

CS101
Fall 2018

Instructor: Ashley Taylor
December 10, 2018

Final Exam

NAME (LAST, FIRST): _____

SUNET: _____@stanford.edu

Problem	1	2	3	4	5	TOTAL
Score						
	22	13	22	22	23	102

Instructions:

- The time for this exam is **3 hours**, 180 minutes total. There are 102 points total, so you should spend a less more than 2 minutes per point, though your pacing may be different on different problems.
- You are only allowed a pencil, pen, and eraser. Any other materials are strictly prohibited. You may not use any digital devices other than a clock for timing. You may ask the proctors for blank scratch paper.
- A **reference sheet** is included as the last page. You may detach it for your reference during the exam.
- For coding problems, we will forgive minor syntax errors (such as missing semicolons, etc.).
- For any problems involving calculations, we will accept an expression (that could be plugged into a calculator) in lieu of the numerical answer.

Please sign *before* beginning the exam:

I agree to abide by the spirit and letter of the Honor Code, and to follow the instructions above.

(Signature)

(Date)

1. Course Themes

- a) (6 points) Programmers must design systems to be *fault-tolerant*, meaning that systems must be able to recover from errors. Describe an error that computers are designed to handle and an error that the TCP/IP protocol is designed to handle. For each error, describe how the corresponding system handles that error.

Computer error: running out of RAM, which is handled by the operating system managing RAM and shutting down programs

TCP/IP error: bits might be missent. This error is recognized by a checksum, and the packets are resent if necessary.

- b) (6 points) As mentioned in class, a *graph* is a collection of *nodes*, or vertices, and *edges*, or arcs. Describe how the internet and Facebook can both be thought of as a graph. For each of the internet and Facebook, what are the nodes, and what are the edges?

internet: nodes are routers, edges are connections between routers

Facebook : nodes are people, edges are friendships

- c) (10 points) Recall that computer scientists use *layers of abstraction* in designing systems. Discuss a specific system that is modelled on layers of abstraction, and describe each layer in that system.

Two options: computers or TCP/IP protocol

See slides for full descriptions

2. Web and the Internet

- a) (6 points) Define CSS, HTML, and JavaScript. For each one, describe one example of that component in action on a website of your choice. For simplicity, you may refer to the CS101 website.

HTML: markup language that websites are written in. CSS: style rules for the website. JavaScript: programming language that allows developers to make animations and other more complicated features (like buttons) on a website. The CS101 webpage is written in HTML. The CSS is the style definition (colors, fonts, etc). An example of JavaScript in use is the ability to run code in the lecture slides.

- b) (2 points) Describe how domain names and IP addresses are hierarchical.

An organization may have one domain name, and all sub organizations will have domain names that branch off the original domain name. IP addresses work the same way. For example, Stanford has domain name stanford.edu, and CS101, a Stanford class, has domain name cs101.stanford.edu.

- c) (5 points) Discuss how the Google search backend works. In your description, you must include:
- i) What information is stored
 - ii) How the information is acquired
 - iii) How Google determines the order of results to display to the user

Google's backend stores words for each website and website metadata. The information is acquired via spidering -- or following links within a webpage to visit a new webpage. Google uses PageRank to determine which websites are referenced the most, and displays those higher in search results.

3. Security

- a) (4 points) Anand has started a food suggestion app, Raivac, an exclusive app for students at universities with red theme colors. A huge outbreak of FOMO causes you and your friends to sign up for Raivac. You create an account, make a profile, and Raivac soon gets hacked. The attacker now has your username, password, and profile information! Unfortunately, Anand didn't take CS101 and never thought about how to keep user information secure. List two things Anand could have done to keep user information secure.

Store hashed passwords. Don't accept HTTP requests from a different website (preventing CSRF).

- b) (3 points) Give a real-world example of a company doing something that violates their users' privacy. Additionally, discuss one way in which a company might protect a user's privacy.

A social network company could sell user profile information to companies interested in demographic data. To protect a user's privacy, a company could end-to-end encrypt all messages sent through their platform.

- c) (3 points) Describe a type of attack a hacker might perform to steal your data. In your answer, you should include:
- i) The name of the attack
 - ii) A short description of the attack
 - iii) Why this type of attack could be bad for a user

A hacker could perform a dictionary attack, where a hacker guesses all passwords in a dictionary to log into someone's account. This attack is bad for a user if the hacker guesses the password correctly, since the user's information could be compromised.

- d) (4 points) Discuss one action you can take to protect your own security with regard to passwords. Discuss one action you can take to protect your own security that does not involve a password.

Use a password manager! Also, log out of your accounts when finished.

- e) (3 points) What does it mean that a messaging service is “end-to-end encrypted”? What is an example of a messaging service with “end-to-end encryption”? What is an example of a messaging service that does not “end-to-end encrypt” your messages?

End-to-end encryption implies that messages are encrypted in transmission from sender to receiver. This means that a messaging service can't decipher the content of the messages. Messenger does not use end-to-end encryption. WhatsApp uses end-to-end encryption.

- f) (5 points) Discuss one concern about bitcoin (or cryptocurrency in general) that has been addressed by blockchain. How does blockchain address that concern? Additionally, discuss a concern about bitcoin that has not been addressed.

Money is decentralized, so we don't have to trust a single authority to handle money transfers. Blockchain makes sure 50% of the network verifies a transaction, instead of just one central figure. A concern about bitcoin that hasn't been addressed is its volatility -- there are no market makers like there are in NASDAQ or NYSE.



4. Potpourri

- a) (4 points) Suppose P vs. NP has been resolved, and it has been shown that $P = NP$. Describe two implications of $P = NP$ on society.

Computers will get much faster, since previously time-consuming problems are now easier to solve. Also, any company that has a scheduling algorithm will be able to optimally schedule events (such as airlines scheduling flights).

- b) (4 points) Marty and Maria are creating their own company and company website. They've heard that several companies use Amazon Web Services (AWS) as an alternative to maintaining their own servers. Discuss one reason why Marty and Maria should use AWS over a local server. Additionally, discuss one reason their company should use its own servers over AWS.

AWS is easier to start up than buying an entire machine and installing all tools. AWS, however, is run by another company, so if Marty and Maria really care about security, they should take matters into their own hands and use their own servers.

- c) (2 points) In class, we discussed MapReduce as an example of a distributed system, which acts as a distributed CPU. List another type of distributed system. Additionally, describe an advantage that a distributed system of that type has over a non-distributed system of the same type.

AWS is a distributed system. Disk memory is distributed across multiple machines. An advantage of AWS is that exactly enough memory can be allocated for an instance, whereas in a local computer, a lot of memory might go to waste.

- d) (2 points) Recall that some programming languages, such as Java, are compiled. Other programming languages, such as Javascript, are interpreted. What is a difference between compiled and interpreted programming languages?

The compiler produces assembly that the computer runs, whereas in an interpreted language, each line of code is run step by step.

- e) (2 points) Imagine you are working on Anand's food suggestion app as a designer. Describe a dark pattern you could incorporate into the application.

We could make it really difficult to deactivate the account. Or, we could require the user to connect their social media accounts to get relevant information.

- f) (2 points) What two factors do computer scientists look at to determine if an algorithm is good (with respect to input size)?

Space and time

- g) (2 points) Recall that computer vision describes the field concerning how computers derive semantic information from photos and videos. Describe one way computer vision is applied to society.

Face classifiers (FaceID on iPhone)

- h) (4 points) Describe a scenario in which a computer scientist might use lossless compression instead of lossy compression. Additionally, describe another scenario in which they might use lossy compression instead of lossless compression.

Using lossless compression is ideal for text documents, because we want to reconstruct text perfectly. Using lossy compression is ideal for large images / videos, because we don't care about reconstructing each pixel perfectly.

5. Code Reading/Comprehension

- a) (8 points) Below is the code from green screen as we talked about in class. Edit it so that the printed image completely merges the two images together **regardless of the color of the pixels in either image**. You can cross out or add code as necessary, but please use arrows to clearly indicate your edits. Hint: you don't need to edit any of the first four lines of code.

```
img = new SimpleImage("yosemite.jpg");
back = new SimpleImage("paris.jpg");
back.setSameSize(image);
img.setZoom(0.5);

for (pixel : img) {

    backPixel = back.getPixel(pixel.getX(), pixel.getY());

    if (pixel.isSimilarTo(63, 121, 248, 70)) {

        pixel.setRGB((backPixel.getRed() + pixel.getRed() / 2),
                    (backPixel.getGreen() + pixel.getGreen() / 2),
                    (backPixel.getBlue() + pixel.getBlue() / 2));

    }

}

print(back);
```

- b) (10 points) Write code to change the bottom right quadrant of an image to have a negative filter. A negative filter replaces each color with its inverse (by subtracting the current red, green, or blue value from 255). For example, purple would become green, and white would become black. Your solution should work on any image of any size.

```
img = new SimpleImage("poppy.jpg");
img.setZoom(0.5);
// write your answer below

for (pixel: img) {

    if (pixel.getX() > img.getWidth() / 2 && pixel.getY() > img.getHeight()
        / 2) {

        pixel.setRGB (255 - pixel.getRed(), 255 - pixel.getGreen(), 255 -
            pixel.getBlue());

    }

}

print(img);
```

- c) (5 points) Below is a description of columns in a spreadsheet of Good Place episodes. Describe how you would find the **total** number of views of all Good Place episodes. Then describe how you would find which character appears as the main character in the show for the most number of episodes. Finally, describe how you would determine which main character is the most popular, as defined by average number of views.

Column	Definition
EPISODE	The episode title, e.g., "Janet(s)"
VIEWS	The number of people who watched that episode
MAIN_CHARACTER	The main character of that episode, e.g., "Chidi"

We do not expect you to provide actual code or function calls, but you may use any of the following functions in your answer. You should describe which function(s) you'd use and which columns/conditions you'd use.

COUNTUNIQUE(range)
 UNIQUE(range)
 SUM(range)

COUNTIFS(condition_range, condition)
 SUMIFS(range, condition_range, condition)
 MINIFS(range, condition_range, condition)
 MAXIFS(range, condition_range, condition)
 AVERAGEIFS(range, condition_range, condition)

To find total number of views: take SUM of VIEWS column

To find which character appears as main character for most number of episodes: make a new column that stores the number of occurrences of each character. Then get the max number in that new column, which will correspond to the answer.

To find which character is the most popular based on the average number of views: make a new column that computes average number of views for all the episodes for a given character. Then, compute the max of this new column, which will correspond to the answer.

```
pixel.getX()
pixel.getY()
pixel.getRed()
pixel.getGreen()
pixel.getBlue()
pixel.setRGB(red, green, blue)
pixel.isSimilarTo(red, green, blue, threshold)
```

```
image = new SimpleImage("image.png")
image.countNeighbors(pixel)
image.getHeight()
image.getWidth()
image.getPixel(row, column)
```

```
for (pixel : image) {
    // your code here
}
```

```
for (neighbor : image.getNeighbors(pixel)) {
    // your code here
}
```

```
if (condition) {
    // your code here
} else {
    // your code here
}
```

```
&& => and
|| => or
! => not
!= => is not equal
== => is equal
```

```
1TB = 1000GB
```


1GB = 1000MB

1MB = 1000KB

1KB = 1000 bytes

Bit: 0 is "off", 1 is "on"