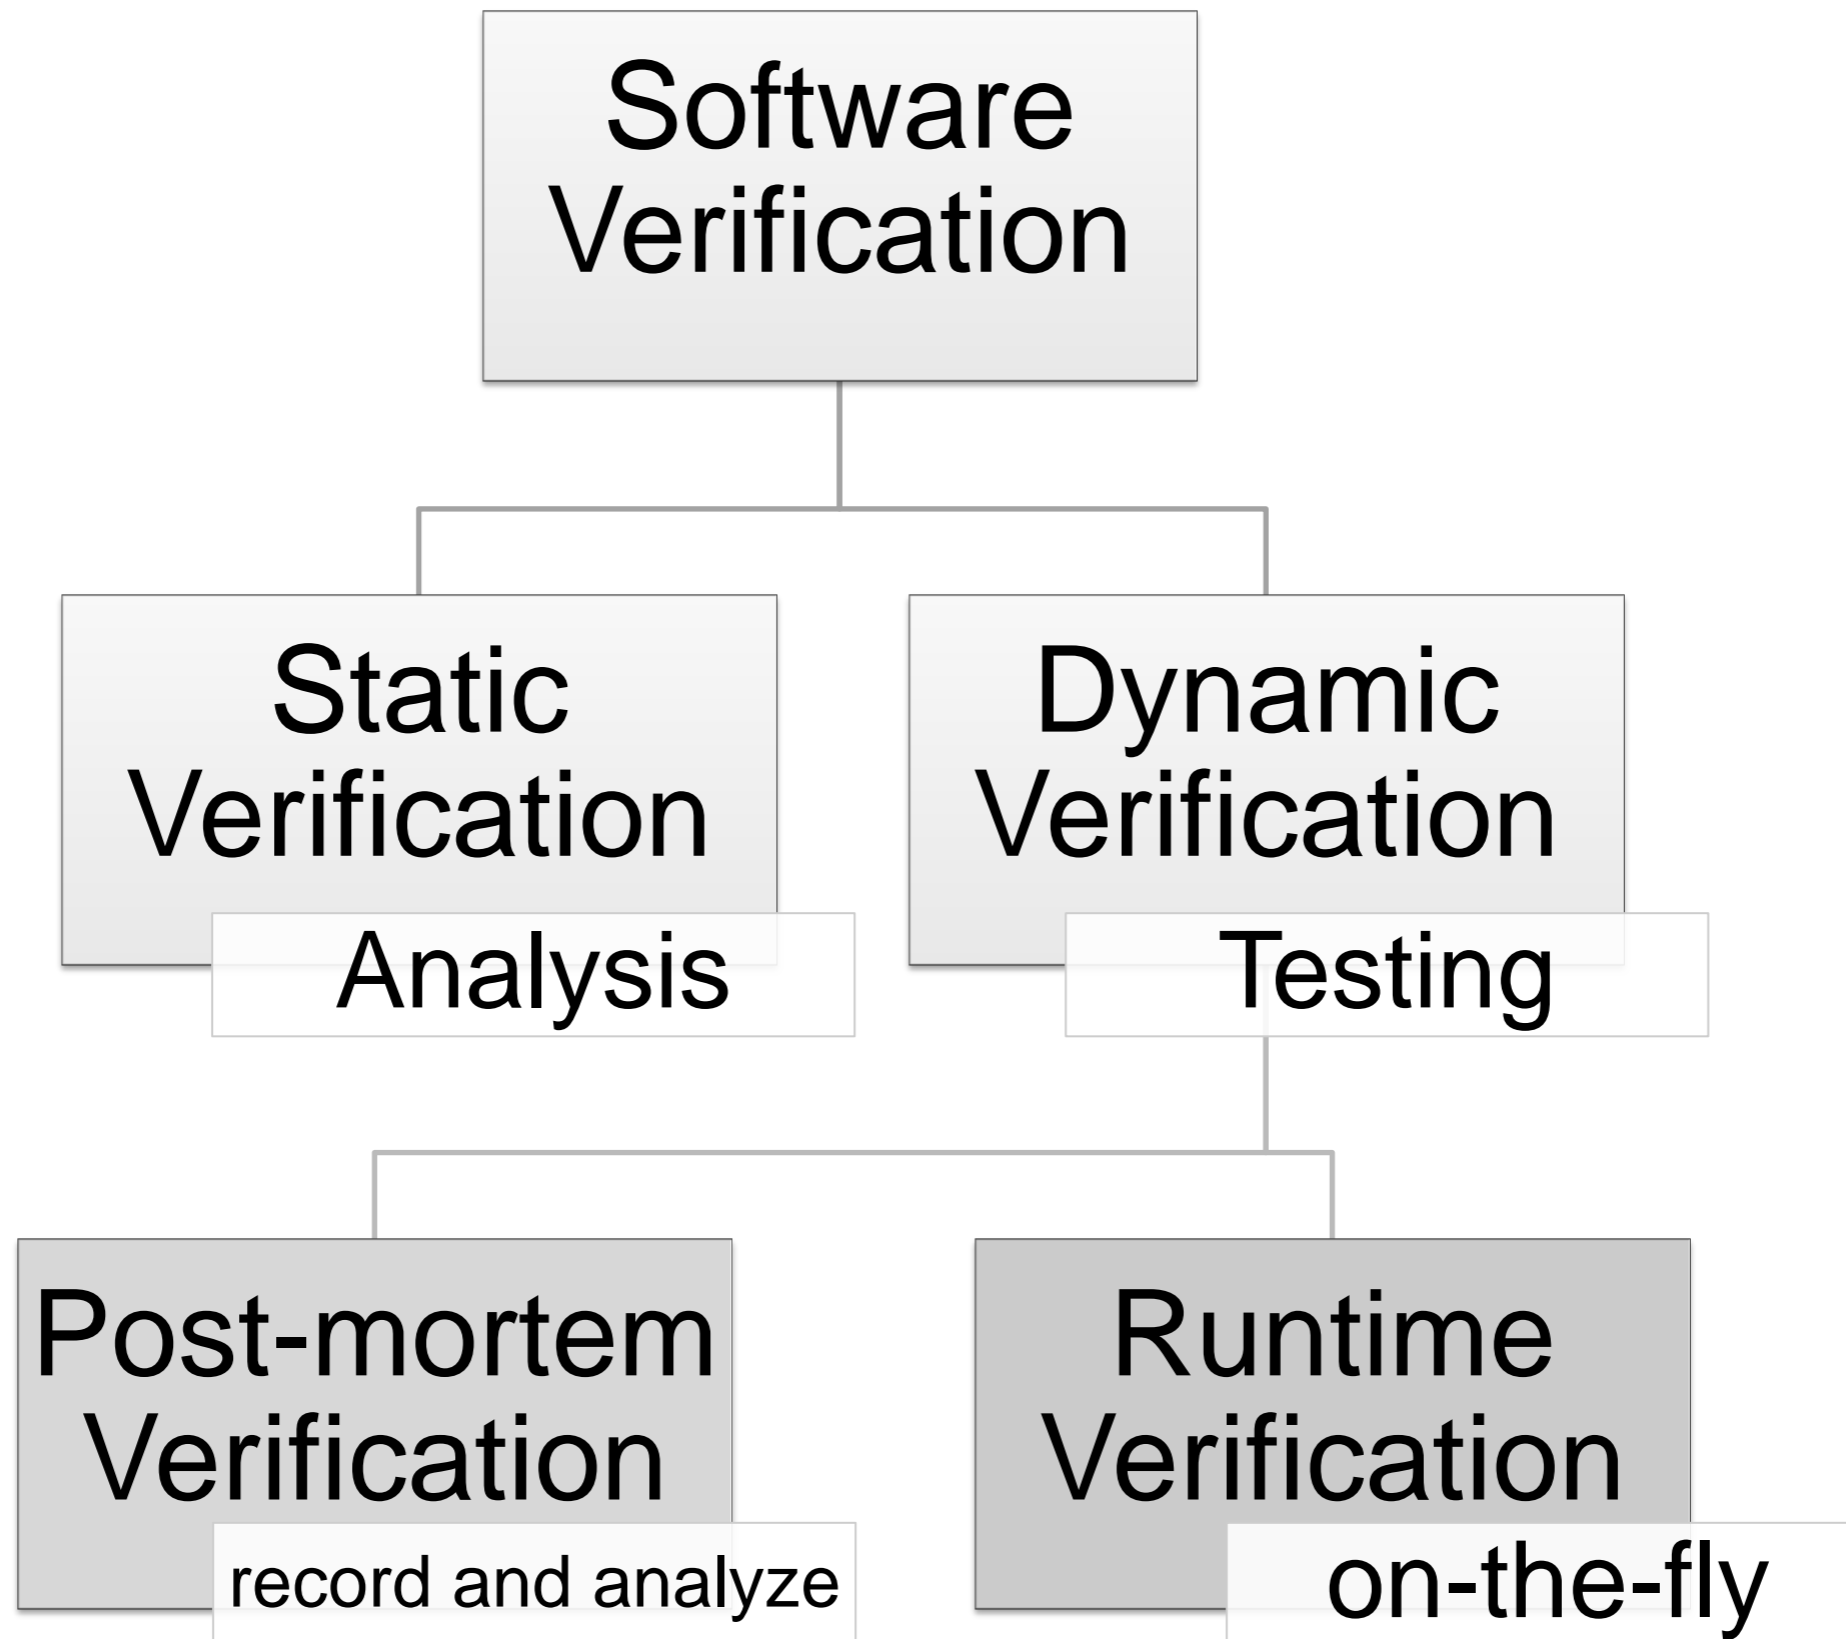


# Kuda: The Split Race Checker

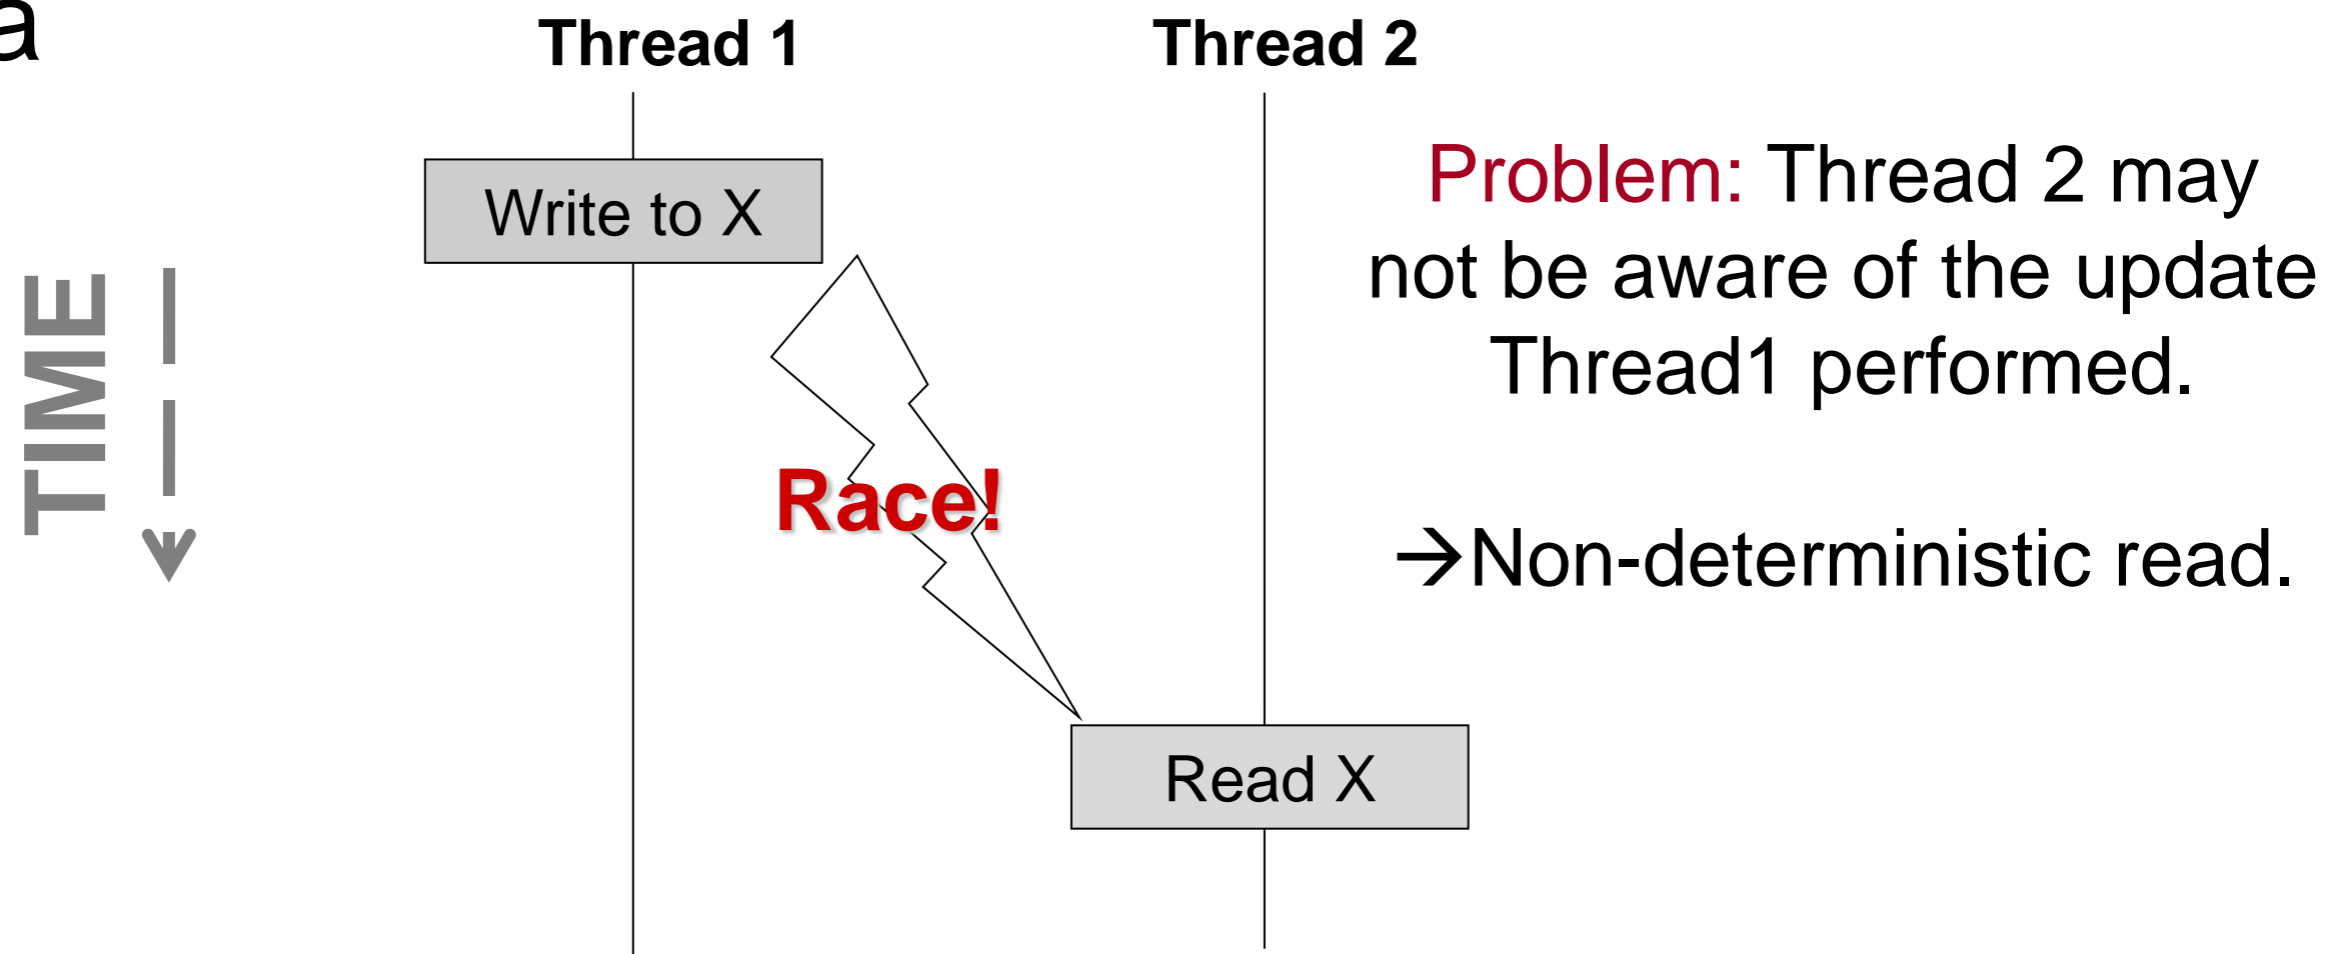
Can Bekar (Koç U), Tayfun Elmas (UC Berkeley),  
Semih Okur (UIUC), Serdar Tasiran (Koç U)



# Data Races

It is **not** standardized for each memory model yet But here it goes:

Conflicting, unsynchronized accesses to shared data



# Verifying for the Race- freedom

How can we prove the absence of an element in a set?

Check each interleaving and also check for each input?

**Check all  
the elements**



**If you don't see any races,  
you are good**

# “Verifying” the Absence of (Data) Races in Runtime

Bad news:

It is not practical to check each interleaving and each input

Good news:

You can *convince* yourself that your coverage is good  
-and interesting- enough 😊

Conclusion:

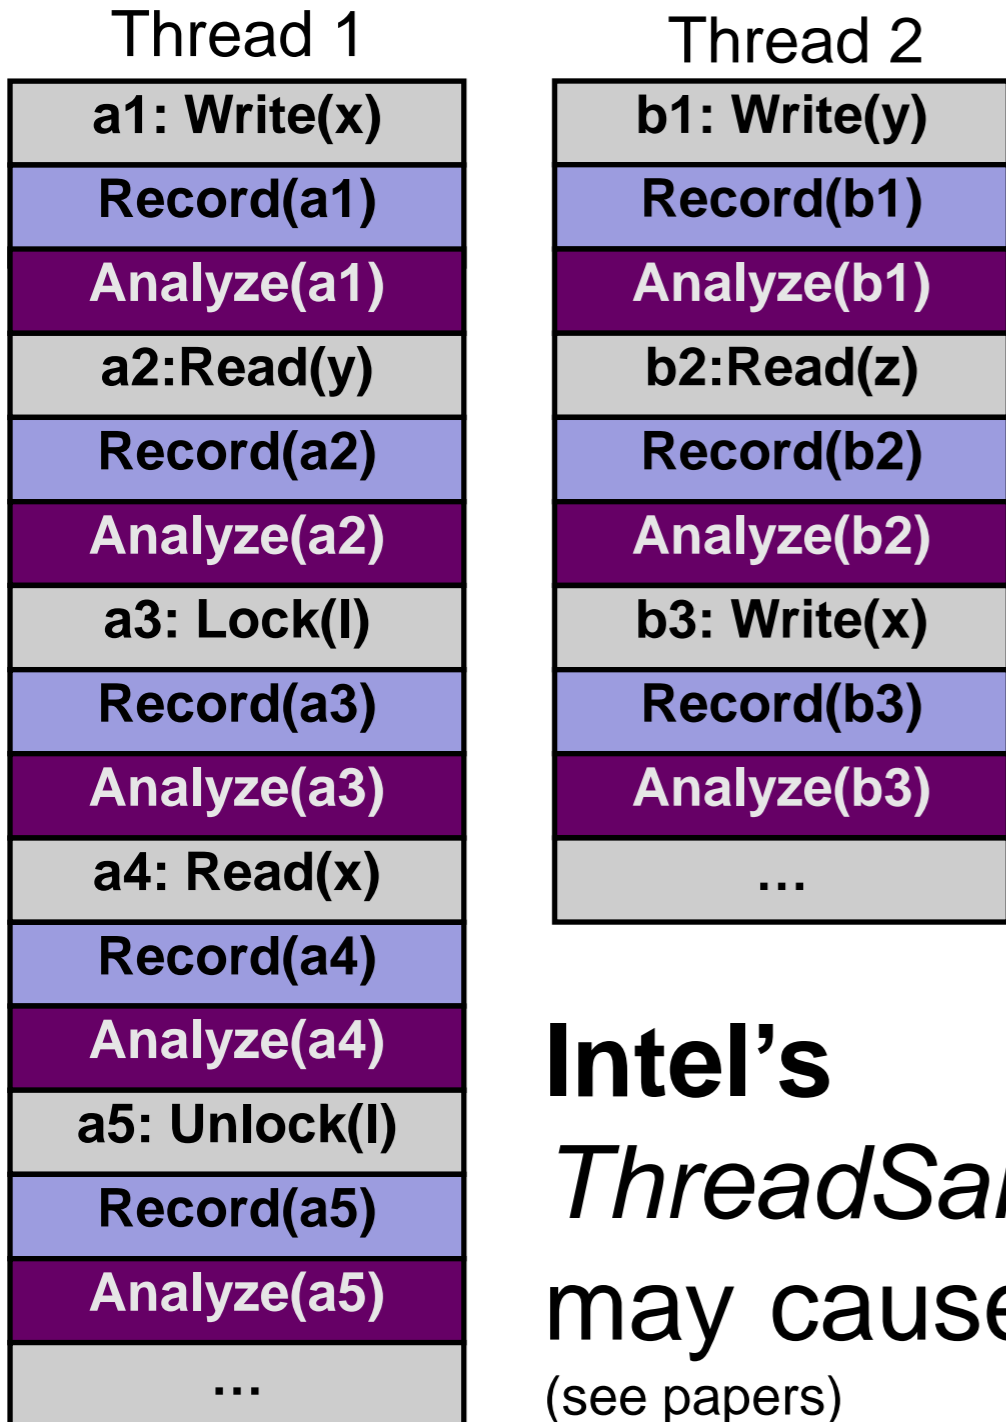
It is not *formal* verification, but bosses like the idea  
i.e. the less bugzilla entries after the software launch, the better

Result:

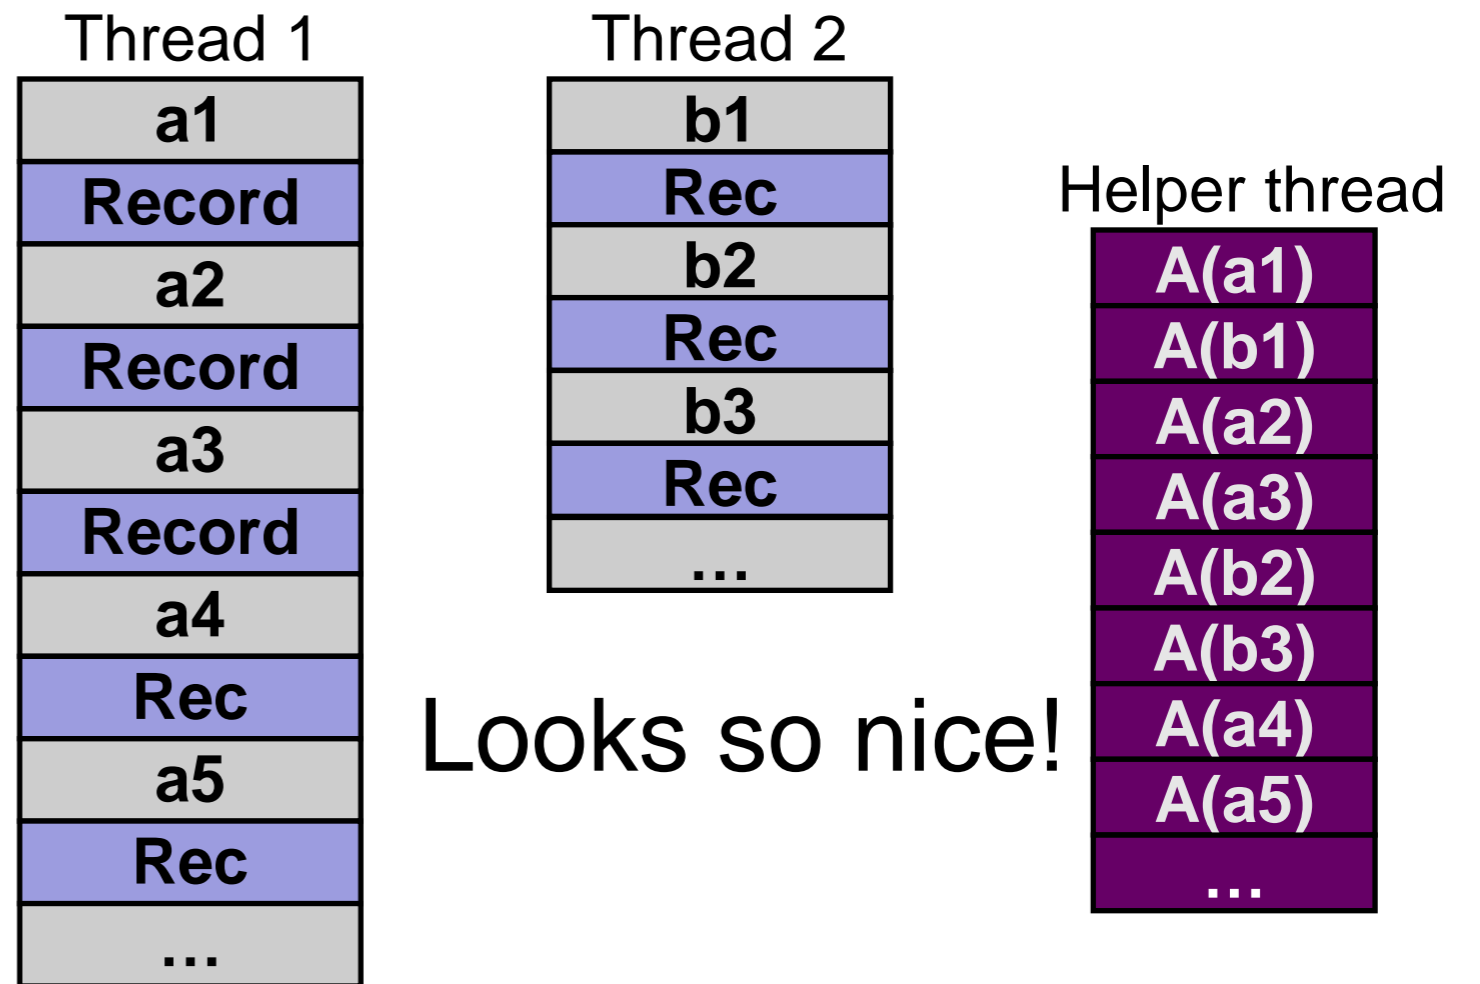
Stress Testing

# On-the-fly Race Checking

## Traditional Race Checking




## Race Checking in Parallel



Looks so nice!

**Intel's ThreadSanitizer and Helgrind** → **Google's Checker, Helgrind** → each tool may cause several hundreds times slowdown!  
(see papers)

# Producer/Consumer Relation on all CPU Records



an event = int4  $\rightarrow$  event{varID,threadID,type,index}

a frame = 1024 events

the buffer = 2048 frames = 16 MB


# Does It Work?

|  |       |
|--|-------|
| Baseline Program (computation only)  | 3.8   |
| + Runtime instrumentation  | sec   |
| + Cheap Algorithm (on CPU)   | 3x    |
| • Eraser – Simple/Fast & Imprecise   | ~400x |
| + Event buffering  | 78x   |
| + Expensive Algorithm (on GPU)   | 83.6x |
| • Goldilocks – Complex & Precise   | 5x    |
| = Precise race detection on GPU is <u>5x</u> faster than imprecise race detection on CPU |       |

(see experiments, 2 slides later)



# Non-blocking, Lock-free Buffer



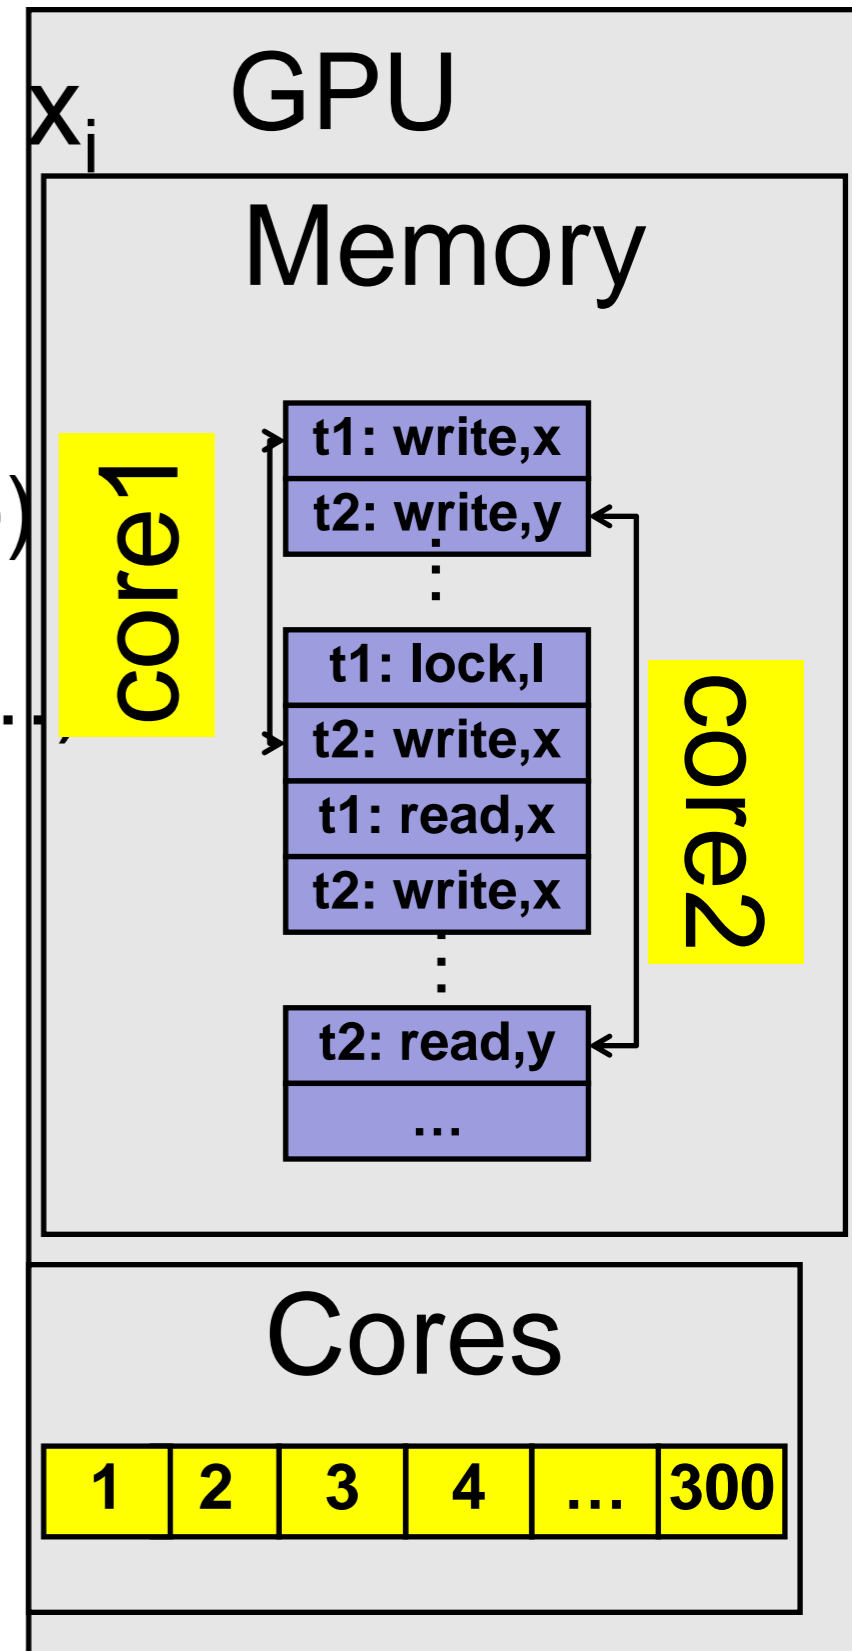
16 MB buffer is enough

Always room for the next event to be recorded (GPU is fast)

# Parallel Race

## Checking

- Suppose  $\text{event}[i]$  accesses  $\text{var } x_i$
- GPU core[i]
  - $\text{handle}(\text{event}[i] \text{ until } \text{next\_event}(x_i))$   
when done →  $\text{handle}(\text{event}[i+N] \dots)$
- Read-only device memory!
- No shared data besides the frame



| Benchmarks     | Normal exe | Instrumented | Eraser@CPU                | Event buffer | Goldilocks@GPU  |
|----------------|------------|--------------|---------------------------|--------------|-----------------|
|                | Reference  | Slowdown     | Slowdown                  | Slowdown     | <b>Slowdown</b> |
| blackscholes   | 1          | 2.1          | 101                       | 14.7         | <b>21.1</b>     |
| bodytrack      | 1          | 2.6          | 251                       | 72.5         | <b>75.6</b>     |
| canneal        | 1          | 1.6          | 47                        | 6.9          | <b>7.4</b>      |
| dedup          | 1          | 3.1          | 429                       | 88           | <b>101.6</b>    |
| fluidanimate   | 1          | 2.5          | 308                       | 83.9         | <b>88.1</b>     |
| raytrace       | 1          | 2            | 123.7                     | 5.7          | <b>6.2</b>      |
| swaptions      | 1          | 3            | 437                       | 117.5        | <b>119.9</b>    |
| x264           | 1          | 7.1          | 645.2                     | 151.7        | <b>155.8</b>    |
| barnes         | 1          | 4            | 499.1                     | 111.5        | <b>116.1</b>    |
| cholesky       | 1          | 3.2          | 216.2                     | 32.4         | <b>33.2</b>     |
| fmm            | 1          | 2.9          | 1455.8                    | 89.6         | <b>98.6</b>     |
| fft            | 1          | 2.1          | 222.7                     | 45.3         | <b>48.1</b>     |
| lu             | 1          | 5.9          | 742.2                     | 190          | <b>201.9</b>    |
| ocean          | 1          | 4            | 301.6                     | 63.4         | <b>68.1</b>     |
| radix          | 1          | 2            | 126.6                     | 31.1         | <b>33.1</b>     |
| raytrace       | 1          | 2            | 122.2                     | 6.9          | <b>7.2</b>      |
| water-n2       | 1          | 3.5          | 707.6                     | 182.1        | <b>191.4</b>    |
| w-spatial      | 1          | 4.6          | 485.2                     | 114.9        | <b>131.5</b>    |
| <b>average</b> | <b>1x</b>  | <b>3.2x</b>  | <sup>11</sup> <b>401x</b> | <b>78x</b>   | <b>83x</b>      |

# Future Work

- Profiling + Pruning of thread local vars
- Compare state-of-the-art RC in C(Fast-track)
- Offloading other concurrent analysis to GPU
- Reduce slowdown due to buffering
  - Custom HW

# Q&A



Our codebase is online:  
[kuda.codeplex.com](http://kuda.codeplex.com)

# Some Links



[msrc.ku.edu.tr](http://msrc.ku.edu.tr)

We need new graduate students



[rv2012.ku.edu.tr](http://rv2012.ku.edu.tr)

International Runtime Verification Conference in  
Istanbul