Noninvasive concurrency with Java STM

Guy Korland[1] and Nir Shavit[1] and Pascal Felber[2]

[1] Computer Science Department
Tel-Aviv University, Israel,
`guykorla@post.tau.ac.il`, `shanir@post.tau.ac.il`
[2] University of Neuchâtel, Neuchâtel, Switzerland,
`pascal.felber@unine.ch`

Abstract. In this paper we present a complete Java STM framework, called
DEUCE, intended as a platform for developing scalable concurrent applications
and as a research tool for designing new STM algorithms. It was not clear
if one could build an efficient Java STM without compiler support. DEUCE
provides several benefits over existing Java STM frameworks: it avoids any
changes or additions to the JVM, it does not require language extensions or
intrusive APIs, and it does not impose any memory footprint or GC overhead.
To support legacy libraries, DEUCE dynamically instruments classes at load
time and uses an original "field-based" locking strategy to improve concur-
rency. DEUCE also provides a simple internal API allowing different STMs
algorithms to be plugged in. We show empirical results that highlight the
scalability of our framework running benchmarks with hundreds of concur-
rent threads. This paper shows for the first time that one can actually design
a Java STM with reasonable performance without compiler support.

1 Introduction

Multicore CPUs have become commonplace, with dual-cores powering almost any
modern portable or desktop computer, and quad-cores being the norm for new
servers. While multicore hardware has been advancing at an accelerated pace, soft-
ware for multicore platforms seems to be at a crossroads.

Currently, two diverging programming methods are commonly used. The first ex-
ploits concurrency by synchronizing multiple threads based on locks. This approach
is well known to be a two-edged sword: on the one hand, locks give the programmer
a fine-grained way to control the applications critical sections, allowing an expert to
provide great scalability; on the other hand, because of the risk of deadlocks, starva-
tion, priority inversions, etc., they impose a great burden on non-expert programmers,
often leading them to prefer the safer but non-scalable alternative of coarse-grained
locking.

The second method is a shared-nothing model usually common in Web based
architectures. The applications usually contain only the business logic deployed in
a container, while the state is saved in an external multi-versioned control system,
such as database, message queue, or distributed cache. While this method removes
the burden of handling concurrency in the application, it imposes a huge performance
impact on data accesses and is not seamlessly applicable to many types of application.

The hope is that transactional memory (TM) will simplify concurrent program-
ming by combining the desirable features of both methods, providing *state-aware
shared memory* with *simple concurrency control*.

In the past several years there has been a flurry of software transactional memory (STM) design and implementation work. The state of the art today however is less than appealing. With the notable exception of transactional C/C++ compilers [9], most of STM initiatives have remained academic experiments, applicable only to "toy applications", and though we have learned much from the process of developing them, they never reached a state that will allow them to be seriously field tested. There are several reasons for this. Among them are the problematic handling of many features of the target language, the large performance overheads, and the lack of support for legacy libraries. Moreover, many of the published results have been based on prototypes whose source code is unavailable or poorly maintained, making a reliable comparison between the various algorithms and designs very difficult.

In this paper, we introduce DEUCE, our novel open-source Java framework for transactional memory. DEUCE has several desired features not found in earlier Java STM frameworks. As we discuss in Section 2, there currently does not exist an *efficient* Java STM framework that delivers a full feature STM that can be added to an existing application *without changes* to its compiler or libraries. It was not clear if one could build an efficient Java STM without compiler support.

DEUCE is intended to fill this gap. It is non-intrusive in the sense that no modifications to the Java virtual machine (JVM) or extensions to the language are necessary. It uses, by default, an original locking design that detects conflicts at the level of individual *fields* without a significant increase in the memory footprint (no extra data is added to any of the classes) and therefore there is no GC overhead. This locking scheme provides finer granularity and better parallelism than former object-based lock designs. In order to avoid performance penalty, DEUCE provides weak atomicity, i.e., it does not guarantee that concurrent accesses to a shared memory location from both inside and outside a transaction are consistent. It also supports a pluggable STM back-end and an open easily extendable architecture, allowing researchers to integrate and test their own STM algorithms.

DEUCE has been heavily optimized for efficiency and, while there is still room for improvements, our performance evaluations on several high-end machines (up to 128 hardware threads on a 16-core Sun Niagara-based machine and a 96-core Azul machine) demonstrate that it scales well. Our benchmarks show that it outperforms the main competing JVM-independent Java STM, the DSTM2 framework [6], in many cases by two orders of magnitude, and in general scales well on many workloads. This paper shows for the first time that one can actually design a Java STM with reasonable performance without compiler support.

The DEUCE framework has been in development for more than a year, and we believe it has reached a level of maturity sufficient to allow it to be used by developers with no prior expertise in transactional memory. It is our hope that DEUCE can help democratize the use of STM among developers, and that its open-source nature will encourage STM them to extend the infrastructure with novel features, as has been the case, for instance, with the Jikes RVM [1].

The rest of the paper is organized as follows. We first overview related work in Section 2. Section 3 describes DEUCE from the perspective of the developer of a concurrent application. In Section 4 we discuss the implementation of the framework, and we show how it can be extended by the means of the STM backends in Section 5. Finally, in Section 6, we evaluate the performance of DEUCE.

2   Related work

Several tools have been proposed to allow the introduction of STM into Java. They differ in their programming interface and in the way they are implemented. One of the first proposals is Harris and Fraser's CCR [5] that used a C-based STM implementation underneath the JVM and a programmatic API to use STM in Java.

DSTM [7] is another pioneering Java STM implementation. It used an explicit API to demarcate transactions and read/write shared data. Transactional objects had to additionally provide some pre-defined functionality for cloning their state.

More recently, the DSTM2 [6] framework was introduced, proposing a set of higher-level mechanisms to support transactions in Java. DSTM2 is a pure Java library that does not require any change to the JVM nor to the compiler. It uses a special annotation to mark an atomic interface. Only a class that implements an annotated interface can participate in a transaction. This class must be created using a special factory and it can only support transactional access to primitive types or other annotated classes. Transactions must be started using a `Callable` object. This design creates an API that requires important refactoring of legacy code and introduces many limitations, such as the lack of support for existing libraries.

LSA-STM [12] is a Java STM that also relies primarily on annotations for TM programming. Transactional objects to be accessed in the context of a transaction must be annotated as `@Transactional` and methods must be declared as `@Atomic`. Additional annotations can be used to indicate that some methods will not modify the state of the target object (read-only) or to specify the behavior in case an exception is thrown within a transaction. Transactional objects are implemented in LSA-STM using inheritance and must support state cloning, as such it does not fully support legacy code and is not transparent nor non-invasive as DEUCE is. Among the limitations of the framework, one can mention that accesses to a field are only supported as part of methods from the owner class, therefore public and static fields cannot be accessed transactionally. Note that this limitation also applies to other Java STM implementations with object-level conflict detection like DSTM and DSTM2.

AtomJava [8] takes a different approach to Java STM design by providing a source-to-source translator based on Polyglot [10], an extensible compiler framework. AtomJava adds a new `atomic` keyword to mark atomic blocks, and performs some major extensions during code transformation, such as adding an extra instance field to every single class. AtomJava relies on this field to maintain an object-based lock schema. It is shown to reduce lock access overhead but imposes a memory overhead which impacts not only transactional objects but all the application objects. The source code translation approach simplifies the tuning and verification of STM instrumentation, but it makes it harder for the programmer to debug her original code. Also, by translating source code, AtomJava imposes a major limitation on users in that it prevents them from using compiled libraries, and renders their source code with atomic blocks incompatible with regular Java compilers (unlike approaches such as DSTM2 and LSA-STM that are based on annotations).

In summary, DEUCE is the first *efficient full-feature* STM framework that can be added to an existing application *without changing its compilation process or libraries*.

3 Concurrent Programming with DEUCE

One of the main goals in designing the DEUCE API was to keep it simple. A survey of the past designs (see previous section) reveals three main approaches: (i) adding a new reserved keyword to mark a block as atomic, e.g., `atomic` in AtomJava [8]; (ii) using explicit method calls to demarcate transactions and access shared objects, e.g., DSTM2 [6] and CCR [5]; or (iii) requiring atomic classes to implement a predefined interface or extend a base class, and to implement a set of methods [7]. The first approach requires modifying the compiler and/or the JVM, while the others are intrusive from the programmer's perspective as they require significant modifications to the code (even though some systems use semi-automatic weaving mechanisms such as aspect-oriented programming to ease the task of the programmer, e.g., [12]).

In contrast, DEUCE has been designed to avoid any addition to the language or changes to the JVM. In particular, no special code is necessary to indicate which objects are transactional, no method needs to be implemented to supporting transactional objects (e.g., cloning the state as in [7]). This allows DEUCE to seamlessly support transactional accesses to compiled libraries. The only piece of information that must be specified by the programmer is, obviously, which part of the code should execute atomically in the context of a transaction.

To that end, DEUCE relies on Java annotations. Introduced as a new feature in Java 5, annotations allow programmers to mark a method with metadata that can be consulted at class loading time. DEUCE introduces new types of annotations to mark methods as atomic: their execution will take place in the context of a transaction.

This approach has several advantages, both technically and semantically. First technically, the smallest code block Java can annotate is a method, which simplifies the instrumentation process of DEUCE and provides a simple model for the programmer. Second, atomic annotations operate at the same level as *synchronized* methods, which execute in a mutually exclusion manner on a given object; therefore, atomic methods provide a familiar programming model.

From a semantic point of view, implementing atomic blocks at the granularity of methods removes the need to deal with local variables as part of the transaction. In particular, since Java doesn't allow any access to stack variables outside the current method, the STM can avoid logging many memory accesses. For instance, in Figure 1, a finer-granularity atomic block would

```
1  public int sum(List list) {
2    int total = 0;
3    atomic {
4      for (Node n : list)
5        total += n.getValue();
6    }
7    return total;
8  }
```

Fig. 1. Atomic block example.

require costly logging of the `total` local variable (otherwise the method would yield incorrect results upon abort) whereas no logging would be necessary when considering atomic blocks at the granularity of individual methods.

To illustrate the use and implementation of DEUCE, we will consider a well-known but non-trivial data structure: the skip list [11]. A skip list is a probabilistic structure based on multiple parallel, sorted linked lists, with efficient search time in $\mathcal{O}(\log n)$.

Figure 2 shows a partial implementation of skip list, with an inner class representing nodes and a method to search for a value through the list. The key observation in this code is that transactifying an application is as easy as adding `@Atomic` annotations to methods that should execute as transactions. No code needs to be changed

within the method or elsewhere in the class. Interestingly, the linked list directly manipulates arrays and accesses public fields from outside their enclosing objects, which would not be possible with DSTM2 or LSA-STM.

```
1   public class SkipList {                    16    @Atomic(retries=64)
2     private static class Node {              17    public boolean contains(int v) {
3       public final int value;                18      Node node = head;
4       public final Node[] forward;           19      for (int i = level; i >= 0; i−−) {
5       // ...                                  20        Node next = node.forward[i];
6       public Node(int level, int v) {         21        while (next.value < v) {
7         value = v;                            22          node = next;
8         forward = new Node[level + 1];        23          next = node.forward[i];
9       }                                        24        }
10      //...                                    25      }
11    }                                          26      node = node.forward[0];
12    private static int MAX_LEVEL = 32;         27      return (node.value == v);
13    private int level;                         28    }
14    private final Node head;                   29    // ...
15    // Continued in next column...             30  }
```

Fig. 2. @Atomic method example.

One can also observe that the `@Atomic` annotation provides one configurable attribute, `retries`, to optionally limit the amount of retries the transaction attempts (at most 64 times in the example). A `TransactionException` is thrown in case this limit is reached. Alternatively one can envision providing timeout instead (Which we might add in future versions).

A DEUCE application is compiled with a regular Java compiler. Upon execution, one needs to specify a Java agent that allows DEUCE to intercept every class loaded and manipulate it before it is loaded by the JVM. The agent is simply specified on the command line as a parameter to the JVM, as:

```
java -javaagent:deuceAgent.jar MainClass args...
```

As will be discussed in the next section, DEUCE instruments every class that may be used from within a transaction, not only classes that have `@Atomic` annotations. If it is known that a class will never be used in the context of a transaction, one can prevent it from being instrumented by providing exclusion lists to the DEUCE agent. This will speed up the application loading time yet should not affect execution speed.

4   DEUCE Implementation

This section describes the implementation of the DEUCE framework. We first give a high-level overview of its main components. Then, we explain the process of code instrumentation. Finally, we describe various optimizations that enhance the performance of transactional code.
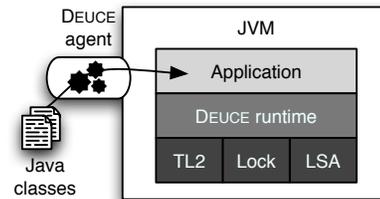

Fig. 3. Main components of the DEUCE architecture.

4.1   Overview

The DEUCE framework is conceptually made of 3 layers, as shown in Figure 3:

1. The application, which consists of user classes written without any relationship to the STM implementation, except for annotations added to atomic methods.
2. The DEUCE runtime that orchestrates the interactions between transactions executed by the application and the underlying STM implementation.

3. The actual STM libraries that implement the DEUCE context API (see Section 5), including a reference single-global-lock implementation ("Lock" in the Figure).

In addition, the DEUCE agent intercepts classes as they are loaded and instruments them before they start executing.

4.2 Instrumentation Framework

DEUCE's instrumentation engine is based on ASM [3], an all-purpose Java bytecode manipulation and analysis framework. It can be used to modify existing classes or to dynamically generate classes, directly in binary form. The DEUCE Java agent uses ASM to dynamically instrument classes as they are loaded, before their first instance is created in the virtual machine. During instrumentation, new fields and methods are added, and the code of transactional methods is transformed. We now describe the different manipulations performed by the Java agent.

Fields. For each instance field in any loaded class, DEUCE adds a synthetic constant field (`final static public`) that represents the relative position of the field in the class. This value, together with the the instance of the class, uniquely identifies a field and is used by the STM implementation to log field accesses.

Static fields in Java are effectively fields of the enclosing class. To designate static fields, DEUCE defines for each class a constant that represents the base class and can be used in combination with the field position, instead of the class instance.

```
1  public class SkipList {
2    public static final long __CLASS_BASE__   = ...
3    public static final long MAX_LEVEL__address__ = ...
4    public static final long level__address__ = ...
5    // ...
6    private static int MAX_LEVEL = 32;
7    private int level;
8    // ...
9  }
```

Fig. 4. Fields address.

For instance, in Figure 4, the `level` field is represented by `level__ADDRESS__` while the `SkipList` base class is represented by `__CLASS_BASE__`.

Accessors. For each field of any loaded class, DEUCE adds synthetic `static` accessors used to trigger field's access events on the local context.

Figure 5 shows the synthetic accessors of class `SkipList`. The *getter* `level__Getter$` receives the current field value and a context, and triggers two events: `beforeReadAccess` and `onReadAccess`; the result of the latter is returned by the getter. The *setter* receives a context and yields a single event:

```
1  public class SkipList {
2    private int level;
3    // ...
4
5    // Synthetic getter
6    public int level__Getter$(Context c) {
7      c.beforeReadAccess(this, level__ADDRESS__);
8      return c.onReadAccess(this, level, level__ADDRESS__);
9    }
10   // Synthetic setter
11   public void level__Setter$(int v, Context c) {
12     c.onWriteAccess(this, v, level__ADDRESS__);
13   }
14   // ...
15 }
```

Fig. 5. Fields accessors.

`onWriteAccess`. The reason for having two events in the getter is technical: the

"before" and "after" events allow the STM backend to verify that the value of the field, which is accessed between both events without using costly reflection mechanisms, is consistent. This is typically the case in time-based STM algorithms like TL2 and LSA that ship with DEUCE.

Duplicate methods. In order to avoid any performance penalty on non transactional code and since DEUCE provides weak atomicity, DEUCE duplicates each method to provide two distinct versions. The first version is identical to the original method: it does not trigger an event upon memory access and, conse-

```
1  private static class Node {
2    public Node[] forward;
3    // ...
4
5    // Original method
6    public void setForward(int level, Node next) {
7      forward[level] = next;
8    }
9
10   // Synthetic duplicate method
11   public void setForward(int level, Node next, Context c) {
12     Node[] f = forward__Getter$(c) {
13     forward__Setter$(f, level, c)
14   }
15   // ...
16 }
```

Fig. 6. Duplicate method.

quently, does not impose any transactional overhead. The second version is a synthetic method with an extra `Context` parameter. In this instrumented copy of the original method, all field accesses (except for `final` fields) are replaced by calls to the associated transactional accessors. Figure 4 shows the two versions of a method of class `Node`. The second synthetic overload replaces `forward[level] = next` by calls to the synthetic getter and setter of the `forward` array. The first call obtains the reference to the array object, while the second one changes the specified array element (note that getters and setters for array elements have an extra index parameter).

```
1  public class SkipList {
2    // ...
3
4    // Original method instrumented
5    public boolean contains(int v) {
6      Throwable throwable = null;
7      Context context =
8        ContextDelegator.getInstance();
9      boolean commit = true;
10     boolean result;
11
12     for (int i = 64; i > 0; --i) {
13       context.init ();
14       try {
15         result = contains(v, context);
16       } catch(TransactionException ex) {
17         // Must rollback
18         commit = false;
19       } catch(Throwable ex) {
20         throwable = ex;
21       }
22       // Continued in next column...

23         // Try to commit
24         if (commit) {
25           if (context.commit()) {
26             if (throwable == null)
27               return result;
28             // Rethrow application exception
29             throw (IOException)throwable;
30           }
31         } else {
32           context.rollback ();
33           commit = true;
34         }
35     } // Retry loop
36     throw new TransactionException();
37   }
38
39   // Synthetic duplicate method
40   public boolean contains(int v, Context c) {
41     Node node = head__Getter$(c);
42     // ...
43   }
44 }
```

Fig. 7. Atomic method.

Atomic methods. The duplication process described above has one exception: a method annotated as `@Atomic` does not need an uninstrumented version. Instead, the original method is replaced by a transactional version that calls the instrumented version from within a transaction that executes in a loop. The process repeats as long as the transaction aborts and the number of retires is not reached. Figure 7 shows the transactional version of method `contains`.

4.3 Optimizations

During instrumentation, we perform several optimizations to improve the performance of DEUCE. First, we do not instrument accesses to final fields as they cannot be modified after creation. This optimization, together with the declaration of final fields whenever possible in the application code, dramatically reduces the overhead.

Second, instead of generating accessor methods, DEUCE actually inlines the code of the getters and setters directly in the transactional code. We have observed a slight performance improvement from this optimization.

Third, we chose to use the `sun.misc.Unsafe` pseudo-standard internal library to implement fast reflection, as it proved to be vastly more efficient than the standard Java reflection mechanisms. Using `sun.misc.Unsafe` even outperformed the approach taken in AtomJava [8], which is based on using an anonymous class per field to replace reflection.

Finally, we tried to limit as much as possible the stress on the garbage collector, notably by using object pools when keeping track of accessed fields (read and write sets) in threads. In order to avoid any contention on the pool, we had each thread keep a separate object pool as part of its context.

Together, the above optimizations helped to significantly decrease the implementation overhead, in some of our benchmarks by almost an order of magnitude (almost ten times faster) as compared to our initial implementation.

5 Customizing Concurrency Control

DEUCE was designed to provide a research platform for STM methods. In order to provide a simple API for researchers to plug in their own STM implementation, DEUCE defines the `Context` API as shown in Listing 8. The API includes an `init` method called once before the transaction starts and then upon each retry, allowing the transaction to initialize its internal data structures. The `atomicBlockId` argument allows the transaction to log information about the specific atomic block (statically assigned in the bytecode).

One of the heuristics we added to the LSA implementation is, following [4], that each one of the atomic blocks will initially be a read-only block. It will be converted to become a writable block upon retry once it encounters a first write. Using this method, read-only blocks can save most of the overhead of logging the fields' access.

```
1  public interface Context {
2    void init(int atomicBlockId);
3    boolean commit();
4    void rollback();
5
6    void beforeReadAccess(Object obj, long field);
7
8    Object onReadAccess(Object obj, Object value, long field);
9    int    onReadAccess(Object obj, int    value, long field);
10   long   onReadAccess(Object obj, long value, long field);
11   // ...
12
13   void onWriteAccess(Object obj, Object value, long field);
14   void onWriteAccess(Object obj, int    value, long field);
15   void onWriteAccess(Object obj, long value, long field);
16   // ...
17 }
```

Fig. 8. Context interface.

Another heuristic is that `commit` is called in case the atomic block finishes without a `TransactionException` and can return `false` in case the commit fails, which in

turn will cause a retry. A `rollback` is called when a `TransactionException` is thrown during the atomic block (this can be used by the business logic to force a retry).

The rest of the methods are called upon field access: a field read event will trigger a `beforeReadAccess` event followed by a `onReadAccess` event. A write field event will trigger an `onWriteAccess` event. DEUCE currently includes two `Context` implementations, `tl2.Context` and `lsa.Context`, implementing the TL2 [4] and LSA [12] STM algorithms. DEUCE also provides a reference implementation based on a single global lock. Since a global lock doesn't log any field access, it doesn't implement the Context interface and doesn't impose any overhead on fields' access.

TL2 Context. The TL2 context is a straightforward implementation of the TL2 algorithm [4]. The general principle is as follows (many details omitted).

TL2 uses a shared array of *revokable versioned locks*, with each object field being mapped to a single lock. Each lock has a version that corresponds to the commit timestamp of the transaction that last updated some field protected by this lock. Timestamps are acquired upon commit from a global time base, implemented as a simple counter.

When reading some data, a transaction checks that the associated timestamp is valid, i.e., not more recent than the time when the transaction started, and keeps track of the accessed location in its read set. Upon write, the transaction buffers the update in its write set.

At commit time, the transaction acquires (using an atomic compare-and-set operation) the locks protecting all written fields, verifies that all entries in the read set are still valid, acquires a unique commit timestamp, writes the modified fields to memory, and releases the locks. If locks cannot be acquires or validation fails, the transaction aborts, discarding buffered updates.

LSA Context. The LSA context uses the LSA algorithm [12] that was developed concurrently with TL2 and is based on a similar design. The main differences are that (1) LSA acquires locks as fields are written, instead of at commit time, and (2) it performs "incremental validation" to avoid aborting transactions that read data that has been modified more recently than the start time of the transaction.

Both TL2 and LSA take a simple approach to conflict management: they simply abort and restart (possibly after a delay) the transaction that encounters the conflict. This strategy is simple and efficient to implement, because transactions unilaterally abort without any synchronization with others, but it can sometimes hamper progress, e.g., by producing livelocks. DEUCE also supports a variant of the LSA context that provides modular contention management strategies (as first proposed in [7]), at the price of some additional synchronization overhead at runtime.

6 Performance Evaluation

We evaluated the performance of DEUCE on a Sun UltraSPARC T2 Plus multicore machine (2 CPUs each with 8 cores at 1.2 GHz, each with 8 hardware threads) and an Azul Vega2 (2 CPUs each with 48 cores).

6.1 DEUCE overheads

We first briefly discuss the overheads introduced by the DEUCE agent when instrumenting Java classes. Table 1 shows the memory footprint and the processing overhead when processing compiled libraries: `rt.jar` is the runtime library containing all Java built-in classes for JDK 6; `cache4j.jar` (version 0.4) is a widely used in-memory cache for Java objects; *JavaGrande* (version 1.0) is a well-known collection of Java benchmarks. Instrumentation times were measured on a x86 Core 2 Duo CPU running at 1.8 GHz. Note that instrumentation was executed serially, i.e., without exploiting multiple cores.

Application	Memory		Instrumentation Time
	Original	Instrumented	
`rt.jar`	50 MB	115 MB	29 s
`cache4j.jar`	248 KB	473 KB	<1 s
JavaGrande	360 KB	679 KB	<1 s

Table 1. Memory footprint and processing overhead.

As expected, the size of the code approximately doubles because DEUCE duplicates each method, and the processing overhead is proportional to the size of the code. However, because Java uses lazy class loading, only the classes that are actually *used* in an application will be instrumented at load time. Therefore, the overheads of DEUCE are negligible considering individual classes are loaded on demand.

In terms of execution time, instrumented classes that are not invoked within a transaction incur no performance penalty as the original method executes, not the transactional version.

6.2 Benchmarks

We tested the three built-in DEUCE STM options: LSA, TL2, and a simple Global Lock. Our LSA version captures many of the properties of the LSA-STM [13] framework. The TL2 [4] form provides an alternative STM implementation that acquires locks at commit time using a write set as opposed to the encounter order and undo log approach used in LSA.

Our experiments included three classical micro-benchmarks: a red/black tree, a sorted linked list and a skip list, on which we performed concurrent insert, delete, and search operations. We controlled the size of the data structures, which remained constant for the duration of the experiments, and the mix of operations.

We also experimented with two real-world applications, the first implements the Lee routing algorithm (as presented in Lee-TM [2]). It is worth noting that adapting the application to DEUCE was trivial, with little more than a change of synchronized blocks into atomic methods. The second application called Cache4J, is an LRU cache implementation, again adapting the application to DEUCE was trivial.

Micro-benchmarks. We began by comparing DEUCE to DSTM2, the only competing STM framework that does not require compiler support. We benchmarked DSTM2 and DEUCE based on the Red/Black tree benchmark provided with the DSTM2 framework. We tried many variations of operation combinations and levels of concurrency on both the Azul and Sun machines. Comparison results for a tree with 10k elements are shown in Figure 9, we found that in all the benchmarks DSTM2

was about 100 times (two orders of magnitude) slower than DEUCE. Our results are consistent with those published in [6], where DSTM2's throughput in operations per second is in the thousands while DEUCE throughput is in the millions. Based on these experiments, we believe one can conclude that DEUCE is the first viable compiler and JVM independent Java STM framework.
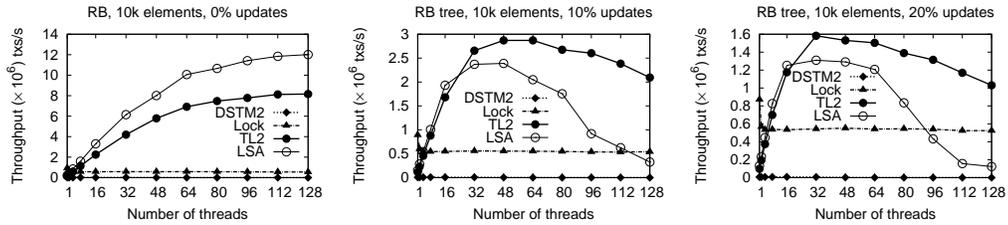

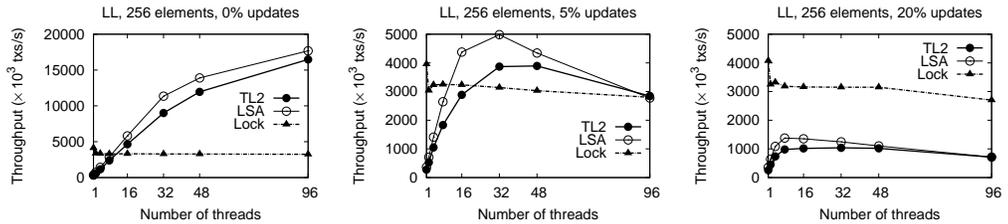**Fig. 9.** The red/black tree benchmark (Sun).
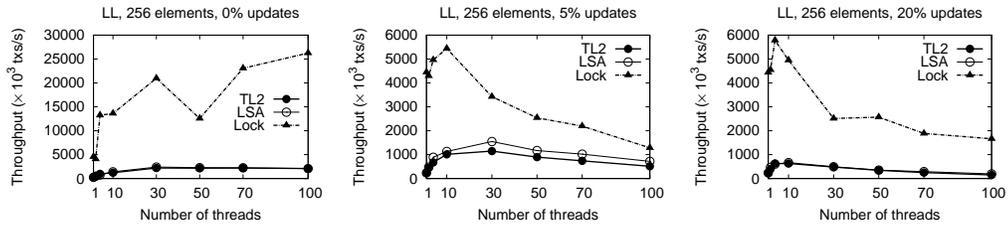

Fig. 10. The linked list benchmark (Sun).

Fig. 11. The linked list benchmark (Azul).
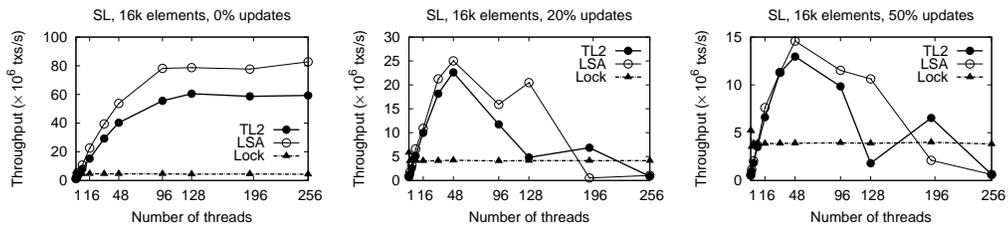
Fig. 12. The skiplist benchmark (Sun).

On the other hand we observe that the DEUCE STMs scale in an impressive way even with 20% updates (up to 32 threads). While DSTM2 shows no scalability. When we investigated deeper we found out that most of DSTM2 overhead rest in two areas. The most significant area is the contention manager which acts as a contention point, the second area is the reflection mechanism used by DSTM2 to retrieve and assign values during the transaction and in commit.
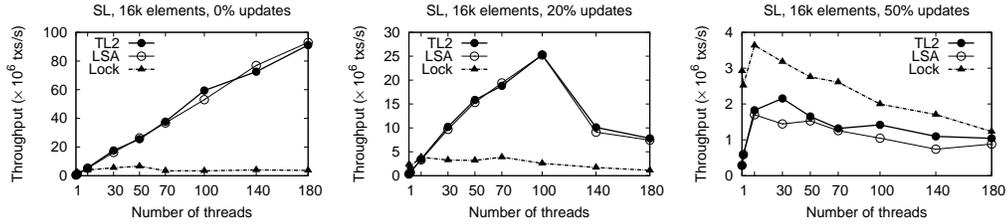
Fig. 13. The skiplist benchmark (Azul).

Then we present the results of the tests with the linked list and the skip list on the Azul machine. The linked list benchmark is known to produce a large degree of contention as all threads must traverse the same long sequence of nodes to reach the target element. Results of experiments are shown in Figure 10 and Figure 11 for a linked list with 256 elements and different update rates. We observe that because there is little potential for parallelism, the single lock, which performs significantly less lock acquisition and tracking work than the STMs, wins in all three benchmarks. The STMs scale a bit until 30 threads when there are up to 5% updates, and then drop due to a rise in overhead with no benefit in parallelism. With 20% updates the max is reached at about 5 threads as after then we have statistically at least 2 concurrent updates and the STMs suffer from repeated transactional aborts.

Next, we consider results from the skip list benchmark. Skip lists allow significantly greater potential for parallelism than linked lists because different threads are likely to traverse the data structure following distinct paths. Results for a list with 16k elements are shown in Figure 12 and Figure 13. We observe that the STMs scale in an impressive way as long as there is an increase in the level of parallelism, and provide great scalability even with 20% updates as long as the abort rate remains at a reasonable level. We added a benchmark with 50% updates to show that there is a limit to their scalability, when we increase the fraction of updates to 50%, the STMs in Figure 13 reach a "melting point" much earlier (at about 10 threads versus 100 threads in the 20% benchmark) and overall the lock wins again because the abort rate is again high and the STMs pay with overhead without a gain in parallelism. We note that typically search structures have about 10% updates.

Real applications. Our last benchmarks demonstrate how simple it's to replace the critical sections with transactions. The first takes a serial version of the Lee routing algorithm and demonstrates how a simple replacement of the critical sections by transactions significantly improves scalability. While the second takes a non-multi-threaded lock based a LRU cache implementation (Cache4J) and shows that it's very simple to replace the critical sections but scalability isn't promised.

Circuit routing is the process of automatically producing an interconnection between electronic components. Lee's routing algorithm is attractive for parallelization since circuits (as shown in [2]) consist of thousands of routes, each of which can potentially be routed concurrently.

The graph in Figure 14 shows execution time (not throughput). As can be seen, DEUCE scales well, with the overall time decreasing even with a large board (Mem-Board).

Cache4J is an LRU lock-based cache implementation. The implementation is based on two internal data structures, a tree and a hash-map. The tree manages the LRU while the hash-map holds the data. The Cache4J implementation is based
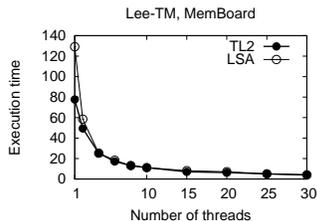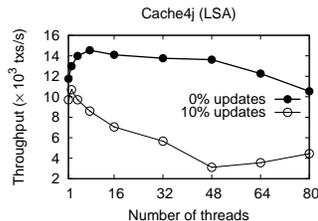
Fig. 14. The Lee Routing benchmark. **Fig. 15.** The `cache4j` benchmark.

on a single global lock which naturally leads to zero scalability. The graph in Figure 15 shows the result of replacing the single lock with transactions. As can be seen DEUCE doesn't scale well, with the overall throughput slightly decreasing. A quick profiling shows that the fact that Cache4J is an LRU cache implies that every get operation also updated the internal tree. This alone makes every transaction an update transaction, the fact that the total throughput remains almost the same even with 80 threads encouraging due to the fact that transactions' biggest advantages, part of scalability, is code simplicity and code robustness.

In conclusion, DEUCE shows the typical performance patterns of an STM, good scalability when there is a potential for parallelism, and unimpressive performance when parallelism is low. The results, however, are very encouraging as we see that the fact that DEUCE is a pure Java library, with all the implied overheads, does not prevent it from showing scalability even for small numbers of threads.

7 Conclusion

We introduced DEUCE, a novel open-source Java framework for transactional memory. As we showed, DEUCE is the first efficient fully featured Java STM framework that can be added to an existing application without changes to its compiler or libraries. It's a proof that one can design an efficient pure Java STM without a compiler support.

Though much work obviously remains to be done, we believe DEUCE is ready for use by developers. It is freely downloadable from `http://code.google.com/p/deuce` under the Apache license.

In conclusion Table 2 shows a general overview comparing the two known Java STMs AtomJava [1] and DTMS2 vs DEUCE. This comparison shows that DEUCE gets the best from both worlds and provides new novel capabilities which yell much better usability and performance than currently exists.

Application	AtomJava	DSTM2	DEUCE
Locks	Object based	Object based	Field based
Instrumentation	Source	Runtime	Runtime
API	Non-Intrusive	Intrusive	Non-Intrusive
Libraries support	no-support	no-support	Runtime & offline support
Fields access	Anonymous classes	Reflection	low level Unsafe
Execution overhead	Medium	High	Medium
Extensible	Yes	Yes	Yes
Transaction context	Atomic block	Need to create a `Callable`	`@Atomic` annotated method

Table 2. AtomJava vs DSTM2 vs DEUCE

---

[1] We don't show comparison benchmark results with AtomJava due technical limitations and bugs found in its current implementation. Although small benchmarks showed DEUCE outperformed AtomJava

References

1. B. Alpern, S. Augart, S. M. Blackburn, M. Butrico, A. Cocchi, P. Cheng, J. Dolby, S. Fink, D. Grove, M. Hind, K. S. McKinley, M. Mergen, J. E. B. Moss, T. Ngo, and V. Sarkar. The jikes research virtual machine project: building an open-source research community. *IBM Syst. J.*, 44(2):399–417, 2005.
2. M. Ansari, C. Kotselidis, I. Watson, C. Kirkham, M. Luján, and K. Jarvis. Lee-tm: A non-trivial benchmark suite for transactional memory. In *Proc. of the International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP)*, pages 196–207, Berlin, Heidelberg, 2008. Springer-Verlag.
3. W. Binder, J. Hulaas, and P. Moret. Advanced java bytecode instrumentation. In *Proc. of the International Symposium on Principles and Practice of Programming in Java (PPPJ)*, pages 135–144, New York, NY, USA, 2007. ACM.
4. D. Dice, O. Shalev, and N. Shavit. Transactional locking ii. In *Proc. of the International Symposium on Distributed Computing (DISC)*, pages 194–208, 2006.
5. T. Harris and K. Fraser. Language support for lightweight transactions. In *Proc. of the ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 388–402, New York, NY, USA, 2003. ACM.
6. M. Herlihy, V. Luchangco, and M. Moir. A flexible framework for implementing software transactional memory. In *Proc. of the ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 253–262, New York, NY, USA, 2006. ACM.
7. M. Herlihy, V. Luchangco, M. Moir, and W. N. Scherer III. Software transactional memory for dynamic-sized data structures. In *Proc. of the Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 92–101, New York, NY, USA, 2003. ACM.
8. B. Hindman and D. Grossman. Atomicity via source-to-source translation. In *Proc. of the Workshop on Memory System Performance and Correctness (MSPC)*, pages 82–91, New York, NY, USA, 2006. ACM.
9. Y. Ni, A. Welc, A.-R. Adl-Tabatabai, M. Bach, S. Berkowits, J. Cownie, R. Geva, S. Kozhukow, R. Narayanaswamy, J. Olivier, S. Preis, B. Saha, A. Tal, and X. Tian. Design and implementation of transactional constructs for C/C++. In *Proc. of the ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 195–212, New York, NY, USA, 2008. ACM.
10. N. Nystrom, M. R. Clarkson, and A. C. Myers. Polyglot: An extensible compiler framework for java. In *Proc. of the International Conference on Compiler Construction (CC)*, pages 138–152, 2003.
11. W. Pugh. Skip lists: a probabilistic alternative to balanced trees. *Commun. ACM*, 33(6):668–676, 1990.
12. T. Riegel, P. Felber, and C. Fetzer. A lazy snapshot algorithm with eager validation. In *Proc. of the International Symposium on Distributed Computing (DISC)*, pages 284–298, Sep 2006.
13. T. Riegel, C. Fetzer, and P. Felber. Time-based transactional memory with scalable time bases. In *Proc. of the Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 221–228, New York, NY, USA, 2007. ACM.