Problem 1

Suppose we use MQMS on a machine with CPU 0 and 1. Each CPU has a scheduling queue.

- CPU will pick the first job to run in its own queue when scheduling.
- All CPU work, no I/O.
- Each queue uses RR policy.
- Time slice is 1ms.
- A new job will be added at the end of the queue with less jobs. (If equal, add to queue 0)
- Work stealing will happen every 5 ms. Each queue will peek at another queue. If that queue has 2 or more jobs more, steal the last job on that queue and put it at the end of itself.
- The running work is still in(at the first of) the queue.
- RR before work stealing before new job arrives. Assume the process of RR includes running a new job which has already in the queue.

Following are the jobs in the workload. At time 0 Job A and B is running.

Job	Arrival time	Length of time to run
A	--	3ms
B	--	10ms
C	1ms	4ms
D	2ms	8ms

Please draw the table to show the state of MQMS. Note: In time 0ms column, fill CPU0&CPU1 with the jobs that run during time [0,1), and the queues at time 0 (after all operations about queues).

Time	0ms	1ms	2ms	3ms	4ms	5ms	6ms	7ms	8ms
CPU0	A	A	C	A	C	C	C	--	--
Q0	A	AC	CA	AC	C	C	C	--	--
CPU1	B	B	B	D	B	D	B	D	B
Q1	B	B	BD	DB	BD	DB	BD	DB	BD
Time	9ms	10ms	11ms	12ms	13ms	14ms	15ms	16ms	17ms
CPU0	--	--	D	D	D	D			
Q0	--	D	D	D	D	D			
CPU1	D	B	B	B	B				
Q1	DB	B	B	B	B				

What if we perform work stealing every 7 ms? Fill in the blank again.

Time	0ms	1ms	2ms	3ms	4ms	5ms	6ms	7ms	8ms
CPU0	A	A	C	A	C	C	C	--	B
Q0	A	AC	CA	AC	C	C	C	B	B
CPU1	B	B	B	D	B	D	B	D	D
Q1	B	B	BD	DB	BD	DB	BD	D	D
Time	9ms	10ms	11ms	12ms	13ms	14ms	15ms	16ms	17ms
CPU0	B	B	B	B					
Q0	B	B	B	B					
CPU1	D	D	D	D					

Q1	D	D	D	D					

Problem 2

1. Assume our system uses LRU page replacement policy. Fill in the following table.

time	0	1	2	3	4	5	6	7	8	9	10	
Referencing page	-	3	4	2	6	4	3	7	4	3	6	
Stack contents	-	3	3	3	4	2	6	4	3	7	4	Page
After	-	-	4	4	2	6	4	3	7	4	3	Absent
reference	-	-	-	2	6	4	3	7	4	3	6	time
Page absent?	F	T	T	T	T	F	T	T	F	F	T	7

2. Assume the hit rate of LRU page replacement is 80%, while random page replacement policy has 70% hit rate. When using random policy, the cost is 100ns for accessing memory and 10ms for accessing disk. Assume LRU will spend additional 10ms too choose which page to evict among millions of pages in each miss. Calculate the AMAT of LRU and random.

AMAT (random) = 100ns*0.7 + 10ms*0.3 = 3.00007ms
AMAT (LRU) = 100ns*0.8 + (10ms+10ms)*0.2 = 4.00008ms