

A fast and precise analysis for data race detection

Pietro Ferrara

Ecole Polytechnique, Paris
F-91128 Palaiseau (France)
Università Ca' Foscari di Venezia
I-30170 Venezia (Italy)
`Pietro.Ferrara@polytechnique.edu`

April 5, 2008

Abstract interpretation

- ▶ Static analysis
- ▶ Mathematical theory to define and soundly approximate the semantics of a program [CC77,CC79]
- ▶ Concrete domain: $\langle A, \sqsubseteq_C, \perp_C, \top_C, \sqcup_C, \sqcap_C \rangle$
- ▶ The abstract domain approximates it $\langle \hat{A}, \sqsubseteq_A, \perp_A, \top_A, \sqcup_A, \sqcap_A \rangle$
- ▶ Abstraction α and concretization γ functions:
 - ▶ $\forall s \in A : s \sqsubseteq_C \gamma(\alpha(s))$
 - ▶ $\forall \hat{s} \in \hat{A} : \alpha(\gamma(\hat{s})) \sqsubseteq_A \hat{s}$
- ▶ Abstract and concrete semantics:

$$\forall s \in A : \mathbb{F}[[s]] \sqsubseteq_C \gamma(\hat{\mathbb{F}}[[\alpha(s)]])$$


Java and Multithreading

- ▶ Java: native support for threads
- ▶ Threads are objects → identified by reference
- ▶ `Thread.start()` launches a new thread
- ▶ `monitorenter` and `monitorexit`
- ▶ Many other ways of synchronizing
 - ▶ `wait` and `notify` on objects
 - ▶ `isAlive()` on threads
 - ▶ ...
- ▶ We focus only on monitors

Data Race Condition


- ▶ “two threads concurrently access the same area of shared memory, they are not synchronized, and at least one is performing a write operation”
- ▶ data race \approx a bug
- ▶ In Java:
 - ▶ Shared memory \rightarrow heap
 - ▶ Monitors \rightarrow objects
 - ▶ Threads \rightarrow objects

Thread1: ... putfield #1



Thread2: ... putfield #1

Thread1: ... putfield #1




Thread2: ... putfield #1

Thread1: ... putfield #1


#8



Thread2: ... putfield #1



Thread1: ... putfield #1



Thread2: ... putfield #1

Thread1:

... putfield #1

#10

Heap

#10

0

1

2

#10

Thread2:

... putfield #1

{#15}



Thread1:

... putfield #1

#10

Heap

#10

0

1


2

#10

Thread2:

... putfield #1

{#12}



- ▶ We need to precisely trace the accesses to the shared memory
- ▶ At low level
- ▶ In practice, formalization of the Java Virtual Machine specification:
 - ▶ Values: references or integers - $\text{Val} = \text{Ref} \cup \mathbb{N}$
 - ▶ Operand stack: stack of values - $\text{Op} = \mathcal{ST}(\text{Val})$
 - ▶ Local Variables: array of values - $\text{LV} = \mathbb{A}\mathbb{R}(\text{Val})$
 - ▶ Frame: an operand stack, a local variables element, a program counter - $\text{F} = \text{Op} \times \text{LV} \times \text{PC}$

- ▶ Single thread state
- ▶ Monitors can be locked more than once
- ▶ Function that traces for each monitor how many times has been locked

Definition (Monitor)

$$L = [\text{Ref} \rightarrow \mathbb{N}]$$

- ▶ Heap: reference \mapsto local variables' element
- ▶ We augment the information contained by the local variables
- ▶ For each value, we trace
 - ▶ The thread that wrote it, identified by its reference
 - ▶ The program counter of the instruction
 - ▶ The stack of called method
 - ▶ The owned monitors
- ▶ Necessary in order to check the data race condition

Definition (Augmented Local Variables)

$$ALV = \mathbb{A}\mathbb{R}(\text{Val} \times \text{Ref} \times \text{PC} \times \mathbb{S}\mathbb{T}(\text{PC}) \times \text{L})$$

Definition (Heap)

$$H = [\text{Ref} \rightarrow \text{ALV}]$$

- ▶ Monothread state:
 - ▶ A stack of frame (one for each method call)
 - ▶ An heap
 - ▶ A monitors' function

Definition (State)

$$\Sigma = \text{ST}(\text{F}) \times \text{H} \times \text{L}$$

Multithreaded state

- ▶ Collects for each thread the trace of its execution
- ▶ The inter-thread order of execution is abstracted away
- ▶ The same approach has been adopted by [FER08]

Definition (Multithreaded State)

$$\Omega = \wp([\text{Ref} \rightarrow \Sigma^{\vec{+}}])$$

- ▶ Partial trace semantics [CC79]
- ▶ At each step
 - ▶ it chooses randomly a thread
 - ▶ it checks if it is going to acquire a monitor already owned by another thread
 - ▶ it takes its last state
 - ▶ it builds up a possible state of the heap following a memory model [FER08]
 - ▶ it executes one step
 - ▶ it stores the result
- ▶ Until a fixpoint is reached

Data Race Condition

- ▶ Check **locally** if a data race happens
- ▶ When a value in the shared memory is accessed:
 - ▶ check if it has been written by another thread
 - ▶ check if it owned a monitor currently acquired by the thread that is performing the access
 - ▶ if it is not the case, a data race is detected

{[#15 ↦ 1]}



Thread1:



#3

... putfield #1



#10

Heap


#10

0
1 (#3, ..., {#15})
2


Thread2:

... putfield #1

Thread1: ... putfield #1



... putfield #1



#3 ≠ #8
{#15} ∩ {#12} = ∅
Data Race!

$$\widehat{\text{Val}} = \widehat{\text{Ref}} \cup \widehat{i}$$

$$\widehat{\text{Op}} = \text{ST}(\widehat{\text{Val}})$$

$$\widehat{\text{LV}} = \text{AR}(\widehat{\text{Val}})$$

$$\widehat{\text{F}} = \widehat{\text{Op}} \times \widehat{\text{LV}} \times \text{PC}$$

- ▶ Alias analysis that traces
 - ▶ When two values **may** point to the same address
 - ▶ check when two threads may access the same area of shared memory
 - ▶ When two values contains **always** the same address
 - ▶ check when two threads are surely synchronized on the same monitor

May-aliasing

- ▶ Abstract reference:
 - ▶ the program counter
 - ▶ the stack of the called methods
 - ▶ the abstract reference of the executing threadwhen the `new` statement allocates it
- ▶ `#mainthread`
- ▶ An abstract reference represents multiple concrete references
- ▶ Used to
 - ▶ Identify threads
 - ▶ Check when two accesses to the shared memory may be on the same location

Definition (May-alias domain)

$$\hat{P} = (\text{PC} \times \mathcal{ST}(\text{PC}) \times \hat{P}) \cup \{\#mainthread\}$$

Must-aliasing

- ▶ Equivalence classes
- ▶ Identify monitors
- ▶ If two threads lock on the same equivalence class they are synchronized

Definition (Monitors)

$$\widehat{L} = [\widehat{E} \rightarrow \mathbb{N}]$$

Definition

$$\widehat{\text{Ref}} = \widehat{E} \times \wp(\widehat{P})$$

Augmented local variables

- ▶ In the concrete, different executions may contain at the same point of computation
 - ▶ different values
 - ▶ written by different threads and different points of the program
- ▶ In abstract
 - ▶ collect together all these values
 - ▶ lub of all possible values
 - ▶ set union of all the threads and program points that may write the value

Definition

$$\widehat{ALV} = \mathbb{A}\mathbb{R}(\widehat{Val} \times \wp(\widehat{P} \times PC \times \mathcal{S}\mathcal{T}(PC) \times \widehat{L}))$$

Definition (Heap)

$$\widehat{H} = [\widehat{P} \rightarrow \widehat{ALV}]$$

Definition (Monothread State)

$$\hat{\Sigma} = \text{ST}(\hat{F}) \times \hat{H} \times \hat{L}$$

Definition (Multithreaded State)

$$\hat{\Omega} = [\hat{P} \rightarrow \hat{\Sigma}^{\mp}]$$

- ▶ Computed through two nested fixpoints
 - ▶ Single thread semantics: approximate the semantics of a single thread given a multithreaded state
 - ▶ values written in parallel by other threads
 - ▶ Multithread semantics: iterate for all the active threads their single thread semantics passing at each iteration the last multithread state
- ▶ This approach has been formalized and proved to be sound on the happens-before memory model by a previous work [FER08]

$\{[2 \mapsto 1]\}$





Thread1:



\hat{a}_1


... putfield #1



Thread2:

... putfield #1


Thread1: ... putfield #1



Thread2: ... putfield #1

$\{18 \mapsto 1\}$

$\{\hat{a}_4, \hat{a}_5\}$



$\hat{a}_2 \neq \hat{a}_1$
 $\{2\} \cap \{18\} = \emptyset$
Data Race!

If a data race occurs

- ▶ Data race \Rightarrow An access to the memory reads or overwrites a value written by another thread
- ▶ We know:
 - ▶ Program points that cause the data race
 - ▶ The call stack followed to arrive at that point
 - ▶ The threads that perform the accesses
 - ▶ Equivalence classes of the monitors owned when the value has been previously written by the other thread
- ▶ Equivalence classes reachable from the local environment
- ▶ Suggest which objects should lock (we can have 0, 1, or more suggestions!)
- ▶ Only an advice

# Application	# Threads	# Statements	Time (sec)
1	3	424	7"
2	5	637	7"
3	7	751	10"
4	11	977	13"
5	13	1.097	17"
6	15	1.310	24"
7	17	1.422	25"
8	19	1.635	36"
9	20	1.742	49"
10	24	2.126	1'07"

Table: Experimental results

- ▶ Encouraging
- ▶ Test our analysis with bigger (i.e. industrial) examples
- ▶ No false alarms

Conclusion

- ▶ New static analysis of data races
- ▶ Combination of a may- and a must- alias analysis
- ▶ Provide useful information when a data race may happen in order to discover missing lock actions
- ▶ Implemented, the experimental results are promising

Future works

- ▶ Extend the analyzer in order to support all the Java bytecode language
- ▶ Apply it to bigger and perhaps industrial programs
- ▶ Investigate other properties on multithreaded programs, like determinism

[FER08] P. Ferrara. Static analysis via abstract interpretation of the happens-before memory model. In Springer, editor, Proceedings of TAP 08, LNCS, 2008.

- [CC77]** Patrick Cousot and Radhia Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In POPL 77, pages 238-252, Los Angeles, California, 1977. ACM Press.
- [CC79]** Patrick Cousot and Radhia Cousot. Systematic design of program analysis frameworks. In POPL 79, pages 269-282, San Antonio, Texas, 1979. ACM Press.

- [DEU94]** A. Deutsch. Interprocedural may-alias analysis for pointers: beyond k-limiting. In ACM Press, editor, Proceedings of PLDI 04, 1994.
- [VEN02]** A. Venet. Nonuniform alias analysis of recursive data structures and arrays. In Springer, editor, SAS, LNCS, 2002.
- [SRW99]** M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3valued logic. In ACM Press, editor, Proceedings of POPL 99, 1999.

Bibliography - Data Race

- [BAN03]** U. Banerjee, B. Bliss, Z. Ma, and P. Petersen. Unraveling data race detection in the intel thread checker. In Proceedings of STMCS 06, 2006.
- [CAL03]** R. OCallahan and Jong-Deok Choi. Hybrid dynamic data race detection. In ACM Press, editor, Proceedings of PPOPP 03, 2003.
- [POZ03]** E. Poznianski and A. Schuster. Efficient on-the-fly data race detection in multithreaded c++ programs. In IEEE Computer Society, editor, Proceedings of IPDPS 03, 2003.
- [SAV97]** S. Savage, M. Burrows, G. Nelson, P. Sobalvarro, and T. Anderson. Eraser: A dynamic data race detector for multithreaded programs. In ACM Press, editor, Proceedings of SOSP 97, 1997.
- [YU05]** Y. Yu, T. Rodeheer, and W. Chen. Racetrack: efficient detection of data race conditions via adaptive tracking. In ACM Press, editor, Proceedings of SOSP 05, 2005.
- [ABA06]** M. Abadi, C. Flanagan, and S. N. Freund. Types for safe locking: Static race detection for java. In ACM Press, editor, Proceedings of TOPLAS 06, 2006.