CS 4410: Operating Systems
Homework 4

Instructions for Homework 4:

- Homework may be done in pairs, or individually. If doing in pairs, **one** of you should upload to gradescope and add your partner to the group assignment in the upper right corner of the screen. (Do **not** just upload the assignment twice or it will be graded twice, which means grading will take longer.)
- The deadline is Wednesday, October 26th at 11:59AM.
- No late submissions will be accepted.
- You must attribute every source used to complete this homework.
- For some of the problems, you will need two integers. Here is the algorithm for computing these integers:
  - If you are working with a partner, let `var` be the lexicographically smaller of the two NetIDs.
  - Let `varInt` be the integral part of `var`. That is, if `var = rst12`, then `varInt = 12`.
  - If `varInt` is a single digit integer, let `varInt` = 13 × `varInt`
  - Let `Int1` be the first digit of `varInt`
  - Let `Int2` be the second digit of `varInt`
- **For all problems that use `Int1` or `Int2`, please write down the parameters (related variables and settings calculated from your NetID) before answering each question.**

1 Deadlines Are Marching

"OMG!", Ted couldn't believe his eyes after reading his email. His research progress report is awaiting, his audio recording for English class is haunting, a pile of papers is staling, and what's more, his TA responsibility is calling. Luckily he remembers that he is TAing CS4410; he comes to you and asks whether you can use your newly-gained knowledge to help him schedule these pending tasks. Each task i is given by a tuple (s_i, l_i, d_i), where

- s_i: when the tasks i arrives,
- l_i: the accumulated time it takes to finish task i,
- d_i: the deadline for task i.

The tasks are listed in the table below. Before answering the following question, please replace the length of $(\texttt{Int1}+1)$th task by $l'_{\texttt{Int1}} = l_{\texttt{Int1}} + (\texttt{Int1} \mod 5) + 1$, and replace $(\texttt{Int2}+1)$th task by $l'_{\texttt{Int2}} = l_{\texttt{Int2}} + (\texttt{Int2} \mod 5) + 1$.

i	s_i	l_i	d_i
0	0	2	5
1	3	2	7
2	5	1	8
3	7	2	14
4	0	1	3
5	1	3	8
6	5	4	12
7	0	8	10
8	1	3	5
9	5	3	15

Write down the scheduling order (in terms of task indexes), average completion time and tail completion time for each of the following scheduling policies.

1. FIFO (FCFS)
2. LIFO
3. SJF
4. SRTF
5. EDF

2 Not a Mathematician

Once again, Ted is in trouble; this time he's trying to understand why the SJF scheduling policy is considered optimal. More precisely, he can't convince himself that the average completion time is optimal using SJF, given the set of pending tasks and length of each. To simplify the problem, Ted thinks it is okay to assume all tasks come at once, meaning the starting time is the same.

1. Give a rigorous proof of the claim that SJF can achieve the optimal average completion time given the same starting time for all tasks.

2. Suppose that n tasks come at once, each of which is of length l_i. Suppose these tasks are already ordered so that $l_i < l_{i+1}\ \forall i \in [0, n)$. What's the average completion time for all these tasks?

3. Ted is tired of SJF and proposes another scheduling policy which appears to be non-deterministic (read: *much cooler*) called "monkey scheduling". Monkey scheduling is very simple and done by shuffling all the tasks evenly, which is to say, to sample a random permutation of all tasks uniformly (the probability for each permutation is equal).

   Ted feels very confident in his monkey scheduling policy and decides to use it for scheduling all the tasks at hand. However, Rachit expresses concern and tells Ted whatever policy Ted will use, he should always put two of the tasks given by Rachit as the first two tasks to finish. Therefore, Ted will only use monkey scheduling for the rest of the tasks. The indexes i_1, i_2 for two tasks given by Rachit is $i_1 = \texttt{Int1}$ and $i_2 = \texttt{Int2}$ (if $i_1 = i_2$, let $i_2 = (i_1 + 1) \mod 10$).

   Suppose there are 10 tasks whose length are $1, 2, 3, 3, 5, 8, 12, 12, 13, 17$ (the first task is indexed 0). Write down the *expectation* of average completion time using monkey scheduling.

4. Under the same requirement given by Rachit, write down the *expectation* of average completion time given n ($n \geq 10$) tasks start all at once, each takes l_i time to finish.

3 Network 101

Consider the following packet switching network. Each link (denoted by an edge) is annotated with bandwidth and distance. Assume the data can be transmitted with the speed of light within each link.

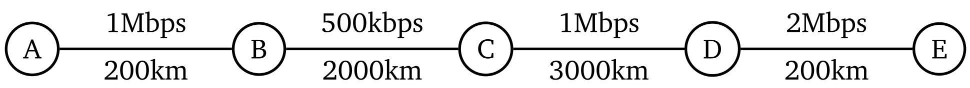

- Packet size: $P = 10000 + \mathtt{Int1} \cdot 1000$ bits,
- File size: $S = 10^6 + \mathtt{Int1} \cdot 10^5$ bytes,
- Number of packets: $N = 20 + \mathtt{Int2}$,
- Header size: $H = 40$ bytes.

1. The first node A sends a P-bit packet to B. How long does it take to arrive at B?
2. Now A wants to send more, for example, a S-byte file to B, so that it first divides the file into 2000 byte chunks, and prepends each chunk with a header containg necessary metadata to reassemble the chunks.
   (a) How long does it take for B to receive the entire file?
   (b) What is the goodput? Goodput is an applicaiton-level throughput and it means the number of payload bits arriving per second.
3. Consider sending N P-bit packets from A to E one after another. Assume each node in between maintains FIFO queues of size 5 for buffering the data in both directions.
   (a) How many packets are dropped?
   (b) Which packets are dropped? (Assume the packets are numbered as $1, 2, \ldots, N$, according to the sending order)
4. E sends an infinite stream of P-bit packet to A at its maximum rate. What is the fraction of lost packets?

4 The Furthest Distance in The World

> The furthest distance in the world
> Is not between death and life
> But when I shoot you in an online game
> Yet the latency always keeps you alive.
>
> — A Tragic Gamer

As the project manager of a video game company, Unisoft, Sherry has received numerous complaints from the gamers that the latency is too high to have a normal gaming experience. Therefore, she comes to you, the engineer in charge of game server networking for help. To diagnose the problem, you consider the transmission of data between a user's game client and your central server goes through a chain of intermediate switches, connected by L links (so that there are $L-1$ intermediate switches between the client and the central server). For each of L links, the bandwidth is B bits per second and there is 3×10^{-3} seconds of propagation delay due to the physical medium. For a packet sent and received for the online game, H bits are reserved for the metadata and the maximum packet size is M, meaning the effective payload is at most $M-H$.

The value of L,B,H,M,S,C for this problem is calculated as follow:

- Number of links: $L = 3 + (\texttt{Int1} \mod 4)$,
- Bandwidth: $B = 1000 + \texttt{Int1} \cdot 100$ bits per second,
- Metadata size: $H = 100 + \texttt{Int1} \cdot 10$ bits,
- Maximum packet size: $M = 2000 + \texttt{Int2} \cdot 100$ bits,
- Message size: $S = (M-H) \cdot (\texttt{Int2} \mod 4 + 1)$ bits,
- Virtual circuit setup packet size: $C = 0.8 \cdot M$ bits.

1. How long does it take for a client to transmit an S-bit message to the central server? Assume the network uses packet-based store-and-forward switching.
2. You soon noticed that it is possible to optimize the switching by letting a switch transmit the packet immediately after receiving the metadata (first H bits) of a packet, even if the packet hasn't been completely received yet. How long does it take to transmit an S-bit message using this technique?
3. A bored but rich gamer, Harold, comes to the company and wants to setup a special route for his gaming to reduced the latency effectively. The special route uses circuit switching after the transmission is initiated. When the client of the gamer wants to send a message to the central server, it first sends a special packet of C bits to initiate the transmission. The switches still pass on the packet using packet-based store-and-forward switching. However, after the receipt of this packet, a virtual "circuit" has been initiated and from now on, the switches will not use store-and-forward switching, but transmit each incoming bit to the next switch immediately. Therefore, the client first sends C bits and then it will transmit the S-bit message without metadata attached. How long does it take to send the message to the server?