

CS 489 / 698: Software and Systems Security

Module 2: Program Security (Attacks)

data races

Meng Xu (*University of Waterloo*)

Spring 2023

Outline

- 1 Why studying data races?
- 2 Intuitive definition
- 3 Formal reasoning
- 4 Data race vs atomicity
- 5 Other form of races

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 in other settings

Data races 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";
```

Data race 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 data race 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;
```

Recall the “nice” properties of memory error

Data race is a common attack vector and building blocks for sophisticated exploitations... just like **memory error**.

- Memory errors have **universally** accepted definitions
 - Once you find a memory error, you do not need to diligently argue that this is a bug and not a feature
- Memory errors often lead to a set of known consequences that are **generally** considered severe (e.g., data leak or denial-of-service)
 - Once you find a memory error, you do not need to construct a working exploit to justify it
- Finding memory errors typically **do not require** program-specific domain knowledge
 - If you have a technique that can find memory errors in one codebase, you can scale it up to millions of codebases

In fact, very few types of vulnerabilities meet these requirements.

⇒ **data race** is one of them!

“s/memory error/data race/g”

- Data races have **universally** accepted definitions
 - Once you find a data race, you do not need to diligently argue that this is a bug and not a feature
- Data races often lead to a set of known consequences that are **generally** considered severe (e.g., data leak or denial-of-service)
 - Once you find a data race, you do not need to construct a working exploit to justify it
- Finding data races typically **do not require** program-specific domain knowledge
 - If you have a technique that can find data races in one codebase, you can scale it up to millions of codebases

Data races can only happen in programs **with data sharing through a concurrency model**, e.g., **multi-threaded** or **distributed** programs.

Data race may lead to 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

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 Why studying data races?
- 2 Intuitive definition**
- 3 Formal reasoning
- 4 Data race vs atomicity
- 5 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.
- 4 Both accesses could **interleave** freely without restrictions such as **synchronization primitives** or **causality relations**.

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](#).

Revisit the example

global var **count** = 0

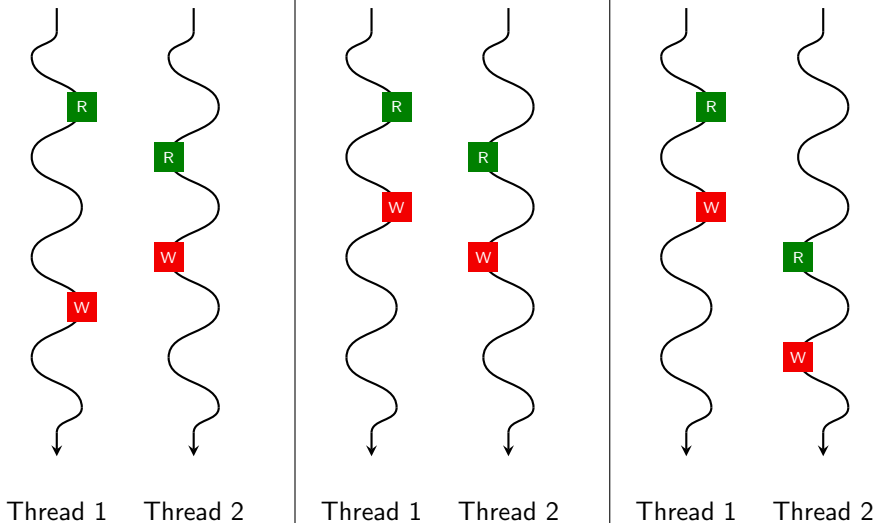
```
for(i = 0; i < x; i++) {  
  count++;  
}
```

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

- 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



Outline

- 1 Why studying data races?
- 2 Intuitive definition
- 3 Formal reasoning**
- 4 Data race vs atomicity
- 5 Other form of races

Revisiting the definition

If we can find, statically or dynamically, a pair of memory access instructions (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:

- 1 A_1 **strictly happens before** A_2 or vice versa due to causality, **OR**
- 2 A_1 and A_2 can only occur when **a common lock** is held, **OR**
- 3 (A_1, A_2) is a **data race**.

Q: Wait... how are locks implemented?

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.
```


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

Dekker's algorithm

Q: Suppose that you are not aware that Dekker's algorithm is implementing a lock, are there data races in Dekker's algorithm?

A: By looking at the code, yes...
However, this is often called a **benign** data race.

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 data race?

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

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

Thread 1

```
1 while (!flag) {};
```

```
2 x--;
```

Thread 2

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
 - $t \leftarrow \max(t, t') + 1$

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:

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

Practice exercise (vector clock)

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

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

Thread 1

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

Thread 2

Prove: the write of x at $x--$ in thread 2 can **never happen before** the read of x in $x++$ in thread 1.

Practice exercise (vector clock)

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

```
1 v = load(&x);  
2 store(&x, v + 1);  
3 store(&r, true);
```

Thread 1

```
1 loop:  
2   c = load(&r);  
3   jump_if_false(c, loop);  
4   v = load(&x);  
5   store(&x, v - 1);
```

Thread 2

Prove: line 5 at thread 2 can **never happen before** line 1 at thread 1.

Outline

- 1 Why studying data races?
- 2 Intuitive definition
- 3 Formal reasoning
- 4 Data race vs atomicity**
- 5 Other form 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?

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)

Extract the commonalities of the two variants

Q: What is common in developers' expectations in the two variants?

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

A: Generalization: state 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?

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?

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 Why studying data races?
- ② Intuitive definition
- ③ Formal reasoning
- ④ Data race vs atomicity
- ⑤ Other form of races

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** ⟩