

# Is Transactional Programming Actually Easier?

Christopher J. Rossbach, Owen S. Hofmann, and Emmett Witchel  
*Department of Computer Science, University of Texas at Austin*  
{rossbach,osh,witchel}@cs.utexas.edu

## Abstract

Chip multi-processors (CMPs) have become ubiquitous, while tools that ease concurrent programming have not. The promise of increased performance for all applications through ever more parallel hardware requires good tools for concurrent programming, especially for average programmers. Transactional memory (TM) has enjoyed recent interest as a tool that can help programmers program concurrently.

The TM research community claims that programming with transactional memory is easier than alternatives (like locks), but evidence is scant. In this paper, we describe a user-study in which 147 undergraduate students in an operating systems course implemented the same programs using coarse and fine-grain locks, monitors, and transactions. We surveyed the students after the assignment, and examined their code to determine the types and frequency of programming errors for each synchronization technique. Inexperienced programmers found baroque syntax a barrier to entry for transactional programming. On average, subjective evaluation showed that students found transactions harder to use than coarse-grain locks, but slightly easier to use than fine-grained locks. Detailed examination of synchronization errors in the students' code tells a rather different story. Overwhelmingly, the number and types of programming errors the students made was much lower for transactions than for locks. On a similar programming problem, over 70% of students made errors with fine-grained locking, while less than 10% made errors with transactions.

## 1 Introduction

Transactional memory (TM) has enjoyed a wave of attention from the research community. The increasing ubiquity of chip multiprocessors has resulted in a high availability of parallel hardware resources, without many concurrent programs. TM researchers position TM as an enabling technology for concurrent programming for the “average” programmer.

Transactional memory allows the programmer to delimit regions of code that must execute atomically and in

isolation. It promises the performance of fine-grain locking with the code simplicity of coarse-grain locking. In contrast to locks, which use mutual exclusion to serialize access to critical sections, TM is typically implemented using optimistic concurrency techniques, allowing critical sections to proceed in parallel. Because this technique dramatically reduces serialization when dynamic read-write and write-write sharing is rare, it can translate directly to improved performance without additional effort from the programmer. Moreover, because transactions eliminate many of the pitfalls commonly associated with locks (e.g. deadlock, convoys, poor composability), transactional programming is touted as being easier than lock based programming.

Evaluating the ease of transactional programming relative to locks is largely uncharted territory. Naturally, the question of whether transactions are easier to use than locks is qualitative. Moreover, since transactional memory is still a nascent technology, the only available transactional programs are research benchmarks, and the population of programmers familiar with both transactional memory and locks for synchronization is vanishingly small.

To address the absence of evidence, we developed a concurrent programming project for students of an undergraduate Operating Systems course at the University of Texas at Austin, in which students were required to implement the same concurrent program using coarse and fine-grained locks, monitors, and transactions. We surveyed students about the relative ease of transactional programming as well as their investment of development effort using each synchronization technique. Additionally, we examined students' solutions in detail to characterize and classify the types and frequency of programming errors students made with each programming technique.

This paper makes the following contributions:

- A project and design for collecting data relevant to the question of the relative ease of programming with different synchronization primitives.
- Data from 147 student surveys that constitute the first (to our knowledge) empirical data relevant to the question of whether transactions are, in fact, easier to use than locks.

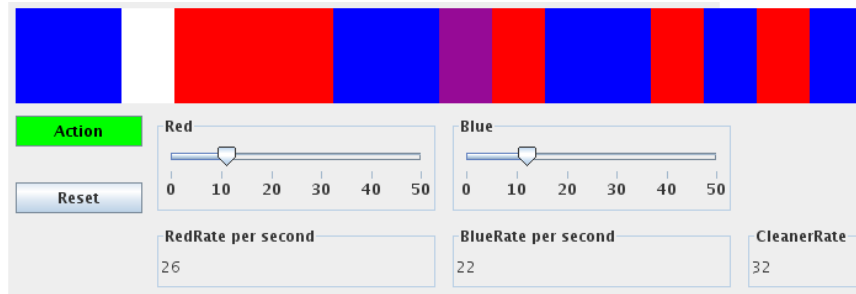


Figure 1: A screen-shot of sync-gallery, the program undergraduate OS students were asked to implement. In the figure the colored boxes represent 16 shooting lanes in a gallery populated by shooters, or *rogues*. A red or blue box represents a box in which a rogue has shot either a red or blue paint ball. A white box represents a box in which no shooting has yet taken place. A purple box indicates a line in which both a red and blue shot have occurred, indicating a race condition in the program. Sliders control the rate at which shooting and cleaning threads perform their work.

- A taxonomy of synchronization errors made with different synchronization techniques, and a characterization of the frequency with which such errors occur in student programs.

## 2 Sync-gallery

In this section, we describe sync-gallery, the Java programming project we assigned to students in an undergraduate operating systems course. The project is designed to familiarize students with concurrent programming in general, and with techniques and idioms for using a variety of synchronization primitives to manage data structure consistency. Figure 1 shows a screen shot from the sync-gallery program.

The project asks students to consider the metaphor of a shooting gallery, with a fixed number of lanes in which *rogues* (shooters) can shoot in individual lanes. Being pacifists, we insist that shooters in this gallery use red or blue paint balls rather than bullets. Targets are white, so that lanes will change color when a rogue has shot in one. Paint is messy, necessitating *cleaners* to clean the gallery when all lanes have been shot. Rogues and cleaners are implemented as threads that must check the state of one or more lanes in the gallery to decide whether it is safe to carry out their work. For rogues, this work amounts to shooting at some number of randomly chosen lanes. Cleaners must return the gallery to its initial state with all lanes white. The students must use various synchronization primitives to enforce a number of program invariants:

1. **Only one rogue may shoot in a given lane at a time.**
2. **Rogues may only shoot in a lane if it is white.**
3. **Cleaners should only clean when all lanes have been shot (are non-white).**
4. **Only one thread can be engaged in the process of cleaning at any given time.**

If a student writes code for a rogue that fails to respect the first two invariants, the lane can be shot with both red and blue, and will therefore turn purple, giving the student instant visual feedback that a race condition exists in the program. If the code fails to respect to the second two invariants, no visual feedback is given (indeed these invariants can only be checked by inspection of the code in the current implementation).

We ask the students to implement 9 different versions of rogues (Java classes) that are instructive for different approaches to synchronization. Table 1 summarizes the rogue variations. Gaining exclusive access to one or two lanes of the gallery in order to test the lane's state and then modify it corresponds directly to the real-world programming task of locking some number of resources in order to test and modify them safely in the presence of concurrent threads.

### 2.1 Locking

We ask the students to synchronize rogue and cleaner threads in the sync-gallery using locks to teach them about coarse and fine-grain locking. To ensure that students write code that explicitly performs locking and unlocking operations, we require them to use the Java `ReentrantLock` class and do not allow use of the `synchronized` keyword. In locking rogue variations, cleaners do not use dedicated threads; the rogue that colors the last white lane in the gallery is responsible for becoming a cleaner and subsequently cleaning all lanes. There are four variations on this rogue type: **Coarse**, **Fine**, **Coarse2** and **Fine2**. In the coarse implementation, students are allowed to use a single global lock which is acquired before attempting to shoot or clean. In the fine-grain implementation, we require the students to implement individual locks for each lane. The Coarse2 and Fine2 variations require the same mapping of locks to ob-

<pre> final int x = 10; Callable c = new Callable&lt;Void&gt; {     public Void call() {         // txn1 code         y = x * 2;         return null;     } } Thread.doIt(c); </pre>	<pre> Transaction tx = new Transaction(id); boolean done = false; while(!done) {     try {         tx.BeginTransaction();         // txn1 code         done = tx.CommitTransaction();     } catch(AbortException e) {         tx.AbortTransaction();         done = false;     } } </pre>
--	---

Figure 2: Examples of (left) DSTM2 concrete syntax, and (right) JDASTM concrete syntax.

jects in the gallery as their counterparts above, but introduce the additional stipulation that rogues must acquire access to and shoot at two random lanes rather than one. The pedagogical value is illustration that fine-grain locking requires a lock-ordering discipline to avoid deadlock, while a single coarse lock does not. Naturally, the use of fine grain lane locks complicates the enforcement of invariants 3 and 4 above.

## 2.2 Monitor implementations

Students must use condition variables along with signal/wait to implement both fine and coarse locking versions of the rogue programs. These two variations introduce dedicated threads for cleaners: shooters and cleaners must use condition variables to coordinate shooting and cleaning phases. In the coarse version (**CoarseCleaner**), students use a single global lock, while the fine-grain version (**FineCleaner**) requires per-lane locks.

## 2.3 Transactions

Finally, the students are asked to implement 3 TM-based variants of the rogues that share semantics with some locking versions, but use transactional memory for synchro-

nization instead of locks. The most basic TM-based rogue, **TM**, is analogous to the Coarse and Fine versions: rogue and cleaner threads are not distinct, and shooters need shoot only one lane, while the **TM2** variation requires that rogues shoot at two lanes rather than one. In the **TM-Cleaner**, rogues and cleaners have dedicated threads. Students can rely on the TM subsystem to detect conflicts and restart transactions to enforce all invariants, so no condition synchronization is required.

## 2.4 Transactional Memory Support

Since sync-gallery is a Java program, we were faced with the question of how to support transactional memory. The ideal case would have been to use a software transactional memory (STM) that provides support for atomic blocks, allowing students to write transactional code of the form:

```

void shoot() {
    atomic {
        Lane l = getLane(rand());
        if(l.getColor() == WHITE)
            l.shoot(this.color);
    }
}

```

Rogue name	Technique	R/C Threads	Additional Requirements
<b>Coarse</b>	Single global lock	not distinct.	
<b>Coarse2</b>	Single global lock	not distinct	rogues shoot at 2 random lanes
<b>CoarseCleaner</b>	Single global lock, conditions	distinct	conditions, wait/notify
<b>Fine</b>	Per lane locks	not distinct	
<b>Fine2</b>	Per lane locks	not distinct	rogues shoot at 2 random lanes
<b>FineCleaner</b>	Per lane locks, conditions	distinct	conditions, wait/notify
<b>TM</b>	TM	not distinct	
<b>TM2</b>	TM	not distinct	rogues shoot at 2 random lanes
<b>TMCleaner</b>	TM	distinct	

Table 1: The nine different rogue implementations required for the sync-gallery project. The technique column indicates what synchronization technique was required. The R/C Threads column indicates whether coordination was required between dedicated rogue and cleaner threads or not. A value of “distinct” means that rogue and cleaner instances run in their own thread, while a value of “not distinct” means that the last rogue to shoot an empty (white) lane is responsible for cleaning the gallery.

No such tool is yet available; implementing compiler support for atomic blocks, or use of a source-to-source compiler such as spoon [1] were considered out-of-scope for the project. The trade-off is that students are forced to deal directly with the concrete syntax of our TM implementation, and must manage read and write barriers explicitly. We assigned the lab to 4 classes over 2 semesters. During the first semester both classes used DSTM2 [14]. For the second semester, both classes used JDASTM [24].

The concrete syntax has a direct impact on ease of programming, as seen in Figure 2. Both examples pepper the actual data structure manipulation with code that explicitly manages transactions. We replaced DSTM2 in the second semester because we felt that JDASTM syntax was somewhat less baroque and did not require students to deal directly with programming constructs like generics. Also, DSTM2 binds transactional execution to specialized lock version first. Students then either completed the fine-grained or TM version second, depending on their assigned group. However, both DSTM2 and JDASTM require explicit read and write barrier calls for transactional reads and writes.

### 3 Methodology

Students completed the sync-gallery program as a programming assignment as part of several operating systems classes at the University of Texas at Austin. In total, 147 students completed the assignment, spanning two sections each in classes from two different semesters of the course. The semesters were separated by a year. We provided an implementation of the shooting gallery, and asked students to write the rogue classes described in the previous sections, respecting the given invariants.

We asked students to record the amount of time they spent designing, coding, and debugging each programming task (rogue). We use the amount of time spent on each task as a measure of the difficulty that task presented to the students. This data is presented in Section 4.1. After completing the assignment, students rated their famil-

ilarity with concurrent programming concepts prior to the assignment. Students then rated their experience with the various tasks, ranking synchronization methods with respect to ease of development, debugging, and reasoning (Section 4.2).

While grading the assignment, we recorded the type and frequency of synchronization errors students made. These are the errors still present in the student’s final version of the code. We use the frequency with which students made errors as another metric of the difficulty of various synchronization constructs.

To prevent experience with the assignment as a whole from influencing the difficulty of each task, we asked students to complete the tasks in different orders. In each group of rogues (single-lane, two-lane, and separate cleaner thread), students completed the coarse-grained first. Students then either completed the fine-grained or TM version second, depending on their assigned group. We asked students to randomly assign themselves to groups based on hashes of their name. Due to an error, nearly twice as many students were assigned to the group completing the fine-grained version first. However, there were no significant differences in programming time between the two groups, suggesting that the order in which students implemented the tasks did not affect the difficulty of each task.

#### 3.1 Limitations

Perhaps the most important limitation of the study is the much greater availability of documentation and tutorial information about locking than about transactions. The novelty of transactional memory made it more difficult both to teach and learn. The concrete syntax of transactions is also a barrier to ease of understanding and use (see §4.2). Lectures about locking drew on a larger body of understanding that has existed for a longer time. It is unlikely that students from one year influenced students from the

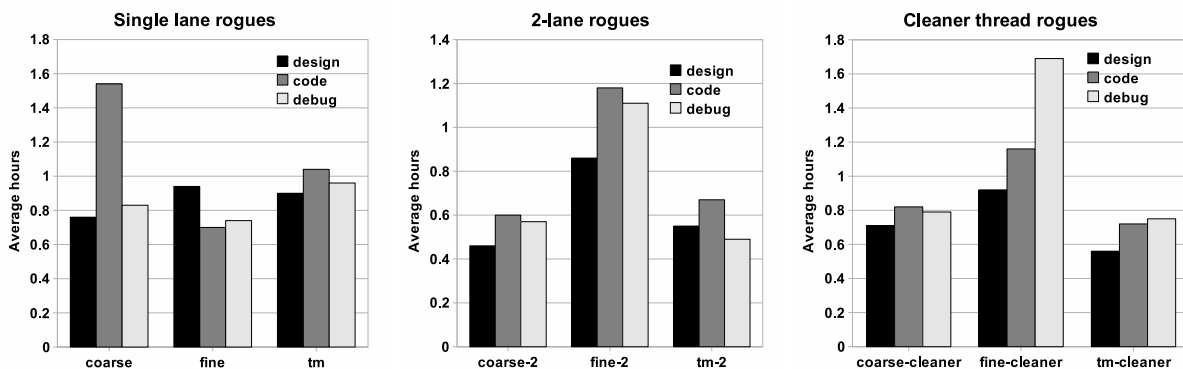


Figure 3: Average design, coding, and debugging time spent for analogous rogue variations.

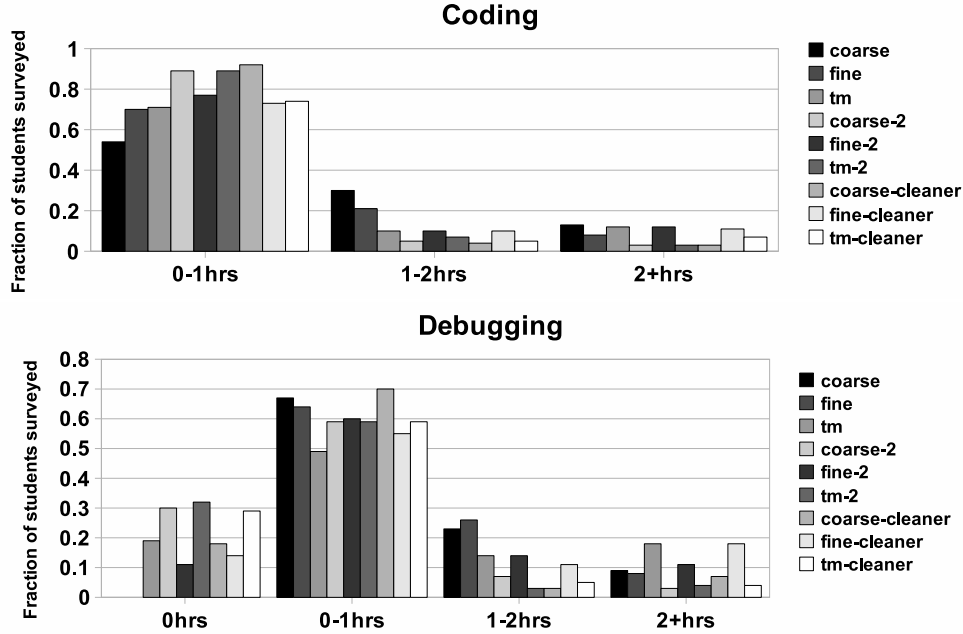


Figure 4: Distributions for the amount of time students spent coding and debugging, for all rogue variations.

next year given the difference in concrete syntax between the two courses.

## 4 Evaluation

We examined development time, user experiences, and programming errors to determine the difficulty of programming with various synchronization primitives. In general, we found that a single coarse-grained lock had similar complexity to transactions. Both of these primitives were less difficult, caused fewer errors, and had better student responses than fine-grained locking.

### 4.1 Development time

Figures 4 and 3 characterize the amount of time the students spent designing, coding and debugging with each synchronization primitive. On average, transactional memory required more development time than coarse locks, but less than required for fine-grain locks and condition synchronization. With more complex synchronization tasks, such as coloring two lanes and condition synchronization, the amount of time required for debugging increases relative to the time required for design and coding (Figure 3).

We evaluate the statistical significance of differences in development time in Table 2. Using a Wilcoxon signed-rank test, we evaluated the alternative hypothesis on each pair of synchronization tasks that the row task required

less time than the column task. Pairs for which the signed-rank test reports a p-value of  $< .05$  are considered statistically significant, indicating that the row task required less time than the column. If the p-value is greater than  $.05$ , the difference in time for the tasks is not statistically significant or the row task required more time than the column task. Results for the different class years are separated due to differences in the TM part of the assignment (Section 2.4).

We found that students took more time to develop the initial tasks while familiarizing themselves with the assignment. Except for fine-grain locks, later versions of similar synchronization primitives took less time than earlier, e.g. the Coarse2 task took less time than the Coarse task. In addition, condition synchronization is difficult. For both rogues with less complex synchronization (Coarse and TM), adding condition synchronization increases the time required for development. For fine-grain locking, students simply replace one complex problem with a second, and so do not require significant additional time.

In both years, we found that coarse locks and transactions required less time than fine-grain locks on the more complex two-lane assignments. This echoes the promise of transactions, removing the coding and debugging complexity of fine-grain locking and lock ordering when more than one lock is required.

## 4.2 User experience

To gain insight into the students’ perceptions about the relative ease of using different synchronization techniques we asked the students to respond to a survey after completing the sync-gallery project. The survey ends with 6 questions asking students to rank their favorite technique with respect to ease of development, debugging, reasoning about, and so on.

A version of the complete survey can be viewed at [2].

In student opinions, we found that the more baroque syntax of the DSTM2 system was a barrier to entry for new transactional programmers. Figure 5 shows student responses to questions about syntax and ease of thinking about different transactional primitives. In the first class year, students found transactions more difficult to think about and had syntax more difficult than that of fine-grain locks. In the second year, when the TM implementation was replaced with one less cumbersome, student opinions aligned with our other findings: TM ranked behind coarse locks, but ahead of fine-grain. For both years, other questions on ease of design and implementation mirrored these results, with TM ranked ahead of fine-grain locks.

## 4.3 Synchronization Error Characterization

We examined the solutions from the second year’s class in detail to classify the types of synchronization errors students made along with their frequency. This involved both a thorough reading of every student’s final solutions and automated testing. While the students’ subjective evaluation of the ease of transactional programming does not

clearly indicate that transactional programming is easier, the types and frequency of programming errors does.

While the students showed an impressive level of creativity with respect to synchronization errors, we found that all errors fit within the taxonomy described below.

1. **Lock ordering (lock-ord)**. In fine-grain locking solutions, a program failed to use a lock ordering discipline to acquire locks, admitting the possibility of deadlock.
2. **Checking conditions outside a critical section (lock-cond)**. This type of error occurs when code checks a program condition with no locks held, and subsequently acts on that condition after acquiring locks. This was the most common error in sync-gallery, and usually occurred when students would check whether to clean the gallery with no locks held, subsequently acquiring lane locks and proceeding to clean. The result is a violation of invariant 4 (§2). This type of error may be more common because no visual feedback is given when it is violated (unlike races for shooting lanes, which can result in purple lanes).
3. **Forgotten synchronization (lock-forgot)**. This class of errors includes all cases where the programmer forgot to acquire locks, or simply did not realize that a particular region would require mutual exclusion to be correct.
4. **Exotic use of condition variables (cv-exotic)**. We encountered a good deal of signal/wait usage on condition variables that indicates no clear understanding of what the primitives actually do. The canonical example of this is signaling and waiting the same con-

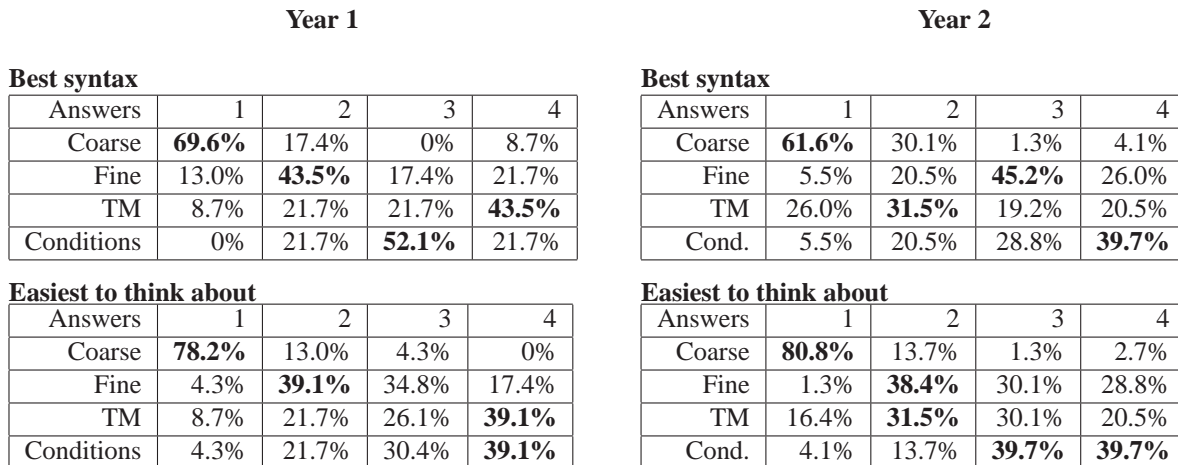


Figure 5: Selected results from student surveys. Column numbers represent rank order, and entries represent what percentage of students assigned a particular synchronization technique a given rank (e.g. 80.8% of students ranked Coarse locks first in the “Easiest to think about category”). In the first year the assignment was presented, the more complex syntax of DSTM made TM more difficult to think about. In the second year, simpler syntax alleviated this problem.

		Coarse	Fine	TM	Coarse2	Fine2	TM2	CoarseCleaner	FineCleaner	TMCleaner
Coarse	Y1	1.00	<b>0.03</b>	<b>0.02</b>	1.00	<b>0.02</b>	1.00	0.95	0.47	0.73
	Y2	1.00	0.33	0.12	1.00	0.38	1.00	1.00	0.18	1.00
Fine	Y1	0.97	1.00	0.33	1.00	0.24	1.00	1.00	0.97	0.88
	Y2	0.68	1.00	0.58	1.00	0.51	1.00	1.00	0.40	1.00
TM	Y1	0.98	0.68	1.00	1.00	0.13	1.00	1.00	0.98	0.92
	Y2	0.88	0.43	1.00	1.00	0.68	1.00	1.00	0.41	1.00
Coarse2	Y1	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>	1.00	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>
	Y2	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>	1.00	< <b>0.01</b>	0.45	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>
Fine2	Y1	0.98	0.77	0.87	1.00	1.00	1.00	1.00	1.00	0.98
	Y2	0.62	0.49	0.32	1.00	1.00	1.00	0.99	0.59	1.00
TM2	Y1	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>	0.99	< <b>0.01</b>	1.00	<b>0.04</b>	< <b>0.01</b>	< <b>0.01</b>
	Y2	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>	0.55	< <b>0.01</b>	1.00	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>
CoarseCleaner	Y1	0.05	< <b>0.01</b>	< <b>0.01</b>	1.00	< <b>0.01</b>	0.96	1.00	< <b>0.01</b>	0.08
	Y2	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>	1.00	< <b>0.01</b>	1.00	1.00	< <b>0.01</b>	0.96
FineCleaner	Y1	0.53	<b>0.03</b>	<b>0.02</b>	1.00	< <b>0.01</b>	1.00	0.99	1.00	0.46
	Y2	0.83	0.60	0.59	1.00	0.42	1.00	1.00	1.00	1.00
TMCleaner	Y1	0.28	0.12	0.08	1.00	<b>0.03</b>	1.00	0.92	0.55	1.00
	Y2	< <b>0.01</b>	< <b>0.01</b>	< <b>0.01</b>	0.99	< <b>0.01</b>	1.00	<b>0.04</b>	< <b>0.01</b>	1.00

Table 2: Comparison of time taken to complete programming tasks for all students. The time to complete the task on the row is compared to the time for the task on the column. Each cell contains p-values for a Wilcoxon signed-rank test, testing the hypothesis that the row task took less time than the column task. Entries are considered statistically significant when  $p < .05$ , meaning that the row task did take less time to complete than the column task, and are marked in bold. Results for first and second class years are reported separately, due to differing transactional memory implementations.

dition in the same thread.

5. **Condition variable use errors (cv-use).** These types of errors indicate a failure to use condition variables properly, but do indicate a certain level of understanding. This class includes use of `if` instead of `while` when checking conditions on a decision to wait, or failure to check the condition at all before waiting.
6. **TM primitive misuse (TM-exotic).** This class of error includes any misuse of transactional primitives. Technically, this class includes mis-use of the API, but in practice the only errors of this form we saw were failure to call `beginTransaction` before calling `endTransaction`. Omission of read/write barriers falls within this class as well, but it is interesting to note that we found no bugs of this form.

7. **TM ordering (TM-order).** This class of errors represents attempts by the programmer to follow some sort of locking discipline in the presence of transactions, where they are strictly unnecessary. Such errors do not result in an incorrect program, but do represent a misunderstanding of the primitive.
8. **Forgotten TM synchronization (TM-forgot).** Like the forgotten synchronization class above (lock-forgot), these errors occur when a programmer failed to recognize the need for synchronization and did not use transactions to protect a data structure.

Table 3 shows the characterization of synchronization for programs submitted in year 2. Figure 6 shows the overall portion of students that made an error on each programming task. Students were far more likely to make an error on fine-grain synchronization than on coarse or TM.

	lock-ord	lock-cond	lock-forgot	cv-exotic	cv-use	TM-exotic	TM-order	TM-forgot
<b>occurrences</b>	11	62	26	11	14	5	4	1
<b>opportunities</b>	134	402	402	134	134	201	201	201
<b>rate</b>	8.2%	6.5%	15.4%	8.2%	10.5%	2.5%	2.0%	0.5%

Table 3: Synchronization error rates for year 2. The occurrences row indicates the number of programs in which at least one bug of the type indicated by the column header occurred. The **opportunities** row indicates the sample size (the number of programs we examined in which that type of bug could arise: e.g. lock-ordering bugs cannot occur in with a single coarse lock). The **rate** column expresses the percentage of examined programs containing that type of bug. Bug types are explained in Section 4.3.




Figure 6: Overall error rates for programming tasks. Error bars show a 95% confidence interval on the error rate. Fine-grained locking tasks were more likely to contain errors than coarse-grained or transactional memory (TM).

About 70% of students made at least one error on the Fine and Fine2 portions of the assignment.

## 5 Related work

Hardware transactional memory research is an active research field with many competing proposals [4–7, 9–11, 15–17, 19–23, 26]. All this research on hardware mechanism is the cart leading the horse if researchers never validate the assumption that transactional programming is actually easier than lock-based programming.

This research uses software transactional memory (which has no shortage of proposals [3, 12–14, 18, 25]), but its purpose is to validate how untrained programmers learn to write correct and perform concurrent programs with locks and transactions. The programming interface for STM systems is the same as HTM systems, but without compiler support, STM implementations require explicit read-write barriers, which are not required in an HTM. Compiler integration is easier to program than using a TM library [8]. Future work research could investigate whether compiler integration lowers the perceived programmer difficulty in using transactions.

## 6 Conclusion

To our knowledge, no previous work directly addresses the question of whether transactional memory actually delivers on its promise of being easier to use than locks. This paper offers evidence that transactional programming really is less error-prone than high-performance locking, even if newbie programmers have some trouble understanding transactions. Students subjective evaluation showed that they found transactional memory slightly harder to use than coarse locks, and easier to use than fine-grain locks and condition synchronization. However, analysis of synchronization error rates in students’ code yields a more dramatic result, showing that for similar programming tasks, transactions are considerably easier to get correct than locks.

## References

- [1] *Spoon*, 2009. <http://spoon.gforge.inria.fr/>.
- [2] *Sync-gallery survey*: <http://www.cs.utexas.edu/users/witchel/tx/sync-gallery-survey.html>, 2009.
- [3] A.-R. Adl-Tabatabai, B. Lewis, V. Menon, B. Murphy, B. Saha, and T. Shpeisman. Compiler and runtime support for efficient software transactional memory. In *PLDI*, Jun 2006.
- [4] Lee Baugh, Naveen Neelakantam, and Craig Zilles. Using hardware memory protection to build a high-performance, strongly atomic hybrid transactional memory. In *Proceedings of the 35th*



- Annual International Symposium on Computer Architecture*. Jun 2008.
- [5] Colin Blundell, Joe Devietti, E. Christopher Lewis, and Milo M. K. Martin. Making the fast case common and the uncommon case simple in unbounded transactional memory. *SIGARCH Comput. Archit. News*, 35(2):24–34, 2007.
- [6] Jayaram Bobba, Neelam Goyal, Mark D. Hill, Michael M. Swift, and David A. Wood. Tokentm: Efficient execution of large transactions with hardware transactional memory. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*. Jun 2008.
- [7] J. Chung, C. Minh, A. McDonald, T. Skare, H. Chafi, B. Carlstrom, C. Kozyrakis, and K. Olukotun. Tradeoffs in transactional memory virtualization. In *ASPLOS*, 2006.
- [8] Luke Dalessandro, Virendra J. Marathe, Michael F. Spear, and Michael L. Scott. Capabilities and limitations of library-based software transactional memory in c++. In *Proceedings of the 2nd ACM SIGPLAN Workshop on Transactional Computing*. Portland, OR, Aug 2007.
- [9] L. Yen et al. Logtm-SE: Decoupling hardware transactional memory from caches. In *HPCA*. 2007.
- [10] Mark Moir et. al. Experiences with a commercial processor supporting htm. *ASPLOS 2009*.
- [11] L. Hammond, V. Wong, M. Chen, B. Hertzberg, B. Carlstrom, M. Prabhu, H. Wijaya, C. Kozyrakis, and K. Olukotun. Transactional memory coherence and consistency. In *ISCA*, 2004.
- [12] T. Harris, M. Plesko, A. Shinnar, and D. Tarditi. Optimizing memory transactions. In *PLDI*, Jun 2006.
- [13] Tim Harris and Keir Fraser. Language support for lightweight transactions. In *OOPSLA*, pages 388–402, Oct 2003.
- [14] M. Herlihy, V. Luchangco, and M. Moir. A flexible framework for implementing software transactional memory. In *OOPSLA*, pages 253–262, 2006.
- [15] M. Herlihy and J. E. Moss. Transactional memory: Architectural support for lock-free data structures. In *ISCA*, May 1993.
- [16] Owen S. Hofmann, Christopher J. Rossbach, and Emmett Witchel. Maximal benefit from a minimal tm. *ASPLOS*.
- [17] Yossi Lev and Jan-Willem Maessen. Split hardware transactions: true nesting of transactions using best-effort hardware transactional memory. In *PPoPP '08: Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming*, pages 197–206, New York, NY, USA, 2008. ACM.
- [18] V. Marathe, M. Spear, C. Heriot, A. Acharya, D. Eisenstat, W. Scherer III, and M. Scott. Lowering the overhead of nonblocking software transactional memory. In *TRANSACT*, 2006.
- [19] A. McDonald, J. Chung, B. Carlstrom, C. Minh, H. Chafi, C. Kozyrakis, and K. Olukotun. Architectural semantics for practical transactional memory. In *ISCA*, Jun 2006.
- [20] K. E. Moore, J. Bobba, M. J. Moravan, M. D. Hill, , and D. A. Wood. Logtm: Log-based transactional memory. In *HPCA*, 2006.
- [21] R. Rajwar and M. Herlihy K. Lai. Virtualizing transactional memory. In *ISCA*, Jun 2005.
- [22] H. Ramadan, C. Rossbach, D. Porter, O. Hofmann, A. Bhandari, and E. Witchel. Metatm/tlinux: Transactional memory for an operating system. In *ISCA*, 2007.
- [23] H. Ramadan, C. Rossbach, and E. Witchel. Dependence-aware transactional memory for increased concurrency. In *MICRO*, 2008.
- [24] Hany E. Ramadan, Indrajit Roy, Maurice Herlihy, and Emmett Witchel. Committing conflicting transactions in an STM. In *PPoPP*, 2009.
- [25] Nir Shavit and Dan Touitou. Software transactional memory. In *Proceedings of the 14th ACM Symposium on Principles of Distributed Computing*, pages 204–213, Aug 1995.
- [26] Arrvindh Shriraman, Sandhya Dwarkadas, and Michael L. Scott. Flexible decoupled transactional memory support. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*. Jun 2008.