

Comparison Of Two Approaches For Data-Race Prevention

Erik Kundt

Freie Universität Berlin

Seminar on Software Engineering 2015

Outline

Motivation

Introduction

- Key contributions
- Ideas and elements
- The type system

Formalization

- Grammars
- Type checker

Examples

- Bounded-Buffer Producer-Consumer (BBPC)
- TStack

Conclusion

- ▶ When an application depends on the sequence or timing of processes or threads for it to operate properly
- ▶ A data-race happens when there are two memory accesses in a program where both:
 - ▶ Target the same location
 - ▶ Are performed concurrently by two threads
 - ▶ Are not reads
 - ▶ Are not synchronization operations

Motivation: Improve design quality

- ▶ Attributes of software quality:
 - ▶ **efficiency**, modifiability, readability, **correctness**, **robustness**, usability
- ▶ Decrease defect rate (here: synchronization)
 - ▶ Counted during inspections and testings (expensive)
 - ▶ Less likely to produce synchronization defects because of reduced complexity
 - ▶ Static formal verification
- ▶ Difficult to find empirical analyses of the root cause of defects in the literature

- ▶ Rust[3]: Rust tasks
 - ▶ Like threads
 - ▶ Developer can choose the low-level details of how they operate
 - ▶
- ▶ Go[4]: goroutine
 - ▶ independent concurrent thread of control
 - ▶ same address space
 - ▶ Go has runtime (dynamic checks)
- ▶ Paradigm shift: Functional programming
 - ▶ High degree of parallelization (Scala)

Motivation: Rust

```
1 // tx is the sending half (tx for transmission), and rx is the receiving
2 // half (rx for receiving).
3 let (tx, rx) = channel();
4
5 // Spawn off an expensive computation
6 spawn(proc() {
7     tx.send(expensive_computation());
8 });
9
10 // Let's see what that answer was
11 println!("{}", rx.recv());
```

A

A Time-Aware Type System For Data-Race Protection and Guaranteed Initialization[1]

Nicholas D. Matsakis and Thomas R. Gross
ETH Zurich, Switzerland
{nmatsaki,trg}@inf.ethz.ch

B

Ownership Types for Safe Programming: Preventing Data Races and Deadlocks[2]

Chandrasekhr Boyapati, Robert Lee and Martin Rinard
Laboratory of Computer Science, MIT, Cambridge
{chandra,rhlee,rinard}@lcs.mit.edu

Outline

Motivation

Introduction

- Key contributions

- Ideas and elements

- The type system

Formalization

- Grammars

- Type checker

Examples

- Bounded-Buffer Producer-Consumer (BBPC)

- TStack

Conclusion

- ▶ Intervals
- ▶ Type system that is aware of intervals and the happens before relation
- ▶ Flexible data-race protection:
 - ▶ lock-free and locked parallel patterns
 - ▶ field-by-field, object-by-object
- ▶ Controlled object initialization:
 - ▶ access control system for fields and methods
- ▶ Language *Inter* and a suitable type checker

- ▶ Ownership types
- ▶ Type system that is aware of ownerships in Java programs
- ▶ Simple data-race and deadlock prevention:
 - ▶ lock levels and partial order among them
 - ▶ lock level polymorphism
- ▶ Language extensions for Race-Free Java
- ▶ Formal rules for type checking
- ▶ Runtime ordering of locks

Outline

Motivation

Introduction

Key contributions

Ideas and elements

The type system

Formalization

Grammars

Type checker


Examples

Bounded-Buffer Producer-Consumer (BBPC)

TStack

Conclusion

- ▶ Time as first-class objects
- ▶ Represents the span of program time in which a certain piece of code executes
- ▶ Statement, method call or asynchronous task
- ▶ Prevent data-races: Partially ordered through a happens-before relation
- ▶ Can be nested inside another interval



```

1  class Subintervals(Interval parent, Lock l) {
2      interval a(this.parent) { /* code for a */ }
3      interval b(this.parent) { /* code for b */ }
4      interval c(this.parent) { /* code for c */ }
5      this.a hb this.c;
6      this.b hb this.c;
7      this.c locks this.l;
8  }

```

- ▶ Every object has an owner
- ▶ Locking discipline specified using type declarations
- ▶ Statically enforceable way of specifying object encapsulation
- ▶ A class definition in Race-Free Java is parameterized by a list of owners

```
1 class BalancedTree {
2     LockLevel l = new;
3     Node<self:l> root = new Node;
4 }
5
6 class Node<self:k> {
7     tree Node<self:k> left;
8     tree Node<self:k> right;
9
10    synchronized void rotateRight() locks(this) {
11        final Node x = this.right; if (x == null) return;
12        synchronized (x) {
13            final Node v = x.left; if (v == null) return;
14            synchronized (v) {
15                final Node w = v.right;
16                v.right = null;
17                x.left = w;
18                this.right = v;
19                v.right = x;
20            }
21        }
22    }
```

Outline

Motivation

Introduction

Key contributions

Ideas and elements

The type system

Formalization

Grammars

Type checker

Examples

Bounded-Buffer Producer-Consumer (BBPC)

TStack

Conclusion


```
1 // Objects which will serve as guards must implement the interface Guard
2 Interface Guard {
3
4     // checks whether the guard g permits i to write to the fields guarded by g
5     boolean permitsWr(Interval current);
6
7     // same as permitsWr but for reads
8     boolean permitsRd(Interval current);
9
10    // checks whether the fields protected by g are immutable for the interval i
11    boolean ensuresImm(Interval current);
12 }
```

- ▶ Intervals and locks implement this interface

Outline

Motivation

Introduction

Key contributions

Ideas and elements

The type system

Formalization

Grammars

Type checker

Examples

Bounded-Buffer Producer-Consumer (BBPC)

TStack

Conclusion

Example

```
cdecl      := class c(tys fs) extends c(paths) { members }
member    := gdecl | fdecl | mdecl | idecl | ldecl | hbdecl
gdecl     := ghost f
fdecl     := ty f guardedBy path
mdecl     := void m(tys xs) reqs { lstmts }
req       := path rel path
idecl     := interval f(path) { lstmts }
ldecl     := path locks path
hbdecl    := path hb path
path      := x.fs
rel       := trel | wcrel | locks
trel      := subOf | inlineSubOf | hb
wcrel     := permitsWr | permitsRd | ensuresImm | eq
lstmt    := x : stmt
```

Example

```
stmt  := x.f = x
      | x = x.f
      | x = new c[fs eq paths] (xs)
      | x.m(xs)
      | assert x wcrel x
ty   := c[fs wcrels paths]
```

Example

```
P      := defn* e
defn   := class cn extends c body
c      := cn | Object
body   := field* meth*
meth   := t mn(arg*) e
field  := [final] t fd = e
arg    := [final] t x
t      := c | int | boolean
e      := new c | x | x = e | e.fd | e.fd = e | e.mn(e*)
        e;e | let ( arg = e ) ine | if (e) then e |
        synchronized (e) in e | fork (x*) e
```

```
cn     ∈ class names
fd     ∈ field names
mn     ∈ method names
x      ∈ variable names
```

Example

defn := class *c*<*owner formal**> extends *c* *body*

c := *cn*<*owner+*> | Object

c := *cn*<*owner+*> | Object

owner := *formal* | self | thisThread | *efinal*

meth := *t mn*(*arg**) accesses |(*efinal**) *e*

efinal := *e*

formal := *f*

f ∈ owner names

Outline

Motivation

Introduction

Key contributions

Ideas and elements

The type system

Formalization

Grammars

Type checker

Examples

Bounded-Buffer Producer-Consumer (BBPC)

TStack

Conclusion

- ▶ Heart of the type checker
- ▶ Records facts deduced by compiler
- ▶ As compiler proceeds, new facts are added; nothing is ever removed
- ▶ Facts in the environment take the form of tuples:
 - ▶ *path rel path* relate two paths
 - ▶ $x : ty$ record the type for a local variable

- ▶ A path is just a local variable name[2]
- ▶ Facts about paths have to be always true as the program executes
- ▶ Will not become invalidated as fields are updated
- ▶ Ensuring that all paths referenced from the environment are *stable*
- ▶ Always stable:
 - ▶ `x` or `this`, variable reassignment is forbidden
 - ▶ `this.parent`, `parent` is constructor argument and immutable
- ▶ Only stable at certain times:
 - ▶ `aResult`: When the class is first created, interval `a` has not yet executed, and so `aResult` is not stable

STMT-LOAD

$$\frac{\mathcal{E} \vdash x_o : c \quad \text{reified}(c, f) = (ty_f; path_g) \quad \Theta = [\text{this} \rightarrow x_o] \quad \mathcal{E} \vdash \Theta(path_g) \text{stableBy } x_l \quad \mathcal{E} \vdash \Theta(path_g) \text{permitsRd } x_l}{\mathcal{E} \vdash x_l : (x_d = x_o.f) \rightarrow \mathcal{E} + (x_d : \Theta(ty_f))}$$

STMT-LOAD-IMMUTABLE

$$\frac{\mathcal{E} \vdash x_o : c \quad \text{reified}(c, f) = (ty_f; path_g) \quad \Theta = [\text{this} \rightarrow x_o] \quad \mathcal{E} \vdash \Theta(path_g) \text{stableBy } x_l \quad \mathcal{E} \vdash \Theta(path_g) \text{ensuresImm } x_l}{\mathcal{E} \vdash x_l : (x_d = x_o.f) \rightarrow \mathcal{E} + (x_d : \Theta(ty_f)) + (x_d \text{eq } x_o.f)}$$

STMT-STORE

$$\frac{\mathcal{E} \vdash x_o : c \quad \text{reified}(c, f) = (ty_f; path_g) \quad \Theta = [\text{this} \rightarrow x_o] \quad \mathcal{E} \vdash \Theta(path_g) \text{stableBy } x_l \quad \mathcal{E} \vdash \Theta(path_g) \text{permitsWr } x_l \quad \mathcal{E} \vdash x_v : \Theta(ty_f)}{\mathcal{E} \vdash x_l : (x_o.f = x_v) \rightarrow \mathcal{E}}$$

- ▶ Set of rules for reasoning about the typing judgement
- ▶ Holds information about class definitions
- ▶ Judgements based on state of environment
- ▶ Ensures that the program is well typed

1. The owner of an object does not change over time
2. The ownership relation forms a forest of rooted trees, where the roots can have self loops.
3. The necessary and sufficient condition for a thread to access to an object is that the thread must hold the lock on the root of the ownership tree that the object belongs to.
4. Every thread implicitly holds the lock on the corresponding `thisThread` owner. A thread can therefore access any object owned by its corresponding `thisThread` owner without any synchronization.

[METHOD]

$$E' = E, \text{arg}_{1..n}, \text{locks}(cn_j.l_j \ j \in 1..k \ [l]_{\text{opt}})$$

$$P; E' \vdash_{\text{final}} e_i : t_i \quad P; E' \vdash \text{RootOwner}(e_i) = r_i$$

$$ls = \text{thisThread}, r_{1..r}$$

$$l_{\min} = \text{LUB}(cn_j.l_j \ j \in 1..k)$$

$$P; E'; ls; l_{\min} \vdash e : t$$

$$P; E \vdash t \text{ mn}(\text{arg}_{1..n}) \text{ accesses}(e_{1..r})$$

$$\text{locks}(cn_j.l_j \ j \in 1..k \ [l]_{\text{opt}}) \{e\}$$

Outline

Motivation

Introduction

- Key contributions

- Ideas and elements

- The type system

Formalization

- Grammars

- Type checker


Examples

- Bounded-Buffer Producer-Consumer (BBPC)

- TStack

Conclusion

- ▶ Two independent parallel tasks
- ▶ Producer task writes data into a buffer, consumer reads it out
- ▶ If producer is too fast, buffer becomes full, producer has to wait for consumer to catch up
- ▶ Reference: synchronous programming



```

1  abstract class Stream {
2      Stream next guardedBy inter;
3      abstract interval inter;
4  }

```



```
1 // this == c[i], prod == p[i]
2 class Consumer(Interval parent, Producer prod) extends Stream {
3     prod.inter hb inter; // p[i] -> c[i]
4     interval inter(parent) {
5         /* consume prod.result */
6         next = new Consumer(parent, (Producer)prod.next);
7     }
8 }
```

```
1 // this == p[i], cons == c[i-2]
2 class Producer(Interval parent, Stream cons) extends Stream {
3   Object result guardedBy inter;
4   cons.inter hb inter; // c[i-2] -> p[i]
5   interval inter(parent) {
6     result = /* produce item */;
7     next = new Producer(parent, cons.next);
8   }
9 }
```

```
1  class DummyConsumer(Interval parent) extends Stream {
2      Stream link guardedBy parent;
3      interval inter(parent) {
4          next = link;
5      }
6  }
7
8  class Start {
9      void main() {
10         join: {
11             DummyConsumer one = new DummyConsumer(join);
12             DummyConsumer two = new DummyConsumer(join);
13             Producer prod = new Producer(join, one);
14             one.link = two;
15             two.link = new Consumer(join, prod);
16         }
17     }
18 }
```

Outline

Motivation

Introduction

- Key contributions

- Ideas and elements

- The type system

Formalization

- Grammars

- Type checker

Examples

- Bounded-Buffer Producer-Consumer (BBPC)

- TStack

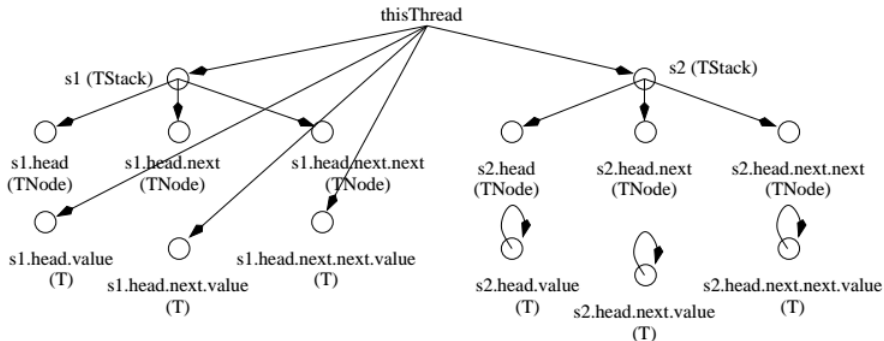
Conclusion

- ▶ Stack of T objects
- ▶ TStack is implemented using a linked list

```
1 // thisOwner owns the TStack object
2 // TOwner owns the T objects in the stack.
3 class TStack<thisOwner, TOwner> {
4     TNode<this, TOwner> head = null;
5
6     T<TOwner> pop() accesses (this) {
7         if (head == null) return null;
8         T<TOwner> value = head.value();
9         head = head.next();
10        return value;
11    }
12    ...
13 }
```

```
1  class TNode<thisOwner, TOwner> {
2      T<TOwner> value;
3      TNode<thisOwner, TOwner> next;
4
5      T<TOwner> value() accesses (this) {
6          return value;
7      }
8
9      TNode<thisOwner, TOwner> next() accesses (this) {
10         return next;
11     }
12     ...
13 }
```

```
1  class T<thisOwner> { int x=0; }  
2  
3  TStack<thisThread, thisThread> s1 = new TStack<thisThread, thisThread>;  
4  TStack<thisThread, self> s2 = new TStack<thisThread, self>;
```



- ▶ Both give a practical approach for giving users the ability to implement lock discipline
- ▶ **Race-Free Java** more practical in comparison to **Inter**
- ▶ Proven that the type systems **protects against data races**
- ▶ Usage of similar technique in Rust proves that approach is suitable

For Further Reading I

-  Matsakis, N. and Gross, T. *A Time-Aware Type System For Data-Race Protection and Guaranteed Initialization*. ETH Zurich, Switzerland, published in OOPSLA/SPLASH 2010
-  Boyapati, C. and Lee, R. and Rinard, M. *Ownership Types for Safe Programming: Preventing Data Races and Deadlocks*. Laboratory of Computer Science, MIT, Cambridge, published in OOPSLA 2002
-  Rust programming language *Rust*. <http://www.rust-lang.org>
-  Go programming language *Go*. <http://www.golang.org>
-  Matsakis, N. *E-Mail conversation*. 03.11.2014