

Catching Race Conditions

An Extremely Difficult Task

- ❑ *Statically* detecting race conditions in a program using multiple semaphores is **NP-hard**.
- ❑ Thus, no efficient algorithms are available. We have to use human debugging skills.
- ❑ It is virtually impossible to catch race conditions *dynamically* because hardware must examine *every* memory access.
- ❑ So, we shall use a few examples to illustrate some subtle race conditions.

Problem Statement

- ❑ Two groups, **A** and **B**, of processes *exchange messages*.
- ❑ Each process in **A** runs a function $T_A()$, and each process in **B** runs a function $T_B()$.
- ❑ Both $T_A()$ and $T_B()$ have an infinite loop and never stop.
- ❑ In the following, *we show execution sequences that cause race conditions. You can always find a correct execution sequence without race conditions.*

Processes in group A

```
T_A()
```

```
{
```

```
while (1) {
```

```
    // do something
```

```
    .....  
    : Ex. Message :  
    .....  
    // do something
```

```
}
```

```
}
```

Processes in group B

```
T_B()
```

```
{
```

```
while (1) {
```

```
    // do something
```

```
    .....  
    : Ex. Message :  
    .....  
    // do something
```

```
}
```

```
}
```

What is *Exchange Message*?

- ❑ When a process in **A** makes a message available, it can continue only if it receives a message from a process in **B** who has successfully retrieves A's message.
- ❑ Similarly, when a process in **B** makes a message available, it can continue only if it receives a message from a process in **A** who has successfully retrieves **B**'s message.
- ❑ *How about exchanging business cards?*

Watch for Race Conditions

- Suppose process A_1 presents its message for B to retrieve. If A_2 comes for message exchange before B retrieves A_1 's, will A_2 's message overwrites A_1 's?
- Suppose B has already retrieved A_1 's message. Is it possible that when B presents its message, A_2 picks it up rather than A_1 ?
- Thus, the messages between A and B must be well-protected to avoid race conditions.

First Attempt

```
sem A = 0, B = 0;  
int Buf_A, Buf_B;
```

```
T_A()
```

```
{
```

```
    int V_a;
```

```
    while (1) {
```

```
        V_a = ..;
```

```
        B.signal();
```

```
        A.wait();
```

```
        Buf_A = V_a;
```

```
        V_a = Buf_B;
```

```
}
```

```
T_B()
```

```
{
```

```
    int V_b;
```

```
    while (1) {
```

```
        V_b = ..;
```

```
        A.signal();
```

```
        B.wait();
```

```
        Buf_B = V_b;
```

```
        V_b = Buf_A;
```

```
}
```

I am ready

Wait for your card!

First Attempt: Problem (a)

Thread A	Thread B
B.signal()	
A.wait()	
	A.signal()
	B.wait()
Buf_A = V_a	
V_a = Buf_B	
	Buf_B = V_b

Buf_B has no value, yet!

Oops, it is too late!

The diagram illustrates a race condition between Thread A and Thread B. Thread A calls B.signal() and then A.wait(). Thread B calls A.signal() and then B.wait(). Thread A then sets Buf_A = V_a and V_a = Buf_B. Thread B then sets Buf_B = V_b. A red dashed arrow points from V_a = Buf_B to Buf_B = V_b, indicating that Buf_B is still zero when V_a is assigned to it. Blue arrows point from the text annotations to the relevant code lines.

First Attempt: Problem (b)

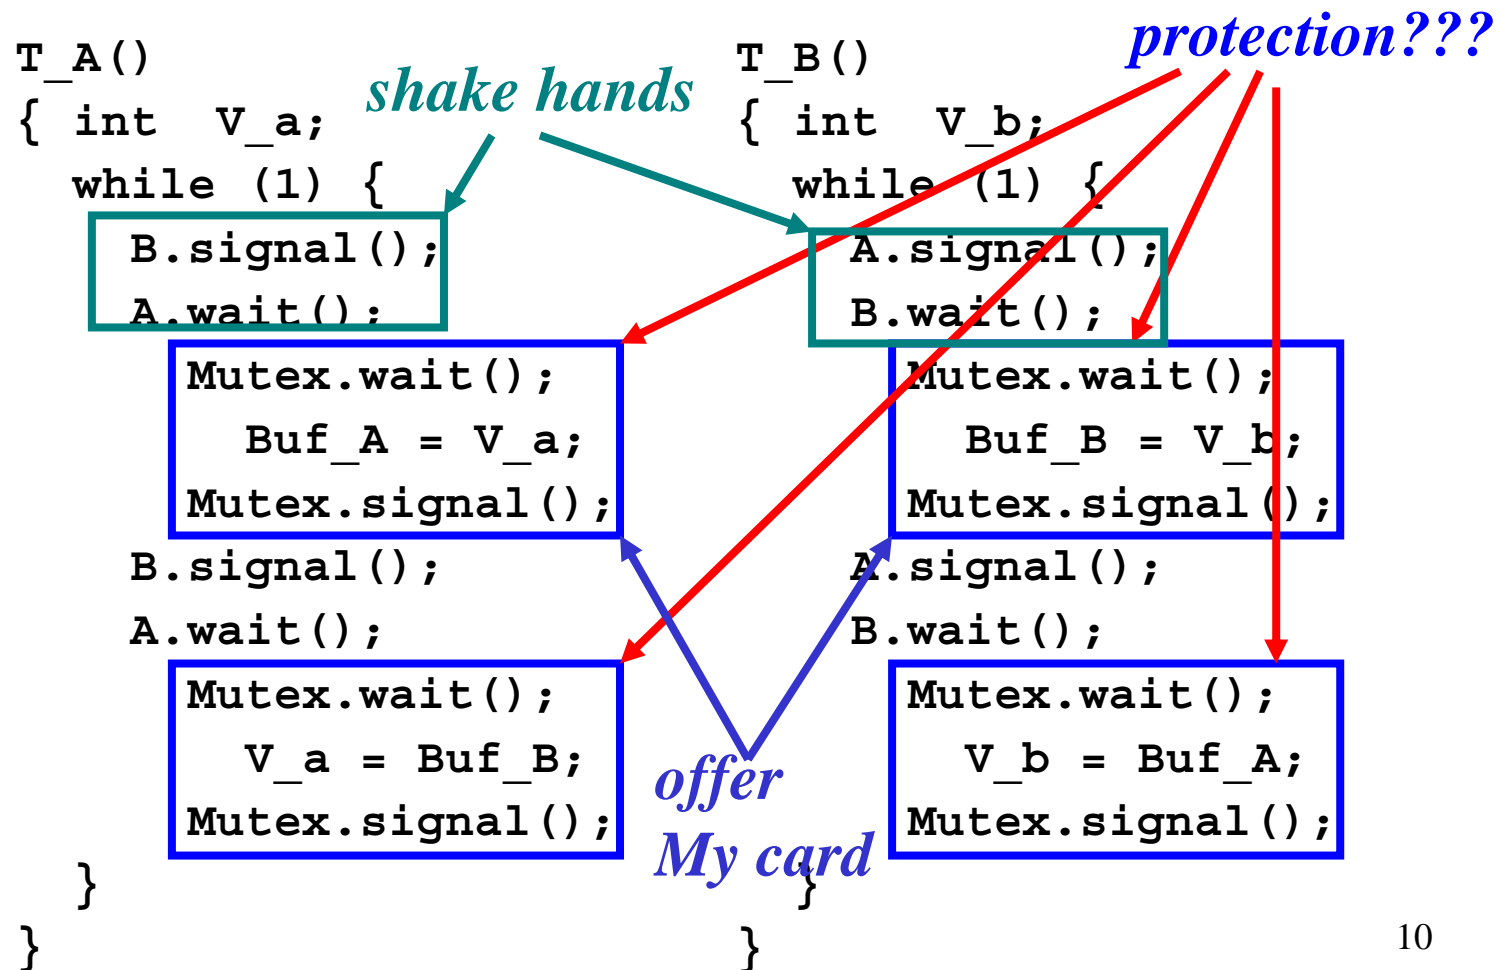
A_1	A_2	B_1	B_2
B.signal()			
A.wait()			
		A.signal()	
		B.wait()	
	B.signal()		
	A.wait()		
		Buf_B = .	
<i>Race Condition</i>			A.signal()
Buf_A = .			
	Buf_A = .		

What did we learn?

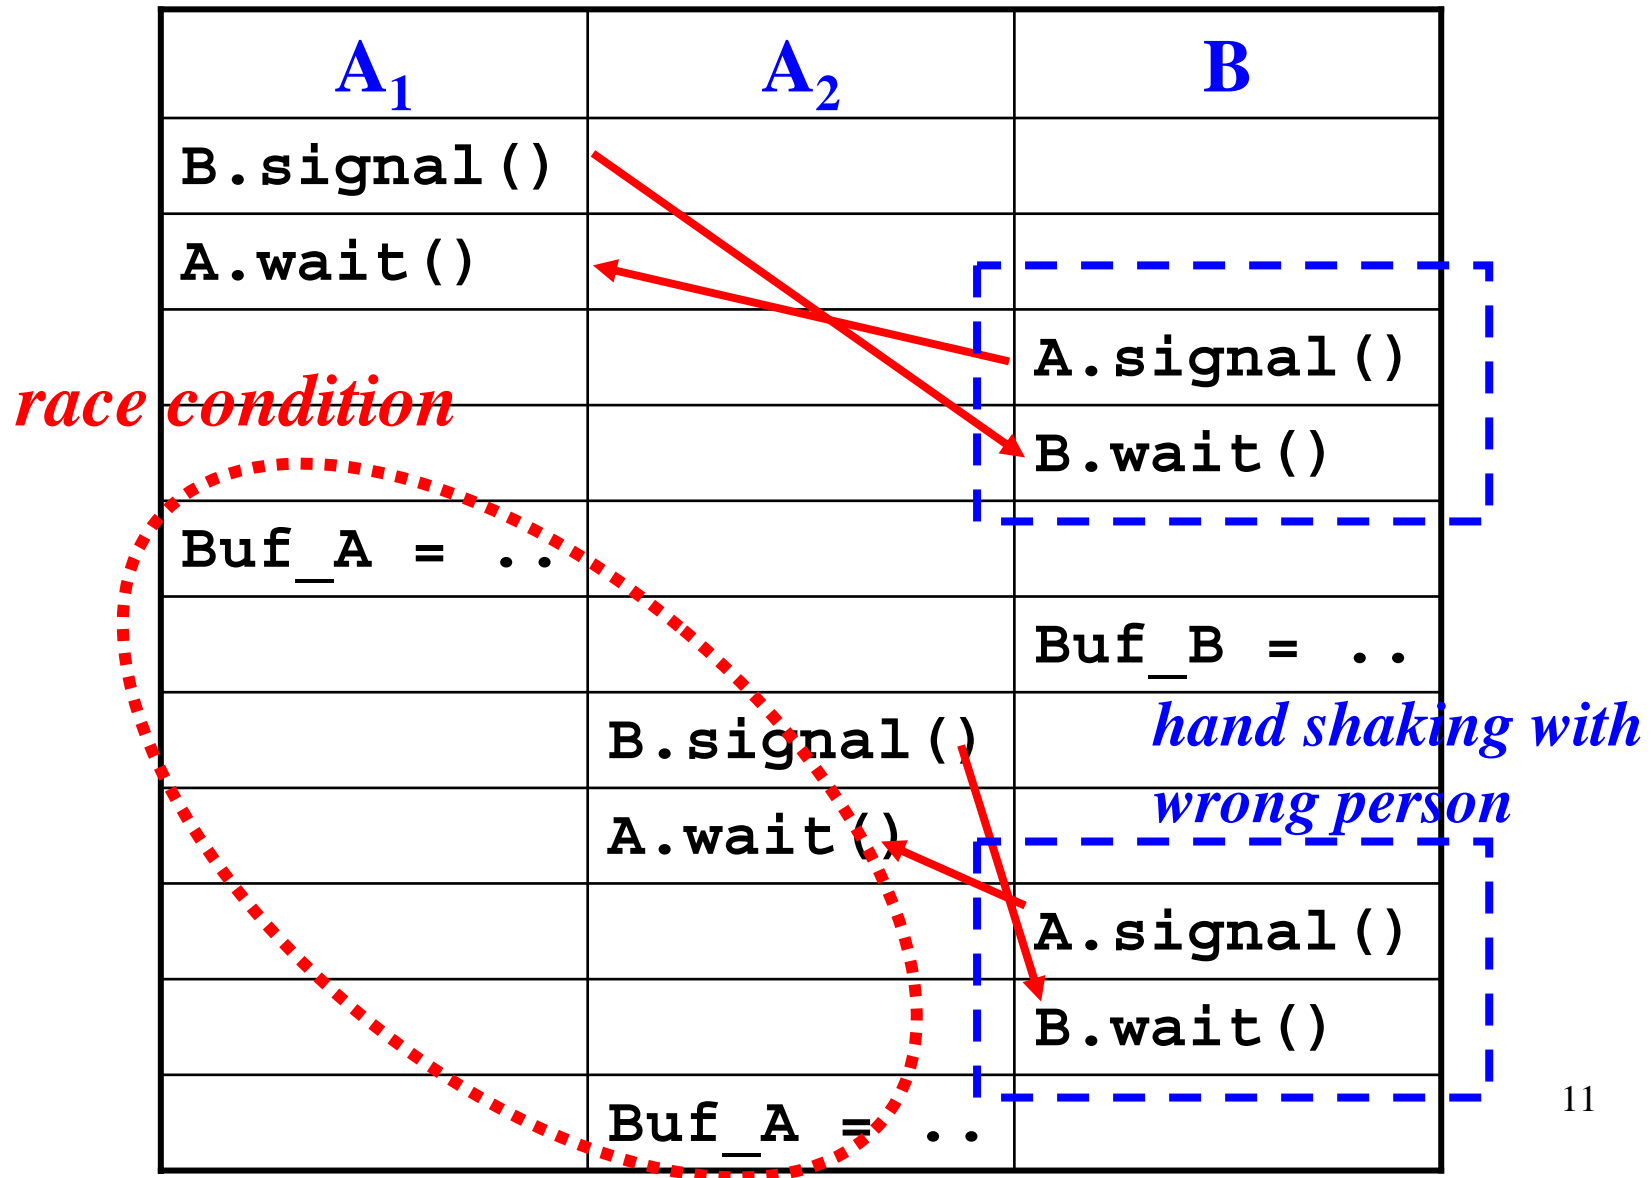
- ❑ If there are shared data items, always protect them properly. Without a proper mutual exclusion, race conditions are likely to occur.
- ❑ In this first attempt, both global variables `Buf_A` and `Buf_B` are shared and should be protected.

Second Attempt

```
sem  A = B = 0;
sem  Mutex = 1;
int  Buf_A, Buf_B;
```



Second Attempt: Problem



What did we learn?

- ❑ Improper protection is no better than no protection, because it gives us an *illusion* that data have been well-protected.
- ❑ We frequently forgot that protection is done by a critical section, which *cannot be divided*.
- ❑ Thus, protecting “here is my card” followed by “may I have yours” separately is not good enough.

Third Attempt

job done \longrightarrow `sem Aready = Bready = 1; \longleftarrow ready to proceed`
`sem Adone = Bdone = 0;`
`int Buf_A, Buf_B;`

<p><code>T_A()</code></p> <p><code>{ int V_a;</code></p> <p><code>while (1) {</code></p> <p><i>only one A can</i> <code>Aready.wait();</code> <i>Pass this point</i></p> <p><i>here is my card</i> <code>Buf_A = ..;</code> <i>let me have</i> <code>Adone.signal();</code> <i>yours</i> <code>Bdone.wait();</code></p> <p><code>V_a = Buf_B;</code></p> <p><code>Aready.signal();</code></p> <p><code>}</code></p> <p><code>}</code></p>	<p><code>T_B()</code></p> <p><code>{ int V_b;</code></p> <p><code>while (1) {</code></p> <p><code>Bready.wait();</code></p> <p><code>Buf_B = ..;</code></p> <p><code>Bdone.signal();</code></p> <p><code>Adone.wait();</code></p> <p><code>V_b = Buf_A;</code></p> <p><code>Bready.signal();</code></p> <p><code>}</code></p> <p><code>}</code></p>
--	---

Third Attempt: Problem

Thread A	Thread B
Buf_A =	
Adone.signal()	
Bdone.wait()	
	Bdone.signal()
	Adone.wait()
... = Buf_B	
Already.signal()	
** loop back **	
Already.wait()	
Buf_A = ...	
	... = Buf_A

ruin the original value of Buf_A

B is a slow thread

race condition

What did we learn?

- ❑ Mutual exclusion for one group may not prevent processes in other groups from interacting with a process in this group.
- ❑ It is common that we protect a shared item for one group and forget other possible, unintended accesses.
- ❑ Protection must be applied *uniformly* to all processes rather than within groups.

Fourth Attempt

job done → `sem Aready = Bready = 1; ← ready to proceed`
`sem Adone = Bdone = 0;`
`int Buf_A, Buf_B;`

	<code>T_A()</code>	<i>wait/signal switched</i>	<code>T_B()</code>
	<code>{ int V_a;</code>		<code>{ int V_b;</code>
	<code>while (1) {</code>		<code>while (1) {</code>
<i>I am the only A</i> →	<code>Bready.wait();</code>		<code>Aready.wait();</code>
	<code>Buf_A = ..;</code>		<code>Buf_B = ..;</code>
<i>here is my card</i> →	<code>Adone.signal();</code>		<code>Bdone.signal();</code>
<i>waiting for yours</i> →	<code>Bdone.wait();</code>		<code>Adone.wait();</code>
	<code>V_a = Buf_B;</code>		<code>V_b = Buf_A;</code>
<i>Job done & next B please</i> →	<code>Aready.signal();</code>		<code>Bready.signal();</code>
	<code>}</code>		<code>}</code>
	<code>}</code>		<code>}</code>

Fourth Attempt: Problem

A_1	A_2	B
Bready.wait()		
Buf_A = ...		
Adone.signal()		Buf_B = ...
		Bdone.signal()
		Adone.wait()
		... = Buf_A
		Bready.signal()
	Bready.wait()	
	<i>Hey, this one is for A₁!!!</i>
	Bdone.wait()	
	... = Buf_B	

What did we learn?

- ❑ We use locks for mutual exclusion.
- ❑ The owner, the one who locked the lock, should unlock the lock.
- ❑ In the above “solution,” `Already` is acquired by a process `A` but released by a process `B`. This is risky!
- ❑ In this case, a pure lock is more natural than a binary semaphore.

Conclusions

- ❑ Detecting race conditions is difficult as it is an **NP-hard** problem.
- ❑ Hence, detecting race conditions is heuristic.
- ❑ Incorrect mutual exclusion is no better than no mutual exclusion.
- ❑ Race conditions are sometimes very subtle. They may appear at unexpected places.
- ❑ Check the **ThreadMentor** tutorial pages for more details and correct solutions.