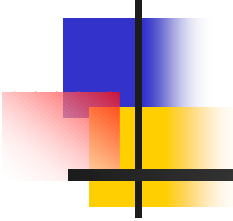


# Scaling Model Checking of Dataraces Using Dynamic Information



---

Ohad Shacham  
Tel Aviv University and IBM

Mooly Sagiv  
Tel Aviv University

Assaf Schuster  
Technion



# Datarace

---

- Happens when two threads access a memory location concurrently
  - At least one access is a write
- Unpredictable results
- Can indicate bugs
- Hard to detect
- Hard to reproduce



# Datarace example

---

```
TicketPurchase(NumOfTickets)
{
    if (NumOfTickets = FreeTickets)
        FreeTickets -= NumOfTickets
    else
        Print "Full";
}
```



# Datarace example

---

{FreeTickets = 4}

## Thread I

TicketPurchase(2)

if (NumOfTickets = FreeTickets)

FreeTickets -= NumOfTickets

## Thread II

TicketPurchase(4)

if (NumOfTickets = FreeTickets)

FreeTickets -= NumOfTickets

{FreeTickets = -2}




# Datarace example

---

```
TicketPurchase(NumOfTickets)
{
    Lock(lockFreeTickets)
    if (NumOfTickets = FreeTickets)
        FreeTickets -= NumOfTickets
    else
        Print "Full";
    Unlock(lockFreeTickets)
}
```

# Therac 25

- A medical radiation machine to treat cancer
- 6 patients got a radiation overdose
  - 4 died
  - 2 injured





# Datarace detection

---

- Static datarace detection tools
  - Racex
  - rccjava
- Dynamic datarace detection tools:
  - Lamport's *happens-before* partial order (Djit)
  - Lock based techniques (Lockset)

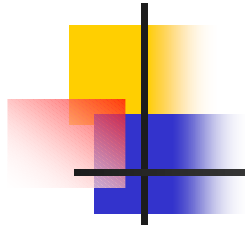


# Difficulties in model checking databases

---

- Infinite state space
- Huge number of interleavings
- Huge transition systems
- Size problem





# Observation

---

Dataraces free programs maintain  
a locking discipline



# Hybrid solution


---

- Dynamically check the locking discipline
- Produce witnesses for data races using a model checker
  - Explore suffixes of the trace

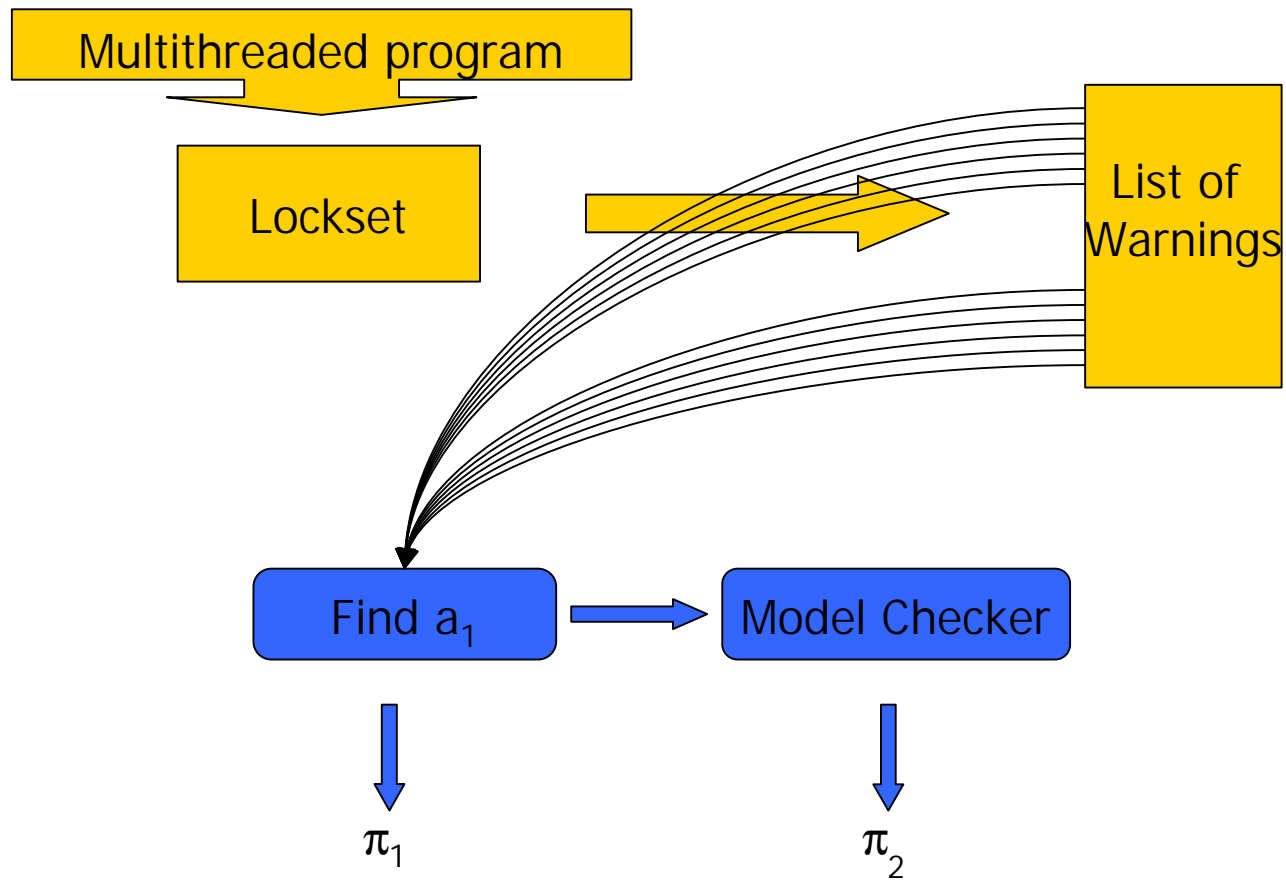


# Basic idea

---



# Algorithm flow





# Lockset invariant

---

Multiple accesses to a specific  
memory location are guarded by a  
unique lock



# Lockset example

---

## Thread I

Lock(lock<sup>x</sup>)      f  
X = 7              {lock<sup>x</sup>}  
Unlock(lock<sup>x</sup>)  
Lock(lock<sup>y</sup>)      f  
Z = Y              {lock<sup>y</sup>}  
Unlock(lock<sup>y</sup>)


## Thread II

Lock(lock<sup>y</sup>)       $\phi$   
Y = 2              {lock<sup>y</sup>}  
Unlock(lock<sup>y</sup>)  
Lock(lock<sup>y</sup>)       $\phi$   
Y = X              {lock<sup>y</sup>}

## C(X)

{lock<sup>x</sup>, lock<sup>y</sup>}  
{lock<sup>x</sup>}

f



# Lockset

---

- Advantage
  - Predict dataraces which may occur in a different thread interleaving
- Disadvantages
  - Spurious dataraces
  - Hard to use
    - Lack of trace



# Lockset strength

---

## Thread I

Lock(lock<sup>x</sup>);      f  
X = 7;              {lock<sup>x</sup>}  
Unlock(lock<sup>x</sup>);  
Lock(lock<sup>y</sup>);      f  
Z = Y;              {lock<sup>y</sup>}  
Unlock(lock<sup>y</sup>);

## Thread II

Lock(lock<sup>y</sup>);       $\phi$   
Y = 2;              {lock<sup>y</sup>}  
Unlock(lock<sup>y</sup>);  
Lock(lock<sup>y</sup>);       $\phi$   
Y = X;              {lock<sup>y</sup>}

## C(X)

{lock<sup>x</sup>, lock<sup>y</sup>}  
{lock<sup>x</sup>}

f





# Our solution

---

- Combine Lockset & Model Checking
  - Provide witnesses for dataraces
    - Rare dataraces
    - Dataraces in large programs

Model Checking  
Provide witnesses for rare DR

+

Lockset  
scale for large programs




# A witness for a datarace

---

Thread I

Thread II



$\pi_1$

$\pi_2$

$m^{a_1} = m^{a_2}$   
 $a_1 = \text{Write} \ ? \ a_2 = \text{Write}$




# Required data from Lockset

---

Thread I

Thread II






# Using Lockset data

---

- Lockset provides for each warning only a single access event  $a_2$



- Find a prior access event  $a_1$  which can take part in a race with  $a_2$



# Using Lockset data

---

$X = 7$      $\{\text{lock}^x\}$

$X=7$

$Z = Y$      $\{\text{lock}^y\}$

$Z=Y$

$Y = 2$      $\{\text{lock}^y\}$

$Y=2$

$Y = X$      $\{\text{lock}^y\}$


$Y=X$

**A Warning on X**



# Prefix


---






# Building a model

---



# Using a model checker








# Using a model checker

---





## Reduce the model checker cost

---

- Reduction in the model size
- Elimination of thread  $t^{a_1}$
- Providing a single new initial configuration
- Heuristically reducing the number of steps that the model checker should carry out

## Thread II

 $\pi_1$  $\pi_2$ C(X)

f

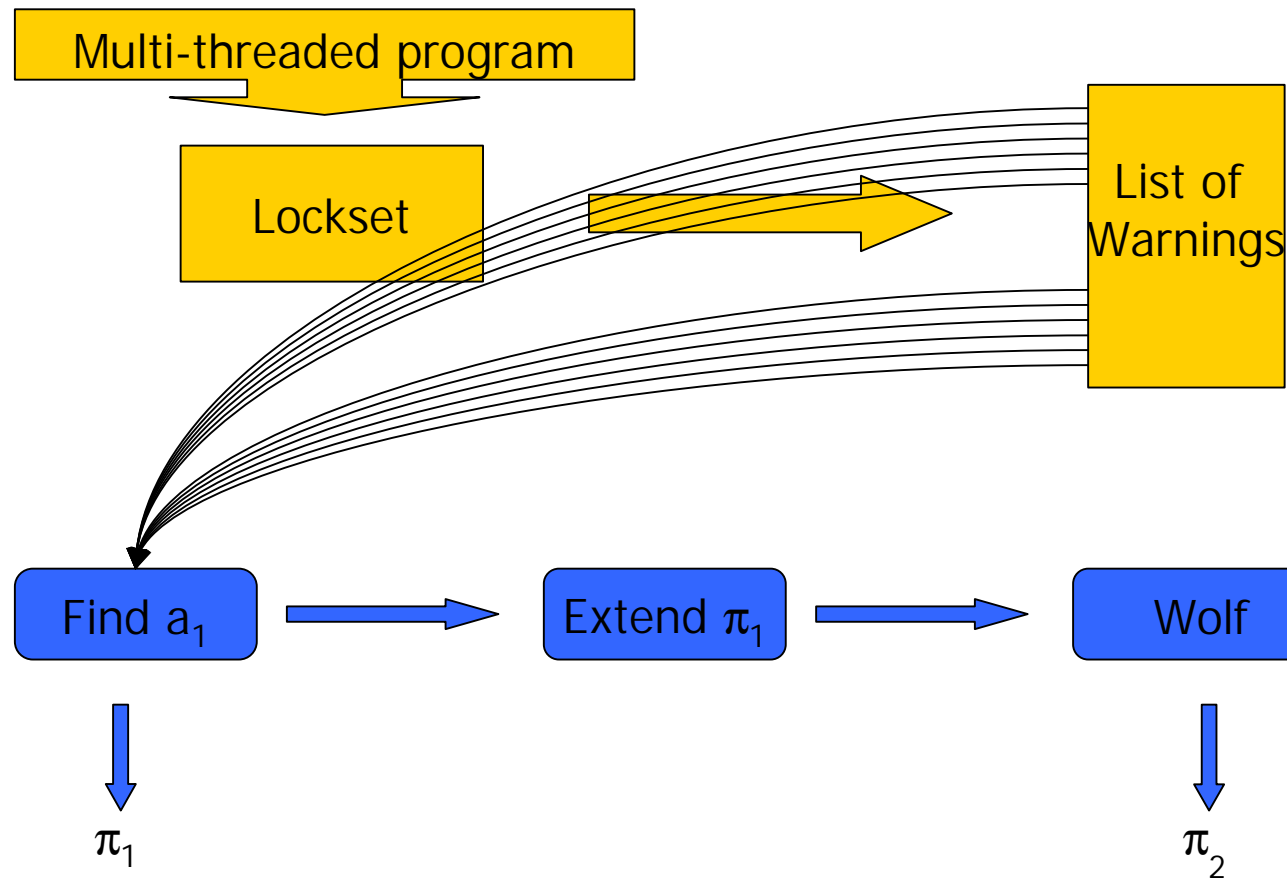


# Prototype implementation

---

- A prototype tool based on IBM tools
- Lockset – The IBM Watson tool
- Wolf – IBM Haifa's software model checker

# Prototype implementation





# Benchmark programs

---

Program	Description	Lines
Tsp	traveling salesman from ETH	706
Our_tsp	Enhanced traveling salesman	708
mtrt	Multithreaded raytracer from specjvm98	3751
Hedc	Web Crawler Kernel from ETH	29948
SortArray	Parallel sort	362
PrimeFinder	Finds prime numbers in a given interval	129
Elevsim	Elevator simulator	150
DQueries	Shared DB simulator	166



# Experimental results

Program	2 threads		3 threads		4 threads	
	Time (sec)	Memory (MB)	Time (sec)	Memory (MB)	Time (sec)	Memory (MB)
our_tsp	35069		Mem Out		Mem Out	
SortArray	569.3	123	1334.93	396	Mem Out	
PrimeFinder	888.7	116	2645.5	143	4547.1	168
ElevSim	33.02		67.92	33	147.9	48
DQueries	140.1	60	201.8	89	585.97	136
Hedc	2.66	11	7.33	12	9	17
tsp	35243	377	Mem Out		Mem Out	



# Conclusion

---

- Hybrid technique which combines dynamic datarace detector and a model checker
- Provide witnesses for dataraces which occur only in rare interleavings
- Helps the user in analyzing the datarace
- No spurious dataraces