

ERASER : A DYNAMIC DATA RACE DETECTOR FOR MULTITHREADED PROGRAMS


BY STEFAN SAVAGE ET AL.



KRISH

9/1/2011

MOTIVATION



- Only parallel programs can benefit from today's multi-core processors
- Multithreading has become a common(important) programming technique
- Synchronization between threads is a challenge
 - Synchronization errors are easy to cause and hard to debug
 - DATA RACE
 - Tracking errors take weeks and months


DATA RACE : NON-DETERMINISM

When two concurrent threads access a shared variable

1. at least one is a write and
2. the threads use no explicit synchronization to prevent simultaneous access

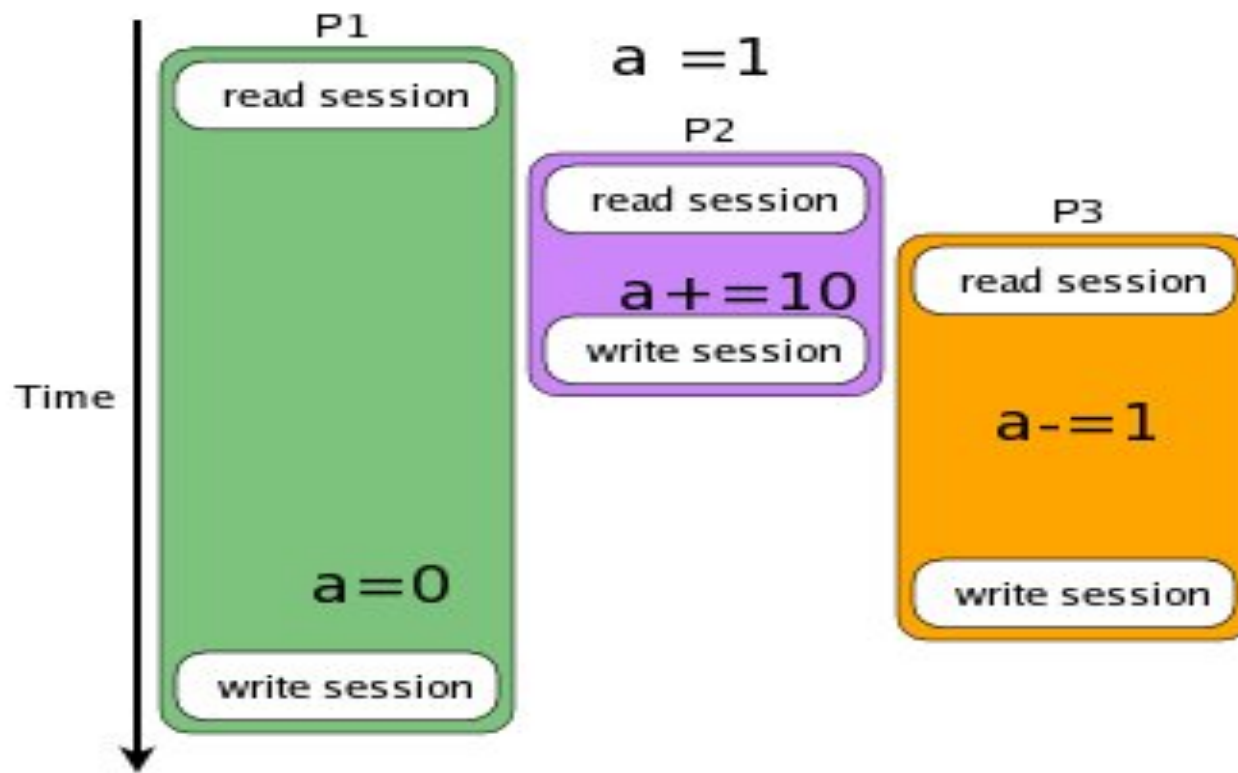
then the execution will depend on interleaving

ASCIIVille by Todd Presta



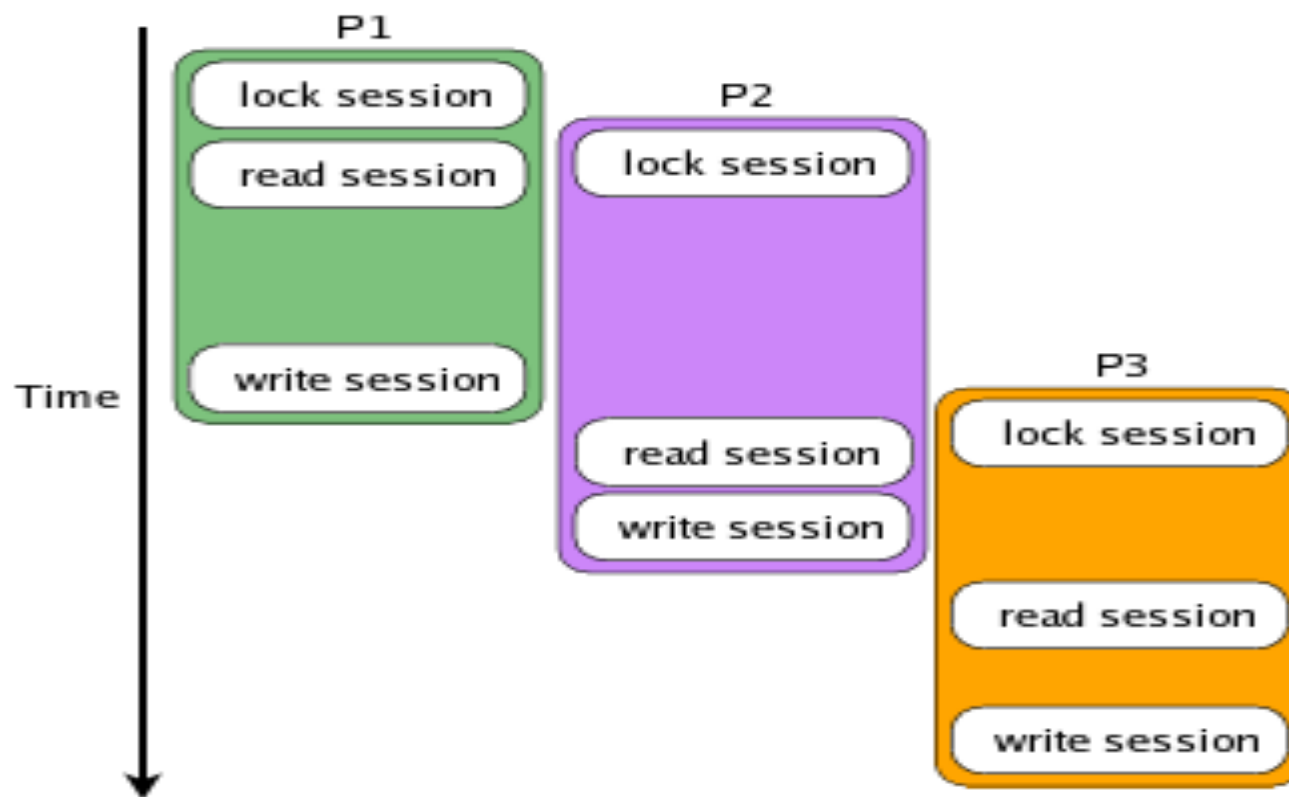
© 2008. Todd Presta. All rights reserved. <http://www.asciiville.com>

DATA RACE : NON-DETERMINISM



'a'-----??????

DATA RACE : NON-DETERMINISM



Locks are used to avoid data race; Let's see a tool that detect's race.

COMPLEXITY IN DATA RACE DETECTION



- Data Race detection is a NP complete problem
- For t threads of n instructions, the number of possible orders is about t^{n*t} .
- A thorough detection will involve examining all the possible order to make sure there exist only one order.
- Practical race detection tools are based on heuristics - so that they can detect maximum number of races, with in limited computation

ERASER: A DYNAMIC RACE DETECTOR FOR MULTI- THREADED PROGRAMS

Stefan Savage

University of Washington

Michael Burrows, Greg Nelson, Patrick Sobalvarro

Digital System Research Center

Thomas Anderson

University of Washington

SOSP' 1997

ACM Transactions on Computer Systems, Nov. 1997.

Cited by 876

OBJECTIVE OF THE WORK

The work presents the theory, implementation and experience of a testing tool that detect dynamic data race in multi-threaded programs




AUTHOR'S CLAIMS



Does not identify all races in the program, but for programs using lock-based synchronization ensures better results than Lamport's (previous) work

OUTLINE

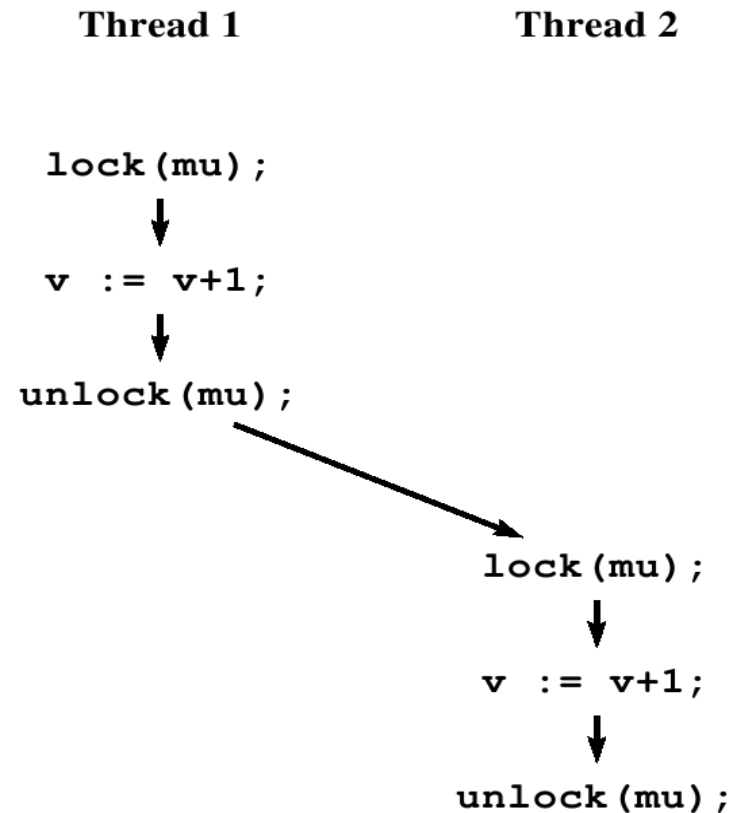


- Previous Work
- The Lockset algorithm
- Eraser implementation
- Experiences

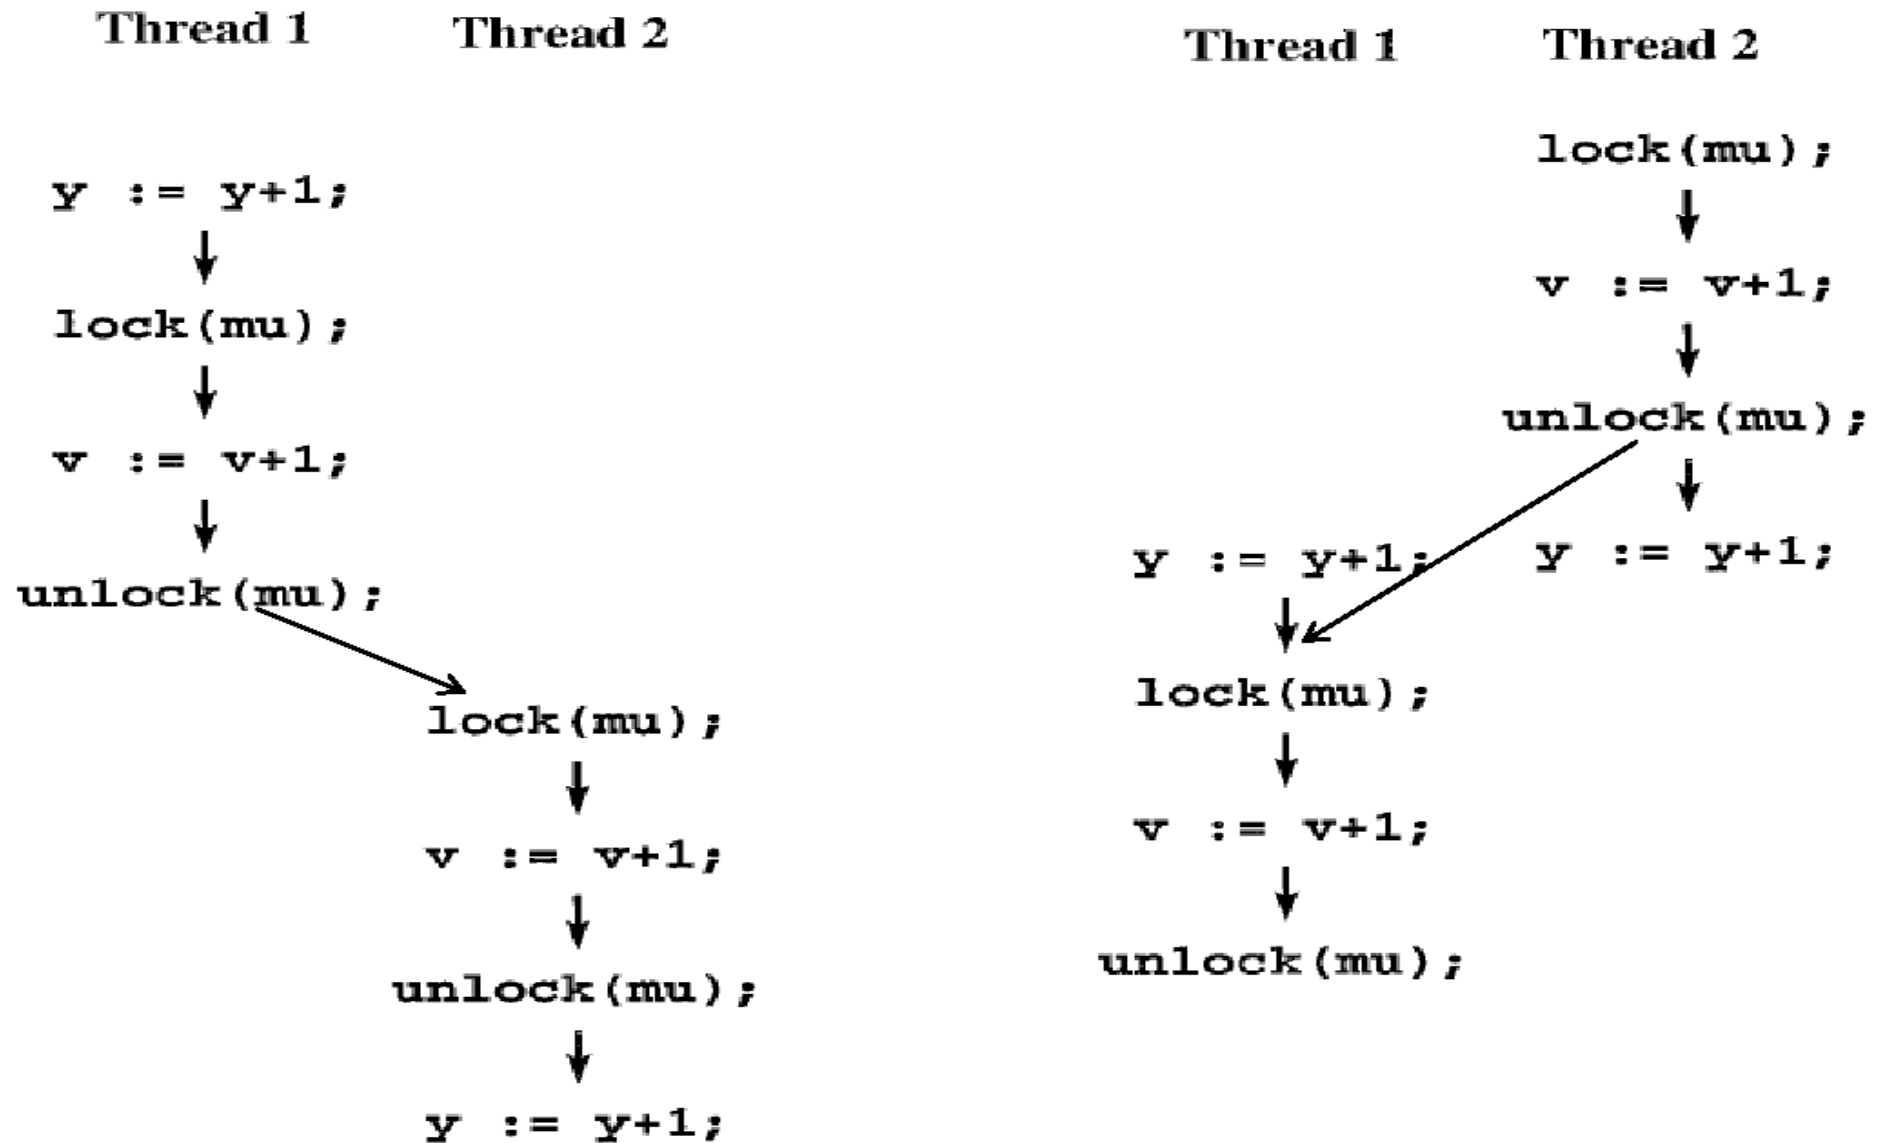
LAMPORT'S HAPPENES-BEFORE

The happens-before order is a **partial order** on all events of all threads in a concurrent execution,

- **Single Thread** - events are ordered by their occurrence.
- **Between threads** - events are ordered by the synchronization objects they access.



HAPPENS-BEFORE FAILS



OTHER ISSUES

- Difficult to implement efficiently - need per-thread information about ordering to all shared memory locations.
- Highly dependent on scheduler - needs large number of test cases.



ERASER'S APPROACH

Eraser uses **binary rewriting techniques** to monitor every shared memory reference and verify that **consistent locking behavior** is observed.

Heuristics – Consistent Locking Behavior



ERASER'S APPROACH


Eraser uses **binary rewriting techniques** to monitor every shared memory reference and verify that **consistent locking behavior** is observed.

'**binary rewriting**' – Observe all the Load and Store instructions.

'**consistent locking behavior**' - All instance of a shared variable **v** is locked by same set of locks **L(v)**

LOCKSET ALGORITHM

- ◆ $v = \text{shared variable}$
- ◆ $C(v) = \text{candidate locks for } v$
- ◆ $\text{locks_held}(t) = \text{set of locks held by thread } t$
- ◆ A lock l is in $C(v)$ if all threads hold l while accessing v
- ◆ A new variable at initialization is supposed to have all possible locks in $C(v)$



TRaversing LOCKSET ALGORITHM

<i>Program</i>	<i>locks_held</i>	<i>C(v)</i>
int v; v := 1024;	{}	{mu1, mu2}
lock(mu1); v := v + 1;	{mu1}	{mu1}
unlock(mu1); lock(mu2); v := v + 1;	{}	{mu1}
unlock(mu2);	{mu2}	{}
	{}	

INCREASING THE COMFORT ZONE

Initialization - Initialization can be done without holding a lock.

Read-shared data - Multiple reads safely accessed without locks.

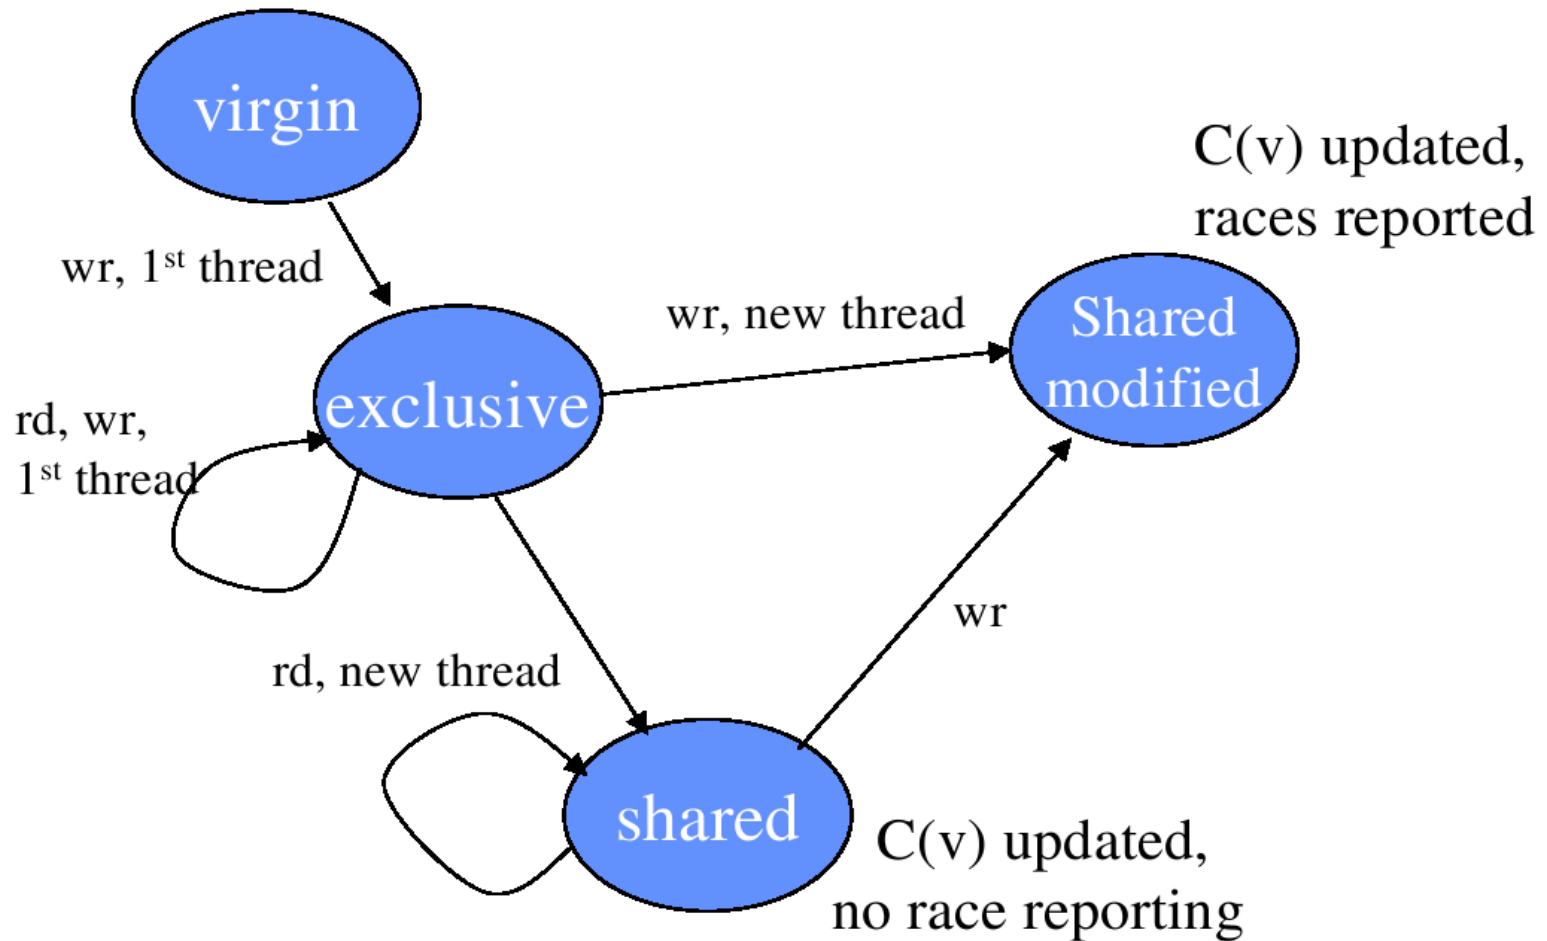
Read-write locks - Multiple readers, but allow only a single writer.



INITIALIZATION AND READ-SHARING

- A variable is considered initialized when it is accessed by second thread
- Simultaneous reads of shared variable are not races
- Report races only after an initialized variable become write shared

PER-LOCATION STATE

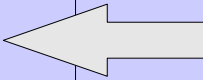


EXTENDING LOCKSET ALGORITHM

- Read / Write locking modes
- Locks held in read mode are removed from $C(v)$ when a write occur
- *On each read of v by thread t ,*
 - *set $C(v) := C(v) \cap \text{locks_held}(t)$;*
 - *if $C(v) = \{\}$, then issue a warning.*
- *On each write of v by thread t ,*
 - *set $C(v) := C(v) \cap \text{write_locks_held}(t)$;*
 - *if $C(v) = \{\}$, then issue a warning.*

TRaversing LOCKSET ALGORITHM

	<i>Program</i>	<i>locks_held</i>	<i>C(v)</i>	<i>State(v)</i>
T1	int v; v := 1024;	{}	{mu1, mu2}	Virgin Exclusive
T2	lock(mu1); v := v + 1; unlock(mu1);	{mu1}	{mu1}	Shared Shared-Modified
T1	lock(mu2); v := v + 1; unlock(mu2);	{mu2}	{}	

Race detected correctly
 

IMPLEMENTATION

Eraser's binary modification involves


1. Calls to storage allocator initializes $C(v)$
2. Each load and store updates $C(v)$
3. Each acquire or release call updates $locks_held(t)$



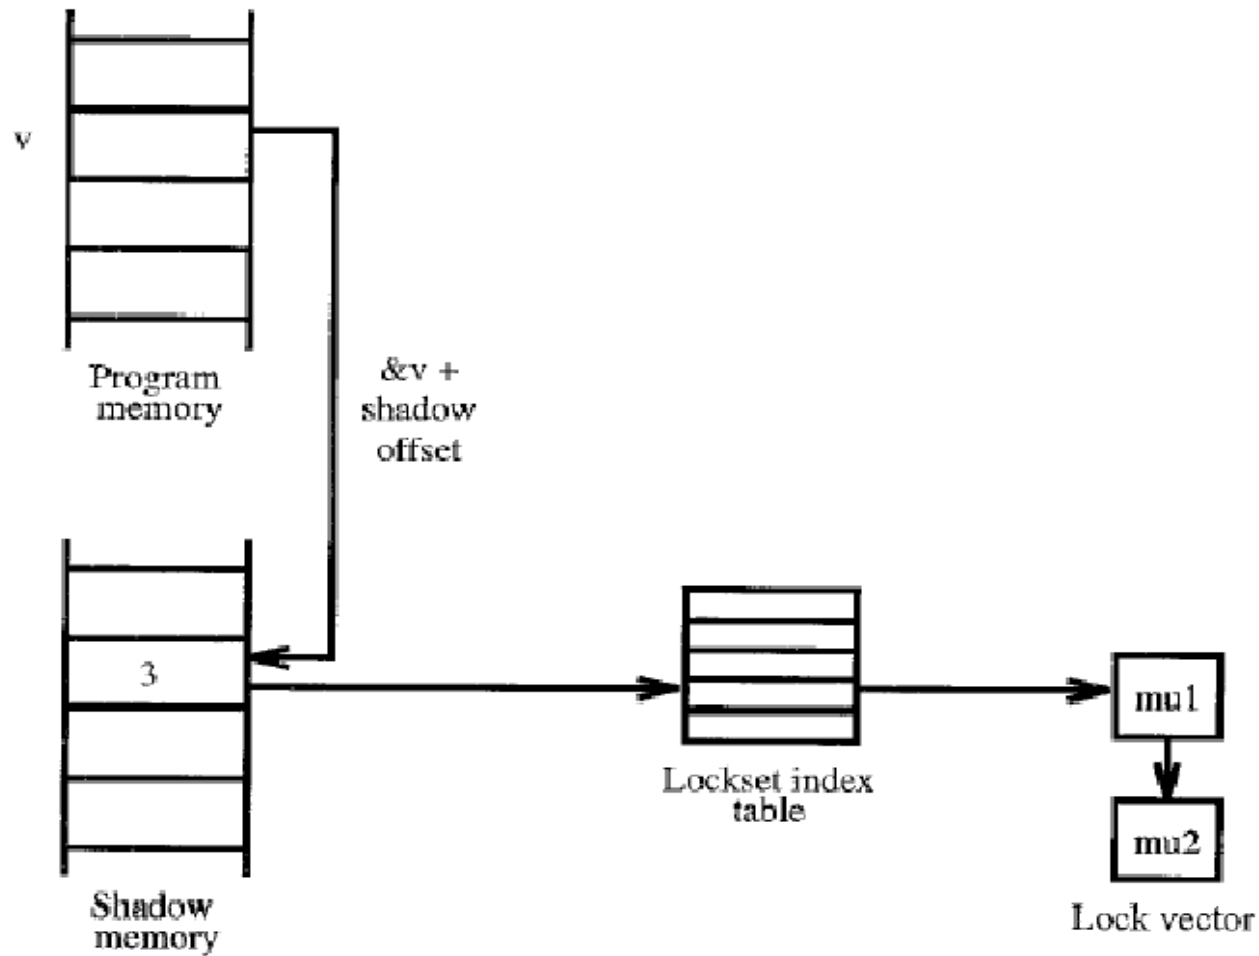
IMPLEMENTATION

Data Structure,

1. Maintains hash table of sets of locks
2. Represents each set of locks with an index.
3. Every shared memory location has shadow memory containing lockset index and state.
4. Shadow memory is located by adding offset to shared memory location address.



IMPLEMENTATION



FALSE POSITIVES

- **Memory reuse**
Caused by memory reset, with out resetting the shadow memory.
- **Private locks**
When Locks that are not part of standard pthread interface are used.
- **Benign races**
True Data race that did not affect the execution of the program.

© Original Artist
Reproduction rights obtainable from
www.CartoonStock.com



"I have a positive attitude...
but it might be a false positive."

ANNOTATIONS TO AVOID FALSE POSITIVES

1. For memory reuse

1. `EraserReuse(address, size)`

2. For private locks

1. `EraserReadLock(lock)`
2. `EraserReadUnlock(lock)`
3. `EraserWriteLock(lock)`
4. `EraserWriteUnlock(lock)`

3. For benign races

1. `EraserIgnoreOn()`
2. `EraserIgnoreOff()`



EXPERIENCE

1. **AltaVista**
 1. *Mhttpd http server* - 5,000 lines of C source code, 100 distinct locks, 9 annotations.
 2. *Ni2 indexing engine* - 20,000 lines of C source code, 900 distinct locks, 10 annotations.
2. **Vesta Cache Server** - 30,000 lines of C++ source code, 10 threads, 26 distinct locks, 10 annotations.
3. **Petal - distributed disk server** - 30,000 lines of C++ source code, 64 threads
4. **Undergrad coursework** – 100 multi threaded programs

EXPERIENCE

- Deliberately introduced race conditions were detected.
- Other data races were also detected.
- False alarms were raised, but use of annotations resolved sizable number of them.



EXPERIENCE



<i>Program</i>	<i>Serious races</i>	<i>Minor races</i>	<i>Benign races</i>
AltaVista		✓	✓
Vesta	✓	✓	
Petal		✓	
Undergrad assignments	✓		

PERFORMANCE

- NP-Complete – Computationally hard problem.
- Implementing Eraser slows down the application by a factor of 10 to 30
- Overhead of making a procedure call at every load and store instruction
- Performance was never a major goal



RECAP :OBJECTIVE OF THE WORK

The work presents the theory, implementation and experience of a testing tool that detect dynamic data race in multi-threaded programs

CONCLUSION

- Uses the locking principles as a heuristics to identify data races
- Successful implementation identifies potential races in enterprise software
- Implementing Eraser slows down the application by a factor of 10 to 30
- Diverse case study was given to support the tool



LET'S DISCUSS

