





Networking & Security



Latency and Bandwidth

- Napoleon's Network: Paris to Toulon, 475 mi
- Latency: 13 minutes (1.6s per mile)
 - What is the delay at each signaling station, how many stations to reach destination
 - At this rate, it would take ~1 hour to get a bit from California
- Bandwidth: 2 symbols per minute (98 possible symbols, so that is ~13 bits per minute)
 - How fast can signalers make symbols
 - At this rate, it would take you about 9 days to get *ps8.zip*

#5

One-Slide Summary

- **Bandwidth** is the throughput of a communication resource, measured in **bits per second**. **Latency** is the time delay between the moment when communication is initiated and the moment the first bit arrives, measured in **seconds**.
- In **circuit switching**, a path through a network is reserved (high quality-of-service, used in telephones). In **packet switching**, each packet is routed individually (internet, postal service).
- The **world wide web** involves simple schemes for retrieving resources (**URL**, **HTTP**) and a simple language for displaying information (**HTML**). HTTP is **stateless**, so long-running sessions store info on the client (**cookies**) or server (database).
- A **dynamic website** generates content by running a program on the client (e.g., Google maps interface) or the server (e.g., rest of PS8).

#2

Measuring Networks

- **Latency**
Time from sending a bit until it arrives
seconds (or seconds per geographic distance)
- **Bandwidth**
How much information can you transmit per time unit
bits per second

#4

Improving Latency

- Fewer transfer points
 - Longer distances between transfer points
 - Semaphores: how far can you see clearly
 - Curvature of Earth is hard to overcome
 - Use wires (electrical telegraphs, 1837)
- Faster transfers
 - Replace humans with machines
- Faster travel between transfers
 - Hard to beat speed of light (semaphore network)
 - Electrons in copper: about 1/3rd speed of light

#6

How many transfer points between here and California?

#7

#8

tracert

```

K:\>tracert www.cs.berkeley.edu
Tracing route to hyperion.cs.berkeley.edu [169.229.60.105]
over a maximum of 30 hops:
  1  3 ms  3 ms  4 ms  128.143.69.1
  2  <1 ms <1 ms <1 ms carruthers-6509a-x.misc.Virginia.EDU [...]
  3  <1 ms <1 ms <1 ms new-internet-x.misc.Virginia.EDU [128.32.100.1]
  4  4 ms  4 ms  4 ms  nwr-nlr13.misc.Virginia.EDU [192.35.48.30]
  5  5 ms  5 ms  5 ms  nlr13-router.networkvirginia.net [192.7.1.1]
  6  18 ms 18 ms 18 ms  atla-wash-64.layer3.nlr.net [216.24.186.20]
  7  43 ms 43 ms 42 ms  hous-atla-70.layer3.nlr.net [216.24.186.8]
  8  73 ms 73 ms 73 ms  losa-hous-87.layer3.nlr.net [216.24.186.1]
  9  72 ms 72 ms 72 ms  hpr-lax-hpr-100.layer3.nlr.net [216.24.186.1]
 10 80 ms 81 ms 81 ms  svl-hpr-lax-hpr-100g.cenic.net [137.164.1.1]
 11 145 ms 81 ms 81 ms  hpr-ucb-ge-svl-hpr.cenic.net [137.164.1.1]
 12 81 ms 81 ms 81 ms  g3-12.inr-201-eva.Berkeley.EDU [128.32.100.1]
 13 81 ms 82 ms 83 ms  evans-soda-br-5-4.EECS.Berkeley.EDU [...]
 14 83 ms 84 ms 83 ms  sbd2a.EECS.Berkeley.EDU [169.229.59.226]
 15 83 ms 84 ms 83 ms  hyperion.CS.Berkeley.EDU [169.229.60.105]

Trace complete.
    
```

UVA

UCB

Atlanta → Houston → LA?

```

>>> cvilleberkeley = 3813 # kilometers
>>> seconds = 84.0/1000
>>> speed = cvilleberkeley / seconds
>>> speed
45392.857142857138
>>> light = 299792.458 # km/s
>>> speed / light
0.15141427321316114


Packets are traveling average at 15% of the speed of light
(includes transfer time through 15 routers)
    
```

#9

#10

Bandwidth

How much data can you transfer in a given amount of time?



#11

Improving Bandwidth

- Faster transmission
 - Train signalers to move semaphore flags faster
 - Use something less physically demanding to transmit
- Bigger pipes
 - Have multiple signalers transmit every other letter at the same time
- Better encoding
 - Figure out how to code more than 98 symbols with semaphore signal
 - Morse code (1840s)

#12

First Use of Internet

- October 1969: First packets on the ARPANet from UCLA to Stanford. Starts to send "LOGIN", but it crashes on the G.
- 20 July 1969:

Live video (b/w) and audio transmitted from moon to Earth, and to millions of televisions worldwide.



#19

Liberal Arts Trivia: Psychology

- This series of social psychology experiments at Yale University measured the willingness of study participants to obey an authority figure who instructed them to perform acts that conflicted with their personal conscience. The scientist devised the experiments to address the question: "Could it be that Eichmann and his million accomplices in the Holocaust were just following orders? Could we call them all accomplices?" Participants played the role of a "teacher" helping a "learner" with a memory study and were instructed to deliver electric shocks until the the "learner" "died".

#20

Liberal Arts Trivia: Medieval Studies

- This English legal charter, originally issued in Latin in 1215, required King John of England to proclaim certain rights (to nobles), respect certain legal procedures, and generally accept that his will could be bound by the law. It notably included the writ of *habeus corpus*, allowing appeal against unlawful imprisonment. It led to the rule of constitutional law today in the English-speaking world.

#21

Okay, so *who* invented the Internet?

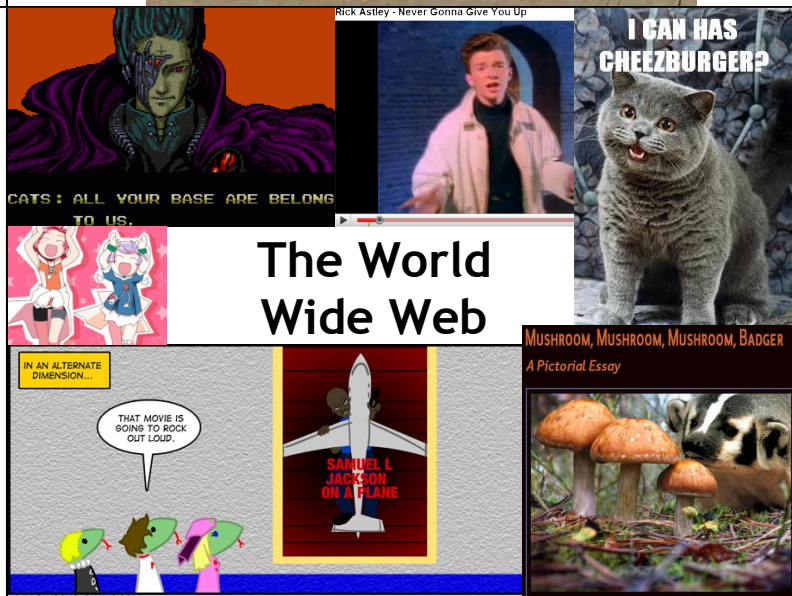


#22

The Modern Internet

- Packet Switching: Leonard Kleinrock (UCLA) thinks he did, Donald Davies and Paul Baran, Edtelcrantz's signalling network (1809)
- Internet Protocol: Vint Cerf, Bob Kahn
- Vision, Funding: J.C.R. Licklider, Bob Taylor
- Government: Al Gore (first politician to promote Internet, 1986; act to connect government networks to form "Interagency Network")

#23



Available within the network will be functions and services to which you subscribe on a regular basis and others that you call for when you need them. In the former group will be investment guidance, tax counseling, selective dissemination of information in your field of specialization, announcement of cultural, sport, and entertainment events that fit your interests, etc. In the latter group will be dictionaries, encyclopedias, indexes, catalogues, editing programs, teaching programs, testing programs, programming systems, data bases, and – most important – communication, display, and modeling programs. **All these will be – at some late date in the history of networking – systematized and coherent; you will be able to get along in one basic language up to the point at which you choose a specialized language for its power or terseness.**

J. C. R. Licklider and Robert W. Taylor, *The Computer as a Communication Device*, April 1968

#25

The World Wide Web

- Tim Berners-Lee, CERN (Switzerland)
- First web server and client, 1990
- Established a *common language* for sharing information on computers
- Lots of previous attempts (Gopher, WAIS, Archie, Xanadu, etc.)

#26

World Wide Web Success

- World Wide Web succeeded because it was **simple!**
 - Didn't attempt to maintain links, just a common way to name things
 - **Uniform Resource Locators (URL)**

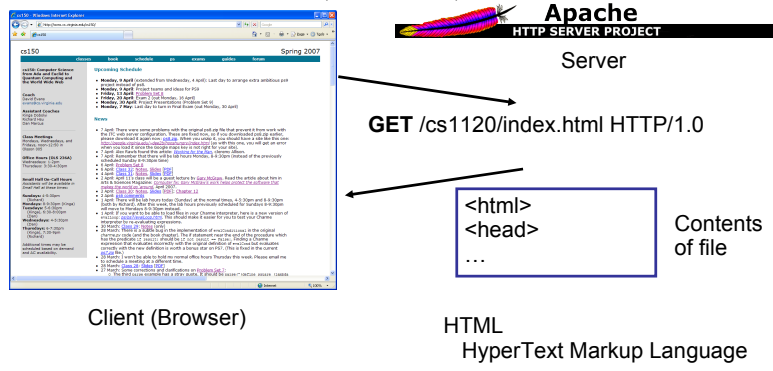
`http://www.cs.virginia.edu/cs1120/index.html`

Service Hostname File Path

HyperText Transfer Protocol

#27

HyperText Transfer Protocol (HTTP)



#28

HTML: HyperText Markup Language

- **HTML** is a language for controlling presentation of web pages
- Uses formatting tags
 - Enclosed between `<` and `>`
- **Not** a universal programming language
 - Proof: no way to make an infinite loop

#29

HTML Grammar Excerpt

```

Document ::= <html> Header Body </html>
Header ::= <head> HeadElements </head>
HeadElements ::= HeadElement HeadElements
HeadElement ::= <title> Element </title>

Body ::= <body> Elements </body>
Elements ::= Element Elements
Elements ::=
Element ::= <p> Element </p>
                Make Element a paragraph.
Element ::= <center> Element </center>
                Center Element horizontally on the page.
Element ::= <b> Element </b>
                Display Element in bold.
Element ::= Text
  
```



Plus text encoding details ...

What is a HTML interpreter?

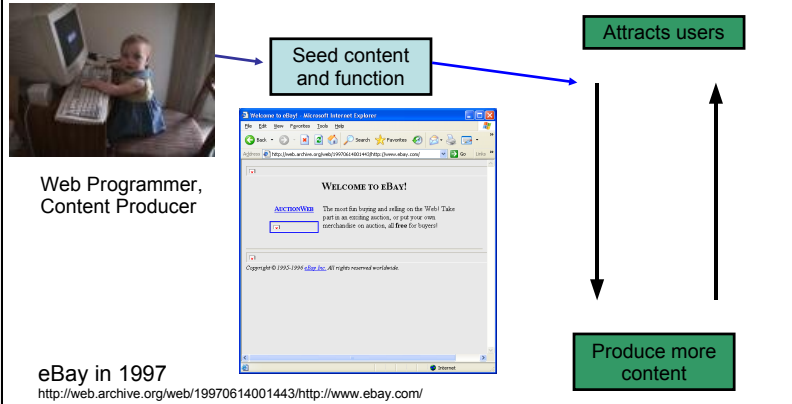
#30

Popular Web Site: Strategy 1 Static, Authored Web Site



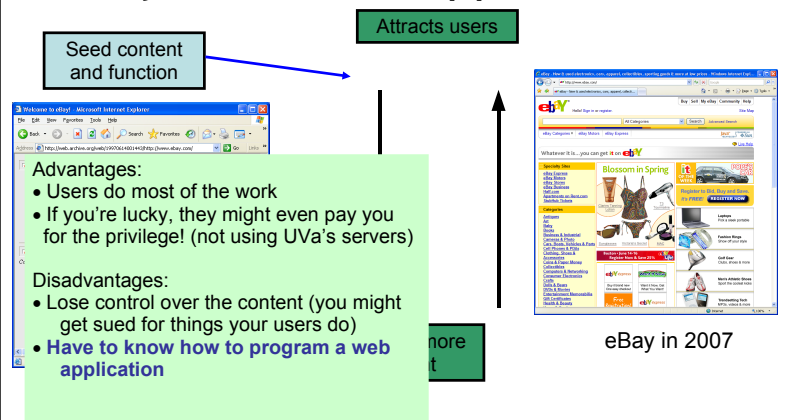
#31

Popular Web Site: Strategy 2 Dynamic Web Applications



#32

Popular Web Site: Strategy 2 Dynamic Web Applications



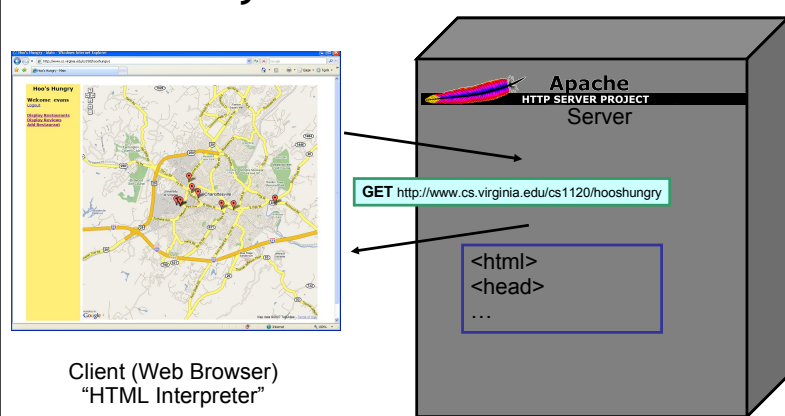
#33

Dynamic Web Sites

- Programs that run on the client's machine
 - Java, JavaScript, Flash, etc.: language must be supported by the client's browser (so they are usually flaky and don't work for most visitors)
 - Used mostly to make annoying animations to make advertisements more noticeable
 - Occasionally good reasons for this: need a fancy interface on client side (like Google Maps)
- Programs that run on the web server
 - Can be written in any language, just need a way to connect the web server to the program
 - Program generates regular HTML - works for everyone
 - (Almost) Every useful web site does this

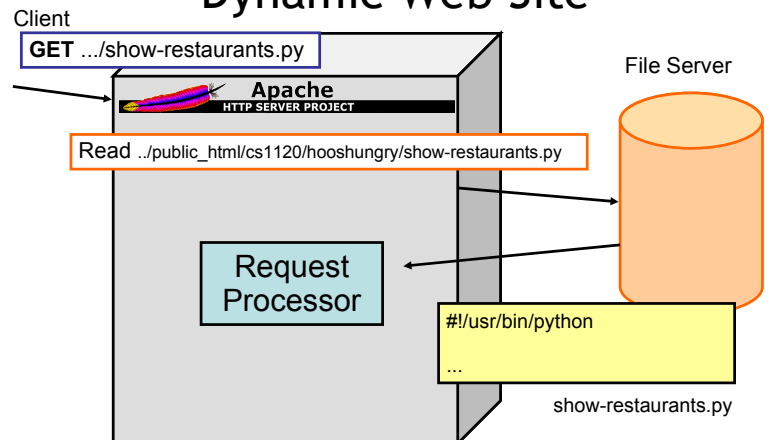
#34

Dynamic Web Site



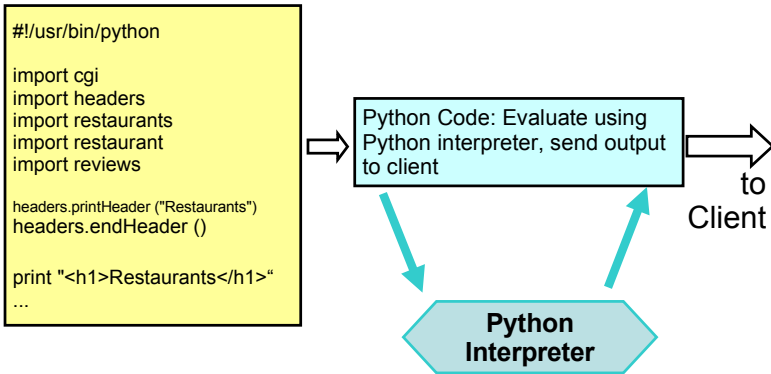
#35

Dynamic Web Site



#36

Processing a GET Request



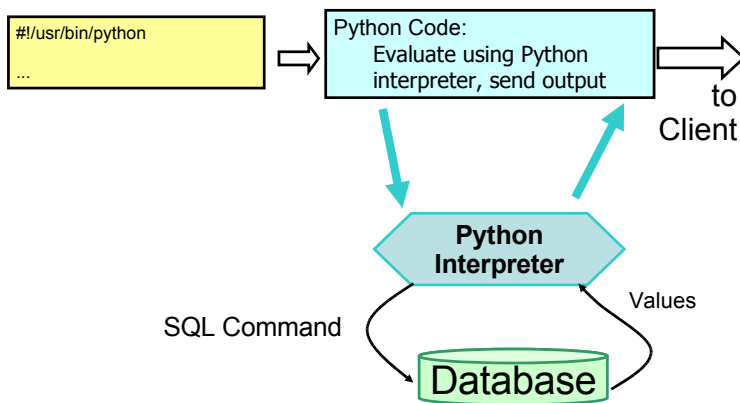
#37

Using a Database

- HTTP is **stateless**
 - No history of information from previous requests
- We probably need some state that changes as people visit the site
- That's what databases are for - store, manipulate, and retrieve data



#38



#39

SQL

- Structured Query Language (**SQL**)
 - (Almost) all databases use it
- Database is tables of fields containing values
- All fields have a type (and may have other attributes like UNIQUE)
- Similar to procedures from PS5

#40

Liberal Arts Trivia: Linguistics (and Sociology)

- This linguistic relativity hypothesis postulates a systematic relationship between the grammatical categories of a language and how the speaker understands and behaves in the world. In essence, it holds that a language's nature influences the habitual thought of its speakers: different languages yield different patterns of thought. Ideas that are prevalent in the culture can be stated concisely (in few words); foreign thoughts are difficult to express.

#41

Liberal Arts Trivia: Chinese History

- This beverage is made solely of the leaves of *Camellia sinensis* that have undergone minimal oxidation during processing. The drink originated in China more than 4000 years ago but has spread (e.g., it is also ubiquitous in Japan).



Liberal Arts Trivia: Engineering, Architecture and Physics

- This bridge's main span famously collapsed on July 1, 1940 due to aeroelastic flutter caused by a 42 mph wind. In 1998 the film of it was selected for preservation in the US Library of Congress as being "culturally, historically or aesthetically significant." The footage is still shown to students as a cautionary tale.



Liberal Arts Trivia: Latin American Studies, Archaeology



- This civilization began as a Cuzco-area tribe around 1200 and grew to absorb other Andean communities, becoming the largest empire in pre-Columbian America. They invented the quipu ("talking knots") for recording decimal numbers in knotted strings of llama hair. They also performed the first successful skull surgery, as well as using coca leaves to deaden pain. Machu Picchu is a World Heritage site associated with this culture.

#44

Secure Programming

cs1120

"Honor System" Programming

All your users are nice and honest
Nothing terribly bad happens if your program misbehaves

Enough to (hopefully) make you dangerous!

cs2220

"Real World" Programming

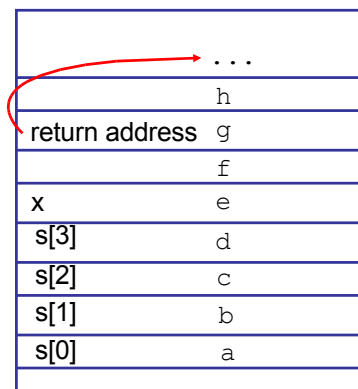
Some users are mean and dishonest
Bad things happen if your program misbehaves

Buffer Overflows

```
int main (void) {
    int x = 9;
    char s[4];

    gets(s);
    printf ("s is: %s\n", s);
    printf ("x is: %d\n", x);
}
```

C Program



Stack

Buffer Overflows

```
int main (void) {
    int x = 9;
    char s[4];

    gets(s);
    printf ("s is: %s\n", s);
    printf ("x is: %d\n", x);
}
```

```
> gcc -o bounds bounds.c
> bounds
abcdeghijkl (User input)
s is: abcdeghijkl
x is: 9
> bounds
abcdeghijklm
s is: abcdeghijklm
x is: 1828716553 = 0x6d000009
> bounds
abcdeghijkln
s is: abcdeghijkln
x is: 1845493769 = 0x6e000009
> bounds
aaa... [a few thousand characters]
crashes!
```

What does this kind of mistake look like in a popular server?

Note: your results may vary (depending on machine, compiler, what else is running, time of day, etc.). This is what makes C fun!

Code Red



Thu Jul 19 00:00:00 2001 (UTC)
Victims: 159

<http://www.caida.org/>
Copyright (C) 2001 UC Regents, Jeff Brown for CAIDA/UCSD

Security in cs1120

Can you have a Buffer Overflow vulnerability in Scheme, Charme, LazyCharme, StaticCharme, or Python?

Security in cs1120

Can you have a Buffer Overflow vulnerability in Scheme, Charme, LazyCharme, StaticCharme, or Python?

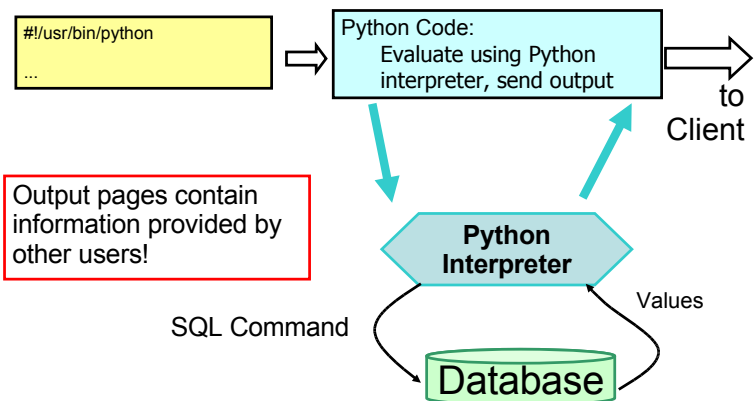
No (unless there is a bug in the underlying implementation)! Memory is managed by the interpreter, so you don't have to allocate it, or worry about how much space you have.

Web Application Security

- Malicious users can send bad input to your application
- **Authentication:** most interesting applications need user logins



Cross-Site Scripting



Cross-Site Scripting Demo

user: weimer
password: \$1\$34254\$fEzeA5iPHLRixnlhpXj0p0
url: https://church.cs.virginia.edu/cs1120/hooshungru/

Enter Review:

```
<script language="javascript">
function button()
{
    while (1) alert("I Own you!")
}
</script>
<BODY onLoad="button()">
```


Preventing Cross-Site Scripting

- Never never never ever trust users!
- Everything you generate from user input needs to be checked and sanitized (remove the tags)

For your ps9 websites, you **may assume** all users are bound by the UVa Honor Code and won't do anything evil. But, don't forget how irresponsible it is to put something like this on the web!


Authentication

NINJAS vs PROFESSORS
A COMPARATIVE ANALYSIS



NINJAS

- Experts in methods of subterfuge
- Employs assortment of lethal weapons
- Can kill you without remorse
- Always shown wearing the same outfit
- Wears a hood
- Hurls Shurikens ✨
- People think they're pretty cool
- Shrouded in mystery



PROFESSORS

- Experts in methods no longer used
- Employs a bunch of lazy peons (you)
- Can kill your career or worse
- Always wears the same outfit
- Wears a hood at graduation
- Hurls when you present your research
- They think they're pretty cool
- Shrouds you in misery

WWW.PHDCOMICS.COM

How would I prove that I am a professor and not a ninja?

How do you authenticate?

- Something you know
 - Password
- Something you have
 - Physical key (email account?, transparency?)
- Something you are
 - Biometrics (voiceprint, fingerprint, etc.)

Serious authentication requires at least 2 kinds

Early Password Schemes

Login does direct password lookup and comparison.

UserID	Password
alyssa	fido
ben	schemer
weimer	Lx.Ly.x


Login: alyssa
Password: spot
Failed login. Guess again.

Login Process

Terminal

Login: alyssa
Password: fido

login sends <"alyssa", "fido">



Password Problems

- Need to store the passwords
 - Dangerous to rely on database being secure

Today

- Need to transmit password from user to host
 - Dangerous to rely on Internet being confidential

Later Class

First Try: Encrypt Passwords

- Instead of storing password, store password encrypted with secret K .
- When user logs in, encrypt entered password and compare to stored encrypted password.

UserID	Password
alyssa	encrypt_K ("fido")
ben	encrypt_K ("schemer")
weimer	encrypt_K ("Lx.Ly.x")

Problem if K isn't so secret: $\text{decrypt}_K(\text{encrypt}_K(P)) = P$

Homework

- PS8 Due Today
- PS9 Description Due Wednesday
- Exam 2 Out Wednesday