

# CSE 224: OVERVIEW AND INTRODUCTION

George Porter  
Jan 10, 2023



# ATTRIBUTION

- These slides are released under an Attribution-NonCommercial-ShareAlike 3.0 Unported (CC BY-NC-SA 3.0) Creative Commons license

# TODO

1. Go to Canvas and take the “Onboarding Survey”
2. Start project 0 (due Jan 17)

# WELCOME!





# CSE 224: NETWORKED SYSTEMS

- Add networking support to software
  - Between two computers
  - Between computer and datacenter (“The Cloud”)
- Develop software that is:
  - Scalable (handles 100s of M to 1+ billion users)
  - Fault-tolerant (survives failures)
  - Evolvable (how to update services without making them unavailable to end users)

# MODERN SOFTWARE INCREASINGLY NETWORKED



# NETWORKED SERVICES DRIVEN BY DATA

```
0 1 0 1 0 1 0 0 0 1 1 1
0 0 1 1 0 1 0 0 1 0 1 1
0 0 1 1 0 1 0 1 1 1 0 0
1 0 1 0 0 1 0 0 1
0 1 1 Data 1 0 0
1 1 0 1 1 0
0 1 1 0 0 0 0 0 1 1 1 1
0 0 0 1 0 1 0 1 1 1 1 1
1 1 1 0 1 0 1 1 1 0 0 0
```

+

**amazon.com**<sup>®</sup>

=

Product  
Recommendations

```
0 1 0 1 0 1 0 0 0 1 1 1
0 0 1 1 0 1 0 0 1 0 1 1
0 0 1 1 0 1 0 1 1 1 0 0
1 0 1 0 0 1 0 0 1
0 1 1 Data 1 0 0
1 1 0 1 1 0
0 1 1 0 0 0 0 0 1 1 1 1
0 0 0 1 0 1 0 1 1 1 1 1
1 1 1 0 1 0 1 1 1 0 0 0
```

+



**Spotify**<sup>®</sup>

=

Custom  
Stations

```
0 1 0 1 0 1 0 0 0 1 1 1
0 0 1 1 0 1 0 0 1 0 1 1
0 0 1 1 0 1 0 1 1 1 0 0
1 0 1 0 0 1 0 0 1
0 1 1 Data 1 0 0
1 1 0 1 1 0
0 1 1 0 0 0 0 0 1 1 1 1
0 0 0 1 0 1 0 1 1 1 1 1
1 1 1 0 1 0 1 1 1 0 0 0
```

+

**Google**

=

Personalized  
Search

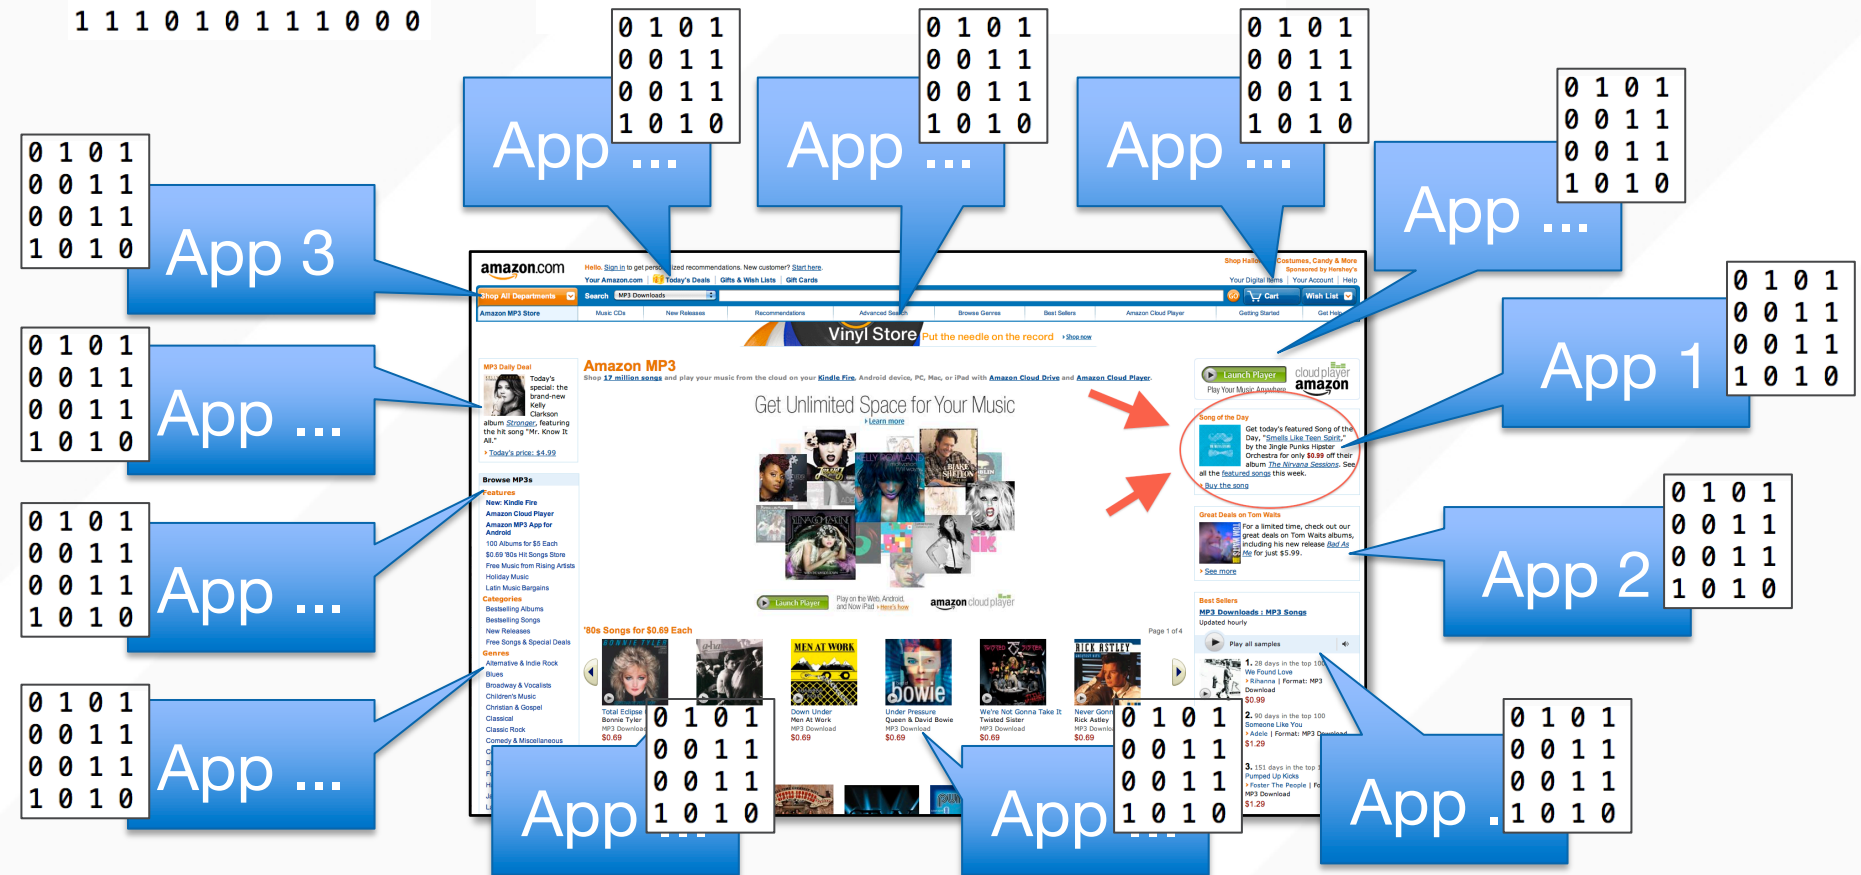
# DATA-DRIVEN, PER-USER CUSTOMIZATION + ML

0	1	0	1	0	1	0	0	0	1	1	1
0	0	1	1	0	1	0	0	1	0	1	1
0	0	1	1	0	1	0	1	1	1	0	0
1	0	1							0	0	1
0	1	1							1	0	0
1	1	0							1	1	0
0	1	1	0	0	0	0	0	1	1	1	1
0	0	0	1	0	1	0	1	1	1	1	1
1	1	1	0	1	0	1	1	1	0	0	0

Data

+ **amazon.com** =

Product  
Recommendations



# MAJOR THEMES OF THE COURSE

- Programming abstractions for communicating over the Internet through various network protocols
- Naming and indexing to find services and connect clients with servers (or clients with other clients)
- Managing scale; scale-out design
- Replicating and updating “mutable” data over the network
- Replicating and caching “immutable” data over the network (think Netflix, Disney+, Youtube, etc)
- Accessing and managing networked storage
- Managing fault tolerance

# HOW CAN YOU WRITE SOFTWARE THAT WORKS DESPITE ADVANCEMENTS IN UNDERLYING TECHNOLOGY?

Think about the first computer you remember using...

- Can we find the oldest example here in class today? The most recent example?

Think about the first network you used (modem? Fiber optics? Mobile network?)

- Can we find the oldest example here in class today? The most recent example?

Discuss with the 3-4 people nearest you for 2 minutes and let's find out!

**THINK ABOUT HOW ONLINE NETWORK SERVICES  
HAVE CHANGED OVER THE PAST 20-ISH YEARS...**

## GOOGLE (1998)



Search the web using Google!

Google Search

I'm feeling lucky

Special Searches

[Stanford Search](#)

[Linux Search](#)

[Help!](#)

[About Google!](#)

[Company Info](#)

[Google! Logos](#)

Get Google!

updates monthly:

your e-mail

Subscribe

[Archive](#)

Copyright ©1998 Google Inc.



# NETFLIX (2002)



**Try Netflix for FREE Today!**

**Rent all the DVD movies you want.**  
For 20 bucks a month. No late fees.

  
World's Largest Selection

  
FREE & FAST HOME DELIVERY

  
NO LATE FEES



**Super Selection!**  
Create a list online of all the movies you want to see

**Free & Fast Home Delivery**  
The movies you select arrive via first-class mail in 2-4 days.

**No Due Dates or Late Fees**  
Keep each DVD as long as you want. Have up to 3 movies on hand.

**Free Shipping!**  
Return one DVD in its prepaid envelope and get another DVD from your list.

  
**\$19.95 PER MONTH**

**One Flat Fee!**  
It's just 20 bucks a month. There are no late fees, no hidden charges, no commitments. If you have any questions, call 1-888-638-3549.

# TWITTER (2007)



Find folks to follow!  search or [Login](#) / [Join Twitter!](#)

A global community of friends and strangers answering one simple question: **What are you doing?** Answer on your phone, IM, or right here on the web!

explore twitter 

Look at what [these people](#) are doing right now...



**claudiof** @ruilmoura: a minha pergunta nao era pa ter a certeza, basta fazeres whols e ves que sim. era mais... "é isto? consolas?" [less than 5 seconds ago](#) from im [in reply to ruilmoura](#)



**mseling** I've started the bad habit of forgetting to eat breakfast and lunch. No wonder i'm always tired. [less than 5 seconds ago](#) from [bt](#)



**cbsiskin** @sharongs It's on its way via email. Not online so can't send a link.. [less than 5 seconds ago](#) from web [in reply to sharongs](#)



**groovesalad** Mushroom Nation - Helsinki (Pancake Mix) [less than 5 seconds ago](#) from web



**2525** weer thuis [less than 5 seconds ago](#) from [twitterrific](#)



**erockenjew** Billions of dollars spent on dirty contractors in Iraq: <http://tinyurl.com/ywkt3c>, but 35 billion over 5 years for poor children is soCI .... [less than 5 seconds ago](#) from [twitterrific](#)



**Becks6736** mein bruder isch jetzt au bei twitter.... sein nickname isch Jesusfan!! Bis dennes ihr hennes!! [less than 10 seconds ago](#) from web



**Levoix** End of my blog : <http://levoix.canalblog.com> [less than 10 seconds ago](#) from web



**rauski** El diseño gráfico es la peor de las parejas y la mejor de las amantes, así pasa... ¿Porqué siempre nos gustan las cabronas? jajajaja [less than 10 seconds ago](#) from web



Please Sign In!

Username or Email

Password

☐ Remember me [Forgot?](#)

[Sign In!](#)

Already using Twitter with your Mobile or IM? [Activate »](#)

Want an account?

**Join for Free!**

It's fast and easy!

Featured!

[FreeBurma](#)

[Amie Street](#)

[WIREDScience](#)

[New Media Expo](#)

[New York Times](#)

[MTV VMAs](#)

[MacRumors Notify](#)


[Lincoln@APEC2007](#)

[Chuck Bartowski](#)

[AIBF](#)

© 2007 Twitter [About Us](#) [Contact](#) [Blog](#) [Explore!](#) [API](#) [Help](#) [Jobs](#) [TOS](#) [Privacy](#)

# YOUTUBE (2004)



# YOUTUBE (2005)

[Home](#) | [My Videos](#) | [My Favorites](#) | [My Messages](#) | [My Profile](#)[Sign Up](#) | [Log In](#) | [Help](#)

# You Tube

Your Digital Video Repository

[Upload Your Videos](#)

[nansheng](#) : [azlan](#) : [wereldband](#) : [ny](#) : [superbike](#) : [japan](#) : [sinceretheory](#) : [jozef](#) : [party](#) : [amazon](#) : [board](#) : [skate](#) : [buckley](#) : [shubs](#) : [falls](#) : [de](#) : [stockshot](#) : [cubbyhole](#) : [burnout](#) : [satellite](#) : [poughkeepsie](#) : [cruise](#) : [heritage](#) : [orgel](#) : [chin](#) : [themed](#) : [mill](#) : [music](#) : [new](#) : [live](#) : [to](#) : [farmer](#) : [mtv](#) : [puenbrouck](#) : [sicily](#) : [fairfield](#) : [musical](#) : [coffeehouse](#) : [bud](#) : [2005](#) : [trip](#) : [jfk](#) : [woordjes](#) : [death](#) : [xlantz](#) : [skill](#) : [olle](#) : [nature](#) : [ads](#) : [dance](#) :  
[See More Tags](#)

Featured Videos

[Denny's](#)  
Added: June 14, 2005  
by [jchyley](#)  
Views: 88 | Comments: 1

[On top of the world!](#)  
Added: May 3, 2005  
by [javed](#)  
Views: 62 | Comments: 0

[Father's Day Special](#)  
Added: June 19, 2005  
by [ATrain](#)  
Views: 46 | Comments: 0

[Jingle en Joe](#)  
Added: June 21, 2005  
by [Lena](#)  
Views: 27 | Comments: 0

[River Otter in Central Park Zoo](#)  
Added: June 28, 2005  
by [IH](#)  
Views: 6 | Comments: 0

[Watch More Videos](#)

[About Us](#) | [Contact Us](#) | [Terms of Use](#) | [Privacy Policy](#) | Copyright © 2005 YouTube, LLC™ | [RSS](#)

# FACEBOOK (2004)



# SURVEY

- When was the web protocol created? When was the first graphical web browser released?
1. 1968 / 1972
  2. 1974 / 1976
  3. 1989 / 1993
  4. 2001 / 2002


# THE DEPLOYMENT OF “THE WEB”

# RISE OF THE WEB





# RISE OF THE WEB




# **WHERE DO NETWORK SERVICES EXECUTE?**

# THE FIRST WEB SERVER (NEXT WORKSTATION, 1991)



# THE RISE OF THE “DATACENTER” (AKA CLOUD COMPUTING)





# DATACENTERS: THE HOME OF ALL THIS COMPUTING AND STORAGE



Microsoft



Google




Facebook



# Google 2012







Microsoft



Google








Facebook





# HARDWARE HAS EVOLVED AS WELL. STARTING WITH CPUS...






# TO GPUS...



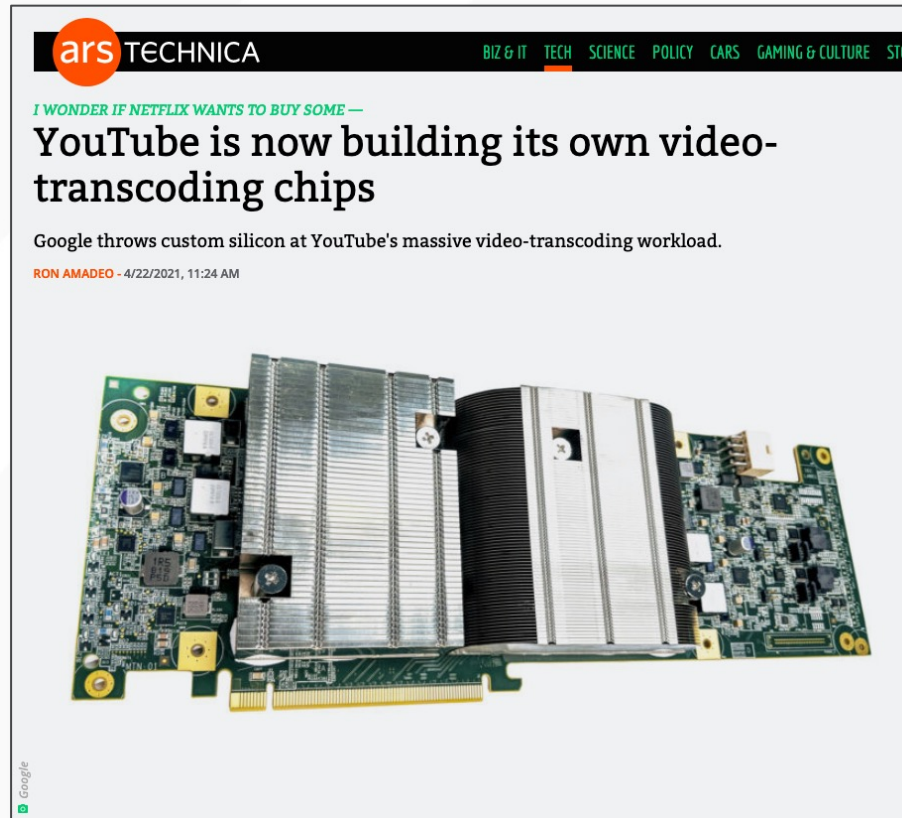
# TO PROGRAMMABLE FPGAS...



# TO CUSTOM DESIGNED CHIPS



# CHUSTOM VIDEO TRANSCODING CHIP





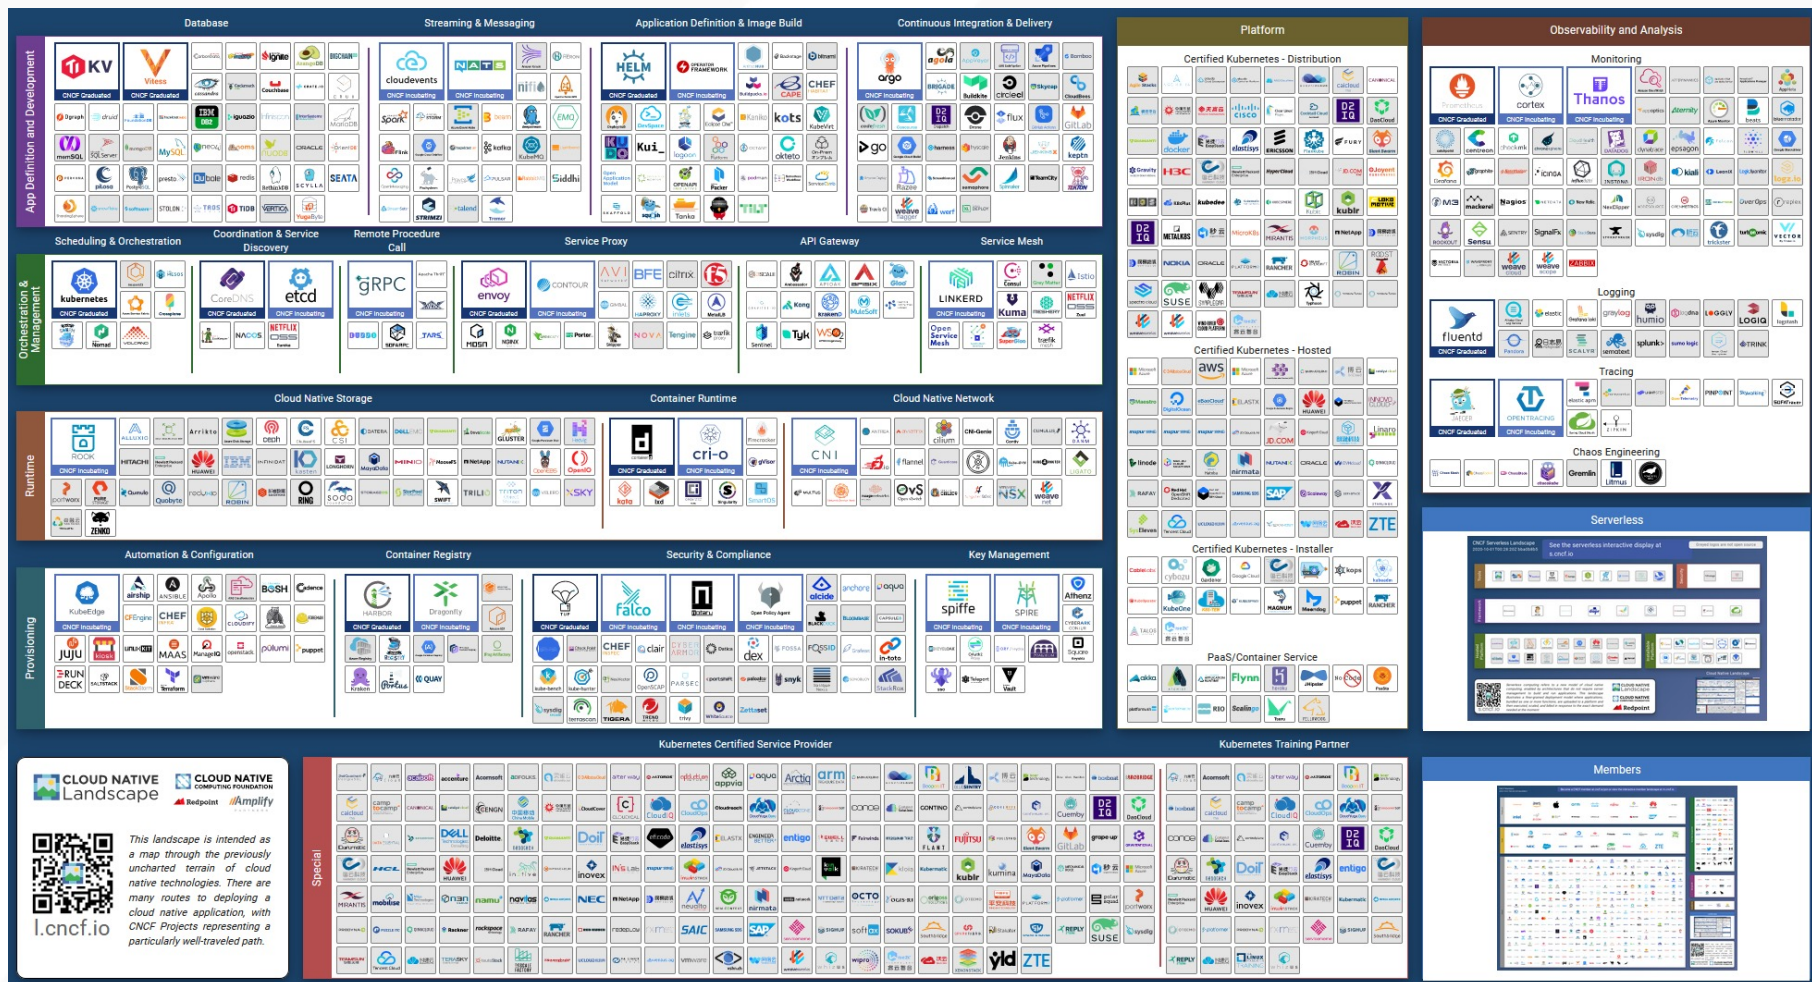
# CLUSTERS OF CUSTOM ASICS FOR AI/MACHINE LEARNING



*Source: google.com*



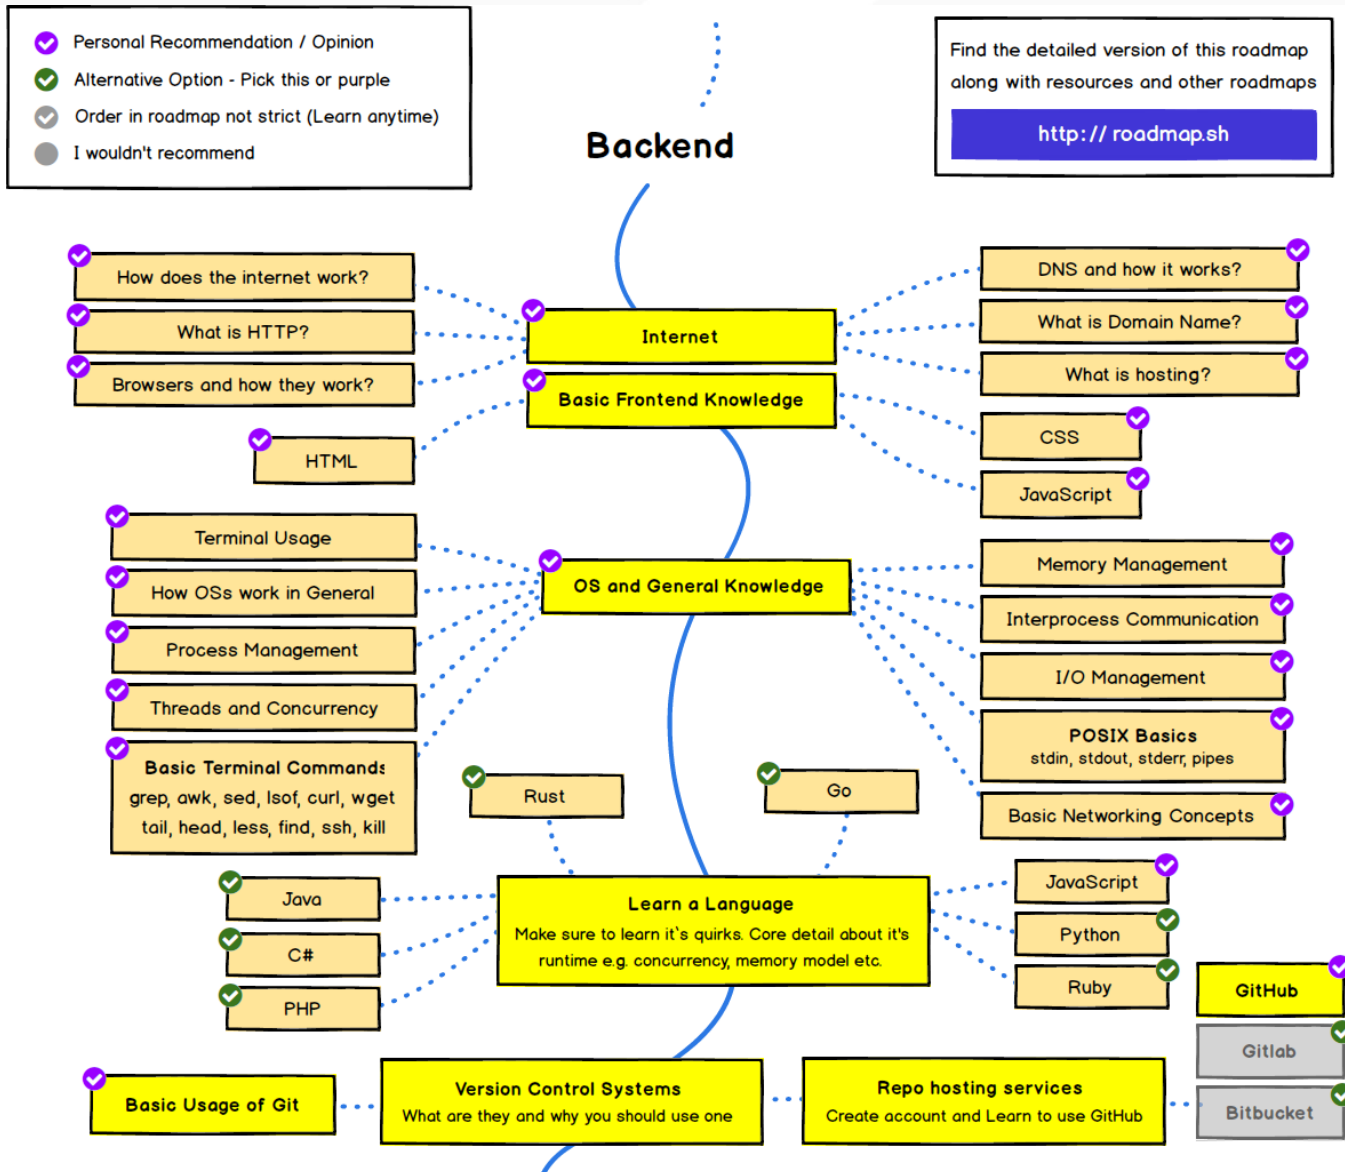
# FULL CLOUD NATIVE LANDSCAPE



# CLOUD NATIVE LANDSCAPE IN A 10-WEEK QUARTER

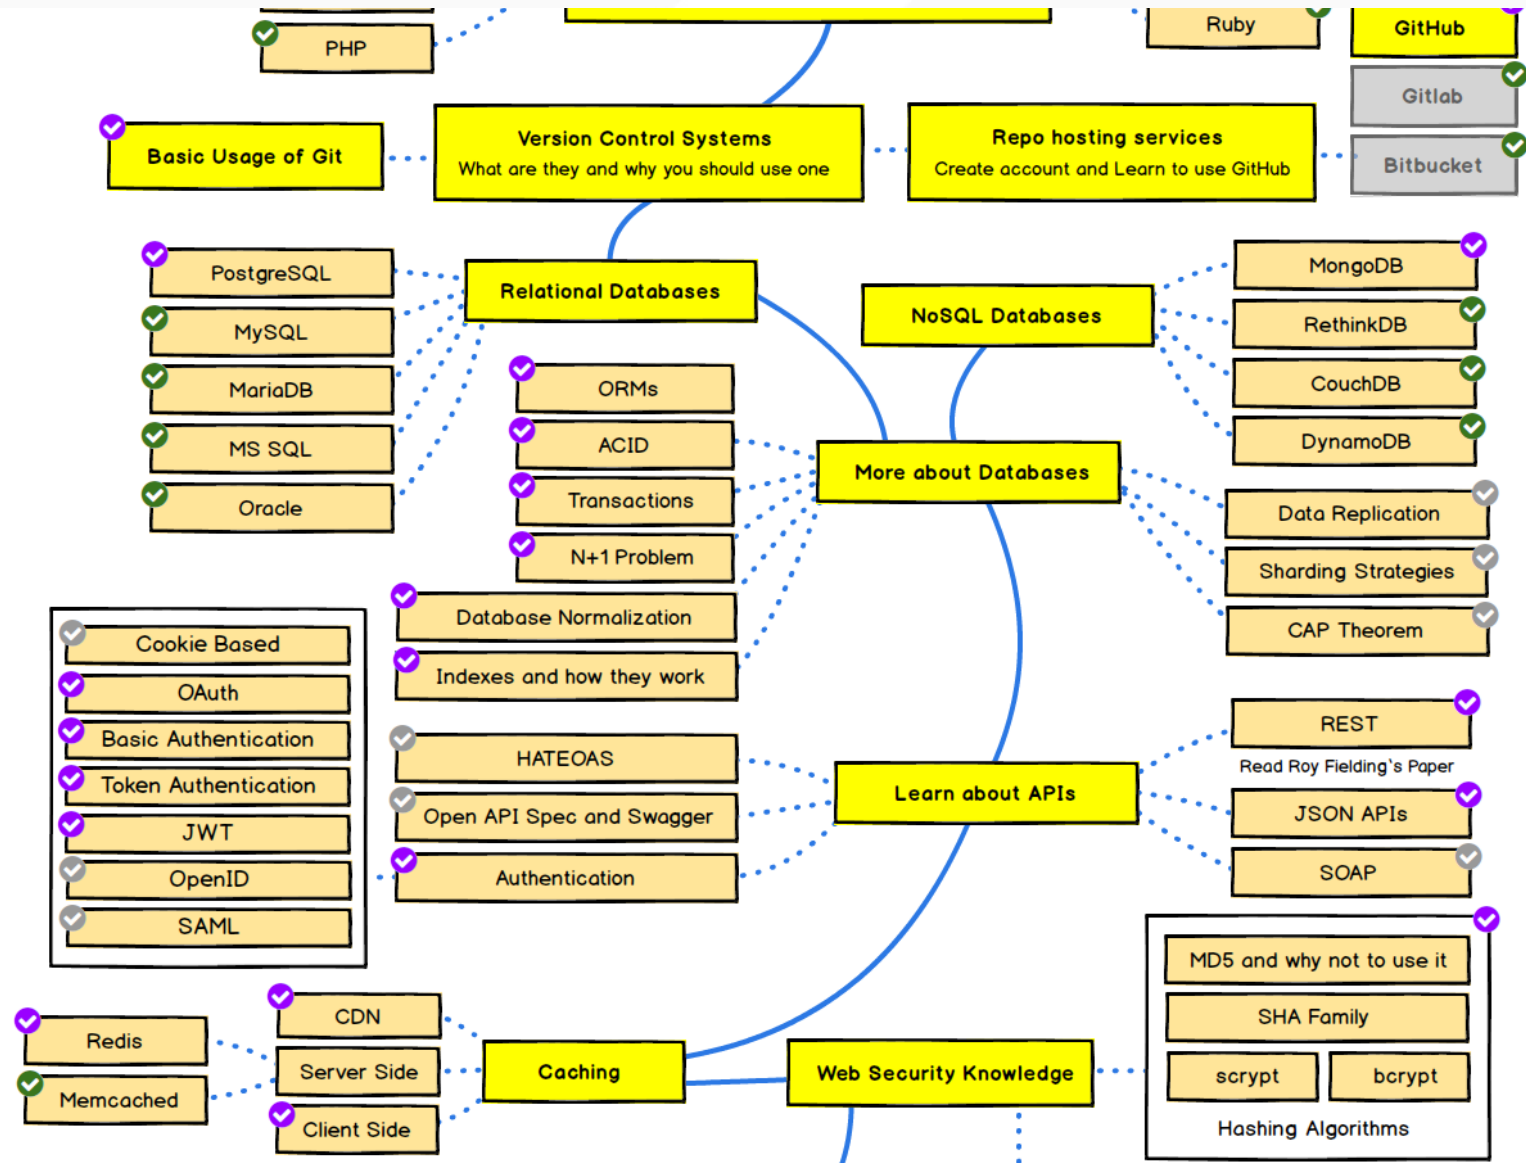


# BACKEND DEVELOPMENT ROADMAP

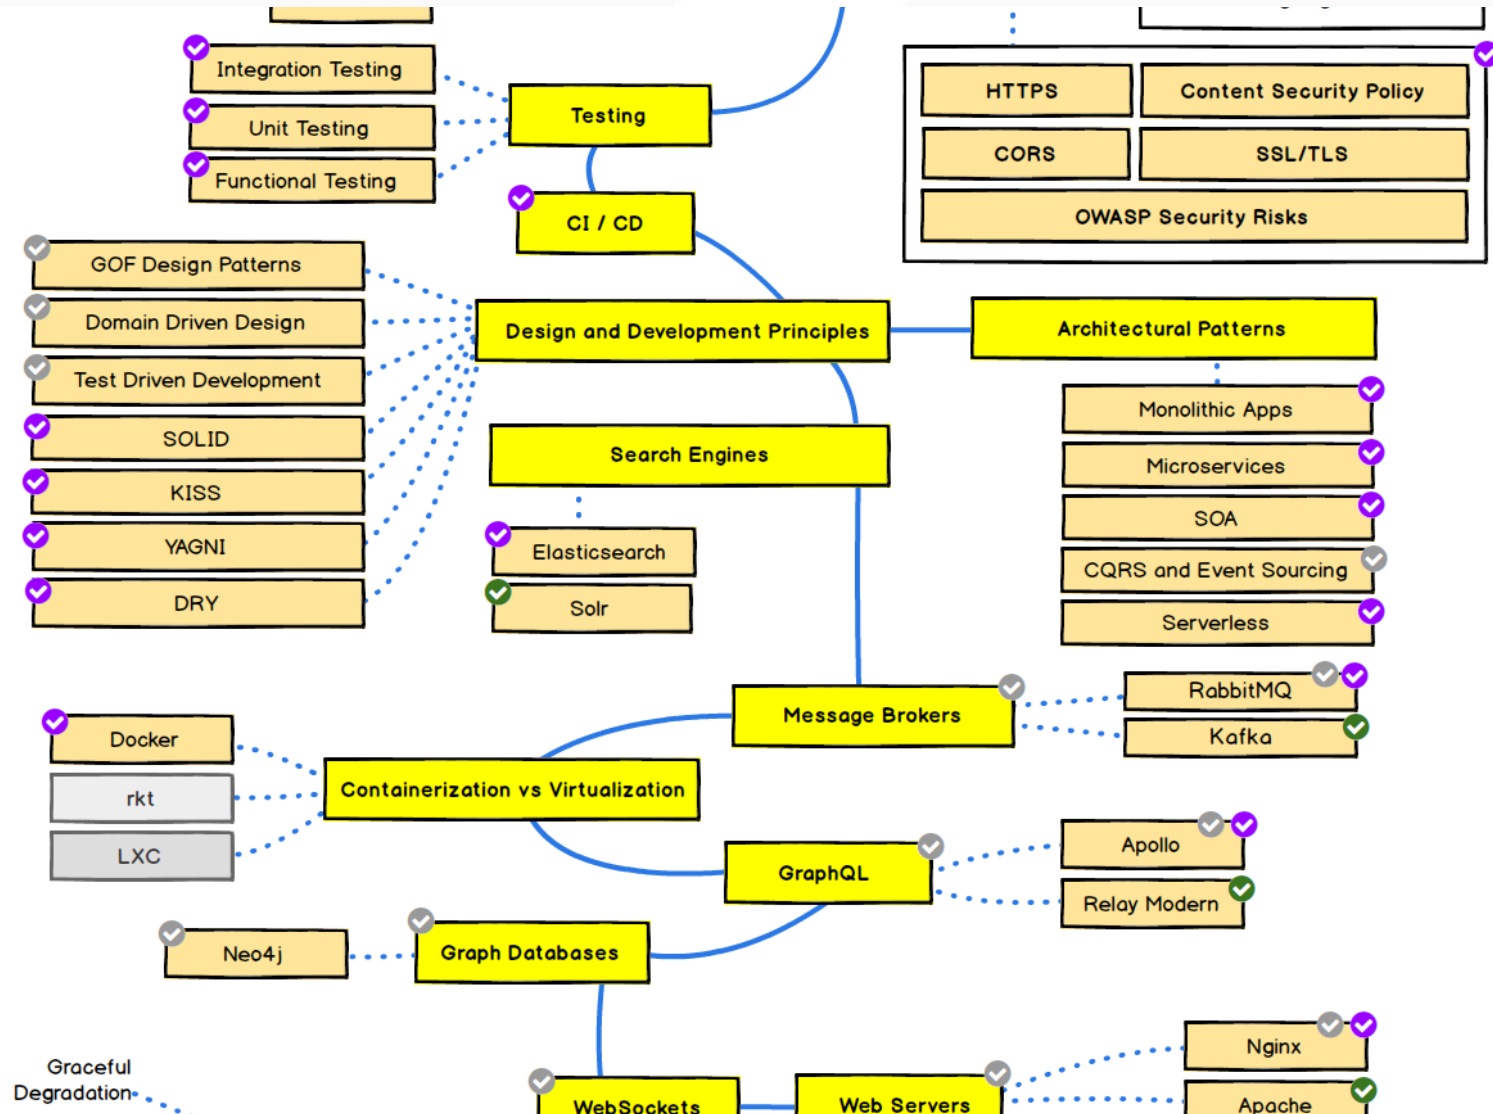




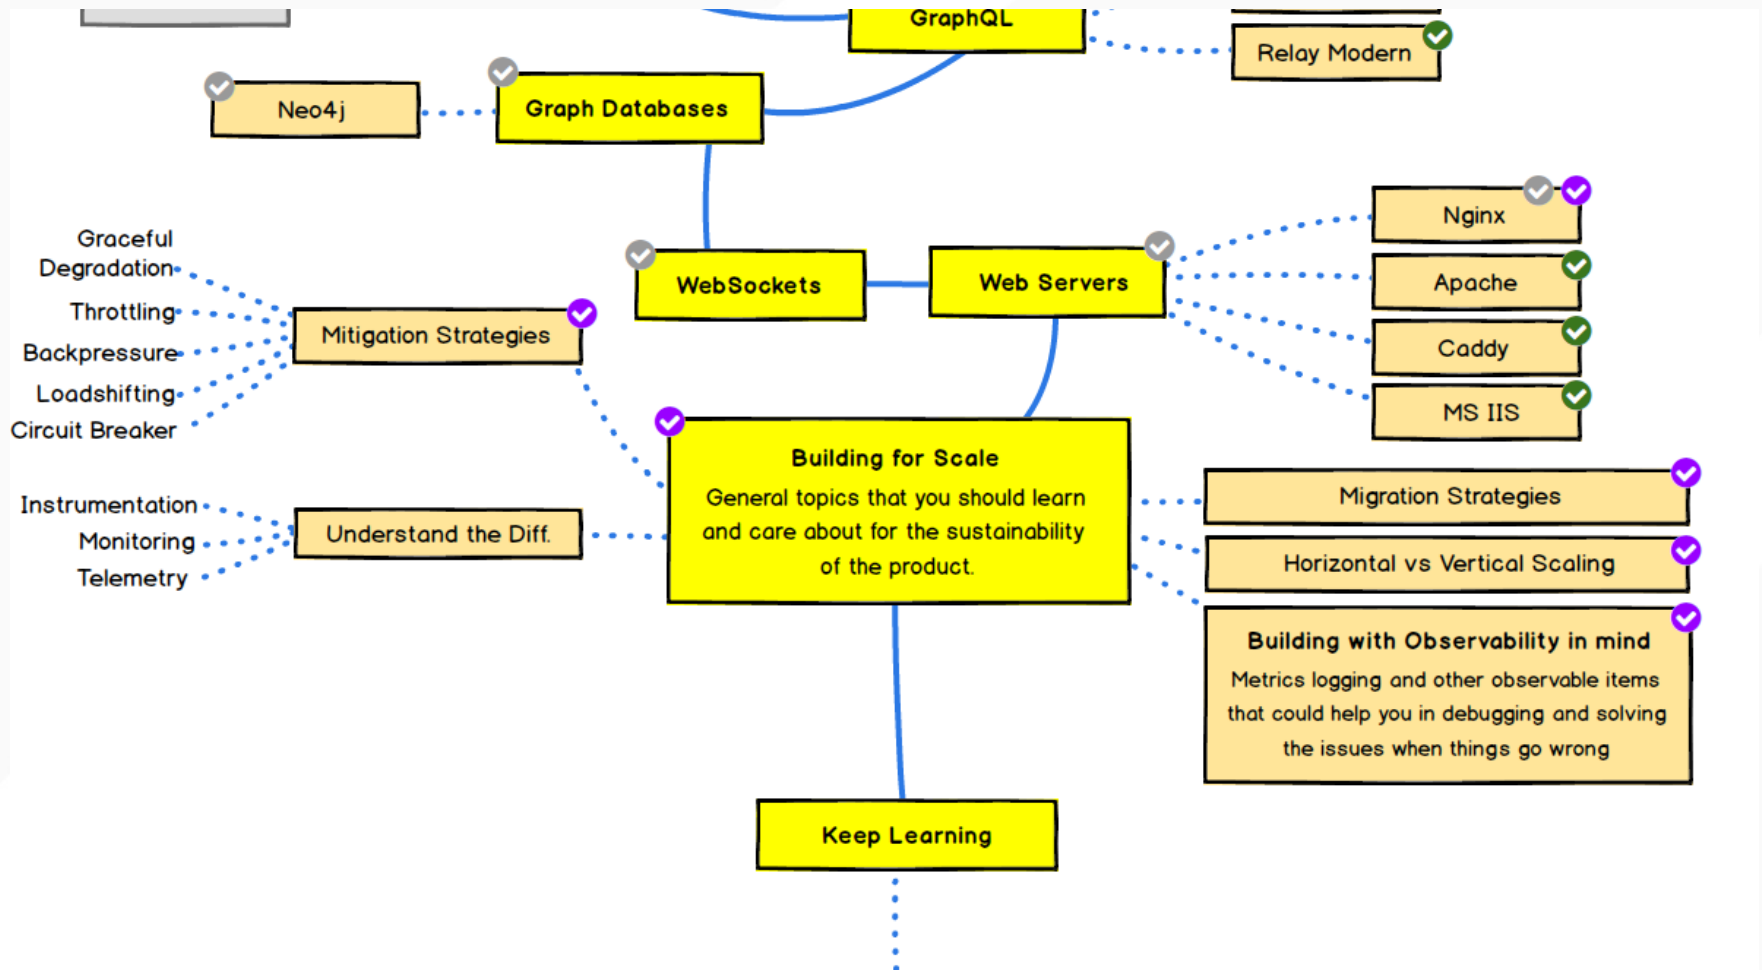
# BACKEND DEVELOPMENT ROADMAP



# BACKEND DEVELOPMENT ROADMAP



# BACKEND DEVELOPMENT ROADMAP



# THE ENVIRONMENTAL IMPACT OF CLOUD COMPUTING




1. LBNL, 2013
2. NRDC report

- Carbon/energy footprint:
  - 1-2% of global energy consumption<sup>1</sup>
  - 140 billion kWh (50 power plants)<sup>2</sup>
  - 100 metric tons of carbon pollution per year<sup>2</sup>

# Google's energy footprint

RENEWABLE ENERGY PURCHASING COMPARED WITH TOTAL ELECTRICITY USE






# SCALING ACROSS TECHNOLOGY IMPROVEMENTS


- Network primitives are designed to scale
- Techniques we learn are directly applicable to global-scale services like Google, Facebook, ...
- Your projects will be tested in small scale
  - Yet could scale immensely with minimal to no modifications

# HOW TO BUILD SUCH LARGE SYSTEMS?



# HOW TO BUILD SUCH LARGE SYSTEMS?

- Systems...
- Built on top of abstractions...
- Built on software...
- Built on hardware...



*We will cover the software abstractions to enable you to write networked software*

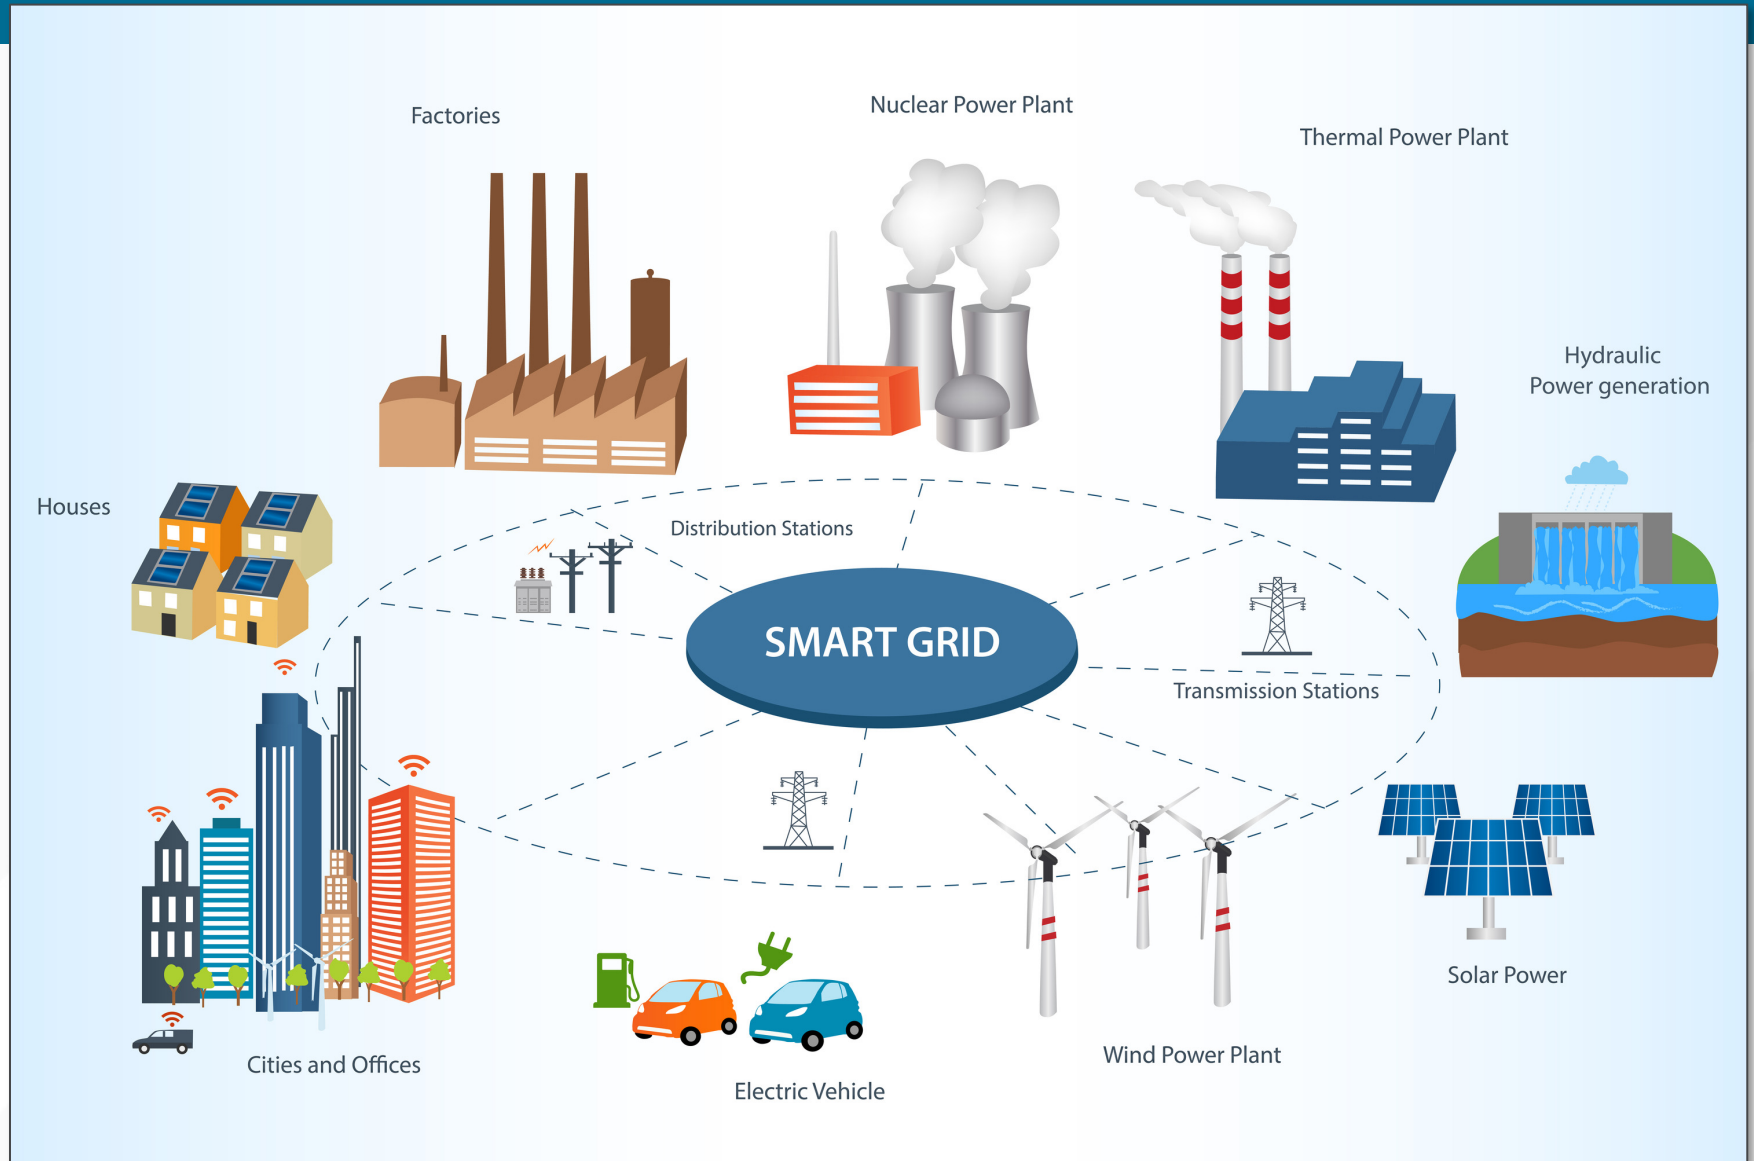
**IT'S NOT JUST WEBSITES AND SOCIAL MEDIA  
THOUGH!**

# SELF-DRIVING CARS AND SMART CITIES





# SMART CITIES AND SMART GRIDS



# CSE 224 VS {221,222A,223B}

- 224: Graduate Networked Systems
  - How to program networked software
  - Socket programming, RPC, protocol design and implementation, consensus and consistency, security, TLS, ...
  - Designed as a *broad survey* of systems thinking
  - Learn through hands-on, programming-based projects
- 224 Target audience:
  - MS “comps” students and BS/MS students
  - Non-systems MS “thesis” and non-systems Ph.D. students
- Note:
  - Cannot receive credit for both 124 and 224
- Research-focused depth sequence
  - 221: Operating Systems
  - 222A: Networking
  - 223B: Distributed systems theory
  - Deep dives into peer-reviewed literature
  - Learn through close readings and in-class discussion of 4 research papers per week
- 221/222A/223B Target audience:
  - Systems MS “thesis” and Systems Ph.D students

# THE CHALLENGE OF NETWORKING


- CS undergraduate curricula includes:
  - Algorithms
  - Programming languages
  - Architecture
  - Data structures
  - Etc...
- How does the network change each of these areas?



# RESOURCES

- Website
  - <https://canvas.ucsd.edu/courses/43955>
  - Gradebook, links to assignments + deadlines, PDFs of lecture slides, in-class demos and exercises
- Piazza discussion board (linked off Canvas)
- Github (for managing your projects)
- Gradescope (for submitting your projects)
- Two books
- TA discussion section (1x week)

# TEACHING ASSISTANTS (PAGE 1)



Aditya Barsainya



Qihuang Chen



Xiuqi Chen



Tanmay Gujar

# TEACHING ASSISTANTS (PAGE 2)



Linfang He



Saketh Khandavalli



Amanda Tomlinson




Abhishek Vijeev

# CLASS MEETINGS


- Mostly putting the material that you read into context
- Live coding demos, activities, some “mini lectures” on algorithms, protocols, etc.
- You are responsible for everything that happens during class
  - Will podcast, but can’t guarantee that system works flawlessly
- Will be asking for feedback on what works and what doesn’t work a lot during the class



# BOOKS




Free if accessed through the  
UCSD library



Free if accessed through the  
UCSD library

# PROGRAMMING SKILLS FOR THIS CLASS

- We'll be using the “Go” language
  - [golang.org](https://golang.org)
  - Designed at Google in 2007
  - Goals: improve programming productivity in an era of multicore, networked machines, and large codebases
  - Kernighan (of ‘C’ fame) co-created
- Why?
  - Simple, readable, no mem allocation (similar to Python)
  - High-performance networking
  - Concurrency/parallelism
  - Static typing and efficient runtime
  - Industry-quality and deployed at massive scale



# CLASS ROADMAP / PROJECTS / GRADING

1. Pre-lecture review question sets [5%]
2. Projects [60%]
  1. [5%] Single-node sort
  2. [10%] Distributed network sort
  3. [15%] Build your own web server
  4. [10%] SurfStore “Dropbox clone”
  5. [5%] Scaled-out SurfStore backend
  6. [15%] Fault-tolerant SurfStore backend
3. Exams [35%]
  1. [15%] Midterm (Thu Feb 9 during class time)
  2. [20%] Final exam (Tue Mar 21, 3-6pm)

# DEFAULT UCSD GRADING SCHEME

## View/Edit Grading Scheme



### eGrades

 [Select Another Scheme](#) 

**Name:**

**Range:**

A+	100 %	to 97.0%
A	< 97.0 %	to 94.0%
A-	< 94.0 %	to 90.0%
B+	< 90.0 %	to 87.0%
B	< 87.0 %	to 84.0%
B-	< 84.0 %	to 80.0%
C+	< 80.0 %	to 77.0%
C	< 77.0 %	to 74.0%
C-	< 74.0 %	to 70.0%
D	< 70.0 %	to 60.0%
F	< 60.0 %	to 0.0%



## COURSE AT A GLANCE

- Basics of networking, sockets API, DNS
- Remote procedure calls w/ Google RPC (gRPC)
- Distributed storage as an application
- Scale-out techniques and methods
- Replicating immutable state via CDNs
- Replicating mutable state with two-phase commit and replicated state machines (+ deep dive on the RAFT protocol)

# IMPORTANT CONTEXT

- In this course you'll learn some things in class (mostly tested via the exams), and you'll learn some things by *doing—working on the projects*.
- We'll cover the big themes and high-level ideas in class, but you'll be learning a lot of the details in the projects
- That's why the projects are NOT designed to be done in one long session or all-nighter—you should work on them a bit each day so you can research what you need to complete them, or to talk to the TAs/myself, etc.
- Start early—start often!

# TAKE THE ONBOARDING SURVEY

- Required by UCSD for some reason or another regarding Federal financial aid
- BUT also, there is a very important question for those of you who need this class to graduate and plan to graduate this term...
- I'm going to export **more** of you this term, not less
- While I'm not going to enforce it, you really need to be physically attending class every time, not just relying on podcasts

# COMPUTING RESOURCES

- You can use the lab computers in the building, or the “ieng6” servers that can be accessed via ssh
- `ssh <username>@ieng6.ucsd.edu`



# TODO

1. Go to Canvas and take the “Onboarding Survey”
2. Start project 0 (due Jan 17)

UC San Diego