

# Sound Dynamic Deadlock Prediction in Linear Time

---

Hünkar Can Tunç


Andreas Pavlogiannis

Umang Mathur


Mahesh Viswanathan




# The Problem




## The Problem



## The Problem



## The Problem



## The Problem


	$t_1$	$t_2$
1	acquire( $l_1$ )	
2		acquire( $l_2$ )
3	request( $l_2$ )	
4		request( $l_1$ )

## The Problem

	$t_1$	$t_2$
1	acquire( $l_1$ )	
2		acquire( $l_2$ )
3	request( $l_2$ )	
4		request( $l_1$ )

(Resource) Deadlock!


## Dynamic Analysis



- Effective method for finding concurrency bugs
- Widely adopted (e.g., ThreadSanitizer, Helgrind)



# Dynamic Analysis




- Effective method for finding concurrency bugs
- Widely adopted (e.g., ThreadSanitizer, Helgrind)
- Traditional techniques:
  - ↪ Analyze the current execution
- **Predictive** techniques:
  - ↪ Analyze the current execution + infer alternates

# Predictive Analysis

	$t_1$	$t_2$
1	acquire( $l_1$ )	
2	acquire( $l_2$ )	
3	release( $l_2$ )	
4	release( $l_1$ )	
5		acquire( $l_2$ )
6		acquire( $l_1$ )
7		release( $l_1$ )
8		release( $l_2$ )


Observed trace  
No deadlock

# Predictive Analysis



Observed trace  
No deadlock

reordering  
→



Reordered trace  
**Deadlock!**

- We study the problem of dynamic deadlock **prediction**
- **Main results:**
  - Complexity characterization
    - ↔ Tradeoff between efficiency and precision is unavoidable
  - Novel algorithms
    - ↔ Strike a good balance between efficiency and precision
  - Empirical evaluation
    - ↔ Outperform state-of-the-art techniques

## State-of-the-art

- **SeqCheck**<sup>1</sup>:
  - Sound but incomplete
  - High polynomial complexity
    - $\hookrightarrow \tilde{O}(\mathcal{N}^4)$
- **Dirk**<sup>2</sup>:
  - Sound and complete
  - Heavyweight SMT solving

---

<sup>1</sup>Yan Cai, Hao Yun, Jinqiu Wang, Lei Qiao, Jens Palsberg. Sound and efficient concurrency bug prediction. ESEC/FSE'21

<sup>2</sup>Christian Gram Kalhauge, Jens Palsberg. Sound deadlock prediction. OOPSLA'18

## State-of-the-art

- **SeqCheck**<sup>1</sup>:
  - Sound but incomplete
  - High polynomial complexity
    - $\hookrightarrow \tilde{O}(\mathcal{N}^4)$
- **Dirk**<sup>2</sup>:
  - Sound and complete
  - Heavyweight SMT solving

## This work

- **Sync-Preserving Deadlocks**:
  - Sound but incomplete
  - (Nearly) Linear time algorithm
    - $\hookrightarrow \tilde{O}(\mathcal{N})$
    - $\hookrightarrow$  Wrt. number of events

---

<sup>1</sup>Yan Cai, Hao Yun, Jinqiu Wang, Lei Qiao, Jens Palsberg. Sound and efficient concurrency bug prediction. ESEC/FSE'21

<sup>2</sup>Christian Gram Kalhauge, Jens Palsberg. Sound deadlock prediction. OOPSLA'18

# Dynamic Deadlock Prediction

## State-of-the-art

- **SeqCheck**<sup>1</sup>:
  - Sound but incomplete
  - High polynomial complexity
    - $\hookrightarrow \tilde{O}(\mathcal{N}^4)$
- **Dirk**<sup>2</sup>:
  - Sound and complete
  - Heavyweight SMT solving

## This work

- **Sync-Preserving Deadlocks**:
  - Sound but incomplete
  - (Nearly) Linear time algorithm
    - $\hookrightarrow \tilde{O}(\mathcal{N})$
    - $\hookrightarrow$  Wrt. number of events

Focus is on identifying real deadlocks!

<sup>1</sup>Yan Cai, Hao Yun, Jinqiu Wang, Lei Qiao, Jens Palsberg. Sound and efficient concurrency bug prediction. ESEC/FSE'21

<sup>2</sup>Christian Gram Kalhauge, Jens Palsberg. Sound deadlock prediction. OOPSLA'18

## Two steps of predictive analysis:

- ① Identify potential buggy events
- ② Check if the potential bug can be realized



## **Two steps of predictive analysis:**

- ① Identify potential buggy events
- ② Check if the potential bug can be realized

# Potential Deadlocks

## - Potential deadlocks:

- Cyclic lock acquisition patterns

↔  $l_1, l_2$

	$t_1$	$t_2$
1	...	
2	acquire( $l_1$ )	
3	acquire( $l_2$ )	
4	...	
5	release( $l_2$ )	
6	release( $l_1$ )	
7	...	
8		...
9		acquire( $l_2$ )
10		acquire( $l_1$ )
11		release( $l_1$ )
12		release( $l_2$ )
13		...

# Potential Deadlocks

## - Potential deadlocks:

- Cyclic lock acquisition patterns  
     $\hookrightarrow l_1, l_2$
- Not protected by a common lock  
     $\hookrightarrow$  No such  $l_3$

	$t_1$	$t_2$
1	acquire( $l_3$ )	
2	acquire( $l_1$ )	
3	acquire( $l_2$ )	
4	...	
5	release( $l_2$ )	
6	release( $l_1$ )	
7	release( $l_3$ )	
8		acquire( $l_3$ )
9		acquire( $l_2$ )
10		acquire( $l_1$ )
11		release( $l_1$ )
12		release( $l_2$ )
13		release( $l_3$ )

# Potential Deadlocks

## - Potential deadlocks:

- Cyclic lock acquisition patterns

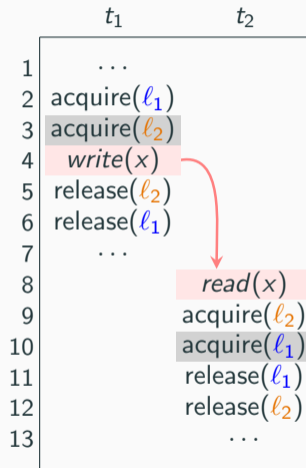
↔  $l_1, l_2$

- Not protected by a common lock

↔ No such  $l_3$

- Necessary but insufficient for an actual deadlock

↔ Control flow/data flow dependencies



- **Our first result:**
  - Identifying potential deadlocks is intractable  
↔ NP-hard

- **Our first result:**

- Identifying potential deadlocks is intractable  
     $\hookrightarrow$  NP-hard

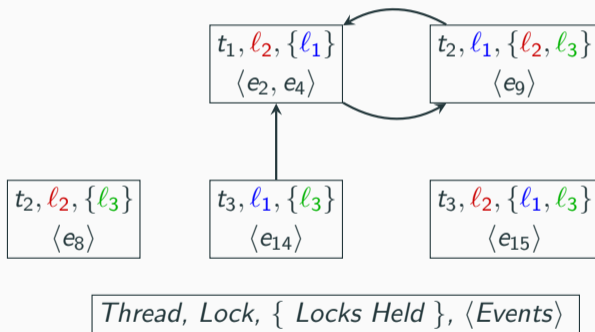
- **Our solution:**

- Abstraction that groups potential deadlocks  
     $\hookrightarrow$  Abstract lock graph

# Abstract Lock Graph

	$t_1$	$t_2$	$t_3$
1	acquire( $l_1$ )		
2	acquire( $l_2$ )		
3	release( $l_2$ )		
4	acquire( $l_2$ )		
5	release( $l_2$ )		
6	release( $l_1$ )		
7		acquire( $l_3$ )	
8		acquire( $l_2$ )	
9		acquire( $l_1$ )	
10		release( $l_1$ )	
11		release( $l_2$ )	
12		release( $l_3$ )	
13			acquire( $l_3$ )
14			acquire( $l_1$ )
15			acquire( $l_2$ )
16			release( $l_2$ )
17			release( $l_1$ )
18			release( $l_3$ )

Observed trace



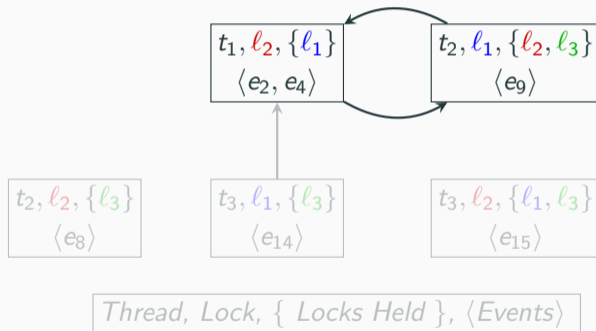
Abstract Lock Graph

# Abstract Lock Graph

	$t_1$	$t_2$	$t_3$
1	acquire( $l_1$ )		
2	acquire( $l_2$ )		
3	release( $l_2$ )		
4	acquire( $l_2$ )		
5	release( $l_2$ )		
6	release( $l_1$ )		
7		acquire( $l_3$ )	
8		acquire( $l_2$ )	
9		acquire( $l_1$ )	
10		release( $l_1$ )	
11		release( $l_2$ )	
12		release( $l_3$ )	
13			acquire( $l_3$ )
14			acquire( $l_1$ )
15			acquire( $l_2$ )
16			release( $l_2$ )
17			release( $l_1$ )
18			release( $l_3$ )

Observed trace

Potential deadlocks:  $\langle e_2, e_9 \rangle, \langle e_4, e_9 \rangle$



Abstract Lock Graph



## Two steps of predictive analysis:

- ① Identify potential buggy events ✓
- ② Check if the potential bug can be realized

# Predicting Deadlocks

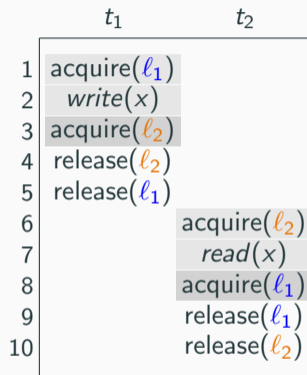
	$t_1$	$t_2$
1	acquire( $l_1$ )	
2	<i>write(x)</i>	
3	acquire( $l_2$ )	
4	release( $l_2$ )	
5	release( $l_1$ )	
6		acquire( $l_2$ )
7		<i>read(x)</i>
8		acquire( $l_1$ )
9		release( $l_1$ )
10		release( $l_2$ )

Observed trace

No deadlock


Potential deadlock → Real deadlock?

# Predicting Deadlocks



Observed trace  
No deadlock

reordering  $\rightarrow$



Witness trace  
**Deadlock!**

Potential deadlock  $\rightarrow$  Real deadlock  $\checkmark$

## Predicting Deadlocks

- **Our second result:** (Given a potential deadlock)
  - Sound and complete deadlock prediction is intractable
    - $\hookrightarrow$  NP-hard

## Predicting Deadlocks

- **Our second result:** (Given a potential deadlock)
  - Sound and complete deadlock prediction is intractable
    - ↔ NP-hard
  
- **General solution:**
  - Consider a restricted problem setting to gain efficiency
    - ↔ Look for a subset of deadlocks

## Predicting Deadlocks

- **Our second result:** (Given a potential deadlock)
  - Sound and complete deadlock prediction is intractable
    - ↪ NP-hard
  
- **General solution:**
  - Consider a restricted problem setting to gain efficiency
    - ↪ Look for a subset of deadlocks
  
- **Challenge:**
  - Restrictions should satisfy the following two properties
    - ↪ Enable efficient analysis
    - ↪ Retain high precision

# Sync-Preserving Deadlocks

- Adapted from data races<sup>1</sup>
- Subset of deadlocks
  - ↔ More conservative restrictions on the allowed reorderings
- Enables efficient analysis

---

<sup>1</sup>Umang Mathur, Andreas Pavlogiannis, Mahesh Viswanathan. Optimal Prediction of Synchronization-Preserving Races. POPL'21

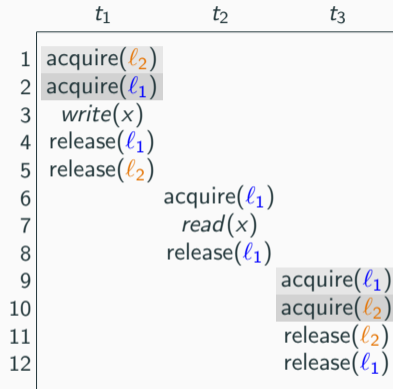
# Sync-Preserving Deadlocks

	$t_1$	$t_2$	$t_3$
1	acquire( $l_2$ )		
2	acquire( $l_1$ )		
3	<i>write(x)</i>		
4	release( $l_1$ )		
5	release( $l_2$ )		
6		acquire( $l_1$ )	
7		<i>read(x)</i>	
8		release( $l_1$ )	
9			acquire( $l_1$ )
10			acquire( $l_2$ )
11			release( $l_2$ )
12			release( $l_1$ )


Observed trace



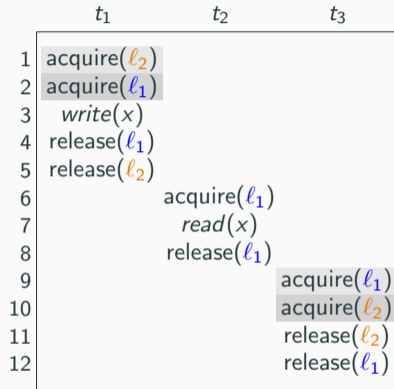
# Sync-Preserving Deadlocks




Observed trace



# Sync-Preserving Deadlocks



Observed trace




Sync-Preserving Deadlock

Order of acquire events on the same lock that occur in the witness are maintained

# Interplay With Abstract Lock Graph

	$t_1$	$t_2$	$t_3$
1	acquire( $l_1$ )		
2	acquire( $l_2$ )		
3	release( $l_2$ )		
4	acquire( $l_2$ )		
5	release( $l_2$ )		
6	release( $l_1$ )		
7		acquire( $l_3$ )	
8		acquire( $l_2$ )	
9		acquire( $l_1$ )	
10		release( $l_1$ )	
11		release( $l_2$ )	
12		release( $l_3$ )	
13			acquire( $l_3$ )
14			acquire( $l_1$ )
15			acquire( $l_2$ )
16			release( $l_2$ )
17			release( $l_1$ )
18			release( $l_3$ )




**Potential deadlocks:**

$\langle e_2, e_9 \rangle, \langle e_4, e_9 \rangle$

# Interplay With Abstract Lock Graph

	$t_1$	$t_2$	$t_3$
1	acquire( $l_1$ )		
2	acquire( $l_2$ )		
3	release( $l_2$ )		
4	acquire( $l_2$ )		
5	release( $l_2$ )		
6	release( $l_1$ )		
7		acquire( $l_3$ )	
8		acquire( $l_2$ )	
9		acquire( $l_1$ )	
10		release( $l_1$ )	
11		release( $l_2$ )	
12		release( $l_3$ )	
13			acquire( $l_3$ )
14			acquire( $l_1$ )
15			acquire( $l_2$ )
16			release( $l_2$ )
17			release( $l_1$ )
18			release( $l_3$ )



Potential deadlocks:

$\langle e_2, e_9 \rangle, \langle e_4, e_9 \rangle$


- Number of checks per cycle:

↪ **Naive approach:** Exponential

↪ **Sync-preserving deadlocks:** Linear

# Interplay With Abstract Lock Graph

	$t_1$	$t_2$	$t_3$
1	acquire( $l_1$ )		
2	acquire( $l_2$ )		
3	release( $l_2$ )		
4	acquire( $l_2$ )		
5	release( $l_2$ )		
6	release( $l_1$ )		
7		acquire( $l_3$ )	
8		acquire( $l_2$ )	
9		acquire( $l_1$ )	
10		release( $l_1$ )	
11		release( $l_2$ )	
12		release( $l_3$ )	
13			acquire( $l_3$ )
14			acquire( $l_1$ )
15			acquire( $l_2$ )
16			release( $l_2$ )
17			release( $l_1$ )
18			release( $l_3$ )




**Potential deadlocks:**

$\langle e_2, e_9 \rangle, \langle e_4, e_9 \rangle$

- Number of checks per cycle:
  - ↪ **Naive approach:** Exponential
  - ↪ **Sync-preserving deadlocks:** Linear
- Time spent per check:
  - ↪ **All deadlocks:** Exponential
  - ↪ **Sync-preserving deadlocks:** Linear

# Interplay With Abstract Lock Graph

	$t_1$	$t_2$	$t_3$
1	acquire( $l_1$ )		
2	acquire( $l_2$ )		
3	release( $l_2$ )		
4	acquire( $l_2$ )		
5	release( $l_2$ )		
6	release( $l_1$ )		
7		acquire( $l_3$ )	
8		acquire( $l_2$ )	
9		acquire( $l_1$ )	
10		release( $l_1$ )	
11		release( $l_2$ )	
12		release( $l_3$ )	
13			acquire( $l_3$ )
14			acquire( $l_1$ )
15			acquire( $l_2$ )
16			release( $l_2$ )
17			release( $l_1$ )
18			release( $l_3$ )



Potential deadlocks:

$\langle e_2, e_9 \rangle, \langle e_4, e_9 \rangle$

- Number of checks per cycle:

↪ **Naive approach:** Exponential

↪ **Sync-preserving deadlocks:** Linear


- Time spent per check:

↪ **All deadlocks:** Exponential

↪ **Sync-preserving deadlocks:** Linear

**Overall  
Linear**

# Interplay With Abstract Lock Graph



Potential deadlocks:

$\langle e_2, e_9 \rangle, \langle e_4, e_9 \rangle$

- Number of checks per cycle:

↪ **Naive approach:** Exponential

↪ **Sync-preserving deadlocks:** Linear

- Time spent per check:

↪ **All deadlocks:** Exponential

↪ **Sync-preserving deadlocks:** Linear

**Overall  
Linear**

How much precision have we lost?

## Experimental Results - Offline

- Implemented *Sync-preserving Offline*
  - ↪ Postmortem analysis
- Compared with *SeqCheck* and *Dirk*
- 48 benchmark traces
  - ↪ Based on standard Java benchmark suites



	<i>Dirk</i>	<i>SeqCheck</i>	<b>Sync-preserving Offline</b>
Total Deadlocks	35	40	<b>40</b>
Total Time	> 1000 <i>minutes</i>	46 <i>minutes</i>	<b>3 minutes</b>





- Implemented *Sync-preserving Offline*
  - ↪ Postmortem analysis
- Compared with *SeqCheck* and *Dirk*
- 48 benchmark traces
  - ↪ Based on standard Java benchmark suites

	<i>Dirk</i>	<i>SeqCheck</i>	<b>Sync-preserving Offline</b>
Total Deadlocks	35	40	<b>40</b>
Total Time	> 1000 <i>minutes</i>	46 <i>minutes</i>	<b>3 minutes</b>

- **False negative analysis:** Only one actual deadlock is missed!
  - ↪ Based on the standard notion of valid reorderings

- Online setting  $\rightarrow$  On-the-fly analysis
- No predictive online method
- Non-predictive online techniques:
  - $\hookrightarrow$  Schedule fuzzing
- **Our work:**
  - $\hookrightarrow$  Prediction + schedule fuzzing

- Implemented *Sync-preserving Online*
- Compared with *DeadlockFuzzer*
- 38 benchmarks
  - ↪ Based on standard Java benchmark suites



	<i>DeadlockFuzzer</i>	<b>Sync-preserving Online</b>
Total Deadlock Hits	2076	<b>7633</b>
Total Unique Deadlocks	42	<b>49</b>

## This work:

- **Complexity characterization:**
  - ↪ Finding potential deadlocks is intractable
  - ↪ Realizing potential deadlocks is intractable

## This work:

- **Complexity characterization:**
  - ↪ Finding potential deadlocks is intractable
  - ↪ Realizing potential deadlocks is intractable
- **Sync-preserving deadlocks:**
  - ↪ Achieves efficiency and high precision
  - ↪ Outperforms state-of-the-art

## This work:

- **Complexity characterization:**
  - ↪ Finding potential deadlocks is intractable
  - ↪ Realizing potential deadlocks is intractable
- **Sync-preserving deadlocks:**
  - ↪ Achieves efficiency and high precision
  - ↪ Outperforms state-of-the-art

Thank you!