

Fall 2016

Anthony D. Joseph

**Midterm Exam #3 Solutions**

November 30, 2016  
CS162 Operating Systems

<b>Your Name:</b>	
<b>SID AND 162 Login:</b>	
<b>TA Name:</b>	
<b>Discussion Section Time:</b>	

General Information:

This is a **closed book and three 2-sided handwritten notes** examination. You have 80 minutes to answer as many questions as possible. The number in parentheses at the beginning of each question indicates the number of points for that question. You should read **all** of the questions before starting the exam, as some of the questions are substantially more time consuming.

Write all of your answers directly on this paper. *Make your answers as concise as possible.* If there is something in a question that you believe is open to interpretation, then please ask us about it!

**Good Luck!!**

<b>QUESTION</b>	<b>POINTS ASSIGNED</b>	<b>POINTS OBTAINED</b>
<b>1</b>	<b>34</b>	
<b>2</b>	<b>35</b>	
<b>3</b>	<b>16</b>	
<b>4</b>	<b>15</b>	
<b>TOTAL</b>	<b>100</b>	

**Solutions NAME:** \_\_\_\_\_

1. (34 points total) Short answer.

a. (12 points) True/False and Why? **CIRCLE YOUR ANSWER.**

i) The Pintos and BSD 4.2 Fast File Systems use a linked-list on-disk data structure to track the allocation of inodes and data blocks.

**TRUE**

**FALSE**

**Why?**

***FALSE.** FFS uses bitmaps (introduced in BSD 4.1) to track the allocation of inodes and data blocks. The correct answer was worth 1 point and the justification was worth an additional 2 points.*

ii) Two processes communicating via a mailbox could be considered an example of a producer-consumer model of process cooperation.

**TRUE**

**FALSE**

**Why?**

***TRUE.** The mailbox provides a queue, which is equivalent to the buffer in producer-consumer. The correct answer was worth 1 point and the justification was worth an additional 2 points.*

iii) A remote procedure call requires network access.

**TRUE**

**FALSE**

**Why?**

***FALSE.** A remote procedure call can occur between two processes on the **same** system. The correct answer was worth 1 point and the justification was worth an additional 2 points.*

**Solutions NAME:** \_\_\_\_\_

- iv) If machines were guaranteed not to crash, we would not need two-phase commit: the coordinator could, in one phase, decide whether a distributed transaction would commit, and then instruct the workers to apply their piece of the transaction.

**TRUE****FALSE****Why?**

**FALSE.** The point to the first phase of 2PC is to see whether all of the machines agree to the proposed transaction. The correct answer was worth 1 points and the justification was worth an additional 2 points.

- b. (4 points) Briefly, in two to three sentences, explain what a DMA controller does and what benefit it provides.  
*A **Direct Memory Access** controller is used to move data between memory and an I/O device. By using a DMA controller, the CPU can offload the reading and writing of data to/from I/O devices..*
- c. (9 points) Abstraction layers.
- i) (2 points) Let us explore how abstraction layers help in the design of an operating system. Give a concrete example of a use of abstraction layers in an operating system.  
*Many examples, including device drivers, filesystems, networking stacks*
- ii) (4 points) Briefly, in one to two sentences, explain the advantages of using the abstraction layers in your example.
- *You can separate different parts of the system software, if each layer uses only the next lower layer.*
  - *It may be easier to prove system correctness, perform debugging, etc.*

**Solutions NAME:** \_\_\_\_\_

- *You can gain some degree of system independence, if only the lowest layers access hardware directly.*
- *In general, better management of large programs, but note that layering is not the same as modularization or component-based designs.*
- *More generality of the functions.*
- *Easier to develop the design concepts.*

iii) (3 points) Briefly, in one to two sentences, explain a disadvantage of using the abstraction layers in your example.  
*Layering can reduce performance*

d. (9 points) Networking: Two computers, A and B, are connected by a network link that runs at 1 Gigabit per second ( $1 \times 10^9$  bits/second). The one-way propagation delay from A to B (and vice versa) is 20 milliseconds. Assume the following conditions:

- The link does not drop or duplicate packets
- Processing time at the two endpoints is zero
- A sends B 625-byte packets – the 625 bytes includes all headers, framing, and inter-packet spacing

i) (3 points) What is the maximum number of 625-byte packets per second that A could in principle send into the wire? *You can leave your answer in unsimplified form.*

$$10^9 \text{ bits/second} \times 1 \text{ byte}/8 \text{ bits} * 1 \text{ packet}/625 \text{ bytes} = 200,000 \text{ packets/second.}$$

**Solutions NAME:** \_\_\_\_\_

- ii) (6 points) Now, consider the above link and the following protocol, and assume that all packets are again 625 bytes. In the protocol, A sends 4,000 packets into the network as quickly as it can and then waits for a one-byte ACK from B. Whenever A receives an ACK from B, A immediately sends another 4000 packets into the network. Meanwhile, B ACKs packet 1, packet 4001, packet 8001, etc. That is, B ACKs every 4,000 packets starting with the *first* one in a burst by A.

What is the long-term throughput of this protocol, expressed as *both* (1) bits per second and (2) a percentage of the link's bandwidth? *You can leave your answer in unsimplified form. Briefly explain your answers.*

*The bandwidth-delay product is  $10^9$  bits/second  $\times$  20 milliseconds = 2.5 Megabytes = 4000 packets of size 625 bytes. If A sent 4000 packets every 20 milliseconds, it would fully use the link. However, from the protocol description, we know that the protocol sends 4000 packets every 40 milliseconds (because the round-trip time is 40 milliseconds, and that's how long it takes for A to get B's ACK). Thus, the throughput is one-half the link's bandwidth, or  $500 \times 10^8$  bits/second.*

**Solutions NAME:** \_\_\_\_\_

## 2. (35 points total) Filesystems.

- a. (4 points) One of the innovations of the BSD FFS is that it tries to allocate large files in long contiguous chunks. Assuming that a disk transfers at a peak rate of 100 MByte/s, and that a combined seek and rotation take, on average, a total of 20 milliseconds. What is the **minimum** size of *each* contiguous run of a large file, in order to achieve 75% of peak transfer rate for large files when they are accessed sequentially?

*The seek/rotation takes 20 milliseconds, so to achieve a 75% peak transfer rate, we need to spend 4/3 of the seek/rotation time transferring data, which is 60 ms:  $60 \times 10^{-3} \text{ seconds} \times 100 \times 10^6 \text{ bytes/seconds} = 6,000 \times 10^3 \text{ bytes} = 6 \text{ Mbytes}$*

- b. (15 points) Large files. We have seen different filesystems that support fairly large files. Now let's see just how large a file various types of filesystems can support. Assume, for all of the questions in this part, that filesystem blocks are 4 KBytes. *Show your solutions in unsimplified form for partial credit.*

- i) (3 points) Consider a really simple filesystem, **directfs**, where each inode only has 10 direct pointers, each of which can point to a single file block. Direct pointers are 32 bits in size (4 bytes). What is the maximum file size for **directfs**?

*The maximum **directfs** file size is  $10 \times 4 \text{ KByte} = 40 \text{ KByte}$*

- ii) (3 points) Consider a filesystem, called **extentsfs**, with a construct called an extent. Extents have a pointer (base address) and a length (in blocks). Assume the length field is 8 bits (1 byte). Assuming that an inode has exactly one extent. What is the maximum file size for **extentsfs**?

*The maximum **extentsfs** file size is  $(2^8 - 1) \times 4 \text{ KByte} = 255 \times 4 \text{ KByte} = 1 \text{ MByte}$*

**Solutions NAME:** \_\_\_\_\_

- iii) (3 points) Consider a filesystem that uses direct pointers, but also adds indirect pointers and double-indirect pointers. We call this filesystem, **indirectfs**. Specifically, an inode within **indirectfs** has 1 direct pointer, 1 indirect pointer, and 1 doubly-indirect pointer field. Pointers, as before, are 4 bytes (32 bits) in size. What is the maximum file size for **indirectfs**?

*The maximum **indirectfs** file size is  $(1 + 1024 + (1024 \times 1024)) \times 4 \text{ KByte} = 4 \text{ GB} + 4100 \text{ KByte}$*

- iv) (3 points) Consider a compact file system, called **compactfs**, tries to save as much space as possible within the inode. Thus, to point to files, it stores only a single 32-bit pointer to the first block of the file. However, blocks within **compactfs** store 4,092 bytes of user data and a 32-bit next field (much like a linked list), and thus can point to a subsequent block (or to NULL, indicating there is no more data). First, draw a picture of an inode and a file that is 10 KBytes in size.



- v) (3 points) What is the maximum file size for **compactfs** (assuming no other restrictions on file sizes)?

*The maximum **compactfs** file size is  $2^{32} \times 4 \text{ KByte} - 2^{32} \times 4 \text{ byte} = 16 \text{ TByte}$*

**Solutions NAME:** \_\_\_\_\_

- c. (5 points) File system crash recovery.

The Linux journaling file system writes the content of all modified disk blocks to the log. Your friend Bob considers such logging to be wasteful since copying the content of modified disk blocks to the log doubles the amount of disk writes for each logged file system operation.

Bob decides to implement a more efficient journaling file system. In particular, he decides to only record an operation's name and parameter in the log file instead of recording the content of all modified blocks. For example, for an operation that creates file `"/d/f"`, the file system would append the transaction record of the form `[create "/d/f"]` to the log. Bob's file system ensures that the corresponding transaction record is written to the log before the modified disk blocks are flushed to disk. Upon crash and recovery, Bob's file system re-executes the logged file system operations and truncates the log.

Bob's new logging mechanism is certainly more efficient since each transaction record is much smaller than that with Linux's logging. Is his design also correct? Specifically, can it recover a file system correctly from crashes? Explain your reasoning and give concrete examples.

*Ben's design is not correct. Performing a filesystem action requires multiple steps and a crash in the middle of the process will leave the filesystem in an inconsistent state. For example, an unlink operation has to update the inode freemap and block freemap.*

- d. (3 points) Your friend suggests doubling the Pintos block size from 512 bytes to 1,024 bytes, since that means you will be able to reach twice as much data from the direct pointers (as a result, medium sized files could fit entirely within the direct region). Why might it be a bad idea to increase the block size?  
*Increasing the minimum disk transfer unit could lead to internal fragmentation.*



**Solutions NAME:** \_\_\_\_\_

- e. (4 points) Suppose you wanted to implement hard links in Pintos. One thing you have to add is a count to ensure you only delete sectors that have no links remaining. Which struct would be the best choice to add this count? Explain your choice.

*The count needs to persist on disk, so you could modify struct `inode_disk`.*

- f. (4 points) The NFS authors had a goal of *transparency*. They wanted applications to be unable to distinguish whether a file system was (a) a remote file system served from an NFS server; or (b) a typical, local Unix file system. They did not succeed (in fact, their goal was impossible). Briefly state one way in which application code can experience different behavior when interacting with a remote NFS file system versus a local Unix file system. Your answer should be in terms of what application code sees, rather than in terms of what a global observer sees. *Latency can be lower (if the server has as large buffer cache for reads or NVRAM for writes) or worse (if the server is on the other side of the world); operations can hang; reads and writes can suddenly fail because another client on another machine deleted the file; close() can return an error (if the server runs out of space when the client flushes its write cache).*

*Note that we did not accept examples where you only said that a file could change due to a remote write, as a local file can also change due to other processes writing the file. We accepted examples where you explicitly stated that there could be inconsistent views of an NFS file due to polling or caching.*

**Solutions NAME:** \_\_\_\_\_

3. (16 points total) Cybersecurity.

a. (8 pts) For each of the following terms, provide a one sentence definition:

i) Authentication

*Proving the identity of an entity to another entity.*

ii) Data integrity

*Ensuring that data is not changed from source to destination or after being written on a storage device.*

iii) Symmetric encryption

*An encipher/decipher algorithm that relies on a single, shared secret key.*

iv) Asymmetric encryption

*An encipher/decipher algorithm that relies on a private key and a public key.*

b. (2 points) Briefly, in one to two sentences, explain the role of a nonce.

*A nonce is used in cryptographic protocols to prevent replay attacks.*

c. (2 points) Briefly, in one to two sentences, explain the role of a HMAC.

*A Hash-based Message Authentication Code is a cryptographically secure hash function that is used to create secure message digests. It takes the input message and a key.*

d. (4 points) Consider two approaches to delegating your access privileges to a server:

- Give your login name and password, so that the server can authenticate to a third party as you, or

**Solutions NAME:** \_\_\_\_\_

- Give the server a capability (e.g., your Kerberos ticket) to act as you. From a security standpoint, which approach is superior? Briefly explain why. *The second approach is superior because while both approaches allow the server to act as you, but the second approach allows you to limit the time period and rights of what can be done as you. With your password, the server can become you whenever it wants. With a capability, the server can do only what you specified and only for as long as the capability is valid.*

**Solutions NAME:** \_\_\_\_\_

4. (15 points total) Queuing Theory.

Consider a single server/single queue (M/M/1) system:

- 5 jobs per second arrive to the server
- The server is capable of executing 8 jobs per second.

*Write your answers in unsimplified form.*

a. (3 pts) Calculate server utilization ( $u$ ).

$$\text{Since } \lambda = 5 \text{ and } \mu = 8, u = \lambda/\mu = 0.625 = 62.5\%$$

b. (3 pts) Calculate how long it takes the server to execute one job ( $T_{ser}$ ).

$$T_{ser} = 1/8 = 0.125 \text{ seconds}$$

c. (3 pts) Calculate the amount of time spent in the queue ( $T_q$ ).

$$T_q = u/(\mu - \lambda) = 0.208333 \text{ seconds}$$

d. (3 pts) Calculate the number of jobs in the queue ( $L_q$ ).

$$L_q = p^2/(1-p) = 1.041667$$

e. (3 pts) Calculate the amount of time a job spends in the system ( $T_{sys}$ ).

$$T_{sys} = T_q + T_{ser} = 0.3333 \text{ seconds}$$