

Explaining Concurrency Bugs with Interpolants

Andreas Holzer², Daniel Schwartz-Narbonne³, Mitra Tabaei¹,
Georg Weissenbacher¹, Thomas Wies³

¹Vienna University of Technology

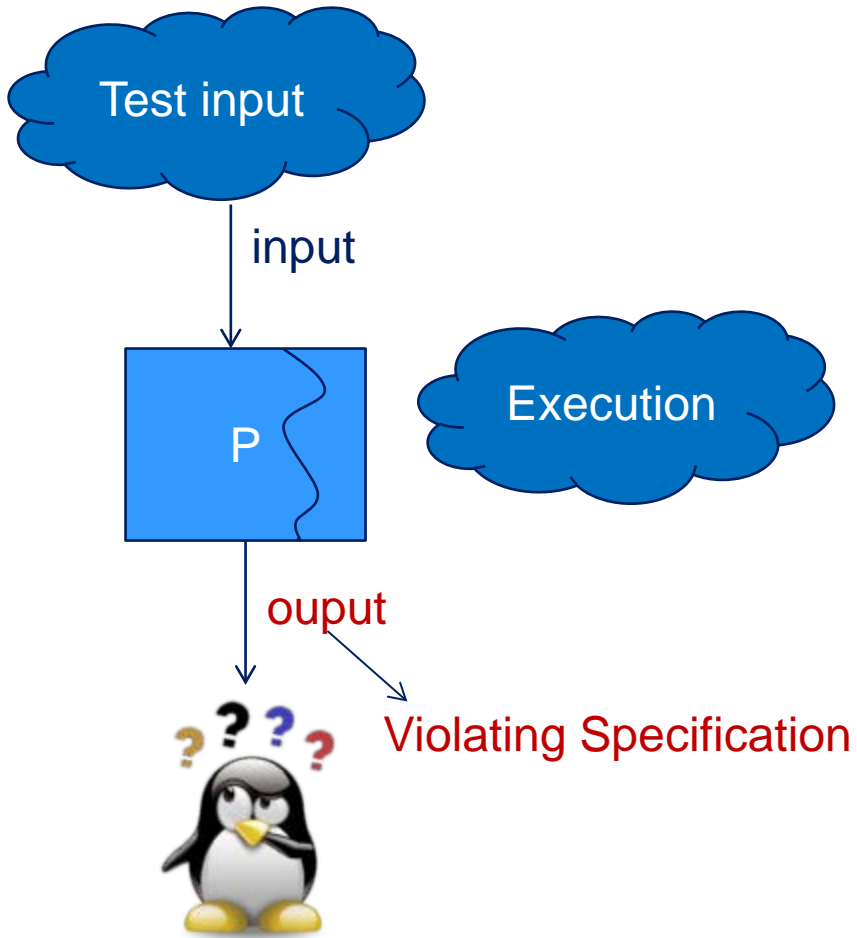
²University of Toronto

³New York University

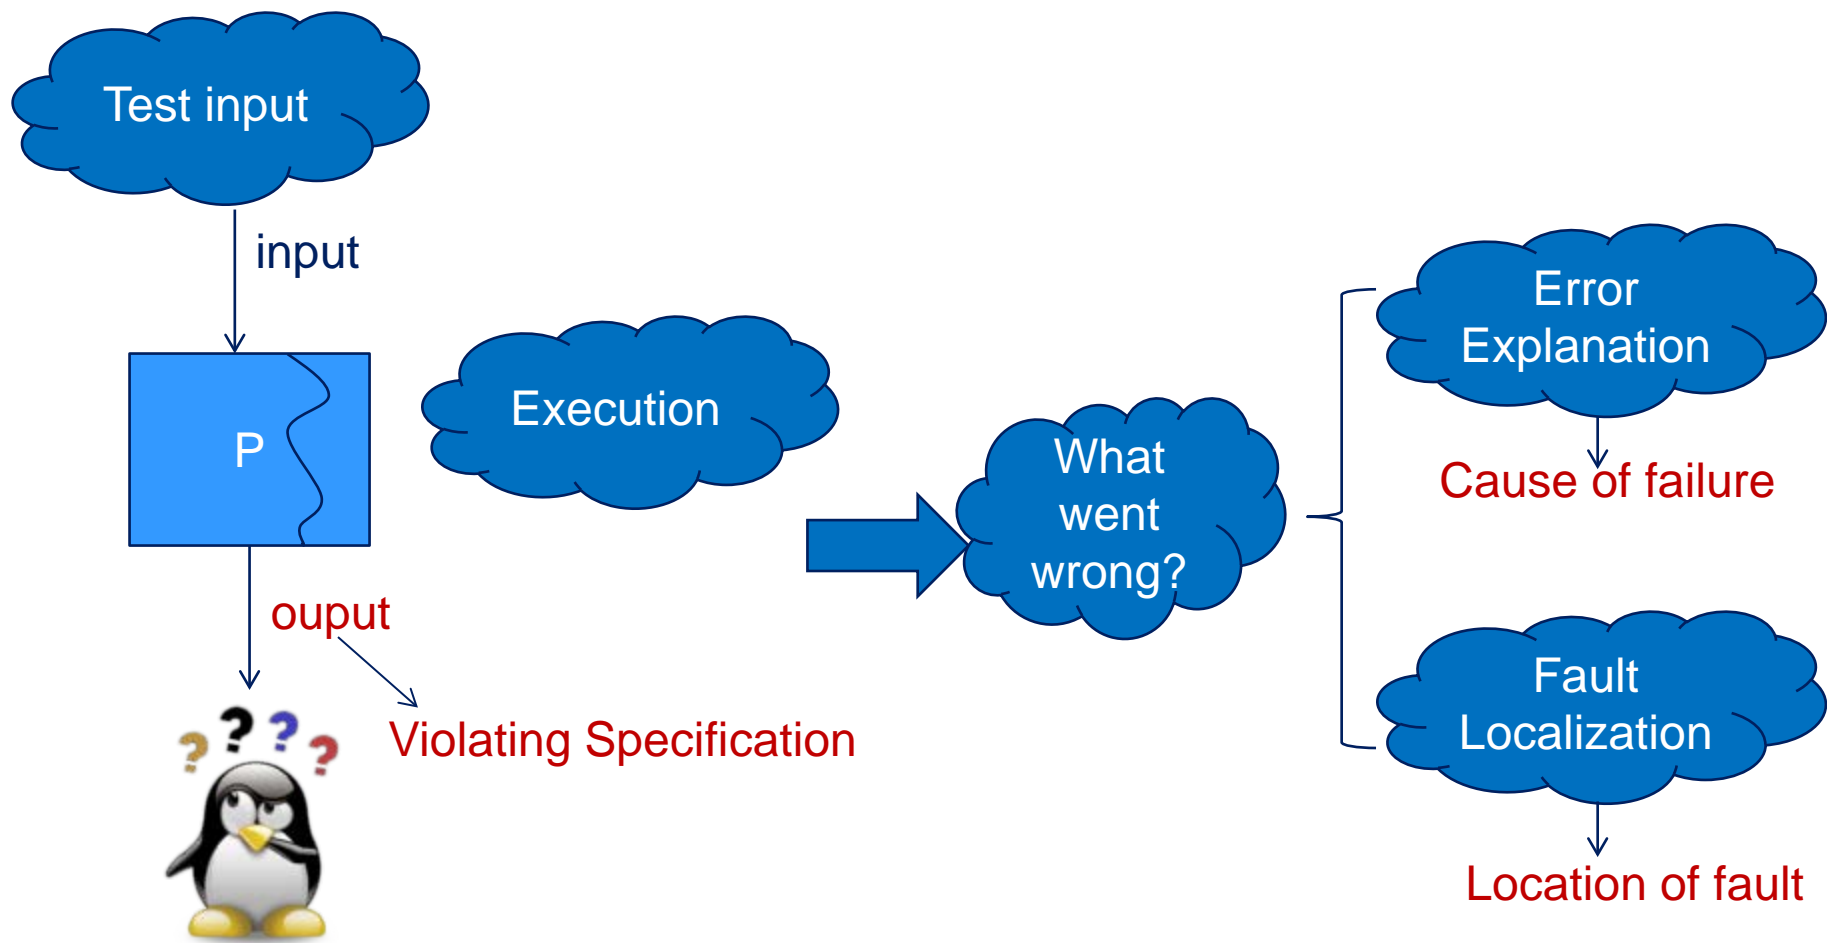
AVM, Austria



Debugging




Debugging



Automatic Debugging Techniques

- Dynamic analysis:
 - Comparison of **failing** and **passing** traces
 - ✦ Quality of test suite



- Symbolic execution analysis:
 - Max-SAT
 - ✦ Cause clue clauses [PLDI11]
 - Interpolation
 - ✦ Error Invariant [FM12]
 - ✦ Flow-sensitive Fault Localization[VMCAI13]
 - ✦ Hybrid Algorithm [VSSTE14]

Overview of our Method

- A concurrency bug explanation technique:
 - Symbolic execution analysis
 - Interpolation
- A general framework for concurrency bug explanation
 - Not relying on specific bug characteristic
 - No given pattern templates or annotations

Outline

- Notion of Interpolant
- Interpolants for debugging sequential traces
 - Encoding control-dependencies
 - ✦ Flow-sensitive slices

- Interpolants for explaining concurrency bugs
 - Encoding:
 - ✦ Locks
 - ✦ Inter-thread data-dependencies
- Empirical Evaluation

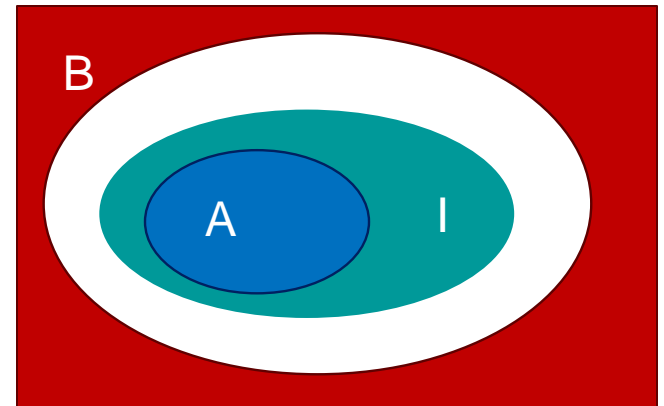
Interpolants

Given: an unsatisfiable conjunction of formulas $A \wedge B$:

$$A \wedge B \equiv \text{false}$$

An **Interpolant** for $A \wedge B$ is a formula I s.t. :

- $A \Rightarrow I$
- $I \wedge B \equiv \text{false}$
- I is only on common variables of A and B




Trace Formula

Failing trace

input


program statement



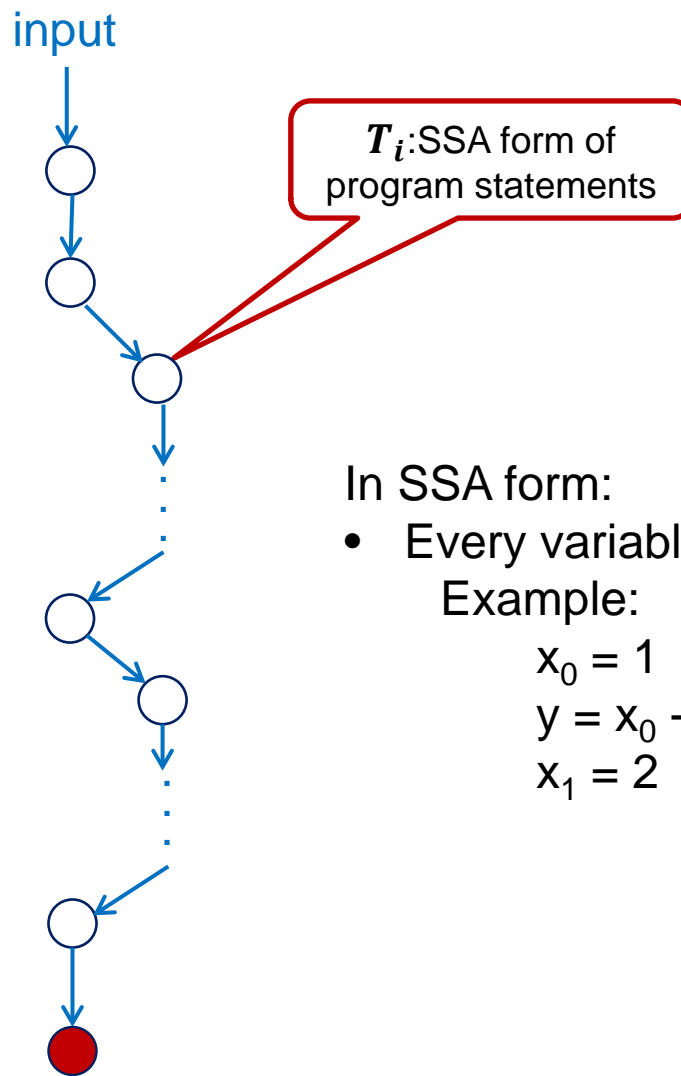
assertion

Trace Formula

Failing trace



Trace in SSA form



In SSA form:

- Every variable is assigned once:

Example:


$$x_0 = 1$$

$$y = x_0 + 5$$

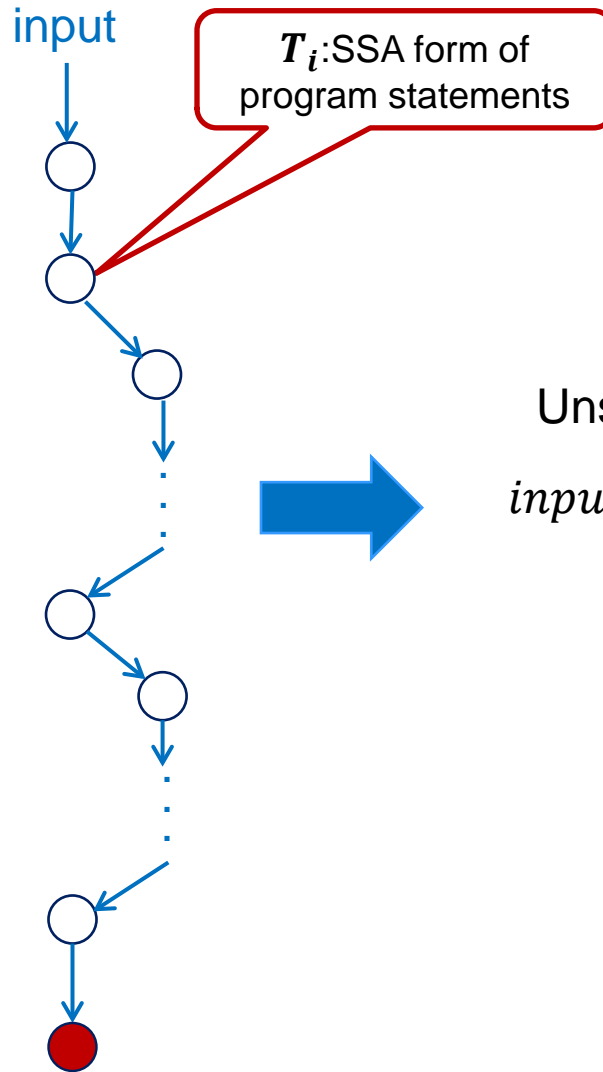
$$x_1 = 2$$

Trace Formula

Failing trace



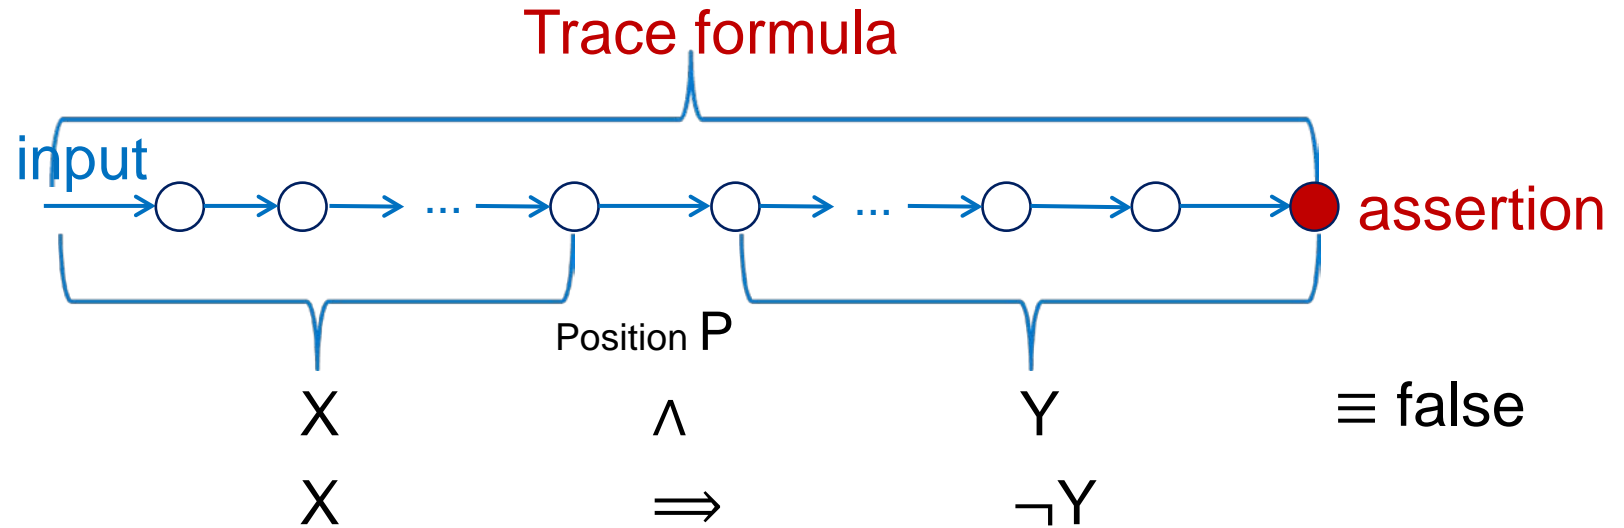
Trace in SSA form



Unsatisfiable trace formula:

$$input \wedge T_1 \wedge T_2 \dots \wedge T_n \wedge assertion$$


Interpolants for Debugging



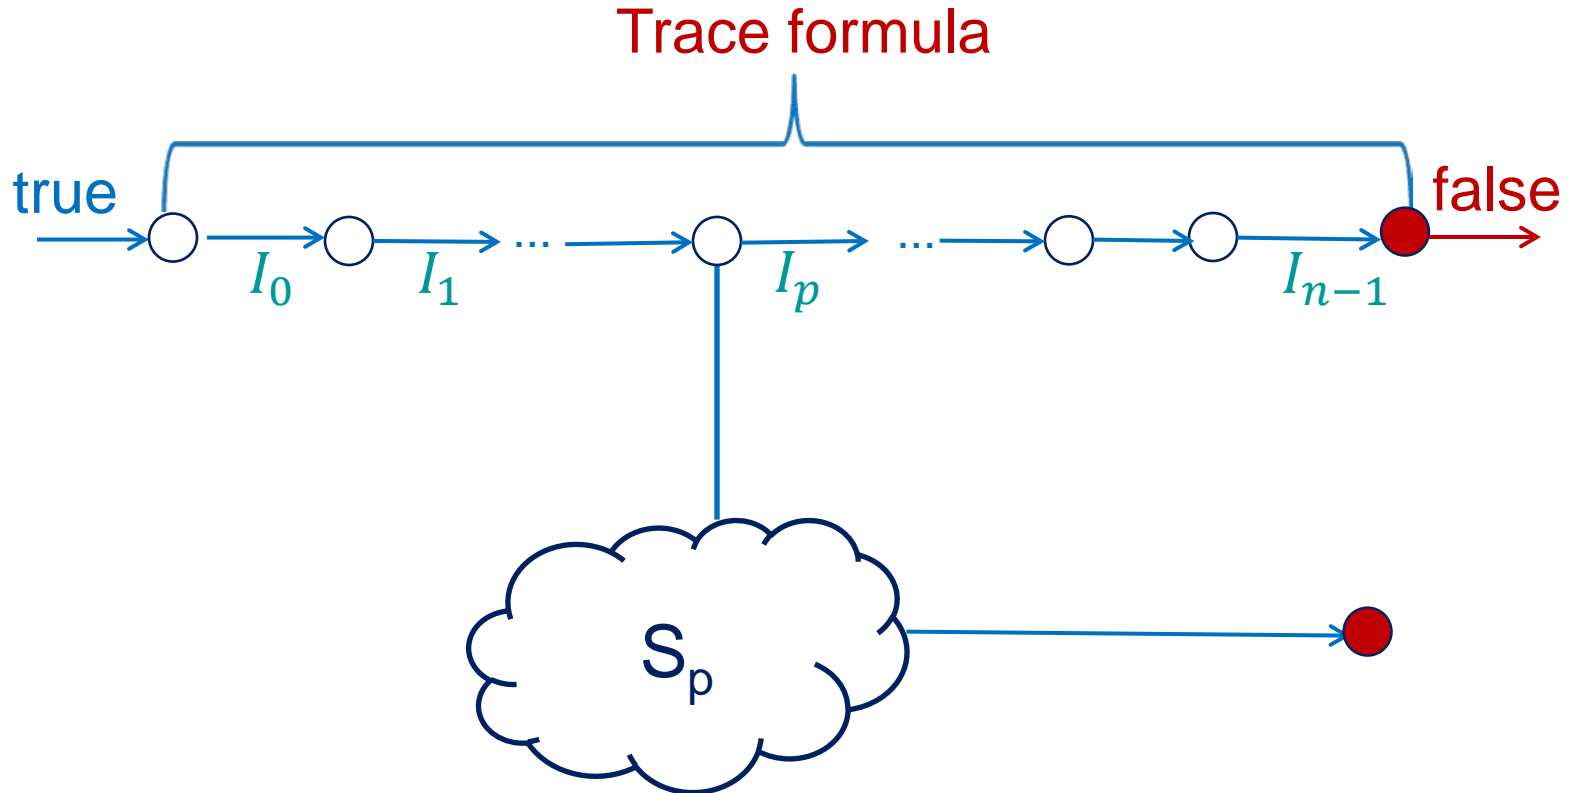
Interpolant at Position P:

$$X \quad \Rightarrow \quad I_P \quad \Rightarrow \quad \neg Y$$

Interpolants for Debugging




Interpolants for Debugging




Over-approximation of reachable states at p

Interpolants for Debugging



Interpolants for Debugging

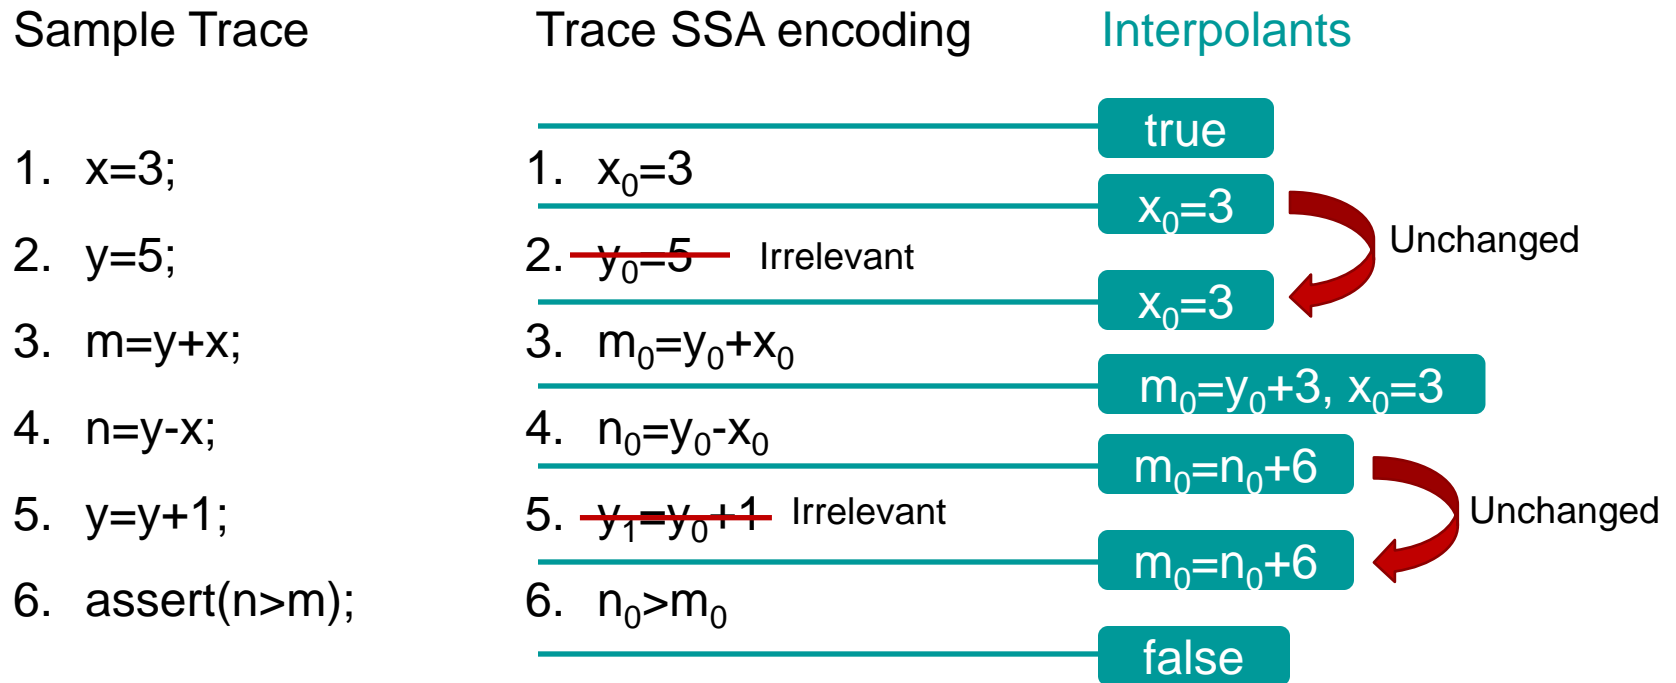


Interpolants for Debugging

Sample Trace	Trace SSA encoding	Interpolants
1. $x=3;$	$x_0=3$	true
2. $y=5;$	$y_0=5$	$x_0=3$
3. $m=y+x;$	$m_0=y_0+x_0$	$x_0=3$
4. $n=y-x;$	$n_0=y_0-x_0$	$m_0=y_0+3, x_0=3$
5. $y=y+1;$	$y_1=y_0+1$	$m_0=n_0+6$
6. $\text{assert}(n>m);$	$n_0>m_0$	$m_0=n_0+6$
		false

- Interpolants
 - contain enough information to understand the failure

Interpolants for Debugging



Error Explanation

Sample Trace	Error Explanation (Slice)	Error Invariants
1. <code>x=3;</code>	1. <code>x=3</code>	true
2. <code>y=5;</code>		<code>x=3</code>
3. <code>m=y+x;</code>	3. <code>m=y+x</code>	<code>m=y+3,x=3</code>
4. <code>n=y-x;</code>	4. <code>n=y-x</code>	
5. <code>y=y+1;</code>		<code>m=n+6</code>
6. <code>assert(n>m);</code>	6. <code>assert(n>m)</code>	false

Error Explanation:

- Slice: Isolating relevant *statements* for assertion violation
- Error Invariants: Revealing the relevant *variables*

Error Explanation

Sample Trace	Error Explanation (Slice)	Error Invariants
1. <code>x=3;</code>	1. <code>x=3</code>	true
2. <code>y=5;</code>		<code>x=3</code>
3. <code>m=y+x;</code>	3. <code>m=y+x</code>	<code>m=y+3,x=3</code>
4. <code>n=y-x;</code>	4. <code>n=y-x</code>	
5. <code>y=y+1;</code>		<code>m=n+6</code>
6. <code>assert(n>m);</code>	6. <code>assert(n>m)</code>	false

Error Explanation:

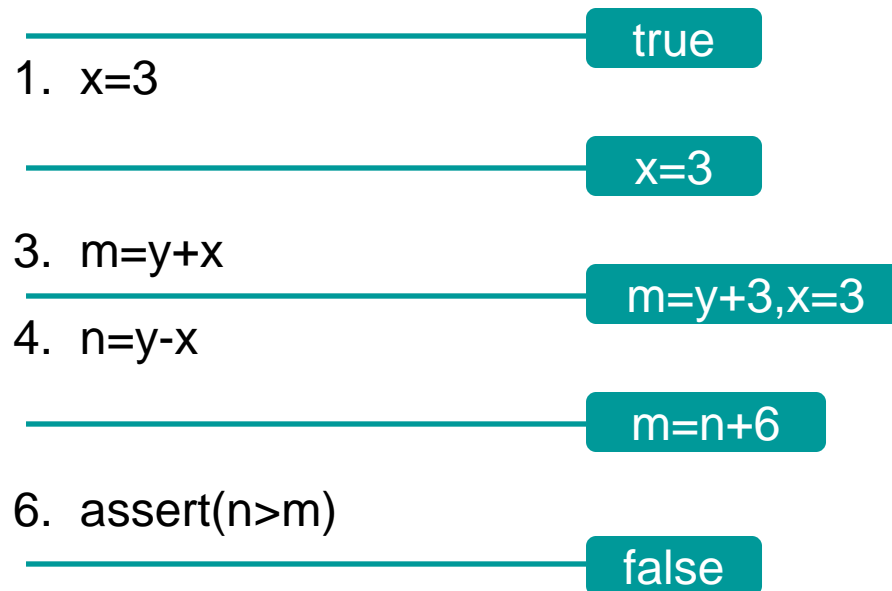
Interpolants are not unique: any formula between WP and SP

- Error Invariants: Revealing the relevant *variables*

Sound Error Explanation Slices

- Soundness of explanation
 - Slice forms an unsatisfiable formula


Sound Slice



Sound Error Explanation Slices

- Soundness of explanation
 - Achieved by *Inductive Interpolant Sequence* [VSSTE 2014]


$$I_{p-1} \wedge T_p \Rightarrow I_p \quad \text{Inductive property}$$



Sound Error Explanation Slices

- Soundness of explanation
 - Achieved by *Inductive Interpolant Sequence* [VSSTE 2014]

$$I_{p-1} \wedge T_p \Rightarrow I_p \quad \text{Inductive property}$$



Interpolants for Debugging

- Generating unsatisfiable trace formula
 - by SSA encoding
- Computing inductive interpolants
 - for each position in the trace
- Excluding statements
 - with stationary surrounding interpolants

Encoding Conditions

Sample Code	Sample Trace	Slice
1. <code>x=1;</code>	1. <code>x₀=1</code>	
2. <code>y=*;</code>	2. <code>y₀=-10</code>	
3. <code>if (y < 0)</code>	3. <code>(y₀<0)∧x₁=0</code>	
4. <code>x=0;</code>	4. <code>x₁ ≠ 0</code>	4. <code>x=0;</code>
5. <code>assert(x!=0);</code>		5. <code>assert(x != 0)</code>

Conditions are required for understanding the failure.

Encoding Conditions

SSA Trace

1. $x_0=1;$	true
2. $y_0=-10;$	true
3. $(y_0 < 0) \wedge x_1=0;$	true
4. $x_1 \neq 0$	$x_1=0$
	false

~~1. $x=1;$~~
~~2. $y=*;$~~
~~3. if ($y < 0$)~~
4. $x=0;$
5. $\text{assert}(x \neq 0);$

Flow-sensitive SSA Trace

1. $x_0=1;$	true
2. $y_0=-10;$	true
3. $(y_0 < 0) \Rightarrow x_1=0;$	$y_0 < 0$
4. $x_1 \neq 0$	$x_1=0$
	false

~~1. $x=1;$~~
2. $y=*;$
3. if ($y < 0$)
4. $x=0;$
5. $\text{assert}(x \neq 0);$

Flow-sensitive Slices

SSA Trace

1. <u>$x_0=1;$</u>	true
2. <u>$y_0=-10;$</u>	true
3. <u>$(y_0 < 0) \wedge x_1=0;$</u>	true
4. <u>$x_1 \neq 0$</u>	$x_1=0$
	false

↓

1. ~~$x=1;$~~

Flow-sensitive SSA Trace

1. <u>$x_0=1;$</u>	true
2. <u>$y_0=-10;$</u>	true
3. <u>$(y_0 < 0) \Rightarrow x_1=0;$</u>	$y_0 < 0$
4. <u>$x_1 \neq 0$</u>	$x_1=0$
	false

↓


1. ~~$x=1;$~~

- Encoding of control-dependency:
 - Conditions are encoded as implications in SSA traces:

$$\left(\bigwedge_{c \in \text{conds}} c \right) \Rightarrow x = e$$


Model of Concurrent Traces

Multi-threaded Programs



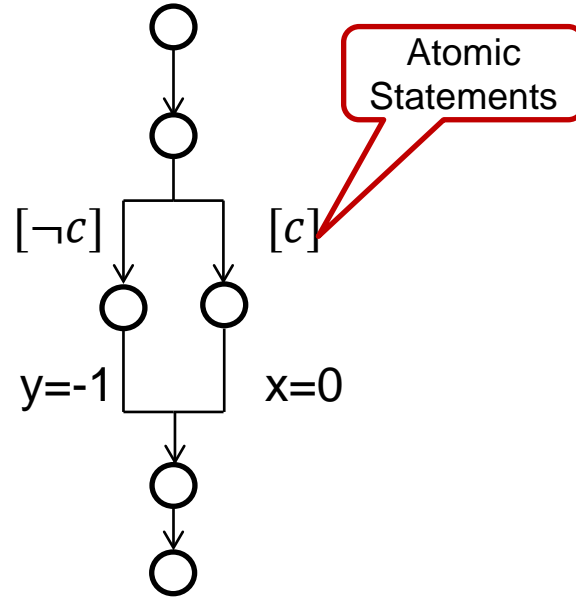
Model of Concurrent Traces

Multi-threaded Programs




modeled as

Thread CFG



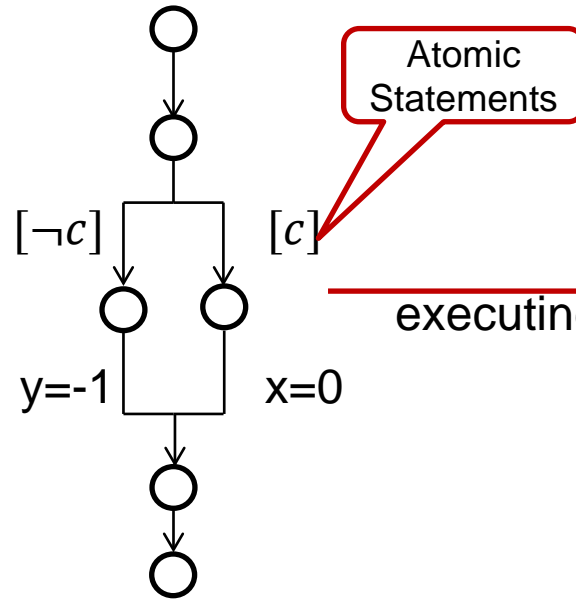
Model of Concurrent Traces

Multi-threaded Programs




modeled as

Thread CFG



executing

Concurrent Trace

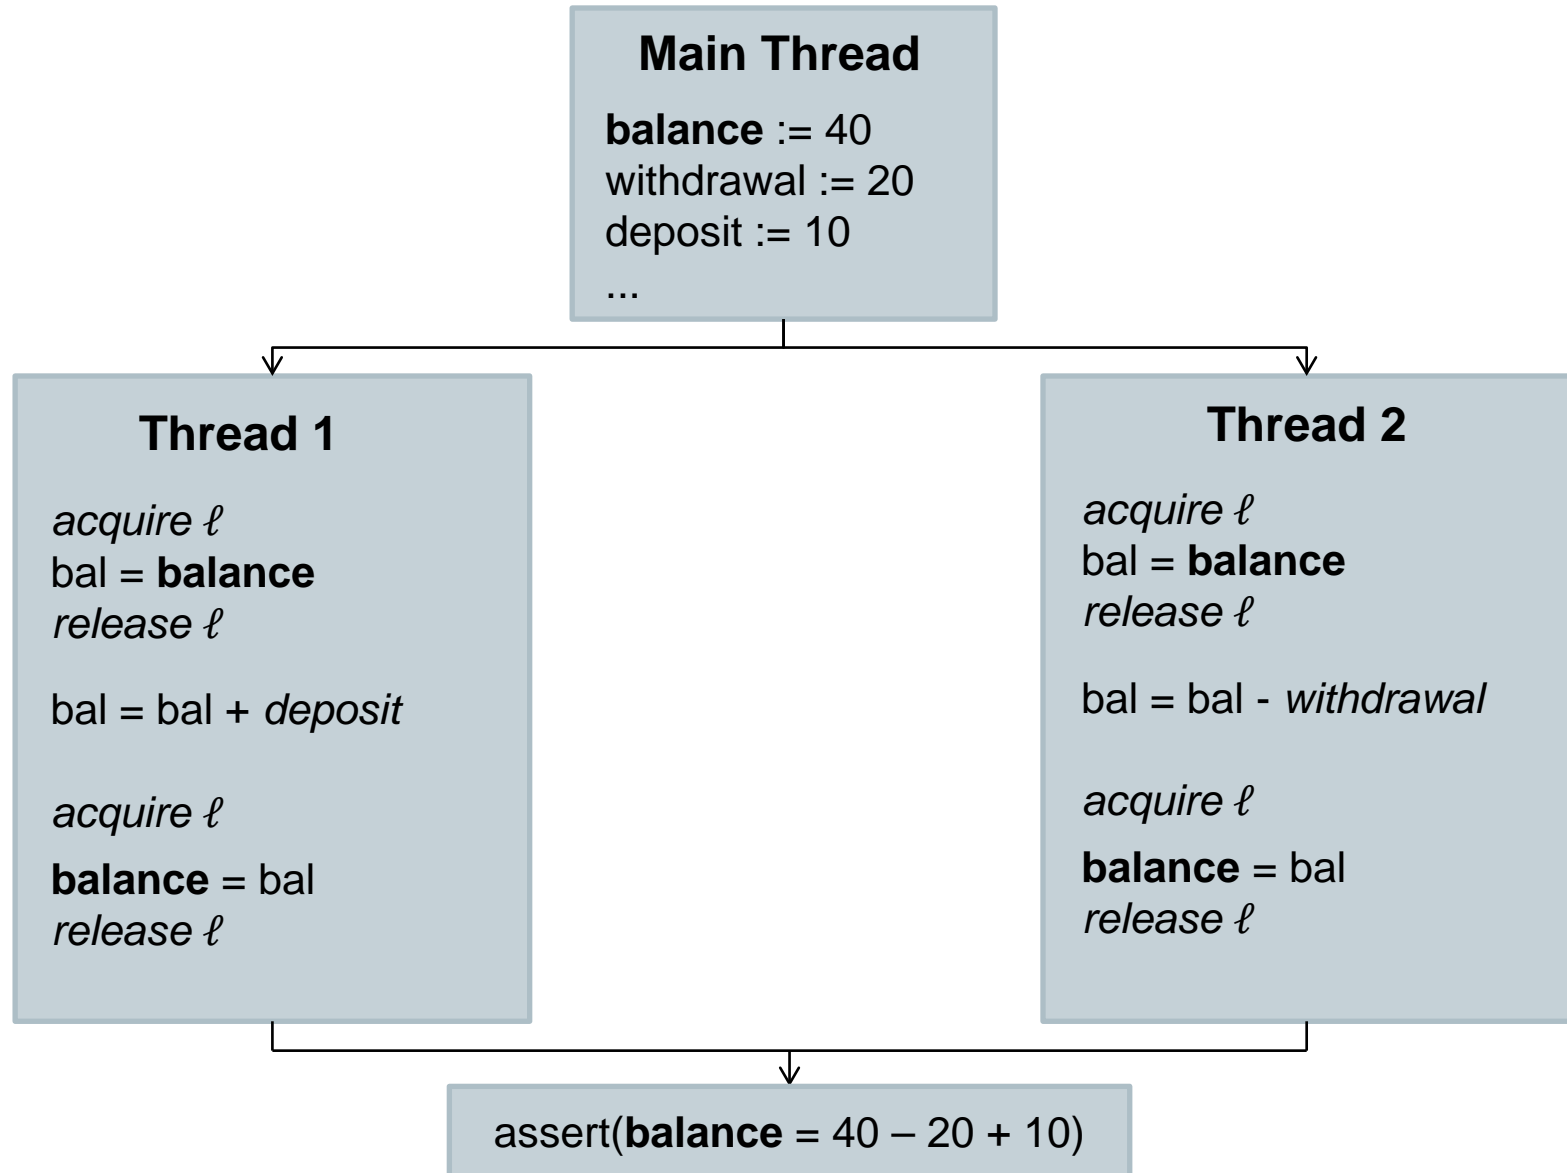


Total order (interleaving semantics)

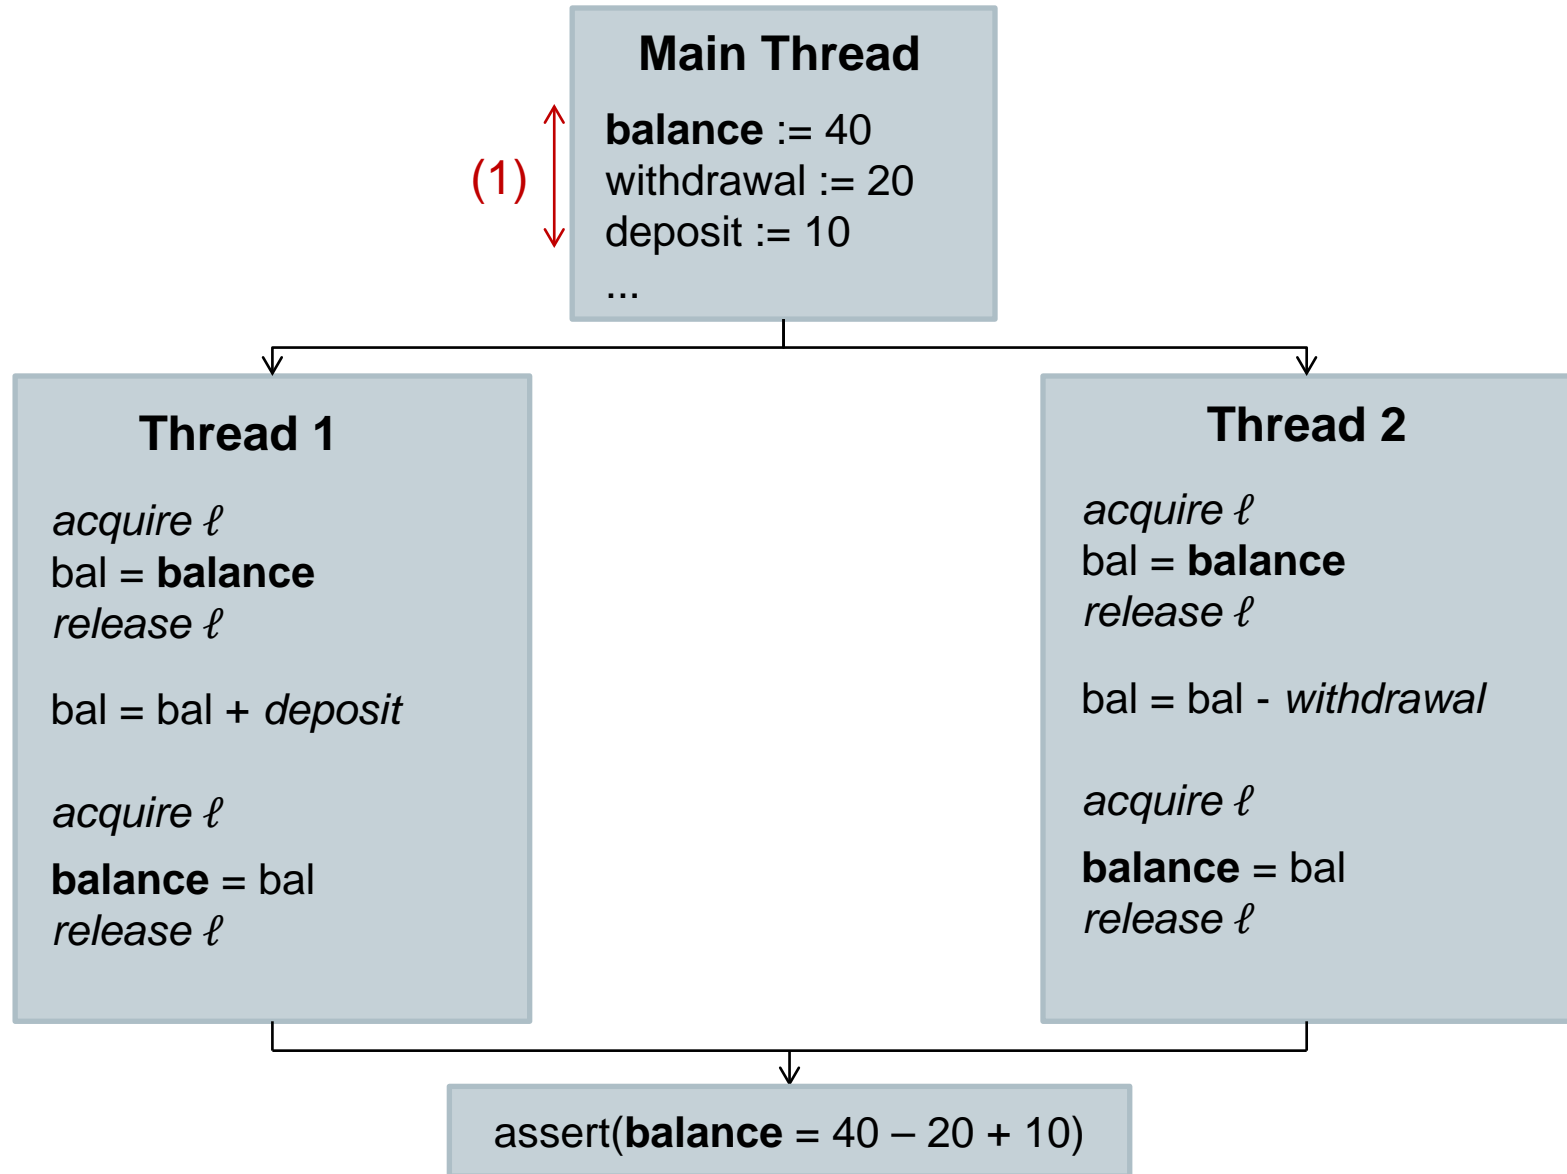
Concurrent Trace Formula

- SSA encoding of variables
- Encoding control-dependency as implication
- Modeling Locks as:
 - Atomic guarded assignments:
acquire ℓ : $(\ell = 0) \triangleright \ell := tid$
release ℓ : $(\ell = tid) \triangleright \ell := 0$
 - Encoding locks as implications (similar to control-dependency)

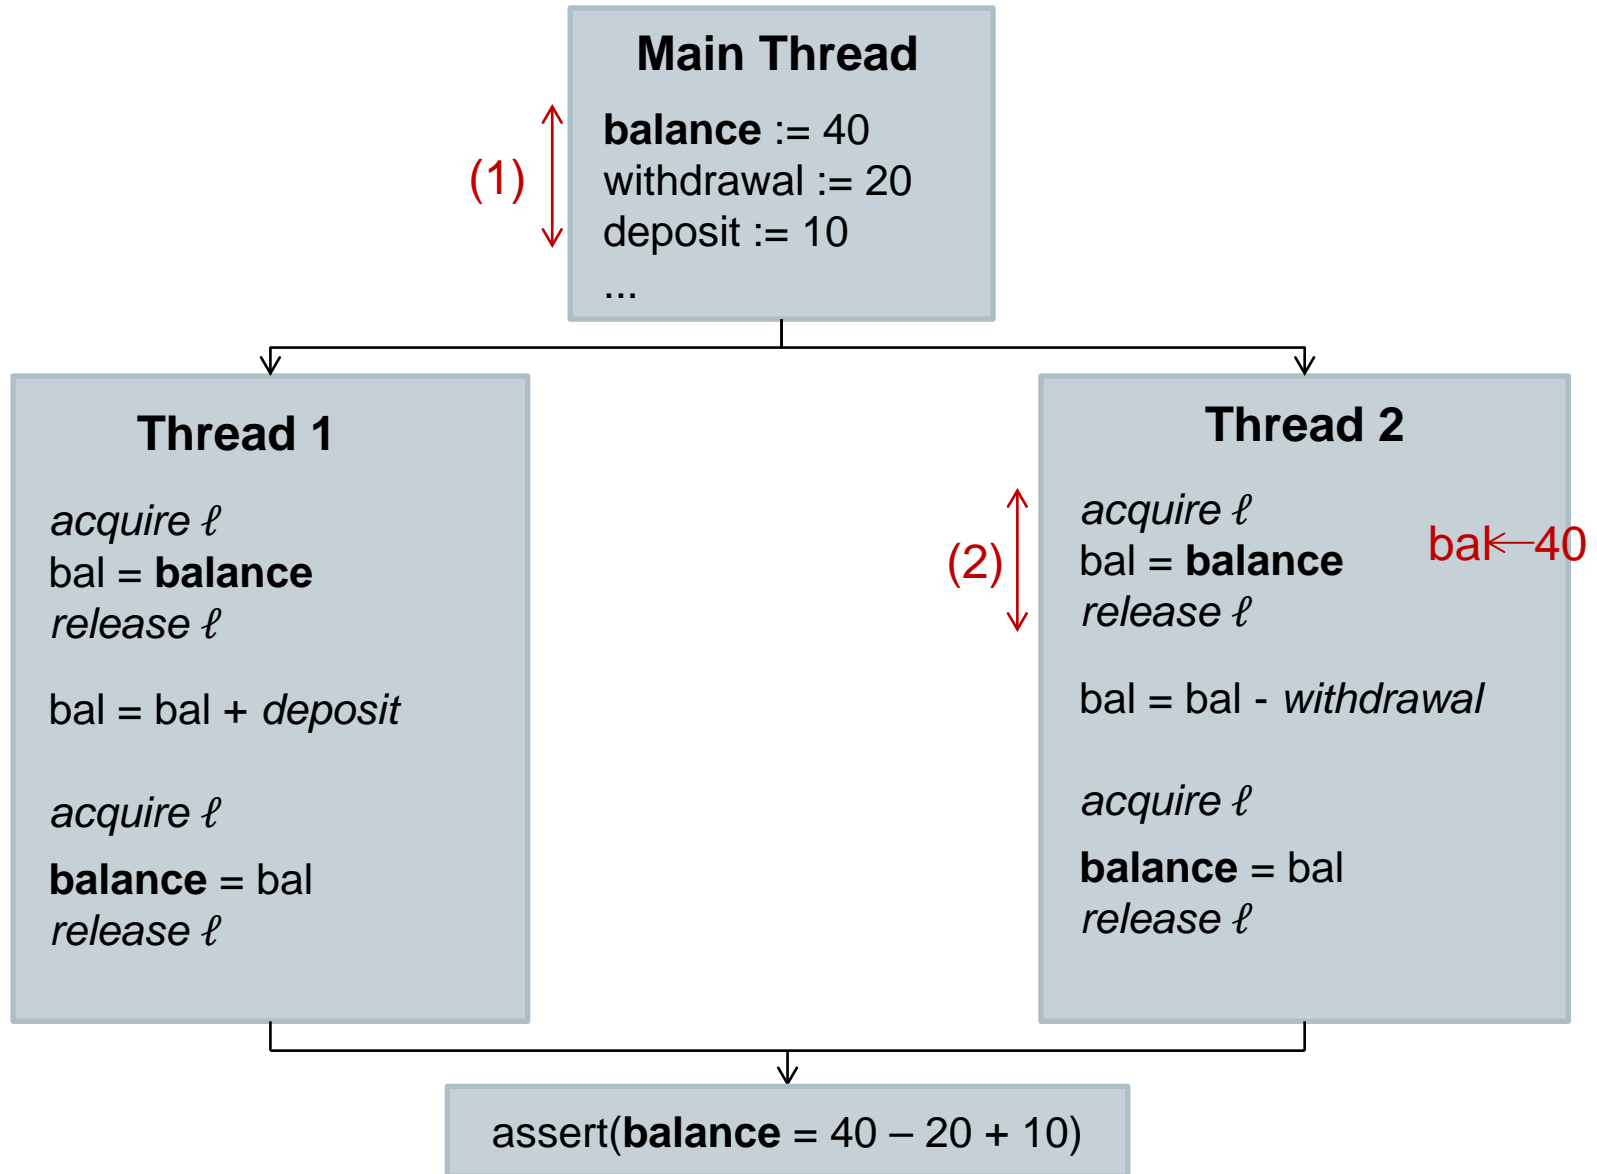
Interpolants for Debugging Concurrent Bugs



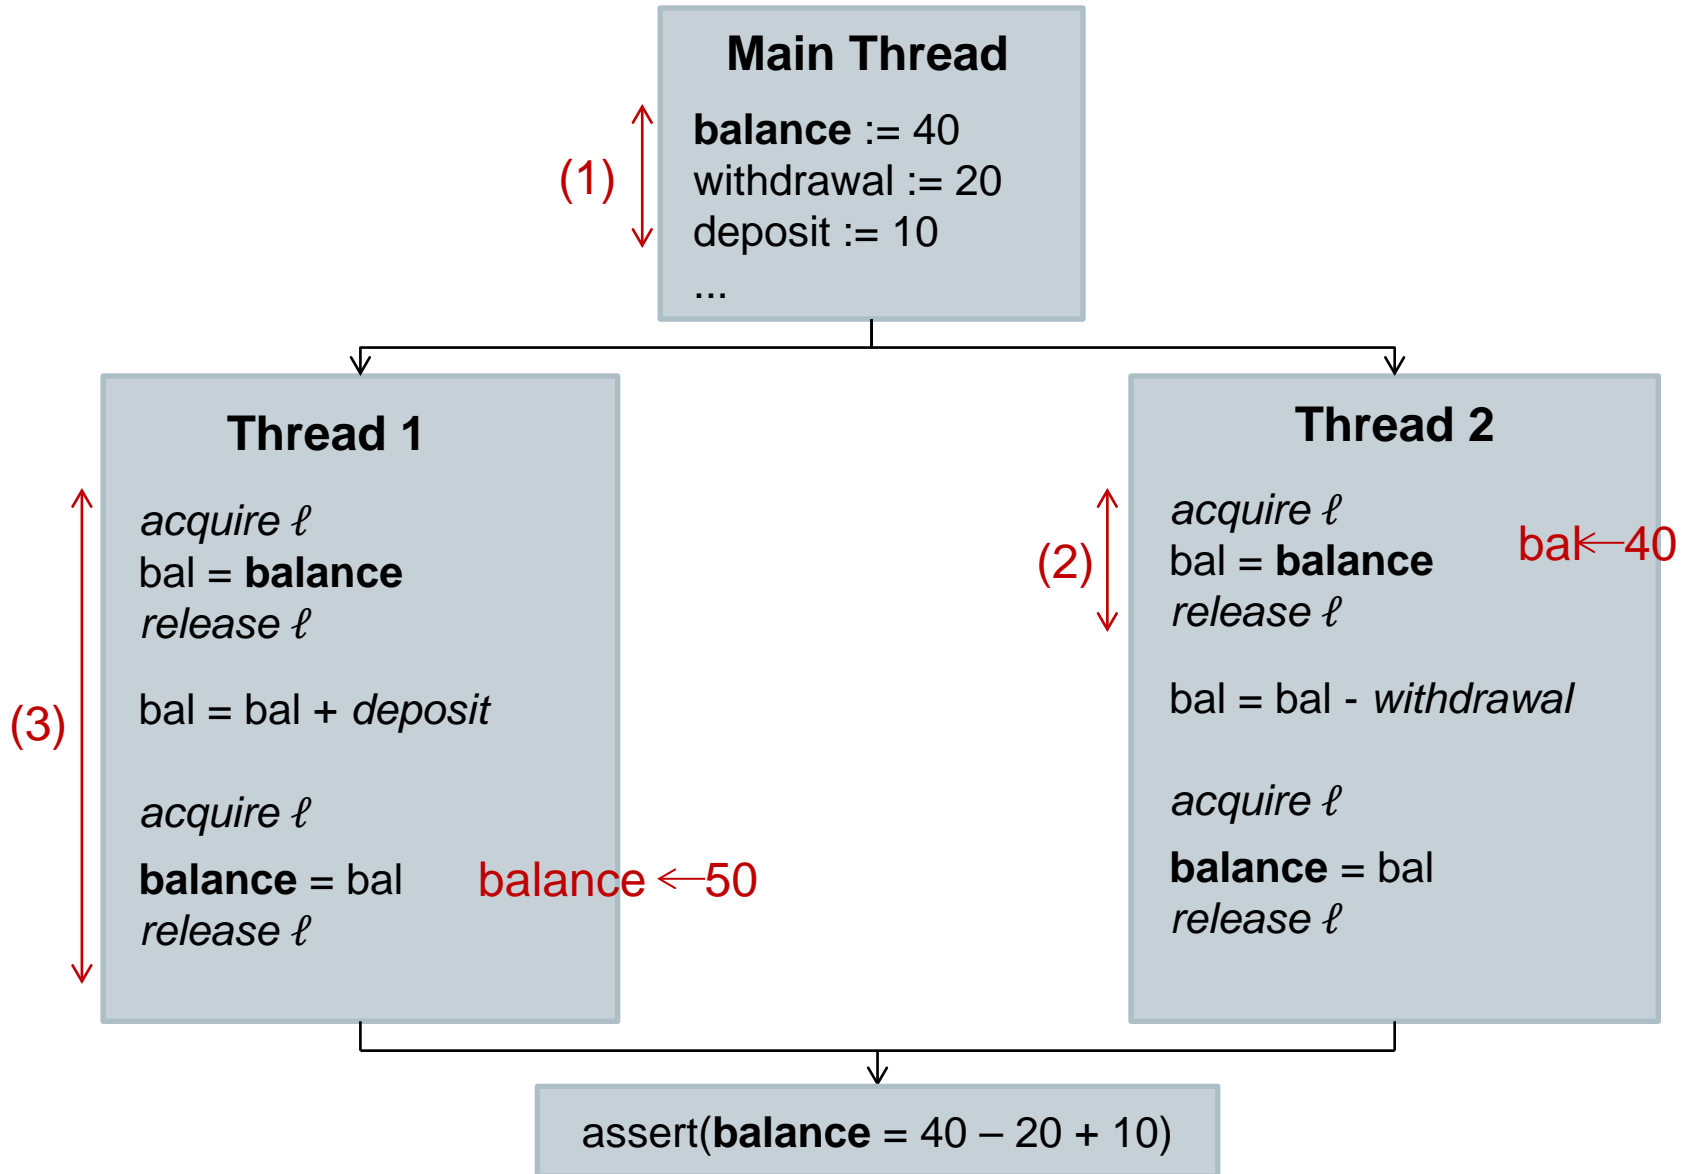
Interpolants for Debugging Concurrent Bugs



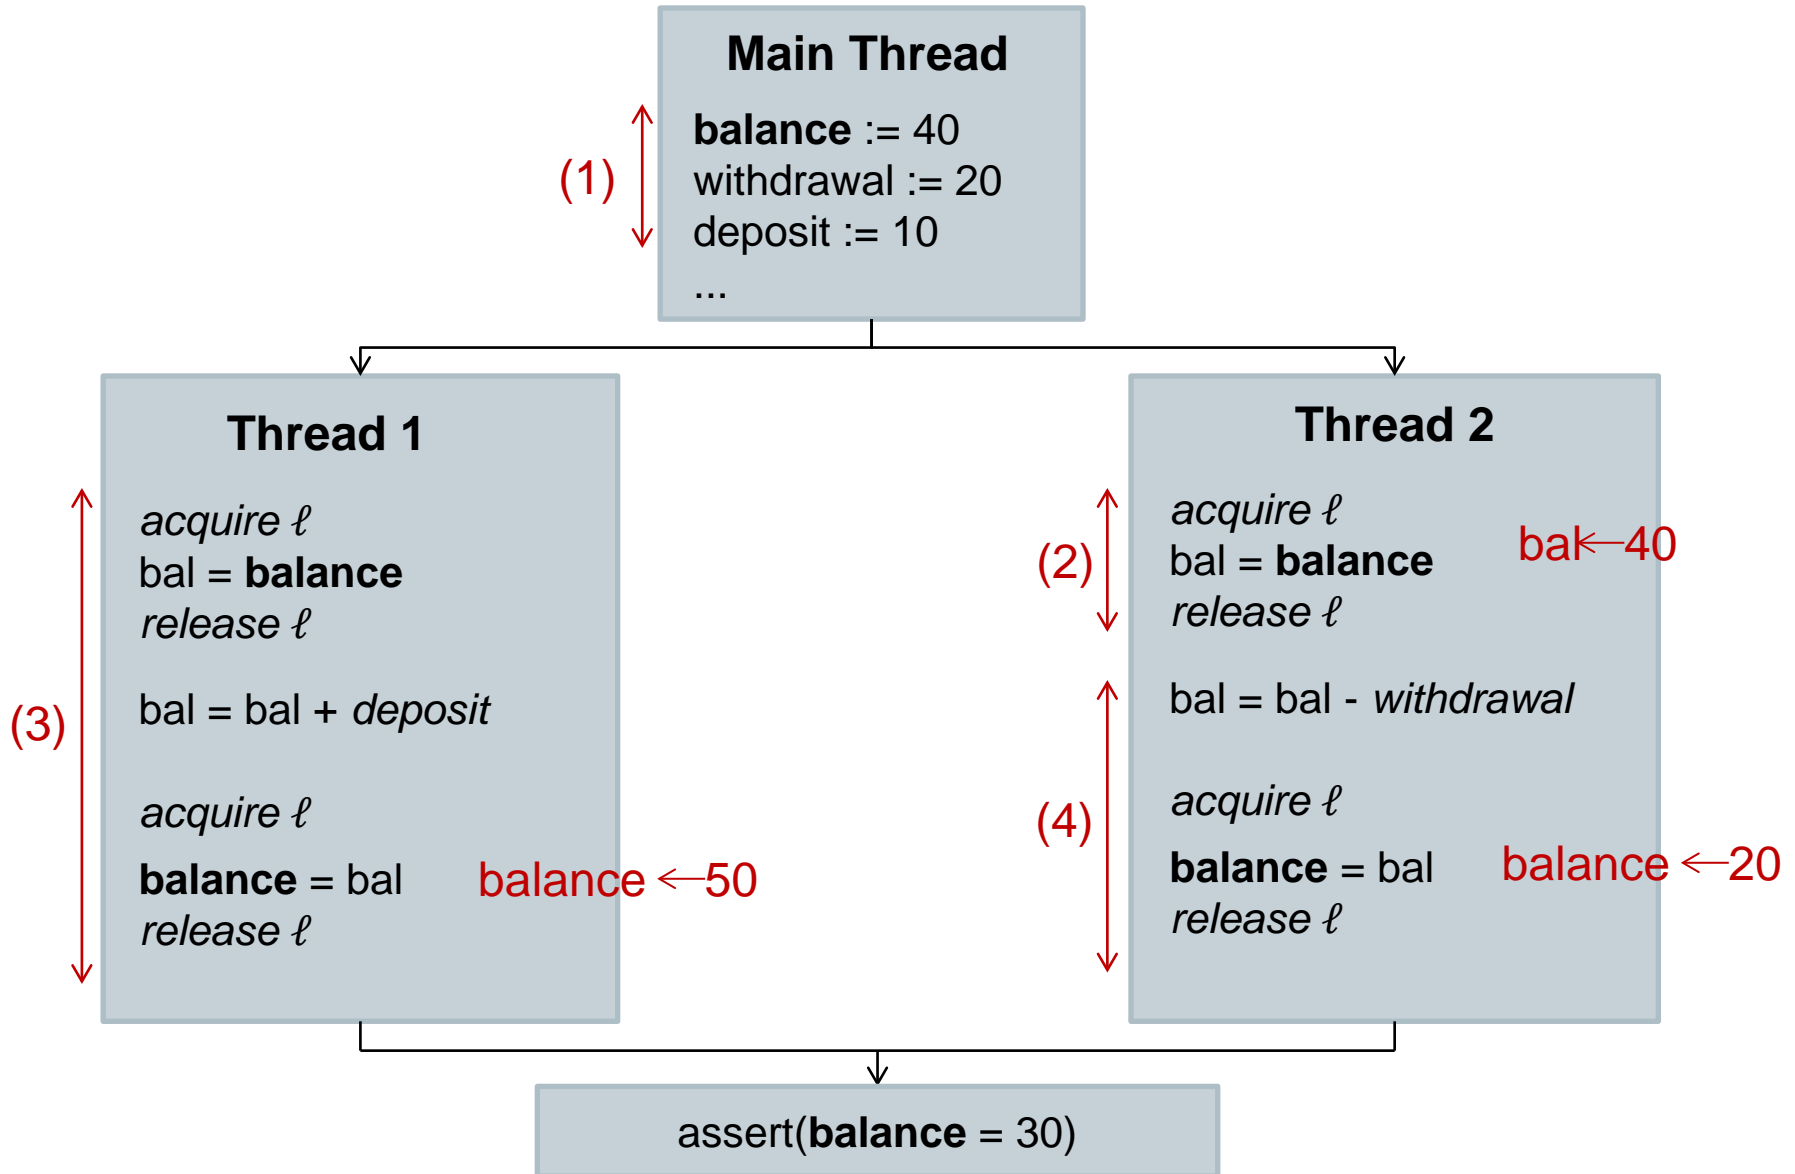
Interpolants for Debugging Concurrent Bugs



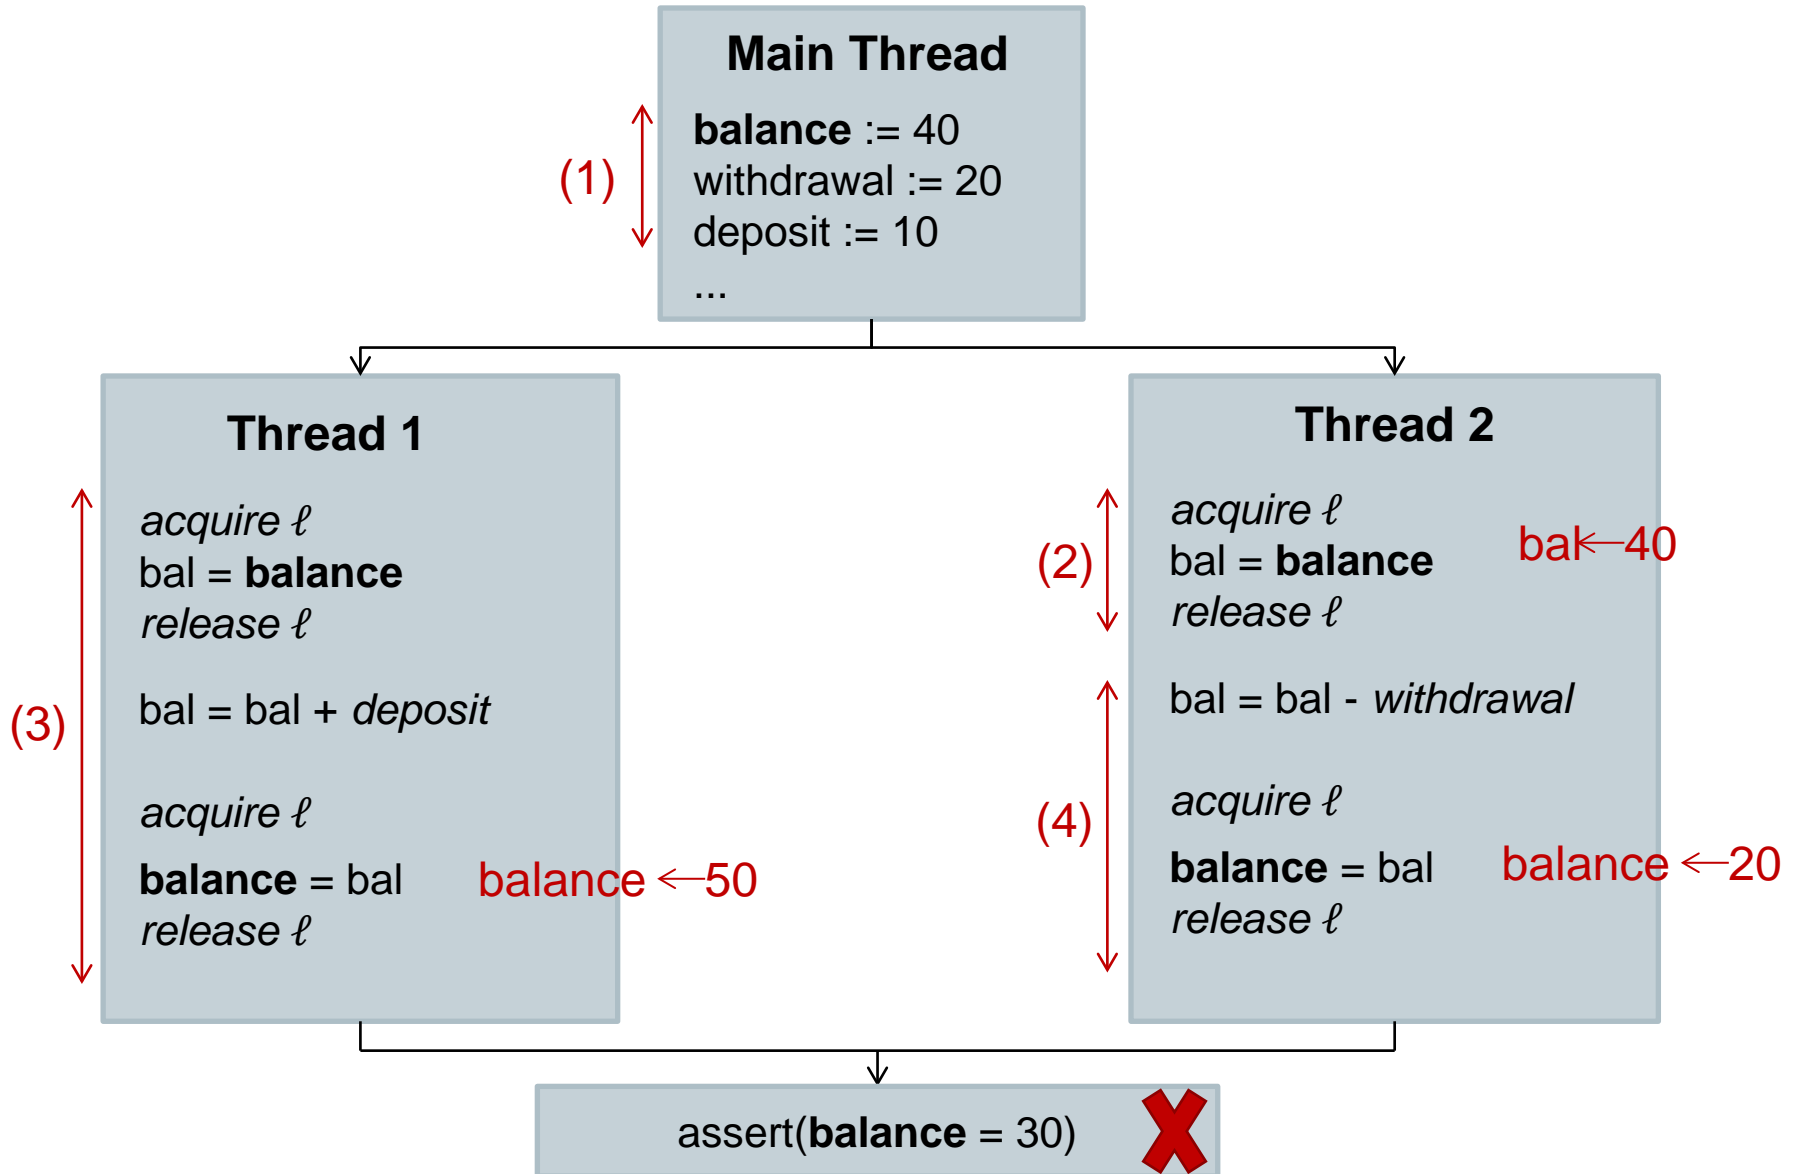
Interpolants for Debugging Concurrent Bugs



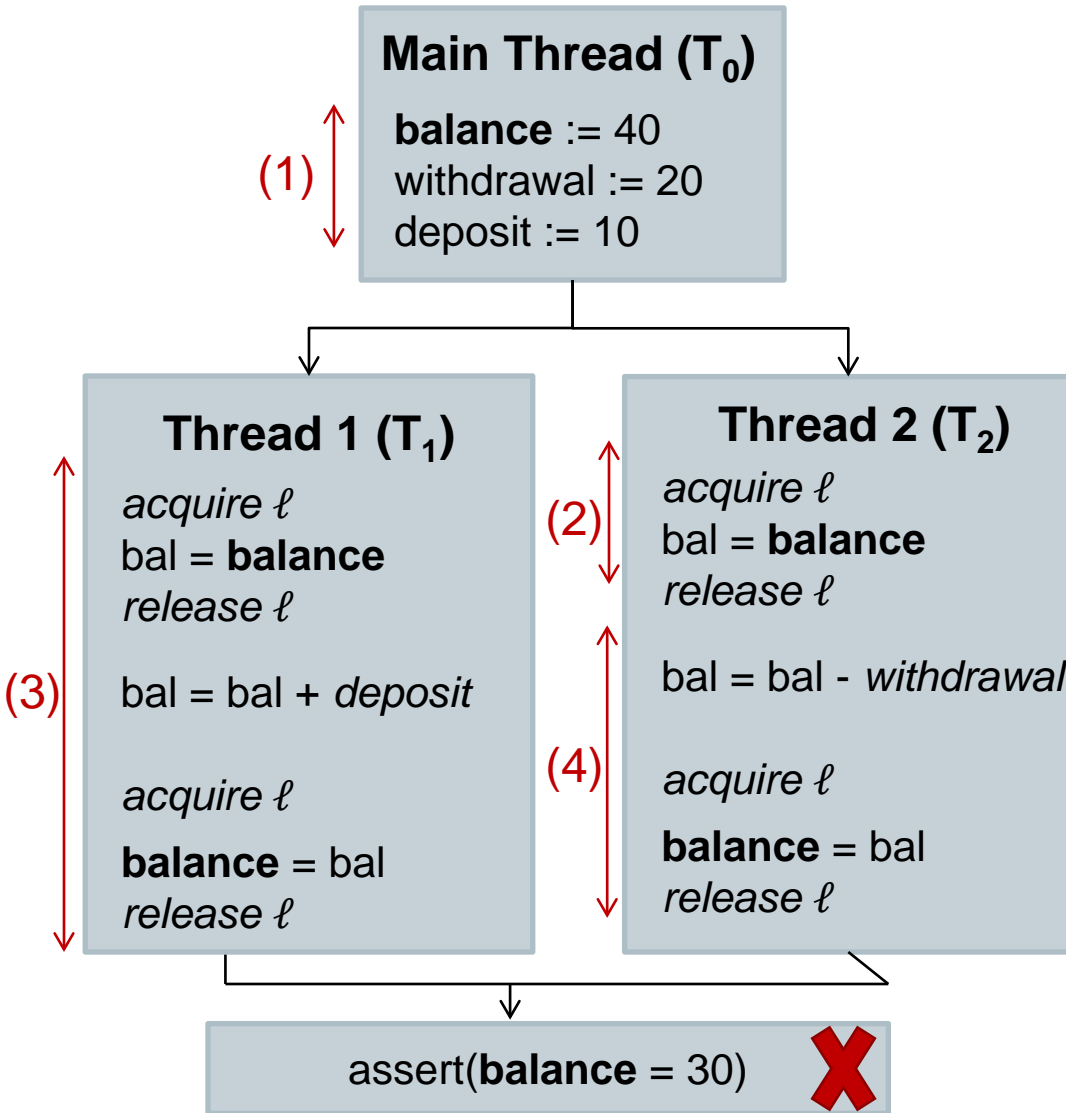
Interpolants for Debugging Concurrent Bugs



Interpolants for Debugging Concurrent Bugs



Interpolants for Debugging Concurrent Bugs



Flow-insensitive Slice

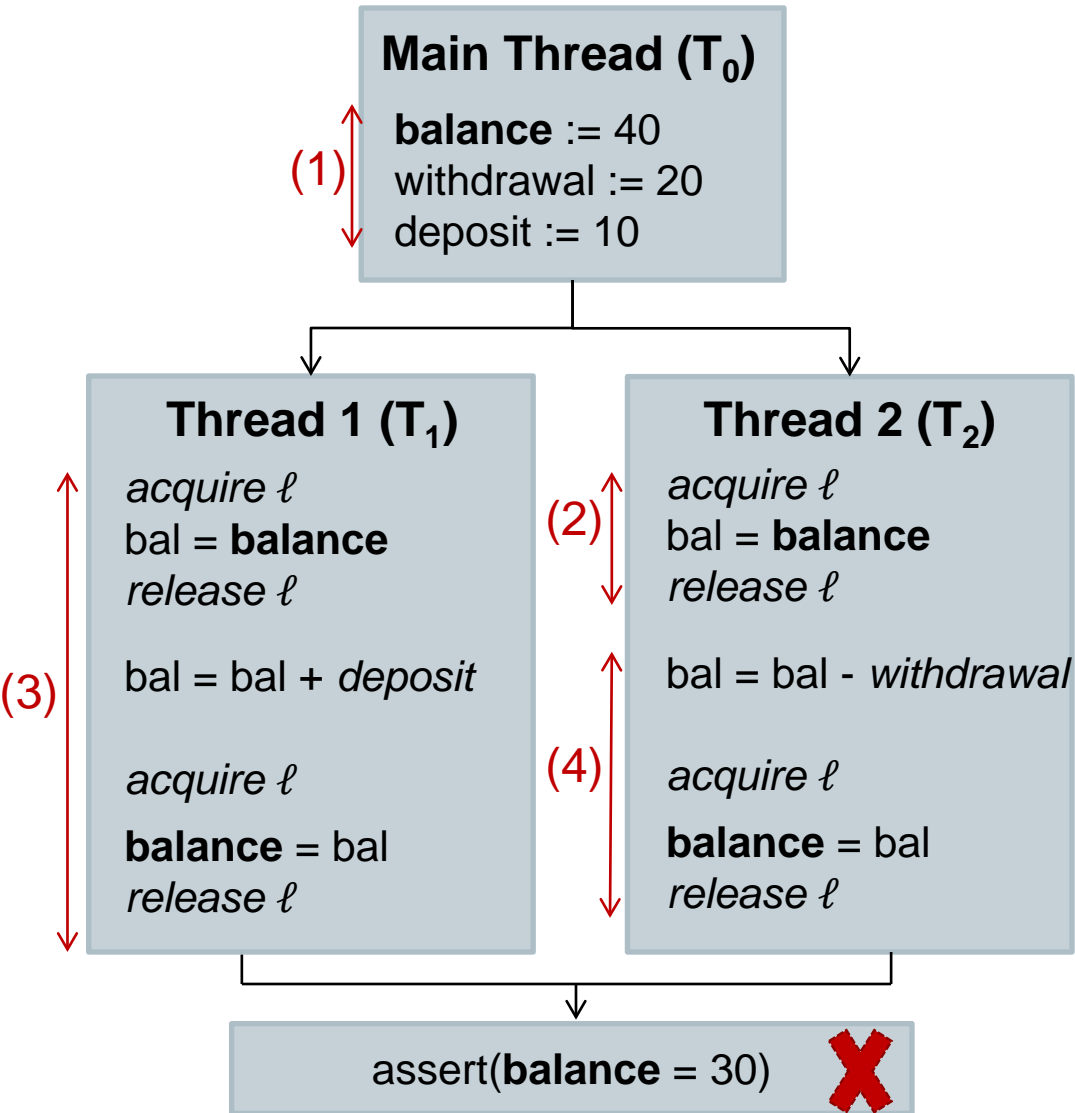
T₀: **balance** := 40
 T₀: withdrawal := 20

T₂: bal = **balance**
 T₂: bal = bal - *withdrawal*
 T₂: **balance** = bal

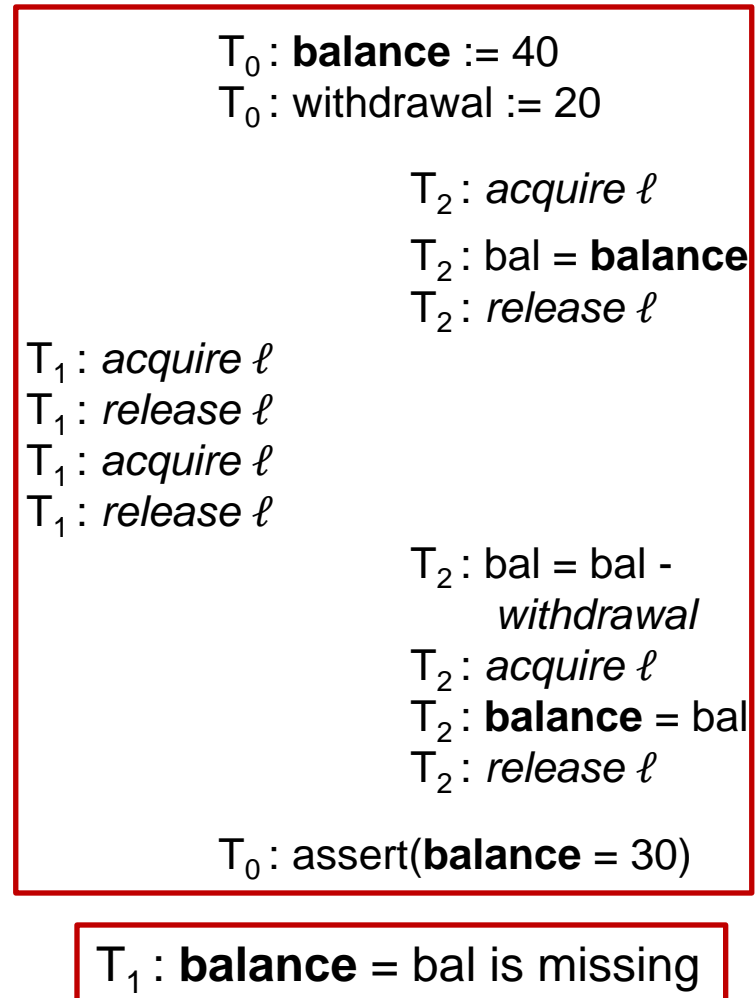
T₀: assert(**balance** = 30)

Ignoring Thread 1 altogether

Interpolants for Debugging Concurrent Bugs



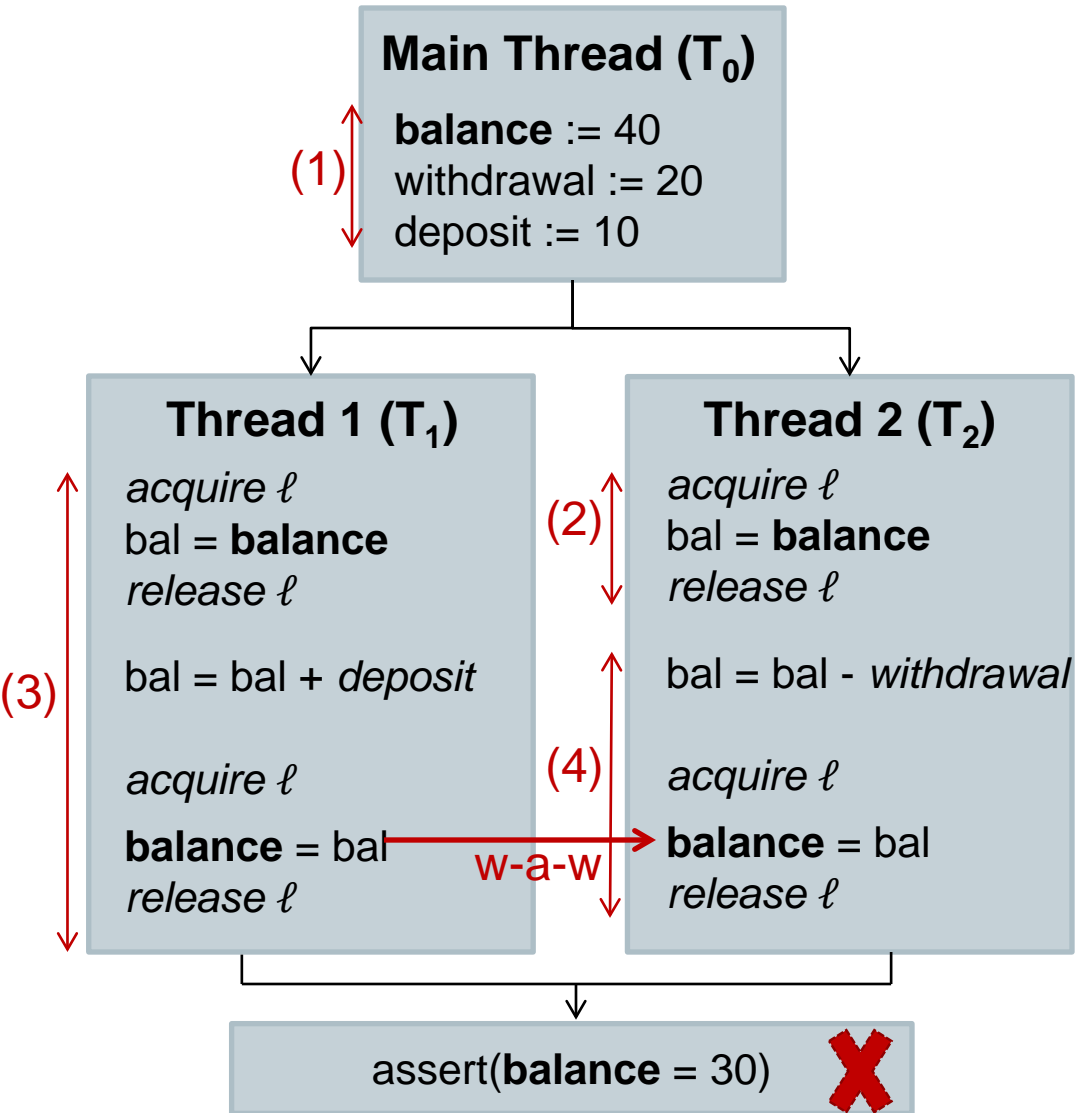
Flow-sensitive Slice



Data dependencies

- Data dependency:
 - Flow of data between statements
- Types of data dependency (in general)
 - Read-after-write
 - $\mathbf{a} = x;$
 - $y = \mathbf{a} + 10;$
 - Write-after-read
 - $x = \mathbf{a};$
 - $\mathbf{a} = y + 10;$
 - Write-after-write
 - ✖ $\mathbf{a} = x + 10;$
 - ✖ $\mathbf{a} = y + 10;$
- Inter-thread data dependencies (in multi threaded programs)
 - Being able to indicate
 - ✖ conflicting accesses or hazards

Interpolants for Debugging Concurrent Bugs



Flow-sensitive Slice

T₀: **balance** := 40
T₀: withdrawal := 20

T₂: acquire ℓ
T₂: bal = **balance**
T₂: release ℓ

T₁: acquire ℓ
T₁: release ℓ
T₁: acquire ℓ
T₁: release ℓ

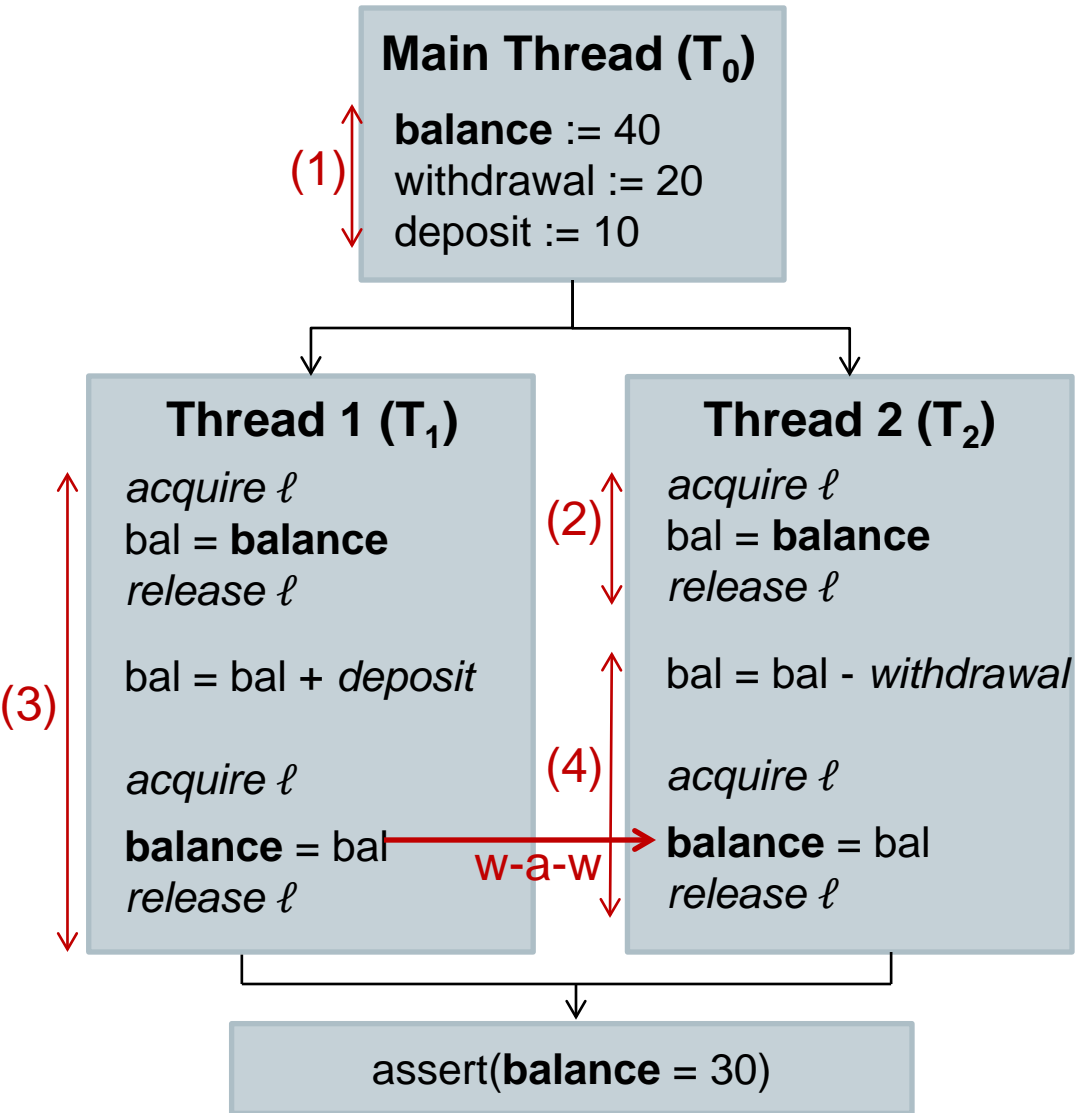
T₂: bal = bal - withdrawal
T₂: acquire ℓ
T₂: **balance** = bal
T₂: release ℓ

T₀: assert(**balance** = 30)

Hazard-sensitive Slices

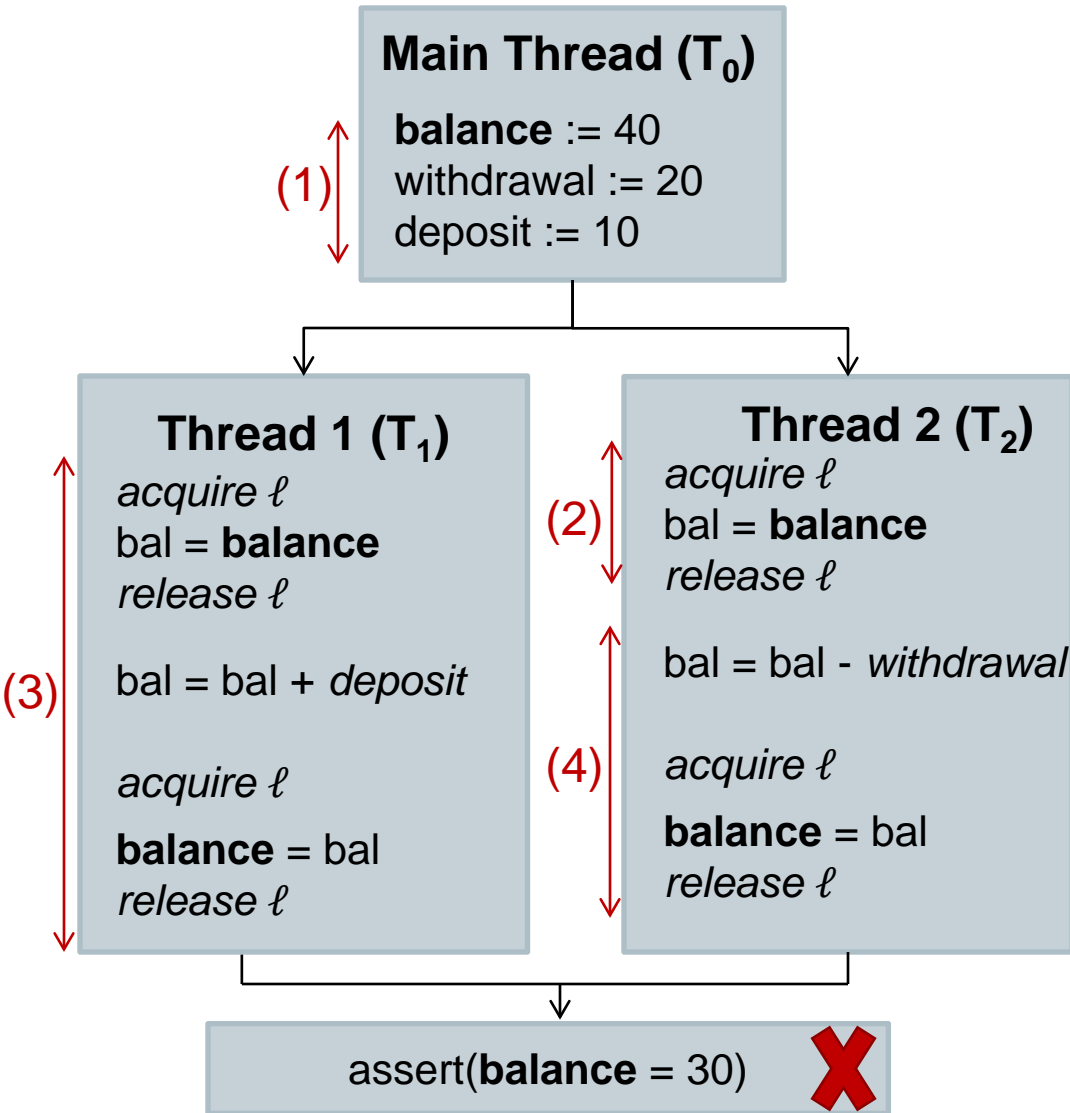
- Encoding inter-thread data dependencies:
 - as implication (using auxiliary variables)
- The resulting slice:
 - Hazard-sensitive slice

Hazard-sensitive Slice

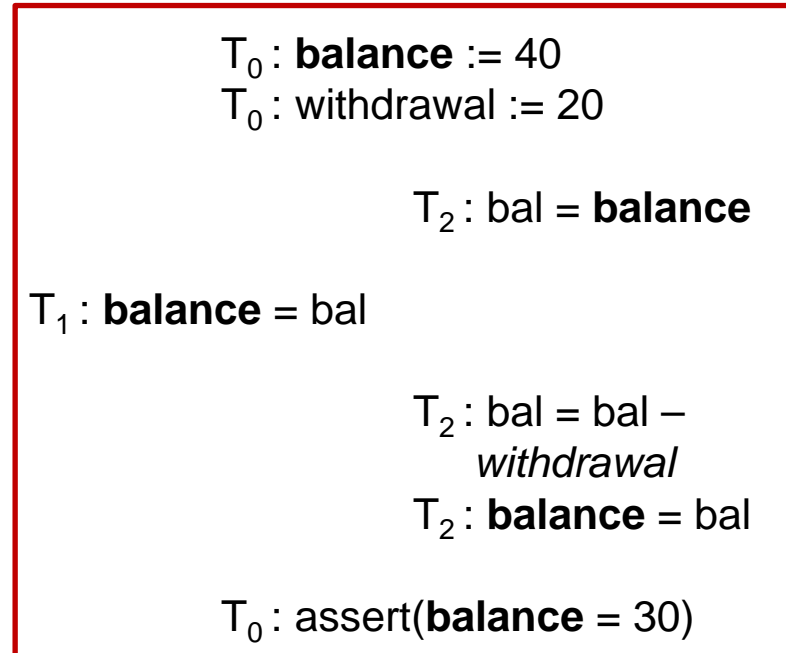


$T_1: v \wedge \text{balance} = \text{bal}$
 $T_2: v \Rightarrow \text{balance} = \text{bal}$

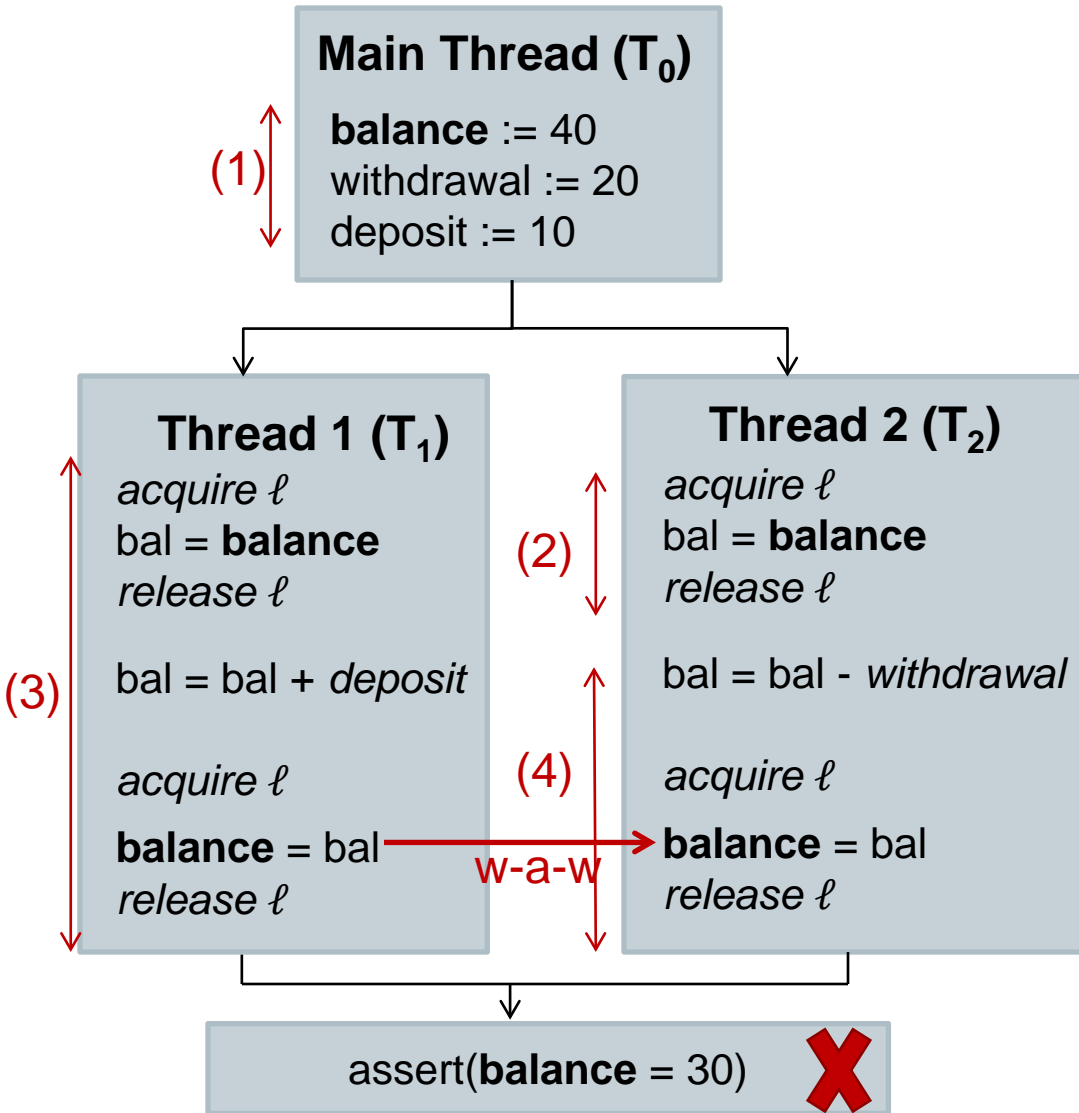
Interpolants for Debugging Concurrent Bugs



Hazard-sensitive Slice



Interpolants for Debugging Concurrent Bugs



Hazard-sensitive Slice

T_0 : **balance** := 40
 $\{balance \leq 40\}$
 T_0 : withdrawal := 20
 $\{balance \leq withdrawal + 20\}$
 T_2 : bal = **balance**
 $\{bal \leq withdrawal + 20\}$
 T_1 : **balance** = bal
 $\{bal \leq withdrawal + 20\}$
 T_2 : bal = bal - withdrawal
 $\{bal \leq 20\}$
 T_2 : **balance** = bal
 $\{balance \leq 20\}$
 T_0 : assert(**balance** = 30)

w-a-w

Fine-Tuning Explanations

- Adding different levels of detail to the explanations
 - Encoding:
 - ✦ control- and inter-thread data-dependency (fs+hs)
 - ✦ control-dependency (fs)
 - ✦ inter-thread data-dependency (hs)
 - ✦ no dependency (\emptyset)
 - Leading to different reductions in number of:
 - ✦ variables
 - ✦ statements

Fine-Tuning Explanations

\emptyset

T_0 : **balance** := 40
 T_0 : withdrawal := 20
 T_2 : bal = **balance**
 T_2 : bal = bal - *withdrawal*
 T_2 : **balance** = bal
 T_0 : assert(**balance** = 30)

fs

T_0 : **balance** := 40
 T_0 : withdrawal := 20
 T_2 : *acquire* ℓ
 T_2 : bal = **balance**
 T_2 : *release* ℓ
 T_1 : *acquire* ℓ
 T_1 : *release* ℓ
 T_1 : *acquire* ℓ
 T_1 : *release* ℓ
 T_2 : bal = bal - *withdrawal*
 T_2 : *acquire* ℓ
 T_2 : **balance** = bal
 T_2 : *release* ℓ
 T_0 : assert(**balance** = 30)

hs



T_0 : **balance** := 40
 T_0 : withdrawal := 20
 T_2 : bal = **balance**
 T_1 : **balance** = bal
 T_2 : bal = bal - *withdrawal*
 T_2 : **balance** = bal
 T_0 : assert(**balance** = 30)

Empirical Evaluation

Quality + Quantity results:

Program	Concurrency bug	Number of traces	Type of slice	Avg. reduction of statements(%)	Avg. reduction of variables
Lock free pool	Linearizability problem	8	fs	61%	34%
			fs+hs	60%	34%
Bank account	Atomicity violation	5	fs+hs	46%	23%
			hs	88%	33%

Conclusion

- A general framework for concurrency bug explanation
 - Interpolation
 - Symbolic execution analysis
 - ✦ Encoding of :
 - Control-dependency
 - Inter-thread data-dependency
 - Implementation
 - ✦ Interpolant computation
 - VERMEER [ICSE15]
 - ✦ Tracing failing concurrent traces
 - ConCrest [FSE13]

Thank you!

