

A=0; B=0;

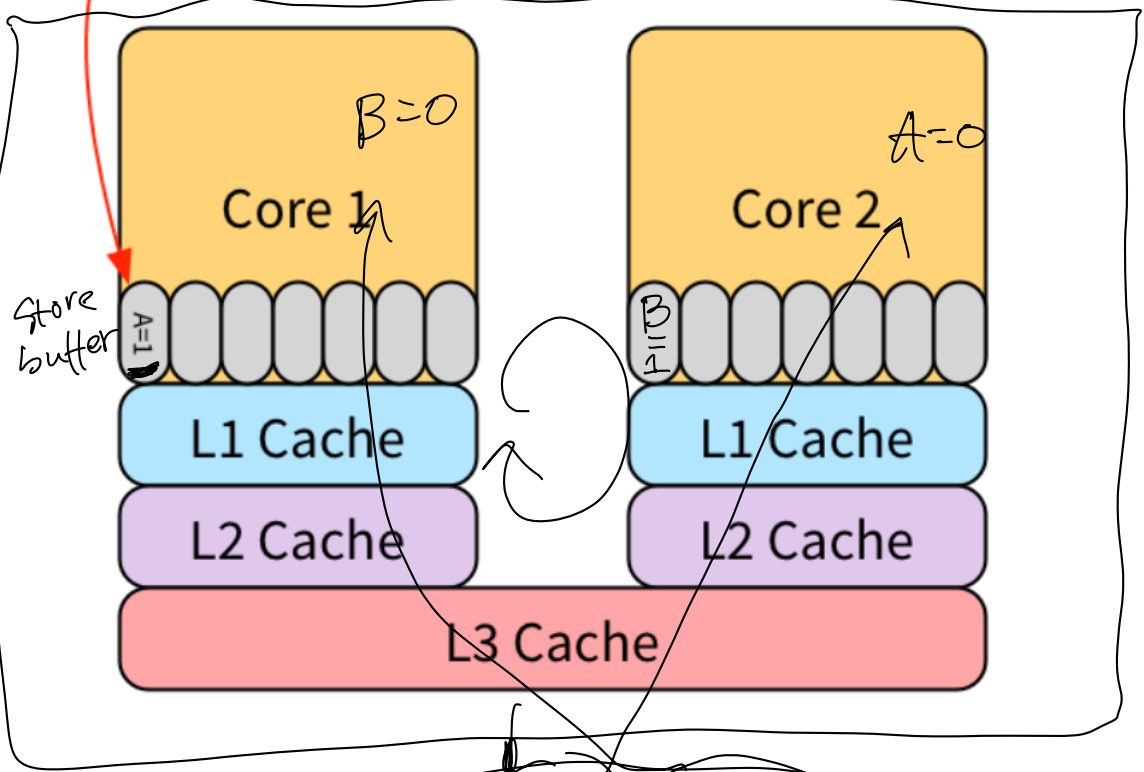
Thread 1

Thread 2

(1) A = 1  
(2) print(B)

(3) B = 1  
(4) print(A)

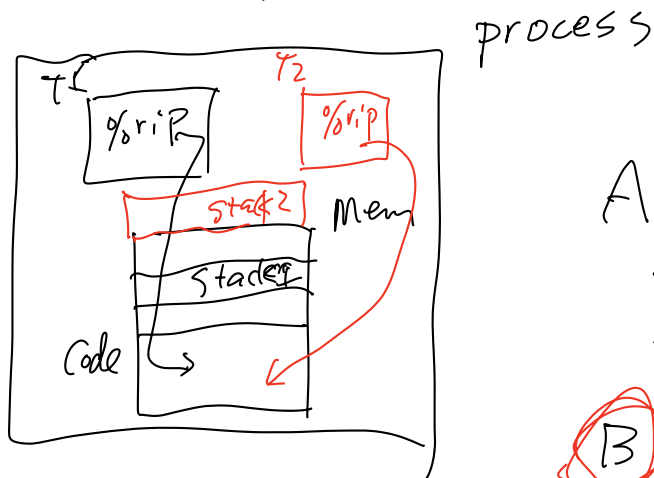
CPU  
TSO



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

- 0. Processes vs. threads
- 1. Intro to concurrency
- 2. Managing concurrency
- 3. Mutexes

Q: What's the difference between processes and threads?


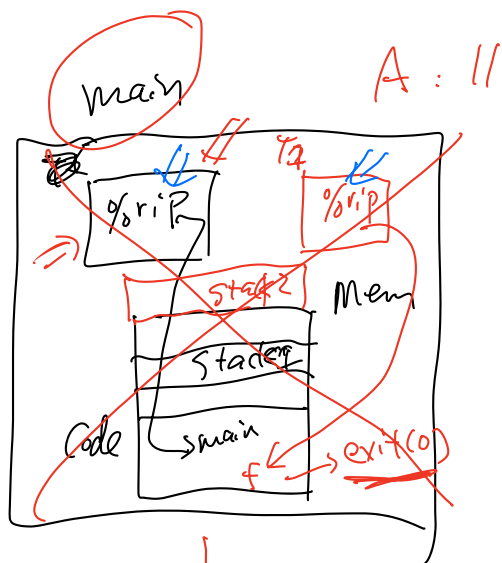


A:  
 this is f  
 this is main

**B**:  
 this is f

C:  
 this is main

D:  
 others



# 1. Intro to Concurrency



# Sequential Consistency

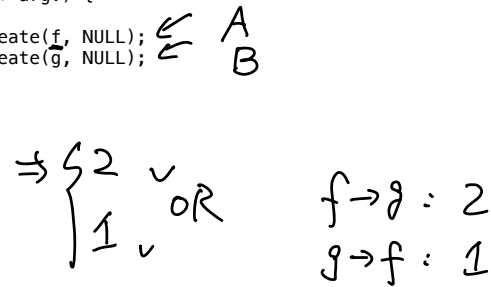
1. Example to illustrate interleavings: say that thread A executes f() and thread B executes g(). (Here, we are using the term "thread" abstractly. This example applies to any of the approaches that fall under the word "thread".)

a. [this is pseudocode]

```

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

```




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

```

37 int x;
38 int y = 12; // global
39
40 f() { x = y + 1; }
41 g() { y = y * 2; }

```

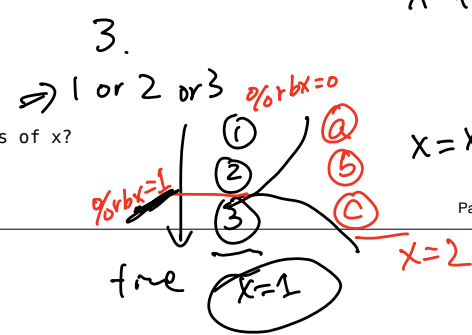


c. Same question as above, but f() and g() are now defined as follows:

```

47 int x = 0; // global
48
49 f() { x = x + 1; }
50 g() { x = x + 2; }


```



```

58
59
60 2. Linked list example
61
62 struct List_elem {
63     int data;
64     struct List_elem* next;
65 };
66
67 List_elem* head = 0;
68
69 insert(int data) {
70     List_elem* l = new List_elem;
71     l->data = data;
72     l->next = head;
73     head = l;
74 }

```




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;

```



address of x

```

X = X + 1 =>
1 movq 0x5000, %rbx # load from address 0x5000 into register
2 addq $1, %rbx # add 1 to the register's value
3 movq %rbx, 0x5000 # store back

```

```

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

```

1. Some concurrent programs. What is the point of these?


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

a. Can both "critical sections" run?

```

10  int flag1 = 0, flag2 = 0;
11
12  int main () {
13      tid id = thread_create (p1, NULL);
14      p2 (); thread_join (id);
15  }
16
17  void p1 (void *ignored) {
18      flag1 = 1;
19      if (!flag2) {
20          critical_section_1 ();
21      }
22  }
23
24  void p2 (void *ignored) {
25      flag2 = 1;
26      if (!flag1) {
27          critical_section_2 ();
28      }
29  }

```




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

```

33  int data = 0, ready = 0;
34
35  void p1 () {
36      data = 2000;
37      ready = 1;
38  }
39  int p2 () {
40      while (!ready) {}
41      use(data);
42  }

```



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

```

46  int a = 0, b = 0;
47
48  void p1 (void *ignored) { a = 1; }
49
50  void p2 (void *ignored) {
51      if (a == 1)
52          b = 1;
53  }
54
55  void p3 (void *ignored) {
56      if (b == 1)
57          use (a);
58  }
59
60

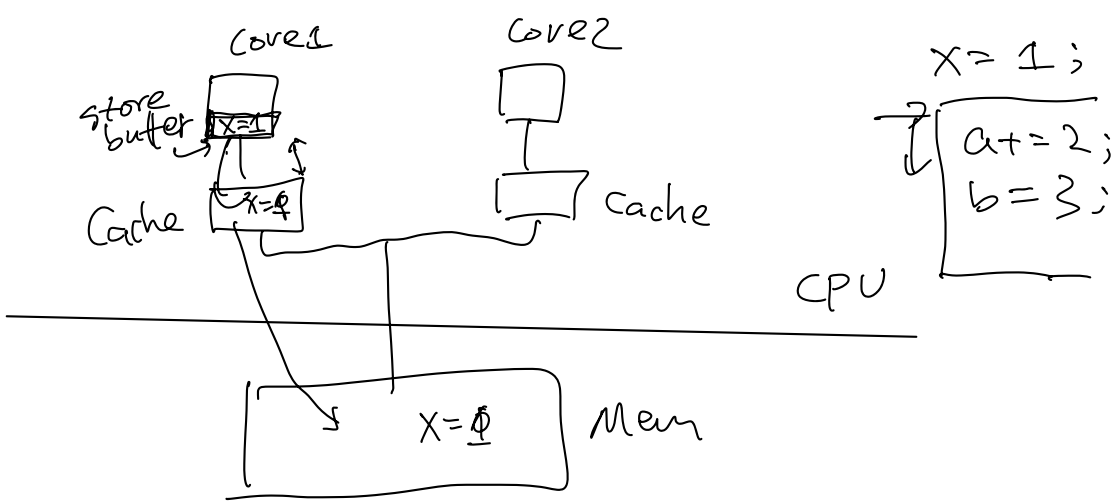
```

2. Protecting the linked list...

```

63  Mutex list_mutex;
64
65  insert(int data) {
66      List_elem* l = new List_elem;
67      l->data = data;
68
69      acquire(&list_mutex);
70
71      l->next = head;
72      head = l;
73
74      release(&list_mutex);
75  }
76
77

```



## Concurrency problem from software

```
int x = 0; // a global variable
```

```
void foo() { // T1
  for (int i=0; i<100; i++) {
    x = 1;
    printf("%d", x);
  }
}
```

```
void bar() { // T2
  x = 0;
}
```

```
|||||||...
||0|1|...
-----
||00...0
```

*bars*

could happen

