

Atomicity Analysis of Concurrent Software

Cormac Flanagan
UC Santa Cruz

Stephen N. Freund
Williams College

Shaz Qadeer
Microsoft Research

Types Against Races

Moore's Law


- Transistors per chip doubles every 18 months
- Single-threaded performance doubles every 2 years
 - faster clocks, deep pipelines, multiple issue
 - wonderful!

Moore's Law is Over

- Sure, we can pack more transistors ...
 - ... but can't use them effectively to make single-threaded programs faster
- Multi-core is the future of hardware
- Multi-threading is the future of software

Programming With Threads

- Decompose program into parallel threads
- Advantages
 - exploit multiple cores/processors
 - some threads progress, even if others block
- Increasingly common (Java, C#, GUIs, servers)



```

*** STOP: 0x00000019 (0x00000000,0xC00E0FF8,0xFFFFFD4,0xC0000000)
BAD_POOL_HEADER

CPUID: GenuineIntel 5.2.c i941:1f SVSVER 0x0000565

Dll Base DateStmp - Name Dll Base DateStmp - Name
01400000 200c1974 - ntoskrnl.exe 00010000 31ee0352 - hml.dll
00001000 31e40624 - advapi32.sys 00006000 31e60274 - CSIPORT.SVS
00241000 31e60624 - advapi32.sys 00250000 31e40424 - Nlsk.sys
c90a0000 31e604f2 - FcAdv.SVS c9c90000 31e60694 - Null.SVS
c0640000 31e60410 - RndObj.SVS c0c90000 31e60724 - Bep.SVS
c0400000 31e60698 - I8042prt.sys c0660000 31e60692 - mouseclass.sys
c0740000 31e60674 - Mouse.sys c0400000 31f50724 - UIDEPORT.SVS
e7fab000 31e60662 - nva_wil.sys c0990000 31e60664 - vga.sys
c7800000 31e60624 - NDIS.SVS 40000000 31f934f7 - win32k.sys
ef8c0000 31f93451 - mui.dll feaf3000 31e60667 - Fastfat.SVS
feb80000 31e60666 - TDI.SVS feaf0000 31e60724 - nbf.sys
eaf10000 31f93482 - GDI32.sys feaf5000 31f50663 - netbt.sys
fc500000 31e60186 - e159k.sys fc560000 31f8f864 -afd.sys
c7300000 31e60674 - netfins.sys fc850000 31e60694 - Parport.sys
c0700000 31e60694 - Parallel.SVS fc950000 31e60694 - ParVdm.SVS
c0300000 31e60614 - Serial.SVS fe400000 31f50810 - adv.sys
fe430000 31f7a1ba - mup.sys fe9d0000 32031abe - svx.sys

Address dump Build [1381] - Name
00000000 00000000 00000000 ffffffff 00000002 - KeCDD.SVS
01471c0 00144000 00144000 ffffffff 00000001 - ntoskrnl.exe
01471d0 00122000 f0032e00 f0280000 c1220004 c1220000 - ntoskrnl.exe
0147304 003023f0 0000023c 00000034 00000000 00000000 - ntoskrnl.exe

Restart and set the recovery options in the system control panel
or the /CRASHDEBUG system start option.
    
```

Concurrency is a problem

- Windows 2000 hot fixes
 - Concurrency errors most common defects among “detectable errors”
 - Incorrect synchronization and protocol errors most common defects among all coding errors
- Windows Server 2003 late cycle defects
 - Synchronization errors second in the list, next to buffer overruns

Security vulnerabilities involving race conditions

- Buffer overruns
- Phishing attacks
 - “A systematic approach to uncover visual ambiguity vulnerabilities”, MSR Technical Report MSR-TR-2006-48 by Chen et al.

Economic Impact

- NIST study


Last year, a study commissioned by the National Institute of Standards and Technology found that software errors cost the U.S. economy about \$59.5 billion annually, or about 0.6 percent of the gross domestic product. More than half the costs are borne by software users, the rest by developers and vendors.

<http://www.nist.gov/director/prog-ofc/report02-3.pdf>

Non-Determinism, Heisenbugs

- Multithreaded programs are non-deterministic
 - behavior depends on interleaving of threads
- Extremely difficult to test
 - exponentially many interleavings
 - during testing, many interleavings behave correctly
 - post-deployment, other interleavings fail
- Complicates code reviews, static analysis, ...

Mars, July 4, 1997
Lost contact due to real-time priority inversion bug



400 horses
100 microprocessors



Bank Account Implementation


```
class Account {
    private int bal = 0;

    public void deposit(int n) {
        int j = bal;
        bal = j + n;
    }
}
```

Bank Account Implementation


```
class Account {
    private int bal = 0;

    public void deposit(int n) {
        int j = bal;
        bal = j + n;
    }
}
```



Bank Account Implementation

A *race condition* occurs if two threads access a shared variable at the same time, and at least one of the accesses is a write



Race Conditions

```
class Ref {
    int i;
    void add(Ref r) {
        i = i
        + r.i;
    }
}
```

Race Conditions

```
class Ref {
    int i;
    void add(Ref r) {
        i = i
        + r.i;
    }
}
```

```
Ref x = new Ref(0);
Ref y = new Ref(3);
```

```
x.add(y);
x.add(y);
```

```
assert x.i == 6;
```

Race Conditions

```
class Ref {
    int i;
    void add(Ref r) {
        i = i
        + r.i;
    }
}
```

```
Ref x = new Ref(0);
Ref y = new Ref(3);
```

```
parallel {
    x.add(y); // two calls happen
    x.add(y); // in parallel
}
```

```
assert x.i == 6;
```

Race condition on x.i

Assertion may fail

Lock-Based Synchronization

```
class Ref {
  int i; // guarded by this
  void add(Ref r) {
    i = i
    + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
```

- Every shared memory location protected by a lock
- Lock must be held before any read or write of that memory location

When Locking Goes Bad ...

- Hesienbugs (race conditions, etc) are common and problematic
 - forget to acquire lock, acquire wrong lock, etc
 - extremely hard to detect and isolate
- Traditional type systems are great for catching certain errors
- *Type systems for multithreaded software*
 - detect race conditions, atomicity violations, ...

Verifying Race Freedom with Types

```
class Ref {
  int i;
  void add(Ref r) {
    i = i
    + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
```

Verifying Race Freedom with Types

```
class Ref {
  int i guarded_by this;
  void add(Ref r) requires this, r {
    i = i
    + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
```

→ check: $this \in \{this, r\}$

Verifying Race Freedom with Types

```
class Ref {
  int i guarded_by this;
  void add(Ref r) requires this, r {
    i = i
    + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
```

→ check: $this \in \{this, r\}$ ✓

→ check: $this[this:=r] = r \in \{this, r\}$ ✓

replace this by r

Verifying Race Freedom with Types

```
class Ref {
  int i guarded_by this;
  void add(Ref r) requires this, r {
    i = i
    + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
```

→ check: $this \in \{this, r\}$ ✓

→ check: $this[this:=r] = r \in \{this, r\}$ ✓

replace formals this, r by actuals x, y

→ check: $\{this, r\}[this:=x, r:=y] \in \{x, y\}$ ✓

Verifying Race Freedom with Types

```

class Ref {
  int i guarded_by this;
  void add(Ref r) requires this, r {
    i = i
    + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;

```

check: $\text{this} \in \{ \text{this}, r \}$ ✓
 check: $\text{this}[\text{this}:=r] = r \in \{ \text{this}, r \}$ ✓

replace formals this,r by actuals x,y

check: $\{ \text{this}, r \} [\text{this}:=x, r:=y] \in \{ x, y \}$ ✓
 check: $\{ \text{this}, r \} [\text{this}:=x, r:=y] \in \{ x, y \}$ ✓

Soundness Theorem:
Well-typed programs are race-free

One Problem ...

```

Object o;
int x guarded_by o;

fork { sync(o) { x++; } }

fork { o = new Object();
      sync(o) { x++; }
}

```

- Lock expressions must be constant

Lock Equality

- Type system checks if lock is in lock set
 - $r \in \{ \text{this}, r \}$
 - same as $r = \text{this} \vee r = r$
- Semantic equality
 - $e_1 = e_2$ if e_1 and e_2 refer to same object
 - need to test whether two program expressions evaluate to same value
 - undecidable in general

Lock Equality

- Approximate (undecidable) semantic equality by syntactic equality
 - two locks expressions are considered equal only if syntactically identical
- Conservative approximation


```

class A {
  void f() requires this { ... }
}


A p = new A();
q = p;
synch(q) { p.f(); }

```

$\text{this}[\text{this}:=p] = p \in \{ q \}$ ✗
- Not a major source of imprecision

RaceFreeJava

- Concurrent extension of CLASSICJAVA [Flatt-Krishnamurthi-Felleisen 99]
- Judgement for typing expressions



Typing Rules

- Thread creation


$$\frac{P; E; \emptyset \vdash e : t}{P; E; ls \vdash \text{fork } e : \text{int}}$$
- Synchronization

$$\frac{P; E \vdash_{\text{final}} e_1 : c \quad \text{lock is constant} \quad P; E; ls \cup \{e_1\} \vdash e_2 : t}{P; E; ls \vdash \text{synchronized } e_1 \text{ in } e_2 : t}$$

Field Access

$$\frac{P; E; ls \vdash e : c \quad e \text{ has class } c \quad P; E \vdash (t \text{ fd guarded by } l) \in c \quad \text{fd } l \text{ is declared in } c \quad P; E \vdash [e/\text{this}]l \in ls \quad \text{lock } l \text{ is held}}{P; E; ls \vdash e.\text{fd} : [e/\text{this}]t}$$

java.util.Vector




```
class Vector {
  Object elementData[] /*# guarded_by this */;
  int elementCount /*# guarded_by this */;

  synchronized int lastIndexOf(Object elem, int n) {
    for (int i = n ; i >= 0 ; i--)
      if (elem.equals(elementData[i])) return i;
    return -1;
  }

  int lastIndexOf(Object elem) {
    return lastIndexOf(elem, elementCount - 1);
  }

  synchronized void trimToSize() { ... }
  synchronized boolean remove(int index) { ... }
}
```

java.util.Vector



```
class Vector {
  Object elementData[] /*# guarded_by this */;
  int elementCount /*# guarded_by this */;

  synchronized int lastIndexOf(Object elem, int n) {
    for (int i = n ; i >= 0 ; i--)
      if (elem.equals(elementData[i])) return i;
    return -1;
  }

  int lastIndexOf(Object elem) {
    return lastIndexOf(elem, elementCount - 1);
  }

  synchronized void trimToSize() { ... }
  synchronized boolean remove(int index) { ... }
}
```

Validation of rccjava

| Program | Size (lines) | Number of annotations | Annotation time (hrs) | Races Found |
|-----------|--------------|-----------------------|-----------------------|-------------|
| Hashtable | 434 | 60 | 0.5 | 0 |
| Vector | 440 | 10 | 0.5 | 1 |
| java.io | 16,000 | 139 | 16.0 | 4 |
| Ambit | 4,500 | 38 | 4.0 | 4 |
| WebL | 20,000 | 358 | 12.0 | 5 |

Basic Type Inference

```
class Ref {
  int i;
  void add(Ref r) {
    i = i + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
```

Basic Type Inference

```
static final Object m = new Object();

class Ref {
  int i;
  void add(Ref r) {
    i = i + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
```

Iterative GFP algorithm:

- [Flanagan-Freund, PASTE'01]
- Start with maximum set of annotations

Basic Type Inference

```

static final Object m = new Object();

class Ref {
  int i guarded_by this, m;
  void add(Ref r) {
    i = i + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
    
```

- Iterative GFP algorithm:
- [Flanagan-Freund, PASTE'01]
 - Start with maximum set of annotations

Basic Type Inference

```

static final Object m = new Object();

class Ref {
  int i guarded_by this, m;
  void add(Ref r) requires this, r, m {
    i = i + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
    
```

- Iterative GFP algorithm:
- [Flanagan-Freund, PASTE'01]
 - Start with maximum set of annotations

Basic Type Inference

```

static final Object m = new Object();

class Ref {
  int i guarded_by this, x;
  void add(Ref r) requires this, r, x {
    i = i + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
    
```

- Iterative GFP algorithm:
- [Flanagan-Freund, PASTE'01]
 - Start with maximum set of annotations
 - Iteratively remove all incorrect annotations

Basic Type Inference

```

static final Object m = new Object();

class Ref {
  int i guarded_by this, x;
  void add(Ref r) requires this, r, x {
    i = i + r.i;
  }
}

Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (x,y) { x.add(y); }
  synchronized (x,y) { x.add(y); }
}
assert x.i == 6;
    
```

- Iterative GFP algorithm:
- [Flanagan-Freund, PASTE'01]
 - Start with maximum set of annotations
 - Iteratively remove all incorrect annotations
 - Check each field still has a protecting lock

Sound, complete, fast

But type system too basic

Harder Example: External Locking

```

class Ref {
  int i;
  void add(Ref r) {
    i = i + r.i;
  }
}

Object m = new Object();
Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (m) { x.add(y); }
  synchronized (m) { x.add(y); }
}
assert x.i == 6;
    
```

- Field *i* of *x* and *y* protected by *external* lock *m*
- Not typable with basic type system
 - *m* not in scope at *i*
- Requires more expressive type system with *ghost parameters*

Ghost Parameters on Classes

```

class Ref {
  int i;
  void add(Ref r) {
    i = i + r.i;
  }
}

Object m = new Object();
Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (m) { x.add(y); }
  synchronized (m) { x.add(y); }
}
assert x.i == 6;
    
```

Ghost Parameters on Classes

```
class Ref<ghost g> {
  int i;
  void add(Ref r) {
    i = i + r.i;
  }
}

Object m = new Object();
Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (m) { x.add(y); }
  synchronized (m) { x.add(y); }
}
assert x.i == 6;
```

- Ref parameterized by external ghost lock g

Ghost Parameters on Classes

```
class Ref<ghost g> {
  int i guarded_by g;
  void add(Ref r) {
    i = i + r.i;
  }
}

Object m = new Object();
Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (m) { x.add(y); }
  synchronized (m) { x.add(y); }
}
assert x.i == 6;
```

- Ref parameterized by external ghost lock g
- Field i guarded by g

Ghost Parameters on Classes

```
class Ref<ghost g> {
  int i guarded_by g;
  void add(Ref r) requires g {
    i = i + r.i;
  }
}

Object m = new Object();
Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (m) { x.add(y); }
  synchronized (m) { x.add(y); }
}
assert x.i == 6;
```

- Ref parameterized by external ghost lock g
- Field i guarded by g
- g held when add called

Ghost Parameters on Classes

```
class Ref<ghost g> {
  int i guarded_by g;
  void add(Ref<g> r) requires g {
    i = i + r.i;
  }
}

Object m = new Object();
Ref x = new Ref(0);
Ref y = new Ref(3);
parallel {
  synchronized (m) { x.add(y); }
  synchronized (m) { x.add(y); }
}
assert x.i == 6;
```

- Ref parameterized by external ghost lock g
- Field i guarded by g
- g held when add called
- Argument r also parameterized by g

Ghost Parameters on Classes

```
class Ref<ghost g> {
  int i guarded_by g;
  void add(Ref<g> r) requires g {
    i = i + r.i;
  }
}

Object m = new Object();
Ref<m> x = new Ref<m>(0);
Ref<m> y = new Ref<m>(3);
parallel {
  synchronized (m) { x.add(y); }
  synchronized (m) { x.add(y); }
}
assert x.i == 6;
```

- Ref parameterized by external ghost lock g
- Field i guarded by g
- g held when add called
- Argument r also parameterized by g
- x and y parameterized by lock m

Type Checking Ghost Parameters

```
class Ref<ghost g> {
  int i guarded_by g;
  void add(Ref<g> r) requires g {
    i = i + r.i;
  }
}

Object m = new Object();
Ref<m> x = new Ref<m>(0);
Ref<m> y = new Ref<m>(3);
parallel {
  synchronized (m) { x.add(y); }
  synchronized (m) { x.add(y); }
}
assert x.i == 6;
```

check: {g} [this:=x,r:=y, g:=m] ∈ {m} ✓

Type Inference with Ghosts

- Type inference is NP-complete
 - iterative GFP algorithm does not work
 - ghost parameters require backtracking search
- RccSAT: Reduce to SAT
 - works up to 30 KLOC
 - precise: 92-100% of fields verified race free