

Finding Races in the Heap

Vesal Vojdani

University of Tartu
Technische Universität München

Mäetaguse Theory Days '09

Data Races

Definition (Race condition)

An **unintended** indeterminism due to the lack of proper ordering constraints in the program.

- Lead to subtle and dangerous bugs.
- Violate Murphy's Law \Rightarrow hard to test!

Famous Examples

- Therac-25 radiation therapy machine
 - Killed 3 people!
 - Race condition only occurred if setup was changed quickly; testers were not as fast.
- North American Blackout of 2003
 - Two processes got write accesses to a shared resource and corrupted it.
 - Alarm subsystem looped indefinitely.
 - The race had a window of only milliseconds!

Outline of Talk

- 1 Coarse-grained locking (static locks)
 - The Lockset Algorithm
- 2 Fine-grained (per-element locking)
 - JAVA: Conditional Must-Not Aliasing
 - C: Existentially Typed Flow (requires annotations)
 - C: **Interprocedural Must-Alias Analysis**
- 3 Medium-grained (per-bucket locking)
 - JAVA: Disjoint Reachability Analysis
 - C: **Region Analysis**
- 4 Unified ApproachTM to Race Detection

The Lockset Analysis

- For each program point
 - Compute set of locks that must be held.
 - `lock(l)` adds the lock that `l` **must** point to.
 - `unlock(l)` removes locks that `l` may point to.
- For each expression `e` in the program
 - Check if `e` **may** point to a shared variable.
 - Write down the access with set of mutexes held.
- Shared var has no common lock \Rightarrow race!

Simple Example (no race)

T_1 : `lock(&l1);`

`v = v + 1;`

`unlock(&l1);`

T_2 : `lock(&l1);`

`v = v + 1;`

`unlock(&l1);`

- List of accesses:

$\langle v, \{l_1\}, \text{write}, \text{file.c} : 2 \rangle$

$\langle v, \{l_1\}, \text{write}, \text{file.c} : 5 \rangle$

- v is protected by $\{l_1\}$.

Simple Example (race!)

T_1 : `lock(&l1);`

`v = v + 1;`

`unlock(&l1);`

T_2 : `lock(&l2);`

`v = v + 1;`

`unlock(&l2);`

- List of accesses:

$\langle v, \{l_1\}, \text{write}, \text{file.c} : 2 \rangle$

$\langle v, \{l_2\}, \text{write}, \text{file.c} : 5 \rangle$

- No common lock!

Complications

- Context-Sensitivity
- Path-Sensitivity
- Synchronization-Sensitivity
- Dynamic Memory Allocation

Dynamic Data, Static Locks

```
p = malloc();  
lock(&l);  
p → d = 5;  
unlock(&l);
```

- List of accesses:
⟨alloc@1.d, {l}, ...⟩
- Blob all elements allocated at that point.
- Can be handled as before.

Dynamic Data, Dynamic Locks

```
p = q = malloc();  
lock(&p → l);  
q → d = 7;  
unlock(&p → l);
```

- List of accesses:
⟨alloc@1.d, {alloc@1.l}, ...⟩
- Is it the same element in the blob?
- Must-equality analysis!

Finding Races in the Heap

- Keep a **symbolic lockset**

$$\{p.l\}$$

- Use address **must-equalities** to match symbolic locksets with accesses.

$$\models p = q$$

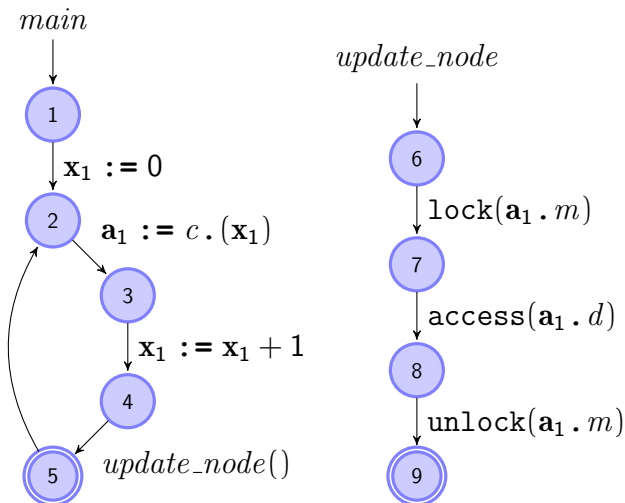
- Use **may points-to** analysis to associate inferred invariant with memory locations.

$$p \mapsto \{\text{alloc}@1\}$$


Per-element locking

```
typedef struct {  
    int datum;  
    char filename[80];  
    pthread_mutex_t mutex;  
} node;  
  
node cache[100];
```

Traverse cache




Valid Must-Equalities



- ① No equalities hold!
- ② No equalities hold!
- ③ $\mathbf{a}_1 \doteq c \cdot (\mathbf{x}_1)$.
- ④ $\mathbf{a}_1 \doteq c \cdot (-1 + \mathbf{x}_1)$.


Abstract domain



How do we compute $E_1 \Rightarrow E_2$?
Is this a complete lattice?

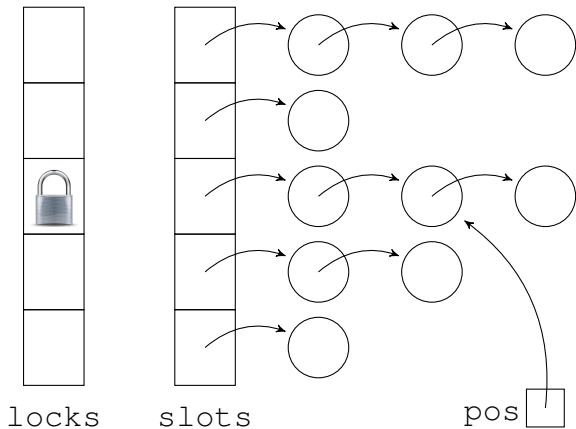
Back to Race detection

update



- We obtain the lockset:
 $\{c \cdot (-1 + \mathbf{x}_1) \cdot m, \mathbf{a}_1 \cdot m\}$
- Lockset must equal held locks at the access!
- Correlate $\mathbf{a}_1 \cdot d$ to $\mathbf{a}_1 \cdot m$.
- Associate invariant with array c .
- Why integer equalities?

Synchronized Hashtable



What do we need?

- Infer the locked addresses
`locks[i]`
- Information about pointers
`pos ∈ slot[i]`
- Disjointness information
`slot[i] ∩ slot[j] = ∅`

Region Analysis

- Heap Abstraction:

$$\Pi \times \mathbb{R}$$

- What are the disjoint regions?
- What region does each pointer belong to?


Region Analysis

- Heap Abstraction:


$$\Pi \times \mathbf{R}$$

- lattice of partitions on \mathbf{A} (owners)
- region mapping $\mathbf{R}: \mathbf{V} \rightarrow \mathcal{P}(\mathbf{A} \cup \{\bullet\})$

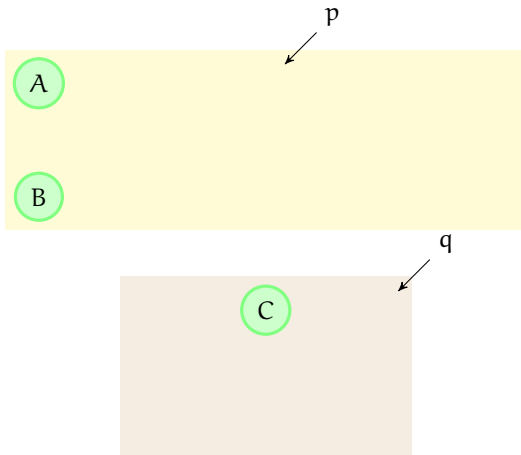
Abstraction



Abstraction



Abstraction



Abstraction

- Partitions: $\{A, B\}, \{C\}$
- Mapping:
 $p \mapsto \{A, B\}$
 $q \mapsto \{C\}$

Analysis

```
p = malloc();  
q = &A;  
p → n = q → n;  
q → n = &B;  
if (★) q = &C;
```

- Partitions:
 $\{A\}, \{B\}, \{C\}$
- Mapping:
 $p \mapsto \emptyset$
 $q \mapsto \emptyset$

Analysis

```
p = malloc();  
q = &A;  
p → n = q → n;  
q → n = &B;  
if (★) q = &C;
```

- Partitions:
 $\{A\}, \{B\}, \{C\}$
- Mapping:
 $p \mapsto \{\bullet\}$
 $q \mapsto \emptyset$

Analysis

```
p = malloc();  
q = &A;  
p → n = q → n;  
q → n = &B;  
if (★) q = &C;
```

- Partitions:
 $\{A\}, \{B\}, \{C\}$
- Mapping:
 $p \mapsto \{\bullet\}$
 $q \mapsto \{A\}$

Analysis

```
p = malloc();  
q = &A;  
p → n = q → n;  
q → n = &B;  
if (★) q = &C;
```

- Partitions:
 $\{A\}, \{B\}, \{C\}$
- Mapping:
 $p \mapsto \{A\}$
 $q \mapsto \{A\}$

Analysis

```
p = malloc();  
q = &A;  
p → n = q → n;  
q → n = &B;  
if (★) q = &C;
```

- Partitions:
 $\{A, B\}, \{C\}$
- Mapping:
 $p \mapsto \{A, B\}$
 $q \mapsto \{A, B\}$

Analysis

```
p = malloc();  
q = &A;  
p → n = q → n;  
q → n = &B;  
if (★) q = &C;
```

- Partitions:
 $\{A, B\}, \{C\}$
- Mapping:
 $p \mapsto \{A, B\}$
 $q \mapsto \{C\}$

Analysis

```
p = malloc();  
q = &A;  
p → n = q → n;  
q → n = &B;  
if (★) q = &C;
```

- Partitions:
 $\{A, B\}, \{C\}$
- Mapping:
 $p \mapsto \{A, B\}$
 $q \mapsto \{A, B, C\}$

Dealing with the arrays

- Allow symbolic index expressions.
- If we need to join two regions with owners $A[e_1]$ and $A[e_2]$
 - If we $\models e_1 = e_2$, just keep one of them.
 - Otherwise, **collapse** the array.
- Partition lattice Π tracks collapsed arrays.
- If an array A has not collapsed in Π , the regions of $A[i]$ and $A[j]$ are disjoint ($i \neq j$).

Back to the example

- The symbolic lockset will be
$$\{\text{locks} . (\mathbf{i})\}$$
- As we access `pos`, we have
 - Points-to analysis: `pos` \mapsto `{alloc@55}`.
 - Region analysis: `pos` \in `slots . (\mathbf{i})`.
- Associate the invariant with the array. Why?

Unified Approach to Race Detection

- Want to deal with dynamic regions.
 - List (instead of array) of regions.
 - Type-based regions for, e.g., per-device structures.
- Where to associate invariants?
 - Static names.
 - Allocation sites.
 - Types.
 - Regions.

Solution

- Under construction . . .
- All the ingredients are there.
- Implementation is making good progress.
- **Where are the theorems?**