CSCI 3310 – Demand Paging Exercises

1. Suppose that **m** is the access time for a resident page (including TLB overhead), **d** is the access time for a nonresident page (including page fault servicing time and all other overhead), and **p** is the probability of a page fault.

> (a) Compute the expected memory access time (**ema**) for an arbitrary page using demand paging in terms of **m**, **d**, and **p**.
>
> (b) Consider a system on which **m** is 100 nanoseconds and **d** is 10 milliseconds (equal to 10 million ns). Suppose that we're willing to tolerate a 10% slowdown in **ema** in order to accommodate demand paging (versus non-demand paging where all pages are resident). What does **p** have to be -- i.e., what proportion of page faults can we tolerate without exceeding this level of overhead?

---

2. Consider the page request reference stream shown below:

A B C A B D A D B C A

Suppose that the machine servicing this reference stream has three frames of physical memory to use. We simulated this scenario using a FIFO (first-in-first-out) page replacement policy and found that this policy resulted in 7 page faults.

> (a) Starting from an empty memory, simulate the behavior of the MIN (aka OPT) page replacement policy on this reference stream. How many page faults occur?
>
> (b) Starting from an empty memory, simulate the behavior of the LRU (least-recently-used) replacement policy on this reference stream. How many page faults occur?

---

3. Consider a reference stream consisting of a repeating pattern, such as shown below:

A B C D A B C D A B C D ...

Again from an empty memory, calculate the number of page faults with three physical frames and LRU replacement (just for the first 12 page requests of the pattern). How do you explain the resulting behavior?