

Sequential Consistency

Exercise 1

For the following executions, please indicate if they are sequentially consistent. All variables are initially set to 0.

1. P1: W(x,1);
P2: R(x,0); R(x,1)
2. P1: W(x,1);
P2: R(x,1); R(x,0);
3. P1: W(x,1);
P2: W(x,2);
P3: R(x,1); R(x,2);
4. P1: W(x,1);
P2: W(x,2);
P3: R(x,2); R(x,1);
5. P1: W(x,1);
P2: W(x,2);
P3: R(x,2); R(x,1);
P4: R(x,1); R(x,2);
6. P1: W(x,1); R(x,1); R(y,0);
P2: W(y,1); R(y,1); R(x,1);
P3: R(x,1); R(y,0);
P4: R(y,0); R(x,0);
7. P1: W(x,1); R(x,1); R(y,0);
P2: W(y,1); R(y,1); R(x,1);
P3: R(y,1); R(x,0);

Exercise 2

P1	P2	P3
x = 1;	y = 1;	z = 1;
print(y,z);	print(x,z);	print(x,y);

All variables are stored in a memory system which offers sequential consistency. All operations, even the print statements, are atomic. Atomicity means, that no operation can overlap with other operations. Are the following sequences legal outputs?

1. 001011
2. 001111
3. 001110

Explain your answer by showing one possible interleaving of the instructions that might lead to the legal output. In case of an illegal output, explain why no possible interleaving exists.

Non-blocking caches

Consider a system with three processors. Each processor has a non-blocking cache (i.e., a processor does not stop upon encountering a cache miss; while the missing cache line is brought in from memory, the processor continues to execute instructions that are not data or control dependent upon a missing load).

The memory locations x and y are originally 0. The processors execute the following sequence:

P1: $W(x,1); R(y,0);$
 P2: $W(y,1); R(y,1); R(x,1);$
 P3: $R(y,1); R(x,0);$

Is this system sequentially consistent?

The x86 Memory Model: TLO+CC

1. Describe the memory model of the x86 architecture. Where does it differ from sequential consistency.
2. Two threads execute the following code (given in AT&T assembly syntax) on a machine using TLO+CC. Is it possible that the registers $eax = 0 \wedge ebx = 0$ after both threads have finished?. Would it be possible in sequential consistency?

Initial: All registers and memory locations are 0	
Thread A	Thread B
<code>movl \$1, 0x42</code>	<code>movl \$1, 0x50</code>
<code>movl 0x50, %eax</code>	<code>movl 0x42, %ebx</code>

In sequential consistency, it is impossible that both threads observe $eax=0 \wedge ebx=0$, however, in TLO+CC this is possible.

First, we will prove that it is impossible in sequential consistency. For this we first translate the above fragments into read and write statements:

Initial: All registers and memory locations are 0	
Thread A	Thread B
$W(a, 1)$	$W(b,1)$
$R(b, x)$	$R(a,x)$
$W(eax, x)$	$W(ebx, x)$

Now we assume that it **is** possible for both threads to finish and have $eax=ebx=0$ and show that this leads to a contradiction. Without loss of generality, we also assume Thread B is the last one to finish.

The first assumption tells us that the processes must execute the following read/write sequences:

$W_A(a, 1) \rightarrow R_A(b, 0) \rightarrow W_A(eax, 0)$
 $W_B(b, 1) \rightarrow R_B(a, 0) \rightarrow W_B(ebx, 0)$

the second assumption translates to

$W_A(eax, 0) \rightarrow W_B(ebx, 0)$

if we combine the three equations, we see that it is necessary to have $W_A(a, 1) \rightarrow R_B(a, 0)$ at the same time $W_B(b, 1) \rightarrow R_A(b, 0)$ which means that our initial assumption was false, and it is $eax=ebx=0$ is impossible in sequential consistency.

In the x86 memory model however, it is legal to reorder the given sequences into

Initial: All registers and memory locations are 0	
Thread A	Thread B
R(b, x)	R(a,x)
W(a, 1)	W(b,1)
W(eax, x)	W(ebx, x)

which allows the following interleaving:

$$R_A(b, 0) \rightarrow R_B(a, 0) \rightarrow W_A(a, 1) \rightarrow W_B(b, 1) \rightarrow W_A(eax, 0) \rightarrow W_B(ebx, 0).$$