

Thread Safety for Third Party Libraries in Relaxed Memory Models

Ridhi Jain, IIIT-Delhi, India



Advisor:

Dr. Rahul Purandare, IIIT-Delhi,
India

In Collaboration with:

Dr. Subodh Sharma, IIT-Delhi, India



INDRAPRASTHA INSTITUTE *of*
INFORMATION TECHNOLOGY **DELHI**

Software crashes can be very expensive

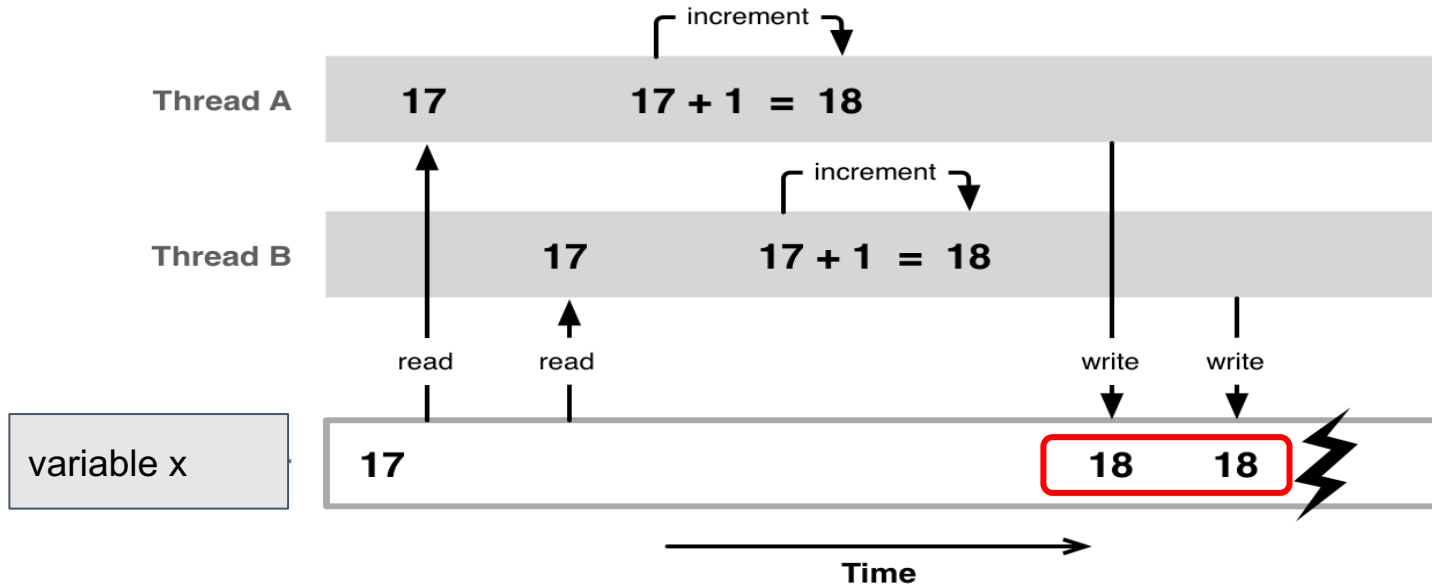
Failure of Mars Pathfinder

Medical Machine Kills, Therac-25

PayPal bug that made a \$92 quadrillion deposit



Example



Data Race on variable x

Weak Memory Models



Weaker than Sequentially Consistent memory model

To enhance performance

Total Store Order: A store followed by load in a same thread on two different memory locations can be reordered

Partial Store Order: A store followed by load or a store in a same thread on two different memory locations can be reordered

Relaxed Memory



Thread 1

```
(1) A = 1  
(2) print(B)
```

Thread 2

```
(3) B = 1  
(4) print(A)
```

```
assert( A==1 || B==1 )
```

Always true for
sequentially consistent
programs

May fail for relaxed
memory models

Problem Statement




Let me use Libraries to save my software construction cost, time and effort!





Problem Statement



- These third party libraries can be faulty.
- The bugs can be even subtle in a multithreaded environment.
- May be even worse in Relaxed Memory models.



Example



libpq

```
17863 # 800 "foo/include/c++/4.9/ostream" 2 3
17864 # 41 "foo/include/c++/4.9/ostream" 2 3
17865
17866 namespace std __attribute__((visibility__("default")))
17867 {
17868
17869 # 40 "foo/include/c++/4.9/ostream" 3
17870     extern ostream cin;
17871     extern ostream cout;
17872     extern ostream cerr;
17873     extern ostream clog;
17874
17875
17876     extern wostream wcin;
17877     extern wostream wcout;
17878     extern wostream wcerr;
17879     extern wostream wclog;
17880
17881
17882
17883
17884     static ios_base::Init __init;
17885
17886
17887 }
17888 # 3 "hello.cpp" 2
17889 using namespace std;
17890 int main()
17891 {
17892     cout<<"hello, world";
17893     return 0;
17894 }
```



Motivation



libpq: Manipulation of same connection object, leads to a crash

libcurl: Has no internal thread synchronization

libpng: Using same instance of a structure might lead to a crash.

Motivation



Library	Version	Thread Safe?
libcrypt		?
Expat		Yes
FreeTDS		?
FreeType		?
GD 1.8.x		?
GD 2.0.x		?
gdbm		No
ImageMagick	5.2.2	Yes
libjpeg	v6b	?
OpenSSL	0.9.6g	Yes

Figure 1: Thread safety status of sample libraries from [Apache HTTP Server 2.x Thread Safety Issues](#)

Research Question



**Can we detect the harmful data races in a third party library
in weaker memory models?**

Questions that Follow



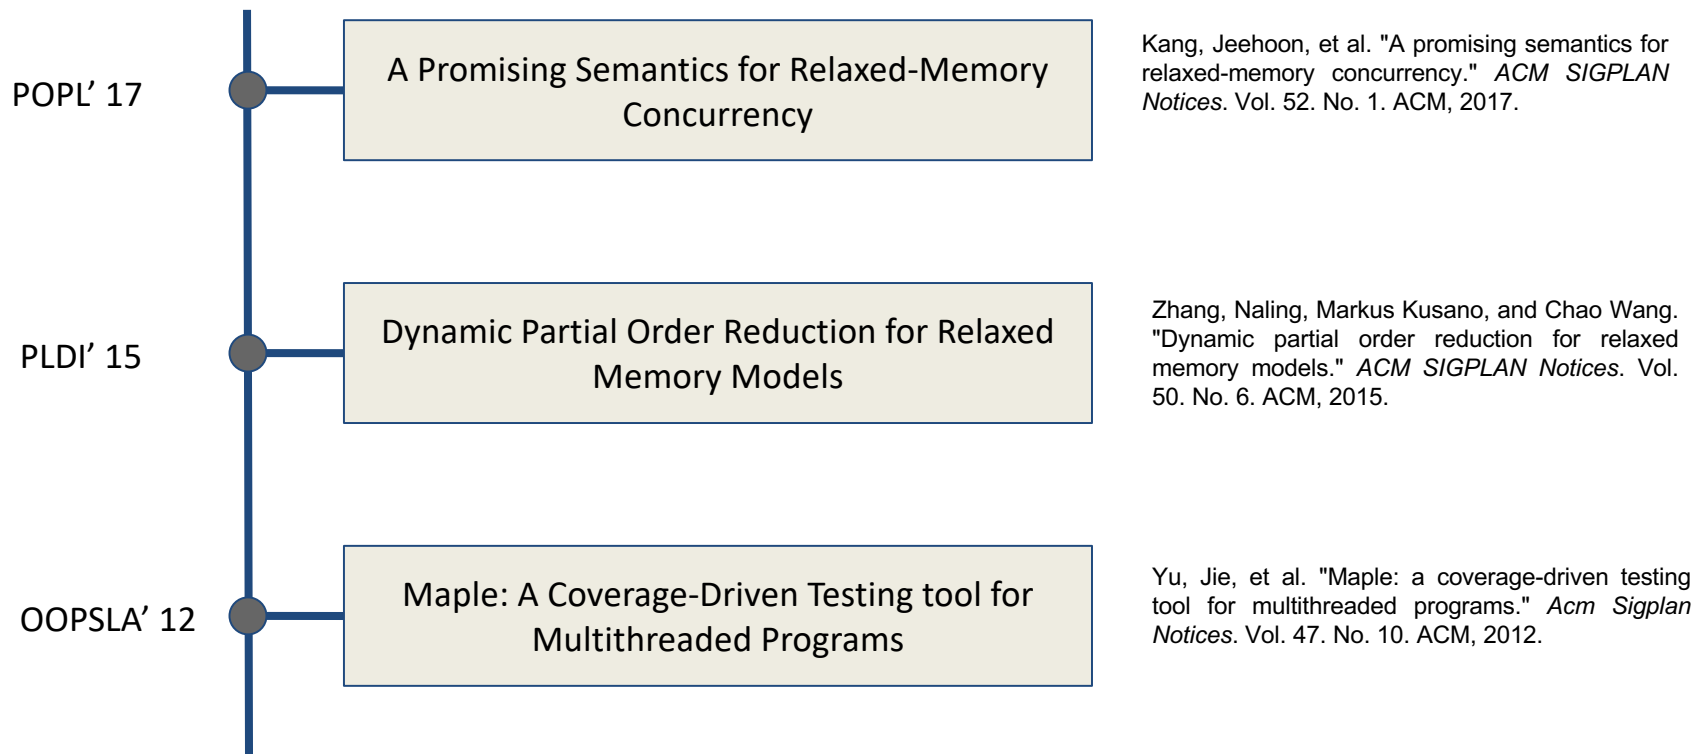
Are all the detected races harmful?

Can we detect data races on stripped binaries?


Can we detect the reorderable instructions in binaries?

Can we reorder stores and loads?

Related Work



System Overview



An Overview of the System

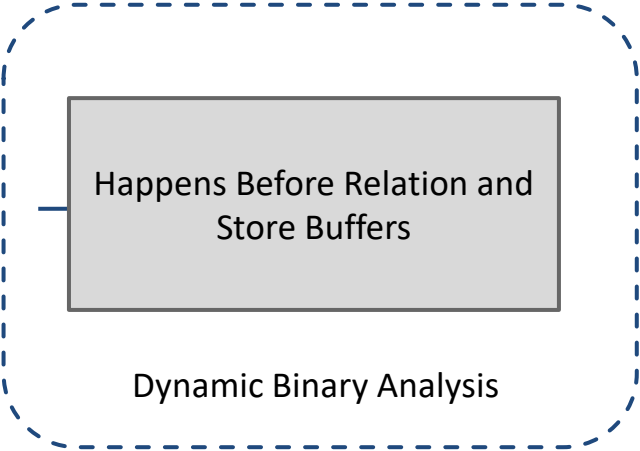
Detection of Data Races



Software Binary




Shared Variable Addresses



```
1_8_r_{1}_{1}_{[]}_{2_7_w}  
1_9_w_{1}_{1}_{[1_10_mov eax, dword ptr [rip+0x2007da]]}_{2_9_w}  
1_10_r_{1}_{1}_{[]}_{2_7_w}  
2_5_r_{2}_{2}_{[]}_{3_7_w}  
2_7_w_{2}_{2}_{[]}_{3_5_r,3_7_w,3_8_r,3_10_r}  
2_8_r_{2}_{2}_{[]}_{3_7_w}  
2_9_w_{2}_{2}_{[]}_{3_9_w}
```

Racing Instructions



Intermediate Results



- Tested on 16 libraries including boost, libcurl and libgpg.
- Detected data-races where **Intel Thread Checker** detected races, with additional reorderable instructions.

```
#include <boost/thread.hpp>
#include <boost/chrono.hpp>
#include <iostream>


void wait(int seconds)
{
    boost::this_thread::sleep_for(boost::chrono::seconds{seconds});
}

boost::mutex mutex;

void thread()
{
    using boost::this_thread::get_id;
    for (int i = 0; i < 5; ++i)
    {
        wait(i);
        mutex.lock();
        std::cout << "Thread " << get_id() << " : " << i << std::endl;
        mutex.unlock();
    }
}

int main()
{
    boost::thread t1(thread);
    boost::thread t2(thread);
    t1.join();
    t2.join();
}
```

Lib boost example



```
1_8_r_{1}_{1}_{[]}_{2_7_w}
1_9_w_{1}_{1}_{[1_10_mov eax, dword ptr [rip+0x2007da]]}_{2_9_w}
1_10_r_{1}_{1}_{[]}_{2_7_w}
2_5_r_{2}_{2}_{[]}_{3_7_w}
2_7_w_{2}_{2}_{[]}_{3_5_r,3_7_w,3_8_r,3_10_r}
2_8_r_{2}_{2}_{[]}_{3_7_w}
2_9_w_{2}_{2}_{[]}_{3_9_w}
```

Racing Instructions

Not all these races are harmful!!




Real Races (Scheduler)




Dynamic Partial Order Reduction is used to to replay the schedules

```
1_8_r_{1}_{1}_{[]}_{2_7_w}  
1_9_w_{1}_{1}_{[1_10_mov eax, dword ptr [rip+0x2007da]]}_{2_9_w}  
1_10_r_{1}_{1}_{[]}_{2_7_w}  
2_5_r_{2}_{2}_{[]}_{3_7_w}  
2_7_w_{2}_{2}_{[]}_{3_5_r,3_7_w,3_8_r,3_10_r}  
2_8_r_{2}_{2}_{[]}_{3_7_w}  
2_9_w_{2}_{2}_{[]}_{3_9_w}
```

**Racing
Instructions**



Scheduler



Real Races

Challenges



Binaries are to be analysed and dynamically controlled.

Debug information might not be available.

The implementation of application invoking library may be faulty.

Scheduling interleavings



Evaluation

Detecting all Data Races

Detecting shared and global variables

Thank You