

CSL862: Minor 2

Oct 25, 2010

Max Marks: 30

State your assumptions, if any. There are **6** total questions. You are allowed to skip 3-marks worth of questions from questions **1-4**. There are no 'bonus' marks for answering more.

1. **Memory Management**

- a. Due to inability to add a balloon driver to the guest OS, ESX decided to use a meta-level page-replacement policy. What policy should ESX use? Arrange the following meta-level page-replacement policies from the one that you expect to perform best to one that you expect to perform worst. Briefly explain. (2)
LRU, MRU, GuestPageReplacementPolicy, Random

b. ESX Memory Management paper has the following statement: "Allocation respects cache coloring by the guest OS; when possible, distinct PPN colors are mapped to distinct MPN colors". What would happen if allocation did *not* respect cache coloring?
(2)

c. On a memory-hungry workload, arrange the following configurations in order of performance. Briefly explain.

- i. VM configured with 1GB RAM with a 750MB balloon
- ii. VM configured with 250 MB RAM with 0MB balloon
- iii. VM configured with 200 MB RAM with 0MB balloon

Assume, balloon size is fixed. (2)

- d. The default value of the idle memory tax rate is 75%. Why not 100%? (1)

2. Nested Page Tables

- a. This is year 2050, and architects have realized that they need first-class support for nested virtualization in their architecture. Hence, they decide to implement Three-Dimensional Page Walks. If a single page walk (within a guest OS or a nested hypervisor) is n levels, how many page entry references are required for a memory access with a cold TLB. (1)

- b. In figure 3 on page-entry re-use, how do you expect the numbers to change if the guest was *not* using large pages. (2)

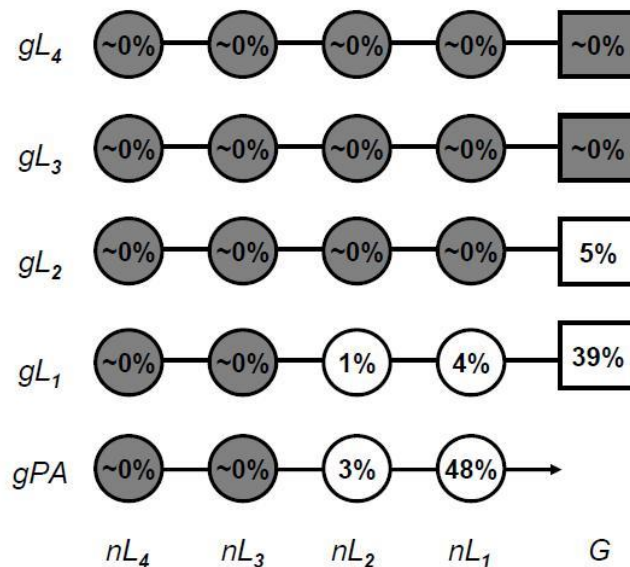


Figure 3. Percentage of all unique page entries for each 2D page walk reference in *MiscServer*

3. Comparison of Virtualization Techniques

Give concrete workloads when the following is true: (4)

- (i) Software (BT) VMM performs better than Hardware VMM without Nested Page Tables
- (ii) Software (BT) VMM performs better than Hardware VMM with Nested Page Tables
- (iii) VMM without NPT support performs better than VMM with NPT support

4. ReVirt for SMPs

- a. What is the main problem when trying to implement ReVirt for SMP (shared-memory multiprocessor) Systems? Give a concrete example, and demonstrate the problem. (2)

- b. Assume you can add support in hardware to efficiently support SMP record/replay of VMs. What support would you like to add. (3)

5. GFS

Describe how you would extend GFS so as to improve fault tolerance using multiple masters instead of one. (4)

6. Map Reduce

Consider a situation where multiple processes are receiving Twitter feeds (each filtered by a different set of keywords).

Suppose it is required to compute and display a continuously varying 'tag cloud' where the strength of each keyword corresponds to the number of tweets containing that word within a given T second interval.

- (i) What data needs to be computed in order to show such information? (2)
- (ii) How can this data be computed using MapReduce, in how many phases? (3)
- (iii) What is the parallel efficiency of your formulation? (2)
- (iv) Extend your solution so as to be able to (i) continuously update the tag-cloud as new tweets arrive, and (ii) allow the user to vary T (within reasonable limits). (3)