolution policy $\mathrm{RP_{LWW}}$ is defined as $\mathrm{RP_{LWW}}(X, \mathcal{T}) \triangleq \{ss\}$ where

$$ss = \lambda k.\mathrm{let}\ \mathcal{T}_k = (\mathcal{T} \cap \{t \mid (\mathsf{w}, k, \_) \in_X t\})\ \mathrm{in}\ \begin{cases} v_0 & \mathrm{if}\ \mathcal{T}_k = \emptyset \\ v & \mathrm{if}\ (\mathsf{w}, k, v) \in_X \max_{\mathrm{AR}_X}(\mathcal{T}_k) \end{cases}$$

D RELATIONSHIP BETWEEN KV-STORES AND ABSTRACT EXECUTION

D.1 KV-Store to Abstract Executions

We introduce the definition of the dependency graph induced an abstract execution:

Definition D.1. Given an abstract execution $\mathcal{X}$ that satisfies the last write wins policy, the dependency graph $\mathrm{graphOf}(\mathcal{X}) \triangleq (\mathcal{T}_\mathcal{X}, \mathrm{WR}_\mathcal{X}, \mathrm{WW}_\mathcal{X}, \mathrm{RW}_\mathcal{X})$ is defined by letting

- $t \xrightarrow{\mathrm{WR}_\mathcal{X}(k)} t'$ if and only if $t = \max_{\mathrm{AR}_\mathcal{X}}(\mathrm{visibleWrites}_\mathcal{X}(k, t'))$,
- $t \xrightarrow{\mathrm{WW}_\mathcal{X}(k)} t'$ if and only $t, t' \in_\mathcal{X} (\mathrm{w},\ k, \_)$ and $t \xrightarrow{\mathrm{AR}_\mathcal{X}} t'$,
- $t \xrightarrow{\mathrm{RW}_\mathcal{X}(k)} t'$ if and only if either $(\mathrm{r}, k, \_) \in_\mathcal{X} t, (\mathrm{w}, k, \_) \in_\mathcal{X} t'$ and whenever $t'' \xrightarrow{\mathrm{WR}_\mathcal{X}(k)} t$, then $t'' \xrightarrow{\mathrm{WW}_\mathcal{X}(k)} t'$.

Note that each abstract execution $\mathcal{X}$ determines a kv-store $\mathcal{K}_\mathcal{X}$, as a result of Def. D.1 and Theorem B.2. Let $\mathcal{K}$ be the unique kv-store such that $\mathcal{G}_\mathcal{K} = \mathrm{graphOf}(\mathcal{X})$, then $\mathcal{K}_\mathcal{X} = \mathcal{K}$. As we will discuss later in this Section, this mapping $\mathcal{K}_{(\_)}$ is NOT a bijection, in that several abstract executions may be encoded in the same kv-store. Because kv-stores abstract away the total arbitration order of transactions.

Upon the relation $\mathcal{K}_\mathcal{X} = \mathcal{K}$, there is a deeper link between kv-store plus views and abstract exertions. This notion, named *compatibility*, bases on the intuition that clients can make observations over kv-stores and abstract executions, in terms of snapshots.

In kv-stores, observations are snapshots induced by views. While in abstract executions, observations correspond to the snapshots induced by the visible transactions. Note that it is under the condition that the abstract execution adopts $\mathrm{RP_{LWW}}$ resolution policy. This approach is analogous to the one used by operation contexts in [Burckhardt et al. 2014]. Thus, a kv-store $\mathcal{K}$ is *compatible* with an abstract execution $\mathcal{X}$, written $\mathcal{K} \simeq \mathcal{X}$ if any observation made on $\mathcal{K}$ can be replicated by an observation made on $\mathcal{X}$, and vice-versa.

Definition D.2. Given a kv-store $\mathcal{K}$, an abstract execution $\mathcal{X}$ is compatible with $\mathcal{K}$, written $\mathcal{X} \simeq \mathcal{K}$, if and only if there exists a mapping $f : \mathcal{P}(\mathcal{T}_\mathcal{X}) \to \mathrm{Views}(\mathcal{K})$ such that

- for any subset $\mathcal{T} \subseteq \mathcal{T}_\mathcal{X}$, then $\mathrm{RP_{LWW}}(\mathcal{X}, \mathcal{T}) = \{\mathrm{snapshot}(\mathcal{K}, f(\mathcal{T}))\}$;
- for any view $u \in \mathrm{Views}(\mathcal{K})$, there exists a subset $\mathcal{T} \subseteq \mathcal{T}_\mathcal{X}$ such that $f(\mathcal{T}) = u$, and $\mathrm{RP_{LWW}}(\mathcal{X}, \mathcal{T}) = \{\mathrm{snapshot}(\mathcal{K}_\mathcal{X}, u)\}$.

The function $\mathrm{getView}(\mathcal{X}, \mathcal{T})$ defines the view on $\mathcal{K}_\mathcal{X}$ that corresponds to $\mathcal{T}$ as the following:

$$\mathrm{getView}(\mathcal{X}, \mathcal{T}) \triangleq \lambda k.\ \{0\} \cup \{i \mid \mathrm{w}(\mathcal{K}_\mathcal{X}(k, i)) \in \mathcal{T}\}$$

Inversely, the function $\mathrm{visTx}(\mathcal{K}, u)$ converts a view to a set of observable transactions:

$$\mathrm{visTx}(\mathcal{K}, u) \triangleq \{\mathrm{w}(\mathcal{K}(k, i)) \mid k \in \mathrm{Key} \wedge i \in u(k)\}$$

Given getView, visTx, Def. D.2, it follows $\mathcal{X} \simeq \mathcal{K}_\mathcal{X}$ shown in Theorem D.3.

THEOREM D.3. *For any abstract execution $\mathcal{X}$ that satisfies the last write wins policy, $\mathcal{X} \simeq \mathcal{K}_\mathcal{X}$.*

PROOF. Given the function $\mathrm{getView}(\mathcal{X}, \cdot)$ from $\mathcal{P}(\mathcal{T}_\mathcal{X})$ to $\mathrm{Views}(\mathcal{K}_\mathcal{X})$, we prove it satisfies the constraint of Def. D.2. Fix a set of transitions $\mathcal{T}$. By the Prop. D.4, the view $\mathrm{getView}(\mathcal{X}, \mathcal{T})$ on $\mathcal{K}_\mathcal{X}$ is a valid view, that is, $\mathrm{getView}(\mathcal{X}, \mathcal{T}) \in \mathrm{Views}(\mathcal{K}_\mathcal{X})$. Given that it is a valid view, the Prop. D.5 proves:

$$\mathrm{RP_{LWW}}(\mathcal{X}, \mathcal{T}) = \{\mathrm{snapshot}(\mathcal{K}_\mathcal{X}, \mathrm{getView}(\mathcal{X}, \mathcal{T}))\} \tag{4.1}$$

The another way round is more subtle, because $\mathcal{T}$ contains any read only transaction. By Prop. D.6, it is safe to erase read only transactions from $\mathcal{T}$, when calculating the view $\text{getView}(X, \mathcal{T})$. Last, by Prop. D.7, we prove the following:

$$\text{RP}_{\text{LWW}}(X, \mathcal{T}) = \text{snapshot}(\mathcal{K}_X, u) \tag{4.2}$$

By Eq. (4.1) and Eq. (4.2), it follows $X \simeq \mathcal{K}_X$.   □

PROPOSITION D.4 (VALID VIEWS). *For any abstract execution X, and $\mathcal{T} \subseteq \mathcal{T}_X$, $\text{getView}(X, \mathcal{T}) \in$ VIEWS($\mathcal{K}_X$).*

PROOF. Assume an abstract execution X, a set of transactions $\mathcal{T} \subseteq \mathcal{T}_X$, and a key k. By the definition of $\text{getView}(X, \mathcal{T})$, then $0 \in \text{getView}(X, \mathcal{T})(k)$, and $0 \leq i < |\mathcal{K}_X(k)|$ for any index i such that $i \in \text{getView}(X, \mathcal{T})(k)$. Therefore we only need to prove that $\text{getView}(X, \mathcal{T})$ satisfies (atomic). Let $j \in \text{getView}(X, \mathcal{T})(k)$ for some key k, and let $t = \text{w}(\mathcal{K}_X(k, j))$. Let also k', i be such that $\text{w}(\mathcal{K}_X(k', i)) = t$. We need to show that $i \in \text{getView}(X, \mathcal{T})(k')$. Note that it $t = t_0$ then $\text{w}(\mathcal{K}_X(k', i)) = t$ only if $i = 0$, and $0 \in \text{getView}(X, \mathcal{T})(k')$ by definition. Let then $t \neq t_0$. Because $\text{w}(\mathcal{K}_X(k, j)) = t$ and $j \in \text{getView}(X, \mathcal{T})$, then it must be the case that $t \in \mathcal{T}$. Also, because $\text{w}(\mathcal{K}_X(k', i)) = t$, then $(\text{w}, k, \_) \in \mathscr{T}_X(t)$. It follows that there exists an index $i' \in \text{getView}(X, t)(k')$ such that $\text{w}(\mathcal{K}_X(k', i')) = t$. By definition of $\mathcal{K}_X$, if $\text{w}(\mathcal{K}_X(k', i')) = t$, then it must be $i' = i$, and therefore $i \in \text{getView}(X, t)(k')$.   □

PROPOSITION D.5 (VISIBLE TRANSACTIONS TO VIEWS). *For any subset $\mathcal{T} \subseteq \mathcal{T}_X$, $\text{RP}_{\text{LWW}}(X, \mathcal{T}) = \{\text{snapshot}(\mathcal{K}_X, \text{getView}(X, \mathcal{T}))\}$.*

PROOF. Fix $\mathcal{T} \subseteq X$, and let $\{\mathcal{K}\} = \text{RP}_{\text{LWW}}(X, \mathcal{T})$. We prove that, for any $k \in$ KEY, $\mathcal{K}(k) = \text{snapshot}(\text{getView}(X, \mathcal{T}))(k)$. There are two different cases:

(1) $\mathcal{T} \cap \{t \mid (\text{w}, k, \_) \in_X t\} = \emptyset$. In this case $\mathcal{K}(k) = v_0$. We know that $\text{graphOf}(X)$ satisfies all the constraints required by the definition of dependency graph ([Cerone et al. 2017]). Together with Theorem B.2 it follows that $\mathcal{K}_X(k, 0) = (v_0, t_0, \_)$. We prove that $\text{getView}(X, \mathcal{T})(k) = \{0\}$, hence

$$\text{snapshot}(\mathcal{K}_X, \text{getView}(X, \mathcal{T}))(k) = \text{val}(\mathcal{K}_X(k, 0)) = v_0$$

Note that whenever $(\text{w}, k, \_) \in_X t$ for some t, then $t \notin \mathcal{T}$. Therefore, whenever $(v, t, \_) = \mathcal{K}_X(k, i)$ for some $i \geq 0$, then $t \notin \mathcal{T}$.

$$\text{getView}(X, \mathcal{T})(k) = \{0\} \cup \{i \mid \text{w}(\mathcal{K}_X(k, i)) \in \mathcal{T}\} = \{0\} \cup \emptyset = \{0\}$$

(2) Suppose now that $\mathcal{T} \cap \{t \mid (\text{w}, k, \_) \in_X t\} \neq \emptyset$. Let then $t = \max_{\text{AR}_X}(\mathcal{T} \cap \{t \mid (\text{w}, k, \_) \in_X t\})$. Then $(\text{w}, k, v) \in_X t$ for some $v \in$ VAL. Furthermore, $\text{RP}_{\text{LWW}}(X, \mathcal{T})(k) = v$. By definition, $t' \in \mathcal{T} \cap \{t \mid (\text{w}, k, \_) \in_X t\}$, then either $t' = t$ or $t' \xrightarrow{\text{AR}_X} t$. The definition of $\text{graphOf}(X)$ gives that $t' \xrightarrow{\text{WW}_X(k)} t$. Because $(\text{w}, k, v) \in_X t$, then there exists an index $i \geq 0$ such that $\mathcal{K}_X(k, i) = (v, t, \_)$. Furthermore, whenever $\text{w}(k, j) = t'$ for some t' and $j > i$, then it must be the case that $t \xrightarrow{\text{WW}_X(k)} t'$, and because $\text{WW}_X(k)$ is transitive and irreflexive, it must be that $\neg(t' \xrightarrow{\text{WW}_X(k)} t)$ and $t \neq t'$: this implies that $t' \notin \mathcal{T}$. It follows that $\max(\text{getView}(X, \mathcal{T})(k)) = i$, hence $\text{snapshot}(\mathcal{K}_X, \text{getView}(X, \mathcal{T})) = \text{val}(\mathcal{K}_X(k, i)) = v$.

                                                                    □

PROPOSITION D.6 (READ-ONLY TRANSACTIONS ERASING). *Let $u \in$ VIEWS($\mathcal{K}_X$), and let $\mathcal{T} \subseteq \mathcal{T}_X$ be a set of read-only transactions in X. Then $\text{getView}(X, \mathcal{T} \cup \text{visTx}(\mathcal{K}_X, u)) = u$.*

PROOF. Fix a key k. Suppose that $i \in \text{getView}(\mathcal{X}, \mathcal{T} \cup \text{visTx}(\mathcal{K}_{\mathcal{X}}, u))(k)$. By definition, $\mathcal{K}_{\mathcal{X}}(k, j) = (\_, t, \_)$ for some $t \in \mathcal{T} \cup \text{visTx}(\mathcal{K}_{\mathcal{X}}, u)$. Because $\mathcal{T}$ only contains read-only transactions, by definition of $\mathcal{K}_{\mathcal{X}}$ there exists no index j such that $\mathcal{K}_{\mathcal{X}}(k, j) = (\_, t', \_)$ for some $t' \in \mathcal{T}$, hence it must be the case that $t \in \text{visTx}(\mathcal{K}_{\mathcal{X}}, u)$. By definition of visTx, this is possible only if there exist a key k' and an index j such that $\mathcal{K}_{\mathcal{X}}(k', u) = (\_, t, \_)$. Because u is atomic by definition, and because $\mathcal{K}_{\mathcal{X}}(k, i) = (\_, t, \_)$, then we have that $i \in u(k)$.

Now suppose that $i \in u(k)$, and let $\mathcal{K}_{\mathcal{X}}(k, i) = (\_, t, \_)$ for some t. This implies that $(\mathsf{w}, k, \_) \in_{\mathcal{X}} t$. By definition $t \in \text{visTx}(\mathcal{K}_{\mathcal{X}}, u)$, hence $t \in \mathcal{T} \cup \text{visTx}(\mathcal{K}_{\mathcal{X}}, u))$. Because $t \in \mathcal{T} \cup \text{visTx}(\mathcal{K}_{\mathcal{X}}, u)$, then for any key k' such that $(\mathsf{w}, k', \_) \in_{\mathcal{X}} t$, there exists an index $j \in \text{getView}(\mathcal{X}, \mathcal{T} \cup \text{visTx}(\mathcal{K}_{\mathcal{X}}, u))$ $\mathcal{K}(k', j) = (\_, t, \_)$; because kv-stores only allow a transaction to write at most one version per key, then the index j is uniquely determined. In particular, we know that $(\mathsf{w}, k, \_) \in_{\mathcal{X}} t$, and $\mathcal{K}_{\mathcal{X}}(k, i) = (\_, t, \_)$, from which it follows that $i \in \text{getView}(\mathcal{X}, \mathcal{T} \cup \text{visTx}(\mathcal{K}_{\mathcal{X}}, u))(k)$. □

PROPOSITION D.7 (VIEWS TO VISIBLE TRANSACTIONS). *Given a view $u \in \text{VIEWS}(\mathcal{K}_{\mathcal{X}})$, there exists $\mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}}$ such that $\text{getView}(\mathcal{X}, \mathcal{T}) = u$, and $\text{RP}_{\text{LWW}}(\mathcal{X}, \mathcal{T}) = \text{snapshot}(\mathcal{K}_{\mathcal{X}}, u)$.*

PROOF. We only need to prove that, for any $u \in \text{VIEWS}(\mathcal{K}_{\mathcal{X}})$, there exists $\mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}}$ such that $\text{getView}(\mathcal{X}, \mathcal{T}) = u$. Then it follows from Prop. D.5 that $\text{RP}_{\text{LWW}}(\mathcal{X}, \mathcal{T}) = \text{snapshot}(\mathcal{K}_{\mathcal{X}}, u)$. It suffices to choose $\mathcal{T} = \bigcup_{k \in \text{KEY}}(\{\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, i)) \mid i > 0 \land i \in u(k)\})$. Fix a key k, and let $i \in u(k)$. We prove that $i \in \text{getView}(\mathcal{X}, \mathcal{T})$. If $i = 0$, then $i \in \text{getView}(\mathcal{X}, \mathcal{T})$ by definition. Therefore, assume that $i > 0$. Let $t = \mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, i))$. It must be the case that $t \in \mathcal{T}$ and $i \in \text{getView}(\mathcal{X}, \mathcal{T})(k)$.

Next, suppose that $i \in \text{getView}(\mathcal{X}, \mathcal{T})(k)$. We prove that $i \in u(k)$. Note that if $i = 0$, then $i \in u(k)$ because of the definition of views. Let then $i > 0$. Because $i \in \text{getView}(\mathcal{X}, \mathcal{T})(k)$, we have that $\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, i)) \in \mathcal{T}$. Let $t = \mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, i))$. Because $i > 0$, it must be the case that $t \neq t_0$. By definition, $t \in \mathcal{T}$ only if there exists an index j and key k', possibly different from k, such that $\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k', j)) = t$ and $j \in u(k')$. Because $t \neq t_0$ we have that $j > 0$. Finally, because u is atomic by definition, $j \in u(k')$ $\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k', j)) = t = \mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, i))$, then it must be the case that $i \in u(k)$, which concludes the proof. □

D.2 KV-Store Traces to Abstract Execution Traces

To prove our definitions using execution test on kv-stores is sound and complete with respect with the axiomatic definitions on abstract executions (§F), we need to prove trace equivalent between these two models.

In this section, we only consider the trace that does not involve P but only committing fingerprint and view shift. In §E, we will go further and discuss the trace installed with P.

Let $\text{ET}_{\top}$ be the most permissive execution test. That is $\text{ET}_{\top} \vdash (\mathcal{K}, u) \rhd \mathcal{F} : (\mathcal{K}', u')$ such that whenever $u(k) \neq u'(k)$ then either $(\mathsf{w}, k, \_) \in \mathcal{F}$ or $(\mathsf{r}, k, \_) \in \mathcal{F}$. We will relate $\text{ET}_{\top}$-traces to abstract executions that satisfy the last write wins resolution policy, i.e. $(\text{RP}_{\text{LWW}}, \emptyset)$.

To bridge $\text{ET}_{\top}$-traces to abstract executions, The absExec (τ) function converse the trace of $\text{ET}_{\top}$ to set of possible abstract executions (Def. D.8). In fact, for any trace τ and abstract execution $\mathcal{X} \in \text{absExec}(\tau)$, the last configuration of τ is $(\mathcal{K}_{\mathcal{X}}, \_)$ (Prop. D.9). We often use $\mathcal{X}_{\tau}$ for $\mathcal{X} \in \text{absExec}(\tau)$.

Definition D.8. Given a kv-store $\mathcal{K}$, a view u, an initial abstract execution $\mathcal{X}_0 = ([], \emptyset, \emptyset)$, an abstract execution $\mathcal{X}$, a set of transactions $\mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}}$, a transaction identifier t and a set of operations

$\mathcal{F}$, the extend function defined as the follows:

$$\text{extend}(\mathcal{X}, t, \mathcal{T}, \mathcal{F}) \triangleq \begin{cases} \text{undefined} & \text{if } t = t_0 \\ (\mathcal{T}_{\mathcal{X}} \uplus \{t \mapsto \mathcal{F}\}, \text{VIS}', \text{AR}') & \text{if } \dagger \end{cases}$$

$$\dagger \equiv\ t = t_{cl}^n \wedge \text{VIS}' = \text{VIS}_{\mathcal{X}} \uplus \{(t', t) \mid t \in \mathcal{T}\}$$
$$\wedge\ \text{AR}' = \text{AR}_{\mathcal{X}} \uplus \{(t', t) \mid t' \in \mathcal{T}_{\mathcal{X}}\}$$

Given a $\text{ET}_\top$ trace τ, let $\text{lastConf}(\tau)$ be the last configuration appearing in τ. The set of abstract executions $\text{absExec}(\tau)$ is defined as the smallest set such that:

- $\mathcal{X}_0 \in \text{absExec}((\mathcal{K}_0, \mathcal{U}_0))$,
- if $\mathcal{X} \in \text{absExec}(\tau)$, then $\mathcal{X} \in \text{absExec}\left(\tau \xrightarrow{(cl, \varepsilon)}_{\text{ET}_\top} (\mathcal{K}, \mathcal{U})\right)$,
- if $\mathcal{X} \in \text{absExec}(\tau)$, then $\mathcal{X} \in \text{absExec}\left(\tau \xrightarrow{(cl, \emptyset)}_{\text{ET}_\top} (\mathcal{K}, \mathcal{U})\right)$,
- let $(\mathcal{K}', \mathcal{U}') = \text{lastConf}(\tau)$; if $\mathcal{X} \in \text{absExec}(\tau)$, $\mathcal{F} \neq \emptyset$, and $\mathcal{T} = \text{visTx}(\mathcal{K}, \mathcal{U}'(cl)) \cup \mathcal{T}_{\text{rd}}$ where $\mathcal{T}_{\text{rd}}$ is a set of *read-only transactions* such that $(\mathsf{w}, k, v) \notin_{\mathcal{X}} t'$ for all keys k and values v and transactions $t' \in \mathcal{T}_{\text{rd}}$, and if the transaction t is the transaction appearing in $\text{lastConf}(\tau)$ but not in $\mathcal{K}$, then $\text{extend}(\mathcal{X}, t, \mathcal{T}, \mathcal{F}) \in \text{absExec}\left(\tau \xrightarrow{(cl, \mathcal{F})}_{\text{ET}_\top} (\mathcal{K}, \mathcal{U})\right)$.

PROPOSITION D.9 (TRACE OF ET TO ABSTRACT EXECUTIONS). *For any $\text{ET}_\top$-trace τ, the abstract execution $\mathcal{X} \in \text{absExec}(\tau)$ satisfies the last write wins policy, and $(\mathcal{K}_{\mathcal{X}}, \_) = \text{lastConf}(\tau)$.*

PROOF. Fix a $\text{ET}_\top$-trace τ. We prove by induction on the number of transitions n in τ.

- Base case: $n = 0$. It means $\tau = (\mathcal{K}_0, \_)$. It follows from Def. D.8 that $\mathcal{X}_\tau = ([], \emptyset, \emptyset)$. This triple satisfies the constraints of Def. C.1, as well as the resolution policy RP_{LWW}. It is also immediate to see that $\text{graphOf}(\mathcal{X}) = ([], \emptyset, \emptyset, \emptyset)$. In particular, $\mathcal{T}_{\text{graphOf}(\mathcal{X})} = \emptyset$, and the only kv-store $\mathcal{K}$ such that $\mathcal{T}_{\mathcal{G}_{\mathcal{K}}} = \emptyset$ is given by $\mathcal{K} = \mathcal{K}_0$. By definition, $\mathcal{K}_{\mathcal{X}_\tau} = \mathcal{K}_0$, as we wanted to prove.

- Inductive case: $n > 0$. In this case, we have that $\tau = \tau' \xrightarrow{(cl, \mu)}_{\text{ET}} (\mathcal{K}, \mathcal{U})$ for some $cl, \mu, \mathcal{K}, \mathcal{U}$. The $\text{ET}_\top$-trace τ' contains exactly $n-1$ transitions, so that by induction we can assume that $\mathcal{X}_{\tau'}$ is a valid abstract execution that satisfies RP_{LWW}. and $\text{lastConf}(\tau') = (\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}')$ for some $\mathcal{U}'$.
  We perform a case analysis on μ. If $\mu = \varepsilon$, then it follows that $\mathcal{K} = \mathcal{K}_{\mathcal{X}_{\tau'}}$, and $\mathcal{X}_\tau = \mathcal{X}_{\tau'}$ by Def. D.8. Then by the inductive hypothesis $\mathcal{X}_\tau$ is an abstract execution that satisfies RP_{LWW}, $\text{lastConf}(\tau) = (\mathcal{K}, \_)$, and $\mathcal{K}_{\mathcal{X}_\tau} = \mathcal{K}_{\mathcal{X}_{\tau'}} = \mathcal{K}$, and there is nothing left to prove.
  Suppose now that $\mu = \mathcal{F}$, for some $\mathcal{F}$. In this case we have that $\mathcal{K} \in \text{update}(\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}'(cl), \mathcal{F}, cl)$. Note that if $\mathcal{F} = \emptyset$, then $\mathcal{K} = \mathcal{K}_{\mathcal{X}_{\tau'}}$ and $\mathcal{X}_\tau = \mathcal{X}_{\tau'}$. By the inductive hypothesis, $\mathcal{X}_\tau$ is an abstract execution that satisfies RP_{LWW}, and $\mathcal{K} = \mathcal{K}_{\mathcal{X}_{\tau'}} = \mathcal{K}_{\mathcal{X}_\tau}$. Assume then that $\mathcal{F} \neq \emptyset$. By definition, $\mathcal{K} = \text{update}(\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}'(cl), \mathcal{F}, t)$ for some $t \in \text{nextTid}(cl, \mathcal{K}_{\mathcal{X}_{\tau'}})$. It follows that t is the unique transaction such that $t \notin \mathcal{K}_{\mathcal{X}_{\tau'}}$, and $t \in \mathcal{K}$ (the fact that $t \in \mathcal{K}$ follows from the assumption that $\mathcal{F} \neq \emptyset$). Let $\mathcal{T} = \text{visTx}(\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}'(cl))$; then $\mathcal{X}_\tau = \text{extend}(\mathcal{K}_{\mathcal{X}_{\tau'}}, t, \mathcal{T}, \mathcal{F})$. Note that $\mathcal{X}_\tau$ satisfies the constraints of abstract execution required by Def. C.1:
  - Because $t \in \text{nextTid}(cl, \mathcal{K}_{\mathcal{X}_\tau})$, it must be the case that $t = t_{cl}^m$ for some $m \geq 1$; we have that $\mathcal{T}_{\mathcal{X}_\tau} = \mathcal{T}_{\mathcal{X}_{\tau'}}[t_{cl}^m \mapsto \mathcal{F}]$, from which it follows that

$$\mathcal{T}_{\mathcal{X}_\tau} = \text{dom}(\mathcal{T}_{\mathcal{X}_\tau}) = \text{dom}(\mathcal{T}_{\mathcal{X}_{\tau'}}) \cup \{t_{cl}^m\} = \mathcal{T}_{\mathcal{X}_{\tau'}} \cup \{t_{cl}^m\}$$

  By inductive hypothesis, $t_0 \notin \mathcal{T}_{\mathcal{X}_{\tau'}}$, and therefore $t_0 \notin \mathcal{T}_{\mathcal{X}_{\tau'}} \cup \{t_{cl}^m\} = \mathcal{T}_{\mathcal{X}}$.

- $VIS_{\mathcal{X}_\tau} \subseteq AR_{\mathcal{X}_\tau}$. Let $(t', t'') \in VIS_{\mathcal{X}_\tau}$. Then either $t'' = t^m_{cl}$ and $t' \in \mathcal{T}$, or $(t', t'') \in VIS_{\mathcal{X}_{\tau'}}$. In the former case, we have that $(t', t^m_{cl}) \in AR_{\mathcal{X}_\tau}$ by definition; in the latter case, we have that $(t', t'') \in AR_{\mathcal{X}_{\tau'}}$ because $\mathcal{X}_{\tau'}$ is a valid abstract execution by inductive hypothesis, and therefore $(t', t'') \in AR_{\mathcal{X}_\tau}$ by definition. This concludes the proof that $VIS_{\mathcal{X}_\tau} \subseteq AR_{\mathcal{X}_\tau}$.
- $VIS_{\mathcal{X}_\tau}$ is irreflexive. Assume $(t', t'') \in VIS_{\mathcal{X}_\tau}$, then either $(t't'') \in VIS_{\mathcal{X}_{\tau'}}$, and because $VIS_{\mathcal{X}_{\tau'}}$ is irreflexive by the inductive hypothesis, then $t' \neq t''$; or $t'' = t^m_{cl}, t' \in \mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}_{\tau'}}$, and because $t^m_{cl} \notin \mathcal{K}_{\mathcal{X}_{\tau'}}$, then $t' \neq t^m_{cl}$.
- $AR_{\mathcal{X}_\tau}$ is total. Let $(t', t'') \in \mathcal{T}_{\mathcal{X}_\tau}$. Suppose that $t' \neq t''$.
(1) If $t' \neq t^m_{cl}, t'' \neq t^m_{cl}$, then it must be the case that $t', t'' \in \mathcal{T}_{\mathcal{X}_{\tau'}}$; this is because we have already argued that $\mathcal{T}_{\mathcal{X}_\tau} = \mathcal{T}_{\mathcal{X}_{\tau'}} \cup \{t^m_{cl}\}$. By the inductive hypothesis, we have that either $(t', t'') \in AR_{\mathcal{X}_{\tau'}}$, or $(t'', t') \in AR_{\mathcal{X}_{\tau'}}$. Because $AR_{\mathcal{X}_{\tau'}} \subseteq AR_{\mathcal{X}_\tau}$, then either $(t', t'') \in AR_{\mathcal{X}_{\tau'}}$ or $(t'', t') \in AR_{\mathcal{X}_\tau}$.
(2) if $t'' = t^m_{cl}$, then it must be $t' \in \mathcal{T}_{\mathcal{X}_{\tau'}}$. By definition, $(t', t^m_{cl}) \in AR_{\mathcal{X}_\tau}$. Similarly, if $t' = t^m_{cl}$, we can prove that $(t'', t^m_{cl}) \in AR_{\mathcal{X}_\tau}$.
- $AR_{\mathcal{X}_\tau}$ is irreflexive. It follows is the same as the one of $VIS_{\mathcal{X}_\tau}$.
- $AR_{\mathcal{X}_\tau}$ is transitive. Assume $(t', t'') \in AR_{\mathcal{X}_\tau}$ and $(t'', t''') \in AR_{\mathcal{X}_\tau}$. Note that it must be the case that $t', t'' \in \mathcal{T}_{\mathcal{X}_{\tau'}}$ by the definition of $AR_{\mathcal{X}}$, and in particular $(t', t'') \in AR_{\mathcal{X}_{\tau'}}$. For t''', we have two possible cases.
(1) Either $t''' \in \mathcal{T}_{\mathcal{X}_\tau}$, from which it follows that $(t'', t''') \in AR_{\mathcal{X}_{\tau'}}$; because of $AR_{\mathcal{X}_{\tau'}}$ is transitive by the inductive hypothesis, then $(t', t''') \in AR_{\mathcal{X}_{\tau'}}$, and therefore $(t', t''') \in AR_{\mathcal{X}_\tau}$.
(2) Or $t''' = t^m_{cl}$, and because $t' \in \mathcal{T}_{\mathcal{X}_{\tau'}}$, then $(t', t^m_{cl}) \in AR_{\mathcal{X}_\tau}$ by definition.
- $SO_{\mathcal{X}_\tau} \subseteq AR_{\mathcal{X}_\tau}$. Let cl' be a client such that $(t^i_{cl'}, t^j_{cl'}) \in AR_{\mathcal{X}_\tau}$. If $cl' \neq cl$, then it must be the case that $t^i_{cl'}, t^j_{cl'} \in \mathcal{T}_{\mathcal{X}_{\tau'}}$, and therefore $(t^i_{cl'}, t^j_{cl'}) \in AR_{\mathcal{X}_{\tau'}}$. By the inductive hypothesis, it follows that $i < j$. If $cl' = cl$, then by definition of $AR_{\mathcal{X}_\tau}$ it must be $i \neq m$. If $j \neq m$ we can proceed as in the previous case to prove that $i < j$. If $j = m$, then note that $t^i_{cl} \in \mathcal{T}_{\mathcal{X}_\tau}$ only if $t^i_{cl} \in \mathcal{K}_{\mathcal{X}_{\tau'}}$. Because $t^m_{cl} \in nextTid(\mathcal{K}_{\mathcal{X}_{\tau'}}, cl)$, then we have that $i < m$, as we wanted to prove.

Next, we prove that $\mathcal{X}_\tau$ satisfies the last write wins policy. Let $t' \in \mathcal{T}_{\mathcal{X}_\tau}$, and suppose that $(r, k, v) \in_{\mathcal{X}_\tau} t'$.

- If $t' \neq t$, then we have that $t \in \mathcal{T}_{\mathcal{X}_{\tau'}}$. We also have that $VIS^{-1}_{\mathcal{X}_\tau}(t') = VIS^{-1}_{\mathcal{X}_{\tau'}}(t')$, $AR^{-1}_{\mathcal{X}_\tau}(t') = AR^{-1}_{\mathcal{X}_{\tau'}}(t')$; finally, for any $t'' \in \mathcal{T}_{\mathcal{X}_{\tau'}}$, $(w, k, v') \in_{\mathcal{X}_\tau} t''$ if and only if $(w, k, v') \in_{\mathcal{X}_{\tau'}} t''$. Therefore, let $t_r = \max_{AR_{\mathcal{X}_\tau}}(VIS^{-1}_{\mathcal{X}_\tau}(t') \cap \{t'' \mid (w, k, \_) \in_{\mathcal{X}_\tau} t''\})$. We have that

$$t_r = \max_{AR_{\mathcal{X}_{\tau'}}}(VIS^{-1}_{\mathcal{X}_{\tau'}}(t) \cap \{t'' \mid (w, k, \_) \in_{\mathcal{X}_{\tau'}} t''\})$$

and because $\mathcal{X}_{\tau'}$ satisfies the last write wins resolution policy, then $(w, k, v) \in_{\mathcal{X}_{\tau'}} t_r$. This also implies that $(w, k, v) \in_{\mathcal{X}_\tau} t_r$.
- Now, suppose that $t' = t$. Suppose that $(r, k, v) \in_{\mathcal{X}_\tau} t'$. By definition, we have that $(r, k, v) \in \mathcal{F}$. Recall that $\tau = \tau' \xrightarrow{(cl, \mathcal{F})}_{ET_\top} (\mathcal{K}, \mathcal{U})$, and $lastConf(\tau') = (\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}')$ for some $\mathcal{U}'$. That is,

$$(\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}') \xrightarrow{(cl, \mathcal{F})}_{ET_\top} (\mathcal{K}, \mathcal{U})$$

which in turn implies that $ET_\top \vdash (\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}'(cl)) \triangleright \mathcal{F} : (\mathcal{K}, \mathcal{U}(cl))$. Let then $r = \max\{i \mid i \in \mathcal{U}'(cl)(k)\}$. By definition of execution test, and because $(r, k, v) \in \mathcal{F}$, then it must be the case that $\mathcal{K}_{\mathcal{X}_{\tau'}}(k, r) = (v, t'', \_)$ for some t''.

We now prove that $t'' = \max_{\mathrm{AR}_{\mathcal{X}_\tau}}(\mathrm{VIS}^{-1}_{\mathcal{X}_\tau}(t) \cap \{t'' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}_\tau} t''\})$. First we have

$$\mathrm{VIS}^{-1}_{\mathcal{X}_\tau}(t) = \mathrm{visTx}\big(\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}'(cl)\big) = \big\{\mathsf{w}(\mathcal{K}_{\mathcal{X}_{\tau'}}(k', i)) \mid k' \in \mathrm{K{\small EY}} \wedge i \in \mathcal{U}'(cl)(k')\big\}$$

Note that $r \in \mathcal{U}'(cl)(k)$, and $t'' = \mathsf{w}(\mathcal{K}_{\mathcal{X}_{\tau'}}(k, r))$. Therefore, $t'' \in \mathrm{VIS}^{-1}_{\mathcal{X}_\tau}(t)$. Because $\mathcal{K} = \mathrm{update}\big(\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}'(cl), \mathcal{F}, t\big)$, it must be the case that $\mathsf{w}(\mathcal{K}(k, r)) = t''$. Also, because $\mathsf{w}(\mathcal{K}_{\mathcal{X}_{\tau'}}(k, r)) = t''$, then $(\mathsf{w}, k, \_) \in_{\mathcal{X}_{\tau''}} t''$, or equivalently $(\mathsf{w}, k, \_) \in \mathcal{T}_{\mathcal{X}_{\tau'}}(t'')$. We have already proved that $\mathrm{VIS}_{\mathcal{X}_\tau}$ is irreflexive, hence it must be the case that $t'' \neq t$. In particular, because $\mathcal{X}_\tau = \mathrm{extend}(\mathcal{X}_{\tau'}, t, \_, \_)$, then we have that $\mathcal{T}_{\mathcal{X}_\tau}(t'') = \mathcal{T}_{\mathcal{X}_{\tau'}}[t \mapsto \mathcal{F}](t'') = \mathcal{T}_{\mathcal{X}_{\tau'}}(t'')$, hence $(\mathsf{w}, k, \_) \in \mathcal{T}_{\mathcal{X}_\tau}(t'')$. Equivalently, $(\mathsf{w}, k, \_) \in_{\mathcal{X}_\tau} t''$. We have proved that $t'' \in \mathrm{VIS}^{-1}_{\mathcal{X}_\tau}(t)$, and $(\mathsf{w}, k, \_) \in_{\mathcal{X}_\tau} t''$.

Now let t''' be such that $t''' \in \mathrm{VIS}^{-1}_{\mathcal{X}_\tau}(t)$, and $(\mathsf{w}, k\_) \in_{\mathcal{X}_\tau} t'''$. Note that $t''' \neq t$ because $\mathrm{VIS}_{\mathcal{X}_\tau}$ is irreflexive. We show that either $t''' = t''$, or $t''' \xrightarrow{\mathrm{AR}_{\mathcal{X}_\tau}} t''$. Because $t''' \in \mathrm{VIS}^{-1}_{\mathcal{X}_\tau}(t)$, then there exists a key k' and an index $i \in \mathcal{U}'(cl)$ such that $\mathsf{w}(\mathcal{K}_{\mathcal{X}_{\tau'}}(k', i)) = t'''$. Because $(\mathsf{w}, k, \_) \in_{\mathcal{X}_\tau} t'''$, and because $t''' \neq t$, then $(\mathsf{w}, k, \_) \in_{\mathcal{X}_{\tau'}} t'''$, and therefore there exists an index j such that $\mathsf{w}(\mathcal{K}_{\mathcal{X}_{\tau'}}(k, j)) = t'''$. We have that $\mathsf{w}(\mathcal{K}_{\mathcal{X}_{\tau'}}(k, j)) = \mathsf{w}(\mathcal{K}_{\mathcal{X}_{\tau'}}(k', i))$, and $i \in \mathcal{U}'(cl)$. By Eq. (atomic), it must be $j \in \mathcal{U}'(cl)$. Note that $r = \max\{i \mid i \in \mathcal{U}'(cl)\}$, hence we have that $j \leq r$. If $j = r$, then $t''' = t''$ and there is nothing left to prove. If $j < r$, then we have that $(t''', t'') \in \mathrm{AR}_{\mathcal{X}_{\tau'}}$, and therefore $(t''', t'') \in \mathrm{AR}_{\mathcal{X}_\tau}$.

Finally, we need to prove that $\mathcal{K} = \mathcal{K}_{\mathcal{X}_\tau}$. Recall $\mathcal{K} = \mathrm{update}\big(\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}'(cl), \mathcal{F}, t\big)$, and $\mathcal{X}_\tau = \mathrm{extend}\big(\mathcal{X}_{\tau'}, \mathrm{visTx}\big(\mathcal{K}_{\mathcal{X}_{\tau'}}, \mathcal{U}'(cl)\big), t, \mathcal{F}\big)$. The result follows then from Prop. D.10.

□

PROPOSITION D.10 (extend MATCHING update). *Given an abstract execution $\mathcal{X}$, a set of transactions $\mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}}$, a transaction $t \notin \mathcal{T}_{\mathcal{X}}$, and a fingerprint $\mathcal{F} \subseteq \mathcal{P}(\mathrm{O{\small PS}})$, if the new abstract execution $\mathcal{X}' = \mathrm{extend}(\mathcal{X}, \mathcal{T}, t, \mathcal{F})$, and the view $u = \mathrm{getView}(\mathcal{K}_{\mathcal{X}}, \mathcal{T})$, then $\mathrm{update}(\mathcal{K}_{\mathcal{X}}, u, \mathcal{F}, t) = \mathcal{K}_{\mathcal{X}'}$.*

PROOF. Let $\mathcal{G} = \mathcal{G}_{\mathrm{update}(\mathcal{K}_{\mathcal{X}}, u, \mathcal{F}, t)}$, $\mathcal{G}' = \mathrm{graphOf}(\mathcal{X}')$. Note that $\mathcal{K}_{\mathcal{X}'}$ is the unique kv-store such that $\mathcal{G}_{\mathcal{K}_{\mathcal{X}'}} = \mathrm{graphOf}(\mathcal{X}') = \mathcal{G}'$. It suffices to prove that $\mathcal{G} = \mathcal{G}'$. Because the function $\mathcal{G}_{\cdot}$ is injective, it follows that $\mathrm{update}(\mathcal{K}_{\mathcal{X}}, u, \mathcal{F}, t) = \mathcal{K}_{\mathcal{X}'}$, as we wanted to prove.

The proof is a consequence of Lemma D.11 and Lemma D.12. Consider the dependency graph $\mathcal{G}_{\mathcal{K}_{\mathcal{X}}}$. Recall that $\mathcal{K}_{\mathcal{X}}$ is the unique kv-store such that $\mathcal{G}_{\mathcal{K}_{\mathcal{X}}} = \mathrm{graphOf}(\mathcal{X})$. We prove that $\mathcal{T}_{\mathcal{G}} = \mathcal{T}_{\mathcal{G}'}$, $\mathrm{WR}_{\mathcal{G}} = \mathrm{WR}_{\mathcal{G}'}$ and $\mathrm{WW}_{\mathcal{G}} = \mathrm{WW}_{\mathcal{G}'}$ (from the last two it follows that $\mathrm{RW}_{\mathcal{G}} = \mathrm{RW}_{\mathcal{G}'}$).

- It is easy to see $\mathcal{T}_{\mathcal{G}} = \mathcal{T}_{\mathcal{G}'}$.

- $\mathrm{WR}_{\mathcal{G}} = \mathrm{WR}_{\mathcal{G}'}$. Let $\mathcal{K} = \mathcal{K}_{\mathcal{X}}$. Suppose that $t' \xrightarrow{\mathrm{WR}_{\mathcal{G}}(k)} t''$ for some t', t''. By Lemma D.12 we have that either $t' \xrightarrow{\mathrm{WR}_{\mathcal{G}_{\mathcal{K}}}(k)} t''$, or $t'' = t$, $(\mathsf{r}, k, \_) \in \mathcal{F}$, $t' = \max_{\mathrm{WW}_{\mathcal{G}_{\mathcal{K}}}(k)}\{\mathsf{w}(k, i) \mid i \in u(k)\}$.

  - If $t' \xrightarrow{\mathrm{WR}_{\mathcal{G}_{\mathcal{K}}}(k)} t''$, then because $\mathcal{G}_{\mathcal{K}} = \mathrm{graphOf}(\mathcal{X})$, we have that $t' \xrightarrow{\mathrm{WR}_{\mathrm{graphOf}(\mathcal{X})}(k)} t''$. Recall that $\mathcal{G}' = \mathrm{graphOf}(\mathrm{extend}(\mathcal{X}, \mathcal{T}, t, \mathcal{F}))$, hence by Lemma D.11 we obtain that $t' \xrightarrow{\mathrm{WR}_{\mathcal{G}'}(k)} t''$.

  - If $t'' = t$, $(\mathsf{r}, k, \_) \in \mathcal{F}$, and $t' = \max_{\mathrm{WW}_{\mathcal{G}_{\mathcal{K}}}(k)}\{\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, i)) \mid i \in u(k)\}$, then we also have that $t' = \max_{\mathrm{WW}_{\mathrm{graphOf}(\mathcal{X})}(k)}(\mathcal{T} \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}} t'''\})$. This is because of the assumption that

$$\{\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, i)) \mid i \in u(k)\} = \{\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k', i)) \mid k' \in \mathrm{K{\small EY}} \wedge i \in u(k')\} \cap \{\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, \_)\}$$
$$= \mathrm{visTx}(\mathcal{K}_{\mathcal{X}}, u) \cap \{\mathsf{w}(\mathcal{K}_{\mathcal{X}}(k, \_)\}$$
$$= \mathcal{T} \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}} t'''\}$$

Again, it follows from Lemma D.11 that $t' \xrightarrow{\text{WR}_{\mathcal{G}'}(k)} t''$.

- $\text{WW}_{\mathcal{G}} = \text{WW}_{\mathcal{G}'}$. The $\text{WW}_{\mathcal{G}} = \text{WW}_{\mathcal{G}'}$ follows the similar reasons as $\text{WR}_{\mathcal{G}} = \text{WR}_{\mathcal{G}'}$.

$\square$

LEMMA D.11 (GRAPH TO ABSTRACT EXECUTION EXTENSION). *Let $\mathcal{X}$ be an abstract execution, $t \notin \mathcal{T}_{\mathcal{X}} \cup \{t_0\}$ be a transaction identifier $\mathcal{T}_{\mathcal{X}}$, and $\mathcal{F} \in \mathcal{T}_{\mathcal{X}}$. Let $\mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}}$ be a set of transaction identifiers. Let $\mathcal{G} = \text{graphOf}(\mathcal{X}), \mathcal{G}' = \text{graphOf}(\text{extend}(\mathcal{X}, t, \mathcal{T}, \mathcal{F}))$. We have the following:*

(1) for any $t' \in \mathcal{T}_{\mathcal{G}'}$, either $t' \in \mathcal{T}_{\mathcal{G}}$ and $\mathcal{T}_{\mathcal{G}}(t') = \mathcal{T}_{\mathcal{G}'}(t')$, or $t' = t$ and $\mathcal{T}_{\mathcal{G}'}(t) = \mathcal{F}$.

(2) $t' \xrightarrow{\text{WR}_{\mathcal{G}'}(k)} t''$ if and only if either $t' \xrightarrow{\text{WR}_{\mathcal{G}}(k)_{\mathcal{G}}} t''$, or $(\mathsf{r}, k, \_) \in \mathcal{F}$, $t'' = t$ and $t' = \max_{\text{WW}_{\mathcal{G}}(k)}(\mathcal{T})$,

(3) $t' \xrightarrow{\text{WW}_{\mathcal{G}'}(k)} t''$ if and only if either $t' \xrightarrow{\text{WW}_{\mathcal{G}}(k)} t''$, or $(\mathsf{w}, k, \_) \in \mathcal{F}$, $t'' = t$, and $(\mathsf{w}, k, \_) \in_{\mathcal{G}} t'$.

PROOF. Fix a key k. Let $\mathcal{X}' = \text{extend}(\mathcal{X}, t, \mathcal{T}, \mathcal{F})$. Recall that $\mathcal{G}' = \text{graphOf}(\mathcal{X}')$.

(1) By definition of extend, and because $t \notin \mathcal{T}_{\mathcal{X}}$, we have that $\mathcal{T}_{\mathcal{X}'} = \mathcal{T}_{\mathcal{X}} \uplus \{t\}$. Furthermore, $\mathcal{T}_{\mathcal{X}'}(t) = \mathcal{F}$, from which it follows that $\mathcal{T}_{\mathcal{G}'}(t) = \mathcal{F}$. For all $t' \in \mathcal{T}_{\mathcal{X}}$, we have that $\mathcal{T}_{\mathcal{X}'}(t') = \mathcal{T}_{\mathcal{X}}(t') = \mathcal{T}_{\mathcal{G}}(t')$.

(2) There are two cases that either the t'' already exists in the dependency graph before, or it is the newly committed transaction.

- Suppose that $t' \xrightarrow{\text{WR}(k)_{\mathcal{G}}} t''$ for some $t', t'' \in \mathcal{T}_{\mathcal{G}}$. By definition, $(\mathsf{r}, k, \_) \in_{\mathcal{X}} t''$, and $t' = \max_{\text{AR}_{\mathcal{X}}}(\text{VIS}_{\mathcal{X}}^{-1}(t'') \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}} t'''\})$. Because $t'' \in \mathcal{T}_{\mathcal{G}} = \mathcal{T}_{\mathcal{X}}$, it follows that $t'' \neq t$. By definition, $\text{VIS}_{\mathcal{X}'}^{-1}(t'') = \text{VIS}_{\mathcal{X}}^{-1}(t)$: also, whenever $t_a, t_b \in \text{VIS}_{\mathcal{X}'}^{-1}(t)$ we have that $t_a, t_b \in \mathcal{T}_{\mathcal{X}}$, and therefore if $t_a \xrightarrow{\text{AR}_{\mathcal{X}'}} t_b$, then it must be the case that $t_a \xrightarrow{\text{AR}_{\mathcal{X}}} t_b$; also, $\mathcal{T}_{\mathcal{X}}(t_a) = \mathcal{T}_{\mathcal{X}'}(t_a)$. As a consequence, we have that

$$\max_{\text{AR}_{\mathcal{X}'}}(\text{VIS}_{\mathcal{X}'}^{-1}(t) \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}'} t'''\}) = \max_{\text{AR}_{\mathcal{X}}}(\text{VIS}_{\mathcal{X}}^{-1}(t) \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}} t'''\}) = t'$$

and therefore $t' \xrightarrow{\text{WR}_{\mathcal{G}'}} t$.

- Suppose now that $(\mathsf{r}, k, \_) \in \mathcal{F}$, and $t' = \max_{\text{WW}(k)_{\mathcal{G}}}(\mathcal{T})$. By Definition, $t' = \max_{\text{AR}_{\mathcal{X}}}(\mathcal{T}) \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}} t'''\}$, and, $\mathcal{T} = \text{VIS}_{\mathcal{X}'}^{-1}(t)$. Because $\mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}}$, we have that for any t_a, t_b, if $t_a \xrightarrow{\text{AR}_{\mathcal{X}}} t_b$, then $t_a \xrightarrow{\text{AR}_{\mathcal{X}'}} t_b$; and $\mathcal{T}_{\mathcal{X}'}(t_a) = \mathcal{T}_{\mathcal{X}}(t_a)$. Therefore,

$$t' = \max_{\text{AR}_{\mathcal{X}'}}(\text{VIS}_{\mathcal{X}'}^{-1}(t) \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}'} t'''\}),$$

from which it follows that $t' \xrightarrow{\text{WR}_{\mathcal{G}'}(k)} t$.

Now, suppose that $t' \xrightarrow{\text{WR}_{\mathcal{G}'}(k)} t''$ for some $t', t'' \in \mathcal{T}_{\mathcal{G}'} = \mathcal{T}_{\mathcal{X}'}$. We have that $(\mathsf{r}, k, \_) \in_{\mathcal{X}'} t''$, $(\mathsf{w}, k, \_) \in_{\mathcal{X}'} t'$, and $t'' = \max_{\text{AR}_{\mathcal{X}'}}(\text{VIS}_{\mathcal{X}'}^{-1}(t'') \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}'} t'''\})$. We also have that $\mathcal{T}_{\mathcal{X}'} = \mathcal{T}_{\mathcal{X}} \uplus \{t\}$. We perform a case analysis on t''.

  - If $t'' = t$, then by definition of extend we have that $\text{VIS}_{\mathcal{X}'}^{-1}(t) = \mathcal{T}$. Note that $\mathcal{T} \subseteq \mathcal{T}_{\mathcal{X}}$, so that for any $t_a, t_b \in \mathcal{T}_{\mathcal{X}}$, we have that $t_a \xrightarrow{\text{AR}_{\mathcal{X}'}} t_b$ if and only if $t_a \xrightarrow{\text{AR}_{\mathcal{X}}} t_b$, and $(\mathsf{w}, k, v) \in_{\mathcal{X}'} t_a$ if and only if $(\mathsf{w}, k, v) \in_{\mathcal{X}} t_a$. Thus, $t' = \max_{\text{AR}_{\mathcal{X}}}(\mathcal{T} \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}} t'''\}) = \max_{\text{WW}_{\mathcal{G}}(k)}(\mathcal{T})$.

  - If $t'' \in \mathcal{T}_{\mathcal{X}}$, then it is the case that $t' = \max_{\text{AR}_{\mathcal{X}'}}(\text{VIS}_{\mathcal{X}'}^{-1}(t'') \cap \{t''' \mid (\mathsf{w}, k, \_) \in_{\mathcal{X}'} t'''\})$. Similarly to the case above, we can prove that $\text{VIS}_{\mathcal{X}'}^{-1}(t'') = \text{VIS}_{\mathcal{X}}^{-1}(t)$, for any $t_a, t_b \in \text{VIS}_{\mathcal{X}}^{-1}(t)$, $(\mathsf{w}, k, v) \in_{\mathcal{X}'} t_a$ implies $(\mathsf{w}, k, v) \in_{\mathcal{X}} t_a$, and $t_a \xrightarrow{\text{AR}_{\mathcal{X}'}} t_b$ implies $t_a \xrightarrow{\text{AR}_{\mathcal{X}}} t_b$,

from which it follows that $t' = \max_{\mathrm{AR}_\chi}(\mathrm{VIS}_\chi^{-1}(t'') \cap \{t''' \mid (\mathrm{w}, k\_) \in_\chi t'''\})$, and therefore $t' \xrightarrow{\mathrm{WR}_{\mathcal{G}}(k)} t''$.

(3) Similar to $\mathrm{WR}(k)_\mathcal{G}$, there are two cases that either the t'' already exists in the dependency graph before, or it is the newly committed transaction.

- Suppose that $t' \xrightarrow{\mathrm{WW}_{\mathcal{G}}(k)} t''$ for some $t', t'' \in \mathcal{T}_\chi$. Then $(\mathrm{w}, k, \_) \in_\chi t'$, $(\mathrm{w}, k, \_) \in_\chi t''$, and $t' \xrightarrow{\mathrm{AR}_\chi} t''$. By definition of extend, it follows that $t' \xrightarrow{\mathrm{AR}_{\chi'}} t''$, and because $t', t'' \in \mathcal{T}_\chi$, hence $t', t'' \neq t$, then $(\mathrm{w}, k, \_) \in_{\chi'} t'$, $(\mathrm{w}, k, \_) \in_{\chi'} t''$. By definition, we have that $t' \xrightarrow{\mathrm{WW}_{\chi'}(k)} t''$.

- Suppose that $(\mathrm{w}, k, \_) \in_\chi t'$, $(\mathrm{w}, k, \_) \in \mathcal{F}$. Because $t' \in \mathcal{T}_\chi$, we have that $t' \neq t$, hence $(\mathrm{w}, k, \_) \in_{\chi'} t'$. By definition, $\mathscr{T}_{\chi'}(t) = \mathcal{F}$, hence $(\mathrm{w}, k, \_) \in_{\chi'} t$. Finally, the definition of extend ensures that $t' \xrightarrow{\mathrm{AR}_{\chi'}} t$. Combining these three facts together, we obtain that $t' \xrightarrow{\mathrm{WW}_{\mathcal{G}'}(k)} t$.

  Now, suppose that $t' \xrightarrow{\mathrm{WW}_{\mathcal{G}'}(k)} t''$ for some $t', t'' \in \mathcal{T}_\chi$. Then $t' \xrightarrow{\mathrm{AR}_{\chi'}} t''$, $(\mathrm{w}, k, \_) \in_{\chi'} t'$, $(\mathrm{w}, k, \_) \in_{\chi'} t''$. Recall that $\mathcal{T}_{\mathcal{G}'} = \mathcal{T}_{\chi'} = \mathcal{T}_\chi \uplus \{t\}$. We perform a case analysis on t''.

  - If $t'' = t$, then the definition of extend ensures that $t' \xrightarrow{\mathrm{AR}_{\chi'}} t$ implies that $t \in \mathcal{T}_\chi$, hence $t' \neq t$. Together with $(\mathrm{w}, k, \_) \in_{\chi'} t'$, this leads to $(\mathrm{w}, k, \_) \in_\chi t'$.

  - If $t'' \in \mathcal{T}_\chi$, then $t'' \neq t$. The definition of extend ensures that $t' \xrightarrow{\mathrm{AR}_\chi} t''$. This implies that $t', t'' \in \mathcal{T}_\chi$, hence $t', t'' \neq t$, and $\mathscr{T}_{\chi'}(t') = \mathscr{T}_\chi(t')$, $\mathscr{T}_{\chi'}(t'') = \mathscr{T}_\chi(t'')$. It follows that $(\mathrm{w}, k, \_) \in_\chi t'$, $(\mathrm{w}, k, \_) \in_\chi t''$, and therefore $t' \xrightarrow{\mathrm{WW}_{\mathcal{G}}(k)} t''$.

  $\square$

LEMMA D.12 (GRAPH TO KV-STORE UPDATE). *Let $\mathcal{K}$ be a kv-store, and $u \in \mathrm{VIEWS}(\mathcal{K})$. Let $t \notin \mathcal{K}$, and $\mathcal{F} \subseteq \mathcal{P}(\mathrm{OPS})$, and let $\mathcal{K}' = \mathrm{update}(\mathcal{K}, u, \mathcal{F}, t)$. Let $\mathcal{G} = \mathcal{G}_\mathcal{K}$, $\mathcal{G}' = \mathcal{G}_{\mathcal{K}'}$; then for all $t', t'' \in \mathcal{T}_{\mathcal{G}'}$ and keys k,*

- $\mathscr{T}_{\mathcal{G}'} = \mathscr{T}_{\mathcal{G}}[t \mapsto \mathcal{F}]$,
- $t' \xrightarrow{\mathrm{WR}_{\mathcal{G}'}(k)} t''$ *if and only if either* $t' \xrightarrow{\mathrm{WR}_{\mathcal{G}}(k)} t''$, *or* $(\mathrm{r}, k, \_) \in \mathcal{F}$ *and*

$$t' = \max_{\mathrm{WW}_{\mathcal{G}}(k)}(\{\mathrm{w}(\mathcal{K}(k, i))\} \, i \in u(k))$$

- $t' \xrightarrow{\mathrm{WW}_{\mathcal{G}'}(k)} t''$ *if and only if either* $t' \xrightarrow{\mathrm{WW}_{\mathcal{G}}(k)} t''$, *or* $(\mathrm{w}, k, \_) \in \mathcal{F}$ *and* $t' = \mathrm{w}(\mathcal{K}(k, \_))$.

PROOF. Fix $k \in \mathrm{KEY}$. Because $t \notin \mathcal{K}$, then $t \notin \mathcal{T}_\mathcal{G}$, and by definition of *update* we obtain that $\{t' \mid t' \in \mathcal{K}'\} = \{t' \mid t' \in \mathcal{K}\} \cup \{t\}$. It follows that $\mathcal{T}_{\mathcal{G}'} = \mathcal{T}_\mathcal{G} \uplus \{t\}$.

(1) Suppose that $(\mathrm{r}, k, v) \in_\mathcal{G} t'$. By definition, there exists an index i such that $\mathcal{K}(k, i) = (v, \_, \{t'\} \cup \_)$. Because $\mathcal{K}' = \mathrm{update}(\mathcal{K}, u, \mathcal{F}, t)$, it is immediate to observe that $\mathcal{K}'(k, i) = (v, \_, \{t'\} \cup \_)$, and therefore $(\mathrm{r}, k, v) \in_{\mathcal{G}'} t'$. Conversely, note that if $(\mathrm{r}, k, v) \in_{\mathcal{G}'} t$, then there exists an index $i = 0, \cdots, |\mathcal{K}'(k)| - 1$ such that $\mathcal{K}'(k, i) = (v, \_, \{t'\} \cup \_)$. it follows that it must be the case that $i \leq |\mathcal{K}(k)| - 1$, and because $t' \neq t$, we have that $\mathcal{K}(k, i) = (v, \_, \{t'\} \cup \_)$. Therefore $(\mathrm{r}, k, v) \in_\mathcal{G} t'$.

Similarly, if $(\mathrm{w}, k, v) \in_\mathcal{G} t'$, then there exists an index $i = 0, \cdots, |\mathcal{K}(k)| - 1$ such that $\mathcal{K}(k, i) = (v, t', v)$. It follows that $\mathcal{K}'(k, i) = (v, t', \_)$, hence $(\mathrm{w}, k, v) \in_{\mathcal{G}'} t'$. If $(\mathrm{w}, k, v) \in \mathcal{F}$, then we have that $\mathcal{K}'(k, |\mathcal{K}'(k)| - 1) = (v, t', \_)$, hence $(\mathrm{w}, k, v) \in_{\mathcal{G}'} t'$. Conversely, if $(\mathrm{w}, k, v) \in_{\mathcal{G}'} t'$, then there exists an index $i = 0, \cdots, |\mathcal{K}'(k)| - 1$ such that $\mathcal{K}(k, i) = (v, t', \_)$. We have two possible cases: either $i < |\mathcal{K}'(k, i)| - 1$, leading to $t' \neq t$ and $\mathcal{K}(k, i) = (v, t', \_)$, or equivalently

$(\mathsf{r}, k, v) \in_{\mathcal{G}} t'$; or $i = |\mathcal{K}'(k, i)| - 1$, leading to $t' = t$, and $\mathcal{K}(k, i) = (v, t, \emptyset)$ for some v such that $(\mathsf{w}, k, v) \in \mathcal{F}$.

Putting together the facts above, we obtain that $\mathscr{T}_{\mathcal{G}'} = \mathscr{T}_{\mathcal{G}}[t \mapsto \mathcal{F}]$, as we wanted to prove.

(2) There are two cases that either the t'' already exists in the dependency graph before, or it is the newly committed transaction.

- Suppose that $t' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(k)} t''$. By definition, there exists an index $i = 0, \cdots, |\mathcal{K}(k)| - 1$ such that $\mathcal{K}(k, i) = (\_, t', \{t''\} \cup \_)$. It is immediate to observer, from the definition of update, that $\mathcal{K}'(k, i) = (\_, t', \{t''\} \cup \_)$, and therefore $t' \xrightarrow{\mathsf{WR}_{\mathcal{G}'}(k)} t''$.

- Next, suppose that $(\mathsf{r}, k, \_) \in \mathcal{F}$, and $t' = \max_{\mathsf{WW}_{\mathcal{G}}(k)}(\{\mathsf{w}(\mathcal{K}(k, i)) \mid i \in u(k)\})$. By Definition, $\mathcal{K}(k, i) = (\_, t', \_)$, where $i = \max(u(k))$. This is because $t' \to \mathsf{WW}_{\mathcal{G}}(k)t''$ if and only if $t' = \mathsf{w}(\mathcal{K}(k, j_1))$, $t'' = \mathsf{w}(\mathcal{K}(k, j_2))$ for some j_1, j_2 such that $j_1 < j_2$. The definition of update now ensures that $\mathcal{K}'(k, i) = (\_, t', \{t\} \cup \_)$, from which it follows that $t' \xrightarrow{\mathsf{WR}_{\mathcal{G}'}(k)} t$.

  Conversely, suppose that $t' \xrightarrow{\mathsf{WR}_{\mathcal{G}'}(k)} t''$. Recall that $\mathcal{T}_{\mathcal{G}'} = \mathcal{T}_{\mathcal{G}} \cup \{t\}$, hence either $t'' \in \mathcal{T}_{\mathcal{G}}$ or $t'' = t$.

  – If $t'' = t$, then it must be the case that there exists an index $i = 0, \cdots, |\mathcal{K}'(k)| - 1$ such that $\mathcal{K}'(k, i) = (\_, t', \{t\} \cup \_)$. Note that if $\mathcal{K}'(k, |\mathcal{K}'(k)| - 1)$ is defined, then it must be the case that $\mathcal{K}'(k, |\mathcal{K}'(k)| - 1) = (\_, t, \emptyset)$, hence it must be the case that $i < |\mathcal{K}'(k)| - 1$. Because $t \notin \mathcal{K}$, then by the definition of update it must be the case that $(\mathsf{r}, k, \_) \in \mathcal{F}$, $\mathcal{K}(k, i) = (\_, t', \_)$ and $i = \max(u(k))$; this also implies that $t' = \max_{\mathsf{WW}(k)} \{\mathsf{w}(\mathcal{K}(k, i)) \mid i \in u(k)\}$.

  – If $t'' \in \mathcal{T}_{\mathcal{G}}$, then it must be the case that $t'' \neq t$. In this case, it also must exist an index $i = 0, \cdots, |\mathcal{K}'(k)| - 1$ such that $\mathcal{K}'(k, i) = (\_, t', \{t''\} \cup \_)$. As in the previous case, we note that $i < |\mathcal{K}'(k)| - 1$, which together with the fact that $t'' \neq t$ leads to $\mathcal{K}(k, i) = (\_, t', \{t''\} \cup \_)$. It follows that $t' \xrightarrow{\mathsf{WR}_{\mathcal{G}}(k)} t''$.

(3) Similar to $\mathsf{WR}(k)_{\mathcal{G}}$, there are two cases that either the t'' already exists in the dependency graph before, or it is the newly committed transaction.

- Suppose that $t' \xrightarrow{\mathsf{WW}_{\mathcal{G}}(k)} t''$. By definition, there exist two indexes i, j such that $\mathcal{K}(k, i) = (\_, t', \_)$, $\mathcal{K}(k, j) = (\_, t'', \_)$ and $i < j$. The definition of update ensures that $\mathcal{K}'(k, i) = (\_, t', \_)$, $\mathcal{K}'(k, j) = (\_, t'', \_)$, and because $i < j$ we obtain that $t' \xrightarrow{\mathsf{WW}_{\mathcal{G}'}(k)} t''$.

- Suppose that $(\mathsf{w}, k, \_) \in \mathcal{F}$. Then $\mathcal{K}'(k, |\mathcal{K}(k)|) = (\_, t, \_)$. Let $t' \in \mathcal{T}_{\mathcal{G}}$; by definition there exists an index $i = 0, \cdots, |\mathcal{K}(k)|$ such that $\mathcal{K}(k, i) = (\_, t', \_)$. It follows that $\mathcal{K}'(k, i) = (\_, t', \_)$, and because $i < |\mathcal{K}(k)|$, then we have that $t' \xrightarrow{\mathsf{WW}_{\mathcal{G}'}(k)} t$.

  Conversely, suppose that $t' \xrightarrow{\mathsf{WW}_{\mathcal{G}'}(k)} t''$. Because $\mathcal{T}_{\mathcal{G}'} = \mathcal{T}_{\mathcal{G}} \cup \{t\}$, we have two possibilities. Either $t'' = t$, or $t'' \in \mathcal{T}_{\mathcal{G}}$.

  – If $t'' = t$, then it must be the case that $(\mathsf{w}, k, \_) \in_{\mathcal{G}'} t$, or equivalently there exists an index $i = 0, \cdots, |\mathcal{K}'(k)| - 1$ such that $\mathcal{K}'(k, i) = (\_, t, \_)$. Because $t \notin \mathcal{K}$, and because for any $i = 0, \cdots, |\mathcal{K}(k)| - 1, \mathcal{K}'(k, i) = (\_, t, \_) \Rightarrow \mathcal{K}(k, i) = (\_, t, \_)$, then it necessarily has to be $i = \mathcal{K}'(k) - 1$. According to the definition of update, this is possible only if $(\mathsf{w}, k, \_) \in \mathcal{F}$. Finally, note that because $t' \xrightarrow{\mathsf{WW}_{\mathcal{G}'}(k)} t$, then there exists an index $j < |\mathcal{K}'(k, i)| - 1$ such that $\mathcal{K}'(k, j) = (\_, t', \_)$. The fact that $j < |\mathcal{K}'(k, i)| - 1$ we obtain that $\mathcal{K}(k, j) = (\_, t', \_)$, or equivalently $t' = \mathsf{w}(\mathcal{K}(k, \_))$.

– If $t'' \in \mathcal{T}_\mathcal{G}$, then there exist two indexes i, j such that $j < |\mathcal{K}'(k, j)| - 1$, $\mathcal{K}'(k, j) = (\_, t'', \_)$, $i < j$, and $\mathcal{K}'(k, i) = (\_, t', \_)$. It is immediate to observe that $\mathcal{K}(k, i) = (\_, t', \_)$, $\mathcal{K}(k, j) = (\_, t'', \_)$, from which $t' \xrightarrow{\mathsf{WW}_\mathcal{G}(k)} t''$ follows.

$\square$

D.3 Abstract Execution Traces to KV-Store Traces

We show to construct, given an abstract execution $\mathcal{X}$, a set of $\mathsf{ET}_\top$-traces $\mathsf{KVtrace}(\mathsf{ET}_\top, \mathcal{X})$ in normal form such that for any $\tau \in \mathsf{KVtrace}(\mathsf{ET}_\top, \mathcal{X})$, the trace τ satisfies $\mathsf{lastConf}(\tau) = (\mathcal{K}_\mathcal{X}, \_)$. We first define the $\mathsf{cut}(\mathcal{X}, n)$ function in Def. D.13 which gives the prefix of the first n transactions of the abstract execution $\mathcal{X}$. The $\mathsf{cut}(\mathcal{X}, n)$ function is very useful for later discussion.

Definition D.13. Let $\mathcal{X}$ be an abstract execution, let $n = |\mathcal{T}_\mathcal{X}|$, and let $\{t_i\}_{i=1}^n \subseteq \mathcal{T}_\mathcal{X}$ be such that $t_i \xrightarrow{\mathsf{AR}_\mathcal{X}} t_{i+1}$. The *cut* of the first n transactions from an abstract execution $\mathcal{X}$ is defined as the follows:

$$\mathsf{cut}(\mathcal{X}, 0) \triangleq ([], \emptyset, \emptyset)$$

$$\mathsf{cut}(\mathcal{X}, i + 1) \triangleq \mathsf{extend}\big(\mathsf{cut}(\mathcal{X}, i), t_{i+1}, \mathsf{VIS}_\mathcal{X}^{-1}(t_{i+1}), \mathscr{T}_\mathcal{X}(t_{i+1})\big)$$

PROPOSITION D.14 (WELL-DEFINED cut). *For any abstract execution $\mathcal{X}$, $\mathcal{X} = \mathsf{cut}(\mathcal{X}, |\mathcal{T}_\mathcal{X}|)$.*

PROOF. This is an instantiation of Lemma D.15 by choosing $i = |\mathcal{T}_\mathcal{X}|$. $\square$

LEMMA D.15 (PREFIX). *For any abstract execution $\mathcal{X}$, and index $i : i \le j \le |\mathcal{T}_\mathcal{X}|$, if $\mathcal{T}_\mathcal{X} = \{t_i\}_{i=1}^n$ be such that $t_i \xrightarrow{\mathsf{AR}_\mathcal{X}} t_{i+1}$, then $\mathsf{cut}(\mathcal{X}, i) = \mathcal{X}_i$ where*

$$\mathscr{T}_{\mathcal{X}_i}(t) = \begin{cases} \mathscr{T}_\mathcal{X}(t) & \text{if } \exists j \le i. \ t = t_j \\ undefined & otherwise \end{cases}$$

$$\mathsf{VIS}_{\mathcal{X}_i} = \left\{ (t, t') \in \mathcal{T}_{\mathcal{X}_i} \ \middle| \ t \xrightarrow{\mathsf{VIS}_\mathcal{X}} t' \right\}$$

$$\mathsf{AR}_{\mathcal{X}_i} = \left\{ (t, t') \in \mathcal{T}_{\mathcal{X}_i} \ \middle| \ t \xrightarrow{\mathsf{AR}_\mathcal{X}} t' \right\}$$

PROOF. Fix an abstract execution $\mathcal{X}$. We prove by induction on $i = |\mathcal{T}_\mathcal{X}|$.

- Base case: $i = 0$. Then $\mathscr{T}_{\mathcal{X}'} = []$, $\mathsf{VIS}_{\mathcal{X}'} = \emptyset$, $\mathsf{AR}_{\mathcal{X}'} = \emptyset$, which leads to $\mathcal{X}' = \mathsf{cut}(\mathcal{X}, 0)$.
- Inductive case: $i = i' + 1$. Assume that $\mathsf{cut}(\mathcal{X}, i') = \mathcal{X}_{i'}$. We prove the following:
  - $\mathscr{T}_{\mathsf{cut}(\mathcal{X}, i)} = \mathscr{T}_{\mathcal{X}_i}$. By definition,

$$\mathscr{T}_{\mathsf{cut}(\mathcal{X}, i)} = \mathscr{T}_{\mathsf{cut}(\mathcal{X}, i')}[t_i \mapsto \mathscr{T}_\mathcal{X}(t_i)]\mathscr{T}_{\mathcal{X}_{i'}}[t_i \mapsto \mathscr{T}_\mathcal{X}](t_i) = \mathscr{T}_{\mathcal{X}_i}$$

  - $\mathsf{VIS}_{\mathsf{cut}(\mathcal{X}, i)} = \mathsf{VIS}_{\mathcal{X}_i}$. Note that, by inductive hypothesis, $\mathcal{T}_{\mathsf{cut}(\mathcal{X}, i')} = \mathcal{T}_{\mathcal{X}_{i'}} = \{t_j\}_{j=1}^{i'}$. We have that

$$\begin{aligned} \mathsf{VIS}_{\mathsf{cut}(\mathcal{X}, i)} &= \mathsf{VIS}_{\mathsf{cut}(\mathcal{X}, i')} \cup \left\{ (t_j, t_i) \in \mathsf{VIS}_\mathcal{X} \ \middle| \ 1 \le j \le i' \right\} \\ &= \mathsf{VIS}_{\mathcal{X}_{i'}} \cup \left\{ (t_j, t_i) \in \mathsf{VIS}_\mathcal{X} \ \middle| \ 1 \le j \le i' \right\} \\ &= \left\{ (t_{j'}, t_j) \in \mathsf{VIS}_\mathcal{X} \ \middle| \ 1 \le j \le i' \right\} \cup \left\{ (t_j, t_i) \in \mathsf{VIS}_\mathcal{X} \ \middle| \ 1 \le j \le i' \right\} \\ &= \left\{ (t_{j'}, t_j) \in \mathsf{VIS}_\mathcal{X} \ \middle| \ 1 \le j \le i' \right\} \\ &= \mathsf{VIS}_{\mathcal{X}_i} \end{aligned}$$

  - $\mathsf{AR}_{\mathsf{cut}(\mathcal{X}, i)} = \mathsf{AR}_{\mathcal{X}_i}$. It follows the same way as the above.

$\square$

Let $\textsc{Client}(X) \triangleq \{cl \mid \exists n.\ t_{cl}^n \in \mathcal{T}_X\}$. Given an abstract execution X, a client cl and an index i : $0 \leq i < |\mathcal{T}_X|$, the function $\mathsf{nextTid}(X, cl, i) \triangleq \min_{\mathsf{AR}_X} \left\{ t_{cl}^j \mid t_{cl}^n \notin \mathcal{T}_{\mathsf{cut}(X,i)} \right\}$. Note that $\mathsf{nextTid}(X, cl, i)$ could be undefined. The conversion from abstract execution tests to ET traces is in Def. D.16.

Definition D.16. Given an abstract execution X and an index $i : 0 \leq i < |\mathcal{T}_X|$, the function $\mathsf{KVtrace}(\mathsf{ET}_\top, X, i)$ is defined as the smallest set such that:

- $(\mathcal{K}_0, \lambda cl \in \textsc{Client}(X).\lambda k.\ \{0\}) \in \mathsf{KVtrace}(\mathsf{ET}_\top, X, 0)$,
- suppose that $\tau \in \mathsf{KVtrace}(\mathsf{ET}_\top, X, i)$ for some i. Let
  - $t = \min_{\mathsf{AR}_X}(\mathcal{T}_X \setminus T_{\mathsf{cut}(X,i)})$,
  - cl, n be such that $t = t_{cl}^n$,
  - $u = \mathsf{getView}\big(X, \mathsf{VIS}_X^{-1}(t_{cl}^n)\big)$,
  - $u' = \mathsf{getView}(X, \mathcal{T})$, where $\mathcal{T}$ is an arbitrary subset of $\mathcal{T}_X$ if $\mathsf{nextTid}(X, cl, i+1)$ is undefined, or is such that $\mathcal{T} \subseteq (\mathsf{AR}_X^{-1})^?(t) \cap \mathsf{VIS}_X^{-1}(\mathsf{nextTid}(cl, i+1))$,
  - $\mathcal{F} = \mathscr{T}_X(t)$,
  - $(\mathcal{K}_\tau, \mathcal{U}_\tau) = \mathsf{lastConf}(\tau)$, and
  - $\mathcal{K} = \mathsf{update}(\mathcal{K}_\tau, u, \mathcal{F}, t)$.

  Then
  $$\left(\tau \xrightarrow{(cl, \varepsilon)}_{\mathsf{ET}_\top} (\mathcal{K}_\tau, \mathcal{U}_\tau[cl \mapsto u]) \xrightarrow{(cl, \mathcal{F})}_{\mathsf{ET}_\top} (\mathcal{K}, \mathcal{U}_\tau[cl \mapsto u']) \right) \in \mathsf{KVtrace}(\mathsf{ET}_\top, X, i+1)$$

Last, the function $\mathsf{KVtrace}(\mathsf{ET}_\top, X) \triangleq \mathsf{KVtrace}(\mathsf{ET}_\top, X, |\mathcal{T}_X|)$.

PROPOSITION D.17 (ABSTRACT EXECUTIONS TO TRACE $\mathsf{ET}_\top$). *Given an abstract execution X satisfying* $\mathsf{RP}_{\mathsf{LWW}}$, *and a trace* $\tau \in \mathsf{KVtrace}(\mathsf{ET}_\top, X)$, *then* $\mathsf{lastConf}(\tau) = (\mathcal{K}_X, \_)$ *and* $\mathcal{K}_X \in \mathsf{CM}(\mathsf{ET}_\top)$.

PROOF. Let X be an abstract execution that satisfies the last write wins policy. Let $n = |\mathcal{T}_X|$. Fix $i = 0, \cdots, n$, and let $\tau \in \mathsf{KVtrace}(\mathsf{ET}_\top, X, i)$. We prove, by induction on i, that $\tau \in \mathsf{CM}(\mathsf{ET}_\top)$, and $\mathsf{lastConf}(\tau) = (\mathcal{K}_{(\mathsf{cut}(X,i)}, \_)$. Then the result follows from Prop. D.14.

- Base case: $i = 0$. By definition, $\tau = (\mathcal{K}_0, \mathcal{U}_0)$, where $\mathcal{U}_0 = \lambda cl \in \textsc{Client}(X).\lambda k.\ \{0\}$. Clearly, we have that $\tau \in \mathsf{CM}(\mathsf{ET}_\top)$.
- Inductive case: $i = i' + 1$. Let $t_i = \min_{\mathsf{AR}_X}(\mathcal{T}_X \setminus \mathcal{T}_{\mathsf{cut}(X,i')})$, and suppose that $t_i = t_{cl}^m$ for some client cl and index m. Fix $u = \mathsf{getView}\big(X, \mathsf{VIS}_X^{-1}(t_i)\big)$, and $\mathcal{F} = \mathscr{T}_X(t_i)$. We prove that there exists a trace $\tau' \in \mathsf{KVtrace}(\mathsf{ET}_\top, X, i')$ and a set $\mathcal{T}$ such that:

(1) if $\mathsf{nextTid}(cl, X, i)$ is undefined then $\mathcal{T} \subseteq \mathcal{T}_X$, otherwise
$$\mathcal{T} \subseteq \mathsf{VIS}_X^{-1}(\mathsf{nextTid}(cl, X, i)) \cap (\mathsf{AR}_X^{-1})^?(t_i)$$

(2) the new trace τ such that
$$\tau = \tau' \xrightarrow{(cl, \varepsilon)} (\mathcal{K}_{\tau'}, \mathcal{U}_{\tau'}[cl \mapsto u]) \xrightarrow{(cl, \mathcal{F})} (\mathcal{K}, \mathcal{U}_{\tau'}[cl \mapsto u'])$$

where $(\mathcal{K}_{\tau'}, \mathcal{U}_{\tau'}) = \mathsf{lastConf}(\tau')$, and $\mathcal{K} = \mathsf{update}(\mathcal{K}_{\tau'}, u, \mathcal{F}, t_i)$, and $u' = \mathsf{getView}(X, \mathcal{T})$. By inductive hypothesis, we may assume that $\tau' \in \mathsf{CM}(\mathsf{ET}_\top)$, and $\mathcal{K}_{\tau'} = \mathcal{K}_{\mathsf{cut}(X,i')}$. We prove the following facts:

(1) $\mathcal{K} = \mathcal{K}_{\mathsf{cut}(X,i)}$. Because of Prop. D.10 and Prop. D.14, we obtain
$$\begin{aligned}
\mathcal{K} &= \mathsf{update}(\mathcal{K}_{\tau'}, u, \mathcal{F}, t_i) \\
&= \mathsf{update}\big(\mathcal{K}_{\mathsf{cut}(X,i')}, \mathsf{getView}\big(X, \mathsf{VIS}_X^{-1}(t_i)\big), \mathscr{T}_X(t_i), t_i\big) \\
&= \mathcal{K}_{\mathsf{extend}\big(\mathsf{cut}(X,i'), \mathsf{VIS}_X^{-1}(t_i), t_i, \mathscr{T}_X(t_i)\big)} \\
&= \mathcal{K}_{\mathsf{extend}(\mathsf{cut}(X,i))}
\end{aligned}$$

(2) $(\mathcal{K}_{\tau'}, \mathcal{U}_{\tau'}) \xrightarrow{(cl, \varepsilon)} (\mathcal{K}_{\tau'}, \mathcal{U}_{\tau'}[cl \mapsto u])$. It suffices to prove that $\mathcal{U}_{\tau'}(cl) \sqsubseteq u$ for any key k. By Lemma D.15 we have that $\mathcal{T}_{\mathrm{cut}(\mathcal{X}, i')} = \{t_j\}_{j=1}^{i'}$, for some $t_1, \cdots, t_{i'}$ such that whenever $1 \le j < j' \le i'$, then $t_j \xrightarrow{\mathrm{AR}_\mathcal{X}} t_{j'}$. We consider two possible cases:

- For all $j : 1 \le j \le i'$, and $h \in \mathbb{N}$, then $t_j \ne t_{cl}^h$. In this case we have that no transition contained in τ' has the form $(\_, \_) \xrightarrow{(cl, \_)} (\_, \_)$, from which it is possible to infer that $\mathcal{U}_{\tau'}(cl) = \lambda k. \{0\}$. Because $u = \mathrm{getView}(\mathcal{X}, \mathrm{VIS}_\mathcal{X}^{-1}(t_i))$, then by definition we have that $0 \in u(k)$ for all keys $k \in \mathrm{KEY}$. It follows that $\mathcal{U}_{\tau'}(cl) \sqsubseteq u$.

- There exists an index $j : 1 \le j \le i'$ and an integer $h \in \mathbb{N}$ such that $t_j = t_{cl}^h$. Without loss of generality, let j be the largest such index. It follows that the last transition in τ' of the form $(\_, \_) \xrightarrow{(cl, \mathcal{F}_j)} (\_, \mathcal{U}_{\mathrm{pre}})$ is such that $\mathcal{U}_{\mathrm{pre}}(cl) = \mathrm{getView}(\mathcal{X}, \mathcal{T}_{\mathrm{pre}})$, for some $\mathcal{T}_{\mathrm{pre}} \subseteq \mathrm{VIS}_\mathcal{X}^{-1}(t_i) \cap (\mathrm{AR}_\mathcal{X}^{-1})^?(t_j)$. This is because $\mathrm{nextTid}(cl, \mathcal{X}, j)$ is defined and equal to t_i. Furthermore, because the trace τ' is in normal form by construction, in τ' a transition of the form $(\_, \_) \xrightarrow{(cl, \varepsilon)}_{\mathrm{ET}_\top} (\_, \_)$ is always followed by a transition of the form $(\_, \_) \xrightarrow{(cl, \mathcal{F}')}_{\mathrm{ET}_\top} (\_, \_)$. Because we assume that the last transition where client cl executes a transaction in τ' has the form $(\_, \_) \xrightarrow{(cl, \mathcal{F}_j)}_{\mathrm{ET}_\top} (\_, \mathcal{U}_{\mathrm{pre}})$, then the latter is also the last transition for client cl in τ' (i.e. including both execution of transactions and view updates). It follows that $\mathcal{U}_{\tau'}(cl) = \mathcal{U}_{\mathrm{pre}}(cl)$, and in particular $\mathcal{U}_{\tau'}(cl) = \mathrm{getView}(\mathcal{X}, \mathcal{T}_{\mathrm{pre}})$. By definition, $\mathcal{T}_{\mathrm{pre}} \subseteq \mathrm{VIS}_\mathcal{X}^{-1}(t_i) \cap (\mathrm{AR}_\mathcal{X}^{-1})^?(t_j) \subseteq \mathrm{VIS}_\mathcal{X}^{-1}(t_i)$. By Lemma D.18, we have that $\mathcal{U}_{\tau'}(cl) = \mathrm{getView}(\mathcal{X}, \mathcal{T}_{\mathrm{pre}}) \sqsubseteq \mathrm{getView}(\mathcal{X}, \mathrm{VIS}_\mathcal{X}^{-1}(t_i)) = u$, as we wanted to prove.

(3) $(\mathcal{K}_{\tau'}, \mathcal{U}_{\tau'}[cl \mapsto u]) \xrightarrow{(cl, \mathcal{F})}_{\mathrm{ET}_\top} (\mathcal{K}, \mathcal{U}_{\tau'}[cl \mapsto u'])$. It suffices to show that $\mathrm{ET}_\top \vdash (\mathcal{K}_{\tau'}, u) \triangleright \mathcal{F} : (\mathcal{K}, u')$. That is, it suffices to show that $u \in \mathrm{VIEWS}(\mathcal{K}_{\tau'})$, $u' \in \mathrm{VIEWS}(\mathcal{K})$, and whenever $(r, k, v) \in \mathcal{F}$, then $\max_<(u(k)) = (v, \_, \_)$. The first two facts are a consequence of Lemma D.19, $\mathcal{K}_{\tau'} = \mathcal{K}_{\mathrm{cut}(\mathcal{X}, i')}$, and $\mathcal{K}_{\mathrm{cut}(\mathcal{X}, i)}$. The last one that if $(r, k, v) \in \mathcal{F}$ then $\max_<(u(k)) = (v, \_, \_)$ follows the fact that $\mathcal{X}$ satisfies the last write wins policy and the fact that $u = \mathrm{getView}(\mathrm{VIS}_\mathcal{X}^{-1}(t_i))$.

□

LEMMA D.18 (MONOTONIC getView). *Let $\mathcal{X}$ be an abstract execution, and let $\mathcal{T}_1 \subseteq \mathcal{T}_2 \subseteq \mathcal{T}_\mathcal{X}$. Then $\mathrm{getView}(\mathcal{X}, \mathcal{T}_1) \sqsubseteq \mathrm{getView}(\mathcal{X}, \mathcal{T}_2)$.*

PROOF. Fix $k \in \mathrm{KEY}$. By definition

$$\begin{aligned}\mathrm{getView}(\mathcal{X}, \mathcal{T}_1)(k) &= \{0\} \cup \{i \mid \mathrm{w}(\mathcal{K}_\mathcal{X}(k, i)) \in \mathcal{T}_1\} \\ &\subseteq \{0\} \cup \{i \mid \mathrm{w}(\mathcal{K}_\mathcal{X}(k, i)) \in \mathcal{T}_2\} \\ &= \mathrm{getView}(\mathcal{X}, \mathcal{T}_2)(k)\end{aligned}$$

then it follows that $\mathrm{getView}(\mathcal{X}, \mathcal{T}_1) \sqsubseteq \mathrm{getView}(\mathcal{X}, \mathcal{T}_2)$. □

LEMMA D.19 (VALID VIEW ON CUT OF ABSTRACT EXECUTION). *Let $\mathcal{X}$ be an abstract execution, with $\mathcal{T}_\mathcal{X} = \{t_i\}_{i=1}^n$ for $n = |\mathcal{T}_\mathcal{X}|$, and $i : 0 \le i < n$ such that $t_i \xrightarrow{\mathrm{AR}_\mathcal{X}} t_{i+1}$. Assuming $\mathcal{T}_0 = \emptyset$, and $\mathcal{T}_i \subseteq (\mathrm{AR}^{-1})^?(t_i)$ for $i : 0 \le i \le n$, then $\mathrm{getView}(\mathcal{X}, \mathcal{T}_i) \in \mathrm{VIEWS}(\mathcal{K}_{\mathrm{cut}(\mathcal{X}, i)})$.*

PROOF. We prove by induction on the index i.

- Base case: $i = 0$. It follows $\mathcal{T}_0 = \emptyset$, and $\mathrm{getView}(\mathcal{X}, \mathcal{T}_0) = \lambda k. \{0\}$. We also have that $\mathcal{K}_{\mathrm{cut}(\mathcal{X}, 0)} = \lambda k. (v_0, t_0, \emptyset)$, hence it is immediate to see that $\mathrm{getView}(\mathcal{X}, \mathcal{T}_0) \in \mathrm{VIEWS}(\mathcal{K}_{\mathrm{cut}(\mathcal{X}, 0)})$.

- Inductive case: $i = i' + 1$. Suppose that for any $\mathcal{T} \subseteq (\mathsf{AR}_X^{-1})^?(t_{i'})$, then $\mathsf{getView}(X, \mathcal{T}) \in \mathsf{Views}(\mathcal{K}_{\mathsf{cut}(X,i)})$. Let consider the set $\mathcal{T}_i$. Note that, because of Prop. D.10, we have that

$$\mathcal{K}_{\mathsf{cut}(X,i)} = \mathcal{K}_{\mathsf{extend}\left(\mathsf{cut}(X,i'),\, t_i,\, \mathsf{VIS}_X^{-1}(t_i),\, \mathscr{T}_X(t_i)\right)} = \mathsf{update}\left(\mathcal{K}_{\mathsf{cut}(X,i')},\, \mathsf{getView}\left(\mathsf{VIS}_X^{-1}(t_i)\right),\, \mathscr{T}_X(t_i),\, t_i\right)$$

   There are two possibilities:

   - $t_i \notin \mathcal{T}_i$, where case $\mathcal{T}_i \subseteq (\mathsf{AR}_X^{-1})^?(t_{i'})$. From the inductive hypothesis we get $\mathsf{getView}(X, \mathcal{T}_i) \in \mathsf{Views}(\mathcal{K}_{\mathsf{cut}(X,i')})$. Note that $\mathcal{K}_{\mathsf{cut}(X,i')}$ only contains the transactions identifiers from t_1 to $t_{i'}$; in particular, it does not contain t_i. Because $\mathcal{K}_{\mathsf{cut}(X,i)} = \mathsf{update}(\mathcal{K}_{\mathsf{cut}(X,i')}, \_, \_, t_i)$, then by Lemma D.20 we have that $\mathsf{getView}(X, t_i) \in \mathsf{Views}(\mathcal{K}_{\mathsf{cut}(X,i)})$.
   - $t \in \mathcal{T}_i$. Note that for any key k such that $(\mathsf{w}, k, \_) \notin \mathscr{T}_X(t_i)$, then $\mathsf{getView}(X, \mathcal{T}_i)(k) = \mathsf{getView}(X, \mathcal{T}_i \setminus \{t_i\})(k)$; and for any key k such that $(\mathsf{w}, k, \_) \in \mathscr{T}_X(t_i)$, then $\mathsf{getView}(X, \mathcal{T}_i)(k) = \mathsf{getView}(X, \mathcal{T}_i \setminus \{t_i\})(k) \cup \{j \mid \mathsf{w}(\mathcal{K}_X(k,i)) = t_j\}$. In the last case, the index j must be such that $j < |\mathcal{K}_{\mathsf{cut}(X,i)}| - 1$, because we know that $t_i \in \mathcal{K}_{\mathsf{cut}(X,i)}$. It follows from this fact and the inductive hypothesis, that $\mathsf{getView}(X, \mathcal{T}_i) \in \mathsf{Views}(\mathcal{K}_{\mathsf{cut}(X,i)})$.

$\square$

LEMMA D.20 (update PRESERVING VIEWS). *Given a kv-store $\mathcal{K}$, a transactions $t \notin \mathcal{K}$, views $u, u' \in \mathsf{Views}(\mathcal{K})$, and set of operations $\mathcal{F}$, then $u \in \mathsf{update}(\mathcal{K}, u', \mathcal{F}, t)$.*

PROOF. Immediate from the definition of update. Note that $t \notin \mathcal{K}$ ensures that u still satisfies (atomic) with respect to the new kv-store $\mathsf{update}(\mathcal{K}, u', \mathcal{F}, t)$. $\square$

E THE SOUND AND COMPLETE CONSTRUCTORS OF THE KV-STORE SEMANTICS WITH RESPECT TO ABSTRACT EXECUTIONS

In this Section we first define the set of ET-traces generated by a program P. Then we prove correctness our semantics on kv-stores, meaning that if a program P executing under the execution test ET terminates in a state $(\mathcal{K}, \_)$, then $\mathcal{K} \in \mathsf{CM}(\mathsf{ET})$.

E.1 Traces of Programs under KV-Stores

The $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P})$ is the set of all possible traces generated by the program P starting from the initial configuration $(\mathcal{K}_0, \mathcal{U}_0)$.

Definition E.1. Given an execution test ET a program P and a state $(\mathcal{K}, \mathcal{U}, \mathcal{E})$, the $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}, \mathcal{K}, \mathcal{U}, \mathcal{E})$ function is defined as the smallest set such that:

- $(\mathcal{K}, \mathcal{U}) \in \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}, \mathcal{K}, \mathcal{U}, \mathcal{E})$
- if $\tau \in \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}', \mathcal{K}', \mathcal{U}', \mathcal{E}')$ and $((\mathcal{K}, \mathcal{U}, \mathcal{E}), \mathsf{P}) \xrightarrow{(cl, \iota)}_{\mathsf{ET}} (\mathcal{K}', \mathcal{U}', \mathcal{E}')$, then

$$\tau \in \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}, \mathcal{K}, \mathcal{U}, \mathcal{E}')$$

- if $\tau \in \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}', \mathcal{K}', \mathcal{U}', \mathcal{E}')$ and $(\mathcal{K}, \mathcal{U}, \mathcal{E}), \mathsf{P}) \xrightarrow{(cl, u, \mathcal{F})} ((\mathcal{K}', \mathcal{U}', \mathcal{E}'), \mathsf{P}')$, then

$$\left((\mathcal{K}, \mathcal{U}) \xrightarrow{(cl, \varepsilon)}_{\mathsf{ET}} (\mathcal{K}, \mathcal{U}[cl \mapsto u]) \xrightarrow{(cl, \mathcal{F})}_{\mathsf{ET}} \tau \right) \in \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}, \mathcal{K}, \mathcal{U}, \mathcal{E})$$

The set of traces generated by a program P under the execution test ET is then defined as $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}) \triangleq \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}, \mathcal{K}_0, \mathcal{U}_0, \mathcal{E}_0)$, where $\mathcal{U}_0 = \lambda cl \in \mathrm{dom}(\mathsf{P}).\lambda k. \{0\}$, and $\mathcal{E}_0 = \lambda cl \in \mathrm{dom}(\mathsf{P}).\lambda a.0$.

PROPOSITION E.2. *For any program P and execution test* ET, $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}) \subseteq \mathrm{conf}(\mathsf{ET})$ *and* $\tau \in \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P})$ *is in normal form.*

PROOF. First, by the definition of Ptraces, it only constructs trace in normal form. It is easy to prove that for any trace τ in $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P})$, by induction on the trace length, the trace is also in $\mathrm{conf}(\mathsf{ET})$. □

COROLLARY E.3. *If a trace in the following form*

$$(\mathcal{K}_0, \mathcal{U}_0, \mathcal{E}_0), \mathsf{P}) \rightarrow_{\mathsf{ET}} \cdots \rightarrow_{\mathsf{ET}} (\mathcal{K}, \mathcal{U}, \mathcal{E}, \lambda cl \in \mathrm{dom}(\mathsf{P}).\mathtt{skip})$$

then $\mathcal{K} \in \mathsf{CM}(\mathsf{ET})$.

PROOF. By the definition of Ptraces, there exists a corresponding trace $\tau \in \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P})$. By Prop. E.2, such trace $\tau \in \mathrm{conf}(\mathsf{ET})$, therefore $\mathcal{K} \in \mathsf{CM}(\mathsf{ET})$ by definition of $\mathsf{CM}(\mathsf{ET})$. □

Similar to $\llbracket \mathsf{P} \rrbracket_{(\mathsf{RP}, \mathcal{A})}$, the function $\llbracket \mathsf{P} \rrbracket_{\mathsf{ET}}$ is defined as the following:

$$\llbracket \mathsf{P} \rrbracket_{\mathsf{ET}} = \left\{ \mathcal{K} \mid (\mathcal{K}_0, \mathcal{U}_0, \mathcal{E}_0), \mathsf{P} \rightarrow^*_{\mathsf{ET}} (\mathcal{K}, \_, \_), \mathsf{P}_f) \right\}$$

where $\mathcal{E}_0 = \lambda cl \in \mathrm{dom}(\mathsf{P}).\lambda \mathtt{x}.0$ and $\mathsf{P}_f = \lambda cl \in \mathrm{dom}(\mathsf{P}).\mathtt{skip}$.

PROPOSITION E.4. *For any program P and execution test* ET: $\llbracket \mathsf{P} \rrbracket_{\mathsf{ET}} = \llbracket \mathsf{P} \rrbracket_{\mathsf{ET}_\top} \cap \mathsf{CM}(\mathsf{ET})$.

PROOF. We prove a stronger result that for any program P and execution test ET, $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}) = \mathsf{Ptraces}(\mathsf{ET}_\top, \mathsf{P}) \cap \mathrm{conf}(\mathsf{ET})$. It is easy to see $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}) \subseteq \mathsf{Ptraces}(\mathsf{ET}_\top, \mathsf{P})$. By Prop. E.2, we know $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}) \subseteq \mathrm{conf}(\mathsf{ET})$. Therefore $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P}) \subseteq \mathsf{Ptraces}(\mathsf{ET}_\top, \mathsf{P}) \cap \mathrm{conf}(\mathsf{ET})$.

Let consider a trace τ in $\mathsf{Ptraces}(\mathsf{ET}_\top, \mathsf{P}) \cap \mathrm{conf}(\mathsf{ET})$. By inductions on the length of trace, every step that commits a new transaction must satisfy ET as $\tau \in \mathrm{conf}(\mathsf{ET})$. It also reduce the program P since $\tau \in \mathsf{Ptraces}(\mathsf{ET}_\top, \mathsf{P})$. By the definition $\mathsf{Ptraces}(\mathsf{ET}, \mathsf{P})$, we can construct the same trace τ so that $\tau \in \mathsf{Ptraces}(\mathsf{ET}, \mathsf{P})$. □

E.2 Adequate of KV-Store Semantic

Our main aim is to prove that for any program P, the set of kv-stores generated by P under ET corresponds to all the possible abstract executions that can be obtained by running P on a database that satisfies the axiomatic definition $\mathcal{A}$. In this sense, we aim to establish that our operational semantics is *adequate*.

More precisely, suppose that a given execution test ET captures precisely a consistency model defined in the axiomatic style, using a set of axioms $\mathcal{A}$ and a resolution policy RP over abstract executions. That is, for any abstract execution X that satisfies the axioms $\mathcal{A}$ and the resolution policy RP, then $\text{KVtrace}(\text{ET}_\top, X) \cap \text{CM}(\text{ET}) \neq \emptyset$; and for any $\tau \in \text{CM}(\text{ET})$, there exists an abstract execution $X \in \text{absExec}(\tau)$ that satisfies the axioms $\mathcal{A}$ and the resolution policy RP.

We now consider the program P. The Prop. E.5 and Prop. E.6 show the connection between reduction steps between the last write win resolution policy $(\text{RP}_{\text{LWW}}, \emptyset)$ and the most permissive execution test $\text{ET}_\top$.

PROPOSITION E.5 (PERMISSIVE EXECUTION TEST TO LAST WRITE WIN). *Suppose that*

$$(\mathcal{K}, \mathcal{U}, \mathcal{E}), \text{P} \xrightarrow{(cl, u, \mathcal{F})}_{\text{ET}_\top} (\mathcal{K}', \mathcal{U}', \mathcal{E}'), \text{P}'$$

Assuming an abstract execution X such that $\mathcal{K}_X = \mathcal{K}$, and a set of read-only transactions $\mathcal{T} \subseteq \mathcal{T}_X$, then there exists an abstract execution X' such that $\mathcal{K}_{X'} = \mathcal{K}'$, and

$$(X, \mathcal{E}), \text{P} \xrightarrow{(cl, \mathcal{T} \cup \text{visTx}(\mathcal{K}, u), \mathcal{F})}_{(\text{RP}_{\text{LWW}}, \emptyset)} (X', \mathcal{E}'), \text{P}'$$

PROOF. Suppose that $(\mathcal{K}, \mathcal{U}, \mathcal{E}), \text{P} \xrightarrow{(cl, u, \mathcal{F})}_{\text{ET}_\top} (\mathcal{K}', \mathcal{U}', \mathcal{E}), \text{P}'$. This transition can only be inferred by applying Rule PSINGLETHREAD, meaning that

- $\text{P}(cl) \mapsto \text{C}$ for some command C,
- $cl \vdash (\mathcal{K}, \mathcal{U}(cl), \mathcal{E}(cl)), \text{C} \xrightarrow{(cl, u, \mathcal{F})}_{\text{ET}_\top} (\mathcal{K}', u', s'), \text{C}'$ for some u', s', and
- $\mathcal{U}' = \mathcal{U}[cl \mapsto u'], \mathcal{E}' = \mathcal{E}[cl \mapsto s']$ and $\text{P}' = \text{P}[cl \mapsto \text{C}']$.

Let X be such that $\mathcal{K}_X = \mathcal{K}$, and let $\mathcal{T} \subseteq \mathcal{T}_X$ be a set of read-only transactions in X. It suffices to show that there exists an abstract execution X' such that $\mathcal{K}_{X'} = \mathcal{K}'$, and

$$cl \vdash (X, \mathcal{E}(cl)), \text{C} \xrightarrow{(cl, \mathcal{T} \cup \text{visTx}(\mathcal{K}, u), \mathcal{F})}_{(\text{RP}_{\text{LWW}}, \emptyset)} (X', s'), \text{C}'.$$

By the ASINGLETHREAD rule, we obtain

$$(X, \mathcal{E})\text{P}, \xrightarrow{(cl, \mathcal{T} \cup \text{visTx}(\mathcal{K}, u), \mathcal{F})}_{(\text{RP}_{\text{LWW}}, \emptyset)} (X', \mathcal{E}'), \text{P}'$$

Now we perform a rule induction on the derivation of the transition

$$cl \vdash (\mathcal{K}, \mathcal{U}(cl), \mathcal{E}(cl)), \text{C} \xrightarrow{(cl, u, \mathcal{F})}_{\text{ET}_\top} (\mathcal{K}', u', ss'), \text{C}'$$

Base case: PCOMMIT. This implies that

- $\text{C} = [\text{T}]$ for some T, and $\text{C}' = \text{skip}$,
- $\mathcal{U}(cl) \sqsubseteq u$,
- let $ss = \text{snapshot}(\mathcal{K}, u)$; then $(\mathcal{E}(cl), ss, \emptyset) \rightarrow^* (s', \_, \mathcal{F})$,
- $\mathcal{K}' = \text{update}(\mathcal{K}, u, \mathcal{F}, t)$ for some $t \in \text{nextTid}(\mathcal{K}, cl)$, and
- $\text{ET}_\top \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$.

Choose an arbitrary set of of read-only transactions $\mathcal{T} \subseteq \mathcal{T}_X$. We observe that $\text{getView}(X, \mathcal{T} \cup \text{visTx}(\mathcal{K}, u)) = u$ since $\mathcal{K}_X = \mathcal{K}$ and Prop. D.6. We can now apply Prop. D.5 and ensure that $\text{RP}_{\text{LWW}}(X, \mathcal{T} \cup \text{visTx}(\mathcal{K}, u)) = \{ss\}$. Let $X' = \text{extend}(X, t, \mathcal{T} \cup \text{visTx}(\mathcal{K}, u))$. Because $\text{getView}(X, \mathcal{T} \cup \text{visTx}(\mathcal{K}, u)) = u, \mathcal{K}_X = \mathcal{K}$, then by Prop. D.10 we have that $\mathcal{K}_{X'} = \text{update}(\mathcal{K}, u, t, \mathcal{F}) = \mathcal{K}'$. To summarise, we

have that $\mathcal{T} \cup \text{visTx}(\mathcal{K}, u) \subseteq \mathcal{T}_{\mathcal{X}}$, $ss \in \text{RP}_{\text{LWW}}(\mathcal{X}, \mathcal{T} \cup \text{visTx}(\mathcal{K}, u))$, $(\mathcal{E}(cl), ss, \emptyset) \to^* (s', \_, \mathcal{F})$ and $t \in \text{nextTid}(\mathcal{T}_{\mathcal{X}}, cl)$. Now we can apply ACOMMIT and infer

$$cl \vdash (\mathcal{X}, \mathcal{E}(cl)), [\mathsf{T}] \xrightarrow{cl, \mathcal{T} \cup \text{visTx}(\mathcal{K}_{\mathcal{X}}, u)}_{(\text{RP}_{\text{LWW}}, \emptyset)} (\mathcal{X}', s'), \mathtt{skip}$$

which is exactly what we wanted to prove.

Base case: PPRIMITIVE,PCHOICE,PITER,PSEQSKIP. These cases are trivial since they do not alter the state of $\mathcal{K}$. Inductive case: PSEQ. It is derived by the I.H.   □

PROPOSITION E.6 (LAST WRITE WIN TO PERMISSIVE EXECUTION TEST). *Suppose that*

$$(\mathcal{X}, \mathcal{E}), \mathsf{P} \xrightarrow{(cl, \mathcal{T}, \mathcal{F})}_{(\text{RP}_{\text{LWW}}, \emptyset)} (\mathcal{X}', \mathcal{E}'), \mathsf{P}'$$

Then for any $\mathcal{U}$ and $u \in \text{VIEWS}(\mathcal{K}_{\mathcal{X}})$ such that $u \sqsubseteq \text{getView}(\mathcal{X}, \mathcal{T})$, the following holds:

$$(\mathcal{K}_{\mathcal{X}}, \mathcal{U}[cl \mapsto u], \mathcal{E}), \mathsf{P} \xrightarrow{(cl, \text{getView}(\mathcal{X}, \mathcal{T}), \mathcal{F})}_{\text{ET}_{\top}} (\mathcal{K}_{\mathcal{X}'}, \mathcal{U}, \mathcal{E}'), \mathsf{P}'$$

PROOF. Suppose that

$$(\mathcal{X}, \mathcal{E}), \mathsf{P} \xrightarrow{(cl, \mathcal{T}, \mathcal{F})}_{(\text{RP}_{\text{LWW}}, \emptyset)} (\mathcal{X}', \mathcal{E}'), \mathsf{P}'$$

Fix a function $\mathcal{U}$ from clients in $\text{dom}(\mathsf{P})$ to views in $\text{VIEWS}(\mathcal{K})$, and a view $u \sqsubseteq \text{getView}(\mathcal{X}, \mathcal{T})$. We show that $(\mathcal{K}_{\mathcal{X}}, \mathcal{U}[cl \mapsto u], \mathcal{E}) \xrightarrow{(cl, \text{getView}(\mathcal{X}, \mathcal{T}), \mathcal{F})}_{\text{ET}_{\top}} (\mathcal{K}_{\mathcal{X}'}, \mathcal{U}, \mathcal{E}'), \mathsf{P}'$.

Note that the transition $\mathcal{X}, \mathcal{E}, \mathsf{P} \xrightarrow{(cl, \mathcal{T}, \mathcal{F})}_{(\text{RP}_{\text{LWW}}, \emptyset)} (\mathcal{X}', \mathcal{E}'), \mathsf{P}'$ can only be inferred using ASIN-GLETHREAD rule, from which it follows that

$$cl \vdash (\mathcal{X}, \mathcal{E}(cl)), \mathsf{P}(cl) \xrightarrow{(cl, \mathcal{T}, \mathcal{F})}_{(\text{RP}_{\text{LWW}}, \emptyset)} (\mathcal{X}', s'), \mathsf{C}'$$

for some s' such that $\mathcal{E}' = \mathcal{E}[cl \mapsto s']$ and C' such that $\mathsf{P}' = \mathsf{P}[cl \mapsto \mathsf{C}']$. It suffices to show that

$$cl \vdash (\mathcal{K}_{\mathcal{X}}, u, \mathcal{E}(cl)), \mathsf{P}(cl) \xrightarrow{(cl, \text{getView}(\mathcal{K}_{\mathcal{X}}, \mathcal{T}), \mathcal{F})}_{\text{ET}_{\top}} (\mathcal{K}_{\mathcal{X}'}, \mathcal{U}(cl), s'), \mathsf{C}'$$

Then by applying PSINGLETHREAD we obtain

$$(\mathcal{K}_{\mathcal{X}}, \mathcal{U}[cl \mapsto u], \mathcal{E}), \mathsf{P} \xrightarrow{(cl, \text{getView}(\mathcal{K}_{\mathcal{X}}, \mathcal{T}), \mathcal{F})}_{\text{ET}_{\top}} (\mathcal{K}_{\mathcal{X}'}, \mathcal{U}, \mathcal{E}'), \mathsf{P}'$$

The rest of the proof is performed by a rule induction on the derivation to inter

$$cl \vdash (\mathcal{X}, \mathcal{E}(cl)), \mathsf{P}(cl) \xrightarrow{(cl, \mathcal{T}, \mathcal{F})}_{(\text{RP}_{\text{LWW}}, \emptyset)} (\mathcal{X}', s'), \mathsf{C}'$$

Base case: ACOMMIT. In this case we have that

- $\mathsf{P} = [\mathsf{T}]$,
- $\mathsf{P}' = \mathtt{skip}$,
- $(\mathcal{E}(cl), ss, \emptyset), \mathsf{T} \to^* (s', \_, \mathcal{F}), \mathtt{skip}$ for an index $ss \in \text{RP}_{\text{LWW}}(\mathcal{X}, \mathcal{T})$, and
- $\mathcal{X}' = \text{extend}(\mathcal{X}, t, \mathcal{T}, \mathcal{F})$ for some $t \in \text{nextTid}(\mathcal{X}, cl)$.

Furthermore, it is easy to see by induction on the length of the derivation $(\mathcal{E}(cl), ss, \emptyset), \mathsf{T} \to^* (s', \_, \mathcal{F}), \mathtt{skip}$, that whenever $(\mathsf{r}, k, v) \in \mathcal{F}$ then $ss(k) = v$. Note that $\text{snapshot}(\mathcal{K}_{\mathcal{X}}, \text{getView}(\mathcal{X}, \mathcal{T})) = ss$ by Prop. D.5. Also, if $(\mathsf{r}, k, v) \in \mathcal{F}$ then $ss(k) = v$, which is possible only if $\mathcal{K}_{\mathcal{X}}(k, \max_{<}(\text{getView}(\mathcal{X}, \mathcal{T})(k))) = (v, \_, \_)$. This ensures that $\text{ET}_{\top} \vdash (\mathcal{K}_{\mathcal{X}}, \text{getView}(\mathcal{X}, \mathcal{T})) \triangleright \mathcal{F} : \mathcal{U}(cl)$. We can now combine all the facts above to apply rule PCOMMIT

$$cl \vdash (\mathcal{K}_{\mathcal{X}}, u, \mathcal{E}(cl)), [\mathsf{T}] \xrightarrow{(cl, \text{getView}(\mathcal{K}_{\mathcal{X}}, \mathcal{T}), \mathcal{F})}_{\text{ET}_{\top}} (\mathcal{K}', \mathcal{U}(cl), s'), \mathtt{skip},$$

where $\mathcal{K}' = \text{update}(\mathcal{K}_{\mathcal{X}}, t, \text{getView}(\mathcal{X}, \mathcal{T}), \mathcal{F})$. Recall that $\mathcal{X}' = \text{extend}(\mathcal{X}, \mathcal{T}, t, \mathcal{F})$. Therefore by Prop. D.10 we have that $\mathcal{K}' = \mathcal{K}_{\mathcal{X}'}$, which concludes the proof of this case.

Base case: APRIMITIVE,ACHOICE,AITER,ASEQSKIP. These cases are trivial since they do not alter the state of $\mathcal{X}$. Inductive case: ASEQ. It is derived by the I.H. □

COROLLARY E.7. *For any program* P,

$$\llbracket P \rrbracket_{ET_{\top}} = \left\{ \mathcal{K}_\mathcal{X} \mid \mathcal{X} \in \llbracket P \rrbracket_{(RP_{LWW}, \emptyset)} \right\}$$

PROOF. It can be derived by Prop. E.6 and Prop. E.5. □

E.3 Soundness and Completeness Constructor

We now show how all the results illustrated so far can be put together to show that the kv-store operational semantics is sound and complete with respect to abstract execution operational semantics.

E.3.1 Soundness. Recall that in the abstract execution operational semantics, a client cl loses information of the visible transactions immediately after it commits a transaction. Yet such information is indirectly presented when the next transaction from the same client is committed. To define the soundness judgement (Def. E.9), we introduce a notation of *invariant* (Def. E.8) to encore constraints on the visible transactions after each commit.

Definition E.8 (Invariant for clients). A *client-based invariant condition*, or simply *invariant*, is a function $I : \text{AbsExecs} \times \text{Client} \to \mathcal{P}(\text{TransID})$ such that for any cl we have that $I(\mathcal{X}, cl) \subseteq \mathcal{T}_\mathcal{X}$, and for any cl' such that $cl' \neq cl$ we have that $I(\text{extend}(\mathcal{X}, t_{cl'}^{\cdot}, \_, \_), cl) = I(\mathcal{X}, cl)$.

Definition E.9 (Soundness judgement). An execution test ET is sound with respect to an axiomatic definition $(RP_{LWW}, \mathcal{A})$ if and only if there exists an invariant condition I such that if assuming that

- a client cl having an initial view u, commits a transaction t with a fingerprint $\mathcal{F}$ and updates the view to u', which is allowed by ET i.e. $\text{ET} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$ where $\mathcal{K}' = \text{update}(\mathcal{K}, u, \mathcal{F}, t)$,
- a $\mathcal{X}$ such that $\mathcal{K}_\mathcal{X} = \mathcal{K}$ and $I(\mathcal{X}, cl) \subseteq \text{visTx}(\mathcal{K}, u)$,

then there exist a set of read-only transactions $\mathcal{T}_{rd}$ such that

- the new abstract execution $\mathcal{X}' = \text{extend}(\mathcal{X}, t, \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{rd}, \mathcal{F})$,
- the view u satisfies $\mathcal{A}$, i.e. $\forall A \in \mathcal{A}$. $\{t' \mid (t', t) \in A(\mathcal{X}')\} \subseteq \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{rd}$,
- the invariant is preserved, i.e. $I(\mathcal{X}', cl) \subseteq \text{visTx}(\mathcal{K}', u')$.

THEOREM E.10 (SOUNDNESS). *If* ET *is sound with respect to* $(RP_{LWW}, \mathcal{A})$, *then*

$$\text{CM(ET)} \subseteq \{\mathcal{K} \mid \exists \mathcal{X} \in \text{CM}(RP_{LWW}, \mathcal{A})). \ \mathcal{K}_\mathcal{X} = \mathcal{K}\}$$

PROOF. Let ET be an execution test that is sound with respect to an axiomatic definition $(RP_{LWW}, \mathcal{A})$. Let I be the invariant that satisfies Def. E.9. Let consider an ET-trace τ. We can assume that τ is in normal form, a trace that every view shift of a client cl is followed by a transaction from cl, and any transaction from cl must be after a view shift of cl. Without lose generality, we can also assume that the trace does not have transitions labelled as $(\_, \emptyset)$. Thus we have that the following trace τ:

$$\tau = (\mathcal{K}_0, \mathcal{U}_0) \xrightarrow{(cl_0, \varepsilon)}_{ET} (\mathcal{K}_0, \mathcal{U}_0') \xrightarrow{(cl_0, \mathcal{F}_0)}_{ET} (\mathcal{K}_1, \mathcal{U}_1) \xrightarrow{(cl_1, \varepsilon)}_{ET} \cdots \xrightarrow{(cl_{n-1}, \mathcal{F}_{n-1})}_{ET} (\mathcal{K}_n, \mathcal{U}_n)$$

For any $i : 0 \leq i \leq n$, let τ_i be the prefix of τ that contains only the first $2i$ transitions. Clearly τ_i is a valid ET-trace, and it is also a $ET_{\top}$-trace. By Prop. D.9, any abstract execution $\mathcal{X}_i \in \text{absExec}(\tau_i)$ satisfies the last write wins policy. We show by induction on i that we can always find an abstract execution $\mathcal{X}_i \in \text{absExec}(\tau_i)$ such that $\mathcal{X}_i \models \mathcal{A}$ and $I(\mathcal{X}_i, cl) \subseteq \mathcal{T}_{cl}^i$ for any client cl and set of transactions $\mathcal{T}_{cl}^i = \text{visTx}(\mathcal{X}_i, \mathcal{U}_i(cl)) \cup \mathcal{T}_{rd}^i$, and read-only transactions $\mathcal{T}_{rd}^i$ in $\mathcal{X}_i$. If so, because

$\mathcal{X}_i$ satisfies the last write wins policy, then it must be the case that $\mathcal{X}_i \models (\mathsf{RP_{LWW}}, \mathcal{A})$. Then by choosing $i = n$, we will obtain that $\mathcal{X}_n \models (\mathsf{RP_{LWW}}, \mathcal{A})$. Last, by Prop. D.9, $\mathcal{K}_{\mathcal{X}_n} = \mathcal{K}_n$, and there is nothing left to prove. Now let prove such $\mathcal{X}_i \in \mathsf{absExec}(\tau_i)$ always exists.

Base case: $i = 0$. Let $\mathcal{X}_0$ be the only abstract execution included in $\mathsf{absExec}(\tau_0)$, that is $\mathcal{X}_0 = ([], \emptyset, \emptyset)$. For any $\mathsf{A} \in \mathcal{A}$, it must be the case that $\mathsf{A}(\mathcal{X}_0) \subseteq \mathcal{T}_{\mathcal{X}_0} = \emptyset$, hence the inequation $\mathsf{A}(\mathcal{X}_0) \subseteq \mathsf{VIS}_{\mathcal{X}_0}$ is trivially satisfies. Furthermore, for the client invariant I we also require that $I(\mathcal{X}_0, \_) \subseteq \mathcal{T}_{\mathcal{X}_0} = \emptyset$; for any client cl we can choose $\mathcal{T}_{cl}^0 = \mathsf{visTx}(\mathcal{K}_{\mathcal{X}_0}, \mathcal{U}_0(cl)) \cup \emptyset = \emptyset$. Therefore $I(\mathcal{X}_0, cl) = \emptyset \subseteq \emptyset = \mathcal{T}_{cl}^0$.

Inductive case: $i' = i + 1$ where $i < n$. By the inductive hypothesis, there exists an abstract execution $\mathcal{X}_i$ such that

- $\mathcal{X}_i \models \mathsf{A}$ for all $\mathsf{A} \in \mathcal{A}$, and
- $I(\mathcal{X}, cl) \subseteq \mathcal{T}_{cl}^i$ for any client cl and set of transactions $\mathcal{T}_{cl}^i = \mathsf{visTx}(\mathcal{K}_i, \mathcal{U}_i(cl))$.

We have two transitions to check, the view shift and committing a transaction.

- the view shift transition $(\mathcal{K}_i, \mathcal{U}_i) \xrightarrow{(cl_i, \varepsilon)}_{\mathsf{ET}} (\mathcal{K}_i, \mathcal{U}_i')$. By definition, it must be the case that $\mathcal{U}_i' = \mathcal{U}_i[cl \mapsto u_i']$ for some u_i' such that $\mathcal{U}_i(cl) \sqsubseteq u_i'$. Let $(\mathcal{T}_{cl}^i)' = \mathsf{visTx}(\mathcal{K}_i, u_i')$; then we have $\mathcal{T}_{cl}^i = \mathsf{visTx}(\mathcal{K}_i, \mathcal{U}_i(cl)) \subseteq \mathsf{visTx}(\mathcal{K}_i, u_i') = (\mathcal{T}_{cl}^i)'$ As a consequence, $I(\mathcal{X}, cl) \subseteq \mathcal{T}_{cl}^i \subseteq (\mathcal{T}_{cl}^i)'$.

- the commit transaction transition $(\mathcal{K}_i, \mathcal{U}_i') \xrightarrow{(cl_i, \mathcal{F}_i)}_{\mathsf{ET}} (\mathcal{K}_{i+1}, \mathcal{U}_{i+1})$. A necessary condition for this transition to appear in τ is that $\mathsf{ET} \vdash (\mathcal{K}_i, \mathcal{U}(cl)) \rhd \mathcal{F}_i : (\mathcal{K}_{i+1}, \mathcal{U}_{i+1}(cl))$. Because I is the invariant to derive that ET is sound with respect to $\mathcal{A}$, and because $I(\mathcal{X}_i, cl_i) \subseteq (\mathcal{T}_{cl}^i)'$, then by Def. E.9 we have the following:
  - there exists a set of read-only transactions $\mathcal{T}_{\mathsf{rd}}$ such that

$$\left\{ t' \mid (t', t_{(cl, i)}) \in \mathsf{A}(\mathcal{X}_{i+1}) \right\} \subseteq \mathcal{T}_{cl}^{i'} \cup \mathcal{T}_{\mathsf{rd}}$$

    where $t_{(cl, i)} \in \mathsf{nextTid}(\mathcal{K}_i, cl)$ and $\mathcal{X}_{i+1} = \mathsf{extend}\left(\mathcal{X}_i, t_{(cl, i)}, (\mathcal{T}_{cl}^i)' \cup \mathcal{T}_{\mathsf{rd}}, \mathcal{F}_i\right)$,
  - $I(\mathcal{X}_{i+1}, cl) \subseteq \mathsf{visTx}(\mathcal{K}_{i+1}, \mathcal{U}_{i+1}(cl))$.

  Because $\mathcal{X}_i \in \mathsf{absExec}(\tau_i)$, by definition of $\mathsf{absExec}(\_)$ we have that $\mathcal{X}_{i+1} \in \mathsf{absExec}(\tau)$ (under the assumption that $\mathcal{F}_i \neq \emptyset$), and because $\mathsf{lastConf}(\tau_{i+1}) = (\mathcal{K}_{i+1}, \_)$, then $\mathcal{K}_{\mathcal{X}_{i+1}} = \mathcal{K}_{i+1}$.

  Now we need to check if $\mathcal{X}_{i+1}$ satisfies $\mathcal{A}$ and the invariant I is preserved.
  - $\mathsf{A}(\mathcal{X}_{i+1}) \subseteq \mathsf{VIS}_{\mathcal{X}}^{i+1}$ for any $\mathsf{A} \in \mathcal{A}$. Fix $\mathsf{A} \in \mathcal{A}$ and $(t', t) \in \mathsf{A}(\mathcal{X}_{i+1})$. Because $\mathcal{X}_{i+1} = \mathsf{extend}\left(\mathcal{X}_i, t_{(cl, i)}, (\mathcal{T}_{cl}^i)' \cup \mathcal{T}_{\mathsf{rd}}, \mathcal{F}_i\right)$, we distinguish between two cases.
    * If $t = t_{(cl, i)}$, then it must be the case that $t' \in (\mathcal{T}_{cl}^i)' \cup \mathcal{T}_{\mathsf{rd}}$, and by definition of extend we have that $(t', t_{(cl, i)}) \in \mathsf{VIS}_{\mathcal{X}_{i+1}}$.
    * If $t \neq t_{(cl, i)}$, then we have that $t, t' \in \mathcal{T}_{\mathcal{X}_i}$. Because $\mathcal{X}_i$ and $\mathcal{X}_{i+1}$ agree on $\mathcal{T}_{\mathcal{X}_i}$, then $(t', t) \in \mathsf{A}(\mathcal{X}_i)$. Because $\mathcal{X}_i \models \mathsf{A}$, then $(t', t) \in \mathsf{VIS}_{\mathcal{X}_i}$. By definition of extend, it follows that $(t', t) \in \mathsf{VIS}_{\mathcal{X}_{i+1}}$.
  - Finally, we show the invariant is preserved. Fix a client cl'.
    * If $cl' = cl$, then we have already proved that $I(\mathcal{X}_{i+1}, cl) \subseteq \mathcal{T}_{cl}^{i+1}$.
    * if $cl' \neq cl$, then note that $\mathcal{U}_i(cl') = \mathcal{U}_i'(cl') = \mathcal{U}_{i+1}(cl')$, and in particular $(\mathcal{T}_i^{cl'})' = \mathsf{visTx}(\mathcal{X}_i, \mathcal{U}_i'(cl')) = \mathsf{visTx}(\mathcal{X}_{i+1}, \mathcal{U}_{i+1}(cl')) = \mathcal{T}_{cl'}^{i+1}$. By the inductive hypothesis we know that $I(\mathcal{X}_i, cl) \subseteq \mathcal{T}_{cl'}^i$, and by the definition of invariant, we have $I(\mathcal{X}_{i+1}, cl) \subseteq \mathcal{T}_{cl'}^i = \mathcal{T}_{cl'}^{i+1}$.

$\square$

COROLLARY E.11. *If* ET *is sound with respect to* $(\mathsf{RP_{LWW}}, \mathcal{A})$, *then for any program* P, $[\![\mathsf{P}]\!]_{\mathsf{ET}} \subseteq \left\{ \mathcal{K}_{\mathcal{X}} \mid \mathcal{X} \in [\![\mathsf{P}]\!]_{(\mathsf{RP_{LWW}}, \mathcal{A})} \right\}$.

Proof.

$$\begin{aligned}
[\![P]\!]_{\mathsf{ET}} &= [\![P]\!]_{\mathsf{ET}_\top} \cap \mathsf{CM}(\mathsf{ET}) \\
&= \{\mathcal{K}_X \mid X \text{ satisfies } \mathsf{RP}_{\mathsf{LWW}}\} \cap \mathsf{CM}(\mathsf{ET}) \\
&\stackrel{Theorem\ E.10}{\subseteq} \{\mathcal{K}_X \mid X \text{ satisfies } \mathsf{RP}_{\mathsf{LWW}} \wedge X \in \mathsf{CM}(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A})\} \\
&= \{\mathcal{K}_X \mid X \in [\![P]\!]_{(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A})}\}
\end{aligned}$$

$\square$

E.3.2 Completeness. The Completeness judgement is in Def. E.12. Given a transaction t_i from client cl, it converts the visible transactions $\mathsf{VIS}_X^{-1}(t_i)$ into view and such view should satisfy the ET. Note that X does not contain precise information about final view after update, yet the visible transactions of the immediate next transaction from the same client cl include those information.

Definition E.12. An execution test ET is *complete* with respect to an axiomatic definition $(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A})$ if, for any abstract execution $X \in \mathsf{CM}(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A})$ and index $i : 1 \le i < |\mathcal{T}_X|$ such that $t_i \xrightarrow{\mathsf{AR}_X} t_{i+1}$, there exist an initial view u_i and a final view u_i' where

- $u_i = \mathsf{getView}(X, \mathsf{VIS}_X^{-1}(t_i))$,
- let $t_i = t_{cl}^n$ for some cl, n;
  - if the transaction $t_i' = \min_{\mathsf{SO}_X} \left\{ t' \,\middle|\, t_i \xrightarrow{\mathsf{SO}_X} t' \right\}$ is defined, then $u' = \mathsf{getView}(X, \mathcal{T}_i)$ where $\mathcal{T}_i \subseteq (\mathsf{AR}_X^{-1})^?(t_i) \cap \mathsf{VIS}_X^{-1}(t_i'))$;
  - otherwise $u' = \mathsf{getView}(X, \mathcal{T}_i)$ where $\mathcal{T}_i \subseteq (\mathsf{AR}_X^{-1})^?(t_i)$,
- $\mathsf{ET} \vdash (\mathcal{K}_{\mathsf{cut}(X, i-1)}, u_i) \triangleright \mathcal{T}_X(t_i) : (\mathcal{K}_{\mathsf{cut}(X, i)}, u_i')$.

THEOREM E.13. *Let* ET *be an execution test that is complete with respect to an axiomatic definition* $(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A})$. *Then* $\mathsf{CM}(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A}) \subseteq \mathsf{CM}(\mathsf{ET})$.

PROOF. Fix an abstract execution $X \in \mathsf{CM}(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A})$. For any $i : 1 \le i < |\mathcal{T}_X|$, suppose that t_i that is the i-*th* transaction follows the arbitrary order, i.e. $t_i \xrightarrow{\mathsf{AR}_X} t_{i+1}$ and let cl_i be the client of the i-*th* step, i.e. $t_i = t_{cl_i}^-$. Because ET is complete with respect to $(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A})$, for any step i we can find an initial views u_i, and a final view u_i' such that

- $u_i = \mathsf{getView}(X, \mathsf{VIS}_X^{-1}(t_i))$,
- there exists a set of transactions $\mathcal{T}_i$ such that $\mathsf{getView}(X, \mathcal{T}_i) = u_i'$, and either $\min_{\mathsf{SO}_X} \left\{ t' \,\middle|\, t_i \xrightarrow{\mathsf{SO}_X} t' \right\}$ is is defined and $\mathcal{T}_i \subseteq (\mathsf{AR}_X^{-1})^?(t_i) \cap \mathsf{VIS}_X^{-1}(t')$, or $\mathcal{T}_i \subseteq (\mathsf{AR}_X^{-1})^?(t_i)$,
- $\mathsf{ET} \vdash (\mathcal{K}_{\mathsf{cut}(X, i-1)}, u_i) \triangleright \mathcal{T}_X(t_i) : (\mathcal{K}_{\mathsf{cut}(X, i)}, u_i')$.

Given above, let $\mathcal{K}_i = \mathsf{cut}(X, i)$ and $\mathcal{F}_i = \mathcal{T}_X(t_i)$. Define the views for clients as

$$\mathcal{U}_0 = \lambda cl \in \{cl' \mid \exists t \in \mathcal{T}_X.\ t = t_{cl'}\}.\ \lambda k.\ \{0\} \quad \mathcal{U}_{i-1}' = \mathcal{U}_i[cl_i \mapsto u_i] \quad \mathcal{U}_i = \mathcal{U}_{i-1}'\big[cl_i \mapsto u_i'\big]$$

and the ke-stores as

$$\mathcal{K}_0 = \lambda k.(v_0, t_0, \emptyset) \quad \mathcal{K}_i = \mathsf{update}(\mathcal{K}_{i-1}, u_i, \mathcal{F}_i, t_i)$$

Now by Prop. D.17 we have that the following sequence of $\mathsf{ET}_\top$-reductions

$$(\mathcal{K}_0, \mathcal{U}_0) \xrightarrow{(cl_1, \varepsilon)}_{\mathsf{ET}_\top} (\mathcal{K}_0, \mathcal{U}_0') \xrightarrow{(cl_1, \mathcal{F}_1)}_{\mathsf{ET}_\top} (\mathcal{K}_1, \mathcal{U}_1) \xrightarrow{(cl_2, \varepsilon)}_{\mathsf{ET}_\top} \cdots \xrightarrow{(cl_n, \mathcal{F}_n)}_{\mathsf{ET}_\top} (\mathcal{K}_n, \mathcal{U}_n)$$

Note that $\mathcal{K}_i = \mathcal{K}_{\text{cut}(\mathcal{X}, i)}$. Because $\text{ET} \vdash (\mathcal{K}_{\text{cut}(\mathcal{X}, i-1)}, u_i) \triangleright \mathcal{F}_i : (\mathcal{K}_i, u_i')$, or equivalently $\text{ET} \vdash (\mathcal{K}_{\text{cut}(\mathcal{X}, i-1)}, \mathcal{U}_{i-1}'(cl_i)) \triangleright \mathcal{F}_i : (\mathcal{K}_{\text{cut}(\mathcal{X}, i-1)}, \mathcal{U}_i(cl_i))$, therefore

$$(\mathcal{K}_0, \mathcal{U}_0) \xrightarrow{(cl_1, \varepsilon)}_{\text{ET}} (\mathcal{K}_0, \mathcal{U}_0') \xrightarrow{(cl_1, \mathcal{F}_1)}_{\text{ET}} (\mathcal{K}_1, \mathcal{U}_1) \xrightarrow{(cl_2, \varepsilon)}_{\text{ET}} \cdots \xrightarrow{(cl_n, \mathcal{F}_n)}_{\text{ET}} (\mathcal{K}_n, \mathcal{U}_n)$$

It follows that $\mathcal{K}_n \in \text{CM(ET)}$ then $\mathcal{K}_n = \mathcal{K}_{\text{cut}(\mathcal{X}, n)} = \mathcal{K}_{\mathcal{X}}$, and there is nothing left to prove. □

COROLLARY E.14. *If* ET *is complete with respect to* $(\text{RP}_{\text{LWW}}, \mathcal{A})$, *then for any program* P,

$$\left\{ \mathcal{K}_{\mathcal{X}} \,\middle|\, \mathcal{X} \in [\![\text{P}]\!]_{(\text{RP}_{\text{LWW}}, \mathcal{A})} \right\} \subseteq [\![\text{P}]\!]_{\text{ET}}$$

PROOF.

$$\left\{ \mathcal{K}_{\mathcal{X}} \,\middle|\, \mathcal{X} \in [\![\text{P}]\!]_{(\text{RP}_{\text{LWW}}, \mathcal{A})} \right\} = \{\mathcal{K}_{\mathcal{X}} \mid \mathcal{X} \text{ satisfies } \text{RP}_{\text{LWW}} \wedge \mathcal{X} \in \text{CM}(\text{RP}_{\text{LWW}}, \mathcal{A})\}$$
$$\overset{Theorem\ E.13}{\subseteq} \{\mathcal{K}_{\mathcal{X}} \mid \mathcal{X} \text{ satisfies } \text{RP}_{\text{LWW}}\} \cap \text{CM(ET)}$$
$$= [\![\text{P}]\!]_{\text{ET}_\top} \cap \text{CM(ET)}$$
$$= [\![\text{P}]\!]_{\text{ET}}$$

□

F THE SOUNDNESS AND COMPLETENESS OF EXECUTION TESTS

We use Defs. E.9 and E.12 to prove the soundness and completeness of execution tests with respect to axiomatic definitions. It is sufficient to match these two definition, then by Cors. E.11 and E.14 we have $\mathsf{CM}(\mathsf{ET}) = \{\mathcal{K}_X \mid X \in \mathsf{CM}(\mathsf{RP}_{\mathsf{LWW}}, \mathcal{A})\}$.

We first prove the Theorem F.1, which states that the least fix point of view matches the constraint on the visibility relation on abstract execution.

THEOREM F.1 (VIEW CLOSURE TO VISIBILITY CLOSURE). *Assume $\mathcal{K}$ and X such that $\mathcal{K} = \mathcal{K}_X$, and $R_{\mathcal{K}}$ and R_X such that $R_{\mathcal{K}} = R_X$. For any $t, \mathcal{F}$, if there is a view $u = \mathrm{getView}\left(\mathcal{K}, \left(R_{\mathcal{K}}^{-1}\right)^* (\mathrm{visTx}(\mathcal{K}, u))\right)$, then the new abstract execution $X' = \mathrm{extend}\left(X, t, \mathcal{F}, \left(R_{\mathcal{K}}^{-1}\right)^* (\mathrm{visTx}(\mathcal{K}, u))\right)$ satisfies $R_X^{-1}(\mathrm{VIS}_{X'}^{-1}(t)) \subseteq \mathrm{VIS}_{X'}^{-1}(t)$. Conversely, If there a new abstract execution $X' = \mathrm{extend}(X, t, \mathcal{F}, \mathcal{T})$ for some $\mathcal{T}$ that satisfies $R_X^{-1}(\mathcal{T}) \subseteq \mathcal{T}$, and if a view $u = \mathrm{getView}(\mathcal{K}, \mathcal{T})$, then the view $u = \mathrm{getView}\left(\mathcal{K}, \left(R_{\mathcal{K}}^{-1}\right)^* (\mathrm{visTx}(\mathcal{K}, u))\right)$.*

PROOF. Assume $\mathcal{K}$ and X such that $\mathcal{K} = \mathcal{K}_X$, and $R_{\mathcal{K}}$ and R_X such that $R_{\mathcal{K}} = R_X$, Assume $t, \mathcal{F}$. Let $\mathcal{T} = \left(R_{\mathcal{K}}^{-1}\right)^* (\mathrm{visTx}(\mathcal{K}, u))$. Assume that a view satisfies $u = \mathrm{getView}(\mathcal{K}, \mathcal{T})$. By the definition of extend, the visible transactions $\mathrm{VIS}_{X'}^{-1}(t) = \mathcal{T}$. Let consider transactions t', t'' such that $t' \xrightarrow{R_X} t'' \xrightarrow{\mathrm{VIS}_{X'}} t$. This means there exists a natural number n such that $t'' \in \left(R_{\mathcal{K}}^{-1}\right)^n (\mathrm{visTx}(\mathcal{K}, u))$. Given that $R_{\mathcal{K}} = R_X$, it follows $t' \in \left(R_{\mathcal{K}}^{-1}\right)^{n+1} (\mathrm{visTx}(\mathcal{K}, u))$, then $t' \in \mathcal{T}$ and so $t' \xrightarrow{\mathrm{VIS}_{X'}} t$.

Assume there a new abstract execution $X' = \mathrm{extend}(X, t, \mathcal{F}, \mathcal{T})$, that satisfies $R_X^{-1}(\mathcal{T}) \subseteq \mathcal{T}$. Assume $u = \mathrm{getView}(\mathcal{K}, \mathcal{T})$. Note that $R_X = R_{\mathcal{K}}$. It suffices to prove $\{t' \in \mathcal{T} \mid t' \text{ has writes}\} = \left\{t' \in \left(R_{\mathcal{K}}^{-1}\right)^* (\mathrm{visTx}(\mathcal{K}, u)) \mid t' \text{ has writes}\right\}$.

- Assume a transaction $t' \in \mathcal{T}$ that has writes. It is easy to see there are k, i such that $i \in u(k)$ and $\mathrm{w}(\mathcal{K}(k, i)) \in \mathcal{T}$. This means $t' \in \mathrm{visTx}(\mathcal{K}, u)$.
- Assume a transaction $t' \in \mathrm{visTx}(\mathcal{K}, u)$, we now prove $\left(R_{\mathcal{K}}^{-1}\right)^n (t') \subseteq \mathcal{T}$ for all n.
  - Base case: n = 0. It trivially holds that $t' \in \mathrm{visTx}(\mathcal{K}, u) \subseteq \mathcal{T}$.
  - Inductive case: n + 1. Assume a transaction $t''' \in \left(R_{\mathcal{K}}^{-1}\right)^{n+1} (t')$. It means there is a $t'' \in \left(R_{\mathcal{K}}^{-1}\right)^n (t')$ such that $t''' \xrightarrow{R_{\mathcal{K}}} t''$. By I.H., $t'' \in \mathcal{T}$. Given $R_{\mathcal{K}} = R_X$ and $R_X^{-1}(\mathcal{T}) \subseteq \mathcal{T}$, it is known that $t''' \in \mathcal{T}$.

□

F.1 Monotonic Read MR

The execution test $\mathsf{ET}_{\mathsf{MR}}$ is sound with respect to the axiomatic definition [Sivaramakrishnan et al. 2015]

$$(\mathsf{RP}_{\mathsf{LWW}}, \{\lambda X.\mathrm{VIS}_X; \mathrm{SO}_X\})$$

We choose an invariant as the following,

$$I(X, cl) = \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \mid n \in \mathbb{N}\}} \mathrm{VIS}_X^{-1}(t_{cl}^n)\right) \setminus \mathcal{T}_{\mathsf{rd}}$$

where $\mathcal{T}_{\mathsf{rd}}$ is all the read-only transactions in $\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \mid n \in \mathbb{N}\}} \mathrm{VIS}_X^{-1}(t_{cl}^n)$. Assume a kv-store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $\mathsf{ET}_{\mathsf{MR}} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$. Also choose

an arbitrary cl, a transaction identifier $t \in \mathsf{nextTid}(\mathcal{K}, cl)$, and an abstract execution $\mathcal{X}$ such that $\mathcal{K}_\mathcal{X} = \mathcal{K}$ and

$$I(\mathcal{X}, cl) \subseteq \mathsf{visTx}(\mathcal{K}, u) \tag{6.1}$$

Let $\mathcal{X}' = \mathsf{extend}(\mathcal{X}, t, \mathcal{F}, \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_\mathrm{rd})$. We now check if $\mathcal{X}'$ satisfies the axiomatic definition and the invariant is preserved:

- $\{t' \mid (t', t) \in \mathsf{VIS}_{\mathcal{X}'}; \mathsf{SO}_{\mathcal{X}'}\} \subseteq \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_\mathrm{rd}$. Suppose that $t' \xrightarrow{\mathsf{VIS}_{\mathcal{X}'}} t'' \xrightarrow{\mathsf{SO}_{\mathcal{X}'}} t$ for some t', t''. We show that $t' \in I(\mathcal{X}, cl)$, and then Eq. (6.1) ensures that $t' \in \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_\mathrm{rd}$.
  Suppose $t'' \xrightarrow{\mathsf{SO}_{\mathcal{X}'}} t$, then $t'' = t_{cl}^n$ for some $n \in \mathbb{N}$. Because $t'' \neq t$ and $\mathcal{T}_{\mathcal{X}'} \setminus \mathcal{T}_\mathcal{X} = \{t\}$, we also have that $t'' \in \mathcal{X}$. By the invariant of $I(\mathcal{X}, cl)$, we have that $\mathsf{VIS}_\mathcal{X}^{-1}(cl) \subseteq I(\mathcal{X}, cl)$: because $t' \xrightarrow{\mathsf{VIS}_{\mathcal{X}'}} t''$ and $t'' \neq t$ we have that $t' \xrightarrow{\mathsf{VIS}_\mathcal{X}} t''$ and therefore $t' \in I(\mathcal{X}, cl)$.
- $I(\mathcal{X}', cl) \subseteq \mathsf{visTx}(\mathcal{X}', u') = \mathsf{visTx}(\mathcal{K}', u')$. In this case, because $\mathsf{ET}_\mathsf{MR} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$, then it must be the case that $u \sqsubseteq u'$. A trivial consequence of this fact is that $\mathsf{visTx}(\mathcal{K}, u) \subseteq \mathsf{visTx}(\mathcal{K}, u')$. Also, because $\mathcal{X}' = \mathsf{extend}(\mathcal{X}, t, \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_\mathrm{rd})$, we have that $\mathsf{visTx}(\mathcal{K}_\mathcal{X}, u) = \mathsf{visTx}(\mathcal{K}_{\mathcal{X}'}, u)$. Finally, note that $\{t_{cl}^n \in \mathcal{X}' \mid n \in \mathbb{N}\} = \{t_{cl}^n \in \mathcal{T}_\mathcal{X} \mid n \in \mathbb{N}\} \cup t$, that for any $t_{cl}^n \in \mathcal{T}_\mathcal{X}$ we have that $\mathsf{VIS}_{\mathcal{X}'}^{-1}(t_{cl}^n) = \mathsf{VIS}_\mathcal{X}^{-1}(t_{cl}^n)$, and that $\mathsf{VIS}_{\mathcal{X}'}^{-1}(t) = \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_\mathrm{rd}$. Using all these facts, we obtain

$$
\begin{aligned}
I(\mathcal{X}', cl) &= \left(\bigcup_{\{t_{cl}^n \in \mathcal{X}' \mid n \in \mathbb{N}\}} \mathsf{VIS}_{\mathcal{X}'}^{-1}(t_{cl}^n) \right) \setminus \mathcal{T}_\mathrm{rd} \\
&= \left(\left(\bigcup_{\{t_{cl}^n \in \mathcal{X} \mid n \in \mathbb{N}\}} \mathsf{VIS}_\mathcal{X}^{-1}(t_{cl}^n) \right) \setminus \mathcal{T}_\mathrm{rd} \right) \cup \left(\mathsf{VIS}_{\mathcal{X}'}^{-1}(t) \setminus \mathcal{T}_\mathrm{rd} \right) \\
&= I(\mathcal{X}, cl) \cup \mathsf{visTx}(\mathcal{K}, u) \\
&\overset{(6.1)}{\subseteq} \mathsf{visTx}(\mathcal{K}, u) \\
&= \mathsf{visTx}(\mathcal{K}_\mathcal{X}, u) \\
&= \mathsf{visTx}(\mathcal{K}_{\mathcal{X}'}, u) \\
&\subseteq \mathsf{visTx}(\mathcal{K}_{\mathcal{X}'}, u')
\end{aligned}
$$

We show that the execution test ET_MR is complete with respect to the axiomatic definition

$$(\mathsf{RP}_\mathsf{LWW}, \{\lambda \mathcal{X}.(\mathsf{VIS}_\mathcal{X}; \mathsf{SO}_\mathcal{X})\})$$

Let $\mathcal{X}$ be an abstract execution that satisfies the definition $\mathsf{CM}(\mathsf{RP}_\mathsf{LWW}, \{\lambda \mathcal{X}.(\mathsf{VIS}_\mathcal{X}; \mathsf{SO}_\mathcal{X})\})$, and consider a transaction $t \in \mathcal{T}_\mathcal{X}$. Assume i-*th* transaction t_i in the arbitrary order, and let $u_i = \mathsf{getView}(\mathcal{X}, \mathsf{VIS}_\mathcal{X}^{-1}(t_i))$. We have two possible cases:

- the transaction $t_i' = \min_{\mathsf{SO}_\mathcal{X}} \left\{ t' \mid t_i \xrightarrow{\mathsf{SO}_\mathcal{X}} t' \right\}$ is defined. In this case let

$$u_i' = \mathsf{getView}\left(\mathcal{X}, (\mathsf{AR}_\mathcal{X}^{-1})^?(t_i) \cap \mathsf{VIS}_\mathcal{X}^{-1}(t_i')\right)$$

Note that $t_i \xrightarrow{\mathsf{SO}_\mathcal{X}} t_i'$, and because $\mathcal{X} \models \mathsf{VIS}_\mathcal{X}; \mathsf{SO}_\mathcal{X}$, it follows that $\mathsf{VIS}_\mathcal{X}^{-1}(t_i) \subseteq \mathsf{VIS}_\mathcal{X}^{-1}(t_i')$. We also have that $\mathsf{VIS}_\mathcal{X}^{-1}(t_i) \subseteq (\mathsf{AR}_\mathcal{X}^{-1})^?(t_i)$ because of the definition of abstract execution. It follows that

$$\mathsf{VIS}_\mathcal{X}^{-1}(t_i) \subseteq (\mathsf{AR}_\mathcal{X}^{-1})^?(t_i) \cap \mathsf{VIS}_\mathcal{X}^{-1}(t_i'),$$

Recall that $u_i = \text{getView}(\mathcal{X}, \text{VIS}_{\mathcal{X}}^{-1}(t_i))$, and $u_i' = \text{getView}(\mathcal{X}, (\text{AR}_{\mathcal{X}}^{-1})^?(t_i) \cap \text{VIS}_{\mathcal{X}}^{-1}(t_i'))$. Thus we have that $u_i \sqsubseteq u_i'$, and therefore $\text{ET}_{\text{MR}} \vdash (\mathcal{K}_{\text{cut}(\mathcal{X},i)}, u_i) \rhd \mathscr{T}_{\mathcal{X}}(t_i) : (\mathcal{K}_{\text{cut}(\mathcal{X},i+1)}, u_i')$.

- the transaction $t_i' = \min_{\text{SO}_{\mathcal{X}}} \left\{ t' \mid t_i \xrightarrow{\text{SO}_{\mathcal{X}}} t_i \right\}$ is not defined. In this case, let

$$u_i' = \text{getView}\Big(\mathcal{X}, (\text{AR}_{\mathcal{X}}^{-1})^?(t_i)\Big)$$

As for the case above, we have that $u_i \sqsubseteq u_i'$, and therefore $\text{ET}_{\text{MR}} \vdash (\mathcal{K}_{\text{cut}(\mathcal{X},i)}, u_i) \rhd \mathscr{T}_{\mathcal{X}}(t_i) :$
$(\mathcal{K}_{\text{cut}(\mathcal{X},i+1)}, u_i', u_i')$.

F.2 Monotonic Write MW

The execution test ET_{MW} is sound with respect to the axiomatic definition [Sivaramakrishnan et al. 2015]

$$(\text{RP}_{\text{LWW}}, \{\lambda \mathcal{X}.(\text{SO}_{\mathcal{X}} \cap \text{WW}_{\mathcal{X}}); \text{VIS}_{\mathcal{X}}\})$$

We pick the invariant as empty set given the fact of no constraint on the view after update:

$$I(\mathcal{X}, cl) = \emptyset$$

Assume a kv-store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $\text{ET}_{\text{MW}} \vdash (\mathcal{K}, u) \rhd \mathcal{F} :$
$(\mathcal{K}', u')$. Also choose an arbitrary cl, a transaction identifier $t \in \text{nextTid}(\mathcal{K}, cl)$, and an abstract execution $\mathcal{X}$ such that $\mathcal{K}_{\mathcal{X}} = \mathcal{K}$ and $I(\mathcal{X}, cl) = \emptyset \subseteq \text{visTx}(\mathcal{K}, u)$. Note that since the invariant is empty set, it remains to prove that there exists a set of read-only transactions $\mathcal{T}_{\text{rd}}$ such that $\mathcal{X}' = \text{extend}(\mathcal{X}, t, \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}}, \mathcal{F})$ and:

$$\forall t'. \ (t', t) \in (\text{SO}_{\mathcal{X}'} \cap \text{WW}_{\mathcal{X}'}); \text{VIS}_{\mathcal{X}'} \Rightarrow t' \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}}$$

which can be derived from Theorem F.1.

The execution test ET_{MW} is complete with respect to the axiomatic definition

$$(\text{RP}_{\text{LWW}}, \{\lambda \mathcal{X}.((\text{SO}_{\mathcal{X}} \cap \text{WW}_{\mathcal{X}}); \text{VIS}_{\mathcal{X}})\})$$

Let $\mathcal{X}$ be an abstract execution that satisfies the definition $\text{CM}(\text{RP}_{\text{LWW}}, \{\lambda \mathcal{X}.((\text{SO}_{\mathcal{X}} \cap \text{WW}_{\mathcal{X}}); \text{VIS}_{\mathcal{X}})\})$, and consider a transaction $t \in \mathcal{T}_{\mathcal{X}}$. Let $\mathcal{K} = \mathcal{K}_{\mathcal{X}}$. Assume i-*th* transaction t_i in the arbitrary order, and let view $u_i = \text{getView}(\mathcal{K}, \text{VIS}_{\mathcal{X}}^{-1}(t_i))$. We also pick any final view such that $u_i' \subseteq \text{getView}(\mathcal{K}, (\text{AR}_{\mathcal{X}}^{-1})^?(t_i))$. It suffices to prove $\text{ET}_{\text{MW}} \vdash (\mathcal{K}_{\text{cut}(\mathcal{X},i-1)}, u_i) \rhd \mathscr{T}_{\mathcal{X}}(t_i) : (\mathcal{K}_{\text{cut}(\mathcal{X},i-1)}, u_i')$. It means to prove the following:

$$u_i = \text{getView}\Big(\mathcal{K}_{\text{cut}(\mathcal{X},i-1)}, \text{lfpTx}\Big(\mathcal{K}_{\text{cut}(\mathcal{X},i-1)}, u, \text{SO} \cap \text{WW}_{\mathcal{K}_{\text{cut}(\mathcal{X},i-1)}}\Big)\Big)$$

which can be derived from Theorem F.1.

F.3 Read Your Write RYW

The execution test ET_{RYW} is sound with respect to the axiomatic definition [Sivaramakrishnan et al. 2015] $(\text{RP}_{\text{LWW}}, \{\lambda \mathcal{X}. \text{SO}_{\mathcal{X}}\})$. We pick an invariant for the ET_{RYW} as the following:

$$I(\mathcal{X}, cl) = \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} (\text{SO}_{\mathcal{X}}^{-1})^?(t_{cl}^n) \right) \setminus \mathcal{T}_{\text{rd}}$$

where $\mathcal{T}_{\text{rd}}$ is all the read-only transactions in $\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} (\text{SO}_{\mathcal{X}}^{-1})^?(t_{cl}^n)$. Assume a kv-store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $\text{ET}_{\text{RYW}} \vdash (\mathcal{K}, u) \rhd \mathcal{F} : (\mathcal{K}', u')$. Also choose an arbitrary cl, a transaction identifier $t_{cl}^n \in \text{nextTid}(\mathcal{K}, cl)$, and an abstract execution $\mathcal{X}$ such that

$\mathcal{K}_X = \mathcal{K}$ and $I(X, cl) \subseteq \text{visTx}(\mathcal{K}, u)$. Let a new abstract execution $X' = \text{extend}\left(X, t_{cl}^n, \mathcal{F}, \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}}\right)$. We need to prove that X' satisfies the constraint and the invariant is preserved:

- $t \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}}$ for all t such that $t \xrightarrow{\text{SO}_{X'}} t_{cl}^n$. Assume a transaction t such that $t \xrightarrow{\text{SO}_{X'}} t_{cl}^n$. It immediately implies that $t = t_{cl}^m$ where $m < n$ and $t_{cl}^m \in X$. Thus we prove that

$$t \in \left(\bigcup_{\left\{ t_{cl}^n \in \mathcal{T}_X \mid n \in \mathbb{N} \right\}} (\text{SO}_X^{-1})^? (t_{cl}^n) \right) \subseteq \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}}$$

- $I(X', cl) \subseteq \text{visTx}(\mathcal{K}_{X'}, u')$. Let $\mathcal{T}_{\text{rd}}' = \mathcal{T}_{\text{rd}}$ if the new transaction t_{cl}^n has writes, otherwise $\mathcal{T}_{\text{rd}}' = \mathcal{T}_{\text{rd}} \cup \left\{ t_{cl}^n \right\}$. First we have

$$I(X', cl) = \left(\bigcup_{\left\{ t_{cl}^m \in \mathcal{T}_{X'} \mid m \in \mathbb{N} \right\}} (\text{SO}_{X'}^{-1})^? (t_{cl}^m) \right) \setminus \mathcal{T}_{\text{rd}}' = \left((\text{SO}_{X'}^{-1})^? (t_{cl}^n) \right) \setminus \mathcal{T}_{\text{rd}}'$$

Note that t_{cl}^n is the latest transaction committed by the client cl. For any transaction $t \in (\text{SO}_{X'}^{-1})^? (t_{cl}^n) \setminus \mathcal{T}_{\text{rd}}'$ that has write, because execution test requires $z \in u'(k)$ for any key k and index z such that $\text{w}(\mathcal{K}_{X'}(k, z)) \xrightarrow{\text{SO}_X} t$, then $t \in \text{visTx}(\mathcal{K}_{X'}, u')$ as what we wanted.

The execution test ET_{RYW} is complete with respect to the axiomatic definition $(\text{RP}_{\text{LWW}}, \{\lambda X.\text{SO}_X\})$. Let X be an abstract execution that satisfies the definition $\text{CM}(\text{RP}_{\text{LWW}}, \{\lambda X.\text{SO}_X\})$. Assume i-*th* transaction t_i in the arbitrary order, and let view $u_i = \text{getView}(X, \text{VIS}_X^{-1}(t_i))$. We construct the final view u_i' depending on whether t_i is the last transaction from the client.

- If the transaction $t_i' = \min_{\text{SO}_X} \left(\left\{ t' \mid t_i \xrightarrow{\text{SO}_X} t' \right\} \right)$ is defined, then $u_i' = \text{getView}(X, \mathcal{T}_i)$ where $\mathcal{T}_i \subseteq (\text{AR}_X^{-1})^? (t_i) \cap \text{VIS}_X^{-1}(t_i')$ for some $\mathcal{T}_i$. Given the definition $\lambda X.\text{SO}_X$, we know $\text{SO}_X^{-1}(t_i') \subseteq \text{VIS}_X^{-1}(t_i')$, so $(\text{AR}_X^{-1})^? (t_i) \cap \text{SO}^{-1}(t_i') = (\text{SO}^{-1})^? (t_i) \subseteq \mathcal{T}_i$. Take j, k such that $\text{w}(\mathcal{K}_{\text{cut}(X,i)}(k, j)) \xrightarrow{\text{SO}^?} t_i'$. By the constraint of X, that is $\text{SO}_X \subseteq \text{VIS}_X$, it follows $\text{w}(\mathcal{K}_{\text{cut}(X,i)}(k, j)) \in \mathcal{T}_i$. Recall $u_i' = \text{getView}(\mathcal{K}_{\text{cut}(X,i)}, \mathcal{T}_i)$. By the definition of getView, it follows $i \in u_i'(k)$. Therefore $\text{ET}_{\text{RYW}} \vdash (\mathcal{K}_{\text{cut}(X,i-1)}, u_i) \triangleright \mathcal{T}_X(t_i) : (\mathcal{K}_{\text{cut}(X,i)}, u_i')$.
- If there is no other transaction after t_i from the same client, we pick $u_i' = \text{getView}(X, \mathcal{T}_i)$ where $\mathcal{T}_i = (\text{SO}_X^{-1})^? (t_i)$, so $\text{ET}_{\text{RYW}} \vdash (\mathcal{K}_{\text{cut}(X,i-1)}, u_i) \triangleright \mathcal{T}_X(t_i) : (\mathcal{K}_{\text{cut}(X,i)}, u_i')$.

F.4 Write Following Read WFR

The write-read relation on X is defined as the following:

$$\text{WR}(X, k) \triangleq \left\{ (t, t') \mid \exists v. \ (\text{w}, k, v) \in_X t \land (\text{r}, k, v) \in_X t' \land t = \max_{\text{AR}} (\text{VIS}^{-1}(t')) \right\}$$

The notation WR_X is defined as $\text{WR}_X \triangleq \bigcup_{k \in \text{KEY}} \text{WR}(X, k)$. Note that for a kv-store $\mathcal{K}$ such that $\mathcal{K} = \mathcal{K}_X$, by the definition of $\mathcal{K} = \mathcal{K}_X$, the following holds:

$$\text{WR}_X = \{ (t, t') \mid \exists k, i. \ \mathcal{K}(k, i) = (\_, t, t' \cup \_) \}$$

Note that such WR_X coincides with $\text{WR}_{\mathcal{G}}$ and $\text{WR}_{\mathcal{K}}$.

The execution test ET_{WFR} is sound with respect to the axiomatic definition [Sivaramakrishnan et al. 2015]

$$(\text{RP}_{\text{LWW}}, \{\lambda X.\text{WR}_X; (\text{SO} \cap \text{RW}_X)^?; \text{VIS}_X\})$$

We pick the invariant as $I(X, cl) = \emptyset$, given the fact of no constraint on the view after update. Assume a kv-store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $\mathsf{ET}_{\mathsf{WFR}} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$. Also choose an arbitrary cl, a transaction identifier $t \in \mathsf{nextTid}(\mathcal{K}, cl)$, and an abstract execution X such that $\mathcal{K}_X = \mathcal{K}$ and $I(X, cl) = \emptyset \subseteq \mathsf{visTx}(\mathcal{K}, u)$. Note that since the invariant is empty set, it remains to prove there is a set of read-only transactions $\mathcal{T}_{\mathsf{rd}}$ such that Let $X' = \mathsf{extend}(X, t, \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathsf{rd}}, \mathcal{F})$ and

$$\forall t'. \; (t', t) \in \mathsf{WR}_{X'}; \mathsf{SO}^?_{X'}; \mathsf{VIS}_{X'} \Rightarrow t' \in \mathsf{visTx}(\mathcal{K}, u)$$

which can be derived from Theorem F.1.

The execution test $\mathsf{ET}_{\mathsf{WFR}}$ is complete with respect to the axiomatic definition

$$(\mathsf{RP}_{\mathsf{LWW}}, \{\lambda X.\mathsf{WR}; \mathsf{WR}_{\mathcal{K}}; (\mathsf{SO} \cap \mathsf{RW}_{X'})^?; \mathsf{VIS}_{X'}\})$$

Assume i-th transaction t_i in the arbitrary order, and let view $u_i = \mathsf{getView}(\mathcal{K}_{\mathsf{cut}(X, i-1)}, \mathsf{VIS}_X^{-1}(t_i))$. We also pick any final view such that $u'_i \subseteq \mathsf{getView}(\mathcal{K}_{\mathsf{cut}(X, i)}, (\mathsf{AR}_X^{-1})^?(t_i))$. Note that there is nothing to prove for u'_i, so it is sufficient to prove the following:

$$u_i = \mathsf{getView}\Big(\mathcal{K}_{\mathsf{cut}(X, i-1)}, \mathsf{lfpTx}\Big(\mathcal{K}_{\mathsf{cut}(X, i-1)}, u, \mathsf{WR}_{\mathcal{K}_{\mathsf{cut}(X, i-1)}}; \mathsf{SO}\Big)\Big)$$

which can be derived from Theorem F.1.

F.5 Causal Consistency CC

The wildly used definition on abstract executions for causal consistency is that VIS is transitive. Yet it is for the sack of elegant definition, while there is a minimum visibility relation given by $(\mathsf{WR}_X \cup \mathsf{SO}_X); \mathsf{VIS}_X \subseteq \mathsf{VIS}_X$ (Lemma F.2).

LEMMA F.2. *For any abstract execution X under last-write-win, if it satisfies the following:*

$$(\mathsf{WR}_X \cup \mathsf{SO}_X); \mathsf{VIS}_X \subseteq \mathsf{VIS}_X \quad \mathsf{SO}_X \subseteq \mathsf{VIS}_X$$

There exists a new abstract execution X' where $\mathcal{T}_X = \mathcal{T}_{X'}$, $\mathsf{AR}_X = \mathsf{AR}_{X'}$, $\mathsf{VIS}_{X'}; \mathsf{VIS}_{X'} \subseteq \mathsf{VIS}_{X'}$, and under last-write-win $\mathcal{T}_X(t) = \mathcal{T}_{X'}(t)$ for all transactions t.

PROOF. To recall, the write-read relation under a key $\mathsf{WR}(X, k)$ is defined as $\mathsf{WR}(X, k) \triangleq \{(t, t') \mid \exists v. \; (\mathsf{w}, k, v)$ Given an X that satisfies the following

$$(\mathsf{WR}_X \cup \mathsf{SO}_X); \mathsf{VIS}_X \subseteq \mathsf{VIS}_X \quad \mathsf{SO}_X \subseteq \mathsf{VIS}_X$$

we erase some visibility relation for each transaction following the order of arbitration AR until the visibility is transitive. Assume the i-th transaction t_i with respect to the arbitration order. Let R_i denote a new visibility for transaction t_i such that $R_i|_2 = \{t_i\}$ and the visibility relation before (including) t_i is transitive. Let $X_i = \mathcal{K}_{\mathsf{cut}(X, i)}$ and $\mathsf{VIS}_i = \bigcup_{0 \le k \le i} R_i$. For each step, says i-$th$ step, we preserve the following:

$$\mathsf{VIS}_i; \mathsf{VIS}_i \subseteq \mathsf{VIS}_i \tag{6.2}$$

$$\forall t. \; (t, t_i) \in R_i \Rightarrow (t, t_i) \in (\mathsf{WR}_i \cup \mathsf{SO}_i) \tag{6.3}$$

- Base case: $i = 1$ and $R_1 = \emptyset$. Assume it is from client cl. There is no transaction committed before, so $\mathsf{VIS}_1 = \emptyset$ and $\mathsf{VIS}_1; \mathsf{VIS}_1 \subseteq \mathsf{VIS}_1$ as Eq. (6.2).
- Inductive case: i-th step. Suppose the (i-1)-th step satisfies Eq. (6.2) and Eq. (6.3). Let consider i-th step and the transaction t_i. Initially we take R_i as empty set. We first extend R_i by closing with respect to WR_i and prove that it does not affect any read from the transaction t_i. Then we will do the same for SO_i.

- WR_i. For any read $(r, k, v) \in t_i$, there must be a transaction t_j that $t_j \xrightarrow{WR(X_i, k), AR} t_i$ and $j < i$. We include $(t_j, t_i) \in R_i$. Let consider all the visible transactions of t_j. Assume a transaction $t' \in VIS_{i-1}^{-1}(t_j)$, thus $t' \in VIS_j^{-1}(t_j) = R_j^{-1}(t_j)$. It is safe to include $(t', t_i) \in R_i$ without affecting the read result, because those transaction t' is already visible for t_i in the abstract execution X: by Eq. (6.3) we know $R_j \subseteq (WR_j \cup SO_j)^+ \subseteq (WR_X \cup SO_X)^+$, and by the definition of $WR(X_i, k)$ we know $WR(X_i, k) \subseteq VIS_X$.

- Given $SO_X \subseteq VIS_X$, we include (t_j, t_i) for some t_j such that $t_j \xrightarrow{SO_X} t_i$. For the similar reason as WR, it is safe to includes all the visible transactions t' for t_j, i.e. $t' \in R_j^{-1}$.

By the construction, both Eq. (6.2) and Eq. (6.3) are preserved. Thus we have the proof.

$\square$

PROPOSITION F.3. *For any abstract execution X under last-write-win, if it satisfies the following:*

$$(WR_X \cup SO_X)^+; VIS_X \subseteq VIS_X \quad SO_X \subseteq VIS_X$$

then

$$\exists R \subseteq AR_X. \ VIS = (WR_X \cup SO_X \cup R)^+$$

By Lemma F.2, the execution test ET_{CC} is sound with respect to the axiomatic definition

$$(RP_{LWW}, \{\lambda X.(SO_X \cup WR_X)^+; VIS_X, \lambda X. SO_X\})$$

We pick an invariant for the ET_{CC} as the union of those for MR and RYW shown in the following:

$$I_1(X, cl) = \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} VIS_X^{-1}(t_{cl}^n) \right) \setminus \mathcal{T}_{rd}$$

$$I_2(X, cl) = \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} (SO_X^{-1})^?(t_{cl}^n) \right) \setminus \mathcal{T}_{rd}$$

where $\mathcal{T}_{rd}$ is all the read-only transactions included in both:

$$\mathcal{T}_{rd} \in \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} VIS_X^{-1}(t_{cl}^n) \right)$$

and

$$\mathcal{T}_{rd} \in \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} (SO_X^{-1})^?(t_{cl}^n) \right)$$

Assume a kv-store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $ET_{CC} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$. Also choose an arbitrary cl, a transaction identifier $t_{cl}^n \in \text{nextTid}(\mathcal{K}, cl)$, and an abstract execution X such that $\mathcal{K}_X = \mathcal{K}$ and $I_1(X, cl) \cup I_2(X, cl) \subseteq \text{visTx}(\mathcal{K}, u)$. We are about to prove there exists an extra set of read-only transactions $\mathcal{T}'_{rd}$ such that the new abstract execution $X' = \text{extend}\left(X, t_{cl}^n, \mathcal{F}, \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{rd} \cup \mathcal{T}'_{rd}\right)$ and:

$$\forall t. \ (t, t_{cl}^n) \in SO_{X'} \Rightarrow t \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{rd} \cup \mathcal{T}'_{rd} \tag{6.4}$$

$$\forall t. \ (t, t_{cl}^n) \in (SO_{X'} \cup WR_{X'}); VIS_{X'} \Rightarrow t \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{rd} \cup \mathcal{T}'_{rd} \tag{6.5}$$

$$I_1(X', cl) \cup I_2(X', cl) \subseteq \text{visTx}(\mathcal{K}_{X'}, u') \tag{6.6}$$

- The invariant I_2 implies Eq. (6.4) as the same as RYW in §F.3.
- Eq. (6.5). Note that $(t, t_{cl}^n) \in (\mathsf{SO}_{\mathcal{X}'} \cup \mathsf{WR}_{\mathcal{X}'}); \mathsf{VIS}_{\mathcal{X}'} \Rightarrow (t, t_{cl}^n) \in (\mathsf{SO}_{\mathcal{X}} \cup \mathsf{WR}_{\mathcal{X}}); \mathsf{VIS}_{\mathcal{X}'}$. Also, recall that $\mathsf{SO}_{\mathcal{X}} = \mathsf{SO}_{\mathcal{K}}$ and $\mathsf{WR}_{\mathcal{X}} = \mathsf{WR}_{\mathcal{K}}$. Let $\mathcal{T}'_{\mathrm{rd}} = \mathsf{lfpTx}(\mathcal{K}, u, \mathsf{SO}_{\mathcal{K}} \cup \mathsf{WR}_{\mathcal{K}})$. This means that $\mathcal{X}' = \mathsf{extend}\big(\mathcal{X}, t_{cl}^n, \mathcal{F}, \mathsf{lfpTx}(\mathcal{K}, u, \mathsf{SO}_{\mathcal{K}} \cup \mathsf{WR}_{\mathcal{K}}) \cup \mathcal{T}_{\mathrm{rd}}\big)$. Let assume $t \xrightarrow{\mathsf{SO}_{\mathcal{K}} \cup \mathsf{WR}_{\mathcal{K}}} t'$ and $t' \in \mathsf{lfpTx}(\mathcal{K}, u, \mathsf{SO}_{\mathcal{K}} \cup \mathsf{WR}_{\mathcal{K}}) \cup \mathcal{T}_{\mathrm{rd}}$. We have two possible cases:
  - If $t' \in \mathsf{lfpTx}(\mathcal{K}, u, \mathsf{SO}_{\mathcal{K}} \cup \mathsf{WR}_{\mathcal{K}})$, by Theorem F.1, we know $t \in \mathsf{lfpTx}(\mathcal{K}, u, \mathsf{SO}_{\mathcal{K}} \cup \mathsf{WR}_{\mathcal{K}})$.
  - If $t' \in \mathcal{T}_{\mathrm{rd}}$, there are two cases:
    * $t' \in \Big(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} \mathsf{VIS}_{\mathcal{X}}^{-1}(t_{cl}^n)\Big)$. By the property of $\mathcal{X}$ (before update) that $(\mathsf{SO} \cup \mathsf{WR}_{\mathcal{X}}); \mathsf{VIS}_{\mathcal{X}} \in \mathsf{VIS}_{\mathcal{X}}$, it is known that $t \in \Big(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} \mathsf{VIS}_{\mathcal{X}}^{-1}(t_{cl}^n)\Big)$, that is, $t \in \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}}$.
    * $t' \in \Big(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} \mathsf{SO}_{\mathcal{X}}^{-1}(t_{cl}^n)\Big)$. Then we know $t \in (\mathsf{SO} \cup \mathsf{WR}_{\mathcal{X}})^{-1} \Big(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} \mathsf{SO}_{\mathcal{X}}^{-1}(t_{cl}^n)\Big)$. By the property of $\mathcal{X}$ (before update) that $\mathsf{SO} \cup \mathsf{WR}_{\mathcal{X}} \in \mathsf{VIS}_{\mathcal{X}}$, it follows:

$$t \in VIS_{\mathcal{X}}^{-1}\left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} \mathsf{SO}_{\mathcal{X}}^{-1}(t_{cl}^n)\right)$$
$$= \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} \mathsf{VIS}_{\mathcal{X}}^{-1}(t_{cl}^n)\right)$$
$$= \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}}$$

- Finally the new abstract execution preserves the invariant I_1 and I_2 because CC satisfies MW and RYW. The proofs are the same as those in §F.1 and §F.3.

The execution test $\mathsf{ET}_{\mathsf{CC}}$ is complete with respect to the axiomatic definition

$$(\mathsf{RP}_{\mathsf{LWW}}, \{\lambda\mathcal{X}.\mathsf{VIS}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}}, \lambda\mathcal{X}.\,\mathsf{SO}_{\mathcal{X}}\})$$

Assume i-th transaction t_i in the arbitrary order, and let view $u_i = \mathsf{getView}\big(\mathcal{X}, \mathsf{VIS}_{\mathcal{X}}^{-1}(t_i)\big)$. We pick final view as $u'_i = \mathsf{getView}\big(\mathcal{X}, (\mathsf{AR}_{\mathcal{X}}^{-1})^?(t_i) \cap \mathsf{VIS}_{\mathcal{X}}^{-1}(t'_i)\big)$, if $t'_i = \min_{\mathsf{SO}}\big\{t' \mid t_i \xrightarrow{\mathsf{SO}} t'\big\}$ is defined, otherwise $u'_i = \mathsf{getView}\big(\mathcal{X}, (\mathsf{AR}_{\mathcal{X}}^{-1})^?(t_i)\big)$. Let the $\mathcal{K} = \mathcal{K}_{\mathsf{cut}(\mathcal{X}, i-1)}$. Now we prove the three parts separately.

- MR. By Prop. F.3 since $\mathsf{VIS}_{\mathcal{X}}; \mathsf{SO}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}$ so it follows as in §F.1.
- RYW. For RYW, since $\mathsf{WR}_{\mathcal{X}}; \mathsf{SO}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}$, the proof is as the same proof as in §F.3.
- $\mathtt{allowed}(\mathsf{WR}_{\mathcal{K}} \cup \mathsf{SO})$. It is derived from Theorem F.1 and $(\mathsf{WR}_{\mathcal{X}} \cup \mathsf{SO}); \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}$.

F.6 Update Atomic UA

Given abstract execution $\mathcal{X}$, we define write-write relation for a key k as the following [Cerone et al. 2015a]:

$$\mathsf{WW}(\mathcal{X}, k) \triangleq \left\{(t, t') \;\middle|\; t \xrightarrow{\mathsf{AR}_{\mathcal{X}}} t' \land (\mathsf{w}, k, \_) \in t \land (\mathsf{w}, k, \_) \in t'\right\}$$

Then, the notation $\mathsf{WW}_{\mathcal{X}} \triangleq \bigcup_{k \in \mathsf{KEY}} \mathsf{WW}(\mathcal{X}, k)$. Note that for a kv-store $\mathcal{K}$ such that $\mathcal{K} = \mathcal{K}_{\mathcal{X}}$, by the definition of $\mathcal{K} = \mathcal{K}_{\mathcal{X}}$, the following holds:

$$\mathsf{WW}_{\mathcal{X}} = \{(t, t') \mid \exists k, i, j.\ t = \mathsf{w}(\mathcal{K}(k, i)) \land t' = \mathsf{w}(\mathcal{K}(k, j)) \land i < j\}$$

Also the WW_X coincides with WW_G and WW_K.

The execution test ET_{UA} is sound with respect to the axiomatic definition $(RP_{LWW}, \{\lambda X.WW_X\})$. We pick the invariant as $I(X, cl) = \emptyset$, given the fact of no constraint on the final view. Assume a kv-store K, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $ET_{UA} \vdash (K, u) \triangleright \mathcal{F} : (K', u')$. Also choose an arbitrary cl, a transaction identifier $t \in \text{nextTid}(K, cl)$, and an abstract execution X such that $K_X = K$ and $I(X, cl) = \emptyset \subseteq \text{visTx}(K, u)$. Let $X' = \text{extend}(X, t, \text{visTx}(K, u), \mathcal{F})$. Note that since the invariant is empty set, it remains to prove the following:

$$\forall t'. \ t' \xrightarrow{WW_{X'}} t \Rightarrow t' \in \text{visTx}(K, u)$$

Assume a transaction t' that writes to a key k as t, i.e. $t' \xrightarrow{WW_{X'}} t$. Since that t' is a transaction already existing in K, we have $w(K(k, i)) = t'$ for some index i and key k. It means $(w, k, \text{val}(K(k, i))) \in \mathcal{F}$. By the execution test of UA, we know $i \in u(k)$ therefore $t' \in \text{visTx}(K, u)$.

The execution test ET_{UA} is complete with respect to the axiomatic definition $(RP_{LWW}, \{\lambda X.WW_X\})$. Assume i-*th* transaction t_i in the arbitrary order, and let view $u_i = \text{getView}(X, \text{VIS}_X^{-1}(t_i))$. We also pick any final view such that $u_i' \subseteq \text{getView}(X, (AR_X^{-1})^?(t_i))$. Note that there is nothing to prove for u_i', so it is sufficient to prove the following:

$$\forall k. \ (w, k, \_) \in \mathcal{T}_X(t_i) \Rightarrow \forall j : 0 \le j < |K_{\text{cut}(X, i-1)}(k)|. \ j \in u_i(k)$$

Let consider a key k that have been overwritten by the transaction t_i. By the constraint of X that $WW_X \subseteq \text{VIS}_X$, for any transaction t that writes to the same key k and committed before t_i, they are included in the visible set $t \in \text{VIS}_X^{-1}(t_i)$. Note that $t \xrightarrow{WW_X} t_i \Rightarrow t \xrightarrow{AR_X} t_i \Rightarrow t \in K_{\text{cut}(X, i-1)}$. Since that the transaction t write to the key k, it means $w(K_{\text{cut}(X, i-1)}(k, j)) = t$ for some index j. Then by the definition of getView, we have $j \in u_i(k)$.

F.7 Consistency Prefix CP

Given abstract execution X, we define read-write read-write relation:

$$RW(X, k) \triangleq \left\{ (t, t') \ \middle| \ t \xrightarrow{AR_X} t' \wedge (r, k, \_) \in t \wedge (w, k, \_) \in t' \right\}$$

It is easy to see $RW(X, k)$ can be derived from $WW(X, k)$ and $WR(X, k)$ as the following:

$$RW(X, k) = \{(t, t') \mid \exists t''. \ (t'', t) \in WR(X, k) \wedge (t'', t') \in WW(X, k)\}$$

Then, the notation $RW_X \triangleq \bigcup_{k \in \text{KEY}} RW(X, k)$. Note that for a kv-store K such that $K = K_X$, by the definition of $K = K_X$, the following holds:

$$RW_X = \{(t, t') \mid \exists k, i, j. \ t \in \text{rs}(K(k, i)) \wedge t' = w(K(k, j)) \wedge i < j\}$$

The RW_X also coincides with RW_G and RW_K.

An abstract execution X satisfies consistency prefix (CP), if it satisfies $AR_X; \text{VIS}_X \subseteq \text{VIS}_X$ and $SO_X \subseteq \text{VIS}_X$. Given the definition, there is a corresponding definition on dependency graph by solve the following inequalities:

$$WR \subseteq VIS$$
$$WW \subseteq AR$$
$$VIS \subseteq AR$$
$$VIS; RW \subseteq AR$$
$$AR; AR \subseteq AR$$
$$SO \subseteq VIS$$
$$AR; VIS \subseteq VIS$$

By solving the inequalities the visibility and arbitration relations are:

$$\text{AR} \triangleq \left((\text{SO} \cup \text{WR}); \text{RW}^? \cup \text{WW} \cup R\right)^+$$

$$\text{VIS} \triangleq \left((\text{SO} \cup \text{WR}); \text{RW}^? \cup \text{WW} \cup R\right)^* ; (\text{SO} \cup \text{WR})$$

for some relation $R \subseteq \text{AR}$. When $R = \emptyset$, it is the smallest solution therefore the minimum visibility required.

LEMMA F.4. *For any abstract execution $\mathcal{X}$, if it satisfies*

$$\left((\text{SO} \cup \text{WR}); \text{RW}^? \cup \text{WW}\right) ; \text{VIS}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}} \qquad \text{SO}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}}$$

then there exists a new $\mathcal{X}'$ such that $\mathcal{T}_{\mathcal{X}} = \mathcal{T}_{\mathcal{X}'}$, under last-write-win $\mathcal{I}_{\mathcal{X}}(t) = \mathcal{I}_{\mathcal{X}'}(t)$ for all transactions t, and the relations satisfy the following:

$$\text{AR}_{\mathcal{X}'} ; \text{VIS}_{\mathcal{X}'} \subseteq \text{VIS}_{\mathcal{X}'} \qquad \text{SO}_{\mathcal{X}'} \subseteq \text{VIS}_{\mathcal{X}'}$$

and vice versa.

PROOF. Assume abstract execution $\mathcal{X}'$ that satisfies $\text{AR}_{\mathcal{X}'} ; \text{VIS}_{\mathcal{X}'} \subseteq \text{VIS}_{\mathcal{X}'}$ and $\text{SO}_{\mathcal{X}'} \subseteq \text{VIS}_{\mathcal{X}'}$. We already show that:

$$\text{AR}_{\mathcal{X}'} = \left((\text{SO}_{\mathcal{X}} \cup \text{WR}_{\mathcal{X}}); \text{RW}_{\mathcal{X}}^? \cup \text{WW}_{\mathcal{X}} \cup R\right)^+$$

$$\text{VIS}_{\mathcal{X}'} = \left((\text{SO}_{\mathcal{X}} \cup \text{WR}_{\mathcal{X}}); \text{RW}_{\mathcal{X}}^? \cup \text{WW}_{\mathcal{X}} \cup R\right)^* ; (\text{SO}_{\mathcal{X}} \cup \text{WR}_{\mathcal{X}})$$

for some relation $R \subseteq \text{AR}_{\mathcal{X}'}$. If we take $R = \emptyset$, we have the proof for:

$$\text{SO} \subseteq \text{VIS}_{\mathcal{X}} \qquad \left((\text{SO}_{\mathcal{X}} \cup \text{WR}_{\mathcal{X}}); \text{RW}_{\mathcal{X}}^? \cup \text{WW}_{\mathcal{X}}\right) ; \text{VIS}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}}$$

For another way, we pick the R that extends $\left((\text{SO}_{\mathcal{X}} \cup \text{WR}_{\mathcal{X}}); \text{RW}_{\mathcal{X}}^? \cup \text{WW}_{\mathcal{X}} \cup R\right)^+$ to a total order.
□

By Lemma F.4 to prove soundness and completeness of ET_{CP}, it is sufficient to use the definition:

$$(\text{RP}_{\text{LWW}}, \left\{\lambda \mathcal{X}. \left((\text{SO} \cup \text{WR}); \text{RW}^? \cup \text{WW}\right) ; \text{VIS}_{\mathcal{X}}, \lambda \mathcal{X}. \text{SO}_{\mathcal{X}}\right\})$$

For the soundness, we pick the invariant as the following:

$$I_1(\mathcal{X}, cl) = \left(\bigcup_{\{t_{cl}^i \in \mathcal{T}_{\mathcal{X}} \mid i \in \mathbb{N}\}} \text{VIS}_{\mathcal{X}}^{-1}(t_{cl}^i)\right) \setminus \mathcal{T}_{\text{rd}}$$

$$I_2(\mathcal{X}, cl) = \left(\bigcup_{\{t_{cl}^i \in \mathcal{T}_{\mathcal{X}} \mid i \in \mathbb{N}\}} (\text{SO}_{\mathcal{X}}^{-1})^?(t_{cl}^i)\right) \setminus \mathcal{T}_{\text{rd}}$$

where $\mathcal{T}_{\text{rd}}$ is all the read-only transactions included in both

$$\mathcal{T}_{\text{rd}} \in \left(\bigcup_{\{t_{cl}^i \in \mathcal{T}_{\mathcal{X}} \mid i \in \mathbb{N}\}} \text{VIS}_{\mathcal{X}}^{-1}(t_{cl}^i)\right)$$

and

$$\mathcal{T}_{\mathrm{rd}} \in \left(\bigcup_{\{t^i_{cl} \in \mathcal{T}_X \ | \ i \in \mathbb{N}\}} (\mathrm{SO}_X^{-1})^? (t^i_{cl}) \right)$$

Assume a key-value store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $\mathrm{ET}_{\mathrm{CP}} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$. Also choose an arbitrary cl, a transaction identifier $t^n_{cl} \in \mathrm{nextTid}(\mathcal{K}, cl)$, and an abstract execution X such that $\mathcal{K}_X = \mathcal{K}$ and $I_1(X, cl) \cup I_2(X, cl) \subseteq \mathrm{visTx}(\mathcal{K}, u)$. We are about to prove that there exists an extra set of read-only transaction $\mathcal{T}'_{\mathrm{rd}}$ such that the new abstract execution $X' = \mathrm{extend}\left(X, t^n_{cl}, \mathcal{F}, \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}} \cup \mathcal{T}'_{\mathrm{rd}}\right)$ and

$$\forall t. \ (t, t^n_{cl}) \in \mathrm{SO}_{X'} \Rightarrow t \in \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}} \cup \mathcal{T}'_{\mathrm{rd}} \tag{6.7}$$

$$\forall t. \ (t, t^n_{cl}) \in \left((\mathrm{SO}_{X'} \cup \mathrm{WR}_{X'}); \mathrm{RW}^?_{X'} \cup \mathrm{WW}_{X'}\right); \mathrm{VIS}_{X'} \Rightarrow t \in \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}} \cup \mathcal{T}'_{\mathrm{rd}} \tag{6.8}$$

$$I_1(X', cl) \cup I_2(X', cl) \subseteq \mathrm{visTx}(\mathcal{K}_{X'}, u') \tag{6.9}$$

- the invariant I_2 implies the Eq. (6.7) where the proof is the same as RYW in §F.3.
- Eq. (6.8). Note that

$$(t, t^n_{cl}) \in \left((\mathrm{SO}_{X'} \cup \mathrm{WR}_{X'}); \mathrm{RW}^?_{X'} \cup \mathrm{WW}_{X'}\right); \mathrm{VIS}_{X'}$$
$$\Rightarrow (t, t^n_{cl}) \in \left((\mathrm{SO}_X \cup \mathrm{WR}_X); \mathrm{RW}^?_X \cup \mathrm{WW}_X\right); \mathrm{VIS}_{X'}$$

Also, recall that $R_X = R_{\mathcal{K}}$ for $R \in \{\mathrm{SO}, \mathrm{WR}, \mathrm{WW}, \mathrm{RW}\}$. Let

$$\mathcal{T}'_{\mathrm{rd}} = \mathrm{lfpTx}\left(\mathcal{K}, u, (\mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}}); \mathrm{RW}^?_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}}\right)$$

Let assume $t \xrightarrow{(\mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}}); \mathrm{RW}^?_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}}} t'$ and $t' \in \mathrm{lfpTx}\left(\mathcal{K}, u, (\mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}}); \mathrm{RW}^?_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}}\right) \cup \mathcal{T}_{\mathrm{rd}}$. We have two possible cases:

- If $t' \in \mathrm{lfpTx}\left(\mathcal{K}, u, (\mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}}); \mathrm{RW}^?_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}}\right)$, by Theorem F.1, we know

$$t \in \mathrm{lfpTx}\left(\mathcal{K}, u, (\mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}}); \mathrm{RW}^?_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}}\right)$$

- If $t' \in \mathcal{T}_{\mathrm{rd}}$, there are two cases:
  * $t' \in \left(\bigcup_{\{t^n_{cl} \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} \mathrm{VIS}_X^{-1}(t^n_{cl})\right)$. Note that t' is a read-only transaction, which means $t \xrightarrow{\mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}}} t'$. By the property of X (before update) that $(\mathrm{SO} \cup \mathrm{WR}_X); \mathrm{VIS}_X \in \mathrm{VIS}_X$, it is known that $t \in \left(\bigcup_{\{t^n_{cl} \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} \mathrm{VIS}_X^{-1}(t^n_{cl})\right)$, that is, $t \in \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}}$.
  * $t' \in \left(\bigcup_{\{t^n_{cl} \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} \mathrm{SO}_X^{-1}(t^n_{cl})\right)$ and it is a read only transaction. Then we know $t \in (\mathrm{SO} \cup \mathrm{WR}_X)^{-1}\left(\bigcup_{\{t^n_{cl} \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} \mathrm{SO}_X^{-1}(t^n_{cl})\right)$. By the property of X (before update) that

$\mathsf{SO} \cup \mathsf{WR}_{\mathcal{X}} \in \mathsf{VIS}_{\mathcal{X}}$, it follows:

$$t \in VIS_{\mathcal{X}}^{-1}\left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} \mathsf{SO}_{\mathcal{X}}^{-1}(t_{cl}^n)\right)$$

$$= \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \mid n \in \mathbb{N}\}} \mathsf{VIS}_{\mathcal{X}}^{-1}(t_{cl}^n)\right)$$

$$= \mathsf{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathsf{rd}}$$

- Since CP satisfies RYW and MR, thus invariants I_1 and I_2 are preserved after update.

The execution test $\mathsf{ET}_{\mathsf{CP}}$ is complete with respect to the axiomatic definition

$$(\mathsf{RP}_{\mathsf{LWW}}, \{\lambda \mathcal{X}.\mathsf{AR}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}}, \lambda \mathcal{X}.\,\mathsf{SO}_{\mathcal{X}}\})$$

Assume i-*th* transaction t_i in the arbitrary order, and let view $u_i = \mathsf{getView}(\mathcal{X}, \mathsf{VIS}_{\mathcal{X}}^{-1}(t_i))$. We pick final view as $u_i' = \mathsf{getView}(\mathcal{X}, (\mathsf{AR}_{\mathcal{X}}^{-1})^?(t_i) \cap \mathsf{VIS}_{\mathcal{X}}^{-1}(t_i'))$, if $t_i' = \min_{\mathsf{SO}}\left\{t' \mid t_i \xrightarrow{\mathsf{SO}} t'\right\}$ is defined, otherwise $u_i' = \mathsf{getView}(\mathcal{X}, (\mathsf{AR}_{\mathcal{X}}^{-1})^?(t_i))$. Let the $\mathcal{K} = \mathcal{K}_{\mathsf{cut}(\mathcal{X}, i-1)}$. Now we prove the three parts separately.

- MR. By Prop. F.3 since $\mathsf{VIS}_{\mathcal{X}}; \mathsf{SO}_{\mathcal{X}} \subseteq \mathsf{AR}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}$ so it follows as in §F.1.
- RYW. For RYW, since $\mathsf{WR}_{\mathcal{X}}; \mathsf{SO}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{AR}_{\mathcal{X}}; \mathsf{AR}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}$, the proof is as the same proof as in §F.3.
- $\mathsf{allowed}\left((\mathsf{SO}; \mathsf{RW}_{\mathcal{K}}^?) \cup (\mathsf{WR}_{\mathcal{K}}; \mathsf{RW}_{\mathcal{K}}^?) \cup \mathsf{WW}_{\mathcal{K}}\right)$ can be derived from Theorem F.1 and

$$(\mathsf{SO}; \mathsf{RW}_{\mathcal{X}}^?) \cup (\mathsf{WR}_{\mathcal{X}}; \mathsf{RW}_{\mathcal{X}}^?) \cup \mathsf{WW}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{AR}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}$$

F.8 Parallel Snapshot Isolation PSI

The axiomatic definition for PSI is

$$(\mathsf{RP}_{\mathsf{LWW}}, \{\lambda \mathcal{X}.\mathsf{VIS}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}}, \lambda \mathcal{X}.\,\mathsf{SO}_{\mathcal{X}}, \lambda \mathcal{X}.\mathsf{WW}_{\mathcal{X}}\})$$

There exist a minimum visibility such that

$$(\mathsf{RP}_{\mathsf{LWW}}, \{\lambda \mathcal{X}.(\mathsf{WR}_{\mathcal{X}} \cup \mathsf{WW}_{\mathcal{X}} \cup \mathsf{SO}); \mathsf{VIS}_{\mathcal{X}}, \lambda \mathcal{X}.\,\mathsf{SO}_{\mathcal{X}}, \lambda \mathcal{X}.\mathsf{WW}_{\mathcal{X}}\})$$

by solve the following inequalities:

$$\mathsf{WR} \subseteq \mathsf{VIS}$$
$$\mathsf{WW} \subseteq \mathsf{VIS}$$
$$\mathsf{SO} \subseteq \mathsf{VIS}$$
$$\mathsf{VIS}; \mathsf{VIS} \subseteq \mathsf{VIS}$$

It is easy to see the former implies to later. For another way round, Lemma F.5.

LEMMA F.5. *For any abstract execution $\mathcal{X}$ under last-write-win, if it satisfies the following:*

$$(\mathsf{WR}_{\mathcal{X}} \cup \mathsf{WW}_{\mathcal{X}} \cup \mathsf{SO}_{\mathcal{X}}); \mathsf{VIS}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}} \quad \mathsf{SO}_{\mathcal{X}} \subseteq \mathsf{VIS}_{\mathcal{X}}$$

There exists a new abstract execution $\mathcal{X}'$ where $\mathcal{T}_{\mathcal{X}} = \mathcal{T}_{\mathcal{X}'}$, $\mathsf{AR}_{\mathcal{X}} = \mathsf{AR}_{\mathcal{X}'}$, $\mathsf{VIS}_{\mathcal{X}}; \mathsf{VIS}_{\mathcal{X}'} \subseteq \mathsf{VIS}_{\mathcal{X}'}$, and under last-write-win $\mathscr{T}_{\mathcal{X}}(t) = \mathscr{T}_{\mathcal{X}'}(t)$ for all transactions t.

PROOF. we erase some visibility relation for each transaction following the order of arbitration AR until the visibility is transitive. Assume the i-*th* transaction t_i with respect to the arbitration order. Let R_i denote a new visibility for transaction t_i such that $R_i|_2 = \{t_i\}$ and the visibility relation

before (including) t_i is transitive. Let $X_i = \mathcal{K}_{\mathrm{cut}(X,i)}$ and $\mathrm{VIS}_i = \bigcup_{0 \leq k \leq i} R_i$. For each step, says i-*th* step, we preserve the following:

$$\mathrm{VIS}_i; \mathrm{VIS}_i \subseteq \mathrm{VIS}_i \tag{6.10}$$

$$\forall t. \ (t, t_i) \in R_i \Rightarrow (t, t_i) \in (\mathrm{WR}_i \cup \mathrm{WW}_i \cup \mathrm{SO}_i) \tag{6.11}$$

- Base case: $i = 1$ and $R_1 = \emptyset$. Assume it is from client cl. There is no transaction committed before, so $\mathrm{VIS}_1 = \emptyset$ and $\mathrm{VIS}_1; \mathrm{VIS}_1 \subseteq \mathrm{VIS}_1$ as Eq. (6.10).
- Inductive case: i-*th* step. Suppose the (i-1)-*th* step satisfies Eq. (6.10) and Eq. (6.11). Let consider i-*th* step and the transaction t_i. Initially we take R_i as empty set. We first extend R_i by closing with respect to WR_i and prove that it does not affect any read from the transaction t_i. Then we will do the same for SO_i and WW_i.

  - WR_i. For any read $(\mathrm{r}, k, v) \in t_i$, there must be a transaction t_j that $t_j \xrightarrow{\mathrm{WR}(X_i, k), \mathrm{AR}} t_i$ and $j < i$. We include $(t_j, t_i) \in R_i$. Let consider all the visible transactions of t_j. Assume a transaction $t' \in \mathrm{VIS}_{i-1}^{-1}(t_j)$, thus $t' \in \mathrm{VIS}_j^{-1}(t_j) = R_j^{-1}(t_j)$. It is safe to include $(t', t_i) \in R_i$ without affecting the read result, because those transaction t' is already visible for t_i in the abstract execution X: by Eq. (6.11) we know $R_j \subseteq (\mathrm{WR}_j \cup \mathrm{SO}_j \cup \mathrm{WW}_j)^+ \subseteq (\mathrm{WR}_X \cup \mathrm{SO}_X \cup \mathrm{WW}_X)^+$, and by the definition of $\mathrm{WR}(X_i, k)$ we know $\mathrm{WR}(X_i, k) \subseteq \mathrm{VIS}_X$.
  - Given $\mathrm{SO}_X \subseteq \mathrm{VIS}_X$ (and $\mathrm{WW}_X \subseteq \mathrm{VIS}_X$ respectively) we include (t_j, t_i) for some t_j such that $t_j \xrightarrow{\mathrm{SO}_X} t_i$ (and $t_j \xrightarrow{\mathrm{WW}_X} t_i$ respectively). For the similar reason as WR, it is safe to includes all the visible transactions t' for t_j, i.e. $t' \in R_j^{-1}$.

  By the construction, both Eq. (6.2) and Eq. (6.3) are preserved. Thus we have the proof.

$\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad \square$

To prove soundness, we pick an invariant for the $\mathrm{ET}_{\mathrm{PSI}}$ as the union of those for MR and RYW shown in the following:

$$I_1(X, cl) = \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} \mathrm{VIS}_X^{-1}(t_{cl}^n) \right) \setminus \mathcal{T}_{\mathrm{rd}}$$

$$I_2(X, cl) = \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} (\mathrm{SO}_X^{-1})^?(t_{cl}^n) \right) \setminus \mathcal{T}_{\mathrm{rd}}$$

where $\mathcal{T}_{\mathrm{rd}}$ is all the read-only transactions included in both

$$\mathcal{T}_{\mathrm{rd}} \in \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} \mathrm{VIS}_X^{-1}(t_{cl}^n) \right)$$

and

$$\mathcal{T}_{\mathrm{rd}} \in \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_X \ | \ n \in \mathbb{N}\}} (\mathrm{SO}_X^{-1})^?(t_{cl}^n) \right)$$

Assume a kv-store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $\mathrm{ET}_{\mathrm{PSI}} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$. Also choose an arbitrary cl, a transaction identifier $t_{cl}^n \in \mathrm{nextTid}(\mathcal{K}, cl)$, and an abstract execution X such that $\mathcal{K}_X = \mathcal{K}$ and $I_1(X, cl) \cup I_2(X, cl) \subseteq \mathrm{visTx}(\mathcal{K}, u)$. We are about to

prove there exists an extra set of read-only transactions $\mathcal{T}'_{\mathrm{rd}}$ such that the new abstract execution $\mathcal{X}' = \mathrm{extend}\big(\mathcal{X}, t^n_{cl}, \mathcal{F}, \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}} \cup \mathcal{T}'_{\mathrm{rd}}\big)$ and:

$$\forall t. \ (t, t^n_{cl}) \in \mathrm{SO}_{\mathcal{X}'} \Rightarrow t \in \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}} \cup \mathcal{T}'_{\mathrm{rd}} \tag{6.12}$$

$$\forall t. \ (t, t^n_{cl}) \in \mathrm{WW}_{\mathcal{X}'} \Rightarrow t \in \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}} \cup \mathcal{T}'_{\mathrm{rd}} \tag{6.13}$$

$$\forall t. \ (t, t^n_{cl}) \in (\mathrm{SO}_{\mathcal{X}'} \cup \mathrm{WR}_{\mathcal{X}'} \cup \mathrm{WW}_{\mathcal{X}'}); \mathrm{VIS}_{\mathcal{X}'} \Rightarrow t \in \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}} \cup \mathcal{T}'_{\mathrm{rd}} \tag{6.14}$$

$$I_1(\mathcal{X}', cl) \cup I_2(\mathcal{X}', cl) \subseteq \mathrm{visTx}(\mathcal{K}_{\mathcal{X}'}, u') \tag{6.15}$$

- The invariant I_2 implies Eq. (6.12) as the same as RYW in §F.3.
- Since PSI also satisfies UA, the Eq. (6.17) can be proven as the same as UA in §F.6.
- Eq. (6.14). Note that

$$(t, t^n_{cl}) \in (\mathrm{SO}_{\mathcal{X}'} \cup \mathrm{WR}_{\mathcal{X}'} \cup \mathrm{WW}_{\mathcal{X}'}); \mathrm{VIS}_{\mathcal{X}'} \Rightarrow (t, t^n_{cl}) \in (\mathrm{SO}_{\mathcal{X}} \cup \mathrm{WR}_{\mathcal{X}} \cup \mathrm{WW}_{\mathcal{X}}); \mathrm{VIS}_{\mathcal{X}'}$$

Also, recall that $\mathrm{SO}_{\mathcal{X}} = \mathrm{SO}_{\mathcal{K}}$, $\mathrm{WR}_{\mathcal{X}} = \mathrm{WR}_{\mathcal{K}}$ and $\mathrm{WW}_{\mathcal{X}} = \mathrm{WW}_{\mathcal{K}}$. Let

$$\mathcal{T}'_{\mathrm{rd}} = \mathrm{lfpTx}(\mathcal{K}, u, \mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}})$$

This means

$$\mathcal{X}' = \mathrm{extend}\big(\mathcal{X}, t^n_{cl}, \mathcal{F}, \mathrm{lfpTx}(\mathcal{K}, u, \mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}}) \cup \mathcal{T}_{\mathrm{rd}}\big)$$

Let assume $t \xrightarrow{\mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}}} t'$ and $t' \in \mathrm{lfpTx}(\mathcal{K}, u, \mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}}) \cup \mathcal{T}_{\mathrm{rd}}$. We have two possible cases:

– If $t' \in \mathrm{lfpTx}(\mathcal{K}, u, \mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}})$, by Theorem F.1, we know

$$t \in \mathrm{lfpTx}(\mathcal{K}, u, \mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}} \cup \mathrm{WW}_{\mathcal{K}})$$

– If $t' \in \mathcal{T}_{\mathrm{rd}}$, there are two cases:
  * $t' \in \Big(\bigcup_{\{t^n_{cl} \in \mathcal{T}_{\mathcal{X}} \ \mid \ n \in \mathbb{N}\}} \mathrm{VIS}^{-1}_{\mathcal{X}}(t^n_{cl})\Big)$. Since t' is a read-only transaction, it means $t \xrightarrow{\mathrm{SO}_{\mathcal{K}} \cup \mathrm{WR}_{\mathcal{K}}} t'$. By the property of $\mathcal{X}$ (before update) that $(\mathrm{SO} \cup \mathrm{WR}_{\mathcal{X}}); \mathrm{VIS}_{\mathcal{X}} \in \mathrm{VIS}_{\mathcal{X}}$, it is known that $t \in \Big(\bigcup_{\{t^n_{cl} \in \mathcal{T}_{\mathcal{X}} \ \mid \ n \in \mathbb{N}\}} \mathrm{VIS}^{-1}_{\mathcal{X}}(t^n_{cl})\Big)$, that is, $t \in \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}}$.
  * $t' \in \Big(\bigcup_{\{t^n_{cl} \in \mathcal{T}_{\mathcal{X}} \ \mid \ n \in \mathbb{N}\}} \mathrm{SO}^{-1}_{\mathcal{X}}(t^n_{cl})\Big)$. Given that t' is a read only transaction, we know $t \in (\mathrm{SO} \cup \mathrm{WR}_{\mathcal{X}})^{-1}\Big(\bigcup_{\{t^n_{cl} \in \mathcal{T}_{\mathcal{X}} \ \mid \ n \in \mathbb{N}\}} \mathrm{SO}^{-1}_{\mathcal{X}}(t^n_{cl})\Big)$. By the property of $\mathcal{X}$ (before update) that $\mathrm{SO} \cup \mathrm{WR}_{\mathcal{X}} \in \mathrm{VIS}_{\mathcal{X}}$, it follows:

$$t \in VIS^{-1}_{\mathcal{X}}\left(\bigcup_{\{t^n_{cl} \in \mathcal{T}_{\mathcal{X}} \ \mid \ n \in \mathbb{N}\}} \mathrm{SO}^{-1}_{\mathcal{X}}(t^n_{cl})\right)$$

$$= \left(\bigcup_{\{t^n_{cl} \in \mathcal{T}_{\mathcal{X}} \ \mid \ n \in \mathbb{N}\}} \mathrm{VIS}^{-1}_{\mathcal{X}}(t^n_{cl})\right)$$

$$= \mathrm{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\mathrm{rd}}$$

- Finally the new abstract execution preserves the invariant I_1 and I_2 because CC satisfies MW and RYW.

The execution test $\mathrm{ET}_{\mathrm{PSI}}$ is complete with respect to the axiomatic definition

$$(\mathrm{RP}_{\mathrm{LWW}}, \{\lambda\mathcal{X}.\mathrm{VIS}_{\mathcal{X}}; \mathrm{VIS}_{\mathcal{X}}, \lambda\mathcal{X}.\,\mathrm{SO}_{\mathcal{X}}, \lambda\mathcal{X}.\,\mathrm{WW}_{\mathcal{X}}\})$$

Assume i-*th* transaction t_i in the arbitrary order, and let view $u_i = \text{getView}\big(X, \text{VIS}_X^{-1}(t_i)\big)$. We pick final view as $u_i' = \text{getView}\big(X, (\text{AR}_X^{-1})^?(t_i) \cap \text{VIS}_X^{-1}(t_i')\big)$, if $t_i' = \min_{\text{SO}}\left\{ t' \,\middle|\, t_i \xrightarrow{\text{SO}} t' \right\}$ is defined, otherwise $u_i' = \text{getView}\big(X, (\text{AR}_X^{-1})^?(t_i)\big)$. Let the $\mathcal{K} = \mathcal{K}_{\text{cut}(X, i-1)}$. Now we prove the three parts separately.

- MR. By Prop. F.3 since $\text{VIS}_X; \text{SO}_X \subseteq \text{VIS}_X; \text{VIS}_X \subseteq \text{VIS}_X$ so it follows as in §F.1.
- RYW. For RYW, since $\text{WR}_X; \text{SO}_X; \text{VIS}_X \subseteq \text{VIS}_X; \text{VIS}_X; \text{VIS}_X \subseteq \text{VIS}_X$, the proof is as the same proof as in §F.3.
- UA. Since $\text{WW}_X \subseteq \text{VIS}_X$, the proof is as the same proof as in §F.6.
- $\texttt{allowed}(\text{WR}_\mathcal{K}\text{WW}_\mathcal{K} \cup \cup\text{SO})$. It is derived from Theorem F.1 and

$$(\text{WR}_X \cup \text{WW}_X \cup \text{SO}_X); \text{VIS}_X \subseteq \text{VIS}_X; \text{VIS}_X \subseteq \text{VIS}_X$$

F.9 Snapshot Isolation SI

The axiomatic definition for SI is

$$(\text{RP}_{\text{LWW}}, \{\lambda X.\text{AR}_X; \text{VIS}_X, \lambda X.\, \text{SO}_X, \lambda X.\, \text{WW}_X\})$$

By a lemma proven in [Cerone and Gotsman 2016], for any X satisfies the SI there exists an equivalent X' with minimum visibility $\text{VIS}_{X'} \subseteq \text{VIS}_X$ satisfying

$$\left(\text{RP}_{\text{LWW}}, \left\{\lambda X.\, \big((\text{SO}_{X'} \cup \text{WW}_{X'} \cup \text{WR}_{X'}); \text{RW}_{X'}^?\big); \text{VIS}_{X'}', \lambda X.\, (\text{WW}_{X'} \cup \text{SO}_{X'})\right\}\right)$$

Under the minimum visibility VIS all the transactions still have the same behaviour as before, meaning they do not violate last-write-win.

   To prove the soundness, we pick the invariant as the following:

$$I_1(X, cl) = \left(\bigcup_{\{t_{cl}^i \in \mathcal{T}_X \mid i \in \mathbb{N}\}} \text{VIS}_X^{-1}(t_{cl}^i)\right) \setminus \mathcal{T}_{\text{rd}}$$

$$I_2(X, cl) = \left(\bigcup_{\{t_{cl}^i \in \mathcal{T}_X \mid i \in \mathbb{N}\}} (\text{SO}_X^{-1})^?(t_{cl}^i)\right) \setminus \mathcal{T}_{\text{rd}}$$

where $\mathcal{T}_{\text{rd}}$ is all the read-only transactions included in both

$$\mathcal{T}_{\text{rd}} \in \left(\bigcup_{\{t_{cl}^i \in \mathcal{T}_X \mid i \in \mathbb{N}\}} \text{VIS}_X^{-1}(t_{cl}^i)\right)$$

and

$$\mathcal{T}_{\text{rd}} \in \left(\bigcup_{\{t_{cl}^i \in \mathcal{T}_X \mid i \in \mathbb{N}\}} (\text{SO}_X^{-1})^?(t_{cl}^i)\right)$$

Assume a kv-store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $\text{ET}_{\text{SI}} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$. Also choose an arbitrary cl, a transaction identifier $t_{cl}^n \in \text{nextTid}(\mathcal{K}, cl)$, and an abstract execution X such that $\mathcal{K}_X = \mathcal{K}$ and $I_1(X, cl) \cup I_2(X, cl) \subseteq \text{visTx}(\mathcal{K}, u)$. We are about to prove there exists an extra set of read-only transaction $\mathcal{T}_{\text{rd}}'$ such that the new abstract execution

$\mathcal{X}' = \text{extend}\Big(\mathcal{X}, t_{cl}^n, \mathcal{F}, \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}} \cup \mathcal{T}_{\text{rd}}'\Big)$ and

$$\forall t. \ (t, t_{cl}^n) \in \text{SO}_{\mathcal{X}'} \Rightarrow t \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}} \cup \mathcal{T}_{\text{rd}}' \tag{6.16}$$

$$\forall t. \ (t, t_{cl}^n) \in \text{WW}_{\mathcal{X}'} \Rightarrow t \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}} \cup \mathcal{T}_{\text{rd}}' \tag{6.17}$$

$$\forall t. \ (t, t_{cl}^n) \in \Big((\text{SO}_{\mathcal{X}'} \cup \text{WW}_{\mathcal{X}'} \cup \text{WR}_{\mathcal{X}'}); \text{RW}_{\mathcal{X}'}^?\Big); \text{VIS}_{\mathcal{X}'} \Rightarrow t \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}} \cup \mathcal{T}_{\text{rd}}' \tag{6.18}$$

$$I_1(\mathcal{X}', cl) \cup I_2(\mathcal{X}', cl) \subseteq \text{visTx}(\mathcal{K}_{\mathcal{X}'}, u') \tag{6.19}$$

- The invariant I_2 implies Eq. (6.16) as the same as RYW in §F.3.
- Since SI also satisfies UA, the Eq. (6.17) can be proven as the same as UA in §F.6.
- Eq. (6.18). Note that

$$(t, t_{cl}^n) \in \Big((\text{SO}_{\mathcal{X}'} \cup \text{WW}_{\mathcal{X}'} \cup \text{WR}_{\mathcal{X}'}); \text{RW}_{\mathcal{X}'}^?\Big); \text{VIS}_{\mathcal{X}'}$$
$$\Rightarrow (t, t_{cl}^n) \in \Big((\text{SO}_{\mathcal{X}} \cup \text{WW}_{\mathcal{X}} \cup \text{WR}_{\mathcal{X}}); \text{RW}_{\mathcal{X}}^?\Big); \text{VIS}_{\mathcal{X}'}$$

Also, recall that $R_{\mathcal{X}} = R_{\mathcal{K}}$ for $R \in \{\text{SO}, \text{WR}, \text{WW}, \text{RW}\}$. Let

$$\mathcal{T}_{\text{rd}}' = \text{lfpTx}\Big(\mathcal{K}, u, \big((\text{SO}_{\mathcal{X}'} \cup \text{WW}_{\mathcal{X}'} \cup \text{WR}_{\mathcal{X}'}); \text{RW}_{\mathcal{X}'}^?\big)\Big)$$

This means that

$$\mathcal{X}' = \text{extend}\Big(\mathcal{X}, t_{cl}^n, \mathcal{F}, \text{lfpTx}\big(\mathcal{K}, u, ((\text{SO}_{\mathcal{X}'} \cup \text{WW}_{\mathcal{X}'} \cup \text{WR}_{\mathcal{X}'}); \text{RW}_{\mathcal{X}'}^?)\big) \cup \mathcal{T}_{\text{rd}}\Big)$$

Let assume $t \xrightarrow{\big((\text{SO}_{\mathcal{X}'} \cup \text{WW}_{\mathcal{X}'} \cup \text{WR}_{\mathcal{X}'}); \text{RW}_{\mathcal{X}'}^?\big)} t'$ and

$$t' \in \text{lfpTx}\Big(\mathcal{K}, u, \big((\text{SO}_{\mathcal{X}'} \cup \text{WW}_{\mathcal{X}'} \cup \text{WR}_{\mathcal{X}'}); \text{RW}_{\mathcal{X}'}^?\big)\Big) \cup \mathcal{T}_{\text{rd}}$$

We have two possible cases:

- If $t' \in \text{lfpTx}\Big(\mathcal{K}, u, \big((\text{SO}_{\mathcal{X}'} \cup \text{WW}_{\mathcal{X}'} \cup \text{WR}_{\mathcal{X}'}); \text{RW}_{\mathcal{X}'}^?\big)\Big)$, by Theorem F.1, we know

$$t \in \text{lfpTx}\Big(\mathcal{K}, u, \big((\text{SO}_{\mathcal{X}'} \cup \text{WW}_{\mathcal{X}'} \cup \text{WR}_{\mathcal{X}'}); \text{RW}_{\mathcal{X}'}^?\big)\Big)$$

- If $t' \in \mathcal{T}_{\text{rd}}$, there are two cases:
  * $t' \in \Big(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \ | \ n \in \mathbb{N}\}} \text{VIS}_{\mathcal{X}}^{-1}(t_{cl}^n)\Big)$. Since t' is a read-only transaction, it means $t \xrightarrow{\text{SO}_{\mathcal{K}} \cup \text{WR}_{\mathcal{K}}} t'$. By the property of $\mathcal{X}$ (before update) that $(\text{SO} \cup \text{WR}_{\mathcal{X}}); \text{VIS}_{\mathcal{X}} \in \text{VIS}_{\mathcal{X}}$, it is known that $t \in \Big(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \ | \ n \in \mathbb{N}\}} \text{VIS}_{\mathcal{X}}^{-1}(t_{cl}^n)\Big)$, that is, $t \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}}$.
  * $t' \in \Big(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \ | \ n \in \mathbb{N}\}} \text{SO}_{\mathcal{X}}^{-1}(t_{cl}^n)\Big)$. Given that t' is a read only transaction, we know $t \in (\text{SO} \cup \text{WR}_{\mathcal{X}})^{-1}\Big(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \ | \ n \in \mathbb{N}\}} \text{SO}_{\mathcal{X}}^{-1}(t_{cl}^n)\Big)$. By the property of $\mathcal{X}$ (before update) that $\text{SO} \cup \text{WR}_{\mathcal{X}} \in \text{VIS}_{\mathcal{X}}$, it follows:

$$t \in VIS_{\mathcal{X}}^{-1}\left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \ | \ n \in \mathbb{N}\}} \text{SO}_{\mathcal{X}}^{-1}(t_{cl}^n)\right)$$
$$= \left(\bigcup_{\{t_{cl}^n \in \mathcal{T}_{\mathcal{X}} \ | \ n \in \mathbb{N}\}} \text{VIS}_{\mathcal{X}}^{-1}(t_{cl}^n)\right)$$
$$= \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}}$$

- Since SI satisfies RYW and MR, thus invariants I_1 and I_2 are preserved, that is, Eq. (6.19).

The execution test ET_{SI} is complete with respect to the axiomatic definition

$$(\text{RP}_{\text{LWW}}, \{\lambda \mathcal{X}.\text{AR}_{\mathcal{X}}; \text{VIS}_{\mathcal{X}}, \lambda \mathcal{X}.\text{SO}_{\mathcal{X}}, \lambda \mathcal{X}.\text{WW}_{\mathcal{X}}\})$$

Assume i-*th* transaction t_i in the arbitrary order, and let view $u_i = \text{getView}(\mathcal{X}, \text{VIS}_{\mathcal{X}}^{-1}(t_i))$. We pick final view as $u_i' = \text{getView}(\mathcal{X}, (\text{AR}_{\mathcal{X}}^{-1})^?(t_i) \cap \text{VIS}_{\mathcal{X}}^{-1}(t_i'))$, if $t_i' = \min_{\text{SO}} \left\{ t' \mid t_i \xrightarrow{\text{SO}} t' \right\}$ is defined, otherwise $u_i' = \text{getView}(\mathcal{X}, (\text{AR}_{\mathcal{X}}^{-1})^?(t_i))$. Let the $\mathcal{K} = \mathcal{K}_{\text{cut}(\mathcal{X}, i-1)}$. Now we prove the three parts separately.

- MR. By Prop. F.3 since $\text{VIS}_{\mathcal{X}}; \text{SO}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}}; \text{VIS}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}}$ so it follows as in §F.1.
- RYW. For RYW, since $\text{WR}_{\mathcal{X}}; \text{SO}_{\mathcal{X}}; \text{VIS}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}}; \text{VIS}_{\mathcal{X}}; \text{VIS}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}}$, the proof is as the same proof as in §F.3.
- UA. Since $\text{WW}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}}$, the proof is as the same proof as in §F.6.
- $\texttt{allowed}\left(\left((\text{SO}_{\mathcal{K}} \cup \text{WW}_{\mathcal{K}} \cup \text{WR}_{\mathcal{K}}); \text{RW}_{\mathcal{K}}^?\right)\right)$. It is derived from Theorem F.1 and

$$\left((\text{SO}_{\mathcal{X}} \cup \text{WW}_{\mathcal{X}} \cup \text{WR}_{\mathcal{X}}); \text{RW}_{\mathcal{X}}^?\right); \text{VIS}_{\mathcal{X}} \subseteq \text{AR}_{\mathcal{X}}; \text{VIS}_{\mathcal{X}} \subseteq \text{VIS}_{\mathcal{X}}$$

F.10 Serialisability SER

The execution test ET_{SER} is sound with respect to the axiomatic definition

$$(\text{RP}_{\text{LWW}}, \{\lambda \mathcal{X}.\text{AR}\})$$

We pick the invariant as $I(\mathcal{X}, cl) = \emptyset$, given the fact of no constraint on the view after update. Assume a kv-store $\mathcal{K}$, an initial and a final view u, u' a fingerprint $\mathcal{F}$ such that $\text{ET}_{\text{SER}} \vdash (\mathcal{K}, u) \triangleright \mathcal{F} : (\mathcal{K}', u')$. Also choose an arbitrary cl, a transaction identifier $t \in \text{nextTid}(\mathcal{K}, cl)$, and an abstract execution $\mathcal{X}$ such that $\mathcal{K}_{\mathcal{X}} = \mathcal{K}$ and $I(\mathcal{X}, cl) = \emptyset \subseteq \text{visTx}(\mathcal{K}, u)$. Let $\mathcal{X}' = \text{extend}(\mathcal{X}, t, \text{visTx}(\mathcal{K}, u), \mathcal{F})$. Note that since the invariant is empty set, it remains to prove there exists a set of read-only transactions $\mathcal{T}_{\text{rd}}$ such that:

$$\forall t'. \ t' \xrightarrow{\text{AR}_{\mathcal{X}'}} t \Rightarrow t' \in \text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}}$$

Since the abstract execution satisfies the constraint for SER, i.e. $\text{AR} \subseteq \text{VIS}$, we know $\text{AR} = \text{VIS}$. Since $\text{visTx}(\mathcal{K}, u)$ contains all transactions that write at least a key, we can pick a $\mathcal{T}_{\text{rd}}$ such that $\text{visTx}(\mathcal{K}, u) \cup \mathcal{T}_{\text{rd}} = \mathcal{T}_{\mathcal{X}}$, which gives us the proof.

The execution test ET_{UA} is complete with respect to the axiomatic definition $(\text{RP}_{\text{LWW}}, \{\lambda \mathcal{X}.\text{AR}_{\mathcal{X}}\})$. Assume i-*th* transaction t_i in the arbitrary order, and let view $u_i = \text{getView}(\mathcal{X}, \text{VIS}_{\mathcal{X}}^{-1}(t_i))$. We also pick any final view such that $u_i' \subseteq \text{getView}(\mathcal{X}, (\text{AR}_{\mathcal{X}}^{-1})^?(t_i))$. Note that there is nothing to prove for u_i', Now we need to prove the following:

$$\forall k, j. \ 0 \leq j < |\mathcal{K}_{\text{cut}(\mathcal{X}, i-1)}(k)| \Rightarrow j \in u_i(k)$$

Because $\text{VIS}^{-1}(t_i) = \text{AR}^{-1}(t_i) = \left\{ t \mid t \in \mathcal{K}_{\text{cut}(\mathcal{X}, i-1)} \right\}$, so for any key k and index j such that $0 \leq j < |\mathcal{K}_{\text{cut}(\mathcal{X}, i-1)}(k)|$, the j-*th* version of the key contains in the view, i.e. $j \in u(k)$.

G PROGRAM ANALYSIS

We give applications of our theory aimed at showing the robustness of a transactional library against a given consistency model. The first application considers a single counter library, and proves that it is robust against Parallel Snapshot Isolation. We present a general robustness conditions for WSI and then show multiple-counter example and a banking example [Alomari et al. 2008] are robust against WSI.

G.1 Single counter

We start by reviewing the transactional code for the increment and read operations provided by a counter object over a key k, denoted as $\mathsf{inc}(k)$ and $\mathsf{read}(k)$, respectively.

$$\mathsf{inc}(k) = \begin{bmatrix} \mathsf{a} := [k]; \\ [k] := \mathsf{a} + 1; \end{bmatrix} \qquad\qquad \mathsf{read}(k) = [\, \mathsf{a} := [k]; \,]$$

Clients can interact with the key-value store only by invoking the $\mathsf{inc}(k)$ and $\mathsf{read}(k)$ operations. A transactional library is a set of transactional operations. For a single counter over key k, we define the transactional library $\mathsf{Counter}(k) = \{\mathsf{inc}(k), \mathsf{read}(k)\}$, while for multiple counters over a set of keys $\mathsf{K} = \{k_i\}_{i \in I}$, respectively, we define $\mathsf{Counter}(\mathsf{K}) = \bigcup_{i \in I} \mathsf{Counter}(k_i)$.

KV-store semantics of a transactional library. Given the transactional code $[\mathsf{T}]$, we define $\mathcal{F}(\mathcal{K}, u, [\mathsf{T}])$ to be the fingerprint that would be produced by a client that has view u over the kv-store $\mathcal{K}$, upon executing $[\mathsf{T}]$. For the $\mathsf{inc}(k)$ and $\mathsf{read}(k)$ operations discussed above, we have that

$$\mathcal{F}(\mathcal{K}, u, \mathsf{inc}(k)) = \{(\mathsf{r}, k, n), (\mathsf{w}, k, n+1) \mid n = \mathsf{snapshot}(\mathcal{K}, u)(k)\}$$

and

$$\mathcal{F}(\mathcal{K}, u, \mathsf{read}(k)) = \{(\mathsf{r}, k, n) \mid n = \mathsf{snapshot}(\mathcal{K}, u)(k)\}$$

Given an execution test ET, and a transactional library $L = \{[\mathsf{T}_i]\}_{i \in I}$, we define the set of valid ET-traces for L as the set $\mathsf{Traces}(\mathsf{ET}, \{[\mathsf{T}_i]\}_{i \in I})$ of ET-traces in which only ET-reductions of the form

$$(\mathcal{K}_0, \mathcal{U}_0) \xrightarrow{(cl_0, \lambda_0)}_{\mathsf{ET}} (\mathcal{K}_1, \mathcal{U}_1) \xrightarrow{(cl_1, \lambda_1)}_{\mathsf{ET}} \cdots \xrightarrow{(cl_{n-1}, \lambda_{n-1})}_{\mathsf{ET}} (\mathcal{K}_n, \mathcal{U}_n),$$

where for any $j = 0, \cdots, n-1$, either $\lambda_j = \varepsilon$ or $\lambda_j = \mathcal{F}(\mathcal{K}_j, \mathcal{U}_j(cl_j), [\mathsf{T}_i])$ for some $i \in I$. Henceforth we commit an abuse of notation and write $(\mathcal{K}, \mathcal{U}) \xrightarrow{(cl, [\mathsf{T}])}_{\mathsf{ET}} (\mathcal{K}', \mathcal{U}')$ in lieu of $(\mathcal{K}, \mathcal{U}) \xrightarrow{(cl, \mathcal{F}(\mathcal{K}, \mathcal{U}(cl), [\mathsf{T}]))}_{\mathsf{ET}} (\mathcal{K}', \mathcal{U}')$. We also let $\mathsf{KVStores}(\mathsf{ET}, \{[\mathsf{T}_i]\}_{i \in I})$ be the set of kv-stores that can be obtained when clients can only perform operations from $\{[\mathsf{T}_i]\}_{i \in I}$ under the execution test ET. Specifically,

$$\mathsf{KVStores}(\mathsf{ET}, \{[\mathsf{T}_i]\}_{i \in I}) \triangleq \left\{ \mathcal{K} \,\middle|\, \left((\mathcal{K}_0, \mathcal{U}_0) \xrightarrow{\cdot}_{\mathsf{ET}} \cdots \xrightarrow{\cdot}_{\mathsf{ET}} (\mathcal{K}, \_) \right) \in \mathsf{Traces}(\mathsf{ET}, \{[\mathsf{T}_i]\}_{i \in I}) \right\}$$

G.1.1 Anomaly of a single counter under Causal Consistency. It is well known that the transactional library consisting of a single counter over a single key, $\mathsf{Counter}(k)$, implemented on top of a kv-store guaranteeing Causal Consistency, leads to executions over the kv-store that cannot be simulated by the same transactional library implemented on top of a serialisable kv-store. For simplicity, let us assume that $\mathsf{KEY} = \{k\}$. Let $\mathcal{K}_0 = [k \mapsto (0, t_0, \emptyset)]$, $\mathcal{K}_1 = [k \mapsto (0, t_0, \{t^1_{cl_1}\}) :: (0, t^1_{cl_1}, \emptyset)$, $\mathcal{K}_2 = [k \mapsto (0, t_0, \{t^1_{cl_1}, t^1_{cl_2}\}) :: (0, t^1_{cl_1}, \emptyset) :: (0, t^1_{cl_2}, \emptyset)$. Let also $u_0 = [k \mapsto 0]$. Then we have that

$$(\mathcal{K}_0, [cl_1 \mapsto u_0, cl_2 \mapsto u_0]) \xrightarrow{(cl_1, \mathsf{inc}(k))}_{\mathsf{ET}_{\mathsf{CC}}} (\mathcal{K}_1, [cl_1 \mapsto \_, cl_2 \mapsto u_0]) \xrightarrow{(cl_1, \mathsf{inc}(k))}_{\mathsf{ET}_{\mathsf{CC}}} (\mathcal{K}_2, \_).$$

By looking at the kv-store $\mathcal{K}_2$, we immediately find a cycle in the graph induced by the relations $\mathsf{SO}_{\mathcal{K}_2}, \mathsf{WR}_{\mathcal{K}_2}, \mathsf{WW}_{\mathcal{K}_2}, \mathsf{RW}_{\mathcal{K}_2}\colon t^1_{cl_1} \xrightarrow{\mathsf{RW}} t^1_{cl_2} \xrightarrow{\mathsf{RW}} t^1_{cl_1}$. Following from Theorem 6.1, then which proves that $\mathcal{K}_2$ is not included in $\mathsf{CM}(\mathsf{ET}_{\mathsf{SER}})$, i.e. it is not serialisable.

G.1.2 Robustness of a Single counter under Parallel Snapshot Isolation. Here we show that the single counter library $\mathsf{Counter}(k)$ is robust under any consistency model that guarantees both write conflict detection (formalised by the execution test $\mathsf{ET}_{\mathsf{UA}}$), monotonic reads (formalised by the execution test $\mathsf{ET}_{\mathsf{MR}}$) and read your writes (formalised by the execution test $\mathsf{ET}_{\mathsf{RYW}}$). Because $\mathsf{ET}_{\mathsf{PSI}}$ guarantees all such consistency guarantees, i.e. $\mathsf{CM}(\mathsf{ET}_{\mathsf{PSI}}) \subseteq \mathsf{CM}(\mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}} \cap \mathsf{ET}_{\mathsf{UA}})$, then it also follows that a single counter is robust under Parallel Snapshot Isolation.

PROPOSITION G.1. *Let* $\mathcal{K} \in \mathsf{KVStores}(\mathsf{ET}_{\mathsf{UA}} \cap \mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}}, \mathsf{Counter}(k))$. *Then there exist* $\{t_i\}_{i=1}^{n}$ *and* $\{\mathcal{T}_i\}_{i=0}^{n}$ *such that*

$$\mathcal{K}(k) = ((0, t_0, \mathcal{T}_0 \uplus \{t_1\}) :: \cdots :: (n-1, t_{n-1}, \mathcal{T}_{n-1} \uplus \{t_n\})) :: (n, t_n, \mathcal{T}_n) \tag{7.1}$$

$$\forall i : 0 \le i \le n.\ \mathcal{T}_i \cap \{t_i\}_{i=0}^{n} = \emptyset \tag{7.2}$$

$$\forall t, t', i, j : 0 \le i, j \le n.\ t \xrightarrow{\mathsf{SO}} t' \wedge t \in \{t_i\} \cup \mathcal{T}_i \Rightarrow \left(\begin{array}{c} (t' = t_j \Rightarrow i < j) \wedge \\ (t' \in \mathcal{T}_j \Rightarrow i \le j) \end{array} \right) \tag{7.3}$$

PROOF. It suffices to prove that the properties (7.1),(7.2), (7.3) given in Prop. G.1, are invariant under $(\mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}} \cap \mathsf{ET}_{\mathsf{UA}})$-reductions of the form

$$(\mathcal{K}, \mathcal{U}) \xrightarrow{(cl, \mathsf{inc}(k))}_{\mathsf{ET}_{\mathsf{UA}} \cap \mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}}} (\mathcal{K}', \mathcal{U}') \tag{7.4}$$

$$(\mathcal{K}, \mathcal{U}) \xrightarrow{(cl, \mathsf{read}(k))}_{\mathsf{ET}_{\mathsf{UA}} \cap \mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}}} (\mathcal{K}', \mathcal{U}). \tag{7.5}$$

To this end, we will need the following auxiliary result which holds for any configuration $(\mathcal{K}, \mathcal{U})$ that can be obtained under the execution test $\mathsf{ET}_{\mathsf{RYW}} \cap \mathsf{ET}_{\mathsf{MR}}$:

$$\forall i, j, n, m, cl, k.\ t^n_{cl} \in \{\mathsf{w}(\mathcal{K}(k, i))\} \cup \mathsf{rs}(\mathcal{K}(k, i)) \\ \wedge t^m_{cl} \in \{\mathsf{w}(\mathcal{K}(k, j))\} \cup \mathsf{rs}(\mathcal{K}(k, j)) \wedge m < n \wedge i \in \mathcal{U}(cl)(k) \Rightarrow j \in \mathcal{U}(cl)(k) \tag{7.6}$$

Suppose that there exist two sets $\{t_i\}_{i=1}^{n}$ and $\{\mathcal{T}_i\}_{i=0}^{n}$ such that $(\mathcal{K}, \{t_i\}_{i=1}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies the properties (7.1)-(7.3). We prove that, for transitions of the form (7.4)-(7.5), there exist an index m and two collections $\{t_i\}_{i=1}^{m}, \{\mathcal{T}_i'\}_{i=0}^{m}$ such that $(\mathcal{K}', \{t_i\}_{i=1}^{m}, \{\mathcal{T}_i'\}_{i=0}^{m})$ satisfies the properties (7.1)-(7.3). We consider the two transitions separately.

- Assume that

$$(\mathcal{K}, \mathcal{U}) \xrightarrow{(cl, \mathsf{inc}(k))}_{\mathsf{ET}_{\mathsf{UA}} \cap \mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}}} (\mathcal{K}', \mathcal{U}')$$

for some $cl, \mathcal{K}', \mathcal{U}'$. Let $n+1 = |\mathcal{K}(k)|$. Because of the definition of $\mathsf{ET}_{\mathsf{UA}}$, we must have that $\mathcal{U}(cl) = [k \mapsto \{0, \cdots, n\}]$. Also, because $\mathcal{K}$ satisfies (7.1), we have that $\mathsf{snapshot}(\mathcal{K}, \mathcal{U}(cl))(k) = n$. In particular, $\mathcal{F}(k, \mathcal{U}(cl), \mathsf{inc}(k)) = \{(\mathsf{r}, k, n), (\mathsf{w}, k, n+1)\}$. Thus we have that

$$\mathcal{K}' \in \mathsf{update}(\mathcal{K}, \mathcal{U}(cl), cl, \{(\mathsf{r}, k, n), (\mathsf{w}, k, n+1)\})$$

Let t_{n+1} be the transaction identifier chosen to update $\mathcal{K}$, i.e.

$$\mathcal{K}' = \mathsf{update}(\mathcal{K}, \mathcal{U}(cl), t_{n+1}, \{(\mathsf{r}, k, n), (\mathsf{w}, k, n+1)\})$$

where $t_{n+1} \in \mathsf{nextTid}(\mathcal{K}, cl)$; let also $\mathcal{T}_{n+1} = \emptyset$. Then we have the following:
- $(\mathcal{K}', \{t_i\}_{i=1}^{n+1}, \{\mathcal{T}_i\}_{i=0}^{n+1})$ satisfies Property (7.1). Recall that $(\mathcal{K}, \{t_i\}_{i=1}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies (7.1), i.e.

$$\mathcal{K}(k) = ((0, t_0, \mathcal{T}_0 \uplus \{t_1\}) :: \cdots :: (n-1, t_{n-1}, \mathcal{T}_{n-1} \uplus \{t_n\})) :: (n, t_n, \mathcal{T}_n).$$

It follows that $\mathcal{K}'(k) = ((0, t_0, \mathcal{T}_0 \uplus \{t_1\}) :: \cdots :: (n-1, t_{n-1}, \mathcal{T}_n \uplus \{t_{n+1}\})) :: (n+1, t_{n+1}, \mathcal{T}_{n+1})$, where we recall that $\mathcal{T}_{n+1} = \emptyset$.

- $(\mathcal{K}', \{t_i\}_{i=1}^{n+1}, \{\mathcal{T}_i\}_{i=0}^{n+1})$ satisfies Property (7.2). Let $i = 0, \cdots, n+1$. If $i = n+1$, then $\mathcal{T}_i = \emptyset$, from which $\mathcal{T}_i \cap \{t_j\}_{j=0}^{n+1} = \emptyset$ follows. If $i < n+1$, then because $(\mathcal{K}, \{t_i\}_{i=1}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies Property (7.2), then $\mathcal{T}_i \cap \{t_j\}_{j=0}^{n} = \emptyset$. Finally, because t_{n+1} was chosen to be fresh with respect to the transaction identifiers appearing in $\mathcal{K}$, and $\mathcal{T}_i \subseteq \mathrm{rs}(\mathcal{K}(k,i))$, then we also have that $\mathcal{T}_i \cap \{t_{n+1}\} = \emptyset$.

- $(\mathcal{K}', \{t_i\}_{i=1}^{n+1}, \{\mathcal{T}_i\}_{i=0}^{n+1})$ satisfies Property (7.3). Let t, t' be such that $t \xrightarrow{\mathrm{SO}} t'$. Choose two arbitrary indexes $i, j = 0, \cdots, n+1$, and assume that $t \in \{t_i\} \cup \mathcal{T}_i$. Note that if $i \leq n$, $j \leq n$, then because $(\mathcal{K}, \{t_i\}_{i=1}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies Property (7.3), then if $t' = t_j$ it follows that $i < j$, and if $t' \in \mathcal{T}_j$ it follows that $i \leq j$, as we wanted to prove. If $t \in \{t_{n+1}\} \cup \mathcal{T}_{n+1}$, then it must be $t = t_{n+1}$ because $\mathcal{T}_{n+1} = \emptyset$. Recall that t_{n+1} is the transaction identifier that was used to update $\mathcal{K}$ to $\mathcal{K}'$, i.e. $\mathcal{K}' = \mathrm{update}(\mathcal{K}, \mathcal{U}(cl), t_{n+1}, \_)$. By definition of update, it follows that $t_{n+1} \in \mathrm{nextTid}(\mathcal{K}, cl)$, and because $t_{n+1} \xrightarrow{\mathrm{SO}} t'$, then t' cannot appear in $\mathcal{K}$. In particular, $t' \notin \{t_j\}_{j=0}^{n+1} \cup \bigcup \{\mathcal{T}_j\}_{j=0}^{n+1}$, hence in this case there is nothing to prove. Finally, if $t' \in \{t_{n+1}\} \cup \mathcal{T}_{n+1}$, then it must be the case that $t' = t_{n+1}$. If $t = t_j$, because $t \xrightarrow{\mathrm{SO}} t'$ and $t' = t_{n+1}$, it cannot be $t = t_{n+1}$, hence it must be $i \leq n < n+1$.

- Suppose that

$$(\mathcal{K}, \mathcal{U}) \xrightarrow[\mathsf{ET_{UA}} \cap \mathsf{ET_{MR}} \cap \mathsf{ET_{RYW}}]{(cl, \mathrm{read}(k))} (\mathcal{K}', \mathcal{U}').$$

As in the previous case, we have that $\mathcal{K}' = \mathrm{update}(\mathcal{K}, \mathcal{U}(cl), t, \{(r, k, i)\})$, where $m = \mathrm{snapshot}(\mathcal{K}, \mathcal{U}(cl))(k)$ - in particular, because $(\mathcal{K}, \{t_i\}_{i=1}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies Property (7.1), then it must be the case that $m = \max_<(\mathcal{U}(cl)(k))$ - and $t \in \mathrm{nextTid}(\mathcal{K}, cl)$. For $i = 0, \cdots, n$, let $\mathcal{T}_i' := \mathcal{T}_i$ if $i \neq m$, $\mathcal{T}_i' = \mathcal{T}_i \cup \{t\}$ if $i = m$. Then we have that $(\mathcal{K}', \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i'\}_{i=0}^{n})$ satisfies properties (7.1)-(7.3). Putting all these facts together, we obtain the following:

- $(\mathcal{K}', \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i'\}_{i=0}^{n})$ satisfies Property (7.1). WIthout loss of generality, suppose that $m < n$. Because $(\mathcal{K}, \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies Property (7.1), we have that

$$\mathcal{K}(k) = ((0, t_0, \mathcal{T}_0 \uplus \{t_1\}) :: \cdots :: (m, t_m, \mathcal{T}_m \uplus \{t_{m+1}\}) :: \cdots :: (n-1, t_{n-1}, \mathcal{T}_{n-1} \uplus \{t_n\})) :: (n, t_n, \mathcal{T}_n),$$

and from the definition of update it follows that

$$\begin{aligned}\mathcal{K}(k) \quad &= ((0, t_0, \mathcal{T}_0 \uplus \{t_1\}) :: \cdots :: (m, t_m, \mathcal{T}_m \cup \{t\} \uplus \{t_{m+1}\}) :: \cdots :: (n-1, t_{n-1}, \mathcal{T}_{n-1} \uplus \{t_n\})) :: (n, t_n, \mathcal{T}_n) \\ &= ((0, t_0, \mathcal{T}_0' \uplus \{t_1\}) :: \cdots :: (m, t_m, \mathcal{T}_m' \uplus \{t_{m+1}\}) :: \cdots :: (n-1, t_{n-1}, \mathcal{T}_{n-1} \uplus \{t_n\})) :: (n, t_n, \mathcal{T}_n)\end{aligned}$$

- $(\mathcal{K}', \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i'\}_{i=0}^{n})$ satisfies Property (7.2). Recall that $m = \max_<(\mathcal{U}(cl)(k))$; let $i = 0, \cdots, n$.
  Let again $i = \max_< \mathcal{U}(cl)(k)$. If $i \neq m$, then $\mathcal{T}_i' = \mathcal{T}_i$, and because $(\mathcal{K}, \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies Property (7.2) we have that $\mathcal{T}_i' \cap \{t_i\}_{i=0}^{n} = \emptyset$. If $i = m$, then we have that $\mathcal{T}_i' = \mathcal{T}_m' = \mathcal{T}_m \cup \{t\}$, where we recall that $t \in \mathrm{nextTid}(\mathcal{K}, cl)$. Because $(\mathcal{K}, \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies Property (7.2), we have that $\mathcal{T}_m \cap \{t_i\}_{i=0}^{n} = \emptyset$. Finally, because $t \in \mathrm{nextTid}(\mathcal{K}, cl)$, then it must be the case that for any $i = 0, \cdots, n$, $t \notin \{\mathrm{w}(\mathcal{K}'(k,i))\}_{i=0}^{m} = \{t_i\}_{i=0}^{m}$, where the last equality follows because we have already proved that $(\mathcal{K}', \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i'\}_{i=0}^{n})$ satisfies Property (7.1).

- $(\mathcal{K}', \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i'\}_{i=0}^{n})$ satisfies Property (7.3). Let t', t'' be such that $t' \xrightarrow{\mathrm{SO}} t''$. Suppose also that $t' \in \{t_i\} \cup \mathcal{T}_i'$ for some $i = 0, \cdots, n$. We consider two different cases:
  * $t' = t_i$. Suppose then that $t'' = t_j$ for some $j = 0, \cdots, n$. Because $(\mathcal{K}, \{t_i\}_{i=0}^{n}, \{\mathcal{T}_i\}_{i=0}^{n})$ satisfies Property (7.3), then it must be the case that $i < j$. Otherwise, suppose that

$t'' \in \mathcal{T}_j'$ for some $j = 0, \cdots, n$. If $j \neq m$, then $\mathcal{T}_j' = \mathcal{T}_j$, and because $(\mathcal{K}, \{t_i\}_{i=0}^n, \{\mathcal{T}_i\}_{i=0}^n)$ satisfies Property (7.3), we have that $i \leq j$. Otherwise, $\mathcal{T}_j' = \mathcal{T}_m' = \mathcal{T}_m \cup \{t\}$. Without loss of generality, in this case we can assume that $t'' = t$ (we have already shown that if $t'' \in \mathcal{T}_j$, then it must be $i \leq j$. Recall that $j = m = \max(\mathcal{U}(cl)(k))$, by the Definition of $\mathsf{ET_{UA}}$ it must be the case that $\mathcal{U}(cl) = [k \mapsto \{0, \cdots, j\}]$. It also follows that $t = t_{cl}^p$ for some $p \geq 0$, and because $t' \xrightarrow{\mathsf{SO}} t'' = t$, then $t' = t_{cl}^q$ for some $q < p$. Because of Property (7.6), and because $t' = t_i = \mathsf{w}(\mathcal{K}(k, i))$, then it must be the case the case that $i \in \mathcal{U}(cl)(k)$, hence $i \leq m = j$.

* $t' \in \mathcal{T}_i'$. We need to distinguish the cases $i \neq m$, leading to $\mathcal{T}_i' = \mathcal{T}_i$, or $i = m$, in which case $\mathcal{T}_i' = \mathcal{T}_m' = \mathcal{T}_m \cup \{t\}$. If either $i \neq m$, or $i = m$ and $t \in \mathcal{T}_m$, then we can proceed as in the case $t' = t_i$. Otherwise, suppose that $i = m$ and $t' = t$. Then, because $t' \xrightarrow{\mathsf{SO}} t''$, and $t \in \mathsf{nextTid}(\mathcal{K}, cl)$, it must be the case that $t = t_{cl}^p$ for some $p \geq 0$, and whenever $t_{cl}^. \in k$, then $t_{cl}^. \xrightarrow{\mathsf{SO}} t$. In particular we cannot have that $t'' \in k$, because $t \xrightarrow{\mathsf{SO}} t''$, which concludes the proof.

– $(\mathcal{K}', \mathcal{U}')$ satisfies Property (7.6).

$\square$

COROLLARY G.2. *Given $\mathcal{K} \in \mathsf{KVStores}(\mathsf{ET_{UA}} \cap \mathsf{ET_{MR}} \cap \mathsf{ET_{RYW}}, \mathsf{Counter})$, then $\mathsf{graphOf}(\mathcal{K})$ is acyclic.*

PROOF. Let $\{t_i\}_{i=1}^n, \{\mathcal{T}_i\}_{i=0}^n$ be such that $(\{t_i\}_{i=1}^n, \{\mathcal{T}_i\}_{i=0}^n)$ satisfies properties (7.1)-(7.3). First, we define a partial order between transactions appearing in $\mathcal{K}$ as the smallest relation $\dashrightarrow$ such that for any t, t', t'' and $i, j = 0, \cdots, n$

$$
\begin{aligned}
t \in \mathcal{T}_i &\Rightarrow t_i \dashrightarrow t, \\
i < j &\Rightarrow t_i \dashrightarrow t_j, \\
t \in \mathcal{T}_i \wedge i < j &\Rightarrow t \dashrightarrow t_j \\
t, t' \in \mathcal{T}_i \wedge t \xrightarrow{\mathsf{SO}} t' &\Rightarrow t \dashrightarrow t' \\
t \dashrightarrow t' \rightarrow t'' &\Rightarrow t \dashrightarrow t''
\end{aligned}
$$

It is immediate that if $t \dashrightarrow t'$ then either $t \in \{t_i\} \cup \mathcal{T}_i, t' = \{t_j\} \cup \mathcal{T}_j$ for some i, j such that $i < j$, or $t = t_i, t' \in \mathcal{T}_i$, or $t, t' \in \mathcal{T}_i$ and $t \xrightarrow{\mathsf{SO_{\mathcal{K}}}} t'$. A consequence of this fact, is that $\dashrightarrow$ is irreflexive.

Next, observe that we have the following:

- whenever $t \xrightarrow{\mathsf{WR_{\mathcal{K}}}} t'$, then there exists an index $i = 0, \cdots, n$ such that $t = t_i$, and either $i < n$ and $t' \in \mathcal{T}_i \cup \{t_{i+1}\}$, or $i = n$ and $t' \in \mathcal{T}_i$: by definition, we have that $t \dashrightarrow t'$;
- whenever $t, \xrightarrow{\mathsf{WW_{\mathcal{K}}}} t'$, then there exist two indexes $i, j : 0 \leq i < j \leq n$ such that $t = t_i, t' = t_j$; again, we have that $t \dashrightarrow t'$,
- whenever $t \xrightarrow{\mathsf{RW_{\mathcal{K}}}} t'$, then there exist two indexes $i, j : 0 \leq i < j \leq n$ such that either $t \in \mathcal{T}_i$ and $t' = t_j$, or $t = t_{i+1}, i + 1 < j$ and $t' = t_j$; in both cases, we obtain that $t \dashrightarrow t'$,
- whenever $t \xrightarrow{\mathsf{SO_{\mathcal{K}}}} t'$, then $t \in \{t_i\} \cup \mathcal{T}_i$ for some $i = 0, \cdots, n$, and either $t' = t_j$ for some $i < j$, or $t' \in \mathcal{T}_j$ for some $i \leq j$; it follows that $t \dashrightarrow t'$.

We have proved that $\dashrightarrow$ is an irreflexive relation, and it contains $(\mathsf{SO_{\mathcal{K}}} \cup \mathsf{WR_{\mathcal{K}}} \cup \mathsf{WW_{\mathcal{K}}} \cup \mathsf{RW_{\mathcal{K}}})^+$; because any subset of an irreflexive relation is itself irreflexive, we obtain that $\mathsf{graphOf}(\mathcal{K})$ is acyclic. $\square$

COROLLARY G.3. $\mathsf{KVStores}(\mathsf{ET_{PSI}}, \mathsf{Counter}(k)) \subseteq \mathsf{KVStores}(\mathsf{ET_{SER}}, \mathsf{Counter}(k))$.

Proof. Let $\mathcal{K} \in \mathsf{KVStores}(\mathsf{ET}_{\mathsf{PSI}}, \mathsf{Counter}(k))$. Because $\mathsf{ET}_{\mathsf{PSI}} \supseteq \mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}} \cap \mathsf{ET}_{\mathsf{UA}}$, we have that $\mathcal{K} \in \mathsf{KVStores}(\mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}} \cap \mathsf{ET}_{\mathsf{UA}}, \mathsf{Counter}(k))$. By Cor. G.2 we have that $\mathsf{graphOf}(\mathcal{K})$ is acyclic. We can now employ the construction outlined in [Cerone et al. 2017] to recover an abstract execution $\mathcal{X} = (\mathcal{T}_{\mathcal{K}}, \mathsf{VIS}, \mathsf{AR})$ such that $\mathsf{SO} \subseteq \mathsf{VIS}$ and $\mathsf{AR} \subseteq \mathsf{VIS}$, and $\mathsf{graphOf}(\mathcal{X}) = \mathsf{graphOf}(\mathcal{K})$. Finally, the results from §F.10 establish that, from $\mathcal{X}$ we can recover a $\mathsf{ET}_{\mathsf{SER}}$-trace in $\mathsf{Traces}(\mathsf{ET}_{\mathsf{SER}}, \mathsf{Counter}(k))$ whose last configuration is $(\mathcal{K}', \_)$, and $\mathsf{graphOf}(\mathcal{K}') = \mathsf{graphOf}(\mathcal{X}) = \mathsf{graphOf}(\mathcal{K})$, leading to $\mathcal{K}' = \mathcal{K}$. It follows that $\mathcal{K} \in \mathsf{KVStores}(\mathsf{ET}_{\mathsf{SER}}, \mathsf{Counter}(k))$. □

G.2 Robust against WSI

Definition G.4. A key-value store $\mathcal{K}$ is WSI safe if $\mathcal{K}$ is reachable from executing an program P from an initial configuration Γ_0, i.e. $\mathsf{ET}_{\mathsf{WSI}} \vdash \Gamma_0, \mathsf{P} \rightarrow (\mathcal{K}, \mathcal{U}), \mathsf{P}'$, and $\mathcal{K}$ satisfies the following:

$$\forall t, k, k', i, j.\ (t \in \mathsf{rs}(\mathcal{K}(k, i)) \Rightarrow t \neq \mathsf{w}(\mathcal{K}(k, i))) \Rightarrow t \neq \mathsf{w}(\mathcal{K}(k', j)) \tag{7.7}$$

$$\forall t, k, i.\ t \neq t_0 \exists j.\ t = \mathsf{w}(\mathcal{K}(k, i)) \Rightarrow t \in \mathsf{rs}(\mathcal{K}(k, j)) \tag{7.8}$$

$$\forall t, k, k', i, j.\ t \neq t_0 \exists k.\ t = \mathsf{w}(\mathcal{K}(k, i)) \wedge t \in \mathsf{rs}(\mathcal{K}(k', j)) \Rightarrow t = \mathsf{w}(\mathcal{K}(k', j)) \tag{7.9}$$

Theorem G.5. *If a key-value store $\mathcal{K}$ is* WSI *safe, it is robust against* WSI.

Proof. Assume a kv-store $\mathcal{K}$ that is WSI safe. Given Def. G.4 that $\mathcal{K}$ is reachable under WSI therefore CC and UA since $\mathsf{CC} \cup \mathsf{UA} \subseteq \mathsf{WSI}$, it is easy to derive the following properties:

$$\forall t, t'.\ t \xrightarrow{\mathsf{WR} \cup \mathsf{SO} \cup \mathsf{WW}} t' \Rightarrow t \neq t' \tag{7.10}$$

$$\forall k, i, j.\ t = \mathsf{w}(\mathcal{K}(k, i)) \wedge t \in \mathsf{rs}(\mathcal{K}(k, j)) \Rightarrow i = j + 1 \tag{7.11}$$

Eq. (7.10) To prove the robustness, it is sufficient to prove that the relation $(\mathsf{WW} \cup \mathsf{WR} \cup \mathsf{RW} \cup \mathsf{SO})^+$ is irreflexive, that is, for any transactions t and t':

$$t \xrightarrow{(\mathsf{WW} \cup \mathsf{WR} \cup \mathsf{RW} \cup \mathsf{SO})^+} t' \Rightarrow t \neq t'$$

We prove that by contradiction. Let assume $t = t'$. By Eq. (7.10), it must be the case that the cycle contains RW, which means there exists t_1 to t_n such that

$$t = t_1 \xrightarrow{R^*} t_2 \xrightarrow{\mathsf{RW}} t_3 \xrightarrow{R^*} \cdots \xrightarrow{R^*} t_{n-2} \xrightarrow{\mathsf{RW}} t_{n-1} \xrightarrow{R^*} t_n = t'$$

where $R = \mathsf{WR} \cup \mathsf{SO} \cup \mathsf{WW}$. We replace some edges from the cycle.

- First, let consider transactions t_i such that $t_i \xrightarrow{\mathsf{RW}} t_{i+1}$. This means $t_i \in \mathsf{rs}(\mathcal{K}(k, x))$ and $t_{i+1} = \mathsf{w}(\mathcal{K}(k, y))$ for some key k and two indexes x, y such that $x < y$. There are two possible cases depending on if t_i wrote the key k.
  - if t_i also wrote any key k', by Lemma G.6, it also wrote the key k and we can replace the edge with a WW edge, that is $t_i \xrightarrow{\mathsf{WW}} t_{i+1}$.
  - if t_i did not wrote any key, we leave the edge the same as before.
  After the first step, any RW edge in the cycle must start from a read only transaction.

$$t_i \xrightarrow{\mathsf{RW}} t_{i+1} \Rightarrow \forall k, i.\ t \neq \mathsf{w}(\mathcal{K}(k, i)) \tag{7.12}$$

- Second, let now consider transactions t_i such that $\cdots \xrightarrow{\mathsf{RW}} t_i \xrightarrow{R^*} t_{i+1} \xrightarrow{\mathsf{RW}} \cdots$. Transaction t_i at least wrote a key but t_{i+1} is a read-only transaction, thus $t_i \neq t_{i+1}$. This means $\cdots \xrightarrow{\mathsf{RW}} t_i \xrightarrow{R^+} t_{i+1} \xrightarrow{\mathsf{RW}} \cdots$.
- Last, by Lemma G.7 we replace all the WW with WR^*.

Let $R' = \text{WR} \cup \text{SO}$. Now we have cycle in the following form:

$$t = t_1' \xrightarrow{R'^*} t_2' \xrightarrow{\text{RW}} t_3' \xrightarrow{R'^+} \cdot \xrightarrow{R'^+} t_{m-2}' \xrightarrow{\text{RW}} t_{m-1} \xrightarrow{R'^*} t_m' = t'$$

for some transactions t_1' to t_m' and $m \leq n$. This means $t \xrightarrow{((\text{WR}\cup\text{SO});\text{RW}^?)^*} t'$. Because $\mathcal{K}$ is reachable under WSI and so CP, it must the case that $t \neq t'$, which contradicts with the assumption. Therefore, the relation $(\text{WW} \cup \text{WR} \cup \text{RW} \cup \text{SO})^+$ is irreflexive. □

LEMMA G.6. *If a key-value store $\mathcal{K}$ is WSI safe, then for any transactions t, t'*

$$t \xrightarrow{\text{RW}} t' \wedge \exists k, i.\ t = \text{w}(\mathcal{K}(k, i)) \Rightarrow t \xrightarrow{\text{WW}} t'$$

PROOF. Assume $t \xrightarrow{\text{RW}} t'$, which means $t \in \text{rs}(\mathcal{K}(k, i))$ and $t' = \text{w}(\mathcal{K}(k, j))$ for a key k and two indexes i, j such that $i < j$. Assume the transaction t also wrote some key k'. Since that $\mathcal{K}$ is WSI safe, t must write key k too, i.e. $t = \text{w}(\mathcal{K}(k, z))$ for some index z. Because the $\mathcal{K}$ is reachable under WSI and therefore UA, this means $z = i + 1$. Since that each version can only have one writer, we have $i < z = i + 1 < j$, therefore $t \xrightarrow{\text{WW}} t'$. □

LEMMA G.7. *If a key-value store $\mathcal{K}$ is WSI safe, then for any transactions t, t'*

$$t \xrightarrow{\text{WW}} t' \Rightarrow t \xrightarrow{(\text{WR})^*} t'$$

PROOF. Assume a kv-store $\mathcal{K}$. Assume a key k and two versions of it, i and j respectively with $i < j$. Assume $t = \text{w}(\mathcal{K}(k, i))$ and $t' = \text{w}(\mathcal{K}(k, j))$. We prove $t \xrightarrow{(\text{WR})^*} t'$ by induction on the distance of the two versions.

- Base case: $j - i = 1$. By WSI safe (Def. G.4), t' must also read the key k, that is, $t' \in \text{rs}(\mathcal{K}(k, z))$ for some z. Because the $\mathcal{K}$ is reachable under WSI and therefore UA, this means that if t' read and writes key k, it must read the immediate predecessor. This means $z = i$ and then $t \xrightarrow{\text{WR}} t'$.
- Inductive case: $j - i > 1$. By the Def. G.4, t' must also read the key k, that is, $t' \in \text{rs}(\mathcal{K}(k, z))$ for some z. Because the $\mathcal{K}$ is reachable under WSI and therefore UA, this means if t' read and writes key k, it must read the immediate previous version with respect to the version it wrote. This means $z = j - 1$. Assume the writer of the z-th version is $t'' = \text{w}(\mathcal{K}(k, j - 1))$. We have $t'' \xrightarrow{\text{WR}} t'$. Applying I.H., we get $t \xrightarrow{(\text{WR})^*} t''$. Thus we have $t \xrightarrow{(\text{WR})^*} t'$.

□

G.3 Multiple counters

We define a multi-counter library on a set of keys K as the following:

$$\text{Counters}(\text{K}) \triangleq \bigcup_{k \in \text{K}} \text{Counter}(k)$$

G.3.1 Anomaly of multiple counters under Parallel Snapshot Isolation. Suppose that the kv-store contains only two keys k, k', each of which can be accessed and modified by clients using the code of transactional libraries $\text{Counter}(k), \text{Counter}(k'), \cdots$. We show that in this case it is possible to have the interactions of two client with the kv-store result in a non-serialisable final configuration.

More formally, suppose that $\textsc{Key} = \{k_1, k_2\}$, and let $\textsc{Counter} = \bigcup_{k \in \textsc{Key}} \textsc{Counter}(k)$. Let also

$$
\begin{aligned}
\mathcal{K}_0 &= [k_1 \mapsto (0, t_0, \emptyset), k_2 \mapsto (0, t_0, \emptyset)] \\
\mathcal{K}_1 &= [k_1 \mapsto (0, t_0, \left\{t^1_{cl_1}\right\})::(1, t^1_{cl_1}, \emptyset), k_2 \mapsto (0, t_0, \emptyset)] \\
\mathcal{K}_2 &= [k_1 \mapsto (0, t_0, \left\{t^1_{cl_1}\right\})::(1, t^1_{cl_1}, \emptyset), k_2 \mapsto (0, t_0, \{t_{cl_2}\})::(1, t^1_{cl_2}, \emptyset) \\
\mathcal{K}_3 &= [k_1 \mapsto (0, t_0, \left\{t^1_{cl_1}\right\})::(1, t^1_{cl_1}, \emptyset), k_2 \mapsto (0, t_0, \{t_{cl_2}\})::(1, t^1_{cl_2}, \left\{t^2_{cl_1}\right\}) \\
\mathcal{K}_4 &= [k_1 \mapsto (0, t_0, \left\{t^1_{cl_1}\right\})::(1, t^1_{cl_1}, \left\{t^2_{cl_2}\right\}), k_2 \mapsto (0, t_0, \{t_{cl_2}\})::(1, t^1_{cl_2}, \left\{t^2_{cl_1}\right\}) \\[6pt]
\mathcal{U}_0 &= [cl_1 \mapsto [k_1 \mapsto \{0\}, k_2 \mapsto \{0\}], cl_2 \mapsto [k_1 \mapsto \{0\}, k_2 \mapsto \{0\}]] \\
\mathcal{U}_1 &= [cl_1 \mapsto [k_1 \mapsto \{0, 1\}, k_2 \mapsto \{0\}], cl_2 \mapsto [k_1 \mapsto \{0\}, k_2 \mapsto \{0\}]] \\
\mathcal{U}_2 &= [cl_1 \mapsto [k_1 \mapsto \{0, 1\}, k_2 \mapsto \{0\}], cl_2 \mapsto [k_1 \mapsto \{0\}, k_2 \mapsto \{0, 1\}]] \\
\mathcal{U}_3 &= \mathcal{U}_2 \\
\mathcal{U}_4 &= \mathcal{U}_2
\end{aligned}
$$

Observe that we have the sequence of $\mathsf{ET_{PSI}}$-reductions

$$
(\mathcal{K}_0, \mathcal{U}_0) \xrightarrow{cl_1, \mathsf{inc}(k_1)}_{\mathsf{ET_{PSI}}} (\mathcal{K}_1, \mathcal{U}_1) \xrightarrow{(cl_2, \mathsf{inc}(k_2))}_{\mathsf{ET_{PSI}}}
$$
$$
(\mathcal{K}_2, \mathcal{U}_2) \xrightarrow{(cl_1, \mathsf{read}(k_2))}_{\mathsf{ET_{PSI}}} (\mathcal{K}_3, \mathcal{U}_3) \xrightarrow{(cl_2, \mathsf{read}(k_1))}_{\mathsf{ET_{PSI}}} (\mathcal{K}_4, \mathcal{U}_4)
$$

and therefore $\mathcal{K}_4 \in \mathsf{KVStores}(\mathsf{ET_{PSI}}, \textsc{Counter})$. On the other hand, for $\mathsf{graphOf}(\mathcal{K}_4)$ we have the following cycle, which proves that $\mathcal{K}_4 \notin \mathsf{KVStrores}(\mathsf{ET_{SER}}, \textsc{Counter})$:

$$
t^1_{cl_1} \xrightarrow{\mathsf{SO}_{\mathcal{K}_4}} t^2_{cl_1} \xrightarrow{\mathsf{RW}_{\mathcal{K}_4}} t^1_{cl_2} \xrightarrow{\mathsf{SO}_{\mathcal{K}_4}} t^2_{cl_2} \xrightarrow{\mathsf{RW}_{\mathcal{K}_4}} t^1_{cl_1}.
$$

G.3.2 Robustness under Weak Snapshot Isolation. It is easy to see a multi-counter libraries is WSI safe, therefore robust under WSI.

THEOREM G.8. *Mulit-counter libraries* $\textsc{Counters}(K)$ *are* WSI *safe.*

PROOF. Assume an initial configuration $\Gamma_0 = (\mathcal{K}_0, \mathcal{U}_0$ and some P_0 where $\mathsf{dom}[prog] \subseteq \mathsf{dom}[\mathcal{U}_0]$. Under WSI, we prove any reachable kv-store $\mathcal{K}_i$ satisfies Eqs. (7.7) to (7.9) by induction on the length of trace.

- Base case: $i = 0$. The formulae Eqs. (7.7) to (7.9) trivially hold given $\mathcal{K}_0$ contains only the initial transaction t_0.
- Inductive case: $i > 0$. Let $\Gamma_i = (\mathcal{K}_i, \mathcal{U}_i)$ be the result of running P_0 for i steps. We perform case analysis for the next transaction t_{i+1}.
  - If t_{i+1} reads a key k, i.e. $\mathsf{read}(k)$, it must start from a view that is closed to the relation $((\mathsf{WW} \cup \mathsf{WR} \cup \mathsf{SO}) \cup \mathsf{WR}; \mathsf{RW} \cup \mathsf{SO}; \mathsf{RW})^*$. Let $\mathcal{K}_i(k, j) = (v, t', \mathcal{T}')$ be the latest version included in the view. Thus the new kv-store $\mathcal{K}_{i+1} = \mathcal{K}_i[k \mapsto \mathcal{K}_i(k)[j \mapsto (v, t', \mathcal{T}' \uplus \{t_{i+1}\})]]$. Given t_{i+1} only read the key k without writing, Eqs. (7.7) to (7.9) trivially holds. For other transactions t that are different from t_{i+1}, they must exist in $\mathcal{K}_i$. By I.H., then we prove that $mkvs_{i+1}$ is WSI safe.
  - If t_{i+1} increments a key k, i.e. $\mathsf{inc}(k)$, it means that all versions of k must be included in the view. Let $\mathcal{K}_i(k, j) = (v, t', \mathcal{T}')$ be the latest version of key k. Thus the new kv-store $\mathcal{K}_{i+1} = \mathcal{K}_i[k \mapsto (\mathcal{K}_i(k)[j \mapsto (v, t', \mathcal{T}' \uplus \{t_{i+1}\})])::(v + 1, t_{i+1}, \emptyset)]$. Given t_{i+1} only read and then rewrites the key k, Eqs. (7.7) to (7.9) trivially holds. For other transactions t that are different from t_{i+1}, they must exist in $\mathcal{K}_i$. By I.H., then we prove that $mkvs_{i+1}$ is WSI safe.

□

G.4 Bank Example

Alomari et al. [2008] presented a bank example and claimed that this example is robust against SI. We find out that the bank example is also robust against WSI. The example bases on relational database with three tables, account, saving and checking. The account table maps customer names to customer IDs (Account($\underline{\text{Name}}$, CustomerID)) and saving and checking map customer IDs to saving balances (Saving($\underline{\text{CustomerID}}$, Balance)) and checking balances (Checking($\underline{\text{CustomerID}}$, Balance)) respectively. We ignore the account table since it is an immutable lookup table. We encode the saving and checking tables together as a kv-store. Each customer is represent as an integer n, that is, $(\_, n) \in$ Account($\underline{\text{Name}}$, CustomerID), its checking balance is associated with key $n_s = 2 \times n$ and saving with $n_c = 2 \times n + 1$.

$$n_c \triangleq 2 \times n \qquad\qquad n_s \triangleq 2 \times n + 1 \qquad\qquad \text{Key} \triangleq \bigcup_{n \in \mathbb{N}} \{n_c, n_s\}$$

If n is a customer, then

$$(n, \text{val}(\mathcal{K}(n_s, |\mathcal{K}(n_s)|))) \in \text{Saving}(\underline{\text{CustomerID}}, \text{Balance})$$

and

$$(n, \text{val}(\mathcal{K}(n_s, |\mathcal{K}(n_c)|))) \in \text{Checking}(\underline{\text{CustomerID}}, \text{Balance})$$

To interact with tables, there are five types of transactions. For brevity we assume balances are integers.

$$\texttt{balance(n)} \triangleq [\texttt{x} := [n_c]; \ \texttt{y} := [n_s]; \ \texttt{ret} := \texttt{x} + \texttt{y}]$$

$$\texttt{depositChecking(n, v)} \triangleq [\texttt{if}\,(\texttt{v} \geq 0)\{\ \texttt{x} := [n_c]; \ [n_c] := \texttt{x} + \texttt{v}; \ \}]$$

$$\texttt{transactSaving(n, v)} \triangleq [\texttt{x} := [n_s]; \ \texttt{if}\,(\texttt{v} + \texttt{x} \geq 0)\{\ [n_s] := \texttt{x} + \texttt{v}; \ \}]$$

balance(n) returns customer n total balance. depositChecking(n, v) deposits v to the checking account of customer n, if v is non-negative, otherwise the transaction does nothing. While transactSaving(n, v) allows a consumer n to deposit or withdraw money from the saving account as long as the saving account afterwards is non-negative.

$$\texttt{amalgamate(n, n')} \triangleq \begin{bmatrix} \texttt{x} := [n_s]; \ \texttt{y} := [n_c]; \ \texttt{z} := [n'_c]; \\ [n_s] := 0; \ [n_c] := 0; \ [n'_c] := \texttt{x} + \texttt{y} + \texttt{z}; \end{bmatrix}$$

$$\texttt{writeCheck(n, v)} \triangleq \begin{bmatrix} \texttt{x} := [n_s]; \ \texttt{y} := [n_c]; \\ \texttt{if}\,(\texttt{x} + \texttt{y} < \texttt{v})\{\ [n_c] := \texttt{y} - \texttt{v} - 1; \ \}\texttt{else}\{\ [n_c] := \texttt{y} - \texttt{v}; \ \} \\ [n_s] := \texttt{x}; \end{bmatrix}$$

amalgamate(n, n') represents moving all funds from consumer n to the checking account of customer n'. Last, writeCheck(n, v) updates the checking account of n. If funds, both saving and checking, from n is greater than the v, the transaction deduct v from the checking account of n. If funds are not enough, the transaction further deducts one pounds as penalty. Alomari et al. [2008] argued that, to make this example robust against SI, writeCheck(n, v) must be strengthened by writing back the balance to the saving account (the last line, $[n_s] := \texttt{x}$), even thought the saving balance is unchanged. The bank bank libraries are defined by

$$\text{Bank} \triangleq \begin{cases} \texttt{balance(n)}, \texttt{depositChecking(n, v)}, \texttt{amalgamate(n, n')}, \\ \texttt{writeCheck(n, v)}, \texttt{writeCheck(n, v)} \end{cases} [n, n' \in \mathbb{N} \wedge v \in \mathbb{Z}]$$

THEOREM G.9. *The bank libraries* Bank *are WSI safe.*

PROOF. Assume an initial configuration $\Gamma_0 = (\mathcal{K}_0, \mathcal{U}_0)$ and some P_0 where $\mathrm{dom}(\mathsf{P}) \subseteq \mathrm{dom}(\mathcal{U}_0)$. Under WSI, we prove any reachable kv-store $\mathcal{K}_i$ satisfies Eqs. (7.7) to (7.9) by induction on the length of trace.

- Base case: $i = 0$. The formulae Eqs. (7.7) to (7.9) trivially hold given $\mathcal{K}_0$ contains only the initial transaction t_0.
- Inductive case: $i > 0$. Let $\Gamma_i = (\mathcal{K}_i, \mathcal{U}_i)$ be the result of running P_0 for i steps. We perform case analysis for the next transaction t_{i+1}.
  - If t_{i+1} is $\mathtt{balance(n)}$, the only possible fingerprint is $\{(\mathsf{r}, n_c, v_c), (\mathsf{r}, n_s, v_s)\}$ for some values v_c and v_s. Since it is a read-only transaction, Eqs. (7.7) to (7.9) trivially hold.
  - If t_{i+1} is $\mathtt{depositChecking(n, v)}$, in the cases of $v < 0$, the fingerprint is empty and Eqs. (7.7) to (7.9) trivially hold. However, in the case of $v \geq 0$, the fingerprint is $\{(\mathsf{r}, n_c, v_c), (\mathsf{w}, n_c, v_c + v$ Because it read and wrote only one key, n_c, Eqs. (7.7) to (7.9) hold.
  - If t_{i+1} is $\mathtt{transactSaving(n, v)}$, there are two cases: either a read-only fingerprint $\{(\mathsf{r}, n_s, v_s)\}$ when saving account has insufficient funds, or a read and write on key n_s, that is $\{(\mathsf{r}, n_s, v_s), (\mathsf{w}, n_s, v_s +$ For both cases it is easy to see Eqs. (7.7) to (7.9) hold.
  - If t_{i+1} is $\mathtt{amalgamate(n, n')}$, the fingerprint is

  $$\big\{(\mathsf{r}, n_s, v_s), (\mathsf{w}, n_s, 0), (\mathsf{r}, n_c, v_c), (\mathsf{w}, n_c, 0), (\mathsf{r}, n'_c, v'_c), (\mathsf{w}, n'_c, v'_c + v_s + v_c)\big\}$$

  Because the transaction always read and then wrote keys it touched, namely n_s, n_c and n'_c, Eqs. (7.7) to (7.9) hold.
  - Last, if t_{i+1} is $\mathtt{writeCheck(n, v)}$, the fingerprint is

  $$\big\{(\mathsf{r}, n_s, v_s), (\mathsf{w}, n_s, v_s), (\mathsf{r}, n_c, v_c), (\mathsf{w}, n_c, v'_c)\big\}$$

  where v'_c can be either $v_c - v$ or $v_c - v - 1$. Similar to $\mathtt{amalgamate(n, n')}$, the transaction always read and then wrote keys it touched, so Eqs. (7.7) to (7.9) hold.

$\square$

H VERIFICATION OF IMPLEMENTATIONS

We verify two protocols, COPS and Clock-SI, that the former is a full replicated implementation for causal consistency and the latter is a shard implementation for snapshot isolation.

H.1 COPS

H.1.1 Code.

Structure. COPS is a fully replicated database protocol for causal consistency. There are two versions, that the simplified version only supports either a single read or a single write per transaction, and the full version supports either multiple reads or a single write pe transactions. Here we verify the full version.

The overall database is modelled as a key-value store. Each key, instead of a single value, is associated with a list of versions, consisting of value, version number `VersionNo` and dependencies `Dep`. COPS relies on the version number to resolve conflict, that is, the write with greater version number wins. Version number composite by time (higher bits) and replica identifier (lower bits). Since the replica identifiers are full ordered, therefore version numbers are full ordered. The dependencies is a set of versions (pairs of keys and versions numbers) that the current version depends on.

```
1  VersionNo :: LocalTime + ID
2  Dep :: Set(Key, VersionNo)
3  KV :: Key -> List(Val, VersionNo, Dep)
```
Code 1. COPS Structure

Under the hood, there are many replicas, where each replica has a unique identifier and contains the full key-value store yet might be out of data. Replica also tracks its local time.

```
1  Replicas :: ID -> (KV, LocalTime)
```
Code 2. COPS Replicas

To track session order, Client interact with a replica via certain API. To call those API, client has the responsibility to provide context `Ctx` which contains versions that has been read from or written to the replica from the same client.

```
1  Ctx :: {
2      readers : List(Key, VersionNo, Dep)
3      writers : List(Key, VersionNo, Dep)
4  }
```
Code 3. Client context

Write. The client call `put` to commit a new write to a key k with value v with the context `ctx`. It extracts the dependencies from the context, by unioning all the versions inside the context, then calls the replica API `ver = put_after(k,v,deps,dep_to_nearest(deps),dep)`, which require to provide the dependencies so that the replica can check whether they are exists. Note `dep_to_nearest(deps)` is for performance by providing the latest versions. The replica returns a version number allocated for the new version for key k, and then the version is added into the context.

```
1  Ctx put(k,v,ctx) {
2      // add up all the read and write dependency
3      deps = ctx_to_dep(ctx);
4      // call replica API
```

```
5        ver = put_after(k,v,deps,dep_to_nearest(deps));
6        // update context
7        ctx.writers += (k,ver,deps);
8        return ctx;
9    }
10
11   Dep ctx_to_dep(ctx) {
12       return { (k,ver) | (k,ver,_) ∈ ctx.readers ∨ (k,ver,_) ∈ ctx.writers }
13   }
14
15   Dep dep_to_nearest(deps) {
16       return { (k,ver) | ∀k',ver',deps'. (k',ver',deps') ∈ deps ⇒ (k,ver) ∉ deps'
             };
17   }
```

Code 4. Client API for write

The put_after waits until all the versions contained in nearest exists, consequently all the versions contained in deps exists. The replica increments the local time and insert the new version with version numbers time ++ id (local time concatenating replica identifier) to the key k. At this point, replica returns the new version number to client and later on it will broad case to other replica[6].

```
1    VersionNo put_after(k,v,deps,nearest,vers){
2        for (k,ver) in nearest { wait until (_,ver,_) ∈ kv(k); }
3
4        time = inc(local_time);
5
6        // appending local kv with a new version.
7        list_isnert(kv[k], (v, (time ++ id), deps) );
8
9        asyn_brordcase(k, v, (time ++ id), deps);
10       return (local_time + id);
11   }
```

Code 5. Replica API for write

When a replica receives a update message, it checks the existence of versions included in the dependencies and then adds the new version to the replica. Last, the replica updates the local time if the new version's time is greater than the local time.

```
1    on_receive(k,v,ver,deps) {
2        for (k',ver') in deps { wait until (_,ver',_) ∈ kv(k'); }
3
4        list_isnert(kv[k],(v,ver,deps));
5        (remote_local_time + id) = ver;
6        local_time = max(remote_local_time, local_time);
7    }
```

Code 6. Receive update message

---

[6]It uses message queue for broad-casting

Read. To read multiple keys ks in a transaction, client calls the `get_trans(ks,ctx)`. Note that between two reads for different keys, the replica might be interleave and schedule other transactions. The challenge here is to ensure all the values are consistent, i.e. overall they satisfy the causal consistency. That is, if the transaction fetches a version v for key k, and this version v depends on anther version v' for another key k', then the transaction should at least fetches v' for the key k', or any later version for the key k'.

The algorithm, in the first phase, optimistically reads the current latest version for each key from the replica via the replica API `rst[k]=get_by_version(k,LATEST)`. In the second phase, it computes the maximum version `ccv[k]` from any dependencies `rst[k].deps` read in the first phase. Such `ccv[k]` is the minimum version that should be fetched. Therefore at the end, it only needs to re-fetched the specifically version `ccv[k]`, if the version fetched in the first phase is older than `ccv[k]`.

```
1   List(Val) get_trans(ks,ctx) {
2       // only guarantee to read up-to-date value
3       // the moment reading the individual key
4       for k in ks { rst[k] = get_by_version(k,LATEST); }
5
6       for k in ks {
7           ccv[k] = max (ccv[k],rst[k].ver);
8           for dep in rst[k].deps
9               if ( dep.key ∈ ks ) ccv[k] = max (ccv[dep.key],dep.vers);
10      }
11
12      for k in ks
13          if ( ccv[k] > rst[k].vers ) rst[k] = get_by_version(k,ccv[k]);
14
15      // update the ctx
16      for (k,ver,deps) in rst { ctx.readers += (k,ver,deps); }
17
18      return to_vals(rst);
19  }
```

Code 7. Reads

The client API `get_by_version(k,ver)` returns the version ver for key k.

```
1   (Val,Version,Dep) get_by_version(k,ver) {
2       if (ver  = LATEST){ ver = max(kv[k].vers); }
3
4       wait until (_,ver,_) ∈ kv(k);
5
6       let (val,ver,deps) from kv[k];
7       return (val,ver,deps);
8   }
```

Code 8. Replica API for read

H.1.2 Verification.

Semantics the code. Let $r \in$ Repls denotes the set of totally ordered replicates. Each replicate can have multiple clients, and each clients can commit a sequence of either read-only transitions or single-write transactions. To model these, we annotate the transaction identifier with replicate r, client cl, local time of the replicate n and read-only transactions counter n', i.e. $t_{(r,cl)}^{(n,r,n')}$. Note that

the (n, r, n') can be treated as a single number that n are the higher bits, r the middle bits and n' the lower bits. For a new single-write transaction, it is allocated with a transaction identifier with larger local time, and for a read-only transactions, it is allocated with a transaction identifier with larger read-only counter. There is a total order among transitions from the same replica and from the same client.

To model the dependencies of each version, We extend version from Def. 3.1 with the set of all versions it dependencies on, $dp \in \mathcal{P}(\text{KEY} \times \text{TRANSID})$. The function $\text{deps}v$ denotes the dependencies set of the version. We use $\bar{\mathcal{K}}$ for key-value store whose versions contain the dependencies.

We use view to model the client context, that is, a version is included in a context if and only if such version is in the view of the client. We also use view to model a replica state, that is, if a replica contains a version if and only if such version in the view of the replica. For readability, we annotate view with either a replica, u_r, or a client, u_{cl}. The view environment is extended with replicas and their views, $\bar{\mathcal{U}} : (\text{REPLS} \times \text{CLIENT}) \xrightarrow{fin} \text{VIEWS}$. We give the following semantics to capture the behaviours of the code.

Write. For purpose of verification, we eliminate code for performance, and put the client API and replica API in the same function (Code 9).

```
1   // mixing the client API and system API
2   put(repl,k,v,ctx) {
3
4       // Dependency for previous reads and writes
5       deps = ctx_to_dep(ctx);
6
7       // increase local time and appending local kv with a new version.
8       inc(repl.local_time);
9       list_isnert(repl.kv[k],(v, (local_time + id), deps));
10
11      // update dependency for writes
12      ctx.writers += (k,(local_time + id),deps);
13
14      // broad case
15      asyn_brordcase(k, v, (time ++ id), deps);
16      return (local_time + id);
17  }
```

Code 9. put

The following is the rule corresponds to Code 9:

PUT

$$
\frac{
\begin{array}{c}
(s, , \emptyset), [k] := \mathsf{x} \to (s', \_, \{(\mathsf{w}, k, v)\}), \mathsf{skip} \\
dp = \left\{ (k', t) \mid \exists i.\ i \in u_{cl}(k') \wedge t = \mathsf{w}(\bar{\mathcal{K}}(k', i)) \vee (k', t) \in \mathsf{deps}(\bar{\mathcal{K}}(k', i)) \right\} \text{ ---> Code 9, line 5} \\
t = \min \left\{ t^{(n', r, 0)}_{(r, cl)} \mid \forall k', i \in u_r(k'), n.\ t^{(n, \_, \_)}_{(\_, \_)} = \mathsf{w}(\bar{\mathcal{K}}(k', i)) \Rightarrow n' > n \right\} \text{ ---> Code 9, line 8} \\
\bar{\mathcal{K}}' = \bar{\mathcal{K}}\left[k \mapsto \bar{\mathcal{K}}(k)::(k, t, \emptyset, dp) \right] \text{ ---> Code 9, line 9} \\
u_r' = u_r\left[k \mapsto u_r(k) \uplus \left\{ |\bar{\mathcal{K}}'(k)| - 1 \right\} \right] \text{ ---> Code 9, line 9} \\
u_{cl}' = u_{cl}\left[k \mapsto u_r(k) \uplus \left\{ |\bar{\mathcal{K}}'(k)| - 1 \right\} \right] \text{ ---> Code 9, line 12}
\end{array}
}{
r, cl \vdash \bar{\mathcal{K}}, u_r, u_{cl}, s, [[k] := \mathsf{x};] \xrightarrow{u_{cl}, \{(\mathsf{w}, k, v)\}} \bar{\mathcal{K}}', u_r', u_{cl}', s', \mathsf{skip}
}
$$

The first premiss is to execute the transaction locally (Fig. 1). Since there is only a write, the snapshot can be any snapshot. The second line computes the dependency set for the new write

operation, by collecting all the writers of versions included in the view u_{cl}. The third line simulates the increment of local time. Even thought we do not directly track the local time of a replica, yet the local time can compute as the maximum time contained in the replica's view u_r. The forth and fifth simulates the updates of the replica's key-value store, and the last premiss simulates the update of client context.

Read. The following is a simplified algorithm by directly taking a list of versions ccv satisfies causal consistency constraint, i.e. the second phase of Code 7, and then read the versions indicated by ccv.

```
1  List(Val) get_trans(ks,ctx) {
2      // only guarantee to read up-to-date value
3      // the moment reading the individual key
4      for k in ks { rst[k] = get_by_version(k,LATEST); }
5
6      for k in ks {
7          ccv[k] = max (ccv[k],rst[k].ver);
8          for dep in rst[k].deps
9              if ( dep.key ∈ ks ) ccv[k] = max (ccv[dep.key],dep.vers);
10     }
11     for k in ks
12         if ( ccv[k] > rst[k].vers ) rst[k] = get_by_version(k,ccv[k]);
13
14     // update the ctx
15     for (k,ver,deps) in rst { ctx.readers += (k,ver,deps); }
16
17     return to_vals(rst);
18 }
```

Code 10. get_trans

The following is the rule for read-only transaction:

GetTrans
$$T = x_1 := [k_1]; \ldots; x_j := [k_j]; \qquad u_0 = u_{cl}$$
$$\text{for } i \text{ in } \{1, \ldots, j\}$$
$$m_i \in u_r(k_i)$$
$$u_i' = u_{i-1}[k_i \mapsto u_{i-1}(k_i) \cup \{m_i\}] \quad \text{---> Code 10, line 4}$$
$$u_i = \lambda k.\, u_i'(k) \cup \left\{ x \mid (k, \mathsf{w}(\bar{\mathcal{K}}(k, x))) \in \mathsf{deps}(\bar{\mathcal{K}}(k_i, m_i)) \right\}$$
$$\text{---> Code 10, lines 7 to 12 and 15}$$

$$u_{cl}' = u_j \qquad (s, \mathsf{getMax}(\bar{\mathcal{K}}, u_{cl}'), \emptyset), T \to (s', \_, \mathcal{F}), \mathsf{skip} \quad \text{---> Code 10, line 17}$$
$$\frac{t_{(r,cl)}^{(n',r,n)} = \max \left\{ t_{(r,cl)}^{(z',r,z)} \mid t_{(r,cl)}^{(z',r,z)} \in \bar{\mathcal{K}} \right\} \qquad \bar{\mathcal{K}}' = \mathsf{update}\left(\bar{\mathcal{K}}, u_{cl}', \mathcal{F}, t_{(r,cl)}^{(n',r,n+1)}\right)}{r, cl \vdash \bar{\mathcal{K}}, u_r, u_{cl}, s, \left[x_1 := [k_1]; \ldots; x_j := [k_j]; \right] \xrightarrow{u_{cl}', \mathcal{F}} \bar{\mathcal{K}}', u_r, u_{cl}', s', \mathsf{skip}}$$

ClientCommit
$$\frac{r, cl \vdash \bar{\mathcal{K}}, \bar{\mathcal{U}}(r), \bar{\mathcal{U}}(cl), s, \mathsf{P}(cl) \xrightarrow{u_{cl}'', \mathcal{F}} \bar{\mathcal{K}}', u_r', u_{cl}', s', \mathsf{C}'}{\bar{\mathcal{K}}, \bar{\mathcal{U}}, \mathcal{E}, \mathsf{P} \xrightarrow{u_{cl}'', \mathcal{F}} \bar{\mathcal{K}}', \bar{\mathcal{U}}\left[r \mapsto u_r' \right]\left[cl \mapsto u_{cl}' \right], \mathcal{E}[cl \mapsto s'], \mathsf{P}[cl \mapsto \mathsf{C}']}$$

The for-loop in the premiss picks a version for key k_i from u_r and adds to client view, which corresponds to the initial $\mathsf{ccv}(k_i)$. Then it add all the new dependencies $\mathsf{deps}(\bar{\mathcal{K}}(k_i, m_i))$ to the view that yields u_i. The view u_i corresponds to $\mathsf{ccv}(k_i)$ after the update against the dependencies. Given the view u_{cl}, the client fetches the version with the maximum writer it can observed for each key, which is computed by getMax function. It is different from snapshot because snapshot fetches the latest version with respect to the position in the list, but later one we will prove getMax and snapshot are equivalent.

$$\mathsf{getMax}(\bar{\mathcal{K}}, u_{cl}) \triangleq \lambda k. \left(\max \left\{ (v, t, \mathcal{T}, dp) \,\middle|\, \exists i. \ (v, t, \mathcal{T}, dp) = \bar{\mathcal{K}}(k, i) \wedge i \in u_{cl}(k) \right\} \right)|_1$$

$$(v, t_{(r, cl)}^{(n, r, n')}, \mathcal{T}, dp) > (v', t_{(r', cl')}^{(n'', r', n''')}, \mathcal{T}', dp') \overset{\text{def}}{\Leftrightarrow} (n, r, n') > (n'', r, n''')$$

The rest part are trivial be picking a new transaction identifier with larger read counter and committing to key-value store.

Last the rule for receiving a update. A replica updates its local state only if all the dependencies has been receive as shown in Code 6.

$$\frac{\begin{array}{c} \text{SYNC} \\ u_r = \bar{\mathcal{U}}(r)\big[k \mapsto \bar{\mathcal{U}}(r)(k) \uplus i\big] \\ \big(\forall k', m, v. \ v = \bar{\mathcal{K}}(k, i) \wedge (k', \mathsf{w}(\bar{\mathcal{K}}(k', m))) \in v|_4 \Rightarrow m \in u_r'(k')\big) \end{array}}{\bar{\mathcal{K}}, \bar{\mathcal{U}}, \mathcal{E}, \mathsf{P} \rightarrow \bar{\mathcal{K}}, \bar{\mathcal{U}}[r \mapsto u_r], \mathcal{E}, \mathsf{P}}$$

The premiss says, the first line, the replica r receive a new version i for key k, and, the second line, only if all the dependencies of the new update already in the replica's view.

COPS key-value store. To verify the algorithm, we prove that for any COPS trace produced by the algorithm, there exists a corresponding causal consistency trace. First we prove that the key-value stores from COPS trace satisfy Def. 3.1.

THEOREM H.1 (WELL-FORMED COPS KEY-VALUE). *Let ignore the dependencies of versions from $\bar{\mathcal{K}}$. Given the initial key-value store $\bar{\mathcal{K}}_0$, initial views $\bar{\mathcal{U}}_0$ and some programs P_0, for any $\bar{\mathcal{K}}_i$ and $\bar{\mathcal{U}}_i$ such that:*

$$\bar{\mathcal{K}}_0, \bar{\mathcal{U}}_0, \mathcal{E}_0, \mathsf{P}_0 \rightarrow^* \bar{\mathcal{K}}_i, \bar{\mathcal{U}}_i, \mathcal{E}_i, \mathsf{P}_i$$

The key-value store $\bar{\mathcal{K}}_i$ satisfies Def. 3.1 and any replica or client view u from $\bar{\mathcal{U}}_i$ is a valid view of the key-value store, i.e. $u \in \text{VIEWS}(\bar{\mathcal{K}}_i)$.

PROOF. We need to prove the $\bar{\mathcal{K}}_i$ satisfies the well-formed conditions, and any view $u_i \text{VIEWS}(\bar{\mathcal{K}}_i)$. We prove it by introduction on the length i.

- Base case: $i = 0$. It holds trivially since each key only has the initial version $(v_0, t_0, \emptyset, \emptyset)$. Since there is only the initial version for each key, it is easy to see that any view u_0 satisfying the well-formed conditions in Def. 3.2.
- Inductive case: $i + 1$. We perform case analysis on the possible $(i + 1)$-*th* step:
  - PUT Assume the client cl of a replica r commits a single-write transaction t that installs a new version for key k. By the premiss of PUT, the new transaction identifier $t = t_{(r, cl)}^{(n', r, 0)}$ where for some n' that is greater than any n from any writers $t_{(\_, \_)}^{(n, \_, \_)}$ that are observable by the replica r. Since the new transaction $t = t_{(r, cl)}^{(n', r, 0)}$ is a single-write transaction which is always installed at the end of the list associated to k, it is sufficient to prove the following:

$$\forall j. \ 0 \leq j < |\bar{\mathcal{K}}_i(k)| \Rightarrow \mathsf{w}(\bar{\mathcal{K}}_i(k, j)) \neq t \tag{8.1}$$

$$\forall j, n. \ t_{(r, cl)}^{(n, r, \_)} \in \left\{ \mathsf{w}(\bar{\mathcal{K}}_i(k, j)) \right\} \cup \mathsf{rs}(\bar{\mathcal{K}}_i(k, j)) \Rightarrow n < n' \tag{8.2}$$

Lemma H.2 implies Eq. (8.1) and Eq. (8.2). Thus the new key-value store $\bar{\mathcal{K}}_{i+1}$ satisfies the well-formed conditions. Now let consider the views, especially the views of the replica u'_r and the client u'_{cl}. Because views different replicas or clients remain unchanged, by I.H. they satisfy $u' \in \text{Views}(\bar{\mathcal{K}}_{i+1})$. The new view for replica $u'_r = u_r[k \mapsto |\bar{\mathcal{K}}_{i+1}(k)| - 1]$ where u_r is the replica's view before updating and the writer of the last version of k is t. Because t is a single-write transaction, so the new view u'_r still satisfies the atomic read. For similar reason, the new view for client u'_{cl} till satisfies atomic read. Therefore we have $u'_r, u'_{cl} \in \text{Views}(\bar{\mathcal{K}}_{i+1})$.

– GetTrans Assume the client cl of a replica r commits a read-only transaction t. Since it is a read-only transaction, it suffice to prove the following:

$$\forall k, j. \ 0 \le j < |\bar{\mathcal{K}}_i(k)| \Rightarrow t \notin \text{rs}(\bar{\mathcal{K}}_i(k, j)) \tag{8.3}$$

Lemma H.2 implies Eq. (8.3). Thus the new key-value store $\bar{\mathcal{K}}_{i+1}$ satisfies the well-formed conditions. Now let consider the views. Since only the client view has changed, it is sufficient to prove that $u'_{cl} \in \text{Views}(\bar{\mathcal{K}}_{i+1})$, where u'_{cl} is the new client view. By Lemma H.3, the new view u'_{cl} satisfies the atomic read constraint. Therefore $u'_{cl} \in \text{Views}(\bar{\mathcal{K}}_{i+1})$.

□

Lemma H.2 (Monotonic local time). *Given a reduction step such that:*

$$\bar{\mathcal{K}}, u_r, u_{cl}, s, \mathsf{C} {\hookrightarrow}^* \bar{\mathcal{K}}', u'_r, u'_{cl}, s', \mathsf{C}'$$

let $t^{(n,r,n')}_{(r,cl)}$ be the new transaction, i.e. $t^{(n,r,n')}_{(r,cl)} \in \bar{\mathcal{K}}' \wedge t^{(n,r,n')}_{(r,cl)} \notin \bar{\mathcal{K}}$. It implies the new transaction is greater than any transaction committed by the same client view u_{cl}, i.e.

$$t^{(n'',r,n''')}_{(r,cl)} \in \bar{\mathcal{K}} \Rightarrow (n, r, n') > (n'', r, n''')$$

Proof. We perform case analysis.

- Put By the premiss of the rule, the new transaction is in the form $t^{(n,r,0)}_{(r,cl)}$, and for any existing transaction $t^{(n',r,\_)}_{(r,cl)}$,

$$t^{(n',r,\_)}_{(r,cl)} \in \bar{\mathcal{K}} \Rightarrow n > n'$$

- GetTrans Let $t^{(n,r,n')}_{(r,cl)}$ be the new transaction. By the premiss we have that for any existing transaction $t^{(n'',r,n''')}_{(r,cl)}$,

$$t^{(n'',r,n''')}_{(r,cl)} \in \bar{\mathcal{K}} \Rightarrow n = n'' \wedge n' > n'''$$

□

Lemma H.3 (Unique writer). *Each version has a unique writer.*

Proof. Because a transaction can write to at most one key, and a client at least observes its own writes (the premiss of Put) we have the proof. □

COPS normal trace. We define COPS semi-normal trace then normal trace. Later, we prove for any COPS trace, there exists an equivalent normal trace.

The dependency of a transaction $\text{deps}(\bar{\mathcal{K}}, t)$ is defined as:

- if t is a single-write transaction:

$$\text{deps}(\bar{\mathcal{K}}, t) \triangleq dp \text{ where } \exists k, j. \text{ lastConf}(\tau)|_1 (k, j) = (\_, t, \_, dp)$$

- if t is a read-only transaction:

$$\text{deps}(\bar{\mathcal{K}}, t) \triangleq \bigcup_{\{v \mid \exists k, j. \ v = \mathcal{K}(k, j) \wedge t \in \text{rs}(v)\}} \text{deps}(v)$$

Given a trace τ, the function $\text{maxVersion}(\tau, t)$ returns the version that the transaction t depends on and is the last one that appears in the trace:

$$\text{maxVersion}(\tau, t) \triangleq t_i \quad \text{the } t_i \text{ is the last one appears in } \tau \text{ such that}$$
$$(k', t_i) \in \text{deps}(\text{lastConf}(\tau)|_1, t) \wedge t_i \notin \bar{\mathcal{K}}_i \wedge t_i \in \bar{\mathcal{K}}_{i+1}$$

Give two COPS's traces τ and τ', $\bar{\mathcal{K}}$ being the final state of τ and $\bar{\mathcal{K}}'$ for τ', if the two traces are equivalent, if and only if,

$$\forall t, \mathcal{F}, dp. \ t \in \bar{\mathcal{K}} \wedge \mathcal{F} = \text{RWset}(\bar{\mathcal{K}}, t) \wedge dp = \text{deps}(\bar{\mathcal{K}}, t)$$
$$\iff t \in \bar{\mathcal{K}}' \wedge \mathcal{F} = \text{RWset}(\bar{\mathcal{K}}', t) \wedge dp = \text{deps}(\bar{\mathcal{K}}', t)$$

where $\text{RWset}(\bar{\mathcal{K}}, t)$ is the read-write set of t. Note that RWset is well-defined by Theorem H.1.

A COPS's semi-normal trace is a trace τ if it is in the form that read-only transactions t_{rd} directly follows its $\text{maxVersion}(\tau, t_{\text{rd}})$ or another read-only transaction t'_{rd} such that $\text{maxVersion}(\tau, t_{\text{rd}}) = \text{maxVersion}\left(\tau, t'_{\text{rd}}\right)$.

COROLLARY H.4. *For any COPS's trace τ, there exists a equivalent semi-normal trace τ' such that* $\text{lastConf}(\tau) = \text{lastConf}(\tau')$.

PROOF. It is easy to prove by induction on the numbers of read-only transactions that are not in the wanted position. Let take the first one for those read-only transactions t_{rd} who does not follows its $\text{maxVersion}(\tau, t_i)$. It is safe to move the reduction step to the right in the trace, until it directly follows its $\text{maxVersion}(\tau, t_{\text{rd}})$ or another read-only transaction t'_{rd} such that $\text{maxVersion}(\tau, t_{\text{rd}}) = \text{maxVersion}\left(\tau, t'_{\text{rd}}\right)$. Assume it is i-*th* step, and assume $(i-1)$-*th* is not $\text{maxVersion}(\tau, t_{\text{rd}})$ nor another read-only transaction t'_{rd} such that $\text{maxVersion}(\tau, t_{\text{rd}}) = \text{maxVersion}\left(\tau, t'_{\text{rd}}\right)$. It means the $(i-1)$-*th* step must be a write that t_{rd} does not depend on. Because if $(i-1)$-*th* step is a read only transaction t''_{rd} it must be the case $\text{maxVersion}(\tau, t_{\text{rd}}) = \text{maxVersion}\left(\tau, t''_{\text{rd}}\right)$, otherwise t_{rd} is not the first read-only transaction that is not in the wanted place. □

A COPS's normal trace is a trace τ if it is semi-normal trace and the single-write transactions step appears in the trace in the order of the writers of those transactions.

THEOREM H.5 (NORMAL TRACE). *For any COPS's trace τ, there exists an equivalent normal trace τ'.*

PROOF. By Cor. H.4, let consider a semi-normal trace τ. It is easy to prove by induction on the single-write transactions that are out of order. Let take the first single-write transaction $t = t^{(n, r, n')}_{(r, cl)}$ that is out of order. Suppose it is the i-*step*. Assume t write to key k with value v. Assume the

preview write-only transaction $t' = t_{(r', cl')}^{(n'', r', n''')}$ and it follows $(n, r, n') < (n'', r', n''')$. The intuition here is to swap the reduction steps of these two transactions and those read-only transactions following them. We assume t and t' write to the same key, otherwise it is safe to swap. For any read-only transactions $t_{rd} \in \mathcal{T}_{rd}$ following t, we knows $t' \notin \mathsf{deps}(\bar{\mathcal{K}}_i, t_{rd})$, otherwise it violate $\mathsf{maxVersion}(\tau, t_{rd}) = t$. Thus it is safe to swap the reduction steps and alter views that includes either or both versions. Note that if a view u includes both t and t', and if t' also write to the same key k with value v', the u will fetch the value from t': $\mathsf{getMax}(\bar{\mathcal{K}}_i, u)(k) = v'$ □

COROLLARY H.6. *For the PUT and GETTRANS, it is safe to replace the function* getMax *with the function* snapshot.

PROOF. By Theorem H.5, each trace τ has an equivalent normal trace τ'. Assume i-*th* step in the trace τ'. Assume the key-value store before is $\bar{\mathcal{K}}_i$. For any versions m and j from the same key in the τ':

$$0 < m < j < |\bar{\mathcal{K}}_i(k)| \Rightarrow \mathsf{w}(\bar{\mathcal{K}}_i(k, m)) < \mathsf{w}(\bar{\mathcal{K}}_i(k, j))$$

thus it is safe to use snapshot. □

COPS normal trace to ET_{CC} **trace.** Given Theorems H.1 and H.5 and Cor. H.6, for any COPS's trace τ, there exists a trace τ' that is a normal trace and satisfies $\mathsf{ET}_\top$. Then by the following theorem (Theorem H.7), we proof the COPS trace satisfies ET_{CC}.

THEOREM H.7 (COPS SATISFYING CAUSAL CONSISTENCY). *Given a trace starting from the initial key-value store* $\bar{\mathcal{K}}_0$, *initial views* $\bar{\mathcal{U}}_0$ *and some programs* P_0, *for any* $\bar{\mathcal{K}}_i$ *and* $\bar{\mathcal{U}}_i$ *such that:*

$$\bar{\mathcal{K}}_i, \bar{\mathcal{U}}_i, \mathcal{E}_i, \mathsf{P}_i \xrightarrow{u_{cl}, \mathcal{F}} \bar{\mathcal{K}}_{i+1}, \bar{\mathcal{U}}_{i+1}, \mathcal{E}_{i+1}, \mathsf{P}_{i+1}$$

then the i-th step satisfies ET_{CC}, *i.e.*

$$\mathsf{ET}_{CC} \vdash (\bar{\mathcal{K}}_i, u_{cl}) \triangleright \mathcal{F} : (\bar{\mathcal{K}}_{i+1}, \bar{\mathcal{U}}_{i+1}(cl))$$

PROOF. We prove this by induction on the length of trace i. We introduce two invariants I_1, I_2. The I_1 states that for any view u, if the view u includes a version, it also includes all the version it depends on, that is,

$$\forall cl, u, k, k', j, m, dp.$$
$$u = \mathsf{co} - \mathsf{dom}(\bar{\mathcal{U}}_i) \wedge j \in u_{cl}(k') \wedge (\_, \_, \_, dp) = \bar{\mathcal{K}}_i(k, j) \wedge (k', \mathsf{w}(\bar{\mathcal{K}}_i(k', m))) \in dp$$
$$\Rightarrow m \in u(k')$$

The I_2 states that for any transaction t, $\mathsf{deps}(\bar{\mathcal{K}}_i, t)$ includes transactions that t depends on with respect to $(\mathsf{SO} \cup \mathsf{WR})$.

$$\forall t, k, j. \ t = \mathsf{w}(\bar{\mathcal{K}}_i(k, j)) \Rightarrow ((\mathsf{SO} \cup \mathsf{WR})^{-1})^*(t) \subseteq \mathsf{deps}(\bar{\mathcal{K}}_i(k, j))$$

if the view u includes a version, Since $\mathsf{ET}_{CC} = \mathsf{ET}_{MR} \cap \mathsf{ET}_{RYW} \cap \mathtt{allowed}(\mathsf{SO} \cup \mathsf{WR})$, we prove the three constraints separately. In each case we need to consider PUT, GETTRANS and SYNC Let t be the new transaction committed to replica r by the client cl. Let u_{cl} and u'_{cl} be the client views before and after respectively, and u_r and u'_r be the replica views before and after respectively.

- ET_{MR}. For PUT, it is easy to see $u_{cl} \sqsubseteq u'_{cl}$ and $u_r \sqsubseteq u'_r$. For GETTRANS, it is known that $u_r = u'_r$. In the premise of GETTRANS, it fetches a version, m_i-*th* version, from the replica for each key k_i and adds all dependencies, which are associated with version, in the client view. This means $u_{cl} \sqsubseteq u'_{cl}$.
- ET_{RYW}. By ET_{MR}, it suffices to only consider if the newly committed transaction is included in the view. We only need to consider PUT. The new views $u'_{cl} = u_{cl}[k \mapsto j]$ and $u'_r = u_r[k \mapsto j]$ where the new version $\bar{\mathcal{K}}_{i+1}(k, j)$ is written by the client cl.

- allowed(SO ∪ WR). Given the invariants I_1 and I_2, let consider Put, GetTrans and Sync.
  - Put. By invariants I_1 and I_2, we know $u_{cl} = \text{getView}\big(\bar{\mathcal{K}}_i, ((\text{SO} \cup \text{WR})^{-1})^*(\text{visTx}(\bar{\mathcal{K}}_i, u_{cl}))\big)$. Now we need to re-establish the I_1 for the new views u'_{cl} and $\bar{\mathcal{U}}'$ and I_2 for the new transaction t.
    * Since it is a single-write transaction, assume it appends $j\text{-}th$ version to key k. This means the new client views $u'_{cl} = u_{cl}[k \mapsto u_{cl}(k) \cup \{j\}]$, Let dp be the dependency for the new version, i.e. $dp = \text{deps}(\bar{\mathcal{K}}_{i+1}(k, j))$. By the premiss, the dp includes all versions included in u_{cl}. This means for any key k' and index m:

      $$(k', \text{w}(\bar{\mathcal{K}}_i(k', m))) \in dp \Rightarrow m \in u'_{cl}(k')$$

      Note that $u'_{cl} \sqsubseteq u'_r$, therefore the invariant I_1 holds.
    * It is enough to only consider the newly committed transaction t. Since it is a write transaction, let assume any transaction t' such that $t' \xrightarrow{\text{SO}} t$. By MR and RYW, transaction t' is in the view u_{cl}, that is, for some k', m:

      $$t' \in \{\text{w}(\bar{\mathcal{K}}_i(k', m))\} \cup \text{rs}(\bar{\mathcal{K}}_i(k', m)) \Rightarrow m \in u_{cl}(k')$$

      Given the definition of dependencies the new version written by t:

      $$dp = \big\{(k'', t'') \,\big|\, \exists z.\ z \in u_{cl}(k'') \wedge t'' = \text{w}(\bar{\mathcal{K}}_i(k'', z)) \vee (k'', t'') \in \text{deps}(\bar{\mathcal{K}}_i(k'', z))\big\}$$

      it follows $(k', m) \in dp$. By I.H. of invariant I_2, we know that dp is close with respect to SO ∪ WR. Thus we know for the new transaction $((\text{SO} \cup \text{WR})^{-1})^*(t) \subseteq \text{deps}(\bar{\mathcal{K}}_i, t)$, that is, we re-establish invariant for $i + 1$.
  - For GetTrans, the client cl commits a read-only transaction. By the premiss of the rule, the client will pick a new view u'_{cl} such that $u_0 = u_{cl} \sqsubseteq u'_{cl} = u_j$ and:

    $$\begin{aligned}
    &\text{for } z \text{ in } \{1, \ldots, j\} \\
    &\quad m_z \in u_r(k_z) \\
    &\quad u'_z = u_{z-1}[k_z \mapsto u_{z-1}(k_z) \cup \{m_z\}] \\
    &\quad u_z = \lambda k.\, u'_z(k) \cup \big\{x \,\big|\, (k, \text{w}(\bar{\mathcal{K}}(k, x))) \in \text{deps}(\bar{\mathcal{K}}(k_z, m_z))\big\}
    \end{aligned}$$

    We prove each iteration preserves the invariant I_1. Suppose the invariants holds after $(z-1)\text{-}th$, let consider $z\text{-}th$. By construction, $u'_z = u_{z-1}[k_z \mapsto u_{z-1}(k_z) \cup \{m_z\}]$, the m_z version of some key k_z is included in u'_z. Then all the versions $\bar{\mathcal{K}}_i(k_z, m_z)$ is is included in u_z as

    $$u_z = \lambda k.\, u'_z(k) \cup \big\{x \,\big|\, (k, \text{w}(\bar{\mathcal{K}}(k, x))) \in \text{deps}(\bar{\mathcal{K}}(k_z, m_z))\big\}$$

    This implies I_1 for u_z. For I_2, we apply I.H. of invariant I_2, since the new transaction is a read-only transaction.
  - For Sync, the premiss of the rule implies that I_1 holds under the new replica view u'_r. For I_2, we apply I.H. of invariant I_2.

  $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\quad \square$

H.2 Clock-SI

H.2.1 Code.

Structure. Clock-SI is a partitioned distributed NoSQL database, which means each server, also called shard, contains part of keys and does not overlap with any other servers. Clock-SI implements snapshot isolation. To achieve that, each shard tracks the physical time. Note that times between shards do not match, but there is a upper bound of the difference.

```
1   Shard :: ID -> ( clockTime )
```

Code 11. Shard

A key maintains a list of values and their versions. A version is the time when such value is committed.

```
1   VersionNo :: Time
2   KV :: Keys -> List( Val , VersionNo )
3   (each key is asscoaited with a shard)
```

Code 12. Key-value store

The idea behind Clock-SI is that a client starts a transaction in a shard, and the shard is responsible for fetching value from other shards if keys are not stored in the local shard. During execution, a transaction tracks the write set.

```
1   WS :: Key -> Val
```

Code 13. Write set

At the end, the transaction commits all the update in the write set, and the local shard acts as coordinator to update keys either locally or remotely. To commit a transaction, Clock-SI use two-phase commits protocol. A transaction has four states:

- `active`, the transaction is still running;
- `prepared`, shards receive the update requests from the coordinator;
- `committing`, shards receive the update confirmations from the coordinator;
- `committed`, the transaction commits successfully.

To implement SI, a transaction also tracks its snapshot time so it knows which version should be fetched. Also a transaction tracks the prepared and committing times, which are used to postpone other transactions' reads if those transactions' snapshots time are greater.

```
1   State :: { active , prepared , committing , committed }
2   Trans :: ( state , snapTime , prepareTime , commitTime , ws )
```

Code 14. Transaction runtime

Start Transaction. Clock-SI proposes two versions, with or without session order. Here we verify the one with session order. To start a transaction, the client contacts a shard and provides the previous committing time. The shard will return a snapshot time, which is greater than the committing time provided, for the new transaction. Note that client might connects to a different shard from last time, which means that the shard might have to wait until the shard local time is greater than the committing time.

```
1   startTransaction ( Trans t, Time ts )
2       wait until ts < getClockTime ();
3       t.snapshotTime = getClockTime ();
4       t.state = active;
```

Code 15. Transaction runtime

From this point, such transaction will always interact with the shard and the shard will act as coordinator if necessary.

Read. A transaction t might read within the transaction if the key has been updated by the same transaction before, that is, read from the write set ws. A transaction t might read from the original shard if the key store in the shard, but it has to wait until any other transactions t' commit successfully who are supposed to commit before the current transaction's snapshot time, i.e. t' are in prepared or committing stage and the corresponding time is less the t snapshot time.

```
1   Read( Trans t, key k )
2       if ( k in t.ws ) return ws(k);
3       if ( k is updated by t' and t'.state = committing
4                   and t.snapshotTime > t'.committingTime )
5           wait until t'.state == committed;
6       if ( k is updated by t'
7                       and t'.state = prepared and t.snapshotTime > t'.preparedTime
8                       and t.snapshotTime > t'.committingTime )
9           wait until t'.state == committed;
10      return K(k,i), where i is the latest version before t.snapshotTime;
```
<div align="center">Code 16. Read from original shard</div>

If the key is not stored in the original shard, the original shard sends a read request to the shard containing the key. Because of time difference, the remote shard's time might before the snapshot time of the transaction. In this case, the shard wait until the time catches up.

```
1   On read k request from a remote transaction t
2       wait until t.snapshotTime < getClockTime()
3       return read(t,k);
```
<div align="center">Code 17. Read from original shard</div>

Commit Write Set. If all the keys in the write set are hosted in the original shard that the transaction first connected, the write set only needs to commit local.

```
1   localCommit( Trans t )
2       if noConcurrentWrite(t) {
3           t.state = committing;
4           t.commitTime = getClockTime();
5           log t.commitTime;
6           log t.ws;
7           t.state = committed;
8       }
```
<div align="center">Code 18. Local Commit</div>

To commit local, it first checks, by noConcurrentWrite(t), if there is any transaction t' that writes to the same key as the transaction new transaction t, and the transaction t' commit after the snapshot of t. Since writing database needs time, it sets the transaction state to committing and log the commitTime, before the updating really happens. During committing state, other transactions will be pending, if they want to read the keys being updated. Last, the state of transaction is set to committed.

To commit to several shards, Clock-SI uses two-phase protocol.

```
1   distributedCommit( Trans t )
2       for p in t.updatedPartitions { send ``prepare t'' to p; }
3       wait receiving ``t prepared'' from all participants, store into prep;
4       t.state = committing;
5       t.commitTime = max(prep);
```

```
6        log t.commitTime;
7        t.state = committed;
8        for p in t.updatedPartitions { send ``commit t'' to p; }
9
10   On receiving ``prepare t''
11       if noConcurrentWrite(t) {
12           log t.ws and t.coordinator ID
13           t.state = prepared;
14           t.prepareTime = getClockTime();
15           send ``t prepared'' to t.coordinator
16       }
17
18   On receiving ``commit t''
19       log t.commitTime
20       t.state = committed
```

Code 19. Distributed Commit

The original shard, who acts as the coordinator, sends ''prepare t'' to shards that will be updated. Any shard receiving ''prepare t'' checks, similarly, if there is any transaction write to the same key committing after the snapshot time. If the check passes, the shard logs the write set and the coordinator shard ID, set the state to prepared, and sends the local time to the coordinator. Once the coordinator receives all the prepared messages, it starts the second phase by setting the state to committing. Then the coordinator picks the largest time from all the prepared messages as the commit time for the new transaction. Since the write set has been logged in the first phase, so here it can immediately set the state to be committed. Last, the coordinator needs to send commit t to other shards so they will log the commit time and set the state to committed. Note that participants have different view for the new transaction from the coordinator, but it guarantees eventually they all updated to committed with the same commit time.

H.2.2 Verification.

Structure. We model the database use key-value store from Def. 3.1, yet here it is necessary to satisfy the well-formed conditions. Transaction identifier t_{cl}^n are labelled with the committing time n. Sometime we also write t_{cl}^c or omit the client label, i.e. t^n and t^c.

Database is partitioned into several *shards*. A shard $r \in \textsc{Shards}$ contains some keys which are disjointed from keys in other shards. The shardOf(k) denotes the shard where the key k locates.

Shards and clients are associated with clock times, $c \in \textsc{ClockTimes} \triangleq \mathbb{N}$, which represent the current times of shards and clients. We use notation $C \in (\textsc{Shards} \cup \textsc{Client}) \xrightarrow{fin} \textsc{ClockTimes}$.

We will use notation [T] to denote the static code of a transaction, and $[T]_c^{\mathcal{F}}$ for the runtime of a transaction, where c is the snapshot time and $\mathcal{F}$ is the read-write set. Note that in the model, we only distinguishes active and committed state, since the prepared and committing are only for better performance.

Start Transaction. To start a transaction, it picks a random shard r as the coordinator, reads the local time $C(r)$ as the snapshot, and sets the initial read-write set to be an empty set. Also the client time is updated to the snapshot time. For technical reason, we also update the shard time to avoid time collision to other transaction about to commit. Note that in real life, all the operations

running in a shard take many time cycles, so it is impossible to have time collision.

$$\frac{\text{StartTrans} \quad c < C(r) \qquad C' = C[r \mapsto C(r) + 1] \quad \text{---> simulate time elapses}}{cl \vdash \mathcal{K}, c, C, s, [\mathsf{T}] \xrightarrow{C(r), c, \emptyset, \perp} \mathcal{K}, C(r), C', s, [\mathsf{T}]_{C(r)}^{\emptyset}}$$

Read. The clock-SI protocol includes some codes related to performance which does not affect the correctness. Clock-SI distinguishes a local read/commit and a remote read/commit, yet it is sufficient to assume all the read and commit are "remote", while the local read and commit can be treated as self communication. Similarly we assume a transaction always commits to several shards.

```
1   On receive ``read(t,k)'' {
2       if ( k in t.ws ) return ws(k);
3
4       asssert( t.snapshotTime < getClockTime() )
5       for t' that writes to k:
6           if(t.snapshotTime > t'.preparedTime
7                       || t.snapshotTime > t'.committingTime )
8               asssert( t.state == committed )
9
10      return K(k,i), where i is the latest version before t.snapshotTime;
11  }
```

Code 20. simplified read

If the key exists in the write set $\mathcal{F}$, the transaction read from the write set immediately.

$$\frac{\text{ReadTrans} \quad k = [\![\mathsf{E}]\!]_s \qquad (\mathsf{w}, k, v) \in \mathcal{F} \quad \text{---> Code 20, line 2}}{cl \vdash \mathcal{K}, c, C, s, [\mathsf{x} := [\mathsf{E}]]_c^{\mathcal{F}} \xrightarrow{cl, c, \mathcal{F} \lessdot (\mathsf{r}, k, v), \perp} \mathcal{K}, c, C, s[\mathsf{x} \mapsto v], [\mathsf{skip}]_c^{\mathcal{F} \lessdot (\mathsf{r}, k, v)}}$$

Otherwise, the transaction needs to fetch the value from the shard. The first premiss says the transaction must wait until the shard local time $C(\text{shardOf}(k))$ is greater than the snapshot time c. If so, by the second line, the transaction fetches the latest version for key k before the snapshot time c.

$$\frac{\begin{array}{c} \text{ReadRemote} \\ k = [\![\mathsf{E}]\!]_s \qquad (\mathsf{w}, k, \_) \notin \mathcal{F} \qquad c < C(\text{shardOf}(k)) \quad \text{---> Code 20, line 4} \\ n = \max \left\{ n' \;\middle|\; \exists j. \; t^{n'} = \mathsf{w}(\mathcal{K}(k, j)) \wedge n' < c \right\} \quad \text{---> Code 20, line 10} \\ \mathcal{K}(k, i) = (v, t^n, \_) \end{array}}{cl \vdash \mathcal{K}, c, C, s, [\mathsf{x} := [\mathsf{E}]]_c^{\mathcal{F}} \xrightarrow{cl, c, \mathcal{F} \lessdot (\mathsf{r}, k, v), \perp} \mathcal{K}, c, C', s[\mathsf{x} \mapsto v], [\mathsf{skip}]_c^{\mathcal{F} \lessdot (\mathsf{r}, k, v)}}$$

Write. Write will not go to the shard until committing time. Before it only log it in the write set.

$$\frac{\text{Write} \quad k = [\![\mathsf{E}_1]\!]_s \qquad v = [\![\mathsf{E}_2]\!]_s}{cl \vdash \mathcal{K}, c, C, s, [[\mathsf{E}_1] := \mathsf{E}_2]_c^{\mathcal{F}} \xrightarrow{cl, c, \mathcal{F} \lessdot (\mathsf{w}, k, v), \perp} \mathcal{K}, c, C, s, [\mathsf{skip}]_c^{\mathcal{F} \lessdot (\mathsf{w}, k, v)}}$$

Commit. We also assume transaction always commit to several shards and the local commit is treated as self-communication.

```
1   commit( Trans t )
2       for p in t.updatedPartitions
3           send ``prepare t'' to p;
4       wait receiving ``t prepared'' from all participants, store into prep;
5       t.state = committing;
6       t.commitTime = max(prep);
7       log t.commitTime;
8       t.state = committed;
9       for p in t.updatedPartitions
10          send ``commit t'' to p;
11
12  On receiving ``prepare t''
13      if noConcurrentWrite(t)
14          log t.ws to t.coordinator ID
15          t.state = prepared;
16          t.prepareTime = getClockTime();
17          send ``t prepared'' to t.coordinator
18
19  On receiving ``commit t''
20      log t.commitTime
21      t.state = committed
```

Code 21. simplified commit

Note that Clock-SI uses two phase commit: the coordinator (the shard that the client directly connects to) distinguishes "committing" state and "committed" state, where in between the coordinator pick the committing time and log the write set, and the participants distinguishes "prepared" state and "committed" state. Such operations are for possible network partition or single shard errors, and allowed a more fine-grain implementations which do not affect the correctness, therefore it suffices to assume they are one atomic step.

COMMIT

$$\forall k, i.\ (\mathrm{w}, k, \_) \in \mathcal{F} \land \mathrm{w}(\mathcal{K}(k, i)) < c \quad \text{---> Code 21, line 13}$$
$$n = \max\left(\{c' \mid \exists k.\ (\_, k, \_) \in \mathcal{F} \land c' = C(\mathrm{shardOf}(k))\} \cup \{c\}\right) \quad \text{---> Code 21, lines 4 to 6}$$
$$\mathcal{K}' = \mathrm{commitKV}(\mathcal{K}, c, t_{cl}^n, \mathcal{F}) \quad \text{---> Code 21, lines 20 and 21}$$
$$\forall r.\ \begin{pmatrix} r \in \{\mathrm{shardOf}(k) \mid (\_, k, \_) \in \mathcal{F}\} \\ \Rightarrow C'(r) = C(r) + 1 \end{pmatrix} \lor (C'(r) = C(r)) \quad \text{---> simulate time elapses}$$

$$\overline{r, cl \vdash \mathcal{K}, c, C, s, [\mathrm{skip}]_c^{\mathcal{F}} \xrightarrow{cl, c, \mathcal{F}, n} \mathcal{K}', n+1, C', s, \mathrm{skip}}$$

To commit the new transaction, it needs to check, by the first premiss, there is no other transactions writing to the same keys after the snapshot time. If it passes, by the second line it picks the maximum time n among all participants as the commit time. The new key-value store $\mathcal{K}' =$

$\text{commitKV}\big(\mathcal{K}, c, t_{cl}^n, \mathcal{F}\big)$, where

$$\text{commitKV}(\mathcal{K}, c, t, \mathcal{F} \uplus \{(\mathsf{r}, k, v)\}) \triangleq \texttt{let } n = \max\left\{ n' \,\middle|\, \exists j.\ t^{n'} = \mathsf{w}(\mathcal{K}(k, j)) \wedge n' < c \right\}$$
$$\text{and } \mathcal{K}(k, i) = (v, t^n, \mathcal{T})$$
$$\text{and } \mathcal{K}' = \text{commitKV}(\mathcal{K}, c, t, \mathcal{F})$$
$$\texttt{in } \mathcal{K}'[k \mapsto \mathcal{K}'(k)[i \mapsto (v, t^n, \mathcal{T} \cup \{t\})]]$$

$$\text{commitKV}(\mathcal{K}, c, t, \mathcal{F} \uplus \{(\mathsf{w}, k, v)\}) \triangleq \texttt{let } \mathcal{K}' = \text{commitKV}(\mathcal{K}, c, t, \mathcal{F})$$
$$\texttt{in } \mathcal{K}'[k \mapsto \mathcal{K}'(k)::(v, t, \emptyset)]$$

Note that commitKV is similar to update by appending the new version to the end of a list. The commitKV also updates versions read by the new transaction using the snapshot time of the transaction. Last, like STARTTRANS we update the client time after the commit time, i.e. $n+1$ and simulate time elapses for all shards updated.

Time Tick. For technical reasoning, we have non-deterministic time elapses.

TIMETICK
$$\mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{r, C(r)+1} \mathcal{K}, C, C'[r \mapsto C(r) + 1], \mathcal{E}, \mathsf{P}$$

CLIENTSTEP
$$\frac{cl \vdash \mathcal{K}, C(cl), C', \mathcal{E}(cl), \mathsf{P}(cl) \xrightarrow{cl, c', \mathcal{F}, c''} \mathcal{K}', c, C'', s, \mathsf{C}}{\mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c', \mathcal{F}, c''} \mathcal{K}, C[cl \mapsto c], C'', \mathcal{E}[cl \mapsto s], \mathsf{P}[cl \mapsto \mathsf{C}]}$$

Verification. Clock-SI allows interleaving, yet for any clock-si trace τ there exists a equivalent trace τ' where transactions do not interleave with others (Theorem H.8). Furthermore, in such trace τ', transactions are reduced in their commit order.

THEOREM H.8 (NORMAL CLOCK-SI TRACE). *A clock-SI trace τ is a clock-SI normal trace if it satisfies the following: there is no interleaving of a transaction,*

$$\forall cl, c, \mathcal{K}_i, C_i, C_i', \mathcal{E}_i, \mathsf{P}_i.$$
$$\tau = \cdots \xrightarrow{cl, c, \_, \bot} \mathcal{K}_i, C_i, C_i', \mathcal{E}_i, \mathsf{P}_i \Rightarrow \cdots$$
$$\Rightarrow \exists c', \mathcal{K}_j, C_j, C_j', \mathcal{E}_j, \mathsf{P}_j.$$
$$\tag{8.4}$$
$$\tau = \cdots \xrightarrow{cl, c, \_, \bot} \mathcal{K}_i, C_i, C_i', \mathcal{E}_i, \mathsf{P}_i \xrightarrow{cl, c, \_, \bot} \_ \xrightarrow{cl, c, \_, \bot} \cdots \xrightarrow{cl, c, \_, c'} \mathcal{K}_j, C_j, C_j', \mathcal{E}_j, \mathsf{P}_j$$

and transactions in the trace appear in the committing order,

$$\forall cl_i, cl_j, c_i, c_j, c_i', cl_j', \mathcal{F}_i, \mathcal{F}_j \mathcal{K}_i, \mathcal{K}_j, C_i, C_j, C_i', C_j', \mathcal{E}_i, \mathcal{E}_j, \mathsf{P}_i, \mathsf{P}_j.$$
$$\tau = \cdots \xrightarrow{cl_i, c_i, \mathcal{F}_i, c_i'} \mathcal{K}_i, C_i, C_i', \mathcal{E}_i, \mathsf{P}_i \Rightarrow \cdots \xrightarrow{cl_j, c_j, \mathcal{F}_j, c_j'} \mathcal{K}_j, C_j, C_j', \mathcal{E}_j, \mathsf{P}_j \Rightarrow c_i' < c_j'$$
$$\tag{8.5}$$

For any clock-SI trace τ, there exists an equivalent normal trace τ' which has the same final configuration as τ.

PROOF. Given a trace τ, we first construct a trace τ' that satisfies Eq. (8.5), by swapping steps. Let take the first two transactions $t_{cl_i}^n$ and $t_{cl_j}^m$ that are out of order, i.e. $n > m$ and

$$\tau = \cdots \xrightarrow{cl_i, c_i, \mathcal{F}_i, n} \mathcal{K}_i, C_i, C_i', \mathcal{E}_i, \mathsf{P}_i \Rightarrow \cdots \xrightarrow{cl_j, c_j, \mathcal{F}_j, m} \mathcal{K}_j, C_j, C_j', \mathcal{E}_j, \mathsf{P}_j$$

By Lemma H.9, the two clients are different $cl_i \neq cl_j$ and thus two steps are unique in the trace. We will construct a trace that $t_{cl_i}^n$ commits after $t_{cl_j}^m$.

- First, it is important to prove that $t_{cl_j}^m$ does not read any version written by $t_{cl_i}^n$. By Lemma H.20, the snapshot time c_j of $t_{cl_j}^m$ is less than the commit time, i.e. $c_j < m$, therefore $c_j < n$. By the READ rule, $c_j < n$ implies the transaction $t_{cl_j}^m$ never read any version written by $t_{cl_i}^n$.

- Let consider any possible time tick for those shard r that has been updated by $t_{cl_j}^m$, that is, $r = \text{shardOf}(k)$ for some key k that $(\mathsf{w}, k, \_) \in \mathcal{F}_i$ and

$$\tau = \cdots \xrightarrow{cl_i, c_i, \mathcal{F}_i, n} \mathcal{K}_i, C_i, C_i', \mathcal{E}_i, \mathsf{P}_i \Rightarrow \cdots \xrightarrow{r, c} \_ \Rightarrow \cdots \xrightarrow{cl_j, c_j, \mathcal{F}_j, m} \mathcal{K}_j, C_j, C_j', \mathcal{E}_j, \mathsf{P}_j \qquad (8.6)$$

Since $c_j < m < n < c$, therefore such time tick will not affect the transaction $t_{cl_j}^m$, which means it is safe to move the time tick step after the $t_{cl_j}^m$.

Now we can move the commit of $t_{cl_i}^n$ and time tick steps similar to Eq. (8.6) after the commit of $t_{cl_j}^m$,

$$\tau' = \cdots \xrightarrow{cl_j, c_j, \mathcal{F}_j, m} \mathcal{K}_j, C_j, C_j', \mathcal{E}_j, \mathsf{P}_j \xrightarrow{cl_i, c_i, \mathcal{F}_i, n} \mathcal{K}_i, C_i, C_i', \mathcal{E}_i, \mathsf{P}_i \xrightarrow{r, c} \cdots$$

We continually swap the out of order transaction until the newly constructed trace τ' satisfying Eq. (8.5).

Now let consider Eq. (8.4). Let take the first transaction t whose read has been interleaved by other transaction or a time tick.

- If it is a step that the transaction t read from local state,

$$\tau = \cdots \xrightarrow{cl, c, \mathcal{F} \lessdot (r, k, v), \perp} \_ \xrightarrow{\alpha} \cdots \xrightarrow{cl, c, \mathcal{F}'', n} \cdots$$

then by READTRANS we know $\mathcal{F} \lessdot (r, k, v) = \mathcal{F}$, and it is safe to swap the two steps as the following

$$\tau' = \cdots \xrightarrow{\alpha} \_ \xrightarrow{cl, c, \mathcal{F} \lessdot (r, k, v), \perp} \cdots \xrightarrow{cl, c, \mathcal{F}'', n} \cdots$$

- If it is a step that the transaction t read from remote, the step might be interleaved by a step from other transaction or time tick step.
  - if it is interleaved by the commit of other transaction $t' = t_{cl'}^m$, that is

  $$\tau = \cdots \xrightarrow{cl, c, \mathcal{F}, \perp} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl', c', \mathcal{F}, m} \cdots \xrightarrow{cl, c, \mathcal{F}'', n} \cdots$$

  where $cl' \neq cl$.
  * if the transaction t' does not write to any key k that is read by t,

  $$\forall k. \ (\mathsf{r}, k, \_) \in \mathcal{F} \Rightarrow (\mathsf{w}, k, \_) \notin \mathcal{F}'$$

  In this case, it is safe to swap the two steps

  $$\tau' = \cdots \xrightarrow{cl', c', \mathcal{F}, m} \_ \xrightarrow{cl, c, \mathcal{F}, \perp} \cdots \xrightarrow{cl, c, \mathcal{F}'', n} \cdots$$

  * if the transaction t' write to a key k that is read by t,

  $$(\mathsf{r}, k, \_) \in \mathcal{F} \wedge (\mathsf{w}, k, \_) \in \mathcal{F}'$$

  Let $r = \text{shardOf}(k)$. By the READREMOTE, we know the current clock time for the shard r is greater than c which is the snapshot time of t, that is, $C'(r) > c$. Then by COMMIT, the commit time of t' is picked as the maximum of the shards it touched, i.e. $m \geq C'(r)$. Now by the READREMOTE and $m \geq c$, it is safe to swap the two steps since the new version of k does not affect the t.

– if it is interleaved by the read of other transaction t', that is

$$\tau = \cdots \xrightarrow{cl,c,\mathcal{F},\perp} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl',c',\mathcal{F}',\perp} \cdots \xrightarrow{cl,c,\mathcal{F}'',n} \cdots$$

Because reads have no side effect to any shard by ReadRemote, it is safe to swap the two steps

$$\tau' = \cdots \xrightarrow{cl',c',\mathcal{F}',\perp} {}_{-} \xrightarrow{cl,c,\mathcal{F},\perp} \cdots \xrightarrow{cl,c,\mathcal{F}'',n} \cdots$$

– if it is interleaved by a time tick step,

$$\tau = \cdots \xrightarrow{cl,c,\mathcal{F} \lessdot (r,k,v),\perp} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{r,c'} \cdots \xrightarrow{cl,c,\mathcal{F}'',n} \cdots$$

∗ if the transaction t does not read from the shard r, it means for any key k,

$$\mathrm{shardOf}(k) \neq r$$

In this case, it is safe to swap the two steps

$$\tau' = \cdots \xrightarrow{r,c'} {}_{-} \xrightarrow{cl,c,\mathcal{F} \lessdot (r,k,v),\perp} \cdots \xrightarrow{cl,c,\mathcal{F}'',n} \cdots$$

∗ if the transaction t read from the shard r, it means that there exists a key k

$$\mathrm{shardOf}(k) = r$$

By the ReadRemote, we know the current clock time for the shard r is greater than the snapshot time of t, that is, $C'(r) > c$. Then by TimeTick, we have $c' > C'(r)$. Now by the ReadRemote and $c' > c$, it is safe to swap the two steps.

$\square$

Lemma H.9 (Monotonic client clock time). *The clock time associated with a client monotonically increases, That is, given a step*

$$\mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \Rightarrow \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}'$$

then for any clients cl,

$$C(cl) \leq C(cl')$$

Proof. It suffices to only check the ClientStep rule which is the only rule updates the client clock time, especially, it is enough to check the client cl that who starts or commits a new transaction.

• Commit. Let c be the clock time before committing, $c = C(cl)$. By the premiss of the rule, the new client time $n + 1$ satisfies that,

$$n = \max\left(\{c' \mid \exists k. \ (\_, k, \_) \in \mathcal{F} \wedge c' = C(\mathrm{shardOf}(k))\} \cup \{c\}\right)$$

It means $c < (n + 1)$.

• StartTrans. Let c be the clock time before taking snapshot, $c = C(cl)$. By the premiss of the rule the new client time $C'(r)$ for a shard r, such that $c < C'(r)$.

$\square$

Lemma H.10 (No side effect local operation). *Any transactional operation has no side effect to the shard and key-value store,*

$$\mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl,c,\mathcal{F},\perp} \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}' \Rightarrow \mathcal{K} = \mathcal{K}'$$

Proof. It is easy to see that StartTrans, ReadTrans, ReadRemote and Write do not change the state of key-value store.

$\square$

Clock-SI also has a notion view which corresponds the snapshot time. The following definition $\text{viewOf}(\mathcal{K}, c)$ extracts the view from snapshot time.

Definition H.11. Given a normal clock-SI trace τ and a transaction t_{cl}, such that

$$\tau = \cdots \xrightarrow{cl, c, \emptyset, \perp} \cdots \xrightarrow{cl, c, \mathcal{F}, \perp} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, c'} \cdots$$

the initial view of the transaction is defined as the following:

$$\text{viewOf}(\mathcal{K}, c) \triangleq \lambda k. \{i \mid \exists t^n.\ \mathsf{w}(\mathcal{K}(k, i)) = t^n \wedge n < c\}$$

Given the view $\text{viewOf}(\mathcal{K}, c)$ for each transaction, we first prove that clock-si produces a well-formed key-value store (Def. 3.1).

LEMMA H.12. *Given any key-value store $\mathcal{K}$ and snapshot time c from a clock-SI trace τ,*

$$\tau = \cdots \xrightarrow{} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \rightarrow cl, c, \mathcal{F}, c' \cdots$$

$\text{viewOf}(\mathcal{K}, c)$ *and* $\text{viewOf}(\mathcal{K}, c')$ *(Def. H.11) produce well-formed views.*

PROOF. It suffices to prove that Eq. (atomic) in Def. 3.2. Assume a key-value store $\mathcal{K}$ and a snapshot time c. Suppose a version i in the view $i \in \text{viewOf}(\mathcal{K}, c)(k)$ for some key k. By Def. H.11, the version is committed before the snapshot time, i.e. $t^n = \mathsf{w}(\mathcal{K}(k, i)) \wedge n < c$. Assume another version $t^n = \mathsf{w}(\mathcal{K}(k', j))$ for some key k' and index j. By Def. H.11 we have $j \in \text{viewOf}(\mathcal{K}, c)(k')$. Similarly $\text{viewOf}(\mathcal{K}, c')$ is a well-formed view. □

Second, given the view $\text{viewOf}(\mathcal{K}, c)$ for each transaction, both commitKV and update produce the same result.

LEMMA H.13. *Given a normal clock-SI trace τ and a transaction t_{cl}, such that*

$$\tau = \cdots \xrightarrow{cl, c, \emptyset, \perp} \cdots \xrightarrow{cl, c, \mathcal{F}, \perp} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, c'} \cdots$$

the following holds:

$$\text{commitKV}\left(\mathcal{K}, c, t_{cl}^{c'}, \mathcal{F}\right) = \text{update}\left(\mathcal{K}, \text{viewOf}(\mathcal{K}, c), \mathcal{F}, t_{cl}^{c'}\right)$$

PROOF. We prove by induction on $\mathcal{F}$.

- Base case: $\mathcal{F} = \emptyset$. It is easy to see that

$$\text{commitKV}\left(\mathcal{K}, c, t_{cl}^{c'}, \emptyset\right) = \mathcal{K} = \text{update}\mathcal{K}, \text{viewOf}(\mathcal{K}, c), \mathcal{F}, t_{cl}^{c'}$$

- Inductive case: $\mathcal{F} \uplus (\mathsf{w}, k, v)$. Because in both functions, the new version is installed at the tail of the list associated with k,

$$\begin{aligned}
&\text{commitKV}\left(\mathcal{K}, c, t_{cl}^{c'}, \mathcal{F} \uplus (\mathsf{w}, k, v)\right) \\
&= \quad \text{commitKV}\left(\mathcal{K}, c, t_{cl}^{c'}, \mathcal{F}\right)[k \mapsto \mathcal{K}(k)::(v, t, \emptyset)] \\
&= \quad \text{update}\left(\mathcal{K}, \text{viewOf}(\mathcal{K}, c), \mathcal{F}, t_{cl}^{c'}\right)[k \mapsto \mathcal{K}(k)::(v, t, \emptyset)] \\
&= \quad \text{update}\left(\mathcal{K}, \text{viewOf}(\mathcal{K}, c), \mathcal{F} \uplus (\mathsf{w}, k, v), t_{cl}^{c'}\right)
\end{aligned}$$

- Inductive case: $\mathcal{F} \uplus (\mathsf{r}, k, v)$. Let $\mathcal{K}(k, i)$ be the version being read. That is, the writer $t^n = \mathsf{w}(\mathcal{K}(k, i))$ is the latest transaction written to the key k before the snapshot time c,

$$n = \max\left\{n' \mid \exists j.\ t^{n'} = \mathsf{w}(\mathcal{K}(k, j)) \wedge n' < c\right\}$$

Let the new version $v = \left(\mathrm{val}(\mathcal{K}(k, i)), \mathrm{w}(\mathcal{K}(k, i)), \mathrm{rs}(\mathcal{K}(k, i)) \uplus \left\{t_{cl}^{c'}\right\}\right)$. By Lemma H.12, it follows $i \in \mathrm{viewOf}(\mathcal{K}, c)(k)$, then by Lemma H.14, the version is the latest one $i = \max(\mathrm{viewOf}(\mathcal{K}, c)(k))$. Therefore we have,

$$
\begin{aligned}
& \mathrm{commitKV}\left(\mathcal{K}, c, t_{cl}^{c'}, \mathcal{F} \uplus (\mathrm{r}, k, v)\right) \\
= \quad & \mathrm{commitKV}\left(\mathcal{K}, c, t_{cl}^{c'}, \mathcal{F}\right)[k \mapsto \mathcal{K}(k)[i \mapsto v]] \\
= \quad & \mathrm{update}\left(\mathcal{K}, \mathrm{viewOf}(\mathcal{K}, c), \mathcal{F}, t_{cl}^{c'}\right)[k \mapsto \mathcal{K}(k)[i \mapsto v]] \\
= \quad & \mathrm{update}\left(\mathcal{K}, \mathrm{viewOf}(\mathcal{K}, c), \mathcal{F} \uplus (\mathrm{r}, k, v), t_{cl}^{c'}\right)
\end{aligned}
$$

$\square$

LEMMA H.14 (STRICTLY MONOTONIC WRITERS). *Each version for a key has a writer with strictly greater clock time than any versions before:*

$$
\forall \mathcal{K}, k, i, j, t^n, t^m. \; \mathrm{w}(\mathcal{K}(k, i)) = t^n \wedge \mathrm{w}(\mathcal{K}(k, j)) = t^m \wedge i < j \Rightarrow n < m
$$

By Theorem H.8, it is sufficient to only consider normal clock-SI trace. Since transactions do not interleave in a normal clock-SI trace, all transactional execution can be replaced by Fig. 1.

THEOREM H.15 (SIMULATION). *Given a clock-SI normal trace τ, a transaction t_{cl}^n from the trace, and the following transactional internal steps*

$$
\mathcal{K}_0, c_0, C_0, s_0, [\mathsf{T}] \xrightarrow{cl, c, \emptyset, \perp} \cdots \xrightarrow{cl, c, \mathcal{F}, n} \mathcal{K}_i, c_i, C_i, s_i, [\texttt{skip}]
$$

for some i, there exists a trace

$$
(s_0, \mathrm{snapshot}(\mathcal{K}_0, \mathrm{viewOf}(\mathcal{K}_0, c)), \emptyset), \mathsf{T} \to^* (s_i, ss_i, \mathcal{F}_i), \texttt{skip}
$$

that produces the same final fingerprint in the end.

PROOF. Given the internal steps of a transaction

$$
\mathcal{K}_0, c_0, C_0, s_0, [\mathsf{T}_0]_c^{\mathcal{F}_0} \Rightarrow \cdots \Rightarrow \mathcal{K}_i, c_i, C_i, s_i, [\mathsf{T}_i]_c^{\mathcal{F}_i}
$$

We construct the following trace,

$$
(s_0, \mathrm{snapshot}(\mathcal{K}_0, \mathrm{viewOf}(\mathcal{K}_0, c)), \mathcal{F}_0), \mathsf{T}_0 \to^* (s_i, ss_i, \mathcal{F}_i), \mathsf{T}_i
$$

Let consider how many transactional internal steps.

- Base case: i = 0. In this case,

$$
\mathcal{K}_0, c_0, C_0, s_0, [\mathsf{T}]_c^{\mathcal{F}_0}
$$

  It is easy to construct the following

$$
(s_0, \mathrm{snapshot}(\mathcal{K}, \mathrm{viewOf}(\mathcal{K}, c)), \mathcal{F}_0), \mathsf{T}_0
$$

- Inductive case: i + 1. Suppose a trace with i steps,

$$
\mathcal{K}_0, c_0, C_0, s_0, [\mathsf{T}_0]_c^{\mathcal{F}_0} \Rightarrow \cdots \Rightarrow \mathcal{K}_i, c_i, C_i, s_i, [\mathsf{T}_i]_c^{\mathcal{F}_i}
$$

  and a trace

$$
(s_0, \mathrm{snapshot}(\mathcal{K}_0, \mathrm{viewOf}(\mathcal{K}_0, c)), \mathcal{F}_0), \mathsf{T}_0 \to^* (s_i, ss_i, \mathcal{F}_i), \mathsf{T}_i
$$

  Now let consider the next step.

– ReadTrans. In this case

$$\mathcal{K}_i, c_i, C_i, s_i, [\mathsf{T}_i]_c^{\mathcal{F}_i} \rightarrow \mathcal{K}_{i+1}, c_{i+1}, C_{i+1}, s_{i+1}, [\mathsf{T}_{i+1}]_c^{\mathcal{F}_{i+1}}$$

such that

$$\mathcal{F}_{i+1} = \mathcal{F}_i \lessdot (\mathsf{r}, k, v) = \mathcal{F}_i \wedge (\mathsf{w}, k, v) \in \mathcal{F}_i$$

for some key k and value v, and

$$\mathsf{T}_i \equiv \mathsf{x} := [\mathsf{E}]; \mathsf{T} \wedge [\![\mathsf{E}]\!]_{s_i} = k \wedge s_{i+1} = s_i[\mapsto v] \wedge \mathsf{T}_{i+1} \equiv \mathsf{skip}; \mathsf{T}$$

for some variable x, expression E and continuation T. Since $(\mathsf{w}, k, v) \in \mathcal{F}_i$, it means $_i(k) = v$ for the local snapshot. By the TPrimitive, we have

$$(s_i, ss_i, \mathcal{F}_i), \mathsf{T}_i \rightarrow \mathsf{x} := [\mathsf{E}]; \mathsf{T} \rightarrow (s_{i+1}, ss_i, \mathcal{F}_{i+1}), \mathsf{T}_{i+1}$$

– ReadRemote. In this case

$$\mathcal{K}_i, c_i, C_i, s_i, [\mathsf{T}_i]_c^{\mathcal{F}_i} \rightarrow \mathcal{K}_{i+1}, c_{i+1}, C_{i+1}, s_{i+1}, [\mathsf{T}_{i+1}]_c^{\mathcal{F}_{i+1}}$$

such that

$$\mathcal{F}_{i+1} = \mathcal{F}_i \lessdot (\mathsf{r}, k, v) = \mathcal{F}_i \uplus \{(\mathsf{r}, k, v)\} \wedge \forall v. \ (\mathsf{w}, k, v') \notin \mathcal{F}_i$$

for some key k and value v, and

$$\mathsf{T}_i \equiv \mathsf{x} := [\mathsf{E}]; \mathsf{T} \wedge [\![\mathsf{E}]\!]_{s_i} = k \wedge s_{i+1} = s_i[\mapsto v] \wedge \mathsf{T}_{i+1} \equiv \mathsf{skip}; \mathsf{T}$$

for some variable x, expression E and continuation T. By the premiss of the ReadRemote, the value read is from the last version before the snapshot time:

$$n = \max \left\{ n' \ \middle| \ \exists j. \ t^{n'} = \mathsf{w}(\mathcal{K}(k, j)) \wedge n' < c \right\} \wedge \mathsf{val}(\mathcal{K}_i(k, n)) = v$$

By the definition of $u_0 = \mathsf{viewOf}(\mathcal{K}_0, c)$ and $\mathsf{snapshot}(\mathcal{K}_0, u_0)$ and the fact that there is no write to the key k, it follows $ss_i(k) = v$. Thus, by the TPrimitive, we have

$$(s_i, ss_i, \mathcal{F}_i), \mathsf{T}_i \rightarrow \mathsf{x} := [\mathsf{E}]; \mathsf{T} \rightarrow (s_{i+1}, ss_i, \mathcal{F}_{i+1}), \mathsf{T}_{i+1}$$

– Write. In this case

$$\mathcal{K}_i, c_i, C_i, s_i, [\mathsf{T}_i]_c^{\mathcal{F}_i} \rightarrow \mathcal{K}_{i+1}, c_{i+1}, C_{i+1}, s_{i+1}, [\mathsf{T}_{i+1}]_c^{\mathcal{F}_{i+1}}$$

such that

$$\mathcal{F}_{i+1} = \mathcal{F}_i \lessdot (\mathsf{w}, k, v) = \mathcal{F}_i \setminus \{(\mathsf{w}, k, v') \mid v' \in \mathsf{Val}\} \uplus \{(\mathsf{r}, k, v)\}$$

for some key k and value v, and

$$\mathsf{T}_i \equiv [\mathsf{E}_1] := \mathsf{E}_2; \mathsf{T} \wedge [\![\mathsf{E}_1]\!]_{s_i} = k \wedge [\![\mathsf{E}_2]\!]_{s_i} = v \wedge \mathsf{T}_{i+1} \equiv \mathsf{skip}; \mathsf{T}$$

for some expressions E_1 and E_2, and continuation T. By the TPrimitive, it is easy to see:

$$(s_i, ss_i, \mathcal{F}_i), \mathsf{T}_i \rightarrow [\mathsf{E}_1] := \mathsf{E}_2; \mathsf{T} \rightarrow (s_{i+1}, ss_i, \mathcal{F}_{i+1}), \mathsf{T}_{i+1}$$

□

By Def. H.11, Lemma H.12, and Theorem H.15, we know for each clock-SI trace, there exists a trace that satisfies $\mathsf{ET}_\bot$. Last, we prove such trace also satisfies $\mathsf{ET}_{\mathsf{SI}}$.

Theorem H.16 (Clock-SI satisfying SI). *For any normal trace clock-SI trace τ, and transaction t_{cl}^n such that*

$$\tau = \cdots \xrightarrow{cl, c, \mathcal{F}, \bot} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}' \rightarrow \cdots$$

the transaction satisfies $\mathsf{ET}_{\mathsf{SI}}$, i.e. $\mathsf{ET}_{\mathsf{SI}} \vdash (\mathcal{K}, \mathsf{viewOf}(\mathcal{K}, c)) \triangleright \mathcal{F} : \mathsf{viewOf}(\mathcal{K}, C''(cl))$

Proof. Recall $\mathsf{ET}_{\mathsf{SI}} = \{(\mathcal{K}, u, \mathcal{F}, u') \mid \dagger\} \cap \mathsf{ET}_{\mathsf{MR}} \cap \mathsf{ET}_{\mathsf{RYW}} \cap \mathsf{ET}_{\mathsf{UA}}$ Note that final view of the client, $C''(cl) = n + 1$. We prove the four parts separately.

- $\{(\mathcal{K}, \mathsf{viewOf}(\mathcal{K}, c), \mathcal{F}, \mathsf{viewOf}(\mathcal{K}, C''(cl))) \mid \dagger\}$. Assume a version $i \in \mathsf{viewOf}(\mathcal{K}, c)(k)$ for some key k. Suppose a version $\mathcal{K}(k', j)$ such that

$$\mathsf{w}(\mathcal{K}(k', j)) \xrightarrow{((\mathsf{SO} \cup \mathsf{WR}_{\mathcal{K}} \cup \mathsf{WW}_{\mathcal{K}}); \mathsf{RW}^?_{\mathcal{K}})^+} \mathsf{w}(\mathcal{K}(k, i))$$

  Let $t^n = \mathsf{w}(\mathcal{K}(k', j))$ and $t^m = \mathsf{w}(\mathcal{K}(k, i))$. By Lemmas H.17 and H.18, we know $n < m$ then $j \in \mathsf{viewOf}(\mathcal{K}, c)(k')$.

- $\mathsf{ET}_{\mathsf{MR}}$. By Commit, we know $c \leq n < C''(cl)$ then $\mathsf{viewOf}(\mathcal{K}, c) \sqsubseteq \mathsf{viewOf}(\mathcal{K}, C''(cl))$.

- $\mathsf{ET}_{\mathsf{MW}}$. By Commit, for any write $(\mathsf{w}, k, v) \in \mathcal{F}$, there is a new version written by the client cl in the $\mathcal{K}'$,

$$\mathsf{w}(\mathcal{K}'(k, |\mathcal{K}'(k)| - 1)) = t^n_{cl}$$

  Since $n < C''(cl)$, it follows $|\mathcal{K}'(k)| - 1 \in \mathsf{viewOf}(\mathcal{K}, C''(cl))(k)$.

- $\mathsf{ET}_{\mathsf{UA}}$. By the premiss of Commit, for any write $(\mathsf{w}, k, v) \in \mathcal{F}$, any existed versions of the key k must be installed by some transactions before the snapshot time of c,

$$\forall k, i. \ (\mathsf{w}, k, \_) \in \mathcal{F} \wedge \mathsf{w}(\mathcal{K}(k, i)) < c$$

It implies that

$$\forall i. \ i \in \mathsf{dom}(\mathcal{K}(k)) \Rightarrow i \in \mathsf{viewOf}(\mathcal{K}, c)(k)$$

$\square$

LEMMA H.17 ($\mathsf{RW}_{\mathcal{K}}$). *Given a normal clock-SI trace τ, and two transactions t^n_{cl} and $t^m_{cl'}$ from the trace*

$$\tau = \cdots \rightarrow \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \cdots \wedge \tau = \cdot \rightarrow \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}' \xrightarrow{cl', c', \mathcal{F}, m} \cdots$$

Suppose the final state of the trace τ is $\mathcal{K}''$. , if $t^n_{cl} \xrightarrow{\mathsf{RW}^?_{\mathcal{K}''}} t^m_{cl'}$ then the snapshot time of t^n_{cl} took snapshot before the commit time of $t^m_{cl'}$, i.e. $c \leq m$.

Proof. By definition of $t^n_{cl} \xrightarrow{\mathsf{RW}^?_{\mathcal{K}''}} t^m_{cl'}$, it follows that

$$t^n_{cl} \in \mathsf{rs}(\mathcal{K}''(k, i)) \wedge t^m_{cl'} = \mathsf{w}(\mathcal{K}''(k, j)) \wedge i < j$$

for some key k and indexes i, j. There are two cases depending on the commit order.

- If t^n_{cl} commits after $t^m_{cl'}$, we have,

$$\tau = \cdot \rightarrow \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}' \xrightarrow{cl', c', \mathcal{F}, m} \cdots \rightarrow \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \cdots$$

  We prove by contradiction. Assume $c > m$. Since it is a normal trace τ (Theorem H.8), it follows $n > m$. Note that both transactions access the key k, and then by Lemmas H.20 and H.21, we have $n > c > m$. Given $c > m$, by ReadRemote the transaction t^n_{cl} should at least read the version written by $t^m_{cl'}$ for the key k. That is,

$$t^n_{cl} \in \mathsf{rs}(\mathcal{K}''(k, i)) \wedge t^m_{cl'} = \mathsf{w}(\mathcal{K}''(k, j)) \wedge i > j$$

  which contradict $t^n_{cl} \xrightarrow{\mathsf{RW}^?_{\mathcal{K}''}} t^m_{cl'}$.

- If t^n_{cl} commits before $t^m_{cl'}$,

$$\tau = \cdots \rightarrow \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \cdots \rightarrow \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}' \xrightarrow{cl', c', \mathcal{F}, m} \cdots$$

  It is trivial that $c \leq m$ by Lemmas H.20 and H.21.

□

LEMMA H.18 (WR$_\mathcal{K}$, WW$_\mathcal{K}$ AND SO$_\mathcal{K}$). *Given a normal clock-SI trace τ, and two transactions t_{cl}^n and $t_{cl'}^m$ from the trace*

$$\tau = \cdots \overset{}{\to} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \cdots \wedge \tau = \cdots \overset{}{\to} \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}' \xrightarrow{cl', c', \mathcal{F}, m} \cdots$$

Suppose the final state of the trace τ is $\mathcal{K}''$., if $t_{cl}^n \xrightarrow{\mathrm{WR}_{\mathcal{K}''}^?} t_{cl'}^m$ then the transaction t_{cl}^n commit before the commit time of $t_{cl'}^m$, i.e. $n < m$. Similarly, $n < m$ for the relations WW$_\mathcal{K}$ and SO$_\mathcal{K}$.

PROOF. • WR$_{\mathcal{K}''}$. Since $t_{cl}^n \xrightarrow{\mathrm{WR}_{\mathcal{K}''}^?} t_{cl'}^m$, it is only possible that the later commits after the former,

$$\tau = \cdots \overset{}{\to} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \cdots \wedge \to \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}' \xrightarrow{cl', c', \mathcal{F}, m} \cdots$$

By Lemma H.19, we know $n < m$.
• WW$_{\mathcal{K}''}$. By the definition of WW$_{\mathcal{K}''}$ and Lemma H.14, we know $n < m$.
• SO$_{\mathcal{K}''}$. By the definition of SO$_{\mathcal{K}''}$ and Lemma H.9, we know $n < m$.

□

LEMMA H.19 (READER GREATER THAN WRITER). *Assume a trace τ and two transactions t_{cl}^n and $t_{cl'}^m$,*

$$\tau = \cdots \overset{}{\to} \mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \cdots \wedge \to \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}' \xrightarrow{cl', c', \mathcal{F}, m} \cdots$$

Assume the final state of key-value store of the trace is $\mathcal{K}''$. If $t_{cl'}^m$ reads a version written by t_{cl}^n

$$\mathsf{w}(\mathcal{K}''(k, i)) = t^n \wedge t^m \in \mathsf{rs}(\mathcal{K}''(k, j))$$

Then, the snapshot times of readers of a version is greater then the commit time of the writer $n < c'$

PROOF. Trivially, $\mathsf{w}(\mathcal{K}'(k, i)) = t^n$. By the READREMOTE, it follows

$$n = \max \left\{ n' \ \middle| \ \exists j. \ t^{n'} = \mathsf{w}(\mathcal{K}(k, j)) \wedge n' < c' \right\}$$

which implies $n < c'$.

□

LEMMA H.20 (COMMIT TIME AFTER SNAPSHOT TIME). *The commit time of a transaction is after the snapshot time. Suppose the following step,*

$$\mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}'$$

then $c < n$.

PROOF. It is easy to see by CLIENTSTEP and then COMMIT that

$$n > n - 1 \max \left(\{ c' \mid \exists k. \ (\_, k, \_) \in \mathcal{F} \wedge c' = C'(\mathsf{shardOf}(k)) \} \cup \{ c \} \right)$$

so $c < n$.

□

LEMMA H.21 (MONOTONIC SHARD CLOCK TIME). *The clock time associated with a shard monotonically increases, Suppose the following step,*

$$\mathcal{K}, C, C', \mathcal{E}, \mathsf{P} \xrightarrow{cl, c, \mathcal{F}, n} \mathcal{K}', C'', C''', \mathcal{E}', \mathsf{P}'$$

then

$$\forall r \in \mathsf{dom}(C'). \ C'(r) \leq C'''(r)$$

Proof. We perform case analysis on rules.

- TimeTick By the rule there is one shard r' ticks time $C'''(r') = C'(r) + 1 > C'(r)$.
- ClientStep. There are further five cases, yet only StartTrans and Commit change the shard's clock times.
  - StartTrans By the rule a new transaction starts in a shard r' and triggering the shard r' ticks time $C'''(r') = C'(r) + 1 > C'(r)$.
  - Commit By the rule the transaction commits their fingerprint $\mathcal{F}$ to those shards r' it read or write, and triggering the shard r' ticks time $C'''(r') = C'(r) + 1 > C'(r)$.

□