

# Transactional Memory for Smalltalk<sup>\*</sup>

Lukas Renggli and Oscar Nierstrasz

Software Composition Group, University of Bern  
{renggli,oscar}@iam.unibe.ch  
scg.unibe.ch

**Abstract.** Concurrency control in Smalltalk is based on locks and is therefore notoriously difficult to use. Even though some implementations provide high-level constructs, these add complexity and potentially hard-to-detect bugs to the application. Transactional memory is an attractive mechanism that does not have the drawbacks of locks, however the underlying implementation is often difficult to integrate into an existing language. In this paper we show how we have introduced transactional semantics in Smalltalk by using the reflective facilities of the language. Our approach is based on method annotations, incremental parse tree transformations and an optimistic commit protocol. We report on a practical case study, benchmarks and further and on-going work.

**Keywords.** Transactional Memory, Concurrent Programming, Language Constructs and Features

## 1 The need for transactions

Smalltalk has inherently weak support for concurrent programming and synchronization. Smalltalk-80 [1] only proposes semaphores as a means for synchronizing processes and achieving mutual exclusion. The ANSI standard of Smalltalk [2] does not refer to synchronization at all.

Only a few current Smalltalk implementations provide more sophisticated synchronization support. VisualWorks Smalltalk provides a reentrant lock implementation that enables a process to enter the same critical section multiple times. Other processes get blocked until the owning process leaves the critical section. Unfortunately this kind of synchronization has its drawbacks and is notoriously difficult to use [3]:

**Deadlocks.** If there are cyclic dependencies between resources and processes, applications may deadlock. This problem can be avoided by acquiring resources in a fixed order, however in practice this is often difficult to achieve.

**Starvation.** A process that never leaves a critical section, due to a bug in the software or an unforeseen error, will continue to hold the lock forever. Other processes that would like to enter the critical section starve.

---

<sup>\*</sup> Proceedings of the 2007 International Conference on Dynamic Languages (ICDL 2007), ACM Digital Library, pp. 207-221. DOI 10.1145/1352678.1352692

**Priority Inversion.** Usually schedulers guarantee that processes receive CPU time according to their priority. However, if a low priority thread is within a critical section when a high priority process would like to enter, the high priority thread must wait.

Squeak Smalltalk [4] includes an implementation of monitors [5], a common approach to synchronize the use of shared data among different processes. In contrast to mutual exclusion with the reentrant lock, with monitors, a process can wait inside its critical section for other resources while temporarily releasing the monitor. Although this avoids deadlock situations, the use of monitors is difficult and often requires additional code to specify guard conditions. Moreover, if the process is preempted while holding the monitor, everybody else is blocked. Beginners are often overwhelmed by the complexity of using monitors as Squeak does not offer method synchronization as found in Java.

Transactional memory [6, 7] provides a convenient way to access shared memory by concurrent processes, without the pitfalls of locks and the complexity of monitors. Transactional memory allows developers to declare that certain parts of the code should run atomically: this means the code is either executed as a whole or has no effect. Moreover transactions run in *isolation*, which means they do not affect and are not affected by changes going on in the system at the same time. Upon *commit* the changes of a transaction are applied atomically and become visible to other processes. Optimistic transactions do not lock anything, but rather conflicts are detected upon commit and either lead to an *abort* or *retry* of the transaction.

Most relational and object databases available in Smalltalk provide database transactions following the ACID properties: Atomicity, consistency, isolation, and durability. However, they all provide this functionality for persistent objects only, not as a general construct for concurrent programming. These implementations often rely on external implementations of transactional semantics. GemStone Smalltalk [8] is a commercially available object database, that directly runs Smalltalk code. As such, GemStone provides transactional semantics on the VM level. Guerraoui *et al.* [9] developed GARF, a Smalltalk framework for distributed shared object environments. Their focus is not on local concurrency control, but on distributed object synchronization and message passing. They state that “*A transactional mechanism should however be integrated within group communication to support multi-server request atomicity.*” [10]. Jean-Pier Briot proposed Actalk [11], an environment where Actors communicate concurrently with asynchronous message passing. The use of an Actor model is intrusive. It implies a shift of the programming paradigm, where there is no global state and therefore no concurrency issues.

In this paper we present an implementation of transactions in Squeak based on source-code transformation. In this way most code is free of concurrency annotations, and transactional code is automatically generated only in the contexts where it is actually needed.

The specific contributions of this paper are:

- The implementation of transactional semantics in Smalltalk, using the reflective capabilities of the language.
- A mechanism to specify context-dependent code using method annotations, for example to intercept the evaluation of primitive methods.
- Incremental, on-the-fly parse tree transformation for different execution contexts.
- Efficient, context-dependent code execution using the default execution mechanisms of the VM.

Section 2 presents some basic usage patterns of our implementation. Section 3 shows the implementation of transactions in Squeak without modifying the underlying VM. Section 4 validates our approach by running a collection of benchmarks and by applying the concept to a real world application. Section 5 concludes this article with some remarks about ongoing and future work.

## 2 Programming with transactions

Transactions offer an intuitively simple mechanism for synchronization concurrent actions. They do not require users to declare specific locks or guard conditions that have to be fulfilled. Moreover transactions can be used without prior knowledge of the specific objects that might be modified. Transactions are global, yet multiple transactions can run in parallel. The commit protocol checks for conflicts and makes the changes visible to other processes atomically.

<pre>tree := BTree new. lock := Semaphore forMutualExclusion.  " writing " lock critical: [ tree at: #a put: 1 ].  " reading " lock critical: [ tree at: #a ].</pre>	<pre>tree := BTree new.  " writing " [ tree at: #a put: 1 ] atomic.  " reading " tree at: #a.</pre>
--	---

**Fig. 1.** Lock-based vs. Transactional accesses of a shared data structure.

On the left side of Figure 1 we see the traditional way of using a semaphore to ensure mutual exclusion on a tree data structure. The key problem is that *all* read and write accesses to the tree must be guarded using the same lock to guarantee safety. A thread-safe tree must be fully protected in all of its public methods. Furthermore, we cannot easily have a second, unprotected interface to the same tree for use in a single-threaded context.

On the right side of Figure 1 we present the code that is needed to safely access the collection using a transaction: the write access is put into a block that tells the Smalltalk environment to execute its body within a transaction. The read access can happen without further concurrency control. As long as all write

accesses occur within the context of a transaction, read accesses are guaranteed to be safe. The optimistic commit protocol of the transaction guarantees safety by (i) ensuring that no write conflicts have occurred with respect to the previous saved state, and (ii) atomically updating the global object state.

To make the code using transactions as simple as possible we provide two methods for running code as part of a transaction. These methods are extensions to the standard Smalltalk library, and do not affect the language syntax or runtime.

- `#atomic` causes the receiving block closure to run as a new transaction. Upon termination of the block, any changes are committed atomically. If a conflict is detected, all modifications are cancelled and a commit conflict exception is raised.
- `#atomicIfConflict:` causes the receiving block to run as a new transaction. Instead of raising an exception if a conflict occurs, the block argument is evaluated. This enables developers to take a specific action, such as retrying the transaction or exploring the conflicting changes.

Further convenience methods can easily be built out of these two methods, for example a method to retry a transaction up to fixed number of times, or only to enter a transaction if a certain condition holds.

### 3 Inside transactions

We introduce transactions to Smalltalk without modifying the underlying Virtual Machine (VM). Our approach is based on earlier proposals in which source code is automatically and transparently transformed to access optimistic transactional object memory, rather than directly accessing objects [12, 13]. The key advantage of this approach is that most source code can be written without embedding any explicit concurrency control statements. Transactional code is automatically generated where it is needed. Furthermore, in contrast to the earlier approaches, we generate the needed transactional code dynamically where and when it is needed, and caching the generating code for future invocations.

In a nutshell, our approach works as follows:

- Every method in the system may be compiled to two versions: one to be executed in the normal execution context, and the other within a transactional context. Contrary to the other approaches we do this incrementally and on the fly using a compiler plugin.
- State access in transactional methods is automatically transformed to use an indirection through the transaction context.
- We use method annotations to control the automatic code transformation or to provide different code. Unlike earlier approaches, we take into account the use of primitives and exception of file-system access by providing alternative code to be used in a transactional context.

- When entering a transactional context we record the transaction (an object) in the current process (also an object).
- All objects touched during a transaction are atomically taken a snapshot of. Each snapshot consists of two copies of the original object: one that reflects the initial state and one that is altered during the transaction. For efficiency reasons immutable objects are excluded from snapshots.
- Upon commit we check for conflicts by atomically comparing the state of the object at the beginning of the transaction to the current version in memory. If no conflict is detected, the changes are committed. In case of a conflict the system is left in the state as it was before the transaction and an exception is raised that provides information for further reflection, namely all the changes, the conflicting changes and the transaction itself.

The key novelties of our approach lie in the use of annotations and reflection to lazily generate the transactional versions of methods, and the ability to provide alternative code to use in place of primitives during transactions.

In the following two sections we describe (1) the compilation to transactional code, and (2) the implementation of the transactional object model.

### 3.1 Compiling to transactional code

We transform methods by changing read and write accesses to make use of transactional object memory. Methods are transformed using a new version of the behavioral reflection framework Geppetto [14] which is based on sub-method reflection [15], allowing us to declaratively reify and transform an abstract syntax tree (AST) before compiling to byte-code.

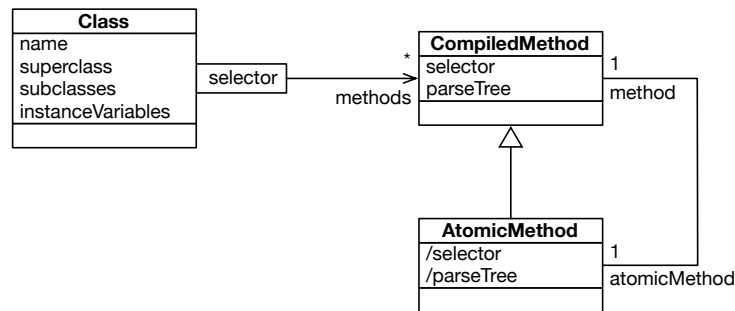


Fig. 2. Static Compilation Model

Whenever the source code of a method is accepted, as seen in Figure 2, our compiler plugin creates an additional compiled method that implements the behaviour to be used within the context of a transaction. The following basic transformations are performed:

1. Reading from instance variables and global variables is transformed to send the messages `#atomicInstVarAt:` or `#atomicValue` respectively. This allows us to implement these two messages to read the current value from within a transactional context instead of directly accessing the variables within the receiving object.
2. Writing to instance and global variables is transformed to send the messages `#atomicInstVarAt:put:` or `#atomicValue:` respectively. Again this allows us to intercept state access and handle it from within the current transaction.
3. Sending a message from inside a transactional method will actually send a different message name, namely we prepend `#__atomic__` to the original selector name.

<pre> <b>BTree&gt;&gt;at: aKey put: anObject</b>   leaf   leaf := root leafForKey: aKey.  leaf insertKey: aKey value: anObject. root := leaf root. ^ anObject </pre>	<pre> <b>BTree&gt;&gt;__atomic__at: aKey put: anObject</b>   leaf   leaf := (self atomicInstVarAt: 1) __atomic__leafForKey: aKey. leaf __atomic__insertKey: aKey value: anObject. self atomicInstVarAt: 1 put: leaf __atomic__root. ^ anObject </pre>
--	---

**Fig. 3.** Original vs. transformed code.

Having applied these three transformations to the code, the two compiled methods are stored in the method dictionary of their owning class. To tell the two methods apart, the atomic version of the method has `#__atomic__` prepended to its name. These methods are hidden, and are only called from generate code within an atomic context. Transactional methods are filtered from the code editors, so they are not visible to the developer and development tools but only to the VM. On the left side of Figure 3 we present the code of a method as the developer implemented it, whereas on the right side we show the same method as it is compiled for the atomic context. All message sends are prepended with `#__atomic__` and instance variable access is dispatched through special methods using the slot index as argument.

A transaction is created by sending the message `#atomic` to a block containing normal (non-transactional) Smalltalk code. The code within such a block is statically transformed to evaluate within a transactional context. We have seen an example for such a call in the introductory example in Figure 1. Methods that send `#atomic` are special, because the code outside this block is compiled normally, whereas we apply the transformations as described above to the inside of the block closure.

Squeak includes a few primitive methods that access and modify state. The most prominent of these are `#at:` and `#at:put:` to access the elements of variable-sized objects. Moreover there are also some central collection and stream methods that are implemented within the VM for efficiency. As primitive operations are written in C and statically compiled into the VM, we cannot use

Geppetto to modify state-access. The only possibility to reify these methods is to replace them with non-primitive methods.

Annotation	Source Code	Transform
(no annotation, default)	method body	yes
<atomic:>	argument	yes
<atomicUseUntransformed>	method body	no
<atomicUseUntransformed:>	argument	no

**Table 1.** Method annotations are used to control how the compiler transforms source code for the transactional context.

We make use of annotations to further control the way in which transactional code may be generated. Table 1 summarizes the effect of the following annotations:

- <atomicUseUntransformed> avoids doing any code transformation. This means the normal and the transactional method will be the same, so no transformation is needed. In the current implementation this is mostly used for exception handling, as this code should continue to work through the boundaries of transactions.
- <atomic:> uses the method identified as its argument as the source for the code transformation. We use this for primitives that are implemented for efficiency reasons only. For example the method `#replaceFrom:to:with:-startingAt:` in the class `Array` calls the primitive 105 and is used to copy elements from one collection to another one. With the method annotation we tell the compiler that it should instead transform and install the method `#atomicReplaceFrom:to:with:startingAt:` that has the same behavior but is implemented in Smalltalk.

**Array>>replaceFrom: start to: stop with: replacement startingAt: repStart**

"Primitive. This destructively replaces elements from start to stop in the receiver starting at index, repStart, in the collection, replacement. Optional."

```
<primitive: 105>
<atomic: #atomicReplaceFrom:to:with:startingAt:>
super replaceFrom: start to: stop with: replacement startingAt: repStart
```

**Array>>atomicReplaceFrom: start to: stop with: replacement startingAt: repStart**

```
| index repOff |
repOff := repStart - start.
index := start - 1.
[ (index := index + 1) <= stop ]
  whileTrue: [ self at: index put: (replacement at: repOff + index) ]
```

- `<atomicUseUntransformed:>` uses the method identified as its argument as untransformed atomic code. We use this mainly in infrastructural code to dispatch primitive requests that access state to the working copy of the receiver. For example indexed slot access is handled through primitives in Squeak. The method `#at:` in the class `Object` calls the primitive 60 to fetch the contents of an indexed element. The method annotation tells the compiler to use `#atomicAt:` instead. This method delegates the request to the current working copy of the object.

**Object>>at: index**

"Primitive. Assumes receiver has indexed slots. Answer the value of an indexable element in the receiver. Fail if the argument index is not an Integer or is out of bounds."

```
<primitive: 60>
<atomicUseUntransformed: #atomicAt:>
self primitiveFail
```

**Object>>atomicAt: index**

```
^ self workingCopy at: index'
```

Compiling all the methods of the system is costly both in time and memory. Most methods available in the system are never called from within a transactional context and therefore do not need to be translated. The dynamic nature of Smalltalk makes it difficult to determine statically the required set of transactional methods, however it allows us to compile methods lazily when they are about to be executed. This produces a slowdown the first time a method is executed within a transactional context, but subsequent invocations are dispatched using the normal mechanisms of the VM and therefore run at full speed.

Most transactional systems prohibit system calls and filesystem access during transactions [16, 17]. Our approach allows replacement code to be specified for use within a transactional context. For example, when deleting a file the action is recorded with a custom change object and atomically applied together with the other changes upon successful commit of the transaction:

**FileDirectory>>deleteFile: aString**

```
<primitive: 'primitiveFileDelete' module: 'FilePlugin'>
<atomicUseUntransformed: #atomicDeleteFile:>
```

**FileDirectory>>atomicDeleteFile: aString**

```
Processor activeProcess currentTransaction
addChange: (CustomChange onApply: [ self deleteFile: aString ])
```

Our model also allows exceptions to be thrown and handled inside the transaction boundaries. An exception that leaves the boundaries of a transaction causes that transaction to abort and the exception to be re-raised in the non-transactional context.



### 3.2 Transactions at runtime

When entering a transaction we create a new transaction object and store it in an instance variable of the current process, as depicted in Figure 4. When leaving a transaction we set the current transaction reference back to `nil`. In this way we can efficiently determine the current transaction from anywhere in our application. Moreover we capture an escape continuation upon entry, to be able to abort the current transaction by doing a non-local jump to the calling context.

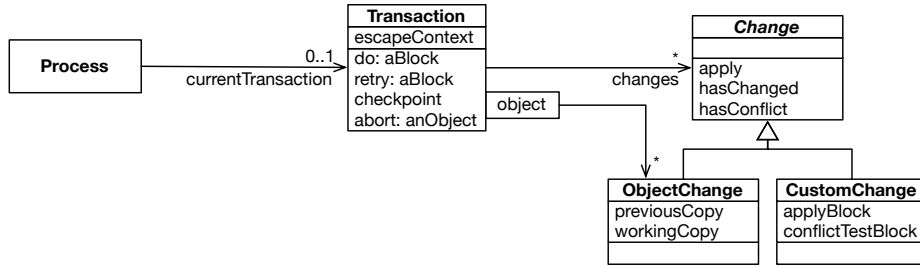


Fig. 4. Dynamic model of transactions at runtime

After having entered a transactional context, all the executed code is in its transformed form. This means that state access goes through special accessor methods and all message sends are redirected to their transactional counterparts. Like that transactional code execution works the same as normal code execution: it shares the same object memory but it uses a different access strategy to access state.

We adopt a conventional optimistic transaction protocol [18]. Whenever an object is touched within the context of a transaction for the first time (read or written), the transaction instantiates a new change object `ObjectChange`. This change object contains references to two copies of the object. The `previousCopy` contains an immutable copy for detecting conflicts. The `workingCopy` is a mutable copy of the object being used during the transaction. The change object knows if it has a conflict (the original object is not the same as the previous copy) and if it has changed (original object is not the same as the working copy).

We also provide a custom change object `CustomChange` that is used to record irreversible actions that should only be applied during the atomic commit phase if there are conflicts. We have seen the use of such a custom change in Section 3.1, where we presented a possible solution for file-deletion within a transactional context.

At the end of the transaction we have to acquire a form of “global lock” on the object memory to be able to check for conflicts and commit the changes. We use `#valueUnpreemptively` implemented on block closures to ensure that

no other process is running at the same time. As a first step we check if any of the changes we gathered during the transaction has a conflict and raise an exception if this is the case. Otherwise we copy the changes from the working copies to the original objects. The time required to hold the lock and to validate and apply the changes linearly depends on the number of objects involved in the transaction.

These are some important properties of our transactional model [19]:

**Repeatable read.** Reads are repeatable. Since data that is read within a transaction is copied, repeated reads from within a transaction are consistent. Changes outside the transaction are not visible after a first read.

**Optimistic write.** Our transactional memory writes optimistically [20]. The transaction boundaries are controlled by working on copies of the objects.

**Lazy version management.** We create copies of objects that are read and written within transactions. This requires an extra redirection for accessing the state and a considerable amount of memory and processing time for copying the involved objects. Aborting a transaction is cheap as no state has to be restored.

**Lazy conflict detection.** Assuming that conflicts are rare, conflicts are checked before committing data. This check happens atomically together with the commit.

**Lazy conflict resolution.** Conflicts are resolved by dropping (or retrying) the transaction that produces the conflict when committing. The changes are eventually collected by the garbage collector.

## 4 Validation

First we assess the cost transactions by means of benchmarks that compare the running time of actions performed with and without transactions. Then we compare the cost of thread-safety realized with semaphores to that of our implementation with transactional memory.

### 4.1 Micro benchmarks

We performed several micro benchmarks to establish the runtime cost of using our implementation of transactional memory for Smalltalk. Table 2 shows the times and ratios of performing basic actions, such as invoking a method or accessing state.  $t_1$  is the time required to perform the action  $10^7$  times outside a transactional context, and  $t_2$  is the time required to perform the same action within a transactional context. The benchmarks were performed on an Apple MacBook Pro, 2.16 GHz Intel Core Duo in Squeak 3.9. The required transactional methods were compiled in advance.

The activation time is the time required to enter a transaction as compared to the time required to evaluate a block closure. The *ratio* indicates that entering a transaction is 31 times slower than entering a block closure. This results from

Operation	$t_1$	$t_2$	ratio
Activation	2.75	85.27	31.03
Method invocation	1.98	1.98	1.00
Special method invocation	1.14	2.00	1.75
Instance variable read	1.03	20.72	20.08
Instance variable write	1.13	21.04	18.60
Indexed variable read	1.11	19.92	17.93
Indexed variable write	1.21	20.22	16.75
Global variable read	1.03	20.89	20.25
Global variable write	1.15	21.72	18.92

**Table 2.**  $t_1$ : time in seconds for  $10^7$  runs in a non-transactional context,  $t_2$ : time in seconds for  $10^7$  runs in a transactional context, ratio:  $t_2/t_1$ , the penalty when used in a transactional context.

the fact that entering a transactions requires several objects to be instantiated to track the changes of the transaction. Moreover the transaction is recorded in the current process and an escape context must be captured to be able to abort a running transaction.

Normal method invocation does not show any speed penalty. In all the benchmarks we assume that the transactional methods are already compiled. For some common selectors, such as `#+`, `#*`, `#=`, `#size`, `#at:put:`, `#new`, `#class`, *etc.*, Squeak uses special byte codes to make the invocation about twice as fast as a common send. In a transactional context these byte codes cannot be used anymore and have to be replaced by normal message sends, resulting in a penalty for special method invocations.

State access within a transactional context is fairly expensive. For instance, indexed and global variable reads and writes produce very similar results: in the current implementation these are about 20 times slower than their non-transactional counterparts. As we have seen in Section 3.2, accessing state of an object requires to lookup the current transaction, the change object and to dispatch the state access to its working copy. This whole procedure involves several message sends that cannot be easily optimized in Smalltalk. Further improvements are possible by writing primitives (or introducing new byte-codes) that can more efficiently dispatch that kind of request.


Here we have been comparing the cost of thread-safe actions to unsafe actions. A fairer comparison would be that between thread-safe actions implemented with semaphores and thread-safe transactions. We discuss this in the following section.

## 4.2 Real world example using transactional memory

We applied our transactional model to Pier, a web-based content management system [21, Chapter 3]. Pier uses a tightly-connected graph of objects to represent pages and their content. Edit operations on pages use the Command design pattern and are executed while holding a global lock. Some operations, such as

adding or removing a page, require the system to walk through the whole object graph to invalidate links. This results in a significant amount of time that all other commands are blocked, even though edit operations rarely conflict with each other.

Before executing a command Pier checks for conflicts on the current page, to avoid that changes of other users accidentally get overridden. It does not check for conflicts that could be caused by the need to update links in other parts of the object model. Pier normally does not lock read operations, such as browsing the web site, as they are very common and would introduce a major bottleneck. In rare cases users could therefore encounter an inconsistent state of a particular page.



**Fig. 5.** Average execution time for non-synchronized, lock-based and transactional execution time  $t$  to complete  $n = 1..100$  concurrent edit operations in Pier.

To assess the effectiveness of transactional memory for Smalltalk, we remove the global lock in Pier and wrap the execution of the command within a transactional context. This means that edit commands can now be evaluated concurrently while still ensuring consistency. Moreover we could remove the manual checks for conflicts as these are now detected and handled by the transaction in a complete manner. Page views are now guaranteed to see a consistent state of the web site, as all the changes are applied atomically through a transaction.

Figure 5 shows the average execution time of an edit command that changes the contents of a single page. Using a script we simulated  $n = 1..100$  concurrent edit operations on different pages, so no conflicts could occur. Interestingly the overhead is just over 100 ms for transactions over locks. The transactions are short, involve only few objects and little state access. Memory requirements are moderate: the edit operation touches 39 objects, whereof the transaction

requires 2 copies of each object to track changes. In this particular use case, a single transaction consumes 2'556 bytes of additional memory.

We believe that the transactional approach would be considerably faster than the lock-based one, if the Squeak VM would exploit multiple CPUs to process concurrent requests.

## 5 Conclusion and future work

Smalltalk VMs traditionally offer poor support for concurrency control. Existing Smalltalk dialects provide only lock-based concurrency control, with the exception of GemStone Smalltalk, which provides transactions only for database code. In this paper we have presented an implementation of optimistic transactions for Squeak Smalltalk without modifying the underlying VM.

Our prototype implementation demonstrates that any Smalltalk can profit from having a transactional model. The implementation can be potentially ported to any of today's available Smalltalk platforms, as it is purely based on parse tree transformation of source code. The fact that the whole implementation is written in Smalltalk makes it an ideal platform to experiment with different transaction policies and implementation strategies. Changes to the transactional runtime system and transactional code can be applied and compiled on the fly, so there is no need to restart or rebuild the system.

Our approach works well with external libraries. New code that is loaded into the Smalltalk environment is transformed lazily within the context of a transaction. Primitive methods, filesystem I/O and exceptions work well together with transactions, as special transformation rules can be specified using method annotations. Contrary to other approaches our implementation integrates well with garbage collection, as the transactions are fully implemented in the object system of Smalltalk.

State access within a transaction is about 20 times slower than usual, which is a big penalty to pay. The integration of transactions with the object model at the VM level would certainly lead to much better performance, however we would also lose the flexibility to be able to quickly change the semantics of the transactional mechanisms. Code not using transactions continues to work exactly as before. The traditional mechanisms used for concurrency control can be even mixed with transactions.

As future work we would like to investigate how to further improve the speed of our model. We would like to investigate other areas of applicability, such as atomic loading of source code. In Smalltalk this is traditionally done in an incremental manner and poses certain problems, for example when the application is supposed to continue running while loading.

Furthermore we would like to see how to apply our approach to other dynamic programming languages, such as Python or Ruby. We expect the implementation in those scripting languages to be much more difficult than in Smalltalk, as both languages have major parts of their library implemented in C. Also they both lack direct support to transform source code using high-level AST representations.

Our approach to implementing optimistic transactions in Smalltalk can be seen as a special case of *context-oriented programming* [22], a programming paradigm that supports context-dependent behaviour. Transactional behaviour is automatically dispatched whenever we enter a transactional context. We believe that this approach can be extended more generally to support other forms of context-dependent concurrency control: instead of littering code with explicit calls to specific concurrency mechanisms, one should be able to simply annotate code with the concurrency properties one would like to ensure, and depending on the runtime context the appropriate behaviour will be automatically selected. We also intend to explore more efficient approaches to implementing contextual behaviour, in particular the use of *scoped reflection* [23] to control the temporal and spatial context in which reflective behaviour is active.

### Acknowledgments

We thank Tudor Gîrba for his careful review of a draft of this paper.

We gratefully acknowledge the financial support of the Swiss National Science Foundation for the project “Analyzing, capturing and taming software change” (SNF Project No. 200020-113342, Oct. 2006 - Sept. 2008).

### References

1. Goldberg, A., Robson, D.: Smalltalk 80: the Language and its Implementation. Addison Wesley, Reading, Mass. (May 1983)
2. American National Standards Institute, Inc.: Draft American National Standard for Information Systems — Programming Languages — Smalltalk. American National Standards Institute (1997)
3. Harris, T., Fraser, K.: Language support for lightweight transactions. In: Object-Oriented Programming, Systems, Languages, and Applications. ACM Press, New York, NY, USA (October 2003) 388–402
4. Ingalls, D., Kaehler, T., Maloney, J., Wallace, S., Kay, A.: Back to the future: The story of Squeak, A practical Smalltalk written in itself. In: Proceedings OOPSLA '97, ACM SIGPLAN Notices, ACM Press (November 1997) 318–326
5. Hansen, P.B.: Monitors and Concurrent Pascal: a personal history. ACM Press, New York, NY, USA (1996)
6. Herlihy, M.P.: Wait-free synchronization. ACM Transactions on Programming Languages and Systems **13**(1) (January 1991) 124–149
7. Herlihy, M.P., Moss, J.E.B.: Transactional memory: Architectural support for lock-free data structures. In: Proceedings of the 20. Annual International Symposium on Computer Architecture. (1993)
8. Butterworth, P., Otis, A., Stein, J.: The GemStone object database management system. Commun. ACM **34**(10) (1991) 64–77
9. Guerraoui, R., Garbinato, B., Mazouni, K.R.: The garf library of dsm consistency models. In: EW 6: Proceedings of the 6th workshop on ACM SIGOPS European workshop, New York, NY, USA, ACM Press (1994) 51–56
10. Guerraoui, R., Felber, P., Garbinato, B., Mazouni, K.: System support for object groups. In: OOPSLA '98: Proceedings of the 13th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications, New York, NY, USA, ACM Press (1998) 244–258

11. Briot, J.P.: Actalk: A testbed for classifying and designing actor languages in the Smalltalk-80 environment. In Cook, S., ed.: Proceedings ECOOP '89, Nottingham, Cambridge University Press (July 1989) 109–129
12. Hindman, B., Grossman, D.: Atomicity via source-to-source translation. In: MSPC '06: Proceedings of the 2006 workshop on Memory system performance and correctness, New York, NY, USA, ACM Press (2006) 82–91
13. Adl-Tabatabai, A.R., Lewis, B.T., Menon, V., Murphy, B.R., Saha, B., Shpeisman, T.: Compiler and runtime support for efficient software transactional memory. In: PLDI '06: Proceedings of the 2006 ACM SIGPLAN conference on Programming language design and implementation, New York, NY, USA, ACM Press (2006) 26–37
14. Röthlisberger, D., Denker, M., Tanter, É.: Unanticipated partial behavioral reflection: Adapting applications at runtime. *Journal of Computer Languages, Systems and Structures* (2007) To appear.
15. Denker, M., Ducasse, S., Lienhard, A., Marschall, P.: Sub-method reflection. *Journal of Object Technology* **6**(9) (September 2007) To appear.
16. Lie, S.: Hardware support for unbounded transactional memory. Master's thesis, Massachusetts Institute of Technology (May 2004)
17. Moore, K.E., Bobba, J., Moravan, M.J., Hill, M.D., Wood, D.A.: LogTM: Log-based transactional memory. In: Proceedings of the 12th International Symposium on High-Performance Computer Architecture. IEEE Computer Society (February 2006) 254–265
18. Lea, D.: *Concurrent Programming in Java, Second Edition: Design principles and Patterns*. 2nd edn. The Java Series. Addison Wesley (1999)
19. Bobba, J., Moore, K.E., Yen, L., Volos, H., Hill, M.D., Swift, M.M., Wood, D.A.: Performance pathologies in hardware transactional memory. In: Proceedings of the 34rd Annual International Symposium on Computer Architecture. International Symposium on Computer Architecture (June 2007)
20. Kung, H.T., Robinson, J.T.: On optimistic methods for concurrency control. *ACM TODS* **6**(2) (June 1981)
21. Renggli, L.: Magritte – meta-described web application development. Master's thesis, University of Bern (June 2006)
22. Costanza, P., Hirschfeld, R.: Language constructs for context-oriented programming: An overview of ContextL. In: Proceedings of the Dynamic Languages Symposium (DLS) '05, co-organized with OOPSLA'05, New York, NY, USA, ACM Press (October 2005)
23. Nierstrasz, O., Denker, M., Gîrba, T., Lienhard, A.: Analyzing, capturing and taming software change. In: Proceedings of the Workshop on Revival of Dynamic Languages (co-located with ECOOP'06). (July 2006)