

I. Architecture [50 %]

關於【計算機結構】的第一部分為是非題或選擇題，請利用答案紙上的第一頁，按照題號順序，標明題號，用一行回答一題。提供您的最終答案即可，請勿寫上計算過程。

第一部分：回答這部分的問題之前，請先閱讀以下這摘錄自科技網站上的文章片段：

Intel pitches ARM alternative to data centers—including Facebook's

by Jon Brodtkin - Dec 12 2012, posted on <http://arstechnica.com>

Intel today (2012.12.12) launched an Atom system-on-chip (SoC) line that combines extremely low power usage with server-class features, including virtualization technology and Error-Correcting Code (ECC) for higher reliability. The Intel Atom S1200 chips are for microservers as well as storage and networking systems that need energy efficiency and enterprise features that—Intel says—you just can't get in ARM chips. Intel called the S1200 "the world's first 6-watt server-class processor," and said microservers using the chip will be able to fit 1,000 nodes into a single rack.

ARM may dominate smartphones and tablets, but Intel hopes to lead the way in bringing smartphone CPUs to data centers. "Right now there are no ARM-based enterprise-class servers," Intel VP Diane Bryant said at a press conference for today's announcement. In addition to the hardware-assisted virtualization and ECC features already mentioned, Bryant noted that the Intel S1200 chips are 64-bit and support the x86 software prevalent in today's data centers.

The prospects of ARM servers from AMD (and some other companies) are intriguing, but ARM isn't a major player in the data center yet. With Atom S1200, Intel hopes to pre-empt ARM's entry into the server market. To prove the Atom chips' usefulness, Intel trotted out partners HP and Microsoft to talk about servers that will use the S1200 SoC and Windows Server's support for the new product line.

Today, Facebook VP of hardware design and supply chain Frank Frankovsky said, "Xeon-class processors have helped us scale Facebook very effectively so far." But not every workload needs what Frankovsky called "brawny cores." Now that "smartphone-class CPUs" for servers are 64-bit and include ECC, they're ready for certain parts of the Facebook infrastructure, he said.

"We've applied what I'll call brawny cores unilaterally across our environment," Frankovsky said. "What's interesting about these smartphone-class CPUs is we can right-size them to the needs of maybe the photo storage tier, for example. Maybe that's a great place to start, where we don't need a brawny core, what we need is maybe a smartphone-class CPU that also includes 64-bit and also includes ECC." Frankovsky noted that you might need two or three times as many Atom (or "wimpy") cores to do the same work Xeon-class processors handle, but ultimately come out ahead when measuring "how much useful work you can get done per watt and per dollar."

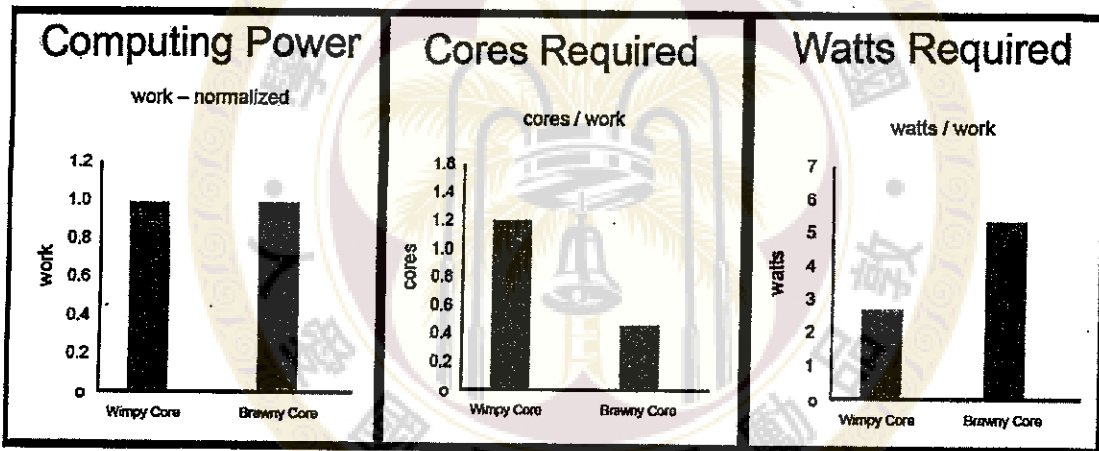
Those advantages only exist for certain types of workloads—the key is different CPUs for different applications. According to HP, compute-intensive applications still require Xeons, but the Atoms will be appropriate for "light scale-out applications" such as static Web serving and in-memory caching.

見背面

Please answer the following questions:

1. [2%] The ARM processor family is already widely used in the data center as enterprise-class servers, and Intel wants to compete in the same market by selling this new Atom S1200 chip. TRUE or FALSE?
2. [2%] Having a 64-bit ISA is critical to data center applications because a 64-bit ISA is much faster than a 32-bit ISA in terms of the computational speed. TRUE or FALSE?
3. [2%] According to the article, Facebook is interested in replacing some brawny cores (e.g. XEON) with ARM or Atom cores for the photo storage tier. TRUE or FALSE?
4. [2%] Hardware-assisted virtualization is important to the server market, because a modern virtual machine monitor can take advantage of the virtualization support to speed up virtual machines. TRUE or FALSE?

Now, to further understand this article, we found the original presentation on the Web, which includes the charts shown below. According to the data on the charts, Facebook has built two systems to perform a specific type of work: one was built with Wimpy cores (Wimpy-Core System, or WCS for short), and the other was built with Brawny cores (Brawny-Core System, or BCS for short).



Please answer the following questions:

5. [2%] In this case, the computing power refers to: (CHOOSE ONE)
 - (A) Latency, measured by the execution time required to get the work done
 - (B) Throughput, measured by the amount of work that can be done in a given time period
 - (C) Power consumption, measured by the cost of power dissipation for the system
 - (D) The power of ISA, measured by the average cycles per Instruction (CPI)
6. [2%] Based on the left-most chart, the two systems perform equally well. According to the article and the data on the other charts, which of the following is false? (CHOOSE ONE)
 - (A) A Brawny Core completes a task much sooner than a Wimpy Core does.
 - (B) The power efficiency of the BCS is about half of the WCS.
 - (C) The WCS has ~3 times more processor cores than the BCS has.
 - (D) Overall, the WCS consumes more power than the BCS.
7. [2%] Let us compare the amount of work that the system can get done per watt and per dollar. Suppose that one Wimpy Core costs one dollar and one Brawny Core costs 3 dollars, which of the ratio of work done per watt and per dollar between WCS and BCS? (CHOOSE ONE)

- (A) The ratio is about 0.9. Facebook should choose BCS.
- (B) The ratio is about 0.9. Facebook should choose WCS.
- (C) The ratio is 0.2. The Facebook should choose BCS.
- (D) The ratio is 0.2. The Facebook should choose WCS.

8. [2%] Which of the following applications may NOT be suitable for the new Intel S1200 chips? (CHOOSE ONE)
- (A) Web serving applications
 - (B) Big-data simple search systems
 - (C) Weather simulation and forecast engine
 - (D) Low-power network storage system

Now, let us try to design a system using the knowledge in the computer architecture course and what we just learned from the article.

As illustrated in the article, availability and speed are well-known metrics for servers, but power consumption is becoming increasingly important. Let us consider a big-data analysis system which runs the Apache Hadoop open-source software. The system is divided into two parts to handle different types of workload: (1) The data are distributed on the disks on the **Filesystem Serving nodes**. (2) The **MapReduce Computing nodes** acquire data from the Filesystem nodes and perform a MapReduce program in parallel.

The Hadoop runtime system handles the I/O details for the program, so the programmer can focus on the computation. The transactions between the Filesystem nodes and the Computing nodes are of 64MB size on average. Suppose the workloads running on the two types of nodes have the following parameters:

Workload	User-Space Instructions per Transaction	OS Instructions per Transaction for Disk Operation	OS Instructions per Transaction for Network Operation
Hadoop MapReduce Computing	200,000,000	0	100,000
Hadoop Filesystem Serving	100,000	30,000,000	30,000,000

Please answer the following questions:

9. [3%] Suppose we build the system with Brawny cores. A Brawny Core can perform 6 billion instructions per second. Assume the disk and network operations are NOT the bottleneck. How fast can a Brawny Core perform the transactions for each type of workload? (CHOOSE ONE)
- (A) 30 transactions/s for MapReduce Computing, 100 transactions/s for Filesystem Serving
 - (B) 30 transactions/s for MapReduce Computing, 600,000 transactions/s for Filesystem Serving
 - (C) 30,000 transactions/s for MapReduce Computing, 100,000 transactions/s for Filesystem Serving
 - (D) 5,000 transactions/s for MapReduce Computing, 10,000 transactions/s for Filesystem Serving

10. [3%] Suppose both the disk subsystem and the network interface are connected to the PCI Express x16 ports on each server. Each lane of a PCIe is 500 MB/sec. Assume the disk subsystem and network interface are NOT the bottleneck. What is the maximum transaction rate for the MapReduce Computing node and the Filesystem Serving node, each with 4 of the Brawny cores mentioned in 3(a)? (CHOOSE ONE)
- (A) 120,000 transactions/s for MapReduce Computing, 400,000 transactions/s for Filesystem Serving
 - (B) 120 transactions/s for MapReduce Computing, 400 transactions/s for Filesystem Serving
 - (C) 120 transactions/s for MapReduce Computing, 125 transactions/s for Filesystem Serving
 - (D) ~8 transactions/s for MapReduce Computing, ~8 transactions/s for Filesystem Serving
11. [3%] Assume the Wimpy Core is capable of 2 billion instructions per second, and it does twice as much performance per watt as the Brawny Core for Filesystem Serving, but only half the performance-per-watt of the Brawny Core for MapReduce Computing. Each Wimpy chip contains 2 Wimpy cores, and each Brawny chip contains 4 Brawny cores. Consider only the power consumption of the processors, which of the following design should be the most energy-efficient and cost-effective? (CHOOSE ONE)
- (A) 2 Brawny chips for a MapReduce Computing node, 2 Wimpy chip for a Filesystem Serving node
 - (B) 1 Brawny chip for a MapReduce Computing node, 1 Brawny chip for a Filesystem Serving node
 - (C) 1 Wimpy chip for a MapReduce Computing node, 1 Brawny chip for a Filesystem Serving node
 - (D) 1 Brawny chip for a MapReduce Computing node, 1 Wimpy chip for a Filesystem Serving node

接次頁

第二部分：

12. [10%] Answer the following questions considering a streaming application with the following address stream (byte address):

0, 4, 8, 12, 16, 20, 24, (i.e., $4*n$, where $n=0, 1, 2, 3...$)

- a. [4%] What are the miss rates of a 32K direct-mapped cache and a two-way associate cache with a 32-byte line?
- b. [2%] What types of cache misses is this streaming application experiencing in the 3C model?
- c. [4%] Considering two mechanisms to improve cache performance, Victim Cache and Stream Buffer:

Victim Cache: adopting a small fully-associative cache to store cache blocks that have been replaced from the main cache. On a cache miss, both the main cache and the victim cache are searched. A hit in the victim cache is considered as a hit and the data are moved into the main cache.

Stream Buffer: prefetching sequentially adjacent cache lines into a separate buffer, when a particular cache line is brought in. If the data is found in the prefetch buffer, this is considered as a hit and moved into the cache and the next cache line is prefetched. Assume a two-entry stream buffer and assume that the cache latency is such that a cache line can be loaded before the computation on the previous cache line can be completed.

What are the cache miss rates for the Victim Cache mechanism with a 4-entry victim cache and the Stream Buffer mechanism with a two-entry stream buffer assuming a 32K direct-mapped main cache?

13. [12%] Assume a pipelined data path of a MIPS processor as shown in Figure 1.

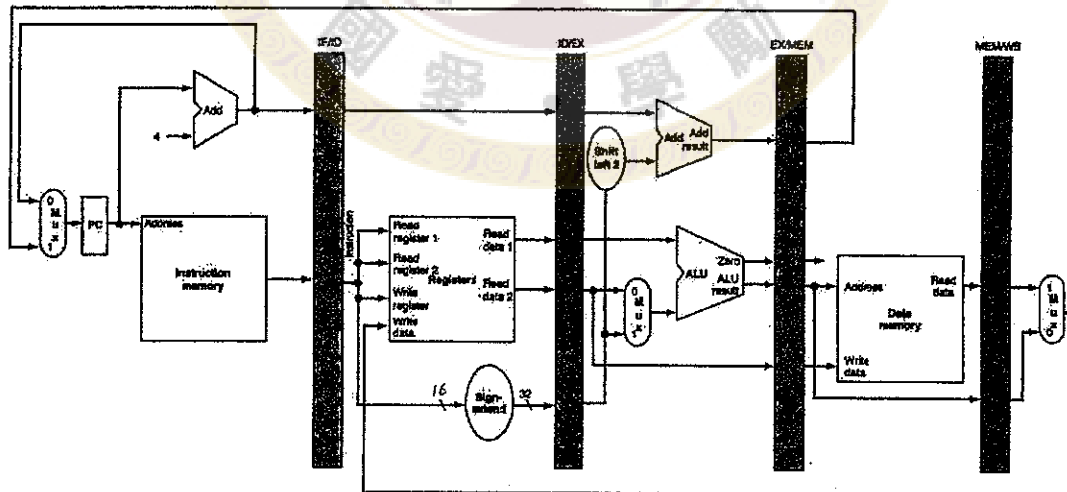


Figure 1: Pipeline Datapath

見背面

Answer this problem set assuming the following instruction mix and access latency of each pipeline stage and pipeline registers:

Add	beq	lw	Sw
45%	30%	15%	10%

Table 1: instruction mix

IF	ID	EX	MEM	WB	Pipeline registers
110ps	120ps	90ps	140ps	60ps	10ps

Table 2: Pipeline latency

- a. [3%] Assuming there is no stall and that 40% of all the conditional branches are taken, in what percentage of clock cycles does the branch adder in the EX stage generate a value that is actually used?
 - b. [3%] Assuming there is no stall, what is the speed-up achieved by pipelining compared to a single-cycle datapath?
 - c. [2%] Assuming there is no stall, what percentages of all cycles do we use the data memory?
 - d. [4%] Assume that the latency of the EX stage becomes 150 ps after adding the forwarding logics. If 20% of the instructions stall for one cycle before adopting forwarding, what percentage of stall cycles need to be eliminated by forwarding in order to get 10% speedup from adding the forward logics compared to the pipeline design without forwarding?
14. [3%] Which of the following statements are correct? (全對才給分)
- (A) The RISC instruction set can lead to more efficient processor implementation than the CISC instruction set.
 - (B) A data access that incurs a TLB miss cannot result in a page fault.
 - (C) Increasing the cache associativity can reduce conflict misses
 - (D) In a virtually-indexed and physically-tagged cache, the tag bits are taken from virtual addresses.

接次頁

II. Operating System [50%]

A **bicycle sharing system (BSS)** is a service in which bicycles are made available for shared use to individuals who do not own them. The central concept of these systems is to provide free or affordable access to bicycles for short-distance trips in an urban area as an alternative to motorized public transport or private vehicles, thereby reducing traffic congestion, noise, and air pollution. As of May 2011 there were around 136 BSSes in 165 cities around the world, made of an estimated fleet of 237,000 bicycles. U-Bike, Taipei BSS of 11 stations and 500 bikes since 2008, will have 162 stations with 5359 bikes by 2018. The number of parking slots in each station is fixed and might range from 10s to 90s. So far, there are over 10000 accesses each day in average. The Taipei U-Bikers can rent a bike from any station and return to any from an unmanned distributed system. They need to wait if there is no available bike or parking slot. We can get the availability in real time from Internet servers. There are some trucks shipping bikes from stations to stations for better utilization. The rental fee, now, is free for the first 30 minutes and 10 NTD for the following each 30 minutes.

Please answer the following questions:

15. [15%] To support the service on all the stations, the server needs to collect all the data including user's transaction, bike availability, etc. and process users' requests. Alice is the system architect for U-Bike service. To support concurrent requests from all the bike stations, she proposes to use a multiple process model to design the server application. On the server, a background process is started to listen to the requests from all the stations. When a request arrives, a new process is created to process the request.
- [3%] Please discuss if Alice should use either `fork()` or `vfork()` for the background daemon process to create new process so as to provide better concurrence in this application? Please compare the pros and cons of these two functions and select the better one.
 - [3%] Copy-on-write is the technique used to in `fork()` and `vfork()` to copy memory regions. Please (1) define technique Copy-on-write and (2) answer which memory segments in a process will be benefit from copy-on-write?
 - [2%] How can Alice know the maximum of concurrent processes being supported in the operating systems?
 - [2%] Suppose each request requires 100ms to process and there is no volunteer yielding during the execution. Moreover, there are 10,000 requests in the ready queue and there are no new requests. The operating system employs Round Robin Scheduling algorithm to schedule all the processes and the time quantum for each process is 50ms. Suppose that there is no other user processes and the CPU time used by kernel processes are negligible. What's the average response time of all the requests?
 - [2%] After a detail investigation, Alice found that the execution time of processing requests evenly distributed in interval [5ms, 200ms] and the exact execution times are

not known when requests arrive. Which of the following scheduling algorithm should be selected to achieve maximum throughput? Candidate scheduling algorithms include Round-Robin Scheduling, First-Come-First-Serve, Shortest-Job First, Multi-level Feedback Queue algorithm.

- f. [3%] Suppose that the system is required to support at least 1,000 transactions per second and each process can create 100 concurrent threads. The context switch overhead for each process is 1ms and that for each thread is 0.1ms. Note that the execution time ranges from 5ms to 200ms. How should Alice revise her design? You should take into account that the execution time of all requests are not equal, the scheduling algorithm used in the operating systems, and the overhead of managing processes and threads.
16. [10%] After Alice implemented the system and started to test, she found that the performance is pretty poor. She was asked to identify the causes and resolve the issue.
- a. [2%] She found that there are large numbers of context switch among threads and processes and no much real works are done. Please name and define such scenario.
- b. [2%] When she checked the logs in the operating system on server, she found that many recently and frequently used memory pages are swapped in and out due to high demanding on U-Bike. Hence, she checked the textbook and decided to take page-fault frequency (PFF) approach to solve the problem. Please describe this technique.
- c. [3%] Unfortunately, PFF approach does not help. She also observed that the execution time of processing a request is not long and can be finished within few context switches when the system load is not high. Moreover, when a process/thread is scheduled, it needs to certain page frames for a certain amount time until it completes. What approach should Alice use to avoid thrashing when the system load is high?
- d. [3%] Alice also finds that the system is short of memory due to large number of processes to process U-bike requests. How can Alice lower the memory usage if she is not allowed to change the number of processes?
17. [10%] How can you schedule the trucks with minimum cost to reduce the average waiting time of renting a bike [5%], and returning a bike [5%]? Hint: during rush hours or not.
18. [10%] How to function the system even when the network is down temporarily, where stations cannot communicate with servers but still work with their own parking slots? Hint: stateful or stateless, timing, synchronization, rollback, what information to keep?
19. [5%] How do you design a booking or reserving service to improve user satisfaction? Hint: for bike and for parking slot.