

Dynamic Partial Order Reduction for Checking Correctness Against Transaction Isolation Levels

Modern applications, such as social networking systems and e-commerce platforms are centered around using large-scale databases for storing and retrieving data. Accesses to the database are typically enclosed in transactions that allow computations on shared data to be isolated from other concurrent computations and resilient to failures. Modern databases trade isolation for performance. The weaker the isolation level is, the more behaviors a database is allowed to exhibit and it is up to the developer to ensure that their application can tolerate those behaviors.

In this work, we propose stateless model checking algorithms for studying correctness of such applications that rely on dynamic partial order reduction. These algorithms work for a number of widely-used weak isolation levels, including Read Committed, Causal Consistency, Snapshot Isolation, and Serializability. We show that they are complete, sound and optimal, and run with polynomial memory consumption in all cases. We report on an implementation of these algorithms in the context of Java Pathfinder applied to a number of challenging applications drawn from the literature of distributed systems and databases.

1 INTRODUCTION

Data storage is no longer about writing data to a single disk with a single point of access. Modern applications require not just data reliability, but also high-throughput concurrent accesses. Applications concerning supply chains, banking, etc. use traditional relational databases for storing and processing data, whereas applications such as social networking software and e-commerce platforms use cloud-based storage systems (such as Azure Cosmos DB [52], Amazon DynamoDB [29], Facebook TAO [20], etc.).

Providing high-throughput processing, unfortunately, comes at an unavoidable cost of weakening the consistency guarantees offered to users: Concurrently-connected clients may end up observing different versions of the same data. These “anomalies” can be prevented by using a strong *isolation level* such as *Serializability* [50], which essentially offers a single version of the data to all clients at any point in time. However, serializability requires expensive synchronization and incurs a high performance cost. As a consequence, most storage systems use weaker isolation levels, such as *Causal Consistency* [9, 42, 44], *Snapshot Isolation* [14], *Read Committed* [14], etc. for better performance. In a recent survey of database administrators [51], 86% of the participants responded that most or all of the transactions in their databases execute at Read Committed level.

A weaker isolation level allows for more possible behaviors than stronger isolation levels. It is up to the developers then to ensure that their application can tolerate this larger set of behaviors. Unfortunately, weak isolation levels are hard to understand or reason about [6, 21] and resulting application bugs can cause loss of business [60].

Model Checking Database-Backed Applications. This paper addresses the problem of *model checking* code for correctness against a given isolation level. *Model checking* [27, 54] explores the state space of a given program in a systematic manner and it provides high coverage of program behavior. However, it faces the infamous state explosion problem, i.e., the number of executions grows exponentially in the number of concurrent clients.

Partial order reduction (POR) [28, 34, 53, 58] is an approach that limits the number of explored executions without sacrificing coverage. POR relies on an equivalence relation between executions where e.g., two executions are equivalent if one can be obtained from the other by swapping consecutive independent (non-conflicting) execution steps. It guarantees that at least one execution from each equivalence class is explored. *Optimal* POR techniques explore exactly one execution from each equivalence class. Beyond this classic notion of optimality, POR techniques may aim

for optimality by avoiding visiting states from which the exploration is blocked. *Dynamic* partial order reduction (DPOR) [32] has been introduced to explore the execution space (and tracking the equivalence relation between executions) on-the-fly without relying on a-priori static analyses. This is typically coupled with *stateless* model checking (SMC) [35] which explores executions of a program without storing visited states, thereby, avoiding excessive memory consumption.

There is a large body of work on (D)POR techniques that address their soundness when checking a certain class of specifications for a certain class of programs, as well as their completeness and their theoretical optimality (see Section 8). Most often these works consider shared memory concurrent programs executing under a strongly consistent memory model.

In the last few years, some works have studied DPOR in the case of shared memory programs running under weak memory models such as TSO or Release-Acquire, e.g. [1, 4, 5, 40]. While these algorithms are sound and complete, they have exponential space complexity when they are optimal. More recently, Kokologianakis et al. [39] defined a DPOR algorithm that has a polynomial space complexity, in addition of being sound, complete and optimal. This algorithm can be applied for a range of shared memory models.

While the works mentioned above concern shared memory programs, we are not aware of any published work addressing the case of database transactional programs running under weak isolation levels. In this paper, we address this case and propose new stateless model checking algorithms relying on DPOR techniques for database-backed applications. We assume that all the transactions in an application execute under the *same* isolation level, which happens quite frequently in practice (as mentioned above, most database applications are run on the default isolation level of the database). Our work generalizes the approach introduced by [39]. However, this generalization to the transactional case, covering the most relevant isolation levels, is not a straightforward adaptation of [39]. Ensuring optimality while preserving the other properties, e.g., completeness and polynomial memory complexity, is very challenging. In the following, we explain the main steps and features of our work.

Formalizing Isolation Levels. Our algorithms rely on the axiomatic definitions of isolation levels introduced by Biswas and Enea [16]. These definitions use logical constraints called *axioms* to characterize the set of executions of a database (e.g., key-value store) that conform to a particular isolation level (this can be extended to SQL queries [17]). These constraints refer to a specific set of relations between events/transactions in an execution that describe control-flow or data-flow dependencies: a program order **po** between events in the same transaction, a session order **so** between transactions in the same session¹, and a write-read **wr** (read-from) relation that associates each read event with a transaction that writes the value returned by the read. These relations along with the events in an execution are called a *history*. A history describes only the interaction with the database, omitting application-side events (e.g., computing values written to the database).

Execution Equivalence. DPOR algorithms are parametrized by an equivalence relation on executions, most often, Mazurkiewicz equivalence [45]. In this work, we consider a weaker equivalence relation, also known as *read-from equivalence* [3, 5, 25, 39–41], which considers that two executions are equivalent when their histories are precisely the same (they contain the same set of events, and the relations **po**, **so**, and **wr** are the same). In general, reads-from equivalence is coarser than Mazurkiewicz equivalence, and its equivalence classes can be exponentially-smaller than Mazurkiewicz traces in certain cases [25].

SMC Algorithms. Our SMC algorithms enumerate executions of a given program under a given isolation level I . They are *sound*, i.e., enumerate only *feasible* executions (admitted by the program under I), *complete*, i.e., they output a representative of each read-from equivalence class, and *optimal*,

¹A session is a sequential interface to the storage system. It corresponds to what is also called a *connection*.

99 i.e., they output *exactly one* complete execution from each read-from equivalence class. For isolation
 100 levels weaker than and including Causal Consistency, they satisfy a notion of *strong optimality*
 101 which says that additionally, the enumeration avoids states from which the execution is “blocked”,
 102 i.e., it cannot be extended to a complete execution of the program. For Snapshot Isolation and
 103 Serializability, we show that *there exists* no algorithm in the same class (to be discussed below) that
 104 can ensure such a strong notion of optimality. All the algorithms that we propose are polynomial
 105 space, as opposed to many DPOR algorithms introduced in the literature.

106 As a starting point, we define a generic class of SMC algorithms, called *swapping based*, general-
 107 izing the approach adopted by [39, 40], which enumerate histories of program executions. These
 108 algorithms focus on the interaction with the database assuming that the other steps in a transaction
 109 concern local variables visible only within the scope of the enclosing session. Executions are
 110 extended according to a generic scheduler function NEXT and every read event produces several
 111 exploration branches, one for every write executed in the past that it can read from. Events in
 112 an execution can be swapped to produce new exploration “roots” that lead to different histories.
 113 Swapping events is required for completeness, to enumerate histories where a read r reads from
 114 a write w that is scheduled by NEXT after r . To ensure soundness, we restrict the definition of
 115 swapping so that it produces a history that is feasible by construction (extending an execution which
 116 is possibly infeasible may violate soundness). Such an algorithm is optimal w.r.t. the read-from
 117 equivalence when it enumerates each history exactly once.

118 We define a concrete algorithm in this class that in particular, satisfies the stronger notion of
 119 optimality mentioned above for every isolation level I which is *prefix-closed* and *causally-extensible*,
 120 e.g., *Read Committed* and *Causal Consistency*. Prefix-closure means that every prefix of a history
 121 that satisfies I , i.e., a subset of transactions and all their predecessors in the causal relation, i.e.,
 122 $(so \cup wr)^+$, is also consistent with I , and causal extensibility means that any pending transaction
 123 in a history that satisfies I can be extended with one more event to still satisfy I , and if this is a
 124 read event, then, it can read-from a transaction that precedes it in the causal relation. To ensure
 125 strong optimality, this algorithm uses a carefully chosen condition for restricting the application of
 126 event swaps, which makes the proof of completeness in particular, quite non-trivial.

127 We show that isolation levels such as Snapshot Isolation and Serializability are *not* causally-
 128 extensible and that there exists no swapping based SMC algorithm which is sound, complete, and
 129 strongly optimal at the same time (independent of memory consumption bounds). This impossibility
 130 proof uses a program to show that any NEXT scheduler and any restriction on swaps would violate
 131 either completeness or strong optimality. However, we define an extension of the previous algorithm
 132 which satisfies the weaker notion of optimality, while preserving soundness, completeness, and
 133 polynomial space complexity. This algorithm will simply enumerate executions according to a
 134 weaker prefix-closed and causally-extensible isolation level, and filter executions according to the
 135 stronger isolation levels Snapshot Isolation and Serializability at the end, before outputting.

136 We implemented these algorithms in the Java Pathfinder (JPF) model checker [59], and evaluated
 137 them on a number of challenging database-backed applications drawn from the literature of
 138 distributed systems and databases.

139 Our contributions and outline are summarized as follows:

- 140 § 3 identifies a class of isolation levels called prefix-closed and causally-extensible that admit
 141 efficient SMC.
- 142 § 4 defines a generic class of swapping based SMC algorithms based on DPOR which are
 143 parametrized by a given isolation level.
- 144 § 5 defines a swapping based SMC algorithm which is sound, complete, strongly-optimal, and
 145 polynomial space, for any isolation level that is prefix-closed and causally-extensible.
 146

$$\begin{array}{ll}
 x \in \text{Vars} & a \in \text{LVars} \\
 \text{Prog} ::= \text{Sess} \mid \text{Sess} \parallel \text{Prog} & \text{Body} ::= \text{Instr} \mid \text{Instr}; \text{Body} \\
 \text{Sess} ::= \text{Trans} \mid \text{Trans}; \text{Sess} & \text{Instr} ::= \text{InstrDB} \mid a := e \mid \text{if}(\phi(\vec{a}))\{\text{Instr}\} \\
 \text{Trans} ::= \text{begin}; \text{Body}; \text{commit} & \text{InstrDB} ::= a := \text{read}(x) \mid \text{write}(x, a) \mid \text{abort}
 \end{array}$$

Fig. 1. Program syntax. The set of global variables is denoted by Vars while LVars denotes the set of local variables. We use ϕ to denote Boolean expressions over local variables, and e to denote expressions over local variables interpreted as values. We use $\vec{\cdot}$ to denote vectors of elements.

§ 6 shows that there exists no swapping based algorithm for Snapshot Isolation and Serializability, which is sound, complete, and strongly-optimal at the same time, and proposes a swapping based algorithm which satisfies “plain” optimality.

§ 7 reports on an implementation and evaluation of these algorithms.

Section 2 recalls the formalization of isolation levels of Biswas and Enea [16, 17], while Sections 8 and 9 conclude with a discussion of related work and concluding remarks. Additional formalization, proofs, and experimental data can be found in the technical report [18].

2 TRANSACTIONAL PROGRAMS

2.1 Program Syntax

Figure 1 lists the definition of a simple programming language that we use to represent applications running on top of a database. A program is a set of *sessions* running in parallel, each session being composed of a sequence of *transactions*. Each transaction is delimited by `begin` and either `commit` or `abort` instructions, and its body contains instructions that access the database and manipulate a set LVars of local variables. We use symbols a, b , etc. to denote elements of LVars.

For simplicity, we abstract the database state as a valuation to a set Vars of *global variables*², ranged over using x, y , etc. The instructions accessing the database correspond to reading the value of a global variable and storing it into a local variable a ($a := \text{read}(x)$), writing the value of a local variable a to a global variable x ($\text{write}(x, a)$), or an assignment to a local variable a ($a := e$). The set of values of global or local variables is denoted by Vals. Assignments to local variables use expressions e over local variables, which are interpreted as values and whose syntax is left unspecified. Each of these instructions can be guarded by a Boolean condition $\phi(\vec{a})$ over a set of local variables \vec{a} (their syntax is not important). Our results assume bounded programs, as usual in SMC algorithms, and therefore, we omit other constructs like `while` loops. SQL statements (SELECT, JOIN, UPDATE) that manipulate relational tables can be compiled to reads or writes of variables that represent fields or rows in a table (see for instance, [17, 55]).

2.2 Isolation Levels

We present the axiomatic framework introduced by Biswas and Enea [16] for defining isolation levels. Isolation levels are defined as logical constraints, called *axioms*, over *histories*, which are an abstract representation of the interaction between a program and the database in an execution.

2.2.1 Histories. Programs interact with a database by issuing transactions formed of `begin`, `commit`, `abort`, `read` and `write` instructions. The effect of executing one such instruction is represented using an *event* (e, type) where e is an *identifier* and type is a *type*. There are five types of events: `begin`, `commit`, `abort`, `read(x)` for reading the global variable x , and `write(x, v)` for writing value v to x . \mathcal{E} denotes the set of events. For a read/write event e , we use $\text{var}(e)$ to denote the variable x .

²In the context of a relational database, global variables correspond to fields/rows of a table while in the context of a key-value store, they correspond to keys.

197 A *transaction log* $\langle t, E, \text{po}_t \rangle$ is an identifier t and a finite set of events E along with a strict
 198 total order po_t on E , called *program order* (representing the order between instructions in the
 199 body of a transaction). The minimal element of po_t is a begin event. A transaction log without
 200 neither a commit nor an abort event is called *pending*. Otherwise, it is called *complete*. A complete
 201 transaction log with a commit event is called *committed* and *aborted* otherwise. If a commit or an
 202 abort event occurs, then it is maximal in po_t ; commit and abort cannot occur simultaneously in
 203 the same transaction log. The set E of events in a transaction log t is denoted by $\text{events}(t)$. Note
 204 that a transaction is aborted because it executed an abort instruction. Histories do not include
 205 transactions aborted by the database because their effect should not be visible to other transactions
 206 and the abort is not under the control of the program. For simplicity, we may use the term *transaction*
 207 instead of transaction log.

208 Isolation levels differ in the values returned by read events which are not preceded by a write on
 209 the same variable in the same transaction. We assume in the following that every transaction in a
 210 program is executed under the same isolation level. For every isolation level that we are aware of,
 211 if a read of a global variable x is preceded by a write to x in po_t , then it should return the value
 212 written by the last write to x before the read (w.r.t. po_t).

213 The set of $\text{read}(x)$ events in a transaction log t that are *not* preceded by a write to x in po_t , for
 214 some x , is denoted by $\text{reads}(t)$. Also, if t does *not* contain an abort event, the set of $\text{write}(x, _)$
 215 events in t that are *not* followed by other writes to x in po_t , for some x , is denoted by $\text{writes}(t)$.
 216 If a transaction contains multiple writes to the same variable, then only the last one (w.r.t. po_t)
 217 can be visible to other transactions (w.r.t. any isolation level that we are aware of). If t contains
 218 an abort event, then we define $\text{writes}(t)$ to be the empty set. This is because the effect of aborted
 219 transactions (its set of writes) should not be visible to other transactions. The extension to sets
 220 of transaction logs is defined as usual. Also, we say that a transaction log t *writes* x , denoted by
 221 t *writes* x , when $\text{writes}(t)$ contains some $\text{write}(x, _)$ event.

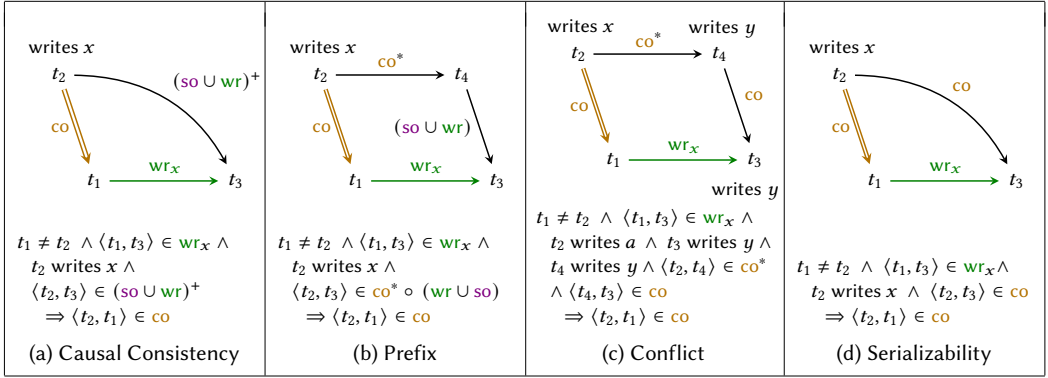
222 A *history* contains a set of transaction logs (with distinct identifiers) ordered by a (partial) *session*
 223 *order* so that represents the order between transactions in the same session. It also includes a
 224 *write-read* relation (also called *read-from*) that defines read values by associating each read to a
 225 transaction that wrote that value. Read events do *not* contain a value, and their return value is
 226 defined as the value written by the transaction associated by the write-read relation. Let T be a
 227 set of transaction logs. For a write-read relation $\text{wr} \subseteq \text{writes}(T) \times \text{reads}(T)$ and variable x , wr_x is
 228 the restriction of wr to reads of x , $\text{wr}_x = \text{wr} \cap (\text{writes}(T) \times \{e \mid e \text{ is a read}(x) \text{ event}\})$. We extend
 229 the relations wr and wr_x to pairs of transactions by $\langle t_1, t_2 \rangle \in \text{wr}$, resp., $\langle t_1, t_2 \rangle \in \text{wr}_x$, iff there
 230 exists a $\text{write}(x, _)$ event w in t_1 and a $\text{read}(x)$ event r in t_2 s.t. $\langle w, r \rangle \in \text{wr}$, resp., $\langle w, r \rangle \in \text{wr}_x$.
 231 Analogously, wr and wr_x can be extended to tuples formed of a transaction (containing a write) and
 232 a read event. We say that the transaction log t_1 is *read* by the transaction log t_2 when $\langle t_1, t_2 \rangle \in \text{wr}$.

233
 234 *Definition 2.1.* A *history* $\langle T, \text{so}, \text{wr} \rangle$ is a set of transaction logs T along with a strict partial *session*
 235 *order* so , and a *write-read* relation $\text{wr} \subseteq \text{writes}(T) \times \text{reads}(T)$ such that

- 236 • the inverse of wr is a total function,
- 237 • if $\langle w, r \rangle \in \text{wr}$, then w and r are a write and respectively, a read, of the same variable, and
- 238 • $\text{so} \cup \text{wr}$ is acyclic (here we use the extension of wr to pairs of transactions).

239
 240 Every history includes a distinguished transaction writing the initial values of all global variables.
 241 This transaction precedes all the other transactions in so . We use h, h_1, h_2, \dots to range over histories.

242 The set of transaction logs T in a history $h = \langle T, \text{so}, \text{wr} \rangle$ is denoted by $\text{tr}(h)$, and $\text{events}(h)$ is the
 243 union of $\text{events}(t)$ for $t \in T$. For a history h and an event e in h , $\text{tr}(h, e)$ is the transaction t in h
 244 that contains e . Also, $\text{writes}(h) = \bigcup_{t \in \text{tr}(h)} \text{writes}(t)$ and $\text{reads}(h) = \bigcup_{t \in \text{tr}(h)} \text{reads}(t)$.



287 Fig. 2. Axioms defining isolations levels (all logical variables representing transactions, e.g., t_1 , are universally quantified). The reflexive and transitive, resp., transitive, closure of a relation rel is denoted by rel^* , resp., rel^+ . Also, \circ denotes the composition of two relations, i.e., $rel_1 \circ rel_2 = \{\langle a, b \rangle \mid \exists c. \langle a, c \rangle \in rel_1 \wedge \langle c, b \rangle \in rel_2\}$.


288 We extend so to pairs of events by $(e_1, e_2) \in \text{so}$ if $(\text{tr}(h, e_1), \text{tr}(h, e_2)) \in \text{so}$. We also define $\text{po} = \bigcup_{t \in T} \text{po}_t$.

289 2.2.2 *Axiomatic Framework.* A history satisfies a certain isolation level if there is a strict total order co on its transactions, called *commit order*, which extends the write-read relation and the session order, and which satisfies certain properties. These properties, called *axioms*, relate the commit order with the so and wr relations in a history and are defined as first-order formulas of the form:

$$290 \quad \forall x, \forall t_1 \neq t_2, \forall t_3. \\ 291 \quad \langle t_1, t_3 \rangle \in wr_x \wedge t_2 \text{ writes } x \wedge \phi(t_2, t_3) \Rightarrow \langle t_2, t_1 \rangle \in \text{co} \quad (1)$$

292 where ϕ is a property relating t_2 and τ (i.e., the read or the transaction reading from t_1) that varies from one axiom to another.³ Note that an aborted transaction t cannot take the role of t_1 nor t_2 in equation 1 as the set $\text{writes}(t)$ is empty. Intuitively, this axiom schema states the following: in order for τ to read specifically t_1 's write on k , it must be the case that every t_2 that also writes k and satisfies $\phi(t_2, \tau)$ was committed before t_1 . The property ϕ relates t_2 and τ using the relations in a history and the commit order. Figure 2 shows two axioms which correspond to their homonymous isolation levels: *Causal Consistency* (CC) and *Serializability* (SER). The conjunction of the other two axioms Conflict and Prefix defines *Snapshot Isolation* (SI). *Read Atomic* (RA) is a weakening of CC where $(\text{so} \cup \text{wr})^+$ is replaced with $\text{so} \cup \text{wr}$. *Read Committed* (RC) is defined similarly. Note that SER is stronger than SI (i.e., every history satisfying SER satisfies SI as well), SI is stronger than CC, CC is stronger than RA, and RA is stronger than RC.

293 For instance, the axiom defining Causal Consistency [42] states that for any transaction t_1 writing a variable x that is read in a transaction t_3 , the set of $(\text{wr} \cup \text{so})^+$ predecessors of t_3 writing x must precede t_1 in commit order ($(\text{wr} \cup \text{so})^+$ is usually called the *causal order*). A violation of this axiom can be found in Figure 3: the transaction t_2 writing 2 to x is a $(\text{wr} \cup \text{so})^+$ predecessor of the transaction t_3 reading 1 from x because the transaction t_4 , writing 1 to y , reads x from t_2 and t_3 reads y from t_4 .



294 Fig. 3. Causal Consistency violation. Boxes group events from the same transaction.

295 ³These formulas are interpreted on tuples $\langle h, \text{co} \rangle$ of a history h and a commit order co on the transactions in h as usual.

This implies that t_2 should precede in commit order the transaction t_1 writing 1 to x , which is inconsistent with the write-read relation (t_2 reads from t_1).

The Serializability axiom requires that for any transaction t_1 writing to a variable x that is read in a transaction t_3 , the set of **co** predecessors of t_3 writing x must precede t_1 in commit order. This ensures that each transaction observes the effects of all the **co** predecessors.

Definition 2.2. For an isolation level I defined by a set of axioms X , a history $h = \langle T, \text{so}, \text{wr} \rangle$ satisfies I iff there is a strict total order **co** s.t. $\text{wr} \cup \text{so} \subseteq \text{co}$ and $\langle h, \text{co} \rangle$ satisfies X .

A history that satisfies an isolation level I is called I -consistent. For two isolation levels I_1 and I_2 , I_1 is *weaker than* I_2 when every I_1 -consistent history is also I_2 -consistent.

2.3 Program Semantics

We define a small-step operational semantics for transactional programs, which is parametrized by an isolation level I . The semantics keeps a history of previously executed database accesses in order to maintain consistency with I .

For readability, we define a program as a partial function $P : \text{SessId} \rightarrow \text{Sess}$ that associates session identifiers in SessId with concrete code as defined in Figure 1 (i.e., sequences of transactions). Similarly, the session order **so** in a history is defined as a partial function $\text{so} : \text{SessId} \rightarrow \text{Tlogs}^*$ that associates session identifiers with sequences of transaction logs. Two transaction logs are ordered by **so** if one occurs before the other in some sequence $\text{so}(j)$ with $j \in \text{SessId}$.


The operational semantics is defined as a transition relation \Rightarrow_I between *configurations*, which are defined as tuples containing the following:

- history h storing the events generated by database accesses executed in the past,
- a valuation map \vec{v} that records local variable values in the current transaction of each session (\vec{v} associates identifiers of sessions with valuations of local variables),
- a map \vec{B} that stores the code of each live transaction (mapping session identifiers to code),
- sessions/transactions P that remain to be executed from the original program.

The relation \Rightarrow_I is defined using a set of rules as expected. Starting a new transaction in a session j is enabled as long as this session has no live transactions ($\vec{B}(j) = \epsilon$) and results in adding a transaction log with a single begin event to the history and scheduling the body of the transaction (adding it to $\vec{B}(j)$). Local steps, i.e., checking the truth value of a Boolean condition or computation with local variables, manipulate the local variable valuations and advance the code as expected. Read instructions of some global variable x can have two possible behaviors: (1) if the read follows a write on x in the same transaction, then it returns the value written by the last write on x in that transaction, and (2) otherwise, the read reads from another transaction t' which is chosen non-deterministically as long as extending the current history with the write-read dependency associated to this choice leads to a history that still satisfies I . Depending on the isolation level, there may not exist a transaction t' the read can read from. For other instructions, e.g., `commit` and `abort`, the history is simply extended with the corresponding events while ending the transaction execution in the case of `abort`.


An *initial* configuration for program P contains the program P along with a history $h = \langle \{t_0\}, \emptyset, \emptyset \rangle$, where t_0 is a transaction log containing only writes that write the initial values of all variables, and empty current transaction code ($B = \epsilon$). An execution of a program P under an isolation level I is a sequence of configurations $c_0 c_1 \dots c_n$ where c_0 is an initial configuration for P , and $c_m \Rightarrow_I c_{m+1}$, for every $0 \leq m < n$. We say that c_n is *I-reachable* from c_0 . The history of such an execution is the history h in the last configuration c_n . A configuration is called *final* if it contains the empty

344
345
346
347
348
349



350 Fig. 4. Explaining the notion of prefix of a history. **init** denotes the transaction log writing initial values.
351 Boxes group events from the same transaction.

352
353
354
355
356
357



358 (a) Extensible history. (b) Non-extensible history.
359 Fig. 5. Explaining causal extensibility. **init** denotes the transaction log
360 writing initial values. Boxes group events from the same transaction.

361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392

Fig. 6. A counter-example to causal extensibility for SI and SER. The **so**-edges from **init** to the other transactions are omitted for legibility.

program ($P = \emptyset$). Let $\text{hist}_I(P)$ denote the set of all histories of an execution of P under I that ends in a final configuration.

3 PREFIX-CLOSED AND CAUSALLY-EXTENSIBLE ISOLATION LEVELS

We define two properties of isolation levels, prefix-closure and causal extensibility, which enable efficient DPOR algorithms (as shown in Section 5).

3.1 Prefix Closure

For a relation $R \subseteq A \times A$, the restriction of R to $A' \times A'$, denoted by $R \downarrow A' \times A'$, is defined by $\{(a, b) : (a, b) \in R, a, b \in A'\}$. Also, a set A' is called R -downward closed when it contains $a \in A$ every time it contains some $b \in A$ with $(a, b) \in R$.

A *prefix* of a transaction log $\langle t, E, \text{po}_t \rangle$ is a transaction log $\langle t, E', \text{po}_t \downarrow E' \times E' \rangle$ such that E' is po_t -downward closed. A *prefix* of a history $h = \langle T, \text{so}, \text{wr} \rangle$ is a history $h' = \langle T', \text{so} \downarrow T' \times T', \text{wr} \downarrow T' \times T' \rangle$ such that every transaction log in T' is a prefix of a different transaction log in T but carrying the same id, $\text{events}(h') \subseteq \text{events}(h)$, and $\text{events}(h')$ is $(\text{po} \cup \text{so} \cup \text{wr})^*$ -downward closed. For example, the history pictured in Fig. 4b is a prefix of the one in Fig. 4a while the history in Fig. 4c is not. The transactions on the bottom of Fig. 4c have a **wr** predecessor in Fig. 4a which is not included.

Definition 3.1. An isolation level I is called *prefix-closed* when every prefix of an I -consistent history is also I -consistent.

Every isolation level I discussed above is prefix-closed because if a history h is I -consistent with a commit order **co**, then the restriction of **co** to the transactions that occur in a prefix h' of h satisfies the corresponding axiom(s) when interpreted over h' .

THEOREM 3.2. *Read Committed, Read Atomic, Causal Consistency, Snapshot Isolation, and Serializability are prefix closed.*

3.2 Causal Extensibility

We start with an example to explain causal extensibility. Let us consider the histories h_1 and h_2 in Figures 5a and 5b, respectively, *without* the events $\text{read}(y)$ and $\text{write}(y, 2)$ written in blue bold font.

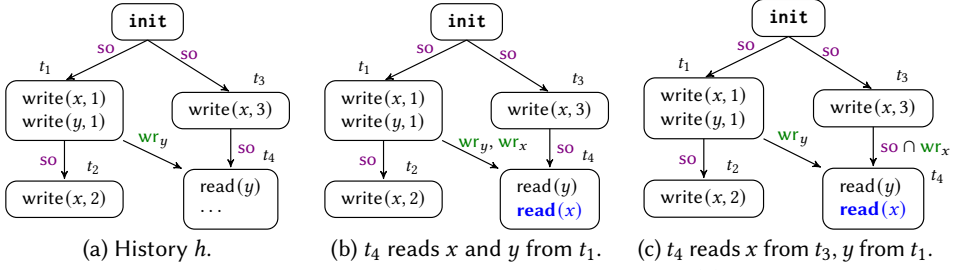


Fig. 7. Two causal extensions of the history h on the left with the $\text{read}(x)$ event written in blue.

These histories satisfy Read Atomic. The history h_1 can be extended by adding the event $\text{read}(y)$ and the wr dependency $\text{wr}(\text{init}, \text{read}(y))$ while still satisfying Read Atomic. On the other hand, the history h_2 can not be extended with the event $\text{write}(y, 2)$ while still satisfying Read Atomic. Intuitively, if the reading transaction on the bottom reads x from the transaction on the right, then it should read y from the same transaction because this is more “recent” than init w.r.t. session order. The essential difference between these two extensions is that the first concerns a transaction which is maximal in $(\text{so} \cup \text{wr})^+$ while the second no. The extension of h_2 concerns the transaction on the right in Figure 5b which is a wr predecessor of the reading transaction. Causal extensibility will require that at least the $(\text{so} \cup \text{wr})^+$ maximal (pending) transactions can always be extended with any event while still preserving consistency. The restriction to $(\text{so} \cup \text{wr})^+$ maximal transactions is intuitively related to the fact that transactions should not read from non-committed (pending) transactions, e.g., the reading transaction in h_2 should not read from the still pending transaction that writes x and later y .

Formally, let $h = \langle T, \text{so}, \text{wr} \rangle$ be a history. A transaction t is called $(\text{so} \cup \text{wr})^+$ -maximal in h if h does not contain any transaction t' such that $(t, t') \in (\text{so} \cup \text{wr})^+$. We define a *causal extension* of a pending transaction t in h with an event e as a history h' such that:

- e is added to t as a maximal element of po_t ,
- if e is a read event and t does not contain a write to $\text{var}(e)$, then wr is extended with some tuple (t', e) such that $(t', t) \in (\text{so} \cup \text{wr})^+$ in h (if e is a read event and t does contain a write to $\text{var}(e)$, then the value returned by e is the value written by the latest write on $\text{var}(e)$ before e in t ; the definition of the return value in this case is unique and does not involve wr dependencies),
- the other elements of h remain unchanged in h' .

For example, Figure 7b and 7c present two causal extensions with a $\text{read}(x)$ event of the transaction t_4 in the history h in Figure 7a. The new read event reads from transaction t_1 or t_3 which were already related by $(\text{so} \cup \text{wr})^+$ to t_4 . An extension of h where the new read event reads from t_2 is not a causal extension because $(t_2, t_4) \notin (\text{so} \cup \text{wr})^+$.

Definition 3.3. An isolation level I is called *causally-extensible* if for every I -consistent history h , every $(\text{so} \cup \text{wr})^+$ -maximal pending transaction t in h , and every event e , there exists a causal extension h' of t with e that is I -consistent.

THEOREM 3.4. *Causal Consistency, Read Atomic, and Read Committed are causally-extensible.*

Snapshot Isolation and Serializability are not causally extensible. Figure 6 presents a counter-example to causal extensibility: the causal extension of the history h that does not contain the $\text{write}(x, 2)$ written in blue bold font with this event does not satisfy neither Snapshot Isolation nor Serializability although h does. Note that the causal extension with a write event is unique. (Note that both h and this causal extension satisfy Causal Consistency and therefore, as expected, this counter-example does not apply to isolation levels weaker than Causal Consistency.)

Algorithm 1 EXPLORE algorithm

```

442 1: function EXPLORE( $P, h_{<}, locals$ )
443
444 2:    $j, e, \gamma \leftarrow \text{NEXT}(P, h_{<}, locals)$ 
445 3:    $locals' \leftarrow locals[e \mapsto \gamma]$ 
446 4:   if  $e = \perp$  and  $\text{VALID}(h)$  then
447 5:     output  $h, locals'$ 
448 6:   else if  $\text{type}(e) = \text{read}$  then
449 7:     for all  $t \in \text{VALIDWRITES}(h, e)$  do
450 8:        $h'_{<} \leftarrow h_{<} \oplus_j e \oplus \text{wr}(t, e)$ 
451 9:       EXPLORE( $P, h'_{<}, locals'$ )
452 10:      EXPLORESWAPS( $P, h'_{<}, locals'$ )
453
454 11:   else
455 12:      $h'_{<} \leftarrow h_{<} \oplus_j e$ 
456 13:     EXPLORE( $P, h'_{<}, locals'$ )
457 14:     EXPLORESWAPS( $P, h'_{<}, locals'$ )

```

Algorithm 2 EXPLORESWAPS

```

1: function EXPLORESWAPS( $P, h_{<}, locals$ )
2:    $l \leftarrow \text{COMPUTEREORDERINGS}(h_{<})$ 
3:   for all  $(\alpha, \beta) \in l$  do
4:     if  $\text{OPTIMALITY}(h_{<}, \alpha, \beta, locals)$  then
5:       EXPLORE( $P, \text{SWAP}(h_{<}, \alpha, \beta, locals)$ )

```

4 SWAPPING-BASED MODEL CHECKING ALGORITHMS

We define a class of stateless model checking algorithms for enumerating executions of a given transactional program, that we call *swapping-based algorithms*. Section 5 will describe a concrete instance that applies to isolation levels that are prefix-closed and causally extensible.

These algorithms are defined by the recursive function EXPLORE listed in Algorithm 1. The function EXPLORE receives as input a program P , an *ordered history* $h_{<}$, which is a pair $(h, <)$ of a history and a total order $<$ on all the events in h , and a mapping $locals$ that associates each event e in h with the valuation of local variables in the transaction of e ($\text{tr}(h, e)$) just before executing e . For an ordered history $(h, <)$ with $h = \langle T, \text{so}, \text{wr} \rangle$, we assume that $<$ is consistent with po , so , and wr , i.e., $e_1 < e_2$ if $(\text{tr}(h, e_1), \text{tr}(h, e_2)) \in (\text{so} \cup \text{wr})^+$ or $(e_1, e_2) \in \text{po}$. Initially, the ordered history and the mapping $locals$ are empty.

The function EXPLORE starts by calling NEXT to obtain an event representing the next database access in some pending transaction of P , or a begin/commit/abort event for starting or ending a transaction. This event is associated to some session j . For example, a typical implementation of NEXT would choose one of the pending transactions (in some session j), execute all local instructions until the next database instruction in that transaction (applying the transition rules IF-TRUE, IF-FALSE, and LOCAL) and return the event e corresponding to that database instruction and the current local state γ . NEXT may also return \perp if the program finished. If NEXT returns \perp , then the function VALID can be used to filter executions that satisfy the intended isolation level before outputting the current history and local states (the use of VALID will become relevant in Section 6).

Otherwise, the event e is added to the ordered history $h_{<}$. If e is a read event, then VALIDWRITES computes a set of write events w in the current history that are valid for e , i.e., adding the event e along with the wr dependency (w, e) leads to a history that still satisfies the intended isolation level. Concerning notations, let h be a history where so is represented as a function $\text{so} : \text{SessId} \rightarrow \text{Tlogs}^*$ (as in § 2.3). For event e , $h \oplus_j e$ is the history obtained from h by adding e to the last transaction in $\text{so}(j)$ as the last event in po (i.e., if $\text{so}(j) = \sigma; \langle t, E, \text{po}_t \rangle$, then the session order so' of $h \oplus_j e$ is defined by $\text{so}'(k) = \text{so}(k)$ for all $k \neq j$ and $\text{so}(j) = \sigma; \langle t, E \cup \{e\}, \text{po}_t \cup \{(e', e) : e' \in E\} \rangle$). This is extended to ordered histories: $(h, <) \oplus_j e$ is defined as $(h \oplus_j e, < \cdot e)$ where $< \cdot e$ means that e is added as the last element of $<$. Also, $h \oplus_j (e, \text{begin})$ is a history where $\langle t, \{(e, \text{begin})\}, \emptyset \rangle$ with t a fresh id is appended to $\text{so}(j)$, and $h \oplus \text{wr}(t, e)$ is defined by adding (t, e) to the write-read of h .

Once an event is added to the current history, the algorithm may explore other histories obtained by re-ordering events in the current one. Such re-orderings are required for completeness. New read events can only read from writes executed in the past which limits the set of explored histories to the scheduling imposed by NEXT. Without re-orderings, writes scheduled later by NEXT cannot be read by read events executed in the past, although this may be permitted by the isolation level.

The function EXPLORESWAPS calls COMPUTEREORDERINGS to compute pairs of sequences of events α, β that should be re-ordered; α and β are *contiguous and disjoint* subsequences of the total order $<$, and α should end before β (since β will be re-ordered before α). Typically, α would contain a read event r and β a write event w such that re-ordering the two enables r to read from w . Ensuring soundness and avoiding redundancy, i.e., exploring the same history multiple times, may require restricting the application of such re-orderings. This is modeled by the Boolean condition called OPTIMALITY. If this condition holds, the new explored histories are computed by the function SWAP. This function returns local states as well, which are necessary for continuing the exploration. We assume that SWAP($h_{<}, \alpha, \beta, locals$) returns pairs ($h'_{<}, locals'$) such that

- (1) h' contains at least the events in α and β ,
- (2) h' without the events in α is a prefix of h , and
- (3) if a read r in α reads from different writes in h and h' (the *wr* relations of h and h' associate different transactions to r), then r is the last event in its transaction (w.r.t. *po*).

The first condition makes the re-ordering “meaningful” while the last two conditions ensure that the history h' is feasible by construction, i.e., it can be obtained using the operational semantics defined in Section 2.3. Feasibility of h' is ensured by keeping prefixes of transaction logs from h and all their *wr* dependencies except possibly for read events in α (second condition). In particular, for events in β , it implies that h' contains all their $(po \cup so \cup wr)^*$ predecessors. Also, the change of a read-from dependency is restricted to the last read in a transaction (third condition) because changing the value returned by a read may disable later events in the same transaction⁴.

A concrete implementation of EXPLORE is called:

- *I-sound* if it outputs only histories in $hist_I(P)$ for every program P ,
- *I-complete* if it outputs every history in $hist_I(P)$ for every program P ,
- *optimal* if it does not output the same history twice,
- *strongly optimal* if it is optimal and never engages in fruitless explorations, i.e., EXPLORE is never called (recursively) on a history h that does not satisfy I , and every call to EXPLORE results in an output or another recursive call to EXPLORE.

5 SWAPPING-BASED MODEL CHECKING FOR PREFIX-CLOSED AND CAUSALLY-EXTENSIBLE ISOLATION LEVELS

We define a concrete implementation of EXPLORE, denoted as EXPLORE-CE, that is *I-sound*, *I-complete*, and *strongly optimal* for any isolation level I that is prefix-closed and causally-extensible. The isolation level I is a parameter of EXPLORE-CE. The space complexity of EXPLORE-CE is polynomial in the size of the program. An important invariant of this implementation is that it explores histories with *at most one* pending transaction and this transaction is maximal in session order. This invariant is used to avoid fruitless explorations: since I is assumed to be causally-extensible, there always exists an extension of the current history with one more event that continues to satisfy I . Moreover, this invariant is sufficient to guarantee completeness in the sense defined above of exploring all histories of “full” program executions (that end in a final configuration).

Section 5.1 describes the implementations of NEXT and VALIDWRITES used to extend a given execution, Section 5.2 describes the functions COMPUTEREORDERINGS and SWAP used to compute

⁴Different *wr* dependencies for previous reads can be explored in other steps of the algorithm.

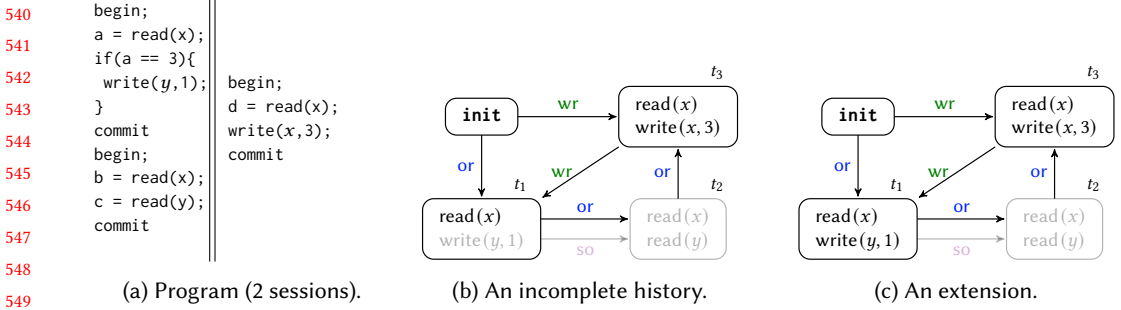


Fig. 8. A program with two sessions (a), a history h (b), and an extension of h with an event returned by NEXT (c). The so-edges from **init** to the other transactions are omitted for legibility. We use edges labeled by **or** to represent the oracle order $<_{or}$. Events in gray are not yet added to the history.

re-ordered executions, and Section 5.3 describes the OPTIMALITY restriction on re-ordering. We assume that the function VALID is defined as simply $VALID(h) ::= true$ (no filter before outputting). Section 5.4 discusses correctness arguments.

5.1 Extending Histories According to An Oracle Order

The function NEXT generates events representing database accesses to extend an execution, according to an *arbitrary but fixed* order between the transactions in the program called *oracle order*. We assume that the oracle order, denoted by $<_{or}$, is consistent with the order between transactions in the same session of the program. The extension of $<_{or}$ to events is defined as expected. For example, assuming that each session has an id, an oracle order can be defined by an order on session ids along with the session order **so**: transactions from sessions with smaller ids are considered first and the order between transactions in the same session follows **so**.

NEXT returns a new event of the transaction that is not already completed and that is *minimal* according to $<_{or}$. In more detail, if j, e, γ is the output of $NEXT(P, h_{<}, locals)$, then either:

- the last transaction log t of session j (w.r.t. **so**) in h is pending, and t is the smallest among pending transaction logs in h w.r.t. $<_{or}$
- h contains no pending transaction logs and the next transaction of sessions j is the smallest among not yet started transactions in the program w.r.t. $<_{or}$.

This implementation of NEXT is deterministic and it prioritizes the completion of pending transactions. The latter is useful to maintain the invariant that any history explored by the algorithm has at most one pending transaction. Preserving this invariant requires that the histories given as input to NEXT also have at most one pending transaction. This is discussed further when explaining the process of re-ordering events in Section 5.2.

For example, consider the program in Figure 8a, an oracle order which orders the two transactions in the left session before the transaction in the right session, and the history h in Figure 8b. Since the local state of the pending transaction on the left stores 3 to the local variable a (as a result of the previous $read(x)$ event) and the Boolean condition in `if` holds, NEXT will return the event $write(y, 1)$ when called with h .

According to Algorithm 1, if the event returned by NEXT is not a read event, then it is simply added to the current history as the maximal element of the order $<$ (cf. the definition of \oplus_j on ordered histories). If it is a read event, then adding this event may result in multiple histories depending on the chosen **wr** dependency. For example, in Figure 9, extending the history in Figure 9b with the $read(x)$ event could result in two different histories, pictured in Figure 9c and 9d, depending on the write with whom this read event is associated by **wr**. However, under CC, the




Fig. 9. Extensions of a history by adding a read event. Events in gray are not yet added to the history.

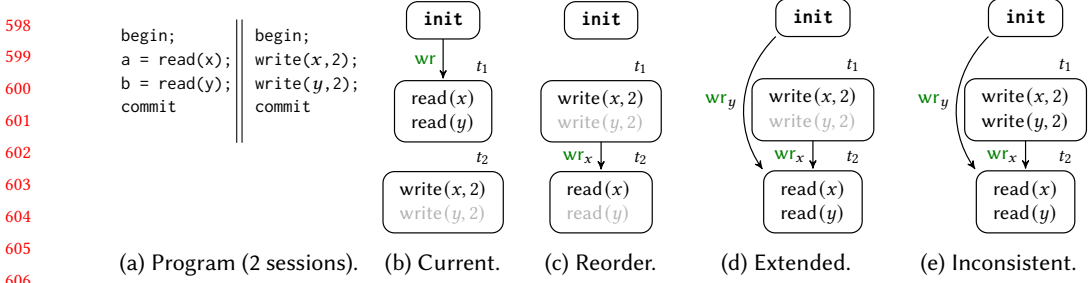


Fig. 10. Example of inconsistency after swapping two events. All **so**-edges from **init** to the other transactions are omitted for legibility. The history order $<$ is represented by the top to bottom order in each figure. Events in gray are not yet added to the history.

latter history is inconsistent. The function **VALIDWRITES** limits the choices to those that preserve consistency with the intended isolation level I , i.e.,

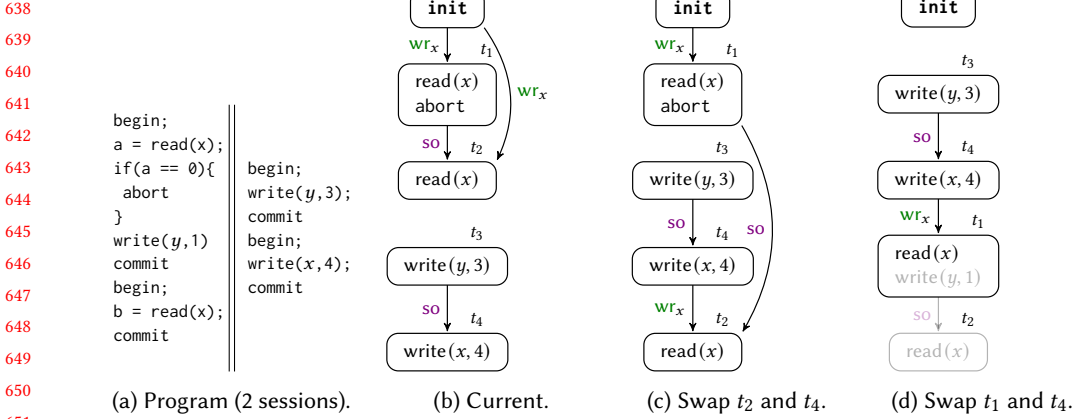
$$\text{VALIDWRITES}(h, e) := \{t \in \text{commTrans}(h) \mid h \oplus e \oplus \text{wr}(t, e) \text{ satisfies } I\}$$

where $\text{commTrans}(h)$ is the set of committed transactions in h .

5.2 Re-Ordering Events in Histories

After extending the current history with one more event, **EXPLORE** may be called recursively on other histories obtained by re-ordering events in the current one (and dropping some other events).

Re-ordering events must preserve the invariant of producing histories with at most one pending transaction. To explain the use of this invariant in avoiding fruitless explorations, let us consider the program in Figure 10a assuming an exploration under Read Committed. The oracle order gives priority to the transaction on the left. Assume that the current history reached by the exploration is the one pictured in Figure 10b (the last added event is $\text{write}(x, 2)$). Swapping $\text{write}(x, 2)$ with $\text{read}(x)$ would result in the history pictured in Figure 10c. To ensure that this swap produces a new history which was not explored in the past, the wr_x dependency of $\text{read}(x)$ is changed towards the $\text{write}(x, 2)$ transaction (we detail this later). By the definition of **NEXT** (and the oracle order), this history shall be extended with $\text{read}(y)$, and this read event will be associated by wr_y to the only available $\text{write}(y, _)$ event from **init**. This is pictured in Figure 10d. The next exploration step will extend the history with $\text{write}(y, 2)$ (the only extension possible) which however, results in a history that does *not* satisfy Read Committed, thereby, the recursive exploration branch being blocked. The core issue is related to the history in Figure 10d which has a pending transaction that is *not* $(\text{so} \cup \text{wr})^+$ -maximal. Being able to extend such a transaction while maintaining consistency is not guaranteed by Read Committed (and any other isolation level we consider). Nevertheless, causal extensibility guarantees the existence of an extension for pending transactions that are



652 Fig. 11. Re-ordering events. All so -edges from $init$ to other transactions are omitted for legibility. The history
653 order $<$ is represented by the top to bottom order in each figure. Events in gray are deleted from the history.

654
655 $(so \cup wr)^+$ -maximal. We enforce this requirement by restricting the explored histories to have at
656 most one pending transaction. This pending transaction will necessarily be $(so \cup wr)^+$ -maximal.

657 To enforce histories with at most one pending transaction, the function `COMPUTEREORDERINGS`,
658 which identifies events to reorder, has a non-empty return value only when the last added event is
659 commit (the end of a transaction)⁵. Therefore, in such a case, it returns pairs of some transaction
660 log prefix ending in a read r and the last completed transaction log t , such that the transaction
661 log containing r and t are *not* causally dependent (i.e., related by $(so \cup wr)^*$) (the transaction
662 log prefix ending in r and t play the role of the subsequences α and respectively, β in the de-
663 scription of `COMPUTEREORDERINGS` from Section 4). To simplify the notation, we will assume that
664 `COMPUTEREORDERINGS` returns pairs (r, t) .

665
666
$$\text{COMPUTEREORDERINGS}(h_{<}) := \{(r, t) \in \mathcal{E} \times T \mid r \in \text{reads}(T) \wedge t \text{ writes } \text{var}(r) \wedge \text{tr}(h, r) < t$$

667
$$\wedge (\text{tr}(h, r), t) \notin (so \cup wr)^* \wedge t \text{ is complete and it includes the last event in } \langle \rangle\}$$

671 For example, for the program in Figure 11a and history h in Figure 11b, `COMPUTEREORDERINGS`(h)
672 would return (r_1, t_4) and (r_2, t_4) where r_1 and r_2 are the read(x) events in t_1 and t_2 respectively.

673 For a pair (r, t) , the function `SWAP` produces a new history h' which contains all the events
674 ordered before r (w.r.t. $<$), the transaction t and all its $(so \cup wr)^*$ predecessors, and the event r
675 reading from t . All the other events are removed. Note that the po predecessors of r from the same
676 transaction are ordered before r by $<$ and they will be also included in h' . The history h' without
677 r is a prefix of the input history h . By definition, the only pending transaction in h' is the one
678 containing the read r . The order relation is updated by moving the transaction containing the read
679 r to be the last; it remains unchanged for the rest of the events.

680
681
$$\text{SWAP}(h_{<}, r, t, \text{locals}) := ((h' = (h \setminus D) \oplus wr(t, r), <'), \text{locals}'), \text{ where } \text{locals}' = \text{locals} \downarrow \text{events}(h')$$

682
$$D = \{e \mid r < e \wedge (\text{tr}(h, e), t) \notin (so \cup wr)^*\}$$
 and $<' = (< \downarrow (\text{events}(h') \setminus \text{events}(\text{tr}(h', r)))) \cdot \text{tr}(h', r)$

684 ⁵Aborted transactions have no visible effect on the state of the database so swapping an aborted transaction cannot produce
685 a new meaningful history.

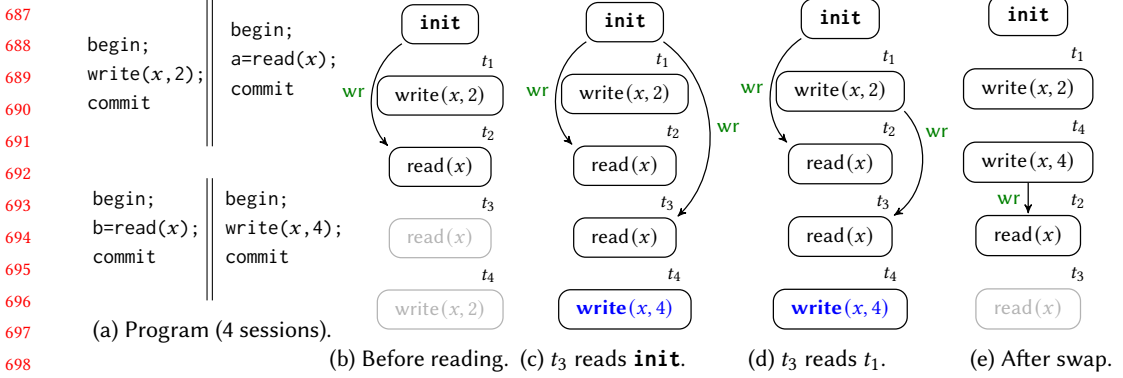


Fig. 12. Re-ordering events versus optimality. We assume an oracle order orders transaction from left to right, top to bottom in the program. All transaction logs are history-ordered top to bottom according to their position in the figure. Events in gray are not yet added to the history.

Above, $h \setminus D$ is the prefix of h obtained by deleting all the events in D from its transaction logs; a transaction log is removed altogether if it becomes empty. Also, $h'' \oplus wr(t, r)$ denotes an *update* of the wr relation of h'' where any pair $(_, r)$ is replaced by (t, r) . Finally, $\langle \cdot \rangle \cdot tr(h', r)$ is an extension of the total order $\langle \cdot \rangle$ obtained by appending the events in $tr(h', r)$ according to program order.

Continuing with the example of Figure 11, when swapping r_1 and t_4 , all the events in transaction t_2 belong to D and they will be removed. This is shown in Figure 11d. Note that transaction t_1 aborted in Figure 11b while it will commit in Figure 11d (because the value read from x changed). When swapping r_2 and t_4 , no event but the commit in t_2 will be deleted (Figure 11c).

5.3 Ensuring Optimality

Simply extending histories according to NEXT and making recursive calls on re-ordered histories whenever they are I -consistent guarantees soundness and completeness, but it does not guarantee optimality. Intuitively, the source of redundancy is related to the fact that applying SWAP on different histories may give the same result.

As a first example, consider the program in Figure 12a with 2 transactions that only read some variable x and 2 transactions that only write to x , each transaction in a different session. Assume that EXPLORE reaches the ordered history in Figure 12b and NEXT is about to return the second reading transaction. EXPLORE will be called recursively on the two histories in Figure 12c and Figure 12d that differ in the write that this last read is reading from (the initial write or the first write transaction). On both branches of the recursion, NEXT will extend the history with the last write transaction written in blue bold font. For both histories, swapping this last write with the first read on x will result in the history in Figure 12e (cf. the definition of COMPUTEREORDERINGS and SWAP). Thus, both branches of the recursion will continue extending the same history and optimality is violated. The source of non-optimality is related to wr dependencies that are removed during the SWAP computation. The histories in Figure 12c and Figure 12d differ in the wr dependency involving the last read, but this difference was discarded during the SWAP computation. To avoid this behavior, SWAP is enabled only on histories where the discarded wr dependencies relate to some “fixed” set of writes, i.e., latest⁶ writes w.r.t. $\langle \cdot \rangle$ that guarantee consistency by causal extensibility (see the definition of $readLatest_I(_, _)$ below). By causal extensibility, a read r can always read

⁶We use latest writes because they are uniquely defined. In principle, other ways of identifying some unique set of writes could be used.

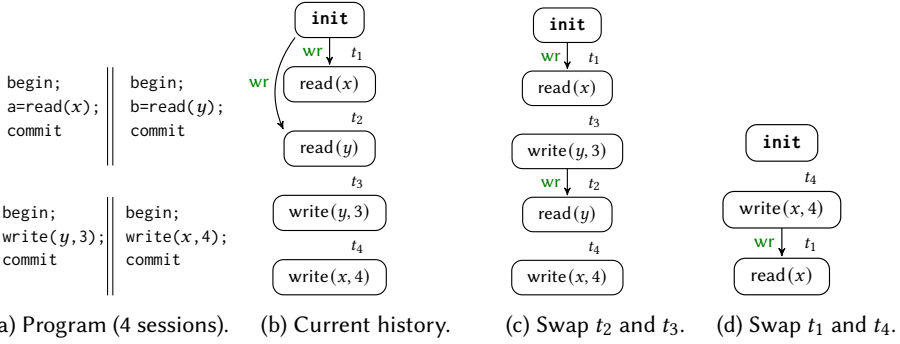


Fig. 13. Re-ordering the same read on different branches of the recursion.

from a write which already belongs to its “causal past”, i.e., predecessors in $(\text{so} \cup \text{wr})^*$ excluding the wr dependency for r . For every discarded wr dependency, it is required that the read reads from the latest such write w.r.t. $<$. In this example, re-ordering is enabled only when the second read(x) reads from the initial write; write($x, 2$) does not belong to its “causal past” (when the wr dependency of the read itself is excluded).

The restriction above is not sufficient, because the two histories for which SWAP gives the same result may not be generated during the same recursive call (for different wr choices when adding a read). For example, consider the program in Figure 13a that has four sessions each containing a single transaction. EXPLORE may compute the history h pictured in Figure 13b. Before adding transaction t_4 , EXPLORE can re-order t_3 and t_2 and then extend with t_4 and arrive at the history h_1 in Figure 13c. Also, after adding t_4 , it can re-order t_1 and t_4 and arrive at the history h_2 in Figure 13d. However, swapping the same t_1 and t_4 in h_1 leads to the same history h_2 , thereby, having two recursive branches that end up with the same input and violate optimality. Swapping t_1 and t_4 in h_1 should not be enabled because the read(y) to be removed by SWAP has been swapped in the past. Removing it makes it possible that this recursive branch explores that wr choice for read(y) again.

The OPTIMALITY condition restricting re-orderings requires that the re-ordered history be I -consistent and that every read deleted by SWAP or the re-ordered read r (whose wr dependency is modified) reads from a latest valid write, cf. the example in Figure 12, and it is not already swapped, cf. the example in Figure 13 (the set D is defined as in SWAP):

OPTIMALITY($h_{<}, r, t, \text{locals}$) := the history returned by SWAP($h_{<}, r, t, \text{locals}$) satisfies I

$$\wedge \forall r' \in \text{reads}(h) \cap (D \cup \{r\}). \neg \text{SWAPPED}(h_{<}, r') \wedge \text{readLatest}_I(h_{<}, r', t)$$

A read r reads from a causally latest valid transaction, denoted as $\text{readLatest}_I(h_{<}, r)$, if reading from any other later transaction t' w.r.t. $<$ which is in the “causal past” of $\text{tr}(h_{<}, r)$ violates the isolation level I . Formally, assuming that t_r is the transaction such that $(t_r, r) \in \text{wr}$ in h ,

$$\text{readLatest}_I(h_{<}, r, t) := t_r = \max_{<} \left\{ \begin{array}{l} t' \text{ writes } \text{var}(r) \wedge (t', \text{tr}(h_{<}, r)) \in (\text{so} \cup \text{wr})^* \text{ in } h' \\ \wedge h' \oplus r \oplus \text{wr}(t', r) \models I \end{array} \right\}$$

where $h' = h \setminus \{e \mid r \leq e \wedge (\text{tr}(h, e), t) \notin (\text{so} \cup \text{wr})^*\}$.

We say that a read r is *swapped* in $h_{<}$ when (1) r reads from a transaction t that is a successor in the oracle order $<_{\text{or}}$ (the transaction was added by NEXT after the read), which is now a predecessor⁷ in the history order $<$, (2) there is no transaction t' that is before r in both $<_{\text{or}}$ and $<$, and which

⁷The EXPLORE maintains the invariant that every read follows the transaction it reads from in the history order $<$.

785 is a $(\text{so} \cup \text{wr})^+$ successor of t , and (3) r is the first read in its transaction to read from t . Formally,
 786 assuming that t is the transaction such that $(t, r) \in \text{wr}$,

$$787 \quad \text{SWAPPED}(h_{<}, r) := t < r \wedge t >_{\text{or}} r \wedge \forall t' \in h. t' <_{\text{or}} \text{tr}(h, r) \Rightarrow (r < t' \vee (t, t') \notin (\text{so} \cup \text{wr})^+)$$

$$788 \quad \wedge \forall r' \in \text{reads}(h). (t, r') \in \text{wr} \Rightarrow (r', r) \notin \text{po}$$

790 Condition (1) states a quite straightforward fact about swaps: r could not have been involved in a
 791 swap if it reads from a predecessor in the oracle order which means that it was added by NEXT after
 792 the transaction it reads from. Conditions (2) and (3) are used to exclude spurious classifications as
 793 swapped reads. Concerning condition (2), suppose that in a history h we swap a transaction t with
 794 respect a (previous) read event r . Later on, the algorithm may add a read r' reading also from t .
 795 Condition (2) forbids r' to be declared as swapped. Indeed, taking $\text{tr}(h, r)$ as an instantiation of t' ,
 796 $\text{tr}(h, r)$ is before r' in both $<_{\text{or}}$ and $<$ and it reads from the same transaction as r' , thereby, being a
 797 $(\text{so} \cup \text{wr})^+$ successor of the transaction read by r' . Condition (3) forbids that, after swapping r and
 798 t in h , later read events from the same transaction as r can be considered as swapped.

799 Showing that I -completeness holds despite discarding re-orderings is quite challenging. Intu-
 800 itively, it can be shown that if some SWAP is *not* enabled in some history $h_{<}$ for some pair (r, t)
 801 although the result would be I -consistent (i.e., $\text{OPTIMALITY}(h_{<}, r, t, \text{locals})$ does not hold because
 802 some deleted read is swapped or does not read from a causally latest transaction), then the algorithm
 803 explores another history h' which coincides with h except for those deleted reads who are now
 804 reading from causally latest transactions. Then, h' would satisfy $\text{OPTIMALITY}(h_{<}, r, t, \text{locals})$, and
 805 moreover applying SWAP on h' for the pair (r, t) would lead to the same result as applying SWAP
 806 on h , thereby, ensuring completeness.

808 5.4 Correctness

809 The following theorem states the correctness of the algorithm presented in this section:

811 **THEOREM 5.1.** *For any prefix-closed and causally extensible isolation level I , EXPLORE-CE is I -sound,
 812 I -complete, strongly optimal, and polynomial space.*

813 I -soundness is a consequence of the VALIDWRITES and OPTIMALITY definitions which guarantee
 814 that all histories given to recursive calls are I -consistent, and of the SWAP definition which ensures
 815 to only produce feasible histories (which can be obtained using the operational semantics defined in
 816 Section 2.3). The fact that this algorithm never engages in fruitless explorations follows easily from
 817 causal-extensibility which ensures that any current history can be extended with any event returned
 818 by NEXT. Polynomial space is also quite straightforward since the **for all** loops in Algorithm 1 have
 819 a linear number of iterations: the number of iterations of the loop in EXPLORE, resp., EXPLORESWAPS,
 820 is bounded by the number of write, resp., read, events in the current history (which is smaller than
 821 the size of the program; recall that we assume bounded programs with no loops as usual in SMC
 822 algorithms). On the other hand, the proofs of I -completeness and optimality are quite complex.

823 I -completeness means that for any given program P , the algorithm outputs every history h in
 824 $\text{hist}_I(P)$. The proof of I -completeness defines a sequence of histories produced by the algorithm
 825 starting with an empty history and ending in h , for every such history h . It consists of several steps:

- 827 (1) Define a *canonical* total order $<$ for every unordered partial history h , such that if the
 828 algorithm reaches $h_{<'}$, for some order $<'$, then $<$ and $<'$ coincide. This canonical order is
 829 useful in future proof steps as it allows to extend several definitions to arbitrary histories
 830 that are not necessarily reachable, such as OPTIMALITY or SWAPPED.
- 831 (2) Define the notion of *or-respectfulness*, an invariant satisfied by every (partial) ordered
 832 history reached by the algorithm. Briefly, a history is *or-respectful* if it has only one pending
 833

transaction and for every two events e, e' such that $e <_{\text{or}} e'$, either $e < e'$ or there is a swapped event e'' in between.

- (3) Define a deterministic function PREV which takes as input a partial history (not necessarily reachable), such that if h is reachable, then $\text{PREV}(h)$ returns the history computed by the algorithm just before h (i.e., the previous history in the call stack). Prove that if a history h is **or**-respectful, then $\text{PREV}(h)$ is also **or**-respectful.
- (4) Deduce that if h is **or**-respectful, then there is a finite collection of **or**-respectful histories $H_h = \{h_i\}_{i=0}^n$ such that $h_n = h$, $h_0 = \emptyset$, and $h_i = \text{PREV}(h_{i+1})$ for each i . The **or**-respectfulness invariant and the causal-extensibility of the isolation level are key to being able to construct such a collection. In particular, they are used to prove that h_i has at most the same number of swapped events as h_{i+1} and in case of equality, h_i contain exactly one event less than h_{i+1} , which implies that the collection is indeed finite.
- (5) Prove that if h is **or**-respectful and $\text{PREV}(h)$ is reachable, then h is also reachable. Conclude by induction that every history in H_h is reachable, as h_0 is the initial state and $h_i = \text{PREV}(h_{i+1})$.

The proof of strong optimality relies on arguments employed for I -completeness. It can be shown that if the algorithm would reach a (partial) history h twice, then for one of the two exploration branches, the history h' computed just before h would be different from $\text{PREV}(h)$, which contradicts the definition of $\text{PREV}(h)$.

In terms of time complexity, the $\text{EXPLORE-CE}(I)$ algorithm achieves polynomial time between consecutive outputs for isolation levels I where checking I -consistency of a history is polynomial time, e.g., RC, RA, and CC.

6 SWAPPING-BASED MODEL CHECKING FOR SNAPSHOT ISOLATION AND SERIALIZABILITY

For EXPLORE-CE , the part of strong optimality concerning *not* engaging in fruitless explorations was a direct consequence of causal extensibility (of the isolation level). However, isolation levels such as SI and SER are *not* causally extensible (see Section 3.2). Therefore, the question we investigate in this section is whether there exists another implementation of EXPLORE that can ensure strong optimality along with I -soundness and I -completeness for I being SI or SER. We answer this question in the negative, and as a result, propose an SMC algorithm that extends EXPLORE-CE by just filtering histories before outputting to be consistent with SI or SER.

THEOREM 6.1. *If I is Snapshot Isolation or Serializability, there exists no EXPLORE algorithm that is I -sound, I -complete, and strongly optimal.*

The proof of Theorem 6.1 defines a program with two transactions and shows that any concrete instance of EXPLORE in Alg. 1 *cannot be both* I -complete and strongly optimal.

Given this negative result, we define an implementation of EXPLORE for an isolation level $I \in \{\text{SI}, \text{SER}\}$ that ensures optimality instead of strong optimality, along with soundness, completeness, and polynomial space bound. Thus, let $\text{EXPLORE-CE}(I_0)$ be an instance of EXPLORE-CE parametrized by $I_0 \in \{\text{RC}, \text{RA}, \text{CC}\}$. We define an implementation of EXPLORE for I , denoted by $\text{EXPLORE-CE}^*(I_0, I)$, which is exactly $\text{EXPLORE-CE}(I_0)$ except that instead of $\text{VALID}(h) ::= \text{true}$, it uses

$$\text{VALID}(h) ::= h \text{ satisfies } I$$

$\text{EXPLORE-CE}^*(I_0, I)$ enumerates exactly the same histories as $\text{EXPLORE-CE}(I_0)$ except that it outputs only histories consistent with I . The following is a direct consequence of Theorem 5.1.

COROLLARY 6.2. *For any isolation levels I_0 and I such that I_0 is prefix-closed and causally extensible, and I_0 is weaker than I , $\text{EXPLORE-CE}^*(I_0, I)$ is I -sound, I -complete, optimal, and polynomial space.*

7 EXPERIMENTAL EVALUATION

We evaluate an implementation of EXPLORE-CE and EXPLORE-CE* in the context of the Java Pathfinder (JPF) [59] model checker for Java concurrent programs. As benchmark, we use bounded-size client programs of a number of database-backed applications drawn from the literature. The experiments were performed on an Apple M1 with 8 cores and 16 GB of RAM.

7.1 Implementation

We implemented our algorithms as an extension of the `DFSearch` class in JPF. For performance reasons, we implemented an iterative version of these algorithms where roughly, inputs to recursive calls are maintained as a collection of histories instead of relying on the call stack. For checking consistency of a history with a given isolation level, we implemented the algorithms proposed by Biswas and Enea [16]. We plan to make our implementation publicly available.

Our tool takes as input a Java program and isolation levels as parameters. We assume that the program uses a fixed API for interacting with the database, similar to a key-value store interface. This API consists of specific methods for starting/ending a transaction, and reading/writing a global variable. The fixed API is required for being able to maintain the database state separately from the JVM state (the state of the Java program) and update the current history in each database access. This relies on a mechanism for “transferring” values read from the database state to the JVM state.

7.2 Benchmark

We consider a set of benchmarks inspired by real-world applications and evaluate them under different types of client programs and isolation levels.

Shopping Cart [56] allows users to add, get and remove items from their shopping cart and modify the quantities of the items present in the cart.

Twitter [30] allows users to follow other users, publish tweets and get their followers, tweets and tweets published by other followers.

Courseware [47] manages the enrollment of students in courses in an institution. It allows to open, close and delete courses, enroll students and get all enrollments. One student can only enroll to a course if it is open and its capacity has not reached a fixed limit.

Wikipedia [30] allows users to get the content of a page (registered or not), add or remove pages to their watching list and update pages.

TPC-C [57] models an online shopping application with five types of transactions: reading the stock of a product, creating a new order, getting its status, paying it and delivering it.

SQL tables are modeled using a “set” global variable whose content is the set of ids (primary keys) of the rows present in the table, and a set of global variables, one variable for each row in the table (the name of the variable is the primary key of that row). SQL statements such as INSERT and DELETE statements are modeled as writes on that “set” variable while SQL statements with a WHERE clause (SELECT, JOIN, UPDATE) are compiled to a read of the table’s set variable followed by reads or writes of variables that represent rows in the table (similarly to [17, 55]).

7.3 Experimental Results

We designed three experiments where we compare the performance of a baseline model checking algorithm, EXPLORE-CE and EXPLORE-CE* for different (combinations of) isolation levels, and we explore the scalability of EXPLORE-CE when increasing the number of sessions and transactions per session, respectively. For each experiment we report running time, memory consumption, and the number of end states, i.e., histories of complete executions and in the case of EXPLORE-CE*, before applying the VALID filter. As the number of end states for a program on a certain isolation level increases, the running time of our algorithms naturally increases as well.

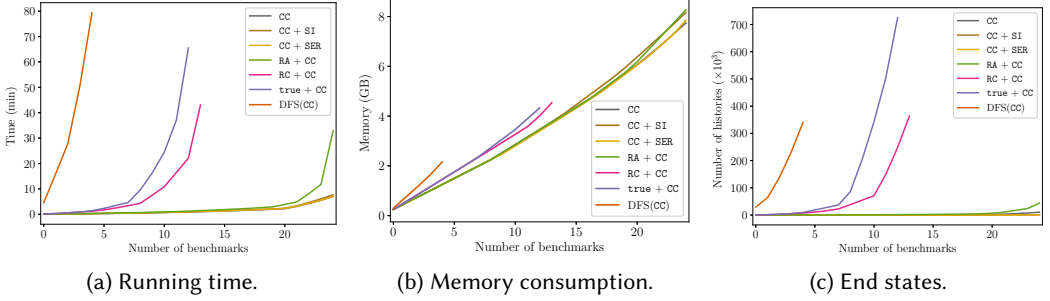


Fig. 14. Cactus plots comparing different algorithms in terms of time, memory, and end states. For readability, we use CC to denote EXPLORE-CE under CC, $I_1 + I_2$ stands for EXPLORE-CE* (I_1, I_2), and true is the trivial isolation level where every history is consistent. Differences between CC, CC + SI and CC + SER are very small and their graphics overlap. Moreover, DFS(CC) denotes a standard DFS traversal of the semantics defined in Section 2.3. These plots exclude benchmarks that timeout (30 mins): 11 12 and 20 benchmarks timeout for (RC, CC), (true, CC) and DFS(CC) respectively.

The first experiment compares the performance of our algorithms for different combinations of isolation levels and a baseline model checking algorithm that performs no partial order reduction. We consider as benchmark five (independent) client programs⁸ for each application described above (25 in total), each program with three sessions and three transactions per session. The running time, memory consumption, and number of end states are reported in Figure 14 as cactus plots [19].

To justify the benefits of partial order reduction, we implement a baseline model checking algorithm DFS(CC) that performs a standard DFS traversal of the execution tree w.r.t. the formal semantics defined in Section 2.3 for CC (for fairness, we restrict interleavings so at most one transaction is pending at a time). This baseline algorithm may explore the same history multiple times since it includes no partial order reduction mechanism. In terms of time, DFS(CC) behaves poorly: it timeouts for 20 out of the 25 programs and it is less efficient even when it terminates. We consider a timeout of 30 mins. In comparison the strongly optimal algorithm EXPLORE-CE(CC) (under CC) finishes in 17 seconds in average. DFS(CC) is also worse in terms of memory consumption. The memory consumption of DFS(CC) is 441MB in average, compared to 317MB for EXPLORE-CE(CC) (JPF forces a minimum consumption of 256MB).

To argue about the benefits of *strong* optimality, we compare EXPLORE-CE(CC) which is strongly optimal with “plain” optimal algorithms EXPLORE-CE* (I_0, CC) for different levels I_0 . As shown in Figure 14(a), in terms of time, EXPLORE-CE(CC) is more efficient than every “plain” optimal algorithm, and the difference in performance grows as I_0 becomes weaker. In the limit, when I_0 is the trivial isolation level true where every history is consistent, EXPLORE-CE* (true, CC) timeouts for 12 out of the 25 programs. The average speedup (average of individual speedups) of EXPLORE-CE(CC) w.r.t. EXPLORE-CE* (RA, CC), EXPLORE-CE* (RC, CC) and EXPLORE-CE* (true, CC) is 2, 31, and 54 respectively. In terms of memory, all algorithms consume around 300 MB in average.

For the SI and SER isolation levels that admit no strongly optimal EXPLORE algorithm, we observe that the overhead of EXPLORE-CE* (CC, SI) or EXPLORE-CE* (CC, SER) relative to EXPLORE-CE(CC) is negligible (the corresponding lines in Figure 14 are essentially overlapping). This is due to the fact that the consistency checking algorithms of Biswas and Enea [16] are polynomial time when the number of sessions is fixed, which makes them fast at least on histories with few sessions.

In our second experiment, we investigate the scalability of EXPLORE-CE when increasing the number of sessions. For each $i \in [1, 5]$, we consider five (independent) client programs for TPC-C

⁸For an application that defines a number of transactions, a client program consists of a number of sessions, each session containing a sequence of transactions defined by the application.




Fig. 15. Evaluating the scalability of EXPLORE-CE(CC) for TPC-C and Wikipedia client programs when increasing their size. These plots include benchmarks that timeout (30 mins): 6 and 10 for 4 and 5 sessions respectively in Figure 15a, and 2 and 5 for 4 and 5 transactions per sessions respectively in Figure 15b.

and five for Wikipedia (10 in total) with i sessions, each session containing three transactions⁹. We take CC as isolation level. The plot in Figure 15a shows average running time and memory consumption for each number $i \in [1, 5]$ of sessions. As expected, increasing the number of sessions is a bottleneck running time wise because the number of histories/executions increases significantly as well. However, memory consumption does not grow with the same trend, as expected from the polynomial space complexity bound.

For our third experiment, we evaluate the scalability of EXPLORE-CE(CC) when increasing the number of transactions per session. We consider five (independent) TPC-C client programs and five (independent) Wikipedia client programs with 3 sessions and i transactions per session, for every $i \in [1, 5]$. Figure 15b shows average running time and memory consumption for each number $i \in [1, 5]$ of transactions per session. Increasing the number of transactions per session is also a bottleneck for the same reasons as before.

8 RELATED WORK

Checking Correctness of Database-Backed Applications. One line of work is concerned with the logical formalization of isolation levels [7, 14, 16, 23, 61]. Our work relies on the axiomatic definitions of isolation levels introduced by Biswas and Enea [16], which have also investigated the problem of checking whether a given history satisfies a certain isolation level. Our SMC algorithms rely on these algorithms to check consistency of a history with a given isolation level.

Another line of work focuses on the problem of finding “anomalies”: behaviors that are not possible under serializability. This is typically done via a static analysis of the application code that builds a static dependency graph that over-approximates the data dependencies in all possible executions of the application [15, 24, 31, 33, 37, 60]. Anomalies with respect to a given isolation level then correspond to a particular class of cycles in this graph. Static dependency graphs turn out to be highly imprecise in representing feasible executions, leading to false positives. Another source of false positives is that an anomaly might not be a bug because the application may already be designed to handle the non-serializable behavior [22, 33]. Recent work has tried to address these issues by using more precise logical encodings of the application [21, 22], or by using user-guided heuristics [33]. Another approach consists of modeling the application logic and the isolation level in first-order logic and relying on SMT solvers to search for anomalies [38, 46, 49], or defining specialized reductions to assertion checking [12, 13]. Our approach, based on SMC, does not generate false positives because we systematically enumerate only valid executions of a program which allows to check for user-defined assertions.

⁹We consider 10 client programs with 5 sessions, and remove sessions one by one to obtain client programs with a smaller number of sessions.

1030 Several works have looked at the problem of reasoning about the correctness of applications
 1031 executing under weak isolation and introducing additional synchronization when necessary [11,
 1032 36, 43, 47]. These are based on static analysis or logical proof arguments. The issue of repairing
 1033 applications is orthogonal to our work.

1034 MonkeyDB [17] is a mock storage system for testing storage-backed applications. While being
 1035 able to scale to larger code, it has the inherent incompleteness of testing. As opposed to MonkeyDB,
 1036 our algorithms enable a systematic and complete exploration of executions and can establish cor-
 1037 rectness at least in some bounded context, and they are designed to avoid redundancy, enumerating
 1038 equivalent executions multiple times. Such guarantees are beyond the scope of MonkeyDB.

1039 **Dynamic Partial Order Reduction.** Abdulla et al. [2] introduced the concept of *source sets* which
 1040 provided the first strongly optimal DPOR algorithm for Mazurkiewicz trace equivalence. Other
 1041 works study DPOR techniques for coarser equivalence relations, e.g., [3, 8, 10, 25, 26]. In all cases,
 1042 the space complexity is exponential when strong optimality is ensured.

1043 Other works focus on extending DPOR to weak memory models either by targeting a specific
 1044 memory model [1, 4, 5, 48] or by being parametric with respect to an axiomatically-defined memory
 1045 model [39–41]. Some of these works can deal with the coarser reads-from equivalence, e.g., [5, 39–
 1046 41]. Our algorithms build on the work of Kokologiannakis et al. [39] which for the first time,
 1047 proposes a DPOR algorithm which is both strongly optimal and polynomial space. The definitions
 1048 of database isolation levels are quite different with respect to weak memory models, which makes
 1049 these previous works not extensible in a direct manner. These definitions include a semantics for
 1050 *transactions* which are collections of reads and writes, and this poses new difficult challenges. For
 1051 instance, reasoning about the completeness and the (strong) optimality of existing DPOR algorithms
 1052 for shared-memory is agnostic to the scheduler (NEXT function) while the strong optimality of our
 1053 EXPLORE-CE algorithm relies on the scheduler keeping at most one transaction pending at a time. In
 1054 addition, unlike TruSt, EXPLORE-CE ensures that no swapped events can be swapped again and that
 1055 the history order $<$ is an extension of $so \cup wr$. This makes our completeness and optimality proofs
 1056 radically different. Moreover, even for transactional programs with one access per transaction,
 1057 where SER and SC are equivalent, TruSt under SC and EXPLORE-CE $^*(I_0, SER)$ do not coincide, for
 1058 any $I_0 \in \{RC, RA, CC\}$. In this case, TruSt enumerates only SC-consistent histories at the cost of
 1059 solving an NP-complete problem at each step while the EXPLORE-CE * step cost is polynomial time at
 1060 the price of not being strongly-optimal. Furthermore, we identify isolation levels (SI and SER) for
 1061 which it is impossible to ensure both strong optimality and polynomial space bounds (at least with
 1062 a swapping-based algorithm), a type of question that has not been investigated in previous work.

1063

1064

1065 9 CONCLUSIONS

1066

1066 We have presented efficient SMC algorithms based on DPOR for transactional programs running
 1067 under standard isolation levels. These algorithms are instances of a generic schema, called swapping-
 1068 based algorithms, which is parametrized by an isolation level. Our algorithms are sound and
 1069 complete, and have a polynomial space complexity. Additionally, we have identified a class of
 1070 isolation levels, including RC, RA, and CC, for which our algorithms are strongly optimal, and we have
 1071 shown that swapping-based algorithms cannot be strongly optimal in the case of the stronger levels
 1072 SI and SER (but just optimal). It is interesting to observe that for the isolation levels we considered,
 1073 there is an intriguing coincidence between the existence of a strongly optimal swapping-based
 1074 algorithm and the complexity for checking if a given history is consistent with that level. Indeed,
 1075 checking consistency is polynomial time for RC, RA, and CC, and NP-complete for SI and SER.
 1076 Investigating further the relationship between strong optimality and polynomial-time consistency
 1077 checks is an interesting direction for future work.

1078

REFERENCES

- [1] Parosh Aziz Abdulla, Stavros Aronis, Mohamed Faouzi Atig, Bengt Jonsson, Carl Leonardsson, and Konstantinos Sagonas. 2017. Stateless model checking for TSO and PSO. *Acta Informatica* 54, 8 (2017), 789–818. <https://doi.org/10.1007/s00236-016-0275-0>
- [2] Parosh Aziz Abdulla, Stavros Aronis, Bengt Jonsson, and Konstantinos Sagonas. 2017. Source Sets: A Foundation for Optimal Dynamic Partial Order Reduction. *J. ACM* 64, 4 (2017), 25:1–25:49. <https://doi.org/10.1145/3073408>
- [3] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, Magnus Lång, Tuan Phong Ngo, and Konstantinos Sagonas. 2019. Optimal stateless model checking for reads-from equivalence under sequential consistency. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 150:1–150:29. <https://doi.org/10.1145/3360576>
- [4] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Carl Leonardsson. 2016. Stateless Model Checking for POWER. In *Computer Aided Verification - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 9780)*, Swarat Chaudhuri and Azadeh Farzan (Eds.). Springer, 134–156. https://doi.org/10.1007/978-3-319-41540-6_8
- [5] Parosh Aziz Abdulla, Mohamed Faouzi Atig, Bengt Jonsson, and Tuan Phong Ngo. 2018. Optimal stateless model checking under the release-acquire semantics. *Proc. ACM Program. Lang.* 2, OOPSLA (2018), 135:1–135:29. <https://doi.org/10.1145/3276505>
- [6] A. Adya. 1999. *Weak Consistency: A Generalized Theory and Optimistic Implementations for Distributed Transactions*. Technical Report. USA.
- [7] Atul Adya, Barbara Liskov, and Patrick E. O’Neil. 2000. Generalized Isolation Level Definitions. In *Proceedings of the 16th International Conference on Data Engineering, San Diego, California, USA, February 28 - March 3, 2000*, David B. Lomet and Gerhard Weikum (Eds.). IEEE Computer Society, 67–78. <https://doi.org/10.1109/ICDE.2000.839388>
- [8] Pratyush Agarwal, Krishnendu Chatterjee, Shreya Pathak, Andreas Pavlogiannis, and Viktor Toman. 2021. Stateless Model Checking Under a Reads-Value-From Equivalence. In *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12759)*, Alexandra Silva and K. Rustan M. Leino (Eds.). Springer, 341–366. https://doi.org/10.1007/978-3-030-81685-8_16
- [9] Deepthi Devaki Akkoorath and Annette Bieniusa. 2016. *Antidote: the highly-available geo-replicated database with strongest guarantees*. Technical Report. <https://pages.lip6.fr/syncfree/attachments/article/59/antidote-white-paper.pdf>
- [10] Stavros Aronis, Bengt Jonsson, Magnus Lång, and Konstantinos Sagonas. 2018. Optimal Dynamic Partial Order Reduction with Observers. In *Tools and Algorithms for the Construction and Analysis of Systems - 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 10806)*, Dirk Beyer and Marieke Huisman (Eds.). Springer, 229–248. https://doi.org/10.1007/978-3-319-89963-3_14
- [11] Valter Balegas, Sérgio Duarte, Carla Ferreira, Rodrigo Rodrigues, Nuno M. Prego, Mahsa Najafzadeh, and Marc Shapiro. 2015. Putting consistency back into eventual consistency. In *Proceedings of the Tenth European Conference on Computer Systems, EuroSys 2015, Bordeaux, France, April 21-24, 2015*, Laurent Réveillère, Tim Harris, and Maurice Herlihy (Eds.). ACM, 6:1–6:16. <https://doi.org/10.1145/2741948.2741972>
- [12] Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. 2019. Checking Robustness Against Snapshot Isolation. In *Computer Aided Verification - 31st International Conference, CAV 2019, New York City, NY, USA, July 15-18, 2019, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 11562)*, Isil Dillig and Serdar Tasiran (Eds.). Springer, 286–304. https://doi.org/10.1007/978-3-030-25543-5_17
- [13] Sidi Mohamed Beillahi, Ahmed Bouajjani, and Constantin Enea. 2019. Robustness Against Transactional Causal Consistency. In *30th International Conference on Concurrency Theory, CONCUR 2019, August 27-30, 2019, Amsterdam, the Netherlands (LIPIcs, Vol. 140)*, Wan J. Fokink and Rob van Glabbeek (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 30:1–30:18. <https://doi.org/10.4230/LIPIcs.CONCUR.2019.30>
- [14] Hal Berenson, Philip A. Bernstein, Jim Gray, Jim Melton, Elizabeth J. O’Neil, and Patrick E. O’Neil. 1995. A Critique of ANSI SQL Isolation Levels. In *Proceedings of the 1995 ACM SIGMOD International Conference on Management of Data, San Jose, California, USA, May 22-25, 1995*, Michael J. Carey and Donovan A. Schneider (Eds.). ACM Press, 1–10. <https://doi.org/10.1145/223784.223785>
- [15] Giovanni Bernardi and Alexey Gotsman. 2016. Robustness against Consistency Models with Atomic Visibility. In *27th International Conference on Concurrency Theory, CONCUR 2016, August 23-26, 2016, Québec City, Canada (LIPIcs, Vol. 59)*, Josée Desharnais and Radha Jagadeesan (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:15. <https://doi.org/10.4230/LIPIcs.CONCUR.2016.7>
- [16] Ranadeep Biswas and Constantin Enea. 2019. On the complexity of checking transactional consistency. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 165:1–165:28. <https://doi.org/10.1145/3360591>
- [17] Ranadeep Biswas, Diptanshu Kakwani, Jyothi Vedurada, Constantin Enea, and Akash Lal. 2021. MonkeyDB: effectively testing correctness under weak isolation levels. *Proc. ACM Program. Lang.* 5, OOPSLA (2021), 1–27. <https://doi.org/10.1145/3485546>

- 1128 [18] Ahmed Bouajjani, Constantin Enea, and Enrique Román-Calvo. 2023. Dynamic Partial Order Reduction for Checking
 1129 Correctness Against Transaction Isolation Levels. arXiv:2303.12606 [cs.PL]
- 1130 [19] Martin Brain, James H. Davenport, and Alberto Griggio. 2017. Benchmarking Solvers, SAT-style. In *Proceedings of the*
 1131 *2nd International Workshop on Satisfiability Checking and Symbolic Computation co-located with the 42nd International*
 1132 *Symposium on Symbolic and Algebraic Computation (ISSAC 2017), Kaiserslautern, Germany, July 29, 2017 (CEUR Workshop*
 1133 *Proceedings, Vol. 1974)*, Matthew England and Vijay Ganesh (Eds.). CEUR-WS.org. <http://ceur-ws.org/Vol-1974/RP3.pdf>
- 1134 [20] Nathan Bronson, Zach Amsden, George Cabrera, Prasad Chakka, Peter Dimov, Hui Ding, Jack Ferris, Anthony Giardullo,
 1135 Sachin Kulkarni, Harry C. Li, Mark Marchukov, Dmitri Petrov, Lovro Puzar, Yee Jiun Song, and Venkateshwaran
 1136 Venkataramani. 2013. TAO: Facebook’s Distributed Data Store for the Social Graph. In *2013 USENIX Annual Technical*
 1137 *Conference, San Jose, CA, USA, June 26-28, 2013*, Andrew Birrell and Emin Gün Sirer (Eds.). USENIX Association, 49–60.
 1138 <https://www.usenix.org/conference/atc13/technical-sessions/presentation/bronson>
- 1139 [21] Lucas Brutschy, Dimitar K. Dimitrov, Peter Müller, and Martin T. Vechev. 2017. Serializability for eventual consistency:
 1140 criterion, analysis, and applications. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming*
 1141 *Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM,
 1142 458–472. <https://doi.org/10.1145/3009837.3009895>
- 1143 [22] Lucas Brutschy, Dimitar K. Dimitrov, Peter Müller, and Martin T. Vechev. 2018. Static serializability analysis for
 1144 causal consistency. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and*
 1145 *Implementation, PLDI 2018, Philadelphia, PA, USA, June 18-22, 2018*, Jeffrey S. Foster and Dan Grossman (Eds.). ACM,
 1146 90–104. <https://doi.org/10.1145/3192366.3192415>
- 1147 [23] Andrea Cerone, Giovanni Bernardi, and Alexey Gotsman. 2015. A Framework for Transactional Consistency Models
 1148 with Atomic Visibility. In *26th International Conference on Concurrency Theory, CONCUR 2015, Madrid, Spain, September*
 1149 *1.4, 2015 (LIPIcs, Vol. 42)*, Luca Aceto and David de Frutos-Escrig (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für
 1150 Informatik, 58–71. <https://doi.org/10.4230/LIPIcs.CONCUR.2015.58>
- 1151 [24] Andrea Cerone and Alexey Gotsman. 2018. Analysing Snapshot Isolation. *J. ACM* 65, 2 (2018), 11:1–11:41. <https://doi.org/10.1145/3152396>
- 1152 [25] Marek Chalupa, Krishnendu Chatterjee, Andreas Pavlogiannis, Nishant Sinha, and Kapil Vaidya. 2018. Data-centric
 1153 dynamic partial order reduction. *Proc. ACM Program. Lang.* 2, POPL (2018), 31:1–31:30. <https://doi.org/10.1145/3158119>
- 1154 [26] Krishnendu Chatterjee, Andreas Pavlogiannis, and Viktor Toman. 2019. Value-centric dynamic partial order reduction.
 1155 *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 124:1–124:29. <https://doi.org/10.1145/3360550>
- 1156 [27] Edmund M. Clarke, E. Allen Emerson, and A. Prasad Sistla. 1983. Automatic Verification of Finite State Concurrent
 1157 Systems Using Temporal Logic Specifications: A Practical Approach. In *Conference Record of the Tenth Annual ACM*
 1158 *Symposium on Principles of Programming Languages, Austin, Texas, USA, January 1983*, John R. Wright, Larry Landweber,
 1159 Alan J. Demers, and Tim Teitelbaum (Eds.). ACM Press, 117–126. <https://doi.org/10.1145/567067.567080>
- 1160 [28] Edmund M. Clarke, Orna Grumberg, Marius Minea, and Doron A. Peled. 1999. State Space Reduction Using Partial
 1161 Order Techniques. *Int. J. Softw. Tools Technol. Transf.* 2, 3 (1999), 279–287. <https://doi.org/10.1007/s100090050035>
- 1162 [29] Giuseppe DeCandia, Deniz Hastorun, Madan Jambani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin,
 1163 Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. 2007. Dynamo: amazon’s highly available key-
 1164 value store. In *Proceedings of the 21st ACM Symposium on Operating Systems Principles 2007, SOSP 2007, Stevenson,*
 1165 *Washington, USA, October 14-17, 2007*, Thomas C. Bressoud and M. Frans Kaashoek (Eds.). ACM, 205–220. <https://doi.org/10.1145/1294261.1294281>
- 1166 [30] Djelle Eddine Difallah, Andrew Pavlo, Carlo Curino, and Philippe Cudré-Mauroux. 2013. OLTP-Bench: An Extensible
 1167 Testbed for Benchmarking Relational Databases. *Proc. VLDB Endow.* 7, 4 (2013), 277–288. <https://doi.org/10.14778/2732240.2732246>
- 1168 [31] Alan D. Fekete, Dimitrios Liarokapis, Elizabeth J. O’Neil, Patrick E. O’Neil, and Dennis E. Shasha. 2005. Making snapshot
 1169 isolation serializable. *ACM Trans. Database Syst.* 30, 2 (2005), 492–528. <https://doi.org/10.1145/1071610.1071615>
- 1170 [32] Cormac Flanagan and Patrice Godefroid. 2005. Dynamic partial-order reduction for model checking software. In
 1171 *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005,*
 1172 *Long Beach, California, USA, January 12-14, 2005*, Jens Palsberg and Martín Abadi (Eds.). ACM, 110–121. <https://doi.org/10.1145/1040305.1040315>
- 1173 [33] Yifan Gan, Xueyuan Ren, Drew Ripberger, Spyros Blanas, and Yang Wang. 2020. IsoDiff: Debugging Anomalies Caused
 1174 by Weak Isolation. *Proc. VLDB Endow.* 13, 12 (July 2020), 27732786. <https://doi.org/10.14778/3407790.3407860>
- 1175 [34] Patrice Godefroid. 1996. *Partial-Order Methods for the Verification of Concurrent Systems - An Approach to the State-*
 1176 *Explosion Problem*. Lecture Notes in Computer Science, Vol. 1032. Springer. <https://doi.org/10.1007/3-540-60761-7>
- 1177 [35] Patrice Godefroid. 1997. Model Checking for Programming Languages using Verisoft. In *Conference Record of*
 1178 *POPL ’97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at the*
 1179 *Symposium, Paris, France, 15-17 January 1997*, Peter Lee, Fritz Henglein, and Neil D. Jones (Eds.). ACM Press, 174–186.
 1180 <https://doi.org/10.1145/263699.263717>

- 1177 [36] Alexey Gotsman, Hongseok Yang, Carla Ferreira, Mahsa Najafzadeh, and Marc Shapiro. 2016. 'Cause I'm strong enough:
1178 reasoning about consistency choices in distributed systems. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT*
1179 *Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav
1180 Bodik and Rupak Majumdar (Eds.). ACM, 371–384. <https://doi.org/10.1145/2837614.2837625>
- 1181 [37] Sudhir Jorwekar, Alan D. Fekete, Krithi Ramamritham, and S. Sudarshan. 2007. Automating the Detection of Snapshot
1182 Isolation Anomalies. In *Proceedings of the 33rd International Conference on Very Large Data Bases, University of Vienna,*
1183 *Austria, September 23-27, 2007*, Christoph Koch, Johannes Gehrke, Minos N. Garofalakis, Divesh Srivastava, Karl Aberer,
1184 Anand Deshpande, Daniela Florescu, Chee Yong Chan, Venkatesh Ganti, Carl-Christian Kanne, Wolfgang Klas, and
1185 Erich J. Neuhold (Eds.). ACM, 1263–1274. <http://www.vldb.org/conf/2007/papers/industrial/p1263-jorwekar.pdf>
- 1186 [38] Gowtham Kaki, Kapil Earanky, K. C. Sivaramakrishnan, and Suresh Jagannathan. 2018. Safe replication through
1187 bounded concurrency verification. *Proc. ACM Program. Lang.*, 2, OOPSLA (2018), 164:1–164:27. <https://doi.org/10.1145/3276534>
- 1188 [39] Michalis Kokologiannakis, Iason Marmanis, Vladimir Gladstein, and Viktor Vafeiadis. 2022. Truly stateless, optimal
1189 dynamic partial order reduction. *Proc. ACM Program. Lang.*, 6, POPL (2022), 1–28. <https://doi.org/10.1145/3498711>
- 1190 [40] Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. 2019. Model checking for weakly consistent libraries.
1191 In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI*
1192 *2019, Phoenix, AZ, USA, June 22-26, 2019*, Kathryn S. McKinley and Kathleen Fisher (Eds.). ACM, 96–110. <https://doi.org/10.1145/3314221.3314609>
- 1193 [41] Michalis Kokologiannakis and Viktor Vafeiadis. 2020. HMC: Model Checking for Hardware Memory Models. In *ASPLOS*
1194 *'20: Architectural Support for Programming Languages and Operating Systems, Lausanne, Switzerland, March 16-20, 2020*,
1195 James R. Larus, Luis Ceze, and Karin Strauss (Eds.). ACM, 1157–1171. <https://doi.org/10.1145/3373376.3378480>
- 1196 [42] Leslie Lamport. 1978. Time, Clocks, and the Ordering of Events in a Distributed System. *Commun. ACM* 21, 7 (1978),
1197 558–565. <https://doi.org/10.1145/359545.359563>
- 1198 [43] Cheng Li, João Leitão, Allen Clement, Nuno M. Prego, Rodrigo Rodrigues, and Viktor Vafeiadis. 2014. Automating
1199 the Choice of Consistency Levels in Replicated Systems. In *2014 USENIX Annual Technical Conference, USENIX ATC*
1200 *'14, Philadelphia, PA, USA, June 19-20, 2014*, Garth Gibson and Nikolai Zeldovich (Eds.). USENIX Association, 281–292.
1201 https://www.usenix.org/conference/atc14/technical-sessions/presentation/li_cheng_2
- 1202 [44] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. 2011. Don't settle for eventual: scalable
1203 causal consistency for wide-area storage with COPS. In *Proceedings of the 23rd ACM Symposium on Operating Systems*
1204 *Principles 2011, SOSP 2011, Cascais, Portugal, October 23-26, 2011*, Ted Wobber and Peter Druschel (Eds.). ACM, 401–416.
1205 <https://doi.org/10.1145/2043556.2043593>
- 1206 [45] Antoni W. Mazurkiewicz. 1986. Trace Theory. In *Petri Nets: Central Models and Their Properties, Advances in Petri*
1207 *Nets 1986, Part II, Proceedings of an Advanced Course, Bad Honnef, Germany, 8-19 September 1986 (Lecture Notes in*
1208 *Computer Science, Vol. 255)*, Wilfried Brauer, Wolfgang Reisig, and Grzegorz Rozenberg (Eds.). Springer, 279–324.
1209 https://doi.org/10.1007/3-540-17906-2_30
- 1210 [46] Kartik Nagar and Suresh Jagannathan. 2018. Automated Detection of Serializability Violations Under Weak Consistency.
1211 In *29th International Conference on Concurrency Theory, CONCUR 2018, September 4-7, 2018, Beijing, China (LIPIcs,*
1212 *Vol. 118)*, Sven Schewe and Lijun Zhang (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 41:1–41:18.
1213 <https://doi.org/10.4230/LIPIcs.CONCUR.2018.41>
- 1214 [47] Sreeja S. Nair, Gustavo Petri, and Marc Shapiro. 2020. Proving the Safety of Highly-Available Distributed Objects.
1215 In *Programming Languages and Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the*
1216 *European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings*
1217 *(Lecture Notes in Computer Science, Vol. 12075)*, Peter Müller (Ed.). Springer, 544–571. https://doi.org/10.1007/978-3-030-44914-8_20
- 1218 [48] Brian Norris and Brian Demsky. 2013. CDSchecker: checking concurrent data structures written with C/C++ atomics.
1219 In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages*
1220 *& Applications, OOPSLA 2013, part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, Antony L. Hosking,
1221 Patrick Th. Eugster, and Cristina V. Lopes (Eds.). ACM, 131–150. <https://doi.org/10.1145/2509136.2509514>
- 1222 [49] Burcu Kulahcioglu Ozkan. 2020. Verifying Weakly Consistent Transactional Programs Using Symbolic Execution.
1223 In *Networked Systems - 8th International Conference, NETYS 2020, Marrakech, Morocco, June 3-5, 2020, Proceedings*
1224 *(Lecture Notes in Computer Science, Vol. 12129)*, Chryssis Georgiou and Rupak Majumdar (Eds.). Springer, 261–278.
1225 https://doi.org/10.1007/978-3-030-67087-0_17
- [50] Christos H. Papadimitriou. 1979. The serializability of concurrent database updates. *J. ACM* 26, 4 (1979), 631–653.
<https://doi.org/10.1145/322154.322158>
- [51] Andrew Pavlo. 2017. What Are We Doing With Our Lives? Nobody Cares About Our Concurrency Control Research.
In *Proceedings of the 2017 ACM International Conference on Management of Data (Chicago, Illinois, USA) (SIGMOD '17)*.
Association for Computing Machinery, New York, NY, USA, 3. <https://doi.org/10.1145/3035918.3056096>

- 1226 [52] Jos Rolando Guay Paz. 2018. *Microsoft Azure Cosmos DB Revealed: A Multi-Modal Database Designed for the Cloud* (1st
1227 ed.). Apress, USA.
- 1228 [53] Doron A. Peled. 1993. All from One, One for All: on Model Checking Using Representatives. In *Computer Aided*
1229 *Verification, 5th International Conference, CAV '93, Elounda, Greece, June 28 - July 1, 1993, Proceedings (Lecture Notes in*
1230 *Computer Science, Vol. 697)*, Costas Courcoubetis (Ed.). Springer, 409–423. https://doi.org/10.1007/3-540-56922-7_34
- 1231 [54] Jean-Pierre Queille and Joseph Sifakis. 1982. Specification and verification of concurrent systems in CESAR. In
1232 *International Symposium on Programming, 5th Colloquium, Torino, Italy, April 6-8, 1982, Proceedings (Lecture Notes*
1233 *in Computer Science, Vol. 137)*, Mariangiola Dezani-Ciancaglini and Ugo Montanari (Eds.). Springer, 337–351. https://doi.org/10.1007/3-540-11494-7_22
- 1234 [55] Kia Rahmani, Kartik Nagar, Benjamin Delaware, and Suresh Jagannathan. 2019. CLOTHO: directed test generation for
1235 weakly consistent database systems. *Proc. ACM Program. Lang.* 3, OOPSLA (2019), 117:1–117:28. <https://doi.org/10.1145/3360543>
- 1236 [56] K. C. Sivaramakrishnan, Gowtham Kaki, and Suresh Jagannathan. 2015. Declarative programming over eventually
1237 consistent data stores. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and*
1238 *Implementation, Portland, OR, USA, June 15-17, 2015*, David Grove and Stephen M. Blackburn (Eds.). ACM, 413–424.
1239 <https://doi.org/10.1145/2737924.2737981>
- 1240 [57] TPC. 2010. . Technical Report. Transaction Processing Performance Council. [http://www.tpc.org/tpc_documents_](http://www.tpc.org/tpc_documents_current_versions/pdf/tpc-c_v5.11.0.pdf)
1241 [current_versions/pdf/tpc-c_v5.11.0.pdf](http://www.tpc.org/tpc_documents_current_versions/pdf/tpc-c_v5.11.0.pdf)
- 1242 [58] Antti Valmari. 1989. Stubborn sets for reduced state space generation. In *Advances in Petri Nets 1990 [10th International*
1243 *Conference on Applications and Theory of Petri Nets, Bonn, Germany, June 1989, Proceedings] (Lecture Notes in Computer*
1244 *Science, Vol. 483)*, Grzegorz Rozenberg (Ed.). Springer, 491–515. https://doi.org/10.1007/3-540-53863-1_36
- 1245 [59] Willem Visser, Corina S. Pasareanu, and Sarfraz Khurshid. 2004. Test input generation with java Pathfinder. In
1246 *Proceedings of the ACM/SIGSOFT International Symposium on Software Testing and Analysis, ISSTA 2004, Boston,*
1247 *Massachusetts, USA, July 11-14, 2004*, George S. Avrunin and Gregg Rothermel (Eds.). ACM, 97–107. <https://doi.org/10.1145/1007512.1007526>
- 1248 [60] Todd Warszawski and Peter Bailis. 2017. ACIDRain: Concurrency-Related Attacks on Database-Backed Web Applica-
1249 tions. In *Proceedings of the 2017 ACM International Conference on Management of Data (Chicago, Illinois, USA) (SIGMOD*
1250 *'17)*. Association for Computing Machinery, New York, NY, USA, 520. <https://doi.org/10.1145/3035918.3064037>
- 1251 [61] ANSI X3. 1992. 135-1992. *American National Standard for Information Systems-Database Language-SQL*. Technical
1252 Report.
- 1253
- 1254
- 1255
- 1256
- 1257
- 1258
- 1259
- 1260
- 1261
- 1262
- 1263
- 1264
- 1265
- 1266
- 1267
- 1268
- 1269
- 1270
- 1271
- 1272
- 1273
- 1274

A AXIOMATIC LEVELS: READ COMMITTED AND READ ATOMIC.

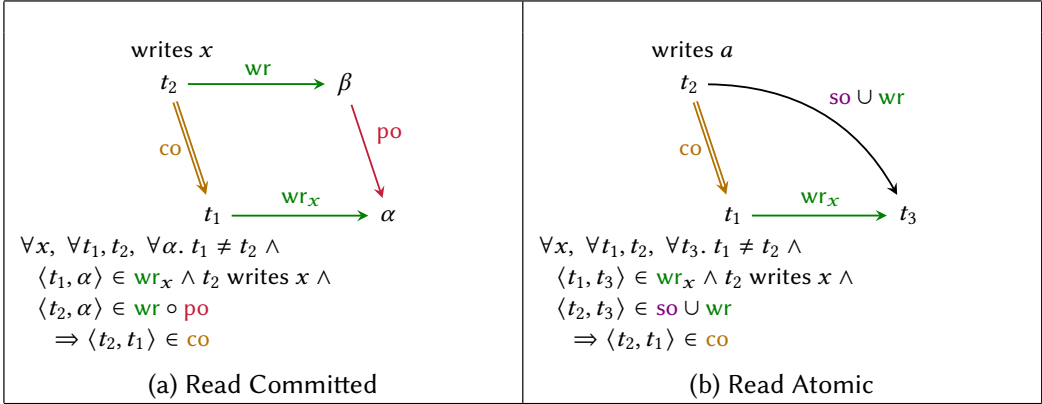


Fig. A.1. Axioms defining isolation levels. The reflexive and transitive, resp., transitive, closure of a relation rel is denoted by rel^* , resp., rel^+ . Also, \circ denotes the composition of two relations, i.e., $rel_1 \circ rel_2 = \{\langle a, b \rangle \mid \exists c. \langle a, c \rangle \in rel_1 \wedge \langle c, b \rangle \in rel_2\}$.

The axioms defined above in Figure A.1 define the homonymous isolation levels *Read Atomic* (also called Repeatable Read in the literature) and *Read Committed*.

B RULES OF THE OPERATIONAL SEMANTICS (SECTION 2.3).

1324	1325	1326	1327	1328	1329	1330	1331	1332	1333	1334	1335	1336	1337	1338	1339	1340	1341	1342	1343	1344	1345	1346	1347	1348	1349	1350	1351	1352	1353	1354	1355	1356	1357						
			SPAWN	$\frac{t \text{ fresh} \quad e \text{ fresh} \quad P(j) = \text{begin}; \text{Body}; \text{commit}; S \quad \vec{B}(j) = \epsilon}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h \oplus_j \langle t, \{\langle e, \text{begin} \rangle\}, \emptyset \rangle, \vec{\gamma}[j \mapsto \emptyset], \vec{B}[j \mapsto \text{Body}; \text{commit}], P[j \mapsto S]}$																																			
				$\frac{\text{IF-TRUE} \quad \psi(\vec{x})[x \mapsto \vec{\gamma}(j)(x) : x \in \vec{x}] \text{ true} \quad \vec{B}(j) = \text{if}(\psi(\vec{x}))\{\text{Instr}\}; B}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h, \vec{\gamma}, \vec{B}[j \mapsto \text{Instr}; B], P}$																																			
				$\frac{\text{IF-FALSE} \quad \psi(\vec{x})[x \mapsto \vec{\gamma}(j)(x) : x \in \vec{x}] \text{ false} \quad \vec{B}(j) = \text{if}(\psi(\vec{x}))\{\text{Instr}\}; B}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h, \vec{\gamma}, \vec{B}[j \mapsto B], P}$																																			
				$\frac{\text{LOCAL} \quad v = \vec{\gamma}(j)(e) \quad \vec{B}(j) = a := e; B}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h, \vec{\gamma}[(j, a) \mapsto v], \vec{B}[j \mapsto B], P}$																																			
				$\frac{\text{WRITE} \quad v = \vec{\gamma}(j)(x) \quad e \text{ fresh} \quad \vec{B}(j) = \text{write}(x, a); B \quad h \oplus_j \langle e, \text{write}(x, v) \rangle \text{ satisfies } I}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h \oplus_j \langle e, \text{write}(x, v) \rangle, \vec{\gamma}, \vec{B}[j \mapsto B], P}$																																			
				$\frac{\text{READ-LOCAL} \quad \text{writes}(\text{last}(h, j)) \text{ contains a write}(x, v) \text{ event} \quad e \text{ fresh} \quad \vec{B}(j) = a := \text{read}(x); B}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h \oplus_j \langle e, \text{read}(x) \rangle, \vec{\gamma}[(j, a) \mapsto v], \vec{B}[j \mapsto B], P}$																																			
				$\frac{\text{READ-EXTERN} \quad \begin{array}{l} \text{writes}(\text{last}(h, j)) \text{ does not contain a write}(x, v) \text{ event} \quad e \text{ fresh} \quad \vec{B}(j) = a := \text{read}(x); B \\ h = \langle T, \text{so}, \text{wr} \rangle \quad t = \text{last}(h, j) \quad \text{write}(x, v) \in \text{writes}(t') \text{ with } t' \in \text{commTrans}(h) \text{ and } t \neq t' \\ h' = (h \oplus_j \langle e, \text{read}(x) \rangle) \oplus \text{wr}(t', e) \quad h' \text{ satisfies } I \end{array}}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h', \vec{\gamma}[(j, a) \mapsto v], \vec{B}[j \mapsto B], P}$																																			
				$\frac{\text{COMMIT} \quad e \text{ fresh} \quad \vec{B}(j) = \text{commit}}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h \oplus_j \langle e, \text{commit} \rangle, \vec{\gamma}, \vec{B}[j \mapsto \epsilon], P}$				$\frac{\text{ABORT} \quad e \text{ fresh} \quad \vec{B}(j) = \text{abort}; B}{h, \vec{\gamma}, \vec{B}, P \Rightarrow_I h \oplus_j \langle e, \text{abort} \rangle, \vec{\gamma}, \vec{B}[j \mapsto \epsilon], P}$																															

Fig. B.1. An operational semantics for transactional programs. Above, $\text{last}(h, j)$ denotes the last transaction log in the session order $\text{so}(j)$ of h , and $\text{commTrans}(h)$ denotes the set of transaction logs in h that are committed

Figure B.1 uses the following notation. Let h be a history that contains a representation of so as above. We use $h \oplus_j \langle t, E, \text{po}_t \rangle$ to denote a history where $\langle t, E, \text{po}_t \rangle$ is appended to $\text{so}(j)$. Also, for an event e , $h \oplus_j e$ is the history obtained from h by adding e to the last transaction log in $\text{so}(j)$ and as a last event in the program order of this log (i.e., if $\text{so}(j) = \sigma; \langle t, E, \text{po}_t \rangle$, then the session order so' of $h \oplus_j e$ is defined by $\text{so}'(k) = \text{so}(k)$ for all $k \neq j$ and $\text{so}'(j) = \sigma; \langle t, E \cup \{e\}, \text{po}_t \cup \{(e', e) : e' \in E\} \rangle$). Finally, for a history $h = \langle T, \text{so}, \text{wr} \rangle$, $h \oplus \text{wr}(t, e)$ is the history obtained from h by adding (t, e) to the write-read relation.

SPAWN starts a new transaction in a session j provided that this session has no live transaction ($\vec{B}(j) = \epsilon$). It adds a transaction log with a single begin event to the history and schedules the body

1373 of the transaction. IF-TRUE and IF-FALSE check the truth value of a Boolean condition of an if
1374 conditional. LOCAL models the execution of an assignment to a local variable which does not impact
1375 the stored history. READ-LOCAL and READ-EXTERN concern read instructions. READ-LOCAL handles
1376 the case where the read follows a write on the variable x in the same transaction: the read returns
1377 the value written by the last write on x in that transaction. Otherwise, READ-EXTERN corresponds to
1378 reading a value written in another transaction t' . The transaction t' is chosen non-deterministically
1379 as long as extending the current history with the write-read dependency associated to this choice
1380 leads to a history that still satisfies I . READ-EXTERN applies only when the executing transaction
1381 contains no write on the same variable. COMMIT confirms the end of a transaction making its writes
1382 visible while ABORT ends the transaction's execution immediately.

1383

1384

1385

1386

1387

1388

1389

1390

1391

1392

1393

1394

1395

1396

1397

1398

1399

1400

1401

1402

1403

1404

1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1418

1419

1420

1421

C PROOF OF THEOREM 3.4

THEOREM 3.4. *Causal Consistency, Read Atomic, and Read Committed are causally-extensible.*

PROOF. Let I be an isolation level in $\{\text{CC}, \text{RA}, \text{RC}\}$. We show that any commit order co justifying that a history h is I -consistent can also be used to justify that a causal extension h' of a $(\text{so} \cup \text{wr})^*$ -maximal pending transaction t in h with an event e is I -consistent as well. We consider a causal extension h' where if e is a read event, then it reads from the last transaction t_w in co such that t_w writes $\text{var}(e)$ and $(t_w, t) \in (\text{so} \cup \text{wr})^+$. Assume by contradiction that this is not the case. Let $\phi_{\text{CC}}(h', t', e') = t' (\text{so} \cup \text{wr})^+ \text{tr}(h', e')$, $\phi_{\text{RA}}(h', t', e') = t' (\text{so} \cup \text{wr}) \text{tr}(h', e')$ and $\phi_{\text{RC}}(h', t', e') = t' (\text{wr} \circ \text{po}) e'$ be sub-formulas of the axioms defining the corresponding isolation level. Then, h' contains transactions t_1, t_2, t_3 such that t_2 writes some variable x , t_3 contains some read event e' , $(t_1, e') \in \text{wr}_x$ and $\phi_I(h', t_2, e')$ but $(t_1, t_2) \in \text{co}$. The assumption concerning co implies that the extended transaction t is one of t_1, t_2, t_3 (otherwise, co would not be a “valid” commit order for h). Since t is $(\text{so} \cup \text{wr})^+$ -maximal in h , we have that $t \notin \{t_1, t_2\}$. If e is *not* a read event, or if e is a read event different from e' , then $t \neq t_3$, as t_1, t_2 and t_3 would satisfy the same constraints in h , which is impossible by the hypothesis. Otherwise, if $e = e'$, then this contradicts the choice we made for the transaction t_w that e reads from. Since $(t_1, t_2) \in \text{co}$ and t_2 writes $\text{var}(e)$, it means that $t_w = t_1$ is not maximal w.r.t. co among transactions that write $\text{var}(e)$ and precede t in $(\text{so} \cup \text{wr})^+$. Both cases lead to a contradiction, which implies that h' is I -consistent, and therefore the theorem holds. \square

D PROOF OF THEOREM 6.1

1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519

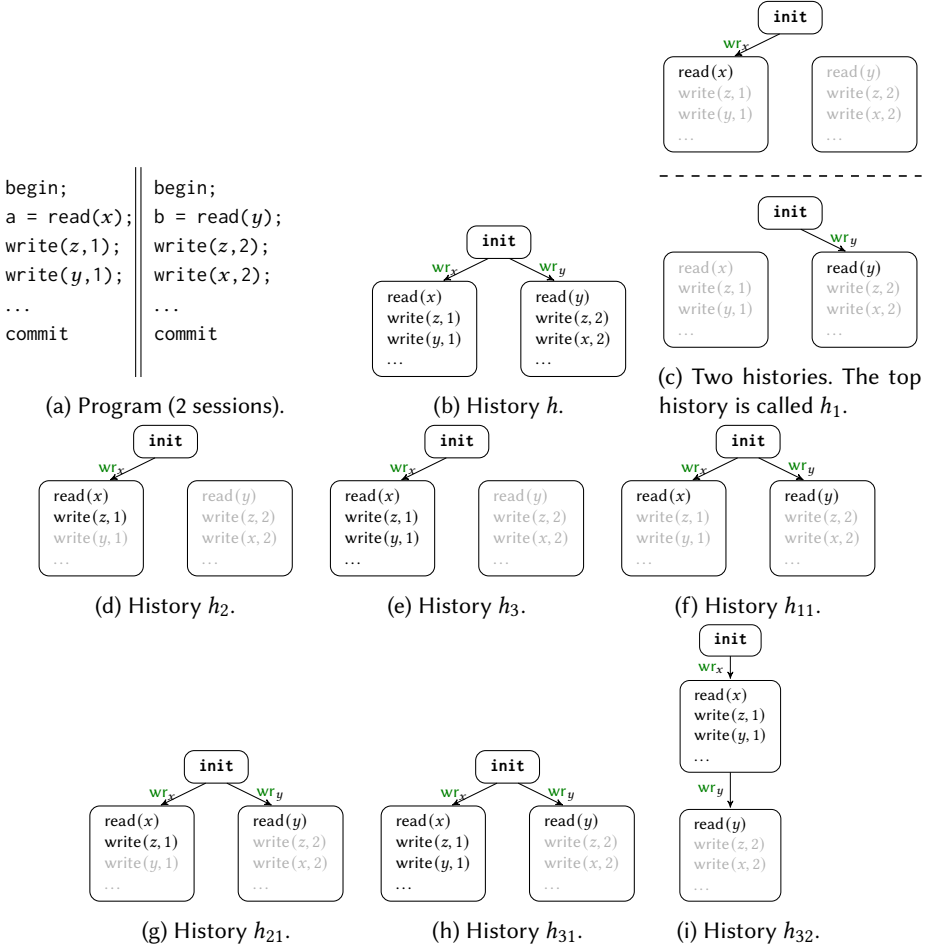


Fig. D.1. A program and some partial histories. Events in grey are not yet added to the history. For h_3 , h_{31} and h_{32} , the number of events that follow $write(y, 1)$ and $write(x, 2)$ is not important (we use black ... to signify that).

THEOREM 6.1. *If I is Snapshot Isolation or Serializability, there exists no EXPLORE algorithm that is I -sound, I -complete, and strongly optimal.*

PROOF. We consider the program in Figure D.1a, and show that any concrete instance of the EXPLORE function in Algorithm 1 can not be both I -complete and strongly optimal. This program contains two transactions, where only the first three instructions in each transaction are important. We show that if EXPLORE is I -complete, then it will necessarily be called recursively on a history h like in Figure D.1b which does not satisfy I , thereby violating strong optimality. In the history h , both Snapshot Isolation and Serializability forbid the two reads reading initial values while the writes following them are also executed (committed).

Assuming that the function NEXT is not itself blocking (which would violate strong optimality), the EXPLORE will be called recursively on exactly one of the two histories in Figure D.1c, depending

1520 on which of the two reads is returned first by NEXT. We will continue our discussion with the
 1521 history h_1 on the top of Figure D.1c. The other case is similar (symmetric).

1522 From h_1 , depending on the order defined by NEXT between $\text{read}(y)$ and $\text{write}(y, 1)$, EXPLORE
 1523 can be called recursively either on h_1 , on h_2 in Figure D.1d, or on h_3 in Figure D.1e before adding
 1524 $\text{read}(y)$. From h_1 and h_2 , EXPLORE explores h_{11} in Figure D.1f and h_{21} in Figure D.1g respectively;
 1525 while from h_3 two alternative histories may be explored: h_{31} and h_{32} in Figure D.1h and Figure D.1i
 1526 respectively.

1527 However, from histories h_{11} , h_{21} or h_{31} EXPLORE will necessarily be called recursively on a history
 1528 h like in Figure D.1b which does not satisfy I , thereby violating strong optimality. Thus, any EXPLORE
 1529 implementation that is strong optimal should only explore h_{32} . In such case, by the restrictions
 1530 on the SWAP function (defined in Section 4), any extension of h_{32} does not allow to explore the
 1531 history where $\text{read}(x)$ reads from $\text{write}(x, 2)$: any outcome of a re-ordering between two contiguous
 1532 subsequences α and β must be prefix of h_e when the events in α are taken out. In particular, for
 1533 any extension h' of h_{32} and pair of contiguous sequences α, β such that $h' \setminus \alpha$ is a prefix of h' , if an
 1534 event from the second transaction belongs to β , $\text{read}(y)$ must also be in β . Therefore, $\text{write}(x, 2)$
 1535 must also be in β , and so $\text{read}(x)$ must be. Analogously, if $\text{read}(x)$ belongs to β , **init** belongs to it.
 1536 Altogether, if β contains any element, then α must be empty; so no swaps can be produced from
 1537 h_{32} . To conclude, in this case EXPLORE violates I -completeness.

□

1538
 1539
 1540
 1541
 1542
 1543
 1544
 1545
 1546
 1547
 1548
 1549
 1550
 1551
 1552
 1553
 1554
 1555
 1556
 1557
 1558
 1559
 1560
 1561
 1562
 1563
 1564
 1565
 1566
 1567
 1568

E PROOF OF THEOREM 5.1

THEOREM 5.1. *For any prefix-closed and causally extensible isolation level I , EXPLORE-CE is I -sound, I -complete, strongly optimal, and polynomial space.*

As explained in Section 5.4, I -soundness, the polynomial space bound, and the part of strong completeness that refers to not engaging in fruitless explorations follow directly from definitions. In the following, we focus on I -completeness and then optimality. For the sake of the proof's readability, we will omit all local states of the algorithm's definition during the proof. Therefore, we consider programs where we can describe all their events.

E.1 Completeness

By definition, EXPLORE-CE is I -complete if for any given program P , it outputs every history in $\text{hist}_I(P)$. Let $h \in \text{hist}_I(P)$. Our objective is to produce a computable path of ordered histories that lead to h (i.e. a (finite) ordered collection of ordered histories such that $h_0 = \emptyset$ and for every n , if $e = \text{NEXT}(h_n)$, either $h_{n+1} = h_n \oplus e$, $h_{n+1} = h_n \oplus \text{wr}(e, t)$ for some $t \in h_n$ or $h_{n+1} = \text{SWAP}(h_n, r, t)$ for some $r, t \in h_n$).

However, the algorithm EXPLORE-CE works with ordered histories. Therefore, we first have to furnish h with a total order called *canonical order* that, if h were reachable, it would coincide with its history order. Secondly, we describe a function PREV defined over the set of all partial histories that, if h is reachable, PREV(h) returns the previous history of h computed by EXPLORE-CE. Then, we prove that there exists a finite collection of histories $H = \{h_i\}_{i=0}^n$ such that $h_n = h$, $h_0 = \emptyset$ and $h_i = \text{PREV}(h_{i+1})$. As it ends in the initial state, we can therefore prove that this collection conforms an actual computable path; which allow us to conclude that h is reachable. Nevertheless, for proving both the equivalence between history order and canonical order and the soundness of function PREV we will define the notion of *or-respectfulness*, an invariant satisfied by every reachable history based on the events' relative positions in the oracle order.

E.1.1 Canonical order.

As mentioned, we need to formally define a total order for every history that coincide on reachable histories with the history order. For achieving it, we analyze how the algorithm orders transaction logs in a history. In particular, we observe that if two transactions t, t' have a $(\text{so} \cup \text{wr})^*$ dependency, the history order in the algorithm orders them analogously. But if they are $(\text{so} \cup \text{wr})^*$ -incomparable, the algorithm prioritizes the one that is read by a smaller read event according *or*. Combining both arguments recursively we obtain a *canonical order* for a history, which is formally defined with the function presented below.

The function CANONICALORDER produces a relation between transactions in a history, denoted \leq^h . In algorithm 3's description, we denote \perp to represent the end of the program, which always exists, and that is *so*-related with every single transaction.

Firstly, we prove our canonical order is well defined for every pair of transactions.

LEMMA E.1. *For every history h , event e and transaction t , $\text{DEP}(h, t, \min_{<_{\text{or}}} \text{DEP}(h, t, e)) \subseteq \text{DEP}(h, t, e)$. Moreover, if $\text{DEP}(h, t, e) \neq t$, the inclusion is strict.*

PROOF. Let $r' = \min_{<_{\text{or}}} \text{DEP}(h, t, e)$. If $\text{DEP}(h, t, r') = t$ the lemma is trivially proved, so let's suppose there exists $r \in \text{DEP}(h, t, r') \setminus t$. Then, $\exists t'$ s.t. $t [\text{so} \cup \text{wr}]^* t' \wedge t' [\text{wr}] r \wedge \text{tr}(h, r) [\text{so} \cup \text{wr}]^+ \text{tr}(h, r')$ and $\exists t''$ s.t. $t [\text{so} \cup \text{wr}]^* t'' \wedge t'' [\text{wr}] r' \wedge \text{tr}(h, r') [\text{so} \cup \text{wr}]^+ \text{tr}(h, e)$; so $\text{tr}(h, r) [\text{so} \cup \text{wr}]^+ \text{tr}(h, r') [\text{so} \cup \text{wr}]^+ \text{tr}(h, e)$. In other words, $r \in \text{DEP}(h, t, e)$. The moreover comes trivially as $r' \notin \text{DEP}(h, t, r')$. \square

Algorithm 3 CANONICAL ORDER

```

1618 1: procedure CANONICALORDER( $h, t, t'$ )
1619 2:   return  $t$  [ $\text{so} \cup \text{wr}$ ] $^* t' \vee$ 
1620 3:      $(\neg(t' [\text{so} \cup \text{wr}]^* t) \wedge \text{MINIMALDEPENDENCY}(h, t, t', \perp))$ 
1621 4: procedure MINIMALDEPENDENCY( $h, t, t', e$ )
1622 5:   let  $a = \min_{<_{\text{or}}} \text{DEP}(h, t, e)$ ;  $a' = \min_{<_{\text{or}}} \text{DEP}(h, t', e)$ 
1623 6:   if  $a \neq a'$  then
1624 7:     return  $a <_{\text{or}} a'$ 
1625 8:   else
1626 9:     return MINIMALDEPENDENCY( $h, t, t', a$ )
1627 10: procedure DEP( $h, t, e$ )
1628 11: return  $\{r \mid \exists t' \text{ s.t. } t [\text{so} \cup \text{wr}]^* t' \wedge t' [\text{wr}] r \wedge \text{tr}(h, r) [\text{so} \cup \text{wr}]^+ \text{tr}(h, e)\} \cup t$ 

```

LEMMA E.2. For every pair of distinct transactions t, t' , MINIMALDEPENDENCY(h, t, t', \perp) always halts.

PROOF. Let's suppose by contrapositive that MINIMALDEPENDENCY(h, t, t', \perp) does not halt. Therefore, there would exist an infinite chain of events $e_n, n \in \mathbb{N}$ such that $e_0 = \perp, e_{n+1} = \min_{\text{or}} \text{DEP}(h, t, e_n) = \min_{\text{or}} \text{DEP}(h, t', e_n)$. Firstly, as h is finite, so are both $\text{DEP}(h, t, e_n)$ and $\text{DEP}(h, t', e_n)$. Moreover, if $e_n \notin t$, $\text{DEP}(h, t, e_{n+1}) \subseteq \text{DEP}(h, t, e_n)$ (and analogously for t'). Therefore, there exist some indexes n_0, m_0 such that $e_{n_0} \in t$ and $e_{m_0} \in t'$. Let $k = \max\{n_0, m_0\}$. Because ; but if $e_n \in t$, $t = \text{DEP}(h, t, e_n)$ and $e_{n+1} = e_n$, so $e_k = e_{n_0}$ and $e_k = e_{m_0}$. Therefore $e_k \in t \cap t'$; so $t = t'$ as transaction logs do not share events; which contradict the assumptions. \square

COROLLARY E.3. The relation \leq^h is well defined for every pair of transactions.

PROOF. As by lemma E.2, we know that MINIMALDEPENDENCY(h, t, t', \perp) always halts if $t \neq t'$; it is clear that CANONICALORDER(h, t, t') also does it. Therefore, the relation is well defined. \square

Now that \leq^h has been proved a well defined relation between each pair of transactions, let us prove that it is indeed a total order.

LEMMA E.4. The relation \leq^h is a total order.

PROOF.

- **Strongly connection** Let t_1, t_2 s.t. $t_1 \not\leq^h t_2$. If $t_2 [\text{so} \cup \text{wr}]^* t_1$, then $t_2 \leq^h t_1$. Otherwise, as $\neg(t_1 [\text{so} \cup \text{wr}]^* t_2)$ and MINIMALDEPENDENCY halts (lemma E.2) either MINIMALDEPENDENCY(h, t_1, t_2, \perp) or MINIMALDEPENDENCY(h, t_2, t_1, \perp) holds. But as $t_1 \not\leq^h t_2$, $t_2 \leq^h t_1$.
- **Reflexivity**: By definition, for every t , $t \leq^h t$.
- **Transitivity**: Let t_1, t_2, t_3 three distinct transactions such that $t_1 \leq^h t_2$ and $t_2 \leq^h t_3$. Clearly, if $t_1 [\text{so} \cup \text{wr}]^* t_3$, $t_1 \leq^h t_3$. However, if $t_3 [\text{so} \cup \text{wr}]^* t_1$, we would find one of the following three scenarios:
 - $t_1 [\text{so} \cup \text{wr}]^* t_2$, which is impossible by strong connectivity as that would mean $t_3 \leq^h t_2$.
 - $t_2 [\text{so} \cup \text{wr}]^* t_3$, which is also impossible by strong connectivity, as $t_2 \leq^h t_1$.
 - $\neg(t_1 [\text{so} \cup \text{wr}]^* t_2)$ and $\neg(t_2 [\text{so} \cup \text{wr}]^* t_3)$. Then, let us call $e_0^i = \perp$ and $e_{n+1}^i = \min_{<_{\text{or}}} \text{DEP}(h, t_i, e_n^i)$ for $i \in \{1, 2, 3\}$. Let's prove by induction that if for every $k < n$ $e_n^1 \notin t^1$, then $e_n^1 = e_n^2 = e_n^3$. Clearly this hold for $n = 0$ and, assuming it holds for every

1667 $k \leq n - 1$, as $t_1 \leq^h t_2$, $t_2 \leq^h t_3$, we know $e_n^1 \leq_{\text{or}} e_n^2 \leq_{\text{or}} e_n^3$ and as $t^3 [\text{so} \cup \text{wr}]^* t^1$,
 1668 if $e_n^1 \notin t^1$, $e_n^3 \leq_{\text{or}} e_n^1$. In other words, they coincide. However, by lemma E.2, we
 1669 know MINIMALDEPENDENCY(h, t^1, t^3, \perp) halts, so there exists some minimal n_0 such
 1670 that $e_{n_0}^1 \in t^1$; so $e_{n_0}^2 \in t_1$. That implies $t^2 [\text{so} \cup \text{wr}]^* t_1$; which is impossible as $t_1 \leq^h t_2$.
 1671 We deduce then that either $t_1 [\text{so} \cup \text{wr}]^* t_3$ or $\neg(t_3 [\text{so} \cup \text{wr}]^* t_1)$. In the latter case, let's
 1672 take the sequence e_n^i , $i \in \{1, 2, 3\}$ defined in the last paragraph. Then, as by lemma E.2
 1673 MINIMALDEPENDENCY(h, t_1, t_3, \perp) halts, there exists a maximum index n_0 such that $e_{n_0}^1 =$
 1674 $e_{n_0}^2 = e_{n_0}^3$. Then $e_{n_0+1}^1 <_{\text{or}} e_{n_0+1}^2$ or $e_{n_0+1}^2 <_{\text{or}} e_{n_0+1}^3$; so $t_1 \leq^h t_3$.
 1675 • **Antisymmetric** Let t_1, t_2 s.t. $t_1 \leq^h t_2$ and $t_2 \leq^h t_1$. If $t_1 [\text{so} \cup \text{wr}]^* t_2$, then $t_1 = t_2$. If not, by
 1676 the symmetric argument, $\neg(t_2 [\text{so} \cup \text{wr}]^* t_1)$. In that situation, by lemma E.2 we know both
 1677 MINIMALDEPENDENCY(h, t_1, t_2, \perp) and MINIMALDEPENDENCY(h, t_1, t_2, \perp) halt and cannot be
 1678 satisfied at the same time. This contradicts that both $t_1 \leq^h t_2$ and $t_2 \leq^h t_1$ hold; so $t_1 = t_2$.
 1679 □

1681 E.1.2 Oracle-respectful histories.

1682
 1683 The second step in this proof is characterizing all reachable histories with some general invariant that
 1684 can be generalized to every total history. For doing so, we will show that for reachable histories
 1685 any history order coincide with its canonical order; so any property based on a history order can
 1686 be generalized to be based on its canonical order.

1687 *Definition E.5.* An ordered history (h, \leq) is *or-respectful* with respect to \leq if it has at most one
 1688 pending transaction log and for every pair of events $e \in P$, $e' \in h$ s.t. $e \leq_{\text{or}} e'$, either $e \leq e'$ or
 1689 $\exists e'' \in h$, $\text{tr}(h, e'') \leq_{\text{or}} \text{tr}(h, e)$ s.t. $\text{tr}(h, e') [\text{so} \cup \text{wr}]^* \text{tr}(h, e'')$, $e'' \leq e$ and $\text{SWAPPED}(h, e'')$; where
 1690 if $e \notin h$ we state $e' \leq e$ always hold but $e \leq e'$ never does. We will denote it by $R^{\text{or}}(h, \leq)$.
 1691

1692 **LEMMA E.6.** *Let p a computable path. Every ordered history (h, \leq_h) in p is or-respectful with respect*
 1693 *to \leq_h .*

1694
 1695 **PROOF.** We will prove this property by induction on the number of histories this path has. The
 1696 base case, the empty path, trivially holds; so let us prove the inductive case: for every path of at
 1697 most length n the property holds. Let p a path of length $n + 1$ and $h_{<}$ the last reachable history of
 1698 this path. As $p \setminus \{h\}$ is a computable path of length n , the immediate predecessor of h in p , $(h_p, <_{h_p})$
 1699 is or-respectful with respect to $<_p$. Let $a = \text{NEXT}(h_p)$.

1700 Firstly, if a is not a read nor a begin event and $h = h_p \oplus a$, as \leq_h is an extension of \leq_{h_p} , a belongs
 1701 to the only pending transaction and or orders transactions completely, we can deduce that h is
 1702 or-respectful with respect to \leq .

1703 In addition, if a is a begin event and $h = h_p \oplus a$, let $e \in P$, $e' \in h$ s.t. $e <_{\text{or}} e'$. If $e \in h_p$ or $e' \neq a$,
 1704 as \leq_h is an extension of \leq_{h_p} and $R^{\text{or}}(h_p, \leq_{h_p})$ holds, the condition for satisfying $R^{\text{or}}(h, \leq)$ holds
 1705 with e and e' . Moreover, as $a = \min_{\text{or}} P \setminus h_p$, there is no event $e \in P \setminus h_p$ s.t. $e \leq_{\text{or}} a$; so $R^{\text{or}}(h, \leq)$
 1706 holds.

1707 Moreover, if a is a read event and $h = h_p \oplus \text{wr}(t, a)$ for some transaction log t , let us call $e \in$
 1708 P , $e' \in h$ s.t. $e <_{\text{or}} e'$. Once again, if $e \in h$ or $e' \neq a$ the property holds; so let's suppose $e \in P \setminus h_p$ and
 1709 $e' = a$. Let $b = \text{begin}(\text{tr}(h, a))$, that also belongs to h_p . Firstly, as $e \leq_{\text{or}} \text{tr}(h, e') = \text{tr}(h, b)$ we know
 1710 that $e \leq_{\text{or}} b$. Secondly, as $R^{\text{or}}(h_p, \leq_{h_p})$, $e \notin h_p$ and $e \leq_{\text{or}} b$; there exists $c \in h_p$, $\text{tr}(h_p, c) \leq_{\text{or}} \text{tr}(h_p, a)$
 1711 s.t. $(\text{tr}(h_p, b), \text{tr}(h_p, c)) \in (\text{so} \cup \text{wr})^*$, $c \leq b$ and $\text{SWAPPED}((h_p, <_{h_p}), c)$. As $\text{tr}(h, a) = \text{tr}(h, b)$ and
 1712 $\text{SWAPPED}((h_p, <_{h_p}), c)$ implies $\text{SWAPPED}(h_{<}, c)$, we conclude $R^{\text{or}}(h, \leq)$.

1713 But if no previous case is satisfied, it is because $h = \text{SWAP}((h_p, <_{h_p}), r, t)$ for some $r, t \in h_p$ s.t.
 1714 $\text{OPTIMALITY}((h_p, <_{h_p}), r, t)$ holds. Let e, e' two events s.t. $e \leq_{\text{or}} e'$. On one hand, if $e \leq e'$, $R^{\text{or}}(h, e)$
 1715

1716 holds. On the other hand, if $e' < e$ and $e' \leq_{h_p} e$, as $R^{or}(h_p, \leq_{h_p})$ holds and no swapped event is
 1717 deleted by $OPTIMALITY((h_p, <_{h_p}), r, t)$'s definition, the property is also satisfied. Finally, if $e' < e$
 1718 and $e \leq_{h_p} e'$, e has to be a deleted event so $e \in P \setminus h$. As $r \leq_{h_p} e$, if $e \leq_{or} a$, as $e \not\leq a$, there
 1719 would exist a $c \in h_p$, $tr(h_p, c) \leq_{or} tr(h_p, e) \leq_{or} tr(h_p, r)$ s.t. $(tr(h_p, r), tr(h_p, c)) \in (so \cup wr)^*$ and
 1720 $SWAPPED(h_{<}, c)$. However, this is impossible as $tr(h_{<}, r)$ has as maximal event r and the algorithm
 1721 preserves at most one pending transaction; so $e \leq_{or} a$. Taking $e'' = r$ the property is witnessed. \square

1722 PROPOSITION E.7. For any reachable history h , $\leq^h \equiv \leq_h$.
 1723

1724 PROOF. For proving this equivalence, we will show that in any computable path and for any
 1725 ordered history (h, \leq_h) , if $t \leq_h t'$, then $t \leq^h t'$, as by lemma E.4 \leq^h is a total order and therefore
 1726 they have to coincide. We will prove this by induction on the number of histories a path has. The
 1727 base case, the empty path, trivially holds; so let us prove the inductive case: for every path of at
 1728 most length n the property holds. Let p a path of length $n + 1$ and $h_{<_h}$ the last reachable ordered
 1729 history of this path. As $p \setminus \{h\}$ is a computable path of length n , the immediate predecessor of h in
 1730 p , $\leq^{h_p} \equiv \leq_{h_p}$. Let $e = NEXT(h_p)$. Firstly, let's note that if h is an extension of h_p , as $R^{or}(h_p, <_{h_p})$, the
 1731 property can only fail while comparing a transaction t with $tr(h, e)$.

- 1732 • h extends h_p and e is a begin: As $DEP(h_p, t, \perp) = DEP(h, t, \perp)$ for every transaction in h_p , if
 1733 $t \leq^{h_p} t'$, then $t \leq^h t'$. Moreover, $DEP(h, tr(h, e), \perp) = \{e\} = \min_{or} P \setminus h_p$. By lemma E.6 h is
 1734 or -respectful, so for every t , $\min_{or} DEP(h, t, \perp) <_{or} e$; which implies $t <^h tr(h, e)$. By lemma
 1735 E.4, \leq^h is a total order, so it coincides with \leq_h .
- 1736 • h extends h_p and e is not a begin: As no transaction depends on $tr(h, e)$ and $tr(h, e) =$
 1737 $\text{last}(h_p)$, if we prove that for every pair of transactions $\text{MINIMALDEPENDENCY}(h_p, t', t'', \perp)$
 1738 $= \text{MINIMALDEPENDENCY}(h, t', t'', \perp)$, the lemma would hold. On one hand,
 1739 $DEP(h, tr(h, e), \perp) = DEP(h_p, tr(h, e), \perp) = tr(h, e)$ and in the other hand, by lemma E.6,
 1740 $\min_{or} DEP(h_p, t, \perp) <_{or} tr(h, e)$. Finally, as $e \notin DEP(h, \hat{t}, e')$, for every $\hat{t} \neq tr(h, e)$, $e' \neq \perp$, for
 1741 every pair of transactions t', t'' , $\text{MINIMALDEPENDENCY}(h_p, t', t'' \perp) =$
 1742 $\text{MINIMALDEPENDENCY}(h, t', t'', \perp)$.
- 1743 • $h = \text{SWAP}(h_p, r, t)$, where $t = tr(h, e)$: As $OPTIMALITY(h_p, r, t)$ is satisfied and h is
 1744 or -respectful, for every event e' and transaction t' in h , $\min_{or} DEP(h_p, t', e') =$
 1745 $\min_{or} DEP(h, t', e')$, so for every pair of transactions $\text{MINIMALDEPENDENCY}(h_p, t', t'', \perp) =$
 1746 $\text{MINIMALDEPENDENCY}(h, t', t'', \perp)$. In particular, this implies $t' \leq^{h_p} t''$ if and only if $t' \leq^h t''$
 1747 for every pair $t', t'' \in h$. Finally, as for every $t' \in h$, $t' \leq^h tr(h, r)$ (because $tr(h, r)$ is
 1748 $(so \cup wr)^+$ -maximal); we conclude that $\leq^h \equiv \leq_h$.
 1749 \square

1750

1751 Proposition E.7 is a very interesting result as it express the following fact: regardless of the
 1752 computable path that leads to a history, the final order between events will be the same. Therefore,
 1753 all possible history orders collapse to one, the canonical one. This result will have a key role during
 1754 both completeness and optimality, as it restricts the possible histories that precede another while
 1755 describing the computable path leading to it. In addition, proposition E.7 together with lemma
 1756 E.6 justify enlarging definition E.5 with a general order as for reachable histories, $R^{or}(h, \leq_h)$ is
 1757 equivalent to $R^{or}(h, \leq^h)$. From what follows, we will simply state h is or -respectful and we will
 1758 denote it by $R^{or}(h)$. Moreover, we will assume every history is ordered with the canonical order.
 1759

1760 COROLLARY E.8. Let h_p a reachable history and let h a immediate successor of h_p whose last event
 1761 r is a read. Then $h_{<} = \text{SWAP}((h_p, <_{h_p}), r, t)$ if and only if $\text{SWAPPED}(h, r)$ does.
 1762

1763 PROOF. \Rightarrow
 1764

1765 Let's suppose that $h_{<} = \text{SWAP}((h_p, <_{h_p}), r, t)$ for some t transaction. As the last event in h
 1766 is r and by definition of SWAP function no event reads from $\text{wr}^{-1}(r)$ in h besides r , to prove
 1767 $\text{SWAPPED}(h, r)$ holds we just need to show that $r <_{\text{or}} t$. By lemma E.6, $\text{R}^{\text{or}}(h_p)$ holds. As $r <_{h_p} t$,
 1768 $\text{OPTIMALITY}((h_p, <_{h_p}), r, t)$ holds and t is $(\text{so} \cup \text{wr})^+$ -maximal, we conclude that $r <_{\text{or}} t$.
 1769 \Leftarrow Let's suppose that $h = h_p \oplus r \oplus \text{wr}(r, t)$ for some transaction t . Let's suppose that $r <_{\text{or}} t$.
 1770 As $\text{R}^{\text{or}}(h_p)$, there exists some event e'' s.t. $\text{tr}(h_p, e'') \leq \text{tr}(h, r)$, $t [\text{so} \cup \text{wr}]^* \text{tr}(h, e'')$ and $e'' \leq r$ so
 1771 $\neg(\text{SWAPPED}(h, r))$.
 1772 □

1772

1773

1774

LEMMA E.9. Any total history is *or-respectful*.

1775

1776

1777

1778

1779

1780

1781

1782

PROOF. Let h be a total history and t, t' a pair of transactions s.t. $t \leq_{\text{or}} t'$. If $t \leq^h t'$, then the
 statement is satisfied; so let's assume the contrary: $t' \leq^h t$. If $(t', t) \in (\text{so} \cup \text{wr})^*$, then for every
 $e \in t, e' \in t' \exists c \in h$ s.t. $\text{tr}(h, c) \leq_{\text{or}} \text{tr}(h, e)$, $(\text{tr}(h, e'), \text{tr}(h, c)) \in (\text{so} \cup \text{wr})^*$, $\text{SWAPPED}(h, c)$ and
 $c \leq^h e$; so the property is satisfied. Otherwise, by definition of MINIMALDEPENDENCY , there exists
 $r' \in h$ s.t. $(t', \text{tr}(h, r')) \in (\text{so} \cup \text{wr})^*$ and $\text{tr}(h, r') \leq_{\text{or}} t$. Moreover, by CANONICALORDER 's definition,
 $\text{tr}(h, r) \leq^h t$. Finally $\text{SWAPPED}(h, r')$ holds as it is the minimum element according *or*. To sum up,
 $\text{R}^{\text{or}}(h)$ holds. □

1783

1784

E.1.3 Previous of a history.

1785

1786

1787

1788

As a third and final step in our proof, we define the function *previous* that, for a every history
 h , if $\text{PREV}(h)$ is reachable, then h is also reachable. Moreover, $\text{PREV}(h)$ will belong to the same
 computable path.

1789

1790

Algorithm 4 PREV

1791

1792

1793

1794

1795

1796

1797

1798

1799

1800

1801

1802

1803

1804

1805

1806

1807

1808

1809

1810

```

1: procedure PREV(h)
2:   if h = ∅ then
3:     return ∅
4:   a ← last(h)
5:   if ¬SWAPPED(h, a) then
6:     return h \ a
7:   else
8:     let t s.t. (t, r) ∈ wr.
9:     return MAXCOMPLETION(h \ a, {e | e ∉ (h \ a) ∧ e <_{or} t})
10: procedure MAXCOMPLETION(h, D)
11:   if D ≠ ∅ then
12:     e ← min_{<_{or}} D
13:     if type(e) ≠ read then
14:       return MAXCOMPLETION(h ⊕ e, D \ {e})
15:     else
16:       let t s.t. readLatest_I(h ⊕ e ⊕ wr(t, e), e) holds
17:       return MAXCOMPLETION(h ⊕ e ⊕ wr(t, e), D \ {e})
18:   else
19:     return h

```

1811

1812

1813

First, we show that the invariant of our algorithm is preserved via PREV .

LEMMA E.10. For every *or*-respectful history h , $\text{PREV}(h)$ is also *or*-respectful.

PROOF. Let suppose $h \neq \emptyset$, $h_p = \text{PREV}(h)$, $a = \text{last}(h)$, $e \in P$ and $e' \in h_p$ s.t. $e \leq_{\text{or}} e'$. We explore different cases depending if e, e' belong to h or not. If $e' \in h_p \setminus h$, $\neg(\text{SWAPPED}(h_p, e))$ and $\neg(\text{SWAPPED}(h_p, e'))$ holds. As $\min_{<_{\text{or}}} \text{DEP}(h, \text{tr}(h, e'), \perp) = \text{begin}(\text{tr}(h, e'))$, we obtain that $\min_{<_{\text{or}}} \text{DEP}(h, \text{tr}(h, e')) \leq_{\text{or}} e' \leq_{\text{or}} \text{begin}(\text{tr}(h, e'))$. Therefore, as $e' \in h_p \in h$, $\neg(\text{tr}(h, e') [\text{so} \cup \text{wr}]^+ \text{tr}(h, e))$, so $e \leq^h e'$. And if $e' \in h$, either $e \leq^h e'$ or $e' \leq^h e$. In the former case, both are in h and therefore, in h_p . As it cannot happen that $e' \in \text{tr}(h, a)$ and $e \leq^{h_p} a$ because $\text{SWAPPED}(h, a)$ and $e \leq_{\text{or}} e'$, we conclude that $e \leq^h e'$ (\leq_{h_p} keeps the relative orders between transactions different from $\text{tr}(h, a)$ and by lemma E.6 they coincide). In the latter case, by $R^{\text{or}}(h)$, there exists e'' that witness it. In particular, $\text{SWAPPED}(h, e'')$ holds, so $e'' \in h_p$. e'' witness $R^{\text{or}}(h_p)$ holds. In the three cases we deduce that $R^{\text{or}}(h_p)$.

□

Next, we have to prove that previous is a sound function, i.e. the composition between EXPLORE-CE and PREV give us the identity. For doing so, in the case a history is a swap, we deduce that both histories should contain the same elements and they read the same; so they have to coincide.

LEMMA E.11. For every consistent history *or*-respectful h , if $\text{PREV}(h)$ is reachable, then h is also reachable.

PROOF. Let suppose $h \neq \emptyset$, $h_p = \text{PREV}(h)$ and $a = \text{last}(h)$. If $\neg \text{SWAPPED}(h, a)$, let $h_n = h_p \oplus a$ if a is not a read and $h_n = h_p \oplus a \oplus \text{wr}(t, a)$, where t is the transaction s.t. $(t, r) \in \text{wr}$, otherwise. Either way, h_n is always reachable and it coincides with h . On the contrary, if $\text{SWAPPED}(h, a)$, a is a read event and it swapped; so let us call t to the transaction s.t. $(t, a) \in \text{wr}$. Firstly, as $\text{SWAPPED}(h, a)$, $a <_{\text{or}} t$, and by lemma E.6, $R^{\text{or}}(h_p)$ holds, so $a <_{h_p} t$ does; which let us conclude $\text{COMPUTERREORDERINGS}(h_p)$ will always return (a, t) as a possible swap pair. In addition, all transactions in h_p are non-pending and $(t, a) \in \text{wr}$, so in particular $\text{last}(h_p)$ is an commit event. If we call $h_s = \text{SWAP}(h_p, a, t)$, and we prove that $h_p \setminus h = h_p \setminus h_s$ holds, then we would deduce $h = h_s$ as $\text{wr}(t, a)$ in both h_p, h_s and $h \subseteq h_p, h_s \subseteq h_p$; which would allow us to conclude h is reachable from h_p .

On one hand, if $e \in h_p \setminus h$, we deduce that $e \notin h$ and $e <_{\text{or}} t$. In particular, $\neg(\text{tr}(h, e) [\text{so} \cup \text{wr}]^* t)$. Moreover, if $e \leq_{\text{or}} a$, by $R^{\text{or}}(h)$, either $e \leq^h a$ or $\exists e'' \in h, e'' \leq_{\text{or}} e$ s.t. $t(a) [\text{so} \cup \text{wr}]^* \text{tr}(h, e'')$, $e'' \leq^h e$ and $\text{SWAPPED}(h, e'')$; both impossible situations as $e \notin h$ and $a = \text{last}(h)$; so $a \leq_{\text{or}} e$. In other words, $e \in h_p \setminus h_s$.

On the other hand, $e \in h_p \setminus h_s$ if and only if $\neg(\text{tr}(h, e) [\text{so} \cup \text{wr}]^* t(w))$ and $a <_{\text{or}} e <_{\text{or}} w$. If e would belong to h then $e \leq^h a$. As h is *or*-respectful and $a \leq_{\text{or}} e$, we deduce there exists a $e'' \in h$ s.t. $\text{tr}(h, e'') \leq_{\text{or}} t(a)$, $\text{tr}(h, e) [\text{so} \cup \text{wr}]^* \text{tr}(h, e'')$ and $\text{SWAPPED}(h, e'')$. Moreover, as $e'' \in h$, $e'' \in h_p$. By corollary E.8 $\text{SWAPPED}(h_p, e'')$ and $\text{OPTIMALITY}(h_p, a, t)$ hold, $e'' \in h_s$ and so e does. This result leads to a contradiction, so $e \notin h$; i.e. $e \in h_p \setminus h$.

□

COROLLARY E.12. In a consistent *or*-respectful history h whose previous history is reachable, if $a = \text{last}(h)$, $\text{SWAPPED}(h, a)$ and t is a transaction such that $(t, a) \in \text{wr}$, h coincides with $\text{SWAP}(\text{PREV}(h), a, t)$.

PROOF. It comes straight away from the proof of lemma E.11.

□

Once proven that PREV is sound, let us prove that for every history we can compose PREV a finite number of times obtaining the empty history. We are going to prove it by induction on the number of swapped events, so we prove first the recursive composition finishes in finite time and then we conclude our claim.

LEMMA E.13. For every non-empty consistent *or*-respectful history h , $h_p = \text{PREV}(h)$ and $a = \text{last}(h)$, if $\text{SWAPPED}(h, a)$ then $\{e \in h_p \mid \text{SWAPPED}(h_p, e)\} = \{e \in h \mid \text{SWAPPED}(h, e)\} \setminus \{a\}$, otherwise $h_p = h \setminus a$.

PROOF. Let $a = \text{last}(h)$ and $h' = h \setminus a$. If $\neg(\text{SWAPPED}(h, a))$, then $h_p = h'$ and the lemma holds trivially. Otherwise, as $h_p = \text{MAXCOMPLETION}(h')$, we will show that every event not belonging to $h_p \setminus h'$ is not swapped by induction on every recursive call to MAXCOMPLETION . Let us call $D = \{e \mid e \notin h' \wedge e <_{\text{or}}\}$. This set, intuitively, contain all the events that would have been deleted from a reachable history h to produce h_p . In this setting, let us call $h_{|D|} = h'$, $D_{|D|} = D$ and $D_k = D_{k+1} \setminus \{\min_{<_{\text{or}}} D_{k+1}\}$, $e_k = \min_{<_{\text{or}}} D_k$ for every $k, 0 \leq k < |D|$ (i.e. $D_k = D_{k+1} \setminus \{e_{k+1}\}$). We will prove the lemma by induction on $n = |D| - k$, constructing a collection of *or*-respectful histories $h_k, 0 \leq k < |D|$, such that each one is an extension of its predecessor with a non-swapped event.

The base case, $h_{|D|}$ is trivial as by its definition it corresponds with h' . Let's prove the inductive case: $\{e \mid \text{SWAPPED}(h_{k+1}, e)\} = \{e \mid \text{SWAPPED}(h', e)\}$. If e_{k+1} is not a read event, $h_k = h_{k+1} \oplus e_{k+1}$, $\text{R}^{\text{or}}(h_k)$ and $\{e \mid \text{SWAPPED}(h_k, e)\} = \{e \mid \text{SWAPPED}(h', e)\}$; as only read events can be swapped. Otherwise, e_{k+1} is a read event. By the isolation level's causal-extensibility there exists a transaction f_{k+1} that writes the same variable as e_{k+1} , $(f_{k+1}, \text{tr}(h, e_{k+1})) \in (\text{so} \cup \text{wr})^*$ and $h_{k+1} \oplus e_{k+1} \oplus \text{wr}(f_{k+1}, e_{k+1})$ is consistent. Moreover, if e_{k+1} reads from any causal dependent element f' , f' in h_{k+1} , it cannot be swapped: as $\text{R}^{\text{or}}(h_{k+1})$ holds, if $e_{k+1} <_{\text{or}} f'$ there must be an event c_{k+1} s.t. $\text{tr}(h, c_{k+1}) \leq_{\text{or}} \text{tr}(h, e_{k+1})$ and $(f', \text{tr}(h, c_{k+1})) \in (\text{so} \cup \text{wr})^*$. Hence, $\{e \mid \text{SWAPPED}(h_{k+1}, e)\} = \{e \mid \text{SWAPPED}(h_{k+1} \oplus e_{k+1} \oplus \text{wr}(f', e_{k+1}), e)\}$.

Let $E_{k+1} = \{t \mid h_{k+1} \oplus e_{k+1} \oplus \text{wr}(t, e_{k+1}) \models I \wedge \{e \mid \text{SWAPPED}(h_{k+1}, e)\} = s\{e \mid \text{SWAPPED}(h_{k+1} \oplus e_{k+1} \oplus \text{wr}(t, e_{k+1}), e)\}\}$ and let $t_{k+1} = \max_{\leq h_{k+1}} \{t \in E_{k+1} \mid (t, \text{tr}(h_{k+1}, e_{k+1})) \in (\text{so} \cup \text{wr})^*\}$. This element is well defined as f_{k+1} belongs to E_{k+1} . Therefore, $h_k = h_{k+1} \oplus e_{k+1} \oplus \text{wr}(t_{k+1}, e_{k+1})$ is consistent and $\{e \mid \text{SWAPPED}(h_k, e)\} = \{e \mid \text{SWAPPED}(h', e)\}$. Moreover, let's remark that as t_{k+1} is the maximum transaction according to $\leq h_{k+1}$ s.t. is consistent and $\{e \mid \text{SWAPPED}(h_k, e)\} = \{e \mid \text{SWAPPED}(h', e)\}$. In addition, by construction, it also satisfies $\text{readLatest}_I(h_k, e_{k+1}, w_{k+1})$. Finally, h_k is also *or*-respectful as e_{k+1} is not swapped and $\text{R}^{\text{or}}(h_{k+1})$ holds.

Thus, after applying induction, we obtain $h_p = h_0$; which let us conclude $\{e \in h_p \mid \text{SWAPPED}(h_p, e)\} = \{e \in h' \mid \text{SWAPPED}(h', e)\} = \{e \in h \mid \text{SWAPPED}(h, e)\} \setminus \{a\}$. \square

LEMMA E.14. For every consistent *or*-respectful history h there exists some $k_h \in \mathbb{N}$ such that $\text{PREV}^{k_h}(h) = \emptyset$.

PROOF. This lemma is immediate consequence of lemma E.13. Let us call $\xi(h) = |\{e \in h \mid \text{SWAPPED}(h, e)\}|$, the number of swapped events in h , and let us prove the lemma by induction on $(\xi(h), |h|)$. The base case, $\xi(h) = |h| = 0$ is trivial as h would be \emptyset ; so let's assume that for every history h such that $\xi(h) < n$ or $\xi(h) = h \wedge |h| < m$ there exists such k_h . Let h then a history s.t. $\xi(h) = n$ and $|h| = m$. $h_p = \text{PREV}(h)$. On one hand, if $h_p = h \setminus a$ then $\xi(h_p) = \xi(h)$ and $|h_p| = |h| - 1$. On the other hand, if $h_p \neq h \setminus a$, $\xi(h_p) = \xi(h) - 1$. In any case, by induction hypothesis on h_p , there exists an integer k_{h_p} such that $\text{PREV}^{k_{h_p}}(h_p) = \emptyset$. Therefore, $k_h = k_{h_p} + 1$ satisfies $\text{PREV}^{k_h}(h) = \emptyset$. \square

PROPOSITION E.15. For every consistent *or*-respectful history h exists $k \in \mathbb{N}$ and some sequence of *or*-respectful histories $\{h_n\}_{n=0}^k$, $h_0 = \emptyset$ and $h_k = h$ such that the algorithm will compute.

PROOF. Let h a history, k the minimum integer such that $\text{PREV}^k(h) = \emptyset$, which exists thanks to lemma E.14 and $C = \{\text{PREV}^{k-n}(h)\}_{n=0}^k$ a set of indexed histories. By the collection's definition and lemma E.10, $h_0 = \text{PREV}^k(h) = \emptyset$, $h_k = \text{PREV}^0(h) = h$ and $\text{R}^{\text{or}}(h_n)$ for every $n \in \mathbb{N}$; so let us prove by induction on n that every history in C is reachable. The base case, h_0 , is trivially achieved; as it is

1912 always reachable. In addition, by lemma E.11, we know that if h_n is reachable, h_{n+1} is it too; which
 1913 proves the inductive step. \square

1914 THEOREM E.16. *The algorithm EXPLORE-CE is complete.*
 1915

1916 PROOF. By lemma E.9, any consistent total history is or-respectful. As a consequence of propo-
 1917 sition E.15, there exist a sequence of reachable histories which h belongs to; so in particular, h is
 1918 reachable. \square

1919 E.2 Optimality

1921 For proving optimality we are going to exploit two properties already studied for completeness:
 1922 or-respectfulness and the canonical order. Then, as algorithm EXPLORE-CE is sound and complete,
 1923 we will prove that any computable path leading to a consistent history is the one computed in the
 1924 completeness' proof.

1925 THEOREM E.17. *Algorithm EXPLORE-CE is strongly optimal.*
 1926

1927 PROOF. As the model is causal-extensible, any algorithm optimal is also strongly optimal. Let
 1928 us prove that for every reachable history there is only a computable path that leads to it from \emptyset .
 1929 Let's suppose there exists a history h that is reached p_1, p_2 by two computable paths. By lemma
 1930 E.7, we know that $\leq_h \equiv \leq^h$. However, \leq^h is an order that does not depend on the computable path
 1931 that leads to h ; so neither does \leq_h . Therefore, we can assume without loss of generality that h is a
 1932 history with minimal value of $\xi(h) = |\{e \in h \mid \text{SWAPPED}(h, e)\}|$ and in case of tie, that is minimal
 1933 with respect $|h|$; values independent of the computable path that leads to h .

1934 We can also assume without loss of generality that the predecessor of h in p_1 is $h_1 = \text{PREV}h$,
 1935 and h_2 is the predecessor of h in p_2 . If we prove h_1 and h_2 are identical, p_1 and p_2 have to also be
 1936 identical and therefore, the algorithm would be optimal. Firstly, if $\text{last}(h)$ is not a swapped read
 1937 event, by the definition of NEXT function $h_2 = h \setminus \text{last}(h) = h_1$. On the contrary, let's suppose
 1938 $r = \text{last}(h)$ is a swapped event that reads from a transaction t . Because $\text{SWAPPED}(h, r)$ holds, from
 1939 h_2 to h it has to have happened a swap between r and w . But by corollary E.12, $h = \text{SWAP}(h_1, r, w)$,
 1940 so $h_1 \upharpoonright_{h \setminus r} = h_2 \upharpoonright_{h \setminus r}$. As h_1, h_2 are both or-respectful, $e \in h_1 \setminus h \iff e \in h_2 \setminus h$. Finally, as
 1941 $\text{OPTIMALITY}(h_i, r, w)$ holds for $i \in \{1, 2\}$, for every read event e in $h_1 \cap h_2$ there exists a transaction
 1942 t_e s.t. $\text{wr}(e, t_e)$ for both histories. \square

1943
 1944
 1945
 1946
 1947
 1948
 1949
 1950
 1951
 1952
 1953
 1954
 1955
 1956
 1957
 1958
 1959
 1960

F EXPERIMENTAL DATA

F.1 Application Scalability

	CC				CC + SI				CC + SER			
	Histories	End states	Mem.	Time	Histories	End states	Mem.	Time	Histories	End states	Mem.	Time
courseware-1	48	48	256MB	00:00:03	27	48	256MB	00:00:03	9	48	256MB	00:00:03
courseware-2	12	12	256MB	00:00:05	11	12	256MB	00:00:05	6	12	256MB	00:00:05
courseware-3	370	370	310MB	00:00:07	54	370	314MB	00:00:09	3	370	308MB	00:00:08
courseware-4	18	18	256MB	00:00:02	5	18	256MB	00:00:02	1	18	256MB	00:00:02
courseware-5	34	34	256MB	00:00:03	1	34	256MB	00:00:04	1	34	256MB	00:00:06
shoppingcart-1	32	32	256MB	00:00:03	3	32	310MB	00:00:03	1	32	256MB	00:00:03
shoppingcart-2	174	174	256MB	00:00:05	20	174	256MB	00:00:06	1	174	256MB	00:00:06
shoppingcart-3	77	77	256MB	00:00:04	34	77	256MB	00:00:04	14	77	256MB	00:00:04
shoppingcart-4	445	445	256MB	00:00:06	226	445	370MB	00:00:08	1	445	370MB	00:00:07
shoppingcart-5	170	170	308MB	00:00:05	68	170	256MB	00:00:05	10	170	256MB	00:00:06
tpcc-1	22	22	407MB	00:00:08	1	22	450MB	00:00:08	1	22	392MB	00:00:07
tpcc-2	88	88	375MB	00:00:15	2	88	411MB	00:00:16	2	88	378MB	00:00:16
tpcc-3	63	63	447MB	00:00:09	5	63	458MB	00:00:10	1	63	444MB	00:00:10
tpcc-4	105	105	420MB	00:00:11	22	105	478MB	00:00:12	1	105	444MB	00:00:12
tpcc-5	1108	1108	444MB	00:01:28	57	1108	450MB	00:01:35	5	1108	566MB	00:01:26
twitter-1	16	16	309MB	00:00:03	10	16	256MB	00:00:03	3	16	308MB	00:00:03
twitter-2	59	59	308MB	00:00:03	5	59	256MB	00:00:04	1	59	256MB	00:00:04
twitter-3	114	114	256MB	00:00:05	36	114	308MB	00:00:06	6	114	256MB	00:00:05
twitter-4	1995	1995	308MB	00:01:17	84	1995	370MB	00:01:27	84	1995	308MB	00:01:17
twitter-5	1995	1995	308MB	00:01:09	84	1995	308MB	00:01:16	84	1995	311MB	00:01:09
wikipedia-1	32	32	308MB	00:00:06	2	32	326MB	00:00:05	2	32	318MB	00:00:05
wikipedia-2	138	138	372MB	00:00:12	3	138	370MB	00:00:11	1	138	308MB	00:00:10
wikipedia-3	156	156	370MB	00:00:09	1	156	444MB	00:00:09	1	156	377MB	00:00:08
wikipedia-4	53	53	323MB	00:00:09	12	53	323MB	00:00:09	1	53	327MB	00:00:09
wikipedia-5	3208	3208	308MB	00:00:52	2	3208	370MB	00:01:00	1	3208	322MB	00:00:51

	RA + CC				RC + CC				true + CC				DFS(CC)		
	Histories	End states	Mem.	Time	Histories	End states	Mem.	Time	Histories	End states	Mem.	Time	End states	Mem.	Time
courseware-1	48	164	256MB	00:00:04	48	3456	312MB	00:00:17	48	9216	310MB	00:00:31	73482	447MB	00:11:30
courseware-2	12	20	256MB	00:00:05	12	96	256MB	00:00:06	12	96	256MB	00:00:06	29304	469MB	00:04:33
courseware-3	370	1841	308MB	00:00:19	20	719429	308MB	TL	20	786434	308MB	TL	61012	308MB	TL
courseware-4	18	32	256MB	00:00:02	18	1984	308MB	00:00:11	18	1984	312MB	00:00:11	93896	308MB	00:11:48
courseware-5	34	120	308MB	00:00:06	34	99048	308MB	00:05:34	34	138480	308MB	00:06:45	46063	523MB	TL
shoppingcart-1	32	80	256MB	00:00:04	32	6912	308MB	00:00:54	32	9216	370MB	00:01:08	126678	444MB	TL
shoppingcart-2	174	1017	308MB	00:00:13	174	78336	316MB	00:05:41	174	221184	370MB	00:12:34	166311	308MB	TL
shoppingcart-3	77	231	256MB	00:00:06	77	4940	313MB	00:00:44	77	8960	444MB	00:01:10	164385	444MB	TL
shoppingcart-4	445	477	256MB	00:00:08	445	734464	370MB	TL	445	858867	444MB	TL	262924	444MB	TL
shoppingcart-5	170	450	308MB	00:00:08	170	15504	308MB	00:00:55	170	117936	308MB	00:04:54	122523	379MB	TL
tpcc-1	22	80	533MB	00:00:12	4	78164	568MB	TL	1	63588	380MB	TL	17908	1409MB	TL
tpcc-2	88	564	533MB	00:00:57	1	77865	716MB	TL	1	131450	533MB	TL	21885	1230MB	TL
tpcc-3	63	216	533MB	00:00:18	5	36618	669MB	TL	5	38861	568MB	TL	20466	1194MB	TL
tpcc-4	105	114	449MB	00:00:12	17	124679	572MB	TL	9	116126	640MB	TL	20190	1174MB	TL
tpcc-5	1109	19463	533MB	00:21:05	1	83644	464MB	TL	1	84325	444MB	TL	25389	1349MB	TL
twitter-1	16	20	256MB	00:00:03	16	2208	308MB	00:00:34	16	4608	308MB	00:00:56	35056	539MB	00:28:45
twitter-2	59	147	256MB	00:00:05	59	1728	308MB	00:00:18	59	1728	321MB	00:00:18	159100	447MB	TL
twitter-3	114	216	308MB	00:00:07	114	1296	308MB	00:00:19	114	1296	374MB	00:00:18	108792	444MB	00:22:47
twitter-4	195	6860	308MB	00:03:37	10	99558	374MB	TL	1	163231	322MB	TL	55198	444MB	TL
twitter-5	195	6860	308MB	00:03:18	84	61498	444MB	TL	84	118514	322MB	TL	55198	444MB	TL
wikipedia-1	32	48	256MB	00:00:05	32	16480	444MB	00:03:12	32	49280	308MB	00:08:13	54172	370MB	TL
wikipedia-2	138	352	371MB	00:00:13	1	125438	540MB	TL	1	122187	489MB	TL	8169	561MB	TL
wikipedia-3	156	380	370MB	00:00:14	156	115200	544MB	00:20:56	156	161280	444MB	00:28:28	69935	568MB	TL
wikipedia-4	53	104	372MB	00:00:11	1	63360	465MB	TL	1	63023	4652B	TL	25044	768MB	TL
wikipedia-5	3208	3807	311MB	00:01:00	32	16480	308MB	00:03:22	15	30862	444MB	TL	1226	563MB	TL

F.2 Session Scalability

	One session			Two sessions			Three sessions			Four sessions			Five sessions		
	Histories	Mem.	Time	Histories	Mem.	Time	Histories	Mem.	Time	Histories	Mem.	Time	Histories	Mem.	Time
tpcc-1	1	256MB	00:00:02	6	256MB	00:00:04	72	447MB	00:00:13	4662	783MB	00:05:41	41371	1386MB	TL
tpcc-2	1	256MB	00:00:02	30	313MB	00:00:06	2071	640MB	00:02:00	28563	1618MB	TL	18122	1103MB	TL
tpcc-3	1	256MB	00:00:02	4	326MB	00:00:04	100	562MB	00:00:18	24373	1152MB	TL	21035	1038MB	TL
tpcc-4	1	256MB	00:00:02	4	342MB	00:00:05	527	582MB	00:00:44	21118	1352MB	TL	25386	1233MB	TL
tpcc-5	1	256MB	00:00:03	5	380MB	00:00:06	335	453MB	00:00:40	19184	1384MB	TL	26262	1598MB	TL
wikipedia-1	1	256MB	00:00:02	27	256MB	00:00:04	5488	371MB	00:04:39	130848	370MB	TL	45018	400MB	TL
wikipedia-2	1	256MB	00:00:02	6	323MB	00:00:03	216	329MB	00:00:08	2984	394MB	00:12:57	31124	1164MB	TL
wikipedia-3	1	256MB	00:00:02	20	351MB	00:00:06	1369	594MB	00:00:32	43146	566MB	TL	33906	692MB	TL
wikipedia-4	1	256MB	00:00:02	4	346MB	00:00:04	43	390MB	00:00:08	4516	640MB	00:01:05	10059	1833MB	TL
wikipedia-5	1	256MB	00:00:02	9	256MB	00:00:04	67	319MB	00:00:07	2438	582MB	00:01:26	40760	689MB	TL

2010 F.3 Transaction Scalability

2011

2012

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

2023

2024

2025

2026

2027

2028

2029

2030

2031

2032

2033

2034

2035

2036

2037

2038

2039

2040

2041

2042

2043

2044

2045

2046

2047

2048

2049

2050

2051

2052

2053

2054

2055

2056

2057

2058

	One transaction			Two transactions			Three transactions			Four transactions			Five transactions		
	Histories	Mem.	Time	Histories	Mem.	Time	Histories	Mem.	Time	Histories	Mem.	Time	Histories	Mem.	Time
tpcc-1	2	313MB	00:00:03	5	311MB	00:00:04	6	341MB	00:00:05	241	393MB	00:00:23	9491	829MB	00:08:34
tpcc-2	6	256MB	00:00:04	223	456MB	00:00:24	1296	562MB	00:02:03	4867	810MB	00:09:18	11085	1195MB	TL
tpcc-3	2	256MB	00:00:03	26	324MB	00:00:07	502	501MB	00:00:40	10352	1482MB	00:07:52	20949	1304MB	TL
tpcc-4	7	256MB	00:00:04	216	461MB	00:00:19	541	537MB	00:00:41	13012	695MB	00:15:58	20609	1157MB	TL
tpcc-5	1	256MB	00:00:03	2	256MB	00:00:05	3	336MB	00:00:05	8	376MB	00:00:06	19	381MB	00:00:08
wikipedia-1	6	256MB	00:00:02	212	308MB	00:00:07	8430	312MB	00:16:45	117794	378MB	TL	107924	444MB	TL
wikipedia-2	4	256MB	00:00:03	8	308MB	00:00:04	36	256MB	00:00:06	851	460MB	00:00:47	8725	256MB	00:17:10
wikipedia-3	6	256MB	00:00:03	216	311MB	00:00:12	7244	536MB	00:05:34	25589	567MB	TL	21133	645MB	TL
wikipedia-4	1	256MB	00:00:02	9	320MB	00:00:04	315	623MB	00:00:43	1031	338MB	00:01:52	2534	839MB	00:05:01
wikipedia-5	2	256MB	00:00:03	6	256MB	00:00:05	12	256MB	00:16:45	57	376MB	00:00:08	228	533MB	00:00:23