

Data races in Java and static analysis

Frédéric Dabrowski

¹INRIA, LANDE

Outline

- 1 Concurrency in Java
- 2 Static Detection of dataraces
 - lock-based typing $\left\{ \begin{array}{l} \text{Flanagan, Abadi \& Freund} \\ \text{Boyapati, Lee \& Rinard} \end{array} \right.$
 - Points-to analysis : Naik & Aiken
(Points-to analysis + Type and effect system)
- 3 Conclusion and ongoing work

Concurrency in Java

Concurrency in Java

Concurrency model

- Thread-based concurrency : shared memory (fields of shared objects)
- lexically scoped locking construct : `synchronized(x){...}`
- Preemptive scheduling (Interleaving semantics)

Concurrency in Java

Interleaving semantics

(small step)
sequential semantics

$$t, Mem \rightarrow_{seq} t', Mem'$$

interleaving semantics

$$\frac{t_i, Mem \rightarrow_{seq} t'_i, Mem'}{\{\dots, t_i, \dots\}, Mem \rightarrow_{inter} \{\dots, t'_i, \dots\}, Mem'}$$

Problem : This semantics is incomplete with respect to the Java Memory Model, **unless you write well-synchronized programs**

Concurrency in Java

Natural hypothesis : sequential consistency

Intuitively, **sequential consistency** means that all executions respect the program order.

```
void mn(){  
    a    a should not observe b  
    ...  
    b  
}
```

Problem : enforcing sequential consistency for all Java programs makes many of the compiler/processor optimizations illegal.

Why ? some optimizations assume well-synchronized programs !

Concurrency in Java


Example : code reordering (cache mechanisms,...)

Original code

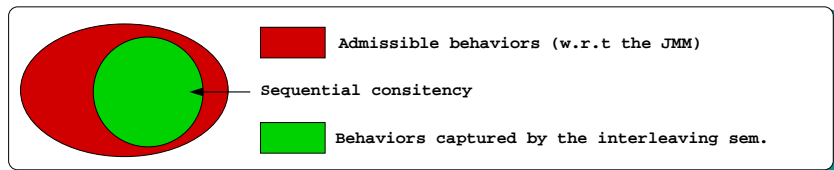
```
C.f = C.g = 0  
1 : x = C.g; | 3 : y = C.f;  
2 : C.f = 1; | 4 : C.g = 1;  
{Perm(1,2,3,4) | 1 < 2, 3 < 4}
```

Optimized code

```
C.f = C.g = 0  
2 : C.f = 1; | 4 : C.g = 1;  
1 : x = C.g; | 3 : y = C.f;  
{Perm(1,2,3,4) | 2 < 1, 4 < 3}
```



Concurrency in Java



Java's memory model is weak memory model

All executions of well-synchronized programs are sequentially consistent ^a.

^aManson, Pugh & Adve : The Java Memory Model (Special Popl issue)

Programs **must** be well-synchronyzed
several static analysis depend on it


Concurrency in Java

Well-synchronized programs

(P_1) : For all execution (w.r.t the interleaving semantics), every conflicting actions a and b are synchronized

`synchronized(C){`
`(1 : x = C.g); (2 : C.f = 1)`
`}`

`synchronized(C){`
`(3 : y = C.f); (4 : C.g = 1)`
`}`



compiler/jvm/jit : (P_1) \Rightarrow every exec. is captured by the inter. sem.

Concurrency in Java

Happens-before relation

\prec_{hb} is the transitive closure of the following rules :

- **sequentiality** $a^t \prec_{hb}^1 b^t$
- **start/join synchronisation** $\begin{cases} t.start() \prec_{hb}^1 a^t \\ a^t \prec_{hb}^1 t.join() \end{cases}$
- **lock-based synchronisation**

$$\begin{cases} \text{unlock}(m) \prec_{hb}^1 \text{lock}(m) \\ \text{write}(x.f) \prec_{hb}^1 \text{read}(x.f) \quad (f \text{ volatile}) \end{cases}$$

Concurrency in Java

Data races

definition (JMM) : a program has a data race if there exists an execution with two conflicting actions not ordered by \prec_{hb} .

alternative definition : a program has a data race if there exists an execution such that, at some point, there is a non deterministic choice (interleaving semantics) among two conflicting actions.

Problem (undecidable) :

Given a program, are all executions of that program race free?

Static detection of data races

RCC Java

[PLDI'00] Type-based race detection for Java (Flanagan and Freund)

- supports classes parameterized by locks of given types (Dependent types)
- introduce a notion of thread local classes
- Fields protected by locks (static fields)
- Encapsulation : self-protected class
- Extend previous work based on a simple thread calculus

Example

```
class A<ghost Object x>{  
    Object y = new Object() guarded_by x;  
    void set(Object z) requires x{  
        this.y = z;  
    }  
}
```

```
class B{  
    final Object z = new Object();  
    A<this.z> x = new A<this.z>();  
    void f(){  
        synchronized(this.z){set(new Object())}  
    }  
}
```

Ownership types

[OOPSLA'02] Ownership types for safe programming : preventing data races and deadlocks (Boyapati, Lee and Rinard)

Basic idea

Write generic code and create different instance with different protection mechanisms.

Ownership types

Ownership types

Each object is owned by $\left\{ \begin{array}{l} \text{another object (final field)} \\ \text{itself (self)} \\ \text{a thread (thisThread)} \end{array} \right.$

```
class C<Owner0, Owner1, ...>    {... new D<Owner1>() ...}
```


Ownership types

Static analysis

- 1 The ownership relation builds a forest of trees
- 2 The fields of an object must be protected by the ancestor of this object (a root, i.e. an object protected by itself or a thread)

Extensions

- support for read-only/single pointer objects

Ownership types

Example

```
class Account<thisOwner>{
    int balance = 0;
    void deposit(int x) requires thisOwner{
        this.balance = this.balance + x;
    }
}

Account <thisThread> a1 = new Account<thisThread>
a1.deposit(10);

final Account<self> a2 = new Account<self>;
fork{synchronized(a2){a2.deposit(10);}}

Account<a2> a3 = new Account<a2>;
```

Type inference

[VMCAI'04] Type Inference for Parameterized Race-Free Java (Agarwal and Stoller)

idea :

- Perform a set of execution
- Extract types from this set
- Check types

problem : incomplete

[SAS'04, SCP'07] Type inference against races (Flanagan and Freund)

- consider parameterization of classes as introduced by Boyapati, Lee and Rinard
- by reduction of the problem of finding a satisfying assignment for a boolean formula (NP-complete)

Bytecode analysis

[Sigplan Not.] A type system for preventing data races and deadlocks in the java virtual machine language
(Permandla, Roberson, Boyapati)

problem : monitorenter/monitorexit replace synchronized blocks

solution : use indexed types to recover structured locking

Indexed types : [TCS'03] A type system for JVM threads (Laneve)
very simple alias analysis

$$i : \text{Load } n, \{ \dots n \mapsto \tau \dots \}, \text{stack} \rightarrow \{ \dots n \mapsto \tau_i \dots \}, \tau_i :: \text{stack}$$

Limitations of type-based approaches

- very strict lock-based discipline
- can't handle other synchronization patterns

Naik and Aiken & al

```


a = new h1[N];
for(i = 0; i < N; i ++){
    a[i] = new h2;
    a[i].f = new h3;
}

```

```

while (*){
    x = a[*];
    fork{
        sync(x){x.f.g = *;}
    }
}

```



Language

$$s ::= \mid x = \text{null} \mid x = \text{new } h \\ \mid x = y \mid x = y.f \mid x.f = y \\ \mid s_1; s_2 \mid \text{if } (*) \text{ then } s_1 \text{ else } s_2 \mid \text{while}^w (*) \text{ do } s$$

h allocation site

$w \in \mathbb{W}$ loop counter

Dynamic semantics

$$\text{Obj} ::= \langle h, \pi \rangle \quad \pi : \mathbb{W} \rightarrow \mathbb{N} \\ C ::= \{ \dots \text{Obj} \triangleright \text{Obj}' \dots \}$$

Conditional Must Not Aliasing

<code>synchronized(x){</code>			<code>synchronized(y){</code>
<code>x.f.g = *;</code>			<code>y.f.g = *;</code>
<code>}</code>			<code>}</code>

Conditional Must Not Aliasing
 $must_not_alias(x, y) \Rightarrow must_not_alias(x.f, y.f)$

Disjoint Reachability

$$h \in DR_C(H) \Leftrightarrow \left\{ \begin{array}{l} \text{Disjoint reachability} \\ \overline{o}_1.h \in H \wedge (\overline{o}_1 \triangleright \overline{o}) \in C^+ \wedge \\ \overline{o}_2.h \in H \wedge (\overline{o}_2 \triangleright \overline{o}) \in C^+ \wedge \\ \overline{o}.h = h \end{array} \right. \Rightarrow \overline{o}_1 = \overline{o}_2$$

		<u>Abstraction</u>
Obj	$::= \langle \hat{h}, \Pi \rangle$	
\hat{h}	$::= h \mid \top$	$\Pi(w) = 0$ w not active
Π	$: \mathbb{W} \mapsto \mathbb{N}_\top$	$\Pi(w) = \top$ w unknown
\mathbb{N}_\top	$= \{0, 1, \top\}$	$\Pi(w) = \Pi'(w) = 1$ same iteration

Judgments : $W, \Pi, \Gamma \vdash s : \Gamma', K$

$h \in DR_K(H)$ $DR_K(H)$ is a safe appr. of $DR_C(H)$

Disjoint reachability

Examples

$h_2 \in DR_K(\{h_1\})?$

```
while1 (*){  
  x = new h1;  
  y = new h2;  
  x.f = y;  
}
```

YES : {⟨h₁, (1)⟩▷
⟨h₂, (1)⟩}

```
while1 (*){  
  x = new h1;  
  while2 (*){  
    y = new h2;  
    x.f = y;  
  }  
}
```

YES : {⟨h₁, (1, 0)⟩▷
⟨h₂, (1, 1)⟩}

```
while1 (*){  
  y = new h2;  
  while2 (*){  
    x = new h1;  
    x.f = y;  
  }  
}
```

NO : {⟨h₁, (1, 1)⟩▷
⟨h₂, (1, 0)⟩}

Disjoint reachability

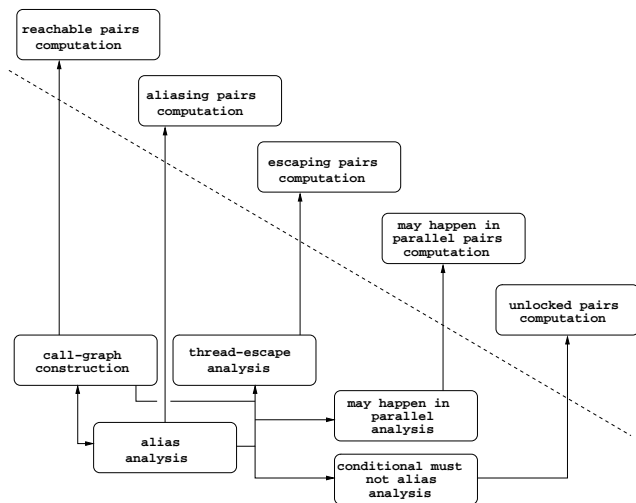
$$W, \Pi, \Gamma \vdash x = \text{new } h : \Gamma[x \mapsto \langle \Pi', h \rangle], \emptyset$$
$$W, \Pi, \Gamma \vdash x = y : \Gamma[x \mapsto \Gamma(y)] \quad W, \Pi, \Gamma \vdash x.f = y, \Gamma, K$$
$$W, \Pi, \Gamma \vdash x = y.f : \Gamma[x \mapsto \langle \lambda w. \top, \top \rangle], \emptyset$$
$$\frac{W, \Pi, \Gamma \vdash s_1 : \Gamma', K_1 \quad W, \Pi, \Gamma' \vdash s_1 : \Gamma'', K_2}{W, \Pi, \Gamma \vdash s_1; s_2 : \Gamma'', K_1 \cup K_2}$$
$$\frac{W, \Pi, \Gamma \vdash s_1 : \Gamma_1, K_1 \quad W, \Pi, \Gamma \vdash s_1 : \Gamma_2, K_2}{W, \Pi, \Gamma \vdash \text{if } (*) \text{ then } s_1 \text{ else } s_2 : \Gamma_1 \sqcup \Gamma_2, K_1 \cup K_2}$$
$$\frac{W \cup \{w\}, \Pi, \Gamma^{w^+} \vdash s : \Gamma, K \quad \Pi(w) \neq 0}{W, \Pi, \Gamma \vdash \text{while}^w (*) \text{ do } s : \Gamma, K}$$

Conflicting pairs elimination

Static analysis

- Call graph construction and context-sensitive points-to analysis
- Type and effect system

Conflicting pairs elimination



Soundness

- reflection
- dynamic loading
- native methods
- libraries ???

Conclusion and ongoing work

- Data races detection is important
- Objects offer a good framework
- Type systems can handle strict lock-based discipline but lack more elaborated synchronisation patterns
- Points-to analysis can give very precise results but is much more complex

Ongoing work

Certification of a static analysis for data race detection in Coq

- 1 Context-sensitive points-to analysis
 - Certification of a result checker in Coq
- 2 Static analysis for data race detection
 - Formalisation of aiken's type and effect system for Java Bytecode
 - Formalisation of successive stages
 - Certification in Coq