

Valor: Efficient, Software-Only Region Conflict Exceptions

Swarnendu Biswas, Minjia Zhang, and

Michael D. Bond

Ohio State University

Brandon Lucia

Carnegie Mellon University

OOPSLA 2015



A Simple C++ Program

```
X* x = NULL;  
bool done= false;
```

Thread T1

```
x = new X();  
done = true;
```

Thread T2

```
if (done) {  
    x->func();  
}
```

A Simple C++ Program



ad T2



Data Races


```
X* x = NULL;  
bool done= false;
```

Thread T1

```
x = new X();  
done = true;
```

Thread T2

```
if (done) {  
    x->func();  
}
```



Data Races are Evil

No semantic guarantees

Lack of semantic guarantees make software unsafe

Complicates language specifications

Challenging to reason about correctness for racy executions

Indicates other concurrency errors

Leads to atomicity, order or sequential consistency violations

Catch-Fire Semantics

C++ treats data races as errors

```
X* x = NULL;  
bool done= false;
```

Thread T1

```
x = new X();  
done = true;
```

Thread T2

```
if (done) {  
    x->func();  
}
```

Catch-Fire Semantics

C++ treats data races as errors

Thread

```
x = new X();  
done = true;
```



2

Catch-Fire Semantics

C++ treats data races as errors

Thread

```
x = new X();  
done = true;
```



A Java Example

Thread T1

```
X = new Object();  
done = true;
```

Thread T2

```
while (!done) {}  
X.compute();
```

A Java Example

Java tries to assign semantics,
which are unsatisfactory


Thread T1

```
X = new Object();  
done = true;
```

Thread T2

```
while (!done) {}  
X.compute();
```

C++ and Java Memory Models



C++ and Java Memory Models



C++ and Java Memory Models

Data-race-free
execution

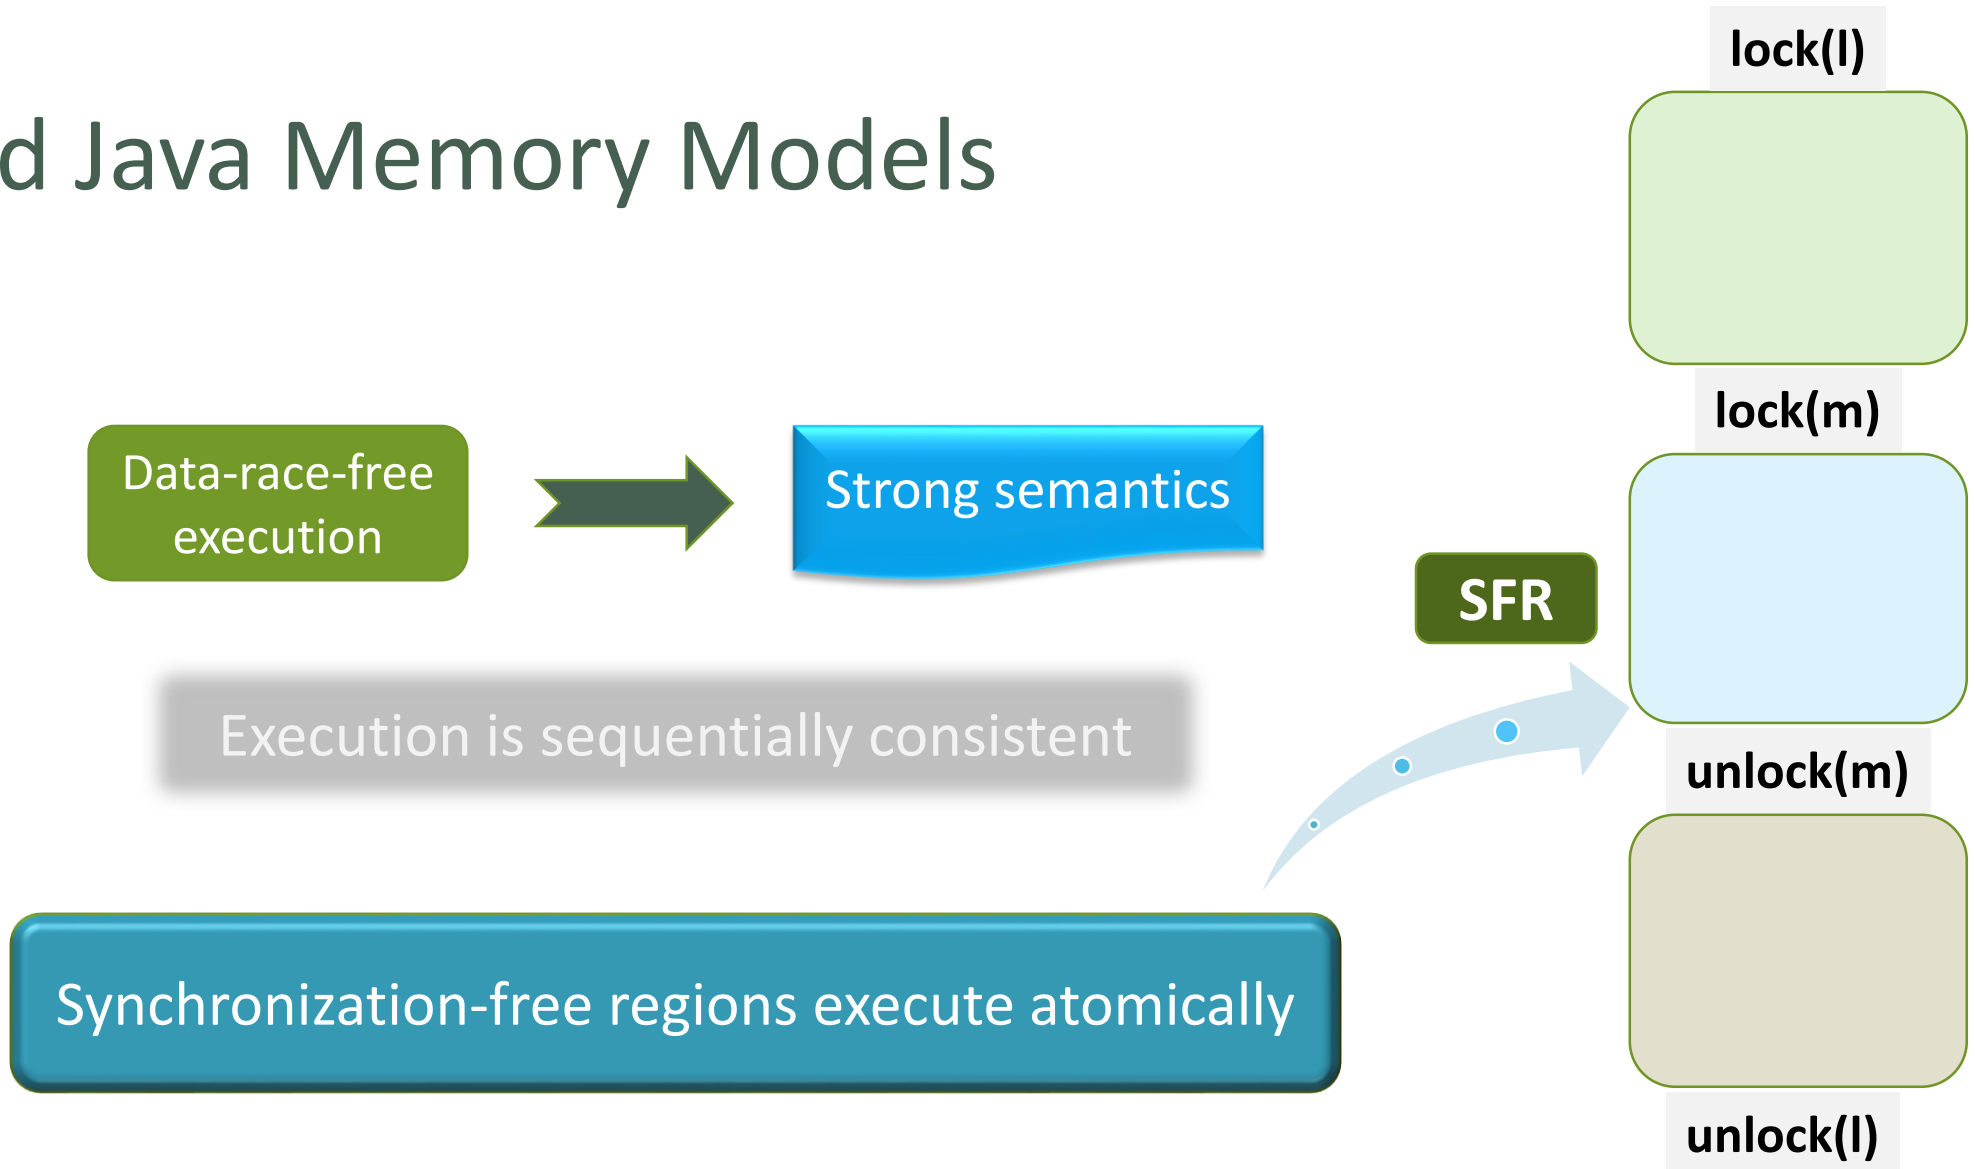


Strong semantics

Execution is sequentially consistent

Synchronization-free regions execute atomically

C++ and Java Memory Models



C++ and Java Memory Models



But what
about data
races?

C++ and Java Memory Models

But what
about data
races?

Racy execution



???

Need for Stronger Memory Models

Adve and Boehm, CACM 2010

“The inability to define reasonable semantics for programs with data races is not just a theoretical shortcoming, but a fundamental hole in the foundation of our languages and systems.”

Need for Stronger Memory Models

Adve and Boehm, CACM 2010

“The inability to define reasonable semantics for programs with data races is not just a theoretical shortcoming, but a fundamental hole in the foundation of our languages and systems.”

“We call upon software and hardware communities to develop languages and systems that enforce data-race-freedom, ...”

Outline

- Programming language memory models and data races
- **Data race and region conflict exceptions model**
- **Valor: Our contribution**
- **Evaluation**

Data Race

Thread T1

```
X = new Object();  
done = true;
```

Thread T2

```
while (!done) {}  
X.compute();
```

Data Race Exceptions

Thread T1

```
X = new Object();  
done = true;
```

EXCEPTION


Thread T2

```
while (!done) {}  
x.compute();
```


REGION CONFLICT EXCEPTIONS MODEL

Region Conflict

Thread T1



Thread T2




Region Conflict

Thread T1



Thread T2



Region Conflict

Thread T1



Thread T2




Region Conflict

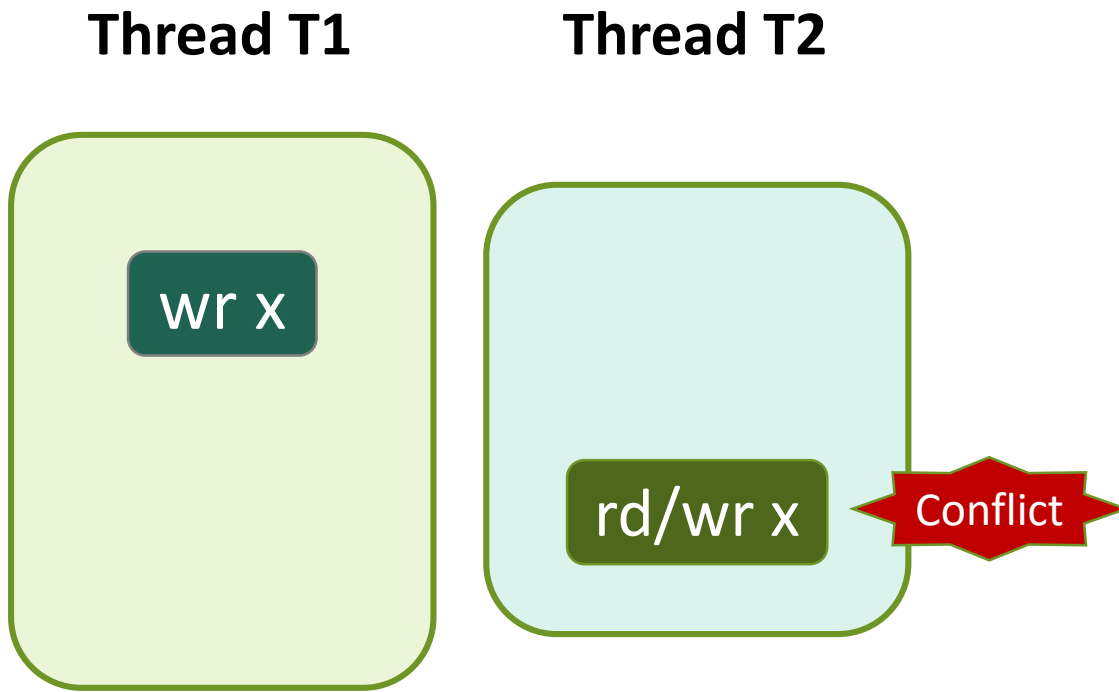
Thread T1



Thread T2

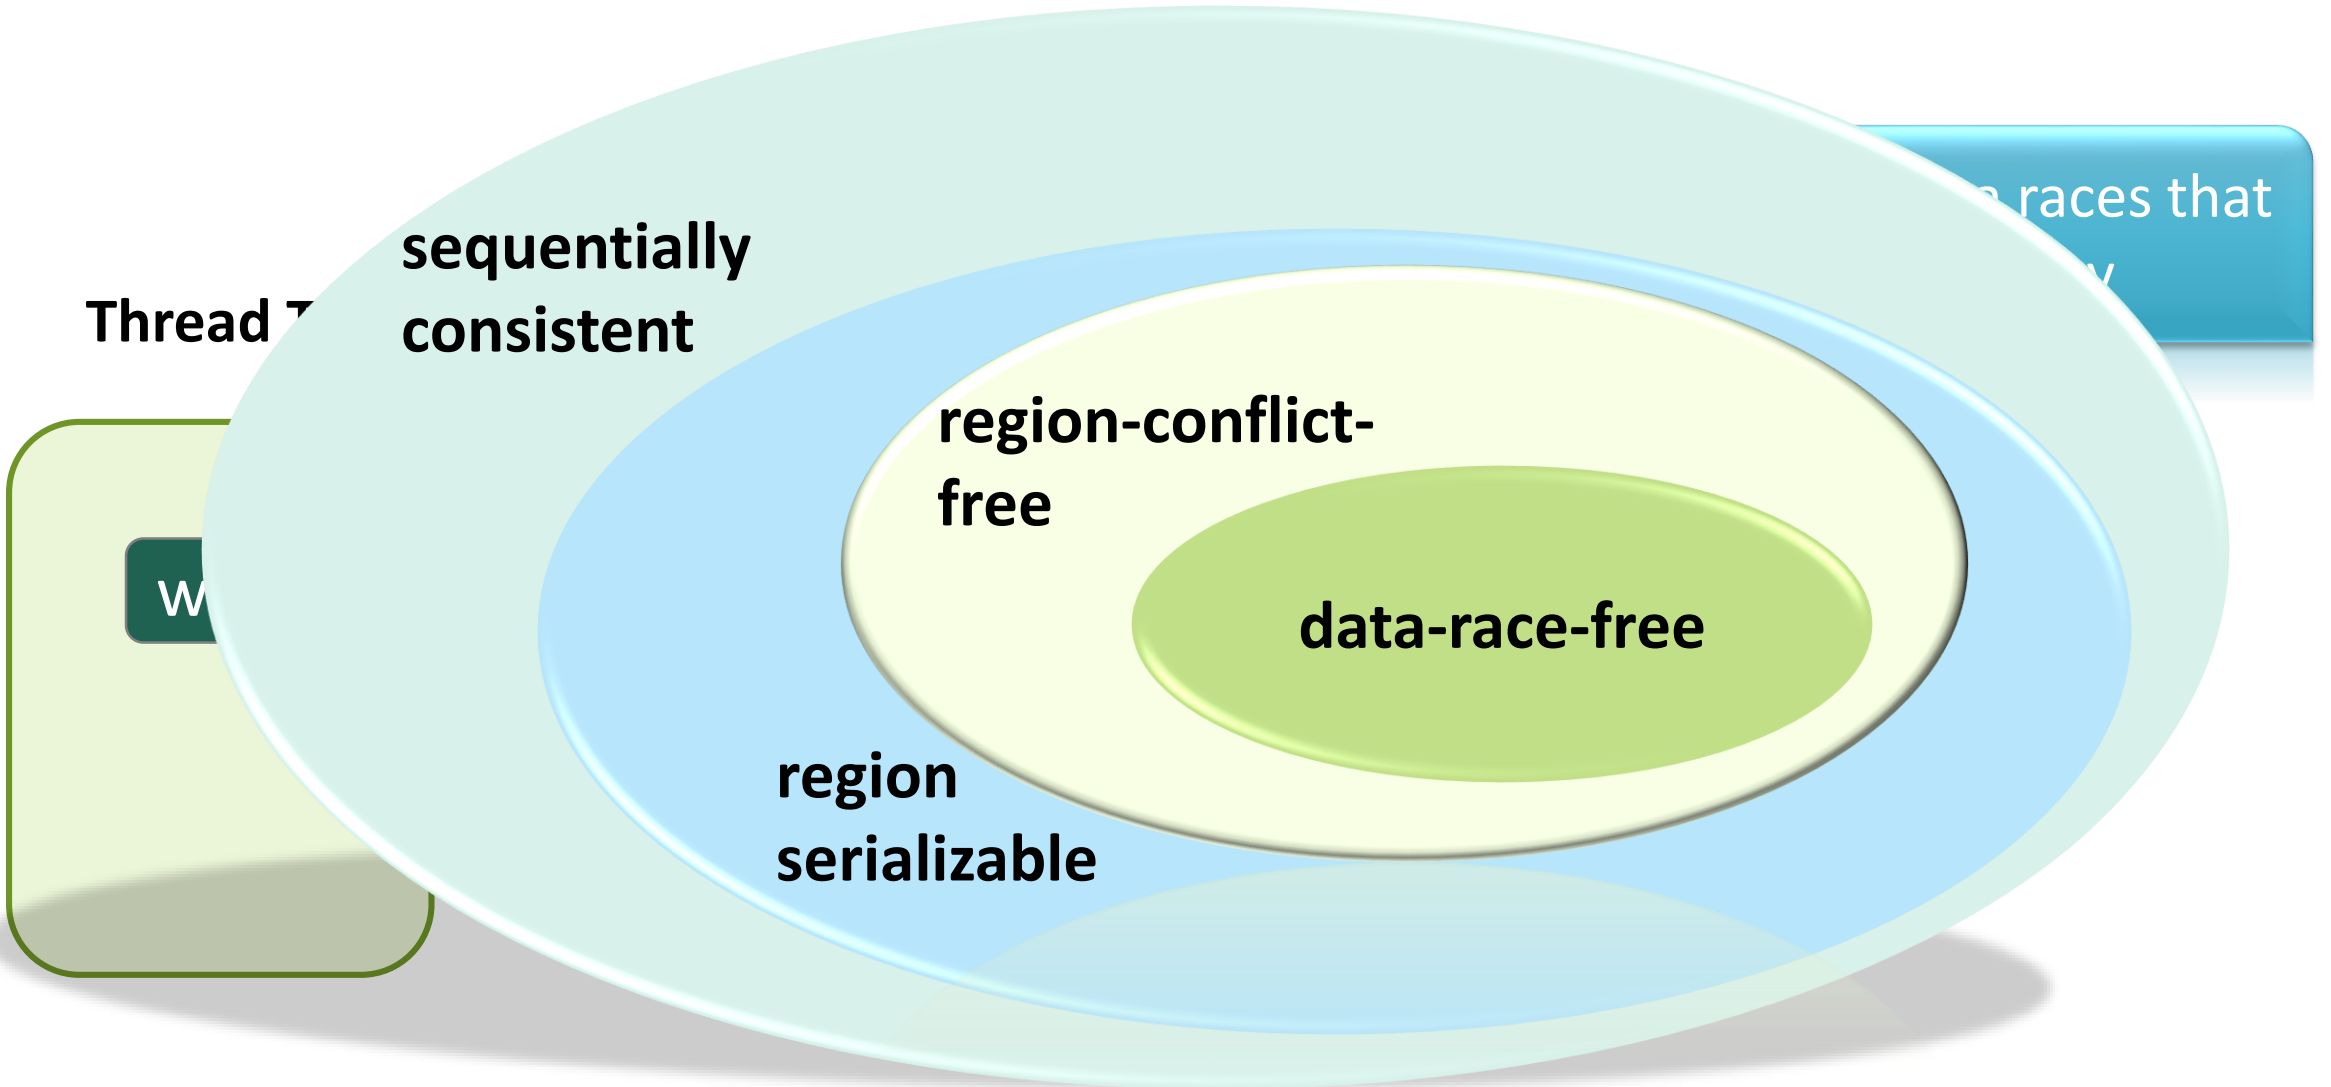


Region Conflict



Reports a subset of true data races that violate region serializability

Execution Models



Region Conflict Exception Model

GOAL

Develop a practical region conflict
detection technique

Region Conflict Detection

Hardware customizations required for good performance

- ❖ Limited by resources and applicability
 - Needs extensive modifications and is unscalable¹
 - Detects serializability violations of bounded regions²

-
1. Lucia et al. Conflict Exceptions: Simplifying Concurrent Language Semantics With Precise Hardware Exceptions for Data-Races. ISCA 2010.
 2. Marino et al. DRFx: A Simple and Efficient Memory Model for Concurrent Programming Languages. PLDI 2010.

Outline

- Programming language memory models and data races
- Data race and region conflict exceptions model
- **Valor: Our contribution**
- **Evaluation**

Valor: Efficient, Software-Only Region Conflict Detector

Elides tracking last readers, only tracks last writer

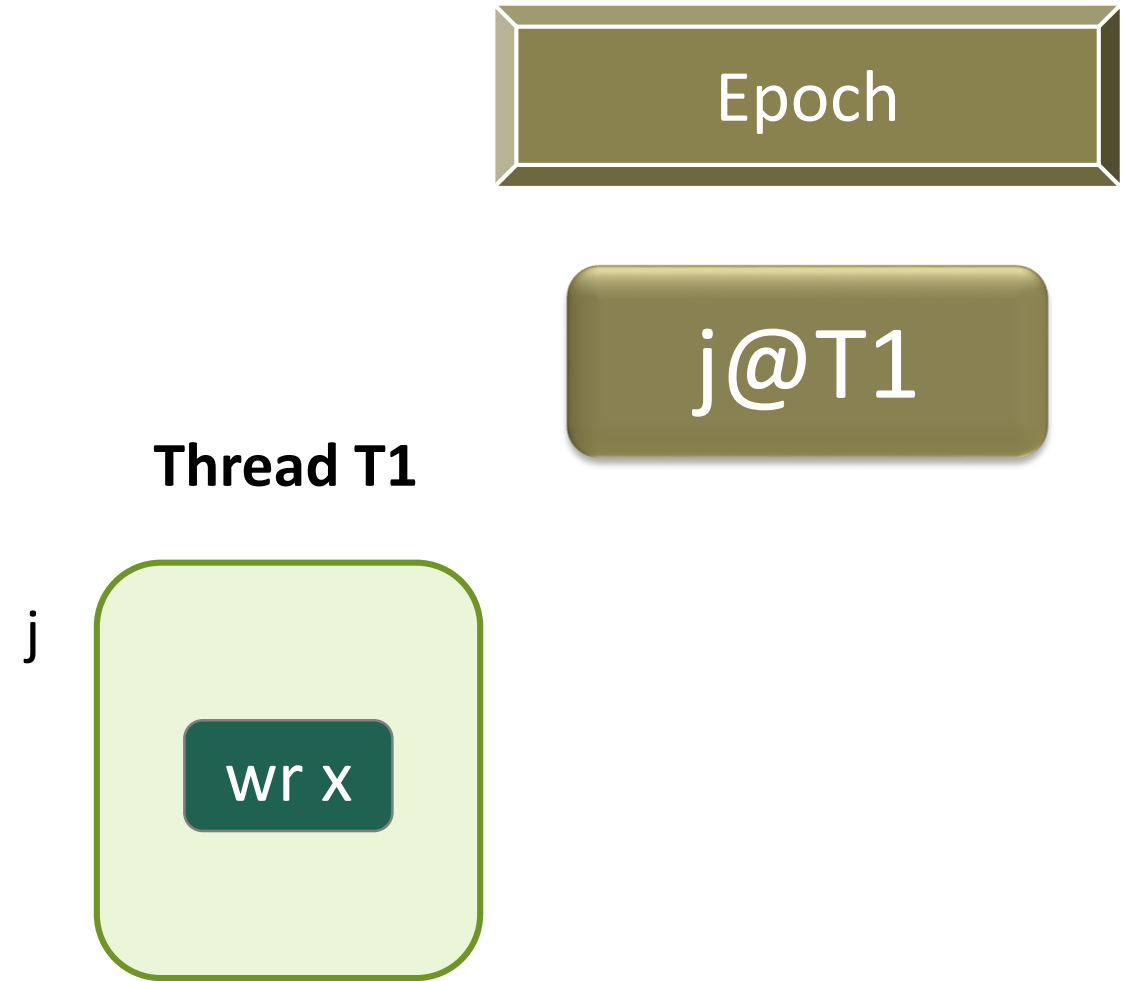
Detect read-write conflicts lazily

Valor: Efficient, Software-Only Region Conflict Detector


- Tracking last writers
- Detecting read-write conflicts lazily
- Impact of lazy conflict detection

Tracking Last Writer


Per-Variable Metadata




Tracking Last Writer




Thread T1



Thread T2





Tracking Last Writer




Thread T1

Thread T2

j



Tracking Last Writer





Track last writer

Thread T1

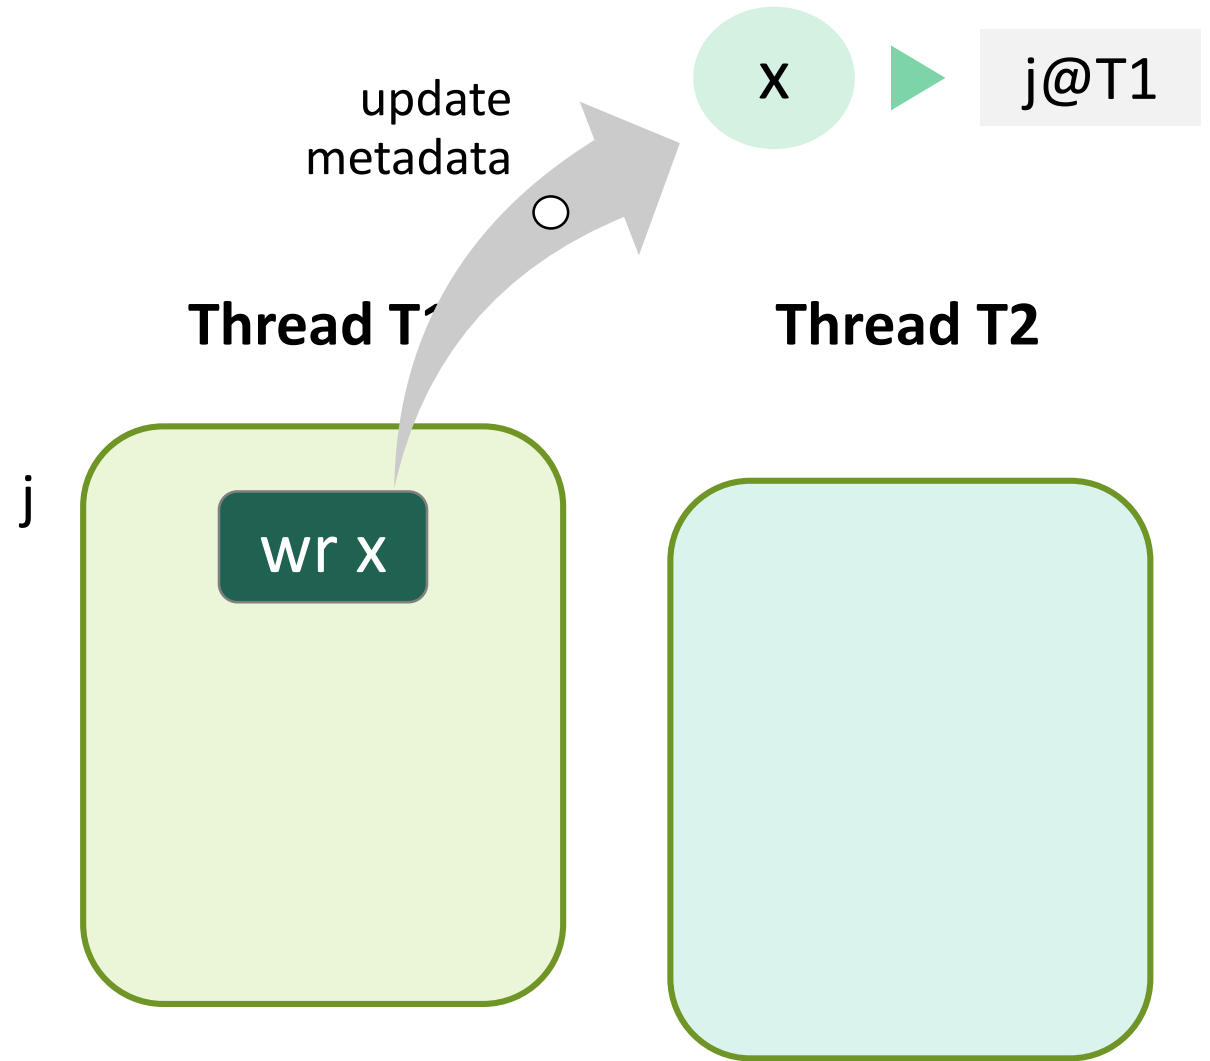
Thread T2

j




Tracking Last Writer

Track last writer



Tracking Last Writer



Track last writer


Thread T1

Thread T2


j

wr x

rd/wr x



Tracking Last Writer

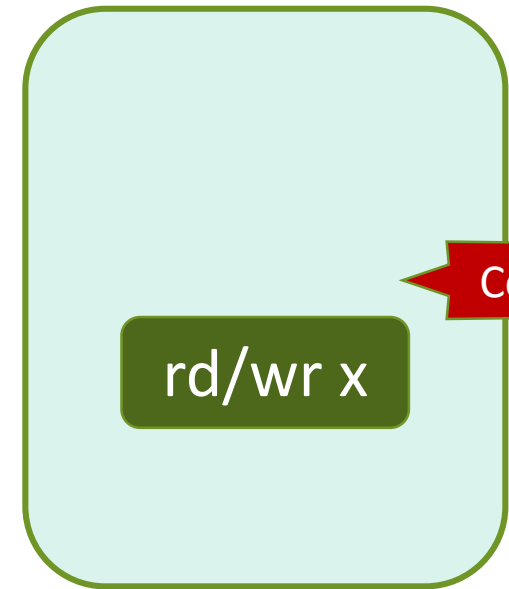



Track last writer


Thread T1

Thread T2

j



Tracking Last Writer



Track last writer


- ❖ Allows precisely detecting write-write and write-read conflicts

Thread T1

j



Thread T2



Conflict

Valor: Efficient, Software-Only Region Conflict Detector


- Tracking last writers
- Detecting read-write conflicts lazily
- Impact of lazy conflict detection

Detecting Read-Write Conflicts


Thread T1



Thread T2




Detecting Read-Write Conflicts



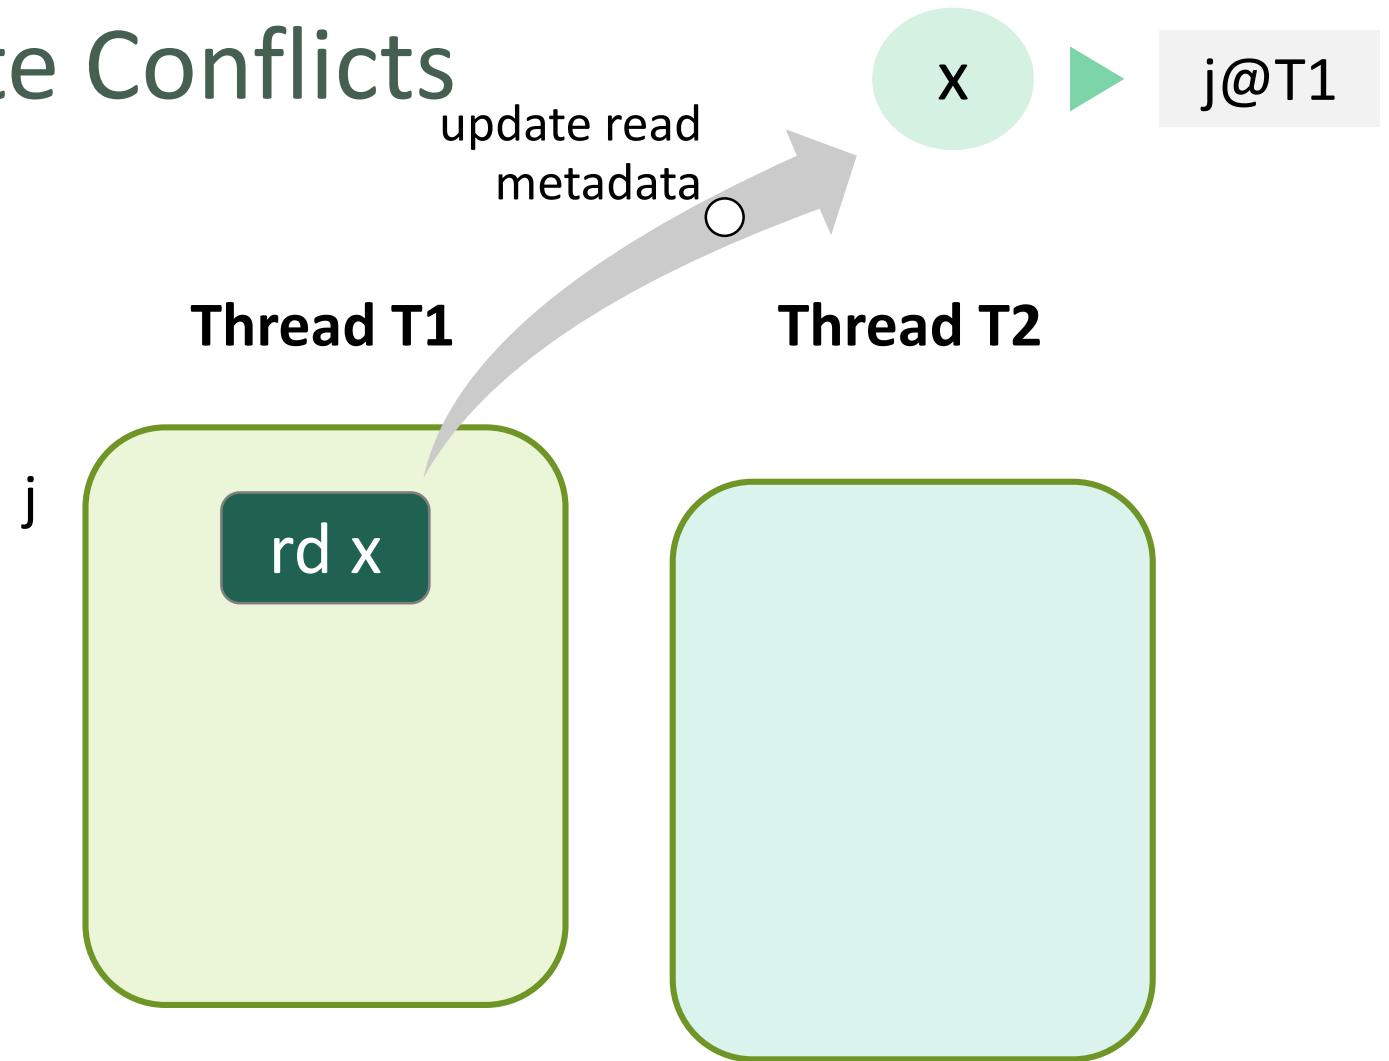
Thread T1

Thread T2


j



Detecting Read-Write Conflicts



Detecting Read-Write Conflicts




Thread T1

Thread T2

j



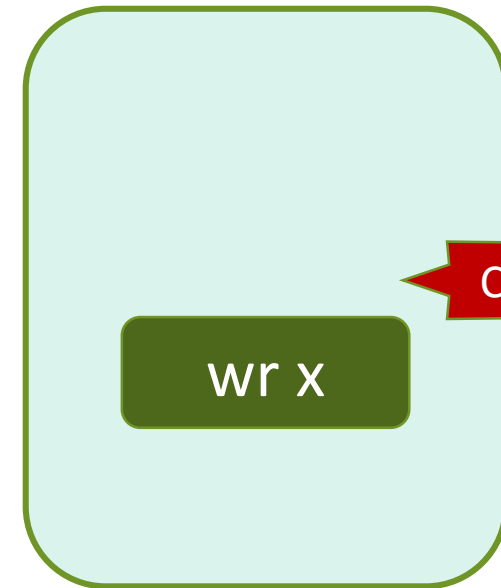
Detecting Read-Write Conflicts




Thread T1

Thread T2

j




Detecting Read-Write Conflicts



Thread T1

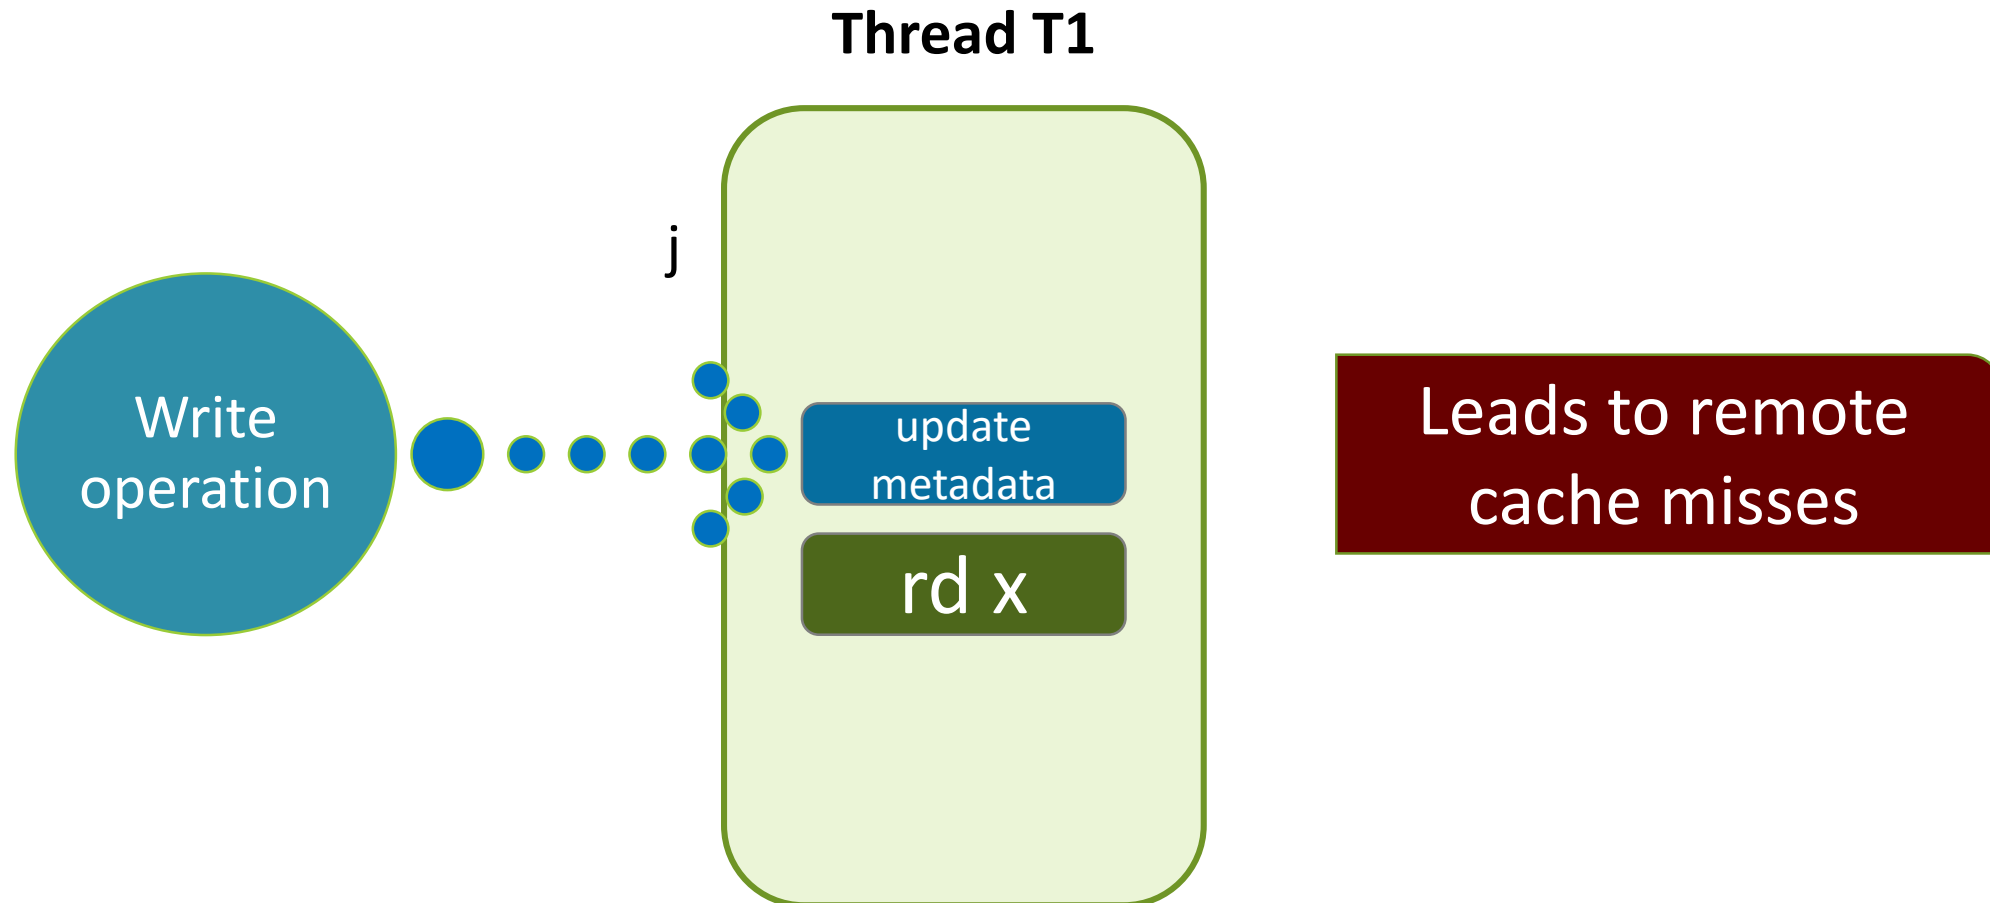
Thread T2

This simple mechanism used in prior work has problems



Conflict

Remote Cache Misses Due to Tracking of Metadata



Metadata Updates

Thread T1


j



Metadata Updates

Thread T1


j



Synchronization on Metadata Updates

Thread T1


j



Synchronization on Metadata Updates

Thread T1

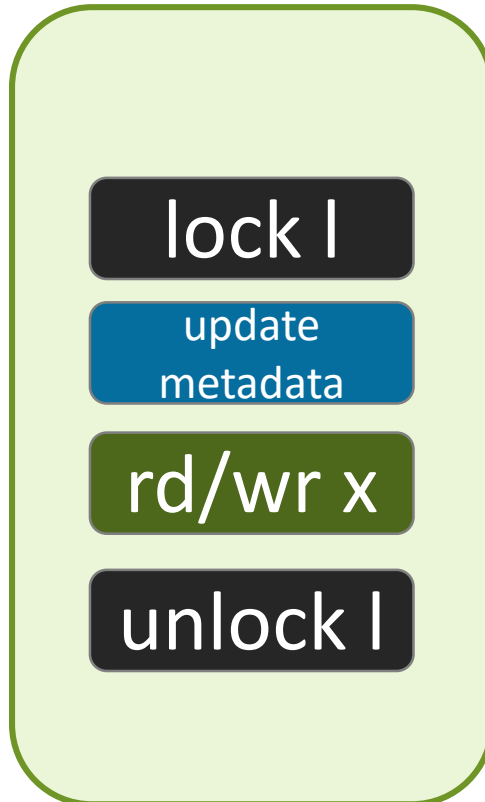
j



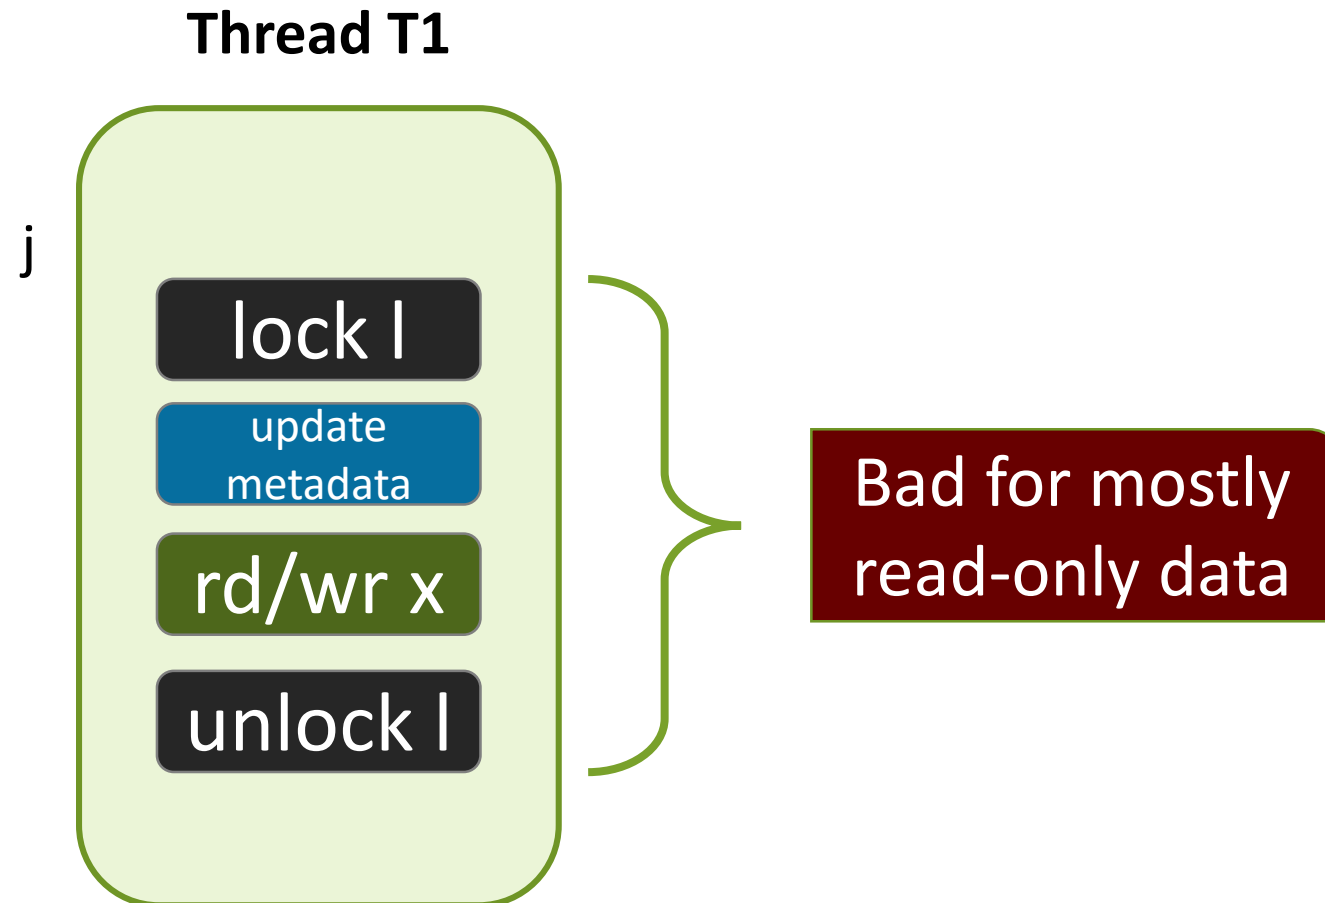
Synchronization on Metadata Updates

Thread T1

j



Synchronization on Metadata Updates




Elide Tracking Last Readers

Elide Tracking Last Readers

Thread T1



Thread T2



Elide Tracking Last Readers

Detect read-write conflicts lazily

Thread T1

rd x

Thread T2

wr x

Elide Tracking Last Readers

Detect read-write conflicts lazily

- ❖ Log read accesses in thread-local buffers
- ❖ Validate reads at region boundaries

Thread T1

rd x

Thread T2

wr x


Per-Variable Metadata



Version, Epoch

$\langle v, j@T1 \rangle$

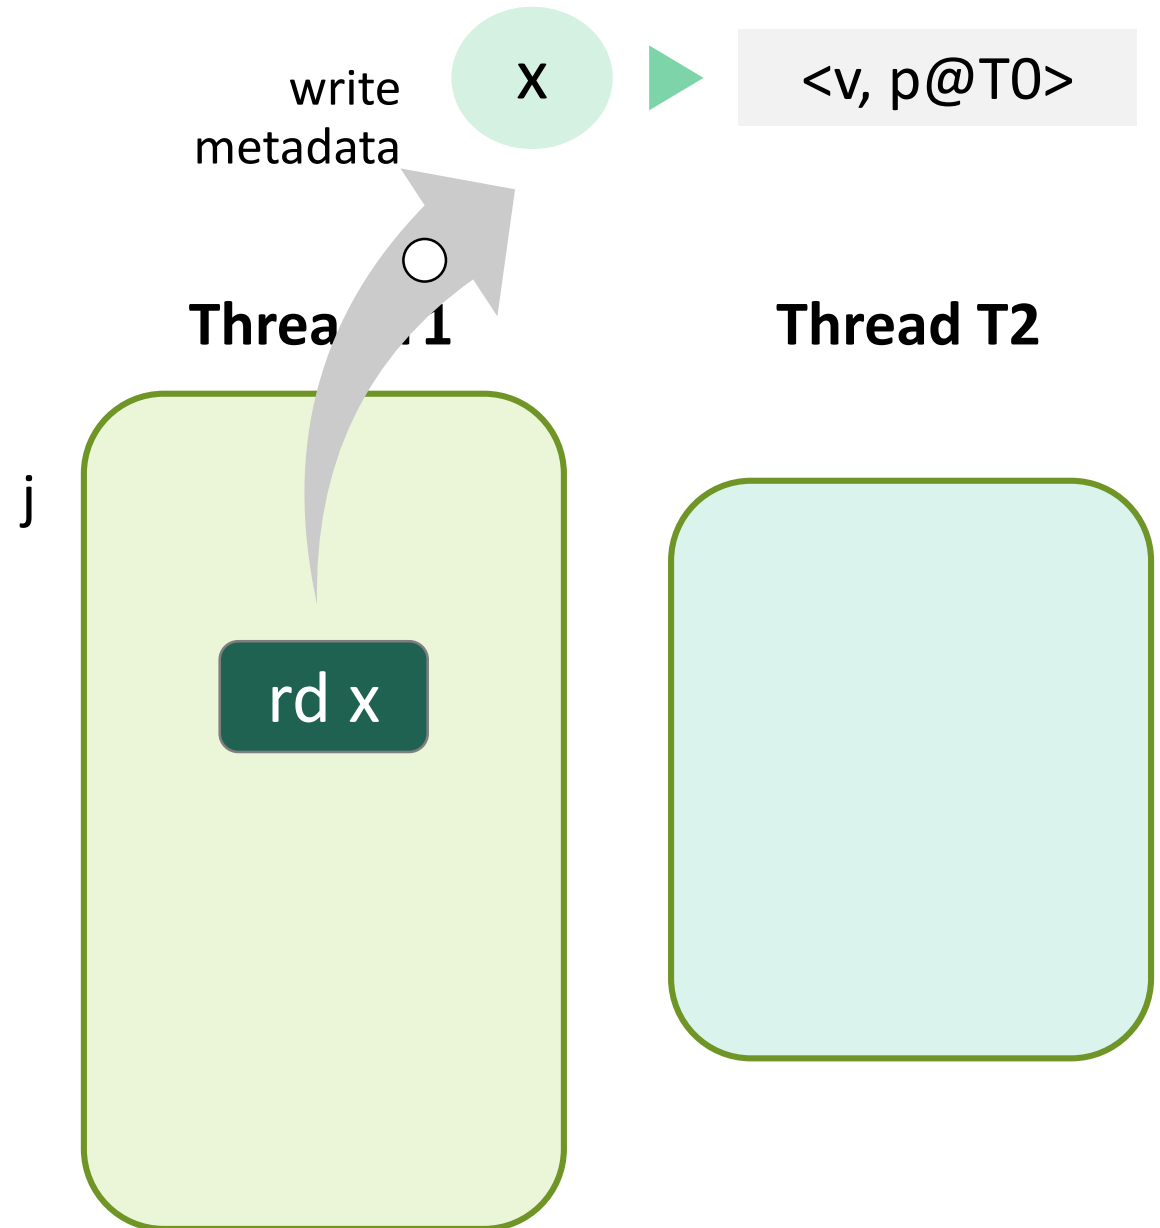
Thread T1



Elide Tracking Last Readers

Detect read-write conflicts lazily

- ❖ Log read accesses in thread-local buffers
- ❖ Validate reads at region boundaries




Elide Tracking Last Readers



Detect read-write conflicts lazily


- ❖ Log read accesses in thread-local buffers
- ❖ Validate reads at region boundaries




Elide Tracking Last Readers




Detect read-write conflicts lazily




Elide Tracking Last Readers



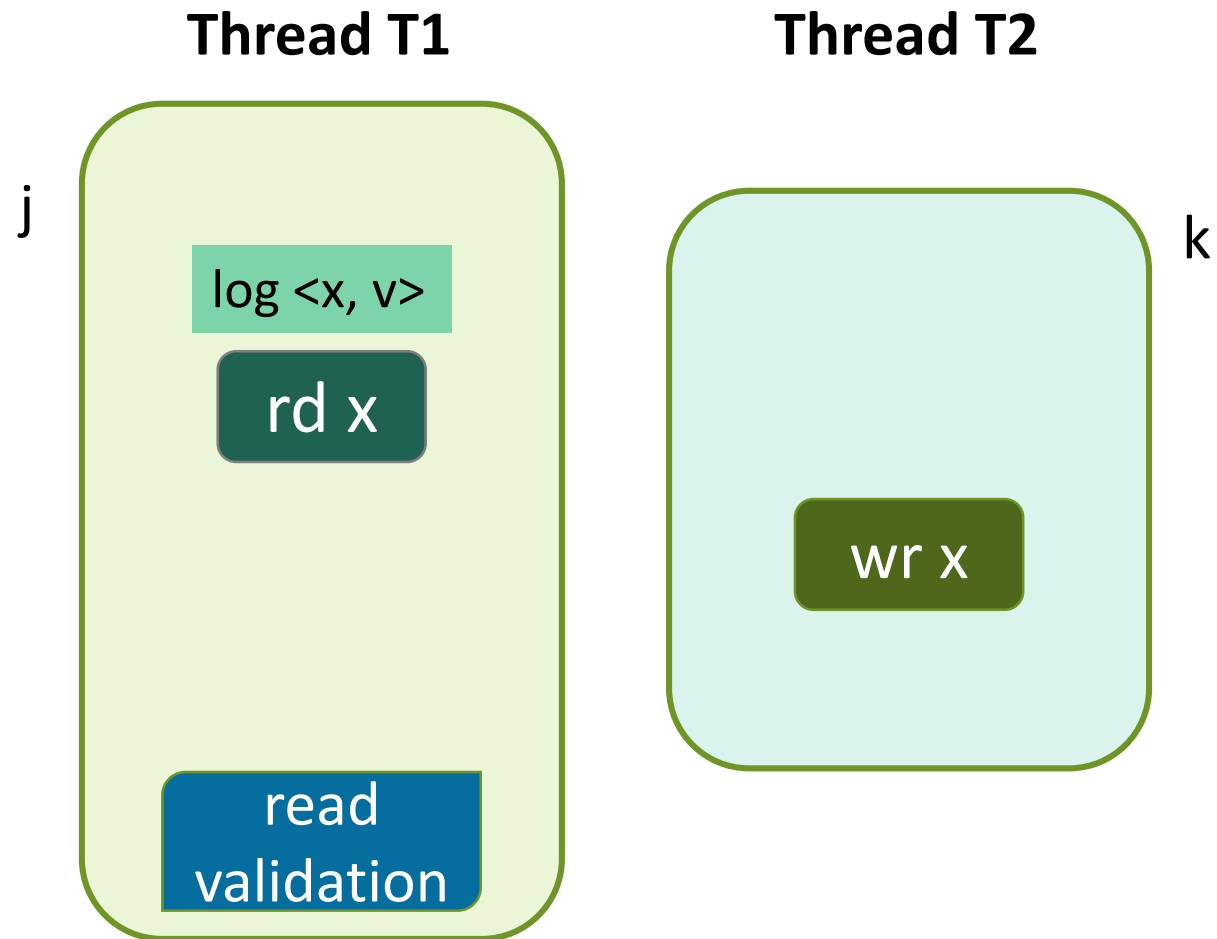
Detect read-write conflicts lazily




Elide Tracking Last Readers



Detect read-write conflicts lazily




Elide Tracking Last Readers




Detect read-write conflicts lazily

T1 is not the last writer

Version read is outdated



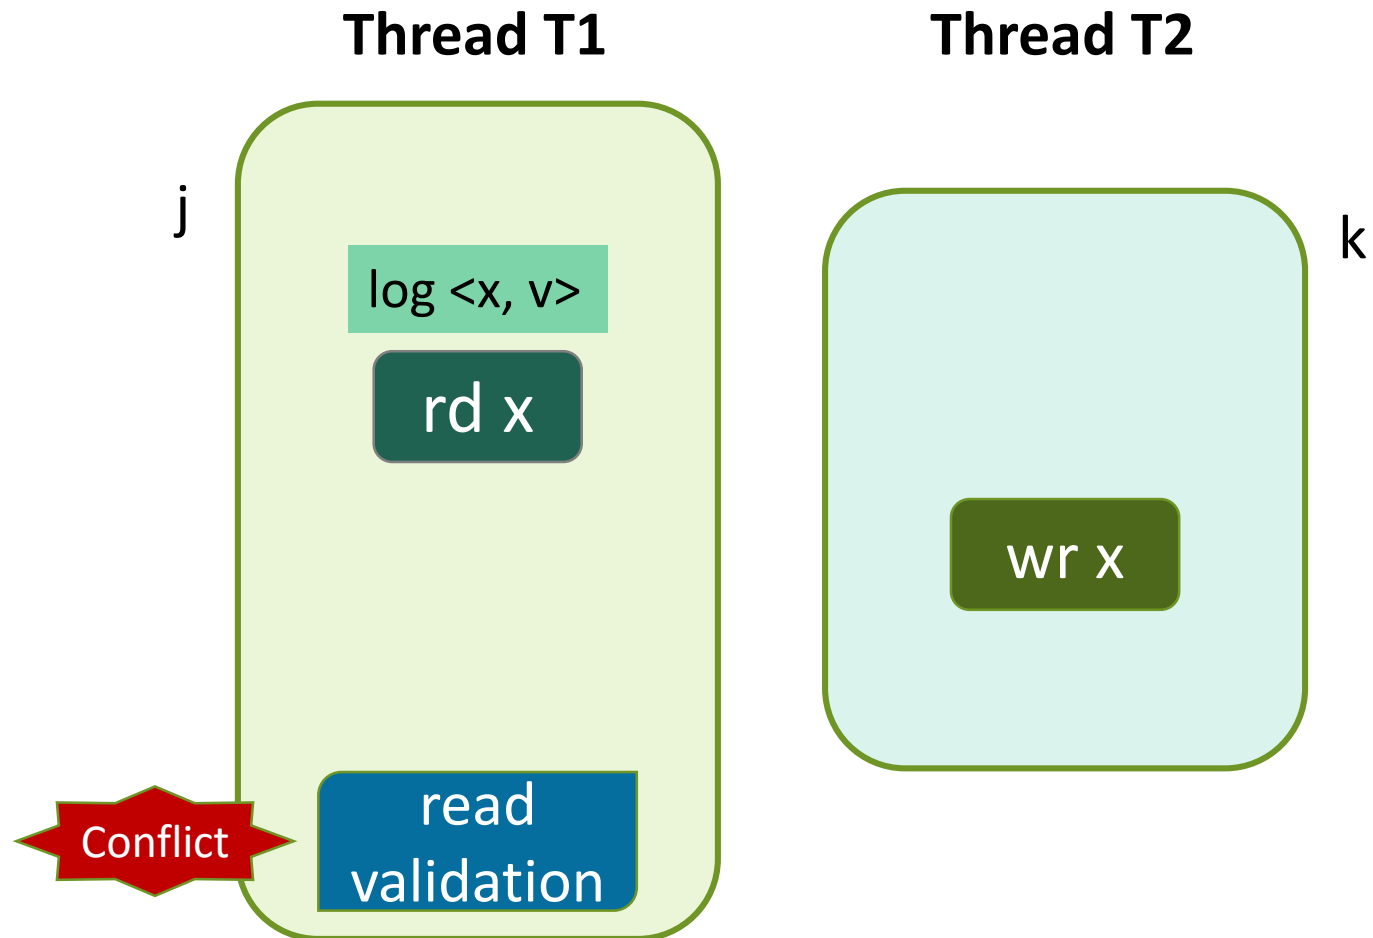
Elide Tracking Last Readers



Detect read-write conflicts lazily

T1 is not the last writer


Version read is outdated



Elide Tracking Last Readers

Avoids

- Remote cache misses
- Synchronization overhead



Valor: Efficient, Software-Only Region Conflict Detector

- Tracking last writers
- Detecting read-write conflicts lazily
- **Impact of lazy conflict detection**

Precise Conflict Detection

Thread T1

Thread T2



Precise vs Lazy Conflict Detection

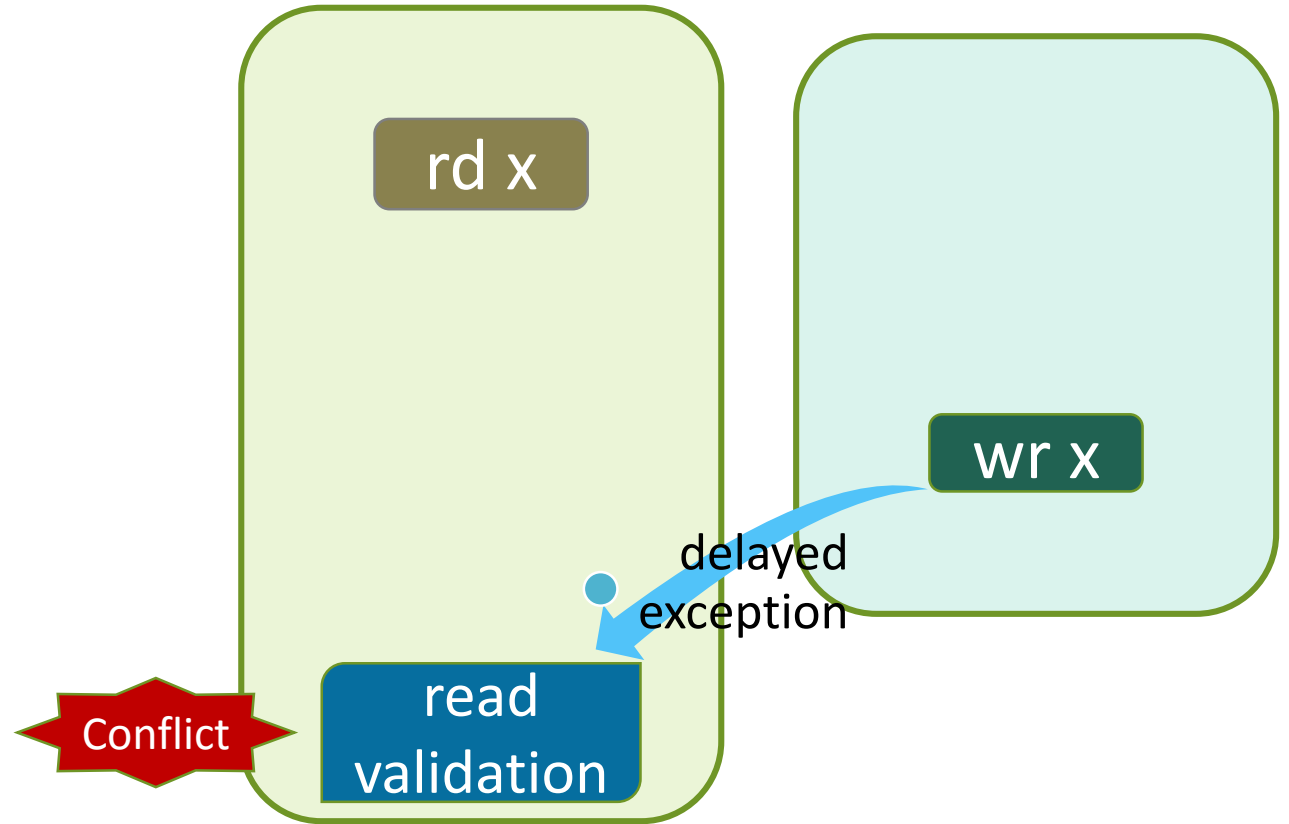
Thread T1

Thread T2



Thread T1

Thread T2



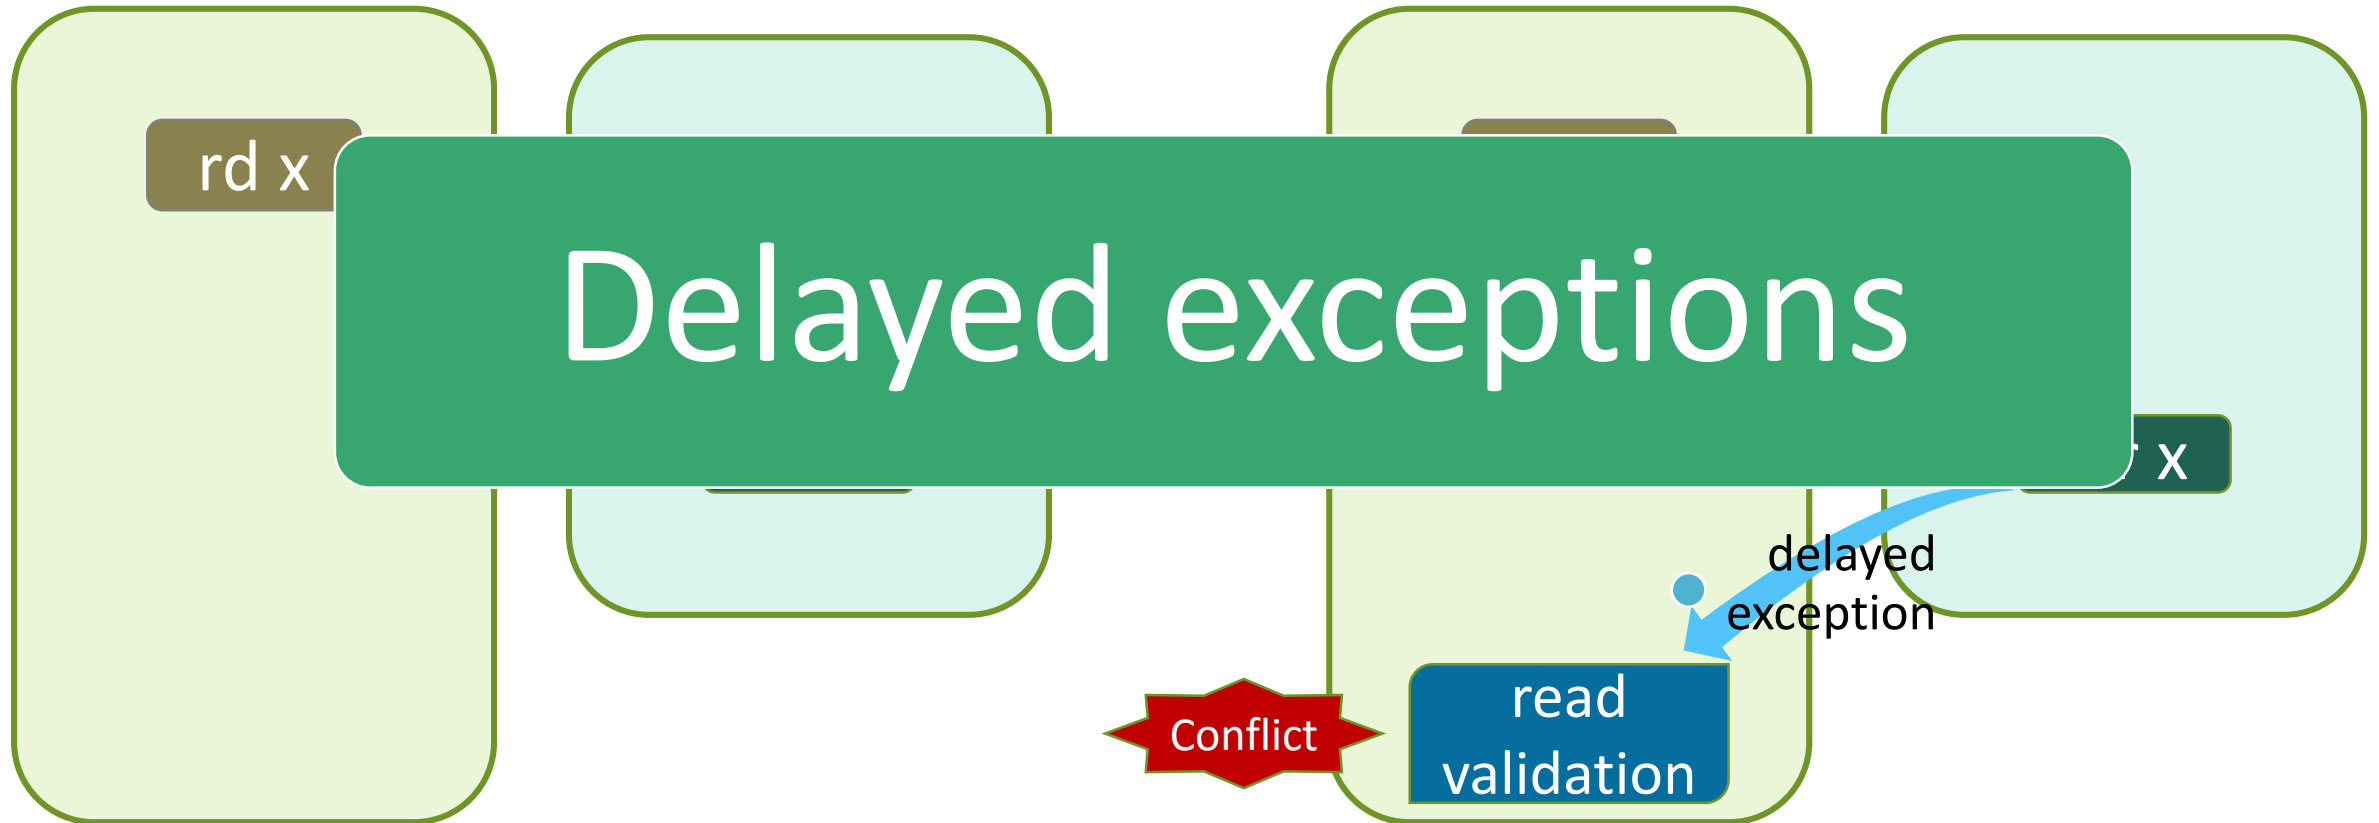
Precise vs Lazy Conflict Detection

Thread T1

Thread T2

Thread T1

Thread T2



Delayed Exceptions

Thread T1

Thread T2

Thread T1

Thread T2

Delayed exceptions

rd x

Do not compromise semantic guarantees

wr x

wr x

Effects should not be externally visible

Conflict

validation

Delayed Exceptions

Thread T1

Thread T2

Thread T1

Thread T2

Debugging might be slightly harder

rd x

Exception will be thrown at the next region boundary from the reader thread

wr x

But does not compromise on soundness and precision

wr x

Conflict

validation

Outline

- Programming language memory models and data races
- Data race and region conflict exceptions model
- Valor: Our contribution
- **Evaluation**

IMPLEMENTATION

Implementation

Developed on top of Jikes RVM 3.1.3



Implementation

Developed on top of Jikes RVM 3.1.3

Implemented FastTrack, state-of-art happens-before analysis based data race detector



Implementation

Developed on top of Jikes RVM 3.1.3

Implemented FastTrack, state-of-art happens-before analysis based data race detector

Shared on Jikes RVM Research Archive and ACM DL



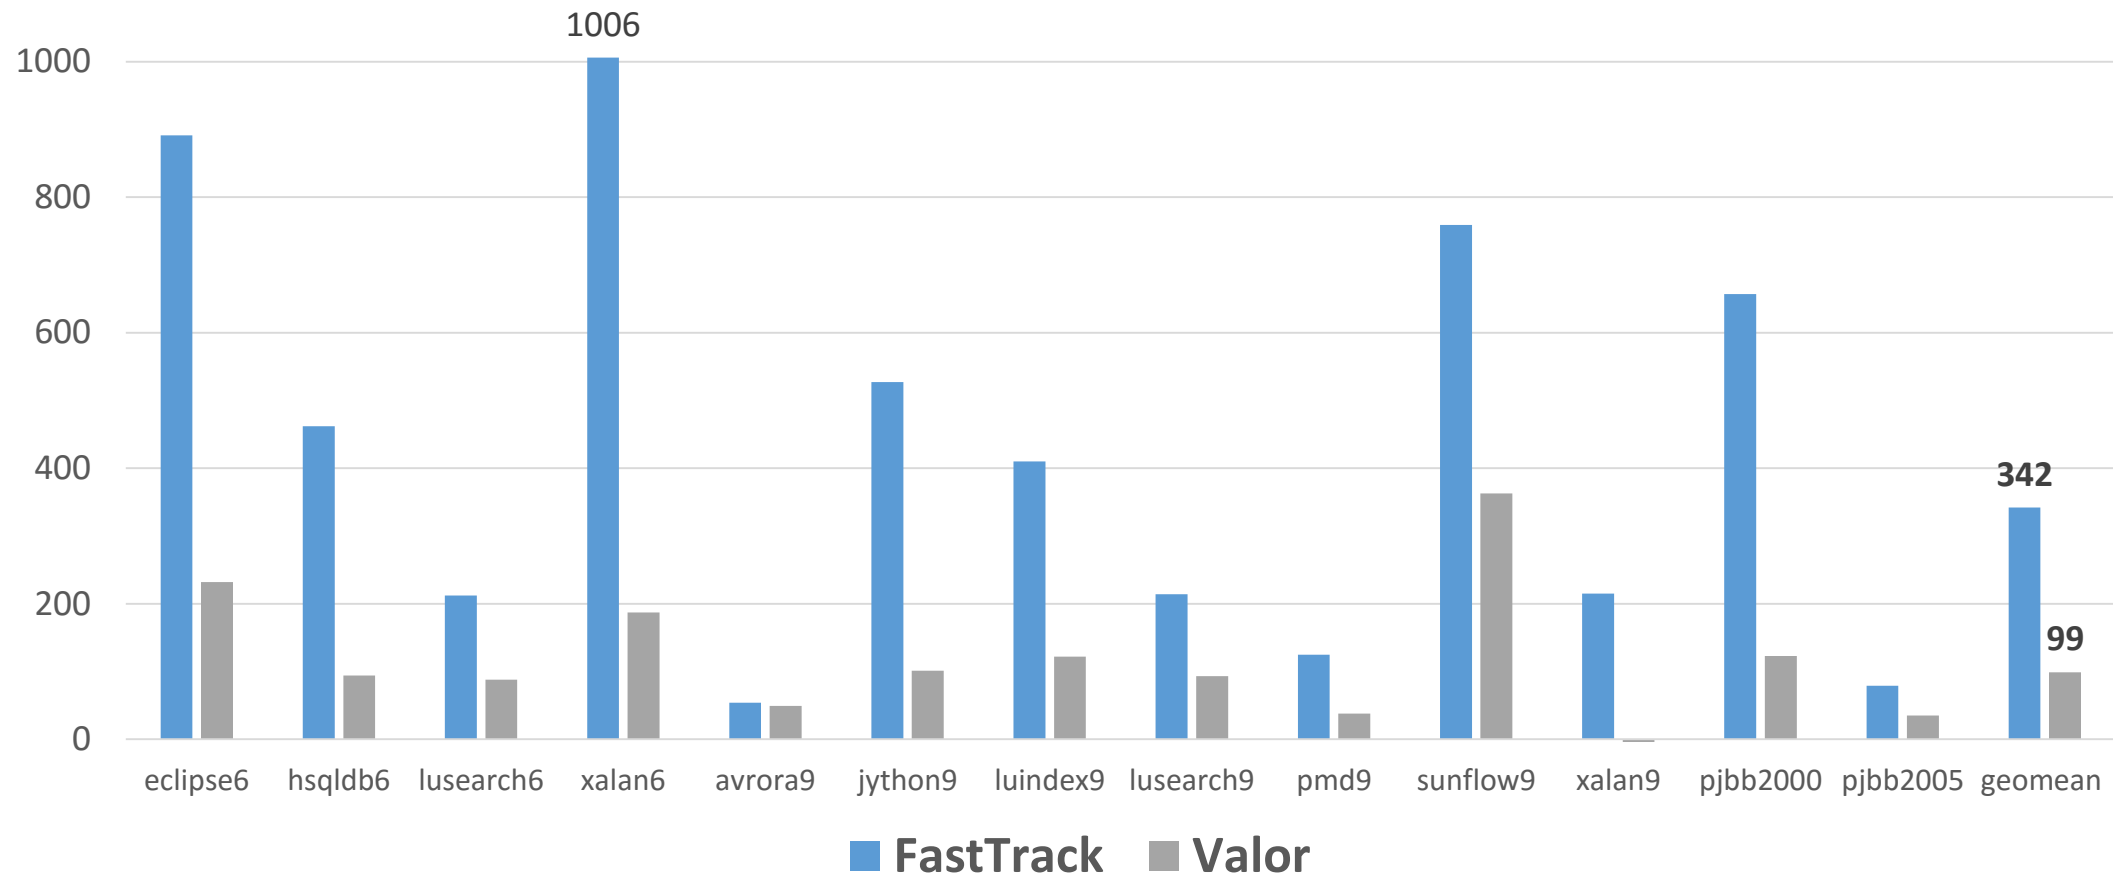
EVALUATION

Experimental Methodology

- **Benchmarks**
 - Large workload sizes of DaCapo 2006 and 9.12-bach suite
 - Fixed-workload versions of SPECjbb2000 and SPECjbb2005
- **Platform**
 - 64-core AMD Opteron 6272


Performance Comparison

Overheads (%) Over Unmodified Jikes RVM



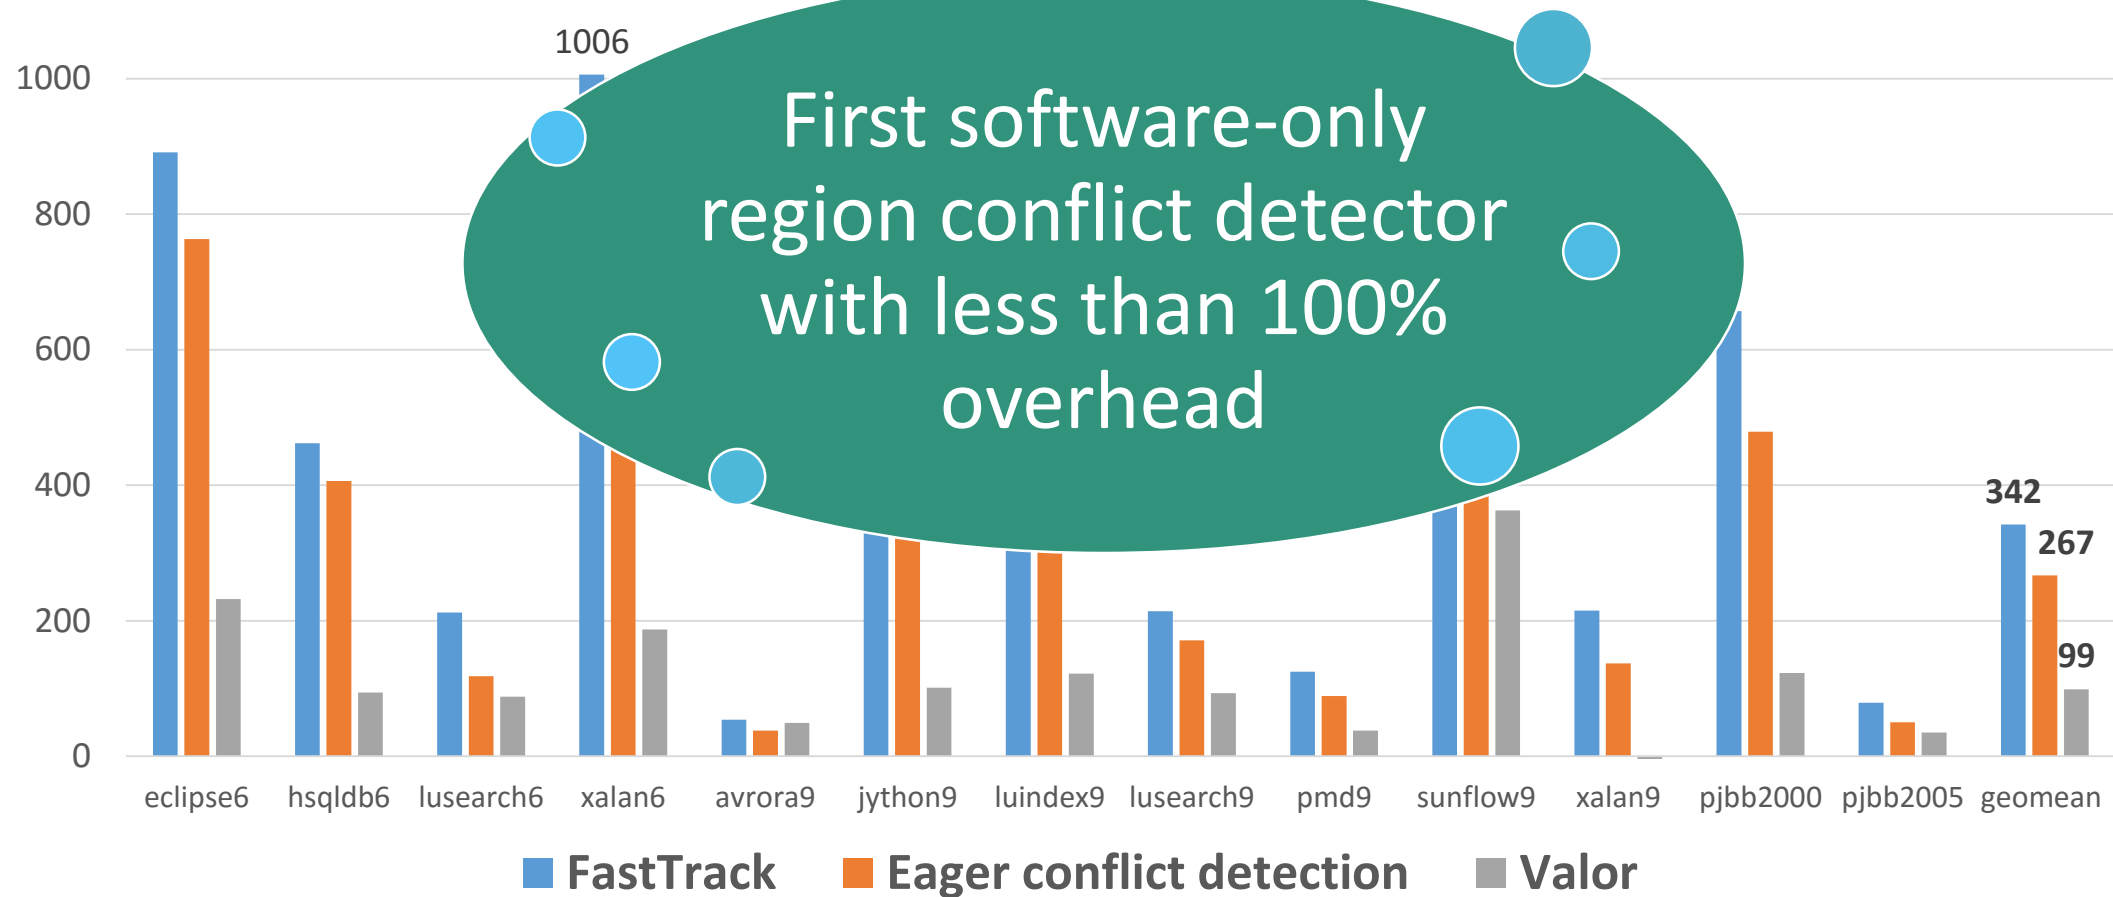
Performance Comparison

Overheads (%) Over Unmodified Jikes RVM



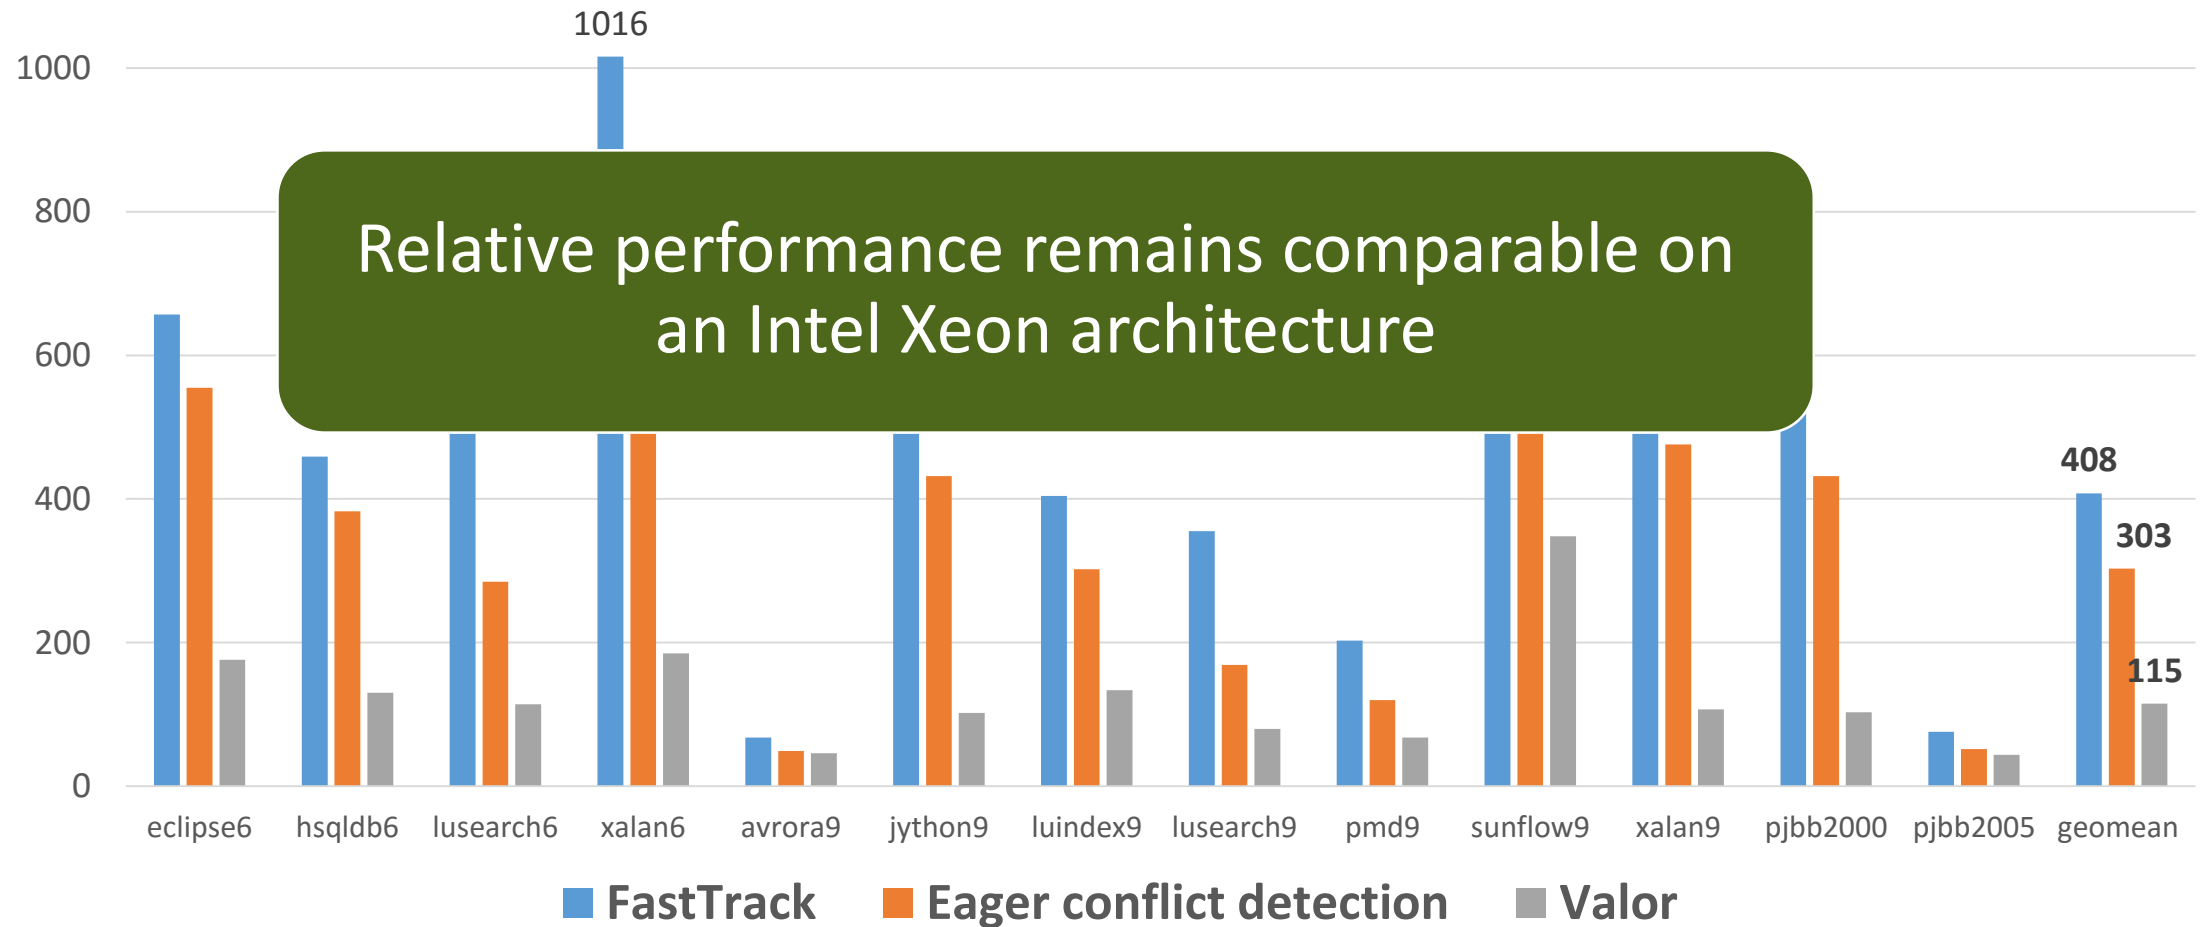
Performance Comparison

Overheads (%) Over Unmodified Jikes RVM

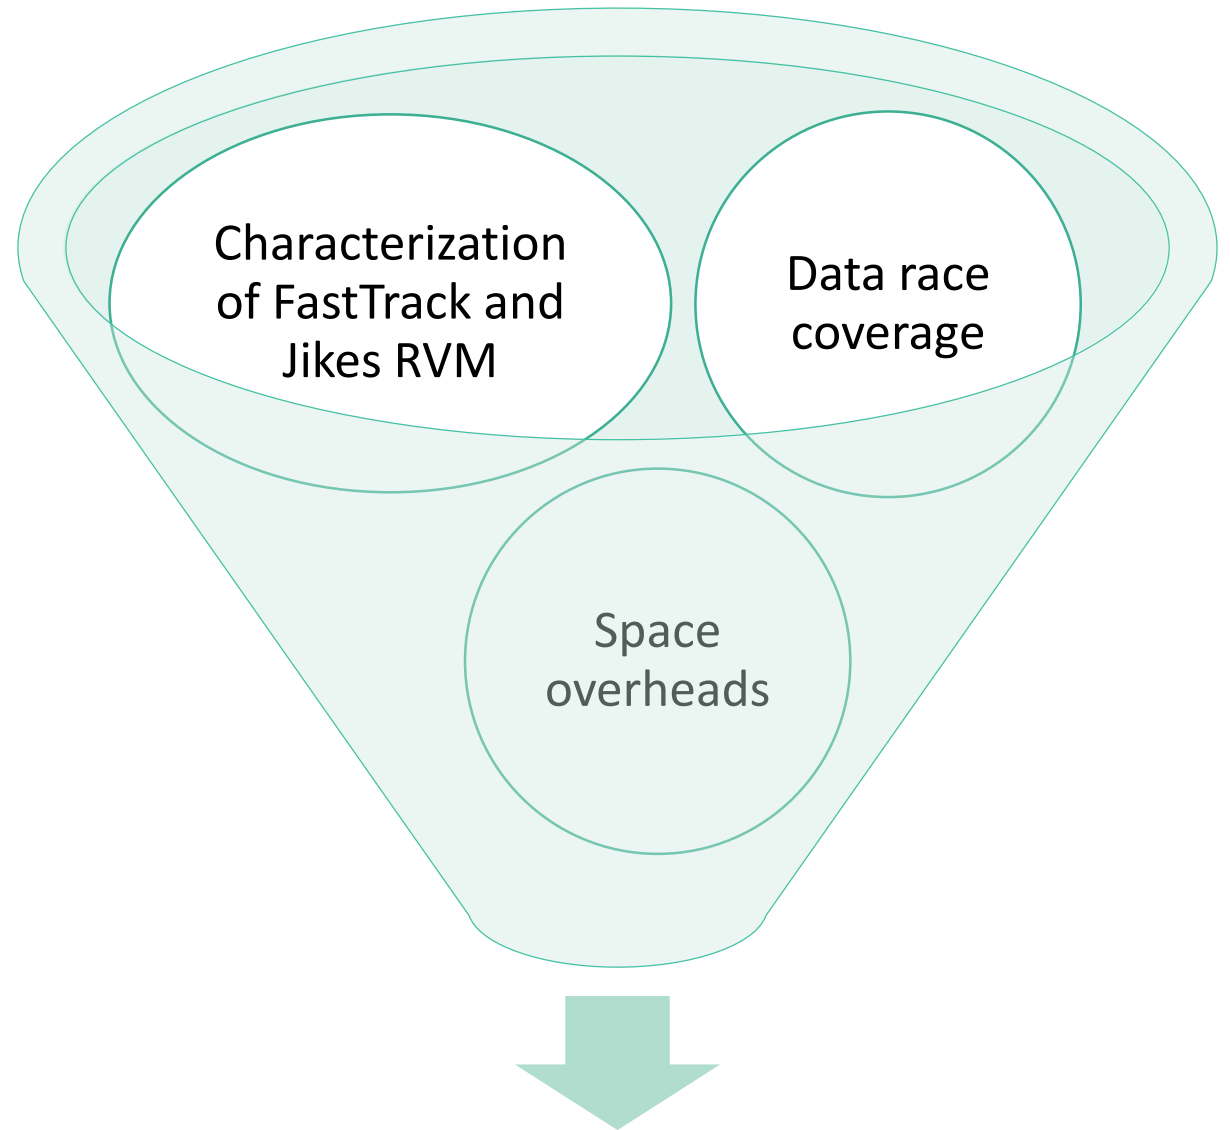


Performance Comparison: Intel Xeon

Overheads (%) Over Unmodified Jikes RVM



Additional Experiments



Please check the paper

Valor: Contributions

Strong execution guarantees in software

Detects all violations of region serializability

Exception-free
execution



Strong semantics

Advances state-of-art

- ❖ Provides strong semantics in software at less than 100% overhead

New Opportunities with Valor

Semantic guarantees Language runtimes could **integrate** this

Debugging Can be used to detect **problematic** data races

Conflict exceptions **All-the-time monitoring** in certain environments

Aggressive optimizations **Reorder and eliminate** redundant loads and stores within synchronization-free regions

Valor: Efficient, Software-Only Region Conflict Exceptions

Swarnendu Biswas, Minjia Zhang, and

Michael D. Bond

Ohio State University

Brandon Lucia

Carnegie Mellon University

OOPSLA 2015

