

Social Networks, File Systems

These exercises are intended to help you master and remember the material discussed in lectures and explored in labs. In future semesters, we may make some or all of these exercises required, but for now they remain optional. We suggest that you do them as we go over the material, but you may also want to use them to review concepts before the exam.

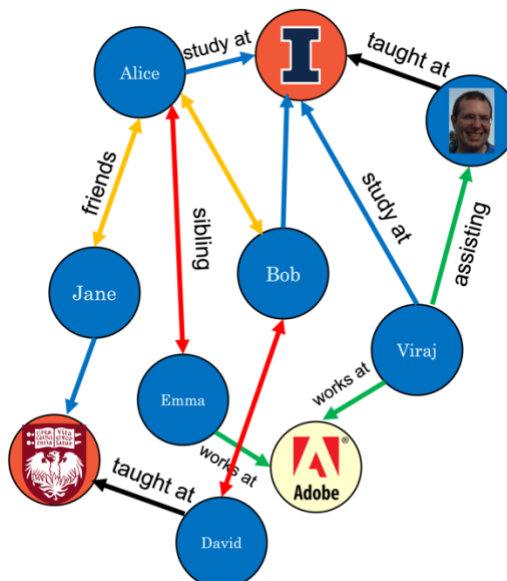
Please note also that some of the exercises are meant to be done with a calculator, while in exams, we just want you to be able to set up the equations correctly. Also, some of the exercises (such as #6) are meant to help you to explore a topic and learn more about it using sources from the Internet and thus won't appear on exams.

Rather than using this version directly, we suggest that you use the version without solutions to solve the problems before looking at the answers. Many studies have shown that people often trick themselves into believing that they know how to solve a problem if they are presented with the answer before they try to solve the problem themselves.

- [L9] Alice and Jane are childhood friends. They both lived in the same city: Chicago. Alice often used to visit her cousin Emma who works as a software engineer at Adobe in the (San Francisco) Bay Area. Inspired by her, Alice decided to enroll in the Software Engineering program at UIUC, while her friend Jane chose to attend law school at the University of Chicago. In freshman year, Alice attended a new course offered by Prof. Lumetta on Digital Information Technologies. Viraj, a PhD student at UIUC and an intern at Adobe, served as assistant to Prof. Lumetta in the course. Here, Alice became friends with one of her classmates in Prof. Lumetta's course—Bob, also a freshman at UIUC. In one of their conversations, Bob mentioned his elder brother David, who is a professor of Law at the University of Chicago. Alice realizes that we indeed live in such an interconnected world!

Draw a social graph based on the story in the previous paragraph. Remember that in a social graph, nodes represent people, places, or organizations such as universities/companies, while their relationships are described using arcs. For computers, such a representation is simpler and easier to understand than the textual description.

Arcs of the graph may look slightly different depending on how you describe the relationships.

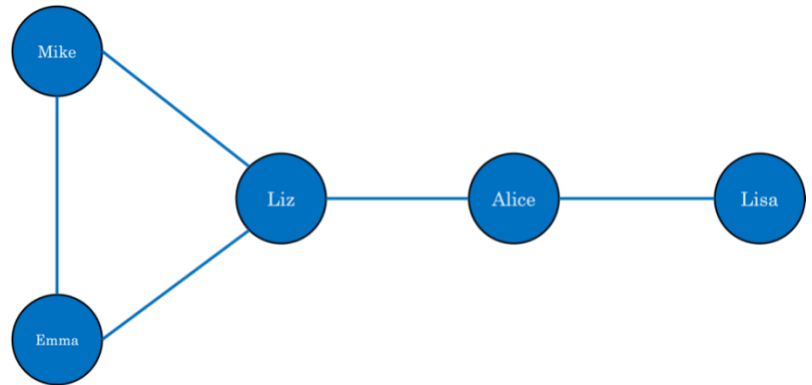


2. [L9] Alice recently joined a new social networking website called IlliniBook. She invited few of her classmates join and marked them as friends on the IlliniBook app.

- A) The part of the IlliniBook social graph related to Alice appears to the right. What is the diameter of this graph?

Hint: a graph's diameter is the longest distance between any pair of nodes.


In the graph shown, the longest distance is 3 hops (between Emma and Lisa, or between Mike and Lisa), so the diameter is 3.



- B) Alice and her friends realized that they didn't have many friends on their new social networking app, so they asked more and more of their classmates to join the new app. After a while, their social network grew a lot, almost double than what it was previously. The new graph is shown to the right.

What is the diameter of the new graph?

Despite having twice as many nodes as the graph in Part A, the longest distance is still 3 hops, so the diameter is still 3.



- C) The number of nodes and arcs doubled when the network grew. How did this growth impact the diameter of the graph? What property of social graphs (*hint: small world graphs*) can explain this observation?

Although the number of nodes and arcs in the graph is larger, the diameter remained constant. In small world graphs such as social graphs, people tend to make more than one connection with existing nodes, so the diameter of the graph grows slowly with growth in the number of nodes and arcs. Typically, the diameter grows logarithmically in the number of people (nodes), so one needs to add a large number of nodes to increase the diameter by a small amount.

3. [L9] In two or three sentences, explain how an attacker might use a social network to abuse a victim's existing trust relationships and/or sidestep a victim's natural deceit detection mechanisms.

(many examples are possible)

By creating several accounts and pretending that they are independent people, an attacker can convince a common "friend" (victim) that the victim is receiving independent opinions from several contacts, providing a much stronger influence to act as the attacker desires.

4. [L9] Bob is supposed to move all of his stuff from his residence hall to a new apartment. He has 100 lbs. of stuff in total, but can comfortably carry only 22 lbs. Bob goes to a store to purchase moving containers. He has two options: purchase a large moving container of 100 lbs. capacity, or purchase multiple small containers, each of 20 lbs. capacity. What option should Bob choose? Facebook made an analogous decision in the design of their TAO architecture. Explain in one or two sentences.

Bob should purchase 5 small containers so that he can distribute and pack his stuff in a way that each small container contains around 20 lbs.—just enough for him to be able to carry comfortably.

Facebook's TAO architecture uses similar analogy as it breaks the social graph into many `shards` (like containers) to make it easier for each computer to carry the load. Shards have variable size, so using only one shard per computer leads to unnecessarily heavy load on some computers, which can (and is) fixed by using many smaller shards per computer and balancing the load.

5. [L9] Give an example of how human deceit detection mechanisms can backfire on a social networks.

(Many examples are possible.)

Humans tend to trust people who demonstrate that they already have the trust of many other humans, assuming that the other humans' deceit detection mechanisms effectively screen against villainous individuals. In a social network, however, methods to prevent automatic account creation are insufficient to stop determined (or well-funded) attackers. Such an attacker can thus create hundreds or thousands (or more) of fake accounts, populate those accounts automatically with images of different "people," and use the accounts to create an aura of trust, high ratings, good recommendations and reviews, or any other form of social support to deceive victims by abusing their trust in large groups of other humans. (Some of Prof. Lumetta's friends, for example, have awarded themselves chili peppers on RateMyProfessor!)

6. [L10] With new technologies and tools being applied to the content of cloud storage systems every day, it becomes difficult to audit them for compliance with the law. Recently, a popular cloud storage system for photos was charged with violating Illinois state privacy laws. Search the Internet to find the story, then summarize it in one or two sentences.

In response to a class-action lawsuit filed over one of its facial recognition features in Google Photos, Google agreed to pay a \$100 million settlement to Illinois residents. The complaint alleges that Google's face grouping tool, which automatically identifies your face in photos and videos uploaded to Photos, violates Illinois' Biometric Information Privacy Act (BIPA).

From the Verge article: <https://www.theverge.com/2022/6/6/23156198/google-class-action-face-grouping-biometric-information-illinois-privacy-act>

7. [L10] Explain each of the following definitions in a sentence or two:

A. consistent

a property provided partially by Internet storage services, ensuring that all users see the same version of data at the same time

B. datacenter

a co-located group of computers and storage used to provide Internet service to many users simultaneously

C. pull model

an approach in which updates to information are only delivered to a client on request, rather than being proactively sent as soon as changes are made