Effective Static Race Detection for Java

Mayur Naik Alex Aiken Stanford University

What is a Race?

A condition in a shared-memory, multi-threaded program in which:

- same memory location accessed by different threads simultaneously (without holding a common lock)
- at least one of the accesses is a write

Why is Race Detection Important?

- Particularly insidious concurrency bug
 - Triggered non-deterministically
 - No fail-fast behavior even in safe languages like Java
- Fundamental in concurrency theory and practice
 - Lies at heart of many concurrency problems
 - atomicity checking, deadlock detection, ...
 - Today's concurrent programs riddled with races
 - "... most Java programs are so rife with concurrency bugs that they work only 'by accident'."

-Java Concurrency in Practice, Addison-Wesley, 2006

Our Result



- 412 bugs in mature Java programs comprising 1.5 MLOC
 - Many fixed within a week by developers

Our Race Detection Approach



Challenges

Same location accessed by different threads simultaneously without common lock held

- Handle multiple aspects •
 - Same location accessed
 - ... by different threads
 - ... simultaneously

- Precision
 - Showed precise *may alias analysis* is central (PLDI'06)
 - low false-positive rate (20%)

- Correlate locks with locations they guard
 - ... without common lock held
- Soundness
 - Devised *conditional must not alias* analysis (POPL'07)
 - Circumvents *must alias analysis*

Our Race Detection Approach



Alias Analysis for Race Detection

```
// Thread 1: // Thread 2:
sync (11) {
    ... e1.f ...
}
```

• Field **f** is race-free if:

e1 and e2 *never* refer to the same value MUSTANOTIANS(\$(\$2,2,2))

k-Object-Stansifivias MayaAylsias Analysis

- Rargenboodyy cafiwoakalysis [Milanova et al. ISSTA'03]
- **Btddptfelm**. CEonotefetwinstestsiticity evaluately!sis
 - Abstract value = set of allocation sites

- Rholp #2 n. CE on the fet we can story changed yes in the test of FA)
 - Context $(k=1) = a \exists b d site on site of this parameter$

foo() {
 el.baz();
 }

bar() {
 e2.baz();
 }

Analyze function **baz** in two contexts

k-Object-Sensitive Analysis: Our Contributions

- No scalable implementations for even k = 1
- Insights:
 - Symbolic representation of relations
 - BDDs [Whaley-Lam PLDI'04, Lhotak-Hendren PLDI'04]
 - Demand-driven race detection algorithm
 - Begin with k = 1 for all allocation sites
 - Increment *k* only for those involved in races
- Allow scalability to k = 7

Our Race Detection Approach



Same location accessed by different threads simultaneously without common lock held

Alias Analysis for Race Detection

```
// Thread 1: // Thread 2:
sync (11) {
    ... e1.f ...
}
```

• Field **f** is race-free if:

e1 and e2 *never* refer to the same value ¬ MAY-ALIAS(e1, e2)

OR

11 and 12 always refer to the same value
 MUST-ALIAS(11, 12)

Must Alias Analysis

• Small body of work

- Much harder problem than may alias analysis

Impediment to many previous race detection approaches
 Folk wisdom: Static race detection is intractable

Insight: Must alias analysis not necessary for race detection!

```
// Thread 1: // Thread 2:
sync (l1) {
    ... e1.f ...
}
```

• Field **f** is race-free if:

Whenever 11 and 12 refer to different values, e1 and e2 also refer to different values

MUST-NOT-ALIAS(11, 12) => MUST-NOT-ALIAS(e1, e2)

Example

```
a = new h0[N];
                                                h0
                                        a[0]
                                                      a[N-1]
for (i = 0; i < N; i++) {
                                             a[i]
    a[i] = new h1;
    a[i].g = new h2;
                                      h1
                                                h1
                                                          h1
                                          . . .
                                                     . . .
}
                                     g
                                               g
                                                          g
                  x2 = a[*];
x1 = a[*];
                                      h2
                                               h2
                                                          h2
                                          . . .
                                                     . . .
                  sync (?) {
sync (?) {
    x1.g.f = ...; x2.g.f = ...;
}
                    }
```

Easy Case: Coarse-grained Locking



Field **f** is race-free if:

MUGT-NOT-ALIAS(alal2> NUSUSYONOALAAS(cg, m2).g)

Example

```
a = new h0[N];
                                                h0
                                        a[0]
                                                      a[N-1]
for (i = 0; i < N; i++) {
                                             a[i]
    a[i] = new h1;
    a[i].g = new h2;
                                                h1
                                      h1
                                                          h1
                                          . . .
                                                     . . .
}
                                               g
                                     g
                                                          g
                  x2 = a[*];
x1 = a[*];
                                     h2
                                               h2
                                                          h2
                                          . . .
                                                     . . .
                  sync (?) {
sync (?) {
    x1.g.f = ...; x2.g.f = ...;
}
                    }
```

Easy Case: Fine-grained Locking



Field **f** is race-free if:

MUST-NOT-ALIAS(1, 1, 2), 2>) US MNOT-NOTAS(21, 3, x2.g)

Example

```
a = new h0[N];
                                                h0
                                        a[0]
                                                      a[N-1]
for (i = 0; i < N; i++) {
                                             a[i]
    a[i] = new h1;
    a[i].g = new h2;
                                                h1
                                      h1
                                                          h1
                                          . . .
                                                     . . .
}
                                               g
                                     g
                                                          g
                  x2 = a[*];
x1 = a[*];
                                     h2
                                               h2
                                                          h2
                                          . . .
                                                     . . .
                  sync (?) {
sync (?) {
    x1.g.f = ...; x2.g.f = ...;
}
                    }
```

Hard Case: Medium-grained Locking



Field **f** is race-free if:

MUS (find g-afidiatin tropic incl has alues inkst to datince has alues 222.g)

Disjoint Reachability



then $\{h2\} \subseteq DR(\{h1\})$

Note: Values abstracted by sets of allocation sites

Conditional Must Not Alias Analysis using Disjoint Reachability



Field **f** is race-free if:

MUSIC (01)-A Sixs(22)12) DR(Sitts(110) - Aits(3(2), e2) - e1 reachable from 11 and e2 reachable from 12

Hard Case: Medium-grained Locking



Field **f** is race-free if:

- $(file)(\underline{sl}$
- trucgachablebfeofrom xndred x2aghradalehable from x2

Experience with Chord

- Experimented with 12 multi-threaded Java programs
 - smaller programs used in previous work
 - larger, mature and widely-used open-source programs
 - whole programs and libraries
- Tool output and developer discussions available at: <u>http://www.cs.stanford.edu/~mhn/chord.html</u>
- Programs being used by other researchers in race detection

Benchmarks

	classes	KLOC	description	time
vect1.1	19	3	JDK 1.1 java.util.Vector	0m28s
htbl1.1	21	3	JDK 1.1 java.util.Hashtable	0m27s
htbl1.4	366	75	JDK 1.4 java.util.Hashtable	2m04s
vect1.4	370	76	JDK 1.4 java.util.Vector	2m02s
tsp	370	76	Traveling Salesman Problem	3m03s
hedc	422	83	Web crawler	9m10s
ftp	493	103	Apache FTP server	11m17s
pool	388	124	Apache object pooling library	10m29s
jdbm	461	115	Transaction manager	9m33s
jdbf	465	122	O/R mapping system	9m42s
jtds	553	165	JDBC driver	10m23s
derby	1746	646	Apache RDBMS	36m03s

Pairs Retained After Each Stage (Log scale)



Classification of Unlocked Pairs

	harmful	benign	false	# bugs
vect1.1	5	12	0	1
htbl1.1	0	6	0	0
htbl1.4	0	9	0	0
vect1.4	0	0	0	0
tsp	7	0	4	1
hedc	170	0	41	б
ftp	212	3	43	32
pool	105	10	13	17
jdbm	91	0	7	2
jdbf	130	0	34	18
jtds	34	14	17	16
derby	1018	0	78	319

Developer Feedback

- 16 bugs in jTDS
 - Before: "As far as we know, there are no concurrency issues in jTDS ..."
 - After: "It is probably the case that the whole synchronization approach in jTDS should be revised from scratch ..."
- 17 bugs in Apache Commons Pool
 - "Thanks to an audit by Mayur Naik many potential synchronization issues have been fixed" -- *Release notes for Commons Pool 1.3*
- 319 bugs in Apache Derby
 - "This looks like *very* valuable information and I for one appreciate you using Derby ... Could this tool be run on a regular basis? It is likely that new races could get introduced as new code is submitted ..."

Related Work

- Static (compile-time) race detection
 - Need to approximate multiple aspects
 - Need to perform must alias analysis
 - Sacrifice precision, soundness, scalability
- Dynamic (run-time) race detection
 - Current state of the art
 - Inherently unsound
 - Cannot analyze libraries
- Shape Analysis
 - much more expensive than disjoint reachability

Summary of Contributions

- Precise race detection (PLDI'06)
 - Key idea: *k*-object-sensitive may alias analysis
 - Important client for may alias analyses
- Sound race detection (POPL'07)
 - Key idea: Conditional must not alias analysis
 - Has applications besides race detection
- Effective race detection
 - 412 bugs in mature Java programs comprising 1.5 MLOC
 - Many fixed within a week by developers

The End



http://www.cs.stanford.edu/~mhn/chord.html