

# The Long Road towards Testing Multi-Threaded Programs with KLEE

Daniel Schemmel, Julian Büning, David Laprell, Klaus Wehrle

RWTH Aachen University

<https://comsys.rwth-aachen.de>

Zoomspace, 2021-06-11

- As core counts rise, multi-threaded programs become more and more prevalent

- As core counts rise, multi-threaded programs become more and more prevalent
- We believe that Symbolic Execution can help test those programs [1]

[1] Symbolic Partial-Order Execution for Testing Multi-Threaded Programs.  
D. Schemmel, J. Büning, C. Rodríguez, D. Laprell, K. Wehrle. CAV 2020.

- As core counts rise, multi-threaded programs become more and more prevalent
- We believe that Symbolic Execution can help test those programs [1]
- Adding support for multi-threaded programs to KLEE initially sounds simple:

[1] Symbolic Partial-Order Execution for Testing Multi-Threaded Programs.  
D. Schemmel, J. Büning, C. Rodríguez, D. Laprell, K. Wehrle. CAV 2020.

- As core counts rise, multi-threaded programs become more and more prevalent
- We believe that Symbolic Execution can help test those programs [1]
- Adding support for multi-threaded programs to KLEE initially sounds simple:
  1. Extend KLEE states to manage multiple thread stacks

[1] Symbolic Partial-Order Execution for Testing Multi-Threaded Programs.  
D. Schemmel, J. Büning, C. Rodríguez, D. Laprell, K. Wehrle. CAV 2020.

- As core counts rise, multi-threaded programs become more and more prevalent
- We believe that Symbolic Execution can help test those programs [1]
- Adding support for multi-threaded programs to KLEE initially sounds simple:
  1. Extend KLEE states to manage multiple thread stacks
  2. Add a concurrency-aware searcher

[1] Symbolic Partial-Order Execution for Testing Multi-Threaded Programs.  
D. Schemmel, J. Büning, C. Rodríguez, D. Laprell, K. Wehrle. CAV 2020.

- As core counts rise, multi-threaded programs become more and more prevalent
- We believe that Symbolic Execution can help test those programs [1]
- Adding support for multi-threaded programs to KLEE initially sounds simple:
  1. Extend KLEE states to manage multiple thread stacks
  2. Add a concurrency-aware searcher
  3. Model the pthread API

[1] Symbolic Partial-Order Execution for Testing Multi-Threaded Programs.  
D. Schemmel, J. Büning, C. Rodríguez, D. Laprell, K. Wehrle. CAV 2020.

- As core counts rise, multi-threaded programs become more and more prevalent
- We believe that Symbolic Execution can help test those programs [1]
- Adding support for multi-threaded programs to KLEE initially sounds simple:
  1. Extend KLEE states to manage multiple thread stacks
  2. Add a concurrency-aware searcher
  3. Model the pthread API
- **Where is the catch?**

[1] Symbolic Partial-Order Execution for Testing Multi-Threaded Programs.  
D. Schemmel, J. Büning, C. Rodríguez, D. Laprell, K. Wehrle. CAV 2020.



- KLEE states are easily extended to track multiple stacks and instruction pointers

- KLEE states are easily extended to track multiple stacks and instruction pointers
- Executing every instruction interleaving causes **immediate state explosion!**

- KLEE states are easily extended to track multiple stacks and instruction pointers
- Executing every instruction interleaving causes **immediate state explosion!**
  - ▶ # of possible paths = #threads<sup>#instructions per thread</sup>

- KLEE states are easily extended to track multiple stacks and instruction pointers
- Executing every instruction interleaving causes **immediate state explosion!**
  - ▶ # of possible paths = #threads<sup>#instructions per thread</sup>
  - ▶ Tiny example:  $8^{20} = 2^{60}$

- KLEE states are easily extended to track multiple stacks and instruction pointers
- Executing every instruction interleaving causes **immediate state explosion!**
  - ▶ # of possible paths = #threads<sup>#instructions per thread</sup>
  - ▶ Tiny example:  $8^{20} = 2^{60}$
  - ▶ Executing any one instruction turns one state into #threads states

- KLEE states are easily extended to track multiple stacks and instruction pointers
- Executing every instruction interleaving causes **immediate state explosion!**
  - ▶ # of possible paths = #threads<sup>#instructions per thread</sup>
  - ▶ Tiny example:  $8^{20} = 2^{60}$
  - ▶ Executing any one instruction turns one state into #threads states
- **A posteriori approaches (searchers, state pruning) alone are insufficient!**

## Thread 1

```
; a += 1  
%1 = load i32, i32* @a, align 4  
%2 = add nsw i32 %1, 1  
store i32 %2, i32* @a, align 4
```

## Thread 2

```
; b += 1  
%1 = load i32, i32* @b, align 4  
%2 = add nsw i32 %1, 1  
store i32 %2, i32* @b, align 4
```

## Thread 1

```
; a += 1
%1 = load i32, i32* @a, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @a, align 4

; x += 1
%3 = load i32, i32* @x, align 4
%4 = add nsw i32 %3, 1
store i32 %4, i32* @x, align 4
```

## Thread 2

```
; b += 1
%1 = load i32, i32* @b, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @b, align 4

; x += 1
%3 = load i32, i32* @x, align 4
%4 = add nsw i32 %3, 1
store i32 %4, i32* @x, align 4
```



# Coarse Thread-Scheduling

## Thread 1

```
; a += 1
%1 = load i32, i32* @a, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @a, align 4
```

```
; x += 1
%3 = load i32, i32* @x, align 4
%4 = add nsw i32 %3, 1
store i32 %4, i32* @x, align 4
```

## Thread 2

```
; b += 1
%1 = load i32, i32* @b, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @b, align 4
```

```
; x += 1
%3 = load i32, i32* @x, align 4
%4 = add nsw i32 %3, 1
store i32 %4, i32* @x, align 4
```



# Coarse Thread-Scheduling

## Thread 1

```
; a += 1  
%1 = load i32, i32* @a, align 4  
%2 = add nsw i32 %1, 1  
store i32 %2, i32* @a, align 4
```

```
; x += 1  
%3 = load i32, i32* @x, align 4  
%4 = add nsw i32 %3, 1  
store i32 %4, i32* @x, align 4
```

## Thread 2

```
; b += 1  
%1 = load i32, i32* @b, align 4  
%2 = add nsw i32 %1, 1  
store i32 %2, i32* @b, align 4
```

```
; x += 1  
%3 = load i32, i32* @x, align 4  
%4 = add nsw i32 %3, 1  
store i32 %4, i32* @x, align 4
```

**Data races are undefined behavior!**  
**[C11 §5.1.2.4/25] [C18 §5.1.2.4/35]**

- Our KLEE fork detects data races as a new bug category

- Our KLEE fork detects data races as a new bug category
- For each byte, we record which thread read/wrote it last

- Our KLEE fork detects data races as a new bug category
- For each byte, we record which thread read/wrote it last
- Efficient if a memory object is used by one thread exclusively

- Our KLEE fork detects data races as a new bug category
- For each byte, we record which thread read/wrote it last
- Efficient if a memory object is used by one thread exclusively
  - ▶ Most memory objects are not shared between threads

- Our KLEE fork detects data races as a new bug category
- For each byte, we record which thread read/wrote it last
- Efficient if a memory object is used by one thread exclusively
  - ▶ Most memory objects are not shared between threads
  - ▶ Symbolic accesses to shared memory objects may require SMT solving

## Thread 1

```
; a += 1  
%1 = load i32, i32* @a, align 4  
%2 = add nsw i32 %1, 1  
store i32 %2, i32* @a, align 4
```

```
; x += 1  
%3 = load i32, i32* @x, align 4  
%4 = add nsw i32 %3, 1  
store i32 %4, i32* @x, align 4
```

## Thread 2

```
; b += 1  
%1 = load i32, i32* @b, align 4  
%2 = add nsw i32 %1, 1  
store i32 %2, i32* @b, align 4
```

```
; x += 1  
%3 = load i32, i32* @x, align 4  
%4 = add nsw i32 %3, 1  
store i32 %4, i32* @x, align 4
```



## Thread 1

```
; a += 1
%1 = load i32, i32* @a, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @a, align 4

call void @lock()
; x += 1
%3 = load i32, i32* @x, align 4
%4 = add nsw i32 %3, 1
store i32 %4, i32* @x, align 4
call void @unlock()
```

## Thread 2

```
; b += 1
%1 = load i32, i32* @b, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @b, align 4

call void @lock()
; x += 1
%3 = load i32, i32* @x, align 4
%4 = add nsw i32 %3, 1
store i32 %4, i32* @x, align 4
call void @unlock()
```

## Thread 1

```
; a += 1
%1 = load i32, i32* @a, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @a, align 4
```

```
call void @lock()
; x += 1
%3 = load i32, i32* @x, align 4
%4 = add nsw i32 %3, 1
store i32 %4, i32* @x, align 4
call void @unlock()
```

## Thread 2

```
; b += 1
%1 = load i32, i32* @b, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @b, align 4
```

```
call void @lock()
; x += 1
%3 = load i32, i32* @x, align 4
%4 = add nsw i32 %3, 1
store i32 %4, i32* @x, align 4
call void @unlock()
```

## Thread 1

```
; a += 1
%1 = load i32, i32* @a, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @a, align 4
```

- call void @lock()  
; x += 1  
%3 = load i32, i32\* @x, align 4  
%4 = add nsw i32 %3, 1  
store i32 %4, i32\* @x, align 4
- call void @unlock()

## Thread 2

```
; b += 1
%1 = load i32, i32* @b, align 4
%2 = add nsw i32 %1, 1
store i32 %2, i32* @b, align 4
```

- call void @lock()  
; x += 1  
%3 = load i32, i32\* @x, align 4  
%4 = add nsw i32 %3, 1  
store i32 %4, i32\* @x, align 4
- call void @unlock()

- If two paths yield the exact same state, only one needs to be pursued further

- If two paths yield the exact same state, only one needs to be pursued further
  - ▶ Happens all the time in parallel programs

- If two paths yield the exact same state, only one needs to be pursued further
  - ▶ Happens all the time in parallel programs
  - ▶ E.g., consider two threads atomically incrementing a counter

- If two paths yield the exact same state, only one needs to be pursued further
  - ▶ Happens all the time in parallel programs
  - ▶ E.g., consider two threads atomically incrementing a counter
- We detect such **cutoff events** by hashing all memory

- If two paths yield the exact same state, only one needs to be pursued further
  - ▶ Happens all the time in parallel programs
  - ▶ E.g., consider two threads atomically incrementing a counter
- We detect such **cutoff events** by hashing all memory
- Also, our partial-order-reduction algorithm **re-executes some states**



- If two paths yield the exact same state, only one needs to be pursued further
  - ▶ Happens all the time in parallel programs
  - ▶ E.g., consider two threads atomically incrementing a counter
- We detect such **cutoff events** by hashing all memory
- Also, our partial-order-reduction algorithm **re-executes some states**
  - ▶ These re-executions must yield the exact same state as the original execution

- If two paths yield the exact same state, only one needs to be pursued further
  - ▶ Happens all the time in parallel programs
  - ▶ E.g., consider two threads atomically incrementing a counter
- We detect such **cutoff events** by hashing all memory
- Also, our partial-order-reduction algorithm **re-executes some states**
  - ▶ These re-executions must yield the exact same state as the original execution

**A deterministic engine is key for our approach!**

- Our KLEE fork can symbolically execute multi-threaded programs

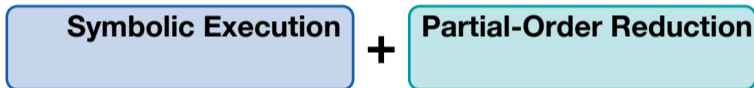
# Conclusion

- Our KLEE fork can symbolically execute multi-threaded programs
- Support for most POSIX threading primitives, including condition variables

- Our KLEE fork can symbolically execute multi-threaded programs
- Support for most POSIX threading primitives, including condition variables
- Relies on recognizing data races as bugs (and does so)

- Our KLEE fork can symbolically execute multi-threaded programs
- Support for most POSIX threading primitives, including condition variables
- Relies on recognizing data races as bugs (and does so)
- Requires deterministic re-execution

- Our KLEE fork can symbolically execute multi-threaded programs
- Support for most POSIX threading primitives, including condition variables
- Relies on recognizing data races as bugs (and does so)
- Requires deterministic re-execution



- Our KLEE fork can symbolically execute multi-threaded programs
- Support for most POSIX threading primitives, including condition variables
- Relies on recognizing data races as bugs (and does so)
- Requires deterministic re-execution

**Symbolic Execution**  
data non-determinism

+

**Partial-Order Reduction**  
thread non-determinism



# Conclusion

- Our KLEE fork can symbolically execute multi-threaded programs
- Support for most POSIX threading primitives, including condition variables
- Relies on recognizing data races as bugs (and does so)
- Requires deterministic re-execution

**Symbolic Execution**  
data non-determinism

+

**Partial-Order Reduction**  
thread non-determinism

Check out our tool:  <https://github.com/por-se/por-se> 

# Conclusion

- Our KLEE fork can symbolically execute multi-threaded programs
- Support for most POSIX threading primitives, including condition variables
- Relies on recognizing data races as bugs (and does so)
- Requires deterministic re-execution

**Symbolic Execution**  
data non-determinism

+

**Partial-Order Reduction**  
thread non-determinism

Check out our tool:  <https://github.com/por-se/por-se> 

Read the extended version of our CAV'20 paper [1]: <https://arxiv.org/abs/2005.06688>

[1] Symbolic Partial-Order Execution for Testing Multi-Threaded Programs.  
D. Schemmel, J. Büning, C. Rodríguez, D. Laprell, K. Wehrle. CAV 2020.



This research is supported by the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme (grant agreement №. 647295 (SYMBIOSYS)).



European Research Council

Established by the European Commission