



IBM Research

# A Type System for Data-Centric Synchronization

**Frank Tip and Jan Vitek<sup>†</sup>**

Mandana Vaziri, Julian Dolby, Christian Hammer<sup>†</sup>

IBM T.J. Watson Research Center and

<sup>†</sup>Purdue University

## Concurrency control is difficult

- Low-level data race occur when threads access a location concurrently (at least one write) with no synchronization
- Locking discipline involves non-local reasoning
- Even if every shared access is protected, data may still end up in an inconsistent state

## A High-Level Data Race

```
class Vector {  
    Object[] data;  
    int count;  
    ...  
    synchronized int size() { ... }  
  
    synchronized boolean addAll(Collection c) {  
        int size = c.size();  
        ...  
        while (it.hasNext())  
            data[count++] = it.next();  
        ...  
    }  
}
```

synchronization is about preserving  
data consistency

so why not associate synchronization  
constraints directly with data?

# Data-centric Synchronization with AJ

# Data-Centric Synchronization: Terminology

```
class BankAccount {  
    int checking, savings;  
    int transferCount;  
  
    void transfer(int amount) {  
        synchronized(this) {  
            checking -= amount;  
            savings += amount;  
        }  
        transferCount++;  
        ...  
    }  
}
```

# Data-Centric Synchronization: Terminology

```
class BankAccount {  
    int checking, savings,  
    int transferCount;  
  
    void transfer(int amount) {  
        synchronized(this) {  
            checking -= amount;  
            savings += amount;  
        }  
        transferCount++;  
        ...  
    }  
}
```

**atomic set S**

# Data-Centric Synchronization: Terminology

```
class BankAccount {  
    int checking, savings,  
    int transferCount;
```

**atomic set S**

**unit of work on S** —  
*preserves consistency  
of S when executed  
sequentially*

```
void transfer(int amount) {  
    synchronized(this) {  
        checking -= amount;  
        savings += amount;  
    }  
    transferCount++;  
    ...  
}
```

```
}
```

# Atomic Sets

```
class Counter {  
    atomicset a;  
    atomic(a) int val;  
  
    Counter() { val = 0; }  
  
    int get() { return val; }  
    void dec() { val--; }  
    void inc() { val++; }  
}
```

# Atomic Sets

```
class Counter {  
    atomicset a;  
    atomic(a) int val;  
  
    Counter() { val = 0; }  
  
    int get() { return val; }  
    void dec() { val--; }  
    void inc() { val++; }  
}
```

Counter c = new Counter();

Thread 1

c.inc();  
c.dec();

Thread 2


c.inc();

---

*c.val has value 1 when both threads have terminated*

# Aliasing

*the atomic set `b` in the object pointed to by `this` is merged with atomic set `a` in the Counter object*



```
class PairCounter {  
    atomicset b;  
    atomic(b) int diff;  
    atomic(b) Counter|a=this.b| low = new Counter|a=this.b| ();  
    atomic(b) Counter|a=this.b| high = new Counter|a=this.b| ();  
  
    void incHigh() {  
        high.inc();  
        diff = high.get() - low.get();  
    }  
}
```

# unitfor

```
class Transfer {  
    void transfer(unitfor(a) Counter from,  
                 unitfor(a) Counter to) {  
        from.dec();  
        to.inc();  
    }  
}
```

## A more realistic example

```
public abstract class AbsList {  
    atomicset a;  
    atomic(a) int size;  
  
    public int size(){ return size; }  
    public abstract ListIterator iterator();  
    public abstract void add(Object o);  
    public abstract boolean addAll(unitfor(a) AbsList c);  
    ...  
}
```

## A more realistic example (continued)

```
class LinkedList extends AbsList {  
    atomic(a) Entry|b=this.a| header;  
  
    public LinkedList(){  
        header = new Entry|b=this.a| (null,null,null);  
        header.next = header.prev = header;  
    }  
    ...  
}  
  
internal class Entry {  
    atomicset b;  
    atomic(b) Object elem;  
    atomic(b) Entry|b=this.b| next;  
    atomic(b) Entry|b=this.b| prev;  
    ...  
}
```

# Implementing AJ

# Implementation

- **source-to-source translator with Eclipse**
  - atomic set annotations entered as Java comments
  - implementation handles Java subset (no generics, inner classes)
  - support alias annotations for arrays `|this.M[]F=this.M|`
  - translation generates standard synchronized block

# Java Collections Framework

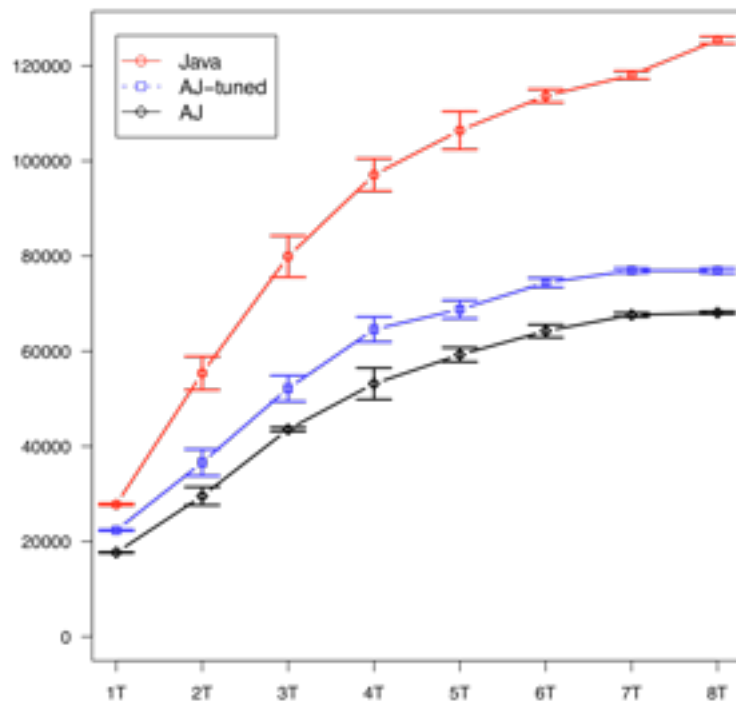
- Selected classes from Java Collections
  - ArrayList, LinkedList, HashMap, HashSet, TreeMap...+ java.util dependencies
  - 63 types and 10,860 LOC
- Introduced atomic sets
  - one atomicset for each of 5 subhierarchies, includes all instance fields
  - alias annotations to relate iterators to their “owner”
  - one class made internal (LinkedList.Entry)

annotation	#
<b>atomicset</b>	<b>0</b>
<b>atomic class</b>	<b>5</b>
<b>atomic</b>	<b>0</b>
<b>unitfor</b>	<b>55</b>
<b>alias</b>	<b>330</b>
<b>array object</b>	<b>24</b>
<b>array element</b>	<b>16</b>
<b>TOTAL</b>	<b>430</b>

- ~1 annotation / 25 LOC

# SPECjbb2005

- Widely used multi-threaded benchmark (8KLOC)
  - inconsistent/redundant synchronization
- Translation into AJ
  - for tuned version refactored some fields to make them final



type	#
<b>atomicset</b>	<b>1</b>
<b>atomic class</b>	<b>14</b>
<b>atomic</b>	<b>25</b>
<b>unitfor</b>	<b>0</b>
<b>alias</b>	<b>8</b>
<b>array object</b>	<b>0</b>
<b>array element</b>	<b>1</b>
<b>TOTAL</b>	<b>49</b>

- ~1 annotation / 160 LOC
- removed 125 occurrences of synchronized
- 25 synchronized remain, related to wait/notify

# Formalizing AJ

## Atomic Set Serializability

- AJ guarantees

### Theorem

*Given a well-formed trace  $T$  and atomic set  $R$ , events of each units of work of  $R$  happen serially.*

## Modeling AJ

To be tractable, select a subset of AJ:

( FeatherweightJava + state + atomicsets ) + threads

or

AJ - *unitfor* - *arrays* - *modifiers* - *generics* - *primitives* - *exceptions*  
 - *finals* - *statics* - *new\_threads* - *interfaces*

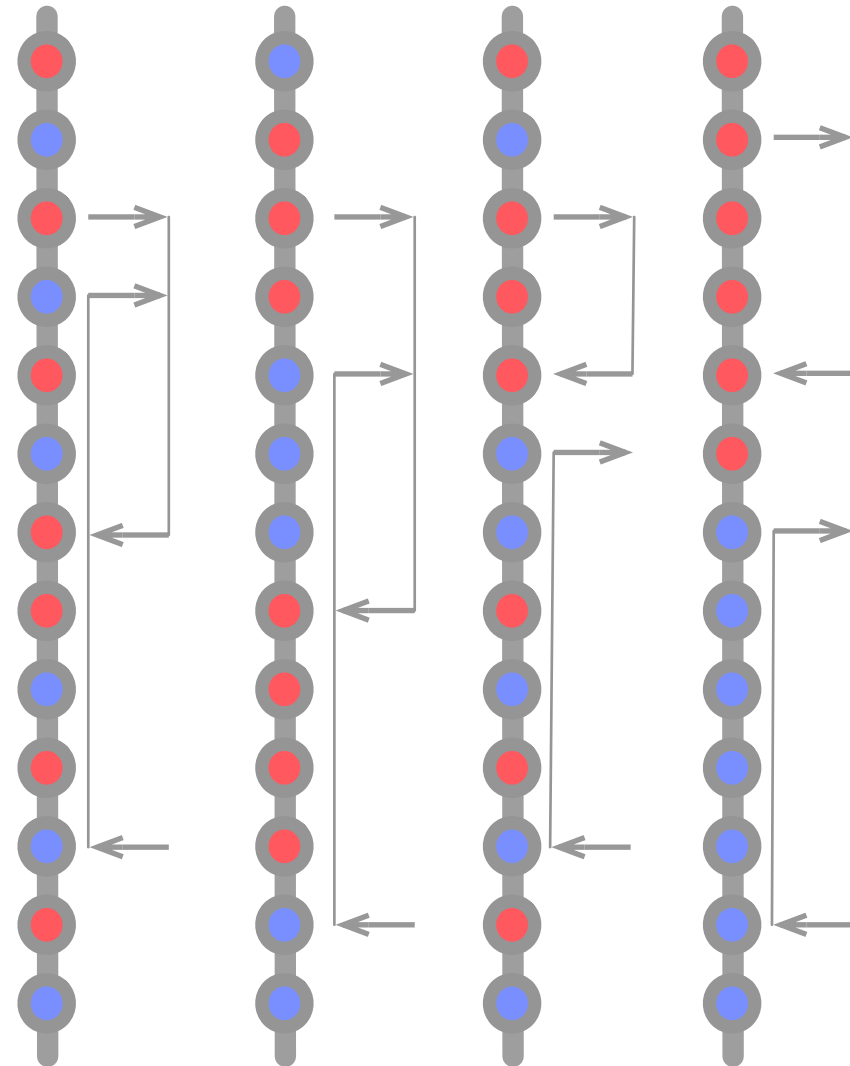
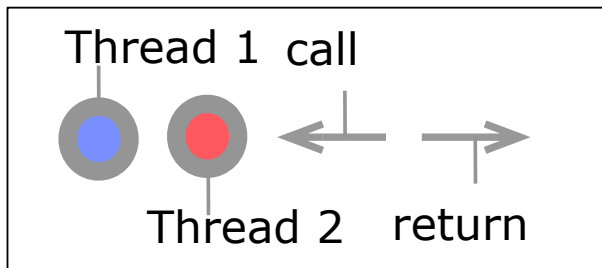
$p ::= \overline{cd}$	<i>program</i>	$\tau ::= C a=this.b  \mid C$	<i>type</i>
$cd ::= \iota \text{ class } C \text{ extends } D \{ as \overline{fd} \overline{md} \}$	<i>class</i>	$\alpha ::= \text{atomic}(a) \mid \epsilon$	
$as ::= \text{atomicset } a \mid \epsilon$		$\iota ::= \text{internal} \mid \epsilon$	
$fd ::= \alpha \tau f$	<i>field</i>		
$md ::= \tau m (\overline{\tau x}) \{ \overline{\tau z}; s; \text{return } y \}$	<i>method</i>	$E ::= [] \mid E[x : \tau]$	<i>type env</i>
$s ::= s; s \mid \text{skip} \mid x = \text{this.f} \mid x = (\tau)y \mid$	<i>statement</i>		
$\text{this.f} = z \mid x = \text{new } \tau () \mid x = y.m (\overline{z})$			

# Abstracting Java executions

- Dynamic Semantics

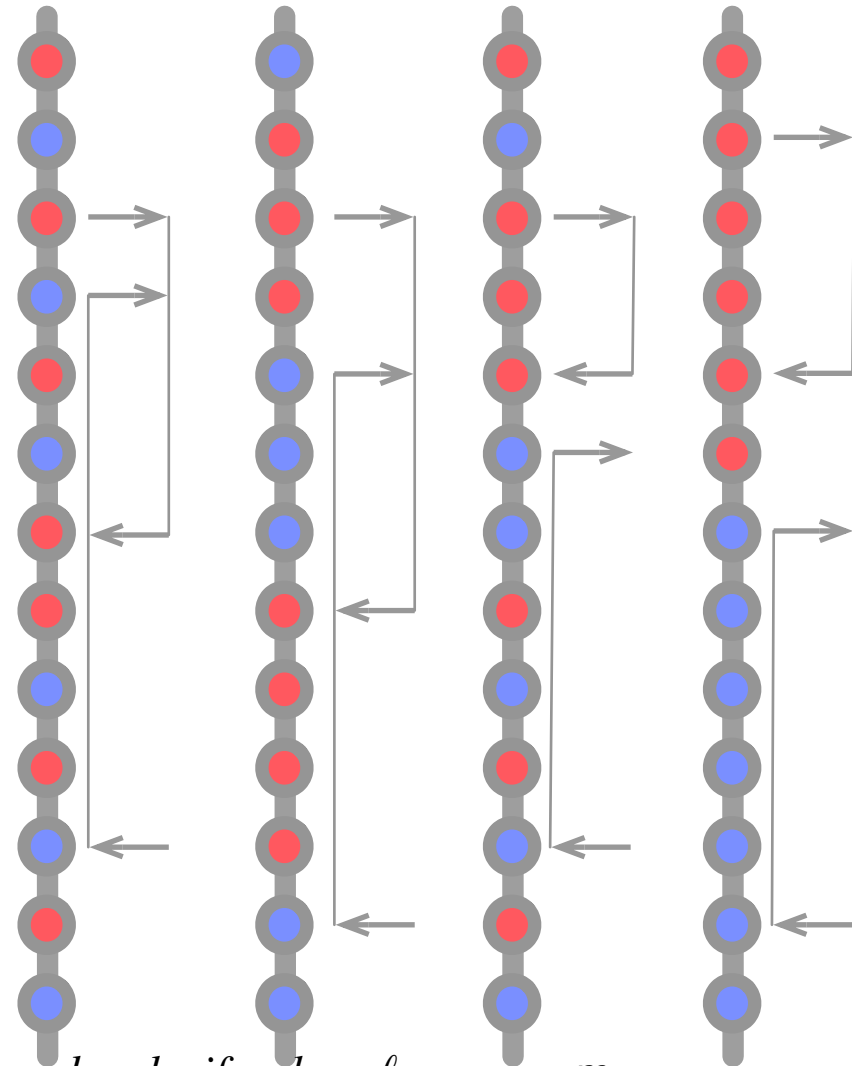
$$H; \overline{T} T \overline{T}' \xrightarrow{\ell}_{\rho} H'; \overline{T} \overline{T}' T'$$

- Scheduling is non-deterministic
- Traces are sequences of events:  
read x.f, write x.f, call x.m, return



## Modeling AJ executions

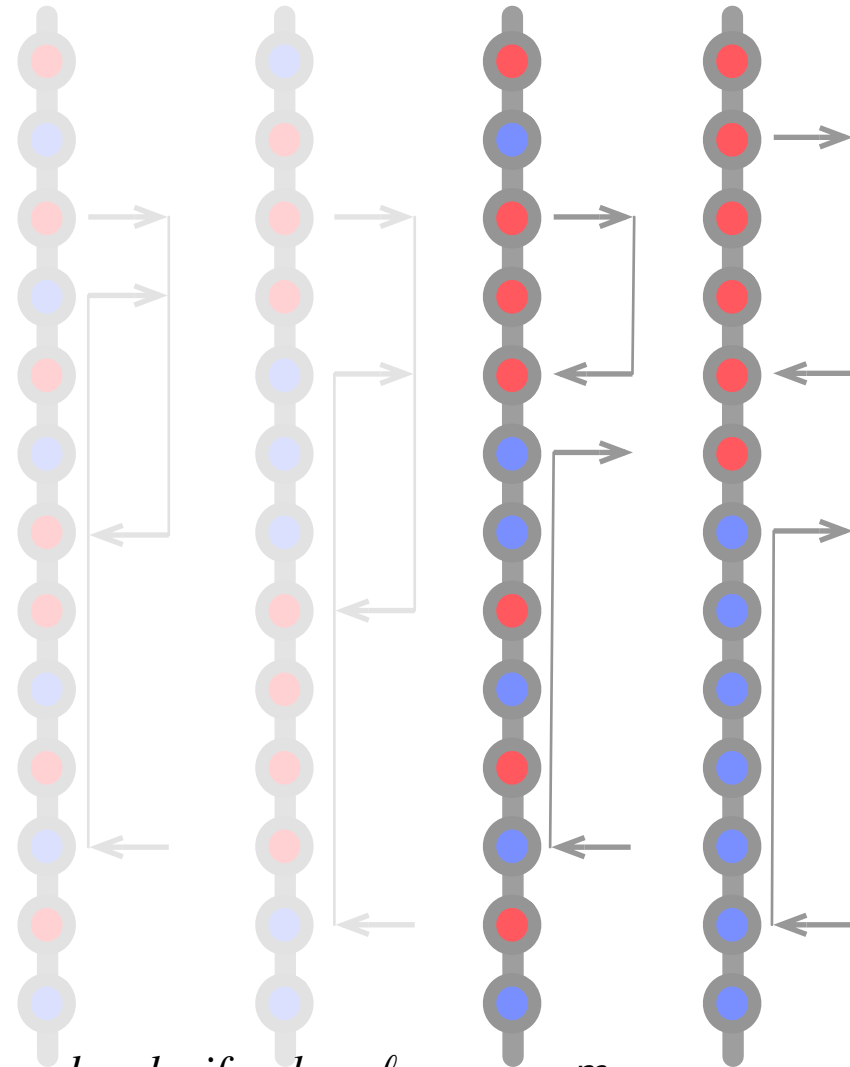
- Units of work on same atomicset are mutually exclusive
- Modeled by restricting valid traces



**Definition** An event  $e = (H, \bar{T}, \ell, \rho)$  is valid if and only if, when  $\ell \Rightarrow r.m$ ,  $H(r) = C|r'|(\bar{r})$  and  $C$  not internal then  $\nexists \rho' S \in \bar{T}.\rho' \neq \rho$  and  $\langle m F s \rangle \in S$  and  $H(F(\text{this})) = D|r'|(\bar{z})$ .

## Modeling AJ executions

- Units of work on same atomicset are mutually exclusive
- Modeled by restricting valid traces

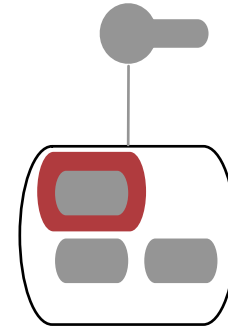


**Definition** An event  $e = (H, \bar{T}, \ell, \rho)$  is valid if and only if, when  $\ell \Rightarrow r.m$ ,  $H(r) = C|r'|(\bar{r})$  and  $C$  not internal then  $\nexists \rho' S \in \bar{T}.\rho' \neq \rho$  and  $\langle m F s \rangle \in S$  and  $H(F(\text{this})) = D|r'|(\bar{z})$ .

# What could go wrong?

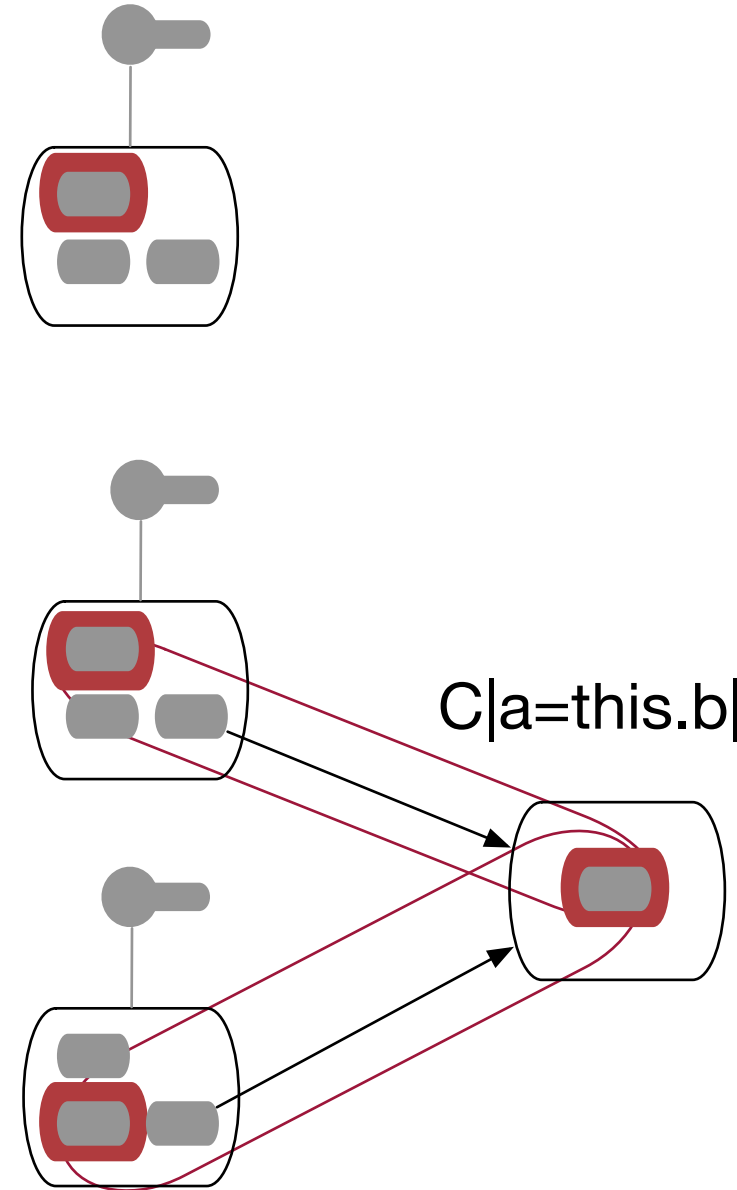
## What could go wrong?

- *A field is accessed outside of the boundaries of a unit of work*



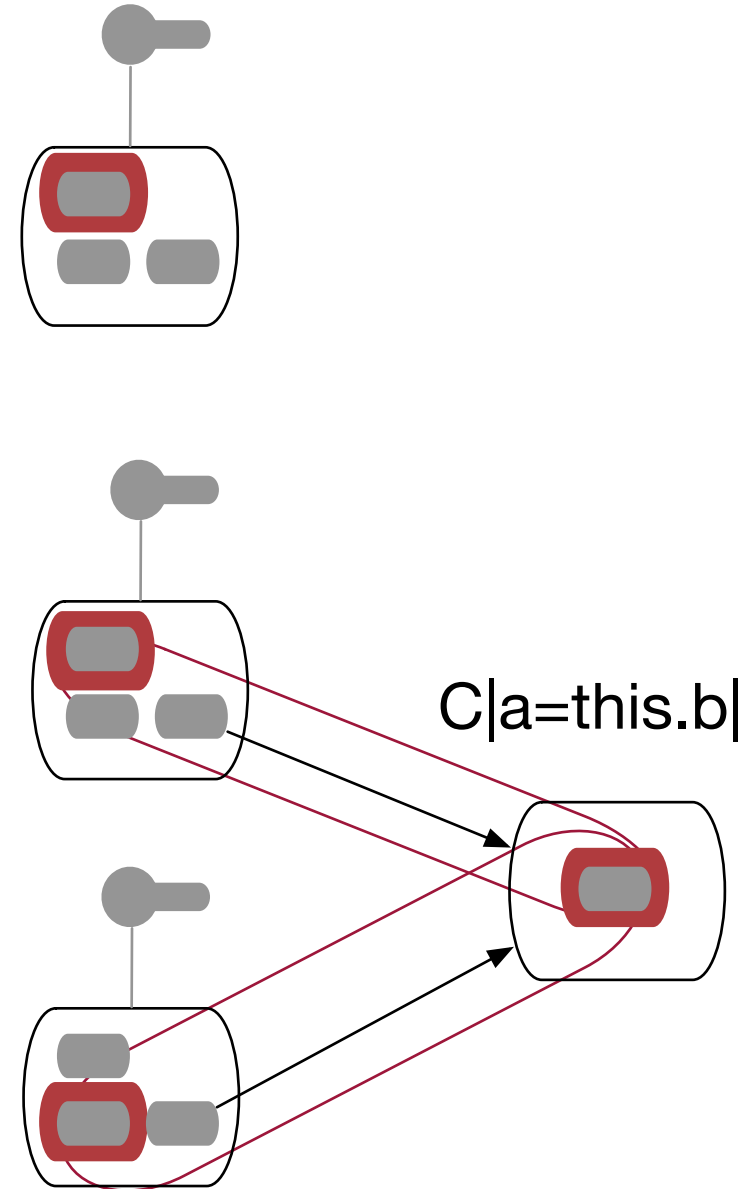
## What could go wrong?

- *A field is accessed outside of the boundaries of a unit of work*
- Could this happen?
  - Alias confusion... an object is referenced from two atomicsets



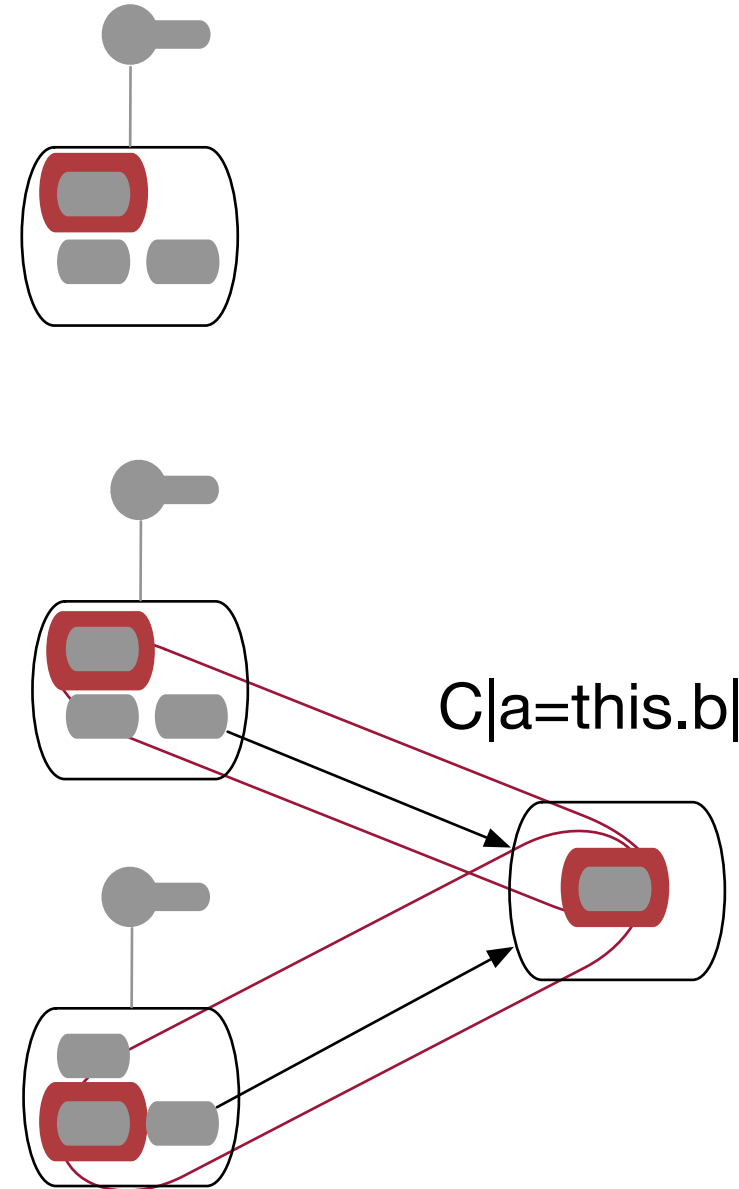
## What could go wrong?

- *A field is accessed outside of the boundaries of a unit of work*
- Could this happen?
  - Alias confusion... an object is referenced from two atomicsets
  - Ownership leak... an object is directly referenced from outside of its atomicset



## What could go wrong?

- *A field is accessed outside of the boundaries of a unit of work*
- Could this happen?
  - Alias confusion... an object is referenced from two atomicsets
  - Ownership leak... an object is directly referenced from outside of its atomicset
- A well-formed configuration is one, where this does not occur



# Technically,

$$\frac{\text{(WF-CONFIGURATION)} \quad H \text{ is WF in } H \quad \bar{T} \text{ is WF in } H \quad \vdash CT}{H; \bar{T} \text{ is WF}}$$

$$\frac{\text{(WF-EMPTY-HEAP)}}{[] \text{ is WF in } H}$$

$$\frac{\text{(WF-NPE-THREAD)}}{\rho \text{ NPE is WF in } H}$$

$$\frac{\begin{array}{l} \text{(WF-THREAD-BOT)} \\ \langle \text{run } F \text{ s} \rangle \text{ is WF in } H \\ \text{not internal}_H(F(\text{this})) \end{array}}{\rho \langle \text{run } F \text{ s} \rangle \text{ is WF in } H} \quad \frac{\begin{array}{l} \text{(WF-THREAD)} \\ \langle \mathbf{m} \ F \ \mathbf{s} \rangle \text{ is WF in } H \quad S \equiv S' \langle \mathbf{m}' \ F' \ \mathbf{x} = \mathbf{y.m}(\bar{\mathbf{z}}'); \mathbf{s}'' \rangle \\ \rho S \text{ is WF in } H \\ (\exists \langle \mathbf{m}'' \ F'' \ \mathbf{s}'' \rangle \in S \langle \mathbf{m} \ F \ \mathbf{s} \rangle, \text{ not internal}_H(F''(\text{this}))) \\ \text{and owner}_H(F''(\text{this})) = \text{owner}_H(F(\text{this})) \end{array}}{\rho S \langle \mathbf{m} \ F \ \mathbf{s} \rangle \text{ is WF in } H}$$

$$\frac{\begin{array}{l} \text{(WF-HEAP)} \\ (\mathbf{C} \text{ has } \mathbf{a} \text{ implies } \omega \neq \epsilon) \quad H' \text{ is WF in } H \\ \text{fields}(\mathbf{C}) = \overline{\alpha \tau \mathbf{f}} \quad \overline{r \mathbf{z}} <_{:r, H} \tau \end{array}}{H'[r \mapsto \mathbf{C}|\omega|(\overline{r \mathbf{z}})] \text{ is WF in } H}$$

$$\frac{\begin{array}{l} \text{(WF-FRAME)} \\ \text{locals}(\mathbf{m}, F) = E \quad E \vdash \mathbf{s} \\ \forall \mathbf{x} \in \text{dom}(F), F(\mathbf{x}) <_{:F(\text{this}), H} E(\mathbf{x}) \end{array}}{\langle \mathbf{m} \ F \ \mathbf{s} \rangle \text{ is WF in } H}$$

## Actually,

- Well-formedness ensures that
  - all threads are in a consistent state
  - heap is well-typed & no aliasing confusion & no ownership leaks
- We prove that AJ programs preserve WF properties

**Theorem 1. Preservation.** *If  $H; \overline{T} T \overline{T}'$  is WF and  $H; \overline{T} T \overline{T}' \xrightarrow{\ell}_{\rho} H'; \overline{T} \overline{T}' T'$ , then  $H; \overline{T} \overline{T}' T'$  is WF.*

**Theorem 2. Progress.** *If  $H; \overline{T} T \overline{T}'$  is WF and  $\text{active}(T)$ , then  $H; \overline{T} T \overline{T}' \xrightarrow{\ell}_{\rho} H'; \overline{T} \overline{T}' T'$ .*

# AJ Type System

- subtyping rules
- viewpoint adaption
- type rule for classes, methods, and statement

## Subtyping

- Simple transitive, reflexive closure of the extends relation
- Alias types are define a distinct relation
- The language requires explicit casts for type changes

$$\frac{}{C <: C} \quad \frac{C \text{ extends } D}{C <: D} \quad \frac{C <: C' \quad C' <: D}{C <: D}$$

$$\frac{C <: D}{C|a = \text{this}.b| <: D|a = \text{this}.b|}$$

## Typing a Class

- The internal annotation of parents must be preserved
- (single atomic set restriction is enforced)
- Must re-check inherited methods

$$\overline{fd} \text{ OK in } \mathbf{C} \quad \text{methods}(\mathbf{C}) = \overline{md'} \quad \overline{md'} \text{ OK in } \mathbf{C} \quad (\mathbf{D} \text{ has } \mathbf{a} \text{ implies } as = \epsilon) \\ (\iota = \text{internal} \text{ implies } \mathbf{C} \text{ has } \mathbf{a}) \quad (\mathbf{D} \text{ is internal} \text{ implies } \iota = \text{internal})$$


---


$$\iota \text{ class } \mathbf{C} \text{ extends } \mathbf{D} \{ as \overline{fd} \overline{md} \} \text{ OK}$$

## Typing a Class

- The internal annotation of parents must be preserved
- (single atomic set restriction is enforced)
- Must re-check inherited methods

$$\overline{fd} \text{ OK in } C \quad \text{methods}(C) = \overline{md'} \quad \overline{md'} \text{ OK in } C \quad (D \text{ has a implies } as = \epsilon)$$

$$(\iota = \text{internal implies } C \text{ has a}) \quad (D \text{ is internal implies } \iota = \text{internal})$$

$$\iota \text{ class } C \text{ extends } D \{ as \overline{fd} \overline{md} \} \text{ OK}$$

## Typing a Class

- The internal annotation of parents must be preserved
- (single atomic set restriction is enforced)
- Must re-check inherited methods

$\overline{fd}$  OK in C     $methods(C) = \overline{md'}$      $\overline{md'}$  OK in C    (D has a implies  $as = \epsilon$ )  
 (  $\iota = \text{internal}$  implies C has a )    (D is internal implies  $\iota = \text{internal}$ )

$\iota$  class C extends D {  $as$   $\overline{fd}$   $\overline{md}$  } OK

## Typing statements

- Casting off an alias type is allowed for non-internal classes
- Casting between alias types must preserve aliases

$$\begin{array}{c}
 E(x) = D|a=this.b| \quad E(y) = C|a=this.b| \\
 C \text{ has } a \quad E(\text{this}) \text{ has } b \quad D <: C \\
 \hline
 E \vdash y = (C|a=this.b|)x
 \end{array}$$

$$\begin{array}{c}
 E(x) = C|a=this.b| \quad C \text{ not internal} \\
 E(y) = C \\
 \hline
 E \vdash y = (C)x
 \end{array}$$

## Typing statements

- Method calls require viewpoint adaption of arguments & return type

$$\text{adapt}(\mathbf{C}, \tau) = \mathbf{C}$$

$$\text{adapt}(\mathbf{C}|\mathbf{a}=\text{this.b}|, \mathbf{D}|\mathbf{b}=\text{this.c}|) = \mathbf{C}|\mathbf{a}=\text{this.c}|$$

$$E(\mathbf{y}) = \tau_y \quad \text{typeof}(\tau_y.\mathbf{m}) = \bar{\tau} \rightarrow \tau \quad E(\bar{\mathbf{z}}) = \bar{\tau}_z$$

$$\bar{\tau}_z = \text{adapt}(\bar{\tau}, \tau_y) \quad \tau' = \text{adapt}(\tau, \tau_y) \quad E(\mathbf{x}) = \tau'$$

---


$$E \vdash \mathbf{x} = \mathbf{y}.\mathbf{m}(\bar{\mathbf{z}})$$

## View point adaption

```

class C { atomicset c;
  B|b=this.c| x;
  x=new B|b=this.c| ();
  x.gee(x)
}

class B { atomicset b;
  void gee(B|b=this.b| x){

```

B|b=this.c|



B|b=this.b|

## Related Work

- *See the paper for a complete list*
- **Atomic set**
  - Vaziri e/a [POPL06] *original proposal*
  - *detection of atomicset serializability violations*
    - Hammer e/a [ICSE08] *dynamic*
    - Kidd e/a [VMCAI09] *static*
- **Data groups**
  - Leino e/a [OOPSLA98] *abstract rep. of groups of fields for modular reasoning*
- **Atomicity and Race-free Types**
  - Flanagan e/a [ESOP99, PLDI00, PLDI03, TOPLAS08,  $\infty$ ]
- **Ownership types**
  - Zhao e/a [JFP06] *lightweight ownership disciplines*
- **Lock inference**
  - Cherem e/a. [PLDI08], McCloskey e/a [POPL06]

## Conclusions

- A type system for data-centric synchronization which
  - guarantees atomic-set serializability
  - enables separate compilation
  - handles unbounded sets
- Annotation overhead competitive with race-freedom and atomicity type systems
- Future
  - improve performance
  - local annotation inference
  - static deadlock prevention