※ 注意：全部題目均請作答於試卷內之「非選擇題作答區」，請標明題號依序作答。

Part I: Computer Architecture

1. (13 points) Recently, security experts revealed two critical hardware vulnerabilities (漏洞) in modern processors, *Meltdown* and *Spectre*. These hardware vulnerabilities allow programs to steal data which is currently processed on the computer. These two vulnerabilities have raised serious security concerns because they may affect many modern processors.

   Let us discuss Meltdown. The following is the abstract from the paper published by the experts who revealed Meltdown:

   The security of computer systems fundamentally relies on memory isolation, e.g., kernel address ranges are marked as non-accessible and are protected from user access. In this paper, we present Meltdown. Meltdown exploits side effects of out-of-order execution on modern processors to read arbitrary kernel-memory locations including personal data and passwords. Out-of-order execution is an indispensable performance feature and present in a wide range of modern processors. The attack is independent of the operating system, and it does not rely on any software vulnerabilities. Meltdown breaks all security assumptions given by address space isolation as well as paravirtualized environments and, thus, every security mechanism building upon this foundation. On affected systems, Meltdown enables an adversary to read memory of other processes or virtual machines in the cloud without any permissions or privileges, affecting millions of customers and virtually every user of a personal computer.

   To understand this paper, the reader needs to be familiar with modern computer architecture. Please answer the following questions:

   (a) [2 points] On a modern system, programs can run concurrently, but programs are typically not permitted to read the memory data from other programs. Please explain what architecture support is needed for this.

   (b) [2 points] As mentioned in the abstract, Meltdown exploits side effects of out-of-order execution on modern processors. Please explain why out-of-order execution is an indispensable performance feature on modern processors.

   (c) [3 points] In practice, CPUs supporting out-of-order execution support running operations speculatively to the extent that the processor's out-of-order logic processes instructions before the CPU is certain whether the instruction. For example, when the CPU encounters a branch instruction, it may execute instructions speculatively before the outcome of the branch instruction is decided. Please explain the benefits and disadvantages of speculative execution.

   (d) [3 points] If an executed instruction causes an exception, diverting the control flow to an exception handler, the subsequent instruction must not be executed anymore. Due to out-of-order execution, the subsequent instructions may already have been partially executed. Please explain what cause an exception and how a modern processor handles the exception correctly for out-of-order execution.

   (e) [3 points] Although the instructions executed out of order do not have any visible architectural effect on registers or memory, they have microarchitectural side effects. During the out-of-order execution, the referenced memory is fetched into a register and is also stored in the cache. An attacker can leverage a microarchitectural side-channel attack such as *Flush+Reload*, which detects whether a specific memory location is cached, to make this microarchitectural state visible. Please explain the importance of caches in modern processors and why cached data may make differences.

見背面

Now you know that knowledge about computer architecture is important to understand how Meltdown works. You can try to read the paper after this examination if you are interested.

2. (12 points) Suppose that you are developing a big data application which needs to process 1 Petabytes (10^{15}) of data. Assume the application consists of mostly independent data parallel computing tasks, and your program only has to add the results of the tasks together in the end. To speed up the application as much as possible, you are thinking about renting a high-performance computing cluster with *N* computers.

   (f) [3 points] First, let us estimate the time required to load the data from storage to the memory. Each independent operation accesses 1 Kilobytes of data. Each data block on the storage is 1 Kilobyte. Suppose each computer has *D* disks. Each disk can transfer data at 250MByte/sec with an average seek time of 5 ms. Please estimate the minimum and maximum time required to read all the 1 Petabyte of data from the disks into the computers.

   (g) [3 points] We would like to accelerate the computation with GPU's. Suppose each computer has 8 GPU cards connected via the PCIe links. Each link can transfer data at about 16GB/sec. Assume that the computation time is smaller than the data transfer time, so the data transfer time dominates, and the computation can be overlapped with the data transfer. Please estimate the time required to have the data processed by the GPU's.

   (h) [3 points] It takes time for the program to add the results together by performing a *reduction* operation. Assume that the time needed to transfer the results from one GPU to another GPU or to the CPU within the same computer is negligible. However, it takes *100* microseconds to transfer the results from one computer to another computer. Please calculate the time required for the reduction operation.

   (i) [3 points] Estimate the minimum total execution time. Note that some of the moving data from storage and sending data to the GPU could be done in parallel. Which part is mostly likely to be the performance bottleneck? What hardware changes can you make to improve the performance?

接次頁

3. (5 points) Meltdown is a hardware vulnerability affecting a wide range of processors. It uses a cache timing attack to read kernel space data by observing the results of speculative operations conditioned on data fetched with invalid privileges [Wikipedia]. Which processor below is NOT affected by Meltdown?

   (a) Intel Core i7-7700 released just last year

   (b) ARM Cortex-A75 with an out-of-order superscalar pipeline

   (c) Intel Itanium, a Very-Long-Instruction-Word processor

   (d) Intel Xeon E5-2620 server-class CPU

4. (5 points)

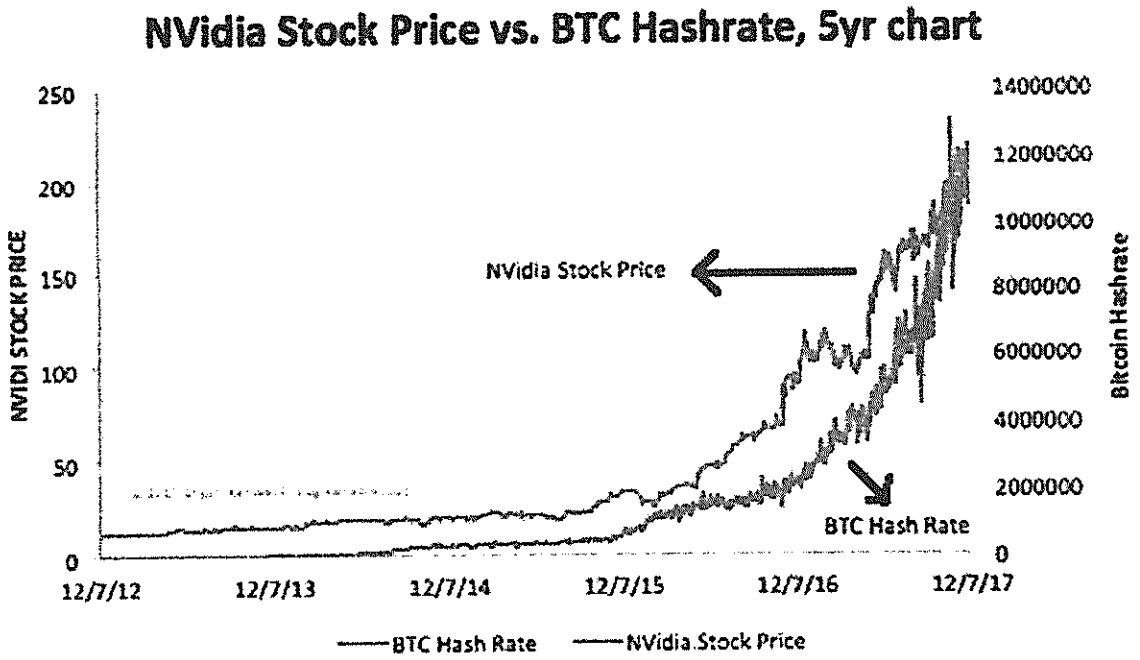

Using his figure, Mr. Kamadolli pointed out the correlation between NVidia's GPU business and Bitcoin (BTC) hashrate. Note that hashrate is a measure of the computational complexity of mining new coins. Which of the following contribute the LEAST to NVidia's booming GPU business?

   (a) Machine learning

   (b) Mining new coins on Ethereum, which is a blockchain that has smaller Market Cap than Bitcoin (BTC)

   (c) Artificial intelligence

   (d) Mining Bitcoin (BTC) in 2017

見背面

5. (4 points) Which statement below is FALSE with respect to Google's TPU (Tensorflow Processing Unit) that is announced in 2016?
   (a) Google's TPU is good at matrix multiplication.
   (b) Google's TPU uses 32-bit floating point processing unit to accelerate deep learning computation.
   (c) Google's TPU does not support native division operations.
   (d) The large memory bandwidth is important for Deep Learning on Google's TPU.

6. (4 points) Which of the following performance improvement techniques does not exploit parallelism?
   (a) Increasing the clock frequency by increasing the pipeline depth
   (b) Reducing the cache access latency by making the cache smaller
   (c) Increasing instruction throughput by adding more functional units
   (d) Improving system performance by adding an additional processor core

7. (2 points) Answer **TRUE** or **FALSE**: Branch prediction is more important for processors with longer pipelines.

8. (2 points) Answer **TRUE** or **FALSE**: Two instructions with data dependencies may not cause data hazard in pipelining execution.

9. (3 points) Answer **TRUE** or **FALSE**: Consider the L1 Cache in Intel i7-6700 (Skylake) with 32 KB, 64 B/line and 8-WAY L1 cache, 4 KB pages mode. If Virtual Address Indexed, Physical Address Tagged cache (VIPT) is used, one RAM data may be repeated placed in different cache slots at the same time.

接次頁

Part II: Operating System

NOTE that in the question, it is intended to provide redundant or miss certain assumption to disguise you. Please make your own assumption if necessary to answer the questions.

10. (25 points) On Jan. 2, 2018, it was announced that *the meltdown vulnerability* of Intel processors may allow normal user programs – from database applications to JavaScript in web browsers – to discern to some extent the layout or contents of protected kernel memory areas. Please answer the following questions.

(a) (5 points) The meltdown vulnerability is caused by shared translation lookaside buffer (TLB) in order to reduce the latency for user/kernel mode switch. Please describe how TLB works while translating virtual address to physical address.

(b) (6 points) The page table shown on the right is for a system with 16-bit virtual and physical addresses and with 4,096-byte pages. Please translate the following 16-bit virtual addresses into their corresponding physical address in hexadecimal.

1. `0xA12C`
2. `0x3A9D`
3. `0xB2D0`

Page #	Page Frame #	Reference Bit
0	4	0
1	5	0
2	-	0
3	2	1
4	-	0
5	-	0
6	9	0
7	12	0
8	-	0
9	-	0
10	15	0
11	-	0
12	1	0
13	3	0
14	-	0
15	-	0

(c) (4 points) On a 64-bit address computer paging physical memory into 4KB page frame, a simple page table can take too much space to map memory pages. How does the operating system do to avoid using huge page table? Please provide **at least two** different approaches. (It can be different on Intel and ARM processors.)

(d) (5 points) Please **describe** the possible method to prevent the security vulnerability on TLB without replacing the processor and **discuss** the performance impacts to the operating systems.

(e) (5 points) The execution of one user process may switch between user mode and kernel model. Please **describe** the procedure of mode switch and **answer** if requesting for the PID of a process, i.e., `getpid()` in POSIX, requires a mode switch.

11. (5 points) In Linux scheduler, describe what performance is improved using red-black tree, and how?

12. (5 points) Describe how the related software components work to know you press some button of mouse and pop out corresponding menu bars?

13. (5 points) What application scenario performs better using FCFS(first come first serve) than SSTF (shortest seek time first) in disk scheduling? What application scenario performs worse?

14. (10 points) Please write codes as compete as possible to implement a counting semaphore using binary semaphores (mutexes).

試題隨卷繳回