

CS 453/698: Software and Systems Security

Module: Other Common Vulnerability Types

Lecture: Race condition and data race

Meng Xu (*University of Waterloo*)

Spring 2025

Outline

- 1 Concepts: race condition vs data race
- 2 Introductory examples
- 3 Atomicity violations
- 4 Bonus: lock implementation
- 5 Other forms of races

What is a race condition?

Wikipedia's definition

A race condition is the condition of a software system where the system's substantive behavior is dependent on the sequence or timing of other uncontrollable events, leading to unexpected or inconsistent results.

It becomes a bug when one or more of the possible behaviors is undesirable.

Wikipedia's definition

*A **race condition** is the condition of a software system where the system's substantive behavior is dependent on the sequence or timing of other uncontrollable events, leading to unexpected or **inconsistent** results.*

*It becomes a **bug** when one or more of the possible behaviors is **undesirable**.*

What is a data race?

Data race definition in C++ standard

When

- an evaluation of an expression writes to a memory location **and**
- another evaluation reads or modifies **the same memory location**,
the expressions are said to **conflict**.

A program that has two **conflicting evaluations** has a **data race** unless:

- both evaluations execute on **the same thread**, **or**
- both conflicting evaluations are **atomic operations**, **or**
- one of the conflicting evaluations **happens-before** another.

Adapted from [a community-backed C++ reference site](#). For the full version, please refer to the related sections in [C++ working draft](#).

An intuitive definition

Intuitively, a *data race* happens when:

- 1 There are two memory accesses from **different threads**.
- 2 Both accesses target the **same memory location**.
- 3 At least one of them is a **write** operation.

An intuitive definition

Intuitively, a *data race* happens when:

- 1 There are two memory accesses from **different threads**.
- 2 Both accesses target the **same memory location**.
- 3 At least one of them is a **write** operation.
- 4 Both accesses could **interleave** freely without restrictions such as **synchronization primitives** or **causality relations**.

Test of your understanding

Q: Based on the definition of race condition and data race, what do you think are the relationship between them?

Outline

- 1 Concepts: race condition vs data race
- 2 **Introductory examples**
- 3 Atomicity violations
- 4 Bonus: lock implementation
- 5 Other forms of races

Introductory case

global var **count** = 0

```
for(i = 0; i < x; i++) {  
    /* do sth critical */  
    .....  
    count++;  
}
```

Thread 1

```
for(i = 0; i < y; i++) {  
    /* do sth critical */  
    .....  
    count++;  
}
```

Thread 2

Q: What is the value of **count** when both threads terminate?

Introductory case

global var **count** = 0

```
for(i = 0; i < x; i++) {  
    /* do sth critical */  
    .....  
    count++;  
}
```

Thread 1

```
for(i = 0; i < y; i++) {  
    /* do sth critical */  
    .....  
    count++;  
}
```

Thread 2

Q: What is the value of **count** when both threads terminate?

Introductory case

```
global var count = 0
global var mutex = 1
```

```
for(i = 0; i < x; i++) {
    /* do sth critical */
    .....
    lock(mutex);
    count++;
    unlock(mutex);
}
```

Thread 1

```
for(i = 0; i < y; i++) {
    /* do sth critical */
    .....
    lock(mutex);
    count++;
    unlock(mutex);
}
```

Thread 2

Q: What is the value of **count** when both threads terminate?

Race conditions in other settings

Race conditions are not tied to a specific programming language, instead, they are tied to **data sharing in concurrent execution**.

Race conditions in other settings

Race conditions are not tied to a specific programming language, instead, they are tied to **data sharing in concurrent execution**.

For example, in the database context:

Q: If two database clients send the following requests concurrently, what will be the result (both try to withdraw \$100 from Alice)?

Client 1

```
SELECT @balance = Balance
  FROM Ledger WHERE Name = "Alice";

UPDATE Ledger SET Balance =
  @balance - 100 WHERE Name = "Alice";
```

Client 2

```
SELECT @balance = Balance
  FROM Ledger WHERE Name = "Alice";

UPDATE Ledger SET Balance =
  @balance - 100 WHERE Name = "Alice";
```


Race conditions in a database setting

One possible interleaving (that messes up the states)

```
SELECT @balance = Balance FROM Ledger WHERE Name = "Alice";  
SELECT @balance = Balance FROM Ledger WHERE Name = "Alice";  
UPDATE Ledger SET Balance = @balance - 100 WHERE Name = "Alice";  
UPDATE Ledger SET Balance = @balance - 100 WHERE Name = "Alice";
```

Race conditions in a database setting

One possible interleaving (that messes up the states)

```
SELECT @balance = Balance FROM Ledger WHERE Name = "Alice";  
SELECT @balance = Balance FROM Ledger WHERE Name = "Alice";  
UPDATE Ledger SET Balance = @balance - 100 WHERE Name = "Alice";  
UPDATE Ledger SET Balance = @balance - 100 WHERE Name = "Alice";
```

Q: How to prevent the race condition in this case?

Race conditions in a database setting

One possible interleaving (that messes up the states)

```
SELECT @balance = Balance FROM Ledger WHERE Name = "Alice";
SELECT @balance = Balance FROM Ledger WHERE Name = "Alice";
UPDATE Ledger SET Balance = @balance - 100 WHERE Name = "Alice";
UPDATE Ledger SET Balance = @balance - 100 WHERE Name = "Alice";
```

Q: How to prevent the race condition in this case?

Interleavings with transactions

```
BEGIN TRANSACTION;
    SELECT @balance = Balance FROM Ledger WHERE Name = "Alice";
    UPDATE Ledger SET Balance = @balance - 100 WHERE Name = "Alice";
COMMIT TRANSACTION;
BEGIN TRANSACTION;
    SELECT @balance = Balance FROM Ledger WHERE Name = "Alice";
    UPDATE Ledger SET Balance = @balance - 100 WHERE Name = "Alice";
COMMIT TRANSACTION;
```

Revisit the example

global var **count** = 0

```
for(i = 0; i < x; i++) {  
    count++;  
}
```


Thread 1

```
for(i = 0; i < y; i++) {  
    count++;  
}
```

Thread 2

Q: Is it a data race?

Free interleavings of memory reads and writes



Revisit the example

global var **count** = 0


```
for(i = 0; i < x; i++) {  
    lock(mutex);  
    count++;  
    unlock(mutex);  
}
```

Thread 1

```
for(i = 0; i < y; i++) {  
    lock(mutex);  
    count++;  
    unlock(mutex);  
}
```

Thread 2

Limited interleavings with locking



Revisiting the definition


Intuitively, a *data race* happens when:

- 1 There are two memory accesses from **different threads**.
- 2 Both accesses target the **same memory location**.
- 3 At least one of them is a **write** operation.
- 4 Both accesses could **interleave** freely without restrictions such as **synchronization primitives** ~~or **causality relations**~~.

Causality relations: an example

```
1 #include <stdio.h>
2 #include <pthread.h>
3
4 int g_x;
5 int g_y;
6
7 void* foo(void* p){
8     printf("Value of g_x: %d\n", g_x);
9     printf("Value of arg: %d\n", *(int *)p);
10    pthread_exit(&g_y);
11 }
12
13 int main(void){
14     int g_x = 1;
15     int arg = 2;
16
17     pthread_t id;
18     pthread_create(&id, NULL, foo, &arg);
19     pthread_join(id, NULL);
20
21     printf("Return value from thread: %d\n", g_y);
22 }
```

Causality relations



Outline

- 1 Concepts: race condition vs data race
- 2 Introductory examples
- 3 Atomicity violations**
- 4 Bonus: lock implementation
- 5 Other forms of races

Revisit the example

global var **count** = 0

```
for(i = 0; i < x; i++) {  
    lock(mutex);  
    t = count;  
    unlock(mutex);  
  
    t++;  
  
    lock(mutex);  
    count = t;  
    unlock(mutex);  
}
```

Thread 1

```
for(i = 0; i < y; i++) {  
    lock(mutex);  
    t = count;  
    unlock(mutex);  
  
    t++;  
  
    lock(mutex);  
    count = t;  
    unlock(mutex);  
}
```

Thread 2

Revisit the example

Q: In this modified example, is there a data race?

Revisit the example

Q: In this modified example, is there a data race?

A: No

Revisit the example

Q: In this modified example, is there a data race?

A: No

Q: But the results are the same with all locks removed?

global var **count** = 0

```
for(i = 0; i < x; i++) {  
    t = count;  
    t++;  
    count = t;  
}
```

```
for(i = 0; i < y; i++) {  
    t = count;  
    t++;  
    count = t;  
}
```

Revisit the example

Q: In this modified example, is there a data race?

A: No

Q: But the results are the same with all locks removed?

global var **count** = 0

```
for(i = 0; i < x; i++) {  
    t = count;  
    t++;  
    count = t;  
}
```

```
for(i = 0; i < y; i++) {  
    t = count;  
    t++;  
    count = t;  
}
```

A: No, depending on how hardware works (e.g., per-bit conflict)

Reading developers' mind

Q: What is developers' expectation in the running example?

Reading developers' mind

Q: What is developers' expectation in the running example?

A: States do not change **for a critical section** during execution.

Reading developers' mind

Q: What is developers' expectation in the running example?

A: States do not change **for a critical section** during execution.

A: Generalization: states remain **integral for a critical section** during execution. No change of states is just one way of remaining integral (assuming state is integral before the critical section).

State integrity example

```
1 struct R { x: int, y: int } g;  
2 [invariant] g.x + g.y == 100;
```

```
1 int add_x(v: int) {  
2   g.x += v;  
3   g.y -= v;  
4 }
```

Thread 1

```
1 int add_y(v: int) {  
2   g.y += v;  
3   g.x -= v;  
4 }
```

Thread 2

State integrity example

```
1 struct R { x: int, y: int } g;  
2 [invariant] g.x + g.y == 100;  
3 lock mutex = unlocked;
```

```
1 int add_x(v: int) {  
2   lock(mutex);  
3   g.x += v;  
4   unlock(mutex);  
5   lock(mutex);  
6   g.y -= v;  
7   unlock(mutex);  
8 }
```

Thread 1

```
1 int add_y(v: int) {  
2   lock(mutex);  
3   g.y += v;  
4   unlock(mutex);  
5   lock(mutex);  
6   g.x -= v;  
7   unlock(mutex);  
8 }
```

Thread 2

Q: Is this the right way of adding locks?

State integrity example

```
1 struct R { x: int, y: int } g;  
2 [invariant] g.x + g.y == 100;  
3 lock mutex = unlocked;
```

```
1 int add_x(v: int) {  
2   lock(mutex);  
3   g.x += v;  
4   unlock(mutex);  
5   lock(mutex);  
6   g.y -= v;  
7   unlock(mutex);  
8 }
```

Thread 1

```
1 int add_y(v: int) {  
2   lock(mutex);  
3   g.y += v;  
4   unlock(mutex);  
5   lock(mutex);  
6   g.x -= v;  
7   unlock(mutex);  
8 }
```

Thread 2

Q: Is this the right way of adding locks?

A: No, as the invariant is not guaranteed

State integrity example

```
1 struct R { x: int, y: int } g;  
2 [invariant] g.x + g.y == 100;  
3 lock mutex = unlocked;
```

```
1 int add_x(v: int) {  
2   lock(mutex);  
3   g.x += v;  
4   g.y -= v;  
5   unlock(mutex);  
6 }
```

Thread 1

```
1 int add_y(v: int) {  
2   lock(mutex);  
3   g.y += v;  
4   g.x -= v;  
5   unlock(mutex);  
6 }
```

Thread 2

Q: Is this the right way of adding locks?

State integrity example

```
1 struct R { x: int, y: int } g;  
2 [invariant] g.x + g.y == 100;  
3 lock mutex = unlocked;
```

```
1 int add_x(v: int) {  
2   lock(mutex);  
3   g.x += v;  
4   g.y -= v;  
5   unlock(mutex);  
6 }
```

Thread 1

```
1 int add_y(v: int) {  
2   lock(mutex);  
3   g.y += v;  
4   g.x -= v;  
5   unlock(mutex);  
6 }
```

Thread 2

Q: Is this the right way of adding locks?

A: Yes, the invariant is guaranteed at each entry and exit of the critical section in both threads

State integrity is hard to capture

However, in practice, the invariant often exists in

- some architectural design documents (which no one reads)
- code comments in a different file (which no one notices)
- folklore knowledge among the dev team
- the mind of the developer who has resigned a few years ago...

Outline

- 1 Concepts: race condition vs data race
- 2 Introductory examples
- 3 Atomicity violations
- 4 Bonus: lock implementation**
- 5 Other forms of races

Common synchronization primitives

Common synchronization primitives

- Lock / Mutex / Critical section
- Read-write lock
- Barrier
- Semaphore

How are synchronization primitives implemented?

- Hardware support
 - Atomic swap
 - Atomic read-modify-write
 - * compare-and-swap
 - * test-and-set
 - * fetch-and-add
 - *

How are synchronization primitives implemented?

- Hardware support
 - Atomic swap
 - Atomic read-modify-write
 - * compare-and-swap
 - * test-and-set
 - * fetch-and-add
 - *
- Software algorithms
 - Dekker's algorithm

Spinlock with atomic swap (xchg)

```
1 locked:                                ; The lock variable. 1 = locked, 0 = unlocked.
2     dd      0
3
4 spin_lock:
5     mov     eax, 1                    ; Set the EAX register to 1.
6     xchg    eax, [locked]             ; Atomically swap the EAX register with
7                                     ; the lock variable.
8                                     ; This will always store 1 to the lock, leaving
9                                     ; the previous value in the EAX register.
0     test    eax, eax                  ; Test EAX with itself. Among other things, this
1                                     ; will set the processor's Zero Flag if EAX is 0.
2                                     ; If EAX is 0, then the lock was unlocked and
3                                     ; we just locked it.
4                                     ; Otherwise, EAX is 1 and we didn't acquire the lock.
5     jnz     spin_lock                 ; Jump back to the MOV instruction if the Zero Flag is
6                                     ; not set; the lock was previously locked, and so
7                                     ; we need to spin until it becomes unlocked.
8     ret                                ; The lock has been acquired, return to the caller.
9
0 spin_unlock:
1     xor     eax, eax                  ; Set the EAX register to 0.
2     xchg    eax, [locked]             ; Atomically swap the EAX register with
3                                     ; the lock variable.
4     ret                                ; The lock has been released.
```

Spinlock with atomic swap (xchg)

Q: Are there data races or race conditions in spinlock implementation?

Spinlock with atomic swap (xchg)

Q: Are there data races or race conditions in spinlock implementation?

A: By looking at the code

- Data race: Yes, but hardware guarantees atomicity
- Race condition: No

Dekker's algorithm

```
1 atomic_bool wants_to_enter[2] = {false, false};
2 int turn = 0; /* or turn = 1 */
```

```
1 // lock
2 wants_to_enter[0] = true;
3 while (wants_to_enter[1]) {
4     if (turn != 0) {
5         wants_to_enter[0] = false;
6         // busy wait
7         while (turn != 0) {}
8         wants_to_enter[0] = true;
9     }
10 }
11
12 /* ... critical section ... */
13
14 // unlock
15 turn = 1;
16 wants_to_enter[0] = false;
```

Thread 1

```
1 // lock
2 wants_to_enter[1] = true;
3 while (wants_to_enter[0]) {
4     if (turn != 1) {
5         wants_to_enter[1] = false;
6         // busy wait
7         while (turn != 1) {}
8         wants_to_enter[1] = true;
9     }
10 }
11
12 /* ... critical section ... */
13
14 // unlock
15 turn = 0;
16 wants_to_enter[1] = false;
```

Thread 2

Dekker's algorithm

Q: Are there data races or race conditions in Dekker's algorithm?

Dekker's algorithm

Q: Are there data races or race conditions in Dekker's algorithm?

A: By looking at the code

- Data race: No (assuming `atomic_bool`)
- Race condition: No

Outline

- 1 Concepts: race condition vs data race
- 2 Introductory examples
- 3 Atomicity violations
- 4 Bonus: lock implementation
- 5 Other forms of races

A more abstract view of race conditions

Q: Why do race conditions happen in the first place?

A more abstract view of race conditions

Q: Why do race conditions happen in the first place?

A: Because two threads in the same process **share memory**

A more abstract view of race conditions

Q: Why do race conditions happen in the first place?

A: Because two threads in the same process **share memory**

We can further generalize this concept by asking:

Q: What else do they share?

Q: What about other entities that may run concurrently?

Example: race over the filesystem

```
1 #include <...>
2
3 int main(int argc, char *argv[]) {
4     FILE *fd;
5     struct stat buf;
6
7     if (stat("/some_file", &buf)) {
8         exit(1); // cannot read stat message
9     }
10
11     if (buf.st_uid != getuid()) {
12         exit(2); // permission denied
13     }
14
15     fd = fopen("/some_file", "wb+");
16     if (fd == NULL) {
17         exit(3); // unable to open the file
18     }
19
20     fprintf(f, "<some-secret-value>");
21     fclose(fd);
22     return 0;
23 }
```

Example: the Dirty COW exploit

CVE-2016-5195

Allows local privilege escalation: `user(1000) → root(0)`.

Existed in the kernel for nine years before finally patched.

Details on the [Website](#).

〈 End 〉