

w/h

Thread 1 $A=0$
 $B=0$

Thread 2

$Q = 0, 0?$

(1) $A = 1$
(2) $\text{print}(B)$

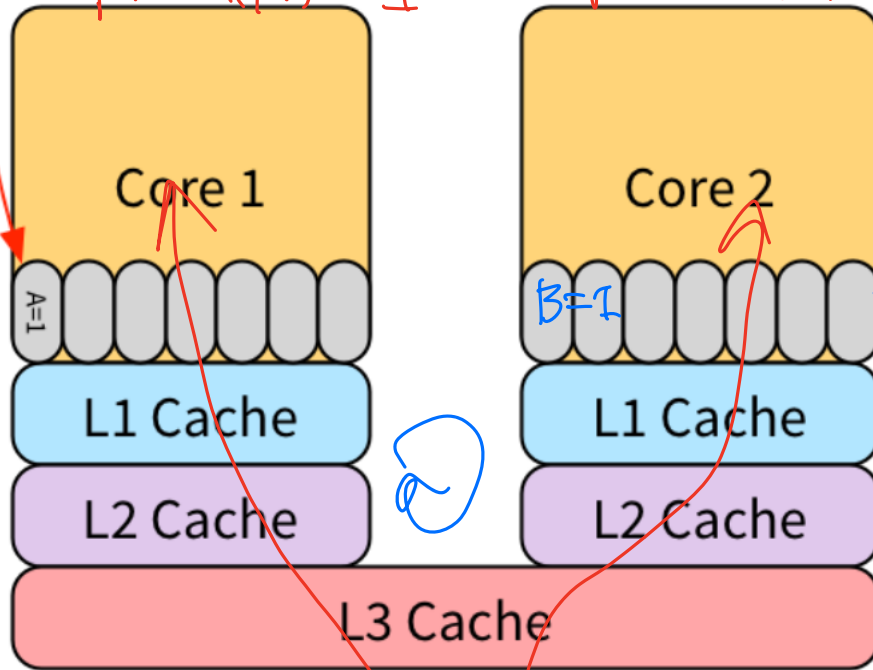
(3) $B = 1$
(4) $\text{print}(A)$

$\text{print}(A) \Rightarrow 1$

$\text{print}(B) \Rightarrow 1$

TSO

ARM



Mem $A=0$
 $B=0$

Borrowed from blog "Memory Consistency Model: A Tutorial", James Bornholt.
<https://www.cs.utexas.edu/~bornholt/post/memory-models.html>

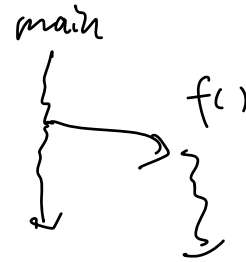
Week 5.b
CS5600
02/08 2023
<https://naizhengtan.github.io/23spring/>

• threads
(vs. process)

- 0. Last time ←
 - 1. Concurrency and consistency model
 - 2. Managing concurrency
 - 3. Mutexes
 - 4. Condition variables
 - 5. (optional) Bakery algorithm
 - 6. (optional) Semaphores
-

```
void *f(void **xx) {  
    sleep(1);  
    printf("this is f\n");  
    exit(0);  
}  
  
int main() {  
    pthread_t tid;  
    pthread_create(&tid, NULL, f, NULL);  
    pthread_join(tid, NULL); // line X  
    printf("this is main\n");  
}
```

Orphan
thread



Q: What you will see?

A: this is f\n
this is main\n

B: this is f\n

C:
this is main\n
this is f\n

1. Example to illustrate interleavings: say that thread A executes f() and thread B executes g(). (Here, we are using the term "thread" abstractly. This example applies to any of the approaches that fall under the word "thread".)

a. [this is pseudocode]

```

7   int x;
9
10  int main(int argc, char** argv) {
11
12      tid tid1 = thread_create(f, NULL);
13      tid tid2 = thread_create(g, NULL);
14
15      thread_join(tid1);
16      thread_join(tid2);
17
18      printf("%d\n", x);
19  }
20
21
22  void f() {
23      x = 1;
24      thread_exit();
25  }
26
27  void g() {
28      x = 2;
29      thread_exit();
30  }

```

What are possible values of x after A has executed f() and B has executed g()? In other words, what are possible outputs of the program above?

b. Same question as above, but f() and g() are now defined as follows

```

39  int y = 12;
40
41  f() { x = y + 1; }
42
43  g() { y = y * 2; }

```

What are the possible values of x?

c. Same question as above, but f() and g() are now defined as follows:

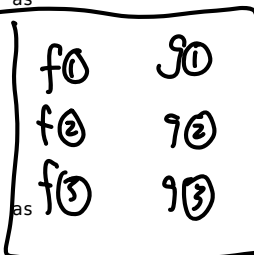
```

50  int x = 0;
51
52  f() { x = x + 1; }
53  g() { x = x + 2; }

```

What are the possible values of x?

$x \in \{1, 2, \text{or } 3\}$



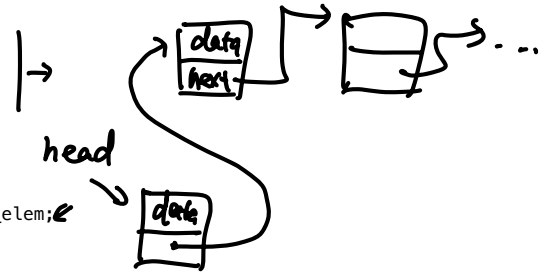
$x = x + 1 \Rightarrow$

2. Linked list example

```

61  struct List_elem {
62      int data;
63      struct List_elem* next;
64  };
65
66  List_elem* head = 0;
67
68  insert(int data) {
69      List_elem* l = new List_elem;
70      l->data = data;
71      l->next = head;
72      head = l;
73  }

```

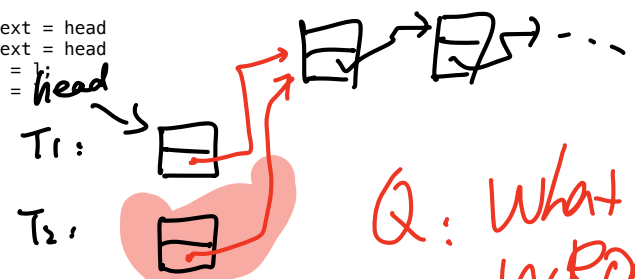


What happens if two threads execute insert() at once and we get the following interleaving?

```

79  thread 1: l->next = head
80  thread 2: l->next = head
81  thread 2: head = l;
82  thread 1: head = l;

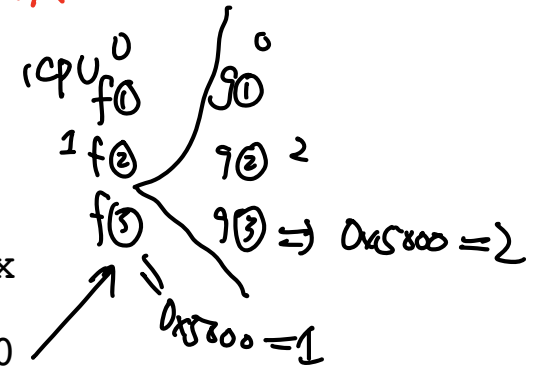
```



LOST DATA

Q: What's wrong?

0x5000: $x = 1$



```

① movq 0x5000, %rbx
② addq $1, %rbx = 1
③ movq %rbx, 0x5000

```

86 3. Producer/consumer example:

```

87  /*
88  89  "buffer" stores BUFFER_SIZE items
90  91  "count" is number of used slots, a variable that lives in memory
92  93  "in" is next empty buffer slot to fill (if any)
94  95  "out" is oldest filled slot to consume (if any)
96  */

```

```

96 void producer (void *ignored) {
97     for (;;) {
98         /* next line produces an item and puts it in nextProduced */
99         nextProduced = means_of_production();
100         while (count == BUFFER_SIZE)
101             ; // do nothing
102         buffer [in] = nextProduced;
103         in = (in + 1) % BUFFER_SIZE;
104         count++;
105     }
106 }

```

```

108 void consumer (void *ignored) {
109     for (;;) {
110         while (count == 0)
111             ; // do nothing
112         nextConsumed = buffer[out];
113         out = (out + 1) % BUFFER_SIZE;
114         count--;
115         /* next line abstractly consumes the item */
116         consume_item(nextConsumed);
117     }
118 }

```

```

121 /*
122 what count++ probably compiles to:
123 reg1 <-- count # load
124 reg1 <-- reg1 + 1 # increment register
125 count <-- reg1 # store
126
127 what count-- could compile to:
128 reg2 <-- count # load
129 reg2 <-- reg2 - 1 # decrement register
130 count <-- reg2 # store
131 */

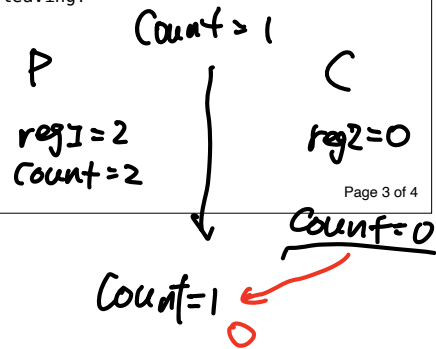
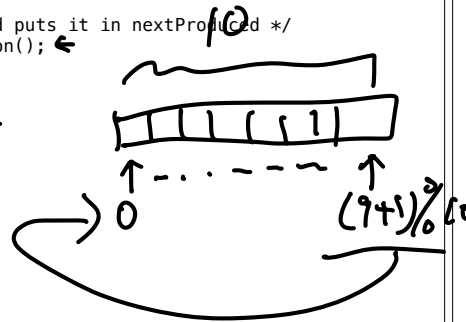
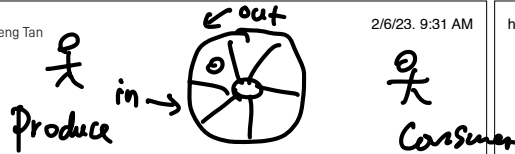
```

132 What happens if we get the following interleaving?

```

133 C | P
134 | |
135 | | reg1 <-- count = 4
136 | | reg1 <-- reg1 + 1 2
137 | | reg2 <-- count
138 | | reg2 <-- reg2 - 1 0
139 | | count <-- reg1
140 | | count <-- reg2
141 | |

```



142

143 4. Some other examples. What is the point of these?

144 [From S.V. Adve and K. Gharachorloo, IEEE Computer, December 1996, 145 66-76. <http://sadve.cs.illinois.edu/Publications/computer96.pdf>]

146 a) Can both "critical sections" run?

```

147 int flag1 = 0, flag2 = 0;
148
149 int main () {
150     tid id = thread_create (p1, NULL);
151     p2 (); thread_join (id);
152 }

```

```

153 void p1 (void *ignored) {
154     flag1 = 1;
155     if (!flag2)
156         critical_section_1 ();
157 }

```

```

158 void p2 (void *ignored) {
159     flag2 = 1;
160     if (!flag1)
161         critical_section_2 ();
162 }

```

163 b. Can use() be called with value 0, if p2 and p1 run concurrently?

```

164 int data = 0, ready = 0;
165
166 void p1 () {
167     data = 2000;
168     ready = 1;
169 }

```

```

170 int p2 () {
171     while (!ready) {}
172     use(data);
173 }

```

174 c. Can use() be called with value 0?

```

175 int a = 0, b = 0;
176
177 void p1 (void *ignored) { a = 1; }
178
179 void p2 (void *ignored) {
180     if (a == 1)
181         b = 1;
182 }

```

```

183 void p3 (void *ignored) {
184     if (b == 1)
185         use (a);
186 }

```

Race conditions (1 CPU) \Rightarrow multi-core

Sequential Consistency \leftarrow Contract \rightarrow [mem]

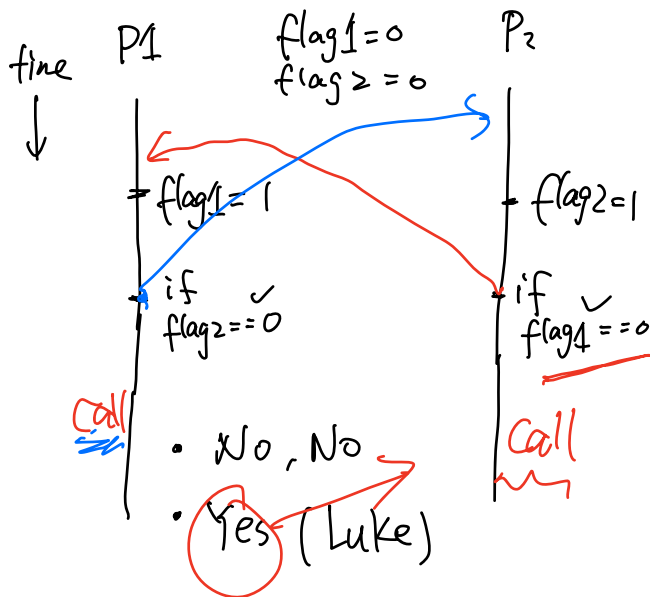
\exists Seq exec of the ops. s.t.

seq exec \equiv What you've seen.



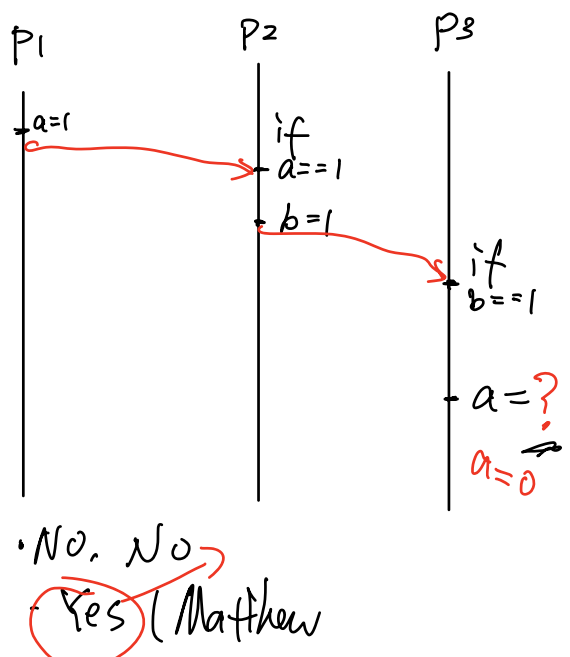
- 6
- R₁ W₂ W₃
 - R₁ W₃ W₂
 - W₂ R₁ W₃
 - W₂ W₃ R₁
 - W₃ R₁ W₂
 - W₃ W₂ R₁

handout 4.a.



4.c

a=0, b=0



```

int x = 0; // a global variable
void foo() {
    for (int i=0; i<100; i++) {
        x = 1;
        printf("%d", x);
    }
}
void bar() {
    x = 0;
}

```

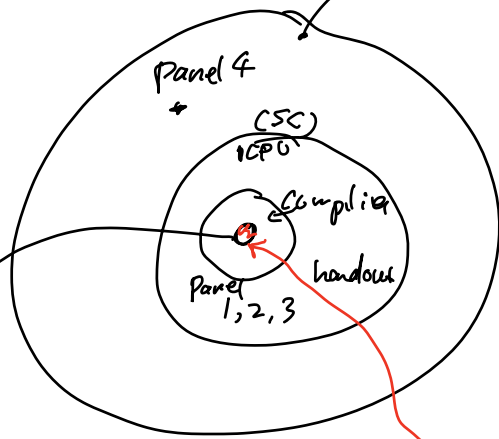
thread

man model

database

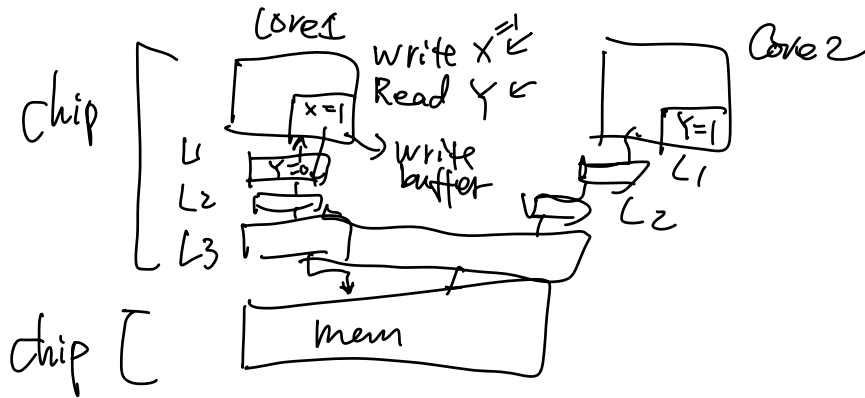
real-world

||||| ...
 10||| ...

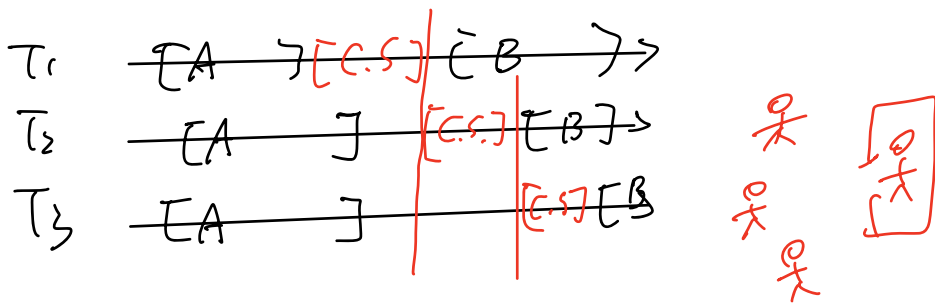
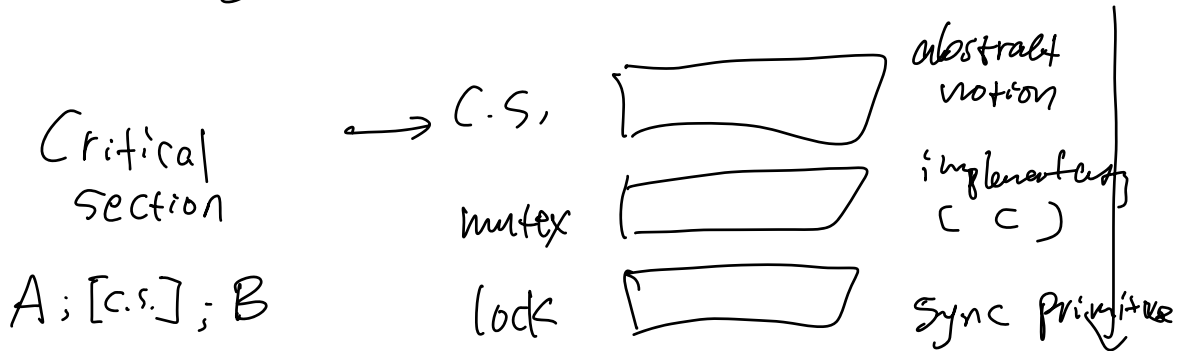


Trade-off; expect program

"Correctness" vs. perf



• managing Concurrency.



- mutual exclusive
- Progressive ←
- bounded waiting ←