

# CS4/591 Midterm

Full Name: \_\_\_\_\_

- This is a closed book, closed note, closed device, closed neighbour exam.
- If you have questions about the exam please quietly approach me at the front so I can assist.
- The exam will end promptly at 11:50am. Please make sure I have your exam before 11:50am.
- There is only one correct answer for each of the multiple-choice questions. If multiple answers are selected the answer will be marked wrong.
- The short answer questions only require a couple of sentences to answer, at most.

## 1 A Brief History of Supercomputing

1. Why was FORTRAN the first successful high level language?
  - DARPA required its use and it was bundled with ASCI Red.
  - Its machine code was as fast as a human programmer's machine code.
  - It supported OpenMP, which made it faster than any other programming language.
  - FORTRAN supported circuit switching.
2. What are Beowulf clusters and why are they important to HPC?

## 2 iDRAC and Linux Installation

3. iDRAC is useful because:
  - It runs in firmware.
  - Allows remote configuration of a server.
  - All HPC servers have iDRAC.
  - All of the above.
4. What distribution and version of Linux are we using in class?

### 3 Vector Operations

5. Vector operations allow CPUs to?
- Compute vector distances efficiently.
  - Switch threads of execution with less overhead.
  - Apply a single instruction to multiple values at once.
  - None of the above.
6. Do vector operations always speed up code execution? Why or why not?

### 4 Devices

7. In userspace device events are handled by:
- Udevd
  - InitDevd
  - Httpd
  - Sshd
8. What does it mean to say “everything is a file in Linux” in the context of devices?

### 5 Filesystems

9. You find an EFI partition on a drive. What does that imply
- It is a GPT drive.
  - It could have an MBR boot block.
  - It is a UEFI compatible partition scheme.
  - All of the above.
10. Give a reason why spinning disks might be more suitable in an HPC environment than SSDs.

### 6 Networking

11. ARP does which of the following:
- Secures routers against hackers.
  - Provides benefits for those over 60 years old.

- Makes packet switching possible.
  - Maps ethernet mac addresses to IP addresses.
12. Which subnet mask allows more host addresses: 255.255.255.0 or 255.255.0.0 and why?
13. Give a difference between infiniband and ethernet and why it matters.

## 7 The Machine Room

14. Which of the following is true about the CARC Machine Room?
- It uses a hot aisle/cold aisle configuration.
  - It is solar powered.
  - It is liquid cooled.
  - It isolates the supercomputers from the outside world.
15. What is the fastest type of network in the CARC machine room?

## 8 Booting and the Kernel

16. Which of the following is true:
- The bootloader is an ELF program.
  - The first thing the Kernel does is load the init program (systemd or initd).
  - The Kernel is read from the MBR boot block.
  - User space programs interact with the Kernel through syscalls.
17. Name three things for which the kernel is responsible.

## 9 System Daemons

18. Which of the following is false:
- Processes are created with fork()-exec().
  - The init process is the ancestor of of all processes.
  - In System V, initd is the init process.
  - Systemd is part of the Kernel.
19. What utility would you use monitor the output logs of systemd units?

## 10 Compiler Optimisation

20. In Homework 3 you compared the speed of programs compiled with Intel and GCC compilers. What did you observe?
21. In Homework 3 you compared different optimisation levels. What did you observe?

## 11 Shared Memory Parallelism

22. Shared memory parallelism technologies such as OpenMP allows for:
- Multiple threads of execution in the same program.
  - Hyperthreading
  - Less context switching by the Kernel.
  - Distribution of computation over the network to multiple nodes.
23. What does this mean: “export OMP\_NUM\_THREADS=4”? How much faster would an OpenMP program run given that command (be careful)?

## 12 Process and Resource Monitoring

24. What is the system load?
- Percentage of CPU in use.
  - Number of processes running.
  - Number of processes in the ready state.
  - Number of processes in the blocked state.
25. Give two commands you could run to monitor the CPU usage of a process.

## 13 SLURM

26. SLURM is
- a way to implement SMP.
  - a job scheduler.
  - a delicious soda.
  - the best way to ensure code is properly parallelised.
27. Explain the difference between “ntasks” and “cpus-per-task” in SLURM.