“In a world of abundance, the only scarcity is human attention.”

Herbert Simon

Lecture G.1:
Graph Parallel Processing

CS205: Computing Foundations for Computational Science
Dr. David Sondak
Spring Term 2021

Contributions from Bill Richmond (AWS)

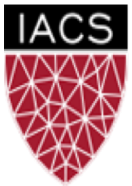

INSTITUTE FOR APPLIED
COMPUTATIONAL SCIENCE
AT HARVARD UNIVERSITY

HARVARD
School of Engineering
and Applied Sciences

Performance Competition Results

Position	Timing (s)
1	0.583
2	1.01
3	1.48

Before We Start

Where We Are

Computing Foundations for Computational and Data Science

How to use modern computing platforms in solving scientific problems

Intro: Large-Scale Computational and Data Science

A. Parallel Processing Fundamentals

B. Parallel Computing

C. Parallel Data Processing

- C.1. Batch Data Processing
- C.2. Dataflow Processing
- C.3. Stream Data Processing

Wrap-Up: Advanced Topics

CS205: Contents

APPLICATION SOFTWARE

APPLICATION PARALLELISM

PARALLEL PROGRAM DESIGN

PROGRAMMING MODEL

Optimization

OpenACC

OpenMP

MPI

Spark

Map-Reduce

PLATFORM

B. BIG COMPUTE

C. BIG DATA

CLOUD COMPUTING

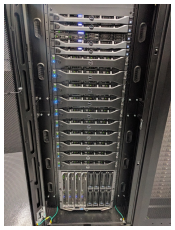

PARALLEL ARCHITECTURES

Context

Graph Parallel Processing

A graph is a mathematical structure that helps us visualize and analyze problems.

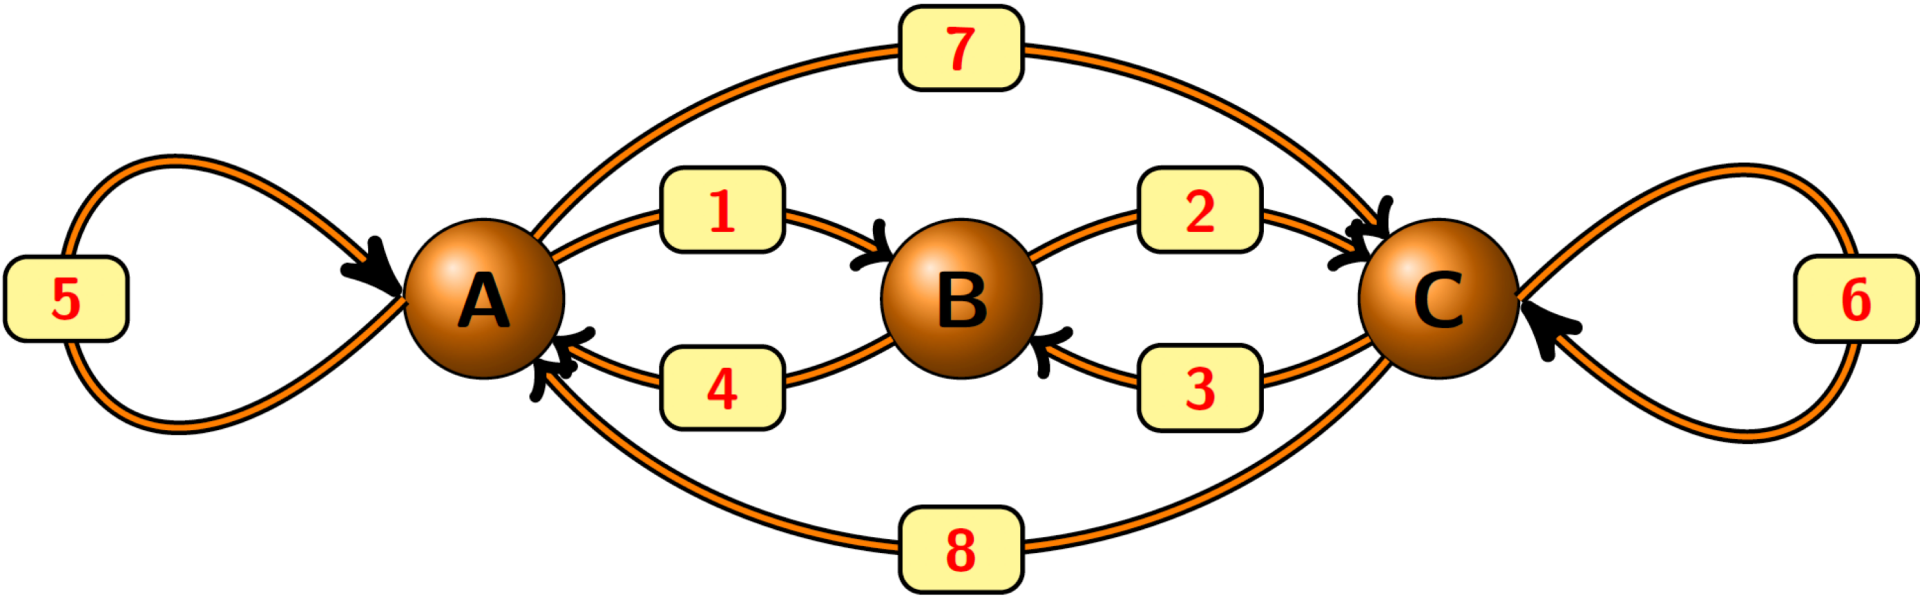

Definitons

Basic Graph Anatomy

A graph G consists of a finite, nonempty set of objects called vertices V and a set of 2-element subsets of V called edges E. A graph is often denoted by $G = (V, E)$.

- **Vertices:** The points in the graph
- **Edges:** The lines connecting the vertices
- **Degree:** The number of vertices
- **Size:** The number of edges
- **Labeled vs. Unlabeled:** Vertices can be labeled or not
- **Multigraph:** Two vertices can be connected by more than one edge
- **Directed graph:** An edge between two vertices is directed from one vertex to another, but not the other way around
- **Weighted:** Edges are labeled with weights

History

The Seven Bridges of Königsberg

Can you walk around Königsberg by crossing each bridge only one time?

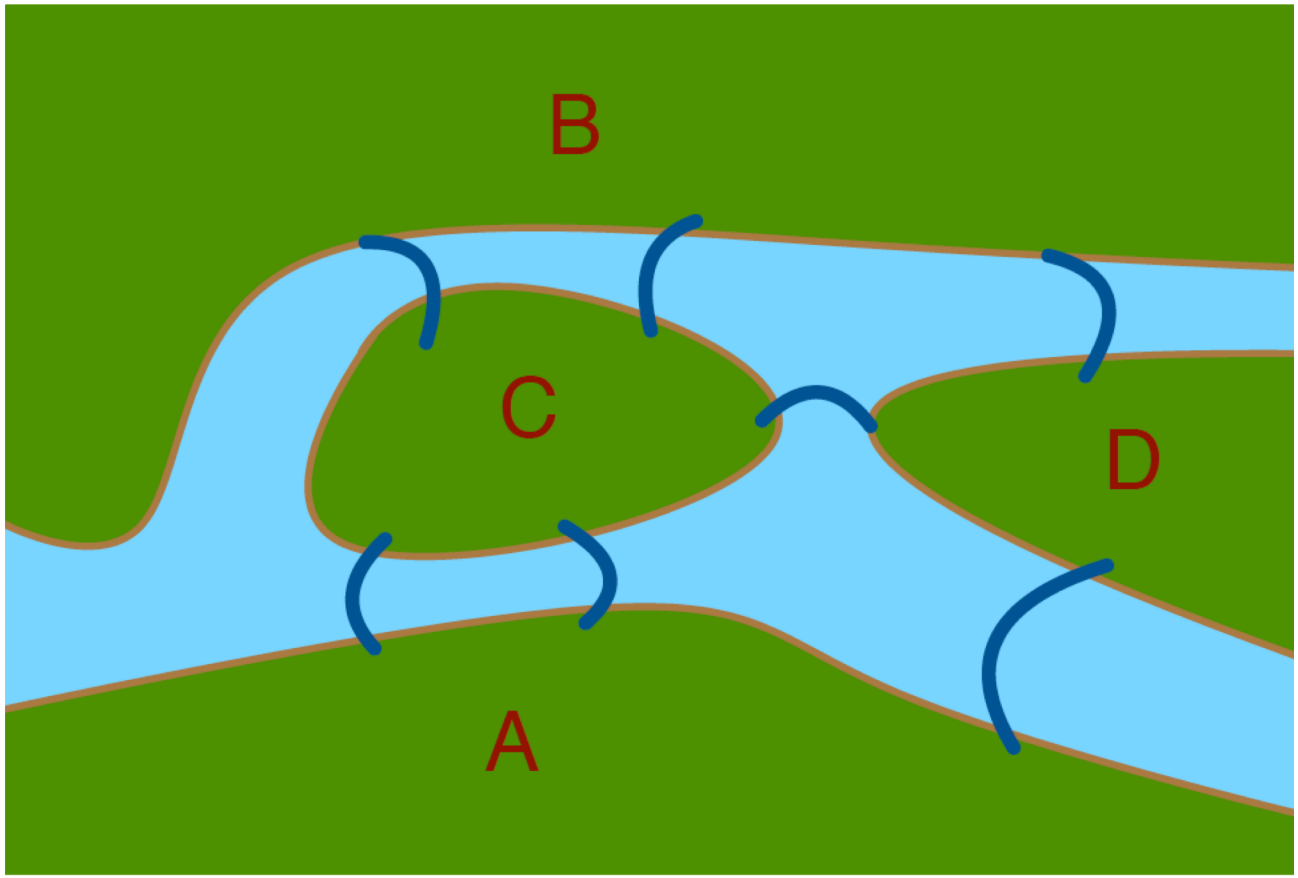

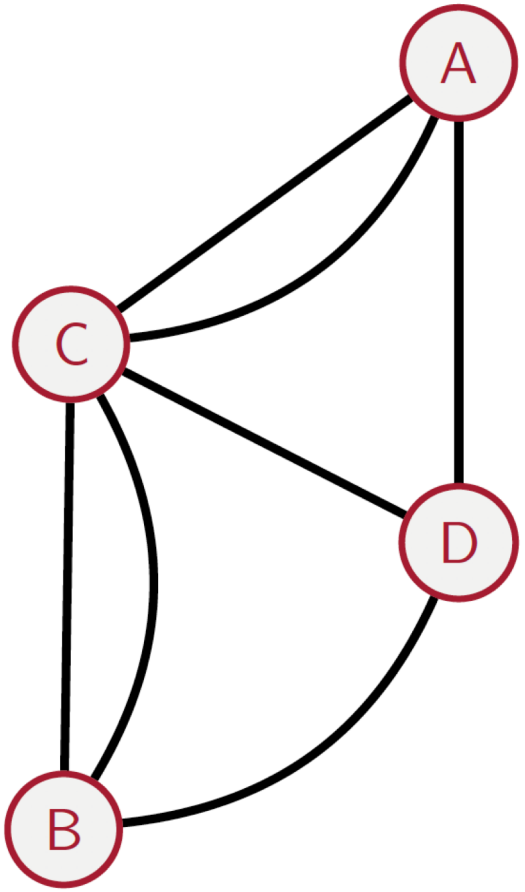

History

Solution to the Seven Bridges of Königsberg

Can you walk around Königsberg by crossing each bridge only one time?

- The number of times you can enter a non-terminal vertex equals the number of times you leave a non-terminal vertex
- Consequently, the number of bridges touching non-terminal land masses must be even
    - Otherwise the bridge would be crossed more than once

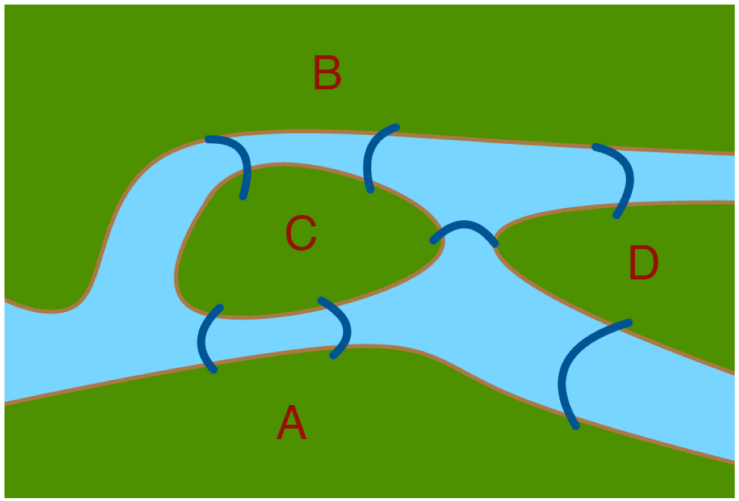

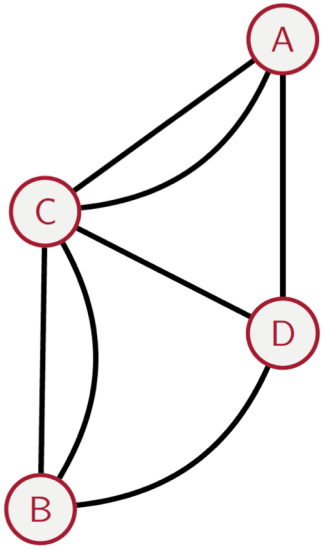

History

Solution to the Seven Bridges of Königsberg

Can you walk around Königsberg by crossing each bridge only one time?

- For the Königsberg configuration, we only have two non-terminal land masses
- In our graph, each land mass is touched by an odd number of bridges

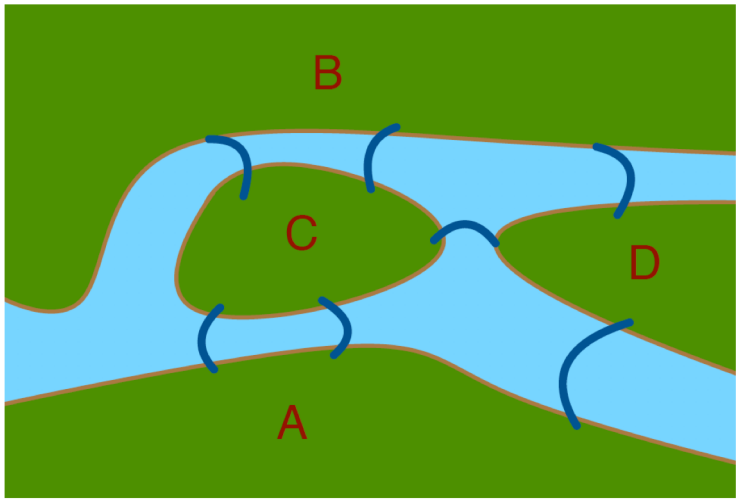

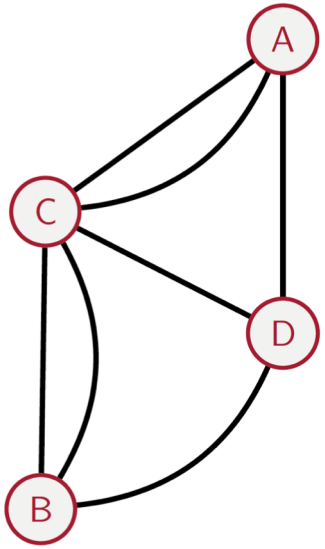

- Answer: **NO**

Modern Graphs

Highly Connected Data

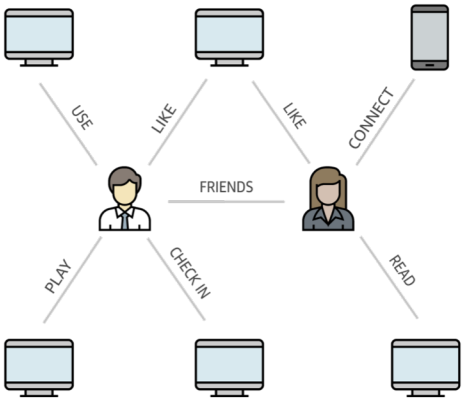

Social Networks

- Build a news feed into your app
- Prioritize showing your users the latest updates from their family, from their friends who's updates they “Like” a lot and from friends that live close to them

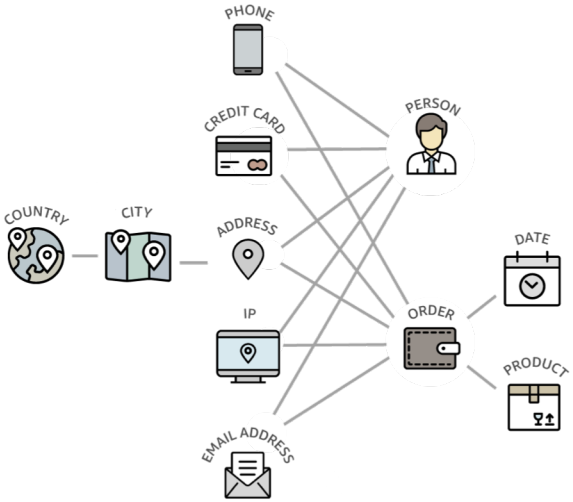

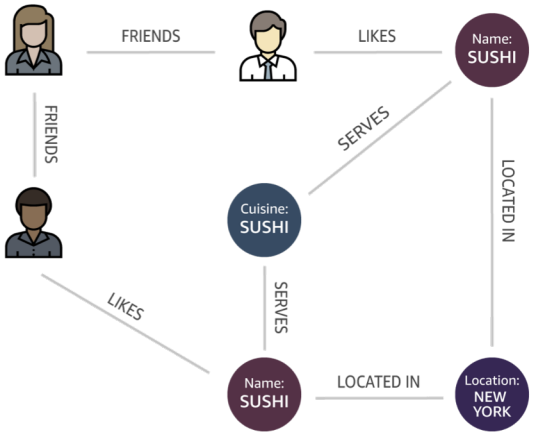

Restaurant Recommendations

- Restaurant recommendation app
- Provide recommendations of restaurants of a certain cuisine like Sushi in a city like New York that at least two of the users friends also like

Retail Fraud Detection

- Retail fraud detection app
- Easily detect relationship patterns like multiple people associated with a personal email address or multiple people sharing the same IP address but residing in different physical addresses.

Natural to express this data as a graph. Very unwieldy to express in a relational model.

Modern Graphs

Use Cases for Highly Connected Data

Social Networking

Recommendations

Knowledge Graphs

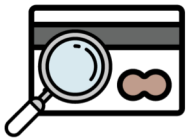

Fraud Detection

Life Sciences

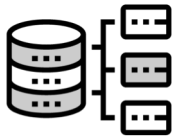

Network & IT Operations

Modern Graphs

Examples of Connected Data Queries

- Which friends and colleagues do we have in common?
- Which applications and services in my network will be affected if a particular network element – a router or switch, for example – fails? Do we have redundancy throughout the network for our most important customers?
- What's the quickest route between two stations on the underground?
- What do you recommend this customer should buy, view, or listen to next?
- Which products, services and subscriptions does a user have permission to access and modify?
- What's the cheapest or fastest means of delivering this parcel from A to B?
- Which parties are likely working together to defraud their bank or insurer?
- Which institutions are most at risk of poisoning the financial markets?

Modern Graphs

Recommendations Based on Relationships

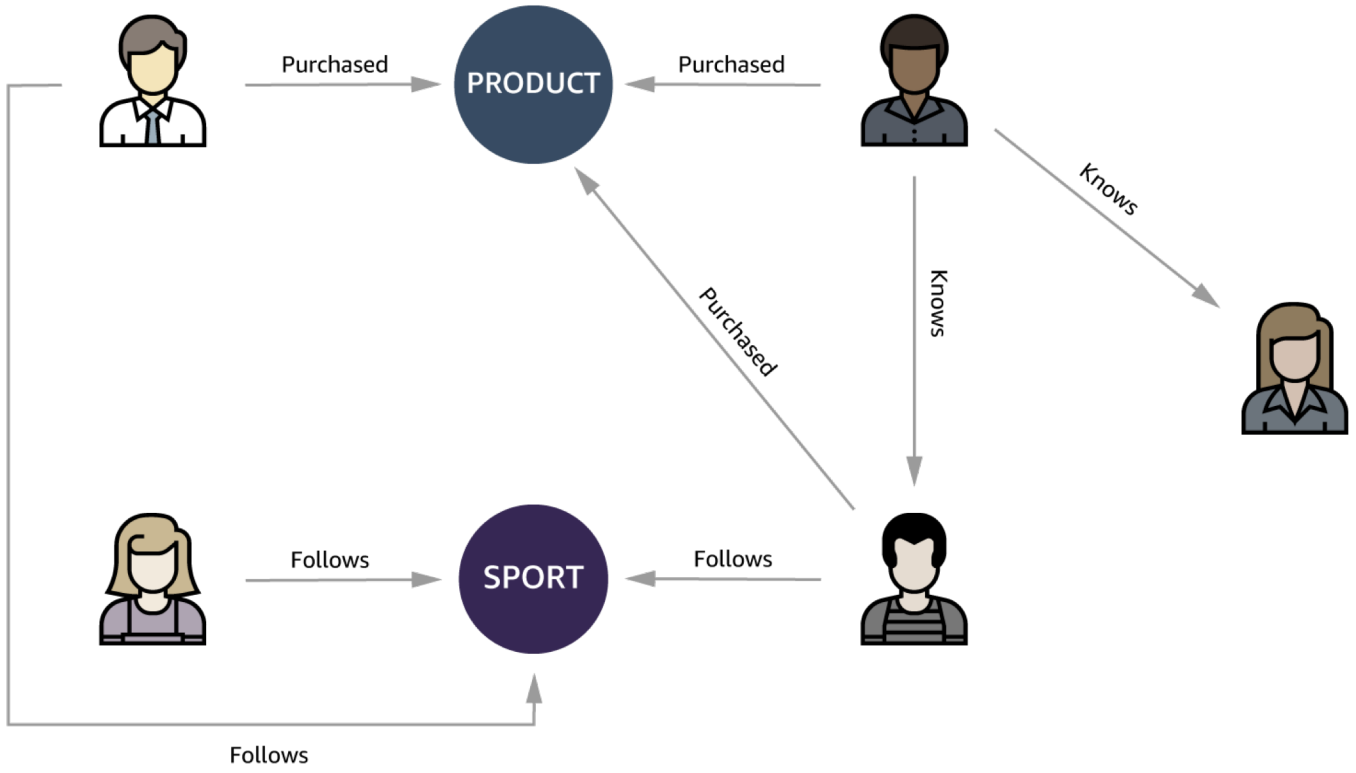

- Store relationships between information (customer interests, friends, purchase history).
- Quickly query it to make personalized, relevant recommendations.

Modern Graphs

Knowledge Graph Applications

Who painted the Mona Lisa?

What museums should Alice visit while in Paris?

What artists have paintings in The Louvre?

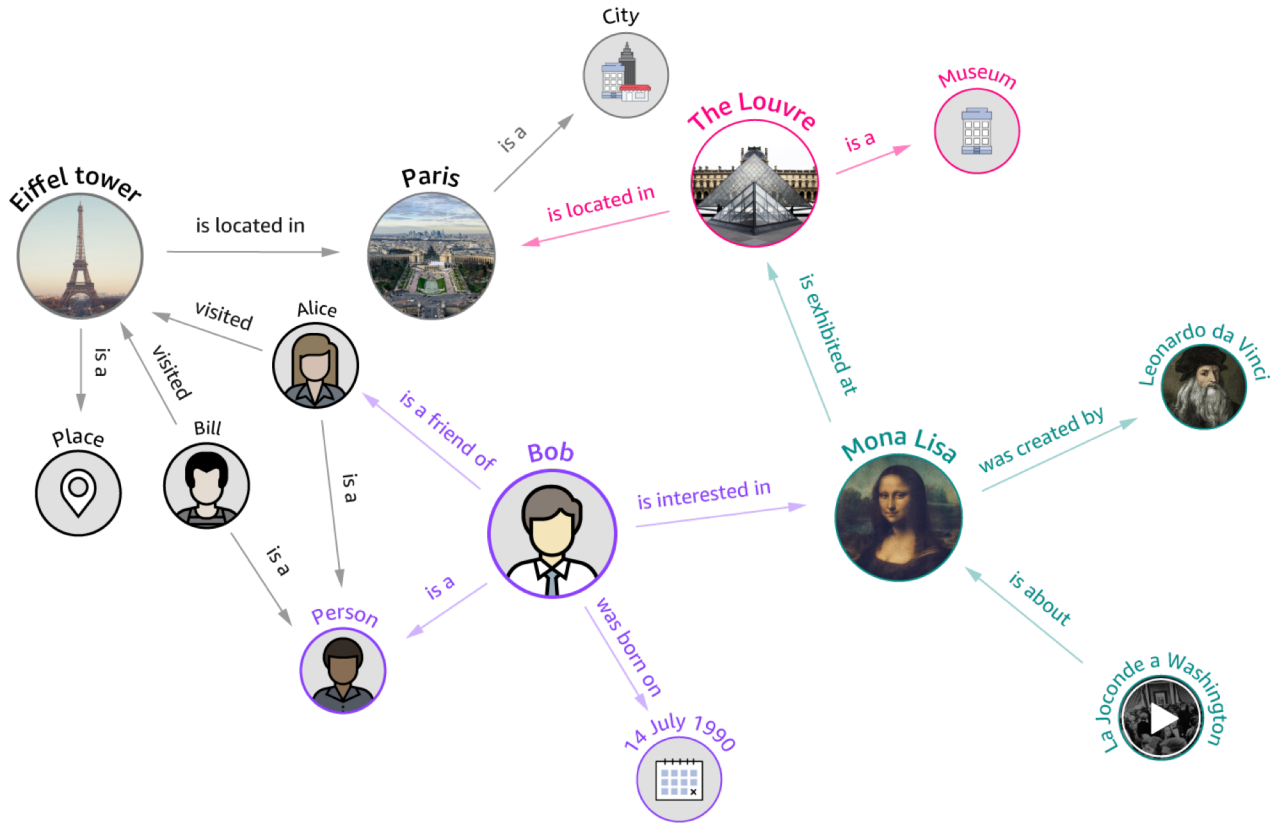

How Large is Large?

Facebook

- 2.50B+ Monthly Active Users
- 1.66B+ Daily Active Users
- 1.15B+ Mobile Daily Active Users
- 94% of advertising revenue comes from mobile ads
- 83 million fake profiles
- Like & Share Buttons are viewed across 10M websites daily
- Age 25 to 34 (30% of users) is the most common age demographic
- 50% of 18-24 year-olds go on Facebook when they wake up
- Highest traffic occurs mid-week between 1 to 3 pm
- A 7pm post will result in more clicks on average than posting at 8pm
- On Thursdays and Fridays, engagement is 18% higher
- One in five page views in the United States occurs on Facebook
- Every 60 seconds on Facebook: 510K comments are posted, 293K statuses are updated, and 136K photos are uploaded

Imagine trying to use all of this related data (and much more – who is friends with who, is interested in what, etc.) to effectively run a business!

How Large is Large?

Siemens Smart Infrastructure

- In 2018, there were eight billion devices connected to the internet; by 2030, there will be about one trillion, according to a report by the World Economic Forum. These connected devices include components of the systems that buildings use for vital functions like fire prevention, security and access, HVAC, lighting, and power. A typical smart office building has about 60 types of sensors generating more than 500 MB of data a day—a volume projected to double every two years.

- Siemens Smart Infrastructure focuses on connecting energy systems, buildings, and industry and is interested in the prospect of increasingly sophisticated sensors generating more and more building data. For example, if an HVAC system could use an access-control system's data, it could automatically increase the air conditioning as a conference room fills up, and then turn it down again after the meeting is over.

- Siemens uses graph databases (Amazon Neptune) to model the complex object dependencies in building-generated datasets.

How Large is Large?

FINRA

- The Financial Industry Regulatory Authority (FINRA) writes and enforces rules governing the activities of more than 3,800 broker-dealers representing more than 600,000 brokers, examines firms for compliance, fosters market transparency, and educates investors.

- Every day, FINRA oversees up to 75 billion market events—99 percent of equities trades and 65 percent of options trades in the United States—applying data analytics to uncover insider trading and other strategies used to gain an unfair advantage.

- Broker-dealers must submit daily electronic data to FINRA, adding up to more than 50,000 files. As soon as data is received, FINRA validates it to ensure it is complete and correctly formatted according to a set of more than 200 rules. The system performs up to half a trillion validations each day. Processing demand varies significantly over time and can double or triple in response to market conditions that drive higher trading volumes.

Graph Algorithms

Some Graph Algorithms

- **Traversal**
  - Visit vertices of a graph in a certain order
    - e.g. depth-first search and breadth-first search
- **Topological Sort**
  - Sort the vertices according to some criterion
- **Strongly Connected Components**
  - Maximal strongly connected subgraphs of the graph G
    - A graph is strongly connected if all vertices are reachable from every vertex
- **Label Propagation**
  - Determine communities in a network
- **PageRank**
  - Rank the relative importance of vertices in a graph
- **Triangle Count**
  - Counts the number of triangles passing through a given vertex

Graph Analysis Packages

GraphX and GraphFrames

GraphX

- GraphX is part of Spark
- Extends the Spark RDD by introducing a Graph abstraction
- Provides a suite of graph operations, builders, and algorithms
- Only works with Scala

GraphFrames

- Not part of Spark
- It *is* designed to be used with Spark
- Graphs are built on Spark dataframes rather than RDDs
- Contains the same functionality as GraphX (and more!)
- Works with Scala, Java, and Python
- Can be challenging to set up

Graph Analysis Packages

GraphFrames Exercises

Spin up an AWS instance (may be good to use the one that you created with Spark in local mode) and `ssh` to it from your terminal.

Launch Pyspark with the following command:

```
pyspark --packages graphframes:graphframes:0.6.0-spark2.3-s_2.11
```

And follow the instructions on the course website:

https://harvard-iacs.github.io/2021-CS205/lectures/G1/

Next Steps

- Wellness day on Thursday
- HWC due on Monday 4/19!
- Final Project (upcoming milestones):
  Project design (4/20 and 4/22) --- Same order as last time
  Project presentation (5/10) --- Stay tuned for details
  More info at:
  https://harvard-iacs.github.io/2021-CS205/

Project Requirements

- Demonstrate the need for big compute and/or big data processing, and what can be achieved thanks to large-scale parallel processing.
- Solve a problem for a non-trivial computation graph and with hierarchical parallelism.
- Be implemented on a distributed-memory architecture with either a many-core or a multi-core compute node, and evaluated on at least 8 compute nodes (note: each compute node on Cannon is a multi-core with 32, or 64 cores or with a many-core GPU with hundreds of cores)
- Use a hybrid parallel program in either, for example: MPI + OpenMP, MPI + OpenACC (or OpenCL), Spark or MapReduce + OpenACC (or OpenCL) or MPI + Spark or MapReduce
- Be evaluated on large data sets or problem sizes to demonstrate both weak and strong scaling using appropriate metrics (throughput, efficiency, iso-efficiency...).

Project Design Presentation

You will have **5, and ONLY 5, minutes** to briefly summarize your proposal answering bellow questions. You have to prepare 4 slides for your proposal. We will enforce the 5-minute time limit.

Show that you **understand** the application.

Specify the **programming model** and **infrastructure** that you will use, and if code exists, provide an analysis/profiling of the existing sequential/parallel code.

Describe the main overheads in the parallelization and techniques that will be applied to mitigate them.

Describe the **numerical complexity** of the algorithm.

Estimate the **theoretical speed-up** and **scalability** expected.

Zoom Presentation Guidelines

Record your video ahead of time and submit on Canvas one day before your presentation.

Each group member should present.

Practice ahead of time!

Make sure your mics are muted when you are not presenting.