

CS 858: Software Security  
Offensive and Defensive Approaches

**Attacks: data race**

Meng Xu (*University of Waterloo*)

Fall 2022

# Outline

- 1 Introduction
- 2 Intuitive definition
- 3 Formal reasoning
- 4 Other form of races

# What is data race?

# What is data race?

```
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?

# What is data race?

```
global var count = 0
global var mutex = ⊥
```

---

```
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?

# Data race combined with memory errors

# Data race combined with memory errors

`p` is a global pointer initialized to NULL

```
if (!p) {  
    p = malloc(128);  
}
```

Thread 1

```
if (!p) {  
    p = malloc(256);  
}
```

Thread 2

**Q:** What are the possible outcomes of this execution?

# Data race combined with memory errors

**p** is a global pointer initialized to NULL

```
if (!p) {  
    p = malloc(128);  
}
```

```
if (p) {  
    free(p);  
    p = NULL;  
}
```

Thread 1

```
if (!p) {  
    p = malloc(256);  
}
```

```
if (p) {  
    free(p);  
    p = NULL;  
}
```

Thread 2

**Q:** What are the possible outcomes of this execution?



# Data race as heisenbug

# Data race as heisenbug

Programs which contain data races usually demonstrate unexpected and even **non-deterministic** behavior.

# Data race as heisenbug

Programs which contain data races usually demonstrate unexpected and even **non-deterministic** behavior.

- The outcome might depend on a specific execution order (a.k.a. **thread interleaving**).
- Re-running the program may not always produce the same results.

# Data race as heisenbug

Programs which contain data races usually demonstrate unexpected and even **non-deterministic** behavior.

- The outcome might depend on a specific execution order (a.k.a. **thread interleaving**).
- Re-running the program may not always produce the same results.

Concurrent programs are hard to debug and even harder to ensure correctness.

# Outline

- 1 Introduction
- 2 Intuitive definition**
- 3 Formal reasoning
- 4 Other form of races

# 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**.

# 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



# Free interleavings without locking



## 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



# Common synchronization primitives

# Common synchronization primitives

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

# 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 i;
5 int retval;
6
7 void* foo(void* p){
8     printf("Value of i: %d\n", i);
9     printf("Value of j: %d\n", *(int *)p);
10    pthread_exit(&retval);
11 }
12
13 int main(void){
14     int i = 1;
15     int j = 2;
16
17     pthread_t id;
18     pthread_create(&id, NULL, foo, &j);
19     pthread_join(id, NULL);
20
21     printf("Return value from thread: %d\n", retval);
22 }
```

# Causality relations





# Wait..., how are synchronization primitives implemented?

# Wait..., how are synchronization primitives implemented?

- Dekker's algorithm
- Atomic swap
- Atomic read-modify-write
  - compare-and-swap
  - test-and-set
  - fetch-and-add
  - .....

# Dekker's algorithm

```
1 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

# Bonus: 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.
```

## Revisiting the definition (again)

If we can find, statically or dynamically, a pair of memory accesses  $(A_1, A_2)$  such that

- they originate from **different threads**,
- both  $A_1$  and  $A_2$  target the **same memory location**, AND
- at least one of them is a **write** operation,

then we conclude that  $(A_1, A_2)$  must be one of the following cases:

- $(A_1, A_2)$  is part of a **synchronization primitive**, OR
- $(A_1, A_2)$  is a **data race**.

## Revisiting the definition (again)

If we can find, statically or dynamically, a pair of memory accesses  $(A_1, A_2)$  such that

- they originate from **different threads**,
- both  $A_1$  and  $A_2$  target the **same memory location**, AND
- at least one of them is a **write** operation,

then we conclude that  $(A_1, A_2)$  must be one of the following cases:

- $(A_1, A_2)$  is part of a **synchronization primitive**, OR
- $(A_1, A_2)$  is a **data race**.

**Q:** Is this definition good enough?

# Is this a data race?

```
1 int x = 0;
2 bool flag = false;
3 lock mutex = unlocked;
```

---

```
1 x++;
2 lock(mutex);
3 flag = true;
4 unlock(mutex);
```

---

Thread 1

---

```
1 while(true) {
2     lock(mutex);
3     if (flag) {
4         break;
5     }
6     unlock(mutex);
7 }
8 x--;
```

---

Thread 2

# Is this a (bad) data race?

```
1 int x = 0;  
2 bool flag = false;
```

---

```
1 x++;  
2 flag = true;
```

---

Thread 1

---

```
1 while (!flag) {};  
2 x--;
```

---

Thread 2



# Outline

- 1 Introduction
- 2 Intuitive definition
- 3 Formal reasoning**
- 4 Other form of races

# How to model concurrency mathematically?

# How to model concurrency mathematically?

- Lamport clock
- Vector clock

# Lamport clock algorithm

Each thread has its own clock variable  $t$

- On initialization:
  - $t \leftarrow 0$
- On write to shared memory  $*ptr = val$ :
  - $t \leftarrow t + 1$
  - store  $t$  alongside  $val$  at memory location  $ptr$
- On read from shared memory  $val = *ptr$ :
  - retrieve the stored clock  $t'$  at memory location  $ptr$
  - $t \leftarrow \max(t, t') + 1$

# Lamport clock algorithm

Each thread has its own clock variable  $t$

- On initialization:
  - $t \leftarrow 0$
- On write to shared memory  $*ptr = val$ :
  - $t \leftarrow t + 1$
  - store  $t$  alongside  $val$  at memory location  $ptr$
- On read from shared memory  $val = *ptr$ :
  - retrieve the stored clock  $t'$  at memory location  $ptr$
  - $t \leftarrow \max(t, t') + 1$

## Properties of Lamport clock:

- $a \rightarrow b \implies L(a) < L(b)$
- $L(a) < L(b) \not\implies a \rightarrow b$

# Vector clock algorithm

Each thread  $i$  has its own clock vector  $t$

- On initialization:
  - $T \leftarrow \langle 0, 0, \dots, 0 \rangle_N$ , assuming  $N$  threads
- On write to shared memory  $*ptr = val$ :
  - $T[i] \leftarrow T[i] + 1$
  - store  $T$  alongside  $val$  at memory location  $ptr$
- On read from shared memory  $val = *ptr$ :
  - retrieve the stored clock  $T'$  at memory location  $ptr$
  - $\forall k \in [0, N) : T[k] = \max(T[k], T'[k])$
  - $T[i] \leftarrow T[i] + 1$

# Properties of the vector clock algorithm

With the following definition on the timestamp ordering:

- $T = T' \iff \forall i \in [0, N) : T[i] = T'[i]$
- $T \leq T' \iff \forall i \in [0, N) : T[i] \leq T'[i]$
- $T < T' \iff T \leq T' \wedge T \neq T'$
- $T \parallel T' \iff T \not\leq T' \wedge T' \not\leq T$

We have:

- $a \rightarrow b \iff V(a) < V(b)$
- $a = b \iff V(a) = V(b)$
- $a \parallel b \iff V(a) \parallel V(b)$

# Homework exercise

```
1 int x = 0;  
2 bool flag = false;
```

---

```
1 x++;  
2 flag = true;
```

---

Thread 1

---

```
1 while (!flag) {};  
2 x--;
```

---

Thread 2

**Prove** (by hand) that the write of `x` at `x--` in thread 2 can **never happen before** the read of `x` in `x++` in thread 1.



# Outline

- 1 Introduction
- 2 Intuitive definition
- 3 Formal reasoning
- 4 Other form of races**

# A more abstract view of data race

**Q:** Why data race can happen in the first place?

# A more abstract view of data race

**Q:** Why data race can happen in the first place?

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

# A more abstract view of data race

**Q:** Why data race can 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?

And the answer to these questions will help define **race condition**.

# 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)`.

Exists in the kernel for nine years before finally patched.

Details on the [Website](#).

⟨ **End** ⟩