



A transactional object calculus

Suresh Jagannathan*, Jan Vitek, Adam Welc, Antony Hosking

Department of Computer Science, Purdue University, West Lafayette, IN 47906, United States

Received 25 May 2004; received in revised form 4 February 2005; accepted 6 March 2005

Available online 28 April 2005

Abstract

A transaction defines a locus of computation that satisfies important concurrency and failure properties. These so-called ACID properties provide strong serialization guarantees that allow us to reason about concurrent and distributed programs in terms of higher-level units of computation (e.g., transactions) rather than lower-level data structures (e.g., mutual-exclusion locks). This paper presents a framework for specifying the semantics of a transactional facility integrated within a host programming language. The TFJ calculus, an object calculus derived from Featherweight Java, supports nested and multi-threaded transactions. We give a semantics to TFJ that is parametrized by the definition of the transactional mechanism that permits the study of different transaction models. We give two instantiations: one that defines transactions in terms of a versioning-based optimistic concurrency model, and the other which specifies transactions in terms of a pessimistic two-phase locking protocol, and present soundness and serializability properties for our semantics.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Transactions; Threads; Object calculus; Nesting; Optimistic concurrency; Two-phase locking

1. Introduction

The integration of transactional facilities into programming languages has been driven by applications ranging from middleware infrastructure for enterprise applications [1] to runtime support for optimistic concurrency. The concept of transactions is well-known in

* Corresponding author.

E-mail addresses: suresh@cs.purdue.edu (S. Jagannathan), jv@cs.purdue.edu (J. Vitek), welc@cs.purdue.edu (A. Welc), hosking@cs.purdue.edu (A. Hosking).

database systems; the main challenge in associating transactions with programming control structures comes from the mismatch between the concurrency model of the programming language and the concurrency model of the transactional facility. Issues related to the interaction commit mechanisms found in transactional models with locking, visibility, and update semantics in programming languages can be subtle and complex. Existing technologies enforce little or no relationship between these models, so that programmers can neither rely on transactions for complete isolation among concurrent threads, nor use concurrent threading to more conveniently program transaction logic. To make matters worse, different systems exhibit subtly (or vastly) different observable behavior with respect to failure and isolation properties. While the differences are often mandated by pragmatic concerns, they make it hard to reason about and compare different systems. Furthermore correctness of an implementation often remains an article of faith.

As a first step towards addressing some of these concerns, we propose a semantic framework in which different transactional mechanisms can be studied and compared formally. Requirements for such a framework are that it be sufficiently expressive to allow the specification of core transactional features, that it provide a way to validate the correctness of the semantics, and that it support features found in realistic programs. We are interested in the impact of design choices on observable behavior (e.g., aborts, deadlocks, livelocks) and on implementation performance (e.g., space and time overhead). Our long-term goal is to leverage this framework to aid in the definition of program analyses and optimization techniques for transactional languages.

This paper introduces TFJ, or *Transactional Featherweight Java*, an object calculus with syntactic support for transactions. The operational semantics of TFJ is given in terms of a stratified set of rewrite rules parametrized over the meaning of the transactional constructs. This allows us to define a variety of transactional semantics within the same core language. In this paper, we study *nested and multi-threaded* transactions with two different concurrency control models (two-phase locking and versioning). We have tried to retain the flavor and economy of Featherweight Java [17], while adding features such as concurrency and state needed to model transactions. The primary contribution of this paper is a formal characterization and proof of correctness of different transactional models when incorporated into the core language. While there has been significant previous work on devising formal notation and specifications [7,20] describing transactional properties, we are unaware of other efforts that use operational semantics to study the interplay of concurrency and serializability among different transactional models. Instead, we explore the structure of program traces induced by the semantics, and prove a simpler permutation lemma on these traces that enforce a localized notion of serializability derived from control and dependencies among transaction operations. The proof of serializability over arbitrary traces follows naturally from the composition of permutable sub-traces.

Correctness of the transactional semantics is given in terms of the traditional ACID properties [15] described below. To prove correctness, we show that all local operations performed by concurrent transactions can be *serialized* (i.e., from the perspective of some transaction T all data accesses performed by T are performed serially with respect to other transactions). The use of a stratified semantics allows us to avoid proving this serialization theorem for different concurrency control models. Instead, we explore the structure of program traces induced by the semantics, and prove a simpler permutation lemma on

these traces that enforce a localized notion of serializability derived from control and dependencies among transaction operations.

This article is an extended version of a paper presented in the 2004 European Symposium on Programming entitled, *A Semantic Framework for Designer Transactions* [24].

Overview. Section 2 gives an informal introduction to atomic transaction and the ACID properties. A taxonomy of the different concurrency control protocols with different observable behavior and performance characteristics appears in Section 3. The Transactional Featherweight Java calculus is introduced in Section 4. Section 5 and Section 6 give, respectively, versioning and strict two phase locking semantics to the calculus. The soundness results are given in Section 7. Observations about the calculus is given in Section 8. Related work is discussed in Section 9 and Section 10 concludes.

2. Atomic transactions

Atomic transactions [15] are a control abstraction that arose in the database community to delimit logical units of data processing. They permit concurrent data access among transactions, with operations on shared data interleaved to the extent that they do not violate the so-called “ACID” transaction semantics, ensuring:

1. *Atomicity:* A transaction is a sequence of operations that is performed *atomically*, either in its entirety or else not at all. If it completes successfully, it *commits*. Otherwise, it *aborts* and has no effects.
2. *Isolation:* A transaction should not make its effects visible to other transactions until it commits. The effect is that concurrent transactions are *serializable* in that they appear to occur one at a time.
3. *Consistency:* A correct execution of a transaction takes shared data from one consistent state (as defined by the application semantics) to another provided that transactions execute atomically and in isolation.
4. *Durability:* Once it has committed, the effects of a transaction survive subsequent system failures.

Transactions may be *nested* [22], with each top-level transaction divided into a number of child transactions. Because transactions commit from the bottom up, child transactions must always commit before their parent. A transaction abort at one level does not necessarily affect a transaction in progress at a higher level. The updates of committed transactions at intermediate levels are visible only within the scope of their immediate predecessors. Effects performed by a parent transaction are always visible to the child.

Concurrency control ensures the consistency property, usually via protocols that guarantee transaction serializability. Executions are consistent if they execute in a way that is serializable. Concurrency protocols are classified as either pessimistic or optimistic. The former detect violations of serializability property early but usually incur significant runtime and space overheads, the latter are less expensive but detect serializability violations as late as the commit time. *Two-phase locking* (2PL) [13] is a protocol that requires all lock operations on behalf of a transaction to precede the *first* unlock operation in the transaction, resulting in two computation phases: a growing phase as locks are acquired




Fig. 1. Threads and transactions may be interleaved in various ways: (a) plain, (b) nested, (c) multi-threaded, (d) multi-threaded and nested.

and a shrinking phase as they are released. In *strict 2PL* a transaction does not release any of its locks until after it commits or aborts. 2PL can be extended to handle nested transactions as follows [22]: a transaction may acquire a lock if all owners are ancestors of the transaction and when a transaction commits, all its locks are handed back to the parent or released if the transaction is top level. Optimistic protocols are based on a simple idea. Every read or write is remembered in a private *journal* (or log) instead of being directly applied to the global data. The transaction proceeds uninterrupted until commit time, but before it is allowed to commit, the controller performs conflict detection. If the serializability property has been violated then the transaction aborts and all the data in its private log is abandoned. There are various ways of performing conflict detection, e.g. comparing initial (remembered at the first access) values of data used by a transaction with values of the same data at commit time. Extending optimistic schemes to incorporate nested transactions is reasonably straightforward and involves modifying the commit procedure. Every transaction keeps a separate journal and at commit time the updates performed by a nested transaction are propagated from its private journal to the journal of its immediate parent. The updates are propagated to the global data set only after a successful commit of the top-level transaction.

3. A taxonomy of protocols

For a given transactional model, such as the nested and multi-threaded one studied here (see Fig. 1), there exist a range of concurrency control protocols with different observable behavior and performance characteristics. While we do not purport to present an exhaustive list, we list some known interesting points in the spectrum.

- **Precise (pessimistic)**: Block a transaction T iff it will definitely have a conflict with a transaction T' . This is in general undecidable without an oracle.

- **Strict 2PL**: Block a transaction T if a *potential* conflict is detected with a transaction T' . Under strict 2PL locks on objects may be acquired before objects are referenced or modified with new values.
- **2PL**: Similar to strict 2PL except that lock acquisition is less aggressive. Locks are acquired only when necessary to ensure atomicity properties. Note that an apparent conflict such as two writes to the same variable may not be a real conflict, e.g. both writes affect the same value to the variable and no intervening read observes the change.
- **Byzantine**: An optimistic scheme in which any transaction can be aborted at any time, this even if the transaction will not have a conflict. All optimistic transactions can lead to livelock.
- **Fail-fast**: Abort a transaction T as soon as a potential conflict has been detected with some transaction T' . Thus, every read and write checks whether serialization invariants on the target object have been violated, and aborts if so.
- **Precise (optimistic)**: Abort a transaction T if, and only if, a conflict has been detected with some transaction T' when T is about to commit. This protocol delays aborts, thus allowing the possibility that a transaction performing a write action which would have led to an abort in a fail-fast scheme is allowed to commit because subsequent actions mask the effect of the write.

The tradeoffs between these protocols range from programming model issues (which kinds of failures should the programmer be exposed to, how quickly should failures be detected, and how should deadlocks and livelocks be resolved) to implementation performance. For example, using memory pages as the unit of granularity for detecting conflicts may be efficient, but may lead to unpredictable or frequent aborts. From a programmer's point of view this is akin to byzantine failures in distributed systems. While a more detailed and formal investigation of these alternatives remains as an open problem, we pick two points in the design space – precise optimism (which we refer to as a *versioning* instantiation) and strict 2PL (which we refer to as a *locking* instantiation) – and give their semantics in the following sections.

4. The Transactional Featherweight Java calculus

Transactional Featherweight Java (TFJ) is a new object calculus inspired by the work of Igarashi et al. [17]. TFJ includes threads and the imperative constructs needed to model transactions. We focus on a simplified variant of TFJ, that is dynamically typed and in which all classes directly inherit the distinguished `Object` class.¹ To introduce the syntax and semantics of TFJ, we start with a simple example.

Consider the classes given in Fig. 2. Class `Updater` encapsulates an update to an object. Class `Runner` is designed to perform an action within a new thread. Class `Transactor` performs two operations within a transaction. In more detail, class `Updater` has two fields, `n` and `v` and an `update` method which assigns `v` to `n`'s `val` field. `Runner` has a

¹ Even though types and inheritance are central features of object-oriented languages they are orthogonal to issues studied in this paper. The interaction of concurrency and inheritance is well-studied [2,21], and no additional novel problems are posed by considering object inheritance in the context of transactions.

```

class Updater {
  n, v;
  init( n, v) { this.n := n; this.v := v; this; }
  update() { this.n.val:= this.v; }
}

class Runner {
  r;
  init( r) { this.r := r; this; }
  run() { spawn this.r.run(); }
}

class Transactor {
  u, r;
  init( r, u) { this.u := u; this.r := r; this; }
  run() {
    onacid;
    this.u.update(); this.r.run(); this.u.n.val;
    commit;
  }
}

```

Fig. 2. Example class definitions.

```

n := new Number();
s1 := new Transactor.init( Noop, new Updater().init( n, new One()));
r1 := new Runner().init( s1);
s2 := new Transactor.init( r1, new Updater().init( n, new Two()));
new Runner().init( s2).run();
n.val := new Three()

```

Fig. 3. A TFJ code fragment using definition of Fig. 2.

run method which starts a new thread and invokes a method on its `r` field within that thread. `Transactor` has a `pick` method which is used to evaluate two expressions in a non-deterministic order; non-determinism is achieved since the order in which arguments are evaluated in a method call is unspecified. It also has a `run` method which starts a new transaction and invokes `update` on field `u` and `run()` on field `r`. The keyword **onacid** marks the beginning of a transaction and **commit** ends the current transaction. All objects have an `init` method which assigns values to their fields and returns `this`.

Fig. 3 gives a TFJ code fragment making use of the above class definitions. Variable `n` is bound to a new object of some class `Number` (whose only feature of interest is that it must have a `val` field; we further assume the existence of classes `One`, `Two`, and `Three` that define specific numbers). `Noop` is an initialized `Runner` object with an uninteresting `run` method. Objects `s1` and `s2` are transactors which will be used to initiate nested transaction (`s2` within `s1`). Two runner objects will be used to create threads t_1 and t_2 .




Fig. 4. Threads and transactions in Fig. 3. Threads t_0 , t_1 , t_2 , execute, respectively, in transactions l_0 , l_1 , l_2 . The transactions are nested, such that l_0 is the outermost transaction. Transaction commits are implicit thread joins.

Evaluating the program of Fig. 3 will result in the creation of two threads (t_1 and t_2) and two new transactions (l_1 and l_2). Thread t_1 executes solely within transaction l_1 , while t_2 starts executing in l_1 , before starting transaction l_2 . We assume that there is a default top-level transaction, l_0 and primordial thread t_0 . Fig. 4 shows the structure of this computation. The threads in a parent transaction can execute concurrently with threads in nested transactions. A design choice in TFJ is that all threads must join (via a commit) for the entire transaction to commit. Alternatives, such as interrupting threads that are not at a commit point when another thread in the same transaction is ready to commit, or silently committing changes while the thread is running are either programmer unfriendly or counter to the spirit of transactions.

The state in this program is defined by the instance of class `Number` that is threaded through the transactions and handed down to `Updaters` for modification. Each invocation of `update()` performs a read and a write of the `val`. One valid interleaving of the operations is, for example:

$$[n := \text{One}]_{l_1} \rightarrow [n]_{l_1} \rightarrow [n := \text{Two}]_{l_2} \rightarrow [n]_{l_2} \rightarrow [n := \text{Three}]_{l_0}$$

This is correct because all of the changes performed by l_1 occur before changes of transactions l_2 and l_0 . An example of an invalid interleaving of these operations is:

$$[n := \text{One}]_{l_1} \rightarrow [n := \text{Two}]_{l_2} \rightarrow [n]_{l_1} \rightarrow [n]_{l_2} \rightarrow [n := \text{Three}]_{l_0}$$

In this schedule, serializability is broken because l_1 reads the value of `n.val` that was changed by l_2 . Thus from l_1 's viewpoint the global state is `n.val = Two()`. Most concurrency control protocols will flag this as a conflict and abort l_1 . We note that in this particular case the conflict is benign as l_1 discards the value it read and thus the state of the system is not affected by it reading a stale value.

4.1. Syntax

The syntax of TFJ is given in Fig. 5. We take metavariables L to range over class declarations, C, D to range over classes, M to range over methods, and f and x to range over fields and parameters, respectively. We also use P for process terms, and e for expressions. We assume that the keyword `this` is included in the set of variables, but that `this` is never

used as the name of an argument to a method. We use v to range over basic expressions, and r and u to range over references. The only values in TFJ are references created by $\text{newC}()$ operations, and the special reference null . Field selection, method call, and assignment are subject to evaluation if their constituent elements are references.

We use over-bar to represent a finite ordered sequence, for instance, \bar{f} represents $f_1 f_2 \dots f_n$. The term $\bar{l}l$ denotes the extension of the sequence \bar{l} with a single element l , and $\bar{l} . \bar{l}'$ for sequence concatenation. We write $\bar{l} \triangleleft \bar{l}'$ if \bar{l} is a prefix of \bar{l}' .

A process term P can be either the empty process 0 , the parallel composition of processes $P \mid P$ or a thread t running expression e , denoted $t[e]$. The thread label t is distinct for every thread.

The calculus has a call-by-value semantics. The expression $C(\bar{r}) \downarrow_i^{r'}$ yields an object identical to $C(\bar{r})$ except in the i th field which is set to r' . The null metavariable is used to represent an unbound reference. By default all objects are null initialized (i.e. $C(\overline{\text{null}})$).

Since TFJ has by-value semantics for invocation, sequencing can be encoded as a sequence of method invocations. For readability, we sometimes write “ $(e_1; e_2)$ ” in examples to indicate sequencing of expressions e_1 and e_2 . The value of a sequence is always the value of the last expression.

An *expression* e can be either a variable x , the `this` pseudo-variable, a field access $e.f$, a method invocation $e.m(\bar{e})$, an object construction $\text{newC}()$, a thread creation `spawn` e , an `onacid` command or a `commit`. The latter three operations are unique to TFJ. The expression `spawn` e creates a new thread of control to evaluate e . The evaluation of e takes place in the same environment as the thread executing `spawn` e . A new transaction is started by executing `onacid`. The dynamic context of `onacid` is delimited by `commit`. Effects performed within the context of `onacid` are not visible outside the transaction until a `commit` occurs. Transactions may be nested. When the `commit` of an inner transaction occurs, its effects are propagated to its parent. Threads may be spawned within the context of a transaction. The local state of the transaction is visible to all threads which execute within it. Transactions may also execute concurrently. For example, in `spawn` e , e may be an expression that includes `onacid` and `commit`; the transaction created by `onacid` executes concurrently with the thread executing the `spawn` operation.

Note that the language does not provide an explicit `abort` operation. Transactions may abort *implicitly* because serialization invariants are violated. Our semantics expresses implicit aborts both in the definition of `commit` and in the treatment of read and write operations that would otherwise expose violations of necessary serializability invariants. Implicit aborts are tantamount to stuck states.

4.2. Reduction

The dynamic semantics of our language shown in Figs. 5 and 6 is given by a two-level set of rewrite rules. The computational core of the language is defined by a reduction relation of the form $\mathcal{E} e \xrightarrow{\alpha} \mathcal{E}' e'$. Here \mathcal{E} is a sequence of *transaction environments*, e is an expression and the action label α determines which reduction was picked. Each transaction environment consists of a *transaction label*, and a binding environment that maps references to objects ($v \mapsto C(\bar{r})$). Action labels for the computational core are selected from the set $\{rd, wr, xt\}$, respectively denoting read, write and extend. In addition

Syntax:

$P ::= 0 \mid P|P \mid t[e]$
 $L ::= \text{class } C \{ \bar{f}; \bar{M} \}$
 $M ::= m(\bar{x}) \{ e; \}$
 $e ::= x \mid e.f \mid e.m(\bar{e}) \mid e.f := e \mid$
 $\quad \text{new } C() \mid \text{spawn } e \mid \text{onacid} \mid \text{commit} \mid \text{null}$
 $v ::= r \mid v.f \mid v.m(\bar{v}) \mid v.f := v$

Field look-up:

$$\frac{CT(C) = \text{class } C \{ \bar{f}; \bar{M} \}}{\text{fields}(C) = (\bar{f})}$$
Method body look-up:

$$\frac{CT(C) = \text{class } C \{ \bar{f}; \bar{M} \} \quad m(\bar{x}) \{ e; \} \in \bar{M}}{\text{mbody}(m, C) = (\bar{x}, e)}$$
Local Computation:

$$\frac{\mathcal{E}', C(\bar{u}) = \text{read}(r, \mathcal{E}) \quad \text{fields}(C) = (\bar{f})}{\mathcal{E} \ r.f_i \xrightarrow{rd\ r} \mathcal{E}' \ u_i} \quad (\text{R-FIELD})$$

$$\frac{\mathcal{E}', C(\bar{f}) = \text{read}(v, \mathcal{E}) \quad \mathcal{E}'' = \text{write}(r \mapsto C(\bar{f}) \downarrow_i^r, \mathcal{E}')}{\mathcal{E} \ r.f_i := r' \xrightarrow{wr\ r'} \mathcal{E}'' \ r'} \quad (\text{R-ASSIGN})$$

$$\frac{\mathcal{E}', C(\bar{f}) = \text{read}(r, \mathcal{E}) \quad \text{mbody}(m, C_0) = (\bar{x}, e)}{\mathcal{E} \ r.m(\bar{f}) \xrightarrow{rd\ r} \mathcal{E}' \ [\bar{f}/\bar{x}, r/\text{this}]e} \quad (\text{R-INVK})$$

$$\frac{r \text{ fresh} \quad \mathcal{E}' = \text{extend}(r \mapsto C(\overline{\text{null}}), \mathcal{E})}{\mathcal{E} \ \text{new } C() \xrightarrow{xt\ r} \mathcal{E}' \ r} \quad (\text{R-NEW})$$
Global Computation:

$$\frac{P = P'' \mid t[e] \quad \mathcal{E} \ e \xrightarrow{\alpha} \mathcal{E}' \ e' \quad P' = P'' \mid t[e']}{\Gamma \ P \xrightarrow{\alpha} \Gamma' \ P'} \quad (\text{G-PLAIN})$$

$$\frac{P = P'' \mid t[E[\text{spawn } e]] \quad P' = P'' \mid t[E[\text{null}]] \mid t'[e] \quad t' \text{ fresh} \quad \Gamma' = \text{spawn}(t, t', \Gamma)}{\Gamma \ P \xrightarrow{sp\ t'} \Gamma' \ P'} \quad (\text{G-SPAWN})$$

$$\frac{P = P'' \mid t[E[\text{onacid}]] \quad P' = P'' \mid t[\text{null}] \quad l \text{ fresh} \quad \Gamma' = \text{start}(l, t, \Gamma)}{\Gamma \ P \xrightarrow{ac} \Gamma' \ P'} \quad (\text{G-TRANS})$$

$$\frac{\Gamma = \Gamma'.t, \mathcal{E} \quad \mathcal{E} = \mathcal{E}'.l : \rho \quad \bar{t} = t_1 \dots t_n = \text{intranse}(l, \Gamma)}{P = P'' \mid t[E[\text{commit}]] \quad P' = P'' \mid t[E[\text{null}]] \quad t_1, \mathcal{E}_1, t_2, \mathcal{E}_2, \dots, t_n, \mathcal{E}_n \in \Gamma \quad \bar{\mathcal{E}} = \mathcal{E}_1, \mathcal{E}_2, \dots, \mathcal{E}_n \quad \Gamma' = \text{commit}(\bar{t}, \bar{\mathcal{E}}, \Gamma)}{\Gamma \ P \xrightarrow{co} \Gamma' \ P'} \quad (\text{G-COMM})$$

$$\frac{P = P' \mid t[r] \quad \Gamma = t, \mathcal{E} \ . \Gamma'}{\Gamma \ P \xrightarrow{ki} \Gamma' \ P'} \quad (\text{G-THKILL})$$

Fig. 5. TFJ syntax and semantics.

Evaluation Contexts:

$$E[\bullet] \quad | \quad E[\bullet].f := e \quad | \quad e.f := E[\bullet] \quad | \quad E[\bullet].m(\bar{e}) \quad | \quad e.m(\bar{r}, E[\bullet], \bar{e})$$

Congruence:

$$\frac{\mathcal{E} e \xrightarrow{\alpha} \mathcal{E}' e'}{\mathcal{E} E[e] \xrightarrow{\alpha} \mathcal{E}' E[e']}$$

Transaction membership:

$$\frac{}{nested(\mathbb{1}, \langle \rangle) = \langle \rangle} \quad \frac{\Gamma = \mathfrak{t}, \mathcal{E} . \Gamma' \quad nested(\mathbb{1}, \Gamma') = \bar{\mathfrak{t}} \quad \bar{\mathbb{1}} . \mathbb{1} . \mathbb{1}' . \bar{\Gamma}' = \ell(\mathfrak{t}, (\mathfrak{t}, \mathcal{E}))}{nested(\mathbb{1}, \Gamma) = \mathfrak{t}\bar{\mathfrak{t}}} \quad \frac{\Gamma = \mathfrak{t}, \mathcal{E} . \Gamma' \quad nested(\mathbb{1}, \Gamma') = \bar{\mathfrak{t}} \quad \mathbb{1} \notin \ell(\mathfrak{t}, (\mathfrak{t}, \mathcal{E})) \vee \bar{\mathbb{1}}\mathbb{1} = \ell(\mathfrak{t}, (\mathfrak{t}, \mathcal{E}))}{nested(\mathbb{1}, \Gamma) = \bar{\mathfrak{t}}}$$

$$\frac{}{intranse(\mathbb{1}, \langle \rangle) = \langle \rangle} \quad \frac{\Gamma = \mathfrak{t}, \mathcal{E} . \Gamma' \quad intranse(\mathbb{1}, \Gamma') = \bar{\mathfrak{t}} \quad \mathbb{1} \in \ell(\mathfrak{t}, (\mathfrak{t}, \mathcal{E})) \quad nested(\mathbb{1}, \Gamma) = \langle \rangle}{intranse(\mathbb{1}, \Gamma) = \mathfrak{t}\bar{\mathfrak{t}}} \quad \frac{\Gamma = \mathfrak{t}, \mathcal{E} . \Gamma' \quad intranse(\mathbb{1}, \Gamma') = \bar{\mathfrak{t}} \quad \mathbb{1} \notin \ell(\mathfrak{t}, (\mathfrak{t}, \mathcal{E}))}{intranse(\mathbb{1}, \Gamma) = \bar{\mathfrak{t}}}$$

Fig. 6. Auxiliary definitions.

to specifying the action on whose behalf a particular reduction is taken, we also specify the action's effects; for example, we write $wr \ vv'$ to denote an action with label wr which has an effect on locations v and v' . A read action affects the location being read, a write action has an effect on both the location being written and the location whose value it reads, and an extend operation has an effect on the newly created location.

A second reduction relation $\xrightarrow{\alpha}_{\mathfrak{t}}$ defines operations over the entire program and has the form $\Gamma \ P \xrightarrow{\alpha}_{\mathfrak{t}} \Gamma' \ P'$ where Γ is a program state composed of a sequence of thread environments $\bar{\mathfrak{t}}, \mathcal{E}$ where each $\mathfrak{t}, \mathcal{E}$ pair represent the association of a thread to its transaction environments. The action label α can be one of the computational core labels or one of $\{sp, ac, co, ki\}$ for, respectively, spawn, onacid, commit, and kill. As with core actions, the actions corresponding to these labels have an effect on the global state; these effects are given in brackets. Thus, a `spawn` action has the effect of creating a new thread with label \mathfrak{t} ; an `onacid` action creates a new transaction with label $\mathbb{1}$; a `commit` operation has an effect on the current transaction; and, a `kill` action has an effect on the current thread.

The metavariable $\mathbb{1}$ ranges over transaction names, and sequences of transaction names are used to represent the nesting structure. As usual, $\xrightarrow{\alpha}_{\mathfrak{t}}^*$ denotes the reflexive and transitive closure of the global reduction relation. The congruence rules given in Fig. 6 are straightforward.

We work up to congruence of processes ($P|P' \equiv P'|P$, $(P_1|P_2)|P_3 \equiv P_1|(P_2|P_3)$, and $P|0 \equiv P$). Congruence over expressions is defined in terms of evaluation contexts. An evaluation context is a term with a “hole” in it. The expression “picked” for evaluation (i.e., that fills the hole) is determined by the structure of these contexts. Observe that the main role of contexts is to enforce order of evaluation (e.g., left-to-right evaluation for method calls), and to ensure that subexpressions of e in `spawn` (e) are not chosen for evaluation prior to the evaluation of the `spawn` action to create a new thread. The other definitions are similar to those used in the specification of FJ. `fields` returns the list of all fields of a class including inherited ones. `mbody` returns the body of the method in a given class.

Let \mathcal{E} be a transaction environment of the form $l_0:\rho_0 \dots l_n:\rho_n$, then $\ell(\mathcal{E})$ extracts the order transaction label sequence, $\ell(\mathcal{E}) = l_0 \dots l_n$ if $\Gamma = \tau, \mathcal{E} . \Gamma'$. We override the definition of ℓ such that $\ell(\tau, (\tau, \mathcal{E} . \Gamma')) = \ell(\mathcal{E})$. The auxiliary function $last(x, \rho)$ is defined to return a one element sequence containing the last value referenced by x in the environment ρ or the empty sequence if there is no binding for x . It is defined inductively to return $\langle \rangle$ if $\rho = \langle \rangle$, $C(\bar{x})$ if $\rho = \rho' . x \mapsto C(\bar{x})$ and $last(x, \rho')$ otherwise. The function $first(x, \rho)$ is similar but returns the first binding for x in the sequence. Finally, $findlast(x, \mathcal{E})$ finds the last binding for x in transaction environment \mathcal{E} .

There are four computational core reduction rules shown in Fig. 5. (R-FIELD) evaluates a field access expression, (R-ASSIGN) evaluates an assignment expression, (R-INVK) evaluates a method invocation expression and (R-NEW) evaluates an object instantiation expression. Notice that TFJ has a call-by-value semantics which requires that arguments be fully evaluated before performing method invocations or field access; the order in which arguments are evaluated in a call is unspecified. These rules are complemented by five global reduction rules. (G-PLAIN) corresponds to a step of computation, (G-SPAWN) corresponds to a thread creation, (G-TRANS) corresponds to the start of a new transaction, (G-COMM) corresponds to the commit of a transaction, and (G-THKILL) is a reclamation rule for threads in normal form. Most of the rules are straightforward. G-PLAIN makes use of a *reflect* operation that must propagate the action performed to other threads executing within this transaction. Notice that (G-COMM) requires that, if some thread τ running in transaction l is ready to commit, all other threads executing in that transaction be ready to commit. The auxiliary predicate $intranse(l, \Gamma)$ given in Fig. 6 returns the set of threads that currently have the transaction label l . Note that if there is any thread running in a nested transaction (e.g., has label ll' , for some l'), $intranse(l, \Gamma)$ will return the empty sequence as nested transactions must commit before their parent transaction. The (G-THKILL) rule takes care of removing threads that have terminated, to prevent blocking a transaction (terminated threads are not ready to commit).

The dynamic semantics leaves open the specification of a number of operations. In particular, the definitions of *read*, *write*, *spawn*, *extend*, *reflect*, *start*, and *join* are left unspecified. A particular incarnation of a transactional semantics must provide a specification for these operations.

Example. To illustrate the rules, consider the term:

```
r.m(onacid, new C().m().f, commit)
```

To be faithful to the calculus syntax, sequencing of operations is expressed by the order of arguments in a method call. This expression creates an instance of class C , evaluates method m to yield an object that contains field f , and returns the reference denoted by that field. These actions are encapsulated within a transaction. The result of the transaction (here, the reference denoted by field f) is then supplied as an argument to method m accessed via reference r .

We sketch the application of the semantic rules for this term. Assume that this term is evaluated by thread τ in program state Γ , and process state P . By the congruence rule and definition of evaluation contexts, we can pick E to be

```
r.m(•, new C().m().f, commit)
```

and apply rule (G-TRANS). Application of this rule yields a new term,

$$r.m(\text{null}, \text{new } C().m().f, \text{commit})$$

and new program state T' defined to be $\text{start}(l, t, T')$; application of start results in the installation of a new transaction label l for thread t in its transaction environment. We can next “pick” \bullet to be $\text{new } C()$ (by the definition of the context grammar) to yield a new context term:

$$r.m(\text{null}, \bullet.m().f, \text{commit})$$

By application of rule G-PLAIN and rule R-NEW, the binding environment associated with l in t 's transaction environment is extended with a binding for fresh reference v to $C(\overline{\text{null}})$. The resulting term is:

$$r.m(\text{null}, v.m().f, \text{commit})$$

Again, by the structure of evaluation contexts, we can pick E to be:

$$r.m(\text{null}, v.\bullet.f, \text{commit})$$

and apply rule G-PLAIN and R-INVK on term $m()$. Evaluating the body of m (not shown) will eventually yield a reference v' to an object containing field f . Thus, the context becomes

$$r.m(\text{null}, v.v'.\bullet, \text{commit})$$

and is further evaluated through the application of rules G-PLAIN and R-FIELD. The value yielded therein (call it u) results in the term:

$$r.m(\text{null}, u, \text{commit})$$

We now apply rule G-COMMIT to yield a new state that reflects the commitment of this transaction, and a new term:

$$r.m(\text{null}, u, \text{null})$$

This expression can be now reduced by application of rule R-INVK.

5. Versioning semantics

In Fig. 7 we define an instantiation of TFJ in which transactions implement sequences of object versions. The versioning semantics extends the notion of transaction environments to be an ordered sequence of pairs, each pair consisting of a transaction label and an environment. The intuition is that every transaction operates using a private *log*; these logs are treated as sequences of pairs, binding a reference to its value. A log thus records effects that occur while executing within the transaction. A given reference may have different binding values in different logs. If $\mathcal{E} = l_1:\rho_1 . l_2:\rho_2$ then a thread t executing with respect to this sequence of transaction environments is evaluating expressions whose effects are recorded in log ρ_2 and which are part of the dynamic context of an `onacid` command with label l_2 . If l_2 successfully commits, bindings in ρ_2 are merged with those in ρ_1 . Once l_2

$$\begin{array}{c}
\frac{\mathcal{E} = \mathcal{E}' . 1:\rho \quad \text{findlast}(\mathbf{r}, \mathcal{E}) = \mathbf{C}(\overline{\mathbf{r}}) \quad \mathcal{E}'' = \mathcal{E}' . 1:(\rho . \mathbf{r} \mapsto \mathbf{C}(\overline{\mathbf{r}}))}{\text{read}(\mathbf{r}, \mathcal{E}) = \mathcal{E}'', \mathbf{C}(\overline{\mathbf{r}})} \quad \frac{\mathcal{E} = \mathcal{E}' . 1:\rho \quad \text{findlast}(\mathbf{r}, \mathcal{E}) = \mathbf{D}(\overline{\mathbf{r}}) \quad \mathcal{E}'' = \mathcal{E}' . 1:(\rho . \mathbf{r} \mapsto \mathbf{D}(\overline{\mathbf{r}}) . \mathbf{r} \mapsto \mathbf{C}(\overline{\mathbf{r}}))}{\text{write}(\mathbf{r} \mapsto \mathbf{C}(\overline{\mathbf{r}}), \mathcal{E}) = \mathcal{E}''} \\
\\
\frac{\mathcal{E} = \mathcal{E}' . 1:\rho \quad \mathcal{E}'' = \mathcal{E}' . 1:(\rho . \mathbf{r} \mapsto \mathbf{C}(\overline{\mathbf{r}}))}{\text{extend}(\mathbf{r} \mapsto \mathbf{C}(\overline{\mathbf{r}}), \mathcal{E}) = \mathcal{E}''} \\
\\
\frac{\Gamma = \mathbf{t}, \mathcal{E} . \Gamma' \quad \Gamma'' = \mathbf{t}', \mathcal{E} . \Gamma}{\text{spawn}(\mathbf{t}, \mathbf{t}', \Gamma) = \Gamma''} \quad \frac{\Gamma = \mathbf{t}, \mathcal{E} . \Gamma' \quad \Gamma'' = \mathbf{t}, (\mathcal{E} . 1:\langle \rangle) . \Gamma}{\text{start}(\mathbf{1}, \mathbf{t}, \Gamma) = \Gamma''} \\
\\
\frac{}{\text{reflect}(\mathbf{t}, \mathcal{E}, \langle \rangle) = \langle \rangle} \quad \frac{\Gamma = \mathbf{t}, \mathcal{E}' . \Gamma' \quad \text{reflect}(\mathbf{t}, \mathcal{E}, \Gamma') = \Gamma''}{\text{reflect}(\mathbf{t}, \mathcal{E}, \Gamma) = \mathbf{t}, \mathcal{E}' . \Gamma''} \quad \frac{\Gamma = \mathbf{t}', \mathcal{E}' . \Gamma' \quad \text{reflect}(\mathbf{t}, \mathcal{E}, \Gamma') = \Gamma'' \quad \text{copy}(\mathcal{E}, \mathcal{E}') = \mathcal{E}'' \quad \Gamma''' = \mathbf{t}', \mathcal{E}'' . \Gamma''}{\text{reflect}(\mathbf{t}, \mathcal{E}, \Gamma) = \Gamma'''} \\
\\
\frac{}{\text{commit}(\langle \rangle, \langle \rangle, \Gamma) = \Gamma} \quad \frac{\mathcal{E} = \mathcal{E}' . 1:\rho \quad \text{readset}(\rho, \langle \rangle) = \rho' \quad \text{writeset}(\rho, \langle \rangle) = \rho'' \quad \text{check}(\rho', \mathcal{E}') \quad \mathcal{E}' = \mathcal{E}'' . 1:\rho''' \quad \text{reflect}(\mathbf{t}, (\mathcal{E}'' . 1:\rho'''.\rho''), \Gamma) = \Gamma'' \quad \text{commit}(\overline{\mathbf{t}}, \overline{\mathcal{E}}, \Gamma') = \Gamma''}{\text{commit}(\mathbf{t}\overline{\mathbf{t}}, \mathcal{E}\overline{\mathcal{E}}, \Gamma) = \Gamma''} \\
\\
\frac{\mathcal{E} = 1:\rho . \mathcal{E}'' \quad \mathcal{E}' = 1:\rho'' . \mathcal{E}'''}{\text{copy}(\mathcal{E}, \mathcal{E}') = 1:\rho . \text{copy}(\mathcal{E}'', \mathcal{E}''')} \quad \frac{\mathcal{E} = 1':\rho' . \mathcal{E}'' \quad \mathcal{E}' = 1:\rho . \mathcal{E}'''}{\text{copy}(\mathcal{E}, \mathcal{E}') = 1:\rho . \text{copy}(\mathcal{E}'', \mathcal{E}''')} \\
\\
\frac{}{\text{check}(\langle \rangle, \mathcal{E})} \quad \frac{\text{findlast}(\mathbf{r}, \mathcal{E}) = \mathbf{C}(\overline{\mathbf{r}}) \quad \text{check}(\rho, \mathcal{E})}{\text{check}(\rho . \mathbf{r} \mapsto \mathbf{C}(\overline{\mathbf{r}}), \mathcal{E})} \\
\\
\mathbf{Mod\ sets:} \\
\\
\frac{}{\text{readset}(\langle \rangle, _) = \langle \rangle} \quad \frac{\rho = \mathbf{u} \mapsto \mathbf{C}(\overline{\mathbf{u}}) . \rho'' \quad \mathbf{u} \notin \overline{\mathbf{r}} \quad \text{readset}(\rho'', \overline{\mathbf{r}}\mathbf{u}) = \rho'}{\text{readset}(\rho, \overline{\mathbf{r}}) = \mathbf{u} \mapsto \mathbf{C}(\overline{\mathbf{u}}) . \rho'} \\
\\
\frac{\rho = \mathbf{u} \mapsto \mathbf{C}(\overline{\mathbf{u}}) . \rho'' \quad \mathbf{u} \in \overline{\mathbf{r}} \quad \text{readset}(\rho'', \overline{\mathbf{r}}) = \rho'}{\text{readset}(\rho, \overline{\mathbf{r}}) = \rho'} \\
\\
\frac{}{\text{writeset}(\langle \rangle, _) = \langle \rangle} \quad \frac{\rho = \mathbf{r} \mapsto \mathbf{C}(\overline{\mathbf{r}}) . \rho'' \quad \text{writeset}(\rho'', \rho') = \rho'' \quad \mathbf{r} \mapsto \mathbf{C}(\overline{\mathbf{r}}) \neq \text{first}(\mathbf{r}, \rho')}{\text{writeset}(\rho, \rho') = \mathbf{u} \mapsto \mathbf{D}(\overline{\mathbf{u}}) . \rho''}
\end{array}$$

Fig. 7. Versioning semantics.

commits, subsequent expressions evaluated by t occur within the dynamic context of an `onacid` command with label l_1 ; effects performed by these expressions are recorded in environment ρ_1 .

Thus, a transaction environment in a versioning semantics defines a chain of nested transactions: every $l:\rho$ element in \mathcal{E} is related to its predecessor in the sequence defined by \mathcal{E} under an obvious static nesting relationship. A locus of computation can be therefore uniquely denoted by a thread t and the transaction label sequence \bar{l} in which t is executing.

When a new thread is created (cf. *spawn*), the global state is augmented to include the new thread; evaluation of this thread occurs in a transaction environment inherited from its parent. In other words, a spawned thread begins evaluation in the environment of its parent extant at the point where the thread was created.

When a thread enters a new transaction (cf. *start*), a new transaction environment is added to its state. This environment is represented as a pair consisting of a label denoting the transaction, and a log used to hold bindings for objects manipulated within the transaction. Initially, the newly created transaction is bound to an empty log.

The essence of the versioning semantics is captured by the *read*, *write*, and *commit* operations. If a *read* operation on reference r occurs within transaction l , the last value for r in the log is returned via the auxiliary procedure *findlast*, and the log associated with l is augmented to include this binding. Thus, the first *read* operation for reference r within transaction l will bind a value for r computed by examining the logs of l 's enclosing transactions, choosing the binding value found in the one closest to l . Subsequent reads of r made within l will find a binding value within l 's log. Thus, this semantics ensures an isolation property on reads: once an object is read within transaction l , effects on that object performed within other transactions are not visible until l attempts to commit its changes.

The *write* operation is defined similarly. Note that *write* augments the log of the current transaction with two bindings, one binding the reference to its value prior to the assignment, and the other reflecting the effect of the assignment. The former binding is needed to guarantee transactional consistency. Consider a write to a reference r in transaction l which has not yet been read or written in l . The effects of this write can be made visible when l attempts to commit only if no other transaction has committed modifications to r in the interim between the time where the write occurred, and l attempts to commit. If this invariant were violated, the desired serialization semantics on transaction would fail to hold. The *extend* operation inserts a new binding in the current transaction's log; since the reference being bound is fresh, there is no existing binding in the parent transaction against which a consistency check must be made upon commit.

The *commit* operation is responsible for committing a transaction. In our versioning semantics, a commit results in bindings for objects written within a transaction's log to be propagated to its parent. In order for a commit of transaction l to succeed, it must be the case that the binding value of every reference read or written in l must be the same as its current value in l 's parent transaction. Satisfaction of this condition implies the absence of a data race between l and its parent or siblings. The *reflect* operation defined in *commit* makes visible the effects of l in all threads executing in l 's parent transaction; when used in a transaction-local action, it propagates the effects of the action to other threads executing within this same transaction.

$$\begin{array}{c}
\frac{\mathcal{E} = \mathbf{1}_L : \rho_L . \mathcal{E}' \quad \text{last}(\mathbf{r}, \rho_L) = \bar{\Gamma} \quad \bar{\Gamma} \triangleleft \bar{\ell}(\mathcal{E})}{\text{checklock}(\mathbf{r}, \mathcal{E}) = \text{true}} \quad \frac{\mathcal{E} = \mathbf{1}_L : \rho_L . \mathcal{E}' \quad \text{last}(\mathbf{r}, \rho_L) = \bar{\Gamma} \quad \bar{\Gamma} \triangleleft \bar{\ell}(\mathcal{E}) \quad \mathcal{E}'' = \mathbf{1}_L : (\rho_L . \mathbf{r} \mapsto \ell(\mathcal{E})) . \mathcal{E}'}{\text{acquirelock}(\mathbf{r}, \mathcal{E}) = \mathcal{E}''} \\
\\
\frac{\mathcal{E} = \mathcal{E}' . \mathbf{1} : \rho \quad \text{findlast}(\mathbf{r}, \mathcal{E}) = \mathbf{C}(\bar{\mathbf{r}}) \quad \text{checklock}(\mathbf{r}, \mathcal{E}) = \text{true}}{\text{read}(\mathbf{r}, \mathcal{E}) = \mathcal{E}'' , \mathbf{C}(\bar{\mathbf{r}})} \quad \frac{\text{findlast}(\mathbf{r}, \mathcal{E}) = \mathbf{D}(\bar{\mathbf{u}}) \quad \mathcal{E}' = \text{acquirelock}(\mathbf{r}, \mathcal{E}) \quad \mathcal{E}'' = \mathcal{E}' . \mathbf{1} : \rho \quad \mathcal{E}''' = \mathcal{E}'' . \mathbf{1} : (\rho . \mathbf{r} \mapsto \mathbf{D}(\bar{\mathbf{u}})) . \mathbf{r} \mapsto \mathbf{C}(\bar{\mathbf{r}})}{\text{write}(\mathbf{r} \mapsto \mathbf{C}(\bar{\mathbf{r}}), \mathcal{E}) = \mathcal{E}'''} \\
\\
\frac{}{\text{commit}(\langle \rangle, \langle \rangle, \Gamma) = \Gamma} \quad \frac{\mathcal{E} = \mathbf{1}_L : \rho_L . \mathcal{E}' \quad \rho'_L = \text{release}(\ell(\mathcal{E}), \rho_L) \quad \mathcal{E}'' = \mathbf{1}_L : \rho'_L . \mathcal{E}' \quad \text{reflect}(\mathbf{t}, \mathcal{E}, \Gamma) = \Gamma' \quad \text{commit}(\bar{\mathbf{r}}, \bar{\mathcal{E}}, \Gamma') = \Gamma''}{\text{commit}(\mathbf{t}\bar{\mathbf{r}}, \mathcal{E}\bar{\mathcal{E}}, \Gamma) = \Gamma''} \\
\\
\frac{\mathcal{E}' . \mathbf{1} : \rho = \text{acquirelock}(\mathbf{r}, \mathcal{E}) \quad \mathcal{E}'' = \mathcal{E}' . \mathbf{1} : (\rho . \mathbf{r} \mapsto \mathbf{C}(\bar{\mathbf{r}}))}{\text{extend}(\mathbf{r} \mapsto \mathbf{C}(\bar{\mathbf{r}}), \mathcal{E}) = \mathcal{E}'''} \quad \frac{\rho_L = \rho'_L : \mathbf{r} \mapsto \bar{\Gamma} \mathbf{1} \quad \rho''_L = \text{release}(\bar{\Gamma} \mathbf{1}, \rho'_L) . \mathbf{r} \mapsto \bar{\Gamma}}{\text{release}(\bar{\Gamma} \mathbf{1}, \rho_L) = \rho''_L} \quad \frac{\rho_L = \rho'_L : \mathbf{r} \mapsto \bar{\Gamma} \quad \bar{\Gamma} \mathbf{1} \neq \bar{\Gamma}}{\text{release}(\bar{\Gamma} \mathbf{1}, \rho_L) = \rho''_L}
\end{array}$$

Fig. 8. Lock-based commitment semantics.

The versioning semantics defined here is akin to an optimistic concurrency protocol in which the validity of reads and writes of references performed within a transaction $\mathbf{1}$ is determined by the absence of modifications to those references in transactions which commit between the time the first read or write of the reference takes place in $\mathbf{1}$ and the time $\mathbf{1}$ commits. For example, consider transaction $\mathbf{1}_1$ that commits \mathbf{r}_1 , transaction $\mathbf{1}_2$ that commits \mathbf{r}_2 and transaction $\mathbf{1}$ that accesses both \mathbf{r}_1 and \mathbf{r}_2 ; a valid serialization of these transactions would commit $\mathbf{1}_1$ prior to the first access of \mathbf{r}_1 in $\mathbf{1}_2$, and would commit $\mathbf{1}_2$ prior to the first access of \mathbf{r}_2 in $\mathbf{1}$. Provided $\mathbf{1}_2$ does not modify \mathbf{r}_1 , no atomicity or consistency invariants on these transactions would be violated.

6. Strict two-phase locking

With slight alteration, the versioning semantics can be modified to support a two-phase locking protocol. Our semantics is faithful to a locking protocol in which locks are first acquired on objects on behalf of a transaction before the objects can be accessed, and released only when commit actions on transactions occur. In our semantics, a reader executing within transaction $\mathbf{1}$ can read an object provided that the lock on that object is held by transaction $\mathbf{1}$ or a prefix of $\mathbf{1}$; every object is initially locked by the transaction in which it was created. A writer executing within transaction $\mathbf{1}$ can acquire exclusive access to an object provided the object is currently locked by $\mathbf{1}$ or a prefix of $\mathbf{1}$. To support locking,

we define a unique transaction label l_L bound to a lock environment ρ_L ; $\rho_L(\mathbf{r})$ maps \mathbf{r} to the transaction label sequence which identifies the transaction that currently has exclusive access to \mathbf{r} . If $\bar{l} = l_1.l_2 \dots l_n$ is such a sequence, then any thread t executing within l_p where $n \leq p$ can acquire a lock for \mathbf{r} . Lock ownership is changed either because (a) the transaction in which a read or write action occurs is a prefix of the transaction which currently owns the lock, or (b) the lock is currently owned by a child transaction which is about to commit, and lock ownership must be transferred to the parent transaction. Unlike the versioning semantics presented earlier, commit actions always succeed since the manner in which locks are acquired ensure that no serializability violations ensue. As a consequence, there is no explicit notion of abort in this definition. Once a lock is acquired, the transaction has exclusive ownership until it commits, or a child attempts to access the object. Transactions implicitly abort if it reaches a stuck state; in this case, a deadlock would be modeled by a global state in which every thread is stuck, executing within a transaction that requires a lock held by another. The modifications necessary to support two-phase locking are shown in Fig. 8.

7. Soundness

Proving the soundness of a particular transactional facility requires relating it to desired serialization characteristics that dictate a transaction's ACID properties. For any abort-free program trace there must be a corresponding trace in which the transactions executed serially, i.e. all concurrent transactions execute atomically wrt one another. The key idea is that we should be able to *reorder* any abort-free sequence of reduction steps into a sequence that yields the same final state and in which reduction steps taken on behalf of different parallel transactions are not interleaved. We proceed to formalize this intuitive definition.

The height of a transaction environment $\mathcal{E} = l_0:\rho_0 \dots l_n:\rho_n$, written $|\mathcal{E}|$, is n . For a state Γ , $\max(\Gamma)$ returns a thread transaction environment t, \mathcal{E} such that \mathcal{E} is the environment with the largest height $|\mathcal{E}|$ in Γ .

Intuitively, a *well-defined* state is a state in which there are no conflicts. Or in other words, a well-defined state is a state in which, if all threads were to attempt to commit, no abort would occur. One way to express the notion of a successful commit is by comparing the state of a transaction's parent at the point the child is started (call it S_i), and its state when the child commits (call it S_f). If the objects modified by the child have the same value in both S_i and S_f , the child can commit its changes since no atomicity or serialization invariants with respect to the parent have been violated. We formalize this intuition thus:

Definition 1 (*Well-defined*). Let $\Gamma = (t, \mathcal{E}) . \Gamma'$. We say that state Γ is well-defined if Γ' is also well-defined and for $\mathcal{E} = l_1:\rho_1 \dots l_n:\rho_n$, we have $\text{first}(\mathbf{r}, \rho_j) = \text{last}(\mathbf{r}, \rho_{j-1})$ if $2 \leq j \leq n$, and $\mathbf{r} \in \text{Dom}(\rho_{j-1}) \cap \text{Dom}(\rho_j)$.

To define soundness properties, we introduce the notion of control and data dependencies. A dependency defines a relation on actions which can be used to impose structure of transition sequences. In other words, a well-defined transition sequence will be

one in which dependencies are not violated, and thus define safe serial orderings. To define dependencies, we first formalize a notion of an action: given a transition $P \Gamma \xrightarrow{\alpha}_{\tau} P' \Gamma'$, we say that the corresponding *action*, written \mathcal{A} is $(\alpha, \tau, \ell(\tau, \Gamma))$.

Definition 2 (Control Dependency). Define a preorder $\overset{c}{\sim}$ on actions such that $\mathcal{A}_1 \overset{c}{\sim} \mathcal{A}_2$ (read \mathcal{A}_1 is control-dependent on \mathcal{A}_2) if the following holds:

1. $\mathcal{A}_1 = (\alpha, \tau, \bar{\Gamma})$ and $\mathcal{A}_2 = (sp \tau, \tau', \bar{\Gamma})$.
2. $\mathcal{A}_1 = (co, \tau, \bar{\Gamma})$ and $\mathcal{A}_2 = (\alpha, \tau', \bar{\Gamma}')$ where $\alpha \in \{rd, wr, xt\}$ and $\bar{\Gamma}' \triangleleft \bar{\Gamma}$.
3. $\mathcal{A}_1 = (\alpha, \tau, \bar{\Gamma})$ and $\mathcal{A}_2 = (ac, \tau', \bar{\Gamma}')$ where $\bar{\Gamma}' \triangleleft \bar{\Gamma}$.

Actions within a thread have an obvious dependency on the *spawn* operation that creates the thread. A commit action has a control dependency with *read* and *write* actions performed within the transaction to be committed or any of its parents; our definition of nested transactions requires that parent actions are visible to their children. In particular, *read*, *write* and *extend* actions performed by a parent transaction are influenced by commit actions performed by a child. Finally, actions performed within a transaction $\mathbb{1}$ depend upon the *onacid* operation that creates the transaction in which they execute.

Definition 3 (Data Dependency). Define a preorder $\overset{d}{\sim}$ on actions such that $\mathcal{A}_1 \overset{d}{\sim} \mathcal{A}_2$ (read \mathcal{A}_1 is data dependent on \mathcal{A}_2) if \mathcal{A}_1 is either $(rd \ r, \tau, \bar{\Gamma})$, $(wr \ rr', \tau, \bar{\Gamma})$ or $(wr \ r'r, \tau, \bar{\Gamma})$, and \mathcal{A}_2 is either $(wr \ rr'', \tau', \bar{\Gamma}')$ or $(xt \ r, \tau', \bar{\Gamma}')$, with $\bar{\Gamma}' \triangleleft \bar{\Gamma}$.

A read or write action in some transaction that effects locations has a data dependency with a write or extend action executed in a parent transaction that effects these same locations since all such effects induced by the parent are visible to its descendants.

The key property for our soundness result is the permutation lemma which describes the conditions under which two reduction steps can be permuted. Let \mathcal{A} and \mathcal{A}' be a pair of actions which are not related under a control or data dependency. We write $\mathcal{A} \overset{\bar{d}}{\sim} \mathcal{A}'$ and $\mathcal{A} \overset{\bar{c}}{\sim} \mathcal{A}'$ to mean action \mathcal{A} has, respectively, no *c*-dependence or *d*-dependence on \mathcal{A}' .

Definition 4 (Independence). Action \mathcal{A} is independent of \mathcal{A}' if $\mathcal{A} \overset{\bar{c}}{\sim} \mathcal{A}'$ and $\mathcal{A} \overset{\bar{d}}{\sim} \mathcal{A}'$.

The permutation lemma states that any pair of actions which have no dependencies with one another can be permuted in a reduction sequence. This result follows intuitively since the reordering of these actions is inconsequential because neither action observes the effects produced by the other. In other words, the definition of data and control dependence ensures that permuting actions which no have dependency relation can never introduce a conflict, or fail to prevent a commit operation as a result of the reordering provided the same commit action would have succeed in the original (unpermuted) sequence. The permutation lemma ensures that the state of the computation is the same irrespective of the order of the reduction steps (though the intermediate steps may differ). This lemma is specific to a particular transactional facility semantics (e.g., optimistic, two-phase locking) and must be proved with respect to this semantics.

Lemma 1 (*Permute*). Assume that Γ and Γ' are well-defined, and let R be the two-step sequence of reductions $P \Gamma \xrightarrow{\alpha}_{\tau} P_0 \Gamma_0 \xrightarrow{\alpha'}_{\tau'} P' \Gamma'$. Let \mathcal{A} be \mathcal{A} is independent of \mathcal{A}' then there exists a two-step sequence R' such that R' is $P \Gamma \xrightarrow{\alpha'}_{\tau'} P_1 \Gamma_1 \xrightarrow{\alpha}_{\tau} P' \Gamma'$.

Proof (*Versioning Semantics*). We prove the lemma by cases on α', τ' . The two most interesting cases are when α' is either *rd* or *wr*. We assume that $\bar{\Gamma} = \ell(\tau, \Gamma)$ and $\bar{\Gamma}' = \ell(\tau', \Gamma_0)$.

Case $\alpha' = rd \ r'$ and $\bar{\Gamma} \triangleleft \bar{\Gamma}'$.

By definition of data dependence α cannot be a *wr* or *xt* on r' . Since nested transactions must commit before the parent, α cannot be a *co* action on $\bar{\Gamma}$. By definition of control dependence, α cannot be an *ac* action on $\bar{\Gamma}'$ or a *sp* action on τ . Since Γ' is well defined, α is not a *ki* of τ' . There are five remaining cases for α, τ :

1. α is *rd* r or α is *rd* r' : by the definition of *read*, a read operation only extends the bindings in the latest transaction environment of τ with the value of r or r' . Thus, reordering two read operations has no visible effect.
2. α is *wr* $r \ r''$ or *xt* r' : *write* and *extend* operations add new bindings to the latest transaction environment of τ , and in the case of *write*, a binding that records the current value of the object. As (α, τ, Γ) is independent of $(\alpha', \tau', \Gamma')$, by definition of data dependence it must be that $r \neq r'$, these effects on the log are not visible to any read on r .
3. α is *sp* or *ki*: A spawn action augments the global state; by the definition of independence, the spawned thread is unrelated to the thread in which α' executes. Similar reasoning applies for thread termination.

Case $\alpha' = wr \ r'$ and $\bar{\Gamma} \triangleleft \bar{\Gamma}'$.

Since *write* and *extend* actions performed by a child transaction are not visible to its parent, a *commit* action is the only action that a child can perform whose effects would be visible to a *read* action undertaken by the parent. By definition of control dependence, however, α cannot be a *co* if (α, τ) is independent of (α', τ') .

Case $\alpha' = wr \ r' r''$ and $\bar{\Gamma} \triangleleft \bar{\Gamma}'$.

By definition of data dependence, α is neither a *wr* or *xt* action on r' or r'' . By definition of control dependence, α is not a *sp* of τ' , an *ac* of $\bar{\Gamma}'$, or a *co* of $\bar{\Gamma}$. The remaining cases for α are:

1. α is *rd* r : by definition of *write*, changes made within a child transaction $\bar{\Gamma}$ are not visible to the parent $\bar{\Gamma}'$. Thus, a write operation to location r' performed by a child transaction can be permuted with any read operation performed by the parent.
2. α is *wr* $r r'''$ or *xt* r : as above, write actions have no dependence with write or extend actions to different locations, and can thus be permuted with those actions as well.
3. α is *ac*: An action that creates a new transaction has no effect on α', τ' provided the transaction being created is not $\bar{\Gamma}'$.
4. α is *sp* or *ki*: Actions that spawn or kill threads have no effect on α' provided the thread they manipulate is not t' .

Case $\alpha' = wr\ r\ r'$ and $\bar{1}' \triangleleft \bar{1}$.

By definition of data dependency, α cannot be $rd\ r$, $wr\ r\ r''$ or $wr\ r'\ r$. By definition of control dependency, α cannot be a co of $\bar{1}'$. Since no other child action is visible to the parent α can be permuted with any other action.

Having considered symmetric cases for read and write actions, note that ac of $\bar{1}$ is not permutable with any action on $\bar{1}$. Similarly, sp , τ is not permutable with any action α_t . Since the pair of actions defining R are independent, such permutations are not considered. All other actions do not induce observable effects outside the transaction in which they occur and are thus permutable with actions performed in another transaction. \square

Proof (Locking semantics). We prove the lemma by cases on α' , τ' . As before, the two most interesting cases are when α' is either rd or wr . We assume that $\bar{1} = \ell(\tau, \Gamma)$ and $\bar{1}' = \ell(\tau', \Gamma_0)$.

Case $\alpha' = rd\ r'$ and $\bar{1} \triangleleft \bar{1}'$.

By definition of data dependence α cannot be a wr or xt on r' . Since nested transactions must commit before the parent, α cannot be a co action on $\bar{1}$. By definition of control dependence, α cannot be an ac action on $\bar{1}'$ or a sp action on τ . Since Γ' is well defined, α is not a ki of τ' . There are five remaining cases for α , τ :

1. α is $rd\ r$ or α is $rd\ r'$: A read operation extends the bindings in the transaction environment of τ with the value of r or r' , and checks the lock environment of the object read. Now, since we assume $P\ \Gamma \xrightarrow{\alpha}_{\tau} P_0\ \Gamma_0 \xrightarrow{\alpha'}_{\tau'} P'\ \Gamma'$, it must be the case that r and r' are locked by transactions $\bar{1}''$ and $\bar{1}'''$ such that $\bar{1}'' \triangleleft \bar{1}$ and $\bar{1}''' \triangleleft \bar{1}$. Since reads do not affect the lock environment, reordering read actions is permissible assuming the validity of the original transitions.
2. α is $wr\ r\ r''$ or $xt\ r$: A *write* and *extend* adds a binding to the lock environment ρ_L such that $\rho_L(r) = 1$. Since $r \neq r'$, the change in the lock environment reflecting the lock on r has no effect on the *read* on r' .
3. α is sp or ki : As with the versioning semantics, a spawn action augments the global state; by the definition of independence, the spawned thread is unrelated to the thread in which α' executes. Similar reasoning applies for thread termination.

Case $\alpha' = rd\ r'$ and $\bar{1}' \triangleleft \bar{1}$.

Since $P\ \Gamma \xrightarrow{\alpha}_{\tau} P_0\ \Gamma_0 \xrightarrow{\alpha'}_{\tau'} P'\ \Gamma'$, it must be the case that $r \neq r'$ or r is locked by transaction $\bar{1}'' \triangleleft \bar{1}'$. In either case, the actions are permutable. Suppose $\alpha = wr\ r'$. Then, $\rho_L(r') = \bar{1}$. But, this would mean that the *read* on r' in $\bar{1}'$ would not succeed by the definition of *checklock*. A similar argument applies for *extend*. Thus, α cannot be a *write* or *read* to r' . Operations on locations other than r' can be permuted since consistency checks performed by *read* only examine locks on its argument. The permutability of α' with respect to *commit* follows from the definition of control dependency.

Case $\alpha' = wr\ r'\ r''$ and $\bar{1} \triangleleft \bar{1}'$.

By definition of data dependence, α is neither a wr nor xt action on r' or r'' . By

definition of control dependence, α is not a *sp* of τ' , an *ac* of $\bar{\Gamma}'$, or a *co* of $\bar{\Gamma}$. The remaining cases for α are:

1. α is *rd* \underline{x} : by definition of data dependence $\underline{x} \neq \underline{x}'$. By definition of *write*, any update is protected by a lock which is not accessed by *read* since $\underline{x} \neq \underline{x}'$.
2. α is *wr* $\underline{r}\underline{r}''$ or *xt* \underline{x} : By definition of data dependence, $\underline{x} \neq \underline{x}'$. As above, write actions have no dependence with write or extend actions to different locations which are protected by different locks, and can thus be permuted with those actions as well.
3. α is *ac*: An action that creates a new transaction has no effect on α' , τ' provided the transaction being created is not $\bar{\Gamma}'$.
4. α is *sp* or *ki*: Actions that spawn or kill threads have no effect on α' provided the thread they manipulate is not τ' .

Case $\alpha' = \text{wr } \underline{r}\underline{r}'$ and $\bar{\Gamma}' \triangleleft \bar{\Gamma}$.

By definition of data dependency, α cannot be *rd* \underline{x} , *wr* $\underline{r}\underline{r}''$ or *wr* $\underline{r}'\underline{r}$. By definition of control dependency, α cannot be a *co* of $\bar{\Gamma}'$. All other actions on other locations manipulate distinct locks that do not affect permutability.

The remaining cases are identical to their treatment under a versioning semantics. \square

Next, we define the notion of a program trace which abstracts program and states.

Definition 5 (Program Trace). Let R be the sequence of reductions $P_0 \Gamma_0 \xrightarrow{\alpha_0} t_0 \dots P_n \Gamma_n \xrightarrow{\alpha_n} t_n P_{n+1} \Gamma_{n+1}$. The trace of the reduction sequence R , written $\text{tr}(R)$, is $(\alpha_0, \tau_0, \bar{\Gamma}_0) \dots (\alpha_n, \tau_n, \bar{\Gamma}_n)$ assuming that $\bar{\Gamma}_i = \ell(\tau_i, \Gamma_i)$ for $0 \leq i \leq n$.

A program trace is *serial* if for all pairs of reduction steps with the same transaction label ($\bar{\Gamma}$), all reductions occurring between the two steps are taken on behalf of that very transaction or nested transactions ($\bar{\Gamma} \triangleleft \bar{\Gamma}'$).

Definition 6 (Serial Trace). A program trace, $\text{tr}(R) = (\alpha_0, \tau_0, \bar{\Gamma}_0) \dots (\alpha_n, \tau_n, \bar{\Gamma}_n)$ is serial iff $\forall i, j, k$ such that $0 \leq i \leq j \leq k \leq n$ and $\bar{\Gamma}_i = \bar{\Gamma}_k$ we have $\bar{\Gamma}_i \triangleleft \bar{\Gamma}_j$.

We now present a soundness theorem which states that any sequence of reductions that ends in a well-defined state can be reordered so that its program trace is serial.

Theorem 1 (Soundness). Let R be a sequence of reductions $P_0 \Gamma_0 \xrightarrow{\alpha_0} t_0 \dots P_n \Gamma_n \xrightarrow{\alpha_n} t_n P_{n+1} \Gamma_{n+1}$. If Γ_{n+1} is well-defined, then there exists a sequence R' such that R' is $P_0 \Gamma_0 \xrightarrow{\alpha'_0} t'_0 \dots P'_n \Gamma'_n \xrightarrow{\alpha'_n} t'_n P_{n+1} \Gamma_{n+1}$ and $\text{tr}(R')$ is serial.

Proof. By induction on traces. We must show that for a serial trace $R = P_0 \Gamma_0 \xrightarrow{\alpha_0} t_0 \dots P_n \Gamma_n$, if we take a step and have $R' = R \xrightarrow{\alpha} \tau P \Gamma$, where Γ is well-defined and R' is not serial, that R' can be permuted to an equivalent serial trace. Let $\text{tr}(R) = (\alpha_0, \tau_0, \bar{\Gamma}_0) \dots (\alpha_n, \tau_n, \bar{\Gamma}_n)$. Since Γ_{n+1} is well-defined, we must show that $(\alpha_n, \tau, \bar{\Gamma}_n)$ is independent of $(\alpha_{n-1}, \tau_{n-1}, \bar{\Gamma}_{n-1})$ in order to generate a permutation of R' that defines an equivalent serial trace. Suppose that $\mathcal{A}_n = (\alpha_n, \tau, \bar{\Gamma}_n)$ is not independent of $\mathcal{A}_{n-1} = (\alpha_{n-1}, \tau_{n-1}, \bar{\Gamma}_{n-1})$. Then, it must be the case that either $\mathcal{A}_n \xrightarrow{d} \mathcal{A}_{n-1}$ or $\mathcal{A}_n \xrightarrow{c} \mathcal{A}_{n-1}$.

Since R' is not serial, $\overline{\Gamma_{n-1}} \not\leq \overline{\Gamma_n}$. But, this contradicts the definition of a control dependency. If $\mathcal{A}_n \xrightarrow{d} \mathcal{A}_{n-1}$, $\Gamma_{n+1} = \tau, \mathcal{E}, \Gamma'$ and $\mathcal{E} = \mathcal{E}'.\overline{\Gamma_{n-1}}:\rho_{n-1}:\overline{\Gamma_n}:\rho_n$, then ρ 's bindings for some object r must be extended, and differ from its bindings in ρ_{n-1} . But, then Γ_{n+1} could not be well defined. This means that \mathcal{A}_n is independent of \mathcal{A}_{n-1} and by Lemma 1, they can be permuted. This process can be simply repeated until a serial trace is achieved. \square

8. Observations

We now informally relate the TFJ semantics and soundness results to the four ACID properties of database theory.

Atomicity requires that updates made by a thread be installed in the shared environment in a single atomic operation. Equivalently, from the point of view of other threads, all effects of a committing transaction become visible at the instant the transaction commits. In the versioning semantics of Section 5, atomicity is guaranteed by the *reflect* operation used when committing. Upon commit, *reflect* propagates all updates performed by the committed transaction instantly to the environments of all threads. In the locking semantics, changes performed by a transaction become visible when it releases its locks to its parent.

Consistency implies that transactions executed in isolation preserve the consistency of the database as defined by the application semantics.

The serialization theorem ensures that transactions preserve the *isolation* property since effects made within one transaction are not visible to another until a commit point.

Finally, *durability* requires that changes made by a transaction be persistent. Since our model does not include volatile data, we can assume that this property holds.

9. Related work

The association of transactions with programming control structures has provenance in systems such as Argus [19,22], Camelot [12] Avalon/C++ [11] and Venari/ML [16], and has also been studied for variants of Java, notably by Garthwaite [14] and Daynès [8–10].

Our permutation lemma is related to Lipton's notion of reductions [18] where left and right movers are actions that can be permuted. More recently Qadeer et al. [23] have applied the same idea to generate method summaries for multi-threaded Java.

There is a large body of work that explores the formal specification of various flavors of transactions [20,7,15]. However, these efforts do not explore the semantics of transactions when integrated into a high-level programming language. Most closely related to our goals is the work of Black et al. [3] and Chothia and Duggan [6]. The former presents a theory of transactions that specify atomicity, isolation and durability properties in the form of an equivalence relation on processes. Like our work, they present a soundness result that captures the intuitive notion of serializable actions. Beyond significant technical differences in the specification of the semantics, our results differ most significantly from theirs insofar as we present a stratified semantics for a realistic kernel language intended to

express different concurrency control models within the same framework. We believe our formulation will be thus more useful to implementations and analyses.

Choithia and Duggan present the pik-calculus and pike-calculus, extension of the pi-calculus that supports various abstractions for distributed transactions and optimistic concurrency. Their work is related to other efforts [5,4] that encode transaction-style semantics into the pi-calculus and its variants. Our work is distinguished from these efforts in that it provides a simple operational characterization and proof of correctness of transactions that can be used to explore different tradeoffs when designing a transaction facility for incorporation into a language.

Haines et al. [16] describe a composable transaction facility in ML that supports persistence, undoability, locking and threads. Their abstractions are very modular and first class, although their implementation does not rely on optimistic concurrency mechanisms to handle commits.

10. Conclusions

This paper develops a semantic framework for specifying nested and multithreaded transactions. This semantic framework is a basis for exploring the design space of transactional programming language semantics. We believe that transactions can improve the performance of programs running on parallel architectures and that improve reliability by eliminating data races. We have introduced the TFJ calculus, an imperative and concurrent object calculus with call-by-value semantics and support for nested and multi-threaded transactions. The TFJ semantics is parametrized by the definition of the transactional facility, and in particular of the concurrency control protocol. This is important as different protocols have different observable behaviors and different performance characteristics. So being able to discuss these protocols in the same semantic framework allows one to compare them and argue about tradeoffs. In this paper we show two instantiations of TFJ with, respectively, a versioning-based optimistic model, and a pessimistic two-phase locking protocol. We have proven a general soundness theorem that relates the semantics of TFJ to a serializability property. The soundness result is parametric as well as it relies on a permutation lemma that must be proven for each instantiation of TFJ. While much work remains to be done to integrate transactions into mainstream languages, we believe that this paper is a step in that direction.

References

- [1] M. Atkinson, M. Jordan, A review of the rationale and architectures of PJama: a durable, flexible, evolvable and scalable orthogonally persistent programming platform, Technical Report TR-2000-90, Sun Microsystems Laboratories, June 2000.
- [2] N. Benton, L. Cardelli, C. Fournet, Modern concurrency abstractions for C#, in: Proceedings of the 16th European Conference on Object-Oriented Programming, Springer-Verlag, 2002, pp. 415–440.
- [3] A. Black, V. Cremet, R. Guerraoui, M. Odersky, An equational theory for transactions, Technical Report CSE 03-007, Department of Computer Science, OGI School of Science and Engineering, 2003.
- [4] R. Bruni, C. Laneve, U. Montanari, Orchestrating transactions in the join calculus, in: 13th International Conference on Concurrency Theory, 2002.

- [5] N. Busi, R. Gorrieri, G. Zavattaro, On the serializability of transactions in JavaSpaces, in: *ConCoord 2001, International Workshop on Concurrency and Coordination*, 2001.
- [6] T. Chothia, D. Duggan, Abstractions for fault-tolerant computing, Technical Report 2003-3, Department of Computer Science, Stevens Institute of Technology, 2003.
- [7] P. Chrysanthis, K. Ramamritham, Synthesis of extended transaction models using ACTA, *ACM Transactions on Database Systems* 19 (3) (1994) 450–491.
- [8] L. Daynès, Implementation of automated fine-granularity locking in a persistent programming language, *Software—Practice and Experience* 30 (4) (2000) 325–361.
- [9] L. Daynès, G. Czajkowski, High-performance, space-efficient, automated object locking, in: *Proceedings of the International Conference on Data Engineering*, IEEE Computer Society, 2001, pp. 163–172.
- [10] L. Daynès, G. Czajkowski, Lightweight flexible isolation for language-based extensible systems, in: *Proceedings of the International Conference on Very Large Data Bases*, 2002.
- [11] D.D. Detlefs, M.P. Herlihy, J.M. Wing, Inheritance of synchronization and recovery in Avalon/C++, *IEEE Computer* 21 (12) (1988) 57–69.
- [12] J.L. Eppinger, L.B. Mummert, A.Z. Spector (Eds.), *Camelot and Avalon: A Distributed Transaction Facility*, Morgan Kaufmann, 1991.
- [13] K.P. Eswaran, J.N. Gray, R.A. Lorie, I.L. Traiger, The notion of consistency and predicate locks in a database system, *Communications of the ACM* 19 (11) (1976) 624–633.
- [14] A. Garthwaite, S. Nettles, Transactions for Java, in: M.P. Atkinson and M.J. Jordan (Eds.), *Proceedings of the First International Workshop on Persistence and Java*, Sun Microsystems Laboratories Technical Report 96-58, November 1996, pp. 6–14.
- [15] J. Gray, A. Reuter, *Transaction processing: concepts and techniques*, in: *Data Management Systems*, Morgan Kaufmann, 1993.
- [16] N. Haines, D. Kindred, J.G. Morrisett, S.M. Nettles, J.M. Wing, Composing first-class transactions, *ACM Transactions on Programming Languages and Systems* 16 (6) (1994) 1719–1736.
- [17] A. Igarashi, B.C. Pierce, P. Wadler, Featherweight Java: a minimal core calculus for Java and GJ, *ACM Transactions on Programming Languages and Systems* 23 (3) (2001) 396–450.
- [18] R.J. Lipton, Reduction: a new method of proving properties of systems of processes, in: *Proceedings of the 2nd ACM SIGACT–SIGPLAN Symposium on Principles of Programming Languages*, ACM Press, 1975, pp. 78–86.
- [19] B. Liskov, R. Scheifler, Guardians and actions: Linguistic support for robust distributed programs, *ACM Transactions on Programming Languages and Systems* 5 (3) (1983) 381–404.
- [20] N. Lynch, M. Merritt, W. Weihl, A. Fekete, *Atomic Transactions*, Morgan-Kaufmann, 1994.
- [21] S. Matsuoka, A. Yonezawa, Analysis of inheritance anomaly in object-oriented concurrent programming languages, in: P. Wegner, G. Agha, A. Yonezawa (Eds.), *Research Directions in Concurrent Object-Oriented Programming*, The MIT Press, 1993, pp. 107–150 (Chapter 4).
- [22] J.E.B. Moss, *Nested Transactions: An Approach to Reliable Distributed Computing*, MIT Press, Cambridge, MA, 1985.
- [23] S. Qadeer, S.K. Rajamani, J. Rehof, Summarizing procedures in concurrent programs, in: *Proceedings of the 31st ACM SIGPLAN–SIGACT Symposium on Principles of Programming Languages*, ACM Press, 2004, pp. 245–255.
- [24] J. Vitek, S. Jagannathan, A. Welc, A.L. Hosking, A semantic framework for designer transactions, in: *Proceedings of the European Symposium on Programming*, Springer-Verlag, 2004, pp. 249–263.