

다중 쓰레드 프로그램의 경쟁상태오류 검출기법 분류

KAIST Provable SW 연구실
홍신, 김문주

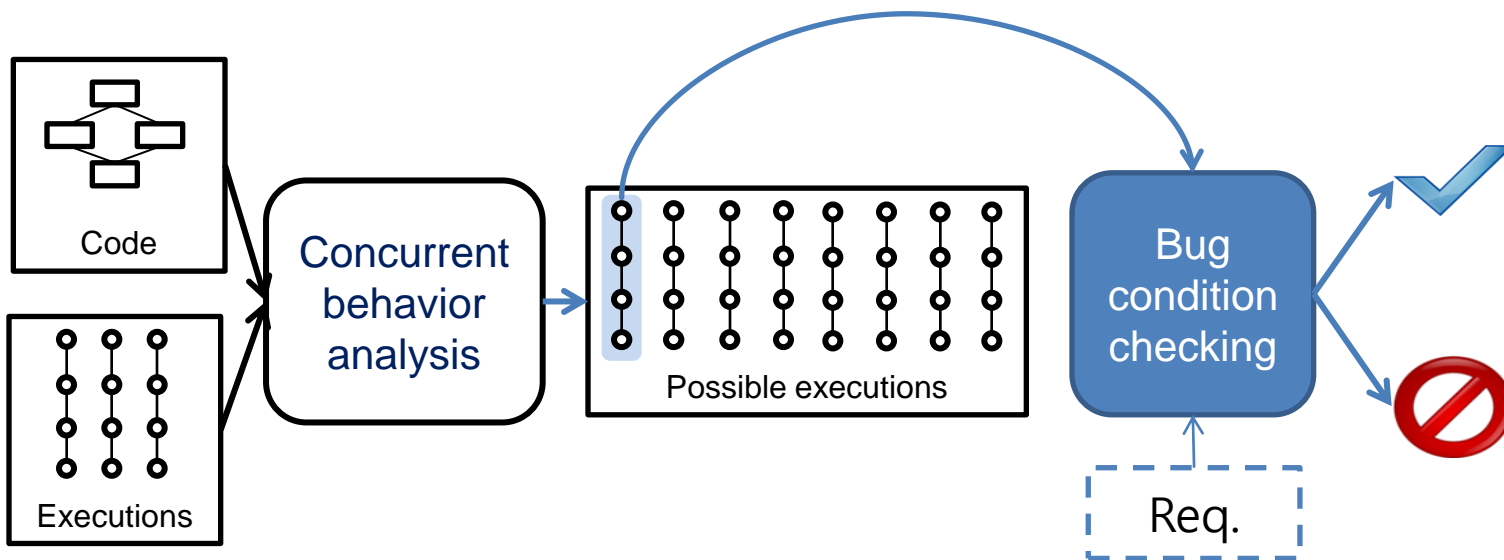
- Many SW applications utilize multi-threaded programming techniques as multi-core hardware become widely spread.
 - Writing correct multi-threaded programs is difficult.
 - Exponential number of execution scenarios
 - Detecting errors by assertion is not effective
- Bug detection techniques specialized for concurrency errors

- Many techniques have evolved
 - Deadlock : [K. M. Chandy *et al.*, TOCS 1983],
[R. Agarwal *et al.* , IBM J. 2010]
- However, for **race bugs**, techniques have used their own definitions and notations without any reliance on a common ground or platform.

	Bug name	Specification	Target program	Bug checking method
Kivati	Atomicity violation	User annotations	C program	Memory access pattern in run-time
LiteRace	Data race	N/A	x86 binary	Execution orders in run-time
Havelund <i>et al.</i>	High-level data race	Analyze code and infer specifications	Java	Relationship between synchronization and variables in run-time execution

→ Develop classification which clarify the relationship between techniques and provide a clear top-down view of race detection techniques

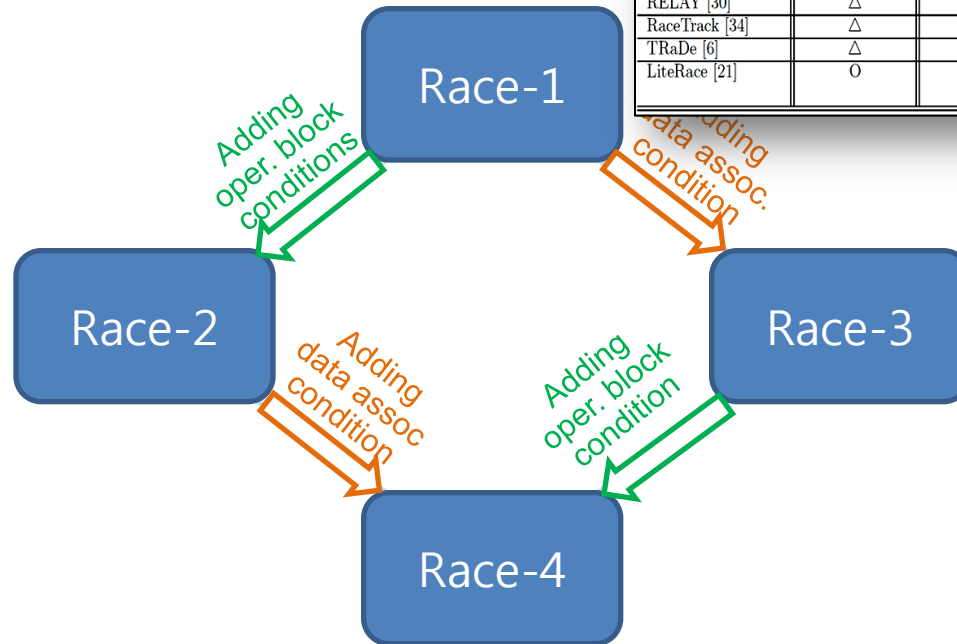
- Provide a formal execution model and specify various bug conditions according to the model
 - Decoupling **what to detect** and **how to detect**
 - **Concurrent behavior analyses**: generate potential executions of a program by information from static/dynamic analyses [F. Chen *et al.* ICSE 2008]
 - **Bug condition checking**: examine a given execution is acceptable/erroneous



- Classify bug conditions according to type of specifications
 - Operation block
 - Data association

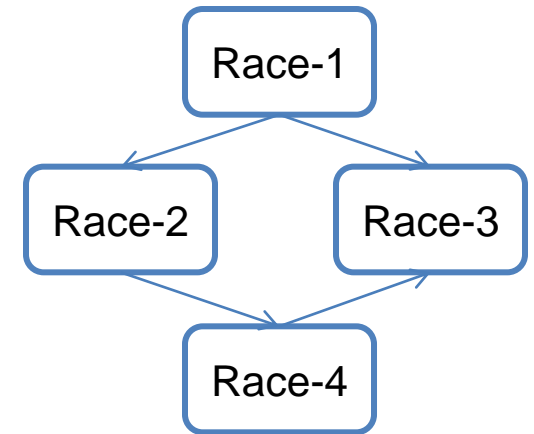
Race-1 techniques

	(A1) <i>thread(p)</i>	(A2) <i>oprd(p)</i>	(A3) <i>conflict(p, q)</i>				(A4) <i>p ≠ q</i>	
	precise thread info.	precise alias info.	W→R	R→W	W→W	R→R	lock	others
Choi et al. [5]	0	0	0	0	0	X	0	thread fork-join
Chord [22]	Δ	Δ	0	0	0	X	0	N/A
Eraser [27]	Δ	0	0	0	0	X	0	N/A
Hybrid Data Race Detection [24]	0	0	0	0	0	X	0	msg. passing
Racer [2]	0	0	0	0	0	X	0	N/A
RacerX [7]	Δ	Δ	0	0	0	X	0	N/A
RccJava [9]	Δ	Δ	0	0	0	0	0	N/A
RELAY [30]	Δ	Δ	0	0	0	X	0	N/A
RaceTrack [34]	Δ	0	0	0	0	X	0	thread fork-join
TRaDe [6]	Δ	0	0	0	0	X	0	N/A
LiteRace [21]	0	0	0	0	0	X	0	atomic operation, msg. passing



Contents

- Execution model for multithreaded program
- Four classes of race detection techniques
 - For each class,
 - bug example
 - bug conditions
 - techniques for checking conditions
- Implications for better race detections



- An execution model of a target program P used for technique D is defined as

$$EM_P(D) = (\underbrace{T, e, \triangleright}_{\text{Program behavior}}, \underbrace{B_{op}, A_{data}}_{\text{Requirements}})$$

- T : a finite set of threads
- e : an interleaved execution
 - a finite sequence of operations p_1, p_2, \dots, p_n where $p_i \in Operation$
 - $thread(p) \in T$
 - $optr(p) \in Operator$
 - $conflict(optr(p), optr(q))$ if the operators of p and q are commutable.
 - $oprd(p) \subseteq V_S$

- An execution model of a target program P used for technique D is defined as

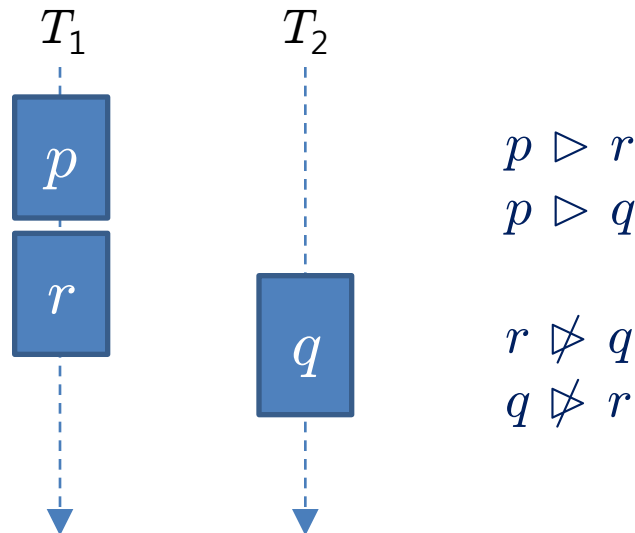
$$EM_P(D) = (T, e, \triangleright, B_{op}, A_{data})$$

- T : a finite set of threads
- e : an interleaved execution
 - a finite sequence of operations p_1, p_2, \dots, p_n where $p_i \in Operation$
 - $thread(p) \in T$
 - $optr(p) \in Operator$
 - $conflict(optr(p), optr(q))$ if the operators of p and q may not be commutable.
 - $oprd(p) \subseteq V_S$

- An execution model of a target program P used for technique D is defined as

$$EM_P(D) = (T, e, \triangleright, B_{op}, A_{data})$$

- Operations in an execution are totally ordered by their start time.
- $\triangleright \subseteq Operation \times Operation$
 $(p, q) \in \triangleright$ if $t_s(p) < t_s(q)$ and $t_e(p) < t_e(q)$



- An execution model of a target program P used for technique D is defined as

$$EM_P(D) = (T, e, \triangleright, B_{op}, A_{data})$$

$B_{op} = \{b_1, b_2, \dots\}$ where $b_i : Operation \times Operation$

An execution of an atomic code region corresponds to a sequence of operation, *operation blocks*.

$(p, q) \in b_i$ indicates that p and q are in the same operation block.

```
class BankAccount {
  int balance ;
  ...
  void withdraw(int amount) {
    if (getBalance() >= amount) {
      lock(m) ;
      balance = balance - amount ;
      unlock(m) ;
    }
  }
}
```

- An execution model of a target program P used for technique D is defined as

$$EM_P(D) = (T, e, \triangleright, B_{op}, A_{data})$$

$A_{data}: V_S \times V_S$ where V_S is a set of shared variables

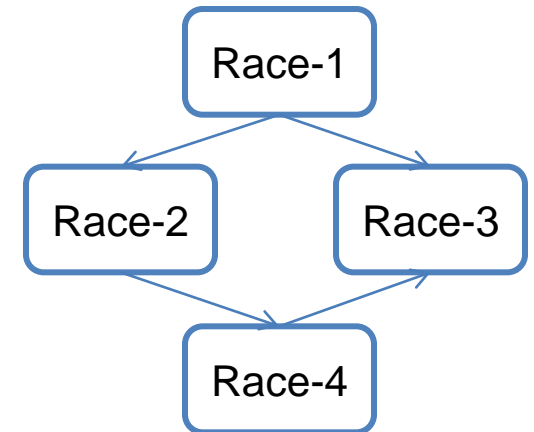
Frequently, variables in a composite data structure have dependencies and there exists relations/invariances on these variables.

```
class BankAccount {  
    int balance ;  
    int debt ;  
    /* invariant: (debt == 0    ^ balance == 0) ∨  
                  (debt > 0    ^ balance == 0) ∨  
                  (balance > 0 ^ debt == 0) */
```

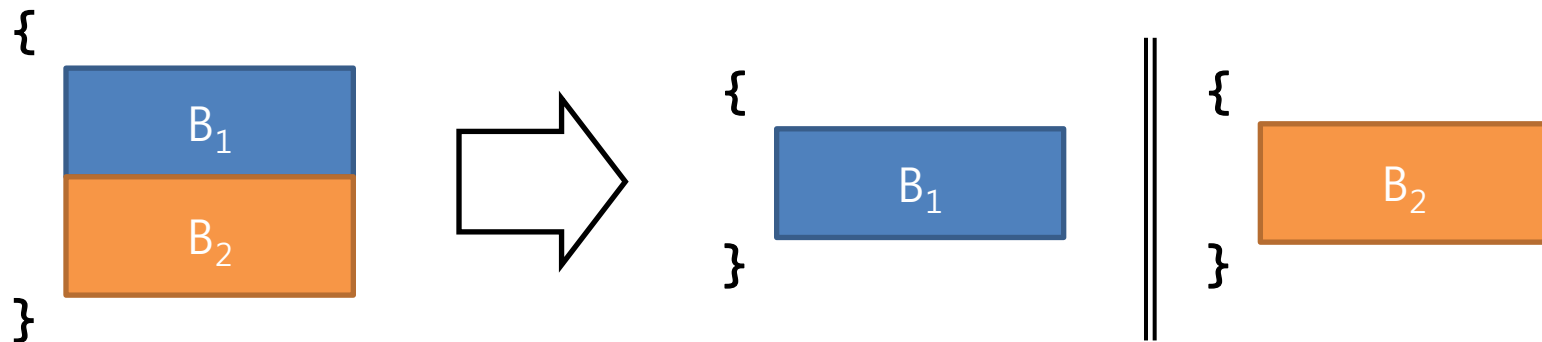
→ (balance, debt) ∈ A_{data}
(debt, balance) ∈ A_{data}

Contents

- Execution model for multithreaded program
- Four classes of race detection techniques
 - For each class,
 - bug example
 - bug conditions
 - techniques for checking conditions
- Implications for better race detections



In parallelization,

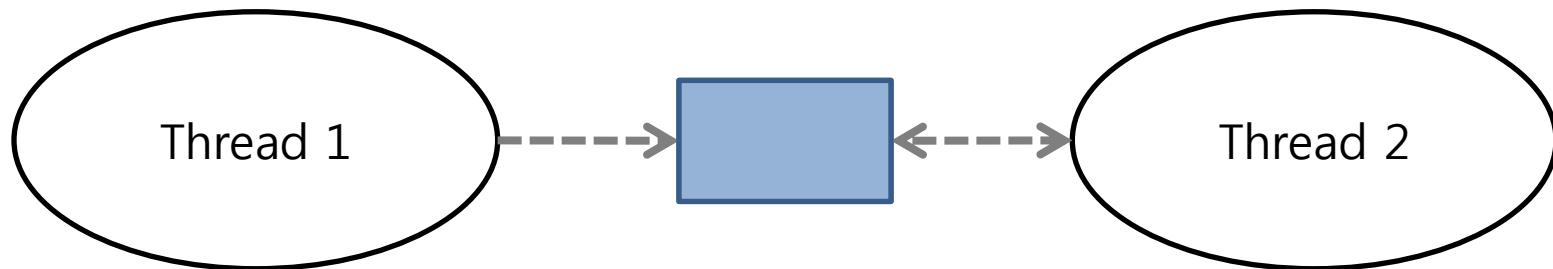


A sufficient condition for safe parallelization:

$$\text{MemoryRead}(B_1) \cap \text{MemoryWrite}(B_2) = \emptyset \text{ and} \\ \text{MemoryWrite}(B_1) \cap \text{MemoryRead}(B_2) = \emptyset$$

- In “What are race conditions?” [R. H. Netzer et al., LOPLAS 1992], a *data race* $\langle a, b \rangle$ over an execution exists if and only if
 - (1) a data conflict exists in a program between a and b ,
 - (2) no temporal ordering between a and b .

- In [Savage et al., SOSP 1997],
a *data race* occurs when there exists two operations
 - (1) executed by two concurrent threads,
 - (2) access a shared variable
 - (3) at least one access is write,
 - (4) no explicit mechanism to coordinate their execution order



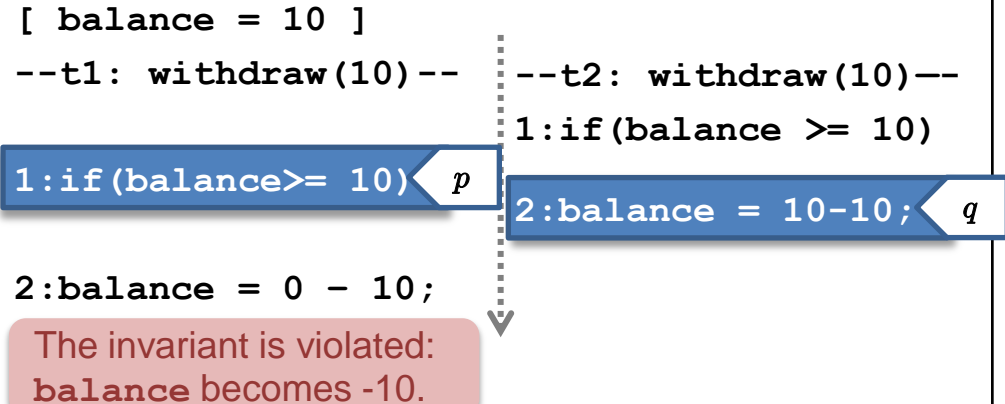
Race-1: Data-race

(3/6)

- Buggy program code

```
class BankAccount_A {
  int balance;
  // balance should be non-negative
  void withdraw(int x) {
1:  if (balance >= x) {
2:    balance = balance - x; }
    ... }
}
```

- Error scenario



- A target program P has a *race-1 bug* if there is an execution σ such that σ has two operations p and q that satisfy the following conditions:

- (A1) $thread(p) \neq thread(q)$
- (A2) $oprd(p) \cap oprd(q) \neq \emptyset$
- (A3) $conflict(optr(p), optr(q))$
- (A4) $p \not\prec q \wedge q \not\prec p$

p and q are commutable?

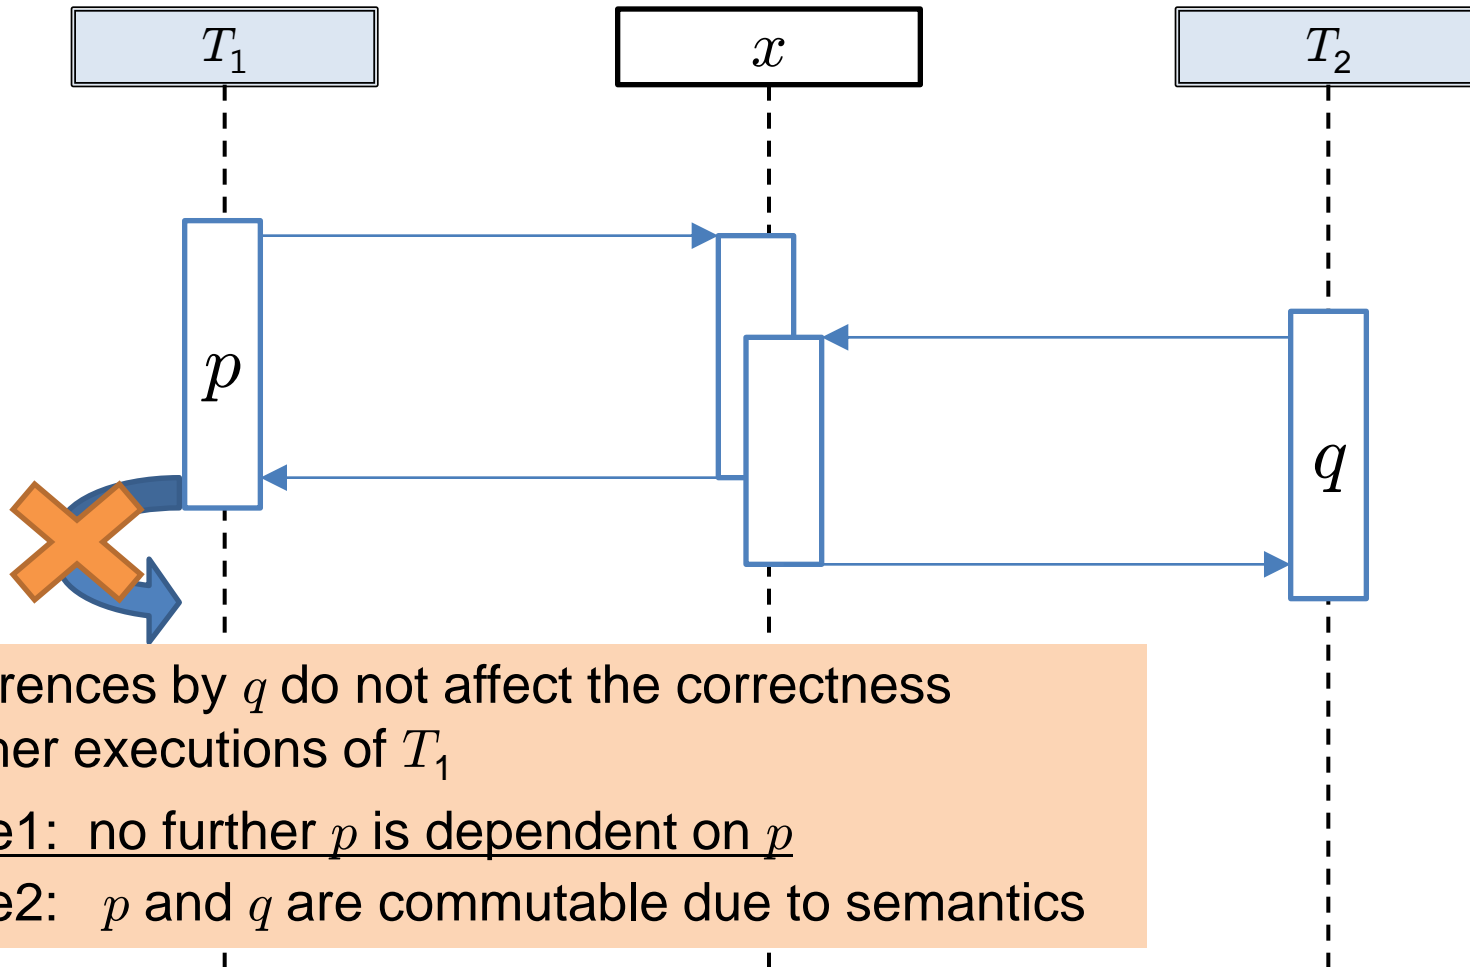
Race-1: Data-race

(4/6)

- Race-1 Detection Techniques

	(A1)	(A2)	(A3)				(A4)
	$thread(p) \neq thread(q)$	$oprd(p) \cap oprd(q) \neq \emptyset$	$conflict(optr(p), optr(q))$				$p \neq q$
			W-R	R-W	W-W	R-R	
Choi <i>et al.</i>	concrete thread id.	concrete mem. addr.	○	○	○	X	tracking lock acq/rel, thread fork/join
Eraser	check shared or non-shared	concrete mem. addr.	○	○	○	X	tracking lock acq/rel
Hybrid data race detection	concrete thread id.	concrete mem. addr.	○	○	○	X	tracking lock acq/rel, message send/receive
Racer	concrete thread id.	concrete mem. addr.	○	○	○	X	tracking lock acq/rel
RaceTrack	check shared or non-shared	concrete mem. addr.	○	○	○	X	tracking lock acq/rel, thread fork/join
TRaDe	check shared or non-shared	concrete mem. addr.	○	○	○	X	tracking lock acq/rel
LiteRace	concrete thread id.	concrete mem. addr.	○	○	○	X	tracking lock acq/rel, message passing, atomic instructions
Chord	static approx.	static alias analysis	○	○	○	X	tracking lock acq/rel
RacerX	static approx.	static alias analysis	○	○	○	X	tracking lock acq/rel
RELAY	static approx.	static alias analysis	○	○	○	X	tracking lock acq/rel
RccJava	static approx.	static alias analysis	○	○	○	○	tracking lock acq/rel

- Limitations – false positive



Interferences by q do not affect the correctness of further executions of T_1

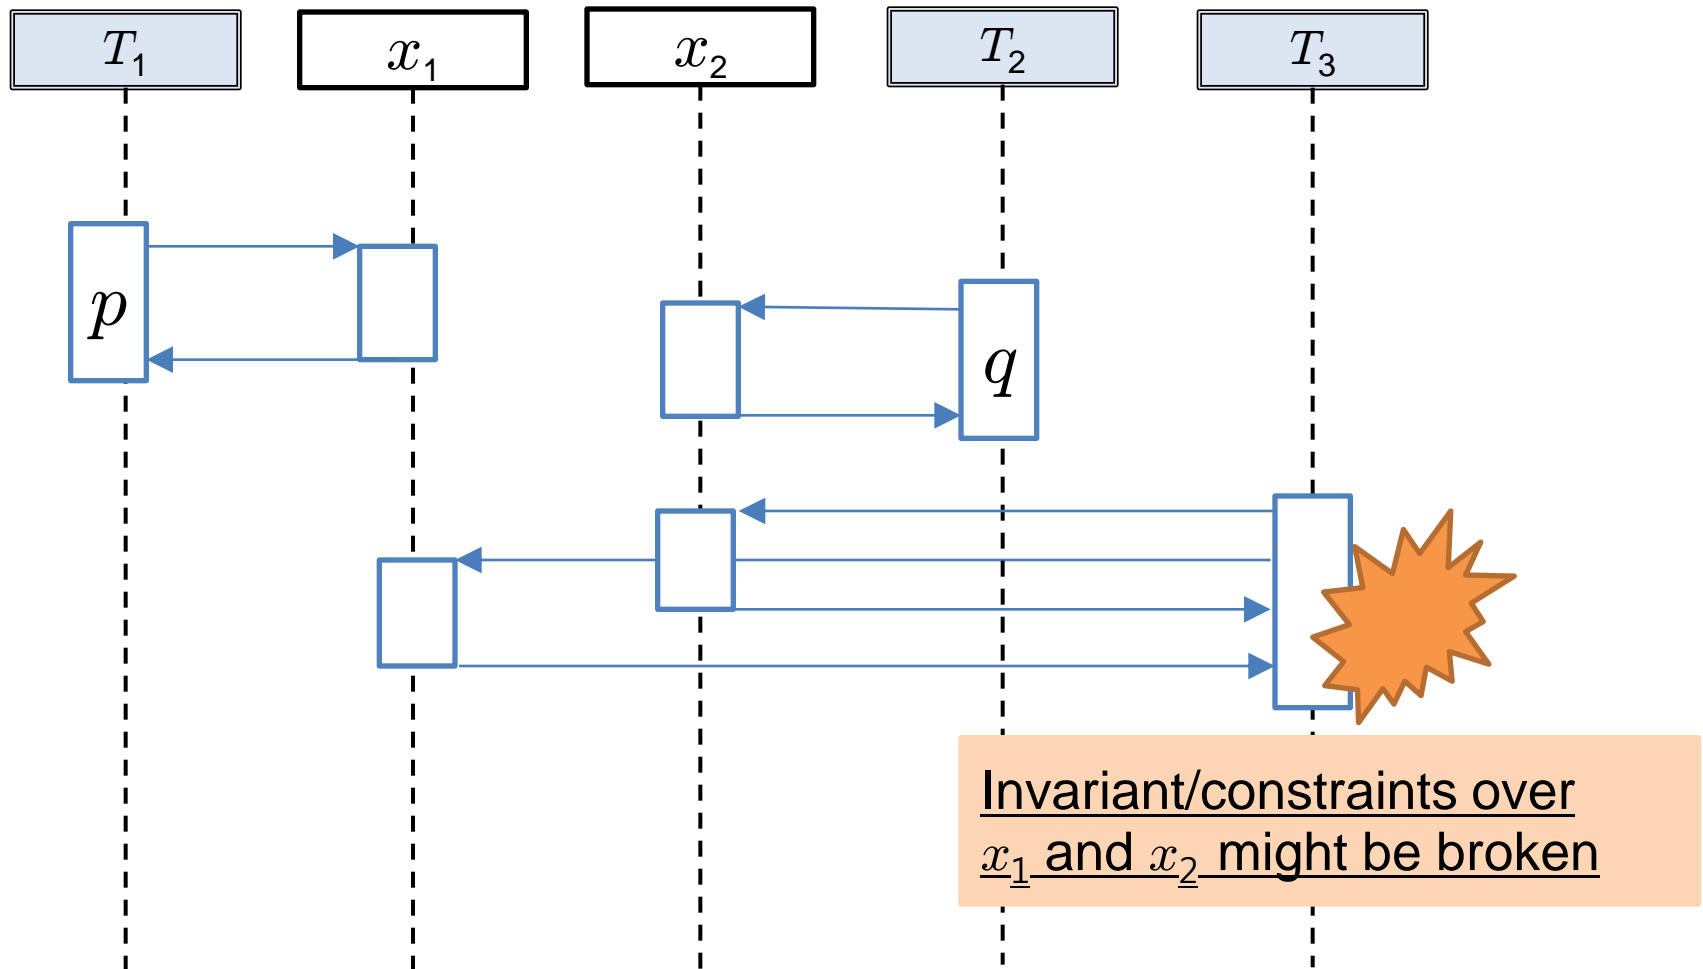
Case1: no further p is dependent on p

Case2: p and q are commutable due to semantics

Race-1: Data-race

(6/6)

- Limitations – false negative



Race-2: Atomic Block Violation (1/5)

- Race-2 techniques check whether or not an **operation block** can interfere with another thread.

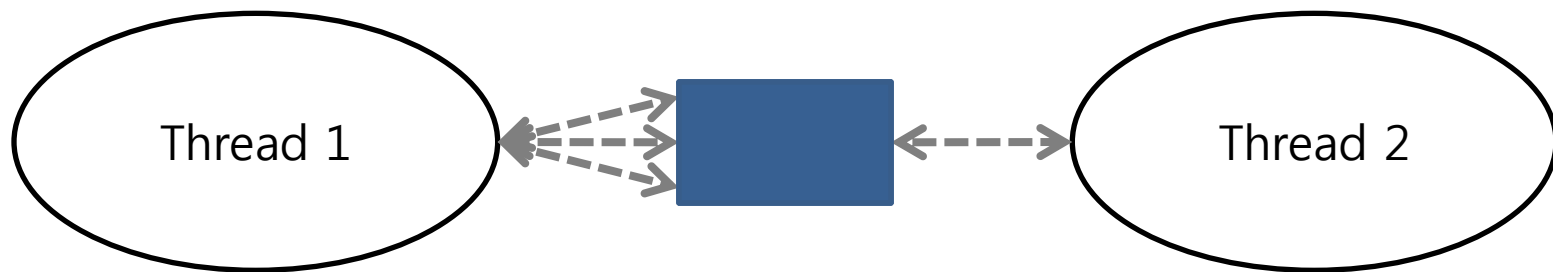
In [Savage et al., SOSP 1997],

From a previous slide for race-1

a *data race* occurs when there exists two operations

- (1) executed by two concurrent threads,
- (2) access a shared variable
- (3) **at least one access is write,**
- (4) no explicit mechanism to coordinate their execution order

a sequence of accesses
(operation block)



Race-2: Atomic Block Violation (2/5)

- A target program P has a *race-2 bug* if there is an execution σ such that σ has three operations p , p' , and q that satisfy the following conditions:

(B1) $thread(p) \neq thread(q)$

(B2) $oprd(p) \cap oprd(q) \neq \emptyset$

(B3) $conflict(optr(p), optr(q))$

(B4) $\exists b_i \in B_{op}. ((p, p') \in b_i)$

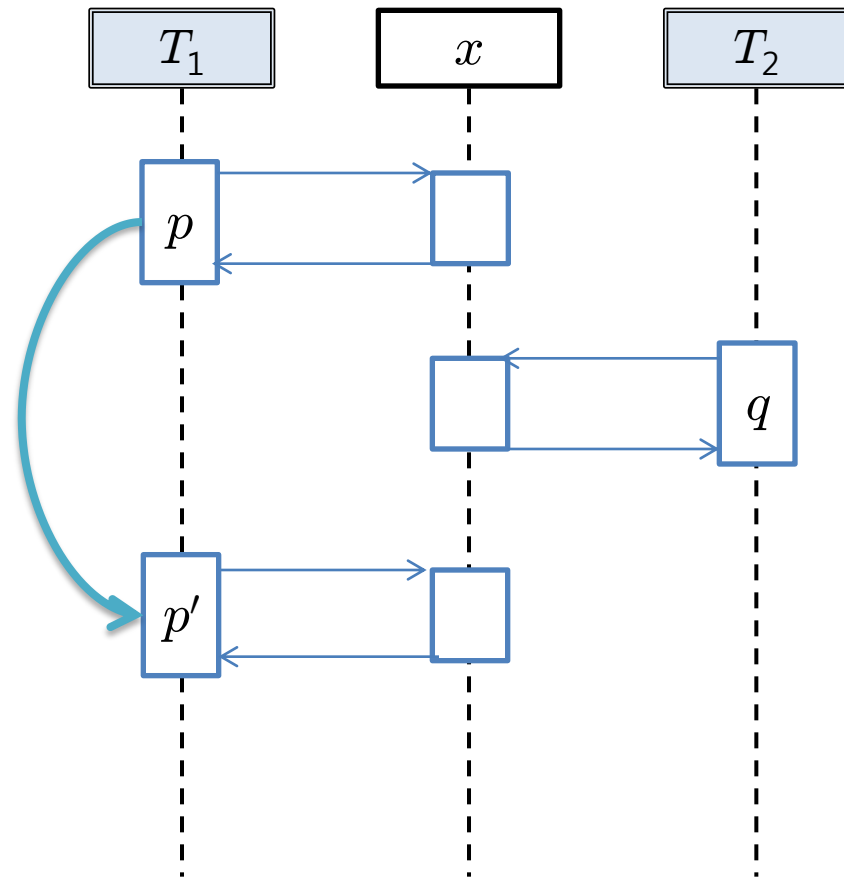
(B5) $oprd(p) \cap oprd(p') \neq \emptyset$

(B6) $conflict(optr(p), optr(p'))$

(B7) $oprd(p') \cap oprd(q) \neq \emptyset$

(B8) $conflict(optr(p'), optr(q))$

(B9) $q \not\triangleright p \wedge p' \not\triangleright q$



Race-2: Atomic Block Violation (3/5)

- Race-2 bug example:

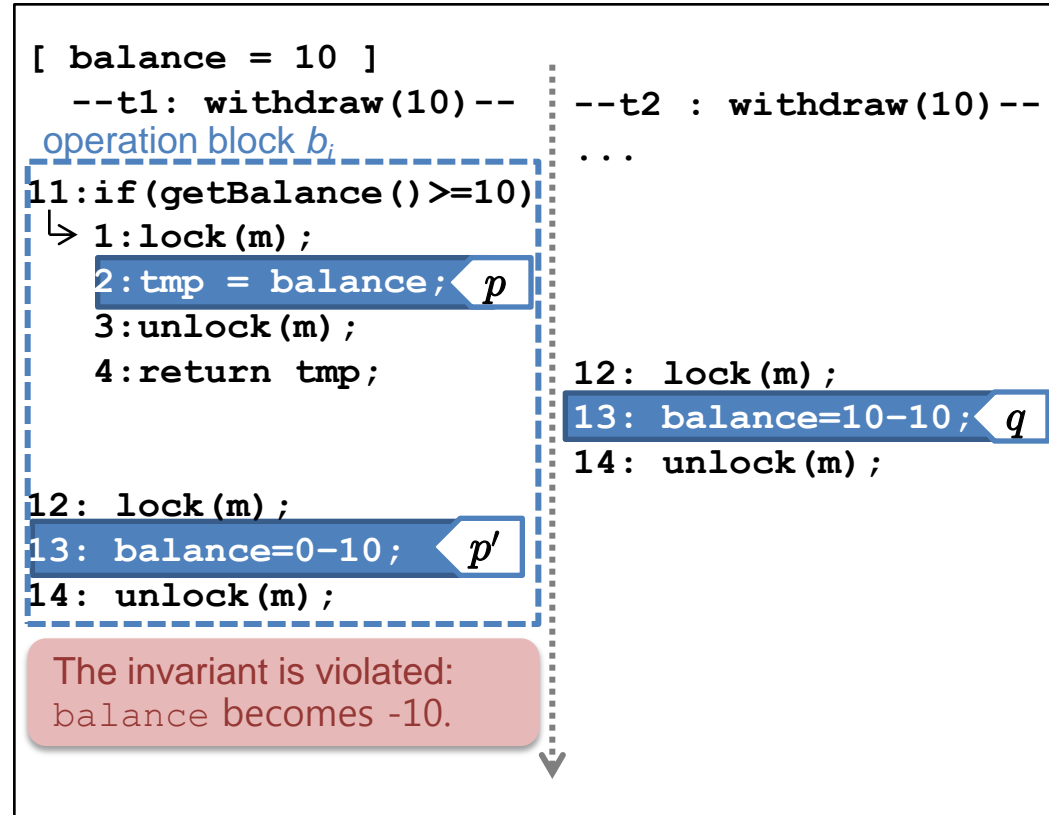
- Buggy program code

```
class BankAccount_B {
    Lock m;
    int balance;
    // balance should be non-negative
    // balance should be synchronized by m

    int getBalance() {
        int tmp;
1:  lock(m);
2:  tmp = balance;
3:  unlock(m);
4:  return tmp; }

    void withdraw(int x) {
        /*@atomic region begins*/
11: if (getBalance() >= x) {
12:     lock(m);
13:     balance = balance - x;
14:     unlock(m); }
        /*@atomic region ends*/
        ...
    }
}
```

- Race-2 error scenario



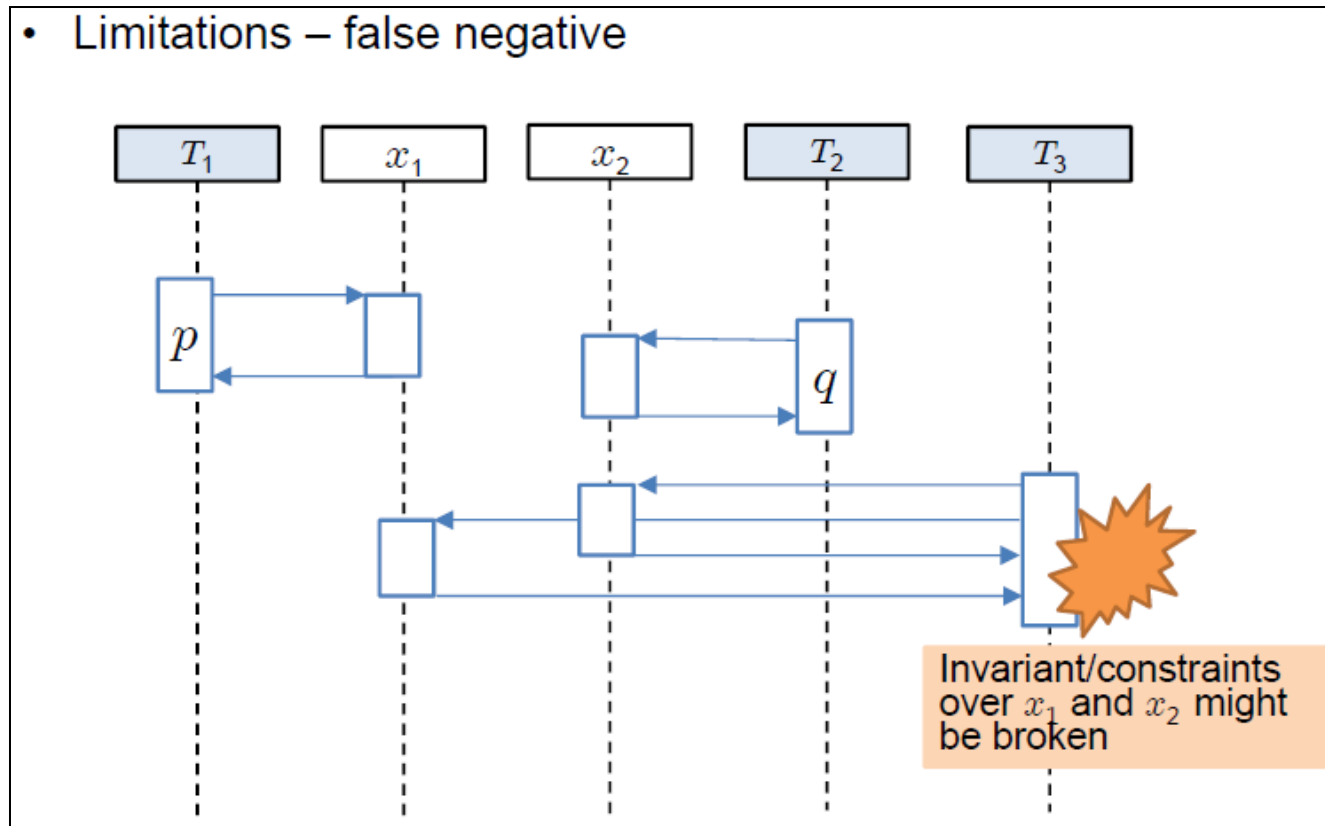
Race-2: Atomic Block Violation (4/5)

- Race-2 Detection Techniques

	(B1)	(B2, B4, B7)	(B6)				(B3, B8)				(B9)
	<i>thread()</i>	<i>oprd()</i>	<i>conflict(optr(p), optr(q))</i> where <i>thread(p) ≠ thread(q)</i>				<i>conflict(optr(p), optr(q))</i> where <i>thread(p) = thread(q)</i>				⚡
			W-R	R-W	W-W	R-R	W-R	R-W	W-W	R-R	
Atomic-Aid	concrete thread id.	concrete mem. addr.	○	○	Cond.	X	○	○	○	○	tracking lock/unlock
AtomRace	concrete thread id.	concrete mem. addr.	○	○	Cond.	X	○	○	○	○	tracking lock/unlock
AVIO	concrete thread id.	concrete mem. addr.	○	○	Cond.	X	○	○	○	○	total execution order
Block-based algorithm	concrete thread id.	concrete mem. addr.	○	○	Cond.	X	○	○	○	○	tracking lock/unlock, message passing
Commit-node	concrete thread id.	concrete mem. addr.	○	○	○	X	○	○	○	○	tracking lock/unlock message passing
HAVE	concrete thread id.	concrete mem. addr.	○	○	○	X	○	○	○	○	tracking lock/unlock message passing.
Kivati	concrete thread id.	concrete mem. addr.	○	○	Cond.	X	○	○	○	○	total execution order
SVD	concrete thread id.	concrete mem. addr.	○	○	○	○	X	X	○	○	total execution order
PENELOPE	concrete thread id.	concrete mem. addr.	○	○	○	X	○	○	○	○	tracking lock/unlock
Velodrome	concrete thread id.	concrete mem. addr.	○	○	○	X	○	○	○	○	tracking lock/unlock, message passing
Atomizer	static apprx.	alias analysis	○	○	○	○	○	○	○	○	tracking lock/unlock

Race-2: Atomic Block Violation (5/5)

- Limitation: false-negative



Race-3: Data Assoc. Violation (1/4)

- A unit of data can be located in two or more distinct variables.
- Race-3 detection techniques look for inconsistent updates of associated variables. [K. Havelund VVEIS03, S. Lu SOSP07, F. Tip ICSE08]

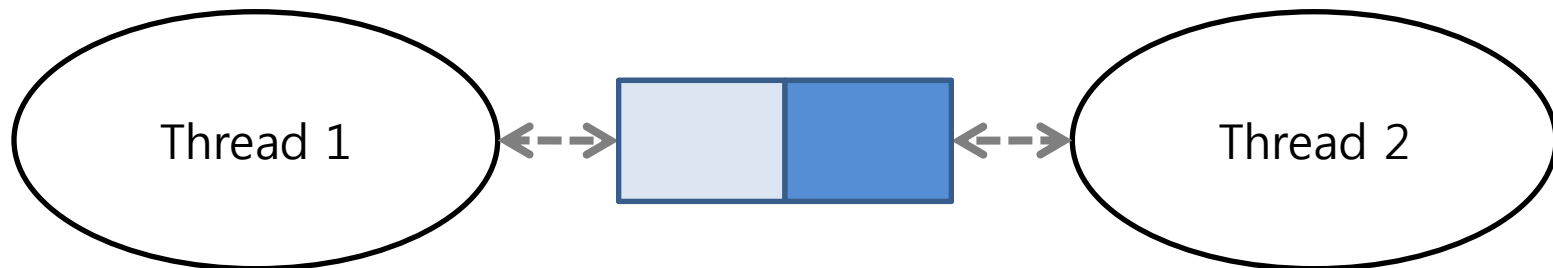
In [Savage et al., SOSP 1997],

From a previous slide for Race-1

a *data race* occurs when there exists two operations

- (1) executed by **two concurrent threads**,
- (2) access **a shared variable**
- (3) **at least one access is write**,
- (4) **no explicit mechanism to coordinate their execution order**

memory area,
set of variables
(associated variables)



Race-3: Data Assoc. Violation (2/4)

- Race-3 bug condition:

$$(C1) \quad \textit{thread}(p) \neq \textit{thread}(q)$$

$$(C2) \quad \exists v_1, v_2 \in V_S. (v_1 \in \textit{oprd}(p) \wedge v_2 \in \textit{oprd}(q) \wedge (v_1, v_2) \in A_{data})$$

$$(C3) \quad \textit{conflict}(\textit{optr}(p), \textit{optr}(q))$$

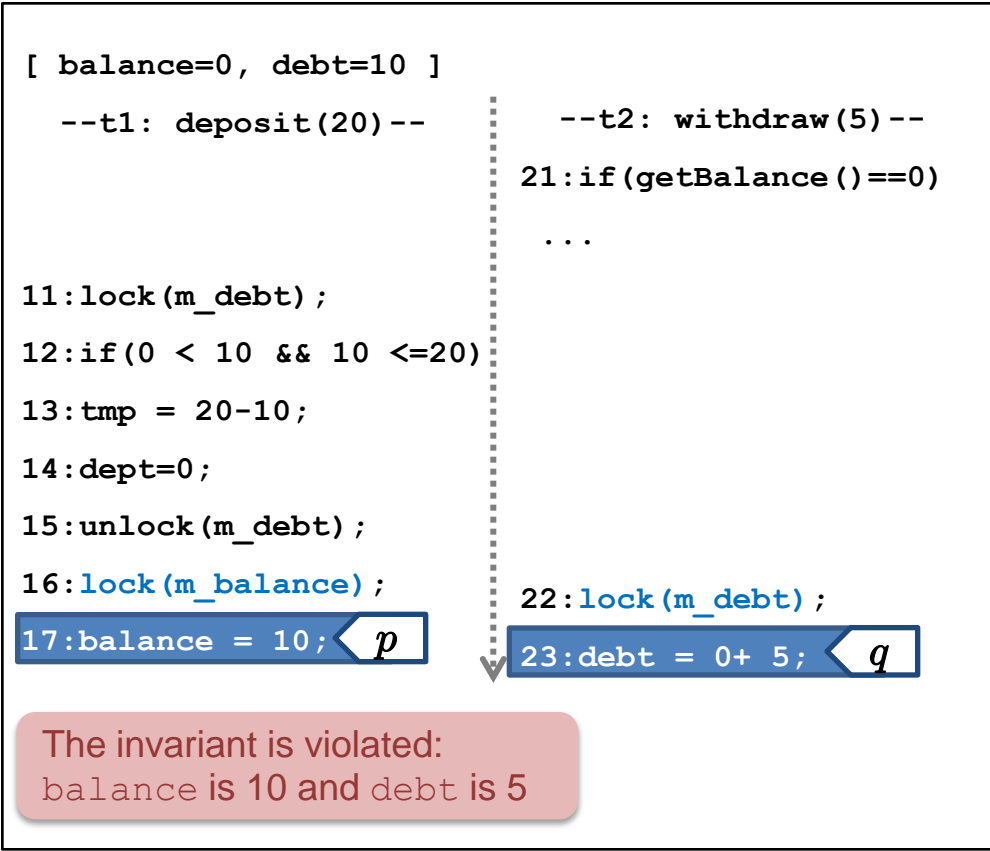
$$(C4) \quad p \not\triangleright q \wedge q \not\triangleright p$$

Race-3: Data Assoc. Violation (3/4)

- Example

```
class BankAccount_C {
  int balance ;
  int debt ;
  /* Invariant:
    (balance == 0  $\wedge$  debt == 0)  $\vee$ 
    (debt > 0  $\rightarrow$  balance == 0)  $\vee$ 
    (balance > 0  $\rightarrow$  debt == 0) */
   $\rightarrow$  (balance, debt)  $\in A_{data} \wedge$ 
      (debt, balance)  $\in A_{data}$ 

  Lock m_balance ;
  Lock m_debt ;
```



Race-3: Data Assoc. Violation (4/4)

- Race-3 detection techniques

	(C1)	(C2)			(C3)	(C4)	
	<i>thread()</i>	A_{data}			<i>conflict()</i>	\nexists	
		Type	transitive	symmetric	Source		
MultiRace	Concrete thread id.	$[a_1, a_2, \dots, a_n]$	O	O	User annotation	W-R, R-W, W-W	tracking lock/unlock
MUVI-Eraser	Heuristics	$\langle (a_1, t_1), (a_2, t_2) \rangle$ where $t_1, t_2 \in \{rd, wr, rd\&wr\}$	X	X	Inferring from executions	W-R, R-W, W-W	tracking lock/unlock
Object data race detection	Static apprx.	$\{ a_1, a_2, \dots, a_n \}$	O	O	User annotation	W-R, R-W, W-W	tracking lock/unlock

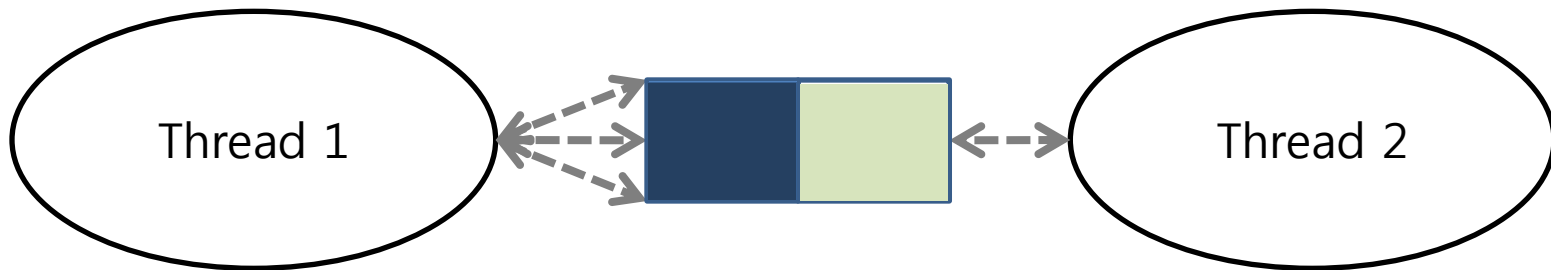
- Race-4 techniques utilize both operation block and data association together to reduce false positives and false negatives.

In [Savage et al., SOSP 1997],

From a previous slide for Race-1

a *data race* occurs when there exists two operations

- (1) executed by **two concurrent threads**,
- (2) access **a shared variable**
- (3) **at least one access is write**,
- (4) **no explicit mechanism to coordinate their execution order**



- Race-4 bug conditions:

Race-B conditions

Race-C conditions

$$(D1) \quad thread(p) \neq thread(q)$$

$$(D2) \quad \exists v_1, v_2 \in V_S. (v_1 \in oprd(p) \wedge v_2 \in oprd(q) \wedge (v_1, v_2) \in A_{data})$$

$$(D3) \quad conflict(optr(p), optr(q))$$

$$(D4) \quad \exists v_3, v_4 \in V_S. (v_3 \in oprd(p') \wedge v_4 \in oprd(q) \wedge (v_3, v_4) \in A_{data})$$

$$(D5) \quad conflict(optr(p'), optr(q))$$

$$(D6) \quad \exists b_i \in B_{op}. (p, p') \in b_i$$

$$(D7) \quad \exists v_5, v_6 \in V_S. (v_5 \in oprd(p') \wedge v_6 \in oprd(q) \wedge (v_5, v_6) \in A_{data})$$

$$(D8) \quad conflict(optr(p), optr(p'))$$

$$(D9) \quad q \not\triangleright p \wedge p' \not\triangleright q$$

Race-4

(3/4)

- Buggy program

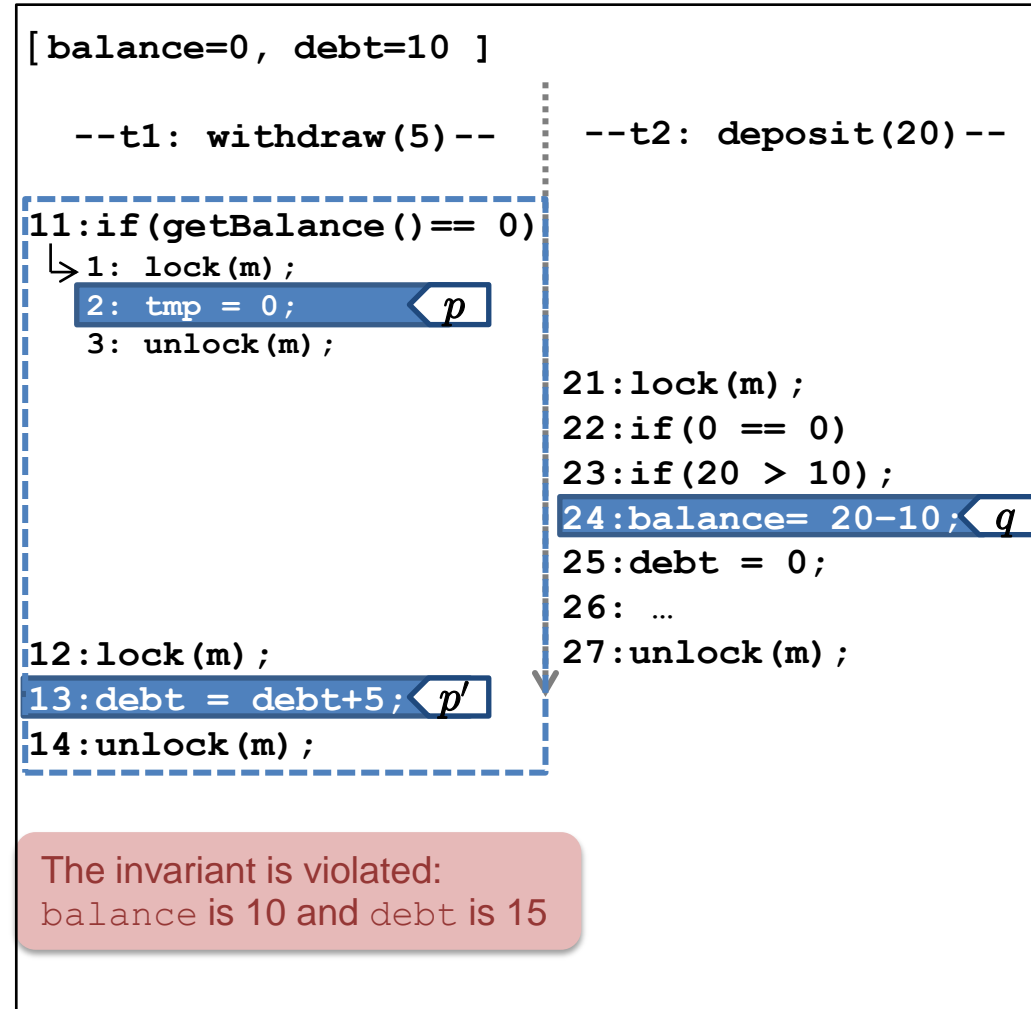
```
class BankAccount_D {
    Lock m;
    int balance, debt;
    /* (balance, debt) ∈ Adata
       (debt, balance) ∈ Adata */

    int getBalance(int x) {
        int tmp;
    1: lock(m);
    2: tmp = balance;
    3: unlock(m);
    4: return tmp;
    }

    int withdraw(int x) {
    11: if (getBalance() == 0) {
    12:     lock(m);
    13:     debt = debt + x;
    14:     unlock(m);
    }
    }

    int deposit(int x) {
    21: lock(m);
    22: if (balance == 0) {
    23:     if (x > debt) {
    24:         balance = x - debt;
    25:         debt = 0;
    26:         ...
    27:     }
    }
    }
}
```

- Error execution scenario

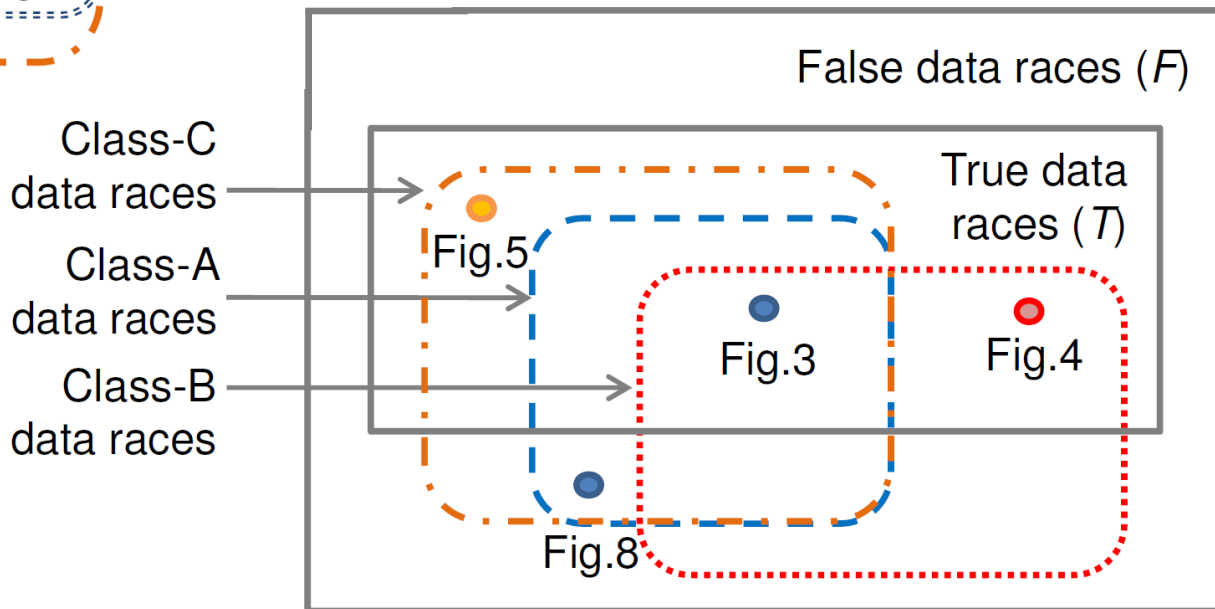
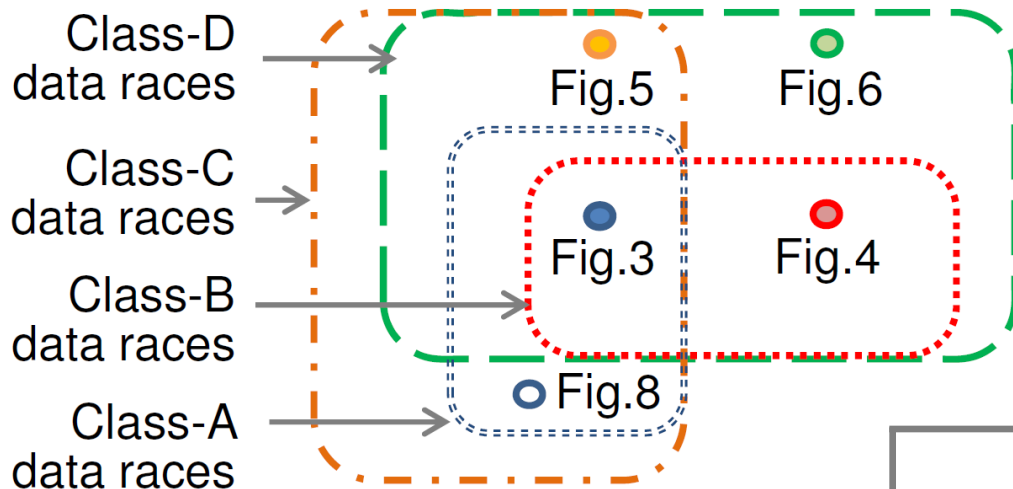


- Race-4 detection techniques

	(D1)	(D2, D4, D7)			(D3,D5)	(D6,D8)	(D9)
	<i>thread()</i>	A_{data}			<i>conflict()</i> where $thread(p) \neq thread(q)$	<i>conflict()</i> where $thread(p) \neq thread(q)$	\nexists
		Type	transi- tive	symm- etric			
Atomic-Set serializability	Concrete thread id.	$\langle a_1, a_2 \rangle$	O	O	W-R, R-W, W-W	W-R, R-W, W-W, R-R	total execution order
ColorSafe	Concrete thread id.	$\langle a_1, a_2 \rangle$	O	O	W-R, R-W, W-W	W-R, R-W, W-W, R-R	total execution order
Method-consistency	Static apprx.	$\langle (a_1, t_1), (a_2, t_2) \rangle$ where $t_1, t_2 \in \{read, update\}$	O	O	W-R, R-W, W-W, R-R	W-R, R-W, W-W, R-R	tracking lock/unlock
MUVI-AVIO	Concrete thread id.	$\langle (a_1, t_1), (a_2, t_2) \rangle$ where $t_1, t_2 \in \{rd, wr, rd\&wr\}$	X	X	W-R, R-W, W-W(cond),	W-R, R-W, W-W, R-R	total execution order
View-consistency	Concrete thread id.	$\langle (a_1, t_1), (a_2, t_2) \rangle$	O	O	W-R, R-W, W-W, R-R	W-R, R-W, W-W, R-R	tracking lock/unlock

Implications to Better Race Detection (1/2)

- Relations in four class of race detections



Implications to Better Race Detection (2/2)

- Static analyses can be applied for much precise race detections
 - Only few work use static analyses for inferring/checking operation block and data associations

Theorem 1. Let $G^{du} = (G, \Sigma, D, U)$ be a def/use graph with $G = (N, E)$, and let $u, v \in N$. If u is semantically dependent on v and this semantic dependence is finitely demonstrated, then u is syntactically dependent on v .²

Theorem 1 is significant because it shows that, given appropriate definitions of control and data dependence, syntactic dependence is a necessary condition for (finitely demonstrated) semantic dependence. Thus, the theorem provides justification for algorithms that use syntactic dependence to approximate semantic dependence. We refer to this desirable relationship between syntactic and semantic dependence as the “*syntactic–semantic* relationship”.

M. J. Harrold, G. Rothermel, and S. Sinha. “Computation of Interprocedural Control Dependence” Proceedings of the ACM International Symposium on Software Testing and Analysis, March 1998.