

## CSL862 Minor 1

Answer all 4 questions

01/09/2011

Max Marks: 20

### 1. Virtualization Techniques [14]

Give one example of a workload for each of the following comparisons. Ideally, you should provide a small snippet of code to specify the workload. The workload need not be anything useful. For each part, assume that all other things (apart from those mentioned) are identical in both configurations being compared. If you think, no workload exists that satisfies these conditions, say “No such workload”. It is possible for a workload to have two or more VMs.

- a. A workload that performs better on BT-based VMMs than on VT-based VMMs (VT-based is same as hardware virtualization). Assume best possible implementations for VT-based and BT-based VMMs. Assume unavailability of Nested Page Tables (e.g. Intel EPTs). [2]

- b. A workload that performs better on VT-based VMMs than BT-based VMMs. Assume best possible implementations for VT-based and BT-based VMMs. Assume unavailability of Nested Page Tables. [2]

- c. A workload that performs better on VT+EPT (VT and Nested Page Tables) than on VT without EPT. [2]

- d. A workload that performs better on VT without EPT than on (VT + EPT) [2]

e. A workload that performs better on bare-metal VMMs (e.g., VMware ESX) than on hosted VMMs (e.g., Qemu/KVM or VMware Workstation). [1]

f. A workload that performs better on para-virtual OS/VMMs (e.g., Xen with modified Linux guest) than on BT-based VMMs. [2]

- g. A workload that works better on VT+EPT than on VT+EPT+ELI. (ELI stands for Exitless Interrupts).  
[3]

## **2. Memory Management [2]**

The VMware ESX Memory Management paper was written for Shadow Page Tables, and described the utility of the following main ideas:

- a. Ballooning
- b. Content-based Page Sharing
- c. Idle Memory Tax

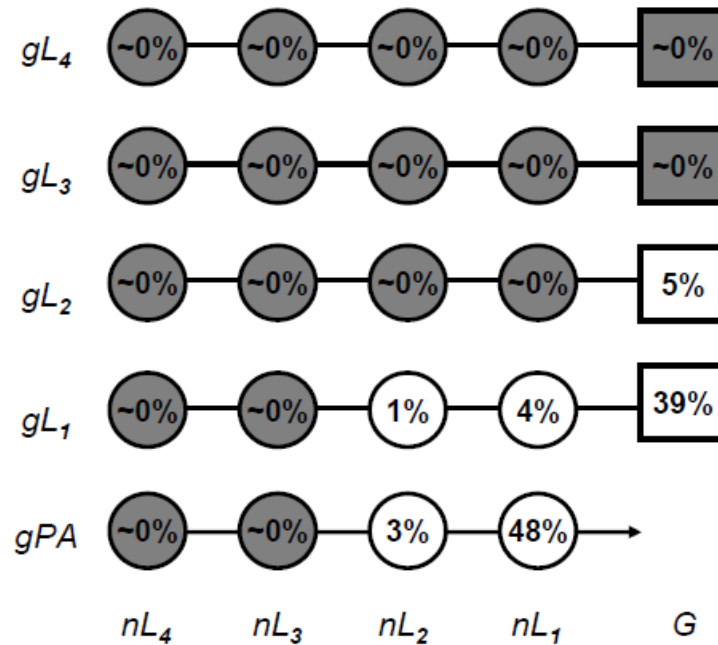
Which of these ideas are still applicable when using Nested Page Tables? Briefly explain.

**3. Paravirtualization [2]**

Xen uses paravirtualization to improve the consolidation ratio obtained through virtualization. Assuming that the guests are mostly idle (i.e., neither CPU-bound, nor I/O bound), could the consolidation ratio be infinity? What are some potential bottlenecks to scaling the consolidation ratio?



#### 4. Two-Dimensional Page Walks [2]



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

This data was generated assuming that the hypervisor uses 4KB pages. If the hypervisor uses huge (2MB) pages, how will this picture change? Also give a rough estimate of the new percentage of all unique page entries in each box in the new picture.