# Solutions for the Second Quiz

COSC 6360

Spring 2014

# First question

- According to Navarro et al., what is the ***purpose*** of ***speculative demotion*** of superpages?

# Answer

- According to Navarro et al., what is the ***purpose*** of ***speculative demotion*** of superpages?
  - To find out if the superpage is still being actively used in its entirety.
  - Done at fixed intervals

# Second question

- Consider a system with 4 GB of RAM, a 1 MB L3 cache, and 4 KB pages. What is the minimum TLB size that could prevent most TLB misses for data already in the cache? (10 points)

# Answer

- Consider a system with 4 GB of RAM, an L3 cache, and 4 KB pages. What is the minimum TLB size that could prevent most TLB misses for data already in the cache?
  - The TLB should cover the whole L3 cache:
    - If cache size is 1 MB:
      - 1 MB/4 KB = 256 entries
    - If cache size is 2 MB:
      - 2 MB/4 KB = 512 entries

# Third question

- In the ARC cache replacement policy, which events result in an update of **target_T1?** (2×10 points)
  - **Target_T1 will increase when**
  - **Target_T1 will decrease when**

# Answer

- In the ARC cache replacement policy, which events result in an update of **target_T1?** (2×10 points)
  - Target_T1 will **increase** when a **page fault** brings in the cache a page in **B1**, the bottom part of L1
  - Target_T1 will **decrease** when a page fault brings in the cache a page in **B2**, the bottom part of L2

# Fourth question

- What is the main limitation of the Nooks ***extension recovery*** mechanism?

# Answer

- What is the main limitation of the Nooks ***extension recovery*** mechanism?
  - It does not work for all extensions:
    - Some extensions cannot be safely killed and restarted

# Fifth question

- Under which conditions do ***MCS locks*** perform best? Why?

# Answer

- Under which conditions do ***MCS locks*** perform best? Why?
  - MCS locks perform best at high contention rates because each contending core spins on a separate location.
    - Result is no cache coherence broadcasts

# Sixth question

- A system of physical clocks consists of two clocks, namely, one that is fast and advances by x minutes every hour and another that neither fast nor slow. Assuming that the clocks are managed by Lamport's physical clock protocol, what will be the time marked by each clock at 4 pm given that:
  - Both clocks indicated the correct time at noon;
  - The processors on which the clocks resides continuously exchanged messages between themselves from ***noon*** to ***two*** pm ; and
  - Message transmission delays are negligible.

# Answer for x = 4 minutes

| Actual time | Fast clock | Regular clock |
|---|---|---|
| Noon | 12:00 pm | 12:00 pm |
| 1 pm | 1:04 pm | 1:04 pm |
| 2 pm | 2:08 pm | 2:08 pm |
| 3 pm | 3:12 pm | 3:08 pm |
| **4 pm** | **4:16 pm** | **4:08 pm** |

# Answer for x = 5 minutes

| Actual time | Fast clock | Regular clock |
|---|---|---|
| Noon | 12:00 pm | 12:00 pm |
| 1 pm | 1:05 pm | 1:05 pm |
| 2 pm | 2:10 pm | 2:10 pm |
| 3 pm | 3:15 pm | 3:10 pm |
| **4 pm** | **4:20 pm** | **4:10 pm** |