

Semantic Atomicity for Multithreaded Programs

Jacob Burnim, George Necula, Koushik Sen
Parallel Computing Laboratory
University of California, Berkeley

- ❖ Difficult to write correct parallel software.
 - **Key:** Interference between parallel threads.
 - **Atomicity** – freedom from harmful interference; a fundamental parallel correctness property.
- ❖ Today: **Semantic atomicity**.
 - Specifying atomicity with respect to user-defined, semantic equivalence.
 - Efficiently testing such specifications.
 - **Overall Goal:** Lightweight, useful specs to help programmers find and fix parallelism bugs.

- ❖ Overview + Motivation
- ❖ **Background: Atomicity**
- ❖ Specifying Semantic Atomicity
- ❖ Testing Semantic Atomicity
- ❖ Experimental Evaluation
- ❖ Conclusion

- ❖ Atomicity a **non-interference** property.
 - Block of code is **atomic** if it behaves **as if** executed all-at-once and without interruption.
 - Interference from other threads is **benign** – cannot change overall program behavior.

- ❖ Atomicity a **non-interference property**.
 - Block of code is **atomic** if it behaves **as if** executed all-at-once and without interruption.

```
int bal = 0;  
  
deposit(int a) {  
    @atomic {  
        int t = bal;  
        bal = t + a;  
    }  
}
```

Atomic **specification**.

Programmer **intends**
that this code is atomic.

Want to **check** specification.
Is the code actually atomic?

- ❖ Atomicity a **non-interference** property.
 - Block of code is **atomic** if it behaves **as if** executed all-at-once and without interruption.

```
int bal = 0;  
  
deposit(int a) {  
    @atomic {  
        int t = bal;  
        bal = t + a;  
    }  
}
```

Thread 1:

deposit(10)

t = 0

bal = 10

Thread 2:

deposit(5)

t = 0

bal = 5

Atomicity specification does **not** hold.

- ❖ Atomicity a **non-interference property**.
 - Block of code is **atomic** if it behaves **as if** executed all-at-once and without interruption.

```
int bal = 0;  
  
deposit(int a) {  
    @atomic {  
        int t = bal;  
        while (!CAS(&bal, t, t+a))  
            t = bal;  
    }  
}
```

With CAS, updates to balance are atomic.

Atomicity specification **does** hold.


- ❖ **Formally:** Two semantics for a program P with specified atomic blocks.
 - **Interleaved:** Threads interleave normally.



- **Serial:** When one thread opens an atomic block, no other thread runs until it closes.



- ❖ Formally, program P is **atomic** iff:
 - **For all interleaved** executions E yielding s_1 , there **exists** a **serial** E' yielding an identical final state.



- ❖ Overview
- ❖ Background: Atomicity
- ❖ **Specifying Semantic Atomicity**
- ❖ Testing Semantic Atomicity
- ❖ Experimental Evaluation
- ❖ Conclusion

```
ConcurrentLinkedQueue q;  
q.add(1); q.add(1);
```

Thread 1:

```
@atomic {  
    q.remove(1);  
}
```

Thread 2:

```
@atomic {  
    q.remove(1);  
}
```

- ❖ Michael & Scott non-blocking queue, in the Java standard library
- ❖ Internally, a linked list with lazy deletion.


Thread 1:

```
@atomic {
    q.remove(1);
}
```

Thread 2:

```
@atomic {
    q.remove(1);
}
```

- ❖ In any **serial** execution:




Thread 1:

```
@atomic {
    q.remove(1);
}
```

Thread 2:

```
@atomic {
    q.remove(1);
}
```

- ❖ But in an **interleaved** execution:




Thread 1:

```
@atomic {
    q.remove(1);
}
```

Thread 2:

```
@atomic {
    q.remove(1);
}
```

- ❖ Traditional atomicity requires:




Thread 1:

```
@atomic {
    q.remove(1);
}
```

Thread 2:

```
@atomic {
    q.remove(1);
}
```

- ❖ Traditional atomicity requires:




Thread 1:

```
@atomic {
    q.remove(1);
}
```

Thread 2:

```
@atomic {
    q.remove(1);
}
```

Replace with user-defined
semantic equivalence.




\forall interleavings

Initial State s_0

\exists serial execution

Final State s_1'

$s_1 \neq s_1'$




Thread 1:

```
@atomic {
    q.remove(1);
}
```

Thread 2:

```
@atomic {
    q.remove(1);
}
```

Replace with user-defined
semantic equivalence.




Thread 1:

```
@atomic {
    q.remove(1);
}
```

Thread 2:

```
@atomic {
    q.remove(1);
}
```

Atomicity predicate: `q.equals(q')`



Thread 1:

```
@atomic {  
    q.remove(1);  
}
```

Thread 2:


```
@atomic {  
    q.remove(1);  
}
```

Atomicity predicate: `q.equals(q')`


Bridge predicate.

- ❖ Burnim, Sen, “Asserting and Checking Determinism for Multithreaded Programs”, FSE 2009, CACM 2010.

❖ Semantic **Atomicity**:



❖ Semantic **Determinism**:



```
int bal = 0;
int conflicts = 0;

deposit(int a) {
    @atomic {
        int t = bal;
        while (!CAS(&bal, t, t+a)) {
            t = bal;
            conflicts += 1;
        }
    }
}
```

With CAS, updates to balance are atomic.

“Performance counter” of # of CAS failures.

Atomicity predicate: **bal == bal'**

```
ConcurrentList list;
```

Thread 1:

```
@atomic {  
    ...  
    list.add(1);  
    ...  
    list.add(2);  
}
```

Thread 2:

```
@atomic {  
    ...  
    list.add(3);  
    ...  
    list.add(4);  
}
```

Atomicity predicate: `eqSets(list,list')`

- ❖ If `list` is [1,3,2,4], an atomicity violation?
 - User must specify **intended** atomicity.


- ❖ Introduction + Motivation
- ❖ Background: Atomicity
- ❖ Specifying Semantic Atomicity
- ❖ **Testing Semantic Atomicity**
- ❖ Experimental Evaluation
- ❖ Conclusion

- ❖ **Interleaved run E is semantically atomic** w.r.t. Ψ iff there exists a **serial** run E' s.t.:
 - The final states of E, E' satisfy $\Psi(s_1, s'_1)$.

$E : s_0$  s_1

Is E semantically atomic w/ respect to Ψ ?




- ❖ **Interleaved run E is semantically atomic** w.r.t. Ψ iff there exists a **serial** run E' s.t.:
 - The final states of E, E' satisfy $\Psi(s_1, s'_1)$.



⋮




Infeasible to try all serial executions.

Can we **restrict** this search?






⋮

1. The final states of E, E' satisfy $\Psi(s_1, s'_1)$.
2. E and E' execute the same atomic blocks.






:

1. The final states of E, E' satisfy $\Psi(s_1, s'_1)$.
2. E and E' execute the same atomic blocks.
3. Non-overlapping atomic blocks appear in the same order in E and E' .





⋮

- ❖ **Def:** Interleaved run E is **semantically serializable** iff exists a serial run E' s.t.:
 1. The final states of E, E' satisfy $\Psi(s_1, s'_1)$.
 2. E and E' execute the same atomic blocks.



❖ **Def:** Interleaved E is **semantically strictly serializable** iff exists a serial run E' s.t.:

1. The final states of E, E' satisfy $\Psi(s_1, s'_1)$.
2. E and E' execute the same atomic blocks.
3. Non-overlapping atomic blocks appear in the same order in E and E' .



- ❖ **Def:** Interleaved E is **semantically strictly serializable** iff exists a serial run E' s.t.:
 1. The final states of E, E' satisfy $\Psi(s_1, s'_1)$.
 2. E and E' execute the same atomic blocks.
 3. Non-overlapping atomic blocks appear in the same order in E and E' .

E has **N** blocks, with **$\leq K$ overlapping**.

\implies

Can check semantic strict serializability by examining **$\leq K!$** serial runs.

- ❖ To test atomicity of program P:
 - Systematically/randomly generate executions E with **$\leq K$ overlapping atomic blocks.**
 - For each E, report a violation if not **semantically strictly serializable**.
- ❖ **Small Scope Hypothesis:** Can find bugs with small # of overlapping atomic blocks.

- ❖ Introduction + Motivation
- ❖ Background: Atomicity
- ❖ Specifying Semantic Atomicity
- ❖ Testing Semantic Atomicity
- ❖ **Experimental Evaluation**
- ❖ Conclusion

- ❖ Wrote semantic atomicity specs for several Java benchmarks.
 - Concurrent data structures and parallel apps.
- ❖ **Setup:** For each benchmark:
 - Generate 200-900 random interleaved runs, with one atomic block interrupted by ≤ 4 others.
 - Check semantic strict serializability of each.
 - To compare, also check conflict-serializability.

Experimental Results I

| Benchmark | LoC | Test Runs | Semantic Atomicity Violations | | Strict Atomicity Violations | |
|---------------------|------------|------------------|--------------------------------------|---------------|------------------------------------|---------------|
| | | | Runs | Static Blocks | Runs | Static Blocks |
| JDK LinkedQueue | 200 | 241 | 7 | 2 | »7 | 4 |
| JDK SkipListMap | 1400 | 487 | 6 | 2 | »7 | 4 |
| JDK CwArrayList | 600 | 222 | 0 | 0 | 0 | 0 |
| lock-free list | 100 | 319 | 57 | 1 | »57 | 2 |
| lazy list-based set | 100 | 231 | 0 | 0 | »0 | 2 |

Experimental Results II

| Benchmark | LoC | Test Runs | Semantic Atomicity Violations | | Strict Atomicity Violations | |
|------------------|------------|------------------|--------------------------------------|---------------|------------------------------------|---------------|
| | | | Runs | Static Blocks | Runs | Static Blocks |
| PJ pi | 150 | 20 | 5 | 1 | 5 | 1 |
| PJ keysearch | 200 | 904 | 0 | 0 | 0 | 0 |
| PJ fractal | 250 | 73 | 0 | 0 | 0 | 0 |
| PJ phylogenetic | 4400 | 603 | 27 | 1 | »27 | 2 |

Application benchmarks from Parallel Java Library
(Kaminsky 2007), use ~15000 LoC from PJ library.

```
ConcurrentLinkedQueue q;
q.add(1); q.add(2);
```

Thread 1:

```
@atomic {
    q.remove(1);
}
@atomic {
    q.add(3);
}
```

Thread 2:

```
@atomic {
    sz = q.size();
}
```

Atomic with respect to:

$$q.equals(q') \wedge (sz == sz')$$

- ❖ **Not atomic:** `q.size()` can return `sz=3`.

```

parallel-for (t in trees) {
    @atomic {
        cost = compute_cost(t);
        synchronized (min_cost) {
            min_cost = min(min_cost, cost);
        }
        if (cost == min_cost) {
            min_tree = t;
        }
    }
}

```

Atomic with respect to:

$$\begin{aligned}
 & \text{min_tree.equals(min_tree')} \\
 & \wedge (\text{min_cost} == \text{min_cost'})
 \end{aligned}$$

Updates to
min_tree not
synchronized.

- ❖ Introduction + Motivation
- ❖ Background: Atomicity
- ❖ Specifying Semantic Atomicity
- ❖ Testing Semantic Atomicity
- ❖ Experimental Evaluation
- ❖ Conclusion

- ❖ **Semantic atomicity.**
 - Generalization for capturing high-level non-interference properties of real, complex code.
 - Testing via strict serializability.
 - Found several unknown atomicity errors.
- ❖ **Overall Goal:** Lightweight specifications for parallel correctness.
 - Easy for programmers to write.
 - With testing, effective in finding real bugs.
 - Determinism [CACM'10, ICSE'10], NDSeq [PLDI '11]

Questions?