

CSL862 Minor 2

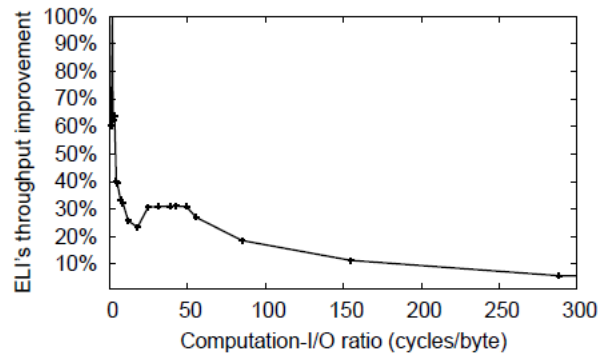
Answer all 7 questions

13/10/2011

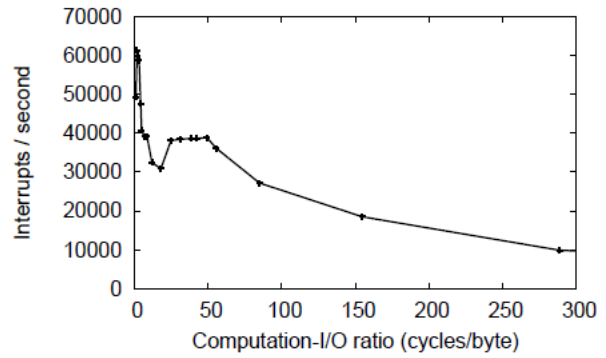
Max Marks: 25

1. ExitLess Interrupts

a. Throughput Improvements: Explain the following figure (Figure 5 from the ELI paper) [3 marks]



(a) ELI's throughput improvement



(b) Baseline interrupt rate

b. Security: Using ELI, can a malicious guest disable interrupts forever? If so, how does ELI recover from this situation. [1]

c. ELI grants the guest direct access to the EOI (end-of-interrupt) register. What can happen if the guest exits to the host without signaling the completion of in-service interrupts? [2]

2. vIC: Interrupt Coalescing for vStorage

- a. Many important applications issue synchronous IOs. Delaying the completion of prior IOs can delay issue of future ones. How does vIC ensure that coalescing in this case does not decrease throughput? [2]

- b. How does vIC upper-bound the average IO latency while coalescing? [2]

3. Protection Strategies for Direct Access to Virtualized I/O Devices

a. In Table 1, why doesn't the software-only approach (last line) not protect against bad device access for inter-guest protection? [1]

b. In Table 3, why is the CPU utilization for implementing protection highest in the "Single-Use" strategy. [2]

4. **vIOMMU**

- a. True/False: Direct device assignment (without vIOMMU) disallows memory overcommitment causing low consolidation ratios. Explain. [2]

- b. For sidecore emulation to be possible, the paper suggests that the device must follow a synchronous register write protocol. Why is this required? [2]

5. Uniprocessor Record/Replay

Arrange the following in increasing order of Uniprocessor Record/Replay overhead. If you expect two workloads to have the same overhead, indicate that using an “=” sign. For example, A>B=C>D

- A. Busyloop
- B. User-level Memsweep on Linux
- C. Kernel-level Memsweep on Linux
- D. “dd” utility to copy one file to another file inside the same ext3 filesystem on fully-replayed disk
- E. “dd” utility to copy one file to another file inside the same ext3 filesystem on output-replayed disk

[2 marks]

6. **Multiprocessor Record/Replay [3]**

- a. True/False: When using SMP-Revirt, one should prefer using blocking locks than spin locks. Briefly explain. [1]

- b. Write a program where using SMP-Revirt on two CPUs will be slower than running the program on one CPU and using uniprocessor Record/Replay. [2]

7. DoublePlay

a. If DoublePlay did not record the order of synchronization operations, would the system still work? What could be the problem? [1]

b. What is forward recovery (in case of a conflict) and why is it a good idea? What are the implementation complications and how are they resolved? [2]