Composing First-Class Transactions

NICHOLAS HAINES, DARRELL KINDRED, J. GREGORY MORRISETT,
SCOTT M. NETTLES and JEANNETTE M. WING
Carnegie Mellon University

We describe the design of a transaction facility for a language that supports higher-order functions. We factor transactions into four separable features: persistence, undoability, locking, and threads. Then, relying on function composition, we show how we can put them together again. Our modular approach towards building transactions enables us to construct a model of concurrent, nested, multi-threaded transactions, as well as other non-traditional models where not all features of traditional transactions are present. Key to our approach is the use of higher-order functions to make transactions first-class. Not only do we get clean composability of transactional features, but also we avoid the need to introduce special control and block-structured constructs as done in more traditional transactional systems. We implemented our design in Standard ML of New Jersey.

Categories and Subject Descriptors: D.3.2 [**Programming Languages**]: Language Classifications—*Applicative Languages*; D.3.3 [**Programming Languages**]: Language Constructs and Features—*Control structures*; *Modules, packages*; *Procedures, functions and subroutines*; D.4.1 [**Operating Systems**]: Process Management—*Concurrency*; *Mutual exclusion*; *Synchronization*; D.4.5 [**Operating Systems**]: Reliability—*Fault-tolerance*

General Terms: Languages, Reliability

Additional Key Words and Phrases: transactions, threads, skeins, persistence, recovery, undoability, serializability, Standard ML, modules

1. INTRODUCTION

Transactions are a well-known and fundamental control abstraction that arose from the database community. They have three properties that distinguish them from normal sequential processes: (1) A transaction is a sequence of operations that is performed *atomically* ("all-or-nothing"). If it completes successfully, it *commits*; otherwise, it *aborts* and has no effects. (2) Concurrent transactions are *serializable* (appear to occur one-at-a-time), supporting the principle of isolation. (3) Effects of committed transactions are *persistent* (survive failures). In our model, transactions can be nested.

This research is sponsored in part by Advanced Research Projects Agency monitored by the Wright Laboratory, Aeronautical Systems Center, Air Force Materiel Command, USAF, under contract number F33615-93-1-1330; and in part by National Science Foundation Fellowships for D. Kindred and J. G. Morrisett. Authors' addresses: N. Haines: Harlequin Ltd., Barrington Hall, Barrington, Cambridge, England CB2 5RG; D. Kindred, J.G. Morrisett, S.M. Nettles and J. Wing: School of Computer Science, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213.

The goal of our work is to provide modular support for transactions in a language that supports higher-order functions. By "modular" we mean

—*Factored.* Each key feature of transactions is supported independently of the others.

—*Composable.* Each individual feature can be composed with any other in a meaningful way. Furthermore, transactions themselves are composable with other features of the language.

As part of the Venari project, we chose to pursue these goals in the context of Standard ML of New Jersey [Milner, Tofte, and Harper 1990]. SML/NJ supports higher-order functions, has a powerful modules facility, is freely available, and has an easily modified implementation. We broke transactions into these four separate features:

—*Persistence.* Effects of a computation can outlive the computation.

—*Undoability.* Effects of a computation on the store can be undone.

—*Threads.* A computation may have multiple threads of control.

—*Locks.* Reader/writer (R/W) locks can be used to synchronize access to shared mutable data.

All but the last of these are useful as independent features and represent significant extensions to the semantics of SML. We package each feature into an SML module; each module exports some key higher-order functions. We then rely on higher-order function application to enable seamless composition of transactional features.

In the rest of this section we describe our modular approach to transactions and contrast it with a more traditional approach taken by the transaction community. In Section 2 we describe our design: the four building blocks in our model of transactions and how they compose. In Section 3, we explain how we express our design in SML. We close with discussions evaluating and summarizing our contributions. Throughout, we discuss related work in relevant sections.

1.1 Our Approach

Essential to our approach is linguistic support for higher-order functions. Given a function `f` we want to be able to create a transactional version of `f` by applying the `transact` function to it. Thus,

```
(transact f) a
```

has the effect of applying `f` to `a` within a transaction. A more typical use is as follows:

```
((transact f) a)
    handle Abort => [some work]
```

The `Abort` exception handler allows some special action to be taken if the transaction aborts. Since `(transact f)` is simply a function and functions are first-class, our approach yields first-class transactions.

Most importantly, we want to be able to treat the function `f` as a black box. We want to be able to "wrap" `transact` around any `f` without changing the source code of `f` (or at most by applying a mechanical transformation to it). Someone else may have written `f`; it might even be multi-threaded. Without being able to simply wrap a transaction around a multi-threaded program, for instance, we would be forced to recode each separate thread in `f` as a concurrent nested transaction of a top-level transaction. This violates one aspect of modularity since the entire program would have to be recoded.

Consider a concrete (and the canonical) example where we want to transfer money from a savings account to a checking account in a bank. The transfer involves withdrawing money from the savings account and depositing it in the checking account. We need to make sure that either both the withdrawal and the deposit succeed, or that neither of them occurs. So, we use a transaction to effect the desired behavior, where the actual work of the transfer is done by the `do_transfer` function:

```
fun transfer (savings, checking, amount) =
    let fun do_transfer () =
        (withdraw (savings, amount);     (***)
         deposit (checking, amount))     (***)
    in
        transact do_transfer ()
    end
```

We wrap a transaction around the `do_transfer` function so that if anything goes wrong, e.g., if `withdraw` raises an exception indicating that `savings` has insufficient funds, the whole transfer will be aborted. According to our semantics for `transact` (Section 3.2), if the transfer aborts, we re-raise the exception that caused the abort.

The `do_transfer` function could even have been multi-threaded. For instance, if the lines marked with `(***)` above appeared as following:

```
(fork (fn () => withdraw (savings, amount));
 deposit (checking, amount))
```

then `do_transfer` could still be turned into a transaction simply by applying `transact` to it.

1.2 The Traditional Approach

In contrast, a more traditional approach supported by transactional systems and languages such as CICS [Helland 1985], R* [Lindsay, Haas, Mohan, Wilms, and Yost 1984], Camelot [Eppinger, Mummert, and Spector 1991], Quicksilver [Haskin, Malachi, Sawdon, and Chan 1988], Argus [Liskov and Scheifler 1983], Arjuna [Shrivastava, Dixon, Hedayati, Parrington, and Wheater 1988], and Avalon/C++ [Detlefs, Herlihy, and Wing 1988], requires separate control constructs like `begin_transaction` and `end_transaction` to delimit a transaction's boundary.

For example, a skeleton of the bank transfer operation in Camelot would appear as in Figure 1 [Eppinger, Mummert, and Spector 1991]. There are several disadvantages to this approach. It requires syntactic extensions to the language to support transactions. Such textual extensions do not compose conveniently, nor can such transactions be manipulated as first-class values. Also the lack of exception handling forces the use of the special `status` variable. The programmer could easily forget to check the status after a transaction, in which case aborts would be ignored. Furthermore it is up to the programmer to propagate aborts in nested

```
BEGIN_TRANSACTION
    ...
    if (savings_balance < amount) {
        ABORT(ERROR_INSUFFICIENT_FUNDS);
    }

    ... transfer money ...

END_TRANSACTION(status)

if (status == ERROR_INSUFFICIENT_FUNDS) {
    ...
}
```

Fig. 1. A Bank Transaction in Camelot

transactions.

2. DESIGN OVERVIEW

Transactions may execute at the *top level* (Figure 2a), be *nested* inside one another (Figure 2b), or execute *concurrently* with each other (Figure 2c). Each may be *multi-threaded* (Figure 2d). The combination of all these kinds of transactions yields concurrent, nested, multi-threaded transactions (Figure 2e). In our pictures, we use a wavy line to denote a thread and a box to delimit the scope of a transaction; time advances from left to right. We appeal to tree terminology in discussing nested transactions: a transaction has a unique parent, a set of children, and sets of ancestors and descendants. A transaction is considered its own ancestor and descendant.

Since we separate the basic transactional features into individual components, we need to introduce terms that distinguish a regular transaction from one that supports some but not all features. A *regular transaction* is persistent, undoable, and locking. We use the term *persist-only transaction* for a computation that supports only persistence; we use the term *persistent transaction* for a computation that supports at least persistence. We use similar terms for *undo* and *locking*. When we say "transaction" unqualified, we mean a transaction of any kind (regular, persist-only, undo-only, locking-only, etc.). We will argue in Section 2.2 that all concurrent transactions need to be locking transactions as well.

In Section 2.1 we consider top-level and nested transactions of each flavor; in Section 2.2, we discuss concurrency, and more generally, different combinations of the features.

2.1 The Pieces

2.1.1 *Persistence.* A persistent value is one that outlives the computation that created it. In particular, a persistent value will survive a "crash". We support a model of persistence popularized by the persistent programming language community [Atkinson, Bailey, Chisolm, Cockshott, and Morrison 1983]: *orthogonal persistence*. In this model, all data reachable by pointer dereferencing from a dis-

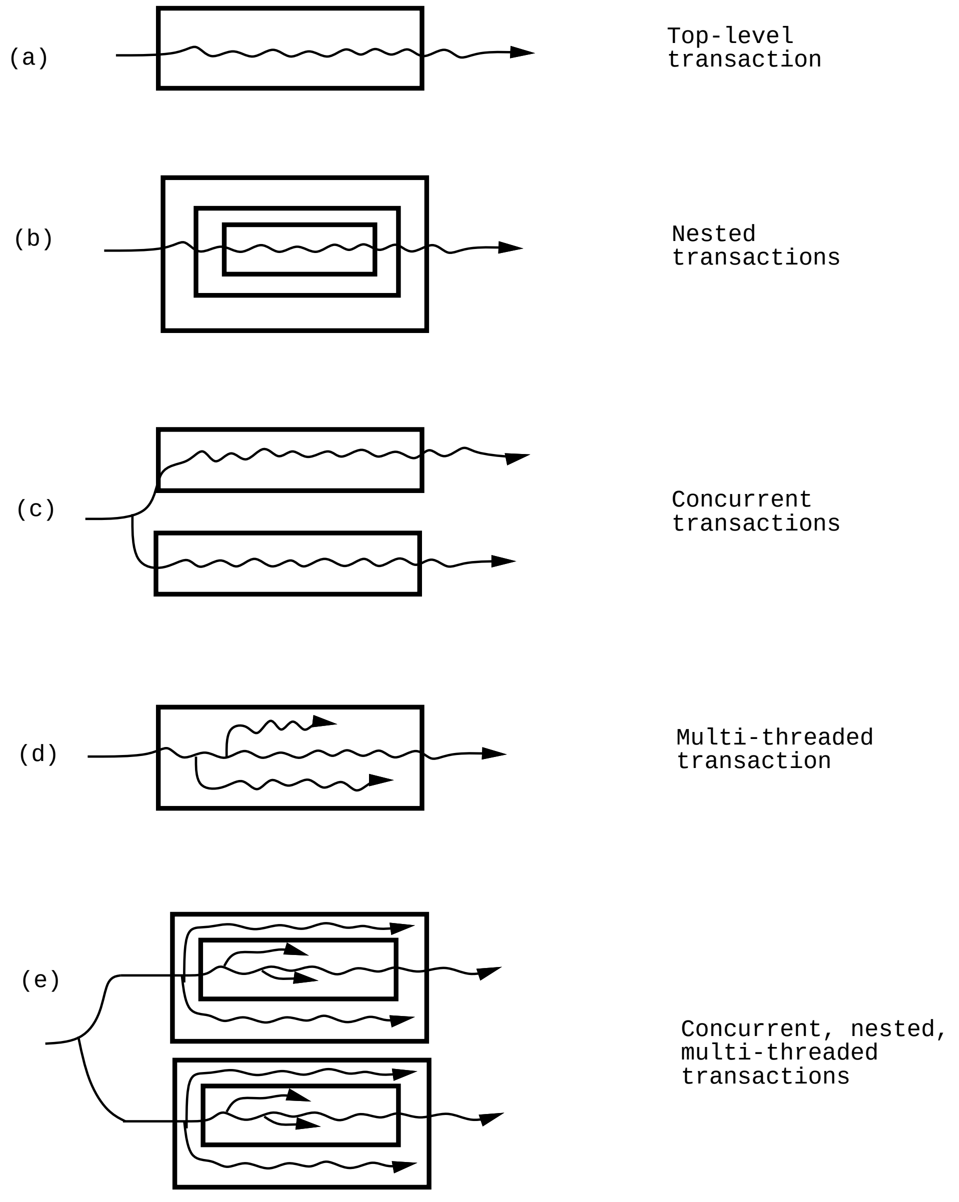

Fig. 2. Nesting, Concurrency, and Multi-Threading

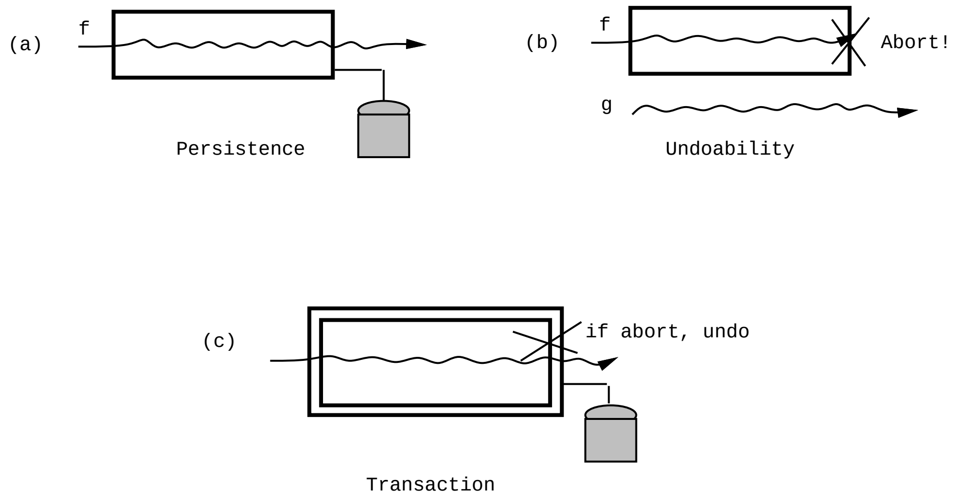

Fig. 3. Persistence and Undoability

tinguished location, the persistent root, are persistent. Figure 3a depicts the execution of a function `f` in a top-level persist-only transaction; when it terminates, all persistent data modified by the transaction are atomically saved to stable storage. If a crash occurs during the execution of `f`, we recover the last committed state from stable storage. All data not reachable from the persistent root are lost. Conceptually, a crash aborts *all* top-level transactions (of any flavor) and terminates all threads, so there is no mechanism for a persist-*only* transaction to abort in isolation.

A variety of approaches can be taken to guarantee that the effects of top-level transactions on stable storage are atomic. For example, our implementation makes the effects of nested transactions permanent only when the enclosing top-level transaction commits. This approach simplifies crash recovery but assumes that the number of modifications done by nested transactions is relatively small. An alternative approach would make nested transactions' effects permanent when they commit, but then crash recovery would have to undo such effects.

2.1.2 *Undoability.* A top-level undo-only transaction has no special effect if it commits. If it aborts then *all* changes it made to the store are undone. Our semantics for undo differs from traditional transactional systems because changes to volatile data are undone in addition to changes to persistent data. Figure 3b depicts the execution of a function `f` whose effects may possibly be undone. At the start (conceptually) a checkpoint of the store is made. If it terminates successfully, then nothing unusual happens; if not, then `f`'s effects are rolled back to the checkpointed state, at which point a possibly different computation `g` can begin.

Undo-only transactions may commit or abort regardless of whether they are nested. However, since a nested transaction's commit is relative to the action of its parent, if the parent aborts then the effects of the (committed) nested transaction must be undone along with the parent's other changes. Thus, when a child transaction commits it hands back ("anti-inherits") to its parent its set of changes to the store.

2.1.3 *Threads.* Threads are lightweight processes that communicate using shared mutable data and synchronize by acquiring and releasing mutual exclusion (mutex) locks. Individual threads may fork and start other computations, thereby providing a way to begin concurrent, nested transactions.

We do not require threads within a transaction to be serializable; thus, they can engage in two-way communication using shared mutable data. Otherwise, we could not wrap `transact` around existing multi-threaded code without modification.

2.1.4 *Locks.* Two-phase reader/writer (R/W) locks are a well-known mechanism for ensuring serializability. Alone, they provide no support for commit or abort. A transaction acquires a R/W lock and holds it until the transaction commits or aborts, thereby avoiding the problem of "cascading aborts". Write locks guarantee that any two concurrent transactions modify disjoint sets of data in the store, unless one is a descendant of the other.

Under Moss's standard locking rules for nested transactions [Moss 1985], transactions acquire locks subject to the following rules:

—A transaction may acquire a read lock if all writers are ancestors of the transaction.

—A transaction may acquire a write lock if all readers *and* writers are ancestors of the transaction.

—When a transaction commits, all its locks are *anti-inherited*, i.e., handed off to the parent or released if the transaction is top-level. If the transaction aborts, all its locks are released.

In our model, however, a parent transaction may run concurrently with its children, so we use a variation of these rules in which we must check that the read (or write) condition holds not only when a lock is acquired, but also every time the transaction reads (or writes) the associated data object. This check is reasonable for SML programs, which use mutable data infrequently, in contrast to imperative languages such as C.

2.2 Putting the Pieces Together

Nesting enables us to construct a top-level regular transaction from an undo-only transaction nested inside a top-level persist-only transaction (Figure 3c). If the undo-only transaction commits then all changes to the stable store are saved by the persist-only transaction. If the undo-only transaction aborts, all changes are rolled back. Thus when the persist-only transaction saves all changes to the stable store, there will be no changes on behalf of the aborted undo-only transaction to save; the net effect is that the stable store is in the same state as at the beginning of the transaction.

More generally, each combination of the different kinds of transactions has a well-defined meaning. For example, an undo-only transaction can have a persist-only transaction nested within it, and vice versa. A transaction can have nested within it concurrent transactions of different flavors.

To support complete "mixing-and-matching" of features, however, we impose two rules, one to deal with concurrency and one to deal with arbitrary nesting:

(1) All accesses to mutable data shared among concurrent transactions (of any flavor) must be coordinated by R/W locks.
(2) Modifications to the persistent store will survive a crash only if the transaction containing the modifications and all its ancestors commit.

The first rule is needed to guarantee the serializability (i.e., non-interference) of transactions. We enforce the rule by checking on each access of a mutable object that the appropriate read or write condition holds.

The second rule gives programmers consistency guarantees regarding the state of stable storage. Our implementation uses this rule to justify delaying writes to stable storage until the commits of top-level transactions. Delaying writes is reasonable for our application domain (short-lived, small transactions) and avoids rolling back partially completed transactions.

Finally, programmers using threads outside of any transaction should not expect strong consistency guarantees; otherwise they should use transactions. Such threads have no interaction with the undo mechanism; their effects cannot be undone. Such threads may modify the persistent store, but since they do so outside of a persistent transaction, programmers cannot expect these changes to be immediately reflected in stable storage. We choose to write such changes to stable storage whenever a top-level persistent transaction completes; we must do such writes at these times because the committing transaction may have depended on the value of persistent data modified by the thread. Other transactional facilities that allow threads to exist outside transactions, e.g., Camelot [Eppinger, Mummert, and Spector 1991] and Encina [Dixon 1993], have similar caveats.

3. EXPRESSING OUR DESIGN IN SML

We are able to express our design in a simple, straightforward, and elegant manner in SML. In the next three sections we first individually describe the SML interfaces for the four transactional building blocks, then show how we put them all together, and finally show how we can use our constructs to implement the bank example. Implementation details for persistence are discussed in greater detail by Nettles and Wing [1992]; for undoability, by Nettles and Wing [1992] and Morrisett [1993]; and for threads in SML, by Cooper and Morrisett [1990]. In Figures 4–7 we show only the portions of the `PERS`, `UNDO`, `RW_LOCK`, and `THREADS` interfaces that are relevant to this paper. The Venari/ML technical report gives further details of these interfaces and examples showing their use [Wing, Faehndrich, Haines, Kietzke, Kindred, Morrisett, and Nettles 1993].

3.1 The Pieces

3.1.1 *Persistence.* The key higher-order function exported by `PERS` is `persist`.[1] The expression `(persist f) a` has the effect of evaluating `f a`. If it is the outermost call of `persist` and `f a` terminates, `f`'s changes to persistent data are saved to disk. If `f` does not terminate, e.g., a crash occurs during its execution, none of `f`'s changes are saved.

[1] We use the underscore character, as in `'_a` and `'_b`, for weak (imperative) type variables [Milner, Tofte, and Harper 1990].

```
signature PERS = sig
    val persist  : ('a -> '_b) -> 'a -> '_b

    val bind     : identifier * 'a -> unit
    val unbind   : identifier -> unit
    val retrieve : identifier -> 'a
    ...
end
```

Fig. 4. PERS Interface

All data reachable from the persistent root are persistent, and thus, recoverable. Any SML value can be made persistent simply by arranging that it be reachable from the persistent root. The other functions of `PERS` allow manipulation of a symbol table that stores bindings between `identifiers` and values; the table itself is reachable. Thus, we can store and retrieve persistent values by name.

Our implementation uses a separate persistent heap to store all values reachable from the persistent root. Modifications to these values may cause values in the volatile heap to become reachable as well. On commit, any newly reachable values must be moved into the persistent heap, and all modifications to persistent values must be written to stable storage. We use the Recoverable Virtual Memory system [Satyanarayanan, Mashburn, Kumar, Steere, and Kistler 1993] to provide an efficient implementation of stable storage based on logging.

```
signature UNDO = sig
    val undoably : ('a -> '_b) -> 'a -> '_b
    exception Restore of exn
    ...
end
```

Fig. 5. UNDO Interface

3.1.2 *Undoability.* `UNDO` exports the `undoably` function, which allows users to make undoable changes to the store, an essential feature of a transaction that may abort. The `undoably` function is a wrapper for any function `f` such that if the exception `Restore` is raised while executing `f`, all of `f`'s effects on the store are undone; `undoably f` behaves exactly like `f` if no exception is raised. The changes undone include those done within any nested transactions.

The semantics of `undoably` is defined only with respect to the store. In particular, a transaction's effects through I/O (e.g., writing to a terminal) are undefined.

We implement undo by logging the location and old value of every mutation. Upon abort we replay the log in reverse order to restore the old values. To anti-inherit changes to the store we splice the child transaction's log onto the parent's log.

In most imperative languages this implementation would have unacceptable performance. In SML/NJ it works well for several reasons. First, assignments are relative rare. Second, the locations of many assignments are already logged to support generational garbage collection [Lieberman and Hewitt 1983; Ungar 1984]. We have simply extended these logs to capture all assignments and to record old values.

Our implementation for both persistence and undoability assumes that concurrent transactions modify disjoint sets of data in the store; this assumption is easily discharged by our first rule (Section 2.2) that concurrent transactions use write locks for accessing data.

3.1.3 *R/W Locks and Safe State*

```
signature RW_LOCK = sig
    eqtype rw_lock

    val create_rw_lock: unit -> rw_lock
    val acquire_read  : rw_lock -> unit
    val acquire_write : rw_lock -> unit
    ...
end
```

Fig. 6. RW_LOCK Interface

A. R/W Locks. We provide R/W locks to enable the programmer to enforce isolation and serializability among concurrent transactions. These locks are associated with mutable objects (see below) and are held per transaction.

A lock is created by a call to `create_rw_lock`. It is acquired for reading or writing by a call to `acquire_read` or `acquire_write` respectively. A thread within a transaction can perform reads and writes on the data protected by a lock, subject to our variation of Moss's rules stated in Section 2.1.4. When a transaction commits, all R/W locks are anti-inherited to the parent transaction (if any), or are released if the transaction is top-level. If a transaction aborts, all locks are released.

B. Safe State

```
signature RW_REF = sig
    type 'a rw_ref
    type rw_lock

    exception Read_Not_Held
    exception Write_Not_Held

    val create_rw_ref  : '_a * rw_lock -> '_a rw_ref
    val rw_get         : 'a rw_ref -> 'a
    val rw_set         : 'a rw_ref -> 'a -> unit
    val lock_of        : 'a rw_ref -> rw_lock
    ...
end
```

The only mutable data types in SML are refs and arrays. Thus, it is easy to provide two structures (the one for refs is shown above) that ensure that a mutable object is only accessed *safely* (i.e., when the appropriate locks are held) [Tolmach and Appel 1991; Morrisett and Tolmach 1993]. *Reader-writer refs* (`RW_REF`) are refs protected by R/W locks; in order for a transaction to access these objects, it must hold the `rw_lock` (for reading or writing, as appropriate).

The above `RW_REF` signature subsumes the SML pervasive `REF` signature. The accessing functions (`rw_get`, `rw_set`) verify that the appropriate read or write conditions hold according to our variation on Moss's locking rules (see Section 2.1.4). If the lock is not held in the appropriate mode, the `Read_Not_Held` or `Write_Not_Held` exception is raised. The `lock_of` function returns the lock associated with a `rw_ref`.

3.1.4 *Threads and Skeins*

```
signature THREADS = sig
    val fork    : (unit -> unit) -> unit

    type mutex
    val create_mutex   : unit -> mutex
    val acquire        : mutex -> unit
    val release        : mutex -> unit
    ...
end
```

Fig. 7. THREADS Interface

A. Threads. The `THREADS` module exports essential functions for creating a thread, and for acquiring and releasing mutex locks. Other functions, not relevant here, support manipulating condition variables and thread state. Our interface is similar to other threads packages for C [Cooper and Draves 1988], Modula-2+ [Rovner 1986], and Modula-3 [Harbison 1992].

The function `create_mutex` creates a new mutex value. The function `acquire` attempts to lock a mutex and blocks the calling thread until it succeeds. At most one thread may hold a given mutex at any time. The function `release` unlocks a mutex, giving other threads a chance to acquire it. Unlike R/W locks, mutex locks are short-term, i.e., they are not held for the duration of a transaction. Programmers have complete control over when to release them. Furthermore, R/W locks are used to coordinate *transactions*, while mutex locks are used to coordinate *threads* within a transaction.

B. Skeins. We introduce a new abstraction, called a *skein*[2], for encapsulating the difficult *control* aspects of a transaction. Conceptually a skein is a "generic" transaction and implements each of the boxes drawn in Figures 2 and 3. Within a skein some SML function (the *body*) is executed. The body itself may fork threads. We assume a barrier synchronization model for skein termination. The skein will

[2] A skein is a collection of threads.

not finish until the body thread returns a value and all other threads have finished; only one thread ever leaves a skein. In this respect, a skein is similar to a Qlisp "heavyweight future" [Goldman, Gabriel, and Sexton 1989]. All held mutexes must be released before return. If any thread (including the body thread) running inside a skein raises an uncaught exception, the skein aborts. Any extant forked threads and child skeins are killed, and the exception is propagated to the outside. A skein also holds R/W locks that are shared among its threads.

A transaction might need to execute certain code within a skein (while the R/W locks are still held), but *after* all threads within that skein have completed or died. Such code might, for example, commit persistent changes to disk or release R/W locks. Thus, our skein abstraction has the following interface:

```
signature SKEIN = sig
    datatype 'a result = Result of 'a
                       | Exception of exn
    exception Abort

    val skein:
        (unit -> unit) ->                   (* initializer *)
        ('_b result -> '_b result) ->       (* completer *)
        ('a -> '_b) ->                      (* body *)
         'a -> '_b                          (* result *)
end
```

The body of a skein is executed in a sub-thread within the skein, while a *control thread* waits for it to complete. The first two arguments to **skein** are (1) an initializer function, which is called in the control thread before the body thread is forked; and (2) a completer function, which is called in the control thread after the body has returned and any extant threads have ended. The completer is applied to the value returned by the body or the exception that caused premature termination, and it returns a **result** value that is in turn presented as the result of the call to **skein**.

If the body of a skein finishes while sub-skeins are still executing, the sub-skeins are terminated, calling their completer functions with the **Abort** exception. The parent skein's completing function is not called until all sub-skeins have completed.

We use skeins to implement multi-threaded transactions of all kinds, e.g., persist-only and undo-only transactions, by passing in appropriate initializer and completer functions.

3.2 Putting It All Together

Putting all these pieces together into a single SML module culminates in our main **VENARI** interface, shown in Figure 8. It provides a way for application programmers to create and manipulate concurrent, nested, multi-threaded transactions. A transaction is a *locking skein* of *threads* whose effects are *undone* if the transaction aborts or made *persistent* if it terminates normally. We require that each transaction access only safe state.

The various features described in the previous sections are all used in **VENARI**'s main function, **transact**, which evaluates its argument within a transaction. We implement **transact** as a special case of a skein, reusing the initializer and com-

```
signature VENARI = sig
    val transact : ('a -> '_b) -> 'a -> '_b

    structure Pers     : PERS
    structure Undo     : UNDO
    structure RW_Lock  : RW_LOCK
    structure RW_Ref   : RW_REF
    structure RW_Array : RW_ARRAY
    structure Skein    : SKEIN
    structure Threads  : THREADS
end
```

Fig. 8. Transaction Interface

pleter functions defined for persist-only and undo-only transactions:

```
val init_transact = Undo.init_undo o Pers.init_pers

val complete_transact = Pers.complete_pers o Undo.complete_undo

val transact = Skein.skein init_transact complete_transact
```

3.3 Implementation of the Bank Example

We give an implementation of a bank account in Figure 9. `L` is the `Venari.RW_Lock` substructure, `R` is `Venari.RW_Ref`, and `V` is `Venari`.

The `account` is a ref to a real (initially 0.0), protected by a R/W lock. Assuming that `amount` is non-negative, the `deposit` function first acquires the lock associated with the account in write mode; it then updates the account's value to the sum of the old value and the new amount. The `withdraw` function is slightly more complicated since it needs to check whether there is sufficient money in the account before the withdrawal occurs. Raising the unhandled `Insufficient_Funds` exception would cause the transaction to abort.

Using this interface we can implement a bank transfer as described in Section 1.1.

4. EVALUATION

In the introduction we stated two goals of our work: factoring transactions into individual features and composing these features with each other and with other features of SML. For the most part, we succeeded in accomplishing both goals and are able to express our results concretely through our Venari/ML interfaces. In this section we evaluate the successes and limitations of our work. We also compare Venari/ML to Avalon/C++.

We achieve composability by making transactions with higher-order functions. Making `transact` higher-order means that `transact` can easily be used as a wrapper function. This kind of composability facilitates code reuse. For example, suppose we have an interface along with a non-transactional implementation. We can implement a transactional version of this interface by wrapping a `transact` around each of the non-transactional functions without any knowledge of their internal structure.

```
functor Account (structure Venari: VENARI)
      : ACCOUNT = struct

    type account = real R.rw_ref

    fun new_account () =
        R.rw_ref (0.0, L.create_rw_lock())

    fun deposit account amount =
        let fun do_deposit () =
            (L.acquire_write (R.lock_of account);
             R.rw_set account ((R.rw_get account) + amount))
        in
            V.transact do_deposit ()
        end

    exception Insufficient_Funds

    fun withdraw account amount =
        let fun do_withdraw () =
            (L.acquire_write (R.lock_of account);
             let val bal = (R.rw_get account)
             in
                 if bal < amount then
                     raise Insufficient_Funds
                 else
                     R.rw_set account (bal - amount)
             end)
        in
            V.transact do_withdraw ()
        end

end
```

Fig. 9. Bank Account Implementation

The one feature of the New Jersey implementation of SML with which our transactional extensions do not interact well is first-class continuations. We use continuations extensively in our implementation of threads. However, we cannot export continuations directly to the user because they do not interact well with SML's exception handling, which we use to deal with aborted transactions in a graceful manner. Unfortunately, when a continuation is invoked, a new exception handler context is installed. Consequently, we cannot guarantee that a computation will pass through our handlers. For example, if continuations were exported to the user, we could not guarantee that a skein's completer function would be called. A solution to this problem would be to implement a mechanism similar to Scheme's `unwind-protect` [Friedman and Haynes 1985; Rees 1992].

We also successfully achieved a factorization of transactions into their component parts. We found that to allow transactions of any type to execute concurrently requires the use of R/W locks. That support for concurrency needs support for synchronization should come as little surprise. More surprising was that we were successful at decoupling the other three features from each other.

One advantage of decoupling transaction features is that each feature can be used independently. For example, undoability is useful for implementing backtracking search. Typical Prolog implementations use an explicit mutation log, called a *trail*, which is used to undo variable bindings when backtracking [Warren 1983]. Undo would allow the elimination of the trail and allow the desired functionality to be expressed directly using `undoably`.

Another benefit of this factorization is that it helped us recognize new abstractions. In our implementation of undo, the ability to save and restore the state of the store is implicit and governed by rules about nesting transactions. Inspired by this work and the semantics of imperative programming languages, Morrisett [1993] has recently proposed and implemented a new programming language feature, *refined first-class stores*. In his system the current store can be captured and saved away like any other first-class value. At any later point during the program's execution the saved store can be restored.

A final benefit to factoring our design has been in factoring our implementation. Our original implementation of the persistence and undo subsystems was factored largely for convenience of implementation. At the same time, our original threads implementation was built completely independently of the transaction system. Surprisingly, adding support for concurrent, multi-threaded transactions has not forced these implementations to merge and become monolithic. Instead the mutation log serves as a common data structure used independently by the undo and persistence subsystems, and is maintained on a per transaction basis. Needless to say the factored nature of the implementation has made it easier to build and maintain.

We have not yet attempted to design and implement support for distributed transactions. If we were to attempt such support, our factored implementation as well as the notion of first-class stores mentioned above would be useful. Committing a distributed transaction requires a two-phase protocol. In the first phase the current state of the transaction must be made persistent in such a way that it can be undone. We can achieve this effect by capturing the store as a first-class value and then making that value persistent. Given support for some kind of distribution, adding distributed transactions should be straightforward.

We deliberately chose not to explore support for other ways to ensure serializability, since this issue has been thoroughly addressed by the database community. Also, we intentionally avoided the hard problems of undoing I/O (as in undoing the dispensing of cash from an ATM machine [Pausch 1988].)

We have also made significant progress in measuring and improving the performance of our system. Recently O'Toole, Nettles, and Gifford [1993] added a concurrent garbage collector for the persistent heap. They show that the performance of both the collector and the persistence subsystem is good—comparable to a simpler system that supports neither orthogonal persistence nor garbage collection. Nettles [1994] is currently completing a more thorough performance evaluation that will allow us to improve the performance of our system substantially.

Avalon/C++ [Detlefs, Herlihy, and Wing 1988] superficially shares part of the factorization of transactional concepts with Venari/ML. Avalon/C++'s `recoverable`, `atomic`, and `subatomic` classes provide functions similar to those provided in Venari/ML's persistence, R/W locks, and threads modules. Avalon/C++ exploits C++'s class inheritance mechanism to achieve composition of features. For example, one way to define an atomic array class would be to inherit from class `atomic`, to use a regular C++ array to represent an atomic array, and to implement atomic array operations in terms of regular C++ array operations plus R/W lock operations inherited from `atomic`. Composition of features in terms of higher-order functions as done in Venari/ML would not be possible in Avalon/C++ because functions cannot be returned as results in C++.

Unlike Venari/ML, Avalon/C++ does not separate the undoability feature from transactions at all, does not separate persistence from R/W or mutex locks, and does not support multi-threaded transactions. Avalon/C++, however, does support distributed transactions. Finally, Avalon/C++ is implemented in a completely different runtime environment, i.e., Camelot; Venari/ML uses RVM, which provides only a small subset of Camelot's functionality.

5. SUMMARY OF CONTRIBUTIONS

The main contribution of our work is to show that transactions can be broken into separable components, each supporting a different aspect of a traditional transactional model: persistence, undoability, locking, and threads. These components can then be composed to build the traditional model or even new models with weaker semantics.

Two technical ideas resulted from pushing hard to achieve our goal of composability. One is the idea of a general-purpose control abstraction, the skein, with which we can build variations of the transactional model as simple special cases. The other is a set of guarantees, captured by our variation of Moss's rules, that gives a reasonable semantics to nested, multi-threaded transactions. Heretofore other systems either permit only a single thread of control to execute with a transaction [Liskov and Scheifler 1983; Detlefs, Herlihy, and Wing 1988] or support multi-threaded transactions with no semantic guarantees [Eppinger, Mummert, and Spector 1991; Dixon 1993]. Except for Humm [1993] we are not aware of any other work that attempts to give nested, multi-threaded transactions such guarantees.

A more concrete contribution is our specific set of extensions to SML/NJ in support of concurrent, nested, multi-threaded transactions. In our design, we exploited

SML's higher-order functions and modules facility. We use its exception handling mechanism to give control to the programmer in case a transaction aborts. Our implementation uses the New Jersey implementation of SML in some critical ways, e.g., its support for continuations and the logs used by its garbage collector. Our current implementation is based on SML/NJ (0.80) and runs in the Mach 2.5 environment.

By adding such extensions to an advanced programming language like SML, we have provided application programmers with some high-level constructs (above the operating-system level) to use transactions unintrusively. By using simple wrapper functions, programmers need not worry about formatting and unformatting data into files in order to achieve persistence; they can undo effects to the store if desired (e.g., for backtracking); and they have explicit control over concurrent access to shared mutable data through mutex and R/W locks.

Acknowledgments

We thank Manuel Faehndrich for his participation in the initial design and implementation of concurrent, multi-threaded transactions, and Karen Kietzke for her help in maintaining our current implementation.

Views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, by the U. S. Government.

REFERENCES

Atkinson, M., Bailey, P., Chisolm, K., Cockshott, W., and Morrison, R. 1983. An approach to persistent programming. *Computer Journal 26,* 4, 360–365.

Cooper, E. and Draves, R. 1988. C Threads. Technical Report CMU-CS-88-154, Carnegie Mellon Computer Science Dept.

Cooper, E. and Morrisett, J. G. 1990. Adding threads to Standard ML. Technical Report CMU-CS-90-186 (Dec.), Carnegie Mellon School of Computer Science.

Detlefs, D. L., Herlihy, M. P., and Wing, J. M. 1988. Inheritance of synchronization and recovery properties in Avalon/C++. *IEEE Computer*, 57–69.

Dixon, G. 1993. Encina (Transarc Corporation). Private communication.

Eppinger, J., Mummert, L., and Spector, A. 1991. *Camelot and Avalon: A Distributed Transaction Facility.* Morgan Kaufmann Publishers, Inc.

Friedman and Haynes 1985. Constraining control. In *Proceedings of ACM Symp. on Prog. Lang.*

Goldman, R., Gabriel, R., and Sexton, C. 1989. Qlisp: An interim report. In *US/Japan Workshop on Parallel Lisp Proceedings*, pp. 161–181.

Harbison, S. 1992. *Modula-3.* Prentice-Hall.

Haskin, R., Malachi, Y., Sawdon, W., and Chan, G. 1988. Recovery management in QuickSilver. *ACM Transactions on Computer Systems 6,* 1 (February), 82–108.

Helland, P. 1985. Transaction monitoring facility. *Database Engineering 8,* 1 (June), 9–18.

Humm, B. G. 1993. An extended scheduling mechanism for nested transactions. In *Int'l IEEE Workshop on Object-Orientation in Operating Systems*, Ashville, NC.

Lieberman, H. and Hewitt, C. 1983. A real-time garbage collector based on the lifetimes of objects. *Commun. ACM 26*, 419–429.

Lindsay, B., Haas, L., Mohan, C., Wilms, P., and Yost, R. 1984. Computation and communication in R*: A distributed database manager. *ACM Trans. on Comp. Systems 2,* 1 (Feb.), 24–38.

Liskov, B. and Scheifler, R. 1983. Guardians and actions: Linguistic support for robust, distributed programs. *ACM Transactions on Programming Language and Systems 5,* 3 (July), 382–404.

Milner, R., Tofte, M., and Harper, R. 1990. *The Definition of Standard ML.* The MIT Press.

Morrisett, J. 1993. Generalizing first-class stores. In *Proceedings of the ACM SIGPLAN Workshop on State in Programming Languages*, pp. 73–87.

Morrisett, J. G. and Tolmach, A. 1993. Procs and locks: A portable multiprocessing platform for Standard ML of New Jersey. In *Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, San Diego*, pp. 198–207.

Moss, J. 1985. *Nested Transactions: An Approach to Reliable Distributed Computing.* The MIT Press.

Nettles, S. M. 1994. Safe and efficient transaction support for heap-based languages. Ph.D thesis in progress.

Nettles, S. M. and Wing, J. 1992. Persistence + Undoability = Transactions. In *Proc. of HICSS-25.* Also CMU-CS-91-173, August 1991.

O'Toole, J., Nettles, S., and D.K.Gifford 1993. Concurrent compacting garbage collection of a persistent heap. In *Proceedings of the 14th ACM Symp. on Operating System Principles.*

Pausch, R. 1988. Adding input and output to the transactional model. Ph. D. thesis, Carnegie Mellon University, Computer Science Department. Published as Technical Report CMU–CS–88–171.

Rees, J. 1992. The Scheme of things: The June 1992 meeting. *LISP Pointers 5,* 4, 40–45.

Rovner, P. 1986. Extending Modula-2 to build large, integrated systems. *IEEE Software 3,* 6, 46–57.

Satyanarayanan, M., Mashburn, H. H., Kumar, P., Steere, D. C., and Kistler, J. J. 1993. Lightweight Recoverable Virtual Memory. In *Proceedings of the 14th ACM Symposium on Operating Systems Principles.*

Shrivastava, S., Dixon, G., Hedayati, F., Parrington, G., and Wheater, S. 1988. A technical overview of Arjuna: A system for reliable distributed computing. Technical report, Computing Laboratory, University of Newcastle upon Tyne.

Tolmach, A. P. and Appel, A. W. 1991. Debuggable concurrency extensions for Standard ML. In *Proc. ACM/ONR Workshop on Parallel and Distributed Debugging*, pp. 120–131. Published as *SIGPLAN Notices* 26(12), December 1991. Also Princeton Univ. Dept. of Computer Science Tech. Rep. CS-TR-352-91.

Ungar, D. 1984. Generational scavenging: A non-disruptive high performance storage management reclamation algorithm. In *ACM SIGPLAN Software Engineering Symposium on Practical Software Development Environments*, Pittsburgh, Pennsylvania, pp. 15–167.

Warren, D. 1983. An abstract PROLOG instruction set. Technical Report 309, AI Center, SRI International, Menlo Park.

Wing, J., Faehndrich, M., Haines, N., Kietzke, K., Kindred, D., Morrisett, J., and Nettles, S. 1993. Venari/ML interfaces and examples. Technical Report CMU-CS-93-123 (March), Carnegie Mellon Computer Science Department.

Received February 1994; revised April 1994; accepted April 1994.