The Complexity of **Dynamic Data Race Prediction**

Umang Mathur University of Illinois at Urbana-Champaign

Andreas Pavlogiannis Aarhus University

Mahesh Viswanathan

University of Illinois at Urbana-Champaign







1



• Concurrency - ubiquitous programming paradigm

1



Concurrency - ubiquitous programming paradigm

1





Concurrency and Challenges

- Concurrency ubiquitous programming paradigm
- Challenging to develop concurrent software
 - Large interleaving space

]





- Concurrency ubiquitous programming paradigm
- Challenging to develop concurrent software
 - Large interleaving space
- Concurrency bugs arise in production-level software
 - Despite rigorous testing







- Concurrency ubiquitous programming paradigm
- Challenging to develop concurrent software
 - Large interleaving space
- Concurrency bugs arise in production-level software
 - Despite rigorous testing
- **Data races** : most common source of concurrency issues







Static analyses



- Analyze source code
- Undecidable problem
- Excessive false alarms





- Analyze executions at runtime
- Typically sound
- Widely adopted TSan, Helgrind, etc.,



Concurrent Program Executions

- Sequences of events
- Event e = <t, op>
 - t is the thread that performs e
 - Op is an operation
- Operations:
 - Read/Write to memory locations
 - Acquire and release of locks
- Well formed-ness
 - At most one thread holds a lock at any time

	t1	t2
1	acq(l)	
2	w(x)	
3	w(y)	
4		r(x)
5	rel(l)	
6		w(y)



Data Race

An execution σ has a data race if

- pair of conflicting events
- consecutive

Data Race

An execution σ has a data race if

- pair of conflicting events
- consecutive

- I. Same memory location
- 2. Different threads
- 3. At least one write

Data Race

An execution σ has a data race if

- pair of conflicting events
- consecutive

- 2. Different threads
- 3. At least one write

t1	t2
acq(l)	
w(x)	
	r(x)
rel(l)	

Data Race

An execution σ has a data race if

- pair of conflicting events
- consecutive

Data Race Detection

- 2. Different threads
- 3. At least one write

t1	t2
acq(l)	
w(x)	
	r(x)
rel(l)	



An execution σ has a data race if

- pair of conflicting events
- consecutive

Data Race Detection Given an execution σ , does σ have a data race?

- Sensitive to thread scheduling
- Miss a lot of races



- 2. Different threads
- 3. At least one write

t1	t2
acq(l)	
w(x)	
	r(x)
rel(l)	



An execution σ has a data race if

- pair of conflicting events
- consecutive

Data Race Detection Given an execution σ , does σ have a data race?

- Sensitive to thread scheduling
- Miss a lot of races



- 2. Different threads
- 3. At least one write

t1	t2
acq(l)	
w(x)	
	r(x)
rel(l)	

t1	t2
acq(l)	
w(x)	
rel(l)	
	r(x)



An execution σ has a data race if

- pair of conflicting events
- consecutive

Data Race Detection Given an execution σ , does σ have a data race?

- Sensitive to thread scheduling
- Miss a lot of races



- Different threads 2.
- 3. At least one write

t1	t2
acq(l)	
w(x)	
	r(x)
rel(l)	

t1	t2
acq(l)	
w(x)	
rel(l)	
	r(x) -





An execution σ has a data race if

- pair of conflicting events
- consecutive

Data Race Detection Given an execution σ , does σ have a data race? • Sensitive to thread scheduling

- Miss a lot of races



- Different threads 2.
- 3. At least one write





An execution σ has a data race if

- pair of conflicting events
- consecutive

Data Race Detection

- Given an execution σ , does σ have a data race? • Sensitive to thread scheduling
 - Miss a lot of races

Given an execution σ , is there a **reordering*** ρ of σ , such that ρ has a data race?



- Different threads 2.
- 3. At least one write





An execution σ has a data race if

- pair of conflicting events
- consecutive

Data Race Detection

- Given an execution σ , does σ have a data race? • Sensitive to thread scheduling
 - Miss a lot of races

Data Race Prediction

Given an execution σ , is there a **reordering*** ρ of σ , such that ρ has a data race?



- Different threads 2.
- 3. At least one write



Any program that generates the observed execution must also generate the reordering

Any program that generates the observed execution must also generate the reordering

Any program that generates the observed execution must also generate the reordering

Reorderings must satisfy some properties -

I. Well formed -

Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap

Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap

t1	t2
acq(l)	
w(x)	
rel(l)	
	acq(l)
	rel(l)

Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap



Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap



Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap



Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap
- 2. Preserve intra-thread ordering



Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap
- 2. Preserve intra-thread ordering



t1	t2
acq(l)	
w(x)	
rel(l)	
	r(x)

Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap
- 2. Preserve intra-thread ordering





Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap
- 2. Preserve intra-thread ordering
- 3. Preserve control flow





Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap
- 2. Preserve intra-thread ordering
- 3. Preserve control flow
 - Every read sees its original write





Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap
- 2. Preserve intra-thread ordering
- 3. Preserve control flow
 - Every read sees its original write







Any program that generates the observed execution must also generate the reordering

- I. Well formed -
 - critical sections on same lock don't overlap
- 2. Preserve intra-thread ordering
- 3. Preserve control flow
 - Every read sees its original write






Which reorderings are allowed?

Any program that generates the observed execution must also generate the reordering

Correct reordering*

Reorderings must satisfy some properties -

- I. Well formed -
 - critical sections on same lock don't overlap
- 2. Preserve intra-thread ordering
- 3. Preserve control flow
 - Every read sees its original write

* Herlihy and Wing, Linearizability: A Correctness Condition for Concurrent Objects, TOPLAS 1990 Smaragdakis et al, Sound predictive race detection in polynomial time, POPL 2012





Data Race Prediction



Data Race Prediction



Data Race Prediction





Data Race Prediction



















Given an execution σ , is there a **correct reordering** of σ with a data race?

2011

Explicit or Symbolic _ Search























Data Race Prediction

Data Race Prediction Input: Trace σ and events e₁ and e₂
[n events, k threads, d memory locations and locks]

Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

[n events, k threads, d memory locations and locks]



Data Race Prediction

• **Input:** Trace σ and events e_1 and e_2 [n events, k threads, d memory locations and locks]

(Easy) Upper Bounds



Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

(Easy) Upper Bounds

NP

• Guess an alternate reordering and check if it is a correct reordering

[n events, k threads, d memory locations and locks]



Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

(Easy) Upper Bounds

- NP
 - Guess an alternate reordering and check if it is a correct reordering
- 2. $O(k^n)$ Enumeration based techniques:
 - At every step, choose which thread to execute

[n events, k threads, d memory locations and locks]



Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

(Easy) Upper Bounds

- NP
 - Guess an alternate reordering and check if it is a correct reordering
- 2. $O(k^n)$ Enumeration based techniques:
 - At every step, choose which thread to execute
- 3. O(SAT(poly(n)) SAT solving based techniques

[n events, k threads, d memory locations and locks]



Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

(Easy) Upper Bounds

- NP
 - Guess an alternate reordering and check if it is a correct reordering
- 2. $O(k^n)$ Enumeration based techniques:
 - At every step, choose which thread to execute
- 3. O(SAT(poly(n)) SAT solving based techniques

[n events, k threads, d memory locations and locks]

• **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Lower Bound



Data Race Prediction

- Input: Trace σ and events e_1 and e_2

(Easy) Upper Bounds

- NP
 - Guess an alternate reordering and check if it is a correct reordering
- 2. $O(k^n)$ Enumeration based techniques:
 - At every step, choose which thread to execute
- 3. O(SAT(poly(n)) SAT solving based techniques

[n events, k threads, d memory locations and locks]

• **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Lower Bound

- Unknown!
- Is it **NP**-hard? Is enumeration unavoidable?
- Is it polynomial time?



Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

[n events, k threads, d memory locations and locks]

• **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Closely Related Work



Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

- I. [Netzer and Miller, '89-'92]: **NP**-hardness of data race prediction.
 - Too strong notion of correct reordering
 - Hardness comes from more powerful synchronization primitives

[n events, k threads, d memory locations and locks]

• **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Closely Related Work



Data Race Prediction • **Input:** Trace σ and events e_1 and e_2 [n events, k threads, d memory locations and locks]

Closely Related Work

- I. [Netzer and Miller, '89-'92]: **NP**-hardness of data race prediction.
 - Too strong notion of correct reordering
 - Hardness comes from more powerful synchronization primitives
- 2. [Gibbons and Korach, '97]: **NP**-hardness of Verifying Sequential Consistency
 - Different problem than data race detection



Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

[n events, k threads, d memory locations and locks]



Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

Extensive study of complexity theoretic questions in data race prediction[†]

[n events, k threads, d memory locations and locks]

Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

General Case

[n events, k threads, d memory locations and locks]

• **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Extensive study of complexity theoretic questions in data race prediction[†]

Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

Extensive study of complexity theoretic questions in data race prediction[†]

General Case

I. <u>Poly-time Upper bound (when k is constant)</u> Algorithm for race prediction **O(kn**^{2(k-1)})

[n events, k threads, d memory locations and locks]
Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2

Extensive study of complexity theoretic questions in data race prediction[†]

General Case

- I. <u>Poly-time Upper bound (when k is constant)</u> Algorithm for race prediction $O(kn^{2(k-1)})$
- Lower bound :W[I] hard in parameter k
 - **NP**-hard when k is not constant
 - Not even FPT in k

[n events, k threads, d memory locations and locks]

• **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Data Race Prediction

- **Input:** Trace σ and events e_1 and e_2 [n events, k threads, d memory locations and locks]
- **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Extensive study of complexity theoretic questions in data race prediction[†]

General Case

- I. <u>Poly-time Upper bound (when k is constant)</u> Algorithm for race prediction $O(kn^{2(k-1)})$
- Lower bound :W[I] hard in parameter k
 - **NP**-hard when k is not constant
 - Not even FPT in k

Special Cases

Data Race Prediction

- Input: Trace σ and events e_1 and e_2 [n events, k threads, d memory locations and locks]
- **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Extensive study of complexity theoretic questions in data race prediction[†]

General Case

- I. <u>Poly-time Upper bound (when k is constant)</u> Algorithm for race prediction $O(kn^{2(k-1)})$
- Lower bound :W[I] hard in parameter k
 - **NP**-hard when k is not constant
 - Not even FPT in k

Special Cases

- <u>Restricting the space of input traces</u> 3.
 - **O(n²)** time algorithm
 - Matching (conditional) lower bound

Data Race Prediction

- Input: Trace σ and events e_1 and e_2 [n events, k threads, d memory locations and locks]
- **Output:** YES iff there is a correct reordering of σ that exhibits data race (e₁, e₂).

Extensive study of complexity theoretic questions in data race prediction[†]

General Case

- Poly-time Upper bound (when k is constant) Algorithm for race prediction $O(kn^{2(k-1)})$
- Lower bound :W[I] hard in parameter k
 - **NP**-hard when k is not constant
 - Not even FPT in k

Special Cases

- <u>Restricting the space of input traces</u> 3.
 - **O(n²)** time algorithm
 - Matching (conditional) lower bound
- 4. <u>Restricting the space of data races to be reported</u>
 - Linear time algorithm





- I. Trace Ideals and Realizability
- 2. Algorithm for Data Race Prediction (General case)
- **3.** Data Race Prediction for Acyclic Communication Topology
- 4. Distance Bounded Data Race Prediction
- 5. Lower Bound (General Case)
- 6. Lower Bound for 2 threads
- 7. Conclusions and Future Work



- 2. Algorithm for Data Race Prediction (General case)
- 3. Data Race Prediction for Acyclic Communication Topology
- 4. Distance Bounded Data Race Prediction
- 5. Lower Bound (General Case)
- 6. Lower Bound for 2 threads
- 7. Conclusions and Future Work

Outline

10

Some Notations

-		
	t1	t2
1	acq(l)	
2	rel(l)	
3	w(x)	
4		r(x)
5		acq(l)
6		rel(l)

Let σ be an execution trace.

-		
	t1	t2
1	acq(l)	
2	rel(l)	
3	w(x)	
4		r(x)
5		acq(l)
6		rel(l)

Let σ be an execution trace.

• TO_{σ} = intra-thread ordering on events of σ



Let σ be an execution trace.

- TO_{σ} = intra-thread ordering on events of σ
- RF_{σ} : reads-from mapping of σ :
 - For a read event $\mathbf{r}, \mathbf{RF}_{\sigma}(\mathbf{r}) = \mathbf{w}$
 - For a release event $rel, RF_{\sigma}(rel) = acq$





Let σ be an execution trace.

- TO_{σ} = intra-thread ordering on events of σ
- RF_{σ} : reads-from mapping of σ :
 - For a read event $\mathbf{r}, \mathbf{RF}_{\sigma}(\mathbf{r}) = \mathbf{w}$
 - For a release event $rel, RF_{\sigma}(rel) = acq$
- TRF_{σ} = smallest partial order that includes TO_{σ} and respects RF_{σ}





Let σ be an execution trace.

- TO_{σ} = intra-thread ordering on events of σ
- RF_{σ} : reads-from mapping of σ :
 - For a read event $\mathbf{r}, \mathbf{RF}_{\sigma}(\mathbf{r}) = \mathbf{w}$
 - For a release event $rel, RF_{\sigma}(rel) = acq$
- TRF_{σ} = smallest partial order that includes TO_{σ} and respects RF_{σ}

Let ρ be a sequence of events with Events(ρ) \subseteq Events(σ). ρ is a correct reordering of an execution trace σ if (i) (ii) ρ respects TO_{σ} (iii) $\mathsf{RF}_{\rho} \subseteq \mathsf{RF}_{\sigma}$



For every lock ℓ , there is at most one unmatched acquire of ℓ in ρ

A feasible trace ideal I of σ is a set of events such that

11

A feasible trace ideal I of σ is a set of events such that

• I is downward-closed under TRF_{σ} , and

- I is downward-closed under TRF_{σ} , and
- for every lock ℓ , there is at most one unmatched acquire of ℓ in I

- I is downward-closed under TRF_{σ} , and
- for every lock ℓ , there is at most one unmatched acquire of ℓ in I





- I is downward-closed under TRF_{σ} , and
- for every lock ℓ , there is at most one unmatched acquire of ℓ in I





- I is downward-closed under TRF_{σ} , and
- for every lock ℓ , there is at most one unmatched acquire of ℓ in I





A <u>feasible trace ideal</u> I of σ is a set of events such that

- I is downward-closed under TRF_{σ} , and
- for every lock ℓ , there is at most one unmatched acquire of ℓ in I

The <u>canonical rf-poset</u> of ideal I is the smallest partial order $\mathcal{P}(I)$ such that





A feasible trace ideal I of σ is a set of events such that

- I is downward-closed under TRF_{σ} , and
- for every lock ℓ , there is at most one unmatched acquire of ℓ in I

The <u>canonical rf-poset</u> of ideal I is the smallest partial order $\mathcal{P}(I)$ such that

• $\mathsf{TRF}_{\sigma} \downarrow_{\mathrm{I}} \subseteq \mathscr{P}(\mathrm{I})$



A feasible trace ideal I of σ is a set of events such that

- I is downward-closed under TRF_{σ} , and
- for every lock ℓ , there is at most one unmatched acquire of ℓ in I

The <u>canonical rf-poset</u> of ideal I is the smallest partial order $\mathcal{P}(I)$ such that

- $\mathsf{TRF}_{\sigma} \downarrow_{\mathrm{I}} \subseteq \mathscr{P}(\mathrm{I})$
- For every lock ℓ with an unmatched acquire **acq**_{unmatched} \in I, and for every ℓ -release event $rel \in I$, we have $rel \leq \mathcal{P}(I) acq_{unmatched}$



A feasible trace ideal I of σ is a set of events such that

- I is downward-closed under TRF_{σ} , and
- for every lock ℓ , there is at most one unmatched acquire of ℓ in I

The <u>canonical rf-poset</u> of ideal I is the smallest partial order $\mathcal{P}(I)$ such that

- $\mathsf{TRF}_{\sigma} \downarrow_{\mathrm{I}} \subseteq \mathscr{P}(\mathrm{I})$
- For every lock ℓ with an unmatched acquire **acq**_{unmatched} \in I, and for every ℓ -release event $rel \in I$, we have $rel \leq \mathcal{P}(I) acq_{unmatched}$

Realizability of Trace Ideals

Given a feasible trace ideal I of σ , check if there is a linearization σ^* of the canonical rf-poset $\mathscr{P}(\mathbf{I})$ such that $\mathsf{RF}_{\sigma^*} \subseteq \mathsf{RF}_{\sigma}$



2. Algorithm for Data Race Prediction (General case)

- 3. Data Race Prediction for Acyclic Communication Topology
- 4. Distance Bounded Data Race Prediction
- 5. Lower Bound (General Case)
- 6. Lower Bound for 2 threads
- 7. Conclusions and Future Work



Proposition

- A pair of conflicting events (e_1 , e_2) is a predictable data race of σ iff
- there exists a feasible trace ideal I of σ such that I is **realizable**
 - and both e_1 and e_2 are **enabled** in I.

Proposition

- A pair of conflicting events (e_1 , e_2) is a predictable data race of σ iff
- there exists a feasible trace ideal I of σ such that I is **realizable**
 - and both e_1 and e_2 are **enabled** in I.

If σ^* is a witness to the realizability of a feasible trace ideal I of σ , then σ^* is a correct reordering of σ



Proposition

- A pair of conflicting events (e_1 , e_2) is a predictable data race of σ iff
- there exists a feasible trace ideal I of σ such that I is **realizable**
 - and both e_1 and e_2 are **enabled** in I.

If σ^* is a witness to the realizability of a feasible trace ideal I of σ , then σ^* is a correct reordering of σ

Event e of σ is <u>enabled</u> in a feasible trace ideal I of σ if $I' = I \cup \{e\}$ is closed under ΤΟσ







Proposition

- A pair of conflicting events (e_1 , e_2) is a predictable data race of σ iff
- there exists a feasible trace ideal I of σ such that I is **realizable**
 - and both e_1 and e_2 are **enabled** in I.

If σ^* is a witness to the realizability of a feasible trace ideal I of σ , then σ^* is a correct reordering of σ

Event e of σ is <u>enabled</u> in a feasible trace ideal I of σ if $I' = I \cup \{e\}$ is closed under ΤΟσ







Proposition

- A pair of conflicting events (e_1 , e_2) is a predictable data race of σ iff
- there exists a feasible trace ideal I of σ such that I is **realizable**
 - and both e_1 and e_2 are **enabled** in I.

If σ^* is a witness to the realizability of a feasible trace ideal I of σ , then σ^* is a correct reordering of σ

Event e of σ is <u>enabled</u> in a feasible trace ideal I of σ if $I' = I \cup \{e\}$ is closed under TOσ














Lemma

Let σ be an execution with n events, k threads, lock nesting depth γ and lock-dependence factor ζ . Let e_1 and e_2 be events of σ . There are $O(\min(n, k \cdot \gamma \cdot \zeta)^{k-2})$ feasible ideals of σ in which both e_1 and e_2 are enabled.

......................

Lemma

Measures data flow between critical sections

Let σ be an execution with n events, k threads, lock nesting depth γ and lock-dependence factor ζ . Let e_1 and e_2 be events of σ .

There are $O(\min(n, k \cdot \gamma \cdot \zeta)^{k-2})$ feasible ideals of σ in which both e_1 and e_2 are enabled.

..................

Lemma

Measures data flow between critical sections

.

Let σ be an execution with n events, k threads, lock nesting depth γ and lock-dependence factor ζ . Let e_1 and e_2 be events of σ .

There are O(min(n, $\mathbf{k} \cdot \boldsymbol{\gamma} \cdot \boldsymbol{\zeta})^{k-2}$) feasible ideals of $\boldsymbol{\sigma}$ in which both e_1 and e_2 are enabled. Lemma

.

Let σ be an execution with n events and k threads and let I be a feasible ideal of σ.
The realizability of I can be checked in time O(k·n^k).

Lemma

Measures data flow between critical sections

Let σ be an execution with n events, k threads, lock nesting depth γ and lock-dependence factor ζ . Let e_1 and e_2 be events of σ .

There are O(min(n, $\mathbf{k} \cdot \gamma \cdot \zeta)^{k-2}$) feasible ideals of σ in which both e_1 and e_2 are enabled.

Lemma

Let σ be an execution with n events and k threads and let I be a feasible ideal of σ . The realizability of I can be checked in time $O(k \cdot n^k)$.

Theorem

Let σ be an execution with n events, k threads, lock nesting depth γ and lock-dependence factor ζ . Let (e₁, e₂) be a conflicting pair of events of σ . The dynamic race prediction problem on can be solved in time O(min(n, $\mathbf{k} \cdot \gamma \cdot \zeta)^{k-2} \cdot \mathbf{k} \cdot \mathbf{n}^{k}$).

- I. Trace Ideals and Realizability
- 2. Algorithm for Data Race Prediction (General case)

3. Data Race Prediction for Acyclic Communication Topology

- 4. Distance Bounded Data Race Prediction
- 5. Lower Bound (General Case)
- 6. Lower Bound for 2 threads
- 7. Conclusions and Future Work



15

Tree Communication Topology

Tree Communication Topology

The <u>communication topology</u> of σ is an undirected graph $G_{\sigma} = (V_{\sigma}, E_{\sigma})$:

- V_{σ} = set of threads in σ
- E_o = {(t, t') | t ≠ t' are threads that perform conflicting accesses or acquire a common lock}

Tree Communication Topology

2

3

4

5

6

8

9

The <u>communication topology</u> of σ is an undirected graph $G_{\sigma} = (V_{\sigma}, E_{\sigma})$:

- V_{σ} = set of threads in σ
- E_o = {(t, t') | t ≠ t' are threads that perform conflicting accesses or acquire a common lock}

t1	t2	t3	t4
	acq(l)		
	w(x)		
	rel(l)		
		r(x)	
acq(l)			
rel(l)			
w(y)			
			r(x)
		w(y)	







Pipeline



Divide-and-conquer





Server-client



Pipeline

Lemma

The realizability of a feasible ideal I of an execution σ (with n events, k threads and d variables) having an acyclic communication topology can be determined in $O(k^2 \cdot d^2 \cdot n^2 \cdot \log n)$ time.



Divide-and-conquer





Server-client



Pipeline

Lemma

The realizability of a feasible ideal I of an execution σ (with n events, k threads and d variables) having an acyclic communication topology can be determined in $O(k^2 \cdot d^2 \cdot n^2 \cdot \log n)$ time.

















- I. Trace Ideals and Realizability
- 2. Algorithm for Data Race Prediction (General case)
- **3.** Data Race Prediction for Acyclic Communication Topology

4. Distance Bounded Data Race Prediction

- 5. Lower Bound (General Case)
- 6. Lower Bound for 2 threads
- 7. Conclusions and Future Work



<u>Set of reversals</u> between σ and its correct reordering ρ :

- <u>Set of reversals</u> between σ and its correct reordering ρ :
 - **Rev**(σ , ρ) = {(e₁, e₂) | e₁ and e₂ are conflicting writes
 - or acquires of same lock such that
 - $e_1 \leq \sigma e_2 \text{ and } e_2 \leq \rho e_1$

- Set of reversals between σ and its correct reordering ρ : $Rev(\sigma, \rho) = \{(e_1, e_2) \mid e_1 \text{ and } e_2 \text{ are conflicting writes}$ or acquires of same lock such that
 - $e_1 \leq \sigma e_2 \text{ and } e_2 \leq \rho e_1$
- <u>Distance</u> between σ and ρ is $\delta(\sigma, \rho) = |\text{Rev}(\sigma, \rho)|$

- <u>Set of reversals</u> between σ and its correct reordering ρ : Rev $(\sigma, \rho) = \{(e_1, e_2) \mid e_1 \text{ and } e_2 \text{ are conflicting writes}\}$ or acquires of same lock such that
 - $e_1 \leq \sigma e_2 \text{ and } e_2 \leq \rho e_1$
- <u>Distance</u> between σ and ρ is $\delta(\sigma, \rho) = |\text{Rev}(\sigma, \rho)|$

	t1	t2
1	acq(l)	
2	rel(l)	
3	w(x)	
4		w(x)
5		acq(l)
6		rel(l)

	t1	t
1		w (
2		acq
3		rel
4	acq(l)	
5	rel(l)	
6	w(x)	



- <u>Set of reversals</u> between σ and its correct reordering ρ : Rev $(\sigma, \rho) = \{(e_1, e_2) \mid e_1 \text{ and } e_2 \text{ are conflicting writes}\}$ or acquires of same lock such that
 - $e_1 \leq \sigma e_2 \text{ and } e_2 \leq \rho e_1$
- <u>Distance</u> between σ and ρ is $\delta(\sigma, \rho) = |\text{Rev}(\sigma, \rho)|$



<u>Set of reversals</u> between σ and its correct reordering ρ : Rev $(\sigma, \rho) = \{(e_1, e_2) \mid e_1 \text{ and } e_2 \text{ are conflicting writes}\}$ or acquires of same lock such that

 $e_1 \leq \sigma e_2 \text{ and } e_2 \leq \rho e_1$

<u>Distance</u> between σ and ρ is $\delta(\sigma, \rho) = |\text{Rev}(\sigma, \rho)|$

Distance-bounded

Race Prediction



<u>Set of reversals</u> between σ and its correct reordering ρ : Rev $(\sigma, \rho) = \{(e_1, e_2) \mid e_1 \text{ and } e_2 \text{ are conflicting writes}\}$ or acquires of same lock such that $e_1 \leq \sigma e_2$ and $e_2 \leq \rho e_1$

<u>Distance</u> between σ and ρ is $\delta(\sigma, \rho) = |\text{Rev}(\sigma, \rho)|$

Distance-bounded

Race Prediction



- Let $\ell \in \mathbb{N}$ be a constant. Given an execution σ and pair of conflicting events
- (e₁, e₂) of σ , the answer to the ℓ -bounded data race prediction problem is

<u>Set of reversals</u> between σ and its correct reordering ρ : Rev $(\sigma, \rho) = \{(e_1, e_2) \mid e_1 \text{ and } e_2 \text{ are conflicting writes}\}$ or acquires of same lock such that $e_1 \leq \sigma e_2 \text{ and } e_2 \leq \rho e_1$

<u>Distance</u> between σ and ρ is $\delta(\sigma, \rho) = |\text{Rev}(\sigma, \rho)|$





Let $\ell \in \mathbb{N}$ be a constant. Given an execution σ and pair of conflicting events (e₁, e₂) of σ , the answer to the ℓ -bounded data race prediction problem is

<u>Set of reversals</u> between σ and its correct reordering ρ : $Rev(\sigma, \rho) = \{(e_1, e_2) \mid e_1 \text{ and } e_2 \text{ are conflicting writes}\}$ or acquires of same lock such that $e_1 \leq \sigma e_2$ and $e_2 \leq \rho e_1$

<u>Distance</u> between σ and ρ is $\delta(\sigma, \rho) = |\text{Rev}(\sigma, \rho)|$





Let $\ell \in \mathbb{N}$ be a constant. Given an execution σ and pair of conflicting events (e₁, e₂) of σ , the answer to the ℓ -bounded data race prediction problem is • YES, if there is a correct reordering σ^* of σ witnessing the race (e₁, e₂)

<u>Set of reversals</u> between σ and its correct reordering ρ : $Rev(\sigma, \rho) = \{(e_1, e_2) \mid e_1 \text{ and } e_2 \text{ are conflicting writes}\}$ or acquires of same lock such that $e_1 \leq \sigma e_2$ and $e_2 \leq \rho e_1$

<u>Distance</u> between σ and ρ is $\delta(\sigma, \rho) = |\text{Rev}(\sigma, \rho)|$





Let $\ell \in \mathbb{N}$ be a constant. Given an execution σ and pair of conflicting events (e₁, e₂) of σ , the answer to the ℓ -bounded data race prediction problem is • YES, if there is a correct reordering σ^* of σ witnessing the race (e₁, e₂)

<u>answer</u> to the ℓ -bounded realizability problem for I is • NO if I is not realizable Distance-bounded Ideal Realizability • YES or NO, otherwise

- Let $\ell \in \mathbb{N}$ be a constant. Given an execution σ and a feasible ideal I of σ , the

 - YES, if there is a witness σ^* of realizability of I such that $\delta(\sigma, \sigma^*) \leq \ell$

<u>answer</u> to the ℓ -bounded realizability problem for I is • NO if I is not realizable Distance-bounded Ideal Realizability • YES or NO, otherwise

- Let $\ell \in \mathbb{N}$ be a constant. Given an execution σ and a feasible ideal I of σ , the

 - YES, if there is a witness σ^* of realizability of I such that $\delta(\sigma, \sigma^*) \leq \ell$



Let $\ell \in \mathbb{N}$ be a constant. Let σ be an execution with n events and k threads and let I be a feasible ideal of σ . The ℓ -bounded realizability of I can be checked in time $O(k\ell^{+O(1)} \cdot n)$.

Lemma









Time complexity of ℓ -bounded data race prediction $= O(k\ell^{+O(1)} \cdot n)$



Outline

- I. Trace Ideals and Realizability
- 2. Algorithm for Data Race Prediction (General case)
- 3. Data Race Prediction for Acyclic Communication Topology
- 4. Distance Bounded Data Race Prediction

5. Lower Bound (General Case)

- 6. Lower Bound for 2 threads
- 7. Conclusions and Future Work
• When k (number of threads) is constant, data race prediction is in **P**.

- When k (number of threads) is constant, data race prediction is in \mathbf{P} .
- Is the problem in **P** even otherwise?

- When k (number of threads) is constant, data race prediction is in \mathbf{P} .
- Is the problem in **P** even otherwise?
- Is the problem in **FPT** with k as a parameter?

- When k (number of threads) is constant, data race prediction is in \mathbf{P} .
- Is the problem in **P** even otherwise?
- Is the problem in **FPT** with k as a parameter?

Can be solved in time $O(f(\mathbf{k}) \cdot \mathbf{n}^c)$?

- When k (number of threads) is constant, data race prediction is in \mathbf{P} .
- Is the problem in **P** even otherwise?
- Is the problem in **FPT** with k as a parameter?

Can be solved in time $O(f(\mathbf{k}) \cdot \mathbf{n}^c)$?

Theorem

Data Race Prediction and ideal

realizability are **W[1]**-hard in the

parameter k (number of threads).

- When k (number of threads) is constant, data race prediction is in \mathbf{P} .
- Is the problem in **P** even otherwise?
- Is the problem in **FPT** with k as a parameter?

Can be solved in time $O(f(\mathbf{k}) \cdot \mathbf{n}^c)$?

Theorem

Data Race Prediction and ideal

realizability are **W[1]**-hard in the

parameter k (number of threads).

- **FPT**-reduction from INDEPENDENT SET(c)
- Reduction takes time O(poly(n+k)) time ${ \bullet }$
- **NP**-hardness follows
- Improves previous result [Gibbons and Korach, '97]



- When k (number of threads) is constant, data race prediction is in \mathbf{P} .
- Is the problem in **P** even otherwise?
- Is the problem in **FPT** with k as a parameter?

Can be solved in time $O(f(k) \cdot n^c)$?

Theorem

Data Race Prediction and ideal

realizability are **W[1]**-hard in the

parameter k (number of threads).



- **FPT**-reduction from INDEPENDENT-SET(c)
- Reduction takes time O(poly(n+k)) time
- **NP**-hardness follows
- Improves previous result [Gibbons and Korach, '97]



Outline

- I. Trace Ideals and Realizability
- 2. Algorithm for Data Race Prediction (General case)
- 3. Data Race Prediction for Acyclic Communication Topology
- 4. Distance Bounded Data Race Prediction
- 5. Lower Bound (General Case)

6. Lower Bound for 2 threads

7. Conclusions and Future Work

Theorem

- Let σ be an execution with n events, $k \ge 2$ threads and $d \ge 9$ variables (and ≥ 1 lock). There is no algorithm that solves the data race prediction problem for in time
 - $O(n^{2-\epsilon})$, for any $\epsilon > 0$, unless the Orthogonal Vectors conjecture fails.



Theorem

- Let σ be an execution with n events, $k \ge 2$ threads and $d \ge 9$ variables (and ≥ 1 lock). There is no algorithm that solves the data race prediction problem for in time
 - $O(n^{2-\epsilon})$, for any $\epsilon > 0$, unless the Orthogonal Vectors conjecture fails.



Theorem

- Let σ be an execution with n events, $k \ge 2$ threads and $d \ge 9$ variables (and ≥ 1 lock).
 - There is no algorithm that solves the data race prediction problem for in time
 - $O(n^{2-\epsilon})$, for any $\epsilon > 0$, unless the Orthogonal Vectors conjecture fails.

SETH implies Orthogonal Vectors Conjecture



Theorem

- Let σ be an execution with n events, $k \ge 2$ threads and $d \ge 9$ variables (and ≥ 1 lock).
 - There is no algorithm that solves the data race prediction problem for in time
 - $O(n^{2-\epsilon})$, for any $\epsilon > 0$, unless the Orthogonal Vectors conjecture fails.

SETH implies Orthogonal Vectors Conjecture

Applies to acyclic topology

- I. Trace Ideals and Realizability
- 2. Algorithm for Data Race Prediction (General case)
- 3. Data Race Prediction for Acyclic Communication Topology
- 4. Distance Bounded Data Race Prediction
- 5. Lower Bound (General Case)
- 6. Lower Bound for 2 threads

7. Conclusions and Future Work







- Dynamic Race Prediction
 - Complexity-theoretic investigation

- Dynamic Race Prediction
 - Complexity-theoretic investigation
- Trace Ideals and Realizability

- Dynamic Race Prediction
 - Complexity-theoretic investigation
- Trace Ideals and Realizability
- General case:
 - O(kn^{2k-2}) time algorithm
 - $O(k \cdot n^k)$ when γ and ζ are constant
 - **W[**]]-hardness
 - Quadratic lower bound for k=2 threads

- Dynamic Race Prediction
 - Complexity-theoretic investigation
- Trace Ideals and Realizability
- General case:
 - O(kn^{2k-2}) time algorithm
 - $O(k \cdot n^k)$ when γ and ζ are constant
 - **W[**]]-hardness
 - Quadratic lower bound for k=2 threads
- Acyclic topology
 - Quadratic time upper and lower bounds

- Dynamic Race Prediction
 - Complexity-theoretic investigation
- Trace Ideals and Realizability
- General case:
 - O(kn^{2k-2}) time algorithm
 - $O(k \cdot n^k)$ when γ and ζ are constant
 - **W[**]]-hardness
 - Quadratic lower bound for k=2 threads
- Acyclic topology
 - Quadratic time upper and lower bounds
- Distance Bounded data race prediction
 - Linear time realizability

Summary

- Dynamic Race Prediction
 - Complexity-theoretic investigation
- Trace Ideals and Realizability
- General case:
 - O(kn^{2k-2}) time algorithm
 - $O(k \cdot n^k)$ when γ and ζ are constant
 - **W[**]]-hardness
 - Quadratic lower bound for k=2 threads
- Acyclic topology
 - Quadratic time upper and lower bounds
- Distance Bounded data race prediction
 - Linear time realizability

Summary

- Dynamic Race Prediction
 - Complexity-theoretic investigation
- Trace Ideals and Realizability
- General case:
 - O(kn^{2k-2}) time algorithm
 - $O(k \cdot n^k)$ when γ and ζ are constant
 - **W[**]]-hardness
 - Quadratic lower bound for k=2 threads
- Acyclic topology
 - Quadratic time upper and lower bounds
- Distance Bounded data race prediction
 - Linear time realizability

- Tighter bounds
 - Is data race prediction
 W[k]-complete for some k?
 - Space-time tradeoffs

Summary

- Dynamic Race Prediction
 - Complexity-theoretic investigation
- Trace Ideals and Realizability
- General case:
 - O(kn^{2k-2}) time algorithm
 - $O(k \cdot n^k)$ when γ and ζ are constant
 - **W[**]]-hardness
 - Quadratic lower bound for k=2 threads
- Acyclic topology
 - Quadratic time upper and lower bounds
- Distance Bounded data race prediction
 - Linear time realizability

- Tighter bounds
 - Is data race prediction
 W[k]-complete for some k?
 - Space-time tradeoffs
- Other algorithmic paradigms
 - Randomized algorithms
 - Parallel algorithms

Summary

- Dynamic Race Prediction
 - Complexity-theoretic investigation
- Trace Ideals and Realizability
- General case:
 - O(kn^{2k-2}) time algorithm
 - $O(k \cdot n^k)$ when γ and ζ are constant
 - **W[**]]-hardness
 - Quadratic lower bound for k=2 threads
- Acyclic topology
 - Quadratic time upper and lower bounds
- Distance Bounded data race prediction
 - Linear time realizability

- Tighter bounds
 - Is data race prediction
 W[k]-complete for some k?
 - Space-time tradeoffs
- Other algorithmic paradigms
 - Randomized algorithms
 - Parallel algorithms
- Practicability constraints -
 - Linear time algorithms
 - (Sub-)linear dependence on d
 - (Sub-)linear space overhead

Thank You!