


Computer Architecture

1. [14 pts] I/O and Storage

Redundant Arrays of Inexpensive Disks, as named by the inventors and commonly referred to as RAID, is a technology that supports the integrated use of two or more hard-drives in various configurations for the purposes of achieving greater performance, reliability through redundancy, and larger disk volume sizes through aggregation.

RAID 0 is a striped set without parity. *RAID 1* is a mirrored set without parity. The two different schemes are illustrated below. A1 stores the first block of file A, A2 stores the second block of file A, etc.



Let the size of the disk be 320GB each.

- (a) [1 pt] What is the total capacity of a *RAID 0* set with 2 disks?
- (b) [1 pt] What is the total capacity of a *RAID 1* set with 2 disks?
- (c) [2 pts] For write operations, which RAID set has the higher throughput?
- (d) [2 pts] For read operations, can the *RAID 1* set deliver the same throughput as the *RAID 0* set?
- (e) [2 pts] Which RAID set has the higher availability?

For the following questions, suppose that you have 4 disks and you are considering the following RAID configuration:

RAID 0: a 4-disk striped set

RAID 0+1: 2 striped sets in a mirrored set

RAID 1+0: 2 mirrored sets in a striped set

- (f) [2 pts] Draw a picture to illustrate how a file is stored on the *RAID 0+1* set.
- (g) [2 pts] Which configuration has the highest write throughput?
- (h) [2 pts] When one of the disks has failed, which configuration has the highest write

見背面


throughput?

2. [11 pts] Multiprocessor Systems

Multiprocessor systems can be found in everyday use. Many new personal computers (PC) are equipped with multi-core processor chips. Because the processor cores on such a system typically share everything on the system bus and the I/O bus, this kind of system is classified as *Symmetric MultiProcessor (SMP)*.

- (a) [1 pt] The multi-core design is popular because most existing applications running on a PC would benefit from the additional processor cores. *True or False?*
- (b) [1 pt] Having multiple cores share the Level-2 cache on the same chip is always a good idea because most of the data in the cache are shared by multiple cores. *True or False?*
- (c) [3 pts] The shared-memory programming model is most natural for developing a parallel application on an SMP system. For that, which of the following can be used? (You can have multiple choices.)

Message-Passing Interface (MPI)
OpenMP
POSIX Threads
XML
HTML
Java



People sometimes pay too little attention to I/O and the operating system (OS) when they develop parallel applications for SMP systems. For example, an application has been fully parallelized to run on a SMP system, but for an 8-processor run, 50% of time a processor has to stall because the processor is waiting to access the disk. There is only one disk in the system. Assume that the OS issues one disk accesses at a time.

- (d) [2 pts] Suppose that the application runs for 100 seconds when only one processor is used, how many seconds would the application requires to run on the 8-processor SMP system? (Note that the 100 seconds is the wall clock time, measured from the beginning to the end of the application).
- (e) [2 pts] For the 8-processor run, suppose that the disk is transferring data during 80% of the execution time. What is the minimum execution time even if the system has unlimited number of processors?
- (f) [2 pts] In order to improve the application performance, the disk is replaced with a 5-disk *RAID 0* system. What is the minimum execution time required for this application, with the new RAID disk?

3. [10 pts] True or False

- (a) VLIW is a software_based approach to exploit ILP
- (b) Pipelining improves both instruction latency and throughput.
- (c) The delayed branch mechanism has lost popularity compared to more expensive but more flexible dynamic branch prediction as processors go to both longer pipelines and wider issue width.
- (d) Two instructions with data dependencies will cause data hazard in pipelining execution.
- (e) A one-bit branch predictor could sometimes achieve higher branch prediction accuracy than a two-bit branch predictor.

4. [10 pts] For the following memory system parameters:

- Cache: virtually-addressed but physically-tagged, 8KB, direct-mapped, 32B block size
- TLB: fully associative, 32 entries
- 32-bit physical address
- Virtual Memory: 8K page, 64-bit virtual address

- (a) (5 pts) How many bits are needed in each tag, index, and block offset fields of the cache?
- (b) (2 pts) How many bits are needed in each tag field of the TLB?
- (c) (3 pts) Why do we want to use a virtually-addressed but physically-tagged cache?

5. [3 pts] Multiple Choices

An embedded processor has both 32 and 16-bit instructions. 16-bit instructions are a subset of the most commonly used 32-bit instructions. Compared to codes written in 32-bit instructions, codes written in 16-bit instructions could potentially

- (a) [1 pt] increase or decrease instruction count?
- (b) [1 pt] increase or decrease code density?
- (c) [1 pt] increase or decrease I-Cache hit rate?

6. [2 pts] List one advantage of variable-length encoding over fixed-length encoding.

System Area (The following four questions count 50 points of your total credit.)

The elections in Taiwan are conducted by stamp-based (or called mark-based) ballots, which are error prone and subject to electoral fraud. It requires well trained supporting staffs, who may give up their rights for voting, and intelligent voters to have a fair election. In 2004 Presidential Election in Taiwan, there were 16 millions legitimate

voters and 13 millions of them cast their votes. Around the country, 13,305 polling stations are set up from 8AM to 4PM and more than 200,000 supporting staffs are hired. After the stations were closed, it took five hours to count all the votes. Among all the cast votes, 337,297 (2.59%) votes were invalid for various reasons.

In order to shorten the time for counting the votes and to reduce the number of invalid votes and the number of supporting staffs, the Central Election Commission (CEC) announced a Call for Proposal to design an Electronic Voting System (EVS) for 2008 Presidential Election. With the EVS, the legitimate voters have to cast their votes using the electronic voting machine (EVM) in polling stations.

You, a senior software architect of Toogle (a software company in Taiwan), are responsible for writing the proposal for competing the project. In the following, please show and argue your system design. NOTE that you should write down your own assumptions if the information given in the quests are not sufficient.

7. [15 pts] Architecture Design:

In November 2004 [1], P. L. Vora etc. published an article on Communication of ACM for the evaluation metrics of electronic voting systems. The metrics include

- Integrity of the votes (both voter verification, "I can check that my vote was captured correctly" and public verification, "anyone can check that all recorded votes were counted correctly");
- Ballot secrecy (both voter privacy and resistance to vote selling and coercion);
- Robustness (including resistance to denial of service attacks);
- Usability and accuracy (including access for the disabled);
- Transparency (both of mechanism and election data).

In addition, the Central Election Commission requires the counting process must be completed within one hour after the polling stations are closed.

Please use a table exempld below to compare the pros and cons for any TWO out of four candidate system architectures listed below for EVSs with stamp-based voting system for the above five metrics. The candidate system architectures are

- Client-server systems: each EVM, i.e., the client, is connected to one of counting servers, i.e., the server, in CEC and there is no valuable data stored on the EVM.
- Peer-to-peer systems: each EVM is connected to one or many of the other EVM or counting servers. The voting results are transmitted to the server in CEC using the Peer-To-Peer communication protocol, which guarantees the results will be received at the end.
- Distributed stand-alone EVM: each EVM accepts the vote without connecting to any other computers and the results are sent to the counting server at the end of the election or at a constant time interval.
- Hierarchical client-server systems: all the devices are logically connected in a hierarchical manner like a tree. The root is the server in CEC and the leave nodes are the EVMs. The results are collected from the leave nodes to the root.

(The text in the table are examples, not answers, except for the header row and column.)

Metric	(System 1)	(System 2)	Stamp-based System
Integrity	Difficult to impose integrity because ...		Poor public verification because of anonymous voting and good private verification because the ballots can be recounted.
Secrecy	The secrecy is poor because ...		
Robustness			
Accuracy			
Transparency			

Reference:

- [1] P.L. Vora, B. Adida, R. Bucholz, D. Chaum, D.L. Dill, D. Jefferson, D.W. Jones, W. Lattin, A.D. Rubin, M.I. Shamos, and M. Yung, "Evaluation of Voting Systems", COMMUNICATIONS OF THE ACM, November 2004, Vol. 47, No. 11, Page 144.

8. [10 pts] Memory management:

In order to increase the throughput, multiple-process or multiple-thread programs are used to serve the counting requests from all the polling stations. Please answer the following questions.

見背面

- (a) [3 pts] How should the data be stored in memory to provide high throughput parallel processing? Please argue your answer. Hint that the number of received votes for each candidate can be either stored in one shared memory page or a number of different memory locations.
- (b) [7 pts] Lock-based approach, transaction-based approach and synchronization with barriers are three popular protocols to assure data consistence. Which one is the best for serving the counting requests? Please define the three terms, compare their pros and cons, and determine the best to use. Note that you should select the best protocol for the EVS.

9. [12 pts] File Systems and I/O systems:

The voting results need to be stored for integrity verification. When each EVM receives the (temporary and final) votes, the data must be stored in a persistent file system (either on the EVM, or the servers in polling stations and CEC). Please answer the following questions

- (a) [6 pts] Should the buffer IO or unbuffer IO be used for manipulating the data? Please compare the two mechanisms for performance and robustness, and explain your choose.
- (b) [6 pts] The available file systems are random access read/write and write-only log-based system. Which one is better for EVS? Please define and compare the two file systems, and explain your selection, based on the given evaluation metrics.

10. [13 pts] Process Scheduling:

When counting the votes, the servers may receive thousands of requests per second. To assure the consistence of the counting results, suppose that your system uses a locking mechanism and each counting request may take up to 200 ms to complete. Please answer the following questions.

- (a) [5 pts] Each request can be served as either a preemptive process or non-preemptive process. Please define the two terms and decide how each request will be served in your system.
- (b) [8 pts] Suppose that there are four process-scheduling algorithms in the operating system of the counting server: (1) round-robin scheduling, (2) shortest-job first scheduling, (3) first-come-first-serve scheduling, and (4) multilevel feedback queue scheduling. Please compare the average response time and worst-case response time for the four algorithms and select one to use in your system. Note that you should make the decision based on the given evaluation criteria.