

Question 10.1: Paging

- a. Explain the basic idea of paging.
- b. How is a virtual address translated to a physical address using a single-level page table?
- c. What is the disadvantage of using single-level page tables?
- d. Calculate the space requirements of a single-level page table for a system with 32-bit virtual addresses and a system with 48-bit virtual addresses (current x86-64). Both use a page size of 4 KiB. Assume that 4 bytes are required for a page table entry.
- e. What alternatives to using single-level page tables exist? What are their respective strengths and weaknesses?

- f. What is the purpose of the valid-bit (or present-bit) that is part of each page table entry?
- g. What is a translation lookaside buffer (TLB)?
- h. Why is reloading the CR3 register (physical address of the top-level page table) expensive? Propose a minor change of the hardware that could reduce this cost.

Question 10.2: Caches

- a. Describe the principle and the benefits of a memory hierarchy. How can a memory hierarchy provide both fast access and large capacity? On what typical program behavior does it depend?
- b. Cache memory is divided into (and loaded in) blocks, also called cache lines. Why is a cache divided into these cache lines? What might limit the size of a cache line?
- c. Enumerate and explain the different kinds of cache misses.
- d. Explain the difference between the write-through and write-back strategy.
- e. How can writes to a data item that is currently not in the cache be handled?
- f. What are the two main problems that can arise with caches that are virtually-indexed and virtually-tagged? What can be done to avoid or solve these problems?
- g. Does using virtually-indexed, physically-tagged caches solve the two problems described in the previous question?