Lecture 20: Scheduling theory

Schedule = detailed assignment of tasks to a processor generated by a scheduler task
Major schedule = allocation of time blocks to a processor such that deadlines & periods of all tasks are met
Cyclic executive = program generated by a real-time kernel that executes the major schedule, along with other kernel responsibilities (eg memory management, I/O, task synchronization, IPC); it runs periodically itself and translates sporadic tasks into periodic ones for CPU time slicing purposes in the major schedule
Frame = length of a clock tick which serves as the most basic time unit of the major schedule
Major cycle time = major schedule period = hyperperiod = least common multiple period time of all tasks (example: T1's period is 10 sec, T2's period is 15 sec => major schedule period for T1 and T2 is 30 sec)
time slice = time that is allocated to one task during a major schedule period
Gantt chart = graphic timeline map of a major schedule, typically shown over 2 or 3 major schedule periods

Example:

task	exe time (e)	period (p)	deadline (d)	priority (P)
A	2	10	10	1
B	3	10	10	2
C	10	20	20	3

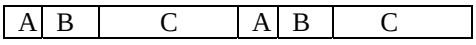

<-><---><----------->
2 3 5 <-frames
^ ^ ^ <-periods
A,B,C A,B A,B,C

Exercise: Does this represent 1 or 2 major cycle periods?
Exercise: Why did C stop after 5 frames the first time?
Exercise: What if the periods and deadlines of A,B,C were 15,20,22, respectively. What would the hyperperiod be?

Two general scheduling techniques:
1)Round-Robin scheduling – Fair allocation of CPU is given to all tasks according to their priority and a guaranteed time slice is given to all tasks in the major schedule period. A fixed priority is set at beginning of scheduling for all tasks.
2)pre-emptive priority scheduling – Kernel ensures that the CPU is allocated to the highest priority task that is ready to run. There are 2 kinds or pre-emptive priority scheduling: priorities are either fixed at beginning of schedule or dynamically set during scheduling.

Theorem: A process is schedulable if there exists enough processor cycles to execute all its computations.

Exercise:

task	exe time (e)	period (p)	deadline (d)	priority (P)
A	30	75	75	1
B	1	5	5	2
C	5	25	25	3

Can round-robin scheduling find a solvable major schedule for these 3 tasks?
Can fixed pre-emptive priority scheduling also find a solution?
What dynamic pre-emptive priority scheduling algorithm might find a solution?
Draw their Gantt charts for the total major cycle period.

Note about execution times (e) in these scheduling theory lectures: it is the sum of execution time, context switching time, scheduler bookkeeping time, and any other pre-emptive interrupt latencies that take place.